Platform Growth is Hard:
Experimental and Quasi-Experimental Evidence from a Carpooling Platform in Singapore

Abhijit Banerjee*
Rema Hanna†
Ruimin He‡
Zoë Hitzig§
Xin Wei Ngiam‡

ABSTRACT. We study platform growth in a field experiment with drivers on a peer-to-peer carpooling platform in Singapore. Motivated by both theory and observed practice, we ask whether temporary subsidies and informational campaigns can have durable effects on the growth of platforms. Drivers were randomly assigned to different promotional strategies, all quite standard: some drivers receive information about favorable market conditions, others receive subsidies for picking up passengers. Our results provide no support for the theories about how promotional strategies encourage platform growth. The information interventions do not work, and in some cases backfire. While subsidies may generate a small initial increase in participation, they reduce participation in the longer term. Estimating drivers' preferences, we find that drivers may adjust their selectivity when they know they are scarce relative to passengers, which might explain these results.

Keywords: two-sided platforms, network effects, platform growth

---

Date: October 2023.
Corresponding author: Hitzig (zhitzig@g.harvard.edu). We thank Aaron Berman, Michael Fryar and Zhaodong Chen for outstanding research assistance. We thank Ting Mao, Edwin Law, Serena Lee and the GrabHitch team for their help completing the project. For helpful comments, we thank Ben Olken, Gabriel Kreindler, Danielle Li and Ben Golub as well as audiences at the Cities and Development Conference at Harvard, and the Market Design and Development conference at Harvard Business School. This project was supported by the JPAL Urban Services Initiative.
* Department of Economics, MIT; NBER; JPAL.
† Kennedy School of Government, Harvard; NBER; JPAL.
‡ Grab.
§ Department of Economics, Harvard.

1. INTRODUCTION

Platforms coordinate a vast and growing portion of economic activity. Young platforms face obstacles to growth: creating a valuable service for one side of the market (e.g. buyers) requires getting the other side (e.g. sellers) on board first. Two-sided platforms thus face the well-known "chicken-and-egg problem" (Caillaud and Jullien 2003), or the problem of coordinating buyers and sellers.

The theory of two-sided platforms (Rochet and Tirole 2003; Parker and Van Alstyne 2005; Armstrong 2006) suggests that price subsidies and promotional campaigns can help emerging platforms overcome the chicken-and-egg problem and drive new market expansion. These strategies create a shock to one side of the market: informational campaigns do so by shifting participants' expectations and beliefs about the size of the market (Katz and Shapiro 1985), while temporary subsidies do so by lowering the cost of participation and thus the reserve price of the marginal participant. Then, the logic goes, a positive shock for one set of participants (e.g. sellers) can increase the number of participants on the other side of the market (e.g. buyers). In two-sided markets, there are always multiple dynamic equilibrium paths—a one-time shock can help move the market to a better growth path. In particular, a one-time shock can set off a positive feedback loop that leverages the complementarities between participation decisions of different participants. As a result, the theory of two-sided platforms suggests, temporary discounts and publicity campaigns play a coordinating role, moving the market to a thicker equilibrium. Many successful platforms—including Airbnb, Uber, Amazon and Paypal—employed such tactics early in their growth (Parker, Van Alstyne, and Choudary 2016).

The empirical support for the theory of two-sided platforms has so far been relatively limited, though there is some recent work based on the data unlocked for research by a few private firms.[1] Using a large-scale field experiment and extensive administrative data, this paper tests whether promotional campaigns can generate durable positive effects in peer-to-peer

[1] See, for example, Cullen and Farronato (2020), Farronato and Fradkin (2022) and Farronato, Fong, and Fradkin (2023).

markets by shifting the cost of entry through subsidies or by coordinating expectations through informational campaigns.

We conduct our experiment in an emerging market for carpooling in Singapore. While carpooling, drivers incur a small cost (e.g. time or hassle) in return for benefits that may be financial (e.g. drivers pass on a part of the petrol cost to the passengers, passengers pay less relative to a taxi), but also may be psychological or pro-social (companionship, altruism, environmentalism). Despite its potential for mutual benefit, carpooling cannot occur without efficient coordination, especially if there is a group of people not necessarily going to the same location at the same time each day (Ostrovsky and Schwarz 2019). In Singapore, GrabHitch was launched in November 2015 to fulfill a coordinating role for passengers and drivers looking to carpool.[2]

We worked with GrabHitch to conduct a large-scale experiment with 11,800 drivers who had expressed interest in using the platform (by signing up on the platform), but were not yet regular users of the system. We randomized drivers into either one of two control conditions—pure control or placebo message—or one of the following treatment groups that correspond to different promotional strategies: a pure publicity intervention that reminded the drivers that having a companion can make the drive more pleasant; a subsidy intervention in which drivers were offered a subsidy per ride which could be either high or low for rides given within a week; and a set of demand information interventions that drew the drivers' attention to the fact that demand had gone up in the recent period (the highlighted increment was either large or small and this was randomized) or that there was a substantial increase in the excess demand in the market. We look at the effect of these treatments on two primary outcomes that capture different aspects of driver participation: whether a driver attempts to provide the service ("makes a plan"), and whether a driver picks up a passenger ("completes a ride").

The results provide no support for the theories of why market-makers should adopt these promotional strategies. The large subsidy had a positive effect on both plans made and rides given during the week of the subsidy, but then both outcomes fall below control in the week and

[2] GrabHitch is a vertical on the mobile application of Grab, a ride-sharing service whose core service resembles Uber or Lyft in the United States.

month after the subsidies expire. The smaller subsidy created a similar pattern as the larger subsidy, although treatment effects were more muted in magnitude and significance. In other words, if anything, the subsidy reduced market participation in the post-subsidy period, which contrasts with the prediction that lowering the cost of participation for marginal users can generate durable growth.

The information treatments failed to promote participation and, in some cases, also backfired. The treatment where drivers were told about a small increase in demand had no effects on behavior; in fact, we can safely rule out any meaningful effects. The demand intervention where drivers were told about a larger increase in demand *reduced* the number of plans made in the week and month after the intervention, and it generated no effect during the weeklong subsidy period. Furthermore, this negative effect was not driven by drivers associating higher demand with a more competitive market: the intervention where the drivers were told about the high level of excess demand had a similar negative effect on plans made in the week and month after the intervention. In sum, drivers did not respond to information about the thickness of the market by increasing participation as the expectations-based theory of platform growth would suggest.

We investigate whether the overall treatment effects masked heterogeneity in responses. We find that female and male drivers respond differently to the treatments, with female drivers responding positively to both information and subsidies during the subsidy period, while men have no statistically significant response during the subsidy period. To understand what explains this difference, we investigate whether the differences in male and female drivers' prior experience with GrabHitch explains their different responses. By creating a sample of men matched to women on prior experience, we find that experience does not explain the differences in responses by gender. Furthermore, we examine whether drivers who had given a ride less recently (more dormant) responded differently from drivers who had given a ride more recently (less dormant) and don't find clear evidence that dormancy led to different responses.

Since our interventions generated a short run increase in the supply of rides, we can use this as a shock at the overall market level in an event-study framework to see if the supply shock leads to a future increase in demand at the market level, as a coordination based theory would suggest. We find no evidence of such an effect.

It is possible that the treatments—and in particular, the demand information treatments–failed because drivers were optimizing on dimensions other than the number of rides, e.g. detour, gender of passenger, neighborhood, etc. Using extensive pre-experiment, administrative data on driver choices among possible passengers, we estimate a model of driver preferences. Using this, we construct a metric of "quality-weighted rides."

Strikingly, we find very different results for the impact of our interventions using this alternative measure. Specifically all the "perverse" effects noted above disappear when we use this alternative metric, which suggests that drivers may react to good news about passenger demand by becoming pickier about the passengers they accept. This is consistent with the price-theoretic framework for matching markets suggested by Azevedo and Leshno (2016), in which the "selectivity'" of one side of the market adjusts endogenously to clear markets in the absence of prices. Note that this explanation is consistent with being pessimistic about promotional interventions.

Moreover, multi-dimensional preferences can explain why men and women react differently to our interventions based on gender-based differences in preferences. For example, if many male drivers joined the platform to meet members of the opposite gender while female drivers are mostly there to make money, it is entirely possible that they may react differently to good news about demand—men may start focusing more on the gender of the passenger while women might decide to offer more rides.

Lastly, the fact that the drivers first react positively to the subsidies but eventually end up participating less might reflect the fact that the drivers were learning about their own preferences (say for taking detours). The subsidy might have made them *think* about making plans and giving rides. In the process, they might have discovered that they were less keen to give more rides at all costs than they had previously imagined.

We also explore other potential explanations for why these interventions failed to generate lasting effects. First, drivers and/or passengers may have pessimistic beliefs. If they are convinced that the platform will eventually fail and shut down, then they have little reason to react to the information or the incentives and increase their durable engagement with the platform. Drivers may still try harder in the subsidy period but with the idea of working less in the immediate future to compensate. However, we do not find evidence for this alternative explanation. Both drivers and passengers are persistent: failures to match do not seem to discourage them from trying again.

Finally we discuss a potential explanation for why subsidies failed to promote participation suggested by previous theories (Benabou and Tirole 2003) and empirical observations (Gneezy and Rustichini 2000): financial incentives may "crowd out" prosocial behavior.

This study contributes to the theoretical and empirical literature on platform economics and peer-to-peer marketplaces. Foundational theoretical work on platform economics focused on how to set prices as a function of user participation decisions (Rochet and Tirole 2003; Caillaud and Jullien 2003; Parker and Van Alstyne 2005; Weyl 2010), and later considered how information disclosure and search influences design considerations (Boudreau & Lakhani, 2009). Empirical studies of peer-to-peer online markets document how geographic heterogeneity generates frictions in growth (Cullen and Farronato 2020), how ranking algorithms influence user choice (Hitsch, Hortascu, and Ariely 2010), how search inefficiencies (Fradkin et al. 2015) and congestion (Arnold, Darmon, & He, 2014; Horton & Zeckhauser, 2010) affect coordination, and how different pricing mechanisms influence equilibrium (Einav & Levin, 2018). For a survey of the literature on online peer-to-peer markets, see Einav, Farronato, and Levin (2016). A closely related experiment specifically tests the expectations-based theory of platform growth in a platform for commercial product-building collaborations (Boudreau 2021).

To this literature, we contribute empirical insights into how obstacles to platform growth may be resistant to traditional promotional strategies.[3] We find that promotional strategies

[3] The study most closely related to our own is a large-scale field experiment in Cohen et al. (2021) conducted with users of Waze. Their experiment similarly sent messages to commuters intended to remind them of the benefits of carpooling. In line with our results, they also find that messages about carpooling can increase ``intent" to carpool, but that this intent does not always translate into an increase in the rides taken. Unlike our study, which focuses on how to sustain drivers' interest in carpooling as the service grows and develops, their study focused on the initial adoption of carpooling. In order to promote initial adoption, their experimental treatments highlight the time-saving aspect of carpooling, specifically targeting

designed for two-sided platforms with conventional prices and surplus-driven services may fail in peer-to-peer contexts in which non-price characteristics of the service are more salient for users. The fact that they fail is somewhat surprising, in light of how common these strategies are with nascent platforms. This finding has direct implications for the design and growth strategies of intermediaries, and it also suggests new avenues for theoretical and empirical work investigating the interaction between platforms and prosocial behavior.

Finally, it is worth noting that our large-scale field experiment offers valuable insights for managers aiming to understand strategies to spur platform growth. In particular, while in theory subsidies and information campaigns can help move two-sided markets to a new equilibrium, our experimental findings highlight how such strategies may be ineffective or even backfire when user preferences are multi-dimensional and therefore unpredictable. Our study therefore suggests that before investing resources on blanket or untargeted subsidies and information campaigns, managers should gather baseline information about the relevant dimensions of users' preferences.

2. Setting and Experimental Design

2.1. Setting.

We partner with Grab, a privately-held South East Asian mobile ride-hailing company with operations in Singapore, Malaysia, Indonesia, Vietnam, Myanmar, Cambodia, Thailand and the Philippines. Grab's smartphone-based app offers many products and services including standard on-demand personal cars for hire (GrabCar and GrabTaxi), on-demand motorbikes (GrabBike), on-demand food delivery (GrabFood) as well as in-app platforms for messaging (GrabChat) and payment (GrabPay).

