Matching Efficiencies in Platform Consolidation: Evidence from the Grab–Uber Merger*

Dingzhe Leng†, Bingzhao Li‡, Juin Kuan Chong§

October 26, 2026

Abstract

Platform mergers can reduce competition while improving matching efficiency. We study the 2018 Grab–Uber merger in Singapore. Trip volumes increase despite higher post-merger prices, contradicting models in the literature that predict declines. We show that consolidation unlocks latent demand by reducing matching failures through scale efficiencies. Incorporating matching probabilities into structural estimation allows the model to accurately capture the direction of post-merger outcomes and reshape merger welfare assessments. Our findings inform antitrust policy that matching probability matters in platform mergers.

*We are grateful to Nano Barahona, Dante Donati, George Gui, Benjamin Handel, Ganesh Iyer, Kinshuk Jerath, Przemyslaw Jeziorski, Sarah Moshary, Olivia Natan, Oded Netzer, Nicola Rosaia, Camilla Schneier, Stephan Seiler, Andrey Simonov, Steve Tadelis, J. Miguel Villas-Boas, and Yihao Yuan for helpful feedback and comments. This paper benefits from participants at 2026 ISMS Marketing Science Conference, marketing brown-bag seminars at UC Berkeley and Columbia Business School, and IO therapy and workshop at UC Berkeley. Bingzhao Li also thanks funding support from Deming Doctoral Fellowships at Columbia Business School. All errors are our own.

†University of California, Berkeley, dingzhe_leng@berkeley.edu

‡Columbia Business School, bl3003@gsb.columbia.edu

§Nanyang Technological University, juinkuan.chong@ntu.edu.sg

1 Introduction

Consolidation among digital platforms involves billions of dollars and affects services used in everyday life (Li and Agarwal 2017; Farronato, Fong, and Fradkin 2024).[1] Platform mergers can increase market power while improving efficiency, leaving their welfare effects ambiguous (Jeziorski 2014; Ivaldi and Zhang 2022; Rosaia 2025). What drives the welfare effects of platform mergers? We show that explicitly accounting for pre-merger matching failures has first-order implications for merger welfare assessments: pooling users across platforms can enable transactions that would otherwise fail.

We study the 2018 Grab–Uber merger in Singapore, where completed trips increased during peak periods despite higher post-merger prices. Yet a demand specification commonly used in the ride-hailing literature predicts the opposite: a decline in completed trips even at the observed post-merger fares and pickup times (Buchholz et al. 2025; Castillo 2025; Rosaia 2025). We argue that observed trips reflect both riders' willingness to take a ride and their probability of being successfully matched. Before the merger, some riders who were willing to take a ride could not secure one because matching took longer than they were willing to wait. Pooling riders and drivers can reduce matching frictions by increasing effective market density (Frechette, Lizzeri, and Salz 2019).[2] In our setting, consolidation can shorten matching delays and increase the probability of securing a driver before the rider abandons the request. We define matching probability as the probability that a ride request is assigned a driver before the rider cancels or the platform's search window expires. An increase in this probability can raise completed trip volume even when higher prices reduce potential demand.

We incorporate matching probability into a rider demand model. Potential demand depends on fares, pickup times, and platform preferences; completed demand equals potential demand multiplied by the matching probability.[3] Our trip-level data, covering approximately two million completed trips per week in April and May 2018,[4] separately measure matching time, the interval from request to driver acceptance, and pickup time, the subsequent wait for the driver to arrive. We use the distribution of matching times to infer the probability of securing a driver before a request is abandoned, while pickup

[1] Recent examples include: the DoorDash–Deliveroo merger closed in October 2025, with an announced equity value of £2.9 billion; and the eBay–Depop merger closed in July 2026 for approximately US$1.4 billion, including adjustments and interest. See DoorDash's transaction terms, DoorDash's completion announcement and Etsy's completion announcement.

[2] The returns to scale in matching vary across settings. Frechette, Lizzeri, and Salz (2019) find increasing returns in taxi matching, whereas Cullen and Farronato (2021) find approximately constant returns on TaskRabbit. Our matching-time data allow us to examine how matching success changes when previously separate rider and driver pools are combined.

[3] The distinction between potential and completed demand parallels the problem of estimating demand under incomplete product availability (Conlon and Mortimer 2013).

[4] The data cover both pre- and post-merger periods. The transaction was announced on March 26, and Uber ceased operating in Singapore on May 7. The April baseline therefore captures continued platform competition after the announcement and before Uber's exit.

time enters rider utility as a time cost. This distinction allows us to recover matching gains even when unsuccessful requests are unobserved.[5] It matters because ignoring matching failures can change the demand parameters estimated from completed trips, with consequences for both post-merger demand predictions and welfare evaluation.

Observations before and after the merger enable us to test the model against post-merger outcomes. We estimate two demand specifications using pre-merger data: the benchmark specification, which assumes full matching, and our specification, which incorporates matching probability.[6] We then compare their predictions for completed trips at the same observed post-merger fares and pickup times. Accounting for matching probability improves predictions of trip changes in all 15 region–time-window cells. For example, during the morning peak, observed trips increase by 2.83%; the benchmark predicts a 10.37% decline, while our model with matching predicts a 1.19% increase. This adjustment also changes the welfare assessment at observed post-merger fares and pickup times: the model with matching probability implies an aggregate consumer-surplus loss of 0.60%, compared with 10.71% under the benchmark.[7]

These post-merger outcomes, however, reflect a short-run equilibrium: driver supply had little time to adjust, and from April 13 the CCCS required the platforms to keep pre-merger fares and surge factors.[8] Merger review turns on the long run, where better matching raises driver earnings and participation, which feeds back into matching and pickup times (Bimpikis, Candogan, and Saban 2019; Besbes, Castro, and Lobel 2021). We therefore embed the validated demand model in an equilibrium model of driver supply and platform pricing (Rochet and Tirole 2003; Armstrong 2006), building on Rosaia (2025), and use the pre-merger estimates to simulate the market after Uber's exit with free driver entry and unrestricted fares. The result shows that accounting for matching reveals lower price elasticities and changes the predictions substantially in such conditions. The average fare is 2.23 SGD per kilometer, compared with 1.75 without matching, despite nearly identical pickup times, and the aggregate consumer-surplus loss also increases from 7.43% to 11.97%, with some markets reversing from gains to losses in equilibrium.[9] The short-run and long-run exercises demonstrate that ignoring matching probabilities distort merger welfare assessments.

We also examine the effects of exclusive driver contracts. Restrictions on exclusivity

[5]Researchers typically do not observe unrealized trips in trip-level data. Even datasets that track app-checking behavior miss some potential riders: many form expectations about matching time and, anticipating a long wait, choose another mode without ever opening the app.

[6]The matching specification uses estimated pre-merger matching probabilities and sets post-merger matching probabilities to one.

[7]Each percentage is relative to the corresponding pre-merger consumer surplus baseline.

[8]On April 13, 2018, CCCS issued interim measures requiring the platforms to maintain pre-merger pricing, including base fares and surge factors. See CCCS's announcement.

[9]This contrasts with the short-run result, where accounting for matching reduced the estimated loss, so the short-run and long-run exercises show that accounting for matching failures can change merger welfare assessments in either direction, depending on market conditions.

address a central concern in platform regulation: exclusive contracts can limit service providers' outside options and restrict competing platforms' access to the supply needed to compete (Armstrong and Wright 2007; Lee 2013).[10] In ride-hailing, exclusivity restrictions also affect matching: a request may remain unmatched even when an idle driver is available nearby on a competing platform. Removing exclusivity could therefore improve matching through the reallocation of existing drivers while preserving platform competition. We simulate the removal of exclusivity by allowing drivers to choose which platform to work for each hour, while holding the total driver population at its pre-merger level.[11] Theoretical studies show that participation across platforms can have conflicting welfare effects (Belleflamme and Peitz 2019; Bernstein, DeCroix, and Keskin 2021); our exercise quantifies how changes in driver exclusivity determine the consequences for riders. The simulation shows that non-exclusive arrangements increase matching probabilities and shorten pickup times, but stronger competition for drivers also raises fares to support higher driver compensation.[12] Removing exclusivity thus produces the same direction of fare and pickup-time changes as a merger, despite intensifying competition for drivers. Stronger competition for drivers can therefore hurt riders through higher prices even as it improves their chances of obtaining service.

Our paper is most closely related to Rosaia (2025), who develops a spatial equilibrium model of competing ride-hailing platforms and evaluates a hypothetical merger in New York City. We complement this work with evidence from an observed merger, allowing us to assess model predictions using pre- and post-merger data. We also extend the demand model to distinguish riders' willingness to use the service from the probability that their requests are successfully matched. The extension allows us to assess how accounting for unsuccessful requests changes demand predictions and merger welfare calculations. The importance of matching extends beyond ride-hailing: fragmentation limits gains from exchange in kidney markets (Agarwal et al. 2019), while thicker housing markets can improve match quality and facilitate transactions (Ngai and Tenreyro 2014). Farronato, Fong, and Fradkin (2024) document matching benefits from combining platform user bases and show that losses from reduced platform differentiation can offset these benefits. These findings underscore the broader relevance of matching for evaluating platform mergers that combine previously separate pools of participants.

The remainder of the paper is organized as follows. Section 2 presents the institutional background, data, and descriptive evidence. Section 3 develops and estimates the demand model and assesses its post-merger predictions. Section 4 evaluates consumer-surplus

[10] Singapore's point-to-point transport regulatory framework, introduced in October 2020, required licensed operators to maintain non-exclusive arrangements with drivers. See LTA's announcement.

[11] Allon, Chen, and Moon (2025) estimate drivers' switching and repositioning decisions across platforms and locations. Our counterfactual focuses on how relaxing exclusive affiliation affects matching failures, equilibrium fares, and rider welfare.

[12] Drivers receive a fixed share of fares in our setting. We keep this assumption in the simulation.

changes. Section 5 presents counterfactual simulations of the long-run post-merger equilibrium and non-exclusive driver arrangements. Section 6 concludes.

2 Institutional Background, Data, and Facts

2.1 Grab–Uber Merger

Before the taxi fleet peaked in 2014, taxi companies were the primary players in chauffeured point-to-point transportation in Singapore, with most trips being street-hailed or phone-booked. Uber Technologies, Inc. is an American multinational transportation company founded in 2009. It entered the Singapore market in February 2013, introducing a service model very similar to the one used in the United States. Uber connected riders to private-hire drivers and used dynamic pricing with upfront credit-card payment. GrabTaxi was originally launched as MyTeksi in Malaysia in 2012 and expanded to Singapore in October 2013, the same year Uber began operations there. Initially, GrabTaxi functioned as a taxi-booking mobile application that relied on licensed taxis and metered fares. In 2015, GrabTaxi launched GrabCar, which allowed private-hire drivers to accept ride requests through the app. The company rebranded as Grab in January 2016. From then on, Uber and Grab aggressively competed to attract drivers and riders through promotions and incentive schemes.

In Singapore, most drivers single-homed on either Uber or Grab. Because of the high upfront costs associated with purchasing a car, especially in obtaining a Certificate of Entitlement,[13] most drivers in Singapore rent from platform-owned fleets under exclusive contracts. According to the Notice of Interim Measures Directions by the Competition and Consumer Commission of Singapore (henceforth CCCS), these exclusivity arrangements effectively prevent drivers from multi-homing, as they are contractually obligated to drive exclusively for one platform.[14] Breaching these exclusivity agreements results in penalties. Even when drivers were not bound by contractual agreements, the costs associated with multi-homing were prohibitively high due to limited attention.

The convenience of online platforms and lower prices than taxis led to a rapid increase in the number of users on these platforms. Grab excelled in catering to local needs, such as offering cash payment options.[15] By 2017, Uber and Grab had become the dominant

[13]Vehicle ownership in Singapore requires a Certificate of Entitlement (henceforth COE), a quota license valid for 10 years. COEs are allocated through an open, uniform-price auction. Due to strict supply quotas, the price of a COE frequently exceeds the market value of the vehicle itself.

[14]If drivers could freely multi-home, this cross-platform access constraint would be attenuated, and the matching gains from consolidation would likely be smaller. Driver multi-homing, however, would not fully replicate an integrated platform, since competing platforms would still observe different rider request pools, run separate dispatch algorithms, and face coordination problems when assigning the same driver across simultaneous requests.

