TØI Report 1875/2022

Bjørn Gjerde Johansen

Institute of Transport Economics
Norwegian Centre for Transport Research

Reducing car use through e-scooters

A nudging experiment

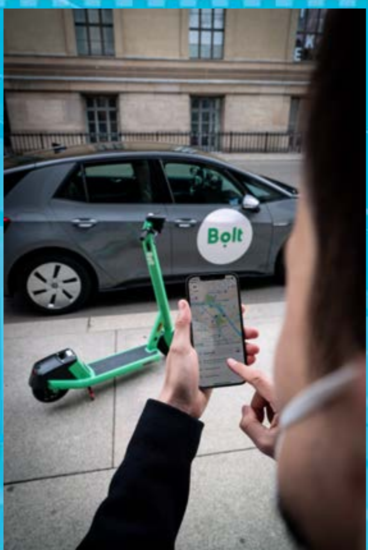

TØI Report
1875/2022

Reducing car use through e-scooters

A nudging experiment

Bjørn Gjerde Johansen

Front-page photography/illustration: Bolt

ISSN 2535-5104 Electronic
ISBN 978-82-480-1921-3 Electronic
Oslo, February 2022

Tittel:	Kan elsparkesykler redusere bilkjøring? Et eksperiment	**Title:**	Reducing car use through e-scooters: A nudging experiment
Forfatter:	Bjørn Gjerde Johansen	**Author:**	Bjørn Gjerde Johansen
Dato:	02.2022	**Date:**	02.2022
TØI-rapport:	1875/2022	**TØI Report:**	1875/2022
Sider:	86	**Pages:**	86
ISSN elektronisk:	2535-5104	**ISSN Electronic:**	2535-5104
ISBN elektronisk:	978-82-480-1921-3	**ISBN Electronic:**	978-82-480-1921-3
Finansieringskilde:	Bolt	**Financed by:**	Bolt
Prosjekt:	5160 – Bolt nudging	**Project:**	5160 – Bolt nudging
Prosjektleder:	Nils Fearnley	**Project Manager:**	Nils Fearnley
Kvalitetsansvarlig:	Askill Harkjerr Halse	**Quality Manager:**	Askill Harkjerr Halse
Fagfelt:	Marked og styring	**Research Area:**	Markets and Governance
Emneord:	Elsparkesykkel Dulting Mikromobilitet	**Keywords:**	E-scooter Nudging Micromobility

Sammendrag:

Sommeren 2021 gjennomførte selskapet Bolt eksperimenter i flere europeiske byer, hvor de endret app-informasjon for å påvirke brukere til å velge delte elsparkesykler istedenfor ridehailing (en taxitjeneste). Dersom app-sesjonen oppfylte visse kriterier, ville brukerne få opp et elsparkesykkelalternativ i ridehailing-delen av appen. Denne rapporten analyserer data fra disse eksperimentene. Vi finner at brukerne som fikk endret app-informasjonen hadde en signifikant høyere sannsynlighet for å velge elsparkesykler i så å si alle eksperimentene. Denne økningen utgjør 0,4-3 prosentpoeng, tilsvarende en 40-200 prosents økning sammenlignet med kontrollgruppen. I Oslo, hvor resultatene er sterkest, finner vi at minst 55 prosent av elsparkesykkelturene eksperimentet genererte erstattet ridehailingturer.

Resultatene demonstrerer at endring av informasjon i multimodale brukergrensesnitt kan være et effektivt virkemiddel for å redusere bilkjøring, uten å medføre kostnader for brukeren.

Summary:

During the 2021 season, Bolt experimented with in-app information in several European cities to “nudge” users from ride hailing to e-scooters: Given a set of criteria, one of the ride hailing alternatives in the app menu would be replaced by an option for e-scooter rental. This report analyses data from these experiments. We find the share of nudged users choosing e-scooter to be significantly higher in virtually all experiments. The e-scooter shares among nudged users are 0.4-3 percentage points higher than for the control groups, constituting a 40-200 percent increase in e-scooter use for the nudged sessions. In Oslo, where results are strongest, at least 55 percent of the e-scooter trips caused by the nudge replaced ride hail trips.

Taken together, results demonstrate that nudging users through changing information in multimodal interfaces can be an effective way of switching users away from cars, at no cost to the user.

Language of report: English

Transportøkonomisk Institutt
Gaustadalleen 21, 0349 Oslo
Telefon 22 57 38 00 - www.toi.no

Institute of Transport Economics
Gaustadalleen 21, N-0349 Oslo, Norway
Telephone +47 22 57 38 00 - www.toi.no

Preface

During the last couple of years, the entry of shared electric scooters (e-scooters) has taken cities throughout the world by storm. Understanding e-scooter demand and substitution patterns to other modes of transport is hence vital to ensure that cities can meet the mobility needs of their citizens in the future.

Bolt is a company offering several app-based mobility services. During the summer season of 2021, the company conducted several experiments by changing in-app information for a randomized set of users in selected European cities. The Institute of Transport Economics (TØI) was subsequently commissioned by Bolt to evaluate the effects of these experiments.

Experiment design, implementation and data collection was done by Bolt prior to TØI's involvement. At TØI, Bjørn Gjerde Johansen has been responsible for data analysis and writing of the report under project management by Nils Fearnley. Nils Fearnley, Alice Ciccone and Paal Wangsness have given useful feedback to both the text and the analyses. Askill Harkjerr Halse has been responsible for the quality control.

The main contact person at Bolt has been Welmoed Neijmeijer, while Dmitry Zamaleev has prepared the datasets. Bolt was given a draft version of the final report for commenting. We would like to thank them for excellent cooperation and constructive feedback, and the opportunity to delve into interesting datasets on their users' mobility patterns.

