Transportation Research Part A 165 (2022) 186–206

Contents lists available at ScienceDirect

Transportation Research Part A

journal homepage: www.elsevier.com/locate/tra

The impact of high-occupancy vehicle lanes on carpooling

Maxime C. Cohen [a], Alexandre Jacquillat [b,*], Avia Ratzon [c], Roy Sasson [c]

[a] Desautels Faculty of Management, McGill University, Montreal, Canada
[b] Sloan School of Management, Massachusetts Institute of Technology, Cambridge, MA, USA
[c] Google/Waze, Tel Aviv, Israel

ARTICLE INFO

Keywords:
HOV lane
Carpooling
Public policy
Natural experiment

ABSTRACT

Since the 1970s, high-occupancy vehicle (HOV) lanes have been a common policy instrument to mitigate traffic congestion. Yet, their effectiveness remains a controversial topic among researchers, policy makers, and the public. In this debate, a key unknown has been the impact of HOV lanes on commuters' carpooling behaviors. This paper brings a new piece of evidence by offering a data-driven assessment of carpooling intent and adoption, using revealed-preferences data. We partner with Waze, a major carpooling platform, and leverage a natural experiment following the introduction of three HOV lanes in Israel in 2019. Using tailored treatment and control groups coupled with econometric analyses, we derive four main findings. First, HOV lanes bring new users to the carpooling platform, which contributes to alleviating the "cold-start" problem in the marketplace. Second, HOV lanes have a positive impact on carpool intent: the number of carpool offers sent by drivers increase manifold following the introduction of the HOV lanes. Third, HOV lanes have a disparate impact on carpool adoption: carpools increase significantly for two out of three HOV lanes. This result underscores the critical impact of HOV lanes design: it seems more beneficial to have round-trip HOV lanes (as opposed to one-way lanes) and two-passenger occupancy requirements (as opposed to three-passenger requirements). Last, HOV lanes have a broader impact, by increasing carpooling on non-HOV routes and shifting the travel behaviors of non-carpoolers. We conclude by discussing policy implications, highlighting collaboration opportunities between policy makers and digital carpooling platforms to enhance the design and operations of HOV lanes.

1. Introduction

Rising mobility needs and private car ownership have led to growing congestion and pollution. In the United States, the economic impact of road congestion was estimated at $88 billion in 2019, or $1,377 per commuter (INRIX, 2020). Infrastructure expansion is very expensive and often infeasible in dense urban areas or highway corridors. Alternative demand management approaches include congestion pricing (like in London, Singapore, and New York City, for example), public transit expansion (like in Seattle and Houston), and investments into alternative modes of transportation such as bikes and scooters. Despite some notable successes, these interventions are not always sufficient to curb traffic congestion, especially outside dense urban areas. Ultimately, traffic congestion remains a complex challenge for governments with no one-size-fits-all solution.

One of the widespread non-monetary congestion mitigation interventions has been the deployment of high-occupancy vehicle (HOV) lanes, namely, dedicated traffic lanes for multi-occupant vehicles. Their primary purposes are to "provide an incentive to use

* Correspondence to: Sloan School of Management, Massachusetts Institute of Technology, USA.
E-mail addresses: maxime.cohen@mcgill.ca (M.C. Cohen), alexjacq@mit.edu (A. Jacquillat), aviar@google.com (A. Ratzon), rsasson@google.com (R. Sasson).

https://doi.org/10.1016/j.tra.2022.08.021
Received 16 January 2022; Received in revised form 14 July 2022; Accepted 25 August 2022

ridesharing and public transportation, remove congestion from normal lanes of travel, and improve overall traffic operations" (U.S. Department of Transportation, 2015). Since the first HOV lane opened in 1969, HOV lanes have been deployed across the six continents. Altogether, HOV lanes amount to a total of nearly 2,000 miles in the United States (U.S. Office of Highway Policy Information, 2014), but are also employed in Canada, Europe, Australia, and other parts of the world (Schijns and Eng, 2006). Current designs are at high variance in terms of directionality (concurrent flow, contraflow, or reversible), temporal use (peak period only vs. full day), passenger requirements (2+, 3+, 4+, 6+), etc.

Despite their prevalence, there is no general consensus on the benefits of HOV lanes in terms of congestion mitigation. In this debate, a critical unknown is the impact of HOV lanes on commuters' carpooling behaviors and the resulting dynamics on carpooling marketplaces. The main contribution of this paper is to provide the first data-driven response to this question by using revealed preferences data. To this end, we leverage real-world data from Waze carpooling platform and a natural experiment setting following the introduction of three HOV lanes in Israel in 2019. Before presenting our contributions in more detail, we first review the current state of knowledge on the traffic impact of HOV lanes as well as the main research gaps in this context.

1.1. Literature review

Policy discussion on HOV lanes. Since their inception, HOV lanes have induced significant debates among researchers, policy-makers, and the public. On the one hand, several surveys indicated strong public support (Lawler, 1991; Martin et al., 2005). Early evidence suggests that HOV lanes in North America have been successful in reducing travel times, especially at peak hours (Wellander and Leotta, 2000; Poppe et al., 1994; Martin et al., 2004; Boriboonsomsin and Barth, 2007).

Yet, voices argue that HOV lanes are not an effective intervention for congestion mitigation. The core argument is that HOV lanes are often underutilized and do not balance the additional congestion created in non-HOV lanes, so that the high infrastructure cost is not justified (Wiseman, 2019). One of the earliest HOV lanes, known as the Santa Monica Freeway Diamond Lane, led to controversy and was reverted to a general-purpose lane after months of protests (Billheimer, 1978). Similarly, HOV lanes were opened in New Jersey in 1994 and in 1998, but ended up being closed due to underutilization and strong political opposition. Similar failures exist outside the United States; for instance, a HOV lane opened near Amsterdam in 1993 and, after public criticism, was opened to general traffic in 1994. Ultimately, whereas HOV lanes are generally viewed as an effective congestion mitigation lever (among others), their effectiveness remains an open question.

HOV lanes are most prevalent in the United States but are also used in other metropolitan areas around the world (Turnbull, 1992). A particularly interesting case is that of Jakarta, Indonesia, where HOV lanes restricted access to some of the busiest streets to vehicles with at least three passengers. Even with this policy, Jakarta was facing high levels of congestion. In addition, the presence of HOV lanes led to the phenomenon of "car jockeys", that is, individuals and families seeking carpools without needing to travel in exchange for payments. These behaviors were seen as defeating the purpose of HOV lanes, so these were suspended and replaced by an "odd-and-even" license plate system.[1] Yet, Hanna et al. (2017) found that the removal of the HOV lanes had a significant impact on congestion, with an increase in average delays by 1–2 min per kilometer. Moreover, the authors identified spillover effects across the city, showing that HOV lanes can have a positive effect on congestion that ripples through a broader geographic area.

In summary, evidence suggests that HOV lanes can have a positive impact on congestion. However, their effect is far from being unequivocal. Ultimately, the impact of HOV lanes depends on travelers' behaviors and on whether HOV lanes incentivize carpooling.

Modeling HOV lanes. The main premise of HOV lanes is that, keeping all else constant, congestion is borne by high-occupancy vehicles as opposed to low-occupancy vehicles, thus reducing total passenger delay. This argument holds as long as total vehicle delay does not increase too much with HOV lanes. An important question in traffic engineering therefore lies in the impact of HOV lanes on vehicle delay. Laval and Daganzo (2006) found that HOV lanes induce a lane-shifting behavior, which has a negative impact on bottleneck capacity—although (Menendez and Daganzo, 2007) argued that lane changes have a minimal effect on general-purpose lanes. Using queuing models, Daganzo and Cassidy (2008) estimated that HOV lanes only increase vehicle delays by 2%, resulting in significant reductions in passenger delays (by 10%). Overall, the effectiveness of HOV lanes depends on factors such as demographics, traffic demand, traffic patterns, and road design (Ben-Akiva and Atherton, 1977b; Mannering and Hamed, 1990; Johnston and Ceerla, 1996; Rodier and Johnston, 1997; Yang, 1998; Dahlgren, 1998; Kwon and Varaiya, 2008; Shewmake, 2012).

This literature evaluates the impact of HOV lanes on traffic congestion given a number of vehicles on the road. However, by introducing a differentiation in travel times between low-occupancy and high-occupancy vehicles, HOV lanes create, at least in theory, an incentive to carpool. The transportation economics literature has proposed analytical models with endogenous numbers of vehicles. Yang and Huang (1999) explored the interplay between HOV lanes and congestion pricing by suggesting a differentiated pricing scheme between HOV lanes and general-purpose lanes, and by designing a second-best uniform pricing scheme when differentiated pricing is impossible to implement. Konishi and Mun (2010) compared the effectiveness of HOV lanes and HOT lanes (which can also be used by solo drivers who are willing to pay a toll), based on a commuter choice model that trades off travel times, prices, and carpool organization costs. Zhong et al. (2020) optimized carpool prices in the presence of HOV and HOT lanes, while endogenizing travel mode choices. Cui et al. (2021) found that free carpools are socially suboptimal and can lead to more cars on the road, hence leading to a higher environmental footprint—suggesting that monetary exchanges between drivers and passengers are beneficial. These studies are based on theoretical models of commuters' behaviors but do not address the carpooling question empirically.

[1] https://voxdev.org/topic/infrastructure-urbanisation/jakarta-s-maligned-congestion-easing-policy-actually-worked, www.theguardian.com/world/2016/apr/04/end-of-the-road-jakartas-passengers-for-hire-targeted-by-carpooling-crackdown

Empirical evidence on carpooling. Gathering evidence on commuters' attitudes towards carpooling is complicated by the lack of easily-accessible information on passenger ridership. One thread of studies has therefore relied on stated-preferences data. Using a mail-based and telephone-based survey of commuters in California, Giuliano et al. (1990) found that the Route 55 HOV in Orange County increased carpools among peak-period commuters, but did not increase carpooling more broadly in California. Parkany (1999) found that the HOT lanes on Route 91, also in Orange County, did not increase carpooling significantly. Tackling the question in the reverse direction, Li et al. (2007) surveyed commuters who did carpool, and found that the availability of a HOV lane was the main factor motivating their decisions. Plotz et al. (2010) used data from the 2001 National Household Travel Survey and found that most HOV trips were taken among family members—a finding echoing the phenomenon of "fampooling" termed by Li et al. (2007).

That said, HOV lanes have also incentivized new types of "casual carpools", or "slugging", among people who may not even know each other (Beroldo, 1990). Matches commonly occur in informal locations where passengers wait for vehicles and where vehicles stop to seek for additional passengers. Spielberg and Shapiro (2000) collected data from a HOV lane with a three-passenger occupancy requirement, and found that slugs often come in pairs—an argument in favor of three-passenger occupancy requirements versus two-passenger occupancy requirements. Using survey data, Burris and Winn (2006) found that slugging is most prominent among regular commuters and young professionals. Mote and Whitestone (2011) showed that slugging became a structured and institutionalized practice, underscoring the role of transportation in the broader social apparatus. Shaheen et al. (2016) found that slugging is primarily motivated by convenience, time savings, and monetary savings, whereas environmental implications are merely a positive byproduct. Finally, Kelly (2007) proposed incentives to extend slugging to areas without HOV lanes.

Using a prescriptive approach, Ma and Wolfson (2013) and Boysen et al. (2021) proposed optimization formulations to build effective carpools along HOV lanes, which result in significantly better travel options than spontaneous carpools in slugging areas. Tsitsokas et al. (2021) and Anderson and Geroliminis (2020) study the closely-related topic of modeling and optimizing dedicated bus lanes. These models rely on underlying assumptions on travelers' mode choice; viewed through this lens, our paper provides empirical evidence to identify travelers' carpooling behaviors along HOV lanes, which can be subsequently used in HOV lane design and in carpool formation.

This literature also falls into the broader body of work identifying the drivers of carpooling behaviors. Using discrete-choice models, Ben-Akiva and Atherton (1977a) found a small response to carpooling incentives. Vanoutrive et al. (2012) observed higher levels of carpooling in less accessible locations and in sectors such as construction, manufacturing, and transport. Using a meta-analysis, Neoh et al. (2017) found that carpooling behaviors were more impacted by the professional environment (e.g., number of employees, partner matching programs, work schedule) than by financial incentives. This result identifies coordination costs as a prominent barrier to carpooling due to increased wait times, detours, and schedule deviations. New opportunities thus arise thanks to on-demand ride-sharing platforms, which suggest carpooling options with limited inconvenience. Using data from Lyft, Brown (2020) reported that one third of ride-hailing trips are shared, although a small fraction of users actually are sharing trips. Using a large field experiment to incentivize carpools, Cohen et al. (2019) found that financial incentives do not impact carpooling significantly but that targeting commuters who can save time via a HOV lane is more effective.

