The Impact of Ride-Hail Surge Factors on Taxi Bookings

Sumit Agarwal
National University of Singapore Business School
ushakri@yahoo.com

Ben Charoenwong
National University of Singapore Business School
bizbgc@nus.edu.sg

Shih-Fen Cheng
School of Information Systems, Singapore Management University
sfcheng@smu.edu.sg

Jussi Keppo
National University of Singapore Business School
keppo@nus.edu.sg

Abstract

We study the role of ride-hailing surge factors on the allocative efficiency of taxis by combining a reduced form estimation with structural analyses using machine-learning-based demand predictions. We estimate an upper-bound of the cross-price elasticity of taxi bookings to surge factors of only 0.26, but incorporating surge factors into a demand-prediction model improves the out-of-sample accuracy by 12-15%. Our structural analyses based on a driver guidance system finds the improved accuracy reduces drivers' vacant roaming times by 9.4% and increases average trips per taxi by 2.3%, suggesting the price information is valuable across platforms, even if elasticities are low.

1 Introduction

The ride-hailing industry has grown from a niche start-up market to the canon of technological disruption, with around US$60 billion in revenue as of 2017 and is projected to grow to $285 billion in revenue by 2030.[1] Moreover, since the entry of Uber, various competitors have emerged, such as Lyft in the United States, Didi in China, and Grab in Southeast Asia. Even conventional taxi companies have adopted or developed smartphone applications (apps) to compete with the ride-hailing industry's convenience. Within three to five taps on a smartphone, a customer can book a ride through another app. So, when the price of one option rises, customers can easily switch to alternative options. In fact, substitution becomes even easier if an app-based flat-rate taxi booking service is available because the price comparison becomes simpler; customers do not need to dynamically check taxi prices. Altogether, this means the future of the conventional transportation industry and the viability of disrupters are in the hands of consumers.

In this paper, we study the effects of ride-hailing surge factor, a common dynamic pricing mechanism defined as a price multiplier relative to regular taxis fares (e.g., an increase in the surge factor represents a decrease in the relative price of taxis). More specifically, we use a dataset on the taxi industry in Singapore to study the impact of ride-hailing surge prices on conventional taxi booking demand and the allocative efficiency of taxis. First, to establish an upper-bound on the cross-price substitution effect from ride-hailing services to taxi bookings, we use an empirical-specification with high-dimensional fixed effects utilizing variation in surge prices and rides in a narrowly-defined geographical zone within a half-hour interval within a given day of the week.[2] Second, to study the informational role of surge prices on taxi bookings and allocative efficiency, we use our detailed taxi data with demand forecasts by a random forest model to simulate a counterfactual guidance system that directs taxis to areas with higher taxi demand and study the decrease in vacant roaming times and the number of rides.

[1] According to a study by Goldman Sachs as reported in a May 27, 2017 MarketWatch article by Caitlin Huston, titled "Ride-hailing industry expected to grow eightfold to $285 billion by 2030". URL: https://www.marketwatch.com/story/ride-hailing-industry-expected-to-grow-eightfold-to-285-billion-by-2030-2017-05-24, last accessed June 10, 2020.

[2] Geographical regions range from around 2 km^2 (1 mi^2) to 20 km^2 (15 mi^2), where most regions in Singapore are closer to the smaller end. The larger regions include areas in Singapore consisting of two water-catchment reservoir areas.

There are three main benefits from the Singapore setting. First, taxis in Singapore already provide similar features as those offered by ride-hailing services, such as cashless payments, location tracking, and booking through a smartphone app. Therefore, to the extent that taxis and ride-hailing services are more easily substitutable in Singapore compared to other countries, our upper bound estimates on the cross-price elasticity of taxi ride demand also provide an upper bound for other countries. Second, the leading ride-hailing service provider in Singapore – Grab – has a ride-hailing feature similar to UberX, which uses a surge factor mechanism to balance supply and demand of rides.[3] Third, the Land and Transport Authority of Singapore maintains a dataset containing the mobility traces of all taxis in Singapore updated every 30 seconds, showing their location, whether a taxi is occupied or available, and if occupied, whether a trip was initiated by app booking or street pick-up. With these three unique characteristics discussed above in the Singapore setting, we identify the cross-price demand elasticity without various confounding factors.

Our first empirical study documents the reduced-form relationships between taxi bookings and surge factors, controlling explicitly for taxi supply within a half-hour interval and within route. Explicitly controlling for the supply measures addresses typical concerns of simultaneity bias when estimating demand elasticities. Our half-hour interval, weekend, and trip fixed effects control for some but not all unobserved variables that affect overall ride demand such as unexpected traffic congestions and local entertainment events. Therefore, our estimated coefficient represents an upper-bound on the cross-price elasticity of taxi booking demand relative to ride-hailing surge factors as any remaining omitted variables would affect the demand of both ride-hailing and taxis, introducing positive bias. Yet, we find that a 10% increase in the surge factor, controlling for taxi supply, raises taxi booking by only 2.6% within the same region, half-hour interval, and day-of-week. Therefore, consumers appear to be relatively price inelastic. But despite this low estimate, we show surge factors are useful for practical applications.[4]

[3] Additional details are available at https://www.grab.com/sg/justgrab/.

[4] In the Appendix, we also study the effect of trip length on the customers' behavior. Although the surge factor applies uniformly across the entire trip, we do not find any statistically significant differences in the price elasticity of customers in short-distance versus long-distance trips. A short-distance trip is defined as a ride that is less than 5 kilometers, or about 3.1 miles. This is about 10% of Singapore's length on the east-west dimension (the longer end). Moreover, we find that the relation between surge factors and taxi bookings is concave (see Appendix B.2), consistent with the idea that consumers perhaps react more to the notification of surge factor than the actual numerical surge factor itself.

Our second empirical study tests whether surge factors affect the demand prediction of taxi bookings. Using a random forest model, we find that the surge factor information is even more important in the demand forecasting than the distance of the ride, the month, and the amount of rainfall. Including the surge factor in addition to other variables improves the accuracy of taxi booking predictions between 12% and 15% in terms of out-of-sample root-mean-squared error. To evaluate the practical implications of this more precise taxi demand prediction, we embed the predictions into a counterfactual structural analysis.

Using a driver guidance system (DGS) for taxi drivers, a realistic structural model that has been implemented into the everyday industry practice, we show that a 15% improvement in the accuracy of demand predictions leads to 9.4% reduction in the average vacant roaming time. The DGS takes a taxi fleet operator's perspective by constructing a microscopic agent-based taxi fleet simulation in which taxis follow roaming guides to maximize the taxi fleet's utilization, and is adopted from the model in Cheng et al. (2018) which has been adopted and deployed by the taxi fleet. The roaming guides in our simulation are generated with demand predictions as inputs. In addition to demand-side information, the DGS also incorporates supply-side information (the current and future distribution of vacant taxis), with the objective of maximizing expected income of guided drivers. The guidance to drivers is delivered via a smartphone app, which is illustrated in Figure 1. The app highlights the recommended zone in red, and when the guided driver enters the recommended zone, additional details such as which streets or taxi stands to visit are displayed. The counterfactual analysis holds all other settings equal and then studies the effects of changes in the demand prediction accuracy.

Incorporating the surge factor to the demand forecasts produces a noticeable increase in the taxi fleet's efficiency and consumer welfare, without necessarily decreasing the number of ride-hailing bookings. The higher counterfactual taxi utilization rate corresponds to a weakly decreasing waiting time for customers as the matching between the riders and taxis improves. Moreover, the simulation documents a 2.3% increase in the number of taxi rides. After all, the surge factor is optimized in response to a constrained ride-hailing service that is partly mitigated with the improved taxi allocation.

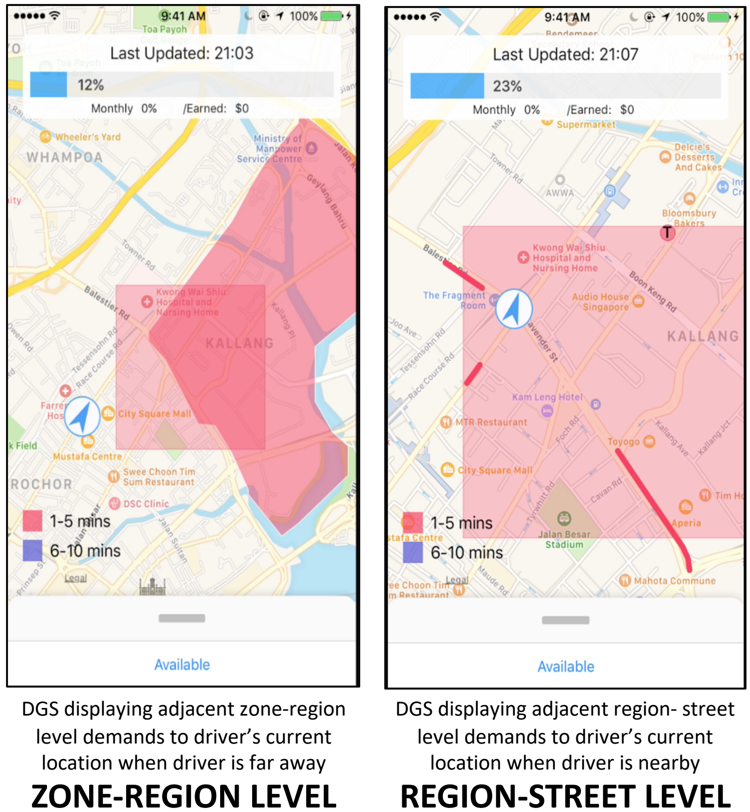

Figure 1: Screenshot of driver guidance system (DGS) app that implements part of our empirical analysis.

Our results highlight the informational role of a platform's surge pricing even for competing services and even if elasticities are low. Altogether, our results contribute to the consumer decision-making and market design literature as well as the demand and revenue prediction literatures using machine learning methods. Our causal estimates of the cross-price elasticity of taxi booking demand and the counterfactual simulations comprise a comprehensive study of microeconomic impacts in business practices. We identify the contribution relative to each literature below.

1.1 Related literature

Our findings contribute to the applied microeconomics literature studying the effect of ride-hailing services like Uber, Lyft, and Grab on the taxi industry, consumer welfare, and market equilibrium. To this end, our research identifies the behavior of consumers who decide between alternative transport options using detailed and high-frequency ride-level data. From a theoretical perspective, Castillo et al. (2017) show that surge factors can be used to avoid a wild goose chase, where drivers are sent to pick up distant customers. Lu et al. (2018) show that a surge factor heat-map attracts drivers and raises their revenues by up to 70%. Meanwhile, Cohen et al. (2016) estimate that Uber increased consumer surplus in the United States in 2015 by $6.8 billion.

Compared to the market design and consumer surplus research, the literature has not reached a consensus on the overall impact of ride-hailing services on the existing transportation industry. For example, Martin and Shaheen (2011) find that ride-hailing has both intensive and extensive margin effects. Current taxi riders may switch to ride-hailing services, while non-taxi riders may switch to both ride-hailing services and taxis. Berger et al. (2017) show using the staggered introduction of Uber in different states in the USA that Uber does not seem to cause adverse employment outcomes for taxi drivers but reduces the earnings potential of incumbent drivers. The empirical effect of ride-hailing services on the incumbent transportation industry has led to a range of diverse legal views regarding their regulation (Posen (2015) and Ross (2015)). Our paper contributes directly to the research studying the impact of ride-hailing services on the incumbent taxi industry by documenting the effects of surge pricing on taxi demand and supply. Our findings also contribute to research studying how platforms compete with each other. According to Crunchbase as of January 30, 2019, the combination of AirBnB, Uber, Grab, and Lyft have raised over US$40 billion cumulatively. Understanding the willingness of customers to substitute across these different platforms is important in analyzing network externalities that are present in these two-sided market settings.

Taxi (or more generally, taxi-like services) demand prediction is well-studied in the areas of data mining and transportation research. This line of research is enabled by the availability of GPS traces of service vehicles (taxis or cars), first low in quality and on an offline basis, but in recent years we have seen great improvement in both the accuracy and the frequency, and also real-time accessibility. For demand prediction, this technological advancement allows researchers to move from classical offline statistical approach such as time series modeling (e.g., Moreira-Matias et al. (2013)) to the state-of-the-art online machine learning approach such as deep learning (e.g., Ke et al. (2017) and Yao et al. (2018)). However, to the best of our knowledge, the cross-platform substitution effect between ride-hailing and taxi services has not been exploited in taxi demand prediction literature.

Finally, our findings also provide implications for the allocation of available taxis based on demand predictions, relating to existing research on the optimal dynamic allocation of scarce resources with uncertain demand.[5] For instance, Mahdian et al. (2007) and Borgs et al. (2014) study pricing and advertisement allocation problems for a firm with unreliable estimates and time-varying capacity levels. Javanmard and Nazerzadeh (2019) and Javanmard et al. (2019) study a dynamic pricing policy when customers have preferences over a high dimensional vector of features. Bian and Liu (2019a,b) analyze optimal ride-sharing allocation for customer-specific tastes, and Barann et al. (2017) study ride-sharing in New York and its effect on traffic, gasoline consumption, and CO_2 emission. Chu et al. (2018) show that information sharing might not always be beneficial for drivers because their equilibrium profits may reduce when ride-hailing platforms give information on the riders' origins and destinations to idle drivers. Related to that paper, Afeche et al. (2018) study demand-side admission control versus supply-side repositioning control for the spatial matching of ride-sharing customers and strategic drivers.

