informs
https://pubsonline.informs.org/journal/msom

MANUFACTURING & SERVICE OPERATIONS MANAGEMENT
Vol. 25, No. 4, July–August 2023, pp. 1263–1284
ISSN 1523-4614 (print), ISSN 1526-5498 (online)

Finalist—2021 *M&SOM* Practice-Based Research Competition

Incentivizing Commuters to Carpool: A Large Field Experiment with Waze

Maxime C. Cohen,[a] Michael-David Fiszer,[b] Avia Ratzon,[b] Roy Sasson[b]

[a] Desautels Faculty of Management, McGill University, Montreal, Quebec H3W 1X 4, Canada; [b] Google/Waze, Tel Aviv 6789155, Israel

Contact: maxime.cohen@mcgill.ca, https://orcid.org/0000-0002-2474-3875 (MCC); mdfz@google.com (M-DF); aviar@google.com (AR); rsasson@google.com (RS)

Received: September 2, 2020
Revised: April 30, 2021
Accepted: July 8, 2021
Published Online in Articles in Advance: December 10, 2021

https://doi.org/10.1287/msom.2021.1033

Abstract. ***Problem definition***: Traffic congestion is a serious global issue. A potential solution, which requires zero investment in infrastructure, is to convince solo car users to carpool. ***Academic/practical relevance***: In this paper, we leverage the Waze Carpool service and run the largest ever digital field experiment to nudge commuters to carpool. ***Methodology***: Our field experiment involves more than half a million users across four U.S. states between June 10 and July 3, 2019. We identify users who can save a significant commute time by carpooling through the use of a high-occupancy vehicle (HOV) lane, users who can still use an HOV lane but have a low time saving, and users who do not have access to an HOV lane on their commute. We send them in-app notifications with different framings: mentioning the HOV lane, highlighting the time saving, emphasizing the monetary welcome bonus (for users who do not have access to an HOV lane), and a generic carpool invitation. ***Results***: We find a strong relationship between the affinity to carpool and the potential time saving through an HOV lane. ***Managerial implications***: Specifically, we estimate that mentioning the HOV lane increases the click-through rate (i.e., proportion of users who clicked on the button inviting them to try the carpool service) and the onboarding rate (i.e., proportion of users who signed up and created an account with the carpool service) by 133%–185% and 64%–141%, respectively, relative to a generic invitation. We conclude by discussing the implications of our findings for carpool platforms and public policy.

History: This paper has been accepted as part of the 2021 *Manufacturing & Service Operations Management* Practice-Based Research Competition.
Supplemental Material: The online appendix is available at https://doi.org/10.1287/msom.2021.1033.

Keywords: carpooling • field experiment • saving commute time via HOV lanes

1. Introduction

Most governments are devoting considerable efforts to alleviate traffic congestion. A study by INRIX conveys that Americans lost an average of 97 hours a year due to congestion, costing them nearly $87 billion in 2018, an average of $1,348 per driver.[1] Thus, it is not surprising that governments allocate substantial investments to reduce traffic congestion. Common policies (beyond expensive investments in infrastructures) include congestion pricing, restricted days, and high-occupancy vehicle (HOV) lanes, just to name a few. An alternative solution is to simply encourage solo car users to carpool.

According to data from the U.S. Census Bureau, the vast majority of Americans go to work by driving alone in their car. Over three-quarters (76.3%) choose to commute this way, with nearly identical numbers for men and women (this figure has been steady in 2015–2017). This translates into 116 million vehicles transporting exactly one person each. This research is in collaboration with Waze, the free community-driven global positioning system (GPS) navigation software app owned by Google. According to Waze data, close to two-thirds of drivers have at least one other regular Waze driver with a perfectly matching commute: that is, driving around the same time from the same origin to the same destination, within less than 500-meters radium---this is mind blowing.

What can persuade commuters to carpool and help reduce traffic congestion? In this paper, we study this question by testing several persuasion factors on a sample of 537,370 U.S. commuters. Specifically, we investigate to what extent highlighting the HOV lane and the resulting reduction in commute time is an effective incentive.

1.1. Setting and Research Questions

Waze launched a service to help users find matches for carpooling.[2] Initially centered on casual carpool requests from riders to Waze drivers with similar routes, the product rapidly shifted to focus on the commute use case, with an emphasis on planning a weekly schedule. Users can register as drivers or riders (or both) and are matched to other users with similar commuting patterns. Subsequently, users can send offers to each other to share a ride. Thus, drivers can accept incoming requests but also invite other users with similar routes to join their carpools. The platform takes care of trade-offs between rider walking and driver detour, handles pricing and payments, and proposes optimized pickup and drop-off locations.

Changing commuters' habits is not easy, but if done at scale, it can have a considerable impact in reducing congestion and pollution. As of October 2018, Waze has access to more than 110 million drivers globally, including more than 30 million in the United States.[3] The number of Waze drivers is abundant enough to provide the density and liquidity necessary for a significant fraction of drivers to leave their car at home---but how can one efficiently convince users to revamp their daily commute? For users who regularly use their car for commuting, it seems very challenging to convince them to start carpooling (i.e., by either opening their car to a stranger or by leaving their car at home and hopping in to someone else's car). Interestingly, however, Waze data show that after a driver does share a ride for the first time, retention (i.e., carpooling again) is high. Cracking the motivations and factors that will lead a driver to make the first step is thus critical.

The Waze Carpool platform is the ideal medium to study how different types of incentives can successfully convince commuters to carpool. For instance, internal econometric models show the extent to which drivers prefer matches that are precisely on route or involve a higher shared mileage and thus, higher cost savings. Overall, motivations can be financial, social, or related to reducing commute time: for instance, by using an HOV lane (more details on HOV lanes are discussed later in this paper).

In this paper, we consider two types of incentives: highlighting time saving and monetary compensation ($10 welcome bonus). We identify three types of users: (i) users who can save a significant commute time by carpooling through the use of an HOV lane, (ii) users who can still use the HOV lane but have a low time saving, and (iii) users who do not have access to an HOV lane on their commute. For each user, we leverage Waze data and algorithms to carefully estimate the potential time saving had this user used the HOV lane for his or her commute (see more details in Section 3.2). Our experimental population comprises 537,370 users across four U.S. states (we explain in greater detail how we select these users in Section 4). We then apply the following set of interventions.

- For commuters with a high time saving---who could save 6–40 minutes if they would carpool and use the HOV lane---we randomly split them into four conditions. Each condition involves sending the user an in-app notification with an invitation to try the carpool service. Specifically, we use the following four framings: (A) mentioning the HOV lane and the potential (high) time saving, (B) mentioning the HOV lane, (C) using a generic carpool invitation, and (D) not sending anything.
- For commuters with a low time saving (i.e., users who could save two to five minutes if they would carpool and use the HOV lane), we randomly split them into the same four conditions.
- For commuters who do not have access to an HOV lane on their commute, we randomly split them into three conditions: (A) mentioning the monetary incentive (receiving a $10 welcome bonus to try the carpool service), (B) using a generic invitation, and (C) not sending anything.

We note that our large-scale field experiment should be seen as three field experiments (one for each type of users), which are run in parallel. Consequently, our goal is not to directly compare the results across different types of users but instead, to provide a comprehensive study that encompasses several types of commuters.

Using the data from our field experiment, we aim to answer the following research questions.

1. What are the most successful framings to convince different types of commuters to carpool?
2. Is highlighting the time saving (for users who can save commute time by using an HOV lane) a successful persuasion factor? If yes, what is the impact on carpool intent and adoption?
3. For users who cannot save time by carpooling and use the HOV lane, is highlighting the monetary welcome bonus effective?

Ultimately, we aim to understand what the successful triggers are that can persuade different types of commuters to join the carpool platform.

1.2. Summary of Results

We conduct several analysis of variance (ANOVA) tests and regression analyses to estimate the impact of our field experiment. We capture users' carpool intent using two metrics, click-through rate (CTR) and onboarding rate (OBR), which are formally defined in Section 4.4. We next summarize our findings.

1.2.1. Mentioning the HOV Lane Is the Most Successful Framing. We consistently observe that highlighting the fact that commuters can use the HOV lane is effective. It increases carpool intent (measured by CTR and OBR) by 64%–185% relative to sending a generic carpool invitation. This result holds for users with both high and low time savings. This effect is further amplified for users with a longer commute distance or a higher time saving.

1.2.2. Mentioning the HOV Lane Is Enough. Highlighting both the HOV lane and the potential time saving does not yield an additional marginal impact relative to only mentioning the HOV lane. One possible conclusion is that mentioning the HOV lane is enough, so that users can directly translate this information into saving commute time. This result holds for users with both high and low time savings.

1.2.3. Quantifying the Economic Impact of Various Factors. Using our preexperiment data, we quantify the economic significance of several factors. For example, we find that an additional 10 minutes in average time saving boosts the carpool intent (captured by OBR) by 18%. We also observe that a fixed commute schedule (captured by a small standard deviation of the leave time) leads to a higher carpool intent. Using our experiment data, we find that users with a higher time saving are more receptive to the carpool offer in terms of CTR and OBR. Interestingly, this occurs for framings A and B (in which the HOV lane is highlighted) but not for framing C (generic carpool invitation). Thus, both the type of user and the framing of the intervention play important roles in incentivizing commuters to carpool.

1.2.4. Importance of Highlighting the Benefit of Carpooling to Targeted Users. Our results suggest that users with a high carpool utility (i.e., high time saving) are more receptive to the carpool service (i.e., higher CTR and OBR)---but only when the benefit of carpooling (in our case, saving time via the HOV lane) is explicitly mentioned in the notification. Thus, targeting high-intent users is effective only when the benefit is explicitly highlighted.

1.2.5. Mentioning the Monetary Incentive Is Not Effective. Highlighting the $10 welcome bonus to non-HOV users does not increase their response (CTR and OBR) to the carpool offer relative to a generic invitation. This finding suggests that highlighting the monetary bonus is not enough to nudge commuters to carpool (we will discuss the limitations of this result in Section 5).

1.2.6. Boosting Carpool Intent but Not Carpool Adoption. We find that our interventions substantially boost carpool intent but do not affect carpool adoption. Consequently, in order to stimulate carpool adoption, additional targeted follow-up interventions are needed.

2. Literature Review

This paper is related to several research streamlines, including field experiments in online platforms, behavioral nudging (e.g., mobile app adoption), and carpooling as a means of reducing congestion.

2.1. Field Experiments in Online Platforms

In recent years, it has become common practice for online platforms to routinely run A/B tests to generate high-quality data and learn users' preferences. Companies like Microsoft, Amazon, Booking.com, Facebook, and Google each conduct more than 10,000 online controlled experiments annually, with many tests engaging millions of users (Kohavi and Thomke 2017). Several researchers have recently used field experiments to compare different interventions targeted to online platforms' users. For example, in the context of ride-sharing, Cohen et al. (2021) examine the results of a field experiment that sent promotions to users who experienced a poor quality of service. In the context of commuting, Jachimowicz et al. (2021) use the data from three field studies and find that lengthy commutes are more aversive for employees with lower trait self-control. The authors conclude that the commute to work can be beneficial when seen as an opportunity to transition into one's work role. Singh et al. (2017) conduct field experiments to compare charity-linked promotions to discount-based promotions in the context of an online taxi-booking platform. To our knowledge, our paper provides the first digital large-scale field experiment with the goal of convincing commuters to carpool.

