How matching platforms respond to demand shocks: Evidence from Singapore's taxi market*

PRELIMINARY WORK. COMMENTS APPRECIATED.

Kwok-Hao Lee,[†] Wenxuan Xu[‡]

May 8, 2026

Abstract

How does real-time information provision shape decentralized supply responses to sudden demand shocks? We study taxi responses to public transit disruptions in Singapore using high-frequency GPS data and quasi-experimental variation from rail service disruptions between 2016 and 2018. Taxi drivers serve booking and street-hail rides, which primarily differ based on the information available before a driver commits. Booking requests contain each passenger's pickup location and destination, which are broadcast to nearby drivers before acceptance. Drivers serving street-hail trips do not have this information before the passenger enters their vehicle. We find that taxi trips rise during disruptions, but the response is concentrated in the information-rich channel. Booking trips increase by about 26%, roughly an order of magnitude larger than the street-hail response; this booking advantage is robust to controlling for local idle taxi supply. A ride-hailing sample from a separate time period shows a similar positive response (7%). We interpret the booking premium as evidence that customer-specific information and the driver acceptance margin are especially valuable when demand shifts abruptly and local search signals are noisy.

1 Introduction

How does real-time information provision shape decentralized supply responses to sudden demand shocks? When demand shifts abruptly, can customer-specific information help self-directed service providers find and serve newly concentrated demand? These questions connect theories of information, search, and organizational adaptation, where prior work has emphasized how organizational structure shapes search and performance under complexity and turbulence (e.g., Csaszar 2012; Levinthal 1997; Siggelkow and Rivkin 2005). Yet clean causal evidence on how real-time information provision changes decentralized supply responses during sudden demand shocks remains scarce. Information systems are often bundled with firm identity, driver quality, pricing, and contractual terms, making their role hard to isolate. This paper isolates the role of information provision by exploiting a natural experiment in which two channels coexist *within the same taxi firm*: one where drivers search without customer-specific information, and one where booking notifications reveal trip terms before acceptance.

We study Singapore's taxi industry during a series of rail disruptions from December 2016 to March 2018. The setting offers a rare identification opportunity: *within the same company*, the same taxis can

*We thank Ivan Png for access to our taxi data and Wan Wei Khor and Sally Yang for excellent research assistance. Special thanks also go to our ridehailing data provider. This research is supported by the Ministry of Education, Singapore (MOE), under the Social Science and Humanities Research Fellowship (SSRC2024-SSHR-001) and Thematic Grant (MOE2022-SSRTG-029). Any opinions or findings herein are those of the authors and do not reflect the views of MOE or our employers.

[†]Corresponding author; NUS Business School and NUS Artificial Intelligence Institute, kwokhao@nus.edu.sg

[‡]NUS Business School, wenxuan.xu@nus.edu.sg

serve both street-hail and booking trips. In *street-hail*, a trip is formed when a driver chooses where to cruise and is flagged down by a passenger; the driver does not choose that match after observing its destination, and once flagged down must serve the hail. In *booking dispatch*, a trip is formed when a driver accepts a notification that includes the customer's pickup location and destination. Thus, the channel distinction is not driver identity or authority, but whether the match is formed through search without customer-specific trip terms or through an information-rich booking notification.

Both channels operate under a common contractual environment: Singapore taxi drivers pay a fixed daily lease to the fleet operator and keep all gross fare, so the lease structure, fare schedule, and firm identity are held constant. What differs is how trip information enters the driver's decision. Rail disruptions then act as quasi-experiments: as unexpected, localized, and time-sensitive shocks to commuter demand, they shift demand toward specific spatial markets at specific times. This lets us test whether booking notifications that reveal pickup locations and destinations before acceptance generate more completed trips than street-hail search, where drivers commit to a location before observing the trip terms of the passenger they will serve. To compare this information-rich matching channel with app-dispatched service that also uses flexible pricing, we complement the taxi analysis with a near-contemporaneous islandwide 10% random sample of ride-hailing trips from other operators in Singapore, using the same spatial event-study design. Together, these results highlight the role of customer-specific trip information in helping drivers respond when demand shifts abruptly toward congested locations.

Our empirical strategy compares taxi trip volumes in areas with disrupted versus undisrupted rail service using a spatial event study. We document three findings. First, aggregate taxi trips show a positive response to disruptions (roughly 6%, marginally significant), and this aggregate effect masks a sharp divergence across booking and street-hail channels. Second, decomposing by channel reveals a sharp booking premium. Within the same firm, booking trips rise by about 26% during disruptions, while street-hail trips show a small and statistically insignificant 3% increase. Because the same fleet can serve both channels, this pattern raises the possibility that booking merely diverts trips from street-hail. The evidence does not support that interpretation: both average point estimates are positive, although the street-hail estimate is small and statistically insignificant, and the booking advantage is robust to directly controlling for local idle taxi supply in each area-hour cell. Third, as an external validity check, ride-hailing trips in the islandwide sample rise by 7% during disruptions, showing that app-dispatched trips also respond positively to sudden transit shocks.

These results contribute to research on information, search, and platform-mediated matching in service markets. Our within-firm design holds the contractual environment fixed (same driver pool, lease structure, fare schedule, firm culture) and identifies the joint effect of customer-specific trip information, destination-visible broadcast, and the driver acceptance margin. The mechanism is related to Grossman (1976)'s broader insight that market institutions can aggregate information dispersed across agents: here, the booking channel directly aggregates and broadcasts realized demand information that individual drivers may not observe while searching. It also relates to Cullen and Farronato (2021), who study matching between supply and demand on peer-to-peer platforms for local, time-sensitive services. In our setting, the platform's role is informational rather than managerial: it broadcasts realized trip requests, while drivers retain control over search and acceptance decisions. In a complementary margin, Rawley and Simcoe (2013) study how computerized dispatch adoption reshapes fleet ownership structure in US taxicabs; we hold ownership structure fixed and study service-provision effects of different trip-information regimes within a firm. Agarwal et al. (2025) study a related information-provision intervention at Singapore's Changi Airport; we discuss the contrast with their setting in Section 7.1. We develop two hypotheses and test each using quasi-experimental variation from rail disruptions, complemented by an external-validity check using ride-hailing data.

2 Theory and Hypotheses

2.1 Information provision and decentralized search under demand shocks

Many service organizations and platform markets rely on agents who must allocate effort before observing realized demand. Taxi markets are a canonical example: drivers choose where to search before knowing which passengers they will meet or where those passengers want to go. This search problem becomes sharper during localized demand shocks. A rail disruption shifts passenger demand across space and time, but the locations, destinations, and urgency of affected passengers are not fully observed by drivers before they decide where to search. Drivers may infer demand from station locations, traffic conditions, past experience, or visible queues, but these signals are local and noisy. As a result, even when a disruption creates additional demand for taxi service, the supply response depends on whether drivers receive information that is specific enough to change their search and acceptance decisions.

This setting highlights a classic information problem. Local decision makers often hold dispersed knowledge about conditions they directly observe, while organizations and platforms can aggregate other forms of information that are not equally available to each participant (Grossman 1976; Hayek 1945). In taxi markets, drivers retain local knowledge about traffic, cruising routes, and their own opportunity costs. The dispatch system, by contrast, observes customer requests submitted through the booking channel. These requests contain trip-specific information that is not available to a street-hail driver before encountering a passenger. The economic question is therefore not whether drivers are centrally directed, but whether timely information about individual requests changes how drivers adapt to a demand shock.

This distinction also connects to organizational-design theories of information and decision rights. When information is dispersed, performance depends on how organizations allocate authority and what information reaches the agents who make operational decisions (Alonso, Dessein, and Matouschek 2008; Csaszar 2012; Dessein 2002). In this setting, the platform does not determine drivers' work hours or search locations. Drivers still decide whether to work and where to position themselves. Instead, booking dispatch broadcasts trip-specific requests, adds an informed acceptance margin, and changes the information available at the moment of search and acceptance.[1]

The institutional distinction between street-hail and booking trips is central. In the street-hail channel, drivers choose where to cruise before observing the destination or trip terms of the passenger they will serve. They can decide to search near a disrupted station, but they must do so based on expectations about passenger arrivals and the distribution of trips they may encounter. In the booking channel, the platform aggregates customer requests and broadcasts each request to nearby drivers through in-cab dispatch screens. The pickup location and destination are visible when the driver decides whether to accept or refuse the request; the first driver to accept forms the match (Castillo, Knoepfle, and Weyl 2025; Zhang et al. 2026). The fare is metered and is not visible ex ante. Thus, booking does not eliminate driver discretion. Instead, it adds a trip-specific acceptance margin and changes the information available at the moment of search and acceptance.

Prior work on taxi and ride-hailing markets emphasizes that search frictions, spatial mismatch, and information affect matching outcomes (Agarwal et al. 2025; Buchholz 2022; Castillo 2025; Castillo, Knoepfle, and Weyl 2025; Zhang et al. 2026). During a transit disruption, these frictions become more important because demand changes sharply and unevenly across locations. Street-hail drivers may know that a disruption has occurred but remain uncertain about which stations generate the most

[1] We use "centralized" only in the informational sense: the platform aggregates booking requests before transmitting them to drivers. The booking channel changes the information and contract-formation opportunities available to drivers, rather than reallocating managerial control over their work hours, search locations, or acceptance decisions.

feasible trips, how quickly passengers will arrive, and whether road congestion will make service costly. Booking notifications reduce some of this uncertainty by revealing a realized passenger request. They therefore make disruption-induced demand more actionable for drivers who are nearby and willing to accept the request.

2.2 Hypothesis 1: Booking notifications and disruption responses

The first implication is that booking trips should respond more strongly than street-hail trips to transit disruptions. Both channels are served by drivers in the same taxi firm and draw on the same broad driver pool, so the comparison holds fixed many features of the firm, vehicle fleet, fare regulation, and local operating environment. The relevant difference is the information and acceptance margin created by booking dispatch. A street-hail driver must decide whether to search in a disrupted area before observing the passenger and destination. A booking driver can observe a concrete request from a nearby passenger and condition acceptance on the pickup and destination information shown on the dispatch screen.

This informational advantage should matter most when the disruption creates additional passenger demand near affected stations. Booking notifications transform latent or difficult-to-observe demand into visible requests. They can also reduce wasted search by helping drivers identify where a passenger is waiting and whether the trip is acceptable before committing to the match. Street-hail drivers can still respond to disruptions, but their response relies more heavily on spatial expectations and local search. Therefore, when transit disruptions increase demand for taxi service in affected areas, the booking channel should show the larger increase in realized trips.

Hypothesis 1 (Booking advantage)**.** *During transit disruptions, booking trips originating in disrupted areas increase more than street-hail trips originating in the same areas, within the same firm.*

2.3 Hypothesis 2: Evening-peak search costs and the booking advantage

The booking advantage should not be uniform across all hours. A natural boundary condition is the evening peak. During evening peak hours, road congestion raises the cost of cruising into disrupted areas and the opportunity cost of unsuccessful search. Pickups may take longer, and local signals such as queues or traffic may be harder to distinguish from ordinary peak-hour congestion. These conditions make street-hail search in disrupted areas more costly and less informative: a driver must enter or remain near the area before learning whether passengers are available and before a street-hail match forms.

Booking notifications are especially valuable under these conditions because they provide request-level information before the driver accepts the trip. The driver observes a specific passenger request, pickup location, and destination, and can decide whether to accept given current traffic conditions, location, and opportunity costs. This does not imply that booking should dominate street-hail in every rush-hour period or that all peak periods generate the same mechanism. The prediction is specific to evening peak hours during disruptions, when the costs of uninformed cruising and the value of trip-specific information are likely to be particularly high.

Hypothesis 2 (Evening-peak moderator)**.** *The booking advantage is larger during evening peak hours, when street-hail search in disrupted areas is more costly and less informative.*

2.4 External benchmark: Ride-hailing platforms

Ride-hailing platforms provide a useful external benchmark because they also intermediate passenger requests through digital interfaces rather than relying on street-hail search. If digital request information helps drivers respond to disruption-induced demand, ride-hailing trips may also increase around disrupted areas. This comparison can help assess whether the main pattern is consistent with a broader role for app-based matching technologies in facilitating supply responses to localized demand shocks.

At the same time, ride-hailing is not a clean mechanism test for the booking channel studied above. Ride-hailing differs from taxi booking along several dimensions, including pricing, driver composition, interface design, regulation, firm identity, and sample period. These differences mean that the ride-hailing comparison is suggestive rather than a clean estimate of the informational mechanism. The benchmark is therefore best interpreted as external evidence on whether another digitally mediated matching technology exhibits a similar disruption response, not as a third hypothesis about the same within-firm channel comparison.

3 Empirical Setting

3.1 Public and point-to-point transportation in Singapore, 2016–2018

Our empirical setting is Singapore, a dense, polycentric city-state where high-capacity public transport is the backbone of daily travel. During our study period (2016-2018), the rail network had expanded to 228.1 km by end-2017 (199.3 km MRT and 28.8 km LRT), providing fast cross-island connectivity that is tightly integrated with an extensive feeder-bus system.[2] Utilization is correspondingly high: in 2017, average daily ridership was about 3.122 million on MRT, 190,000 on LRT, and 3.952 million on public buses.[3] These magnitudes are central for our design: an MRT disruption is not a niche event, but a system-level shock that can shift large numbers of commuters into alternative modes over short horizons.

Point-to-point (P2P) transport, consisting of street-hail/dispatch taxis and app-dispatched private-hire vehicles (PHVs), provides the main outside option when rail service is degraded. On an average day in 2017, the P2P sector served roughly 755,000 trips across taxis and private-hire vehicles, a scale comparable to a major rail line and large enough to matter for both congestion and welfare during disruptions.[4] Importantly, P2P supply is institutionally segmented by (i) vehicle type and licensing (taxis vs chauffeur-driven PHVs), and (ii) platform access and dispatch technology.