The rest of the paper is organized as follows. Section 2 introduces the taxi and ride-hailing market in Singapore. Section 3 explains the used data and empirical methodology, and then Section 4 gives the main results. Section 5 extends the empirical analysis to taxi bookings forecasting and its policy implications. Finally, Section 6 concludes.

2 Taxi Booking and Ride-Hailing in Singapore

The taxi market in Singapore is tightly regulated and more than 99% of taxis are owned by seven taxi companies. The largest group has almost 60% market share. Only citizens 30-year old and above can drive taxis, and intended drivers must pass a vocational license test and then rent a taxi from one of the taxi companies. The rental cost is set by individual operators and covers

[5] Additional works like Liu and Shen (2020) study the impact of demand predictions on the optimization of last mile delivery, Glaeser et al. (2019) use customer location data to optimize the location of retail pick-ups, Ferreira et al. (2016) study demand prediction for an e-commerce platform optimizing inventory, Bimpikis et al. (2019) study optimal pricing in a ride-hailing network, and Seamans and Zhu (2014) and Zervas et al. (2017) study incumbent responses to platform competitors in the newspaper and hotel industries.

all vehicle-related expenses, just like car rentals. Drivers pay for the variable costs such as fuel, parking, and road tolls. The taxi industry grew steadily in terms of aggregated fleet size until the end of 2014, after which the fleet decreased due to the emergence of ride-hailing firms such as Uber and Grab. Overall, the proportions of rides served by the companies correspond to the number of cars they own.

Surcharges were introduced in 1994 and updated since to encourage the supply of taxis at strategic times and locations. Taxis have a 25% surcharge on top of the metered fare on Mondays to Friday between 6:00 am to 9:30 am and Mondays to Sundays and Public Holidays from 6:00 pm to 12:00 am. They also have a 50% surcharge on metered fare from 12:00 am to 5:59 am on all days. Further, rides originating from high-demand tourist destinations such as Marina Bay Sands, Resorts World Sentosa, and Singapore Expo have a fixed dollar surcharge; rides originating from the Central Business District and airport also have a surcharge. Taxi "peak hour" pricing is decided by individual operators and approved by the Public Transport Council while the Land and Transport Authority (LTA) decides the "area-based" surge pricing (e.g., airports, Marina Bay Sands, and Central Business District). Although it is not clear how much control the Public Transport Council exercises over the pricing decisions, the only differences in prices are the initial fixed-cost of the taxi and, thus, all remaining rates and peak hour pricing are the same across all operators. These surcharges can be viewed as "static" price surges to spatiotemporally balance supply and demand. That is, the taxi surcharges do not respond dynamically and geographically to unexpected demand-supply imbalances.

Compared to the taxi prices, the fares of ride-hailing service are set dynamically by the ride-hailing app's algorithm to equilibrate the supply and demand of ride-hailing rides. The price is represented as a surge factor - the relative price between the ride-hailing and the standard taxi. When Grab demand is high relative to Grab supply, the surge factor is greater than one, meaning that ride-hailing prices are higher than the taxi prices. On the other hand, when Grab supply exceeds Grab demand, the surge factor is equal to or less than one, meaning that ride-hailing prices are lower than the taxi prices.

Since 2014, customers have been able to book taxis through the taxi company's smartphone app. Not all taxi operators have their own smartphone taxi booking app, but some do. Moreover, customers can also book taxis through a phone call. All taxi stands have a unique identification number that can be supplied when ordering a taxi through the phone.

However, in recent years, most smaller taxi companies ended their support for their own booking apps and instead started to collaborate with Grab to provide booking service to customers. As of April 2017 (the beginning of our sample), five out of the seven taxi fleet operators had formally reached an agreement with Grab to allow their taxis to be part of the Grab fleet via the JustGrab service. A taxi trip booked via the JustGrab service is priced dynamically within the Grab service platform using Grab's surge pricing scheme. In our main analysis, we focus on the taxi bookings of a taxi company (which owns two taxi operators) that forbids its drivers from participating in the Grab service (for confidentiality, we cannot disclose the name of the taxi company).

In Singapore, private car drivers participating in ride-hailing services are required by the Land and Transport Authority to obtain a ride-hailing drivers' license from the LTA for S$400. If they do not comply with this rule, they face a S$10,000 fine. In addition, a ride-hailing company may have further required training. In April 2017, there were 41,297 Uber and Grab private-car drivers, which was 56% more than the total taxi population of 26,476.[6]

3 Data and Methodology

3.1 Data

The data used in this study can be broken down into those pertaining to prices and those pertaining to quantities. For prices, we derive location-specific surge factors from the price of the JustGrab service from selected location pairs based on publicly available pricing information. For quantities, we use two datasets on taxi rides that were either derived from or obtained directly from the LTA.

For surge factor collection, we choose representative origin-destination pairs that historically produced the most numbers of taxi booking requests with some manual adjustments to avoid

[6] Tan, Christopher. May 24, 2017. "Private-hire cars outnumber taxis by a mile". The Straits Times. URL: https://www.straitstimes.com/singapore/transport/private-hire-cars-outnumber-taxis-by-a-mile, last accessed June 10, 2020.

similar locations being queried too frequently (e.g., the airport).[7] The query is run at each origin-destination pair at fixed intervals, between 30 to 60 minutes per pair based on the historical ride frequencies. The locations are based on the postal districts used in Singapore, which are the first two digits of the six-digit zip code administered by the Singapore Postal service.[8] We note that customers are only shown the price of the whole ride and not a multiplicative surge factor (unlike Uber, which shows both as of September 2019). However, because we can calculate baseline price for a typical taxi ride using published pricing formula, we can extract the surge factors. This procedure yields a dataset of route-level surge factors at every thirty-minute interval throughout the whole sample period.

The first dataset is derived from the raw taxi mobility traces and contains the number of taxi trips between an origin and destination zone across all operators and trip types at every half-hour interval. We also observe whether the ride was initiated through a booking (either through the phone or app), street pick-up, limo services, or other means. For our main analysis, we focus on the taxi bookings of the two taxi operators that forbid their drivers from participating in the Grab service as the outcome variable (in the following sections, we refer to this outcome variable as simply the "taxi bookings"). As a placebo test, we also consider the impact of surge factors on street pick-ups.[9] The half-hour interval also allows us to control precisely for time-of-day effects such as rush hour and the end of school hours. The dataset ranges from April to August 2017, with the exception of July, where the LTA experienced a database malfunction of several taxi operators in certain areas, resulting in severe loss of data. For data consistency, we drop July from our sample.[10]

The second dataset contains a panel of all taxis in Singapore at every half-hour interval. In this dataset, we observe the anonymized taxi identifier, the operator of each taxi, whether the

[7]These pairs are plotted in Figure A.1 in the Appendix. The origin and destination zones are fairly representative of the level of transportation and economic activity in Singapore, which is more concentrated towards the south of Singapore, in the Central Business District.

[8]The postal districts in Singapore were introduced in 1995 to split the island into regions with roughly equal populations in each area for administrative and mailing purposes. Most of the regions are closer to the smaller range of 1-square mile rather than those up to 15-square miles.

[9]In our raw data, there are three additional categories that we do not consider: third-party bookings, unknown and special pickup. We group these into "Other", which consists of around 5% of rides. Because we do not know the exact reasons behind these classifications, we exclude them from the main analyses.

[10]Including the few days of July that we have data for does not affect our main results.

taxi is available or hired, and its location. These micro data allow us to split total taxi supply into excess supply (taxis with the "Available" status) and hired supply (taxis with the "Hired" status) by each operator in each zone within each half-hour interval. We construct both our taxi supply and proxy for Grab supply using this data based on whether an operator participates in JustGrab. The supply of taxis is based on the operators that forbid their taxis from participating in the JustGrab service and the proxy for JustGrab supply is based on the other taxi operators that permit their taxis to participate as JustGrab cars. Our proxy for Grab supply is not perfect because taxi drivers for participating operators choose whether to drive as a Grab driver or not. However, since the participating operators do not have alternative means of booking rides through a smartphone application, the taxi drivers most likely choose to drive as a Grab driver.[11]

Thus, even without direct rides data from Grab, we can study the relation between Grab and taxi bookings controlling for both taxi and Grab supply due to the overlap of Grab cars and certain taxi operators in Singapore. Since the taxis that operate in the Grab platform are able to observe the surge factors, their driving behavior should highly correlate with those of private cars. However, the lack of private car data means our proxy for Grab supply likely underestimates actual Grab supply. But although the point estimate capturing the relation between taxi bookings and our measure of Grab supply will be biased, it can serve as a valid control variable.

We merge surge factors to taxi rides by postal districts and half-hour intervals. Surge factor information is acquired by using the Grab app for specific origin postal codes and destination postal codes repeatedly. To merge the Grab surge factor dataset to the taxi trips dataset, for each half-hour interval, we aggregate surge factors by averaging the surge factor across all available postal codes into a postal district. Lastly, to study whether taxi supply responds to surge factors and public signals, we also use rainfall data from the National Environmental Agency of Singapore as an exogenous shock in our analysis. The weather station locations in terms of latitude and longitude are merged into the nearest postal code, which are then merged to the corresponding postal district. Figure A.2 in the Appendix shows the locations of the weather stations.

[11]Sign up is free and anecdotal evidence suggests that Grab drivers can deny as many jobs as they want with the only "penalty" being that they are ineligible for Grab's distinguished driver awards which contain a cash prize and require a minimum acceptance rate of 90%.

3.2 Reduced-Form Empirical Specification

A simple regression of taxi bookings on the surge factor suffers from simultaneity bias if taxi supply responds to the surge factor, and the standard practice is to find instruments that affect supply and not demand. However, rather than this standard instrumental variables approach, we adopt a direct approach of regressing quantities on prices, controlling for supply. We exploit a unique feature of most taxi markets: the relative price of taxi bookings to ride-hailing services is set in the ride-hailing services market, because taxi prices are according to a fixed pricing schedule and do not dynamically respond to changes in market conditions every half-hour. The surge factor captures the relative price of ride-hailing services relative to standard taxi bookings and is defined as

$$\text{surge factor} = \frac{P_{Grab}}{P_{Taxi}}.$$

The demand of car rides can be split into Grab demand, taxi booking demand, and taxi street pick-up demand. All components of demand are subject to similar demand shocks that affect the rides customers want to take, such as events or rain. The surge factor changes the relative price of Grab versus taxis. However, the relative price of booking a taxi versus street pick-up does not change dynamically. We believe the main margin of substitution for Grab demand is taxi booking demand, given that those booking rides are typically more time constrained and delay-sensitive (Taylor (2018)). Although the surge factor does not impact the prices that the taxi drivers receive, it introduces endogeneity due to its impact on relative prices of transportation options and consequently the residual demand for taxi bookings. Figure 2 shows the link between the Grab and taxi booking market, where the residual demand defined above is plotted against the surge factor.

The taxi market is not necessarily in equilibrium since nominal prices of taxis are fixed in the short term and do not respond to market conditions. Therefore, at any given time, the taxi booking market may have excess supply or excess demand. The same is not true for the Grab market since by construction, the surge factors adjust to clear the JustGrab market, matching supply with demand. Figure A.3 shows the cyclicality in the surge factor throughout the day that arises from the cyclicality from supply and/or demand. The supply of taxis are from the operators

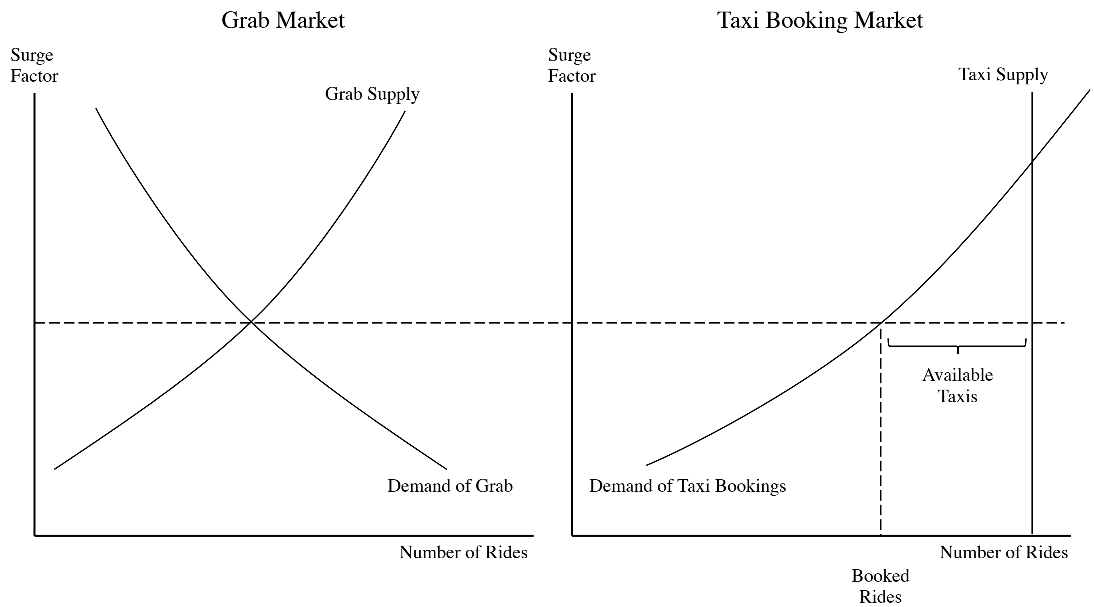

Figure 2: These figures show the interplay between Grab and taxi booking markets.

that do not participate in the JustGrab service, and come from the same pool of taxis available for booking and street pick-ups.[12]

The key feature of our empirical strategy lies in how the relative prices of taxis and ride-hailing cars are set and our data on taxi supply. Three features of our unique setting permit empirical identification of the cross-price elasticity. First, the taxi operators we consider do not permit their taxis to participate in the JustGrab service and, therefore, should not respond to the surge factor (we test this). Second, our high frequency taxi data permits a specification that includes high-dimensional fixed effects, controlling for origin zone by time interval by day-of-week fixed effects to account factors affecting both supply and demand. Given that the postal districts used in our study are around 1 to 5 square miles that were designed to split Singapore into regions with roughly equal portions of people, we believe the fixed effects do not distort our analysis. Third, to further alleviate the concern that perhaps taxi supply reacts to surge factors through other unobservables,

[12]In Table A.3 in the Appendix, we study the impact of surge prices on street pick-ups and find some evidence that street pick-ups decrease in response to a higher surge factor, as more taxis choose to respond to bookings due the booking fee. Also, taxi supply does not differ based on the destination because in Singapore, taxis are not permitted to accept or refuse a ride based on the destination when they are available. Such practice is a finable offence.

we explicitly control for taxi supply in a region at each point in time. Any remaining shocks to demand would introduce an upward bias in the estimated cross-price elasticity, as the demand of taxi bookings and ride-hailing services would be positively correlated.