2.2. Behavioral Nudging

A common lever to incentivize people to take specific actions is the use of nudges: that is, altering the environment to favor the desired outcome (see, e.g., Thaler and Sunstein 2009, Halpern 2015). Nudges are typically easy and inexpensive to implement. Governments and public agencies have leveraged behavioral nudging to address several policy problems, such as increasing retirement savings, college enrollment, energy conservation, and adult influenza vaccinations (see Benartzi et al. 2017 and the references therein). In the context of transportation, Choudhary et al. (2022) use telematics technology (i.e., sensors in mobile devices) to nudge drivers to drive safely. More generally, several papers use behavioral nudging for mobile app adoption. For example, Ghose et al. (2019) collaborate with a large telecom provider to examine how contextual targeting via a public transit app affects user redemptions of coupons. Sun et al. (2019) use a field experiment to investigate the impact of offering incentives or information on customers' mobile app adoption and purchase behavior. To our knowledge, in the context of stimulating commuters to carpool (and hence, reduce traffic congestion), no large nudging strategy was ever deployed. Waze has the unique ability to nudge commuters to alter their driving habits for the better. Waze's ability to nudge commuters does not incur a significant cost (it is nearly free) and can be carefully designed by selecting the appropriate nudging action for the right set of users. In this paper, we test several nudging strategies on users who can save commute time by carpooling and using an HOV lane. Our results lead to important behavioral insights that may help reduce traffic congestion and at the same time, help commuters make better commuting choices.

2.3. Carpooling

There is a vast literature on carpooling---a comprehensive review is beyond the scope of this paper. One of the first empirical papers on this topic is Teal (1987). The author uses data from the 1977–1978 Nationwide Personal Transportation Survey to study the characteristics of carpoolers and to offer explanations on why commuters carpool. Ferguson (1997) uses data from the same survey between 1970 and 1990 to explain the carpooling decline in America. The decline is attributed to several factors, such as increasing household vehicle availability, falling fuel costs, and higher educational attainments among commuters. More recently, Shaheen et al. (2016) examine the motivations and characteristics of casual carpoolers in the San Francisco Bay area by conducting interviews and surveys (with a total of 519 respondents). As expected, casual carpoolers' motivations include convenience, time savings, and financial savings. The recent work by Ostrovsky and Schwarz (2019) studies the complementary nature of carpooling and self-driving cars, focusing on market equilibrium. The authors claim that congestion pricing will play an essential role in boosting the adoption of self-driving cars and carpooling.

Closer to our paper, there is an extensive literature on carpooling in the context of congestion pricing and HOV lanes (see, e.g., Giuliano et al. 1990, Yang and Huang 1999, Konishi and Mun 2010). These studies investigate the extent to which policies such as HOV lanes are effective in terms of increasing carpooling. Giuliano et al. (1990) compare data from the Route 55 HOV in California with a control group of freeway commuters. The authors show that only the carpooling rate for peak-hours commuters has increased. The authors also conclude that travel time savings must be high to attract new carpoolers, which is one of the motivations of our paper. Yang and Huang (1999) propose a theoretical model for carpooling behavior and optimal congestion pricing in a multilane highway. They show that in the absence of HOV lanes, a uniform toll for all vehicles (independent of the number of occupants) should be charged. With HOV lanes, however, the optimal strategy requires differentiating the toll per vehicle across segregated lanes. Finally, Small et al. (2006) and Konishi and Mun (2010) study the trade-off between HOV and HOT (high-occupancy toll) lanes. Small et al. (2006) empirically analyze the behavior of motorists traveling on California State Route 91 and show the importance of considering customer heterogeneity. Konishi and Mun (2010) develop a model to examine under which conditions introducing HOV lanes is socially beneficial and whether converting HOV lanes to HOT lanes improves road efficiency.

Another related line of research is the efficacy of incentives to increase carpooling. Vanoutrive et al. (2012) analyze the popularity and determinants of carpooling in Belgium in the context of the workplace. The authors observe higher levels of carpooling at less accessible locations and in sectors such as construction, manufacturing, and transport. More recently, Neoh et al. (2017) synthesize 22 existing empirical studies (with over 79,000 observations) to create a review of the carpooling literature. Their analysis identifies 24 carpooling factors, including the number of employees, partner-matching programs, gender, and a fixed work schedule. Although there is an extensive literature on incentives and motivations for carpooling, our study is the first large-scale field experiment (with more than half a million users). The scale of our data allows us to sharpen our current understanding on carpooling and ultimately, inform policy making.

2.4. Structure of Paper

Section 3 discusses the data and setting considered in this paper. Section 4 outlines our experimental design. Section 5 presents our various econometric results. Finally, Section 6 reports our conclusions and discusses the practical implications of our findings.

3. Data and Setting

As mentioned, this paper is in the context of the Waze Carpool service. Having access to more than 110 million users, Waze is seeking to persuade drivers who use its service (called Wazers) to become carpoolers. We call the Waze drivers who have installed the Waze Carpool app *onboarded* drivers. This means that these drivers have signed up for the carpool service by entering their home and work addresses as well as their preferred commute times.

3.1. Population of Users

Our goal is to investigate which factors can entice Wazers to become carpoolers. We first use a large historical data set to examine several correlations between onboarding to the carpool service and drivers' attributes. In the context of this paper, we are interested in drivers who have an HOV lane on their daily commute.

We consider a random sample of U.S. Wazers who live in an area with an HOV lane. Specifically, we focus on four U.S. states that have a high HOV occupancy: California (CA), Georgia (GA), Massachusetts (MA), and Washington (WA). The list of HOVs under consideration can be found in the online appendix. We restrict our attention to representative Wazers who are active commuters. To this end, we focus on a random sample of users who completed at least 10 navigations[4] in the last 30 days and at least two navigations with a similar origin and destination (e.g., home and work locations or the most common navigation) on weekday morning hours (from 6 a.m. to 12 p.m.). This allows us to focus on the times when the HOV lane restriction is relevant. We also restrict our sample to users who have a commute distance between 5 and 100 kilometers and a duration between 5 and 100 minutes to avoid unrepresentative outliers. For each user, we use the data from all morning weekday navigations in the last 30 days to compute three key variables.

1. Average leave time. For each navigation, we record the time at which the user left the origin location (i.e., the time when the user started his or her commute) and compute the average. We then transform the time information to a continuous number. For example, if the user left the origin location on average at 7:45 a.m., the average leave time will be 7.75.

2. Standard deviation of leave time. Similarly, we compute the standard deviation of the leave time (in hours) by using all the morning weekday navigations in the last 30 days.

3. Average time saving. For each navigation, we leverage the Waze data and algorithms to compute the estimated travel duration both with and without the HOV lane (assuming that there is an HOV lane on the route). Because this is the most important variable in the context of this paper, we next explain how is it calculated in more detail.

3.2. Time-Saving Computation

First, Waze needs to identify the geolocation of the HOV lanes. This is done by relying on two approaches. (1) It relies on a large community of map editors who voluntarily keep the data of the maps fresh and up to date. Many of the HOV lanes, their exact locations, and the hourly constraints related to the HOV restriction are all marked by map editors. (2) For some regions, Waze relies on unsupervised learning methods to identify whether there are two distinct speeds on the same segment at the same time (using a statistical test on the histogram of the speed values). In this context, a segment is defined as a small portion of a specific road or highway. Because Waze has data from a large number of users, the speed values can easily be analyzed for each segment and each time window. After Waze identifies a segment with two distinct speed values, it provides evidence about the presence of an HOV lane. Note that Waze estimates speed values using a combination of proprietary unsupervised and supervised methods.

Second, Waze needs to compute the potential time saving for each user on the HOV lanes examined in this paper. Whenever a user navigates with Waze, the app automatically computes for the requested route both the ETA (estimated time of arrival) if the user will use the HOV lane and the ETA if the user will use the regular lane. In Figure 1, we illustrate this point for a user who requested a route from San Jose, CA to Mountain View, CA at 10 a.m. on a specific day.[5] For this request, the Waze app displays three possible routes (two without an HOV lane and one with an HOV lane). In this example, the user can save nine minutes by using the HOV lane (i.e., time saving = 9 minutes). We define the time saving per user as the average difference between the time that the app computed if the user was to reach its destination by using the HOV lane minus the time computed if the user was to reach its destination without using the HOV lane (the difference is calculated based on the exact same origin and destination for the two alternatives at the same time of departure). To calculate the average of these differences for each user, we relied only on the commute times (i.e., navigations that took place between 6 a.m. and 12 p.m. on weekdays). It is worth mentioning that Waze data and algorithms offer a unique opportunity to compute the HOV time-saving metric at scale.

Depending on the value of the average time saving, we split the users into three categories.

1. Users with a positive time saving. These users can potentially save time by using the HOV lane. They currently cannot use it given that they are solo commuters, but if they would carpool by riding with a passenger, they may shorten their commute time.

2. Users with a negative time saving. These users cannot save time by using the HOV lane. For example, the fastest route from home to work does not include an HOV lane, so taking the itinerary with the HOV lane will take longer (e.g., by incurring a detour).

3. Users who do not have access to an HOV lane in their commute. These users live in a market where there is at least one HOV lane in their neighborhood, but the HOV lane is not located on their daily commute (and hence, not displayed as one of the route options).

As mentioned, in Figure 1, we query the Waze app for a navigation from San Jose to Mountain View at 10 a.m. The system suggests three possible routes (see the top left of Figure 1). The optimal route has an ETA

Figure 1. (Color online) Example of Time Saving by Using the HOV Lane (Date Accessed: July 23, 2019)

of 30 minutes. The third route suggests using the HOV lane (on US-101 N) for an ETA of 21 minutes. Thus, the estimated time saving in this example is nine minutes.

We eliminate users with a time saving lower than −40 and higher than 40 minutes to avoid outliers (these users represent a negligible fraction of our sample). Ultimately, we use a random sample with 806,790 users. The average time-saving distribution for these users is shown in Figure 2. We also report the time-saving distribution for each market (i.e., state) and each one-hour leave time window (i.e., the average time at which users start their commute) in the online appendix. As we can see, the time-saving distribution is centered around zero, so that users are split into the ones who can save commute time by using the HOV lane and the ones who cannot. The online appendix shows that as the morning progresses, the variance in time saving decreases.