As of end-2017, Singapore had 23,140 licensed taxis and 46,903 chauffeur-driven private-hire cars (distinct from self-drive rental cars).[5] Market structure on the platform side in 2017 featured active competition between two large ride-hailing platforms, Grab and Uber, alongside smaller and hybrid taxi-booking aggregators. Grab operated a large two-sided ride-hailing marketplace and, beginning in 2017, expanded its ability to dispatch taxis through products such as JustGrab, which pooled participating taxi fleets with Grab's private-hire supply under upfront (and potentially dynamic) pricing.[6] Uber remained an active ride-hailing platform throughout 2017 (exiting Southeast Asia only in 2018), competing head-

[2] https://www.lta.gov.sg/content/dam/ltagov/who_we_are/statistics_and_publications/statistics/pdf/Rail_Length_2023.pdf

[3] https://datamall.lta.gov.sg/content/dam/datamall/datasets/Facts_Figures/Public%20Transport/yearly_ave_daily_pt_ridership.csv

[4] https://data.gov.sg/datasets/d_ba615ec4cc5d9f5b7800ad82057f36f1/

[5] https://www.lta.gov.sg/content/dam/ltagov/who_we_are/statistics_and_publications/statistics/pdf/MVP01-1_MVP_by_type.pdf

[6] https://www.channelnewsasia.com/today/voices/why-comfortdelgro-needs-up-its-game-catch-up-5559801

to-head on prices and driver/passenger acquisition during our window.[7] In parallel, traditional taxi operators maintained their own dispatch channels while also multi-homing onto third-party apps; for example, ComfortDelGro, the largest taxi operator, launched a taxi-booking service on Ryde in May 2017, illustrating the broader shift toward platform-mediated matching even for the incumbent taxi fleet.[8]

This combination of (a) a heavily used rail backbone that occasionally experiences nontrivial disruptions and (b) a large, platform-mediated P2P sector with meaningful capacity makes Singapore a clean setting to study how dispatch and matching technologies shape the system's resilience and the distribution of service during transit shocks.

3.2 Rail (MRT) Disruptions in Singapore

Rail (MRT) disruptions in Singapore during our study period are best characterized as infrequent but high-impact shocks that are largely unanticipated by commuters and P2P drivers. While some service changes reflect scheduled engineering works—typically pre-announced and concentrated in off-peak periods—the majority of operational incidents (e.g., signaling faults, power failures, rolling stock issues) arise with little notice and can propagate across segments of a line. Given the scale of MRT usage, even localized failures can affect many commuters within short time windows, generating sharp and unplanned shifts in demand toward substitute modes such as taxis and ride-hailing.

Historically, MRT reliability has been a salient public issue. High-profile breakdowns in the early 2010s and subsequent clusters of incidents heightened media scrutiny and public sensitivity to service reliability. Although aggregate reliability metrics improved following substantial maintenance and renewal efforts, more recent episodes during our study period reinforced the perception that disruptions, while infrequent, are salient and disruptive events. Media coverage and public discourse tend to frame unplanned breakdowns as acute failures that impose widespread inconvenience, amplifying their perceived severity relative to their frequency.

This background is important for our empirical design. At short horizons, commuters and drivers have limited ability to anticipate or strategically respond to disruptions, so observed adjustments largely reflect responses to exogenous variation in public transit(rail) quality and access. Disruptions occur at a much higher frequency than underlying changes in preferences or economic conditions, and their precise timing and location are difficult to predict at the individual level. This generates plausibly exogenous, high-frequency variation suitable for identifying short-run substitution patterns.

That said, this interpretation requires care. First, scheduled disruptions may induce anticipation, such as trip rescheduling or pre-emptive mode switching, and are therefore treated separately or excluded. Second, even unplanned incidents may exhibit systematic temporal patterns, for example being more likely during peak load or on specific lines, which could correlate with demand conditions; we address this using rich time fixed effects and event-time comparisons. Third, information frictions are incomplete: while real-time updates are disseminated through apps and announcements, these typically lag the initial disruption and are unevenly received, so the initial shock remains salient for a large share of commuters.

Overall, by focusing on unplanned, short-notice disruptions and conditioning on predictable temporal variation, MRT incidents provide a credible source of plausibly exogenous shocks to urban transport supply. This setting allows us to study how demand reallocates across modes and how the private transit modes absorb sudden increases in travel demand.

[7] https://lkyspp.nus.edu.sg/docs/default-source/case-studies/revisiting-the-sharing-economy-%28updated-092017%29.pdf?sfvrsn=eaa8950b_0

[8] https://www.todayonline.com/singapore/comfortdelgro-ties-ryde-taxi-bookings

4 Data

The three core data sets comprise taxi activity itineraries, a sample of ridehailing trips, and a list of rail disruptions during the sample period.

4.1 Taxi activity data

Our taxi activity microdata are from a large taxi company in Singapore; the data span December 2016 to May 2017, then August 2017 to March 2018 (Png and Wang 2025). In this data set, each observation is a (vehicle, activity spell) pair. In each observation, we see the anonymized driver ID, vehicle status (e.g., offline, searching, "on call" en route to picking up a passenger, waiting, passenger on board, arrived, payment made, on break), the start and end times, and location coordinates of the spell, the distance traveled during the spell, and the whether the spell belongs to a particular job. A new spell is considered to have begun once the status of the vehicle has changed, or if 5 minutes have elapsed, whichever occurs first. Details are available in Appendix A.1.

4.2 Ridehailing trip data

We also have a 10% sample of ridehailing trips in Singapore from a confidential data provider. These trips span April 2018 to October 2018. Each observation is a completed trip. We observe the time at which the booking is made and/or canceled; when and where the passenger boards and disembarks; and the fare paid. Details are available in Appendix A.2.

4.3 Rail disruption data

We complement these proprietary data with two public data sources. The first is rail disruption data from `mrtdown.org`, a community-run transit monitoring website that collects publicly dispersed disruption information (social media) by local transit providers. Each observation corresponds to one disruption episode. We see the start and end date and time of the disruption, a classification of the type of disruption (train fault, signal fault, etc), a list of affected lines and stations, and the plain text of these social media updates by the transit operator. A total of 97 disruption events occurred during our sample period. These disruptions are summarized in Table 1.

Variable	N	Mean	Std. Dev.	Min	Max
Duration (hours)	97	2.91	5.66	0.06	31
Stations affected (count)	97	11.10	7.71	1	34
Expected (indicator)	97	0.21	0.41	0	1
No service (indicator)	97	0.19	0.39	0	1
Inferred delay (minutes)	68	18.90	8.93	5	45

Table 1: *Summary statistics of MRT disruptions in Singapore during the sample period December 2016 to May 2017, then August 2017 to March 2018. Data are from `mrtdown.org`.*

Further details on data construction are available in Appendix A.

4.4 Descriptive statistics

We begin by visualizing the train disruptions in our `mrtdown.org` data in Figure 1. Each disruption event is a dot. The height of each bar beneath each dot is our proxy for the severity of each disruption: we multiply the total number of affected stations by the duration of the disruption. There are around 600 disruptions in our full sample, 97 of which occurred during our sample period and are visualized. Most disruptions are minor and last under 30 minutes. However, there are several severe disruptions: a signal fault affecting the whole of the Downtown Line on 1 March 2018, and a train derailment affecting Queenstown through Boon Lay (more than 25%) of the East-West Line from 25 to 30 September 2024.

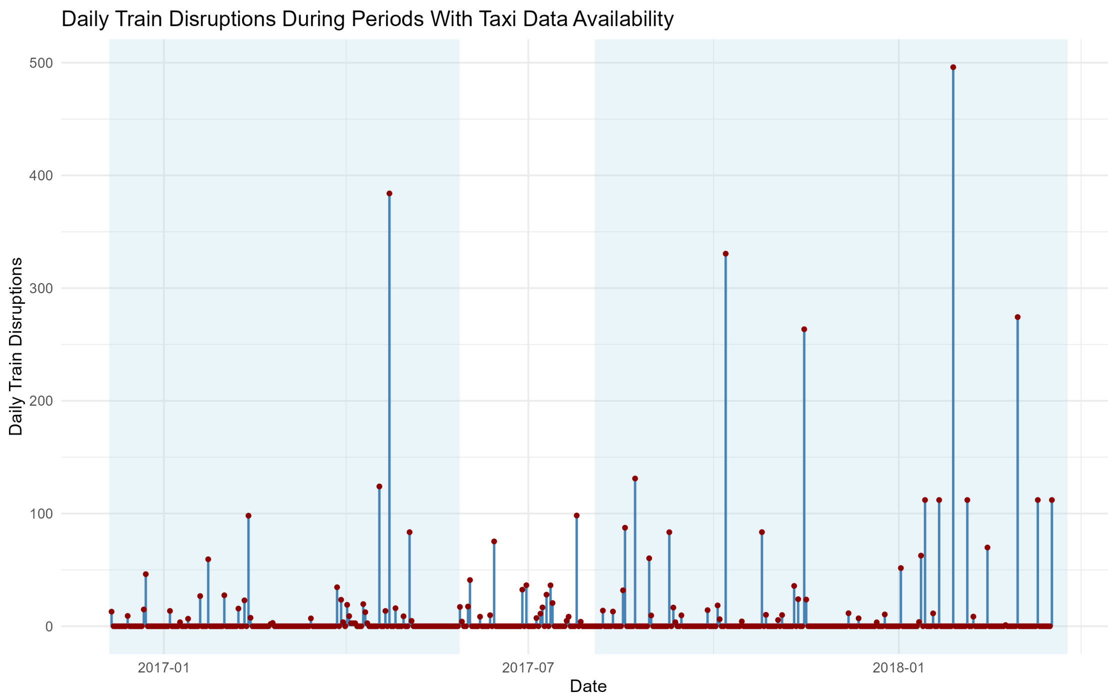

Figure 1: *Disruption events by week from December 2016 to March 2018. Each disruption event is a dot. The height of each bar represents our measure of the severity of each disruption: the total number of affected stations multiplied by the duration of the disruption. The 97 disruption events in the blue shaded area are covered in our sample period. Data are from `mrtdown.org`.*

To illustrate the spatial variation that our empirical strategy exploits, Figure 2 maps the percentage deviation in taxi trips from baseline during a single disruption event: a Downtown Line fault on Tuesday 31 January 2017, 18:43–20:26. The baseline for each planning area is the average trip count for that area on the same day of week and hour of day across all non-disrupted periods. Planning areas containing disrupted stations (outlined in black) show elevated taxi pickups relative to baseline, consistent with displaced commuters substituting into taxis. The spatial pattern previews the treatment variation underlying our formal estimates.

5 Empirical Strategy

Our analysis revolves around the causal effects of transit disruptions on taxi trips through booking and street-hail channels, as well as ridehailing trips. For instance, we want to measure whether taxi and

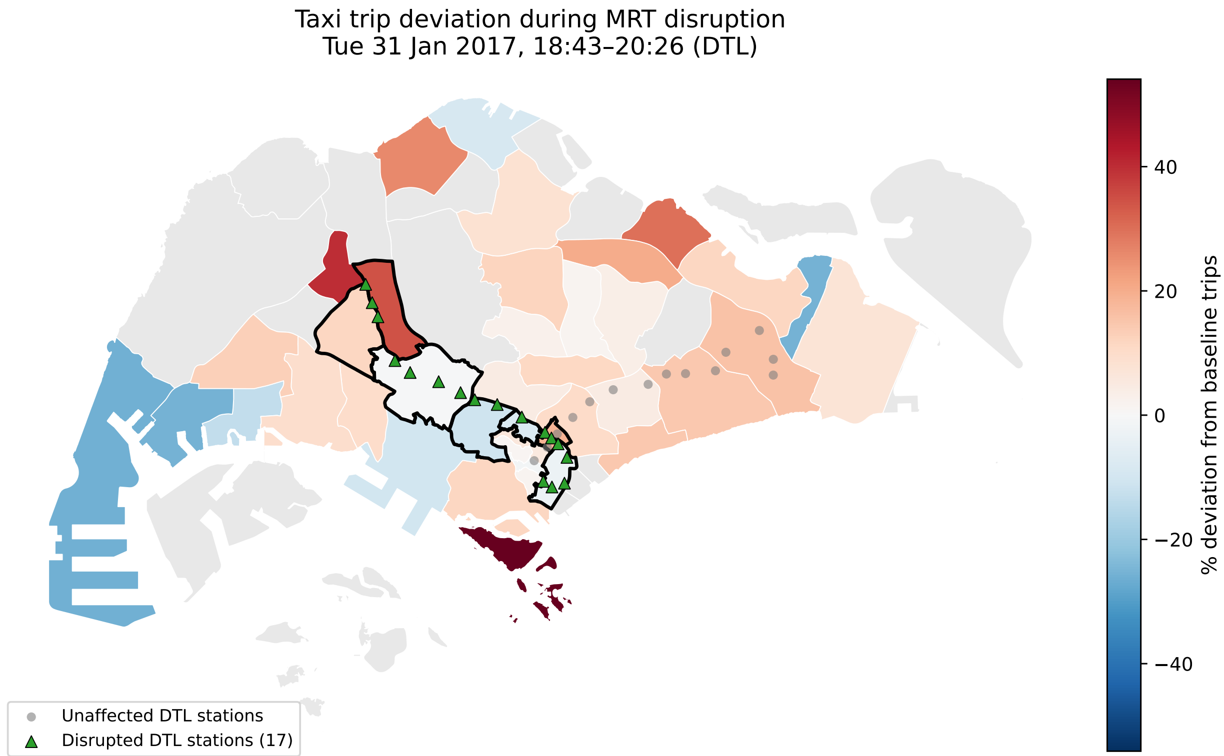

Figure 2: *Taxi trip deviation from baseline during a Downtown Line disruption, Tuesday 31 January 2017, 18:43–20:26. Each planning area is shaded by the percentage deviation in taxi trips relative to the average for that planning area on the same day of week and hour of day across non-disrupted periods. Green triangles mark disrupted stations; grey circles mark unaffected Downtown Line stations. Black outlines indicate planning areas containing at least one disrupted station. Planning areas with fewer than 50 baseline trips per hour during the disruption window are greyed out. Trip data are from a large taxi company in Singapore.*

ridehailing trips rise, as a reasonable observer would expect, in areas where a transit disruption occurs. To estimate these disruption effects, we would ideally randomize the location, severity, and length of each disruption, and the spatial configurations of taxis and private-hire vehicles at the time of the disruption. Because this experiment is infeasible, we instead exploit "as-good-as-random" disruptions across space and time.

Below, we describe how we specify the research design and the formal econometric framework for our spatial event study analysis.

5.1 Spatial event study design

We adopt a spatial event study design. Our core spatial unit is the planning area; Singapore is partitioned into 55 of them.