Our high frequency data on taxi rides, ride-hailing surge factors, and full market-wide taxi supply permits an upper-bound estimation of cross-price elasticity of demand controlling for a myriad of potential shocks to demand and explicitly controlling for changes in taxi supply. Looking within routes, day of week, and half-hour intervals allows us to compare rides in the same route across days at exactly the same half-hour interval. The large cross section and high frequency of the data allow us to isolate local market factors both in terms of space and time using the following specification:

$$\begin{aligned}\log \text{Taxi Bookings}_{o,t,i} =& \alpha_{o,day(t),i} + \gamma_t + \beta \text{Surge Factor}_{o,t,i} + \\ &+ \phi \log \text{Taxi Supply}_{o,t,i} + \varepsilon_{o,t,i},\end{aligned}$$

where o indexes an origin zone, t indexes a date, i denotes a half-hour interval, and $day(t)$ is the day-of-the week of a given date. The outcome variable is the log number of taxi bookings through a phone booking or the taxi company's mobile application. We do not include street pick-ups, which may represent a different market for rides. We assume that the main margin of substitution is between Grab and booking taxis, not Grab and flagging taxis on the street or lining up at taxi stands. Explicitly controlling for taxi supply permits shocks like road conditions that affect both Grab and taxi supply, and not demand, to still identify the cross-price elasticity of surge factors and taxi bookings. Although we believe the surge factor is determined only based on the ride-hailing supply and demand at the origin and not taxi supply, our empirical specification also accounts for any potential reverse causality. In fact, comparing the point estimate with and without the taxi supply controls is indicative of these potential endogeneities. Any remaining demand shifters that affect both Grab ride-hailing demand and taxi-booking demand would induce a positive bias in our estimates, meaning that our estimates would be an upper-bound on the true cross-price elasticity. Figure A.4 in the Appendix shows surge factors are autocorrelated, so we cluster standard errors by starting zone as well as time to allow for cross-sectional correlations of surge factors, demand, and supply as well.[13]

In addition to the total number of rides, we also study the relation between response times and the surge factor, which provides more insight into whether the responses of taxi bookings may be constrained by taxi supply. If taxi supply is unable to accommodate for any substitution between ride-hailing and taxis, then the estimated relation between taxi bookings and surge factors would be biased downwards. Relatedly, in our follow-up analyses, we study the relation between surge factors and taxi supply using the following specification:

$$\log \text{Taxi Supply}^{K}_{o,t,i} = \alpha_{o,day(t),timeofday(t),i} + \beta \text{Surge Factor}_{o,t,i} + \varepsilon_{o,t,i},$$

where K is the taxi status (available or hired) in zone o in a half-hour interval i on day t. We use total taxi supply, which comprises both taxis that are available and taxis that are hired (so, total taxi supply is the sum of excess supply and realized, used supply). We later also break out supply into excess supply and hired supply, and study both separately to see which responds to the surge factor.

3.3 Variable Construction

The raw surge factors are relative to basic taxi fares only and do not account for peak hour or area pricing on actual metered taxi fares. Therefore, the raw surge factors tend to overstate the actual surge price, which might introduce a negative bias on our estimate of the elasticity of substitution, and we use surge factors adjusted by the actual metered taxi fares.[14] In Singapore, there are taxi stands at every mall which provide the pricing schedule for taxis, including all multiplicative and lump-sum surcharges. Moreover, every taxi is required by the LTA to have a standardized sticker of the fare schedule on the window.

Since most taxi riders are aware of the surcharge schedule (both multiplicative and lump sum), as mentioned above, we adjust the raw surge factors to get the actual relative surge factor. For

[13] Related to this, Cachon et al. (2017) study how surge pricing affects the self-scheduling for rides.

[14] For robustness, we also estimate the elasticity of substitution using raw surge factors rather than the adjusted surge factors and find the elasticity of substitution estimate drops from 0.398 to 0.319.

the multiplicative surcharges from 9:00pm to 5:59am, we scale the taxi rates by the surcharge. For area-based lump-sum surcharges, we first convert reported surge factors to actual prices based on standard taxi fares (both the fixed start-up cost and the variable rate cost) and divide by the taxi fare including the lump-sum charge.[15] We also incorporate a lump-sum charge for the booking fee for taxi bookings made through the taxi app.[16] Overall, the adjustment affects 23% of the prices in our sample, with 15% of the observations having a difference in adjusted and unadjusted surge factors by more than 0.3.

The adjusted surge factors still show cyclicality through the day (see Figure A.3 in Appendix A). Because the data spans multiple horizons across many days and narrowly-defined regions in Singapore, we have a large heterogeneity across any half-hour intervals. In our sample, the average surge factor is 0.90, meaning that Grab is 10% cheaper than taxis on average. The variance in the surge factor is 0.24, with a minimum observed surge factor of 0.67 and a maximum of 2.40. In our sample, over 15% of Grab rides have a surge factor above 1.00.

Our raw taxi supply data contains the geographical zone that a specific taxi was in and the number of minutes in each half-hour interval spent in either "Available", "Hired", "Busy", "Changing Shift", or "Other" statuses. Similar to Cramer and Krueger (2016), we construct the total taxi supply in a zone as the sum of all taxi-minutes in each status in half-hour intervals for each day, scaled by the 30-minute interval as

$$\text{Taxi Supply}^K_{o,t,i} = \frac{1}{30}\sum_j h^K_{o,t,i,j}$$

where K is the taxi status and $h^K_{o,t,i,j}$ is the number of minutes taxi j was in status K in zone o in a half-hour interval i on day t. Taxis in the "Available" status are considered excess supply, those in "Hired" are considered hired supply, and those in "Busy", "Changing Shift", or "Other" statuses are not considered in taxi supply. For the rest of our paper, unless explicitly qualified as either excess or hired supply, all references to taxi supply refer to the total taxi supply in a zone in

[15] This adjustment is possible because we query origin-destination zones. If we only had origin prices and trips, we would not be able to do this adjustment.

[16] Although there is no shown booking fee in the Grab platform, the minimum price of a JustGrab for any distance is S$6 while the minimum cost of a street-hail taxi is S$3.20.

a half-hour interval. Note that just because a taxi is in a zone but is hired does not mean it is not contributing to the supply in that area. It contributes to the used supply, but not excess supply. We think of each half-hour by geography as a market. For example, if the time interval from 9:00 am to 9:30 am in zone A had two taxis that spent 12 minutes and 28 minutes respectively in the "Available" status, we calculate an excess taxi supply of 1.33 (= 40/30).

4 Substitution Between Ride-hailing and Taxi Bookings

We find that all else being equal, a 10% change in the relative price of Grab and taxis increases taxi bookings by around 2.6% within the same half-hour interval going to the same-route. We further argue that the estimates capture a causal relationship not driven by confounding factors, reverse causality, or simultaneous causality in the subsections following the main estimates.

Table 1 shows our empirical estimates of the impact of surge factor on taxi bookings when explicitly controlling for both taxi supply and Grab supply within the same half-hour interval and the same geographical zone. Column (1) suggests that a 10% increase in surge factors, holding constant the local taxi and Grab supply, leads to an increase of around 2.6% in taxi bookings within the same half-hour interval. The point estimate on surge factors does not change from Column (1) to (2), which is consistent with the surge factor not being affected by taxi supply.

Table 1: Surge Factors and Taxi Bookings. This table shows the responses of the log number of taxi bookings and response time in response to the level of the surge factor explicitly controlling for both taxi and Grab supply. Observations are at the half-hour by start-zone level. All regressions include zone by time interval by day-of-week fixed effects. Standard errors (in parentheses) are clustered by zone and half-hour interval.

$y_{o,t,i} =$	log Taxi Bookings		$\frac{\text{Response Time}}{\text{Taxi Booking}}$	
	(1)	(2)	(3)	(4)
Surge Factor	0.257***	0.256***	−0.160***	−0.147***
	(0.276)	(0.027)	(0.025)	(0.022)
Taxi Supply	0.008		−0.293**	
	(0.053)		(0.113)	
Observations	57,486	57,486	57,429	57,429
R^2	0.525	0.525	0.449	0.433

*p<0.1; **p<0.05; ***p<0.01

Columns (3) and (4) show the relation between surge factors and the response times per taxi booking in minutes per taxi booking. In this specification, response time is represented as the average fraction of a 30-minute interval that a booked taxi spends to pick up the customer who made a booking. In both Columns (3) and (4), we find that a higher surge factor predicts a lower response time. Although seemingly counterintuitive, since surge factors correspond to high demand or congestion periods, the findings are consistent with an unconstrained taxi supply in an interval-zone that can respond to bookings more quickly as more riders substitute to taxi bookings. In Column (3), we also find that a higher taxi supply corresponds to a decrease in response time, as there are more available cars to respond to bookings in a more time-efficient manner.

Additional analyses in the Appendix show our results are robust to various specifications, alternative variable constructions, and different subsamples. Table A.2 shows that the results are robust to excluding the weekend, and falsification tests drawing from the surge factor or same route do not mechanically generate our results. Moreover, Figure A.5 shows that the point estimates of the elasticity are stable across months from April through August. We also find a negative relation between taxi street pickups and surge factors, likely because taxi drivers prefer booked rides rather than street pick-ups due to the additional booking fee that they receive. Although taxi drivers are legally not permitted to decline street-hailing customers based on their destination, the driver may ignore street hails. Effectively, the taxi drivers can easily switch from supply taxi rides to street pick-ups or bookings. Street pick-ups only offer a fixed pricing schedule, and the customers observe the trip price at the end of the trip. Further, on average, taxi drivers earn more when accepting bookings compared to street pick-ups due to the additional booking fee. This result is not inconsistent with customers switching between on-demand bookings made through a smartphone, and is consistent with our empirical specification to include all of taxi supply in the area. [17]

In the following subsections, we first show that the estimated substitution does not seem to be driven by constrained taxi supply. Second, we show that taxi supply does not respond to changes in the surge factor, even when taking into account relevant exogenous shocks like rainfall. Jointly,

[17]Unfortunately, the street pick-up specification is not a good placebo since street pick-ups also include rides by tourists and others that may not have a smartphone or the Grab app, who are likely to be unaware of switching between street pick-ups and Grab. Moreover, unlike street pick-ups, the taxi booking app and Grab app both show the trip price at the time of booking and, therefore, one could argue that they are close substitutes.

these follow-on analyses rule out alternatives due to simultaneous or reverse causality. Appendix Section B shows additional robustness and cross-sectional heterogeneity in responses. Appendix Section B.1 studies the impact of demand shocks through time and the impact of journey distance on the cross-price elasticity. Appendix Section B.2 studies whether the elasticity of substitution differs based on trip distance and Section B.3 studies whether the elasticity is affected by local area income.

4.1 Is Taxi Supply Constraining Consumer Substitution?

In this subsection, we study whether the availability of taxis restricts the consumers' ability to substitute between ride-hailing and taxi booking. We use a specification that interacts the surge factor with excess taxi supply, defined based on the number of taxis in the "Available" status in the previous half-hour interval. If the substitution is constrained, leading to a lower coefficient estimate relative to their true cross-price elasticity, we would expect the interaction term to impact the substitution effect that we estimated before. We use an empirical specification similar to those shown in Table 1, but also include the interaction term of surge factors with the amount of excess taxi supply from the previous period in the same zone.

In Table 2, we find that the interaction of the excess taxi supply in the previous half-hour interval with the surge factor is not statistically significant overall, in the subset of surge factors above one, on weekdays, and on weekends. The excess taxi supply is as an indicator of whether the taxi supply is greater than the median number of taxi supply in that same region. Overall, our results suggest that the estimated cross-price elasticity is not affected by the lack of available taxis to substitute to when surge factors are high. We conclude that, indeed, consumers do not appear very sensitive to the relative price of ride-hailing services.