[15]Over half of the Southeast Asian population preferred cash transactions over credit card payments during the initial years.

ride-hailing platforms in Singapore, capturing market shares of 20% and 50% in terms of fleet sizes, respectively.[16] Taxis accounted for the remaining 30% market share.

On Monday, March 26, 2018, Grab announced its acquisition of Uber's operations and assets in Southeast Asia, with Uber acquiring a 27.5% stake in Grab.[17] Following the announcement, Uber began redirecting its riders to download the Grab app and informed them that Uber services would no longer be available in Southeast Asia. Grab described the deal as "the largest-ever of its kind" in Southeast Asia.[18]

On March 27, 2018, CCCS initiated an investigation into the transaction under Section 62 of the Competition Act. CCCS noted that Grab and Uber overlapped in two markets: the provision of ride-hailing services through their respective platforms and the rental of private-hire vehicles through their affiliated subsidiaries. As previously mentioned, most platform drivers do not own their cars. Instead, they rent vehicles from Lion City Holdings Pte. Ltd., a subsidiary of Uber, and Grab Rentals Pte. Ltd., a subsidiary of Grab. The cars and Certificates of Entitlement held by these subsidiaries represent valuable assets that are difficult for potential entrants to obtain.

In its final decision, CCCS prohibited Uber from selling its private-hire vehicle rental business to Grab. Accordingly, in this paper we focus only on the part of the transaction that was ultimately carried out: Grab's acquisition of Uber's ride-hailing operations, which was implemented through the shutdown of the Uber platform. Uber ultimately ceased its ride-hailing service in Singapore on May 7, 2018.[19]

2.2 Data

Our primary dataset comprises trips completed on Singapore's two ride-hailing platforms, Uber and Grab.[20] Because of the merger, Uber shut down its app in Singapore on May 7, 2018. Our data cover April, the last full month in which both platforms operated, and May, when Uber exited and Grab became the sole platform. Each observation identifies a chauffeured point-to-point trip and contains information about the driver, platform, pickup and dropoff locations, time, fare and surge multiplier. Time-related variables

[16] Page 19 of the "Notice of Interim Measures Directions" by the Competition and Consumer Commission of Singapore, although Uber contested the accuracy of these market share figures. The CCCS also noted that taxi services, including street-hailing, and online ride-hailing services may be considered part of the same relevant market.

[17] One concern is that the transaction merely reflected Uber's imminent exit from an unprofitable market, in which case post-merger changes could partly reflect pre-existing deterioration rather than consolidation. Contemporaneous accounts did describe Uber's regional retrenchment from costly subsidy competition. CCCS nevertheless concluded that Uber would not have simply exited Singapore absent the transaction and noted that the parties had anticipated potential competition concerns before completion.

[18] `https://www.grab.com/sg/press/business/grab-merges-with-uber-in-southeast-asia/`

[19] On April 13, 2018, CCCS issued Interim Measures Directions requiring Uber to delay the shutdown of its app until May 7, 2018.

[20] All datasets mentioned in this paper are either from the Land Transport Authority (LTA) or from publicly available official government data in Singapore.

include the rider's request time, the driver's acceptance time, and the pick-up and drop-off times. We focus exclusively on the comparable services, UberX and JustGrab, which account for over 80% of trips on each platform.[21]

Table 1 presents summary statistics comparing the representative weeks in April and May, respectively. We restrict the sample to specific weeks, rather than full months, to exclude the transition period and avoid potential bias from anticipatory switching behavior.[22] The data reveal three key patterns: significant heterogeneity in trip volume across time and space, an average increase in pricing, and a substantial improvement in platform matching efficiency.

[21]Before Uber exited Southeast Asia, its platform offered four main product categories: UberX, UberXL, UberPool, and UberExec. During the same period, the Grab app offered five ride services: GrabTaxi, JustGrab, GrabShare, GrabCar, and GrabCarPlus. Except for GrabTaxi, all Grab services employ dynamically determined prices set by the platform. In both ecosystems, the majority of trips are concentrated in the standard private-hire tier: over 80% of Uber trips are UberX, and JustGrab similarly accounts for more than 80% of Grab trips. For the purposes of our empirical analysis, we focus exclusively on UberX and JustGrab trips.

[22]Our observed market shares for the first week of April align with official pre-merger statistics, which indicate that Uber and Grab held 20% and 50% of the market, respectively, while taxis accounted for the remaining 30%. This consistency confirms that the market still reflected pre-merger conduct during this week. For the post-merger period, we use data from the second week of May after verifying that market conditions remained stable in subsequent days, as shown in section 2.1 in Appendix. We therefore treat these two weeks as representative of the pre- and post-merger equilibria for the purpose of Table 1.

Table 1: Summary Statistics for Representative Pre- and Post-Merger Weeks

	Num Trips		Unit Price (SGD/km)		Total Fare (SGD)		Accept Time (min)		Pickup Time (min)		Travel Time (min)		Travel Dist. (km)	
	Pre	Post	Pre	Post	Pre	Post	Pre	Post	Pre	Post	Pre	Post	Pre	Post
Panel A: Average														
	1,959,820	1,995,139	1.54	1.62	12.07	12.80	0.43	0.30	6.34	5.45	17.69	18.29	10.65	11.02
Std. Dev.	–	–	0.95	1.03	5.83	5.66	0.61	0.24	3.24	2.87	8.53	8.78	7.55	7.63
Panel B: By Pickup Region														
Central	963,024	983,478	1.70	1.83	11.53	12.41	0.44	0.30	6.33	5.32	16.62	17.13	9.24	9.52
East	322,700	334,433	1.35	1.39	12.10	12.74	0.40	0.29	6.24	5.29	18.11	18.70	11.74	12.19
West	259,734	259,582	1.39	1.41	13.00	13.59	0.40	0.30	6.32	5.63	19.11	19.95	12.51	13.09
Northeast	256,491	259,242	1.40	1.43	12.19	12.67	0.43	0.31	6.64	5.85	18.23	18.79	10.96	11.29
North	157,871	158,404	1.44	1.47	13.50	14.19	0.43	0.32	6.15	5.63	20.12	21.01	13.47	14.08
Panel C: By Time Window														
7am-10am	429,576	441,733	1.51	1.59	12.64	13.39	0.42	0.29	6.68	5.98	18.77	19.68	10.88	11.29
11am-4pm	652,288	644,568	1.51	1.62	10.61	11.49	0.37	0.28	5.77	4.85	16.84	17.26	9.75	10.15
5pm-8pm	466,918	485,907	1.64	1.75	12.27	13.11	0.43	0.31	6.74	5.60	18.65	19.34	10.39	10.68
9pm-12am	276,574	281,584	1.53	1.57	13.10	13.66	0.51	0.34	6.45	5.62	16.87	17.31	11.72	12.06
1am-6am	134,464	141,347	1.49	1.41	14.43	14.10	0.56	0.33	6.41	5.66	16.69	16.93	12.99	13.28

Notes: Panel A reports means and standard deviations of the entire city; Panels B and C report means only. The first two columns are the total number of trips instead of mean values. All means/std. dev. are computed excluding zero values. Pre is the representative pre-merger week (April 4–10 in 2018, pooling two platforms); Post is the representative post-merger week (May 9–15 in 2018, with Grab only), identified by request date. The two weeks set as representative weeks before and after the merger are fixed throughout the paper. All statistics are computed at the individual-trip level. Price per km is total fare divided by trip distance. Accept time is time elapsed from passenger request to driver acceptance. Pickup time is duration from driver acceptance to passenger pickup. Travel time is duration from passenger pickup to dropoff.

First, trips shifted notably toward peak hours and core business regions. Trip volume increased by 2.83% during the morning peak (7 am–10 am) and by 4.07% during the evening peak (5 pm–8 pm), whereas off-peak midday trips (11 am–4 pm) declined by 1.18%. Geographically, completed trips in the Central and East regions expanded by 2.12% and 3.64%, respectively. Trips remained nearly unchanged in the West and increased modestly in the Northeast and North. We observe a slight aggregate increase of 1.8% in total trips with growth concentrated during peak periods and high-demand regions. Second, unit prices and total fares increase on average. On average, the unit price per kilometer rose by 5.19% (from \$1.54 to \$1.62),[23] and the total fare per trip increased by 6.05% (from \$12.07 to \$12.80).[24] Travel time and distance increased by 3.39% and 3.47%, respectively, which could be attributed to shifts in rider composition, changes in trip purposes, or increased traffic congestion. Finally, the platform achieved marked gains in operational efficiency. The average matching time dropped by 30.23%, falling from 0.43 to 0.30 minutes, and the standard deviation of matching time narrowed significantly from 0.61 to 0.24 minutes.

We also obtain complete administrative records for all taxi trips in Singapore over the same period. Because most taxi trips originate from street-hailing or phone bookings, there is no centralized matching, and pickup times are not observed. Otherwise, the taxi data include the same information available for ride-hailing trips. We treat taxis as the riders' outside option. In addition, we use auxiliary postcode-level mapping data. In Singapore, each building is assigned a unique postcode. This allows us to map trip origins and destinations to specific locations. For residential buildings, we link postcodes to rental price data to proxy for riders' income.

2.3 Evidence on Platform Matching

Table 1 shows that average matching time fell by 30% after the merger, while its standard deviation fell by more than half. This section examines the distribution underlying these averages. We distinguish two margins of matching. The intensive margin concerns the distance between the rider and the assigned driver, which determines pickup time. Existing models of ride-hailing demand commonly account for this margin (Buchholz et al. 2025; Castillo 2025; Rosaia 2025). However, this margin is conditional on a successful match: if no driver is assigned, pickup time is not well defined. We therefore introduce the extensive margin: whether a request is assigned a driver at all. Although failed requests are absent from completed-trip data, the distribution of matching times among completed trips provides information about this margin, because a request with a long

[23]Unless otherwise noted, all monetary values are expressed in Singapore dollars (SGD). The average exchange rate as of April 2018 was approximately 1 SGD $\approx$ 0.74 USD.

[24]The exception is 1am–6am, when the average unit price falls from SGD 1.49 to SGD 1.41 per kilometer and the average total fare falls from SGD 14.43 to SGD 14.10.

matching time came close to being canceled.

We focus on the distribution of matching times on JustGrab, the platform observed in both periods. Figure 1 plots the distribution of matching times in three representative markets, each defined by a pickup planning area and a request hour: Woodlands at 8am, a residential area during the morning peak; Downtown Core at 6pm, the central business district during the evening peak; and Bedok at 12pm, a residential area during non-peak hours.

Three patterns emerge. First, pre-merger matching times exhibit a long right tail. Most requests are matched within about half a minute, but the cumulative distribution then flattens, with a small share of requests in both peak markets waiting more than two minutes. By comparison, pickup times do not exhibit a comparable right tail, as shown in figure 5 in the appendix. Second, matching times fall after the merger in both peak markets, with reductions concentrated in the upper part of the distribution. The two curves nearly coincide at very short matching times but diverge as matching time increases. The horizontal distance between them measures the reduction in matching time at a given percentile: it is small at low percentiles and increases toward the upper end of the distribution. Third, in the non-peak market, the two distributions are nearly identical. These patterns are consistent with consolidation expanding the pool of drivers available to each request, particularly in peak markets.

Figure 2 shows that these patterns hold more broadly across peak markets.[25] For each combination of planning area and request hour during the morning and evening peaks, the figure plots the 1st through 99th percentiles of post-merger matching times against the corresponding pre-merger percentiles. Points on the 45-degree line indicate no change, while those below it indicate shorter matching times after the merger. At low and intermediate percentiles, most curves lie close to the 45-degree line, with some slightly above it. At high percentiles, nearly all curves lie well below it: matching times of about two minutes before the merger correspond to roughly 1.1 to 1.5 minutes afterward. The merger thus shortened matching times mainly by eliminating long waits.

[25]The largest reductions occur where and when demand peaks. During the morning peak, the markets with the largest average reductions are residential areas in the North and West, such as Woodlands, Sembawang, Bukit Panjang, and Bukit Batok, where commuting trips originate. During the evening peak, they are central areas, such as Outram, Singapore River, Tanglin, and Downtown Core, where commuters begin their trips home.