To further visualize the variation in time saving as a function of the leave time, we plot for each one-hour interval the proportion of users who save on average more than two minutes, between two minutes and minus two minutes, and less than minus two minutes (see Figure 3). During the early morning commute hours (i.e., peak times), more than 35% of users with an HOV lane on their route have a meaningful time saving by taking the HOV lane. This percentage then linearly declines with time until 11 a.m. We note that the percentage of users with a positive time saving is underestimated because of the fact that users are likely to optimize their commute based on traffic

Figure 2. (Color online) Time-Saving Distribution

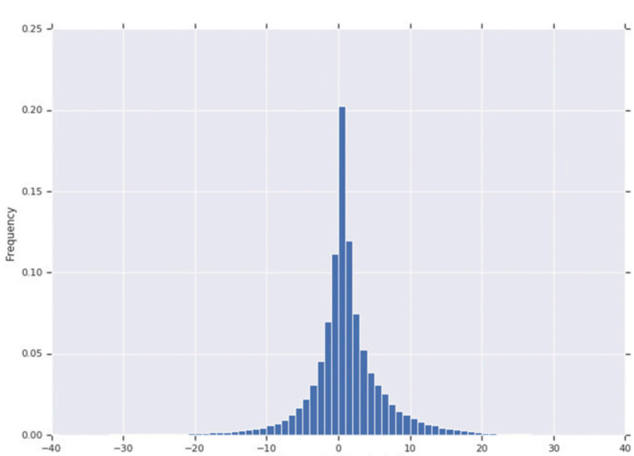

Figure 3. (Color online) Time-Saving Evolution with Time

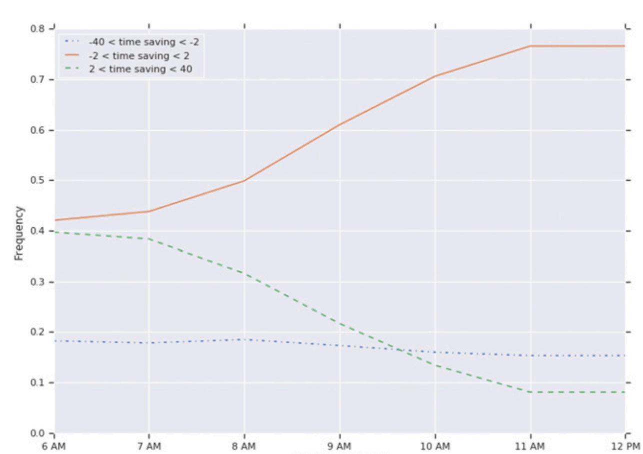

conditions (e.g., they commute later or earlier than they would like to avoid traffic).

3.3. Preexperiment Analysis

We next investigate to what extent the likelihood of onboarding to the carpool service correlates with the time saving. We start by plotting the OBR (that is, the ratio between onboarded drivers and the total number of drivers) as a function of the average time saving (see Figure 4).[6] We find that the OBR increases with the average time saving. This suggests that users with a higher time saving are more likely to be interested in the carpool service. More specifically, users with a time saving of at least 10 minutes are twice as likely to have onboarded relative to users with a negative (or close to zero) time saving. We then compare the average time saving for onboarded and non-onboarded drivers: t statistic = −18.8 with a p-value (much) less than 1%.

Finally, to account for additional covariates that may affect the individual decision maker to carpool, we estimate six regression specifications. We consider two types of models.

1. A linear probability ordinary least square model

$$Y_i = \alpha X_i + \epsilon_i, \tag{1}$$

where i corresponds to the user and Y_i is an indicator variable that captures whether user i has onboarded. The independent variables X_i include the device type (Android versus iOS), the market (CA, GA, MA, WA), the average distance of the commute in kilometers (from origin to destination) in the last 30 days, the average leave time (i.e., the average time at which the user started the morning commute), the number of days since the user joined the Waze app, and the number of times the user used the Waze app during the last 30 days (called number of sessions). Finally, ϵ_i is a stochastic independent and identically distributed

Figure 4. (Color online) Onboarding Rate (Normalized) as a Function of Time Saving

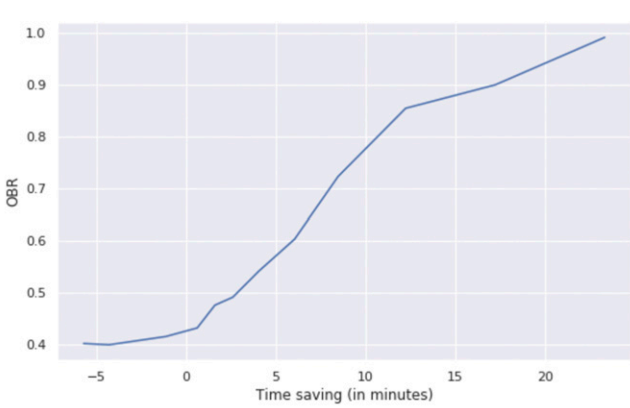

(i.i.d.) Gaussian term, and α is the estimated vector of coefficients associated with the independent variables.

2. A logit probability model

$$\text{logit}(Y_i) \sim f(X_i; \alpha), \tag{2}$$

where the link function $f(\cdot)$ is logit.

The estimated coefficients of Equations (1) and (2) are reported in Table 1.[7] Note that we only consider users with a positive or negative time saving and remove the ones who do not have access to an HOV lane on their commute---so that we remain with 574,559 observations. In the first two columns, we control for the independent variables mentioned. In the third and fourth columns, we also include the standard deviation of the leave time. This variable aims to capture the flexibility in the commute time. In the last two columns, we use the range (i.e., maximum minus minimum) of the leave time instead of the standard deviation (users with a single observation are assigned a standard deviation and a range equal to zero).

The results of Table 1 suggest the following.

- The average distance and number of sessions are positively correlated with OBR. This validates the intuition that commuters who are more active or have a longer commute distance are more likely to be interested in the carpool service.
- Commuters from CA and WA have a significantly higher OBR relative to users from GA and MA. Using the estimates from the first column of Table 1 together with the average values of the control variables, we find that commuters from CA and WA have an OBR that is 30% higher relative to GA users. This supports the fact that carpooling is a more widespread practice on the West Coast, where several HOV lanes can be found, and commuters typically have a longer commute that cannot easily be completed using public transportation.
- More importantly, our estimates confirm that the average time saving is positively correlated with OBR, hence validating our intuition that users who can save more time by using the HOV lane are more interested in the carpool service. More precisely, our estimates suggest that an additional 10 minutes in time saving boosts the OBR by 18% (again, using the estimates from the first column of Table 1 with the average values of the control variables).
- Finally, we find that the standard deviation of the leave time has a negative effect on OBR. This is an interesting finding as one could have posited two opposite hypotheses. (1) A high standard deviation in leave time means more variation in the starting work hour (e.g., unpredictable meeting times), so that such users are less likely to carpool. (2) A high standard deviation in leave time can translate into more flexibility in work hours, so that such users are more likely to

Table 1. Regression Estimates for Preexperiment Analysis (Dependent Variable: Normalized OBR)

	OLS Base	Logistic Base	OLS With STD	Logistic With STD	OLS With range	Logistic With range
log(*avg_distance*)	0.008***	0.27***	0.008***	0.288***	0.008***	0.286***
	(0.0008)	(0.034)	(0.0008)	(0.034)	(0.0008)	(0.034)
log(*avg_leave_time*)	−0.074***	−2.77***	−0.068***	−2.494***	−0.068***	−2.508***
	(0.004)	(0.132)	(0.004)	(0.138)	(0.004)	(0.138)
log(*days_since_joined*)	0.012***	0.504***	0.012***	0.496***	0.012***	0.498***
	(0.0004)	(0.02)	(0.0004)	(0.02)	(0.0004)	(0.02)
log(*sessions_30d*)	0.022***	0.874***	0.022***	0.86***	0.022***	0.862***
	(0.0008)	(0.028)	(0.0008)	(0.028)	(0.0008)	(0.028)
avg_time_saving	0.002***	0.038***	0.002***	0.036***	0.002***	0.036***
	(0.00008)	(0.02)	(0.00008)	(0.02)	(0.00008)	(0.02)
log(*std_leave_time* + 1)			−0.006***	−0.202**		
			(0.002)	(0.084)		
is(*std_leave_time_NA*)			−0.008***	−0.282***		
			(0.002)	(0.046)		
log(*range_leave_time* + 1)					−0.004**	−0.11**
					(0.002)	(0.05)
is(*range_leave_time*_NA)					−0.008***	−0.286***
					(0.002)	(0.05)
GA	−0.03***	−1.328***	−0.03***	−1.33***	−0.03***	−1.33***
	(0.002)	(0.056)	(0.002)	(0.056)	(0.002)	(0.056)
MA	−0.034***	−1.498***	−0.034***	−1.472***	−0.034***	−1.474***
	(0.002)	(0.098)	(0.002)	(0.098)	(0.002)	(0.098)
WA	0.002	0.036	0.002	0.034	0.002	0.034
	(0.002)	(0.056)	(0.002)	(0.056)	(0.002)	(0.056)
Constant	0.034***	−8.448***	0.024***	−8.808***	0.026***	−8.774***
	(0.01)	(0.354)	(0.01)	(0.36)	(0.01)	(0.358)
Observations	574,559	574,559	574,559	574,559	574,559	574,559
R^2	0.007		0.007		0.007	
Log likelihood		−70,961.090		−70,940.280		−70,940.800
Residual standard error	0.164 (df = 574,549)		0.164 (df = 574,547)		0.164 (df = 574,547)	
F statistic	445.013*** (df = 9; 574,549)		367.856*** (df = 11; 574,547)		367.655*** (df = 11; 574,547)	

Notes. Ordinary least squares (OLS) corresponds to Equation (1), whereas logistic corresponds to Equation (2). The first two columns are the base models. In the third and fourth columns, we add the standard deviation of the leave time. In the fifth and sixth columns, we replace the standard deviation with the range. STD, standard deviation (of leave time); df, degrees of freedom.

$^{**}p < 0.05$; $^{***}p < 0.01$.

carpool. The results in Table 1 are in favor of the first hypothesis.

All our statistical tests strongly support that a high time saving is positively correlated with the carpool intent (captured by the OBR). This observation motivates us to further study the relationship between carpool intent and time saving and to examine whether it is causal. The Waze platform offers a great opportunity to run randomized controlled experiments, as we discuss next.

3.4. Experiment Data

Following the process outlined in Section 3.3, we consider users across four U.S. states: CA, GA, MA, and WA. We further focus on users who either have a positive time saving or no access to an HOV lane (i.e., we remove the users with a negative time saving). We also remove all users who already onboarded to the Waze Carpool service. Finally, we carefully ensure that the users in our field experiment are not part of another experiment run concurrently. We remain with a sample of 537,370 users. We call this set of users our experimental population.

4. Experimental Design

As discussed, our goal is to encourage Waze users via an app notification to try the carpool service: that is, to sign up with the platform and ultimately complete a carpool ride. Our experiment allows us to investigate how users respond to various types of incentives (i.e., different framings).

4.1. Experimental Population

Each user in the experimental population belongs to one of three categories: HOV users with a high time saving (called H users), HOV users with a low time saving (called L users), and non-HOV users (called N users). We consider the random sample of users mentioned in Section 3.4. N users are simply the ones who do not have access to an HOV lane on their commute.

Thus, all the remaining users have an HOV lane on their commute and an average positive time saving. We split these remaining users depending on the value of their average time saving: H users (L users) correspond to the top 27th percentile (bottom 73%).[8] The numbers of each user type are reported in Table 2 (the number of impressions will be discussed in Section 4.2).

We highlight that all the interventions used in our field experiment were beneficial to the users because they received an offer to try the carpool service while highlighting different benefits of carpooling. In addition, our analyses are aggregated over hundreds of thousands of users and are meant to be interpreted as statistical averages.

4.2. Treatments

Our field experiment involves several treatments or interventions. Each treatment entails sending an app notification to the user---called an encouragement---to advertise the carpool service. The encouragement is "sent" programmatically to all users at the same time: in our case, on June 10, 2019. Users will see the encouragement displayed on the homepage of the Waze app the first two times they open the app during the period that our experiment is running (June 10 to July 3). As a result, it is possible that users will see the encouragement at different times (we will later explicitly control for such time effects). We then split the users into several conditions as follows.