4.2 Does Taxi Supply Respond to Surge Factors?

In this subsection, we directly test whether taxi supply is correlated with surge factors. There are two main possibilities on the relation of surge factors and taxi supply. First, if non-participating taxi drivers observe ride-hailing surge factors, they could gauge areas with high demand. There-

Table 2: Elasticity Depending on Excess Taxi Supply. This table shows the response of the log number of available taxis in response depending on the previous half-hour's taxi excess supply. The High Surge sample is a subset of data with above median surge factor value, and the Projected sample comes from instrumenting for surge factors with the excess supply in each origin destination in each half-hour interval. The observations are at the starting zone by half-hour interval level. All regressions include zone by time interval by day-of-week fixed effects. Standard errors (in parentheses) are clustered by zone and half-hour interval.

	$y_{o,d,t,i} = \log\left(\text{Taxi Bookings}_{o,d,t,i}\right)$			
	(1)	(2)	(3)	(4)
Surge Factor$_t$	0.240***	0.097**	0.260***	0.213***
	(0.029)	(0.038)	(0.036)	(0.034)
$1_{\{\text{Excess Taxi Supply}_{t-1} > \text{Median}\}}$	-0.051	-0.041	-0.056	-0.043
	(0.032)	(0.080)	(0.044)	(0.047)
Surge Factor$_t \times 1_{\{\text{Excess Taxi Supply}_{t-1} > \text{Median}\}}$	-0.011	-0.063	-0.012	-0.014
	(0.031)	(0.059)	(0.041)	(0.047)
Taxi Supply$_t$	0.017	-0.100	0.091	-0.083*
	(0.045)	(0.071)	(0.062)	(0.049)
Sample	All	Surge > 1	Weekday	Weekend
Observations	57,474	12,447	34,766	22,708
R^2	0.527	0.668	0.535	0.516

*p<0.1; **p<0.05; ***p<0.01

fore, if this is the case then taxi supply increases in the regions and times of rising surge factors. Alternatively, if taxi drivers do not observe the surge factors then we would expect no impact of surge factors on the total taxi supply for that fleet. Second, the surge factor algorithm may increase prices when taxi supply is low. Since these two possibilities counteract, the overall result can be in either direction. However, we argue that the latter case is unlikely since to our knowledge, Grab does not have a live data feed of the taxi drivers from the two operators that are not participating in the JustGrab service.

In this specification, we study how taxi supply responds to surge factors as well as rainfall, an exogenous variable that can affect both ride booking demand and supply. Taxi drivers are unable to observe surge factors but drivers near a rainy area can observe the rain and know whether an area has high demand. In fact, because rain is market-relevant and common information among all the drivers in the area, it may affect the allocation of taxi supply. On one hand, taxi supply may increase in anticipation of a predicted increase in demand in a rainy area. On the other hand, rain may also raise concerns that road conditions are more dangerous or congested.

Column (1) of Table 3 suggests that taxi supply increases by 13% during rainy periods, but the total supply of taxis (excess supply plus the hired taxis) does not appear to respond to surge factors within a zone in a specific half-hour interval. [18] Columns (2) and (3) show a higher surge factor is associated with more hired taxis and a direct reduction of excess supply of taxis, which is the number of taxis in the "Available" status. As consumers switch from Grab to taxis, the excess supply of taxis in a zone appears to completely absorb the substitution.

Table 3: Taxi Supply and Surge Factors. This table shows the response of the log taxi supply in response to the level of the surge factor. 1Rain is an indicator taking the value of one when there is any rain detected by the weather stations in the region. Rainfall is measured in centimeters. All regressions include zone by time interval by day-of-week fixed effects. Standard errors (in parentheses) are clustered by zone and half-hour interval.

$y_{o,t,i} =$	log Taxi Supply	log Excess Taxi Supply	log Hired Taxi Supply
	(1)	(2)	(3)
Surge Factor	-0.024	-0.448***	0.212***
	(0.043)	(0.039)	(0.041)
1{Rain}	0.133***	0.044***	0.142***
	(0.035)	(0.030)	(0.034)
log(1+Rainfall)	-0.003	-0.034	0.016
	(0.025)	(0.020)	(0.027)
Observations	57,427	57,427	57,427
R^2	0.339	0.459	0.375

*p<0.1; **p<0.05; ***p<0.01

Therefore, in a half-hour interval, higher demand matches with the available taxis, but the total taxi supply in the area does not change. Together with results from Tables 1 and 2, we conclude there appears to be a pool of available taxis that absorbs the entire increase in demand in a given half-hour interval.

Finally, the supply of taxis available come from the same pool of taxis as those available for street pick-ups. In Appendix Table A.3, we show that as surge price rises, street bookings decrease. That surge factors have no relation with taxi supply but is associated with an increase in taxi bookings and a small decrease in street pick-ups, our findings suggest the substitution of taxi supply from street pick-ups to bookings. This is consistent with the fact that all else equal, a taxi would prefer a booking due to the additional booking fee of S$3 to S$5 that they receive. Therefore,

[18] To our knowledge, Grab does not have a live feed of taxi supply and thus this is likely not due to reverse causality of taxi supply affecting surge prices. The surge factor more likely captures unfavorable road conditions like higher congestion.

the impact of the substitution between ride-hailing and taxi bookings on overall consumer welfare is complex as consumers who are more waiting-time sensitive (those in the ride-hailing and taxi booking markets) may decrease the consumer welfare of those using street pick-ups. As a result, we focus on the perspective of the taxi operator.

5 Demand Prediction and Structural Analyses

In the previous section, we identified the causal relation between ride-hailing surge prices and taxi bookings. One implication of our estimates is that around 6% of the current revenues from taxi bookings are due to this substitution. However, the implication of this substitution effect on taxi booking demand prediction is not obvious, since the cross-price demand elasticity is fairly small. To tackle this problem head-on, this section illustrates how much the substitution effect improves the predictive accuracy of taxi bookings. This prediction problem is of interest to policy makers as well as ride-hailing and taxi companies that seek to minimize congestion and to improve allocation of taxi supply.

5.1 Predicting Taxi Bookings with Machine Learning

We consider two sampling methods of our data in this prediction exercise. In the first sample, we leave out August 2017 as the test data and use data from April through June 2017 as the training sample. The second sample is built by randomly sampling 75% of the observations in each of the four months for the training data and using the remaining 25% of the data as the testing data. Although the first sample closely approximates the real world datasets used in predicting taxi bookings by policy makers and taxi and ride-hailing companies, comparing the results of these two samples provides a robustness check for the stability of the model over our full data period.

To allow flexibility in the interactions of variables, we use a random forest model based on the data that is available to policy makers (the LTA in this case). We forecast future taxi bookings and compare models based on the LTA data only with the corresponding models that include the surge factor. This allows us to isolate the marginal improvement of the relative price between the ride-hailing and taxi bookings in predictability. Typical guidelines for machine learning applications trade off between computational time and robustness. Since computational time is not a constraint in our counterfactual analysis, we use a larger forest with 500 trees for robustness.[19] Moreover, since our samples are representative across different categories of start zones, end zones, and time intervals, we use a simple random sample method for each out-of-bag sample used to evaluate each tree.

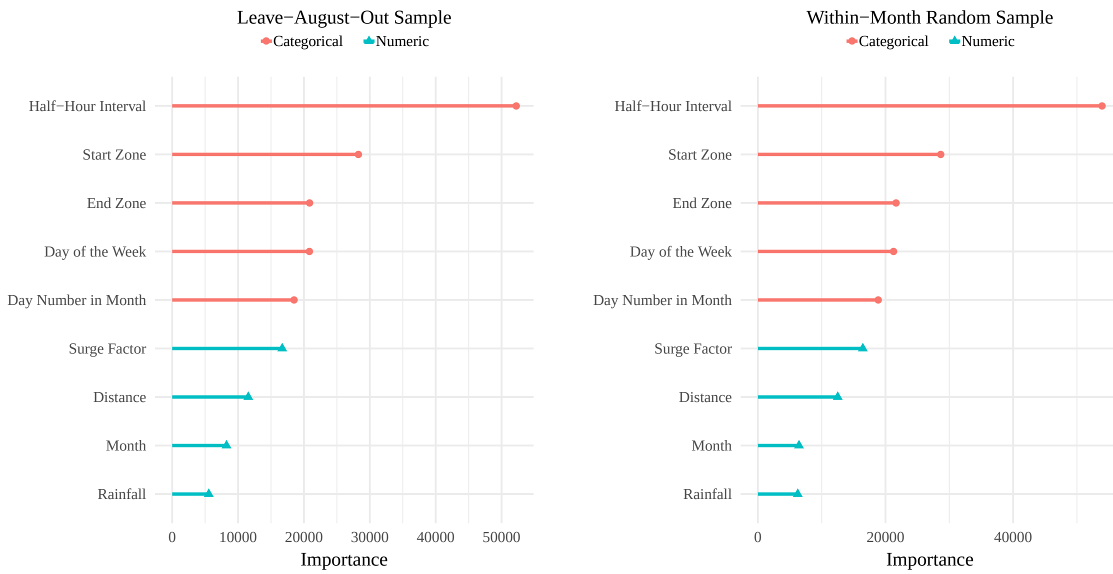

Figure 3: Variable Importance of Taxi Bookings Predictive Model. These figures show the variable importance ranking based on node impurity for each of our two samples. The left panel is for the sample excluding August and the right panel is for the within-month 75% random sample.

Figure 3 shows the variable importance plot for both the samples, splitting variables into categorical and numerical variables by color. In both, the surge factor is the 6th most important variable. The most important variables, expectedly, are the time, day of week, and destination variables. In the training sample using all the data before August, the proportion of variance captured by the random forest model is 49.62% for the model with the surge factor and 48.98% without. Likewise, the in-sample improvement using the random subsample is also small from 47.51% without the surge factor to 47.89% with the surge factor. Thus, in this R^2 sense, the in-sample improvement due to the surge factor is small.

However, surge factor information increases out-of-sample forecast accuracy by 12% to 15%.

[19]Figure A.9 in the Appendix shows the sensitivity of in-sample root-mean squared errors based on the number of trees used.

Table 4: Accuracy of Taxi Bookings Predictive Model. This table shows the accuracy of the random forest model in predicting the number of taxi bookings with different samples (both in-sample and out-of-sample). RMSE stands for root-mean-squared error.

	In-Sample RMSE		
Model Test Sample	No Surge Factor	With Surge Factor	Improvement (%)
August	1.675	1.664	0.661
20% of Each Month	1.732	1.726	0.348

	Out-of-Sample RMSE		
Model Test Sample	No Surge Factor	With Surge Factor	Improvement (%)
August	1.286	1.089	15.319
20% of Each Month	1.488	1.306	12.231

Table 4 shows the in-sample and out-of-sample prediction accuracies of the random forest model for both the leave-August-out and within-month random samples. For the leave-August-out sample, we see that although including the surge factor improves the in-sample accuracy by less than 1%, it raises the out-of-sample accuracy by over 15%. Likewise with the within-month random sample, the model with surge factor data improves the in-sample accuracy by less than 0.5%, but the out-of-sample accuracy rises by over 12%. While the performance is similar with or without the surge factor, the factor substantially raises the out-of-sample prediction accuracy by eliminating systematic errors in the forecasts. Since out-of-sample accuracy is typically the main metric in forecasting problems, our results suggest that the surge factor helps policy makers as well as taxi and ride-hailing companies in forecasting the future taxi bookings.

Although in practice one could query the Grab surge factor in real time, for robustness, we also consider including the half-hour lagged surge factor. Appendix Table A.6 shows this specification, documenting similar out-of-sample performances in the root mean-squared error of predicted taxi bookings of 11.5% and 15% for the out-of-sample periods corresponding to August or a random sample of 20% of each month respectively. This similar out-of-sample performance is likely due to the autocorrelation in the surge factor. In addition, the autocorrelation in the surge factor means there is a persistent mis-match between Grab supply and demand, therefore the autocorrelated component of the surge factor remains highly relevant for predicting taxi bookings.

However, having better taxi demand prediction itself does not directly contribute to the increase in social welfare. In the section below, we operationalize the demand prediction information to study impact on social welfare by providing guidance to taxi drivers. Such guidance should take into account both the predicted demand occurrence in different parts of the cities and the current and predicted future locations of vacant taxis, and generate a suggested region for the driver to go to in the upcoming time period.

5.2 Structural Analyses and Simulation

To shed light on how the surge price information impacts social welfare, we focus on the reduction of congestion and taxi vacancy times. A more accurate taxi booking demand prediction could help policy makers and ride-hailing and taxi companies to anticipate demand better and reduce congestion. Given a taxi fleet size, decreasing congestion and taxi vacancy times increases both consumer welfare and driver welfare. Therefore, by focusing only on the latter, we will underestimate the impact of the surge factor information on social welfare.