Oslo, February 2022

Institute of Transport Economics

Bjørne Grimsrud
Managing Director

Kjell Werner Johansen
Research Director

Contents

Summary

Reducing car use through e-scooters: A nudging experiment

TØI Report 1875/2022
Author: Bjørn Gjerde Johansen
Oslo 2022 86 pages English language

During the 2021 season, Bolt experimented with in-app information in several European cities to "nudge" users from ride hailing to e-scooters: If a set of criteria were met, e-scooter rental would appear as the second alternative in the ride hailing part of the app. This report analyses data from these experiments. We find that nudged users are significantly more likely to choose e-scooters in virtually all experiments. The e-scooter shares among nudged users are 0.4-3 percentage points higher, constituting a 40-200 percent increase in e-scooter use compared to the control group. In Oslo, where results are strongest, at least 55 percent of the e-scooter trips caused by the nudge replaced ride hail trips.

Taken together, results demonstrate that nudging users through changing information in multimodal interfaces can be an effective way of switching users away from cars, at no cost to the user.

The app interface and the experiments

Bolt manages ride hail and e-scooter services in various cities. By downloading and signing up to an app, the user can choose from two different mode options: either searching for e-scooters nearby or scheduling a ride hailing trip by submitting a destination.[1] Since both e-scooters and ride hail trips are offered through the same platform, switching between the two is less cumbersome. Furthermore, app data on users' travel behaviour present an opportunity to study the interface between shared e-scooters and ride hailing trips.

During the summer of 2021, Bolt conducted several similar experiments among users in selected European cities: Krakow in Poland, Brno and Ostrava in the Czech Republic, Lisbon in Portugal, Madrid in Spain, Bordeaux in France, Gothenburg and Stockholm in Sweden, Oslo in Norway and Valletta in Malta.[2] The purpose of the experiments was to see whether users could be "nudged" from booking a ride hail trip into renting an e-scooter. The "nudge" consisted of giving ride hail users information about e-scooters in a more accessible part of the app, by inserting an e-scooter option in the ride hailing search menu. By scrolling down, the user still had access to the same ride hailing alternatives. The ride hailing interface for a "nudged" user is displayed in Figure S1.

Users randomized into a control group saw no difference in how the app functioned, while users in the treatment group were nudged provided that their search session met three criteria: (1) A ride hail search was initiated; (2) an e-scooter was available within 300 meters of the user; and (3) the destination for the trip was less than 2 or 3 kilometres away, depending on the experiment. These criteria were meant to identify the ride hailing trips that most conveniently could be replaced by e-scooter trips. Data from the experiments contain information from 10 different cities, consisting of 12.6 million search sessions from about 1.1 million users. 4.5 percent of these search sessions met the nudge criteria.

[1] Bolt also offers other mobility services, but not in the cities analysed here.

[2] Results from Bordeaux, Brno and Ostrava are not presented due to limited sample sizes.

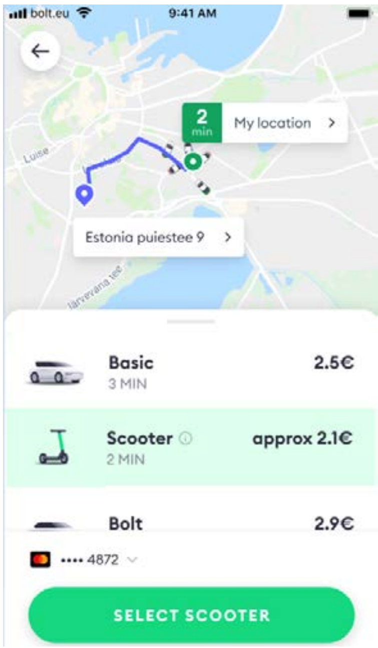

Figure S1: The ride hail interface of the Bolt app in case the user was nudged. An "e-scooter" option appears as the second alternative instead of a ride hail option.

Summary of results

We find that nudging significantly increased the number of e-scooter trips, and reduced the number of ride hailing trips. For Oslo, where results are strongest, we document that 55 percent of e-scooter trips generated by nudging replaced ride hailing. The remaining 45 percent were conducted by users that would otherwise have closed the app without booking a trip. This illustrates that in-app information is able to affect users' transport behaviour in a way that significantly reduces car trips to a larger extent that what has previously been documented in the literature. Hence, nudging can be an effective tool to influence travel behaviour without having to resort to traditional regulatory measures such as taxes or restrictions, where the associated user cost is higher. While the behavioural change is initiated by the nudge experiment, it is facilitated from the fact that the app interface is multimodal. This suggests that interfaces where several modes of transport are integrated can play an important role for mode shifts: centralising mode specific information in one app improves accessibility for the user and in turn allow nudges to influence travel behaviour.

The main findings from the report are summarized below. The first section presents the effect of nudging on e-scooter behaviour, while the next section discusses substitution between e-scooter and ride hailing.

Additional e-scooter information increases e-scooter utilization

The direct effect of being nudged is found by considering outcomes of relevant search sessions, i.e. the 4.5 percent of sessions that met the nudge criteria. By comparing the behaviour of nudged individuals to those in the control group, we can identify whether and to what extent the additional e-scooter information from the nudge directly affected travel behaviours. Figure S2 displays the share of users in the treatment group (red bars) and control group (blue bars) that chose e-scooters (left panel) and ride hailing (right panel) for each experiment.