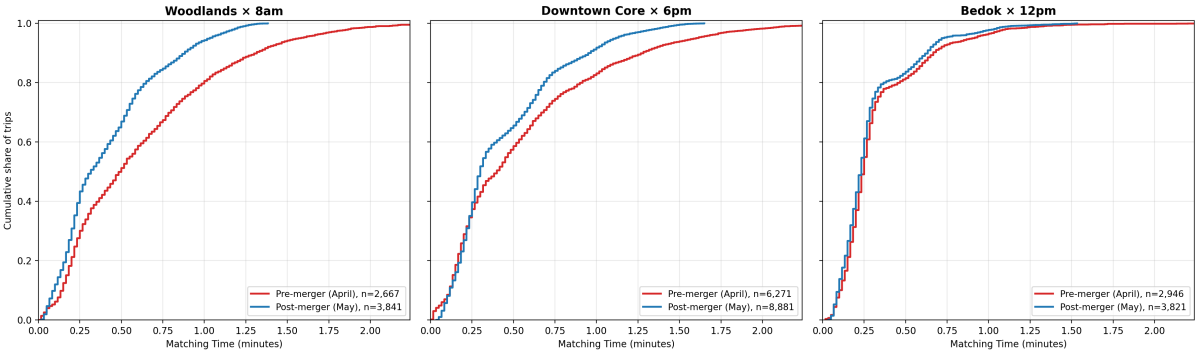


Figure 1: Distribution of Matching Time in Three Markets Before and After the Merger

Notes: Each panel plots the empirical cumulative distribution function of matching time, the time from a rider's request to driver acceptance, in one market defined by pickup planning area and request hour. The sample consists of completed JustGrab trips requested on weekdays during the representative pre-merger week and post-merger week, with five weekdays in each. Legends report the number of trips.

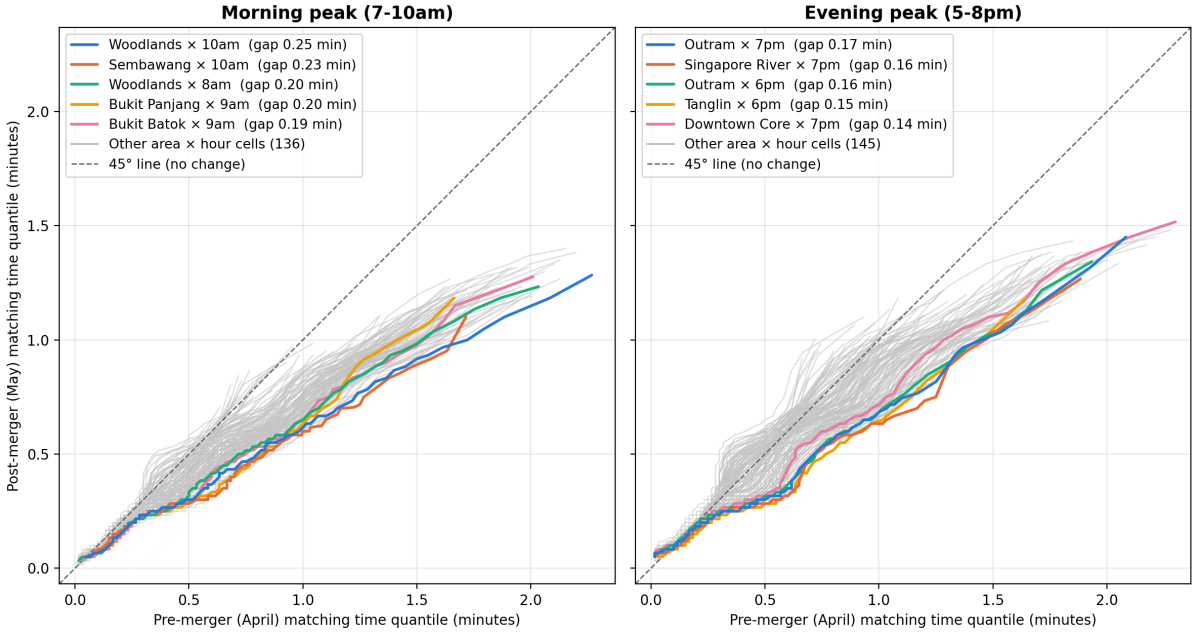


Figure 2: Quantile–Quantile Plots of Matching Time in Peak Markets

Notes: Each line corresponds to one market, defined by pickup planning area and request hour, and plots the 1st to 99th percentiles of the post-merger matching-time distribution (vertical axis) against the same percentiles of the pre-merger distribution (horizontal axis). Points below the dashed 45-degree line indicate shorter matching times after the merger. The sample consists of completed JustGrab trips requested on weekdays during the representative pre-merger week and post-merger week: 141 markets during the morning peak (7–10am) and 150 during the evening peak (5–8pm). In each panel, the five markets with the largest average gap, defined as the mean difference between pre- and post-merger matching times across the 1st to 99th percentiles, are highlighted in color, and the legend reports each gap in minutes.

Taken together, these findings show that the merger shortened matching times primar-

ily by eliminating long waits, with the largest reductions in peak markets. This matters for completed trips because long waits lead to cancellations. Completed-trip data, however, include only successfully matched requests, so changes in the frequency of failed requests cannot be observed directly. The next section therefore develops a model to infer the probability of a successful match and incorporates this probability into demand estimation.

3 Demand Model

We model rider demand in two stages. First, each potential rider chooses among requesting a trip on Grab, requesting a trip on Uber, and taking an outside option. Second, the chosen platform searches for a nearby idle driver, and the request results in a completed trip only if a driver accepts it before it is canceled. Because our data include only completed trips, observed trip volume equals potential demand multiplied by the probability of a successful match.

We present the model in Section 3.1 and estimate its parameters in Section 3.2 using only pre-merger data. In Section 3.3, we assess the model's predictive performance by predicting post-merger demand conditional on observed post-merger market characteristics. This provides an out-of-sample validation of the model. Finally, we examine the welfare implications by evaluating post-merger rider surplus using the estimated demand model.

3.1 Model Setup

We build on the demand model of Rosaia (2025). We divide the city into I regions and define a route as an origin–destination pair (i, j), where $i, j \in \{1, \ldots, I\}$. The index $x \in \{(\text{Monday}, 1 \text{ am}), \ldots, (\text{Sunday}, 12 \text{ am})\}$ denotes the hour of the week.[26] We model demand at the route–hour level and refer to each route–hour cell (x, i, j) as a market. Each market contains a mass M_{xij} of potential riders who choose between requesting a trip on one of the two ride-hailing platforms and taking the outside option. We index the platforms by $m \in N_{RH} = \{1, 2\}$, where $m = 1$ denotes Grab and $m = 2$ denotes Uber, and denote the outside option by $m = 0$. We use taxis as the outside option because they are the closest substitute observed at the same route–hour level.[27]

Riders evaluate each platform based on its fare p^m_{xij} and pickup time w^m_{xi}, defined as the time from driver acceptance to rider pickup. Pickup time varies by origin but not by

[26] Table 3 in Appendix provides a complete overview of the model's notation.

[27] Other modes, including public transit, private vehicles, walking, and cycling, are not observed at the route–hour level. The outside option should therefore be interpreted as the observed non-platform substitute rather than the full set of transportation alternatives.

destination because it depends on the proximity of idle drivers to the pickup location.[28] Matching time, defined as the time from a rider's request to driver acceptance, enters the model differently. As Table 1 shows, average matching time is an order of magnitude shorter than average pickup time, making its direct contribution to a trip's time cost negligible. Nevertheless, the distribution shown in Figure 1 has a long right tail, and requests that remain unmatched for too long are canceled before a driver accepts them. We therefore include pickup time explicitly in rider utility, while matching time affects demand only through the probability of a successful match.[29]

Riders choose among the alternatives according to a nested logit model, in which Grab and Uber form a ride-hailing nest and the outside option forms its own nest. The utility of rider k from the outside option is normalized to

$$u^0_{kxij} = \epsilon^0_{kxij}, \tag{1}$$

and the utility from platform $m \in N_{RH}$ is

$$u^m_{kxij} = \zeta^m_{xij} + \xi_{kxij} + (1-\rho_d)\epsilon^m_{kxij}, \tag{2}$$

where ζ^m_{xij} is the mean utility of platform m, common to all riders in the market. We impose the standard nested-logit distribution on the composite utility shocks, with dissimilarity parameter $1-\rho_d$ for the ride-hailing nest. The shocks ϵ^m_{kxij} are i.i.d. type-I extreme value across riders and alternatives. The term ξ_{kxij} is the demand shock to rider k for ride-hailing, also common to both platforms. The nesting parameter $\rho_d \in [0,1)$ measures how strongly riders' tastes are correlated within the ride-hailing nest. A larger ρ_d implies that, when one platform's offer worsens, more of its riders switch to the other platform rather than to the outside option.

The mean utility of platform m is

$$\zeta^m_{xij} = v_{xij} + \nu_{xij}\mathbf{1}\{m=1\} - \alpha_{xij}p^m_{xij} - \beta_{xij}w^m_{xi}. \tag{3}$$

Here v_{xij} is the value of a ride-hailing trip relative to the outside option, common to both platforms.[30] It absorbs route–hour characteristics such as in-vehicle travel time, trip purpose, and the attractiveness of alternatives that we do not observe. The term ν_{xij} is riders' relative preference for Grab over Uber, which captures platform loyalty and horizontal differentiation. The coefficients α_{xij} and β_{xij} measure riders' sensitivity

[28] Apps display an estimated pickup time before a rider submits a request. This estimate may differ from the realized pickup time for an individual trip. We use average realized pickup time within each platform–origin–hour cell as a proxy for the pickup time riders expect when choosing a platform. This measure averages over idiosyncratic trip-level variation.

[29] We abstract from the direct effect of anticipated matching failure on riders' initial platform choices.

[30] Note that it is different from ξ_{kxij} in equation (2), which also captures the demand shock to both platforms but varies by riders.

to fares and pickup times, respectively.[31]

Let

$$D_{xij} = \sum_{m' \in N_{RH}} \exp\left(\frac{\zeta^{m'}_{xij}}{1-\rho_d}\right), \tag{4}$$

so that $(1-\rho_d)\log D_{xij}$ is the inclusive value of the ride-hailing nest. Potential demand for platform m, the number of requests that riders submit to it, is

$$\widetilde{q}^{m}_{xij} = M_{xij}\frac{\exp\left(\frac{\zeta^{m}_{xij}}{1-\rho_d}\right)}{D_{xij}}\frac{D^{1-\rho_d}_{xij}}{1+D^{1-\rho_d}_{xij}}. \tag{5}$$

After receiving a request, the platform searches for a nearby driver. Riders wait passively and cannot influence the matching process. Because waiting requires continued attention, riders cancel requests that remain unmatched for too long.[32] Rather than explicitly modeling the platform's dispatch algorithm, we take a reduced-form approach and define the matching probability as the probability that matching time does not exceed the time a request is allowed to remain open. Let T^m_{xi} denote the random matching time of a request submitted on platform m from origin i at hour x, and let c^m_{xi} denote an effective cancellation threshold, common to requests within a platform–origin–hour cell, that captures rider patience and the platform's timeout.[33] A request is matched if and only if $T^m_{xi} \leq c_{xi}$. The matching probability is therefore:

$$mp^m_{xi}(s^m_{xi}, x, i) = \Pr\left(T^m_{xi} \leq c_{xi} \mid s^m_{xi}, x, i\right), \tag{6}$$

where s^m_{xi} is the number of idle drivers on platform m in origin i at hour x. Like pickup time, matching probability depends on the origin but not the destination because the platform searches among drivers near the pickup location. Its complement, $1 - mp^m_{xi}$, is the share of requests that fail to be matched.[34]

Completed trips, which are what we observe in the data, equal potential demand multiplied by the matching probability:[35]

$$q^m_{xij} = mp^m_{xi}(s^m_{xi}, x, i)\,\widetilde{q}^m_{xij}. \tag{7}$$

[31]The model describes effective substitution between platform offers. It does not require every rider to compare both apps before each trip.

[32]After submitting a request, riders must continue monitoring the app until a driver accepts. This attention is costly and discourages riders from keeping requests open for long periods.

[33]Platforms cancel requests that remain unassigned beyond a timeout, even if riders are willing to keep waiting.

[34]Matching failures refer only to requests that riders actually submit. A rider who compares several platforms or modes but submits a request to only one does not generate unmatched requests on the others. Such comparisons affect platform choice, not the matching probability conditional on a request.