We focus on one vertical called "GrabHitch." GrabHitch has been in operation since November 2015 in Singapore, and at the time of our experiment was the only non-profit peer-to-

---

commuters who could save time on their commute by using high occupancy vehicle (HOV) lanes. Meanwhile, our study targeted dormant drivers who had already given one ride in the past, and whose beliefs about the value of the platform may be influenced by their experiences when the platform was less developed.

peer carpooling platform of its kind in Singapore.[4] The "Hitch" vertical is a social carpooling service that aims to match passengers and casual drivers who are traveling the same approximate route. Importantly, drivers are not professional drivers—they sign up to take passengers at their convenience and may not complete more than two rides per day. Most are not professional drivers: few (about 8%) sign up to drive for Grab's main ride-hailing system (GrabCar).

The platform automatically charges a small fare, paid to the drivers, intended to help offset petrol costs. For passengers, empirically observed prices for a given route on GrabHitch are 20-40 percent lower than the commercial rideshares or municipal taxis. Pricing is based solely on distance between origin and destination, and there is no surge pricing based on demand. Advertising for GrabHitch stresses the social, environmental and economic benefits of peer-to-peer car-sharing. For example, GrabHitch marketing emphasizes the potential to "expand your social network" while saving money and making Singapore a "car-light and friendly city."[5]

Signing up as a driver for GrabHitch is simple relative to the much lengthier processes required for becoming a professional driver on other verticals. Drivers must be over 18, have a valid license with at least 1 year of driving experience, and have a private car with valid auto insurance. Upon submitting documentation, Grab verifies the driver information and runs a background check, usually approving drivers within two days.

Once approved, drivers can propose "plans" on the app, entering their origin, destination, and desired departure time.[6] Meanwhile, passengers submit their "bookings" on the app, keying in their origins, destinations, and desired start times for their trips. Passengers can also indicate whether there are pets or other passengers traveling with them, and they can optionally choose

[4] Though Hitch was a non-profit vertical at the time of our experiment, one year later, in July 2018, Grab introduced a 10% platform fee for all Hitch rides.

[5] Given that it is generally seen as a luxury to own a car in Singapore due to the high costs of car ownership and strict government policies to keep cars off the road, driving for GrabHitch is not likely to be a significant source of income for drivers. Any prospective car owner must first bid for a Certificate of Entitlement (of which there is a restricted, fixed number). Between taxes, the COE, and the car itself, a regular compact car in 2017 was reported to cost 105,000 SGD (77,000 USD). Meanwhile, the median fare on GrabHitch is 10 SGD (7 USD). In the extreme case of a driver who picked up 2 trips per day every day, the driver would make 560 SGD (415 USD) per month, while median monthly income in Singapore in 2017 was 9023 SGD (6650 USD) (Channel News Asia, 2023).

[6] Drivers also have the option to indicate whether they would like to make a recurring or non-recurring plan.

to only be matched to drivers of the same gender. While the interface allows for passengers to make bookings as few as 15 minutes in advance and as many as 7 days in advance, GrabHitch recommends, for example, booking "the night before for a morning commute or 2 hours ahead of your evening ride home."

After the driver enters her plan, she sees a list of passenger candidates. The list includes the passenger's name, booking details, and photo (if provided). The list is dynamic: As more passengers make bookings that are compatible with the driver's plan, new passenger bookings will appear. This dynamic list is constructed using a proprietary algorithm based on passenger and driver trip details. The driver sees the eligible bookings—as many as fit on her smartphone screen (usually a maximum of 5-15 bookings fit on the screen)—ordered by the compatibility score. The driver can also scroll down to see the full set of compatible bookings.

The driver can then select a passenger from the list of candidates at any point up to 15 minutes before the trip. When selected, the passenger receives a notification along with the driver's name, mobile phone number, license plate, car type, car color, and photo (if available). The matched pair can communicate through an in-app messaging service or via SMS. If all goes smoothly, the driver picks up the passenger at her requested origin at the requested time. After completing a trip, passengers typically pay their drivers through the built-in payment platform or in cash. Drivers can complete a maximum of two GrabHitch rides per day due to legal restrictions on non-commercial driving in Singapore.[7]

2.2. Experimental Design.

The experiment aimed to understand how information and subsidies affect the probability that a driver offers a service on the platform (i.e. makes a plan), as well as the ultimate probability that a driver delivers a service on the platform (i.e. gives a ride). We focus on dormant drivers, defined as those who had not completed a ride in the 60 days prior to the experiment. We chose dormant drivers since market growth depends not just on attracting participation but also maintaining

[7] It is possible for either the driver or the passenger to cancel the ride after the parties have been matched. Cancelling the ride—even after the match is confirmed by both parties—incurs no cost to either party. However, if a user cancels frequently, they can be barred from using the service in the future.

participation; understanding the behavior of those who join (thus demonstrating a need for the service) but subsequently drop out can give us insights into the role of information and subsidies in overcoming coordination failures. We included the entire universe of dormant drivers for the study, generating a sample size of 11,883 drivers to be randomly assigned to treatment groups.

The first set of treatments mimic typical promotions that offer subsidies to try the product or advertise particular features of it. Specifically, we randomly offered small and large bonuses per ride for a promotion period. The amounts and timeframes were determined by GrabHitch to be consistent with Grab's practices. We offered $S4 (small bonus) or $S8 (large bonus) for any ride given on the day the SMS was sent or in the 5 days thereafter. The average fare for a ride on GrabHitch is $S10, so these bonus payments yield large payments to the driver (1.4x and 1.6x the average fare). Moreover, the promotion duration is similar to the duration of other promotions on GrabHitch and of other experiments with rideshare drivers in the literature; for example, Chen et al. (2022) exogenously varies guaranteed wages for Uber drivers over subsidy periods of three and five days.[8]

Second, we randomized some users into receiving information about the "social" benefit of carpooling, i.e. making new friends. This "Companionship" treatment group received messages that read: "We miss you! Offer a ride to a fellow commuter today and meet new friends!"

The next set of treatments are the "density treatments." We provide drivers in this treatment arm with information about how the passenger side of the market has changed since they last took a ride. Specifically, drivers received messages that read "We miss you! Offer a ride to a fellow commuter today! The number of passengers on Hitch has increased X% since you last completed a ride!" where X was randomly assigned to be either 48 percent ("High density treatment") or 24 percent ("low density treatment). These figures were calculated from the administrative data, with a base period of the last two months (i.e. 24 percent) or the last nineteen months (i.e. 48 percent).[9]

---

[8] Chen et al. (2022) offered "subsidies" to drivers in the form of 1.1-1.4x fare, and found that drivers who received the promotion increased their hours worked by one percent during the promotion period.

[9] We choose these two base periods since no drivers had completed rides within the last two months, and the furthest rides were saw in the data were about 19 months.

The density treatments not only signal an increase in passengers, but also an increase in drivers due to network effects. Thus, it is possible that drivers might infer from the density messages that the market is now too competitive, and it is not worth it for them to join. In order to isolate the effect of information about *excess demand* from demand, we also included a SMS message ("excess demand" treatment) that provided information about the growth of unmatched passengers in the last 60 days prior to the experiment: "Since your last ride, the monthly number of unmatched passengers has grown by 814,643!"

To evaluate the effect of these various treatments, we randomly assigned drivers to one of two control groups. First, we sent a placebo message that included just the first sentence in the treatment arms: "We miss you! Offer a ride to a fellow commuter today!" This treatment allows us to separate out the effect of the reminder of the existence of GrabHitch from the content of each message, and it is the primary control group that we compare each treatment to. Second, we also randomize some drivers into a "pure" control group that received no text messages, so we can measure the overall treatment effect of the messages.[10]

2.3. Randomization and Experimental Implementation.

We classified drivers into two strata: (1) "more dormant stratum" that consisted of drivers who had *neither made a plan nor completed a ride* in the 60 days before the experiment, and (2) "less dormant stratum" that consisted of drivers who *had made at least one plan* but *had not completed a ride* in the 60 days prior to the experiment. Within these two strata, we randomized drivers into roughly equal sized bins across the experimental treatments (see Table A1 for sample selection details, and Table A2 for a balance check across gender and pre-treatment driver history).[11]

The experiment launched on July 11, 2017. GrabHitch sent the SMSs to the drivers.

2.4. Data Collection

---

[10] The exact text of messages sent to drivers can be found in Appendix Table A4. From this point onward, we refer to Figure/Table x in the Appendix as Figure/Table Ax.

[11] The drivers in our sample are similar to drivers who joined recently. Table A11 compares drivers in our sample to drivers who joined in the month before our intervention, and find that they are similar in terms of prior experience and gender distribution.

We worked with Grab's data under strict confidentiality conditions, with our access limited to the purposes of this study. We had access to a dataset of all plans made by drivers, bookings made by passengers, as well as ultimate matches and rides completed from the inception of GrabHitch up to a 5 weeks after our experiment ended. For each of these, the data fields included request times, pick-up and drop-off locations, detours to pick up passengers, prices, and the gender of the drivers and passengers. No personally-identifiable data was used. In addition, we also had access to data on the backend universe of candidate bookings that were compatible with each driver plan, i.e. the choice set of passengers shown to drivers for a given ride.

3. Basic Market Description

Before we turn to our experimental results, we describe the evolution of the market from GrabHitch's launch up until our experiment.

The total number of completed GrabHitch rides in Singapore from its inception to just before our intervention is 4,493,303 (Table 1, Panel A). The number of completed rides steadily increased over time, peaking at 217,717 rides in the week of June 12, 2017 (Figure A1). Note that the number of completed rides is subject to market-wide shocks, such as promotions.

We observe some differences in the level of activities across neighborhoods, but not large differences in patterns over time. We used GrabHitch's geographic boundaries, which classifies the city into 54 areas. The area with the most activity in terms of passenger requests is Tampines, a residential suburb. Other active areas include the Woodlands, Houjang, Bedok, Downtown Core, Bukit Merah, and Jurong West. One might think that the overall trend in market growth and development may be masking neighborhood differences in growth trajectories---small initial differences in market conditions could lead to large differences due to economies of scale. However, graphing completed rides by neighborhood for the ten most active neighborhoods in Figure A2, we do not find any observable evidence of this.[12]

[12] In order to find the "top 10 neighborhoods," we ordered the 54 neighborhoods by the total number of completed rides that had their pick-up point in that neighborhood, over our entire observation period.

Most bookings tend to occur during commuting hours. Figure A3 shows the frequency of passenger bookings and driver plans, by time of day. There are sharp peaks in passenger bookings around 8:30am and 6pm. The commuting peaks are less pronounced for drivers.

We next examine key facts about the market participants, defined as anyone who has entered a booking as a passenger or plan as a driver, regardless of whether they have ever actually completed a ride.[13] The total number of participants in the market includes 111,248 drivers and 730,324 passengers. Only about 18 percent of the drivers are female, despite a more even split in the gender of passengers.

Drivers appear more active than passengers. On average, drivers made 122 plans, and completed 57 rides. In contrast, passengers made an average of 18.2 bookings and took about six rides. Note that the platform had somewhat low success rates: of all of the drivers who have attempted to give a ride, 29.6 percent of them have never actually completed a ride. Of all of the passengers who had ever attempted to take a ride, 41 percent have not succeeded in taking a ride.

GrabHitch believed that they could do better in matching drivers and passengers.[14] To see why there is potential scope to do so, Figure 1 shows the number of active drivers per week and the number of active passengers per week (two week moving average) from the launch of GrabHitch up to our intervention. The number of active drivers and active passengers both steadily increase over time. But, note that drivers cannot complete more than two trips a day (on average drivers make 5.87 plans per week and complete 2.49 rides per week). So, when the

[13] One can be both a driver and a passenger. We treat one individual who participates in both sides of the market as distinct when they are in their passenger vs. driver role.

[14] The quality of GrabHitch for a given passenger (driver) depends on the number of drivers (passengers) already using the service. This can lead to inefficient equilibria in which both sides of the market stay out, as in static models where users make a one-time entry decision like that of Caillaud and Jullien (2003) and Rochet and Tirole (2002). However, the static models treat platform growth as an event rather than a process – more generally, the role of the platform as it grows is to carry out a "balancing act" (Rochet and Tirole, 2004), ensuring that at each moment in time there are enough sellers to meet demand and enough buyers to meet supply. Coordination "failure" can occur at different stages of growth – this may be because of equilibrium multiplicity or because the platform fails to reach a "critical mass" of users, as posited by the explicitly dynamic models of Ellison and Fudenberg (2003) and Evans and Schmalensee (2009). Because of equilibrium multiplicity, firms like Grab act "as if" there is a coordination failure, in hopes of moving the market to a better equilibrium.

number of passengers increases but the number of drivers stays the same, the quality of the platform as a whole decreases because passengers are less likely to get a ride. Although the platform as a whole was growing steadily, it was not clear whether the platform was maintaining or increasing the quality of the service over time.