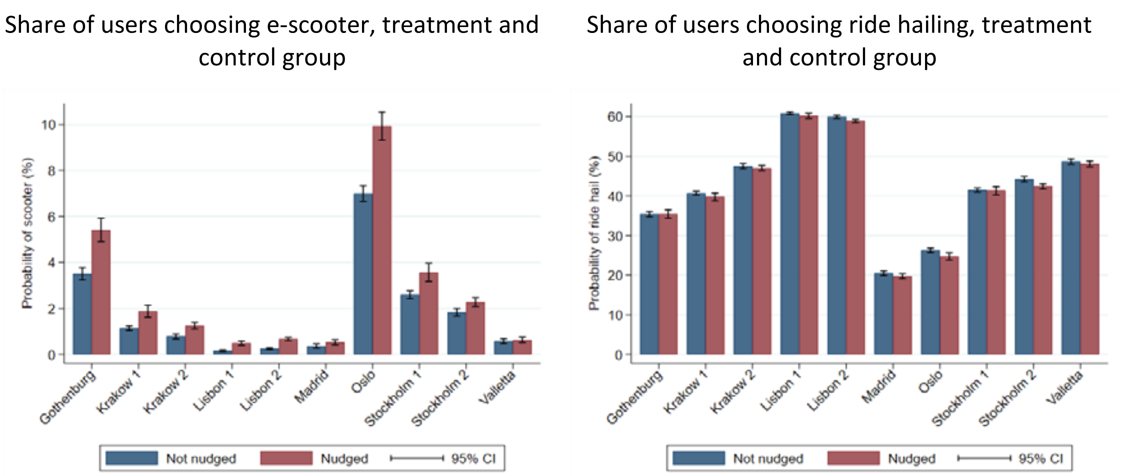

Figure S2: Outcomes of relevant search sessions, treatment and control group.

The trend is that the e-scooter share is higher and the ride hailing share lower in the treatment group compared to the control group, indicating that nudging affected travel behaviour in the intended way. The modal split is also distinctly different across experiments, not only among the treated users but also among the control group. However, it is difficult to compare the e-scooter effect to the ride hailing effect due to the difference in scale – a larger share of the relevant search sessions ended in ride hailing trips than e-scooter trips, because the nudge criteria limit relevant search sessions to users that are searching for ride hailing in the first place. To elucidate the effect of nudging, Figure S3 displays differences between those that are nudged and those that are not.

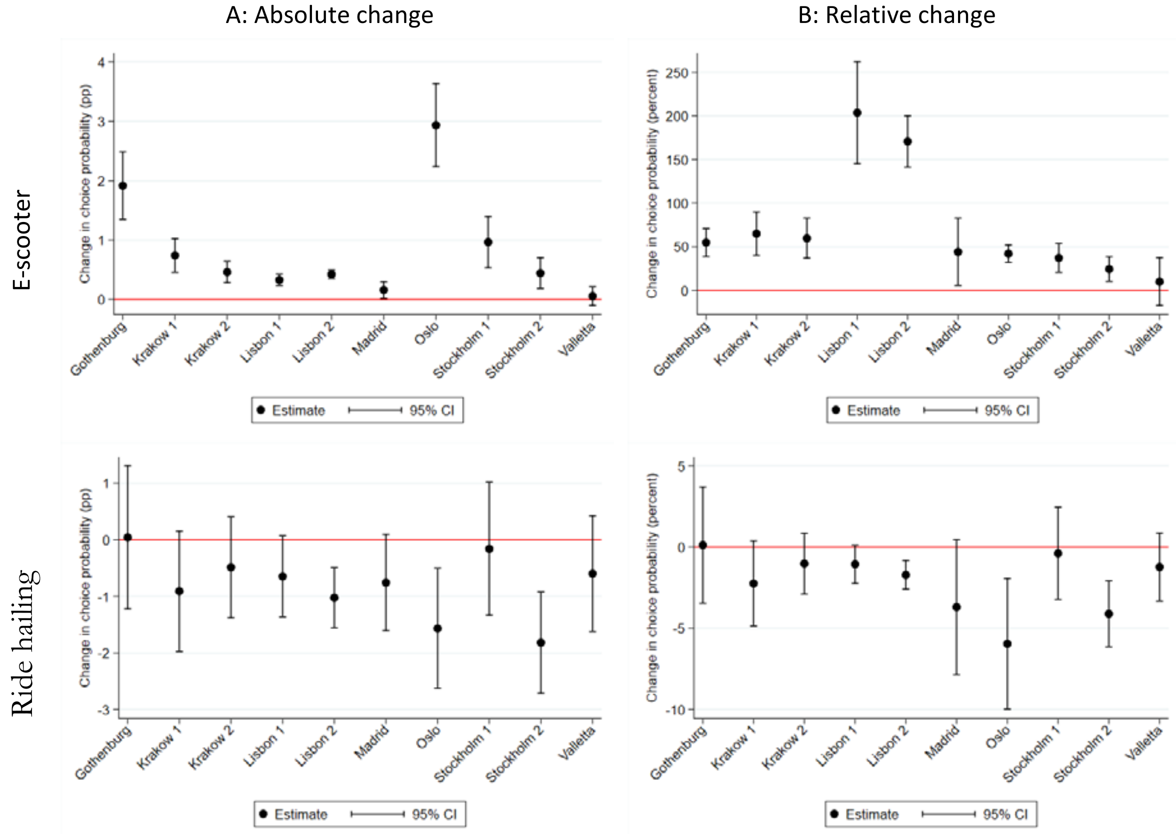

Figure S3: Effect of nudging on probability of choosing e-scooter (top) and ride hailing (bottom), including 95 percent confidence intervals. Absolute difference (left) and difference relative to share in control group (right).

Left panels display absolute differences, while the right panels display the relative size of these differences compared to outcomes for the control group. To explain the difference between the panels we use Oslo as an example, where 10 percent of nudged users and 7 percent of users in the control group chose e-scooter. An estimate of 3 percentage points in the top left panel for Oslo means that the nudging process caused an additional (10-7=) 3 percent of the relevant search sessions to result in e-scooters being chosen. The top right panel shows that this is a (3/7≈) 40 percent increase compared to the outcome for the control group.