- Commuters with a high time saving (H users) are split into four conditions (HA, HB, HC, HD). Specifically, we use the following four framings for the encouragement: (A) mentioning the HOV lane and the potential (high) time saving, (B) mentioning the HOV lane, (C) using a generic carpool invitation, and (D) not sending anything. The encouragements along with the text used in the three treated conditions are shown in Figure 5.
- Commuters with a low time saving (L users) are split into the same four conditions (LA, LB, LC, LD); the exact messages are shown in the online appendix.
- Commuters without access to an HOV lane (N users) are split into three conditions (NA, NB, NC): (A) mentioning the monetary incentive (receiving a $10 welcome bonus to try the carpool service), (B) using a generic carpool invitation, and (C) not sending anything (the screenshots of the encouragements are shown in the online appendix).[9]

Our experiment was live between June 10 and July 3, 2019: that is, a total of 24 days. During this period, when any of the 537,370 users open the Waze app for the first two times, they will see the encouragement inviting them to try the carpool service. If a user was shown the encouragement, we call it an *impression*. Then, the user can either click on the "Try Waze Carpool" button or on the exit button (see Figure 5). The user can also ignore the encouragement. Several reasons may lead to a driver not seeing the encouragement, including cases when the driver did not open the app during the experiment period, the driver reinstalled (or updated) the app, and the driver changed mobile device. Note that by definition, users in the control condition are not sent an encouragement and thus, will not see an impression. Overall, around 78% of our users saw the impression. Fortunately, this number is constant across all conditions and user types, as seen in Figure 6.

Table 2. Number of Users in Our Field Experiment

User type	Number of users	Number of impressions
H users	84,798	49,504
L users	224,518	131,778
N users	228,054	120,021
Total	537,370	301,303

4.3. Balancing Groups

To ensure that our experiment is properly randomized (i.e., there is no selection bias between the conditions for each user type), we randomly split the users from each category (H, L, N) into the different conditions. Instead of simply splitting users randomly, we performed 1,000 random splits for each population of users (H, L, and N separately) and selected the most balanced split: that is, the split that yields the highest p-values when we compare all the balancing variables listed in Table 3. Note that for such a large number of users, all the 1,000 random splits are nearly as good, so that this procedure only serves as an extra safety net. Ultimately, our goal is to have a balanced sample with respect to the relevant variables for each category of users. In Table 3, a session is defined as any interaction between the user and the Waze app, such as ETA check, navigation, and driving when the app is open. A navigation with an HOV lane captures a navigation in which at least one of the proposed itineraries has an HOV lane.

We also trust the randomization to balance the different conditions for each type of users across other relevant variables. As a sanity check, we examine the balance with respect to the variables in Table 3 as well as the following attributes: market (i.e., state), average leave time, device (Android versus iOS), and percentage of users who saw an impression. In Figure 7, we visualize the balancing by reporting the average number of days since joining date, number of navigations in the last 30 days, distance driven in the last 30 days, and (estimated) time saving across the different conditions for each type of users. The plots for the other variables can be found in the online appendix.

Figure 5. (Color online) Copy of the Messages Used in Our Field Experiment for H Users

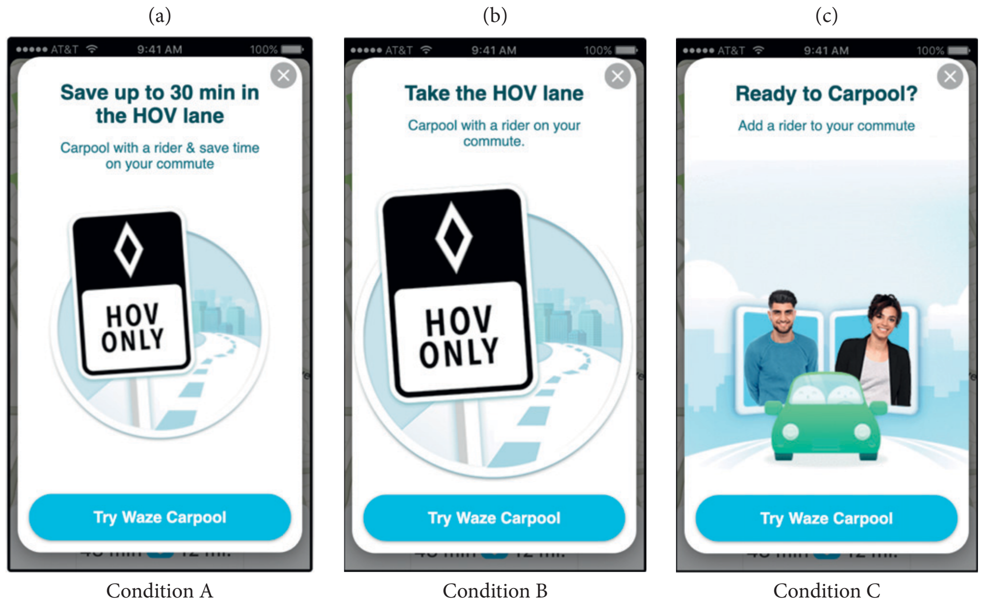

Notes. (a) Condition A. (b) Condition B. (c) Condition C.

As we can see from all the figures, for each user type, all the relevant variables are well balanced across conditions (i.e., there are no statistically significant differences between conditions for each user type). Note that the different conditions are perfectly balanced for each user type but not necessarily balanced across different user types. This is to be expected, given that different types of users admit inherent differences (e.g., HOV users may systematically have a longer commute time relative to non-HOV users, as confirmed by Figure 7(c)). Our regression analyses in later sections will control for the variation of these variables, beyond the random assignments to treatment groups within each category of users.

Recall that N users come from similar geographical locations as H and L users. As explained in Section 3, we focus on users who live in an area with an HOV lane across four U.S. states. However, because N users are fundamentally different from H and L users, it is natural to observe certain differences in several variables. Finally, the average time saving for H users is significantly higher relative to L users (see Figure 7(d)), as expected. The averages are 9.6 and 2.1 minutes for H and L users, respectively (the ranges of values are [6, 40] and [2, 5], respectively).

4.4. Performance Metrics

To measure the treatment effect, we consider two performance metrics related to carpool intent (or affinity): the CTR and the OBR. The CTR is defined as the number of users who clicked on the "Try Waze Carpool" button from the encouragement divided by the number of users who were shown an impression. This is a common metric often used in the context of online advertising (see, e.g., Richardson et al. 2007). A higher CTR typically indicates a higher intent or interest for the service. As discussed before, the OBR is the number of users who signed up (or onboarded) with the Waze Carpool service (by creating an account and completing their home and work addresses) and hence, captures the conversion rate. We consider three versions of the OBR, depending on the normalization. (1) OBR1 is defined as the number of onboarded users divided by the number of users who were shown an impression, (2) OBR2 is defined as the number of onboarded users divided by the number of users who clicked, and (3) absolute OBR is defined as the absolute number of onboarded users. All three OBR metrics are normalized and expressed in percentages relative to their maximal value. The latter metric allows us to include users from the control condition. Indeed, users in the control condition did not receive an encouragement by design, and hence, the concepts of impression and click are not relevant for such users. Our four dependent variables are summarized in Table 4. In Section 5.4, we will consider additional dependent variables that capture carpool adoption (as opposed to carpool intent).

We highlight that our large-scale field experiment should be seen as three field experiments (one for each type of users), which are run in parallel at the same time and in the same environment. Consequently, our goal is not to directly compare the results across different types of users but instead, to examine

Figure 6. (Color online) Average Number of Impressions Across the Different Conditions for Each User Type

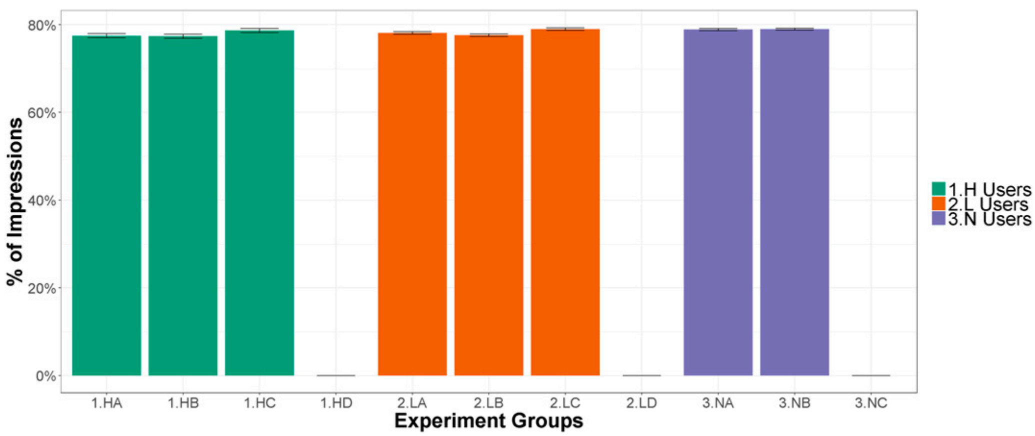

the results for each type of user in order to provide a comprehensive study.

5. Results

In this section, we report the results of our field experiment for each condition, user type, and the performance metrics described in Table 4. First, we show the results of one-way ANOVA tests (see, e.g., Maxwell and Delaney 2004). Second, we estimate several regression specifications to showcase the robustness of our results when controlling for various factors. Third, we refine our findings by investigating potential heterogenous treatment effects. Finally, we conclude our study by considering metrics that capture carpool adoption.

5.1. Basic Results

We start by presenting one-way ANOVA tests on each performance metric by pooling the observations across all three types of users. The results for CTR are reported in Figure 8 ($F(7,301,300) = 200.3$, $p < 0.01$). For each figure, we include the 95% confidence interval corresponding to the average value. We uncover the following findings.

• Highlighting the fact that commuters can use the HOV lane significantly boosts the CTR relative to a generic carpool invitation. Specifically, conditions A and B for H users increase the CTR by 174% and 185%, respectively, relative to a generic message. For L users, these numbers become 139% and 133%, respectively. These effects remain statistically significant when performing a separate ANOVA test for each type of users (H and L).

• Mentioning the HOV lane is enough in the sense that highlighting the potential time saving does not yield an additional marginal effect. Specifically, conditions A and B are not statistically different from each other. One possible implication is that mentioning the HOV lane is enough, so that users can directly translate this information into saving commute time (note that it is still possible that the framing of the message in condition B was not well executed). This result holds both for H and L users (using either the pooled or separate sample).

• Users with a high time saving (i.e., H users) are more receptive to the carpool offer and hence, have a higher CTR than L users. Interestingly, this result holds for framings A and B (in which the HOV lane is highlighted) but not for framing C (generic message).

• Highlighting the $10 incentive for N users (condition NA) does not increase the CTR relative to sending a generic message (i.e., N users in A and B conditions are not statistically different in terms of CTR). This finding suggests that highlighting the $10 welcome bonus is not enough to nudge commuters to carpool in our experiment. However, this finding bears some limitations. First, most of the compensation comes from completing rides rather than from the welcome bonus. Second, our treatment only considered mentioning a few words about $10 in a small font (as opposed to a full-screen HOV picture).