Summary. The above discussion has outlined two key takeaways, which we address in this paper:

1. There is no general evidence that HOV lanes lead to a drastic reduction in the number of vehicles on the road. This argument lies at the core of the criticism of HOV lanes as a congestion mitigation lever. This can be primarily explained by the coordination costs associated with carpooling, and the disproportionate impact of idle delay as compared to in-vehicle delay (Dahlgren, 1998). In this regard, analytics-powered marketplaces provide opportunities to trim these coordination costs and make carpooling a more attractive travel option. At the same time, carpooling marketplaces potentially face a "cold-start problem" to build a sufficient mass of drivers and riders that can create attractive carpooling options.
2. There is no empirical evidence on the impact of HOV lanes on carpooling from revealed-preferences data. The literature has either treated vehicle demand as constant, used theoretical models of carpooling behaviors, or relied on surveys to characterize commuters' behaviors. Menendez and Daganzo (2007) explicitly recognize that it is "difficult to quantify" whether a HOV lane induces carpooling because traffic sensors record vehicle-based information (e.g., count and speed) but not passenger-based information (e.g., number of occupants).

1.2. Empirical setting and contributions

Our paper contributes to the literature by providing the first data-driven assessment of the impact of HOV lanes on carpooling intent and adoption, using revealed preferences data and a natural experiment setting. We partner with Waze, a major navigation software company that operates a carpooling platform. Users can register as drivers or riders and send ride-sharing offers. The platform matches riders and drivers, handles payments, and proposes optimized pick-up and drop-off locations. As such, the Waze platform generates carpooling data that provide visibility into vehicle occupancy, which we leverage to characterize drivers' and riders' carpooling behaviors.

A key challenge to address our research question is that most HOV lanes were built decades ago, which makes it hard to isolate their impact from other confounding factors. Therefore, we focus on three recent HOV lanes in Israel, which provides visibility into carpooling behaviors before and after their introduction—and allows us to cast this event as a natural experiment (Dunning, 2012). Specifically, motivated by the low vehicle occupancy rates (1.2 people per vehicle on average, as opposed to 1.6–1.8 on average among OECD countries), the Israeli government introduced the country's first three HOV lanes in October, 2019: (i) a 30-kilometer stretch from Netanya to Herzliya; (ii) the reverse 30-kilometer stretch from Herzliya to Netanya; and (iii) a 22-kilometer stretch on the Ayalon highway. The former two HOV lanes have a two-passenger requirement, whereas the latter has a three-passenger requirement. These HOV lanes have been subject to heated debates in Israeli politics: the Israel ministry of transportation qualifies the policy as a success,[2] whereas several media articles argue that they lengthen travel times and create chaos.[3] This paper contributes to this question by estimating the impact of the HOV lanes on carpooling.

We design an experimental setting to isolate the impact of the HOV lanes. One complication here is the dichotomy between road segments and origin–destination routes—HOV lanes are designed at the segment level but commuters travel from various origins to various destinations. To circumvent this challenge, we create treatment and control groups at the route level. Our treatment group comprises a sample of origin–destination pairs that are most exposed to the new HOV lanes. We then select controlled routes that are "similar" to the treated ones in terms of distance and navigation patterns, but are not exposed to the HOV lanes. As such, the experimental population is homogeneous prior to the introduction of the HOV lanes. We then validate that the treatment did take place: commuters are indeed using the HOV lanes, and the HOV lanes reduce their travel times (with a median time reduction of up to 8.1, 5.7, and 15.7 min for the three HOV lanes).

Our empirical results can then be classified into the following four categories:

– The HOV lanes bring new users to the carpooling platform, which contributes to alleviating the "cold-start" problem in the marketplace. We quantify this phenomenon through the quick ratio metric, defined as the ratio of new users over lost users. We find that, early on, new drivers place carpooling offers on the platform but the rider mix remains mainly unchanged, lead to no significant changes in the mix of carpools. Later on, however, HOV lanes brought new drivers as well as new riders to the platform, ultimately creating new successful carpooling matches. As such, HOV lanes can help carpooling marketplaces scale both on the demand side and the supply side, and can thus alleviate the cold-start problem.
– The HOV lanes have a positive impact on carpooling intent. For all three HOV lanes, the number of offers from drivers increases significantly after the policy change. This shows that drivers are indeed reacting to the new HOV lanes by reaching out to potential riders. In other words, the cost of carpooling coordination can be outweighed by the resulting time savings induced by the HOV lanes. At the same time, the number of offers sent by riders did not change significantly. This is consistent with the fact that HOV lanes serve as a supply-driven policy, targeted primarily to drivers.
– The three HOV lanes have disparate effects on carpool completions. Specifically, the number of completed carpools increases (by up to hundreds of percent) for the Netanya–Herzliya and Herzliya–Netanya lanes. In contrast, the number of carpool rides does not increase significantly on the Ayalon lane. This disparity can be explained by two factors. First, the former two lanes go in both directions, which strengthens the benefits of carpooling for round-trip commuters. Second, the latter lane requires three occupants (as opposed to two occupants), which increases the coordination costs for drivers. These results underscore the role of HOV design and market dynamics on the ultimate impact of HOV lanes on carpooling.
– The HOV lanes have indirect effects on carpooling. We find that the carpool rates also increase on routes that are unaffected by the HOV lanes. This effect can be explained by the growth of the Waze carpooling platform, but also potentially by the fact that the new HOV lanes raised public awareness on carpooling opportunities—even on routes where they do not have a direct impact on travel times. Ultimately, this paper therefore uncovers three impacts of HOV lanes on carpooling behaviors: the direct travel time savings on the routes under consideration, the effectiveness of the carpooling marketplace as a means of matching riders and drivers, and the general awareness of HOV lanes and carpooling opportunities within the population. We conclude by discussing the policy implications of these findings.

In summary, the main contribution of this paper is to provide data-driven evidence on the positive impact of HOV lanes on carpooling intent and adoption, using revealed-preferences data. In addition, this paper highlights two critical roles that mobility platforms in infrastructure policy: (i) providing new sources of commuting data to assess infrastructure investments (by getting visibility into vehicle occupancy, in our case), and (ii) developing digital technologies to enhance the impact of infrastructure investments (by lowering the barriers to carpooling, in our case).

2. Setting and data

Our research question is how the HOV lanes impacted the marketplace dynamics on carpooling platforms and, in particular, carpooling intent and carpooling adoption. Note that this question focuses on the within-platform behaviors, as opposed to the broader congestion patterns. Let us provide in this section some background on the three HOV lanes under consideration.

2.1. Empirical setting

We consider the three HOV lanes introduced in Israel in October 2019:

1. **Netanya–Herzliya (HN):** This HOV lane connects Netanya (a city in the North of the country) to Herzliya (a city in the Tel Aviv area, in the middle of the country). Its primary goal is to ease traffic for morning commuters. The highway switched from a configuration with three general-purpose lanes to a configuration with two general-purpose lanes and one HOV lane (Fig. 1(a)). This HOV lane is a 2+, that is, vehicles with at least two occupants can use it (including public transportation vehicles).

[2] https://www.calcalistech.com/ctech/articles/0,7340,L-3771641,00.html
[3] https://en.globes.co.il/en/article-israel-carpool-lanes-proving-a-disaster-1001304806

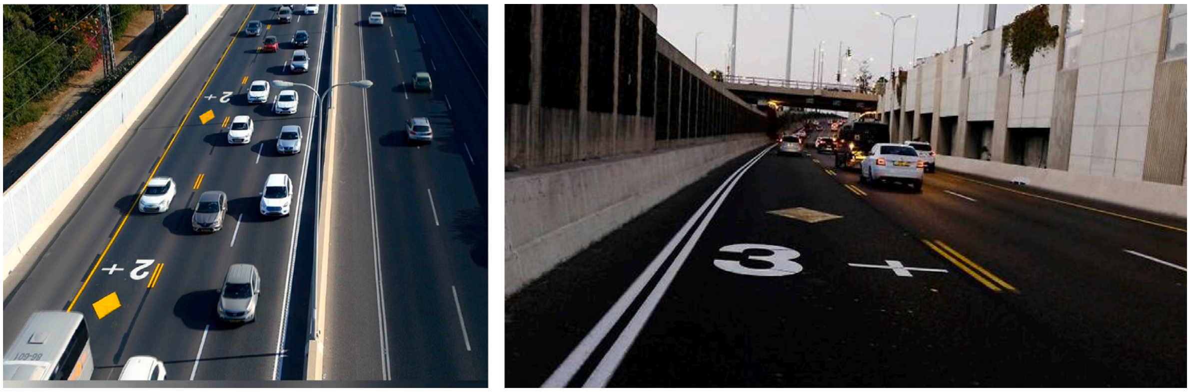

(a) HN lane (source: Migdalor News, 2020) **(b)** MK lane (source: Ynet News, 2020)

Fig. 1. Pictures of the HOV lanes.

2. **Herzliya–Netanya (NH):** This HOV lane covers the opposite direction of HN, targeted to evening commuters. This HOV lane is also a 2+ and has the same structure as HN.
3. **Mavo Ayalon–Kibbutz Galuyot (MK):** This HOV lane connects Rishon LeTsiyon (a city in the South of the country) to the middle of the country along the Ayalon highway. Before the introduction of the HOV lane, the highway had two general-purpose lanes and one reserved for public transportation during morning hours, and three general-purpose lanes outside of the morning peak. Since the introduction of the HOV lane, the highway operates under a traditional configuration with two general-purpose lanes and one HOV lane (Fig. 1(b)). This HOV lane is a 3+, that is, vehicles with at least three occupants can use it.

By design, these three HOV lanes were located along high-traffic commuting segments in order to impact tens of thousand commuters daily. All HOVs are free of charge for vehicles that satisfy the occupancy requirement.

All HOV lanes were introduced on the same day. In fact, the policy materialized in two phases: (1) a *soft enforcement* (on October 6, 2019) and (2) a *hard enforcement* (on November 3, 2019). During the soft enforcement period, drivers who violate the occupancy requirement of the HOV lanes were given a warning without a monetary penalty. During the hard enforcement period, violators who are caught have to pay a fine of 500 Shekels (approximately $150).[4]

Accordingly, we divide our data into three periods. For each one, we focus on high-traffic commuting days, so we remove weekends (Friday and Saturday, in Israel) and Jewish holidays.

– **Before.** This is the period without HOV lanes: September 1, 2019 to October 5, 2019.
– **Soft.** This is the period of soft enforcement: October 6, 2019 to November 2, 2019.
– **Hard.** This is the period of hard enforcement: November 3, 2019 to January 2, 2020.

Extensive media coverage created strong public awareness for the HOV lanes. Fig. 2 shows Google Trends data in Israel on four queries related to the HOV lanes: "Nativ Plus" (the Hebrew term for "HOV lanes"), "Carpool", "Waze", and "Bus Lines". Note that the "Nativ Plus" queries spiked twice, once just before the introduction of the HOV lanes, and once more just before the hard enforcement period. Interestingly, as soon as people started talking about "Nativ Plus", we also observe a spike in the number of queries for "Carpool".

2.2. Data and metrics

To measure the impact of HOV lanes, we leverage carpooling data from Waze. Waze launched its carpooling platform in Israel in 2015 and operates in the United States, Brazil, Mexico, and Israel. The company reached 1 million carpools by the end of 2019. We have access to carpool data before and after the introduction of the HOV lanes, which allows us to assess the impact of each HOV lane on commuting times, carpool intent, and carpool adoption.

We divide the entire country of Israel into small "cells". Each cell is a X-by-X square kilometer. A route is defined as a navigation request (through the Waze app) originating from one cell and ending in another cell. We focus on navigation requests originating from distinct users, that is, if the same user places several consecutive requests, we count those as a single one.

Namely, we consider the following four variables:

– **Driver offers:** Number of offers sent by drivers via the carpool platform.
– **Rider offers:** Number of offers sent by riders via the carpool platform.
– **Completed carpools:** Number of completed rides via the carpool platform.
– **Time saving:** Difference between the time it takes to complete a navigation with the HOV lane and without the HOV lane.

[4] https://www.jpost.com/Israel-News/Drivers-violating-carpool-lane-laws-will-be-fined-606677

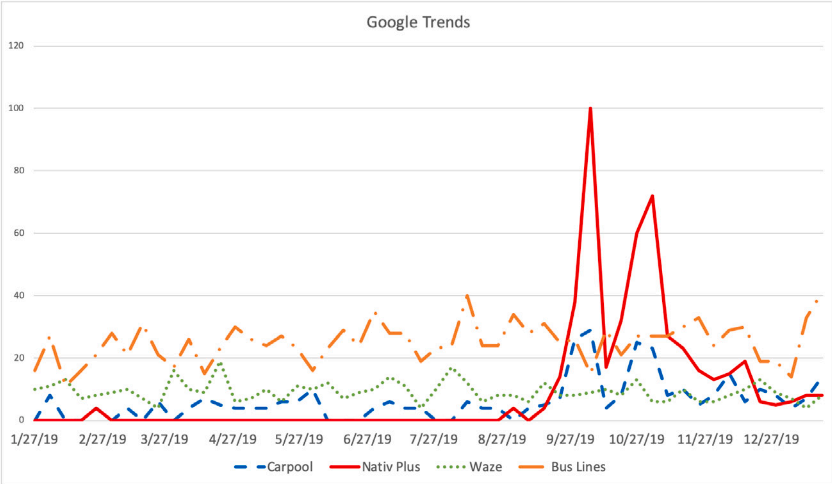

Fig. 2. Google Trends results in Israel for search queries related to HOV lanes.
Source: Google Trends, 2020.