The top row of Figure S3 shows that nudged users have a higher chance of booking an e-scooter compared to the control group, and that the treatment effect is statistically significant for all experiments except Valletta. The absolute effect is largest in Oslo (3 percentage points, corresponding to a 40 percent increase). The relative effect is largest in the Lisbon experiments, where about three times as many nudged users chose e-scooter (a 200 percent increase corresponding to 0.4 percentage points). The bottom row illustrates that the increased number of e-scooter trips are mirrored by a reduction in the number of ride hailing trips, although the estimated effects are statistically more uncertain.

For most cities (Lisbon and Valletta being the outliers) nudging increases the share of users choosing e-scooters by 40-60 percent. The fact that this pattern is fairly stable indicates that whichever (observed or unobserved) factors that are making e-scooters more popular to begin with are also increasing the effect of nudging. Further analyses indicate that one of the most important observed factors is the density of e-scooters. Both the effect of nudging as well as the share of users in the control group that chose e-scooters to begin with is higher when e-scooters are available in close proximity. This pattern holds true both across cities and within experiments, which highlights the importance of density and visibility for e-scooters to be considered a viable option by users.

The effect of nudging in terms of e-scooter trips created is modest when only looking at the search sessions in which the nudging took place. For example, we estimate that the direct effect was 283 additional e-scooter trips in Lisbon and 273 in Oslo. This is not because the effect of nudging is small, but because only a small share of search sessions was considered to be relevant for nudging in the first place. However, users that were nudged at least once were more likely to conduct additional e-scooter searches and e-scooter trips subsequently. We find that about 3,800 additional e-scooter trips were conducted in Oslo and 1,400 in Lisbon over the duration of the experiment, as a result of nudging. This shows that people that are first presented with the additional e-scooter information are more likely to alter their behavior in the long term as a result. These predictions however are statistically more uncertain than when just considering user sessions relevant for nudging.

Additional e-scooter information reduces car use

Whether the reduction in ride hailing trips (bottom row of Figure S3) is smaller or larger than the increase in e-scooter trips varies across experiments. In the cities where e-scooters are most popular among the control group (Oslo, Gothenburg and Stockholm 1) nudging seems to increase the number of completed trips (ride hailing and e-scooter taken together). In remaining experiments, the reduction in ride hailing trips is larger than the increase in e-scooter trips. There are two likely mechanisms: First, users who would not consider e-scooters in the first place might see the additional e-scooter information as an annoyance. Second, users that are nudged might miss out on the available ride hailing options if they do not scroll down in the app.

In Oslo, where effects are largest in magnitude and most precisely estimated, nudging reduces the share of relevant search sessions resulting in ride hail trips by 1.6 percentage points. This constitutes about 55 percent of the additional e-scooter trips created. The remaining 45 percent of e-scooter trips caused by nudges were users that otherwise would have closed the app without booking a trip. Previous research on substitution patterns for rented e-scooters in Europe indicates that only 5-10 percent of e-scooter trips replace car trips. Hence, the in-app information distributed through the nudge is able to substitute car trips to e-scooter trips to a much larger extent that what has previously been documented.

In cities where nudging reduces the total number of trips, the share of deterred ride hail trips is even larger. An example is the second experiment in Lisbon, where the reduction in ride hail trips is about twice as large as the increase in e-scooter trips. Hence, for every additional e-scooter trip two relevant search sessions are deterred away from ride hailing. However, we do not know what the users did instead – they might have booked a car from a competing ride hailing service. This makes it more difficult to draw clear conclusions regarding substitution patterns for experiments in which nudging reduces the total number of trips.

Looking at the number of ride hailing trips per user over the whole duration of the experiment, differences between the treatment group and the control group are more noisily estimated. However, for Oslo where the effect is strongest, we estimate that users exposed to nudging in average are travelling 0.9 fewer kilometres by ride hailing, and 1.8 additional kilometres by e-scooter (while the e-scooter estimate is clearly significant, the ride hailing estimate is more uncertain). This indicates that 50 percent of the additional e-scooter kilometres caused by the nudge is diverted from ride hailing, which matches results from the relevant search sessions well.

Sammendrag

Kan elsparkesykler redusere bilkjøring? Et eksperiment

TØI rapport 1875/2022
Forfatter: Bjørn Gjerde Johansen
Oslo 2022 86 sider

Sommeren 2021 gjennomførte selskapet Bolt eksperimenter i flere europeiske byer, hvor de endret app-informasjon for å påvirke brukere til å velge delte elsparkesykler istedenfor ridehailing (en taxitjeneste). Dersom app-sesjonen oppfylte visse kriterier, ville brukerne få opp et elsparkesykkelalternativ i ridehailing-delen av appen. Denne rapporten analyserer data fra disse eksperimentene. Vi finner at brukerne som fikk endret app-informasjonen hadde en signifikant høyere sannsynlighet for å velge elsparkesykler i så å si alle eksperimentene. Denne økningen utgjør 0,4-3 prosentpoeng, tilsvarende en 40-200 prosents økning sammenlignet med kontrollgruppen. I Oslo, hvor resultatene er sterkest, finner vi at minst 55 prosent av elsparkesykkelturene eksperimentet genererte erstattet ridehailingturer.

Resultatene demonstrerer at endring av informasjon i multimodale brukergrensesnitt kan være et effektivt virkemiddel for å redusere bilkjøring, uten å medføre kostnader for brukeren.

Brukergrensesnittet og eksperimentene

Bolt tilbyr «ridehailing» og elsparkesykkelutleie i flere byer.[3] Ved å laste ned en mobil-applikasjon og opprette en profil, kan brukere enten søke etter tilgjengelige elsparkesykler i nærheten, eller bestille en bil ved å legge inn en destinasjon. Siden både elsparkesykler og ridehailing blir tilbudt gjennom den samme plattformen, gir app-data om brukernes reise-atferd en unik mulighet til å studere substitusjonsmønstre mellom de to transportformene.