• Interestingly, users who can save time by using the HOV lane but for whom the HOV was not mentioned in the notification (i.e., HC and LC users) are not different from N users (i.e., non-HOV users). However, when the HOV is explicitly mentioned in the notification, the

Table 3. Balancing Variables in Our Field Experiment

Variable
Number of sessions in the last 30 days
Number of navigations with an HOV lane in the last 30 days
Days since joined Waze
Distance driven in the last 30 days (in kilometers)
Total navigation time in the last 30 days (in minutes)

Figure 7. (Color online) Comparing Four Key Variables Among the Different Conditions for Each User Type

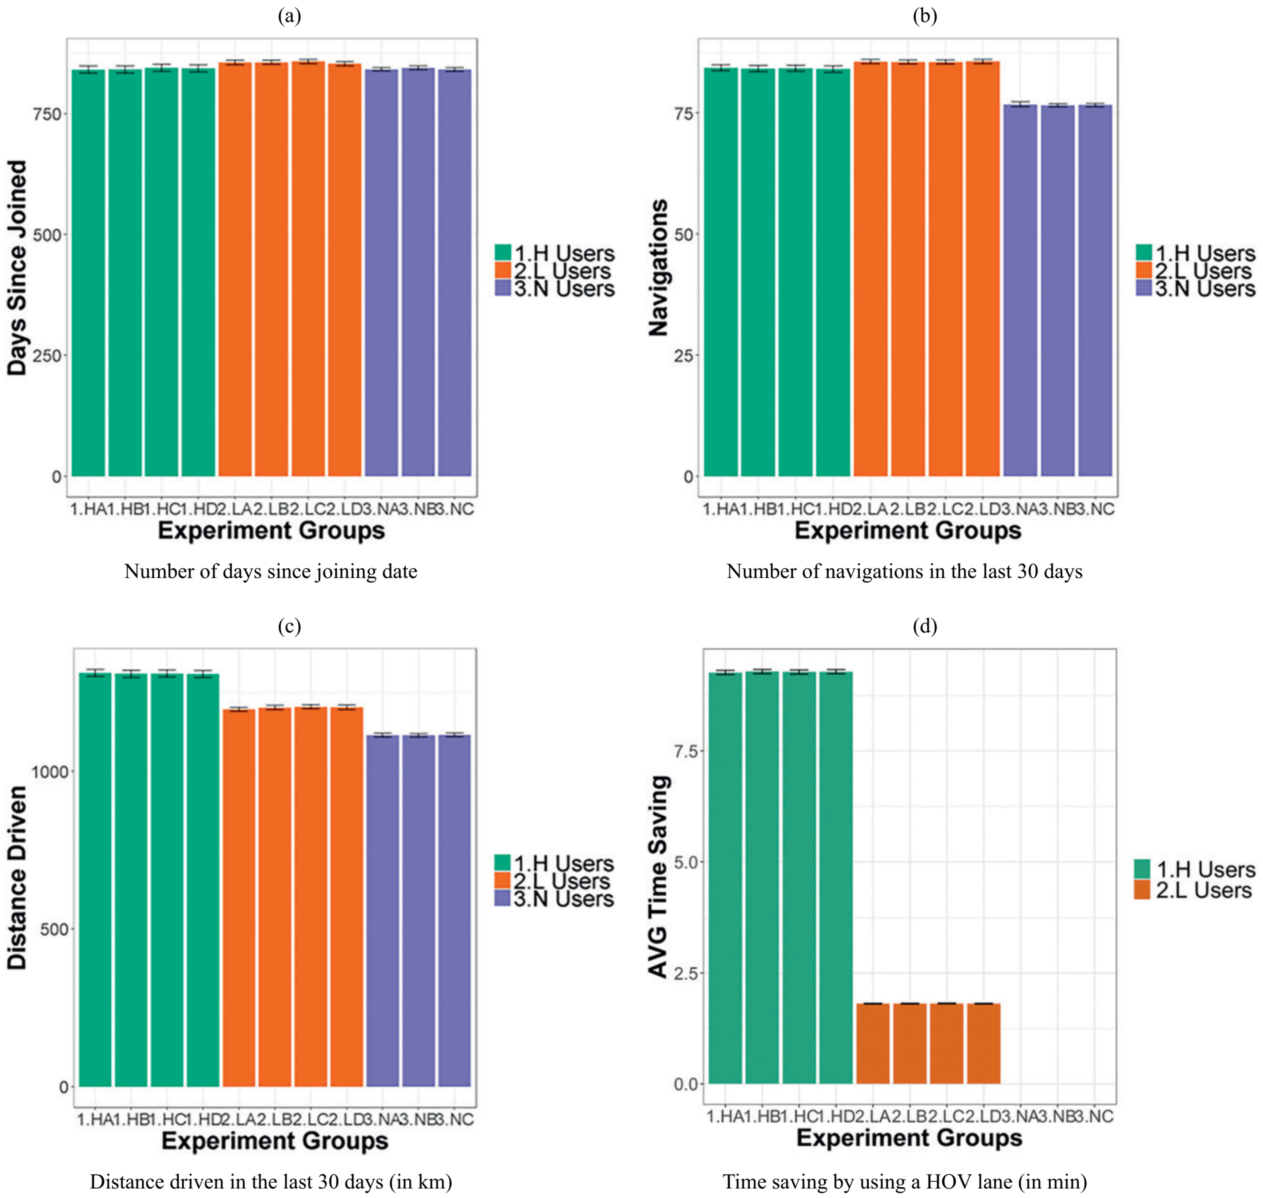

Notes. (a) Number of days since joining date. (b) Number of navigations in the last 30 days. (c) Distance driven in the last 30 days (in kilometers). (d) Time saving by using an HOV lane (in minutes).

CTR increases significantly. This reinforces the importance of explicitly highlighting the benefit of carpooling to targeted high-intent users. As a result, both the type of user and the framing of the intervention play important roles in incentivizing commuters to carpool.

We next consider the OBR metrics. As discussed in Section 4.4, we consider three variants of the OBR,

Table 4. Our Four Dependent Variables

Dependent variable	Description
CTR	Users who clicked divided by users who were shown an impression
OBR1	Users who onboarded divided by users who were shown an impression
OBR2	Users who onboarded divided by users who clicked
Absolute OBR	Users who onboarded

depending on the normalization. In Figure 9(a), we use OBR1: that is, the onboarded rate normalized by the number of impressions. One can see that the results for OBR1 follow the same pattern as the results for CTR and are statistically significant ($F(7, 301, 300) = 15.08$, $p < 0.01$). In Figure 9(b), we consider OBR2 where the onboarded rate is normalized by the number of clicks. In this case, although all the results follow the same qualitative behavior as the two other metrics, a few pairwise comparisons are losing their statistical significance (the pooled sample F statistic is still significant at the 1% level: $F(7, 6, 519) = 4.228$, $p < 0.01$). The fact that our results also hold for OBR2 suggests that even after controlling for users who clicked, the carpool intent of users in conditions A and B remains higher. This allows us to mitigate the potential concern that users may click on the banner from a lack of understanding or to make it disappear. Finally, the same qualitative results also hold for the absolute OBR and are statistically significant ($F(10, 537, 365) = 15.42$, $p < 0.01$). The details are omitted for conciseness.

The results for the OBR metrics confirm the insights that we draw for CTR. Specifically, highlighting the fact that commuters can use the HOV lane significantly boosts the OBR relative to a generic carpool invitation. For example, conditions A and B for H users increase OBR1 by 141% (in both cases) relative to a generic message. For L users, the effect amounts to 64% and 93% for conditions A and B, respectively.

We next perform a difference-in-differences analysis between groups to account for any potential inherent differences between users who can benefit more from the HOV lane (H users) and others (L and N users). We consider six models to compare the following differences: HA − HC versus LA − LC, HB − HC versus LB − LC, HA − HC versus NB − NA, HB − HC versus NB − NA, LA − LC versus NB − NA, and LB − LC versus NB − NA. This analysis retrieves the same numbers as in the ANOVA tests while providing a few refinements. The results for CTR are reported in Table 5 (we obtain similar qualitative results for OBR1 and absolute OBR, whereas for OBR2, we lose statistical significance for some of the coefficients). Comparing the differences between groups allows us to draw the following insights.

- The H and L coefficients are not statistically significant. This means that the boost in CTR cannot be attributed to the fact that users have a high (or low) time saving from using the HOV lane. It thus suggests that there are no differences in CTR between H and L users.
- The A and B coefficients are positive and statistically significant. This confirms that the treatments used in A and B conditions are more effective than using a generic carpool invitation. For example, users in condition A (see the first column of Table 5) have a $[1 + (0.019 - 0.013)/0.013] = 146\%$ higher CTR relative to users in condition C.
- The interaction terms H × A and H × B (in the first two columns) are positive and statistically significant. This confirms that the boost in CTR is driven by the intervention (i.e., receiving a notification that highlights the HOV lane) and not by the difference in time saving between H and L users. Specifically, H users in condition A have an additional lift in CTR of $(0.005/0.019) = 26\%$ relative to L users in condition A.
- All interactions terms in the last four columns are positive and statistically significant with an economic impact that ranges between 128% and 178%. This suggests that highlighting the HOV lane for H and L users (relative to a generic message) is more impactful than explicitly highlighting the monetary incentive for N users (also relative to a generic message).
- The results for OBR1 are in the same spirit (the online appendix).

Figure 8. (Color online) ANOVA Test for CTR

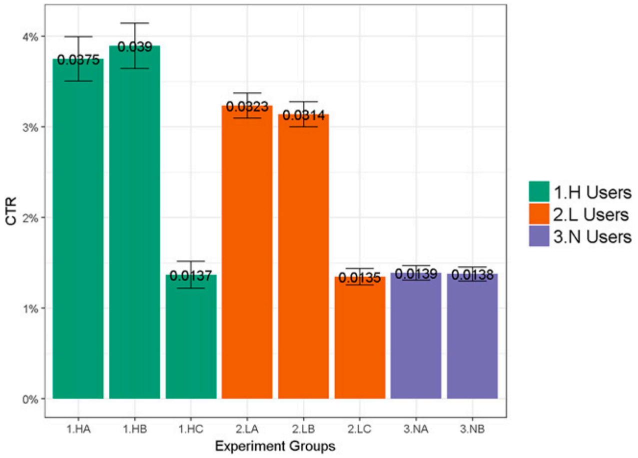

5.2. Regression Analysis

We next test the robustness of the ANOVA test results by estimating several regression specifications. As in Section 3.3, we consider two models for each of the four dependent variables.