For implementation, the policy makers or companies have to collect both the predicted demands and the whereabout of vacant taxis, and compute policies on how to position vacant taxis to optimize spatiotemporal demand-supply matching. An academic prototype of this idea, the Driver Guidance System (DGS) (Jha et al., 2018), has been developed and its field trial results with around 500 taxi drivers show that by following centrally generated recommendations, taxi drivers experience, on average, 28.9% reduction in their vacant roaming time from January to May 2018 (Cheng et al., 2018). Further, if the drivers hypothetically follow the DGS guidance all the time, the number of fetched trips could rise by around 13.5% (we provide details on how this is estimated in Appendix D). To illustrate the information needed for the implementation, we present a stylized model assuming all taxi trips have the same distance and fare, and total taxi supply and demand are balanced. The locations of the demand quantities are random, but could be inferred by observing the surge prices at different locations. This stylized coordination problem for the taxi company can be written as follows:

$$\max_{m_1, m_2, \ldots, m_n} \sum_{i=1}^{n} E\left[\min\{m_i, D_i\}\right] \quad \text{such that } \sum_{i=1}^{n} m_i = m,$$

where n is the number of locations, m is the number of taxis that equals the realized total residual taxi demand quantity, and m_i and D_i are the number of taxis and demand quantity respectively. Let us denote surge price in location i by s_i.

Without the surge price information, D_i is a random variable (but the aggregate demand quantity is fixed and known). Therefore, in this case, the probability for all the customers to find matching taxis is less than one, i.e., $Pr(\text{number of rides} = m) < 1$. However, if the surge price information is incorporated, D_i is inferred and, thus, is no longer a random variable. In other words, instead of having $E[\min\{m_i, D_i\}|s_i]$ in the objective function above, we have $\min\{m_i, D_i\}$. Therefore, with the surge price information, the optimal allocation of taxis is simply $m_i = D_i$ for all $i \in \{1, 2, ..., n\}$, which gives that the number of rides equals m. In this case, all the rides are matched, which maximizes the welfare of consumers within the ride-hailing and taxi-booking market.[20]

Assuming that a taxi driver's objective is revenue maximizing, where the related customer welfare is measured in terms of the numbers of ride-hailing and booked taxi rides, we also see that the impact of taxi companies using surge factors on Grab cars is minimal, as surge factors were already determined due to the matching of Grab supply and demand. A high surge factor implies high Grab demand relative to supply, suggesting that Grab cars will not be able to service all the demand. Of course, this analysis is single-period and does not consider, for instance, the travel time experienced by taxis. To extend this stylized analysis to a more realistic setting, we conduct detailed agent-level simulations, where we evaluate how the improvement of taxi booking demand prediction would positively impact close-to-reality taxi fleet operations.

To incorporate taxi booking demand prediction and observe how it affects taxi driver's decision making process, we adopt a DGS following (Jha et al., 2018), and use a realistic microscopic taxi operation simulation (Cheng and Nguyen, 2011) to evaluate the resulting gains in the social welfare (computed as all taxis' income) if demand predictions were to be improved. Our simulation is mi-

[20]However, as noted before in Section 4.2, the overall consumer welfare in all related markets comprises those from riders of both ride-hailing services, taxi bookings, and taxi street-pick-ups. Since we document some negative spillover effects of the relation between ride-hailing and taxi bookings on street pick-ups, the overall welfare is difficult to quantify. Therefore, we take the objective to maximize the expected revenues of the drivers, and the welfare is measured in terms of the numbers of rides of ride-hailing services and booked taxis.

croscopic and agent-based, where each taxi is modeled as an agent; we also model the service modes of street hails, taxi queues, and booking. As the simulation runs, it also generates a data stream in real-time to feed to the demand-supply matching engine that computes taxi-specific movement recommendations (based on the demand predictions using the simulated data stream). By introducing a parameter that deliberately controls the demand prediction accuracy, we then observe the impact of improving demand prediction on the guidance effectiveness. As the computation of driver guidance depends on the demand prediction as well, we conduct a simulation study to test how much the improvement in the demand prediction raises the effectiveness of the spatiotemporal demand-supply matching. In other words, this counterfactual simulation uses the existing simulation from (Jha et al., 2018) and exogenously improves the demand prediction accuracy as a way to gauge the economic impact.

The reason why the performance of DGS depends on the accuracy of demand prediction is that there are three major components in the optimization formulation: the immediate movement cost, the expected future revenue, and the expected future movement cost. Better demand prediction should contribute to the estimation accuracy of the expected future revenue. We design the simulation to demonstrate the potential benefits of having better taxi demand prediction. For the simulation, we make the following assumptions:

1. We assume that all taxis in the simulation follow the generated guidance; according to the simulation results in (Jha et al., 2018), having all taxis following the guidance generates greatest social welfare. This assumption allows us to estimate the upper bound on the gains that could result from demand prediction.

2. We use the Singapore map grid in Figure 4 to setup the simulation. To provide appropriate granularity, we define the grid of 1km-by-1km to be the minimal geographic unit for demand and supply prediction, and the target of the recommendation. Based on this grid definition, we then calculate realistic geographical features such as traveling distances and cost between grid regions.

3. The demand pattern in the simulation is derived from historical weekdays. For each grid i,

Figure 4: Grid map for Singapore used in the simulation.

each time period t (30 minutes), we calculate the Poisson arrival rate (λ_{it}) as the average number of trips occurring in grid i, at time t. During the actual simulation, Poisson arrivals are simulated in each grid region with the assumption that the arrival rate is λ_{it}. The demand predictions are generated every 5 minutes, for the horizon of 6 time periods (30 minutes). We pre-generate all demand arrivals before the simulation, and store them as the ground truth. During the simulation, the simulated demand arrivals are read from this pre-generated source. The demand predictions are computed on the fly during the simulation, using only realized and observable information. To improve the prediction accuracy, we compare the predicted demand against the pre-generated ground truth and adjust the predicted value so that its difference from the ground truth is improved by α. In other words, if d_{it} is the actual demand (unknown to the driver), $\hat{d}_{it}$ is the predicted demand using a deep learning neural network model using only point-in-time observable variables, and the adjusted demand prediction is given by:

$$\hat{d}'_{it} = d_{it} + (\hat{d}_{it} - d_{it}) \times (1 - \alpha),$$

where the original Forecast Error ($\hat{d}_{it} - d_{it}$) is from the empirical distribution of actual forecast errors relative to the original demand prediction. This $(1-\alpha)$ term corresponds to a decrease in root-mean-squared error of $\alpha\%$ relative to the baseline model.[21]

[21] Relative to an original forecast error $var(\hat{d}_{it} - d_{it}) = \sigma^2$, the variance of the adjusted error term is $var((\hat{d}_{it} - d_{it}) \times (1-\alpha)) = \sigma^2 \times (1-\alpha)^2$. The relative root-mean-squared error is $\sqrt{(1-\alpha)^2\sigma^2}/\sqrt{\sigma^2} - 1 = -\alpha$, a decrease of $\alpha\%$.

4. The guidance from the DGS indicates on the grid region in which a taxi should stay in. The actual movements along the streets are decided by historical frequency: when a simulated taxi reaches a road intersection, the simulator queries the historical frequency on which road segment to turn, and stochastically decides which road segment the taxi should drive to. The constraint is that the choice at the street-level should ensure that the grid-level decision is maintained.

We compare the predicted demands against the corresponding demands. More specifically, we alter the predictions so that they are α percent closer to the corresponding simulated demands (α is a parameter for the experiment, representing the desired improvement percentage of demand prediction accuracy). As the simulated demand scenarios are all pre-generated, we conduct paired performance comparisons of different demand prediction accuracies. To reflect the improvement in demand prediction accuracies, we have compared the cases with $\alpha = 0$ (the baseline without improvement) and $\alpha = 0.15$. We generate 20 demand scenarios and show that with $\alpha = 0.15$, we reduce the vacant roaming time of guided taxis by close to 9.4%, which leads to a 2.3% increase in the number of trips. The improvement is statistically significant with $p-$value close to 0.

Extrapolating these simulation results to the real-world field trials reported in Cheng et al. (2018) using the same DGS mechanism, we further discover that the vacant roaming time of taxis could be lowered by 35.6% if taxis receive and follow guidance generated with $\alpha = 0.15$. If this reduction in the vacant roaming time is realized in the real-world dataset, it could lead to an increase of 17.2% in the number of trips. The details for this back-of-the-envelope calculation is available in Appendix Section D which shows the counterfactual simulated quantities based on the datasets and results of Cheng et al. (2018).

6 Conclusion

We estimate an upper-bound on the substitution between ride-hailing services and taxi bookings of 0.26. Despite the ease through which consumers can switch between mobile apps for ride-hailing services and taxi bookings, the demand for bookings does not appear elastic. A back-of-the-envelope calculation based on the average distance of travel suggests that a consumer who switches from ride-hailing to booking a taxi when surge factors are higher than one saves about 18%. The estimated inelastic demand may be due to marketing tactics in maintaining brand loyalty, inattentive customers, or alternative mechanisms. For example, ride-hailing companies provide discounts and other incentives for customers to remain on their platform. We leave the impact of other specific features of ride-hailing services for future research.

However, despite the low cross-price elasticity, we find that surge factor information appears statistically useful for predicting taxi booking demand, improving prediction accuracy by 15%. Our machine-learning results suggest incorporating surge factor information in real-time across different regions can potentially increase allocative efficiency.

Our structural analyses and simulations suggests that a 15% improvement in the accuracy of demand predictions leads to a 9.4% reduction in the average vacant roaming time. The average number of trips per taxi, which we used as a welfare measure in the previous subsection, also rises by 2.3%. Compared against non-guided taxis, this amounts to a reduction of 35.6% on the vacant roaming time and a 17.2% increase in the number of trips.

References

P. Afeche, Z. Liu, and C. Maglaras. Ride-hailing networks with strategic drivers: The impact of platform control capabilities on performance. *Working Paper*, 2018.

B. Barann, D. Beverungen, and O. Müller. An open-data approach for quantifying the potential of taxi ridesharing. *Decision Support Systems*, 99:86–95, 2017.

T. Berger, C. Chen, and C. B. Frey. Drivers of disruption? Estimating the Uber effect. *Working Paper*, pages 1–11, 2017.

Z. Bian and X. Liu. Mechanism design for first-mile ridesharing based on personalized requirements part i: Theoretical analysis in generalized scenarios. *Transportation Research Part B: Methodological*, 120:147–171, 2019a.

Z. Bian and X. Liu. Mechanism design for first-mile ridesharing based on personalized requirements part ii: Solution algorithm for large-scale problems. *Transportation Research Part B: Methodological*, 120:172–192, 2019b.

K. Bimpikis, O. Candogan, and D. Saban. Spatial pricing in ride-sharing networks. *Operations Research*, 67(3):744–769, 2019. ISSN 15265463. doi: 10.1287/opre.2018.1800.

C. Borgs, O. Candogan, J. Chayes, I. Lobel, and H. Nazerzadeh. Optimal multiperiod pricing with service guarantees. *Management Science*, 60(7):1792–1811, 2014.

G. P. Cachon, K. M. Daniels, and R. Lobel. The role of surge pricing on a service platform with self-scheduling capacity. *Manufacturing and Service Operations Management*, 19(3):368–384, 2017. ISSN 15265498. doi: 10.1287/msom.2017.0618.

J. C. Castillo, D. Knoepfle, and G. Wyle. Surge Pricing Solves the Wild Goose Chase. *Working Paper*, (January), 2017.

S.-F. Cheng and T. D. Nguyen. Taxisim: A multiagent simulation platform for evaluating taxi fleet operations. In *2011 IEEE/WIC/ACM International Conferences on Web Intelligence and Intelligent Agent Technology*, volume 2, pages 14–21, 2011.

S.-F. Cheng, S. S. Jha, and R. Rajendram. Taxis strike back: A field trial of the driver guidance system. In *Seventeenth International Conference on Autonomous Agents and Multiagent Systems*, pages 577–584, 2018.

L. Y. Chu, Z. Wan, and D. Zhan. Harnessing the double-edged sword via routing: Information provision on ride-hailing platforms. *Working Paper*, 2018.

P. Cohen, R. Hahn, J. V. Hall, S. Levitt, and R. Metcalfe. Using Big Data to Estimate Consumer Surplus: The Case of Uber. *NBER Working Paper*, pages 1–43, 2016.

J. Cramer and A. B. Krueger. Disruptive change in the taxi business: The Case of Uber. *American Economic Review Papers & Proceedings*, 106(5):177–182, 2016. ISSN 00028282. doi: 10.1257/aer.p20161002.

K. J. Ferreira, B. H. A. Lee, and D. Simchi-Levi. Analytics for an online retailer: Demand forecasting and price optimization. *Manufacturing and Service Operations Management*, 18(1):69–88, 2016. ISSN 15265498. doi: 10.1287/msom.2015.0561.

C. K. Glaeser, M. Fisher, and X. Su. Optimal retail location: Empirical methodology and application to practice. *Manufacturing and Service Operations Management*, 21(1):86–102, 2019. ISSN 15265498. doi: 10.1287/msom.2018.0759.

A. Javanmard and H. Nazerzadeh. Dynamic pricing in high-dimensions. *Journal of Machine Learning Research*, 20(9):1–49, 2019.

A. Javanmard, H. Nazerzadeh, and S. Shao. Multi-product dynamic pricing in high-dimensions with heterogenous price sensitivity. *Working Paper*, 2019.

S. S. Jha, S.-F. Cheng, M. Lowalekar, W. H. Wong, R. R. Rajendram, T. K. Tran, P. Varakantham, N. Truong Trong, and F. Abd Rahman. Upping the game of taxi driving in the age of Uber. In *Thirtieth Annual Conference on Innovative Applications of Artificial Intelligence*, pages 7779–7785, 2018.