Over sommeren 2021 gjennomførte Bolt flere lignende eksperimenter i utvalgte europeiske byer: Krakow i Polen, Brno og Ostrava i Tsjekkia, Lisboa i Portugal, Madrid i Spania, Bordeaux i Frankrike, Göteborg og Stockholm i Sverige, Oslo i Norge og Valletta i Malta.[4] Formålet med eksperimentene var å se hvorvidt ridehailbrukere kunne bli «dultet» over fra bil til elsparkesykkel. Dulting («nudging» på engelsk) er et begrep fra adferdsøkonomien som innebærer å påvirke folks handlinger uten bruk av tvang, straff eller økonomiske incentiver. Dultet besto i å gjøre informasjon om elsparkesykler lettere tilgjengelig, ved å legge til et elsparkesykkelalternativ i søkemenyen for ridehail. Brukerne som ble dultet hadde fortsatt tilgang på den samme informasjonen som før, men måtte bla nedover en liste for å få opp like mange ridehailalternativer. Brukergrensesnittet for ridehailsøk for en dultet bruker er vist i Figur S1.

For brukerne som utgjorde kontrollgruppen var det ingen forskjell i hvordan Bolt-applikasjonen virket. Brukere som ble randomisert inn i behandlingsgruppen derimot, ble dultet hver gang søkesesjonen deres oppfylte tre kriterier: (1) brukeren opprettet et ridehail-søk; (2) en elsparkesykkel var tilgjengelig innenfor en radius av 300 meter; og (3) destina-sjonen brukeren hadde tastet inn for ridehailsøket var mindre enn 2/3 kilometer unna,

[3] «Ridehailing» er en taxi-tjeneste hvor biler bestilles og betales for gjennom en app-løsning. Vi bruker konsekvent begrepet ridehail eller ridehailing her for å skille tilbudet fra tradisjonelle taxitjenester. Bolt tilbyr også andre mobilitetsløsninger, men ikke i byene vi ser på i denne rapporten.

[4] Det er ikke gjennomført analyser for Bordeaux, Brno eller Ostrava på grunn av et lavt antall observasjoner.

avhengig av eksperimentet. Formålet med disse kriteriene var å utpeke reiser hvor ride-hailing kunne erstattes av elsparkesykkel på mest beleilig måte. Resultatene fra eksperimentene er et datasett bestående av 12,6 millioner søkesesjoner, fordelt mellom 1,1 millioner brukere i 10 forskjellige byer. 4,5 prosent av disse søkesesjonene oppfyller kriteriene for dulting.

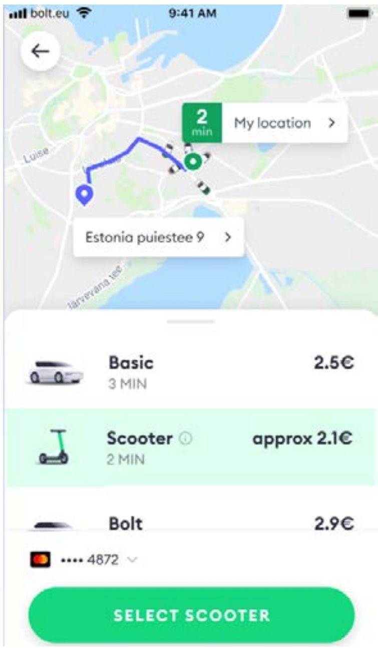

Figur S1: Brukergrensesnittet for å bestille en ridehailtur i Bolt-appen dersom brukeren ble dultet. Elsparkesykkel dukker opp som det andre alternativet på lista over biler.

Sammendrag av resultater

Resultatene viser at dulting i stor grad øker antall elsparkesykkelturer og reduserer antall ridehailturer. I Oslo, hvor resultatene er sterkest, vil 55 prosent av elsparkesykkelturene som er forårsaket av dultingen erstatte ridehailing. For det første illustrerer dette at informasjon brukere får gjennom mobilapplikasjonen er i stand til å endre reisevaner. For det andre viser det en høyere grad av substitusjon fra bilturer til elsparkesykler enn tidligere forskning har tydet på – noe som er lovende dersom målsetningen er å redusere bilkjøring. Dette viser at dulting kan være et effektivt virkemiddel som kan komplementere tradisjonelle reguleringstiltak (for eksempel avgifter og forbud), der de tilhørende brukerkostnadene typisk er høyere.

Atferdsendringen som er forårsaket av dulting, er mulig fordi brukergrensesnittet er multimodalt. Eksperimentene viser dermed også potensialet som ligger i applikasjoner som integrerer flere transportformer i samme grensesnitt. Ved å kombinere transportmiddel-spesifikk informasjon i den samme applikasjonen, kan substitusjon mellom ulike transportformer foregå mer sømløst for brukeren, samtidig som det muliggjør dulting (strategisk plassering av informasjon) som virkemiddel.

Hovedfunnene i rapporten er gjengitt under. Den første delen presenterer resultater om hvordan dulting øker sannsynligheten for at brukeren velger elsparkesykkel, mens den andre delen diskuterer substitusjon mellom elsparkesykkel og ridehailing.

Informasjon om elsparkesykler øker bruken

Direkteeffekten av å bli dultet finner vi ved å se på utfallene av relevante søkesesjoner, altså de sesjonene som oppfylte kriteriene for dulting. Ved å sammenligne utfallene hos de som ble dultet mot utfallene i kontrollgruppen, kan vi måle hvorvidt, og i hvilken grad, den ekstra informasjonen brukerne fikk økte sannsynligheten for at elsparkesykkelturer ble gjennomført. Figur S2 viser andelen av relevante søkesesjoner hos behandlingsgruppen (røde søyler) og kontrollgruppen (blå søyler) hvor brukeren valgte elsparkesykkel (venstre panel) og ridehailing (høyre panel) i hvert eksperiment.