Furthermore, in Figure 2, we show how one measure of service quality – "match rates" (the ratio of drivers/passengers who succeed at completing a ride relative to the number of drivers/passengers who attempt to get a ride) – relate to the number of users in the market. Each observation in Figure 2 is a week in the administrative data from the launch of GrabHitch up to our intervention – the number of active users (drivers + passengers) in a given week is on the x-axis and the match rate for passengers (Panel A) and drivers. We see that as the number of active users per week increases, the match rates for passengers and drivers both increase. This is consistent with there being a coordination element in market, but also with other sources of correlated heterogeneity. Moreover, the match rate for passengers levels off at some point – in the weeks when there are ~150k to ~200k users, the passenger match rate is lower than it is when there are ~95k to 125k users. Although it is not clear what the ideal equilibrium is in this market, the plateau around a low match rate (~.3) is certainly consistent with the possibility of improving coordination.

There is considerable heterogeneity in market conditions across different neighborhoods at different times (Figure A5) and it is very possible that the market is localized (i.e. drivers like to pick up passengers who live near them). Thus, information about the market may be less useful if participants learn from their local neighborhood conditions and this differs from the market as a whole.

Two key facts emerge. Using the passenger to driver ratio as a measure of market density, we can examine what actual matches look like locally with high density. Overall passengers' match efficiency decreases with passenger density (Figure A6, Panel A) while drivers' match efficiency increases with passenger density (Figure A6, Panel B). However, a closer look at match rates plotted against deciles of passenger density shows that while driver match rates increase over the lowest five deciles (starting at 0.13), they then level off around the sixth decile (at a mean match rate of 0.33) before decreasing slightly to 0.29 in the 10th decile (Figure A7).

Finally, we examine detours, trip lengths, and fares, as they provide a sense of how much drivers are willing to change their plans in order to pick up a passenger. Overall, drivers plan longer trips than passengers book: The median driver plan is 21.9 kilometers (Table 1, Panel B) while the median passenger booking is 14.3 kilometers (Table 1, Panel C). For reference, a trip from the center of the largely residential neighborhood of Tampines to the Downtown Central Business District is about 20 kilometers. For completed rides, the median absolute detour that drivers take to pick up their passenger—defined as the distance from the passenger's pick-up point to the driver's origin—is 2.6 kilometers, or about 18 percent of the median distance of completed rides. The average fare paid by the passenger is 10.10 SGD or about $7.50 in USD.

4. DO TEMPORARY BONUSES OR INFORMATIONAL MESSAGES CHANGE BEHAVIOR?

4.1. Experimental Results.

We now turn to the results of our field experiment with dormant drivers on GrabHitch.

We estimate the impact of the different information treatments (T_d) on two key outcomes (y_d): whether the driver d makes a plan, i.e. whether the treatment moves a driver to enter the market, and whether the driver ultimately provides gives a ride to a passenger. Specifically, we estimate using OLS:

$$y_d = \beta_0 + \sum_T \beta_T T_d + \beta_\delta \delta_d + \varepsilon_d \, ,$$

where δ_d is an indicator variable for whether the driver is more or less dormant (i.e. the strata). Standard errors are Huber-White robust standard errors.[15] We estimate the treatment effects for three different time periods. First, we examine what happens when the information has just been received and the subsidies are still active (July 11-16). We then examine what happens after the subsidy period. Specifically, we examine the week after the subsidy period ends (July 17-22) and the month after (July 17-August 15). To clarify the timeline, see Figure 3.

[15] As we randomized treatment, our coefficient estimates capture the causal effect of the treatments. However, as an additional robustness check, we include the total number of plans made prior to the experiment as an additional control variable to correct for any small sample imbalances. Table A6 shows that this specification provides near identical results.

For this analysis, we drop the pure control group from the estimation sample and we compare the information treatments (T_d) to the placebo message treatment (so the sample has 10,401 observations). Therefore, each β_T captures the independent impact of receiving the content of each treatment message, as distinct from the impact of receiving a message. Note, however, in Table 5 we estimate the impact of receiving any message treatment relative the pure control and we find that receiving a message increases the probability of making a plan during the subsidy period by 26 percent (p-value <0.05), but has no observable lasting effect in the subsequent periods or on the probability of taking a ride. This confirms that the messages were being read by drivers, enough so to impact their behavior.

Table 2 provides our findings of the impact of the treatments on drivers' choices to enter a plan and to give a ride.[16] We have two key sets of findings—on the effects of the subsidies and the effects of various informational treatments.

First, we examine the subsidy results. Recall that the theory of two-sided platforms suggests that temporary bonuses to one side of the market promote durable growth. If anything, we find the opposite. While drivers take advantage of the bonuses while the bonuses are available, they are then less likely to use the service afterwards, when the bonus is no longer available. The larger bonus (S$8) increased activity during the subsidy period: it increased the probability that the driver entered a plan during the subsidy period by 24 percent relative to the control group, as well as the probability that they gave a ride by almost 45 percent (significant at the 10 percent level). However, in the post-subsidy week and month, overall, drivers in the larger bonus group reduced the probability that they entered a plan relative to the control group by about 24 percent in the week after (significant at the 10 percent level) and by 13 percent in the month after (significant at the 10 percent level). In addition, drivers who received the large subsidy were 34 percent less likely to complete a ride in the week after the subsidy period.

[16] While Table 2 considers whether the treatment affected intensive margin decisions (whether the driver undertook a plan or ride), Table A5 considers the effect of treatment on both intensive and extensive margin outcomes. Specifically, we examine the number of plans made and the number of rides taken. We find similar effects that the excess demand information reduces the number of plans made, as well as rides taken. However, while we find that the large bonus affects the decision to enter the market, we find no observable impact on the number of rides taken.

The small bonus effects were more muted, but also showed declines after the subsidy period: The small bonuses (S$4 per ride) had no noticeable effect on making a plan nor giving a ride during the subsidy period, and led to a 14 percent decrease in the probability that a driver would make a plan over the course of the month after the subsidy ended. In short, rather than inducing long-term participation, the temporary subsidies appeared to simply shift activity to the subsidy period.

Second, we consider the effects of our informational treatments. Overall drivers do not respond to information about the thickness of the market by increasing participation as the expectations-based theory of platform growth suggests. The low density information treatment did not have any statistically significant impact on plans made or rides given in any subsidy period. The high density treatment did not have a statistically significant impact on plans or rides during the subsidy period, but it did have a statistically significant *negative* impact on plans made in the week and month after the subsidy period ended. In the week and month after the subsidy period ended, 21 percent less likely to make a plan (significant at the 10 percent level) and 18 percent less likely to make a plan (significant at the 1 percent level) , respectively.

It is possible that the density treatments also conveyed information that supply had increased alongside demand—in which case drivers may believe the market is too congested.[17] Thus, we turn to the excess demand treatment, which isolated information on the number of unmatched passengers. Again, we find that, if anything, the excess demand information led to *decreases* in the probability of entering a plan relative to the control group. In the first week after the information was sent, those who received the excess demand information were 30 percent *less* likely to take a ride (significant at the 5 percent level). Note that this effect is materially large, as the percent change is equivalent to the size of the large bonus treatment.

One possible alternative explanation of why information about market thickness did not induce participation is that drivers may respond to the information by being pickier about the

[17] As already mentioned the number of passengers and drivers do co-move positively up to a point, so the worry that a bigger market also has more drivers is not unwarranted.

quality of their rides (e.g. how much of a detour it is, the gender of their passenger, etc.) rather than just increasing rides.[18] Thus, we next examine two measures of quality in Table 4.

First, we examine a very simple measure of quality: the length of the detour that driver makes if they take a ride (Table 3, Columns 1 - 3). We use a *relative* measure of the detour—the detour is the total distance that the passenger's booking adds to a driver's initial plan, divided by the length of the driver's initial plan (in kilometers). Note that we only observe the detour if someone takes a ride, and we know the treatments induce different probabilities of taking a ride across treatments, so these results are suggestive rather than causal.[19] Those who receive the low- or high- density treatments do not take different detours relative to those in the control. However, those in the excess demand group appear picker than the control: They give rides that require detours that are 0.128 kilometers shorter than the detours taken by the control group (0.413), representing a 31.6 percent decrease.

Second, we make use of the detailed administrative data on the driver's choice set of passengers for a given plan to construct a new outcome called "quality-weighted rides." This measure weights a ride given by a driver by a proxy for the idiosyncratic "quality" of that passenger's booking. Our quality weights take into account six observable measures: (i) relative detour (the distance that the passenger booking adds to the driver's entered plan, divided by the driver's entered plan, both in kilometers); (ii) passenger gender; (iii) an interaction term for the driver gender and the passenger gender; (iv) the number of seats requested by the passenger; (v) the difference between the start time of the passenger's booking and the driver's plan (in minutes); (vi) whether the passenger has a photo.[20] Details about the construction of quality weights are provided in Appendix A.

---

[18] This idea relates to predictions that arise from the price-theoretic framework for matching markets presented in Azevedo and Leshno (2016)—when drivers know that they are scarce relative to passengers, they adjust their "selectivity" much like suppliers in standard frameworks would increase their prices.

[19] In Table A6, we also examine detours across the full sample. In particular, we code the detour taken as "0" if the driver did not complete a ride, and so the measure captures both the extensive margin of detour and the intensive margin of whether the driver would take a detour. The results are similar in that those in the drivers who received the excess demand treatment are less likely to take a detour, and so are perhaps choosing high quality matches.

[20] Anecdotal evidence is consistent with the importance of these factors, see for example: https://cnalifestyle.channelnewsasia.com/trending/how-to-get-grab-hitch-booking-driver-accept-pick-up-

We find no observable impact of the information about density or excess demand on quality weighted rides (Table 3, Columns 4 - 6). Particularly striking is that while the excess demand treatment led to a 48 percent decrease in rides taken in the week the information was received, we find no observable impact on quality adjusted rides. This fact suggests that when drivers did give a ride, they tended to give rides of high quality. While not conclusive, the evidence about driver detours and quality adjusted rides suggests that one reason why drivers may have reduced their entry behavior is that they became pickier about the quality of their ride, knowing that there were plenty of passengers to match with. This idea is consistent with the price-theoretic framework for matching markets presented in Azevedo and Leshno (2016).

4.2. Heterogeneity in Treatment Effects

The weak overall effects of the intervention could be masking heterogeneity in responses. We examine two dimensions of heterogeneity here: gender and dormancy. We find that female and male drivers respond differently to the treatments, and that these differences are not explained by differences in experience. Meanwhile, we do not find meaningful differences between drivers who have been more versus less dormant.

Table 3 reports treatment effects separated by gender of the driver. Most strikingly, women respond positively to both the larger information treatment and the larger bonus during the subsidy period (Table 3, Panel B, Column 1), while no treatment had a statistically significant effect on plans made by men during the subsidy period (Table 3, Panel A, Column 1). In particular, the high density messages led to an 84 percent increase in the probability that a female driver made a plan (p value < 0.05). The larger bonus treatment led to a 66 percent increase in the probability that a female driver made a plan (p value < 0.05). These effects did not translate to statistically significant increases in the number of rides given during the subsidy period (Table 3, Panel B, Column 4).

To further examine the gender heterogeneity in treatment effects reported in Table 3, we examine how men and women differ on other dimensions. One salient difference between men

secrets-10256366. Table A7 confirms that detour, passenger gender, booking-plan time difference and number of seats are all important predictors of a driver choosing a passenger.

and women in our sample is that men are more experienced in terms of the number of plans they have made prior to the experiment, and the number of trips they have made prior to the experiment. On average, men made 31.3 plans and gave 19.4 rides prior to the intervention while women made 21.5 plans and gave 11.9 rides prior to the intervention (Table A9). We include the number past plans a driver made as a control and run the split gender treatment effects in Table A7, and find that controlling for past plans is not qualitatively different from the results on gender heterogeneity without this control, reported in Table 3. Furthermore, Table A10 constructs a sample of men matched to women on their past experience (number of prior plans and rides) via coarsened exact matching, and reports treatment effects for this sample of men matched to women. In the sample of men matched to women of experience, there are still no positive responses during the subsidy period, as there are for women in Table 3. These additional analyses suggest that there is some other difference between men and women that drives their differences in responses.

Recall that our sample was stratified into more dormant and less dormant drivers. We consider whether the more dormant drivers respond differently to the intervention than do the less dormant drivers. Table A8 shows separate treatment effects for the more dormant and less dormant drivers. We cannot conclude that the effects were qualitatively different between more dormant and less dormant drivers. The less dormant drivers have similar signed coefficients, but they are less precisely estimated in part due to a smaller sample size.