J. Ke, H. Zheng, H. Yang, and X. M. Chen. Short-term forecasting of passenger demand under on-demand ride services: A spatio-temporal deep learning approach. *Transportation Research Part C: Emerging Technologies*, 85:591–608, 2017.

S. Liu and Z.-j. M. Shen. On-Time Last Mile Delivery : Order Assignment with Travel Time Predictors. *Management Science*, pages 1–51, 2020.

A. Lu, P. I. Frazier, and O. Kislev. Surge pricing moves uber's driver-partners. In *Nineteenth ACM Conference on Economics and Computation*, pages 3–3, 2018.

M. Mahdian, H. Nazerzadeh, and A. Saberi. Allocating online advertisement space with unreliable estimates. In *Eighth ACM Conference on Electronic Commerce*, pages 288–294, 2007.

E. Martin and S. Shaheen. The impact of carsharing on public transit and non-motorized travel: An exploration of North American carsharing survey data. *Energies*, 4(11):2094–2114, 2011. ISSN 19961073. doi: 10.3390/en4112094.

L. Moreira-Matias, J. Gama, M. Ferreira, J. Mendes-Moreira, and L. Damas. Predicting taxi–passenger demand using streaming data. *IEEE Transactions on Intelligent Transportation Systems*, 14(3):1393–1402, 2013.

H. A. Posen. Ridesharing in the sharing economy: Should regulators impose regulations on Uber? *Iowa Law Review*, 101(1):405–433, 2015. ISSN 00210552.

H. Ross. Ridesharing's House of Cards: O'Connor V. Uber Technologies, Inc. and the Viability of Uber's Labor Model in Washington. *Washington Law Review*, 90:1431–1469, 2015.

R. Seamans and F. Zhu. Responses to entry in multi-sided markets: The impact of craigslist on local newspapers. *Management Science*, 60(2):476–493, 2014. ISSN 00251909. doi: 10.1287/mnsc.2013.1785.

T. A. Taylor. Manufacturing & Service Operations Management On-Demand Service Platforms. *On-Demand Service Platforms*, 20(4):704–720, 2018. ISSN 1526-5498. doi: 10.1287/msom.2017.0678.

H. Yao, F. Wu, J. Ke, X. Tang, Y. Jia, S. Lu, P. Gong, J. Ye, and Z. Li. Deep multi-view spatial-temporal network for taxi demand prediction. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.

G. Zervas, D. Proserpio, and J. W. Byers. The rise of the sharing economy: Estimating the impact of airbnb on the hotel industry. *Journal of Marketing Research*, 54(5):687–705, 2017. ISSN 00222437. doi: 10.1509/jmr.15.0204.

Online Appendix

Appendix A Data

Figure A.1 shows the origins as circles and destinations as squares around Singapore that we sampled. The origin and destination zones are fairly representative of the level of transportation and economic activity in Singapore, which is more concentrated towards the south of Singapore, in the Central Business District.

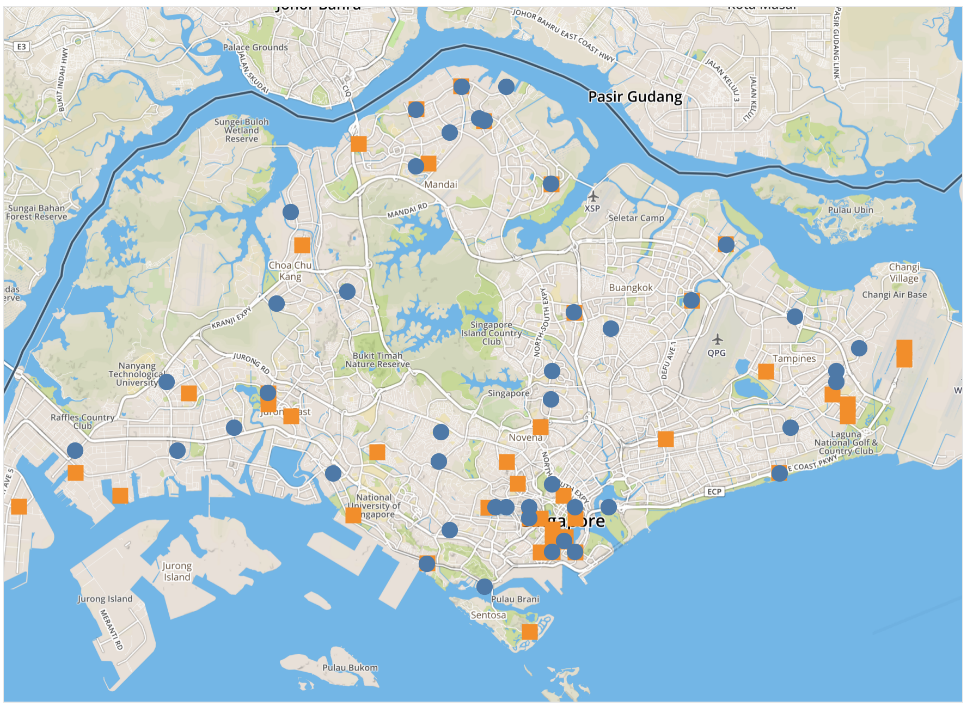

Figure A.1: Origins and Destinations Queried for Prices on the Grab App. Origins and destinations are denoted as circles and squares respectively.

Empirically, as shown in Figure A.4, surge factors are autocorrelated. The gradual decrease in autocorrelation and truncation in partial autocorrelation after one hour (2 lagged half-hour intervals) suggests an AR(2) process for the surge factor.

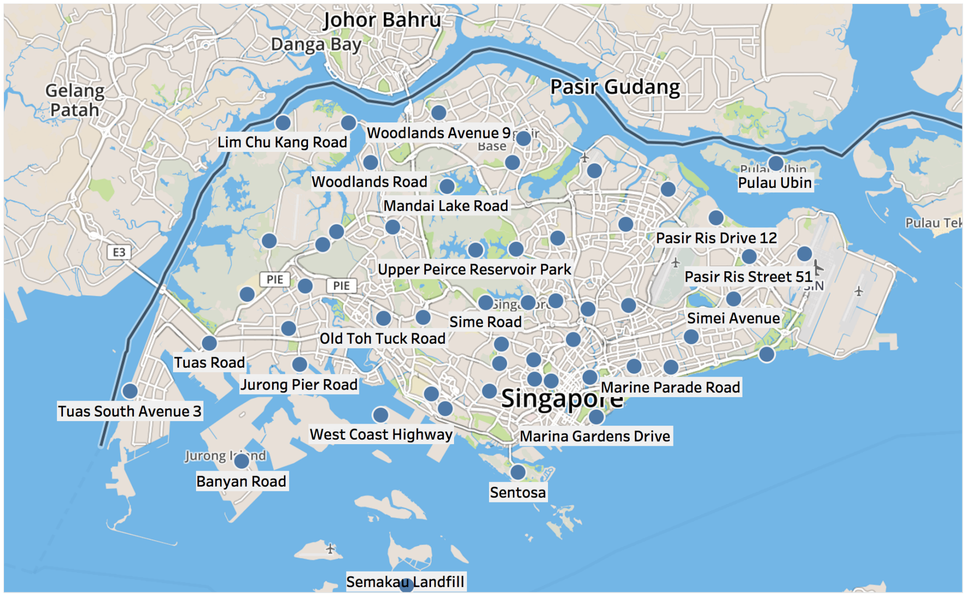

Figure A.2: Locations of All Weather Stations Around Singapore.

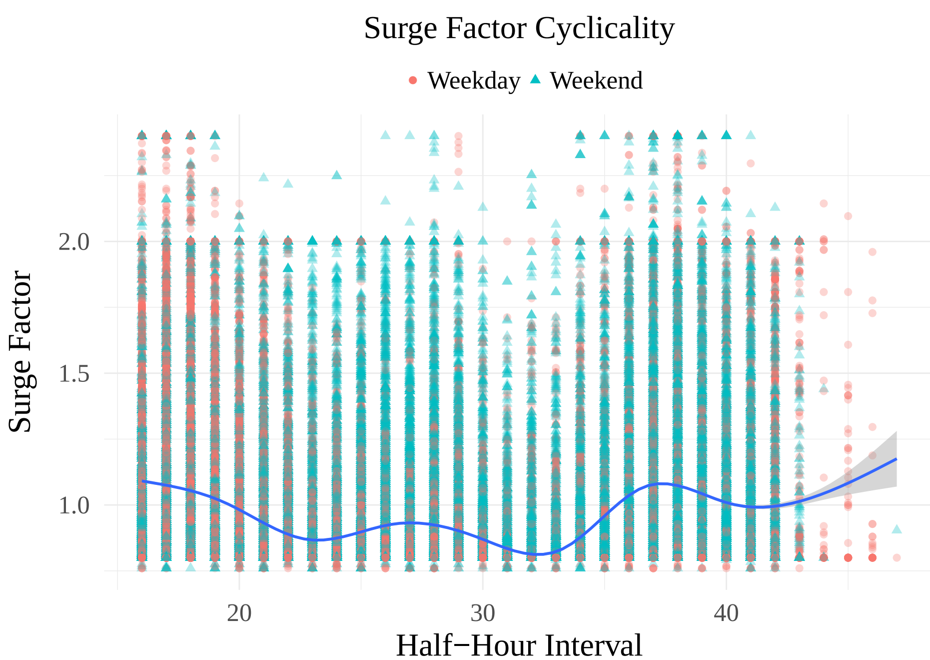

Figure A.3: Surge Factors During the Day. This figure shows a scatter plot of surge factors across different half-hour intervals in the day. The line in the plot is a non-parametric line of best fit estimated using a LOESS algorithm which uses a tri-cube weighting function to categorize local areas. We find some seasonality in surge factors across different hours of the day, corresponding to intuitive notions of peak hours such as morning, lunch, and evening.

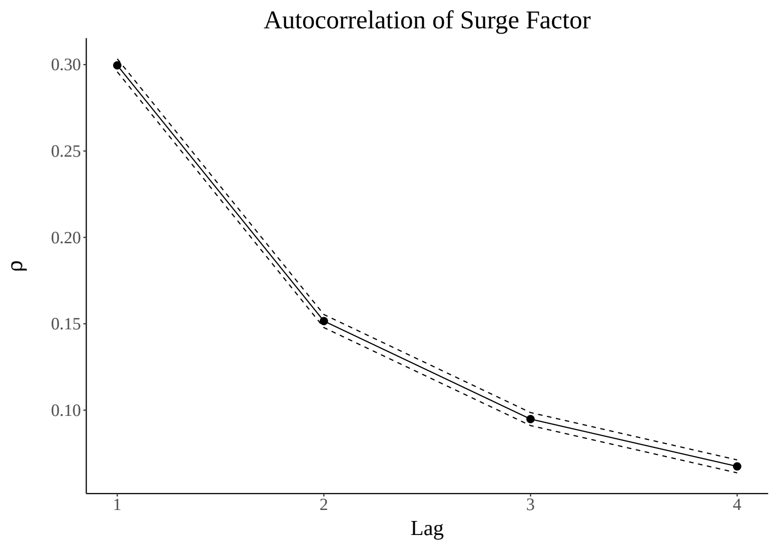

Figure A.4: Autocorrelation of Surge Factors. This figure shows the autocorrelation of surge factors estimated from the equation $SurgeFactor_{o,t,i} = \alpha_{o,i} + Weekend_t + \rho_\tau SurgeFactor_{o,t-\tau,i} + \varepsilon_{o,t,i}$ where o is the origin, t is the date, i is the half-hour interval, and τ is the lag, from 1 through 4 (corresponding to a half-hour lag through two-hour lag). The dotted lines represent the 95% confidence interval.

Appendix B Robustness

We run additional analyses for robustness. First, Table A.1 tests whether the relation between surge factors are linear. Second, we also conduct additional falsification tests in Table A.2 by dropping weekends, randomly sampling surge factors, randomly sampling from surge factors within the same time interval, and randomly sampling from surge factors within the same route. None of these specifications generate the same results, suggesting that our estimated cross-price elasticity is not generated by confounding correlations across areas or across time intervals.

Table A.1: Linearity of Price Elasticity. This table shows the linearity of the price elasticity as well as tests for differences in elasticities of short and long trips. The observations are at the starting zone by half-hour interval by end zone level. All regressions include zone by time interval by day-of-week fixed effects. Standard errors (in parentheses) are clustered by zone and half-hour interval.

	$y_{o,d,t,i} = \log\left(\text{Taxi Bookings}_{o,d,t,i}\right)$				
	(1)	(2)	(3)	(4)	(5)
Surge Factor_t	0.181***	0.497***	0.513***		0.167***
	(0.030)	(0.120)	(0.129)		(0.033)
Surge Factor_t^2		-0.121***			
		(0.038)			
$1\{\text{Surge Factor}_t > 1\}$			0.382***	0.104***	
			(0.109)	(0.017)	
Surge $\text{Factor}_t \times 1\{\text{Surge Factor}_t > 1\}$					-0.046
					(0.085)
Long Distance			-0.398***		
			(0.120)		
Surge $\text{Factor}_t \times$ Long Distance					0.031
					(0.040)
Taxi Supply_t	-0.013	-0.012	-0.013	-0.014	-0.012
	(0.035)	(0.035)	(0.035)	(0.035)	(0.035)
Observations	85,554	85,554	85,554	85,554	85,554
R^2	0.405	0.405	0.406	0.404	0.405

*p<0.1; **p<0.05; ***p<0.01

Table A.2: Robustness Checks. This table shows the response of the log number of taxi bookings in response to the level of the surge factor. Column (1) shows the baseline result, column (2) shows the elasticity of taxi bookings when excluding weekends, column (3) runs the same specification as the baseline result with a randomly assigned surge factor, column (4) repeats the same falsification test sampling from surge factors within the same time interval across routes, column (5) repeats the same falsification test sampling from surge factors within the same route across time, and column (6) tests whether the cross price elasticity of bookings with respect to surge factor depends on being a weekend. The observations are at the starting zone by half-hour interval level. All regressions include zone by time interval by day-of-week fixed effects. Standard errors (in parentheses) are clustered by zone and half-hour interval.