1. A linear probability ordinary least square model

$$Y_i^k = \beta^k TC_i^k + \gamma^k X_i^k + \epsilon_i^k, \tag{3}$$

where k corresponds to the user type (H, L, or N), i represents the user, and Y_i^k is one of our dependent variables for user i from group k. The independent variables are divided in two categories. (1) TC_i^k stands for treatment conditions and includes a categorical variable for each condition (A, B, C, D), and (2) X_i^k represents the control variables. Specifically, we control for the device type (Android versus iOS), the market (CA, GA, MA, WA), the average distance in the last 30

Figure 9. (Color online) ANOVA Test for OBR1 and OBR2 (Normalized)

(a) (b)

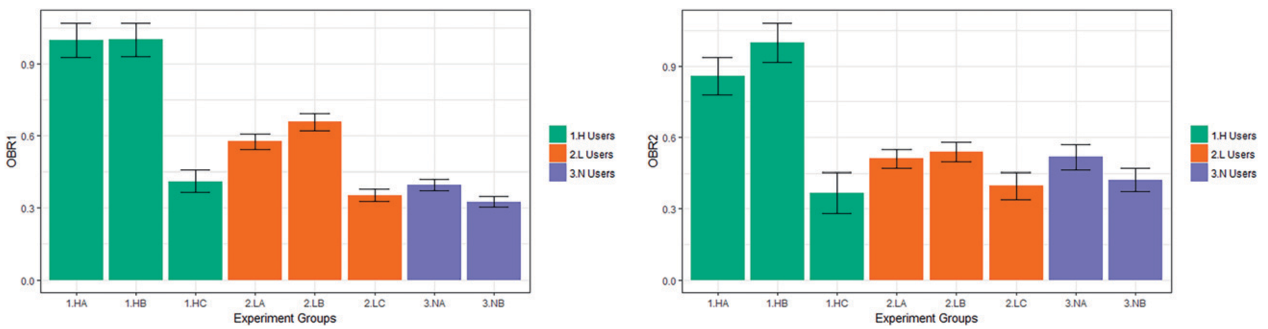

OBR1 OBR2

Notes. (a) OBR1. (b) OBR2.

days, the average leave time, the number of days since joined, and the total number of sessions in the last 30 days. Finally, ϵ_i^k is a stochastic i.i.d Gaussian term, β^k is the estimated vector of treatment effects for group k, and γ^k is the estimated vector of coefficients associated with the control variables.

2. A logit probability model

$$\text{logit}(Y_i^k) \sim f(TC_i^k, X_i^k; \beta^k, \gamma^k), \tag{4}$$

where the link function $f(\cdot)$ is logit.

To showcase the robustness of our results, we estimate the specifications in Equations (3) and (4) under various configurations. We first estimate a separate regression for each user type (H, L, N). We then consider the pooled sample while controlling for the user type (we also consider pooling H and L users). We estimate each model specification without controls, with partial controls, and with all the controls. Finally, we also estimate a model with impression time fixed effects (i.e., the day when the user saw the impression) to account for the context in which the notification is observed by the user, such as weekend versus weekday.

For conciseness, we only report the regression results for CTR and OBR1 using separate samples. The

Table 5. Regression Estimates for CTR when Comparing the Differences

	OLS HA − HC vs. LA − LC	OLS HB − HC vs. LB − LC	OLS HA − HC vs. NB − NA	OLS HB − HC vs. NB − NA	OLS LA − LC vs. NB − NA	OLS LB − LC vs. NB − NA
H	0.0002	0.0002	−0.0002	−0.0002		
	(0.001)	(0.001)	(0.001)	(0.001)		
L					−0.0004	−0.0004
					(0.001)	(0.001)
A	0.019***		−0.0001		−0.0001	
	(0.001)		(0.001)		(0.001)	
$H \times A$	0.005**		0.024***			
	(0.002)		(0.002)			
B		0.018***		−0.0001		−0.0001
		(0.001)		(0.001)		(0.001)
$H \times B$		0.007***		0.025***		
		(0.002)		(0.002)		
$L \times A$					0.019***	
					(0.001)	
$L \times B$						0.018***
						(0.001)
Constant	0.013***	0.013***	0.014***	0.014***	0.014***	0.014***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Observations	121,297	121,011	153,122	153,103	208,217	207,950
R^2	0.005	0.005	0.003	0.004	0.003	0.003
Residual standard error	0.151 (df = 121,293)	0.151 (df = 121,007)	0.127 (df = 153118)	0.127 (df = 153,099)	0.131 (df = 208,213)	0.131 (df = 207,946)
F statistic	185.191*** (df = 3; 121,293)	186.392*** (df = 3; 121,007)	171.222*** (df = 3; 153,118)	190.873*** (df = 3; 153,099)	231.254*** (df = 3; 208,213)	210.239*** (df = 3; 207,946)

Note. df, degrees of freedom.
$^{**}p < 0.05$; $^{***}p < 0.01$.

results for H, L, and N users using the CTR metric are reported in Figures 10–12, whereas the results for OBR1 are relegated to the online appendix. We estimated several additional specifications (for CTR using all users and using H and L users combined) and found consistent results. Each regression table includes four columns where we report the estimates for OLS and logistic models with and without controls. In each case, the baseline is set to users who received a generic message (i.e., HC, LC, NB). We also use California users as our baseline. Here are our main findings.

- Highlighting the monetary welcome bonus to N users does not have a statistically significant impact relative to a generic carpool invitation. This result holds across all four performance metrics. As mentioned, this result bears limitations (e.g., mentioning the $10 bonus in small text) but suggests that highlighting the welcome monetary bonus is not a critical driver to nudge commuters to carpool (at least in the way it was framed in our experiment).
- As before, mentioning the HOV lane in the encouragement has a significant effect in boosting the CTR and OBR1 metrics (it also consistently holds for OBR2 and absolute OBR). It then suggests that the combination of identifying users who can greatly benefit from carpooling and explicitly mentioning the benefit (in our case, the option of using the HOV lane) is successful at boosting carpool intent.
- HC and LC users are not significantly different from NB users. The estimation results for this analysis (using the CTR metric) can be found in Figure 13 and provide evidence that the three types of users (H, L, and N) are somewhat comparable in terms of reactivity to carpool messages or incentives (hence, confirming the results from Table 5). Ultimately, this means that the nudge is effective, only when explicitly highlighting the benefit of carpooling. In other words, sending a generic message to high-intent users is not enough.
- The magnitude of the effect for A and B conditions is very similar (for both H and L users). If we switch the baseline to one of these two conditions, we retrieve the finding that conditions A and B are not statistically different from each other.
- Users in California are generally more responsive to the carpool encouragement (in terms of CTR) relative to the three other states. Similarly, users with a higher average driven distance in the last 30 days are more receptive to the encouragement for both H and L users.

The findings are robust across both model specifications and to the inclusion of impression time fixed effects. We also obtain the same qualitative insights when using different ways of pooling the data (e.g., pooling all users and combining H and L users).

We also find the same qualitative results for OBR1, OBR2, and absolute OBR when using separate pools of users (the regression tables are omitted because of space limitations). We note that the results for absolute OBR can provide a different perspective, as now, the baseline of the regression can be set either to users who received a generic message or to users in the control condition. In addition of retrieving all our previous insights (along with statistical significance), we also find that relative to NC users (i.e., non-HOV users who did not receive any treatment), HA, HB, LA, and LB users have a higher CTR. This confirms that not mentioning the HOV lane to high-intent users and highlighting the welcome bonus to low-intent (non-HOV) users are not effective interventions.

In conclusion, the insights discussed in Section 5.1 continue to hold even after controlling for various factors related to the user and to external attributes. The fact that our results are robust across multiple dependent variables and under a multitude of model specifications strengthens the validity of our findings.

We next perform one last robustness test. Instead of controlling for the market, we control for the zip code (home or work). Explicitly controlling for the zip code can help control for some time-invariant unobserved heterogeneity among users. For example, home zip codes can somewhat capture the heterogeneity in income levels, and work zip codes can account for the type of profession. Overall, we have 3,510 unique home zip codes and 3,029 unique work zip codes. The results are presented in Table 6 using the pooled sample (the baseline group is set to NB). The first four columns are the same models as before (i.e., controlling for market fixed effects) while also controlling for the standard deviation of the leave time. The last three columns show the estimates for a fixed effects linear model (FELM) when we control for the zip code. In the fifth column, we control for the home zip code. In the sixth column, we control for the work zip code, and in the seventh column, we control for both zip codes. As we can see from Table 6, all our results still hold when controlling for zip codes. Similarly, when using the sample with only H and L users, we obtain consistent results. Finally, most of the results are consistent for OBR1, OBR2, and absolute OBR (omitted for conciseness).

5.3. Heterogenous Treatment Effects

In this section, we estimate the models from Equations (3) and (4) for the CTR metric while adding interaction terms between the treatment conditions (A, B, C) and one of the following variables: (a) distance, (b) average leave time, (c) average time saving, (d) number of navigations with an HOV lane in the last 30 days, (e) days since joined, (f) number of sessions in the last 30 days (i.e., frequency of usage), and (g) standard deviation of leave time (to capture the schedule flexibility). We consider both separate samples (i.e., H and L users

Figure 10. Regression Estimates Using CTR for H Users

	Dependent variable:			
	formula1 *OLS* (1)	formula2 *OLS* (2)	formula1 *logistic* (3)	formula2 *logistic* (4)
factor(experiment_group)1.HA	0.024*** (0.002)	0.024*** (0.002)	1.034*** (0.078)	1.035*** (0.078)
factor(experiment_group)1.HB	0.025*** (0.002)	0.025*** (0.002)	1.074*** (0.078)	1.075*** (0.078)
log(avg_distance_km)		0.005*** (0.002)		0.164*** (0.063)
log(avg_leave_time)		0.003 (0.007)		0.113 (0.244)
log(days_since_waze_ob)		-0.001 (0.001)		-0.047* (0.029)
log(sessions_30d)		0.001 (0.001)		0.029 (0.047)
factor(market)GA		-0.010*** (0.003)		-0.390*** (0.109)
factor(market)MA		-0.011*** (0.003)		-0.416*** (0.107)
factor(market)WA		-0.013*** (0.003)		-0.533*** (0.118)
Constant	0.014*** (0.001)	-0.002 (0.018)	-4.279*** (0.067)	-4.837*** (0.622)
Observations	49,504	49,504	49,504	49,504
R^2	0.005	0.006		
Adjusted R^2	0.005	0.006		
Log Likelihood			-6,531.937	-6,502.542
Akaike Inf. Crit.			13,069.880	13,025.080
Residual Std. Error	0.170 (df = 49501)	0.170 (df = 49494)		
F Statistic	115.684*** (df = 2; 49501)	31.883*** (df = 9; 49494)		
Note:				*p<0.1; **p<0.05; ***p<0.01

separately) and the pooled sample (i.e., H and L users together). We next summarize our findings.

First, we observe that the treatment effect is amplified for H users with a higher average time saving (the estimated regression coefficients for the pooled sample can be found in the online appendix). Interestingly, this finding holds for H users but not for L users. Indeed, all L users have a similar (low) time saving (i.e., no substantial variation across users), whereas the range of time saving for H users is wider. Specifically, we find that an additional 20% in time saving increases the CTR by 12.15% and 17.5% for HA and HB users, respectively, relative to LC users.

Second, we find that the treatment effect is also amplified by the driven distance in last 30 days (the regression table is omitted because of space limitations). Specifically, we find that an additional 400 kilometers in the last 30 days (which can be seen as 10 extra kilometers in the one-way commute distance) for users in conditions (HA, HB, LA, LB) increases the CTR by 26.8%, 14.5%, 14.5%, and 21.7%, respectively, relative to LC users. If we use separate data samples, we find that the effect is significant only for HA users.

Third, we did not find that the magnitude of the treatment effect is moderated by any of the other variables we considered (average leave time, number of days since joined, standard deviation of the leave time, and number of sessions). Regarding the standard deviation of the leave time (which captures the flexibility of users in their commute time), the fact that the effect is nonsignificant is interesting. Two plausible explanations come to mind. (i) The treatment effect is not moderated by the flexibility in the leave time, but the OBR variable in general is correlated with this variable (see Table 1 in Section 3.3). (ii) The statistical power of our experiment is not strong enough to identify this heterogenous treatment effect (i.e., the number of onboarded drivers in each condition is not large enough to convey statistical significance).