We collapse our trip data by (planning area, pickup hour, calendar date). In our main specification, treated markets are observations with disruptions; controls are without. Given a location i, hour of day h, calendar date t, and mode $m \in \{\text{all taxi}, \text{street-hail}, \text{booking}, \text{ridehailing}\}$, our outcome of interest is the count of mode-m trips, $n_{iht}^{(m)}$.[9]

Our primary estimating equation is

$$y_{iht}^{(m)} = \beta^{(m)} \times \mathbb{1}\{\text{Disrupted}\}_{iht} + \gamma_i^{(m)} + \gamma_h^{(m)} + \gamma_t^{(m)} + \varepsilon_{iht}^{(m)}, \quad (1)$$

where $y_{iht}^{(m)} = \ln(1 + n_{iht}^{(m)})$, $\mathbb{1}\{\text{Disrupted}\}_{iht}$ is an indicator for a train disruption, $\gamma^{(m)} = (\gamma_i^{(m)}, \gamma_h^{(m)}, \gamma_t^{(m)})$ is a vector of location, hour-of-day, and day-of-week fixed effects, and $\varepsilon_{iht}^{(m)}$ is an error term. This specification yields coefficients with a direct percentage-change interpretation: a disruption increases mode-m trips by approximately $100 \times \beta^{(m)}$ percent.

In the baseline specification, we pool all taxi modes ($m =$ all taxi). Our hypothesis tests then compare mode-specific coefficients: Hypothesis 1 compares $\hat{\beta}^{(\text{booking})}$ with $\hat{\beta}^{(\text{street-hail})}$. For Hypothesis 2, we augment Equation (1) with interactions between $\mathbb{1}\{\text{Disrupted}\}_{iht}$ and disruption-type indicators. As an external-validity check (Section 2.4), we separately estimate $\hat{\beta}^{(\text{ridehailing})}$ using data from a ridehailing platform.

5.2 Identification

As is standard, identification arises from comparing units with different planning areas, across different hours of day, across different calendar dates, and between disrupted and undisrupted markets. The key identifying assumption is:

Assumption 1 (Parallel trends)**.** *Absent the disruption, treated and control markets display similar trip patterns, both in aggregate and across modes (taxi booking, taxi street-hail, ridehailing) after controlling for planning area, day-of-week and hour-of-day fixed effects.*

One possible violation is that treated markets differ systematically from control markets in ways that make them more sensitive to underlying demand fluctuations even in the absence of a disruption. For example, planning areas with more rail stations may have stronger peak-hour commuting patterns, different reliance on street-hail versus bookings, or greater exposure to weather, events, and road congestion. In that case, treated-control differences around disruption episodes could reflect differential secular intraday demand patterns rather than the causal effect of the transit disruption itself.

Another concern could arise from endogenous disruption timing. If rail disruptions are more likely to occur during periods or in locations with unusually high passenger flows, then treated markets may already be on different demand trajectories before the disruption begins. This explanation is especially relevant when comparing across hail modes, because commuters facing tight time constraints may disproportionately substitute into bookings or ridehailing during busy periods even absent a disruption. Under this scenario, estimated treatment effects could partly capture mode-specific demand reallocation associated with predictable rush-hour conditions rather than the disruption itself.

To address these concerns, we first inspect covariate balance across planning areas by disruption exposure tercile. Table A.1 shows that in the urbanized planning areas in the middle and upper terciles of exposure, where identification operates, income and mean fare are comparable across terciles, while high-exposure areas are somewhat less populous and have lower search duration than medium-exposure areas, and mechanically have more MRT stations and smaller land area. These level differences are absorbed by planning area fixed effects.

[9]Some (planning area, hour, date) cells contain zero trips (approximately 20% in the taxi data and 11% in the ridehailing data). Following Chen and Roth (2024), in Section 6.5, we estimate the log-linear specification on cells with positive trips, assessing the role of zeros through separate robustness checks.

Next, we test Assumption 1 using a stacked event-time design. For each of the 97 disruption events in our taxi sample, we define a ± 4-hour window around the disruption start hour. Treated planning areas are those disrupted in the focal event; control planning areas are all planning areas not disrupted during the event window. This allows a planning area treated in one event to serve as a control in another event, so long as it is not concurrently disrupted. The resulting control pool is substantially larger and more comparable to treated areas than restricting to never-treated areas alone, and the design remains valid whenever disruptions at different stations are not systematically correlated with unobserved demand shocks at non-disrupted stations. Stacking all events, we estimate:

$$y_{iet} = \sum_{k \neq -1} \beta_k \cdot \mathbf{1}[\text{event time} = k] \cdot \text{Treated}_i + \alpha_e + \gamma_i + \delta_t + \varepsilon_{iet} \tag{2}$$

where $k \in \{-4, \ldots, +3\}$ indexes hours relative to disruption onset, α_e are event fixed effects, γ_i are planning area fixed effects, δ_t are event-(time-of-day) fixed effects, and we omit $k = -1$ as the reference period. Standard errors are clustered at the planning area level. The coefficients β_k trace out the differential trajectory of treated versus control areas relative to the hour immediately before the disruption begins.

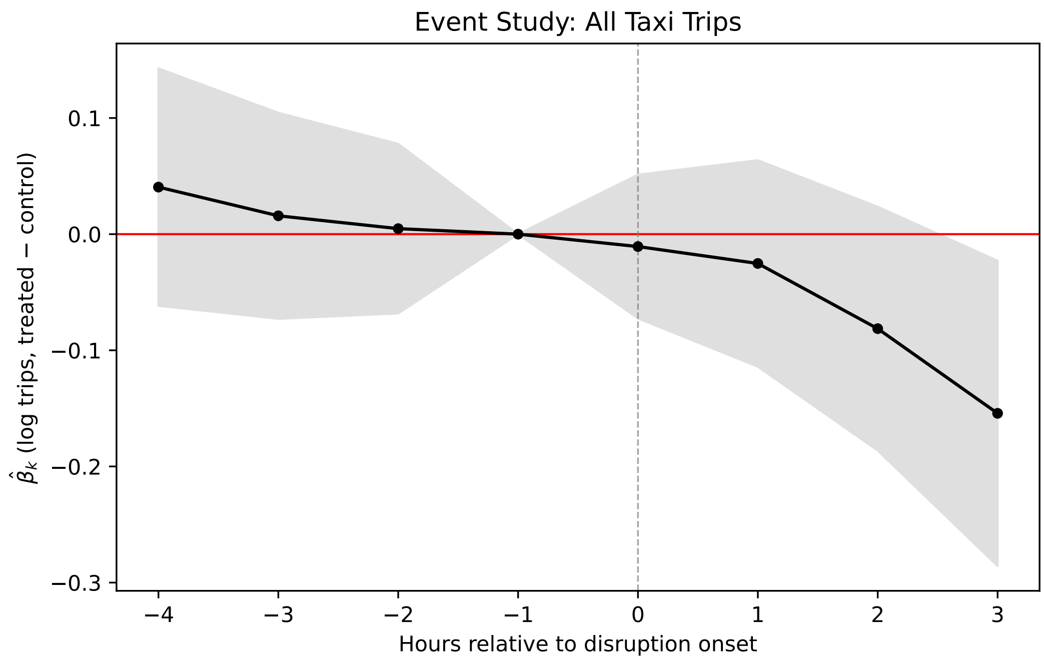

Figure 3: *Event study estimates of the differential effect of MRT disruptions on taxi trips in treated versus control planning areas. For each event, control areas are those not disrupted by any event during the event window. Each point plots the estimated $\hat{\beta}_k$ from the stacked event-time regression, with the shaded band showing the 95% confidence interval. The reference period is $k = -1$. The dashed vertical line marks disruption onset ($k = 0$). Pre-period coefficients ($k = -4$ to $k = -2$) are individually and jointly insignificant (F = 0.364,* p *= 0.78), supporting the parallel trends assumption.*

Figure 3 displays the results for aggregate taxi trips. The pre-period coefficients ($k = -4$ through $k = -2$) are individually insignificant and close to zero; a joint F-test fails to reject the null of no pre-trend ($F = 0.364$, $p = 0.78$). This finding supports the identifying assumption that treated and control areas follow parallel aggregate trajectories before disruption onset; we do not separately test mode-specific pre-trends.

The findings above show the importance of comparing “like-for-like”: observing the same planning area in both undisrupted and disrupted states provides crucial identifying variation. To operationalize these comparisons, in our main analysis, we compare treated and control markets with rich fixed effects that control for planning area, hour-of-day and day-of-week fixed effects.

6 Results

We begin by assessing the average effect of a train disruption on taxi trips. We estimate Equation (1) with $m=$ all on our taxi data.

6.1 Aggregate effect: Transit disruptions do not significantly increase taxi trips

The results from our main specification are displayed in Table 2. Our preferred specification with location, hour-of-day, and day-of-week fixed effects is in Column 2. We find that, on average, treated markets see a roughly 6% increase in trips (β = 0.063) relative to control markets, marginally significant at the 10% level. This aggregate effect is consistent with taxis absorbing displaced commuters during rail disruptions. Columns 3 and 4 interact the disruption indicator with whether the disruption was unexpected or involved a full service outage, respectively. Neither interaction term is statistically significant, suggesting that the aggregate taxi response does not differ markedly by disruption type.

6.2 Test of Hypothesis 1: Dispatch vs. street-hail within firm

The aggregate result in Table 2 pools all taxi trips regardless of how the ride was initiated. Hypothesis 1 predicts that booking dispatch increases service provision relative to street-hail during disruptions. To test this, we estimate Equation (1) separately for booking trips ($m=$ booking) and street-hail trips ($m=$ street-hail). Tables 3a and 3b present the results.

The contrast is striking. For taxi bookings (Table 3a), the preferred specification (Column 2) shows an about 26% increase in trips during disruptions (β = 0.228), statistically significant at the 1% level. Column 3 splits the disruption indicator into planned and unexpected components: the planned main effect is small and insignificant (β = 0.038), while unexpected disruptions add an interaction increment of 0.244 log points (significant at the 5% level), implying a total unexpected-disruption effect of about 0.282 log points, or roughly 33%. The booking response is therefore concentrated in unexpected disruptions, consistent with the H2 prediction that the booking channel’s information advantage matters most when cruising drivers’ local knowledge is least informative.

For street-hail trips (Table 3b), the picture is markedly different. The preferred specification (Column 2) shows an about 3% increase (β = 0.026) that is not statistically significant. Neither the unexpected nor the no-service interaction is significant. Street-hail drivers, who independently choose where to cruise without customer-specific trip information, do not systematically increase service provision in disrupted areas.

The within-firm comparison is sharp: holding the contractual environment fixed (same driver pool, lease structure, fare schedule, firm culture), $\hat{\beta}^{(\text{booking})}=0.228$ while $\hat{\beta}^{(\text{street-hail})}=0.026$. To test this difference formally, we stack the booking and street-hail panels, interact the disruption indicator with a booking-mode dummy, and absorb mode-specific planning area, hour, and day-of-week fixed effects. The interaction coefficient is 0.202 (SE 0.028, $p<0.001$), confirming that booking dispatch generates a statistically significantly larger trip response than street-hail, consistent with Hypothesis 1.

ln(1 + trips)	(1) Baseline	(2) All FE	(3) Unexpected	(4) No Service	(5) Search Pool
Disrupted	1.2593***	0.0633*	0.0395	0.0618*	0.0771***
	(0.2590)	(0.0327)	(0.0291)	(0.0351)	(0.0276)
Unexpected			0.0306		
			(0.0476)		
No Service				0.0191	
				(0.0489)	
Search Minutes					0.0000***
					(0.0000)
Observations	571,200	571,200	571,200	571,200	571,200
Adjusted R^2	0.001	0.934	0.934	0.934	0.936
Fixed Effects:					
Hour FE		✓	✓	✓	✓
Day-of-Week FE		✓	✓	✓	✓
Planning Area FE		✓	✓	✓	✓

Table 2: *Log-linear OLS estimates of the effect of transit disruptions on aggregate taxi trips. "Unexpected" excludes planned disruptions. "No Service" indicates that at least one rail service is completely non-operational along some stretch of stations. "Search Pool" controls for the total taxi-minutes spent searching or waiting in the planning area during that cell-hour. The dependent variable is $\ln(1 + \textit{trips})$ of taxi trips originating from a given planning area. Each observation is a (planning area, hour, date) cell. Standard errors are two-way clustered at the planning area and disruption event level. Trip data are from a large taxi company in Singapore, covering December 2016 to March 2018. Disruption data are from `mrtdown.org`. Asterisks ($*, **, ***$) indicate statistical significance at the 10%, 5% and 1% level respectively.*

6.3 Test of Hypothesis 2: Evening-peak heterogeneity by dispatch channel

Hypothesis 2 predicts that the booking advantage should be larger when street-hail search in disrupted areas is more costly and less informative. The evening peak provides a natural setting for this boundary condition: road congestion raises the cost of cruising into affected areas, while local signals about passenger availability may be harder to separate from ordinary peak-hour traffic conditions.

To assess this time-of-day heterogeneity, we augment Equation (1) with interactions between the disruption indicator and pickup-hour dummies, yielding 24 hour-specific treatment effects. We then decompose the response by dispatch channel. Figure 4a plots the aggregate hour-specific estimates and confidence intervals.

The aggregate estimates reveal substantial within-day heterogeneity. Point estimates are positive in several morning and mid-day hours, but the aggregate response weakens during the evening peak. Because the confidence intervals are wide for many individual hours, we do not interpret each hourly coefficient as a separate precise estimate. Instead, the hour-by-hour pattern motivates the more informative comparison below: whether booking and street-hail respond differently when evening-peak search conditions make street-hail matching more costly.