[35]We abstract from cancellations after driver acceptance, so that a successfully matched request results in a completed trip.

Idle driver supply affects completed trips through the matching probability. When idle drivers near the origin are plentiful, mp_{xi}^m is close to one, and completed trips closely track potential demand. When requests outnumber nearby idle drivers, matching times lengthen and the matching probability falls. Such failures can persist even with dynamic pricing: riders and idle drivers arrive stochastically in real time, relocating drivers across regions is costly, and prices cannot adjust instantly to clear every local imbalance between requests and idle drivers. These frictions provide a rationale for allowing a positive probability of matching failure even when platforms adjust prices dynamically.

3.2 Demand Estimation

We estimate the demand model using only pre-merger data from April 2018. This choice serves two purposes. First, it recovers rider preferences under the pre-merger market structure, before the merger changes prices, pickup times, and matching. Second, it reserves the post-merger data for an out-of-sample evaluation of the model's performance in Section 3.3.

Estimation proceeds in two steps. First, we estimate matching probabilities from the distribution of matching times among completed trips. This step uses only matching-time data and does not depend on the demand parameters. Second, taking the estimated matching probabilities as given, we estimate the preference parameters from riders' choices between Grab and Uber.

Matching Probability A request is completed only if a driver accepts it before the threshold c_{xi}^m is reached. The empirical challenge is that failed requests are not observed, so the matching probability cannot be computed as the share of requests that succeed. We instead infer it from the matching times of completed trips. Because only requests matched before the threshold appear in the data, observed matching times form a right-truncated sample from the latent matching-time distribution. Given this distribution and the threshold, the matching probability is the probability mass below the threshold.

This inference rests on two assumptions. First, all requests on a platform within a cell share a common threshold. Second, matching times for all requests within a cell are drawn from the same distribution, independently of rider characteristics.[36] To keep notation light, we continue to index cells by (m, x, i).

We assume that latent matching times in each cell follow a log-normal distribution with parameters $(\mu_{xi}^m, \sigma_{xi}^m)$, which accommodates the positive support and long right tail documented in Figure 1. Let mt^n denote the observed matching time of completed trip

[36]The common threshold is a reduced-form approximation to cancellation behavior within a cell. It abstracts from within-cell differences in rider patience and search costs. The second assumption is natural because nearby requests submitted at the same time face the same local driver pool and dispatch system, while drivers have limited ability to select among individual riders before accepting a trip.

n, and let $\mathcal{N}_{xi}^m$ be the set of completed trips in the cell with $mt^n \leq c_{xi}^m$. We estimate $(\mu_{xi}^m, \sigma_{xi}^m)$ by maximizing the right-truncated log-normal log-likelihood

$$\max_{\mu_{xi}^m, \sigma_{xi}^m} \sum_{n \in \mathcal{N}_{xi}^m} \left[\log \left(\frac{1}{mt^n \sigma_{xi}^m} \phi \left(\frac{\log mt^n - \mu_{xi}^m}{\sigma_{xi}^m} \right) \right) - \log \Phi \left(\frac{\log c_{xi}^m - \mu_{xi}^m}{\sigma_{xi}^m} \right) \right], \tag{8}$$

where ϕ and Φ are the standard normal density and distribution functions. The second term rescales the density because only requests with matching times below the threshold are observed. The estimated matching probability is the fitted probability that matching time falls below the threshold:

$$\widehat{mp}_{xi}^m = \Phi \left(\frac{\log c_{xi}^m - \widehat{\mu}_{xi}^m}{\widehat{\sigma}_{xi}^m} \right). \tag{9}$$

The threshold is not directly observed. Under the common-threshold model, maximum likelihood would set the threshold to the largest observed matching time, making the estimate sensitive to extreme observations and measurement error. To limit the influence of extreme matching times, we calibrate c_{xi}^m to the 95th percentile of observed matching times within each platform–origin–hour cell during the representative pre-merger week, capped at two minutes.[37] Under this calibration, requests whose matching times approach the threshold are interpreted as close to failing. We assess sensitivity to alternative thresholds in section 6.1 in Appendix.

Figure 3 summarizes the estimates by pickup region and time window. The left column shows the pre-merger estimates that are also used in demand estimation. Matching failures concentrate in the commuting peaks and follow the direction of commuting flows. During the morning peak, when commuters leave residential areas, matching probabilities are lowest in the North and Northeast, at 85% and 89%, respectively. During the evening peak, when commuters leave the city center, the matching probability is lowest in the Central region, at 89%. By contrast, the Central region has the highest matching probability during non-peak hours, at about 97%.

The right column reports estimates for post-merger trips in May obtained using the same truncated-likelihood estimator. We hold thresholds at their pre-merger levels so that changes reflect shifts in the matching-time distribution. These post-merger estimates are used only for descriptive analysis and do not enter demand estimation. After the merger, matching probabilities rise in every region and time window, exceeding 94% throughout. The gains are largest where pre-merger failures were most common: about 11 percentage points in the North during the morning peak and more than 8 percentage

[37] A similar issue arises in reservation-wage estimation: Flinn and Heckman (1982) discuss the sensitivity of estimators based on the lowest accepted wages to measurement error. We use a high-percentile cutoff as a robust proxy for the endpoint and exclude completed trips with matching times above this cutoff from the likelihood.

points in the Central region during the evening peak, compared with less than 1 percentage point in the Central region during non-peak hours. This pattern is consistent with consolidation improving matching most in markets where separate driver pools constrained driver availability. Because post-merger matching probabilities are close to one, we approximate them by one when predicting post-merger demand in Section 3.3.

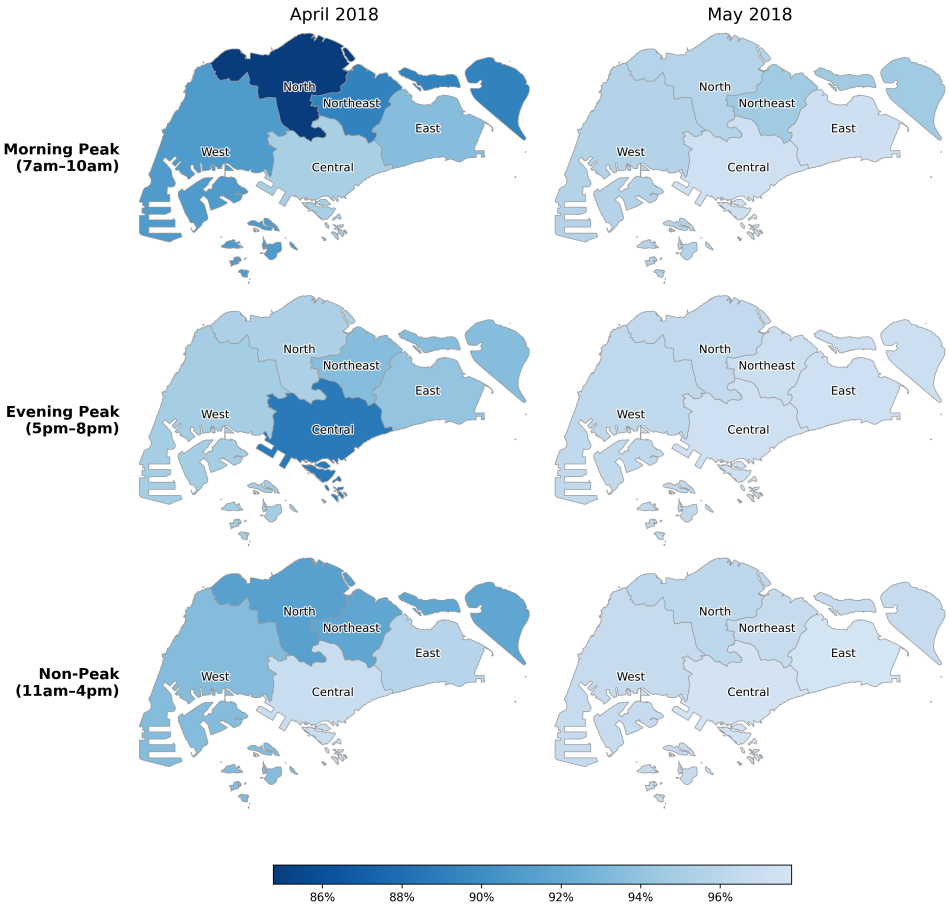


Figure 3: Estimated Matching Probability by Region Before and After the Merger

Notes: Each panel shows the estimated matching probability by pickup region. Rows correspond to the morning peak (7–10am), evening peak (5–8pm), and non-peak hours (11am–4pm). Each value averages hour-level estimates over all hours in the window and all seven days of the representative week, weighting each region–hour cell by its pre-merger market size $\sum_j M_{xij}$ in both columns. The left column combines Grab and Uber, weighting each platform's estimate by its potential demand, $\sum q^m / \sum \left(q^m / \widehat{mp}^m\right)$. The right column covers Grab, the only remaining platform, and evaluates the fitted post-merger matching-time distributions at the pre-merger Grab thresholds. Darker shading indicates a lower matching probability.

Nested Logit with Matching Frictions We next estimate the preference parameters ν_{xij}, α_{xij}, and β_{xij} from riders' choices between Grab and Uber. Endogeneity arises because prices and pickup times are correlated with unobserved demand shocks. We distinguish between shocks common to both platforms, v_{xij}, and platform-specific shocks, ν_{xij}. We assume that these shocks are additively separable and that v_{xij} is identical

across platforms within each pickup region and hour.[38] This assumption reflects the idea that unobserved route-level characteristics and demand shocks, such as weather, traffic conditions, and other idiosyncratic factors, shift demand on both platforms in the same way. Using relative market shares within each route therefore allows us to difference out v_{xij}.[39] Grab's share of completed ride-hailing trips is therefore

$$\frac{q^1_{xij}}{q^1_{xij}+q^2_{xij}} = \frac{\exp(\Lambda_{xij})}{1+\exp(\Lambda_{xij})}, \qquad \Lambda_{xij} = \frac{\nu_{xij} - \alpha_{xij}\Delta p_{xij} - \beta_{xij}\Delta w_{xi}}{1-\rho_d} + \log\frac{\widehat{mp}^1_{xi}}{\widehat{mp}^2_{xi}}, \quad (10)$$

where $\Delta p_{xij} = p^1_{xij} - p^2_{xij}$ and $\Delta w_{xi} = w^1_{xi} - w^2_{xi}$ are the differences in fares and pickup times between Grab and Uber. The log ratio of matching probabilities enters the log of Grab's completed-trip share additively. Without it, the model would attribute a platform's lower share of completed trips caused by matching failures to weaker rider preferences for that platform or to differences in fares and pickup times.

Our identification argument follows Rosaia (2025). Demand shocks common to both platforms, captured by v_{xij}, cancel out of Equation (10). Endogeneity remains a concern only if riders' relative preference for Grab is correlated with Δp and Δw. We address this concern by allowing ν to vary by pickup region and time. We allow platform preferences and sensitivities to fares and pickup times to vary by pickup region, time-of-day window, and weekday or weekend status. Identification uses variation in relative fares and pickup times across observations within these parameter groups. The identifying assumption is that, within these cells, the remaining variation in Δp and Δw is not driven by shocks to riders' relative preference for Grab. Exclusive driver affiliations provide a source of platform-specific supply variation: a shortage of idle drivers need not affect the two platforms' fares and pickup times equally.

We estimate the parameters by maximum likelihood using Equation (10). The relative-share equation identifies the preference parameters up to the scale $1-\rho_d$. We select $\rho_d = 0.45$ through two-fold cross-validation of matching-adjusted Grab shares relative to market size. [40] This value implies a moderate correlation in riders' tastes within the ride-hailing nest. We measure M_{xij} as the largest number of riders observed on route ij at hour x across the weeks of April, counting completed ride-hailing and taxi trips. We also ensure that market size exceeds inferred ride-hailing requests, so that the outside-option share remains positive. The Appendix details this adjustment. Finally, we recover

[38] Demand estimation uses the five regions defined by Singapore's Urban Redevelopment Authority; see `https://en.wikipedia.org/wiki/Regions_of_Singapore`. Finer spatial units, such as planning areas, produce many platform–route–hour cells with zero trips, making relative-share estimates unstable. Aggregating to five regions trades off spatial precision against reliable variation in prices, pickup times, and platform market shares.