4.3. Did Our Intervention Improve Coordination Overall?

To explore the question of complementarities further, we ask whether the package of interventions had any effect on the number of rides given in the week following the messages. As we report in Table 5, the experimental treatments as a whole increased drivers' probabilities of making a plan by nearly 26 percent.[21] We thus can see if this shock to the supply side led to subsequent changes in the market.

[21] Although we cannot tell whether drivers open the SMS messages sent to them, the positive effect of receiving any message indicates that many drivers are reading the messages.

We implement an event study design. Figure 4 shows driver and passenger activity and match rates in the days prior to and after the intervention in the market as a whole. That is, each observation in Figure 4 is the match rate or activity level in all of Singapore on the day indicated on the x-axis.[22]

To avoid the seasonal noise driving our results, we chose to focus our event study on a small window around the intervention: 10 days before and 6 days after. This window is shown in Figure 4 in light grey dotted lines. The period of 6 days after the intervention coincides with the period during which the subsidy was active. Overall, we find no change after the experiment.[23]

One potential explanation for why our intervention—despite increasing the probability that a dormant driver makes a plan by nearly 26 percent—did not lead to overall changes in the market is that some drivers and passengers, once connected will continue to coordinate their rides off the platform. If this is the case, then the Hitch platform's coordinating role lapses after the first ride, and the driver and passenger can subsequently carpool without coordinating through Hitch. Though we do not have data that can speak directly to whether successful matches "exit" the platform, institutional details and prior data on user behavior suggest that this phenomenon is unlikely to be driving our results. Recall that Grab, at the time of our experiment, was running Hitch as a non-profit vertical---since the company wasn't taking a fee for coordinating the ride, there is no monetary gain to the driver or the passenger from moving offline. Furthermore, over 95 percent of all historical rides on Hitch are paid for through the in-app payment service GrabPay rather than paying by cash. So, this fact suggests that coordinating a carpooling ride through Grab rather than off-platform offers an additional convenience of automated hassle-free cashless payments. In addition, the platform offers a legal structure for frictionless contracting between the driver and the passenger.

---

[22] The driver match rate in neighborhood i in week t is simply the total number of driver plans originating in neighborhood i in week t divided by the total number of completed plans in the same neighborhood and week. The passenger match rate is defined similarly.

[23] In addition to the graphical evidence, Table A14 shows the coefficients on our event study specification. Effects on activity outcomes reported in Panel A of Table A14 (plans, bookings, completed rides, number of driver and passengers) are positive but very small in magnitude and not statistically significant. The same is true for match rates, with the exception of passenger match rates, which decreased (Panel B, Table A14).

5. MECHANISMS

We next explore reasons why our interventions, for the most part, failed to move the market to a new equilibrium in this setting.

5.1. User Persistence.

First, we examine whether our treatments failed because individuals have pessimistic beliefs about the long run prospects of the market. To look at this we examine how drivers and passengers react when they fail to be matched despite trying several times using extensive administrative data from the entire GrabHitch platform since its launch up until the launch of our experiment.

In particular, we ask: what is the probability that a driver enters a plan into the app again––i.e. tries again—conditional on having an unsuccessful attempt—i.e. enters a plan into the app that does not lead to a ride? We provide a measure of persistence in Figure 5, where we graph the probability that a driver makes an additional attempt after failing X times. Drivers appear to be very persistent: 93.7 percent of drivers try again after their first attempt is unsuccessful, while 87.7 percent of people try again after two unsuccessful attempts. In fact, after 10 unsuccessful tries (and no successful ones), 54.9 percent of drivers try again, and even after 20 unsuccessful tries, 35.0 percent of drivers try again.

We construct a similar plot for passengers to help put the driver findings into context. We find that passengers are less persistent than drivers: 76.6 percent try again after the first failure, but only 31.4 percent try again after 5 failures in a row. By 10 failures in a row, only 8.6 percent try again. There are a number of reasons why passengers may be less persistent. For example, passengers have other rideshare verticals available to them while drivers do not. In addition, drivers and passengers receive different information about why they failed to get a match.

The evidence of persistence suggests that the drivers and to a lesser extent the passengers are actually quite optimistic about the possibilities of this platform. Indeed this persistence may be one reason the interventions do not have a effect at the margin. Though that cannot explain the perverse reactions we find.

5.2. Did Incentives Crowd Out Prosocial Behavior?

Financial incentives have psychological effects on their recipients which may imply that monetary incentives can actually deter individuals from engaging in the incentivized activity (Benabou and Tirole, 2006). Drivers' motivations to provide rides may come from a prosocial or altruistic impulse. Thus, it may be the case that our intervention crowded out prosocial incentives. Though we do not test this directly, a number of empirical studies in the literature (reviewed in Gneezy et al., 2011) find crowd-out effects in settings similar to ours.[24] However even if drivers are prosocial it is not clear why the density information or information on excess demand should have negative effect on the number of rides—the fact that passengers want rides should typically encourage pro-social drivers.

6. CONCLUSION

When developing new markets, platforms often use subsidies and information campaigns to promote participation and growth. Platforms pursue these tactics with the theory of two-sided platforms in mind: inducing more participation of one side of the market can spur growth of the market as a whole through complementarities in user decisions. Through an experiment with drivers on an emerging carpooling platform in Singapore, we find no evidence to support the theory behind such strategies. Subsidies—which theoretically lower the cost of participation for the marginal user—induce short term participation but decrease participation in the longer term. Information about demand—which, in theory, can shift expectations and beliefs—did not motivate participation and in some cases backfired. Two potential reasons why the two-sided

[24] For example, Meier (2007) finds dynamic effects consistent with ours: Price incentives for charitable giving increased donations while the incentives were in place, but then decreased donations relative to the baseline when the incentives expired. This effect is similar to one we observe, in which drivers who received financial incentives participated more while the subsidy was active but then decreased their participation relative to the baseline when the subsidy lapsed.

platform logic may break down in this setting have to do with expectations: users are persistent so they may not react much to market signals, and users have heterogeneous beliefs. Further, we find suggestive evidence that drivers get pickier about which passengers they choose when they know they are scarce relative to passengers, which blocks the putative feedback loop driven by complementarities in supply and demand decisions. Prior studies suggest that monetary incentives may crowd out prosocial behavior.

Our findings help to inform managerial decisions regarding platform growth. We helped our industry partner to understand whether subsidies and information campaigns could help to grow their platform. In the end, the intervention did not have the intended effects, and the results of the experiment led them to think differently about their growth strategies. More broadly, beyond carpooling, the evidence from our field experiment shows that information campaigns and subsidies—which in theory spur growth by moving the market to a better equilibrium—may be ineffective or may even backfire in settings where participants have preferences that are high dimensional and therefore not easy to manipulate in the desired direction. In particular, in the carpooling setting we study, drivers have preferences over the "quality" of a ride which is informed by many dimensions of the ride, including the passenger's route, the passenger's gender and so on. Before spending resources on (potentially expensive) interventions, managers should examine the relevant dimensions of users' preferences that may go beyond price and quantity.

A valuable avenue for future research is to better understand, theoretically and empirically, why these strategies fail in settings like the one we studied, and what strategies could be better tailored to these settings. For example, it would be valuable to better understand *why* users are so persistent, and thus unlikely to react to information. In addition, understanding how to better match users along non-price dimensions using estimated models of user choice could be enriched by models of prosocial behavior, like that in Benabou and Tirole (2006). These research avenues, together, could generate insights into how platforms can deploy growth strategies that involve user-specific subsidies to the *quality* or other non-monetary dimensions, analogous to subsidies in settings where prices move users in and out of the market.

Armstrong, Mark. 2006. "Competition in Two-Sided Markets." *The RAND Journal of Economics* 37 (3): 668–91.

Azevedo, Eduardo M., and Jacob D. Leshno. 2016. "A Supply and Demand Framework for Two-Sided Matching Markets." *Journal of Political Economy* 124 (5): 1235–68.

Benabou, Roland, and Jean Tirole. 2003. "Intrinsic and Extrinsic Motivation." *The Review of Economic Studies* 70 (3): 489–520.

Boudreau, Kevin J. 2021. "Promoting Platform Takeoff and Self-Fulfilling Expectations: Field Experimental Evidence." *Management Science* 67 (9): 5953–67.

Caillaud, Bernard, and Bruno Jullien. 2003. "Chicken & Egg: Competition among Intermediation Service Providers." *RAND Journal of Economics*, 309–28.

Chen, Kuan-Ming, Ning Ding, John List, and Magne Mogstad. n.d. "Reservation Wages and Workers' Valuation of Job Flexibility: Evidence from a Natural Field Experiment."

Cullen, Zoë, and Chiara Farronato. 2020. "Outsourcing Tasks Online: Matching Supply and Demand on Peer-to-Peer Internet Platforms." *Management Science*.

Einav, Liran, Chiara Farronato, and Jonathan Levin. 2016. "Peer-to-Peer Markets." *Annual Review of Economics* 8: 615–35.

Farronato, Chiara, Jessica Fong, and Andrey Fradkin. 2023. "Dog Eat Dog: Balancing Network Effects and Differentiation in a Digital Platform Merger." *Management Science*.

Farronato, Chiara, and Andrey Fradkin. 2022. "The Welfare Effects of Peer Entry: The Case of Airbnb and the Accommodation Industry." *American Economic Review* 112 (6): 1782–1817.

Fradkin, Andrey, Elena Grewal, Dave Holtz, and Matthew Pearson. 2015. "Bias and Reciprocity in Online Reviews: Evidence From Field Experiments on Airbnb." In , 641–641. Electronic Commerce. ACM. https://doi.org/10.1145/2764468.2764528.

Gneezy, Uri, and Aldo Rustichini. 2000. "Pay Enough or Don't Pay at All." *The Quarterly Journal of Economics* 115 (3): 791–810. https://doi.org/10.1162/003355300554917.

Hitsch, G.J., Ali Hortascu, and Dan Ariely. 2010. "Matching and Sorting in Online Dating." *American Economic Review* 100 (1): 130–63.

Katz, Michael L., and Carl Shapiro. 1985. "Network Externalities, Competition, and Compatibility." *The American Economic Review* 75 (3): 424–40.

Ostrovsky, Michael, and Michael Schwarz. 2019. "Carpooling and the Economics of Self-Driving Cars." In *Proceedings of the 2019 ACM Conference on Economics and Computation*, 581–82.

Parker, Geoffrey G., and Marshall W. Van Alstyne. 2005. "Two-Sided Network Effects: A Theory of Information Product Design." *Management Science* 51 (10): 1494–1504.

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

Rochet, Jean-Charles, and Jean Tirole. 2003. "Platform Competition in Two-Sided Markets." *Journal of the European Economic Association* 1 (4): 990–1029.

Weyl, E. Glen. 2010. "A Price Theory of Multi-Sided Platforms." *American Economic Review* 100 (4): 1642–72.

Tables and Figures

Platform Growth is Hard:

Experimental and Quasi-Experimental Evidence from a Carpooling Platform in Singapore

List of Tables

List of Figures

Table 1: **Summary statistics**

Panel A	
Market Level Statistics	
Total Drivers	111,248
Total Passengers	730,324
Total Completed Rides	4,493,903
Percent Drivers w> 0 Rides	71.4
Percent Passengers w> 0 Rides	59.0
Percent Drivers Female	18.2
Percent Passengers Female	54.1
Panel B	
Driver Level Statistics	
Mean Total Plans	122.3
Mean Total Rides	57.4
Median Plan Dist. (km)	21.9
Median Detour Completed Rides	2.6
Mean Fare Completed Rides	10.1
Panel C	
Passenger Level Statistics	
Mean Total Bookings	18.2
Mean Total Rides	6.2
Mean Booking Dist. (km)	14.3

NOTE. Summary statistics of the GrabHitch market in Singapore, between January 1, 2016 and July 10, 2017. Panel A reports market level statistics. Panel B reports statistics about the universe of drivers, while Panel C reports statistics within the universe of passengers.