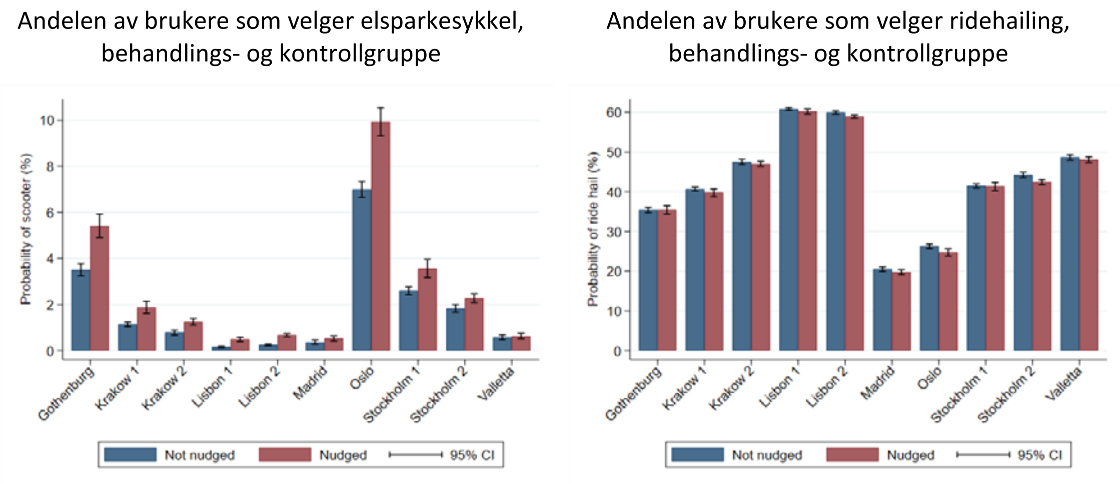

Figur S2: Utfall av relevante søkesesjoner, gjennomsnitt behandlings- og kontrollgrupper samt 95 prosents konfidensintervall.

Trenden er at elsparkesykkelandelen er høyere, mens ridehailandelen er lavere, i behandlingsgruppa sammenlignet med kontrollgruppa. Dette indikerer at eksperimentene hadde ønsket effekt. Det er også betydelige forskjeller i disse andelene på tvers av eksperimenter, hovedsakelig fordi transportmidlenes pris og tilgjengelighet varierer på tvers av byer. Grunnen til at ridehailandelen er høyere enn elsparkesykkelandelen hos begge gruppene er, som nevnt over, fordi vi kun ser på app-sesjoner hvor brukeren søkte etter ridehailing i utgangspunktet.

For å tydeliggjøre effekten av dulting, viser Figur S3 forskjellen mellom de som har blitt dultet og kontrollgruppen. De venstre panelene viser absolutte forskjeller (i prosentpoeng), mens de høyre panelene viser hvor store forskjellene er relativt til andelen i kontrollgruppen (i prosent). For å klargjøre forskjellen mellom de venstre og de høyre panelene, bruker vi elsparkesykler i Oslo som eksempel. Her var elsparkesykkelandelen 10 prosent hos brukere som ble dultet, og 7 prosent hos kontrollgruppen. Som vist øverst til venstre, utgjør dette en 3 prosentpoengs økning i elsparkesykkelturer i relevante søkesesjoner, forårsaket av dultingen. Øverst til høyre ser vi at dette tilsvarer en (3/7≈) 40 prosents økning relativt til det opprinnelige nivået i kontrollgruppen.