Figure 11. Regression Estimates Using CTR for L Users

	Dependent variable:			
	formula1	formula2	formula1	formula2
	OLS	*OLS*	*logistic*	*logistic*
	(1)	(2)	(3)	(4)
factor(experiment_group)2.LA	0.019^{***}	0.019^{***}	0.897^{***}	0.896^{***}
	(0.001)	(0.001)	(0.049)	(0.049)
factor(experiment_group)2.LB	0.018^{***}	0.018^{***}	0.866^{***}	0.866^{***}
	(0.001)	(0.001)	(0.050)	(0.050)
log(avg_distance_km)		0.006^{***}		0.224^{***}
		(0.001)		(0.038)
log(avg_leave_time)		0.015^{***}		0.596^{***}
		(0.003)		(0.140)
log(days_since_waze_ob)		-0.002^{***}		-0.091^{***}
		(0.0005)		(0.019)
log(sessions_30d)		0.003^{***}		0.103^{***}
		(0.001)		(0.030)
factor(market)GA		-0.007^{***}		-0.288^{***}
		(0.001)		(0.049)
factor(market)MA		0.001		0.048
		(0.003)		(0.122)
factor(market)WA		-0.006^{***}		-0.247^{***}
		(0.002)		(0.066)
Constant	0.013^{***}	-0.032^{***}	-4.295^{***}	-6.146^{***}
	(0.001)	(0.009)	(0.041)	(0.376)
Observations	131,778	131,778	131,778	131,778
R^2	0.003	0.004		
Adjusted R^2	0.003	0.004		
Log Likelihood			-15,513.490	-15,447.060
Akaike Inf. Crit.			31,032.980	30,914.120
Residual Std. Error	0.158 (df = 131775)	0.158 (df = 131768)		
F Statistic	200.712^{***} (df = 2; 131775)	59.321^{***} (df = 9; 131768)		
Note:				$^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$

Overall, the findings are not particularly surprising and confirm our intuition that users who can benefit the most from using an HOV lane or have a longer commute distance are responding better to the HOV-explicit encouragements tested in our experiment.

5.4. Actual Carpool Adoption

So far, we measured the impact of our field experiment in terms of carpool intent (i.e., clicking on the notification and onboarding to the carpool platform). A more impactful metric would be to consider carpool adoption. Before doing so, we next describe the carpool conversion process.

5.4.1. Carpool Conversion Process. After a user downloads the Waze Carpool app, she or he first needs to complete the onboarding process (i.e., entering home and work locations, preferred commute times, and several other personal details). This step of the conversion process is captured by our OBR variables. After a driver onboards to the Waze Carpool platform, the next relevant action is to send an offer to a rider. More precisely, each driver can see a list of potential "matched" riders (i.e., users with similar commuting patterns who are interested in finding a ride). Each of these riders incurs a different detour and may display a different price. Then, the onboarded driver may decide to send an offer to specific matched riders. An offer entails to propose a ride for a specific commute (time and origin-destination pair) at a specific price (computed by the platform). Sending an offer to a rider is a clear signal of carpool intent and is captured by the total sent offer (TSO) variables defined here. Finally, the last step in the funnel is to complete a carpool ride (by actually taking a rider for the commute). This step is captured by our total completed ride (TCR) variables (also defined here). In some way, the action of onboarding can be seen as the beginning of the carpool funnel, sending

Figure 12. Regression Estimates Using CTR for N Users

	Dependent variable:			
	formula1	formula2	formula1	formula2
	OLS	*OLS*	*logistic*	*logistic*
	(1)	(2)	(3)	(4)
factor(experiment_group)3.NA	0.0001	0.0001	0.009	0.009
	(0.001)	(0.001)	(0.049)	(0.049)
log(avg_distance_km)		-0.0001		-0.007
		(0.001)		(0.044)
log(avg_leave_time)		0.005^{*}		0.373^{*}
		(0.003)		(0.215)
log(days_since_waze_ob)		-0.002^{***}		-0.110^{***}
		(0.0004)		(0.027)
log(sessions_30d)		0.0004		0.031
		(0.001)		(0.043)
factor(market)GA		-0.001		-0.054
		(0.001)		(0.071)
factor(market)MA		0.0003		0.024
		(0.001)		(0.065)
factor(market)WA		-0.003		-0.224
		(0.003)		(0.249)
Constant	0.014^{***}	0.012	-4.271^{***}	-4.458^{***}
	(0.0005)	(0.008)	(0.035)	(0.554)
Observations	120,021	120,021	120,021	120,021
R^2	0.00000	0.0002		
Adjusted R^2	-0.00001	0.0001		
Log Likelihood			-8,754.661	-8,744.225
Akaike Inf. Crit.			17,513.320	17,506.450
Residual Std. Error	0.117 (df = 120019)	0.117 (df = 120012)		
F Statistic	0.030 (df = 1; 120019)	2.659^{***} (df = 8; 120012)		
Note:				*p<0.1; **p<0.05; ***p<0.01

an offer to a rider is the intermediate step, and completing a carpool ride is the ultimate milestone.

5.4.2. Variables and Results. As discussed, our interventions had a substantial impact on the carpool intent (captured by CTR and OBR). The next natural question is whether our experiment could affect carpool adoption. Accordingly, we consider two types of variables related to carpool adoption: TSO---number of users who sent at least one offer during the next 30 or 180 days after being exposed to the field experiment---and TCR---number of users who completed at least one carpool ride (both 30 and 180 days after being exposed to the field experiment). When computing these metrics, we consider two options depending on the sample of users: (1) focusing on users who clicked on the notification from our field experiment or (2) including all the users who received the notification in our field experiment. There is a clear trade-off between these two options. On the one hand, the second option includes a much larger sample size (168,471 for L users and 63,606 for H users) but may suffer from interferences (e.g., other field experiments and campaigns run after our intervention). As a result, we increase the number of observations at the expense of being less accurate and not carefully isolating the effect of our intervention. Note that if users did not click on the notification but later on tried the carpool service, the effect is probably not driven by our field experiment. On the other hand, the first option has a smaller sample size (3,386 for L users and 1,483 for H users) but allows us to better isolate the effect of our intervention. In addition, it allows us to examine the progression of users in the carpool conversion process, which is an interesting analysis in itself. We report the estimation tables only for the first option, but the majority of the results also hold for the second option.

The results using separate samples (i.e., H and L users separately) for all four dependent variables

Figure 13. Regression Estimates Using CTR for All Users Who Received a Generic Message (HC, LC, NB)

	Dependent variable:			
	formula1	formula2	formula1	formula2
	OLS	*OLS*	*logistic*	*logistic*
	(1)	(2)	(3)	(4)
factor(experiment_group)1.HC	-0.0002	-0.001	-0.011	-0.041
	(0.001)	(0.001)	(0.075)	(0.088)
factor(experiment_group)2.LC	-0.0003	-0.001	-0.025	-0.058
	(0.001)	(0.001)	(0.054)	(0.068)
log(avg_distance_km)		0.002^{***}		0.128^{***}
		(0.001)		(0.048)
log(avg_leave_time)		0.008^{***}		0.618^{***}
		(0.003)		(0.210)
log(days_since_waze_ob)		-0.002^{***}		-0.153^{***}
		(0.0004)		(0.027)
log(sessions_30d)		0.001^{**}		0.096^{**}
		(0.001)		(0.043)
factor(market)GA		-0.001		-0.045
		(0.001)		(0.070)
factor(market)MA		0.0005		0.037
		(0.001)		(0.076)
factor(market)WA		-0.004^{**}		-0.303^{**}
		(0.002)		(0.131)
Constant	0.014^{***}	-0.001	-4.265^{***}	-5.413^{***}
	(0.0005)	(0.007)	(0.035)	(0.546)
Observations	120,021	120,021	120,021	120,021
R^2	0.00000	0.0005		
Adjusted R^2	-0.00001	0.0004		
Log Likelihood			-8,694.772	-8,666.139
Akaike Inf. Crit.			17,395.540	17,352.280
Residual Std. Error	0.116 (df = 120018)	0.116 (df = 120011)		
F Statistic	0.108 (df = 2; 120018)	6.466^{***} (df = 9; 120011)		
Note:				*p<0.1; **p<0.05; ***p<0.01

(TSO30, TSO180, TCR30, TCR180) when setting the baseline as HC or LC can be found in the online appendix. In all cases, we find that the impact of our interventions was not statistically significant. Two potential reasons can explain this finding. First, it is possible that the magnitude of the effect is too small to be identified in our sample. Second and much more likely, our intervention is not powerful enough to trigger carpool adoption (which requires some additional efforts from the user).[10] Although sending an app notification can substantially boost carpool intent (as we saw in this paper), it is not enough of a nudge to prompt carpool adoption. Thus, after users have successfully onboarded, it can be valuable to send them additional notifications to remind them about the benefits of carpooling. Note that we did not find any statistically significant effect in terms of heterogeneous treatment effects for the carpool adoption variables.

The fact that our interventions could affect carpool intent but not carpool adoption is an important insight. Although this is not an ideal outcome, we believe that it is still important to report it. This finding calls for future research on better understanding the carpool conversion process. First, it shows the critical distinction between intent and adoption. Second, it can save time to future researchers who plan to run field experiments in the context of carpooling. Finally, this result suggests that in order to stimulate carpool adoption, additional follow-up interventions are needed. After all, our intervention was only a one-time message, and thus, it makes sense that this was not enough to fully convert commuters to carpoolers. One alternative is to employ a sequential nudging approach (e.g., by sending periodic reminders at specific times with an adaptive framing, depending on the past reactions of users; this is left for future research).

Table 6. Regression Estimates for CTR Using the Pooled Sample (Including Zip Code Fixed Effects)

	OLS Base	OLS With controls	Logistic Base	Logistic With controls	FELM Home zip code	FELM Work zip code	FELM Both zip codes
HA	0.024***	0.021***	1.026***	0.898***	0.018***	0.021***	0.018***
	(0.001)	(0.001)	(0.054)	(0.059)	(0.001)	(0.001)	(0.001)
HB	0.025***	0.022***	1.066***	0.936***	0.019***	0.022***	0.019***
	(0.001)	(0.001)	(0.053)	(0.059)	(0.001)	(0.001)	(0.001)
HC	−0.0001	−0.003**	−0.008	−0.137*	−0.006***	−0.003**	−0.006***
	(0.001)	(0.001)	(0.075)	(0.079)	(0.001)	(0.001)	(0.001)
LA	0.019***	0.016***	0.873***	0.752***	0.014***	0.016***	0.015***
	(0.001)	(0.001)	(0.044)	(0.05)	(0.001)	(0.001)	(0.001)
LB	0.018***	0.015***	0.843***	0.723***	0.013***	0.015***	0.3013***
	(0.001)	(0.001)	(0.044)	(0.05)	(0.001)	(0.001)	(0.001)
LC	−0.0003	−0.003***	−0.023	−0.144**	−0.004***	−0.003***	−0.004***
	(0.001)	(0.0001)	(0.054)	(0.059)	(0.001)	(0.001)	(0.001)
NA	0.0001	0.0001	0.009	0.008	0.0001	0.0004	0.0001
	(0.001)	(0.001)	(0.049)	(0.049)	(0.001)	(0.001)	(0.001)
log(*avg_distance*)		0.003***		0.12***	0.003***	0.003***	0.003***
		(0.001)		(0.026)	(0.001)	(0.001)	(0.001)
log(*avg_leave_time*)		0.005**		0.227**	0.006**	0.004*	0.005**
		(0.002)		(0.11)	(0.002)	(0.002)	(0.002)
log(*std_leave_time* + 1)		0.00004		0.017	0.0002	−0.0001	0.0002
		(0.001)		(0.072)	(0.001)	(0.001)	(0.001)
is(*std_leave_time_NA*)		0.005***		0.266***	0.005***	0.005***	0.005***
		(0.001)		(0.039)	(0.001)	(0.001)	(0.001)
log(*days_since_joined*)		−0.002***		−0.07***	−0.001***	−0.002***	−0.001***
		(0.0003)		(0.014)	(0.0003)	(0.0003)	(0.0003)
log(*sessions_30d*)		0.001***		0.059***	0.001***	0.001**	0.001**
		(0.0005)		(0.022)	(0.0005)	(0.0005)	(0.0005)
GA		−0.006***		−0.275***			
		(0.001)		(0.037)			
MA		−0.003***		−0.15***			
		(0.001)		(0.045)			
WA		−0.008***		−0.348***			
		(0.001)		(0.056)			
Constant	0.014***	0.006	−4.271***	−4.652***			
	(0.001)	(0.006)	(0.035)	(0.277)			
Observations	301,303	301,303	301,303	301,303	301,303	301,303	301,303
R^2	0.005	0.006			0.017	0.017	0.027
Log likelihood			−30,796.900	−30,594.600			
Residual standard error	0.145 (df = 301,295)	0.145 (df = 301,285)			0.145 (df = 298,439)	0.145 (df = 298,708)	0.145 (df = 296,218)
F statistic	200.312*** (df = 7; 301,295)	107.640*** (df = 17; 301,285)					