Figure 1: **Number of active users over time**

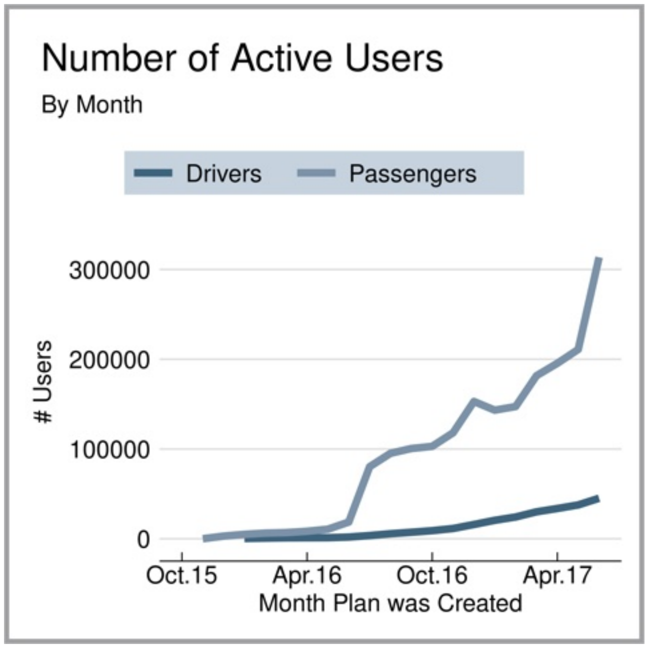

NOTE. This graph shows the number of active drivers and passengers in each week from the inception of GrabHitch up to the day before our intervention. A driver is considered "active" in a given week if she enters a plan in that week. A passenger is considered "active" in a given week if she makes a booking that week.

Figure 2: **Number of users vs. match rates**

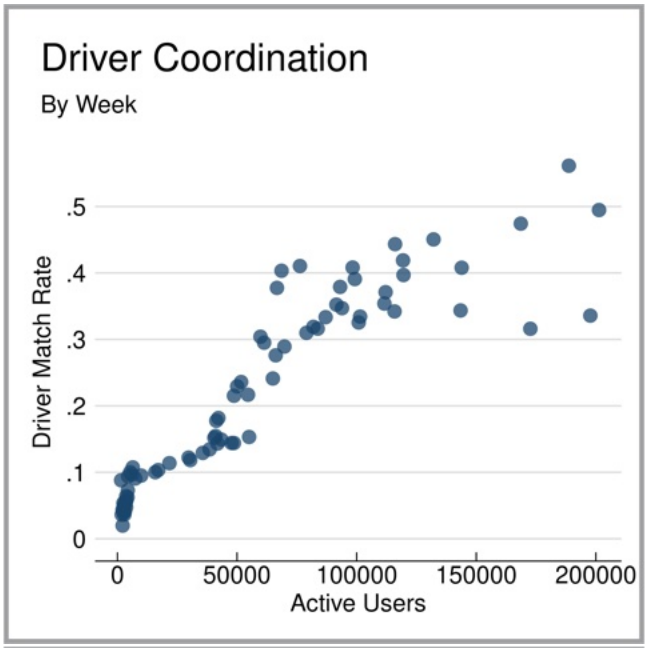

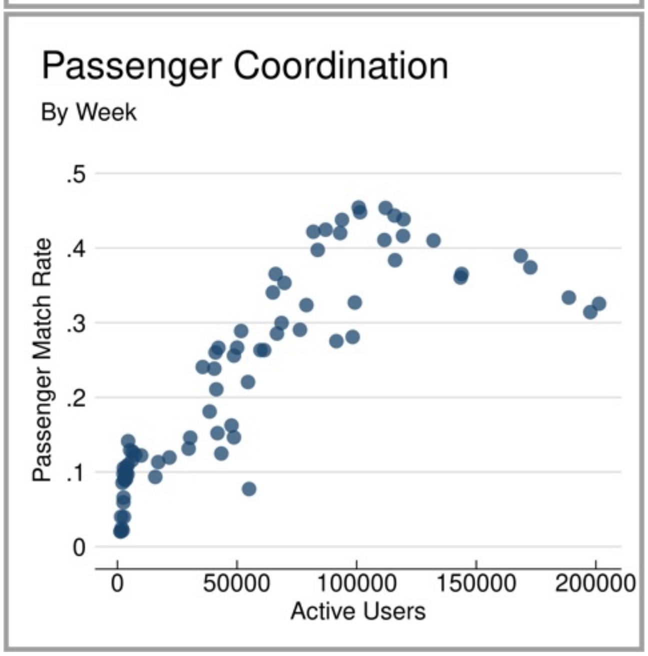

Figure 3: **Timeline of intervention**

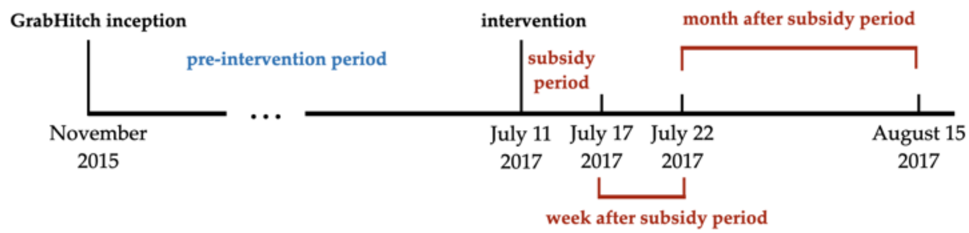

Table 2: **Treatment effects: Made plan or gave ride**

	(1) Made Plan (During)	(2) Made Plan (Week After)	(3) Made Plan (Month After)	(4) Gave Ride (During)	(5) Gave Ride (Week After)	(6) Gave Ride (Month After)
Density Low	0.011	-0.008	-0.004	0.002	0.000	0.007
	(0.009)	(0.010)	(0.014)	(0.006)	(0.006)	(0.009)
Density High	0.002	-0.018*	-0.035***	0.001	-0.007	-0.001
	(0.009)	(0.010)	(0.013)	(0.006)	(0.006)	(0.009)
Excess Demand	-0.001	-0.026***	-0.042***	-0.012**	-0.008	-0.001
	(0.009)	(0.009)	(0.013)	(0.005)	(0.006)	(0.009)
Small Bonus	0.008	-0.019**	-0.028**	-0.002	-0.006	-0.009
	(0.009)	(0.010)	(0.014)	(0.006)	(0.006)	(0.009)
Large Bonus	0.016*	-0.021**	-0.027**	0.011*	-0.010*	-0.003
	(0.009)	(0.009)	(0.014)	(0.006)	(0.006)	(0.009)
Companionship	0.007	-0.012	-0.022*	-0.003	-0.006	-0.003
	(0.009)	(0.010)	(0.014)	(0.005)	(0.006)	(0.009)
Observations	10401	10401	10401	10401	10401	10401
Control Mean	0.066	0.085	0.199	0.025	0.029	0.066

Note. Each column in this table comes from a separate OLS regression of respective outcome on the treatment and strata fixed effects. The control group is drivers who received message with no informational content. Huber-White robust estimates of the standard errors are reported in parentheses. Asterisks are based on standard p-values ($^{*}p < 0.1,^{**} p < 0.05,^{***} p < 0.01$).

Panel A: Male drivers

Table 3: **Treatment effects: Gender heterogeneity**

	(1) Made Plan (During Subsidy)	(2) Made Plan (Week After)	(3) Made Plan (Month After)	(4) Gave Ride (During Subsidy)	(5) Gave Ride (Week After)	(6) Gave Ride (Month After)
Density Low	0.011	-0.006	-0.012	0.003	-0.006	0.002
	(0.011)	(0.011)	(0.016)	(0.007)	(0.007)	(0.011)
Density High	-0.012	-0.025**	-0.047***	-0.004	-0.014**	-0.010
	(0.010)	(0.011)	(0.015)	(0.006)	(0.007)	(0.010)
Excess Demand	-0.003	-0.025**	-0.046***	-0.013**	-0.014**	-0.004
	(0.010)	(0.011)	(0.015)	(0.006)	(0.007)	(0.010)
Small Bonus	0.009	-0.014	-0.024	-0.001	-0.008	-0.007
	(0.010)	(0.011)	(0.015)	(0.006)	(0.007)	(0.010)
Large Bonus	0.009	-0.020*	-0.042***	0.010	-0.018***	-0.009
	(0.010)	(0.011)	(0.015)	(0.007)	(0.007)	(0.010)
Companionship	0.005	-0.015	-0.030**	-0.004	-0.010	-0.008
	(0.010)	(0.011)	(0.015)	(0.006)	(0.007)	(0.010)
Observations	8169	8169	8169	8169	8169	8169
Control Mean	0.067	0.086	0.202	0.025	0.034	0.070

Panel B: Female drivers

	(1) Made Plan (During Subsidy)	(2) Made Plan (Week After)	(3) Made Plan (Month After)	(4) Gave Ride (During Subsidy)	(5) Gave Ride (Week After)	(6) Gave Ride (Month After)
Density Low	0.012	-0.016	0.024	0.002	0.019*	0.023
	(0.019)	(0.020)	(0.029)	(0.012)	(0.010)	(0.019)
Density High	0.052**	0.008	0.003	0.021	0.021*	0.033
	(0.022)	(0.022)	(0.030)	(0.015)	(0.011)	(0.020)
Excess Demand	0.000	-0.033*	-0.033	-0.009	0.011	0.011
	(0.018)	(0.019)	(0.029)	(0.011)	(0.010)	(0.019)
Small Bonus	0.003	-0.039**	-0.041	-0.009	0.001	-0.016
	(0.019)	(0.019)	(0.029)	(0.012)	(0.008)	(0.017)
Large Bonus	0.041**	-0.025	0.028	0.010	0.016	0.020
	(0.021)	(0.020)	(0.030)	(0.013)	(0.010)	(0.019)
Companionship	0.010	-0.000	-0.002	-0.001	0.008	0.015
	(0.020)	(0.022)	(0.030)	(0.012)	(0.009)	(0.019)
Observations	2160	2160	2160	2160	2160	2160
Control Mean	0.062	0.084	0.186	0.025	0.009	0.053

Note. Each column in this table comes from a separate OLS regression of respective outcome on the treatment and strata fixed effects—Panel A includes only drivers who report gender "male" and Panel B includes only drivers who report gender "female". The control group is drivers who received message with no informational content. Huber-White robust estimates of the standard errors are reported in parentheses. Asterisks are based on standard p-values ($^{*}p < 0.1$,$^{**}p < 0.05$,$^{***}p < 0.01$).

Table 4: **Treatment effects: Quality of rides given**

	Relative Detour			Quality Weighted Rides		
	(1)	(2)	(3)	(1)	(2)	(3)
	During	Week After	Month After	During	Week After	Month After
Density Low	-0.006	-0.014	-0.003	-0.065	0.131	0.091
	(0.061)	(0.052)	(0.033)	(0.180)	(0.303)	(0.839)
Density High	0.015	0.018	0.005	0.071	-0.205	-0.156
	(0.065)	(0.053)	(0.037)	(0.210)	(0.244)	(0.844)
Excess Demand	-0.128**	0.008	-0.030	-0.143	-0.148	0.395
	(0.053)	(0.056)	(0.035)	(0.182)	(0.250)	(0.895)
Small Bonus	-0.047	0.072	-0.029	0.455	-0.050	0.201
	(0.063)	(0.057)	(0.037)	(0.449)	(0.249)	(0.864)
Large Bonus	-0.020	0.066	0.008	0.513*	-0.219	-0.485
	(0.057)	(0.066)	(0.036)	(0.280)	(0.249)	(0.792)
Companionship	-0.023	-0.054	-0.034	-0.129	-0.203	0.195
	(0.074)	(0.043)	(0.035)	(0.171)	(0.234)	(0.865)
Observations	170	171	517	10401	10401	10401
Control Mean	0.412	0.346	0.384	0.491	0.692	3.216

NOTE. Each column in this table comes from a separate OLS regression of respective outcome on the treatment and strata fixed effects. Huber-White robust estimates of the standard errors are reported in parentheses. Asterisks are based on standard p-values ($^{*}p < 0.1, ^{**}p < 0.05, ^{***}p < 0.01$).

Table 5: **Treatment effects: Effect of any message on driver activity**

Panel A: Made Plan and Gave Ride

	(1) Made Plan (During Subsidy)	(2) Made Plan (Week After)	(3) Made Plan (Month After)	(4) Gave Ride (During Subsidy)	(5) Gave Ride (Week After)	(6) Gave Ride (Month After)
Message	0.015**	0.001	-0.014	0.003	0.006	0.006
	(0.006)	(0.007)	(0.010)	(0.004)	(0.004)	(0.007)
Observations	11883	11883	11883	11883	11883	11883
Control Mean	0.057	0.070	0.191	0.022	0.018	0.059

Panel B: Number of Plans Made and Number of Rides Given

	(1) # Plans (During Subsidy)	(2) # Plans (Week After)	(3) # Plans (Month After)	(4) # Rides (During Subsidy)	(5) # Rides (Week After)	(6) # Rides (Month After)
Message	0.431**	0.197	0.179	0.228	0.118	0.136
	(0.196)	(0.189)	(0.146)	(0.283)	(0.309)	(0.203)
Observations	11883	11883	11883	11883	11883	11883
Control Mean	0.217	0.271	1.955	0.059	0.060	0.349

Note. Each column in this table comes from a separate OLS regression of respective outcome on the treatment (received any message at all) and strata fixed effects. The control group is drivers who received no message. Huber-White robust estimates of the standard errors are reported in parentheses. Panel A shows binary outcomes—*whether* a given driver made a plan or gave a ride. Panel B shows counts—*how many* plans and rides were given by drivers. Asterisks are based on standard p-values ($^{*}p < 0.1$,$^{**}p < 0.05$,$^{***}p < 0.01$).