By design, these variables capture the different steps in the "carpool funnel". Specifically, the number of offers from drivers and riders captures carpool intent, whereas the number of completed carpools measures carpool adoption. Finally, the time savings metric assesses effectiveness of the HOV lanes in terms of reducing commuting time.

All variables are standardized and defined over routes. We compare each one during the three periods (before, soft, and hard) by computing the daily averages. To capture temporal dynamics, we separate morning commutes (6AM–11AM) and evening commutes (3PM–8PM).

This dataset enables us to link, for the first time to our knowledge, the introduction of HOV lanes to carpooling behaviors—by providing visibility into passenger-level occupancy information as opposed to vehicle-level information. Still, it is worth mentioning that our data has restrictions. First, our data does not include every vehicle on the road. In theory, it could provide a partial representation of the impact of HOV lanes on carpooling behaviors. However, this is a minor concern given the prevalence of the Waze carpooling platform in Israel—as of February 2020, we record over 4 million monthly Waze users in Israel out of 9 million people and 3.5 million registered vehicles.[5] Moreover, the penetration of Waze is even higher among carpoolers than general commuters. Therefore, our data provide a good representation of the overall picture. Second, our data may suffer from biases; for instance, the carpool platform may be primarily used by early adopters or some specific types of users. This is also a minor concern, since our analysis compares user behaviors within the platform across treated vs. controlled routes.

2.3. Route selection

Ideally, we would define "HOV routes" and "non-HOV routes" and compare travel times and carpooling behaviors among the two groups. This analysis, however, is complicated by the dichotomy between the HOV lanes, implemented at the segment level, and the commuting routes, defined from an origin to a destination across the country. In order to isolate the impact of the HOV lanes, we need to carefully select treatment and control groups, that is, commuting routes that are heavily impacted by the new HOV lanes versus commuting routes that are not.

For each route, we determine the path from origin to destination from the Waze navigation system. We divide each HOV lane into disjoint segments, each approximately one kilometer long. We then consider all the routes that are at least 10-minute long and whose path overlap with at least 50% of the HOV segments.[6] Finally, we retain the seven most popular routes in terms of number of navigation requests during the 40 days prior to the introduction of the HOV lanes. This process yields routes that are directly affected by the HOV lanes (significant commuting time and high overlap with the HOV segments) and that are popular among commuters (large number of navigation requests). Ultimately, we obtain 21 routes (7 times 3), which we call the *treated routes*.

Similarly, we select controlled routes to control for seasonality effects and for external variation (e.g., national advertisement campaigns). Specifically, we consider two types of controlled routes:

1. **Control 1:** For each treated route, we seek a route with the same destination that is not affected by the HOV lanes. We identify all routes with the same destination that have a weak overlap with the HOV lanes (less than 5% of the segments) and that are not too similar to the treated route (e.g., two neighboring origin cells). We then select the closest route among that set, in terms of distance and number of navigation requests over the preceding 40 days.

[5] https://www.ceicdata.com/en/israel/number-of-registered-motor-vehicles

[6] Since we consider small enough cells and long enough routes, all requests within the same route were associated with the same paths on highways, hence with the same level of overlap with HOV lanes. Therefore, the level of overlap can be defined at the route level, as opposed to the request level.

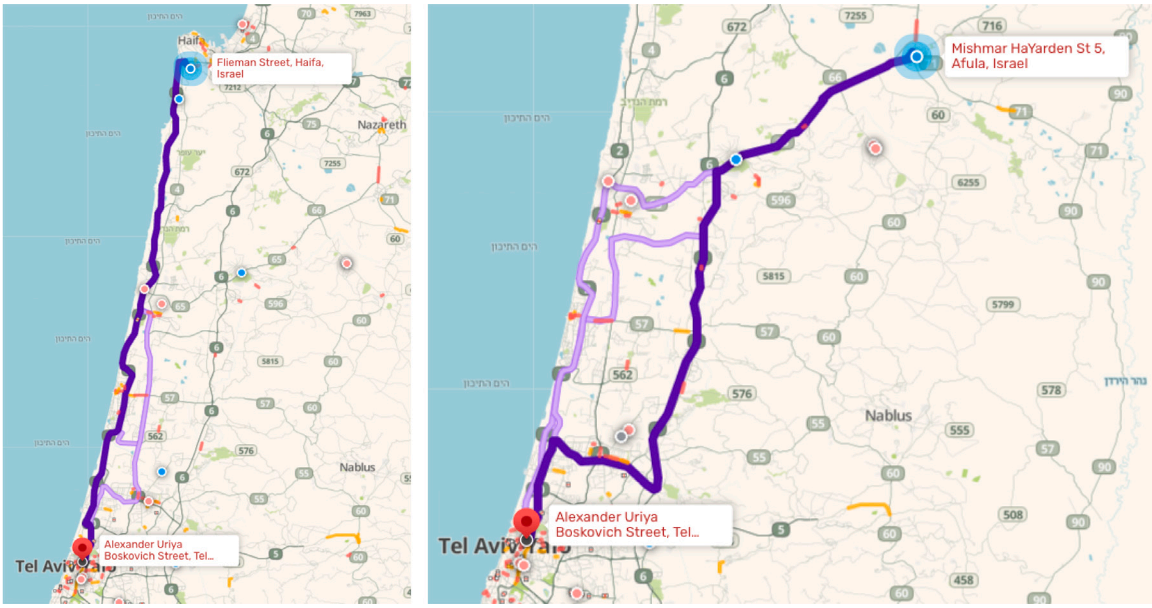

(a) Treated route from Haifa to Tel Aviv **(b)** Control 1 route from Afula to Tel Aviv

Fig. 3. Example of a treated route and the corresponding route from Control 1 for the NH HOV lane.

2. **Control 2:** For each treated route, we seek a route with a different origin and a different destination that is not affected by the HOV lanes. We identify all routes across the country that have a weak overlap with the HOV lanes (less than 5% of the segments). Again, we select the closest route among that set, in terms of distance and number of navigation requests.

The routes from Control 1 capture local awareness, whereas those from Control 2 capture general awareness across the country. An example of a treated route and the corresponding route from Control 1 are shown in Fig. 3. Both routes have the same destination in Tel Aviv but different origins. The treated route goes through the NH HOV lane, but the controlled route does not.

Note that the identification of controlled routes does not use the carpool data. This ensures that the study is not biased towards carpooling behavior, which is the main purpose of our analysis.

It is possible that we could not find a good match for certain treated routes. If none of the routes have a distance and a number of navigation requests within 40% of the treated route, we do not select a corresponding controlled route. This situation occurred four times (two Control 1 routes are missing for HN, one Control 2 route is missing for NH, and one Control 2 route is missing for MK). As a result, we have a total of 21 treated routes and our 38 controlled routes.

3. Descriptive statistics

Table 1 reports the summary statistics on our 59 routes, aggregated across all three HOV lanes for each of the three periods. Since each route is defined at the cell-to-cell level, each value represents the average over a Xkm square cell. In addition to the three carpool variables (driver offers, rider offers, and completed carpools), we report the number of navigation requests (from Waze) and the ride value (price paid to the driver for the ride excluding bonuses and referrals). We normalize the number of navigation requests so that treated routes are assigned a value of 100%.

By design, the average number of navigations is similar across all three groups. We also observe that the carpool metrics increase substantially after the introduction of the HOV lanes. This suggests that the HOV lanes have a positive impact on carpooling among treated routes but also more broadly (we will revisit this idea in Section 4).

Fig. 4 plots the time saving for each HOV lane. This figure aggregates all navigations across the 21 treated routes defined in Section 2.3. We aggregate the soft and hard periods into an overall period following the introduction of the HOV lanes, and compute the difference between the average travel times (in minutes) with and without using the HOV lane. The estimated times (with and without the HOV lane) are calculated in real-time for each navigation request using Waze's data and algorithms. Note that the time saving is zero by definition before the introduction of HOV lanes and for controlled routes. In Fig. 4, we aggregate the observations using a 10-minute time window and plot the 40-th and 60th percentiles as a function of the time of the day.

Note that commuters can expect to save significant commuting times by using the HOV lanes, especially during rush hours. This was, in fact, the original goal of the policy. As mentioned before, NH and MK are mostly beneficial during morning hours, whereas HN is more beneficial during evening hours. Quantitatively, the NH, HN, and MK HOV lanes led to maximum median time savings

Table 1
Summary statistics. The table reports daily averages, and standard deviations in parentheses.

Period	Variable	Treated	Control 1	Control 2
Before	Driver offers	60.7 (6.93)	62.95 (9.24)	51.85 (9.64)
	Rider offers	81.98 (16.54)	94.45 (20.7)	74.94 (19.1)
	Completed carpools	6.45 (0.8)	6.69 (1.1)	7.09 (1.31)
	# navigations (normalized)	100%	98.78%	88.61%
	Ride value (Shekels)	9.12 (2.15)	8.28 (1.14)	9.04 (1.96)
Soft	Driver offers	215.75 (18.2)	105 (8.09)	114.25 (11.65)
	Rider offers	193.5 (12.52)	243.75 (40.56)	205.5 (27.81)
	Completed carpools	16.75 (1.2)	18.25 (2.16)	22.5 (1.99)
	# navigations (normalized)	112.53%	111.51%	95%
	Ride value (Shekels)	9.74 (3.01)	8.29 (1.02)	9.4 (2.37)
Hard	Driver offers	390.5 (22.23)	202.83 (16.29)	178.46 (17.32)
	Rider offers	286.54 (20.69)	239.25 (27.52)	222.71 (24.55)
	Completed carpools	22.96 (1.36)	22.92 (2.42)	26.06 (2.67)
	# navigations (normalized)	122.4%	125.42%	106.28%
	Ride value (Shekels)	8.76 (2.26)	8.32 (1.59)	9.16 (2.4)

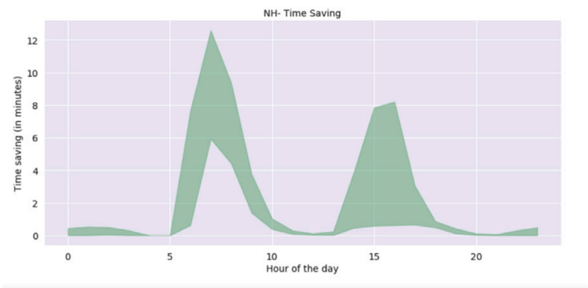

(a) NH HOV lane

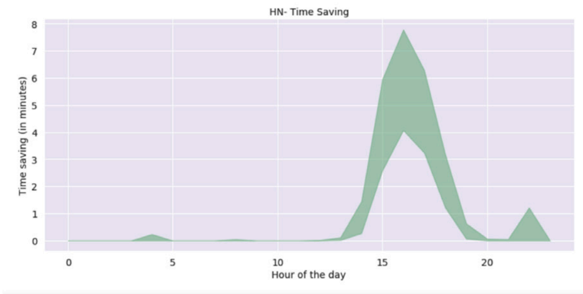

(b) HN HOV lane

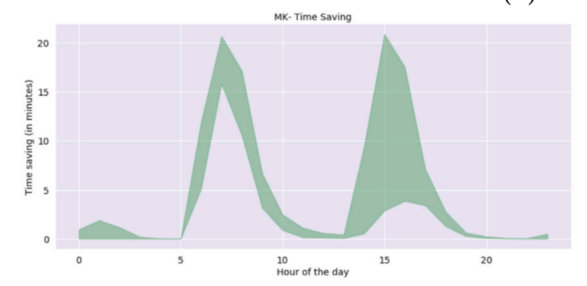

(c) MK HOV lane

Fig. 4. Average time saving across the day for the three HOV lanes.

of 8.1, 5.7, and 15.7 min, respectively, and a reduction in total commuting times spanning 20%–50%. To grasp the magnitude of these time savings, a 15-minute reduction on each commute, five days a week, would amount to saving five full days annually.

These observations suggest that the treatment is effective: the HOV lanes can enhance the travel experience for commuters who satisfy the occupancy requirement. Vice versa, commuters who cannot use the HOV lanes may experience some frustration associated with congestion and long commuting times. As a result, the discrepancy between high-occupancy vehicles and low-occupancy vehicles can potentially increase in carpooling demand. We study this question next.