Notes. OLS corresponds to Equation (3), logistic corresponds to Equation (4), and FELM corresponds to a fixed effects linear model when we control for the zip code. df, degrees of freedom.
$^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

6. Conclusions and Implications

The results presented in this paper bear interesting implications on the design of carpooling platforms. First, our results confirm that commute time saving is a strong value proposition for carpooling (Giuliano et al. 1990). The ability to "bundle" stops along a route to gather enough passengers and be eligible to use the HOV lane may outweigh small detours to pick up riders. Second, if monetary incentives do not have the desired effect, targeting should focus on users who have a high utility. We have shown that a simple notification can be enough when targeted to the right set of users and when using the right framing. Future research will allow us to scale such efforts via personalized machine learning models. It will allow us to identify several groups of users to target and for each group, to deploy the best framing or motivating factor. Either way, it is important to explicitly tell users why they are targeted and why they can benefit from carpooling. Indeed, we have shown that although drivers seem to intuitively know how much time they might save by taking the HOV lane, it is still valuable to emphasize the link between the desired course of action and the eligibility to use the HOV lane. Third, our experiment has been limited to showing success on carpool intent but unfortunately, did not boost carpool adoption. Connecting the call to action to the ability to offer a ride would involve notifying drivers early enough for them to adjust their departure time. A smart notification system that sends such notifications in advance based on drivers' learned routines would allow the platform to match drivers precisely when needed.

We conclude this section by highlighting some of the ramifications our results have on policy. Our study ultimately suggests that users are receptive to saving commute time. Stating the obvious, for HOV lanes to be effective, it is important that they actually lead to meaningful time savings. Time saving should be taken into account when deciding where to locate an HOV lane and when setting its restrictions (e.g., two versus three occupants). When HOV lanes save time, we can generate intent through simple messaging; however, the ability to offer rides ahead of time may be necessary to generate actual adoption. This would involve informing users of potential time savings early or close to their home before carpooling is no longer an option. More importantly, adoption will exist when car owners do not just offer a ride but when some drivers are also willing to leave their car at home—and for this to happen, they need to be highly likely to find a driver. To solve this "chicken-and-egg" problem and bootstrap the marketplace, government agencies could lock in supply by making drivers eligible to use the HOV lane just for committing to pick up riders, even if none materialize. Finally, we inferred that both the type of user and the framing of the intervention play important roles in converting commuters to carpool. Government agencies or large corporations could use the results presented in this paper to incentivize their workers to carpool together by highlighting the potential time saving via HOV lanes.

In this paper, we mainly focused on time saving from using the HOV lane, but many other interventions are also relevant. For example, firms can offer carpoolers to cover specific "painful expenses" (e.g., tolls, car insurance, parking fees). A concrete example is of a large company that can incentivize its workers to carpool together by offering a parking spot to employees who carpool. Cities and government agencies can also adopt the same approach by offering convenient parking spots to carpoolers. Another type of intervention is to highlight the environmental benefits (Byerly et al. 2018), such as the CO_2 reductions and the impact of carpooling on congestion and pollution. For example, companies can run a campaign on Earth Day to sensitize users. A third alternative is to emphasize the social aspects of a carpool ride by meeting new interesting people (e.g., colleagues for carpooling programs run by large companies).

Convincing even a small portion of the 116 million U.S. solo commuters to carpool can have significant environmental and societal impacts. In the same vein, given that there are more than 110 million Waze drivers, if we could convince 1% of them to carpool, this will reduce the number of vehicles on the road by more than 1 million. Thus, sharpening our understanding on the key drivers that can nudge commuters to carpooling is vital. This paper provides a first step toward addressing this question. Our focus in this work was on converting commuters to carpool drivers (i.e., taking a stranger in their car). As discussed, an equally important research question, which is left for future research, is to investigate what determinants and motivating factors can convince commuters to leave their car at home and become carpool riders.

Acknowledgments

The authors thank Yarin Aidelman, Rapha Cohen, Baek Jung Kim, Joshua Loftus, Peleg Samuels, and Hal Varian for insightful feedback and discussions, which have helped to improve this paper.

Endnotes

[1] See http://inrix.com/press-releases/scorecard-2018-us/.

[2] See https://www.waze.com/carpool.

[3] See https://www.timescall.com/2018/10/11/google-owned-waze-aims-to-end-traffic-with-new-carpooling-app/.

[4] A navigation occurs when the user completes a drive from an origin to a destination searched via the Waze app.

[5] See https://www.waze.com/livemap.

[6] We normalize all our figures related to onboarding rates (because of a nondisclosure agreement). A conversion factor has been applied so that the highest number is assigned the value of 1.0, and all other values are adjusted by the same normalizing factor—maintaining the same relationship between different points.

[7] In the regression tables for the onboarding rate, we have scaled all the estimated parameters (along with their standard errors) by a positive constant to avoid revealing business sensitive information. We use a different constant for OLS and logistic models.

[8] We used the cutoffs 27 and 73 as those numbers lead to a clear separation of the average time-saving value.

[9] Note that all users (regardless of our experiment) receive the same $10 monetary incentive as a welcome bonus. The only manipulation we use is to highlight the incentive in the encouragement versus not mentioning it.

[10] We performed a statistical power analysis to tease out between these two reasons. Overall, this analysis shows that if the base rate is above 5% and the minimum relative detectable effect is higher than 50%, our experiment would detect the effect. However, if the base rate is below 3%, then it is definitely possible that our experiment does not have enough statistical power to detect the effect.

References

Benartzi S, Beshears J, Milkman KL, Sunstein CR, Thaler RH, Shankar M, Tucker-Ray W, Congdon WJ, Galing S (2017) Should governments invest more in nudging? *Psych. Sci.* 28(8):1041–1055.

Byerly H, Balmford A, Ferraro PJ, Hammond Wagner C, Palchak E, Polasky S, Ricketts TH, Schwartz AJ, Fisher B (2018) Nudging pro-environmental behavior: Evidence and opportunities. *Frontiers Ecology Environ.* 16(3):159–168.

Choudhary V, Shunko M, Netessine S, Koo S (2022) Nudging drivers to safety: Evidence from a field experiment. *Management Sci.* 68(6):4196–4214.

Cohen MC, Fiszer MD, Kim BJ (2021) Frustration-based promotions: Field experiments in ride-sharing. *Management Sci.*, ePub ahead of print August 31, https://doi.org/10.1287/mnsc.2021.3991.

Ferguson E (1997) The rise and fall of the American carpool: 1970–1990. *Transportation* 24(4):349–376.

Ghose A, Kwon HE, Lee D, Oh W (2019) Seizing the commuting moment: Contextual targeting based on mobile transportation apps. *Inform. Systems Res.* 30(1):154–174.

Giuliano G, Levine DW, Teal RF (1990) Impact of high occupancy vehicle lanes on carpooling behavior. *Transportation* 17(2): 159–177.

Halpern D (2015) *Inside the Nudge Unit: How Small Changes Can Make a Big Difference* (Random House, New York).

Jachimowicz JM, Cunningham JL, Staats BR, Gino F, Menges JI (2021) Between home and work: Commuting as an opportunity for role transitions. *Organ. Sci.* 32(1):64–85.

Kohavi R, Thomke S (2017) The surprising power of online experiments. *Harvard Bus. Rev.* 95(5):74–81.

Konishi H, Mun S-i (2010) Carpooling and congestion pricing: HOV and hot lanes. *Regional Sci. Urban Econom.* 40(4):173–186.

Maxwell SE, Delaney HD (2004) *Designing Experiments and Analyzing Data: A Model Comparison Perspective*, vol. 1 (Psychology Press, Hove, United Kingdom).

Neoh JG, Chipulu M, Marshall A (2017) What encourages people to carpool? An evaluation of factors with meta-analysis. *Transportation* 44(2):423–447.

Ostrovsky M, Schwarz M (2019) Carpooling and the economics of self-driving cars. *Proc. 2019 ACM Conf. Econom. Comput.* (ACM), 581–582.

Richardson M, Dominowska E, Ragno R (2007) Predicting clicks: Estimating the click-through rate for new ads. *Proc. 16th Internat. Conf. World Wide Web* (ACM), 521–530.

Shaheen SA, Chan ND, Gaynor T (2016) Casual carpooling in the San Francisco bay area: Understanding user characteristics, behaviors, and motivations. *Transportation Policy* 51(2016): 165–173.

Singh J, Teng N, Netessine S (2017) Philanthropic campaigns and customer behavior: Field experiments on an online taxi booking platform. *Management Sci.* 65(2):913–932.

Small KA, Winston C, Yan J, Baum-Snow N, Gómez-Ibáñez JA (2006) Differentiated road pricing, express lanes, and carpools: Exploiting heterogeneous preferences in policy design. *Brookings-Wharton Papers Urban Affairs* 2006(2006):53–96.

Sun T, Shi L, Viswanathan S, Zheleva E (2019) Motivating effective mobile app adoptions: Evidence from a large-scale randomized field experiment. *Inform. Systems Res.* 30(2):523–539.

Teal RF (1987) Carpooling: Who, how and why. *Transportation Res. Part A: General* 21(3):203–214.

Thaler RH, Sunstein CR (2009) *Nudge: Improving Decisions About Health, Wealth, and Happiness* (Penguin, London).

Vanoutrive T, Van De Vijver E, Van Malderen L, Jourquin B, Thomas I, Verhetsel A, Witlox F (2012) What determines carpooling to workplaces in Belgium: Location, organisation, or promotion? *J. Transport Geography* 22(2012):77–86.

Yang H, Huang HJ (1999) Carpooling and congestion pricing in a multilane highway with high-occupancy-vehicle lanes. *Transportation Res. Part A: Policy Practice* 33(2):139–155.