Figure 4: **Event study: Activity before vs after intervention**

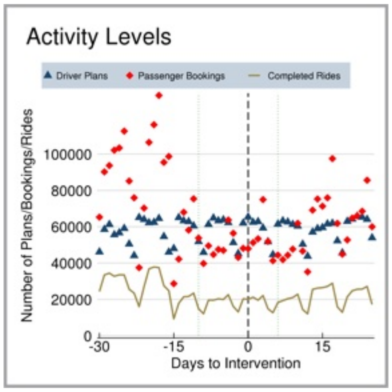

(a)

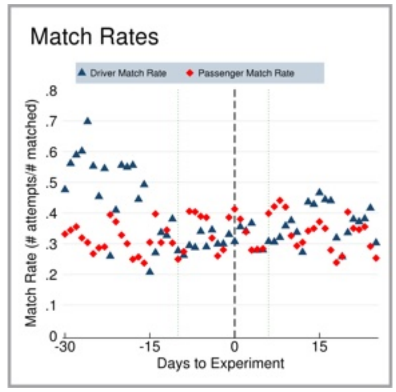

(b)

Note. Panel A shows driver and passenger activity, and the number of completed rides, in the entire market in the 30 days before and after our intervention. Panel B shows driver match rates (number of completed plans/number of driver plans) and passenger match rates (number of completed plans/number of passenger bookings) in the entire market before and after our intervention. In both panels, the black dashed line indicates the date of our intervention, while the light green dotted lines indicate the start and end days of the period used in our event study specification.

Figure 5: **User persistence**

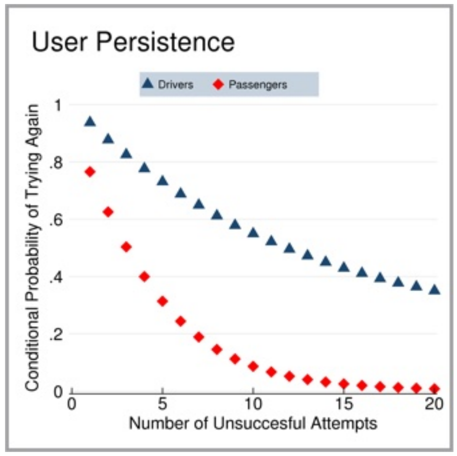

NOTE. This figure reports the conditional probability that drivers and passengers make X+1 attempts given that X attempts have been unsuccessful. These conditional probabilities are calculated off the universe of administrative data from January 2016 until the day before our intervention in July 2017. Passenger persistence is marked in red, while driver persistence is marked in blue.

Appendix

Platform Growth is Hard:

Experimental and Quasi-Experimental Evidence from a Carpooling Platform in Singapore

October 2, 2023

List of Tables

List of Figures

A Construction of the Quality-Weighted Ride Metric

We constructed our measure of quality weighted rides as follows. We looked at a database that shows the compatible passenger bookings shown to each driver when they enter a plan. Then, we analyzed driver choices. Note that when a driver chooses a passenger booking, they do not necessarily complete the ride because passengers could still cancel or the driver may change her mind.

For each driver, we looked at how six factors influence their choice of passenger. The six factors we look at are: relative detour—that is distance that the passenger booking adds to the driver's entered plan, divided by the driver's entered plan (both in kilometers); passenger gender; an interaction term for the driver gender and the passenger gender; the number of seats requested by the passenger; the difference between the start time of the passenger's booking and the driver's plan (in minutes); whether the passenger has a photo.

We trained the data on histories of driver choices before our intervention, using all of drivers in our sample. For every driver plan p and booking i we estimated the following OLS model

$$\begin{aligned} y_{ip} = \alpha_p &+ \beta_1(\text{relative detour})_{ip} + \beta_2\ (\text{passenger female})_{ip} \\ &+ \beta_3(\text{passenger female})_{ip} \times (\text{driver female})_{ip} + \beta_4(\text{booking-plan time difference})_{ip} \\ &+ \beta_5(\#\text{ of seats requested})_{ip} + \beta_6(\text{passenger has photo})_{ip} + \varepsilon_{ip} \end{aligned} \tag{1}$$

and a Logit equivalent (letting $\mathbf{x}_{ip}$ be the vector of passenger characteristics used in Equation 1)

$$\mathbb{P}(y_{ip} = 1 | x_{ip}, \alpha_p) = \frac{e^{\alpha_p + \beta \mathbf{x}_{ip}}}{1 + e^{\alpha_p + \beta \mathbf{x}_{ip}}} \tag{2}$$

As shown in Table A13, relative detour, passenger gender, booking-plan time difference, and number of seats requested are all strong predictors of driver choice.

After estimating the coefficients above, we predict a fitted quality "score" for each passenger candidate in the message treatment groups. So, for each passenger booking i that was ultimately given a ride by driver in our sample (excluding pure control) who made plan p, we estimated the plan-booking-specific probability that that booking would be chosen by that driver. In other words, for each passenger booking p we estimated $\hat{y}_{ip}$ for OLS and $\mathbb{P}(\hat{y}_{ip} = 1 | x_{ip}, \alpha_p))$ for Logit. Then, we normalize these values by dividing by the maximum score in the entire set of drivers who gave rides. Then, our "quality-weighted rides" outcome is simply the number of rides given by a driver weighted by this quality score, which is on a 0-100 scale.

To define this measure more formally, let $\mathcal{D}$ be the set of drivers who complete a ride in our sample. Let P_d be the set of plans logged by driver d and let $C(P_d) \subseteq P_d$be the subset of plans completed by driver d, with generic element p_d. Define $\hat{Y}$ as

$$\max_{d \in \mathcal{D}, p_d \in C(P_d)} \hat{y}_{ip} = \hat{Y}. \tag{3}$$

and q_d as,

$$q_d = \frac{\hat{y}_{ip}}{\hat{Y}} \times 100. \tag{4}$$

Our quality weighted rides outcome for driver d is Q_d, which is defined using Equation 3 and Equation 4. It is the sum of driver d's completed rides weighted by the quality score of each completed ride, i.e.

$$Q_d = \sum_{p_d \in P_d} q_d \mathbf{1}[p_d \in C(P_d)] \tag{5}$$

where 1 is the indicator function.

B Supplementary Material

Table A1: Sample selection

	More Dormant	Less Dormant	Total
Control (No Message)	1,121	361	1,482
Control (Message)	1,131	364	1,495
Density Low	1,133	349	1,482
Density High	1,127	366	1,493
Excess Demand	1,127	358	1,485
Small Bonus	1,124	360	1,484
Large Bonus	1,111	355	1,466
Companionship	1,131	365	1,496
Total	9,005	2,878	11,883

NOTE. This table lists the total number of drivers from each stratum (more dormant, less dormant) assigned to each treatment arm.

Table A2: Balance check

	(1) Separate Regressions						(2) Joint Regression
	Density Low	Density High	Excess Demand	Small Bonus	Large Bonus	Companionship	Treatment on Variables
Driver Gender	-0.023	0.014	0.015	0.025*	-0.001	0.024	-0.019
	(0.015)	(0.015)	(0.015)	(0.015)	(0.015)	(0.015)	(0.047)
Prior Plans	-1.301	0.957	-2.087	-1.868	-0.562	-0.918	-0.000
	(3.577)	(3.570)	(3.575)	(3.575)	(3.586)	(3.568)	(0.000)
Prior Rides	-0.302	1.277	-0.368	-0.637	0.066	-0.417	0.001
	(0.871)	(0.869)	(0.871)	(0.871)	(0.873)	(0.869)	(0.001)
Observations	10401	10401	10401	10401	10401	10401	10401
F-test (p-value)							0.653

Table A3: Minimum detectable effects for Table 2

Outcome	MDE
Made plan (during subsidy period)	0.025
Made plan (week after subsidy period)	0.029
Made plan (month after subsidy period)	0.041
Gave ride (during subsidy period)	0.016
Gave ride (week after subsidy period)	0.017
Gave ride (month after subsidy period)	0.023

NOTE. Assuming power=0.8.

Table A4: Messages to treatment groups

Density Low: "We miss you! Offer a ride to a fellow commuter today! The number of passengers on Hitch has increased 24% since you last completed a ride!"

Density High: "We miss you! Offer a ride to a fellow commuter today! The number of passengers on Hitch has increased 48% since you last completed a ride!"

Excess Demand: "We miss you! Offer a ride to a fellow commuter today! Since your last ride, the monthly number of unmatched passengers has grown by 814,643!"

Small Bonus: "We miss you! Offer a ride to a fellow commuter on 11-16 July & earn S$4 EXTRA per ride! T&C apply."

Large Bonus: "We miss you! Offer a ride to a fellow commuter on 11-16 July & earn S$8 EXTRA per ride! T&C apply."

Companionship: "We miss you! Offer a ride to a fellow commuter today and meet new friends!"

Control (Message): "We miss you! Offer a ride to a fellow commuter today!"

NOTE. Text (verbatim) of messages sent via SMS to drivers in the experimental sample.

Figure A1: Completed rides over time

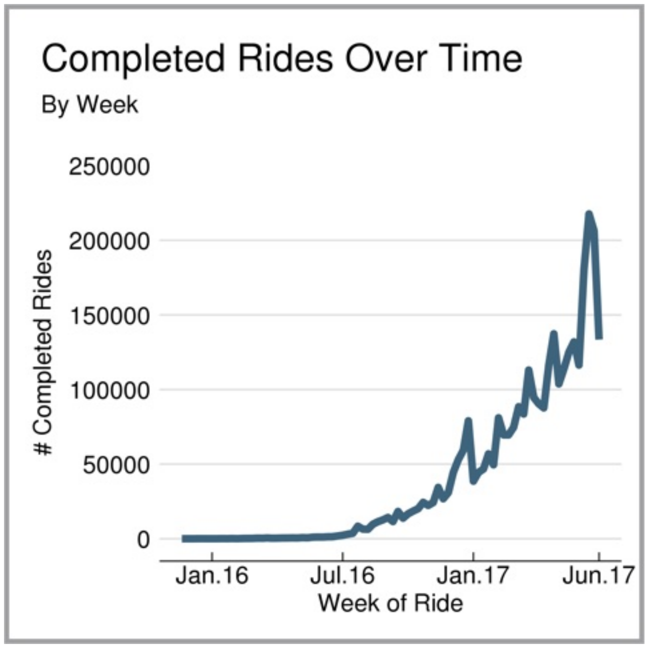

Note. This figure shows the number of completed rides per week from the inception of GrabHitch to the month after our intervention.

Figure A2: Completed rides over time: Top 10 neighborhoods

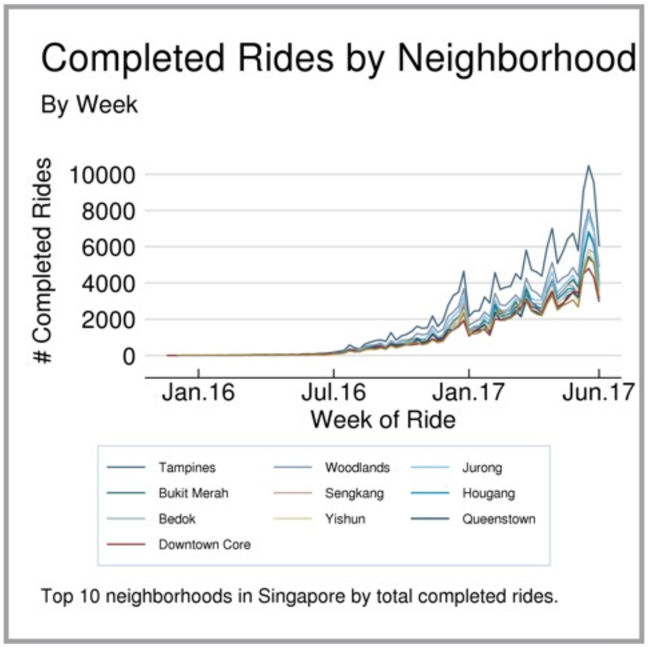

NOTE. This figure shows the number of completed rides per week from the inception of GrabHitch to the month after our intervention in the top 10 neighborhoods. The top 10 neighborhoods are the 10 neighborhoods with the highest number of total completed rides on the platform before our intervention.