Figures 4b and 4c plot $\hat{\beta}^{(\text{booking})}$ and $\hat{\beta}^{(\text{street-hail})}$ separately across hours. The channel contrast is the key pattern. Booking point estimates are positive in most hours, including the late evening. Street-hail point estimates are positive mainly during mid-morning off-peak hours and turn negative during the

ln(1 + trips)	(1) Baseline	(2) All FE	(3) Unexpected	(4) No Service	(5) Search Pool
Disrupted	1.1886***	0.2276***	0.0382	0.2272***	0.2101***
	(0.1013)	(0.0375)	(0.1009)	(0.0411)	(0.0324)
Unexpected			0.2440**		
			(0.1077)		
No Service				0.0055	
				(0.1067)	
Search Minutes					-0.0001***
					(0.0000)
Observations	571,200	571,200	571,200	571,200	571,200
Adjusted R^2	0.002	0.850	0.850	0.850	0.854
Fixed Effects:					
Hour FE		✓	✓	✓	✓
Day-of-Week FE		✓	✓	✓	✓
Planning Area FE		✓	✓	✓	✓

(a) *Trips fulfilled by booking dispatch.*

ln(1 + trips)	(1) Baseline	(2) All FE	(3) Unexpected	(4) No Service	(5) Search Pool
Disrupted	1.2319***	0.0259	0.0027	0.0237	0.0492
	(0.2679)	(0.0418)	(0.0326)	(0.0452)	(0.0300)
Unexpected			0.0299		
			(0.0578)		
No Service				0.0274	
				(0.0516)	
Search Minutes					0.0001***
					(0.0000)
Observations	571,200	571,200	571,200	571,200	571,200
Adjusted R^2	0.001	0.935	0.935	0.935	0.941
Fixed Effects:					
Hour FE		✓	✓	✓	✓
Day-of-Week FE		✓	✓	✓	✓
Planning Area FE		✓	✓	✓	✓

(b) *Trips fulfilled by street-hail.*

Table 3: *Log-linear OLS estimates of the effect of transit disruptions on taxi trips, differentiated by booking and street-hail modes. "Unexpected" excludes planned disruptions. "No Service" indicates that at least one rail service is completely non-operational along some stretch of stations. "Search Pool" controls for the total taxi-minutes spent searching or waiting in the planning area during that cell-hour. The dependent variable is $\ln(1 + \textit{trips})$ of trips of that type (booking/street hail) originating from a given planning area. Each observation is a (planning area, hour, date) cell. Standard errors are two-way clustered at the planning area and disruption event level. Trip data are from a large taxi company in Singapore, covering December 2016 to March 2018. Disruption data are from `mrtdown.org`. Asterisks ($*, **, ***$) indicate statistical significance at the 10%, 5% and 1% level respectively.*

evening peak.

This divergence is consistent with Hypothesis 2 in its evening-peak form. When evening congestion makes street-hail search in disrupted areas costly and local signals less informative, booking requests provide pre-acceptance information about pickup location and destination. Drivers can therefore evaluate a concrete request before committing to the match, whereas street-hail drivers must first enter or remain near the disrupted area before learning whether a feasible match will form. The individual hourly estimates remain imprecise, so we interpret the evening-peak divergence as the primary evidence rather than as a precise ranking of every hour of the day.

This interpretation is also consistent with the robustness evidence reported below. The Chen–Roth decomposition shows that the average booking response is concentrated on the intensive margin, while street-hail does not expand on the extensive margin. In addition, the search-pool exercise shows that disruptions reduce local idle search minutes after fixed effects, so the booking response is not simply a passive accumulation of idle taxis near disrupted stations. A reproducible PM-peak search-pool robustness table will be added separately before we make a PM-specific robustness claim.

6.4 External validity: Ridehailing platforms

We now turn to an out-of-sample test of whether the booking advantage extends beyond one taxi company. As discussed in Section 2.4, ridehailing platforms combine app-based dispatch with flexible surge pricing, offering an independent setting in which to assess whether dispatch-based channels increase service provision during disruptions. We use a separate 10% random sample of ridehailing trips from a confidential data provider, spanning April to October 2018, and estimate Equation (1) with $m =$ ridehailing using the same spatial event study design.

Table 4 presents the results. The preferred specification with fixed effects (Column 2) shows a roughly 7% increase in ridehailing trips during disruptions, positive though imprecisely estimated ($p \approx 0.39$). Column 3 interacts the disruption indicator with whether the disruption was unexpected; the coefficient on planned disruptions is roughly 12% (β = 0.113, insignificant), while the interaction term for unexpected disruptions is roughly -11% (β = -0.109, insignificant). The PPML specification, which may better suit this smaller sample, yields a significant 9% estimate (β = 0.086; Appendix Table A.7), corroborating the positive direction.

The central finding is that ridehailing, like taxi bookings, shows a positive response to disruptions. Both dispatch-based channels—taxi booking under regulated fares ($\hat{\beta}^{(\text{booking})} = 0.228$) and ridehailing under flexible pricing ($\hat{\beta}^{(\text{ridehailing})} = 0.066$), increase service provision, while street-hail shows a smaller and statistically insignificant response ($\hat{\beta}^{(\text{street-hail})} = 0.026$). This suggests that the dispatch-based channel result is not specific to one firm's dispatch technology or driver pool.

We note that the booking point estimate is actually *larger* than the ridehailing estimate, though the two are not directly comparable: they come from different firms, time periods, and data samples. This comparison therefore cannot cleanly identify the incremental effect of surge pricing. We return to the pricing question in Section 7.3.

6.5 Robustness

To assess the robustness of our findings, we consider the following alternative specifications:

Chen-Roth (2024) extensive-intensive margin decomposition. Following Chen and Roth (2024), we decompose the log-linear effect into an extensive margin (linear probability model for $\mathbb{P}(n > 0)$) and an intensive margin (OLS on $\ln n \mid n > 0$). This decomposition addresses concerns about

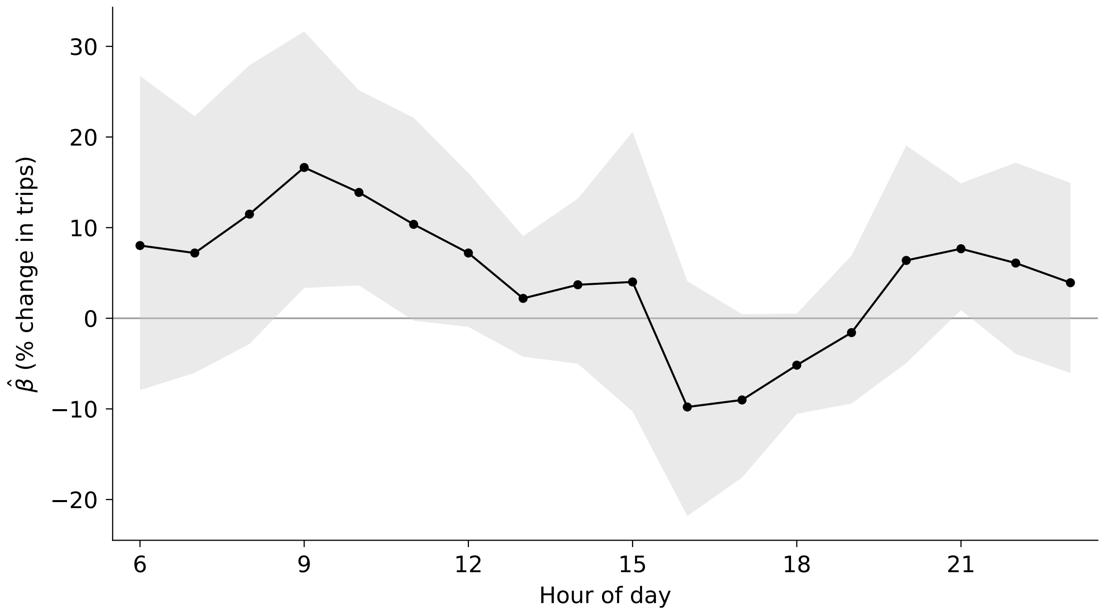

(a) *All taxi trips.*

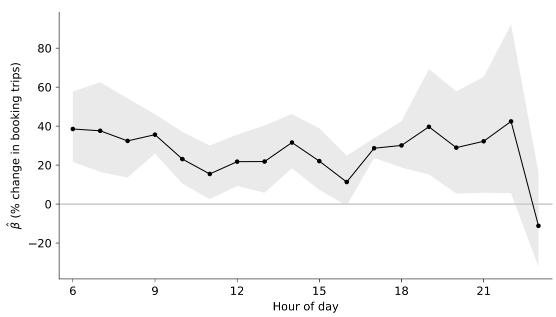

(b) *Booking trips.*

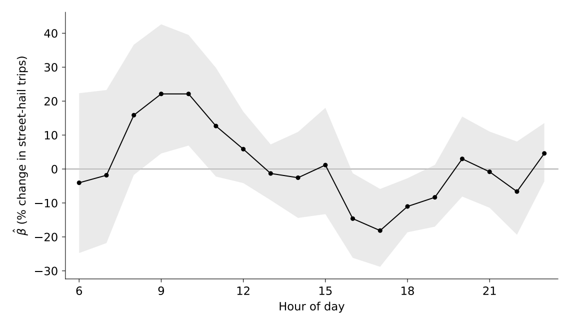

(c) *Street-hail trips.*

Figure 4: *Heterogeneity in trip responses by hour of day and dispatch channel. Trip data are collapsed by (planning area, pickup hour, date) cells. Within a pickup hour, disrupted cells are compared with undisrupted cells to obtain a percentage difference in the number of trips, relative to the undisrupted baseline, in the event of a train disruption. 95% confidence intervals are shaded in gray. Trip data are from a large taxi company in Singapore. Disruption data are from* `mrtdown.org`.

the $\ln(1+Y)$ transformation and clarifies whether disruptions activate new service in cells that would otherwise have zero trips.

Tables A.4a–A.6b in Appendix C report the results. The booking extensive margin is small and insignificant (-0.007): disruptions do not open up new booking service in cells that would otherwise have zero booking trips. Instead, the booking response operates entirely through the intensive margin (0.205, significant at the 1% level): booking dispatch scales up trip volume in cells that already have baseline booking activity. The street-hail extensive margin is negative and significant (-0.008, 1% level): disruptions actually reduce the probability of any street-hail trip in a cell. Neither channel expands on the extensive margin; the divergence between booking and street-hail is

ln(1 + trips)	(1) Baseline	(2) All FE	(3) Unexpected
Disrupted	1.3788***	0.0659	0.1126
	(0.1316)	(0.0773)	(0.1317)
Unexpected			-0.1091
			(0.1382)
Observations	287,616	287,616	287,616
Adjusted R^2	0.001	0.934	0.934
Fixed Effects:			
Hour FE		✓	✓
Day-of-Week FE		✓	✓
Planning Area FE		✓	✓

Table 4: *Log-linear OLS estimates of the effect of transit disruptions on ridehailing trips. "Unexpected" excludes planned disruptions. The dependent variable is ln(1 + trips) of ridehailing trips originating from a given planning area. Each observation is a (planning area, hour, date) cell. Standard errors are two-way clustered at the planning area and disruption event level. Trip data are from a 10% random sample of ridehailing trips in the Singaporean market, April to October 2018. Disruption data are from* `mrtdown.org`*. Asterisks ($*, **, ***$) indicate statistical significance at the 10%, 5% and 1% level respectively.*

entirely intensive. We apply the same decomposition to ridehailing trips (Table A.8): the extensive margin is small and insignificant (-0.004), while the intensive margin is 0.105 (significant at the 5% level), indicating that ridehailing's response operates entirely through increased trips in cells that already have positive ridehailing activity.

PPML specification. As a check on the log-linear OLS specification, we re-estimate the model using Poisson pseudo-maximum likelihood (PPML, Correia, Guimarães, and Zylkin 2020; Silva and Tenreyro 2006),

$$\mathbb{E}\left[n_{iht}^{(m)} \,\middle|\, \mathbf{X}\right] = \exp\left(\beta^{(m)} \times \mathbb{1}\{\text{Disrupted}\}_{iht} + \gamma_i^{(m)} + \gamma_h^{(m)} + \gamma_t^{(m)}\right), \tag{3}$$

which accommodates zero-valued cells without requiring a log transformation (Correia, Guimarães, and Zylkin 2020; Silva and Tenreyro 2006). PPML delivers consistent semi-elasticity estimates under correct specification of the conditional mean, without requiring the data to follow a Poisson distribution (Dingel and Tintelnot forthcoming; Silva and Tenreyro 2006). However, PPML weights more heavily high-count cells, which may overstate the influence of downtown areas with many alternative transportation options.

Tables A.2–A.3b in Appendix C report the results. The qualitative pattern is similar: booking trips show a larger and more precisely estimated point estimate than street-hail trips. Under PPML, the booking coefficient is 0.158 (SE 0.032), highly significant ($p < 0.01$), corroborating the OLS result and addressing the concern that the OLS booking effect could be driven by the log transformation in low-count cells. For ridehailing, PPML yields a significant estimate of about 9% (β = 0.086; Table A.7), stronger than the imprecise OLS point estimate. The stronger PPML estimate suggests that the ridehailing response is concentrated in high-activity cells, consistent with PPML's greater weight on high-count observations and with the intensive-margin pattern in the Chen-Roth decomposition (Table A.8).

Dose-response by distance from disrupted station. To approximate a dose-response curve, we draw concentric 100-meter rings emanating from each disrupted station and assign each taxi trip to its nearest disrupted station within 1 km. We estimate heterogeneous treatment effects by ring distance, comparing each station-ring-hour cell to itself on non-disrupted hours. The estimating equation is

$$
\begin{aligned}
y_{srht} = \delta \times \mathbb{1}\{\text{Disrupted}\}_{sht} + \sum_{r=0}^{4} \beta_r \times \mathbb{1}\{\text{Disrupted}\}_{sht} \times \mathbb{1}\{\text{Ring} = r\} \\
+ \gamma_{sr} + \gamma_h + \gamma_t + \varepsilon_{srht},
\end{aligned}
\tag{4}
$$

where s indexes the nearest disrupted station, $r \in \{0, \ldots, 5\}$ indexes the distance ring (0–100m, ..., 400–500m, 500m–1km), and h and t index hour-of-day and day-of-week. The omitted category is $r = 5$ (500m–1km), so δ captures the disruption effect at the control ring and each β_r captures the differential effect at ring r relative to it. Station-by-ring fixed effects γ_{sr} absorb permanent differences in baseline trip intensity across rings, so the total effect at each ring—$\delta + \beta_r$ for treated rings, δ for the control ring—represents the change in trips during disrupted hours relative to non-disrupted hours within the same station-ring cell. We report these total effects in Figure A.4. Standard errors are clustered at the station level.