[39] The rider-specific demand shock common to both platforms, ξ_{kxij}, also cancels.

[40] We fit the model on one of the first two weeks of April, predict trip shares in the other week, and choose the value that minimizes the average out-of-sample prediction error. In section 6.2 in Appendix, we discuss the choice of ρ_d and report robustness to alternative values.

the common value v_{xij} by inverting the nested-logit share equations to match potential demand, measured as $\widetilde{q}^m_{xij} = q^m_{xij}/\widehat{mp}^m_{xi}$.

Tables 2 and 3 report Grab's own-price and own-pickup-time elasticities on weekdays.[41] These elasticities refer to potential demand, with matching probabilities held fixed. The estimates reveal substantial heterogeneity across regions and times of day. Price elasticities range from -1.424 in the West between 1am and 6am to -0.056 in the East between 5pm and 8pm. Riders are generally more sensitive to prices late at night: in four of the five regions, price elasticities are largest in absolute value between 9pm and 6am; in the Central region, for example, the price elasticity is -1.362 between 9pm and 12am.

Sensitivity to pickup time follows a different pattern. It is highest during the morning peak in the residential regions, where pickup-time elasticities range from -0.493 in the North to -0.897 in the West. In the West, East, and Northeast, the morning-peak pickup-time elasticity exceeds the price elasticity in absolute value; in the West, for example, the corresponding elasticities are -0.897 and -0.809. A natural explanation is that fixed work start times make waiting for a driver more costly for morning commuters. In the midday and late-night windows, by contrast, pickup-time elasticities are small, with absolute values below 0.1 in 12 of the 15 region–window cells.

3.3 Model Fit Assessment

We assess the model's ability to predict post-merger completed trips out of sample, conditional on observed post-merger fares and pickup times. Using parameters estimated only from pre-merger data, we compare predicted trips with observed trips in the post-merger week. We compare two demand models. Our proposed model, described in Subsection 3.1, incorporates pre-merger matching probabilities in estimation. The benchmark model uses the same nested logit specification but omits the matching term as in the existing literature (Castillo 2025; Rosaia 2025).[42]

For each route–hour cell, we hold fixed the pre-merger values of v_{xij}, the common value of ride-hailing relative to the outside option, and ν_{xij}, riders' relative preference for Grab over Uber.[43] We replace fares and pickup times with their observed post-merger values.[44] Because Uber has exited, the ride-hailing nest contains only Grab, and the predictive model reduces to binary logit. In our main specification, we set the post-merger matching

[41] The corresponding tables for Uber, and the elasticity estimates of Grab on weekends, can be found in section 5.1.3 in Appendix.

[42] We estimate each specification separately using pre-merger data and select its nesting parameter through the same cross-validation procedure. The selected values are $\rho_d = 0.45$ for the proposed model and $\rho_d = 0.21$ for the benchmark model.

[43] Changes in common demand conditions or relative platform preferences after the merger may contribute to prediction errors in both models.

[44] We evaluate both models on the same route–hour cells with observed post-merger trips. All post-merger fares and pickup times are available.

Table 2: Own-Price Elasticities of Grab Demand: Weekdays

	7am-10am	11am-4pm	5pm-8pm	9pm-12am	1am-6am
Central	-0.552	-0.320	-0.531	-1.362	-0.658
	(0.035)	(0.044)	(0.049)	(0.066)	(0.070)
East	-0.506	-0.063	-0.056	-0.196	-0.424
	(0.063)	(0.050)	(0.055)	(0.063)	(0.081)
West	-0.809	-0.308	-0.318	-0.692	-1.424
	(0.111)	(0.124)	(0.087)	(0.130)	(0.172)
Northeast	-0.327	-0.401	-0.588	-1.069	-0.691
	(0.097)	(0.091)	(0.130)	(0.175)	(0.165)
North	-0.841	-0.318	-0.323	-0.941	-0.824
	(0.139)	(0.086)	(0.080)	(0.132)	(0.153)

Notes: The table reports own-price elasticities of potential demand for Grab on weekdays, holding matching probabilities fixed. Rows are pickup regions; columns are time-of-day windows of the request hour. Each elasticity is evaluated at the mean Grab fare $\bar{p}^1$ in the cell as $E^p = -\alpha\,\bar{p}^1\left[\frac{1}{1-\rho_d} - \frac{\rho_d}{1-\rho_d}s_{1|RH} - s_1\right]$, where $s_{1|RH}$ is Grab's share of ride-hailing requests, s_{RH} is the ride-hailing share of the market, and $s_1 = s_{1|RH}\,s_{RH}$ is Grab's share of the market. Shares are computed from matching-adjusted requests, $q^m/\widehat{mp}^m$. The price and pickup-time coefficients are constrained to be positive in estimation. Standard errors, in parentheses, are computed by the delta method from robust standard errors of the coefficients, holding fares and shares fixed. The nesting parameter is $\rho_d = 0.45$.

Table 3: Own-Pickup-Time Elasticities of Grab Demand: Weekdays

	7am-10am	11am-4pm	5pm-8pm	9pm-12am	1am-6am
Central	-0.242	-0.225	-0.317	-0.372	-0.030
	(0.036)	(0.028)	(0.034)	(0.081)	(0.099)
East	-0.551	-0.019	-0.096	-0.116	-0.027
	(0.078)	(0.040)	(0.036)	(0.044)	(0.068)
West	-0.897	-0.049	-0.343	-0.018	-0.026
	(0.174)	(0.110)	(0.081)	(0.147)	(0.238)
Northeast	-0.504	-0.088	-0.092	-0.017	-0.025
	(0.124)	(0.069)	(0.075)	(0.098)	(0.150)
North	-0.493	-0.022	-0.032	-0.017	-0.023
	(0.184)	(0.133)	(0.125)	(0.092)	(0.177)

Notes: The table reports own-pickup-time elasticities of potential demand for Grab on weekdays, holding matching probabilities fixed. Rows are pickup regions; columns are time-of-day windows of the request hour. Each elasticity is evaluated at the mean Grab pickup time $\bar{w}^1$ in the cell as $E^w = -\beta\,\bar{w}^1\left[\frac{1}{1-\rho_d} - \frac{\rho_d}{1-\rho_d}s_{1|RH} - s_1\right]$, with shares defined as in Table 2. The price and pickup-time coefficients are constrained to be positive in estimation. Standard errors, in parentheses, are computed by the delta method from robust standard errors of the coefficients, holding pickup times and shares fixed. The nesting parameter is $\rho_d = 0.45$.

probability to one, approximating the high estimated probabilities reported in Figure 3. Applying the same full-matching assumption to both models isolates the contribution of accounting for pre-merger matching failures in demand estimation. Any improvement in prediction therefore comes from the pre-merger demand estimates, without using post-merger matching-time distributions as an additional input. We assess sensitivity to the remaining matching failures below.

Figure 4 reports the results by time window. The benchmark model underpredicts post-merger trips in every window. It predicts declines of 10.37% during the morning peak and 6.79% during the evening peak, whereas observed trips increase by 2.83% and 4.07%, respectively. During non-peak hours, it predicts a decline of 10.95%, far larger than the observed decline of 1.18%. The proposed model predicts a 1.19% increase during the morning peak, a 4.44% increase during the evening peak, and a 2.50% decline during non-peak hours. It thus captures the direction of the observed changes in all three windows and predicts their magnitudes more accurately than the benchmark model.

We assess sensitivity to the full-matching approximation by using estimated post-merger matching probabilities while holding demand parameters, fares, and pickup times unchanged. The proposed model then predicts a 2.40% decline during the morning peak, a 1.37% increase during the evening peak, and a 5.10% decline during non-peak hours.

Although it no longer predicts the observed morning expansion, it remains closer to the observed changes than the benchmark model in all three time windows and all 15 region–window cells. The improvement in predictive accuracy therefore persists when we allow for remaining post-merger matching failures. We report these results in figure 14 and 15 in Appendix.

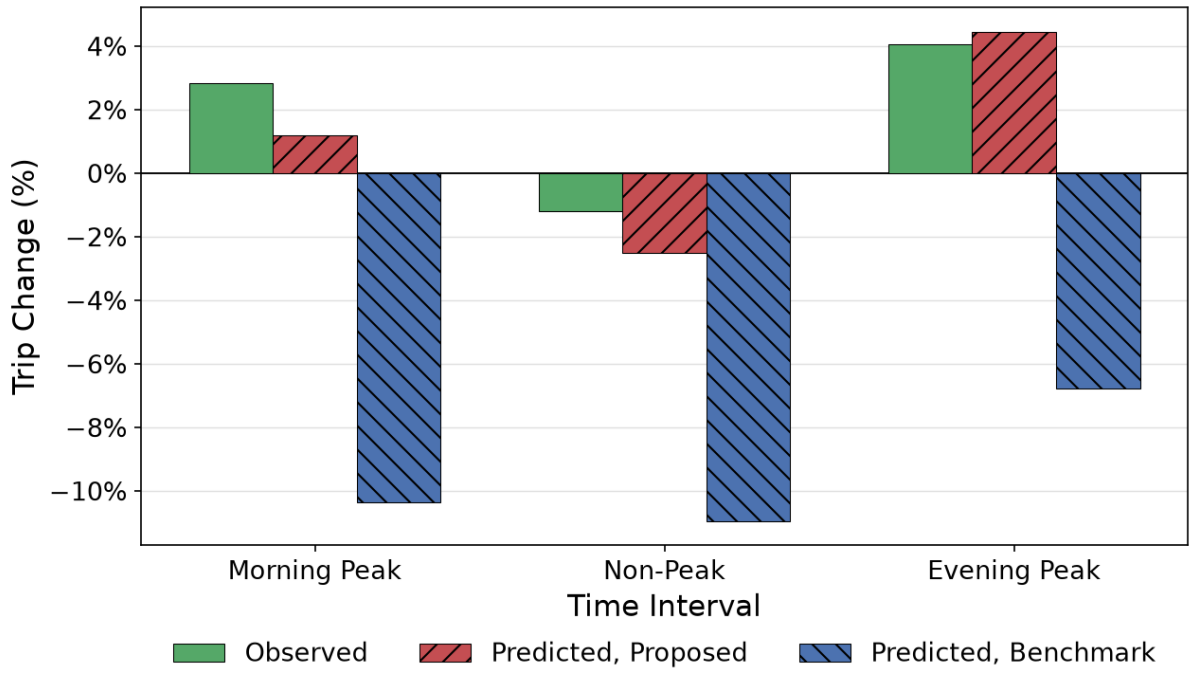


Figure 4: Observed and Predicted Changes in Completed Trips by Time Window

Notes: Each group compares the observed percentage change in completed ride-hailing trips from the representative pre-merger week (April 4–10, 2018; Grab and Uber combined) to the post-merger week (May 9–15, 2018; Grab only) with the changes predicted by the proposed model, estimated with pre-merger matching probabilities, and by the benchmark model, estimated without them. Both predictions use observed post-merger fares and pickup times and set the post-merger matching probability to one. Changes are computed over route–hour cells with completed trips in the post-merger week, relative to pre-merger trips in the same cells, and include all seven days of each week. Morning peak: 7–10am; non-peak: 11am–4pm; evening peak: 5–8pm.

Figure 5 breaks down the predictions by pickup region. During the morning peak, observed trips increase by approximately 3% to 7% in the four residential regions. The proposed model predicts an increase in all four, whereas the benchmark model predicts declines of 5% to 9%. During the evening peak, the largest observed increase is in the Central region, where trips rise by 8.11%. The proposed model predicts a 7.42% increase, whereas the benchmark model predicts a 5.47% decline. During non-peak hours, observed changes are small in every region. The benchmark model predicts declines of 7% to 12%, while the proposed model's predictions are much closer to zero, although they do not always match the direction of the observed changes. Across all 15 region–window

cells, the proposed model's predictions are closer to the observed changes than those of the benchmark model.[45] The improvement is particularly pronounced in markets with low pre-merger matching probabilities, as documented in Section 3.2: residential regions during the morning peak and the Central region during the evening peak.