Figure A3: Frequency of plans and bookings by time requested

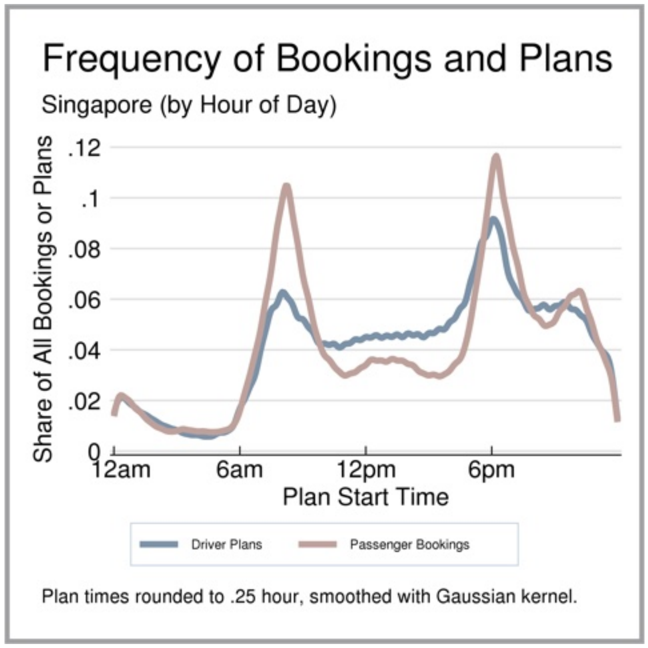

NOTE. This figure shows the frequency of passenger bookings and driver plans by hour of day. Plan and booking times are rounded to the quarter-hour, and the distribution is smoothed with a Gaussian kernel.

Table A5: Treatment effects: Number of plans made and rides given

	(1) # Plans (During Subsidy)	(2) # Plans (Week After)	(3) # Plans (Month After)	(4) # Rides (During Subsidy)	(5) # Rides (Week After)	(6) # Rides (Month After)
Density Low	-0.060	0.002	0.017	-0.235	0.112	0.059
	(0.240)	(0.235)	(0.215)	(0.301)	(0.292)	(0.241)
Density High	-0.080	-0.190	-0.161	-0.184	0.065	0.006
	(0.254)	(0.237)	(0.218)	(0.299)	(0.332)	(0.240)
Excess Demand	-0.593**	-0.255	-0.303	-1.147***	-0.240	-0.164
	(0.247)	(0.227)	(0.193)	(0.369)	(0.308)	(0.232)
Small Bonus	-0.115	-0.151	-0.230	-0.087	-0.021	-0.210
	(0.262)	(0.233)	(0.216)	(0.329)	(0.310)	(0.259)
Large Bonus	-0.032	-0.221	-0.164	0.076	0.001	-0.116
	(0.235)	(0.245)	(0.206)	(0.279)	(0.319)	(0.235)
Companionship	-0.295	-0.362*	-0.195	-0.939***	-0.337	-0.219
	(0.239)	(0.216)	(0.211)	(0.295)	(0.309)	(0.224)
Observations	10401	10401	10401	10401	10401	10401
Control Mean	0.373	0.376	2.641	0.096	0.070	0.431

Note. Each column in this table comes from a separate Poisson regression of respective outcome on the treatment and strata fixed effects. The control group is drivers who received message with no informational content. Huber-White robust estimates of the standard errors are reported in parentheses. Asterisks are based on standard p-values ($^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$).

Table A6: Treatment effects: Made plan or gave ride (with driver history)

	(1) Made Plan (During Subsidy)	(2) Made Plan (Week After)	(3) Made Plan (Month After)	(4) Gave Ride (During Subsidy)	(5) Gave Ride (Week After)	(6) Gave Ride (Month After)
Density Low	0.011	-0.008	-0.003	0.003	0.000	0.007
	(0.009)	(0.010)	(0.014)	(0.006)	(0.006)	(0.009)
Density High	0.002	-0.018*	-0.035***	0.001	-0.007	-0.001
	(0.009)	(0.010)	(0.013)	(0.006)	(0.006)	(0.009)
Excess Demand	-0.001	-0.026***	-0.042***	-0.012**	-0.008	-0.001
	(0.009)	(0.009)	(0.013)	(0.005)	(0.006)	(0.009)
Small Bonus	0.008	-0.018*	-0.027**	-0.002	-0.006	-0.009
	(0.009)	(0.009)	(0.014)	(0.006)	(0.006)	(0.009)
Large Bonus	0.016*	-0.021**	-0.027**	0.011*	-0.010*	-0.003
	(0.009)	(0.009)	(0.014)	(0.006)	(0.006)	(0.009)
Companionship	0.007	-0.012	-0.022	-0.003	-0.006	-0.003
	(0.009)	(0.010)	(0.014)	(0.005)	(0.006)	(0.009)
Observations	10401	10401	10401	10401	10401	10401
Control Mean	0.066	0.085	0.199	0.025	0.029	0.066

NOTE. Each column in this table comes from a separate OLS regression of respective outcome on the treatment, strata fixed effects, and drivers' number of previous plans. The control group is drivers who received message with no informational content. Huber-White robust estimates of the standard errors are reported in parentheses. Asterisks are based on standard p-values ($^{*}p < 0.1,^{**} p < 0.05,^{***} p < 0.01$).

Panel A: Male drivers

Table A7: Treatment effects: Gender heterogeneity (with driver history)

	(1) Made Plan (During Subsidy)	(2) Made Plan (Week After)	(3) Made Plan (Month After)	(4) Gave Ride (During Subsidy)	(5) Gave Ride (Week After)	(6) Gave Ride (Month After)
Density Low	0.011	-0.007	-0.012	0.003	-0.006	0.002
	(0.010)	(0.011)	(0.016)	(0.007)	(0.007)	(0.011)
Density High	-0.012	-0.025**	-0.047***	-0.004	-0.014**	-0.010
	(0.010)	(0.011)	(0.015)	(0.006)	(0.007)	(0.010)
Excess Demand	-0.002	-0.025**	-0.045***	-0.013**	-0.014**	-0.004
	(0.010)	(0.011)	(0.015)	(0.006)	(0.007)	(0.010)
Small Bonus	0.009	-0.014	-0.024	-0.001	-0.008	-0.007
	(0.010)	(0.011)	(0.015)	(0.006)	(0.007)	(0.010)
Large Bonus	0.009	-0.021*	-0.042***	0.010	-0.018***	-0.009
	(0.010)	(0.011)	(0.015)	(0.007)	(0.007)	(0.010)
Companionship	0.005	-0.015	-0.031**	-0.004	-0.010	-0.008
	(0.010)	(0.011)	(0.015)	(0.006)	(0.007)	(0.010)
Observations	8169	8169	8169	8169	8169	8169
Control Mean	0.067	0.086	0.202	0.025	0.034	0.070

Panel B: Female drivers

	(1) Made Plan (During Subsidy)	(2) Made Plan (Week After)	(3) Made Plan (Month After)	(4) Gave Ride (During Subsidy)	(5) Gave Ride (Week After)	(6) Gave Ride (Month After)
Density Low	0.014	-0.014	0.026	0.002	0.019*	0.022
	(0.019)	(0.020)	(0.029)	(0.012)	(0.010)	(0.019)
Density High	0.053**	0.009	0.004	0.021	0.021*	0.033
	(0.022)	(0.022)	(0.030)	(0.015)	(0.011)	(0.020)
Excess Demand	0.002	-0.031	-0.031	-0.009	0.011	0.011
	(0.018)	(0.019)	(0.029)	(0.011)	(0.010)	(0.019)
Small Bonus	0.004	-0.038**	-0.040	-0.008	0.001	-0.016
	(0.019)	(0.019)	(0.029)	(0.012)	(0.008)	(0.017)
Large Bonus	0.043**	-0.023	0.031	0.011	0.016	0.020
	(0.021)	(0.020)	(0.030)	(0.013)	(0.010)	(0.019)
Companionship	0.013	0.002	0.001	-0.001	0.008	0.015
	(0.020)	(0.022)	(0.030)	(0.012)	(0.009)	(0.019)
Observations	2160	2160	2160	2160	2160	2160
Control Mean	0.062	0.084	0.186	0.025	0.009	0.053

Note. Each column in this table comes from a separate OLS regression of respective outcome on the treatment, strata fixed effects, and drivers' number of previous plans. The control group is drivers who received message with no informational content. Huber-White robust estimates of the standard errors are reported in parentheses. Asterisks are based on standard p-values ($^{*}p < 0.1,^{**} p < 0.05,^{***} p < 0.01$).

Panel A: More dormant drivers

Table A8: Treatment effects: Strata heterogeneity (more versus less dormant drivers)

	(1) Made Plan (During Subsidy)	(2) Made Plan (Week After)	(3) Made Plan (Month After)	(4) Gave Ride (During Subsidy)	(5) Gave Ride (Week After)	(6) Gave Ride (Month After)
Density Low	0.017**	-0.006	0.002	0.005	-0.000	0.008
	(0.008)	(0.010)	(0.014)	(0.006)	(0.006)	(0.009)
Density High	0.007	-0.021**	-0.026*	0.004	-0.002	0.002
	(0.007)	(0.009)	(0.014)	(0.006)	(0.006)	(0.009)
Excess Demand	0.005	-0.019**	-0.032**	-0.009*	-0.006	-0.003
	(0.007)	(0.009)	(0.013)	(0.005)	(0.006)	(0.009)
Small Bonus	0.009	-0.021**	-0.024*	-0.002	-0.002	-0.008
	(0.008)	(0.009)	(0.014)	(0.005)	(0.006)	(0.009)
Large Bonus	0.012	-0.029***	-0.027**	0.007	-0.010*	-0.007
	(0.008)	(0.009)	(0.014)	(0.006)	(0.005)	(0.009)
Companionship	0.006	-0.019**	-0.030**	-0.004	-0.004	-0.006
	(0.007)	(0.009)	(0.013)	(0.005)	(0.006)	(0.009)
Observations	7884	7884	7884	7884	7884	7884
Control Mean	0.029	0.058	0.132	0.017	0.022	0.050

Panel B: Less dormant drivers

	(1) Made Plan (During Subsidy)	(2) Made Plan (Week After)	(3) Made Plan (Month After)	(4) Gave Ride (During Subsidy)	(5) Gave Ride (Week After)	(6) Gave Ride (Month After)
Density Low	-0.007	-0.013	-0.020	-0.006	0.002	0.002
	(0.029)	(0.028)	(0.037)	(0.016)	(0.016)	(0.024)
Density High	-0.015	-0.009	-0.062*	-0.008	-0.022	-0.012
	(0.028)	(0.027)	(0.036)	(0.015)	(0.014)	(0.023)
Excess Demand	-0.022	-0.047*	-0.074**	-0.022	-0.013	0.008
	(0.028)	(0.026)	(0.036)	(0.014)	(0.015)	(0.024)
Small Bonus	0.005	-0.012	-0.040	-0.005	-0.019	-0.013
	(0.029)	(0.027)	(0.036)	(0.016)	(0.015)	(0.023)
Large Bonus	0.027	0.004	-0.026	0.027	-0.007	0.011
	(0.029)	(0.028)	(0.036)	(0.018)	(0.016)	(0.025)
Companionship	0.010	0.008	0.002	-0.003	-0.014	0.008
	(0.029)	(0.028)	(0.036)	(0.016)	(0.015)	(0.024)
Observations	2517	2517	2517	2517	2517	2517
Control Mean	0.179	0.168	0.407	0.049	0.049	0.118

NOTE. Each column in this table comes from a separate OLS regression of respective outcome on the treatment and strata fixed effects—Panel A includes only drivers in our sample who are "more dormant" (have not given a ride in the last XX days) and Panel B includes only drivers in our sample who are "less dormant" (have not given a ride in the last XX days). The control group is drivers who received message with no informational content. Huber-White robust estimates of the standard errors are reported in parentheses. Asterisks are based on standard p-values ($^{*}p < 0.1,^{**}p < 0.05,^{***}p < 0.01$).