We conclude by showing the relationship between travel time savings and carpooling behaviors. Fig. 5 plots the number of driver offers and completed carpools as a function of the time savings, across the 21 treated routes. In this figure, we normalize the numbers so that the highest number is assigned the value of 1.0, maintaining the same relationship between different points. Note the positive correlation between the travel time savings and carpool intent (i.e., driver offers) as well as carpool adoption (i.e., completed carpools). These findings are consistent with the results from (Cohen et al., 2019) using data from US commuters. To test the robustness of this finding, Table 2 reports several regression specifications by considering both dependent variables (driver offers and completed carpools), adding HOV lanes fixed effects, and controlling for several factors such as travel distance, the period (before, soft, hard), and the number of navigation requests.

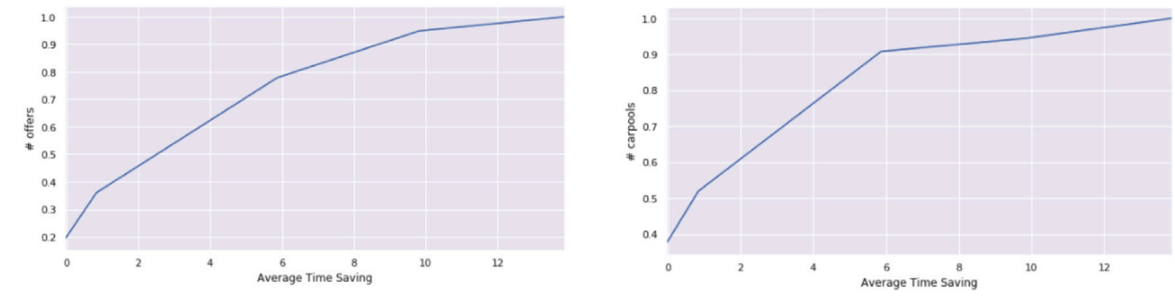

(a) Driver offers (b) Completed carpools

Fig. 5. Positive correlation between the time saving from the HOV lane and carpool intent and adoption.

Table 2
Regression estimates showing the positive correlation of the time saving with the carpool variables.

	Driver offers	Completed carpools
log(num_navs)	0.6329***	0.3016***
	(0.067)	(0.066)
log(median_km_suggested_with_HOV)	−2.1363***	0.1448
	(0.816)	(0.797)
log(median_km_suggested_without_HOV)	−0.7726	−1.2884
	(0.818)	(0.799)
avg_time_saving	0.0216**	0.0161*
	(0.009)	(0.009)
MK	−2.5201***	−1.6838***
	(0.234)	(0.229)
NH	−0.8812***	−1.1154***
	(0.188)	(0.184)
Hard	4.6642***	1.9281***
	(0.294)	(0.288)
Soft	3.8606***	1.6748***
	(0.302)	(0.296)
Constant	8.5248***	3.6029***
	(0.567)	(0.554)
Observations	990	990
R^2	0.359	0.125
F Statistic	78.40*** (df = 7;982)	19.98*** (df = 7;982)

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

4. Results

We break down the analysis into four parts. First, we present an aggregated analysis on the impact of the three HOV lanes in the entire country. Second, we identify the impact of HOV lanes on the marketplace dynamics in the Waze platform, by comparing the *quick ratio* metric across all treated routes vs. all controlled routes. Third, we perform a disaggregate statistical analysis based on the 21 treated routes and 38 controlled routes. Finally, we conduct several robustness tests, based on alternative regression specifications and an alternative city-level aggregation.

4.1. Country-level analysis

Before presenting the results at the disaggregate level, we first investigate the growth rate of carpooling for the entire country during the period of interest. These descriptive analytics provide model-free evidence of the impact of the HOV lanes without any assumptions or any route selection. Fig. 6 plots the compounded growth rate and the overall growth rate of carpool rides in Israel between mid August 2019 and mid January 2020 (all numbers are relative to the starting point—August 18 2019). We separate the carpool rides into two categories: the rides that overlap with one of the three HOV lanes (labeled as "HOV routes") and all other rides ("Non-HOV routes"). The vertical lines indicate important events over the period of interest, including the soft enforcement period, the hard enforcement period, major holidays, and the beginning of the school year. Table 3 summarizes the average week-over-week growth before and after the introduction of the HOV lanes, for both HOV and non-HOV routes.

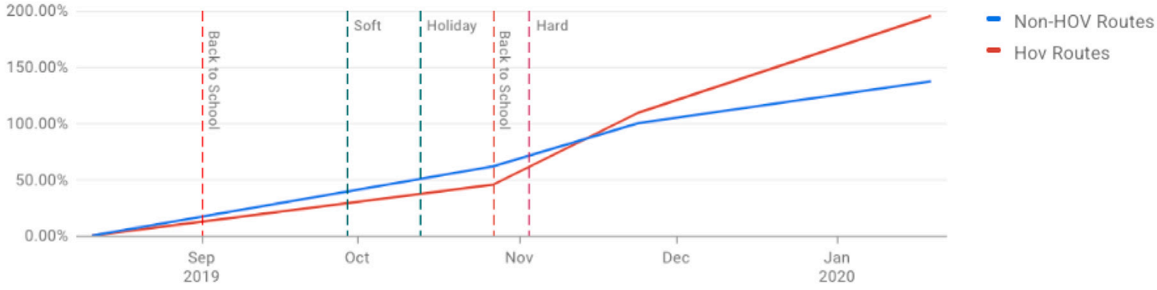

(a) Compounded growth rate

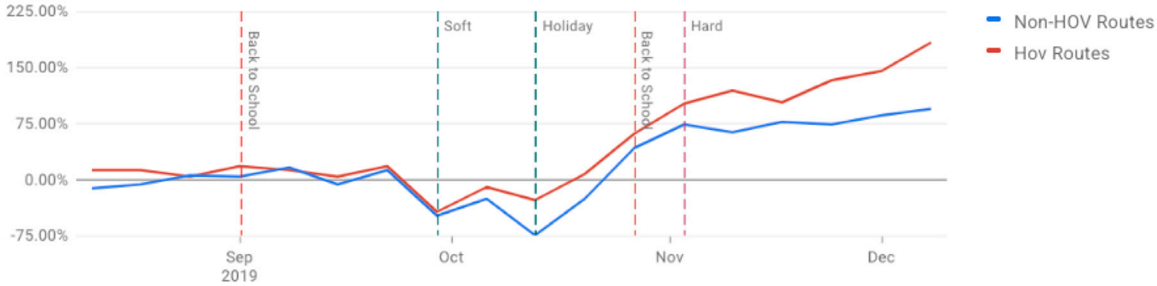

(b) Overall growth rate from first week

Fig. 6. Weekly growth rate for carpool rides in Israel during the period of interest.

Table 3
Week-over-week growth in carpool rides before and after the introduction of the HOV lanes.

	Before Sep 1–Nov 2	Right After Nov 3–Nov 30	Long term Dec 1–Jan 12
Non-HOV routes	2.68%	5.88%	3.55%
HOV routes	1.95%	10.38%	5.72%

The first observation is that carpooling has been growing significantly throughout the country during the period of observation. In fact, the number of carpools was increasing even before the introduction of the HOV lanes. This observation suggests that the Waze platform or carpooling behaviors (or both) were on the rise even before the introduction of HOV lanes.

A second observation is that the growth rate is much stronger on HOV routes than on non-HOV routes after the introduction of the HOV lanes. Before the introduction of the HOV lanes, non-HOV routes and HOV routes exhibit similar growth patterns—in fact, the overall growth rate for non-HOV routes is slightly higher than that for HOV routes. Right after the introduction of the HOV lanes, however, we observe a striking gap between the two types of routes. Between November 1, 2019 and December 1 2019, the slope of the growth for HOV routes becomes much steeper than the slope for non-HOV routes. Numerically, we find a relative increase of carpooling growth of 76.5% [=(10.38-5.88)/5.88] for HOV routes relative to non-HOV routes. In the longer term (i.e,. between December 1 2019 and January 12, 2020), the relative increase is 61.1%. Overall, these insults indicate that HOV lanes, and the resultant travel time savings, have a positive impact on carpooling behaviors. Note, interestingly, that this effect seems to happen at the beginning of the hard enforcement period, as opposed to the soft enforcement period. This can be explained by the stronger impact of hard HOV enforcement on carpooling behaviors, or by a lag in public awareness about HOV lanes or carpooling opportunities.

A third observation is that the introduction of HOV lanes seems to also induce more carpools on non-HOV routes. Even though the increase is less significant than on HOV routes, carpooling behaviors seem to also shift on non-HOV routes. In other words, HOV lanes have a strong and direct impact on carpooling behaviors on HOV routes, as well as a weaker and indirect impact on non-HOV routes. Therefore, by increasing awareness about carpooling opportunities locally, HOV lanes seem to also lead to a higher incidence of carpooling globally, even on routes they do not have a direct impact on travel times.

4.2. Marketplace dynamics

Before turning to the route-level analysis, we first investigate the impact of the HOV lanes on the marketplace dynamics within the Waze platform. One of the core challenges of two-sided platforms, known as the ‘‘cold-start’’ problem, stems from demand-supply mismatches when demand and/or supply is limited. Namely, matching riders with drivers traveling along similar origin–destination routes at similar times can present challenges if the rider pool is too small or if the driver pool is too small (or both). This problem is particularly salient for high-growth platforms like Waze and platforms dealing with emerging behaviors like carpooling. A key question is therefore whether the introduction of HOV lanes can alleviate this cold-start problem, by attracting new users to the platform and thus helping the marketplace grow to a critical mass.

We address this question by leveraging the *quick ratio* metric, defined as the ratio of new and resurrected users over the number of churned users. This metric originated in finance as a liquidity ratio to measure a company's propensity to pay its liabilities in the short term, and has been used in marketing to evaluate product growth. Specifically, at each point in time, the platform's users are divided into four categories: existing users (who were active before and during the period under consideration), new users (who joined the platform during the period under consideration), resurrected users (who were active at one point in time, became inactive, and became active again during the period under consideration), and churned users (who were active before but left the platform during the period under consideration). The quick ratio compares the number of new and resurrected users—an indicator of customer acquisition on the platform—with the number of churned users—an indicator of customer loss. The number of users is growing on the platform if the quick ratio is larger than one, and shrinking otherwise. As such, the quick ratio provides visibility into the dynamism of the platform, as a proxy for the propensity of the platform to avoid the cold-start problem in matching drivers and riders.

Fig. 7 investigates the impact of the introduction of the HOV lanes on the quick ratio, for the number of drivers placing an offer (Fig. 7(a)), the number of riders placing an offer (Fig. 7(b)), and the number of users involved in a carpool (Fig. 7(c)). Note, first, that the quick ratio is highest following the start of the school year and religious holidays, confirming that the metric captures changes in the customer mix. When it comes to the impact of HOV lanes, the main observation is that, following the soft and hard enforcement periods, the quick ratio is significantly larger among the treated routes than among the controlled routes—and significantly larger than one. In other words, the HOV policy seems to induce new users to use the carpooling platform, hence contributing to the scalability of the platform.

These plots provide a few additional insights on the marketplace dynamics following the introduction of HOV lanes. First, the impact is strongest for drivers placing on offer, with a quick ratio 3–5 times higher among treated routes after the soft enforcement period and 2–3 times higher after the hard enforcement period. In comparison, for riders placing an offer, the quick ratio remains similar between treated and controlled routes. This finding is consistent with the fact that HOV lanes serve primarily as a policy instrument to target drivers (e.g., commuters interested in saving time), as opposed to riders (e.g., commuters who do not own a car or otherwise need a ride). Second, the quick ratio increases among all controlled routes following the hard enforcement period. This result confirms the ripple effect of HOV lanes and the surrounding publicity campaigns on carpooling behaviors across the country, even on routes where the HOV lanes did not have a direct effect on travel times. Third, the temporal dynamics underscore the cold start problem and the role of policy interventions as a remedy. Namely, the soft enforcement period generated a surge of new drivers on the platform, but the mix of riders remained similar. As a result, the new driver offers did not translate into new carpools—as indicated by a quick ratio around one. After the hard enforcement period, in contrast, the platform witnessed an influx of new drivers as well as of new riders—which, again, may be due to the stronger enforcement or to the lagged impact of publicity campaigns. In sharp contrast with the soft enforcement period, these patterns led to carpools comprised of new users, reflected by a quick ratio of 6 on treated routes and of 3 on controlled routes. In summary, HOV lanes mainly serve as a supply-driven instrument; by itself, the resulting surge in supply did not lead to significant changes in the carpooling mix; however, HOV lanes also generated activity on the demand side of the platform, ultimately creating new successful matches and thus alleviating the cold-start problem on the platform.

4.3. Route-level analysis

We now compare carpooling outcomes across the 21 treated routes and 38 controlled routes, across the three periods of interest. Let us proceed one HOV lane at a time. For robustness purposes, we repeat all analyses using bootstrapping, by randomly sampling randomly routes (with replacement) across 100 independent iterations. We obtain the same qualitative results and insights.