Figure A.4 in Appendix C displays the results. For all trips, the effect peaks at the station doorstep (0–100m: +0.18 log points), declines through 100–500m (range +0.04 to +0.10), and rises again at 500–1000m (+0.09). This bimodal pattern reflects the composition of two sharply divergent channels. Street-hail trips spike at 0–100m (+0.19, significant at the 1% level) and at 100–200m (+0.09, significant at the 1% level), then attenuate, with a small, weakly significant negative estimate at 400–500m (-0.029, significant at the 10% level) before settling near zero at 500–1000m. Booking trips show the opposite spatial gradient: near zero at 0–100m, rising from 100m outward, and peaking at 500–1000m (+0.23, significant at the 1% level). This result is consistent with the booking channel connecting riders and drivers who are not co-located: stranded commuters away from the congested station area can request bookings that are broadcast to drivers at moderate distance. We note that this finding can be reconciled with our aggregate results: rings farther from the disruption epicenter have a larger area, implying that the increase in bookings at moderate range outweighs the increase in street-hail near the epicenter.

The decomposition is stark: street-hail requires physical proximity to the disruption epicenter and is the most spatially concentrated channel, while booking operates at distance and generates the largest effects well away from the station. The booking-versus-street-hail decomposition reinforces the information-provision interpretation: booking dispatch connects riders and drivers who are not co-located, while street-hail pickup requires co-location.

We caveat that many city-center stations are within 500 meters of each other, creating potential interference across rings; and that we cannot distinguish net new taxi demand from spatial reallocation of existing demand within the 1 km radius.

Controlling for local idle taxi supply. A natural concern is that the booking-street-hail gap reflects differences in the local availability of taxis rather than the organizational form of matching. We construct a "search pool" variable measuring the total taxi-minutes spent searching or waiting in a planning area during each cell-hour, and include it as a control (Column 5 of Tables 2 and 3, and the corresponding PPML and Chen-Roth tables in Appendix C). The disrupted coefficient for bookings is stable after conditioning on the search pool (0.228 to 0.210), suggesting that the booking advantage is not driven by local supply availability. Notably, the sign of the search pool coefficient flips across channels: it is negative for bookings and positive for street-hail. More idle taxis nearby increases the probability of a physical encounter (helping street-hail) but reduces reliance on booking dispatch (substituting away from bookings when supply is locally abundant). This pattern is consistent with the theoretical distinction between search-based and dispatch-based matching technologies.

We also examine whether disruptions increase the local idle taxi supply itself. Table A.9 in the Appendix regresses the search pool measure on the disruption indicator. After conditioning on fixed effects, disruptions reduce local search minutes (OLS estimate -0.16 log points, significant at the 1% level; the analogous PPML estimate is also negative but imprecise), indicating that the trip increases we document coexist with—and are likely produced by—faster matching of idle taxis rather than passive accumulation of idle taxis near disrupted stations.

Vehicle-level evidence. Following Agarwal et al. (2025) (Equation 9), we run a within-driver regression of $\ln(1 + \text{fare})$ on the disruption indicator at the (vehicle $\times$ date $\times$ hour) level. Our preferred specification (Spec D) restricts to the first hour of unexpected disruptions, absorbs vehicle, pre-hour planning-area, date, and hour-of-day fixed effects, and two-way clusters standard errors at the vehicle $\times$ event level. The aggregate coefficient is -0.009 (SE 0.007, insignificant); the channel decomposition gives booking $+0.003$ (n.s.) and street-hail -0.012 (n.s.). See Tables A.10 and A.11 in the Appendix. The street-hail point estimate remains negative, qualitatively consistent with an information/search-friction interpretation, though no longer statistically distinguishable from zero in this specification. These estimates are consistent with greater matching intensity in disrupted areas rather than higher per-hour fares. We do not claim to decompose matching intensity from the fare response, which would require matched acceptance, availability, and on-duty-hour measures beyond our current sample.

7 Discussion

7.1 Interpreting the booking advantage

Our results support Hypothesis 1 (Section 2.2): booking dispatch increases service provision during transit disruptions substantially more than street-hail. We interpret this gap as evidence that customer-specific information changes how decentralized drivers respond to localized demand shocks. When a rail disruption concentrates excess demand near specific stations, cruising drivers observe congestion, slower speeds, longer pickup times, and local competition for fares, but they do not observe a specific passenger request before choosing where to search. A cruising driver without trip-specific information may rationally avoid the affected area because entering it means searching under heightened uncertainty about whether a feasible match will form.

The booking channel mitigates this search friction by aggregating realized customer requests and broadcasting them to nearby drivers. Drivers observe the pickup location and destination before deciding whether to accept, although the fare remains metered and is not visible ex ante. A driver who accepts a booking request from a disrupted area therefore does so with more information about the match than a street-hail driver, who must first enter or remain near the disrupted area before learning whether a passenger and feasible destination will appear. In the language of Grossman (1976), the booking channel aggregates demand information that is otherwise dispersed across passengers. More directly, it makes realized demand visible and actionable before the driver commits to the match, improving the allocation of drivers to passengers relative to decentralized search alone. Our within-firm design is critical for this interpretation: because street-hail and booking trips originate from the same fleet, with a common driver pool, lease structure, fare schedule, and firm culture, the differential response identifies the joint effect of request aggregation, destination-visible broadcast, and the driver acceptance margin, rather than confounding firm-level characteristics.

Our setting differs from the information-provision experiment studied by Agarwal et al. (2025) at Singapore's Changi Airport in three ways. First, *signal content*: our booking channel broadcasts realized individual-trip demand with destinations visible to drivers before they accept; Agarwal et al.'s signboards broadcast aggregate supply-side queue counts and a noisy demand proxy (flight arrivals), with no individual-trip content. Second, *acceptance margin*: in our setting, accepting a booking offer creates a match; Agarwal's signboards have no comparable acceptance margin—drivers see aggregate cues and decide whether to cruise to a particular terminal. Third, *scope*: our design covers the entire city across all planning areas and hours, versus four terminals at one airport. Thus, our booking channel bundles request aggregation, destination-visible broadcast, and driver acceptance. Our design identifies the joint contribution of these features to the booking premium; disentangling the separate role of each margin remains an open question.

Because booking and street-hail draw on the same fleet, a natural concern is that the booking-street-hail gap reflects within-fleet substitution (drivers switching from cruising to dispatched pickups) rather than a genuine channel effect. Three features of our results speak against this interpretation. First, *both* channels show positive point estimates in response to disruptions (bookings about +26%, street-hail about +3%, though the latter is not statistically significant). If the booking increase came at the expense of street-hail, we would observe street-hail trips declining in disrupted areas. The sign pattern is instead consistent with the booking channel capturing a larger share of shock-induced demand, not with reallocation from one channel to the other. Second, the aggregate point estimate is positive (roughly 6%) and marginally significant, consistent with net expansion of service provision during disruptions rather than zero-sum reshuffling of a fixed trip pie across channels. Third, when we directly control for the local idle taxi supply (the "search pool" of taxi-minutes spent searching or waiting in the planning area), the booking disruption coefficient is essentially unchanged (0.228 to 0.210; Section 6.5), ruling out the possibility that the booking response simply reflects more available taxis in the area.

We therefore interpret our within-firm comparison as identifying the joint value of request information and the acceptance margin: the booking channel more effectively converts localized demand shocks into completed trips than street-hail. To the extent that some within-fleet reallocation occurs, we view this as part of the mechanism through which the channel operates, not as a confound. The booking channel makes specific unserved demand visible and actionable to drivers who choose to accept; this is precisely what distinguishes it from street-hail.

A notable feature of our results is that the marginally significant aggregate effect (roughly 6%) masks a sharp divergence across channels: booking trips rise significantly, by about 26%, while street-hail trips show a small and statistically insignificant 3% increase. The marginally significant aggregate motivates the decomposition: without separating by channel, an analyst might conclude that taxis respond only weakly to transit disruptions, a conclusion that obscures the channel-level dynamics we document.

7.2 Boundary conditions: When does the booking advantage attenuate?

Our findings show that the booking channel is especially effective for spatially concentrated, time-sensitive demand shocks when street-hail search is costly and local signals are incomplete. But the theoretical framework in Section 2 implies that this advantage is contingent: street-hail should perform better when local search provides sufficiently informative signals or when the booking channel's request information is unavailable.

Several boundary conditions merit discussion. First, for *distributed shocks* that affect many locations simultaneously rather than concentrating demand near specific stations, the booking channel's information advantage may be less valuable because there is no focal point of unserved demand. Drivers' local knowledge of neighborhood-level conditions may instead be more useful. Second, *long-duration and long-run maintenance disruptions* could erode the booking advantage as cruising drivers accumulate information about the disruption through local cues and other channels, effectively narrowing the information gap, though we do not directly test this. Third, in settings with *weak dispatch infrastructure* or low platform penetration, the request-information channel may simply be unavailable. In those settings, street-hail may remain the primary mode of matching even when localized demand shocks create substantial search frictions.

Our own results provide direct evidence on one boundary condition. The booking advantage is clearest during evening peak hours during disruptions (Hypothesis 2, Section 2.3), when road congestion raises the cost of uninformed cruising and makes local street-hail signals less informative. By contrast, off-peak periods and highly visible service outages may narrow the information gap between booking and street-hail. This interpretation is consistent with the weaker heterogeneity patterns outside the evening peak, but we view these cases as suggestive boundary conditions rather than separately identified mechanisms.

7.3 Matching versus pricing

Our design identifies the joint effect of the booking channel's request information and acceptance margin on service provision during disruptions. A natural follow-up question is whether flexible pricing amplifies the booking advantage. Ridehailing platforms combine app-based dispatch with surge pricing, while taxi bookings operate under regulated fixed fares. If the pricing channel were quantitatively important, we might expect $\hat{\beta}^{(\text{ridehailing})} > \hat{\beta}^{(\text{booking})}$.

In practice, the point estimates go in the opposite direction: $\hat{\beta}^{(\text{booking})} = 0.228$ versus $\hat{\beta}^{(\text{ridehailing})} = 0.066$. However, we stress that this comparison cannot isolate the pricing margin. The two estimates come from different firms, different time periods (2016–2018 versus 2018), different data samples (full fleet data versus a 10% trip sample), and different market conditions. Any of these confounds could account for the ordering. The comparison is therefore best interpreted as consistent with information- and acceptance-based matching being an important channel; it does not rule out a separate role for pricing.

Cleanly separating matching from pricing would require within-platform variation in dispatch authority or pricing rules, for instance a platform experiment that randomizes surge pricing across otherwise identical markets. We regard the matching-versus-pricing decomposition as an important direction for future work.

7.4 Connections to organizational resilience theory

Our findings connect to and extend several literatures. First, we contribute to research on *organizational resilience* (e.g., Duchek 2020; Sutcliffe and Vogus 2003; Van der Vegt et al. 2015). This literature has documented firm-level adaptation to crises largely through qualitative case studies and surveys, identifying adaptive capacity as a hallmark of resilient organizations but leaving open which specific design choices generate it. We provide causal, high-frequency evidence that a concrete design choice—embedding an information-aggregation and broadcast channel within a firm—produces measurable resilience in the form of increased service provision during demand shocks. Our quasi-experimental design resolves a standard confound in this literature by holding the contractual environment fixed (driver pool, lease structure, fare schedule, firm culture) while varying the information regime and pre-acceptance choice.

Second, we speak to the platform design and governance literature (e.g., Boudreau 2010; Gawer 2021; Parker, Van Alstyne, and Choudary 2016). Prior work emphasizes that platforms shape market outcomes not only through prices, but also through rules, interfaces, and allocation technologies. Our setting isolates one such design margin: whether a driver observes a specific booking request or searches independently for street-hail passengers. Because booking and street-hail trips coexist within the same taxi system, we compare responses across channels while holding fixed many features of the market environment, including regulated fares and the broader institutional setting. This within-market comparison frames dispatch as an information architecture: the platform aggregates dispersed customer requests and transmits selected signals to drivers, who retain control over search and acceptance decisions. The evidence shows that this architecture matters most when demand changes faster than drivers can learn from local search.

Third, we contribute to the *gig economy* literature on driver labor supply and search (e.g., Buchholz, Shum, and Xu forthcoming; Camerer et al. 1997; Chen et al. 2019; Farber 2015). This work documents how drivers respond to fares and demand fluctuations under normal market conditions. We show that the information regime generates resilience even when the contractual environment (fares, lease) is held fixed: request aggregation and broadcast is a lever for adaptive capacity distinct from price incentives. The information friction we document is not a failure of incentives (street-hail fares do not change during disruptions) but a failure of real-time customer-specific demand information: cruising drivers lack this information before choosing where to search, while the booking channel provides it before the driver accepts.

Fourth, we provide the first quasi-experimental evidence on how dispatch mode shapes the extent to which point-to-point transportation functions as a "safety valve" during public transit failures (Almagro et al. 2024). Our results show that this complementarity depends critically on organizational structure: booking-dispatched rides increase substantially during disruptions, while street-hail shows a much smaller response. The policy implication is that the resilience of a city's transportation network depends not only on the *availability* of backup modes but on how those modes are *organized*.

7.5 Limitations

Several limitations of our study suggest directions for future research. First, our quasi-experimental design identifies the *direction* and *heterogeneity* of channel-level responses but does not quantify the welfare gains from the booking channel's request information and acceptance margin. A structural approach that models driver search and matching outcomes under alternative information regimes could quantify these margins and evaluate counterfactual designs.

Second, our evidence comes from a single city with a particular regulatory and institutional environment. Singapore is dense, well-regulated, and has high platform penetration. Whether the booking advantage generalizes to settings with different labor market institutions, fare structures, or spatial configurations remains an open question, particularly for cities in developing countries where search frictions may be more severe but dispatch infrastructure is weaker.

Third, our within-firm comparison uses taxi booking data from one company, while our external-validity check draws on ridehailing data from a different provider observed in a different time period. Although both dispatch-based channels show a positive response to disruptions, the cross-firm comparison cannot isolate the pricing margin (Section 7.3). Future work with matched data from multiple platforms operating simultaneously would strengthen the decomposition of matching versus pricing channels.