Table A9: Sample summary: Gender and prior experience

	Male	Female
Prior Plans	31.32	21.50
	(104.57)	(64.89)
N	8169	2160
Prior Rides	19.37	11.88
	(43.15)	(24.46)
N	2525	613

Table A10: Treatment effects: Men matched to women on experience

	(1) Made Plan (During Subsidy)	(2) Made Plan (Week After)	(3) Made Plan (Month After)	(4) Gave Ride (During Subsidy)	(5) Gave Ride (Week After)	(6) Gave Ride (Month After)
Density Low	0.020	-0.021	-0.043	0.015	-0.015	-0.007
	(0.020)	(0.023)	(0.033)	(0.014)	(0.014)	(0.022)
Density High	0.008	-0.027	-0.085***	0.009	-0.005	-0.007
	(0.019)	(0.022)	(0.032)	(0.013)	(0.015)	(0.022)
Excess Demand	0.019	-0.038*	-0.079**	-0.017*	-0.031***	-0.044**
	(0.020)	(0.021)	(0.032)	(0.009)	(0.012)	(0.019)
Small Bonus	0.009	-0.019	-0.043	0.004	-0.010	-0.015
	(0.019)	(0.022)	(0.033)	(0.012)	(0.014)	(0.021)
Large Bonus	0.021	-0.024	-0.078**	0.014	-0.031**	-0.034*
	(0.020)	(0.023)	(0.032)	(0.014)	(0.012)	(0.020)
Companionship	0.016	-0.016	-0.046	0.007	-0.010	-0.030
	(0.019)	(0.022)	(0.032)	(0.012)	(0.014)	(0.020)
Observations	2159	2159	2159	2159	2159	2159
Control Mean	0.055	0.092	0.226	0.021	0.038	0.082

NOTE. This table constructs a sample of men matched to women via coarsened exact matching on prior plans. Each column in this table comes from a separate OLS regression of respective outcome on the treatment and strata fixed effects. The control group is drivers who received message with no informational content. Huber-White robust estimates of the standard errors are reported in parentheses. Asterisks are based on standard p-values ($^{*}p < 0.1$,$^{**}p < 0.05$,$^{***}p < 0.01$).

Table A11: Recently joined drivers versus sample

	Sample	Recent
Plans in First Month	29.88	27.56
	(114.09)	(49.69)
Median	6	11
% Women	.21	.23
Observations	11,883	11,609

Note. “Recent” is composed of drivers who made their first plan in the 30 days prior to our intervention. “Sample” is composed of drivers who were included in our experiment. Prior plans computed in the first month after joining.

Table A12: Driver's relative detour

	(1) During Subsidy	(2) Week After	(3) Month After
Density Low	-0.001	-0.000	0.002
	(0.002)	(0.002)	(0.004)
Density High	0.000	-0.002	-0.000
	(0.002)	(0.002)	(0.004)
Excess Demand	-0.005***	-0.002	-0.001
	(0.002)	(0.002)	(0.003)
Small Bonus	-0.002	0.001	-0.003
	(0.002)	(0.002)	(0.003)
Large Bonus	0.003	-0.001	0.001
	(0.002)	(0.002)	(0.004)
Companionship	-0.001	-0.002	-0.001
	(0.002)	(0.002)	(0.003)
Observations	10401	10401	10401
Control Mean	0.007	0.007	0.019

NOTE. Each column in this table comes from a separate OLS regression of driver's relative detour distance on the treatment and strata fixed effects. Relative detour is the distance from the passenger's pickup location to the driver's start location, divided by the driver's total trip distance. Huber-White robust estimates of the standard errors are reported in parentheses. Asterisks are based on standard p-values ($^{*}p < 0.1,^{**}p < 0.05,^{***}p < 0.01$).

Table A13: Construction of quality metric

	(1) Logit	(2) OLS
Relative Detour	-6.5048***	-0.1056***
	(0.2499)	(0.0037)
Passenger Female	0.5408***	0.0108***
	(0.0681)	(0.0013)
Passenger Female $\times$ Driver Female	-0.1620	-0.0023
	(0.2108)	(0.0042)
Plan-Booking Time Difference	-0.0472***	-0.0008***
	(0.0023)	(0.0000)
Number of Seats Requested	-0.3678***	-0.0058***
	(0.0621)	(0.0011)
Passenger Has Photo	-0.0629	-0.0010
	(0.0567)	(0.0012)
Observations	14291	78677
R^2		0.020
Pseudo R^2	0.213	
Overall Match Rate	0.1342	0.0246

Note. Panel 1 shows the coefficient of a logit regression on plan-booking compatibility characteristics. Panel 2 shows the coefficient of an OLS regression on the same characteristics. These coefficients are used in the construction of our quality-weighted ride metric described in Appendix 1.

Table A14: Event study coefficients

Panel A: Activity Levels

	(1) # Plans	(2) # Bookings	(3) # Completed	(4) # Drivers	(5) # Passengers
During Subsidy Period	551.27	4092.90	556.700	403.43	3049.96
	(4023.52)	(5053.95)	(1883.00)	(1897.41)	(3330.76)
Observations	16	16	16	16	16
Control Mean	57389.63	50907.94	18090.56	30945.69	43590.44

Panel B: Match Rates

	(1) Pax Match Rate	(2) Dax Match Rate	(3) Pax Assign Rate	(4) Dax Assign Rate
During Subsidy Period	-0.006	0.019	0.000	0.027
	(0.031)	(0.018)	(0.038)	(0.024)
Observations	16	16	16	16
Control Mean	0.333	0.310	0.466	0.439

Note. Each column in this table comes from a separate OLS regression of respective outcome on a dummy for whether the day is after the subsidy period. Panel A reports outcomes that have to do with activity levels in the entire market, while Panel B reports match rates in the entire market. Huber-White robust estimates of the standard errors are reported in parentheses. Asterisks are based on standard p-values ($^{*}p < 0.1,^{**} p < 0.05,^{***} p < 0.01$).

Figure A4: Event study: Activity before vs after intervention (by quartiles)

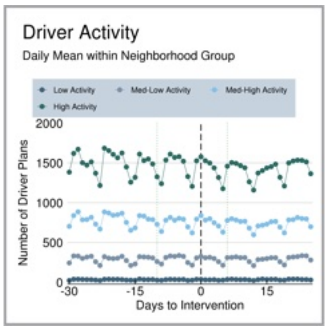

(a)

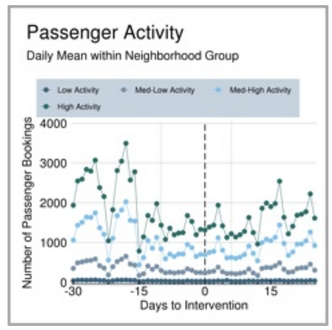

(b)

(c)

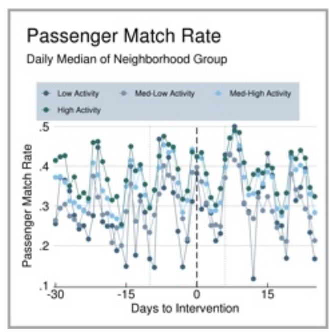

(d)

Note. These figures show market statistics by neighborhood quartiles: driver activity (Panel A), passenger activity (Panel B), driver match rates (Panel C), and passenger match rates (Panel D). Quartiles are defined using the number of completed rides in each neighborhood in the 60 days before our intervention. The black dashed line indicates the date of our intervention, while the light green dotted lines indicate the start and end days of the period used in our event study specification.

Figure A5: Histogram of match rates (area-week observations, 1 year before intervention)

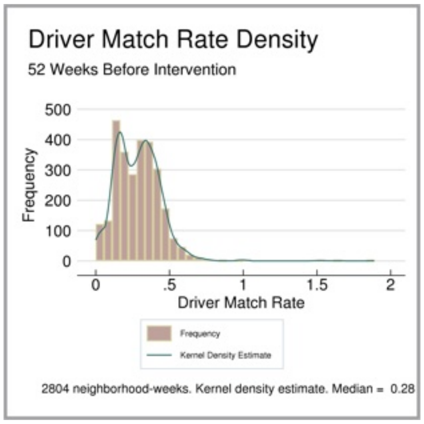

(a)

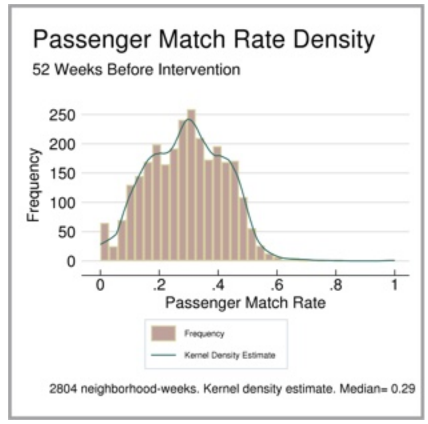

(b)

Note. These figures show a histogram of driver match rates (Panel A) and passenger match rates (Panel B) in all neighborhood-weeks in Singapore in the 52 weeks before our intervention. The green lines show kernel density estimates of the distribution.

Figure A6: Passenger and driver match rates (week-neighborhood, 52 weeks before intervention)

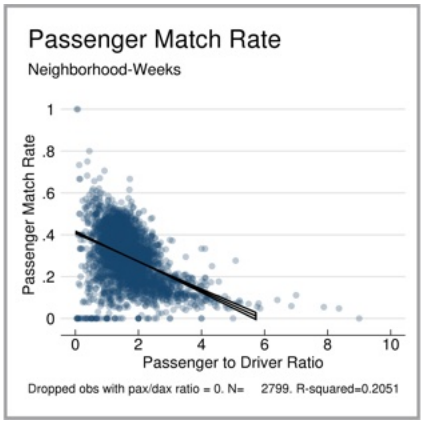

(a)

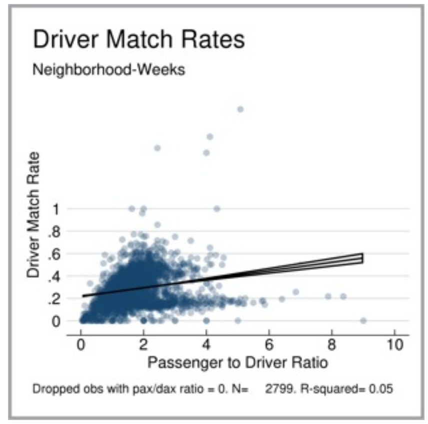

(b)

Note. These figures show a scatter plot of passenger-to-driver ratios versus driver match rates (Panel A) and versus passenger match rates (Panel B) in all neighborhood-weeks in Singapore in the 52 weeks before our intervention. The black lines are best-fit linear regressions.

Figure A7: Match rates by decile of passenger/driver ratio

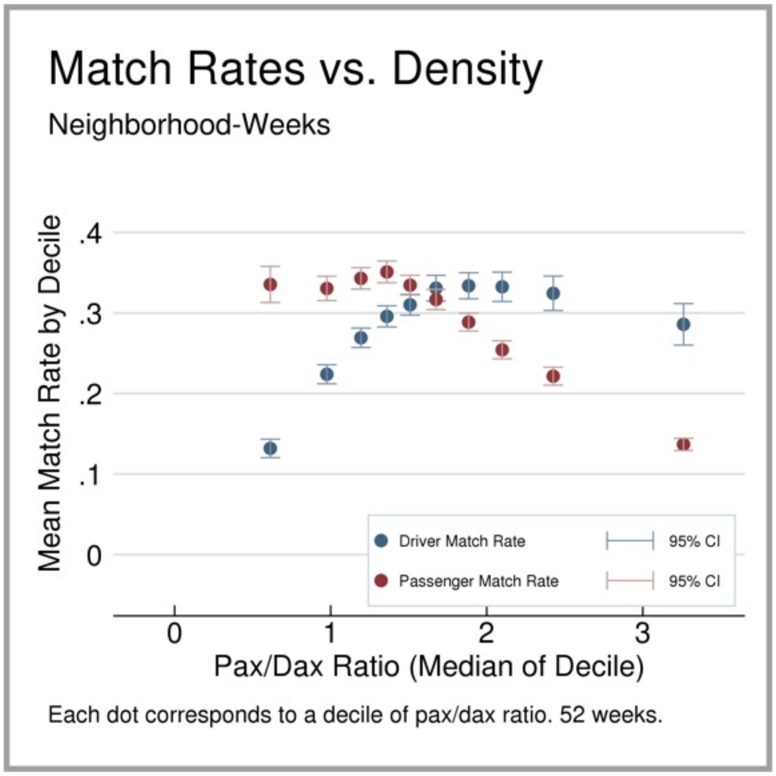

NOTE. This figure breaks all neighborhood-week observations into deciles of passenger-to-driver ratio. The median passenger-to-driver ratio value is plotted against the mean match rate within that decile. Bars around markers indicate 95% confidence intervals.