4.3.1. Netanya–Herzliya (NH):

The NH HOV lane is primarily designed for morning commuters. Fig. 8 compares the number of driver offers, the number of rider offers, and the number of completed carpools across the three periods (before, soft, and hard). Each variable is aggregated across the seven treated routes corresponding to NH and the associated routes in Control 1 and Control 2. Since each route is defined at the cell level (starting in a Xkm square cell and ending in another cell), each variable represents the average over a Xkm square cell.

Note that the introduction of HOV lanes leads to a 100% increase in the number of offers sent by the drivers across the treated routes. This increase induces a sharp increase in the number of completed carpools as well, by an estimated 900% (this number is also driven by the relatively low proportion of carpools across treated routes in the NH HOV lane, prior to the treatment). These two metrics also increase for the controlled routes. Such an increase can be potentially attributed to growth and seasonality, as well

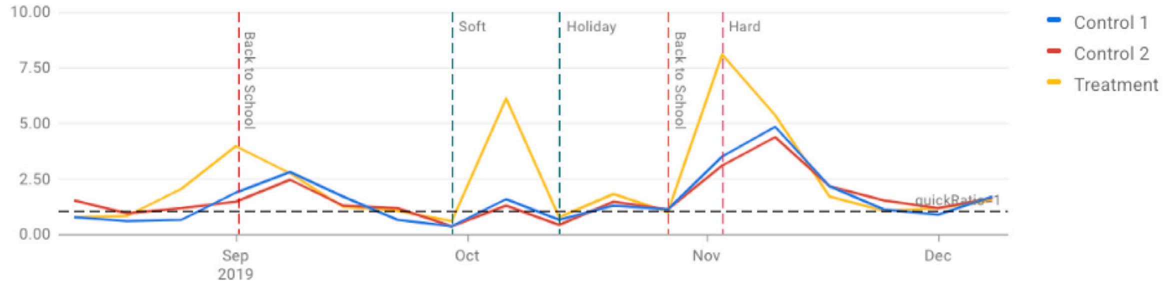

(a) Drivers

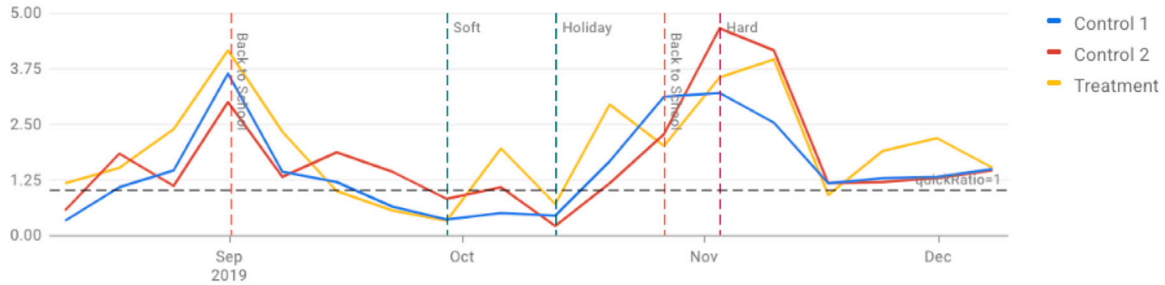

(b) Riders

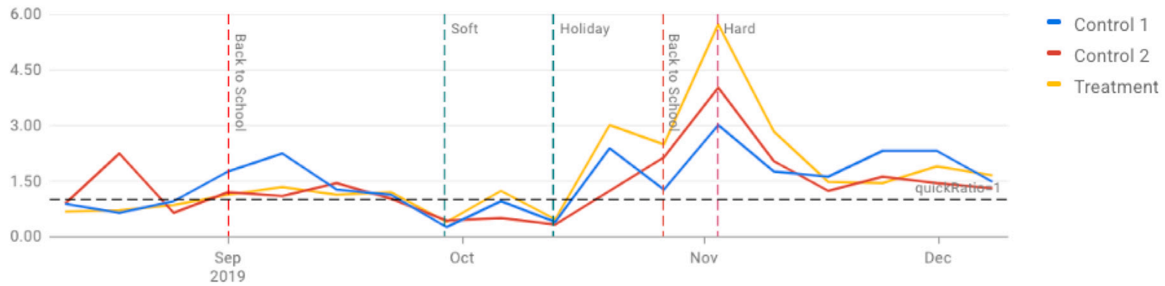

(c) Carpools

Fig. 7. Quick ratio for the number of drivers, riders, and carpools across the HOV routes ("treatment") and non-HOV routes ("Control 1" and "Control 2").

as an increased level of public awareness about carpooling—confirming the ripple effect of HOV lanes on routes where they do not have a direct impact on travel times. Nonetheless, the increase is much higher in relative terms (with respect to the period before the policy change) for the treated routes.[7] Together, these results suggest a positive impact of the HOV lanes on carpool intent (captured by the number of offers sent by the drivers) and carpooling adoption (captured by the number of completed carpools). Furthermore, we find that the number of offers sent by the riders does not significantly increase, confirming that HOV lanes are primarily a supply-driven instrument.

[7] Note that the number of carpools was smaller in treated routes than in controlled routes in the pre-experimental period. It was not possible to design treatment and control groups with similar features throughout the entire carpooling funnel, while also satisfying all constraints (e.g., similar distance, similar number of navigation requests on Waze). As the figure shows, the control routes are similar to the treated routes at the upstream level of the funnel, with similar numbers of offers from both drivers and riders. However, these did not translate into similar numbers of completed carpools across the three groups for the NH HOV lane (it did for the other two HOV lanes).

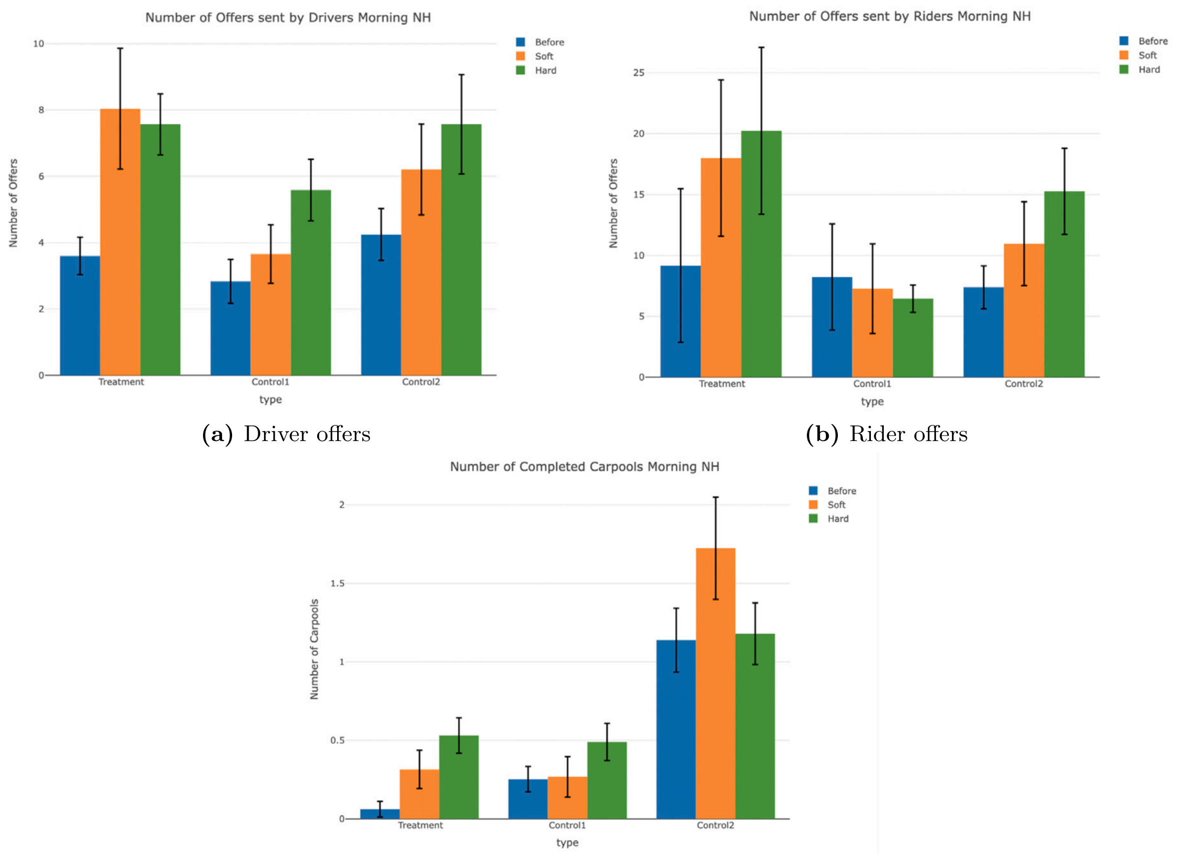

(a) Driver offers **(b)** Rider offers

(c) Completed carpools

Fig. 8. Comparing the number of driver offers, rider offers, and completed carpools for NH during morning hours (averaged over a Xkm square cell).

4.3.2. Herzliya–Netanya (HN):

For HN, the impact is expected to occur during evening hours. Fig. 9 reports the corresponding results, using the same nomenclature as Fig. 8. Again, after introducing the HOV lane, the number of offers sent by drivers increases by 200% and the number of completed carpools raises by 100% for the treated routes. These metrics also increase for the second type of controlled routes. However, once again, the increase is much higher in relative terms (with respect to the period before the policy change) for the treated routes. Interestingly, we find that the impact for HN routes is lower relative to NH. This is not surprising since it is typically harder to find a carpool back from work to home. This can be explained by two factors. One is that schedules are less predictable during evening hours, as compared to morning hours when most people have a pre-determined start time. Another is that commuters have more time and more flexibility to arrange a carpool the night before a morning commute (when they are typically at home) than the day before an evening commute (when they are typically at work).

4.3.3. Mavo Ayalon–Kibbutz Galuyot (MK):

For MK, the impact is expected during morning hours. Fig. 10 shows the corresponding results, again using the same nomenclature as Fig. 8. In this case, the number of offers sent by the drivers increases by 200% after the hard enforcement period for the treated routes (it does not significantly increase during the soft enforcement). However, we find that the number of completed carpools is not significantly affected by the HOV lane.

This result draws a sharp distinction between the two previous HOV lanes and the MK one, and can be explained by two factors. First, recall that the NH and HN HOV lanes complement each other, by running in opposite directions. As a result, it may be easier for commuters to arrange carpools for round-trip commutes, with clear benefits in both directions. In contrast, the MK HOV lane is one-directional, so riders may face the risk of finding a carpool in the morning but not for the trip back in the evening (this is particularly relevant because the Waze platform only offers one-way carpools, as opposed to round-trip carpools). Stated differently, the one-way HOV lane leads to risk and uncertainty for the rider, which contributes to creating barriers towards carpooling adoption. Second, the MK HOV lane has a 3+ occupancy requirement, as opposed to a 2+ occupancy requirement for the other two. As a result, to use the HOV lane, drivers would have to find two riders, as opposed to a single one. From a user perspective, this higher

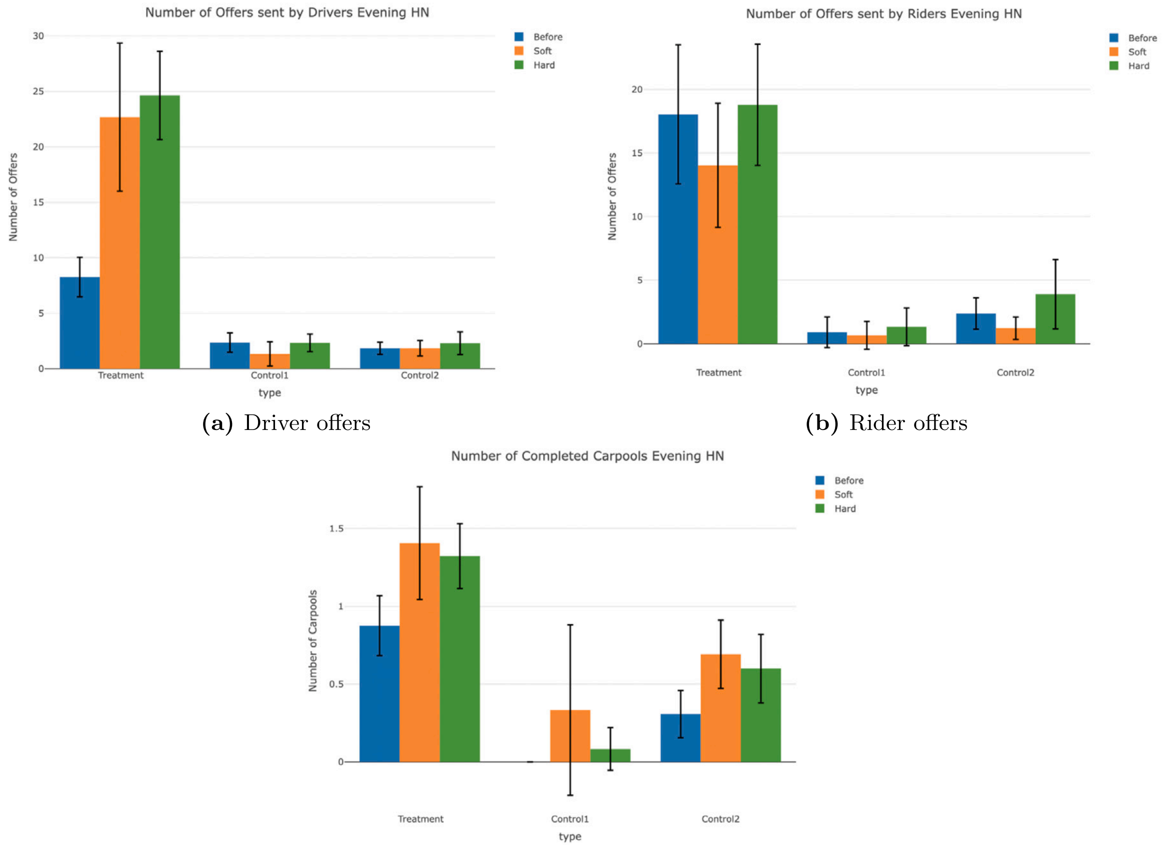

(a) Driver offers

(b) Rider offers

(c) Completed carpools

Fig. 9. Comparing the number of driver offers, rider offers, and completed carpools for HN during evening hours (averaged over a Xkm square cell).

occupancy requirement increases the transaction costs and decreases the matching potential on the platform. As a result, fewer drivers are successful in finding appropriate matches on the platform.