8 Conclusion

This paper provides causal evidence that a booking channel that aggregates and broadcasts customer-specific trip information produces measurable service-provision effects consistent with organizational resilience to localized demand shocks. Using quasi-experimental variation from rail disruptions in Singapore, we show that booking trips rise by about 26% during disruptions, roughly an order of magnitude larger than the small and statistically insignificant street-hail response of about 3%. We interpret this divergence as evidence that customer-specific request information changes how decentralized drivers respond to localized demand shocks. Cruising drivers must choose where to search without observing a specific passenger request or destination, while booking dispatch makes realized requests, including pickup location and destination, visible to nearby drivers before they decide whether to accept. The within-firm design holds the contractual environment fixed (driver pool, lease structure, fare schedule, firm culture) and identifies the joint effect of request aggregation, destination-visible broadcast, and the driver acceptance margin. A benchmark using ride-hailing data from a separate platform and time period provides suggestive external evidence that another digitally mediated matching technology responds in the same direction.

Our findings carry direct implications for practitioners. For *platform operators*, timely request information and dispatch interfaces are important levers for service resilience; showing drivers trip details before acceptance is complementary to pricing algorithms and other incentive tools. For *transit agencies and regulators*, the organization of backup transportation matters: policies that facilitate dispatch integration, such as open booking APIs or integration between transit operators and ride-hailing platforms, can improve a network's absorptive capacity. For *organizational design* more broadly, our results sharpen the contingency theory: the value of platform-mediated information provision is highest when local search signals are least informative about rapidly shifting demand—precisely the conditions created by unexpected, spatially concentrated shocks.

The information/search-friction logic extends well beyond urban transportation. In logistics, emergency response, and healthcare staffing, dispersed workers must be rapidly reallocated when demand shifts unexpectedly. In each of these settings, autonomous agents relying only on local signals may underrespond or misallocate effort relative to what aggregated demand information would support. As platforms that aggregate and broadcast demand information become more widespread, the organizational choice between information-rich and information-poor allocation channels will increasingly shape resilience across a wide range of service markets.

References

Agarwal, Sumit, Shih-fen Cheng, Jussi Keppo, Long Wang, and Yang Yang (2025). "Information Provision and Search Frictions: Evidence from the Taxi Industry in Singapore". In: *Review of Economics and Statistics*, pp. 1–47.

Almagro, Milena, Felipe Barbieri, Juan Camilo Castillo, Nathaniel G Hickok, and Tobias Salz (2024). *Optimal urban transportation policy: Evidence from Chicago*. Tech. rep. National Bureau of Economic Research.

Alonso, Ricardo, Wouter Dessein, and Niko Matouschek (2008). "When Does Coordination Require Centralization?" In: *American Economic Review* 98.1, pp. 145–179.

Boudreau, Kevin (2010). "Open platform strategies and innovation: Granting access vs. devolving control". In: *Management Science* 56.10, pp. 1849–1872. DOI: 10.1287/mnsc.1100.1215.

Buchholz, Nicholas (2022). "Spatial equilibrium, search frictions, and dynamic efficiency in the taxi industry". In: *The Review of Economic Studies* 89.2, pp. 556–591.

Buchholz, Nicholas, Matthew Shum, and Haiqing Xu (forthcoming). “Rethinking reference dependence: Wage dynamics and optimal taxi labor supply”. In: *Journal of Political Economy*.

Camerer, Colin, Linda Babcock, George Loewenstein, and Richard Thaler (1997). “Labor supply of New York City cabdrivers: One day at a time”. In: *The Quarterly Journal of Economics* 112.2, pp. 407–441.

Castillo, Juan Camilo (2025). “Who benefits from surge pricing?” In: *Econometrica* 93.5, pp. 1811–1854.

Castillo, Juan Camilo, Dan Knoepfle, and E Glen Weyl (2025). “Matching and pricing in ride hailing: Wild goose chases and how to solve them”. In: *Management Science* 71.5, pp. 4377–4395.

Chen, Jiafeng and Jonathan Roth (2024). “Logs with Zeros? Some Problems and Solutions”. In: *The Quarterly Journal of Economics* 139.2, pp. 891–936.

Chen, M Keith, Peter E Rossi, Judith A Chevalier, and Emily Oehlsen (2019). “The value of flexible work: Evidence from Uber drivers”. In: *Journal of Political Economy* 127.6, pp. 2735–2794.

Correia, Sergio, Paulo Guimarães, and Thomas Zylkin (2020). “Fast Poisson estimation with high-dimensional fixed effects”. In: *The Stata Journal* 20.1, pp. 95–115.

Csaszar, Felipe A (2012). “Organizational structure as a determinant of performance: Evidence from mutual funds”. In: *Strategic Management Journal* 33.6, pp. 611–632. DOI: 10.1002/smj.1969.

Cullen, Zoë B and Chiara Farronato (2021). “Outsourcing tasks online: Matching supply and demand on peer-to-peer internet platforms”. Working paper.

Dessein, Wouter (2002). “Authority and Communication in Organizations”. In: *The Review of Economic Studies* 69.4, pp. 811–838.

Dingel, Jonathan I and Felix Tintelnot (forthcoming). “Spatial Economics for Granular Settings”. In: *Econometrica*.

Duchek, Stephanie (2020). “Organizational resilience: a capability-based conceptualization”. In: *Business Research* 13.1, pp. 215–246. DOI: 10.1007/s40685-019-0085-7.

Farber, Henry S (2015). “Why you can’t find a taxi in the rain and other labor supply lessons from cab drivers”. In: *The Quarterly Journal of Economics* 130.4, pp. 1975–2026.

Gawer, Annabelle (2021). “Digital platforms’ boundaries: The interplay of firm scope, platform sides, and digital interfaces”. In: *Long Range Planning* 54.5, p. 102045. DOI: 10.1016/j.lrp.2020.102045.

Grossman, Sanford J (1976). “Information and competitive price systems”. In: *American Economic Review* 66.2, pp. 246–253.

Hayek, Friedrich A (1945). “The Use of Knowledge in Society”. In: *American Economic Review* 35.4, pp. 519–530.

Levinthal, Daniel A (1997). “Adaptation on rugged landscapes”. In: *Management Science* 43.7, pp. 934–950. DOI: 10.1287/mnsc.43.7.934.

Parker, Geoffrey G, Marshall W Van Alstyne, and Sangeet Paul Choudary (2016). *Platform Revolution: How Networked Markets Are Transforming the Economy—and How to Make Them Work for You*. New York: W. W. Norton & Company.

Png, Ivan PL and Song Ivy Wang (2025). “Gambler’s Fallacy Biases Sequential Decision-Making: Evidence from Taxi Drivers”. In.

Rawley, Evan and Timothy S. Simcoe (2013). “Information Technology, Productivity, and Asset Ownership: Evidence from Taxicab Fleets”. In: *Organization Science* 24.3, pp. 831–845. DOI: 10.1287/orsc.1120.0758.

Siggelkow, Nicolaj and Jan W Rivkin (2005). “Speed and search: Designing organizations for turbulence and complexity”. In: *Organization Science* 16.2, pp. 101–122. DOI: 10.1287/orsc.1050.0116.

Silva, J M C Santos and Silvana Tenreyro (2006). “The Log of Gravity”. In: *The Review of Economics and Statistics* 88.4, pp. 641–658.

Sutcliffe, Kathleen M and Timothy J Vogus (2003). “Organizing for resilience”. In: *Positive Organizational Scholarship: Foundations of a New Discipline*. Ed. by Kim S Cameron, Jane E Dutton, and Robert E Quinn. San Francisco: Berrett-Koehler, pp. 94–110.

Van der Vegt, Gerben S, Peter Essens, Margareta Wahlström, and Gerard George (2015). “Managing risk and resilience”. In: *Academy of Management Journal* 58.4, pp. 971–980. DOI: 10.5465/amj.2015.4004.

Zhang, Xueli, Wei Miao, Junhong Chu, and Ivan Png (2026). “The design of centralized matching systems on two-sided platforms: Evidence from the ride-hailing market”. In: *Marketing Science*. Forthcoming.

Appendix

A Data Construction

A.1 Taxi data

We use processed taxi microdata from Png and Wang (2025), which spans two distinct time periods: 1 December 2016 till 31 May 2017 and 1 August 2017 to 31 March 2018. The combined data set includes 91.9 million trips, with “almost all drivers were male, somewhat older than 55, and [having] worked over 13 years with the company.” Other day-level and driver-level summary statistics have been redacted upon request by the company.

The unique identifier of each vehicle is `vehicle_id` and that of each driver is `driver_id`. The data comprise:

1. **Activity data**, which comprises vehicle positions (latitude/longitude) and statuses (`new_veh_status`) whenever the vehicle's status changes, or at a time frequency of (approximately) five minutes without a status change. Each row corresponds to one (driver, vehicle, activity, interval) 4-tuple; calendar time is embedded in the interval variable.
   - Within each driver, each continuous span without a status change is labeled by a unique `activity_id`, and each update within an activity (including the start of said activity) is identified by an `interval_id`.
   - A field `new_distance_km` returns the distance covered by a taxi during the interval.
2. **Job data**, which comprises all jobs served by each vehicle over time. Each row is identified by a `job_no` attached to a certain `job_driver_id`. We see the job status (cancelled, completed, failed, no show), start and end positions and times, the type of the job (booking/street hail), distance travelled, time taken, and fare paid.

The vehicle statuses are in one of the following categories:

- *Idle.* A driver is idle if the driver is on the road and available to serve trips. A flag of **SEARCH** is applied if he is outside 50 meters of a taxi stand, and has speed above 10 kilometers an hour; else a flag of **WAIT** is applied.
- *Serving a passenger.* **POB** indicates a passenger is on board and the taxi meter is running; **PAYMENT** indicates that the passenger is making payment and the taxi meter is paused. **STC** implies the taxi is soon to clear the current job and will be ready for new bookings (this last flag is not used in our analysis).
- *Serving a booking.* **ONCALL** implies the taxi has been assigned to a passenger and is en route to picking them up; **ARRIVED** means the driver has arrived at his pick-up point and is waiting for the passenger to board. **NOSHOW** is a rare flag indicating that the passenger has cancelled/defaulted on their booking, or will do so soon.

- *Inactivity.* **OFFLINE** implies the driver is inactive and is logged off from their mobile data terminal. **BUSY** indicates that the driver is unable to serve trips for reasons other than taking a break (e.g., changing shifts or transporting family). **BREAK** indicates the driver has taken a break at a particular location.

A.2 Ridehailing data

We obtained a 10% sample of ridehailing trips in the Singaporean market from a confidential data provider. These trips span April 2018 to October 2018. Each observation is a completed trip. The data contain the following fields, subject to confidentiality restrictions that prevent us from disclosing the identity of the data provider or sharing the raw microdata.

Trip identifiers. Each row is indexed by a unique trip identifier. We observe the company, a masked vehicle serial number, and a masked driver serial number. A binary indicator records whether the trip was dispatched through a third-party booking platform.

Pickup and dropoff. For each trip, we observe the pickup and dropoff timestamps (in seconds since 1 January 1970), geographic coordinates (longitude and latitude, in arc-minutes), nearest street address, six-digit postal code, and planning area. Planning areas are the spatial units used in our event study design.

Fare structure. Total fares are decomposed into a flag-down component, a distance-based component, and several surcharges: Electronic Road Pricing (ERP), location-based surcharges, peak-hour surcharges, midnight surcharges, and a booking fee. A separate field records the fare for fixed-fare trips. Trip distance is recorded in kilometers.

A.3 Train disruption data, `mrtdown.org`

We obtain train disruption data from the GitHub repository (`https://github.com/foldaway/mrtdown-data`) powering the crowdsourced train disruption listings website, `mrtdown.org`. Each disruption is labelled an "issue" and is stored as a JSON file with title of the form "`YYYY-MM-DD-(plaintext-issue-description)`". We see the start and end dates and times of the disruption, the type of the disruption (e.g., power fault, track fault, train fault), as well as a list of updates from when the disruption was first announced on traditional news media or social media, till when the disruption was cleared. Finally, we see a list of stations affected by the disruption.

B Supplementary Analysis: How do taxi drivers search?

Even though our design-based analysis on consummated trips suggests that taxi supply may not respond adequately to rail disruptions, a deeper analysis of driver search could reveal why.

In Figure A.1, we plot driver dwell patterns by planning area for 1 August 2017 (a Tuesday), evaluated on a 2/35 driver-subsample chunk of the data (approximately 700 drivers). Broadly speaking, we observe a strong within-day rhythm across the city: in most planning areas, more drivers are active overnight and in the early morning, drifting downwards through the day with a common low around late afternoon to evening, then rebound again towards midnight.

Three descriptive facts matter for our subsequent analysis. First, considering the downtown core planning area, we find that many drivers are active there shortly after midnight, but not shortly before. By inspecting taxi fare structures, we observe that midnight is exactly the time when a peak hour surcharge of 25% increases to a midnight surcharge of 50%. In our conversations with taxi drivers, we note that a substantial fraction "work the night shift": they come online shortly before or after midnight and work till 6am, at which point the midnight surcharge transitions “back” to the peak hour surcharge of 25%. Thus, it seems important for us to account for a driver’s decision of when to start and end his shift.

Second, there is large spatial heterogeneity in taxi dwell levels. Some areas seem well-served by taxis, for instance, Bukit Merah, Queenstown, and the Downtown Core; while sparsely populated areas see very few taxis, including the water catchments and Lim Chu Kang.

Third, taxis seem to relocate to different sets of planning areas at different times of day. For instance, Changi, where the city’s civilian airport is, appears to have a daytime hump. On the other hand, outer ring residential towns, for instance Woodlands, Yishun, and Sengkang, seem to have prominent midday peaks, suggestive of lunch travel or post-lunch change of work location.