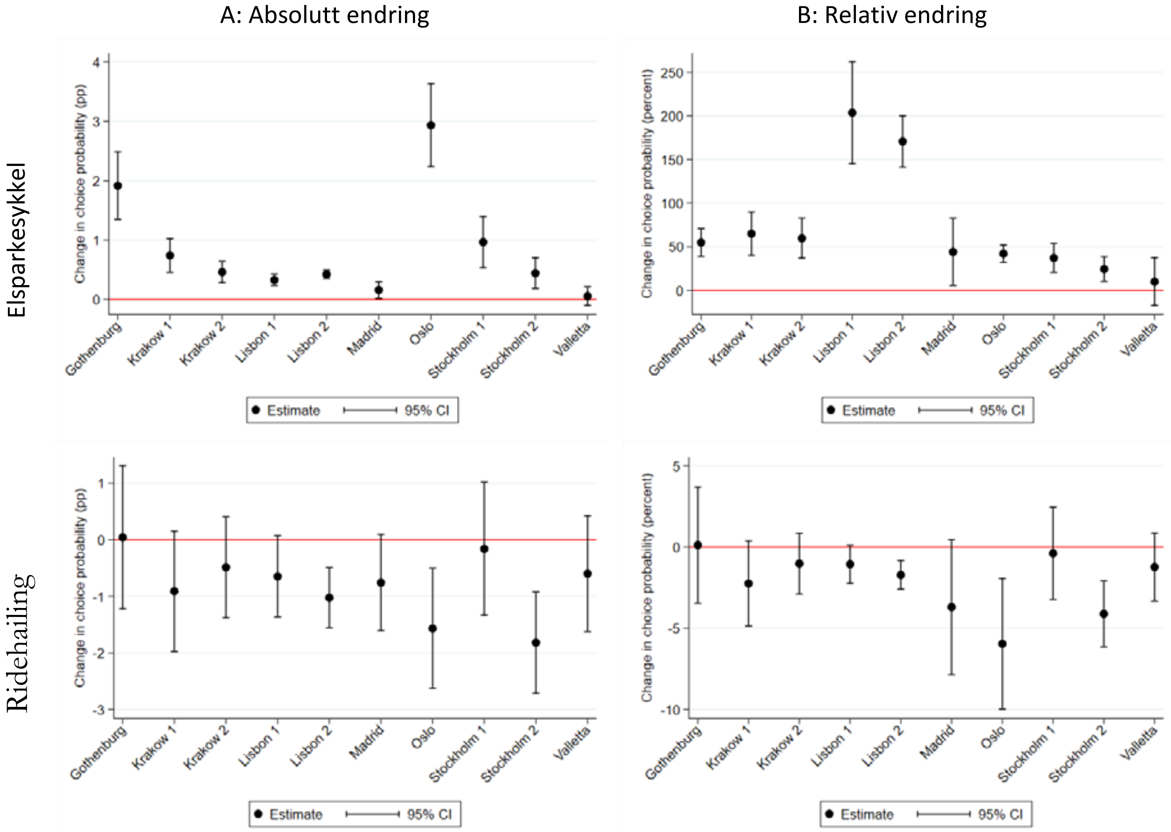

Figur S3: Effekten av dulting på sannsynligheten for å velge elsparkesykkel (øverst) og ridehailing (nederst), inkludert 95 prosents konfidensintervaller. Absolutte endringer (venstre) og relativt til andelen i kontrollgruppen (høyre).

Den øverste raden i Figur S3 viser altså at brukere som har blitt dultet har en høyere sannsynlighet for å velge elsparkesykler sammenlignet med kontrollgruppen, og at effekten er statistisk signifikant for alle eksperimenter bortsett fra Valletta. Den absolutte endringen er størst i Oslo (3 prosentpoeng, tilsvarende en 40 prosents økning). Den relative endringen er størst for eksperimentene i Lisboa, hvor tre ganger så mange i behandlingsgruppen velger elsparkesykkel sammenlignet med kontrollgruppen (en 200 prosents økning, tilsvarende 0,4 prosentpoeng). Den nederste raden viser en motsvarende reduksjon i andelen ridehailturer, selv om denne er mindre presist estimert (større konfidensintervall).

For de fleste byene (bortsett fra Lisboa og Valletta) vil dulting øke andelen elsparkesykkelturer med 40-60 prosent. Dette mønsteret er relativt stabilt på tvers av eksperimenter, noe som indikerer at de samme (observerte eller uobserverte) faktorene som gjør elsparkesykler populære hos kontrollgruppen også bidrar til å øke effekten av dultingen. Videre analyser viser at en av de viktigste faktorene er tettheten av elsparkesykler. Vi finner at både effekten av dulting og andelen som velger elsparkesykkel i kontrollgruppen er betydelig høyere når det er tilgjengelige elsparkesykler i nærheten – dette gjelder både på tvers av byer og innad i eksperimenter.

Den totale økningen i antall elsparkesykkelturer som følge av dultingen i relevante søkesesjoner er beskjeden. I Lisboa førte dultingen til at ytterligere 283 relevante søkesesjoner endte opp som elsparkesykkelturer. Det tilsvarende tallet for Oslo er 273. Dette er ikke fordi effekten av dultingen er liten, men fordi kun en liten andel av søkesesjonene oppfylte kriteriene for dulting. Vi finner imidlertid at brukere som først har blitt utsatt for dulting har økt sannsynlighet for å søke etter elsparkesykler og gjennomføre elsparkesykkelturer i etterkant. Når vi ser på alle app-sesjonene, estimerer vi at dulting førte til omtrent 3 800 flere elsparkesykkelturer i Oslo og 1 400 i Lisboa. Dette viser at dultingen har fått flere brukere til å endre reisevaner i etterkant. Disse prediksjonene er imidlertid statistisk sett mer usikre.

Informasjon om elsparkesykler reduserer bilbruk

Hvorvidt reduksjonen i ridehailturer (den nederste raden i Figur S3) er større eller mindre enn økningen i elsparkesykkelturer varierer mellom eksperimentene. I eksperimenter hvor elsparkesykler er mer populære hos kontrollgruppen (Oslo, Göteborg og Stockholm 1), ser det ut til at dulting øker sannsynligheten for å ende søkesesjonen med en tur (enten ride-hailing eller elsparkesykkel). I de resterende eksperimentene er imidlertid reduksjonen i ridehailturer større enn økningen i elsparkesykkelturer. Det er to sannsynlige mekanismer for dette. For det første kan det hende brukere som ikke ville vurdert elsparkesykkel i utgangspunktet blir irriterte av den ekstra informasjonen som dultingen medfører. For det andre kan det hende de går glipp av ridehailinformasjonen de opprinnelig søkte etter dersom de ikke blar ned i lista over alternativer.

I Oslo, hvor effekten både er størst og mest presist estimert, reduserer dulting sannsynligheten for en ridehailtur med 1,6 prosentpoeng i en relevant søkesesjon. Dette utgjør omtrent 55 prosent av økningen for elsparkesykkel. De resterende 45 prosentene av elsparkesykkelturer forårsaket av dulting ble gjennomført av brukere som i utgangspunktet ville lukket appen uten å gjennomføre en tur. Tidligere forskning på substitusjonsmønstre i utleiemarkedet for elsparkesykler i Europa indikerer at kun 5-10 prosent av turene erstatter bilturer. Med andre ord har dette eksperimentet vist at app-informasjon om elsparkesykler har potensiale for å erstatte bilturer i høyere grad enn det som tidligere har blitt dokumentert.

I byer hvor dulting reduserer det totale antallet gjennomførte turer, er nedgangen i ridehailturer relativt til økningen i elsparkesykkelturer enda større. Et eksempel er det andre eksperimentet i Lisboa, hvor to ridehailturer blir unngått for hver ekstra elsparkesykkeltur. Vi vet imidlertid ikke hva disse brukerne ender opp med å gjøre istedenfor – de kan for eksempel ha endt opp med å bestille en ridehailtur fra et konkurrerende selskap. Dette gjør det vanskeligere å belyse de faktiske substitusjonsmønstrene i eksperimenter hvor dultingen reduserer det totale antallet gjennomførte turer.