A consequence of this discussion is that drivers are more likely to cancel a carpool on the treated routes, if they can find one rider but not two riders. This explanation is supported by Fig. 11, which shows that the number of cancellations is higher during the hard period for the treated routes. In contrast, the difference in the number of cancellations before the policy and after the hard enforcement is not statistically significant for the two other HOV lanes (NH and HN).

In conclusion, the one-way 3+ HOV lane is less impactful than the round-trip 2+ HOV lanes in terms of completed carpools. While the carpool intent did increase, we do not observe a statistically significant effect on the completed carpools. This is especially interesting given that the median time saving for MK was the highest—in fact, the one-way time savings induced by the MK HOV lane could be even larger than the round-trip time savings induced by the NH and HN HOV lanes combined (see Section 3). Ultimately, these results show that the design of HOV lanes (e.g., 2+ vs. 3+ occupancy requirement, one-way vs. round-trip HOV lane), the design of carpooling platforms (e.g., one-way vs. round-trip features, information sharing on travel time savings) and the market dynamics (e.g., balancedness of drivers and riders) play an important role to translate an increased number of offers from the drivers into a higher number of completed carpools in actuality. In our case, the routes overlapping with MK do not seem to have enough rider demand, so the increased carpooling intent on the driver side does not driver higher carpool adoption.

4.4. Regressions

We test the robustness of the previous results by estimating a regression model, as opposed to simply running a two-sided t-test. This regression analysis controls for various factors. Namely, we consider a log-transformed ordinary least square specification:

$$\log(Y_t^k + 1) = \alpha \text{period\_type} + \beta \text{route\_type} + \gamma \text{period\_type} \times \text{route\_type} + \delta \text{Controls} + \epsilon_t^k, \tag{1}$$

where each observation is at the day-route (t, k) level and Y_t^k is one of our carpool dependent variables (driver offers, rider offers, and completed carpools) on day t and route k. We then include three independent variables: (1) period_type, which determines the period (before, soft, hard) of observation (t, k), (2) route_type which captures the route type (treated, Control 1, Control 2), and (3)

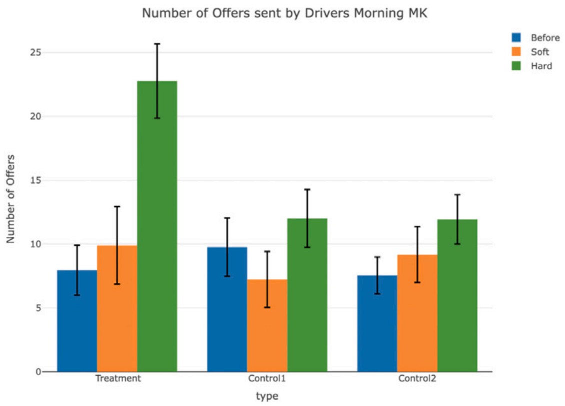

(a) Driver offers

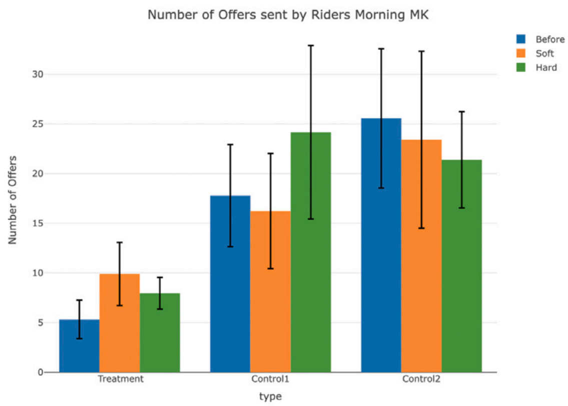

(b) Rider offers

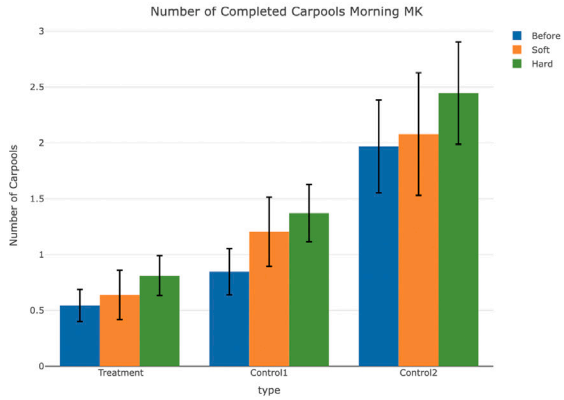

(c) Completed carpools

Fig. 10. Comparing the number of driver offers, rider offers, and completed carpools for MK during morning hours (averaged over a Xkm square cell).

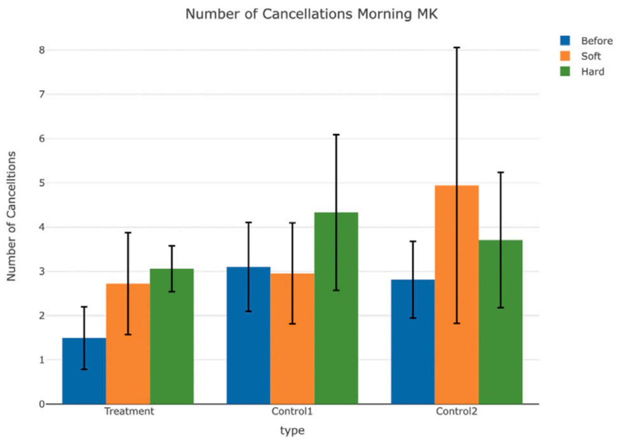

Fig. 11. Total number of canceled rides for MK during the morning hours (averaged over a Xkm square cell).

the interaction between these two variables. We add a vector of control variables, capturing both route-specific features (e.g., travel distance in kilometers, time savings from the HOV lane in minutes, number of navigation requests on Waze) and temporal features (e.g., day of the week). Finally, ϵ_i^k is a stochastic Gaussian term, reflective of idiosyncratic noise. We set the baseline of this regression as the period before introducing the HOV lanes and for the routes in Control 1.

Table 4
Regression estimates of the interaction coefficient γ between treatment and the hard period.

	Driver offers		Rider offers		Completed Carpools	
	(1)	(2)	(1)	(2)	(1)	(2)
NH Morning	0.3584***	0.2044**	0.4314**	0.4269***	0.1594**	0.0487
	(0.1201)	(0.1033)	(0.1702)	(0.1613)	(0.0636)	(0.0584)
HN Evening	0.4165***	0.3457***	0.6006***	0.4740***	0.2398***	0.2445***
	(0.1336)	(0.1167)	(0.159)	(0.1532)	(0.064)	(0.0573)
MK Morning	0.6923***	1.2437***	−0.1259	−0.127	0.0327	0.2212***
	(0.1393)	(0.1554)	(0.1619)	(0.2132)	(0.0769)	(0.0846)

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.
(1): specification without controls; (2): specification with controls.

Table 5
Regression estimates for NH during morning hours.

	Driver offers		Rider offers		Completed Carpools	
	(1)	(2)	(1)	(2)	(1)	(2)
Constant	1.5962***	1.8318***	1.3703***	1.903***	0.4512***	0.4605***
	(0.053)	(0.0769)	(0.0752)	(0.1201)	(0.0281)	(0.0435)
Hard	0.2933***	0.4553***	0.4957***	0.1117	0.0992**	0.1526***
	(0.0771)	(0.1016)	(0.1092)	(0.1586)	(0.0408)	(0.0575)
Soft	−0.0324	0.0771	0.3335	0.2217	0.0075	0.132
	(0.1573)	(0.1791)	(0.223)	(0.2796)	(0.0833)	(0.1013)
Control 2	−0.1772**	−0.2787***	−0.1405	−0.2462**	0.036	0.0684*
	(0.0835)	(0.0733)	(0.1183)	(0.1144)	(0.0442)	(0.0415)
Treatment	−0.302***	−0.4774***	−0.5928***	−0.6739***	−0.3907***	−0.3684***
	(0.0855)	(0.0731)	(0.1212)	(0.1141)	(0.0453)	(0.0414)
Hard ×Control 2	−0.1019	0.1975*	0.1874	0.8064***	0.0655	0.2046***
	(0.1185)	(0.1127)	(0.168)	(0.176)	(0.0627)	(0.0638)
Soft ×Control 2	0.0664	0.4093	0.3372	0.69*	0.2461*	0.4218***
	(0.2455)	(0.2585)	(0.3479)	(0.4037)	(0.1299)	(0.1463)
Hard ×Treatment	0.3584***	0.2044**	0.4314**	0.4269***	0.1594**	0.0487
	(0.1201)	(0.1033)	(0.1702)	(0.1613)	(0.0636)	(0.0584)
Soft ×Treatment	0.5632**	0.5035**	0.8407***	0.7809**	0.1786	0.0787
	(0.2251)	(0.2126)	(0.319)	(0.3319)	(0.1191)	(0.1203)
Monday	–	0.0302	–	−0.1207	–	−0.0019
	–	(0.0658)	–	(0.1027)	–	(0.0372)
Tuesday	–	−0.0345	–	−0.2398**	–	−0.0696*
	–	(0.0647)	–	(0.1011)	–	(0.0366)
Wednesday	–	0.0563	–	−0.2862***	–	0.0142
	–	(0.0654)	–	(0.1022)	–	(0.037)
Thursday	–	−0.0099	–	−0.2703***	–	−0.0164
	–	(0.0663)	–	(0.1036)	–	(0.0375)
log(Distance)	–	0.0287	–	−0.1557***	–	−0.0395**
	–	(0.0343)	–	(0.0536)	–	(0.0194)
log(Time savings)	–	0.0879**	–	0.1785***	–	0.1064***
	–	(0.0407)	–	(0.0635)	–	(0.023)
log(Navigations)	–	−0.0323	–	0.1967***	–	0.0212
	–	(0.0281)	–	(0.0439)	–	(0.0159)
Observations	1550	1550	1550	1550	1550	1550
R^2	0.0526	0.3669	0.0752	0.2535	0.1151	0.3127
F Statistic	10.70***	37.75***	15.65***	22.12***	25.05***	29.63***

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.
(1): specification without controls; (2): specification with controls.

We are especially interested in the interaction coefficient γ between the treated routes and the hard period, reported in Table 4, indicating the different in the slope of carpooling intent and carpooling adoption between controlled and treated routes. The complete regression tables can be found in Tables 5–7 for the three HOV lanes, without and with control variables.

First and foremost, we note that the coefficient corresponding to the interaction term of the hard enforcement period and the treatment group is positive and statistically significant for seven out of nine specifications. This suggests that carpool intent and

Table 6
Regression estimates for HN during evening hours.