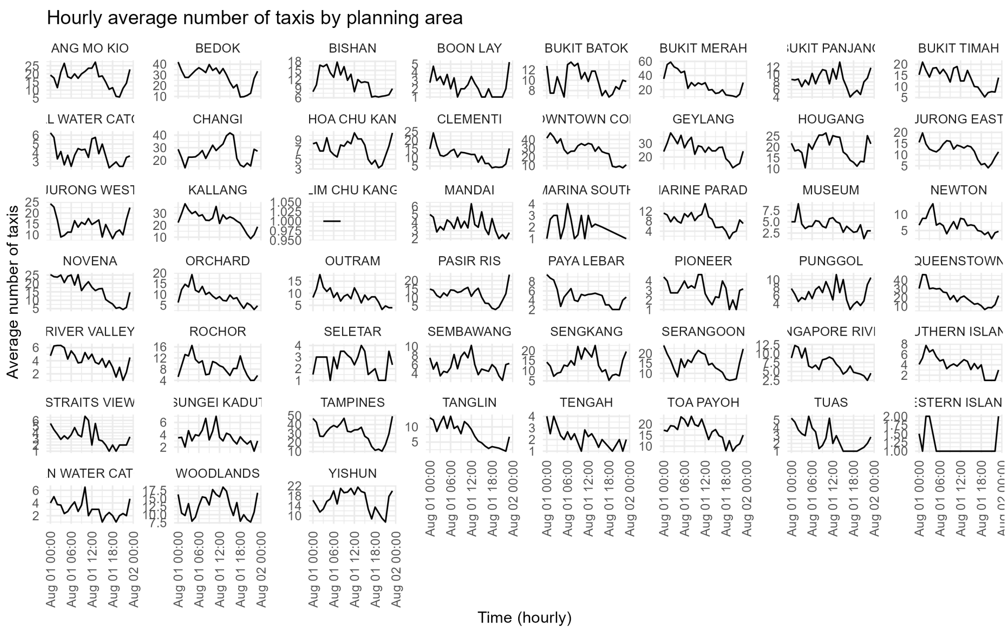

Figure A.1: *Driver dwell patterns on 1 August 2017, evaluated on a 2/35 driver-subsample chunk (approximately 700 drivers). Activity data are collapsed by (planning area, hour) cells, and are from a large taxi company in Singapore.*

C Supplementary Tables and Figures

Variable	T1 (Low)	T2 (Medium)	T3 (High)	Norm. Diff (T3−T2)
Population (thousands)	4.7	125.6	99.7	-0.29
	(12.8)	(87.3)	(93.2)	
Median HH income (SGD)	16716.3	10503.4	9991.5	-0.14
	(5246.4)	(4069.5)	(3002.5)	
MRT stations	0.2	3.7	4.9	0.30
	(0.7)	(2.7)	(4.6)	
Land area (km^2)	19.3	13.3	7.5	-1.00
	(22.2)	(5.2)	(6.2)	
Baseline trips/cell	32.6	316.1	312.8	-0.02
	(50.7)	(187.4)	(244.3)	
Mean fare (SGD)	18.6	14.6	14.1	-0.33
	(3.2)	(1.5)	(1.2)	
Search duration (min/cell)	531.7	4539.3	3654.6	-0.37
	(501.8)	(2679.1)	(2114.1)	
Planning areas	19	19	17	

Table A.1: *Planning area characteristics by disruption exposure tercile. Planning areas are ranked by the fraction of (planning area, hour, date) cell-hours in which at least one MRT disruption occurs, then split into terciles (T1 = lowest exposure, T3 = highest). Population and income data are from the 2020 Census of Population; income is the median monthly household income from work, interpolated from grouped data and excluding households with no employed person. MRT station counts are from the station–subzone crosswalk. Land area is the sum of subzone areas from the URA Master Plan 2014 planning area boundaries. Baseline trips per cell is the mean trip count across non-disrupted (planning area, hour, date) cells, based on our taxi data. Mean fare is the average fare (SGD) across non-disrupted cells with at least one trip. Search duration is the mean total taxi-minutes spent in SEARCH or WAIT status per cell-hour in non-disrupted cells. Normalized differences between T3 (highest exposure) and T2 (medium exposure) are computed as* $(\bar{x}_{T3} - \bar{x}_{T2})/\sqrt{(s_{T3}^2 + s_{T2}^2)/2}$. *Standard deviations in parentheses.*

Trips (Poisson)	(1) Baseline	(2) All FE	(3) Unexpected	(4) No Service	(5) Search Pool
Disrupted	0.5152***	0.0359	0.0464	0.0365	0.0450**
	(0.1435)	(0.0281)	(0.0535)	(0.0305)	(0.0207)
Unexpected			-0.0122		
			(0.0669)		
No Service				-0.0070	
				(0.0670)	
Search Minutes					0.0000***
					(0.0000)
Observations	571,200	571,200	571,200	571,200	571,200
Pseudo R^2	0.001	0.892	0.892	0.892	0.898
Fixed Effects:					
Hour FE		✓	✓	✓	✓
Day-of-Week FE		✓	✓	✓	✓
Planning Area FE		✓	✓	✓	✓

Table A.2: *PPML estimates of the effect of transit disruptions on aggregate taxi trips. "Unexpected" excludes planned disruptions. "No Service" indicates that at least one rail service is completely non-operational along some stretch of stations. "Search Pool" controls for the total taxi-minutes spent searching or waiting in the planning area during that cell-hour. The dependent variable is the count of trips originating from a given planning area. Each observation is a (planning area, hour, date) cell. Standard errors are two-way clustered at the planning area and disruption event level. Trip data are from a large taxi company in Singapore, covering December 2016 to March 2018. Disruption data are from `mrtdown.org`. Asterisks ($*, **, ***$) indicate statistical significance at the 10%, 5% and 1% level respectively.*

Trips (Poisson)	(1) Baseline	(2) All FE	(3) Unexpected	(4) No Service	(5) Search Pool
Disrupted	0.7168***	0.1584***	0.1836	0.1700***	0.1355***
	(0.0083)	(0.0322)	(0.1355)	(0.0335)	(0.0288)
Unexpected			-0.0283		
			(0.1421)		
No Service				-0.1462	
				(0.1228)	
Search Minutes					-0.0001***
					(0.0000)
Observations	571,200	571,200	571,200	571,200	571,200
Pseudo R^2	0.002	0.826	0.826	0.826	0.847
Fixed Effects:					
Hour FE		✓	✓	✓	✓
Day-of-Week FE		✓	✓	✓	✓
Planning Area FE		✓	✓	✓	✓

(a) *Trips fulfilled by booking dispatch (Poisson).*

Trips (Poisson)	(1) Baseline	(2) All FE	(3) Unexpected	(4) No Service	(5) Search Pool
Disrupted	0.4428***	-0.0106	0.0162	-0.0133	0.0054
	(0.1544)	(0.0413)	(0.0555)	(0.0445)	(0.0259)
Unexpected			-0.0316		
			(0.0677)		
No Service				0.0366	
				(0.0642)	
Search Minutes					0.0001***
					(0.0000)
Observations	571,200	571,200	571,200	571,200	571,200
Pseudo R^2	0.001	0.889	0.889	0.889	0.907
Fixed Effects:					
Hour FE		✓	✓	✓	✓
Day-of-Week FE		✓	✓	✓	✓
Planning Area FE		✓	✓	✓	✓

(b) *Trips fulfilled by street-hail (Poisson).*

Table A.3: *PPML estimates of the effect of transit disruptions on taxi trips, differentiated by booking and street-hail modes. "Unexpected" excludes planned disruptions. "No Service" indicates that at least one rail service is completely non-operational along some stretch of stations. "Search Pool" controls for the total taxi-minutes spent searching or waiting in the planning area during that cell-hour. The dependent variable is the count of trips of that type (booking/street hail) originating from a given planning area. Each observation is a (planning area, hour, date) cell. Standard errors are two-way clustered at the planning area and disruption event level. Trip data are from a large taxi company in Singapore, covering December 2016 to March 2018. Disruption data are from `mrtdown.org`. Asterisks ($*, **, ***$) indicate statistical significance at the 10%, 5% and 1% level respectively.*

1[trips > 0] (LPM)	(1) Baseline	(2) All FE	(3) Unexpected	(4) No Service	(5) Search Pool
Disrupted	0.0740***	-0.0070***	0.0055	-0.0067**	-0.0072***
	(0.0036)	(0.0027)	(0.0050)	(0.0028)	(0.0027)
Unexpected			-0.0160**		
			(0.0064)		
No Service				-0.0030	
				(0.0046)	
Search Minutes					-0.0000
					(0.0000)
Observations	571,200	571,200	571,200	571,200	571,200
Adjusted R^2	0.000	0.638	0.638	0.638	0.638
Fixed Effects:					
Hour FE		✓	✓	✓	✓
Day-of-Week FE		✓	✓	✓	✓
Planning Area FE		✓	✓	✓	✓

(a) *Aggregate trips: extensive margin (LPM).*

ln(trips) | trips > 0	(1) Baseline	(2) All FE	(3) Unexpected	(4) No Service	(5) Search Pool
Disrupted	0.9656***	0.0540	0.0447	0.0529	0.0678**
	(0.2171)	(0.0339)	(0.0304)	(0.0366)	(0.0286)
Unexpected			0.0120		
			(0.0489)		
No Service				0.0146	
				(0.0513)	
Search Minutes					0.0000***
					(0.0000)
Observations	528,750	528,750	528,750	528,750	528,750
Adjusted R^2	0.001	0.910	0.910	0.910	0.913
Fixed Effects:					
Hour FE		✓	✓	✓	✓
Day-of-Week FE		✓	✓	✓	✓
Planning Area FE		✓	✓	✓	✓

(b) *Aggregate trips: intensive margin (OLS, trips > 0).*

Table A.4: *Chen-Roth (2024) decomposition of the effect of transit disruptions on aggregate taxi trips. Panel A reports extensive-margin (LPM) estimates; Panel B reports intensive-margin (OLS on log trips conditional on positive trips) estimates. "Search Pool" controls for the total taxi-minutes spent searching or waiting in the planning area during that cell-hour. Each observation is a (planning area, hour, date) cell. Standard errors are two-way clustered at the planning area and disruption event level. Trip data are from a large taxi company in Singapore, covering December 2016 to March 2018. Disruption data are from `mrtdown.org`. Asterisks ($*, **, ***$) indicate statistical significance at the 10%, 5% and 1% level respectively.*

1[bookings > 0] (LPM)	(1) Baseline	(2) All FE	(3) Unexpected	(4) No Service	(5) Search Pool
Disrupted	0.1080***	-0.0067	0.0173*	-0.0060	-0.0068
	(0.0331)	(0.0050)	(0.0090)	(0.0054)	(0.0050)
Unexpected			-0.0310**		
			(0.0122)		
No Service				-0.0097	
				(0.0072)	
Search Minutes					-0.0000
					(0.0000)
Observations	571,200	571,200	571,200	571,200	571,200
Adjusted R^2	0.000	0.606	0.606	0.606	0.606
Fixed Effects:					
Hour FE		✓	✓	✓	✓
Day-of-Week FE		✓	✓	✓	✓
Planning Area FE		✓	✓	✓	✓

(a) *Booking trips: extensive margin (LPM).*

ln(bookings) | trips > 0	(1) Baseline	(2) All FE	(3) Unexpected	(4) No Service	(5) Search Pool
Disrupted	0.9002***	0.2050***	0.0634	0.2065***	0.1858***
	(0.0973)	(0.0408)	(0.1021)	(0.0443)	(0.0363)
Unexpected			0.1824		
			(0.1121)		
No Service				-0.0188	
				(0.1119)	
Search Minutes					-0.0001***
					(0.0000)
Observations	504,473	504,473	504,473	504,473	504,473
Adjusted R^2	0.001	0.810	0.810	0.810	0.816
Fixed Effects:					
Hour FE		✓	✓	✓	✓
Day-of-Week FE		✓	✓	✓	✓
Planning Area FE		✓	✓	✓	✓

(b) *Booking trips: intensive margin (OLS, trips > 0).*

Table A.5: *Chen-Roth (2024) decomposition of the effect of transit disruptions on taxi trips, differentiated by booking and street-hail modes. Panel A reports extensive-margin (LPM) estimates; Panel B reports intensive-margin (OLS on log trips conditional on positive trips) estimates. "Search Pool" controls for the total taxi-minutes spent searching or waiting in the planning area during that cell-hour. Each observation is a (planning area, hour, date) cell. Standard errors are two-way clustered at the planning area and disruption event level. Trip data are from a large taxi company in Singapore, covering December 2016 to March 2018. Disruption data are from* `mrtdown.org`*. Asterisks (*, **, ***) indicate statistical significance at the 10%, 5% and 1% level respectively.*

1[street-hail > 0] (LPM)	(1) Baseline	(2) All FE	(3) Unexpected	(4) No Service	(5) Search Pool
Disrupted	0.0972***	-0.0078***	0.0010	-0.0079***	-0.0082***
	(0.0039)	(0.0027)	(0.0073)	(0.0027)	(0.0028)
Unexpected			-0.0112		
			(0.0080)		
No Service				0.0015	
				(0.0069)	
Search Minutes					-0.0000**
					(0.0000)
Observations	571,200	571,200	571,200	571,200	571,200
Adjusted R^2	0.000	0.633	0.633	0.633	0.633
Fixed Effects:					
Hour FE		✓	✓	✓	✓
Day-of-Week FE		✓	✓	✓	✓
Planning Area FE		✓	✓	✓	✓

(a) *Street-hail trips: extensive margin (LPM).*

ln(street-hail) | trips > 0	(1) Baseline	(2) All FE	(3) Unexpected	(4) No Service	(5) Search Pool
Disrupted	0.8705***	0.0201	0.0041	0.0181	0.0440
	(0.2244)	(0.0430)	(0.0333)	(0.0467)	(0.0311)
Unexpected			0.0207		
			(0.0573)		
No Service				0.0250	
				(0.0543)	
Search Minutes					0.0001***
					(0.0000)
Observations	514,868	514,868	514,868	514,868	514,868
Adjusted R^2	0.001	0.911	0.911	0.911	0.918
Fixed Effects:					
Hour FE		✓	✓	✓	✓
Day-of-Week FE		✓	✓	✓	✓
Planning Area FE		✓	✓	✓	✓

(b) *Street-hail trips: intensive margin (OLS, trips > 0).*

Table A.6: *Chen-Roth (2024) decomposition of the effect of transit disruptions on street-hail trips. Panel A reports extensive-margin (LPM) estimates; Panel B reports intensive-margin (OLS on log trips conditional on positive trips) estimates. "Search Pool" controls for the total taxi-minutes spent searching or waiting in the planning area during that cell-hour. Each observation is a (planning area, hour, date) cell. Standard errors are two-way clustered at the planning area and disruption event level. Trip data are from a large taxi company in Singapore, covering December 2016 to March 2018. Disruption data are from* `mrtdown.org`. *Asterisks* ($*, **, ***$) *indicate statistical significance at the 10%, 5% and 1% level respectively.*