Model/Sample:	(1) Baseline	(2) No Weekend	(3) Random Surge Factor	(4) Random Surge Factor from Same Time	(5) Random Surge Factor from Same Route	(6) Full
Surge Factor	0.256***	0.281***	-0.015	0.013*	-0.015	0.279***
	(0.027)	(0.030)	(0.012)	(0.008)	(0.009)	(0.030)
Surge Factor×Weekend						−0.048
						(0.037)
Taxi Supply	0.032	0.109*	0.028	0.028	0.027	0.032
	(0.008)	(0.061)	(0.047)	(0.047)	(0.007)	(0.047)
Observations	57,474	34,901	57,474	57,474	57,474	57,474
R^2	0.525	0.532	0.521	0.521	0.521	0.525
Note:						*p<0.1; **p<0.05; ***p<0.01

Despite controlling for taxi and Grab supply, if the estimated correlations are due to confounding factors driving both the supply and demand of taxi bookings, a higher surge factor between ride-hailing services and taxis should increase street pick-ups as well. We find that surge factors are actually related to a decrease in street pick-ups. Hence, any confounding factors driving our results would need to consistently impact the supply and demand for taxi bookings but have a slightly negative impact on street-side taxi rides. Since such a confounder is unlikely, we conclude that we are indeed identifying a substitution effect from ride-hailing services to taxi bookings when the relative price of ride-hailing services is high.

Table A.3: Public Signal Results. This table shows the response of the log number of taxi bookings in response to the level of the surge factor. Column (1) shows the baseline result, column (2) shows the interaction of surge factor and cyclical zones (zone cyclicality is subsumed by zone by time interval fixed effects), column (3) shows the interaction of surge factor with rain, and column (4) shows both. Column (5) shows the effect on street pick-ups. The observations are at the starting zone by half-hour interval level. All regressions include zone by time interval by day-of-week fixed effects. Standard errors (in parentheses) are clustered by zone and half-hour interval.

	log Taxi Bookings$_t$				log Taxi Street Pick-ups$_t$
	(1)	(2)	(3)	(4)	(5)
Surge Factor	0.306***	0.256***	0.277***	0.313***	-0.064***
	(0.032)	(0.022)	(0.031)	(0.026)	(0.025)
Cyclicality×Surge Factor	0.032		0.020		0.032
	(0.042)		(0.043)		(0.055)
Rain	0.199***			0.196***	-0.044
	(0.046)			(0.043)	(0.047)
Rain×Surge Factor	-0.066			-0.063	-0.053
	(0.043)			(0.042)	(0.043)
Observations	53,760	53,760	53,760	53,760	68,309
R^2	0.400	0.412	0.385	0.400	0.606
Note:					*p<0.1; **p<0.05; ***p<0.01

During rainy time periods and zones, the surge factor is 16% higher than normal, bringing the surge factor of Grab from an average of 0.9 to 1.05, slightly higher than taxis in the standard pricing period. In periods of heavy rain, defined as rainfall in the top 10% of rainfall of about 10 cm (4 inches), the surge factor increases by about 30%, from 0.9 to 1.2.

We also test whether publicly available signals and market conditions affect the substitution. Table A.3 reports our baseline results using the raw surge factors. Our reduced-form estimate in Column (2) suggests that unconditionally, an increase of 10% in the surge factor leads to a 2.86%

Figure A.5: Price Elasticity across Months. This figure shows elasticity estimates for each month across our sample.

increase in taxi bookings. When controlling for the cyclicality of surge factors in a zone and whether there is rain in a half-hour interval, we find that the cross-price elasticity increases to 3.06%, shown in Column (1). We also include the interaction of zone cyclicality and rain with surge factors, since rain and cyclicality are common knowledge to both taxi and Grab supply.[22] Drivers anticipating more riders in an area and time with rain may anticipate higher demand without having information on surge factors. However, we find no statistically significant interaction between zone cyclicality or rain with surge factors.

B.1 Transitory Demand Shocks

Rather than immediately switching from booking a ride-hailing service car to booking a taxi, consumers can also instead wait for surge factors to decrease. In this subsection, we consider two specifications: (*i*) a distributed lag model that isolates past demand shocks by controlling for taxi and grab supplies at different lags within a region, and (*ii*) an impulse response representation that studies how the impact of an isolated demand shock on taxi bookings propagates over time.

In the first analysis, we find that surge factors from half an hour ago still increase current rides

[22]Rain is fairly difficult to predict in Singapore at only about 30 minutes to 1 hour ahead.

but the impact is short-lived beyond that, controlling for the supply factors. Comparing Column (2) to Column (1) in Table A.4, we find the point estimate on contemporaneous surge factors decrease from 0.396 to 0.362 with a positive and statistically significant coefficient on the first lagged interval term of 0.072. Including additional lags does not change the contemporaneous or the first lagged interval term much. By Table A.4, we conclude that only surge factors half an hour in the past continue to have a statistically significant effect on the current taxi bookings.

Table A.4: Lagged Surge Factor. This table shows the response of the log number of taxi bookings in response to the level of the surge factor, distributed up to two-half hour lags. The observations are at the starting zone by half-hour interval level. All regressions include zone by time interval by day-of-week fixed effects. Standard errors (in parentheses) are clustered by zone and half-hour interval.

	log(Taxi Bookings)		
	(1)	(2)	(3)
Surge Factor$_t$	0.284***	0.275***	0.264***
	(0.033)	(0.061)	(0.062)
Surge Factor$_{t-1}$		0.010	-0.039
		(0.060)	(0.068)
Surge Factor$_{t-2}$			0.071*
			(0.035)
Taxi Supply$_t$	0.028	0.028	0.029
	(0.051)	(0.051)	(0.052)
Observations	36,860	36,860	36,860
R^2	0.507	0.507	0.508

*p<0.1; **p<0.05; ***p<0.01

In the second analysis, we analyze the impact of surge factors on future surge factors, as well as future taxi bookings. Figure A.6 shows that an increase in the surge factor controlling for taxi and Grab supply impacts both future bookings and future surge prices (consistent with A.4 in the Appendix). However, the lagged impact of the surge factor on taxi bookings decreases to zero after one hour, consistent with the surge factor being a transitory shock. More specifically, the impact of a demand shock on taxi rides decays from 0.4% to zero over the next two and a half hours. The impact of a demand shock on future surge factors also decays from 0.46 half an hour after the shock down to 0.1 over the next two and a half hours.

Our transitory intertemporal results are consistent with consumers who self-select into the taxi or ride-hailing market and who are time inelastic. After all, given the small size of Singapore, any

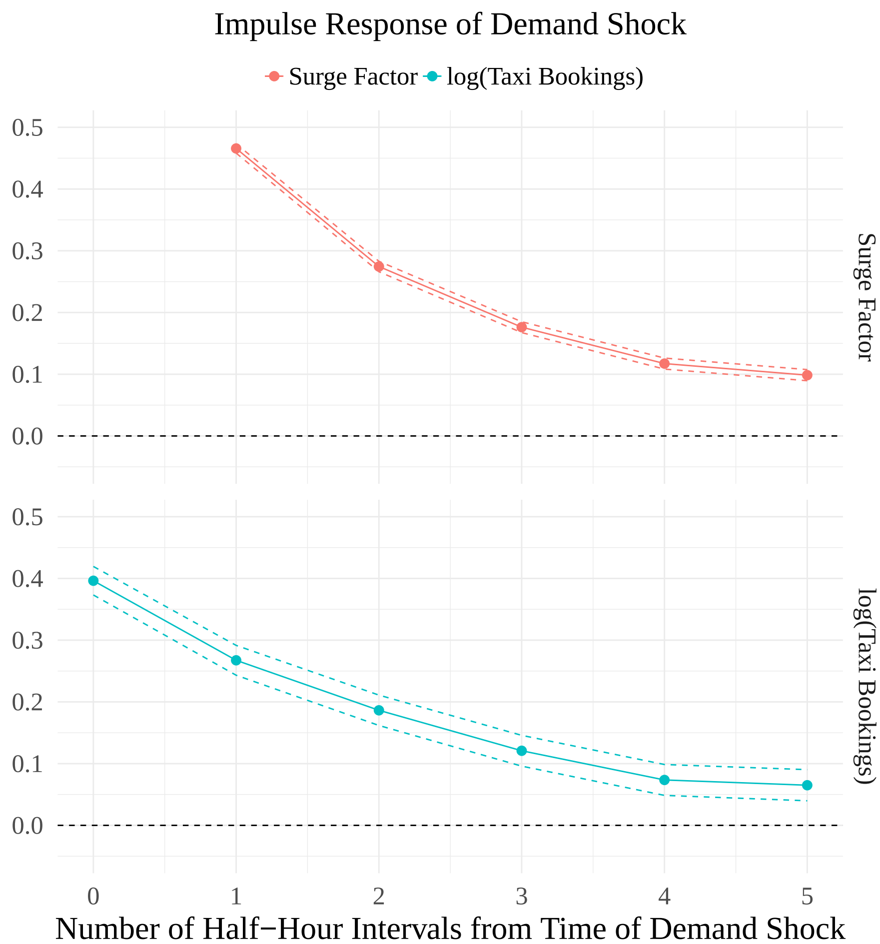

Figure A.6: Impulse Response of Taxi Bookings. This figure shows the impulse response of log taxi bookings for each period from t to $t + 5$ due to a surge factor shock in period t controlling for taxi and Grab supply at time period t. The dotted lines show two-standard error bands.

two destinations on the island are reachable with alternative forms of transportation within one and a half hours. Thus, those taking private-hire cars or taxis are likely not willing to wait for a taxi or a Grab for over half an hour. However, partly because current surge factors affect future surge factors, we find that the total effect of a 10% increase in surge factors due to demand on taxi bookings is about 8% over the next two and a half hours.

In these following two subsections, we consider two additional cross-sectional tests to explore whether local average income of residents in an area and the distance of trips affect our estimated cross-price elasticity. Areas with lower incomes should show a higher substitution between ride-hailing services and taxi bookings when surge factors are high. Moreover, since the surge factor is applied uniformly across the whole trip, consumers taking a long trip should be more sensitive to the surge factor since a longer trip, for any given surge factor, means a larger dollar impact of the surge charge.

B.2 Distance

Taxi fares include both a fixed cost and variable cost. Moreover, booking taxis increases the fixed cost component of taking taxis. On the other hand, the marginal benefit of using a taxi meter is that their price per unit of distance is fixed for a given point in time.[23] Customers going on longer trips reap a larger absolute gain from using the taxi meter compared to ride-hailing service because surge prices are applied over the whole trip and the fixed cost component of the taxi fare is averaged over a longer distance. Therefore, we expect customers going on longer journeys to be more price elastic.

We test this hypothesis in the data by grouping rides into short and long distances. We define short distance as journeys between two zones that are less than 5 kilometers away from each other. Figure A.7 shows that there seems to be no difference in the relation between surge factors and taxi bookings based on distance. In untabulated regression results, we find no statistically significant evidence that the price elasticity of longer distances differ from price elasticity of shorter distances. The results persist whether we consider the distance between two zones both as an indicator of a long trip or as a continuous variable.

B.3 Impact of Income

Out of the 34 regions for which we have sufficient data, all statistically significant results point to higher surge factors increasing taxi bookings. Although 33 out of 34 regions have a positive point estimate, there is large heterogeneity in cross price elasticities across regions, from a low of -0.07 to a high of 1.57 using our baseline specification. A potential explanation of the price elasticity is income. Areas with higher income may have less price sensitivity due to their higher disposable income.

[23] Recall that taxis in Singapore also have "static" surge factors, either 25% or 50% of metered fare depending on the time of day. We label this static because it does not adjust to changing market conditions on a real-time basis.