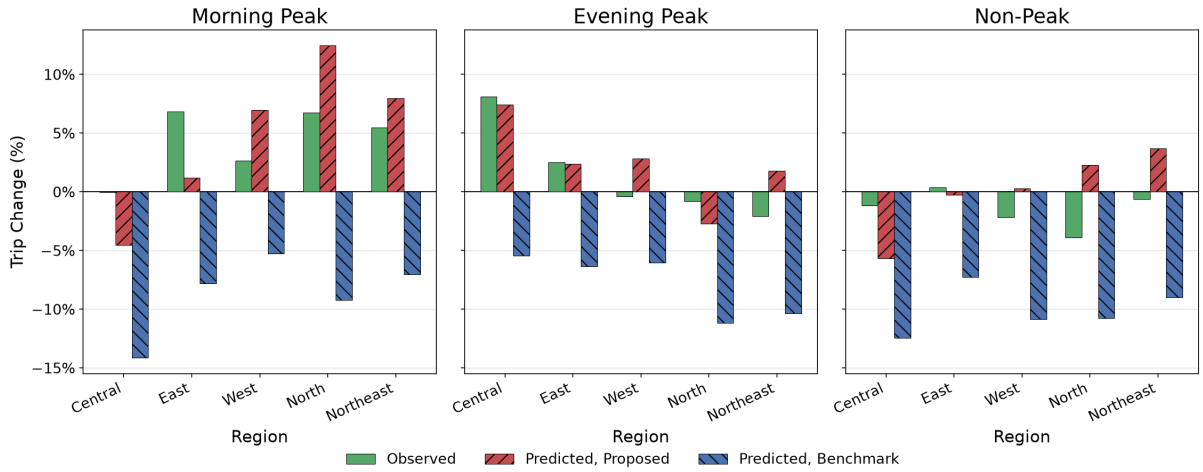


Figure 5: Observed and Predicted Changes in Completed Trips by Pickup Region

Notes: The panels report the morning peak (7–10am), evening peak (5–8pm), and non-peak hours (11am–4pm). Within each panel, each group compares the observed percentage change in completed ride-hailing trips originating in a pickup region with the changes predicted by the proposed and benchmark models. Samples, baselines, and predictions are defined as in Figure 4.

Accounting for pre-merger matching failures substantially improves predictions of post-merger completed trips: the proposed model is closer to the observed changes in all 15 region–time-window cells. Both models use the same observed post-merger fares and pickup times, so the improvement reflects differences in the demand estimates obtained with and without accounting for pre-merger matching failures. We next examine how these differences affect the assessment of the merger's welfare effects.

4 Consumer Surplus

We use the estimated demand models to evaluate changes in consumer surplus between the representative pre- and post-merger weeks. Both models are evaluated at the same observed trip quantities, fares, and pickup times in each period. Using these common market outcomes allows us to assess how accounting for pre-merger matching failures changes the valuation of realized trips. Within each model, we retain the pre-merger demand estimates and fitted platform utility components, so that both periods are evaluated using the same estimated preferences.

[45] We acknowledge that the proposed model overstates the morning expansion in the North and understates it in the East.

The proposed model distinguishes consumers' demand for rides from the number of requests that result in completed trips. Before the merger, we recover implied request quantities by dividing each platform's completed trips by its matching probability. We integrate the estimated inverse-demand system and subtract fare payments at these request quantities to obtain potential consumer surplus jointly across the two platforms. We then weight this surplus by the ratio of completed trips to requests, assigning the average potential surplus per request to each completed trip. The benchmark model assumes that all requests are successfully matched.

Both models set the post-merger matching probability to one. With Grab as the sole remaining platform, consumer surplus in route–hour market k under model s is

$$CS^s_{k,\text{post}} = \int_0^{q_{k,\text{post}}} P^s_k(z; w_{k,\text{post}})\, dz - p_{k,\text{post}} q_{k,\text{post}}, \tag{11}$$

where $q_{k,\text{post}}$, $p_{k,\text{post}}$, and $w_{k,\text{post}}$ denote observed completed trips, fares, and pickup times, respectively. The inverse-demand function P^s_k is constructed from the pre-merger demand estimates and evaluated at the observed post-merger pickup time. Observed trip quantities therefore determine the range over which willingness to pay is integrated, and observed fare payments are deducted to obtain consumer surplus.

For each group reported in Table 4, we sum consumer-surplus levels across the relevant route–hour markets before calculating the percentage change relative to the same model's pre-merger total. Each reported change thus compares aggregate surplus across periods within the estimated demand model.

Accounting for pre-merger matching failures substantially reduces the estimated decline in aggregate consumer surplus. Using observed post-merger quantities, the proposed model implies a decline of 0.60%, compared with 10.71% under the benchmark model. The aggregate point estimate under the proposed model is therefore close to zero. During peak hours, the point estimates have opposite signs: the proposed model implies gains of 0.75% during the morning peak and 0.19% during the evening peak, whereas the benchmark model implies losses of 11.37% and 8.35%, respectively. The proposed model also implies gains of 4.42% from 9 pm to midnight and 6.66% from 1 am to 6 am, whereas the benchmark model implies losses of 9.50% and 16.97%, respectively. Both models imply aggregate losses during midday non-peak hours, from 11 am to 4 pm.

Figure 6 also reveals substantial spatial heterogeneity. Under the proposed model, consumer surplus increases in the residential regions of the East, West, North, and Northeast during the morning peak and in the Central region during the evening peak. Using observed post-merger quantities, consumer surplus in the Central evening market increases by 6.93% under the proposed model but declines by 5.36% under the benchmark model. These gains are specific to particular region–time combinations: consumer surplus in the Central region still declines when aggregated over the full week. Thus, accounting

Table 4: Percentage Changes in Consumer Surplus at Observed Market Outcomes

	Demand Specification	
	Proposed	Benchmark
Panel A: Overall		
Overall	-0.60	-10.71
Panel B: By Pickup Region		
Central	-2.77	-14.49
East	-0.51	-8.93
West	-1.02	-11.39
Northeast	5.93	-7.03
North	-1.05	-13.83
Panel C: By Time Window		
7am-10am	0.75	-11.37
11am-4pm	-3.39	-11.56
5pm-8pm	0.19	-8.35
9pm-12am	4.42	-9.50
1am-6am	6.66	-16.97

Notes: The comparison reflects differences in the demand estimates obtained with and without accounting for pre-merger matching failures, including the separately selected nesting parameters. Entries report percentage changes in consumer surplus between the representative pre-merger week (April 4–10, 2018) and post-merger week (May 9–15, 2018), using observed trip quantities, fares, and pickup times. Post-merger matching probabilities are set to one. The proposed model uses demand estimates obtained with estimated pre-merger matching probabilities; the benchmark model uses demand estimates obtained with pre-merger matching probabilities set to one. Each percentage change is measured relative to the corresponding model's pre-merger consumer-surplus baseline.

for pre-merger matching failures can reverse the sign of the estimated consumer-surplus change in particular region–time markets, even though the aggregate estimate remains slightly negative.

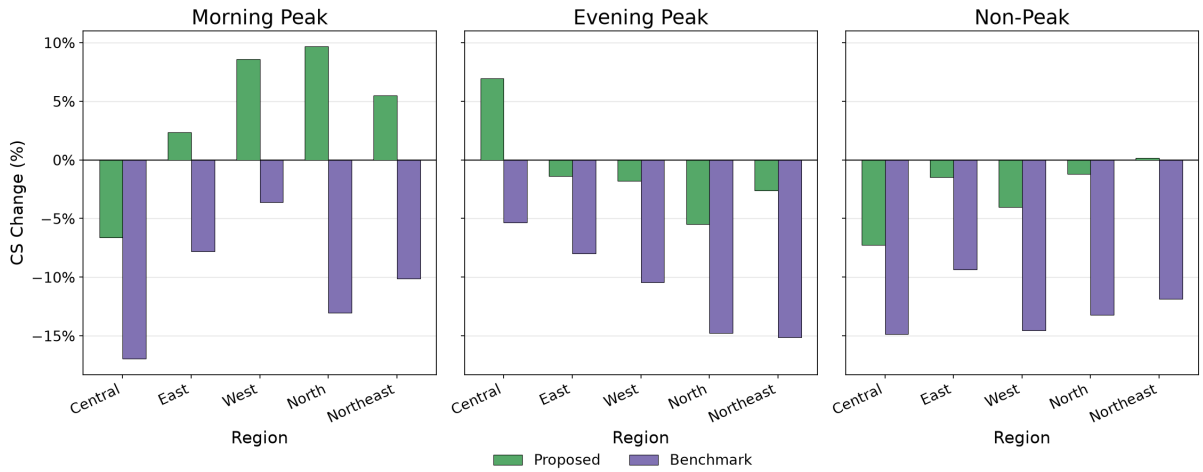


Figure 6: Percentage Changes in Consumer Surplus by Pickup Region and Time Window

Notes: Panels report estimated percentage changes in consumer surplus during the morning peak (7–10 am), evening peak (5–8 pm), and midday non-peak hours (11 am–4 pm). Green bars correspond to the proposed model and purple bars to the benchmark model. Each change is measured relative to the same model's pre-merger consumer surplus in the corresponding region and time window. Market inputs and matching-probability assumptions are defined in Table 4. Bars show point estimates.

These estimates pertain to a short-run post-merger equilibrium in which costly driver entry limits supply adjustment and regulatory constraints restrict pricing. Section 5 examines the long-run equilibrium after driver participation has fully adjusted and these pricing constraints are removed.

5 Counterfactual Simulations

The preceding sections evaluate demand predictions and consumer surplus at observed post-merger outcomes. Two features of these outcomes motivate a long-run full equilibrium simulation. First, our data cover the period shortly after Uber's exit, when driver entry and working hours had limited time to adjust. Assessing the merger's long-run effects therefore requires allowing driver supply to respond fully. Second, regulatory restrictions limited Grab's pricing discretion during the sample period, so observed fares do not represent unconstrained post-merger pricing.

We therefore use pre-merger estimates to simulate a long-run equilibrium with free driver entry and fares charged without these regulatory restrictions. Combining our demand model with the driver-supply and platform-optimization framework of Rosaia (2025), we develop a full equilibrium model which endogenously determines fares, driver

participation, pickup times, matching probabilities, and completed trips. Because the exercise uses only pre-merger estimates, it can also inform merger assessment before post-merger outcomes are observed. We further examine whether removing driver exclusivity can improve matching while preserving platform competition. The details of the model and simulation are in Appendix 7.

5.1 Long-run Post-merger Equilibrium

We simulate the market after Uber's exit, with Grab as the sole platform, under the proposed and benchmark models. Both simulations allow the driver population to adjust through free entry and retain the platform objective estimated under the corresponding specification. The benchmark model assumes successful matching throughout. The proposed model accounts for pre-merger matching failures and allows post-merger matching probability to respond to idle-driver supply:

$$mp_{xi}(s_{xi}) = \Phi\left(\frac{\log c_{xi} - \mu_{xi}(s_{xi})}{\sigma_{xi}(s_{xi})}\right), \tag{12}$$

where s_{xi} is the simulated number of idle drivers and the parameters of the matching-time distribution, $\mu_{xi}(s_{xi})$ and $\sigma_{xi}(s_{xi})$, depend on idle supply through a relationship estimated from pre-merger Grab trips. The cancellation threshold c_{xi} is calibrated from pre-merger matching times and held fixed in the simulation. We maintain the estimated relationship between idle-driver supply and the matching-time distribution after consolidation, evaluating it at the merged platform's equilibrium idle supply. Table 15 in the Appendix reports the matching-function estimates. Under this specification, a change in driver supply shifts the matching-time distribution and therefore the share of requests that are matched, so matching probability is determined jointly with fares, trips, pickup times, and driver supply.

Figure 7 compares simulated changes in completed trips under the two models. Under the proposed model, morning-peak trips increase by approximately 33% in the West, 22% in the North, and 9% in the Northeast; evening-peak trips increase by approximately 5% in the Central region. Trips decline in every region during non-peak hours. The benchmark generally predicts larger peak-period expansions, including approximately 53% in the West during the morning peak.