	Driver offers		Rider offers		Completed Carpools	
	(1)	(2)	(1)	(2)	(1)	(2)
Constant	1.156***	0.6663***	0.6267***	0.5275***	0.246***	0.0663
	(0.0885)	(0.1066)	(0.1053)	(0.14)	(0.0424)	(0.0524)
Hard	0.3390***	0.7753***	−0.1386	−0.1564	−0.1201**	−0.0614
	(0.1158)	(0.1229)	(0.1378)	(0.1614)	(0.0555)	(0.0604)
Soft	0.5700***	0.8839***	0.2091	0.2417	0.0511	0.0907
	(0.2062)	(0.1885)	(0.2454)	(0.2476)	(0.0988)	(0.0926)
Control 2	−0.1426	0.3324***	−0.0108	0.2895**	−0.0218	0.0736
	(0.1146)	(0.1032)	(0.1364)	(0.1355)	(0.0549)	(0.0507)
Treatment	0.3960***	0.3791***	0.9661***	0.8871***	0.2219***	0.1844***
	(0.0996)	(0.0876)	(0.1186)	(0.1151)	(0.0478)	(0.043)
Hard ×Control 2	−0.2929*	−0.5317***	0.1226	0.2687	0.1354*	0.1115
	(0.1567)	(0.1404)	(0.1866)	(0.1844)	(0.0751)	(0.069)
Soft ×Control 2	−0.6750**	−0.6932**	0.1445	0.2034	0.2479	0.1516
	(0.319)	(0.2691)	(0.3797)	(0.3534)	(0.1529)	(0.1322)
Hard ×Treatment	0.4165***	0.3457***	0.6006***	0.4740***	0.2398***	0.2445***
	(0.1336)	(0.1167)	(0.159)	(0.1532)	(0.064)	(0.0573)
Soft ×Treatment	−0.1706	−0.191	−0.0932	−0.3463	0.0463	0.0436
	(0.2445)	(0.2119)	(0.291)	(0.2783)	(0.1172)	(0.1041)
Monday	–	−0.0085	–	0.0679	–	0.0655**
	–	(0.0651)	–	(0.0855)	–	(0.032)
Tuesday	–	−0.104	–	−0.052	–	0.0611*
	–	(0.0647)	–	(0.0849)	–	(0.0318)
Wednesday	–	0.0797	–	−0.1058	–	0.0365
	–	(0.0657)	–	(0.0863)	–	(0.0323)
Thursday	–	0.0446	–	−0.1977**	–	0.0232
	–	(0.0665)	–	(0.0873)	–	(0.0327)
log(Distance)	–	−0.0565*	–	−0.0681	–	0.0252
	–	(0.0339)	–	(0.0445)	–	(0.0166)
log(Time savings)	–	−0.0945	–	0.1703**	–	−0.0086
	–	(0.0645)	–	(0.0847)	–	(0.0317)
log(Navigations)	–	−0.015	–	0.0689**	–	−0.0295***
	–	(0.0218)	–	(0.0286)	–	(0.0107)
Observations	1636	1636	1636	1636	1636	1636
R^2	0.1871	0.4367	0.2204	0.3407	0.1103	0.3519
F Statistic	46.80***	59.13***	57.49***	39.42***	25.20***	41.42***

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.
(1): specification without controls; (2): specification with controls.

adoption generally grow faster (from the pre-treatment period to the hard enforcement period) in treated routes than in controlled routes. In fact, the impact on the number of driver offers is consistently positive and statistically significant. However, for the MK HOV lane, the impact on rider offers and on completed carpools are weaker relative to the NH and HN HOV lanes. This quantitative result corroborates the discussion from Section 4.3. These results are robust for both specifications, without and with controls. As expected, adding the control variables leads to a stronger model fit, as indicated by the higher R^2 values—spanning 0.25–0.55. It also provides some insights on the drivers of carpooling behaviors; for instance, carpooling seems to decrease with the travel distance (the inconvenience of carpooling is smaller on shorter trips) and to increase with the number of navigations on the Waze platform (Waze active users are more likely to be prone to carpooling). That said, we reiterate that our primary finding from these regressions is the robustness of our main insight—namely, that HOV lanes increase carpooling intent and adoption.

4.5. City-level analysis

Instead of focusing on routes defined at the cell level (i.e., navigation requests from one Xkm square cell to another), one can instead conduct an analysis at the city level. Accordingly, we now define route between an origin city and a destination city (e.g., Haifa to Tel Aviv). We focus on the most popular routes by restricting the daily number of navigation requests (using Waze) to be at least 600. This restriction led us to the 35 most popular city-to-city routes in the entire country. We then use the data from the period after the introduction of the HOV lanes to compute the average time saving (by using the HOV lane) for each of the 35

Table 7
Regression estimates for MK during morning hours.

	Driver offers		Rider offers		Completed Carpools	
	(1)	(2)	(1)	(2)	(1)	(2)
Constant	1.8471***	1.979***	1.7421***	1.2689***	0.5364***	0.223***
	(0.0623)	(0.0775)	(0.0724)	(0.1063)	(0.0344)	(0.0422)
Hard	0.3293***	0.2281*	0.6467***	0.5798***	0.1515***	0.1028
	(0.0978)	(0.1358)	(0.1136)	(0.1863)	(0.054)	(0.0739)
Soft	−0.0606	−0.0256	0.5049**	0.5358	0.1408	0.3327**
	(0.1984)	(0.2617)	(0.2306)	(0.3589)	(0.1095)	(0.1425)
Control 2	0.0463	0.1524**	0.1741*	0.2354***	0.2094***	0.2123***
	(0.0845)	(0.0596)	(0.0983)	(0.0818)	(0.0467)	(0.0325)
Treatment	−0.0853	−0.1979***	−0.6218***	−0.7542***	−0.1593***	−0.3134***
	(0.0955)	(0.0676)	(0.111)	(0.0927)	(0.0527)	(0.0368)
Hard ×Control 2	0.1313	0.0854	0.0915	−0.0294	0.096	0.0111
	(0.1315)	(0.1502)	(0.1528)	(0.206)	(0.0726)	(0.0818)
Soft ×Control 2	0.1759	−0.271	−0.0715	−0.1381	−0.0526	−0.2911*
	(0.2621)	(0.3123)	(0.3046)	(0.4284)	(0.1447)	(0.1701)
Hard ×Treatment	0.6923***	1.2437***	−0.1259	−0.127	0.0327	0.2212***
	(0.1393)	(0.1554)	(0.1619)	(0.2132)	(0.0769)	(0.0846)
Soft ×Treatment	0.1732	0.257	−0.1124	−0.732*	−0.1094	−0.2728*
	(0.2726)	(0.2998)	(0.3168)	(0.4113)	(0.1505)	(0.1633)
Monday	–	0.011	–	0.0218	–	0.0177
	–	(0.0652)	–	(0.0895)	–	(0.0355)
Tuesday	–	0.0753	–	−0.1322	–	−0.0688*
	–	(0.0656)	–	(0.09)	–	(0.0357)
Wednesday	–	0.0449	–	−0.1468	–	−0.0002
	–	(0.0658)	–	(0.0902)	–	(0.0358)
Thursday	–	0.0494	–	−0.0435	–	−0.0525
	–	(0.065)	–	(0.0891)	–	(0.0354)
log(Distance)	–	−0.2531***	–	−0.0406	–	−0.0961**
	–	(0.0779)	–	(0.1069)	–	(0.0424)
log(Time savings)	–	−0.0798*	–	0.0841	–	0.0174
	–	(0.0478)	–	(0.0656)	–	(0.0261)
log(Navigations)	–	0.1065***	–	0.0619	–	0.0528**
	–	(0.0379)	–	(0.052)	–	(0.0207)
Observations	1880	1880	1880	1880	1880	1880
R^2	0.0846	0.5442	0.1131	0.3805	0.0834	0.4991
F Statistic	21.62***	80.68***	29.83***	41.51***	21.29***	67.33***

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.
(1): specification without controls; (2): specification with controls.

routes. For 12 out of 35 routes, we find a time saving higher than 3 min. We then assign these 12 routes to the treatment condition and the remaining 23 routes to the control condition. Using the routes defined at the city level, we repeat the same analysis as before (while bootstrapping over 100 iterations) and find that most of our qualitative results and insights still hold (we omit the details for conciseness).

5. Conclusion and policy implications

This paper conducted a data-driven evaluation of the impact of HOV lanes on carpooling. To this end, we leveraged a partnership with Waze, a major carpooling platform. Carpooling platforms offer two major opportunities to ease barriers associated with carpooling usage and assessment. First, by providing visibility into users' travel patterns and by optimizing matches using analytics, carpooling platforms alleviate the coordination costs associated with carpooling, hence making it a more attractive commuting option. Second, carpooling platforms provide passenger-level data that carry information about vehicle occupancy, as opposed to vehicle-level information from traffic sensors. These two features enable us to report the first piece of data-driven evidence on the positive impact of HOV lanes on carpooling behaviors, using revealed-preferences information.

Specifically, we exploited an empirical setting leveraging the recent introduction of three HOV lanes in Israel, casting this event as a natural experiment. We used carpool transactions from Waze before and after the introduction of HOV lanes to identify their impact on carpooling intent (number of carpool offers) and adoption (number of completed carpools). We developed treatment and control groups at the route level, which enabled us to perform a direct comparison of carpooling behaviors between routes that were directly impacted by the HOV lanes and those that were not. We identified a clear impact of the HOV lanes on travel times, cutting commuting times by 20%–50% on average for carpools—hence, confirming that the treatment actually took place.

We derived four main results. First, HOV lanes bring new users onto the platform—new drivers at first, and new riders later on. As such, HOV lanes contribute to scaling the carpooling marketplace and thus to alleviating the cold-start problem. Second, the HOV lanes have a positive impact on carpool intent: the number of offers sent by drivers increase significantly following the introduction of the HOV lanes to take advantage of the shorter travel times. The number of rider offers does not exhibit the same increase, consistent with the fact that the HOV policy is primarily targeted to drivers. Third, the HOV lanes have disparate impacts on carpool adoption: completed carpools increase significantly for two out of three HOV lanes, but not for the third one. This can be explained by two particularities of the third HOV lane: (i) it is one-directional (hence, it places riders at risk of not finding a carpool on the way back), and (ii) it has a three-passenger occupancy requirement, as opposed to a two-passenger requirement (hence, it increases the coordination costs of carpooling to take advantage of the HOV lane). Fourth, the HOV lanes have indirect effects on carpooling, even on routes where they do not directly lead to travel time savings.

Obviously, this study is not without limitations. Namely, our carpool data is only a partial representation of the entire picture. Some commuters may carpool without using the Waze platform, which falls beyond the scope of our analysis. Moreover, our analysis focused on the carpooling impact of HOV lanes, thus leaving a more comprehensive assessment of HOV lanes for future research (which would require bringing together all commuters, including carpoolers, solo drivers, and public transportation users). Finally, our analysis focused on the short-term impact, but it is also important to consider the longer-term effects of HOV lanes.

Policy implications. Yet, our paper has important policy implications on the planning, design, and operations of HOV lanes, which we summarize below:

– *HOV lanes increase carpool intent and can also increase carpool adoption.* For decades, HOV lanes have been part of governments' arsenals in the fight against traffic congestion. Over that period of time, they have generated controversy and debates among researchers, policy makers, and the general public regarding their effectiveness. Most of the available evidence focused on their overall level of utilization and their impact on traffic congestion. However, a critical piece of the puzzle has been missing—the impact of HOV lanes on commuters' carpooling behaviors. This paper leverages new platform-based data to provide a novel piece of evidence in response to this question.
  It is important to note, however, that this evidence was gleaned in an environment with a very high penetration of Waze. As such, the Israeli context indicates both a high propensity for carpooling (due to the high proportion of commuters on the platform) and low carpooling coordination costs (due to the high likelihood of finding a match on the platform). This observation raises the question of the generalizability of the results. First, HOV lanes may have an even stronger impact on carpooling behavior in countries where carpooling levels are low to start with, although these may not translate into the same impact on congestion. Second, HOV lanes may have a weaker impact on carpooling behavior in countries with a lower penetration of carpooling platforms—thus also indicating the role of mobility platforms to strengthen the impact of congestion mitigation policies, as discussed below.
– *HOV lanes have a broader impact on carpooling behaviors.* Indeed, HOV lanes can have positive ripple effects on carpooling in other parts of a jurisdiction. In our case, the three HOV lanes raised the carpool rates on routes that they were not directly affected by the HOV lanes. This boost can be explained by the increased awareness of the public about the opportunity of carpooling, either organically or through dedicated advertisement campaigns.
– *All HOV lanes are not equal.* We found that carpooling adoption critically depends on the design of the HOV lane. Specifically, we show that two HOV lanes had a stronger impact on carpool adoption than the third one. This result can be interpreted in two ways: (i) round-trip HOV lanes are more effective than one-way HOV lanes in driving carpooling adoption, and (ii) less stringent occupancy requirements also lead to a stronger carpooling adoption. These findings can be embedded into more comprehensive evaluations of alternative HOV designs that would compare the required investments for one-way vs. round-trip HOV lanes, evaluate the impact of occupancy requirements on the number of vehicles on the road, and assess the impact of HOV lanes on travel times for non-carpoolers.
– *Market dynamics play an important role in shaping carpooling adoption.* Carpool intent does not necessarily translate into carpool adoption. In our case, we found that all three HOV lanes increased carpool intent, but only two of them led to stronger carpool adoption. The latter featured lower rider demand, higher matching costs (due to the higher-occupancy requirements) and ultimately higher cancellations from drivers. This underscores that the balancedness of the market – between drivers and riders – is an important factor underlying the effectiveness of HOV lanes and their impact on carpooling behaviors.
– *HOV lanes can improve the scalability of carpooling marketplaces.* HOV lanes are a supply-driven policy instruments. Unsurprisingly, we find that they bring new drivers to the carpooling platform. Yet, the surge in supply does not induce significant changes in the carpooling mix by itself. Over time, however, HOV lanes make the demand side of the platform more liquid, which help generate new carpools. As such, HOV lanes can therefore contribute to alleviating the cold-start problem on carpooling marketplaces.
– *The design of on-demand carpooling platforms also plays an important role in boosting carpooling and further enhancing the effectiveness of HOV lanes.* Access to real-time data and efficient matching platforms can allow commuters to find attractive carpooling options. According to Waze data, two thirds of drivers have at least one other regular Waze driver with a "perfectly matching commute", that is, driving from the same origin to the same destination within less than 500 meters radius from each. At the same time, this untapped pool of potential carpoolers also presents opportunities to design features that would drive

up carpooling penetration, especially following the introduction of HOV lanes. For instance, including a round-trip carpooling feature could alleviate the friction due to one-directional HOV lanes. Similarly, quantifying and communicating the travel time savings associated with the use of the HOV lane could further increase demand—especially rider demand.