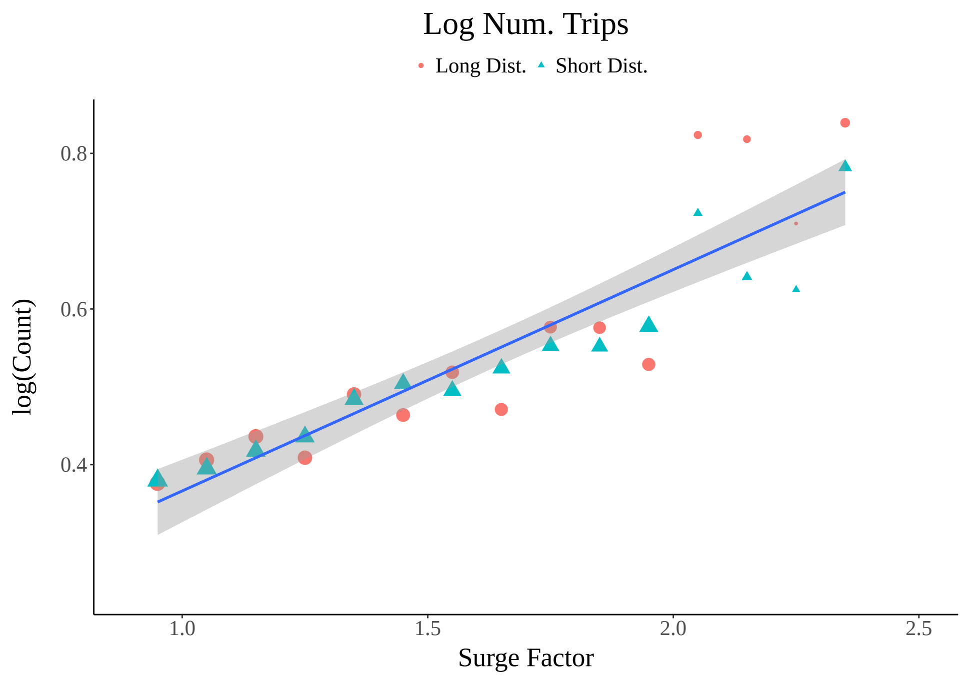

Figure A.7: Taxi Bookings and Surge Factor. The scatter plot for the log number of taxi bookings as a function of surge factor. Each data point represents an average log number of taxi rides where surge factor is split into bins with 0.1 increments. The grey bands represent the 95% confidence interval.

For this analysis, we rely on a proprietary data set for 2010 income for households with a credit card from a regional bank. We do not find that the heterogeneity in cross-price elasticities is related to income. Table A.5 shows that a region with 10 percentage points more than average income is only 0.02 percentage points less elastic, a less than 1% reduction in the cross-price elasticity compared to a region with the average income. This suggests that whether a consumer shows fickle fingers and switches easily between transportation apps is not related to their income level. This could be because transportation costs are not large compared to the average Singaporean's overall consumption bundle as well as the easiness to switch between smartphone apps.

Table A.5: Cross-price Elasticity and Income. This table tests whether the cross-price elasticity of taxi rides to surge factors are related to income. Income is defined as the deviation from the average income level in Singapore in 2016. The income variable is absorbed by the origin-destination zone by time fixed effect. The observations are at the starting zone by half-hour interval level. All regressions include zone by time interval by day-of-week fixed effects. Standard errors (in parentheses) are clustered by zone and half-hour interval.

	log Taxi Bookings$_t$
	Baseline
Surge Factor	0.299***
	(0.027)
Surge Factor×Income	-0.0002
	(0.029)
Observations	53,760
R^2	0.412
Note:	*p<0.1; **p<0.05; ***p<0.01

Appendix C Sensitivity of Random Forest Models

This Appendix shows the sensitivity of the reported random forest models. In the main text, we show the model performance for forests with 500 trees. Figure A.9 shows the root-mean squared error performance of the models for different numbers of trees where the errors are calculated based on out-of-bag evaluations. In both samples, we see that with over 100 trees, the in-sample predictive power of the full data including surge prices outperforms the data without surge prices.

C.1 Using Lagged Surge Factor in Random Forest Model

In this subsection, we consider a robustness test of using the surge factor information from the previous half-hour interval in the demand prediction. In this case, we actually find that including the surge factor variable decreases in-sample predictability based on an out-of-bag cross-validated in-sample root mean-squared error. In the leave-August-out specification, we find a decrease in the in-sample performance of 2.4% and in the leave a randomly selected 20% of the sample out of each month, we find a decrease in the in-sample performance of -4.8%.

Meanwhile, as with the contemporaneous surge factor specification, out-of-sample performance improves by around 15% for the leave-August-out specification and 12% for the leave-20%-of-each-month specification. This is likely due to the autocorrelation in the surge factor shown above, and

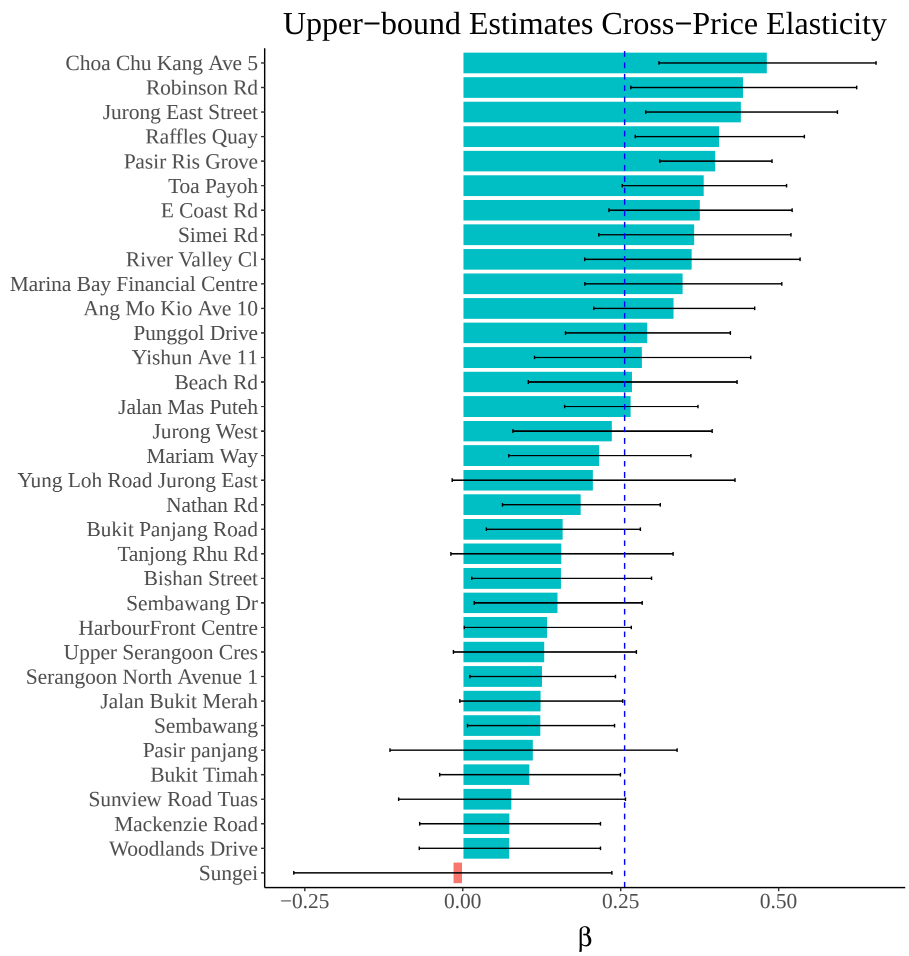

Figure A.8: Elasticity Estimates in Different Regions. This figure shows elasticity estimates across different starting zones in our sample.

that the autocorrelation in the surge factor suggests a persistent demand-supply mismatch in the Grab ride-hailing market, which means also that they can spillover to the taxi booking market.

Appendix D Computing Performance Improvement in the Policy Simulation

In our main text we have included and correlated results from two sets of experiments. The first set of experiments comes from a real-world field trial reported in Cheng et al. (2018). A total of 500 taxi drivers were recruited and instructed to install and use the Driver Guidance System (DGS). As the usage of DGS is voluntary, the authors tracked the trip-fetching performance in

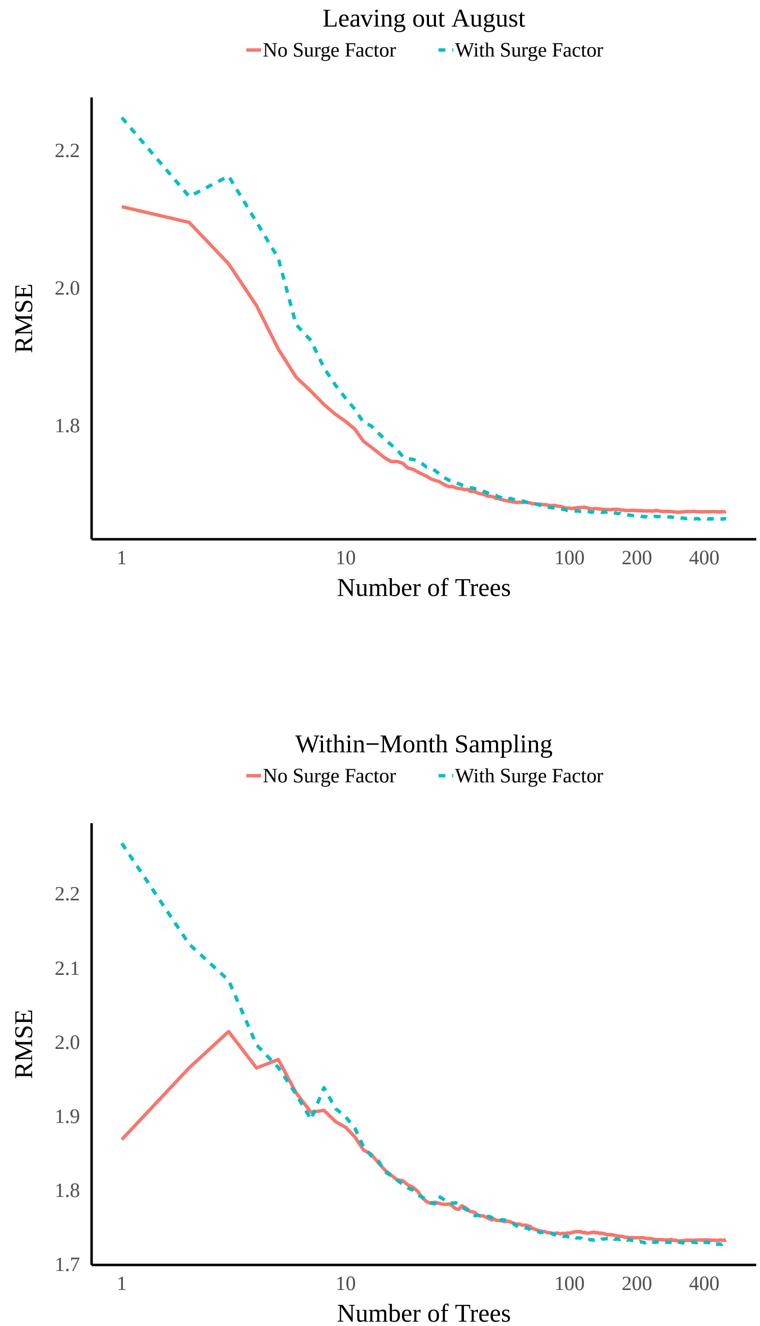

Figure A.9: Root Mean Squared Error Based on the Number of Trees. The figures show the in-sample root-mean squared errors based on the number of trees used for each of the two samples. The "Leaving out August" sample refers to the in-sample using data from April through June. The "Within-Month Sampling" sample refers to the in-sample using a randomly sampled 75% of the full data.

terms of vacant roaming time depending on the DGS usage. Their findings (with latest data) are summarized below:

Table A.6: Accuracy of Taxi Bookings Predictive Model. This table shows the accuracy of the random forest model in predicting the number of taxi bookings with different samples (both in-sample and out-of-sample). RMSE stands for root-mean-squared error.

	In-Sample RMSE		
Model Test Sample	No Surge Factor	With Surge Factor	Improvement (%)
August	1.417	1.451	-2.399
20% of Each Month	1.420	1.448	-4.789

	Out-of-Sample RMSE		
Model Test Sample	No Surge Factor	With Surge Factor	Improvement (%)
August	1.197	1.019	14.870
20% of Each Month	1.288	1.139	11.568

Table A.7: The actual performance of taxi drivers who used the DGS against those who didn't (data is collected from January 2018 to May 2018, when the DGS App is used most frequently).

Trip Type		Non-DGS	DGS
Street Hail Trips	Avg. Roaming Time (min)	10.10	7.09
	Avg. Service Time (min)	15.06	15.26
Booked Trips	Avg. Roaming Time (min)	7.98	5.89
	Avg. Response Time (min)	5.84	5.69
	Avg. Service Time (min)	19.19	18.98
Weighted Average	Avg. Roaming Time (min)	9.44	6.71
	Avg. Response Time (min)	1.82	1.82
	Avg. Service Time (min)	16.35	16.46

From Table A.7, we can see that by following the DGS guidance, drivers on average could reduce their vacant roaming time from 9.44 minutes to 6.71 minutes, a drop of 28.9%.

For each taxi trip, the total time a taxi driver spends on identifying and serving it is defined as:

$$\text{Avg. Roaming Time} + \text{Avg. Response Time} + \text{Avg. Service Time.}$$

The expected number of trips that could be served per unit time is thus:

$$\frac{1}{\text{Avg. Roaming Time} + \text{Avg. Response Time} + \text{Avg. Service Time}}.$$

For DGS and non-DGS trips, the expected numbers of trips are thus 0.341 and 0.3 respectively. Therefore, DGS guidance, when followed all the time, leads to a 13.5% increase in the expected number of trips.

To extrapolate the impact of a 9.4% drop in the vacant roaming time on the DGS performance, when compared against non-DGS cases, we first compute the expected reduction in vacant roaming time as:

$$1 - (1 - 0.289) \cdot (1 - 0.094) = 35.6\%.$$

We next look at the expected increase in the numbers of trips if the vacant roaming time is reduced by 35.6% based on the number in Table A.7. The expected number of trips is thus 0.352, when compared against 0.3 fetched by the non-DGS cases, which amounts to an increase of around 17.2%.