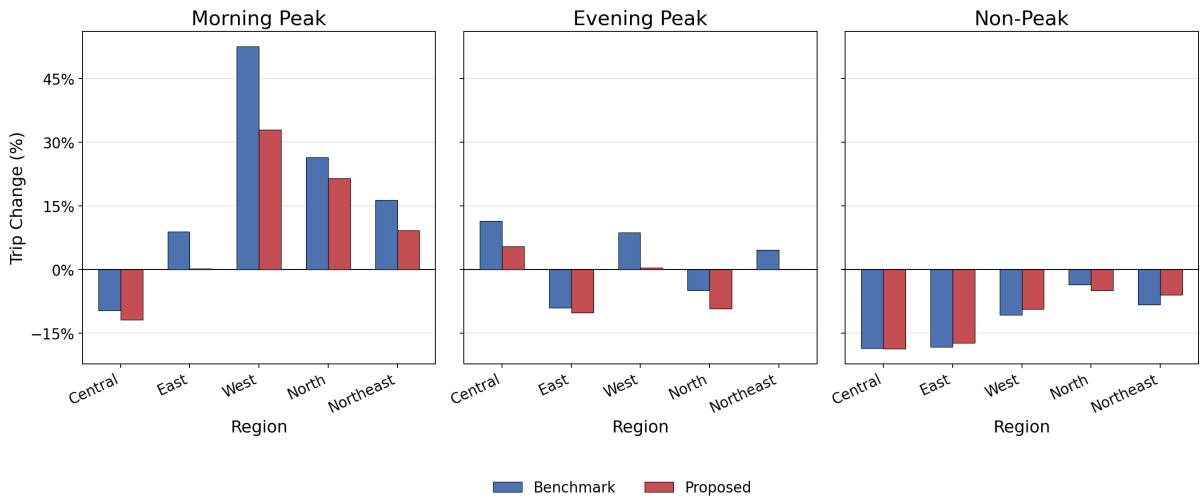


Figure 7: Simulated Long-run Post-Merger Completed Trips

Notes: Each bar reports the percentage change in completed trips originating in a pickup region, relative to observed pre-merger trips. Both simulations allow free driver entry. Blue bars represent the benchmark model, which assumes successful matching throughout. Red bars represent the proposed model, which incorporates pre-merger matching failures and allows post-merger matching probability to respond to idle-driver supply through Equation (12). Morning peak: 7–10 am; evening peak: 5–8 pm; non-peak: 11 am–4 pm. The same legend applies to Figure 8.

Table 5 shows that the proposed model predicts an average fare of 2.23 SGD per kilometer, compared with 1.54 before the merger and 1.62 in the observed post-merger market. The substantially higher simulated fare is consistent with regulatory restrictions limiting Grab's pricing discretion during the sample period. The benchmark predicts a smaller increase, to 1.75 SGD per kilometer, despite nearly identical average pickup times: 3.78 minutes under the benchmark and 3.82 under the proposed model. Accounting for matching failures therefore substantially changes the predicted pricing response, even when pickup-time predictions are similar. The driver population also adjusts substantially, increasing from 45,489 before the merger to approximately 67,600 under the proposed model. The predicted reductions in pickup times therefore occur alongside driver entry as well as the consolidation of the two driver pools.

Table 5: Simulated Long-run Post-Merger Equilibrium: Unit Price and Pickup Time

	Unit Price (SGD/km)			Pickup Time (minutes)		
	Data	Simulation		Data	Simulation	
	Pre merger	Benchmark	Proposed	Pre merger	Benchmark	Proposed
Panel A: Average						
	1.54	1.75	2.23	6.34	3.78	3.82
Panel B: By Pickup Region						
Central	1.70	1.53	1.87	6.33	4.27	4.26
East	1.35	3.19	4.46	6.24	3.41	3.50
West	1.39	1.31	1.58	6.32	3.25	3.31
Northeast	1.40	1.65	2.01	6.64	3.45	3.53
North	1.44	1.37	1.71	6.15	3.17	3.22
Panel C: By Time Window						
7am-10am	1.51	1.60	2.00	6.68	3.66	3.81
11am-4pm	1.51	2.09	2.73	5.77	3.79	3.79
5pm-8pm	1.64	1.84	2.40	6.74	3.55	3.69
9pm-12am	1.53	1.48	1.77	6.45	3.98	3.96
1am-6am	1.49	1.17	1.39	6.41	4.66	4.06

Notes: "Data, pre-merger" reports observed values during April 4–10, 2018, combining Grab and Uber. The benchmark model assumes successful matching throughout. The proposed model incorporates pre-merger matching failures and determines post-merger matching probability from idle-driver supply through Equation (12). Both simulations allow free driver entry. Unit price is total fare divided by trip distance. Observed values are weighted by observed completed trips; simulated values are weighted by simulated completed trips under the corresponding specification.

Aggregate consumer surplus declines by 11.97% under the proposed model and 7.43% under the benchmark, each relative to its own pre-merger baseline. Figure 8 shows substantial heterogeneity. Under the proposed model, morning-peak consumer surplus increases by approximately 54% in the West, 24% in the North, and 8% in the Northeast, while evening-peak surplus increases by approximately 4% in the Central region. Consumer surplus declines in every region during non-peak hours. The benchmark predicts evening-peak gains in the West and Northeast, whereas the proposed model predicts losses. Accounting for matching failures therefore increases the estimated aggregate consumer-surplus loss when fares and driver entry adjust endogenously, whereas it reduces the estimated loss at observed post-merger outcomes in Section 4. This contrast highlights the importance of evaluating matching improvements together with the accompanying pricing and driver-supply responses.

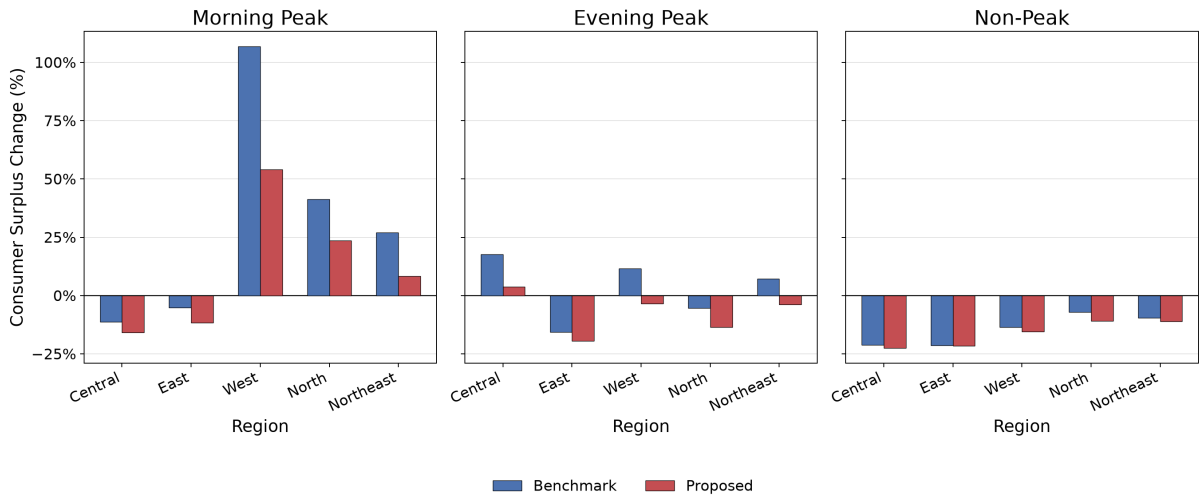


Figure 8: Simulated Long-run Post-Merger Consumer Surplus

Notes: Each bar reports the percentage change in consumer surplus for a pickup region and time window, relative to the same model's pre-merger consumer-surplus baseline. Specifications and time windows are defined in Figure 7.

5.2 Removing Driver Exclusivity under Platform Competition

Exclusive driver affiliations segment supply across platforms: a rider on one platform may be unable to secure a match even when an idle driver is available on the other. This segmentation can prevent the market from using its existing driver supply efficiently and provides one channel through which consolidation can improve matching. It raises a natural policy question: can removing exclusivity deliver matching gains while preserving competition between platforms?[46] We therefore simulate a counterfactual market in which Grab and Uber continue to compete but drivers can work for either platform.

In this counterfactual, drivers choose a platform when becoming active and can reconsider their choice after completing a trip or a relocation decision. Their choices depend on the expected payoffs from working on each platform. A driver is available to one platform in a given state, but affiliation can change at subsequent decision points without a switching cost.[47] Unlike in the post-merger simulation, we hold the total driver population at its pre-merger level. Drivers can nevertheless adjust their participation, working time, location, and platform choice. The exercise therefore examines how removing exclusivity changes the use of an existing driver pool.

Each platform retains its own demand and matching technology, and its matching probability responds to its own idle-driver supply through Equation (12). Platforms choose rider fares and driver payments to maximize their objectives, taking driver re-

[46]The question is directly relevant to merger policy. In its final decision on the Grab–Uber merger, the competition authority directed Grab to remove all exclusivity obligations on drivers and taxi fleets in Singapore. See CCCS's announcement of its infringement decision on September 24, 2018.

[47]Appendix 7 describes the timing of driver decisions and the equilibrium conditions.

sponses into account. Fares, compensation, participation, allocation, and matching therefore adjust jointly. We use the proposed demand model and compare the resulting equilibrium with the observed pre-merger market under exclusive driver affiliations.

Removing exclusivity strengthens competition for drivers because drivers can redirect their availability toward the platform offering a higher expected payoff. Each platform therefore has stronger incentives to offer attractive earnings opportunities. In the simulated equilibrium, drivers' expected net payoffs increase, and these changes are accompanied by higher fares. The supply response operates along two margins. First, drivers in the existing pool can adjust when they become active and how long they remain at work. Second, active drivers can switch platforms in response to differences in expected payoffs, directing supply toward platforms with greater earning opportunities. These adjustments can alleviate platform-specific shortages of idle drivers, improve matching probabilities, and shorten pickup times. Higher fares can therefore coexist with better service even when competition between platforms is preserved.

Figure 9 reports the changes in completed trips. During the morning peak, trips expand most strongly in the West, North, and Northeast, while declining in the Central region. During the evening peak, the Central, West, and Northeast regions experience increases, whereas the East and North experience declines. During non-peak hours, trips fall in the Central, East, and West regions, rise slightly in the North, and remain roughly unchanged in the Northeast. Aggregating across regions, completed trips increase by 11.3% during the morning peak and 5.42% during the evening peak, but decline by 5.76% between 11AM and 4PM. The non-exclusivity equilibrium generates geographically concentrated expansions during peak periods alongside contractions in most regions during non-peak hours.

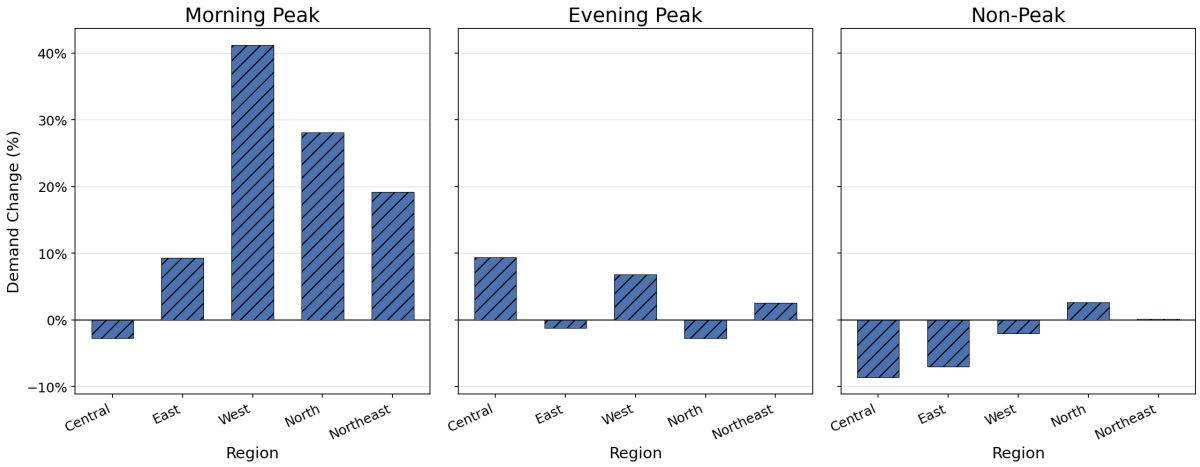


Figure 9: Simulated Completed Trips without Driver Exclusivity

Notes: Each bar reports the percentage change in completed trips, combining Grab and Uber, relative to the observed pre-merger market under exclusive driver affiliations. The simulation uses the proposed model with endogenous matching probabilities and holds the total driver population fixed. Drivers choose a platform when becoming active and can reconsider their choice after completing a trip or a relocation decision.