In addition, on-demand carpooling operations offer opportunities to enhance the coordination of driver and rider operations, once a match has been determined. In the absence of carpooling platforms, drivers and riders usually plan carpools offline, or seek matches online in designated pickup locations (see Beroldo, 1990; Mote and Whitestone, 2011, for a discussion on slugging). By aggregating mobility data across a large pool of users, carpooling platforms can instead optimize pickup and dropoff locations to make carpools more convenient for both riders and drivers (Ma and Wolfson, 2013; Boysen et al., 2021). This mirrors vehicle–customer coordination in ride-sharing (Zhang et al., 2022) and micro-transit (Veve and Chiabaut, 2022). Mobility platforms thus provide a dedicated technology to systematically optimize carpooling operations, which can ultimately further spur carpooling adoption (Martin et al., 2021).

– *There exist important collaboration opportunities between governmental agencies and on-demand platforms to enhance the design and operations of HOV lanes.* At the strategic level, HOV lanes require extensive infrastructure investments, with unclear impacts on downstream traffic patterns. In this regard, carpooling platforms can provide data sources to identify promising locations for HOV lanes towards carpooling adoption and congestion mitigation, and to guide key decisions regarding occupancy requirements. At the tactical level, the impact of HOV lanes remains impeded by frictions in carpooling markets. As such, policy makers and on-demand platforms could work together to incentivize commuters to use HOV lanes, via carpooling. One opportunity would be to share information on travel time savings resulting from the use of the HOV lanes. Vice versa, HOV lanes provide stronger incentives for travelers to find carpools, therefore increasing demand for carpooling platforms—especially in countries with a lower penetration of such platforms than Israel. Another opportunity would involve designing incentive schemes to boost carpooling adoption. For instance, HOV lanes primarily incentivize one side of the market (drivers), so it could be effective to combine this policy with rider-side coupons to alleviate market asymmetries. Yet another opportunity lies in the design of dedicated public transit options to complement one-directional HOV lanes, so riders could find carpools for their morning commutes (from home to work) and would have access to an attractive travel option for their evening commutes (from work to home).

This paper comes at a time where traffic congestion and greenhouse gas emissions remain foremost policy concerns. Our results show that strategic infrastructure investments (HOV lanes, in our case) coupled with enhanced technology (carpooling platforms, in our case) can align commuting behaviors with broader societal goals. As such, this paper contributes to growing evidence on new policy instruments to alleviate traffic congestion and enhance mobility at scale.

CRediT authorship contribution statement

Maxime C. Cohen: Conceptualization, Methodology, Project administration, Writing – original draft, Writing – review & editing. **Alexandre Jacquillat:** Conceptualization, Methodology, Writing – original draft, Writing – review & editing. **Avia Ratzon:** Investigation, Software. **Roy Sasson:** Conceptualization, Supervision.

Data availability

The data that has been used is confidential.

Acknowledgment

We thank Yotam Givati for excellent research assistance on data collection and analysis.

References

Anderson, P., Geroliminis, N., 2020. Dynamic lane restrictions on congested arterials. Transp. Res. A 135, 224–243.
Ben-Akiva, M., Atherton, T.J., 1977a. Choice-model predictions of car-pool demand: Methods and results. Transp. Res. Rec. (637).
Ben-Akiva, M., Atherton, T.J., 1977b. Methodology for short-range travel demand predictions: Analysis of carpooling incentives. J. Transp. Econ. Policy 224–261.
Beroldo, S., 1990. Casual carpooling in the san Francisco Bay Area. Transp. Q. 44 (1).
Billheimer, J.W., 1978. The Santa Monica Freeway Diamond Lanes: Evaluation Overview. (663).
Boriboonsomsin, K., Barth, M., 2007. Evaluating air quality benefits of freeway high-occupancy vehicle lanes in Southern California. Transp. Res. Rec. 2011 (1), 137–147.
Boysen, N., Briskorn, D., Schwerdfeger, S., Stephan, K., 2021. Optimizing carpool formation along high-occupancy vehicle lanes. European J. Oper. Res. 293 (3), 1097–1112.
Brown, A.E., 2020. Who and where rideshares? Rideshare travel and use in los angeles. Transp. Res. A 136, 120–134.
Burris, M.W., Winn, J.R., 2006. Slugging in Houston—Casual carpool passenger characteristics. J. Public Transp. 9 (5), 2.
Cohen, M., Fiszer, M.-D., Ratzon, A., Sasson, R., 2019. Incentivizing commuters to carpool: A large field experiment with waze. Available at SSRN 3458330.
Cui, S., Li, K., Yang, L., Wang, J., 2021. Slugging: Casual carpooling for urban transit. Manuf. Serv. Oper. Manag..
Daganzo, C.F., Cassidy, M.J., 2008. Effects of high occupancy vehicle lanes on freeway congestion. Transp. Res. B 42 (10), 861–872.
Dahlgren, J., 1998. High occupancy vehicle lanes: Not always more effective than general purpose lanes. Transp. Res. A 32 (2), 99–114.
Dunning, T., 2012. Natural Experiments in the Social Sciences: A Design-Based Approach. Cambridge University Press.
Giuliano, G., Levine, D.W., Teal, R.F., 1990. Impact of high occupancy vehicle lanes on carpooling behavior. Transportation 17 (2), 159–177.

Hanna, R., Kreindler, G., Olken, B.A., 2017. Citywide effects of high-occupancy vehicle restrictions: Evidence from "three-in-one" in Jakarta. Science 357 (6346), 89–93.
INRIX, 2020. Global traffic scorecard.
Johnston, R.A., Ceerla, R., 1996. The effects of new high-occupancy vehicle lanes on travel and emissions. Transp. Res. A 30 (1), 35–50.
Kelly, K.L., 2007. Casual carpooling-enhanced. J. Public Transp. 10 (4), 6.
Konishi, H., Mun, S.-i., 2010. Carpooling and congestion pricing: HOV and HOT lanes. Reg. Sci. Urban Econ. 40 (4), 173–186.
Kwon, J., Varaiya, P., 2008. Effectiveness of California's high occupancy vehicle (HOV) system. Transp. Res. C 16 (1), 98–115.
Laval, J.A., Daganzo, C.F., 2006. Lane-changing in traffic streams. Transp. Res. B 40 (3), 251–264.
Lawler, D., 1991. Hov lanes receive support in new jersey opinion poll findings. AASHTO Q. Magaz. 70 (4).
Li, J., Embry, P., Mattingly, S.P., Sadabadi, K.F., Rasmidatta, I., Burris, M.W., 2007. Who chooses to carpool and why? Examination of texas carpoolers. Transp. Res. Rec. 2021 (1), 110–117.
Ma, S., Wolfson, O., 2013. Analysis and evaluation of the slugging form of ridesharing. In: Proceedings of the 21st ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems. pp. 64–73.
Mannering, F.L., Hamed, M.M., 1990. Commuter welfare approach to high occupancy vehicle lane evaluation: An exploratory analysis. Transp. Res. A 24 (5), 371–379.
Martin, P.T., Lahon, D., Stevanovic, A., 2005. Review of the Effectiveness of the High Occupancy Vehicle (HOV) Lanes Extension. University of Utah Traffic Lab, Salt Lake City, UT.
Martin, P.T., Perrin, J., Wu, P., Lambert, R., 2004. Evaluation of the effectiveness of high occupancy vehicle lanes. Report to Utah Department of Transportation.
Martin, S., Taylor, S.J., Yan, J., 2021. Trading flexibility for adoption: Dynamic versus static walking in ridesharing. Available At SSRN 3984476.
Menendez, M., Daganzo, C.F., 2007. Effects of HOV lanes on freeway bottlenecks. Transp. Res. B 41 (8), 809–822.
Mote, J.E., Whitestone, Y., 2011. The social context of informal commuting: Slugs, strangers and structuration. Transp. Res. A 45 (4), 258–268.
Neoh, J.G., Chipulu, M., Marshall, A., 2017. What encourages people to carpool? An evaluation of factors with meta-analysis. Transportation 44 (2), 423–447.
Parkany, E., 1999. Can high-occupancy/toll lanes encourage carpooling? Case study of carpooling behavior on the 91 express lanes. Transp. Res. Rec. 1682 (1), 46–54.
Plotz, J., Konduri, K.C., Pendyala, R.M., 2010. To what extent can high-occupancy vehicle lanes reduce vehicle trips and congestion? Exploratory analysis using national statistics. Transp. Res. Rec. 2178 (1), 170–176.
Poppe, M.J., Hook, D.J., Howell, K.M., 1994. Evaluation of high-occupancy-vehicle lanes in phoenix, arizona. Transp. Res. Rec. 1446.
Rodier, C.J., Johnston, R.A., 1997. Travel, emissions, and welfare effects of travel demand management measures. Transp. Res. Rec. 1598 (1), 18–24.
Schijns, S., Eng, P., 2006. High occupancy vehicle lanes–worldwide lessons for European practitioners. WIT Trans. Built Environ. 89.
Shaheen, S.A., Chan, N.D., Gaynor, T., 2016. Casual carpooling in the San Francisco Bay Area: Understanding user characteristics, behaviors, and motivations. Transp. Policy 51, 165–173.
Shewmake, S., 2012. Can carpooling clear the road and clean the air? Evidence from the literature on the impact of HOV lanes on VMT and air pollution. J. Plan. Literature 27 (4), 363–374.
Spielberg, F., Shapiro, P., 2000. Mating habits of slugs: Dynamic carpool formation in the I-95/I-395 corridor of Northern Virginia. Transp. Res. Rec. 1711 (1), 31–38.
Tsitsokas, D., Kouvelas, A., Geroliminis, N., 2021. Modeling and optimization of dedicated bus lanes space allocation in large networks with dynamic congestion. Transp. Res. C 127, 103082.
Turnbull, K.F., 1992. International high-occupancy vehicle facilities. Transp. Res. Rec. (1360).
U.S. Department of Transportation, 2015. High-occupancy vehicle lanes.
U.S. Office of Highway Policy Information, 2014. High occupancy vehicle (HOV) lanes - by state.
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. Transp. Geogr. 22, 77–86.
Veve, C., Chiabaut, N., 2022. Demand-driven optimization method for microtransit services. Transp. Res. Rec. 2676 (3), 58–70.
Wellander, C., Leotta, K., 2000. Are high-occupancy vehicle lanes effective? Overview of high-occupancy vehicle facilities across North America. Transp. Res. Rec. 1711 (1), 23–30.
Wiseman, Y., 2019. High occupancy vehicle lanes are an expected failure. Int. J. Control Autom. 12 (11), 21–32.
Yang, H., 1998. When should carpool lanes be introduced in a multi-lane highway? J. Adv. Transp. 32 (2), 242–252.
Yang, H., Huang, H.-J., 1999. Carpooling and congestion pricing in a multilane highway with high-occupancy-vehicle lanes. Transp. Res. A 33 (2), 139–155.
Zhang, W., Jacquillat, K., Wang, S., 2022. Routing optimization with vehicle-customer coordination. In: Working Paper.
Zhong, L., Zhang, K., Nie, Y.M., Xu, J., 2020. Dynamic carpool in morning commute: Role of high-occupancy-vehicle (HOV) and high-occupancy-toll (HOT) lanes. Transp. Res. B 135, 98–119.