Trips (Poisson)	(1) Baseline	(2) All FE	(3) Unexpected
Disrupted	0.4413***	0.0860**	0.1680***
	(0.0869)	(0.0350)	(0.0394)
Unexpected			-0.1621***
			(0.0628)
Observations	287,616	287,616	287,616
Pseudo R^2	0.000	0.911	0.911
Fixed Effects:			
Hour FE		✓	✓
Day-of-Week FE		✓	✓
Planning Area FE		✓	✓

Table A.7: *PPML estimates of the effect of transit disruptions on ridehailing trips. "Unexpected" excludes planned disruptions. The dependent variable is the count of ridehailing trips originating from a given planning area. Each observation is a (planning area, hour, date) cell. Standard errors are two-way clustered at the planning area and disruption event level. Trip data are from a 10% random sample of ridehailing trips in the Singaporean market, April to October 2018. Disruption data are from* `mrtdown.org`. *Asterisks (∗, ∗∗, ∗ ∗ ∗) indicate statistical significance at the 10%, 5% and 1% level respectively.*

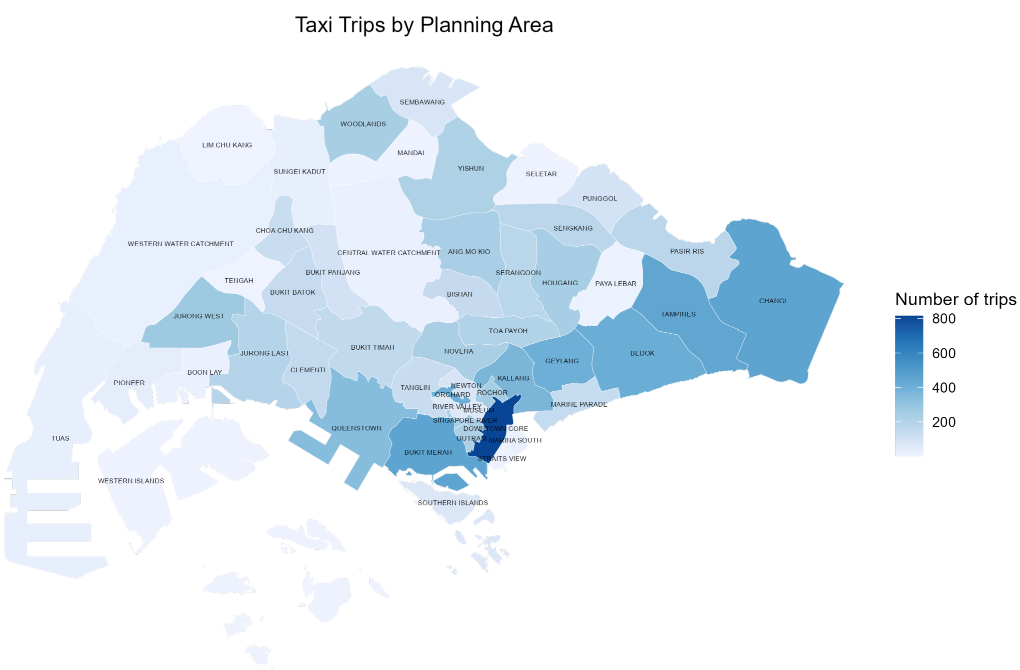

Figure A.2: *Daily average number of taxi trips per hour served by the drivers in our sample, by planning area, from December 2016 to March 2018, except for June and July 2017. Data are from a large taxi company in Singapore.*

1[trips > 0] (LPM)	(1) Baseline	(2) All FE	(3) Unexpected
Disrupted	0.0813***	-0.0044	-0.0033
	(0.0098)	(0.0124)	(0.0214)
Unexpected			-0.0025
			(0.0222)
Observations	287,616	287,616	287,616
Adjusted R^2	0.000	0.585	0.585
Fixed Effects:			
Hour FE		✓	✓
Day-of-Week FE		✓	✓
Planning Area FE		✓	✓

(a) *Ridehailing trips: extensive margin (LPM).*

ln(trips) | trips > 0	(1) Baseline	(2) All FE	(3) Unexpected
Disrupted	1.0630***	0.1049**	0.1924***
	(0.0969)	(0.0411)	(0.0404)
Unexpected			-0.2019***
			(0.0582)
Observations	260,951	260,951	260,951
Adjusted R^2	0.001	0.923	0.923
Fixed Effects:			
Hour FE		✓	✓
Day-of-Week FE		✓	✓
Planning Area FE		✓	✓

(b) *Ridehailing trips: intensive margin (OLS, trips > 0).*

Table A.8: *Chen-Roth (2024) decomposition of the effect of transit disruptions on ridehailing trips. Panel A reports extensive-margin (LPM) estimates; Panel B reports intensive-margin (OLS on log trips conditional on positive trips) estimates. Each observation is a (planning area, hour, date) cell. Standard errors are two-way clustered at the planning area and disruption event level. Trip data are from a 10% random sample of ridehailing trips in the Singaporean market, April to October 2018. Disruption data are from `mrtdown.org`. Asterisks ($*, **, ***$) indicate statistical significance at the 10%, 5% and 1% level respectively.*

ln(1 + search minutes)	(1) OLS	(2) OLS + FE
Disrupted	0.8397***	-0.1601***
	(0.0658)	(0.0584)
Observations	571,200	571,200
Adjusted R^2	0.001	0.903
Fixed Effects:		
Hour FE		✓
Day-of-Week FE		✓
Planning Area FE		✓

(a) *Search pool: OLS.*

Search minutes (Poisson)	(1) OLS	(2) OLS + FE
Disrupted	0.2288	-0.0806
	(0.1400)	(0.0552)
Observations	571,200	571,200
Pseudo R^2	0.000	0.844
Fixed Effects:		
Hour FE		✓
Day-of-Week FE		✓
Planning Area FE		✓

(b) *Search pool: Poisson PPML.*

Table A.9: *Effect of transit disruptions on local idle taxi supply. The dependent variable is the search pool measure: total taxi-minutes spent in SEARCH or WAIT status in a given planning area during a cell-hour. Panel A reports OLS estimates with $\ln(1 + \text{search minutes})$ as the dependent variable; Panel B reports Poisson PPML estimates. Each observation is a (planning area, hour, date) cell. Standard errors are two-way clustered at the planning area and disruption event level. Asterisks $(*, **, ***)$ indicate statistical significance at the 10%, 5% and 1% level respectively.*

ln(1+fare)	(A) Full panel	(B) PA×Hour FE	(D) First-hour unexp.	(D') Date×Hour FE
Disrupted	0.0335***	0.0372***	-0.0087	-0.0088
	(0.0105)	(0.0086)	(0.0067)	(0.0067)
Observations	83,408,441	83,408,434	661,656	661,656
Adjusted R^2	0.085	0.093	0.123	0.123
Fixed Effects:				
Vehicle FE	✓	✓	✓	✓
Pre-hour PA FE	✓		✓	✓
PA x Hour FE		✓		
Date FE	✓	✓	✓	
Hour-of-Day FE	✓		✓	
Date x Hour FE				✓
Clusters:				
Vehicles	17,353	17,353	17,298	17,298
Disruption events	97	97	78	78

Table A.10: *Vehicle-level DiD of transit disruptions on taxi fares. The dependent variable is $\ln(1 + \text{fare})$ at the (vehicle, date, hour) level. Specifications follow Agarwal et al. (2025) (Equation 9), adapted to a setting with time-varying treatment. Spec A is the full on-duty panel with vehicle, pre-hour planning-area, date, and hour-of-day fixed effects. Spec B replaces the spatial and time controls with planning-area $\times$ hour-of-day fixed effects. Spec D restricts to the first hour of unexpected disruptions (our headline causal specification). Spec D' replaces date and hour-of-day fixed effects with date $\times$ hour-of-day fixed effects. Standard errors are two-way clustered at the vehicle and disruption event level. Asterisks ($*, **, ***$) indicate statistical significance at the 10%, 5% and 1% level respectively.*

ln(1+fare) — Spec D	(D) All fares	(D) Booking	(D) Street-hail
Disrupted	-0.0087	0.0027	-0.0123
	(0.0067)	(0.0097)	(0.0076)
Observations	661,656	661,656	661,656
Adjusted R^2	0.123	0.174	0.120
Fixed Effects:			
Vehicle FE	✓	✓	✓
Pre-hour PA FE	✓	✓	✓
PA x Hour FE			
Date FE	✓	✓	✓
Hour-of-Day FE	✓	✓	✓
Date x Hour FE			
Clusters:			
Vehicles	17,298	17,298	17,298
Disruption events	78	78	78

Table A.11: *Vehicle-level DiD of transit disruptions on taxi fares by dispatch channel. The dependent variable is $\ln(1 + \mathit{fare})$ at the (vehicle, date, hour) level, decomposed into fares earned via booking dispatch and fares earned via street-hail. All three columns use Spec D (first hour of unexpected disruptions), with vehicle, pre-hour planning-area, date, and hour-of-day fixed effects. Standard errors are two-way clustered at the vehicle and disruption event level. Asterisks ($*, **, ***$) indicate statistical significance at the 10%, 5% and 1% level respectively.*

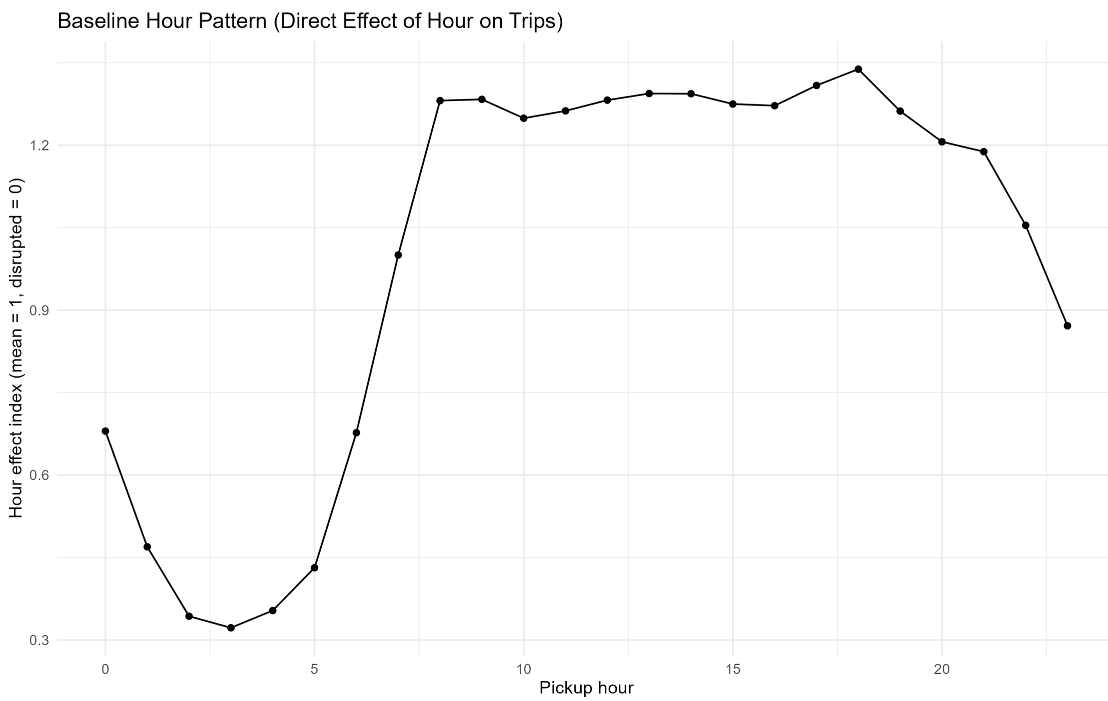

Figure A.3: *Point estimates of the effect of the hour of the day on the number of taxi trips taken, with the mean effect normalized to 1. Estimates are obtained by running the regression implied by Equation (1). Data are from a large taxi company in Singapore.*

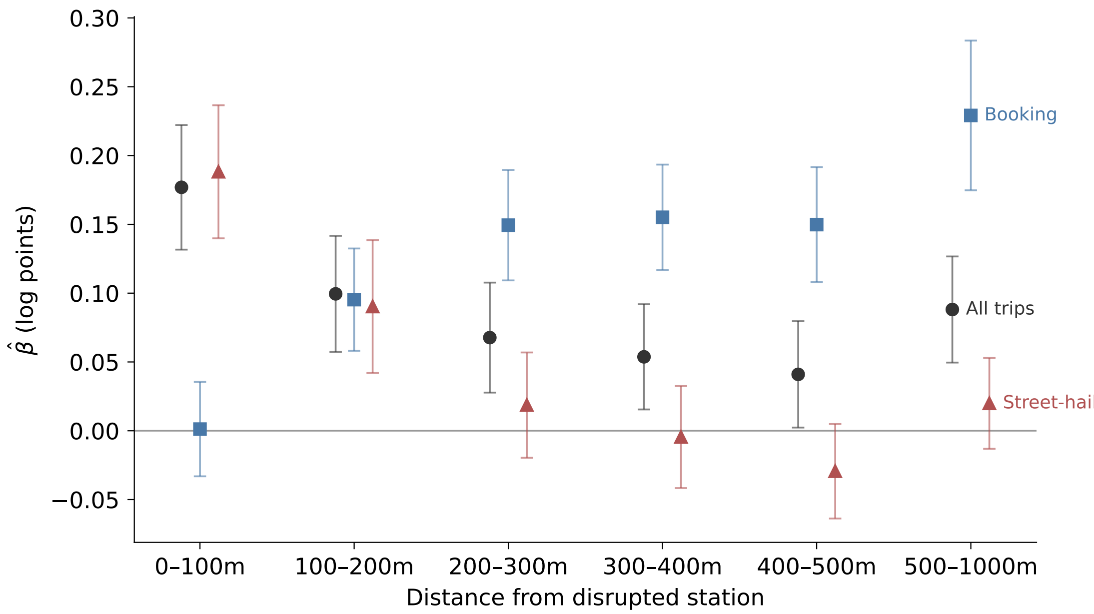

Figure A.4: *Ring difference-in-differences: taxi trips by mode and distance from disrupted station. Each point is the estimated total disruption effect at a given distance ring, comparing each station-ring-hour cell to itself on non-disrupted hours. The specification includes station-by-ring, hour-of-day, and day-of-week fixed effects, with standard errors clustered at the station level. Circles denote all trips, squares denote booking trips, and triangles denote street-hail trips. Vertical bars show 95% confidence intervals.*