Table 6 reports observed pre-merger and simulated unit prices, pickup times, and matching probabilities, combining Grab and Uber. Average unit price increases from 1.54 to 2.08 SGD per kilometer, while average pickup time falls from 6.34 to 4.07 minutes. Average matching probability rises from 91.48% to 99.60%, corresponding to a decline in the average matching-failure probability from 8.52% to 0.40%. Matching improves in every reported pickup region and time window, with simulated averages exceeding 99% except between 1 am and 6 am. The improvements are particularly large in markets with substantial pre-merger matching failures, including peak and overnight periods. These outcomes are consistent with more flexible driver participation and allocation across platforms. The improvement in matching occurs without an expansion of the total driver population, highlighting the potential gains from changing how existing driver capacity is used.

Figure 10 shows substantial heterogeneity in consumer-surplus changes. During the morning peak, riders in the West, North, and Northeast gain, with the largest increase in the West. During the evening peak, gains are limited to the Central and West regions, while consumer surplus declines in every region during midday non-peak hours. These outcomes reflect the opposing consequences of non-exclusivity for riders: more flexible driver allocation improves the probability of securing a ride and shortens pickup times, while the accompanying fare increases raise the cost of service. Improved matching therefore does not translate into consumer-surplus gains in every market. The results show how stronger competition for drivers can coexist with losses for riders, while also demonstrating that some matching gains can be achieved without consolidating the platforms.

Table 6: Counterfactual of Non-exclusive Drivers with Platform Competition

	Unit Price (SGD/km)		Pickup Time (minutes)		Matching Probability (%)	
	Data Exclusive Drivers	Simulation No Exclusivity	Data Exclusive Drivers	Simulation No Exclusivity	Data Exclusive Drivers	Simulation No Exclusivity
Panel A: Average						
All	1.54	2.08	6.34	4.07	91.48	99.60
Panel B: By Pickup Region						
Central	1.70	1.85	6.33	4.65	90.72	99.27
East	1.35	3.64	6.24	3.68	93.09	99.88
West	1.39	1.56	6.32	3.42	93.02	99.98
Northeast	1.40	1.90	6.64	3.61	91.15	99.96
North	1.44	1.68	6.15	3.29	91.19	99.99
Panel C: By Time Window						
7am-10am	1.51	1.92	6.68	4.04	92.55	99.75
11am-4pm	1.51	2.46	5.77	3.99	95.30	99.85
5pm-8pm	1.64	2.31	6.74	4.33	91.38	99.61
9pm-12am	1.53	1.66	6.45	4.02	88.35	99.44
1am-6am	1.49	1.29	6.41	3.80	81.76	98.64

Notes: Each pair compares observed outcomes under exclusivity with simulated outcomes without it, while maintaining competition between Grab and Uber. Unit prices and pickup times are weighted by completed trips in the corresponding scenario. Simulated unit prices divide route-level fares by the average April trip distance on each route and hour. Within each pickup-region–hour cell, platform-specific matching probabilities are combined using potential-demand weights; these cell-level probabilities are then averaged using market-size weights.

The welfare consequences depend on the joint responses of fares, driver compensation, and the allocation of the existing driver pool.

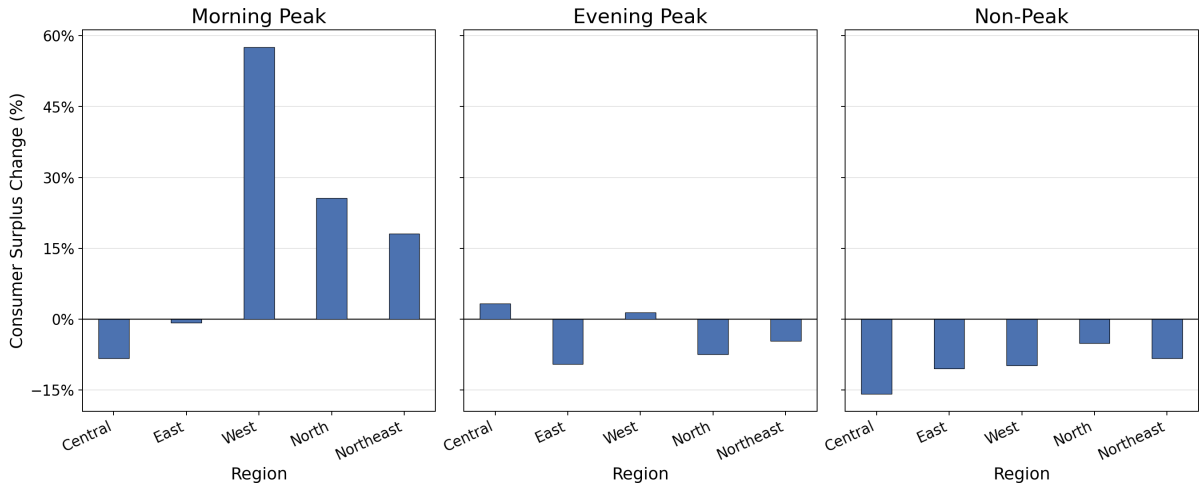


Figure 10: Simulated Consumer Surplus without Driver Exclusivity

Notes: Each bar reports the percentage change in rider consumer surplus for a pickup region and time window, relative to the proposed model's pre-merger consumer-surplus baseline under exclusive driver affiliations. The simulation maintains competition between Grab and Uber and holds the total driver population fixed.

6 Concluding Remarks

This paper identifies matching uncertainty as a central determinant of platform demand and a key channel through which mergers affect consumer welfare. Consumers respond not only to how long they must wait for service, but also to whether they can secure a match with a service provider. Fragmentation across competing platforms heightens matching uncertainty, which consolidation substantially reduces by pooling users on both sides of the market. In the empirical context of a ride-hailing platform merger, we incorporate matching probability into demand estimation, and show how the merger stimulates rider demand through improved matching efficiency and service availability. Granular trip-level data enable us to measure matching time and estimate matching probability, and validate the model against held-out post-merger outcomes. Our findings reveal that improvement in matching probability is crucial to explain demand changes after the merger, and capture gains in consumer surplus. Our framework can also be generalized beyond ride-hailing markets. For example, food-delivery platforms must find a courier for every order, and on-demand staffing platforms must fill every posted shift with an available worker, in which transactions depend on successful matching between two sides of the market.

Our findings generate implications for both the literature and policymaking. First, demand models in the ride-hailing literature that assumes no matching uncertainty can yield biased demand estimates and distort welfare assessments. When estimated from data of completed transactions, these models may conflate consumers' willingness to request with the likelihood that their requests are fulfilled, attributing unfulfilled demand to weak preferences for the service. This misspecification matters when evaluating mergers that alter the probability of successful matching. In our setting, omitting this margin underpredicts post-merger demand and understates the consumer-surplus gains from consolidation, potentially reversing the sign of the estimated welfare effect in markets with severe matching frictions.

Second, our framework helps regulators evaluate the long-run welfare implications of platform mergers and other competition policies, a central consideration in policy design. Since all structural parameters are estimated from pre-merger data, regulators can apply it during merger review to assess changes in matching, prices, and welfare in the long-run post-merger equilibrium, which informs their decision on approving or blocking the merger. Beyond merger review, we apply the framework to evaluate driver exclusivity arrangements under platform competition, which is also common in anti-trust policy practice. We find that removing exclusivity and allowing drivers to work on either platform would improve matching but also raise prices, yielding mixed effects on rider welfare. These findings highlight the importance of assessing how competition policies affect matching and pricing on both sides of the market, since intensifying competition

on one side need not improve welfare on the other.

Our paper is subject to several limitations that suggest directions for future research. First, we mainly study matching uncertainty from the rider's side, but assume it way on the driver side. However, drivers also face supply side uncertainty about how long they will remain idle before receiving a request, and this uncertainty likely shapes their decisions of market entry and idling. Drivers being idle for a long time will exit the labor market, which is also a kind of matching failure. A merger that pools demand may reduce drivers' idle time until being matched with a rider. A model in which matching probabilities on both sides are determined jointly would capture this feedback and allow a fuller account of the merger's effects on driver welfare. Second, on the supply side, we assume all drivers are ex ante homogeneous and make market-entry decisions without frictions. However, drivers are not homogeneous in skills. More experienced drivers are likely to strategically position themselves in profitable regions for the next match, reducing spatial imbalance and hence reducing the impact on surge pricing post-merger. Licensing requirements may limit market entry reducing supply expansion. While incorporating driver heterogeneity and regulatory constraints could enrich the supply-side welfare analysis, the limitations do not undermine the key insights in this paper. We leave these extensions to future research with richer driver-side data.

References

Agarwal, Nikhil, Itai Ashlagi, Eduardo Azevedo, Clayton R Featherstone, and Ömer Karaduman. 2019. "Market failure in kidney exchange." *American Economic Review* 109 (11):4026–4070.

Allon, Gad, Daniel Chen, and Ken Moon. 2025. "Measuring strategic behavior by gig economy workers: Multihoming and repositioning." *Cornell SC Johnson College of Business Research Paper* .

Armstrong, Mark. 2006. "Competition in two-sided markets." *The RAND journal of economics* 37 (3):668–691.

Armstrong, Mark and Julian Wright. 2007. "Two-sided markets, competitive bottlenecks and exclusive contracts." *Economic Theory* 32 (2):353–380.

Belleflamme, Paul and Martin Peitz. 2019. "Platform competition: Who benefits from multihoming?" *International Journal of Industrial Organization* 64:1–26.

Bernstein, Fernando, Gregory A DeCroix, and N Bora Keskin. 2021. "Competition between two-sided platforms under demand and supply congestion effects." *Manufacturing & Service Operations Management* 23 (5):1043–1061.

Besbes, Omar, Francisco Castro, and Ilan Lobel. 2021. "Surge pricing and its spatial supply response." *Management Science* 67 (3):1350–1367.

Bimpikis, Kostas, Ozan Candogan, and Daniela Saban. 2019. "Spatial pricing in ride-sharing networks." *Operations Research* 67 (3):744–769.

Buchholz, Nicholas, Laura Doval, Jakub Kastl, Filip Matejka, and Tobias Salz. 2025. "Personalized pricing and the value of time: Evidence from auctioned cab rides." *Econometrica* 93 (3):929–958.

Castillo, Juan Camilo. 2025. "Who benefits from surge pricing?" *Econometrica* 93 (5):1811–1854.

Conlon, Christopher T and Julie Holland Mortimer. 2013. "Demand estimation under incomplete product availability." *American Economic Journal: Microeconomics* 5 (4):1–30.

Cullen, Zoë and Chiara Farronato. 2021. "Outsourcing tasks online: Matching supply and demand on peer-to-peer internet platforms." *Management Science* 67 (7):3985–4003.

Farronato, Chiara, Jessica Fong, and Andrey Fradkin. 2024. "Dog eat dog: Balancing network effects and differentiation in a digital platform merger." *Management Science* 70 (1):464–483.

Flinn, Christopher and James Heckman. 1982. "New methods for analyzing structural models of labor force dynamics." *Journal of econometrics* 18 (1):115–168.

Frechette, Guillaume R, Alessandro Lizzeri, and Tobias Salz. 2019. "Frictions in a competitive, regulated market: Evidence from taxis." *American Economic Review* 109 (8):2954–2992.

Ivaldi, Marc and Jiekai Zhang. 2022. "Platform mergers: Lessons from a case in the digital TV market." *The Journal of Industrial Economics* 70 (3):591–630.

Jeziorski, Przemysław. 2014. "Effects of mergers in two-sided markets: The US radio industry." *American Economic Journal: Microeconomics* 6 (4):35–73.

Lee, Robin S. 2013. "Vertical integration and exclusivity in platform and two-sided markets." *American Economic Review* 103 (7):2960–3000.

Li, Zhuoxin and Ashish Agarwal. 2017. "Platform integration and demand spillovers in complementary markets: evidence from Facebook's integration of Instagram." *Management Science* 63 (10):3438–3458.

Ngai, L Rachel and Silvana Tenreyro. 2014. "Hot and cold seasons in the housing market." *American Economic Review* 104 (12):3991–4026.

Rochet, Jean-Charles and Jean Tirole. 2003. "Platform competition in two-sided markets." *Journal of the european economic association* 1 (4):990–1029.

Rosaia, Nicola. 2025. "Competing platforms and transport equilibrium." *Econometrica* 93 (6):2235–2271.