Når vi ser på alle gjennomførte turer over hele eksperimentperioden, ikke bare de relevante søkesesjonene, er forskjellene mellom behandlings- og kontrollgruppen mindre presist estimert. I Oslo, hvor effekten er sterkest, ser det ut til at brukere i behandlingsgruppen som har blitt eksponert for dulting i snitt reiser 0,9 færre kilometer med ridehail og 1,8 flere kilometer med elsparkesykkel. Dette stemmer godt overens med resultatene fra de relevante søkesesjonene, og indikerer at halvparten av elsparkesykkelkilometerne forårsaket av dultingen erstatter ridehail.

1 Introduction

Bolt is a company that operates shared electric scooters (e-scooters) as well as ride hailing services in many European cities. Orders are managed by the customer through a smart phone application. During the 2021 season, Bolt experimented with in-app information to prompt users to make a different travel mode choice than initially intended. Given a set of criteria, including travel distance and availability of an e-scooter nearby, the app would suggest e-scooter rental in addition to ride hailing alternatives to the user. According to Bolt, the intent of these experiments was to uncover how they could promote sustainable mobility. This report presents analyses of data collected from these experiments.

1.1 Background

Private car use is associated with several negative externalities which cause enormous efficiency losses to the economy. This includes among other local pollutants, greenhouse gas emissions and congestion, especially during peak hours in urban areas. Van Essen et al (2019) estimate that the total external cost of passenger cars for EU28 countries in 2016 was 625.2 billion Euros, or 4.2 percent of the countries' combined GDP. This corresponds to 7.8 Euro-cent per passenger kilometre without congestion and 12.0 Euro-cent per passenger kilometre with congestion. Rødseth et al (2019) indicate that this cost can be twice as high per passenger kilometre in densely populated areas. As urban areas densify and greenhouse gasses accumulate in the atmosphere, the external cost per passenger kilometre for road transport in cities is expected to increase.

Moreover, greenhouse gas emission cuts in the transport sector are necessary in order to reach the CO_2 mitigation goals set by the Paris agreement, as CO_2 emissions from transport constitute almost 25 percent of global emissions (IEA, 2020). However, the current European policy mixes for road transport do not seem sufficient to achieve this (Axsen et al., 2020).

The concept of "micromobility" has been proposed as an avenue to reduce car transport and achieve more sustainable transport systems in urban areas. Micromobility is a commonly used, but not universally defined, concept referring to shared or personally owned small and lightweight vehicles that are typically electric and used for short distance trips, such as e-bikes and e-scooters. According to Yanocha and Allan (2019), micro-mobility services in combination with walking, cycling and public transport can ease the demand for private car ownership in cities and reduce private driving significantly as a result.

E-scooter rental through app-based sharing systems is a form of micromobility service that has grown in popularity in recent years. Several e-scooter sharing systems are operated in most major European cities, and the demand for shared e-scooters is expected to grow by more than 25 percent per year in the period 2019-2025.[5] However, current research indicates that there is limited substitution from cars to e-scooters when it comes to e-

[5] See https://www.psmarketresearch.com/market-analysis/europe-electric-scooters-and-motorcycles-market (accessed January, 2022).

scooter rental services in Europe.[6] According to a Norwegian survey among e-scooter users, only 8 percent of e-scooter trips replace car based modes of transport (Fearnley et al, 2020a; 2020b). In Vienna, almost 90 percent of renters state that shared e-scooter trips would never replace car trips (Laa and Leth, 2020). A mode choice model estimated for Switzerland predicts that 10 percent of shared e-scooter trips replace car trips (Reck et al, 2022). In all studies, e-scooter rental is most likely to replace walking trips. This poses the important question: under what conditions is e-scooter rental a viable alternative to driving?

1.2 The Bolt app and the nudging experiment

Bolt is a company that administrates a multimodal service: in addition to e-scooter rental, the user is able to order ride hailing trips through the same app interface. Other services such as shared e-bikes and car sharing are also offered through the same application, but not in the cities that were subject to this experiment. The ride hailing service allows users to plot their destination, and then see the location of available drivers on a map, as well as the price and the waiting time for each car. By choosing a particular car, it will pick up and deliver the user at the specified destination, while payment is done through the app interface.

The e-scooter rental part of the app works in a similar fashion – the user can see the location of nearby e-scooters, unlock an e-scooter and pay after the trip is completed from the app interface. However, the user is not required to specify a destination before booking. Instead, a trip can be completed and paid for as long as the e-scooter is parked within a designated area, specified on the map within the app.

The fact that the platform is multimodal means that app data on users' behaviour present a unique opportunity to study the interface between shared e-scooters and car based trips. Furthermore, the opportunities for substitution should be enhanced from the fact that e-scooters and ride hail trips are offered through the same platform, since it makes switching between the two less cumbersome.

During the summer of 2021, Bolt conducted several similar experiments among users in selected European cities: Krakow in Poland, Brno and Ostrava in the Czech Republic, Lisbon in Portugal, Madrid in Spain, Bordeaux in France, Gothenburg and Stockholm in Sweden, Oslo in Norway and Valletta in Malta. The purpose of the experiments was to see whether users could be "nudged" from booking a ride hail trip into renting an e-scooter. The "nudge" consisted of replacing the second ride hail option with an e-scooter option in the ride hailing part of the app, in case the trip was less than a certain length and an e-scooter was available nearby. The ride hailing interface for a "nudged" user is displayed in Figure 1.1.

[6] North American evidence suggests larger degrees of car substitution (Wang et al., 2021).

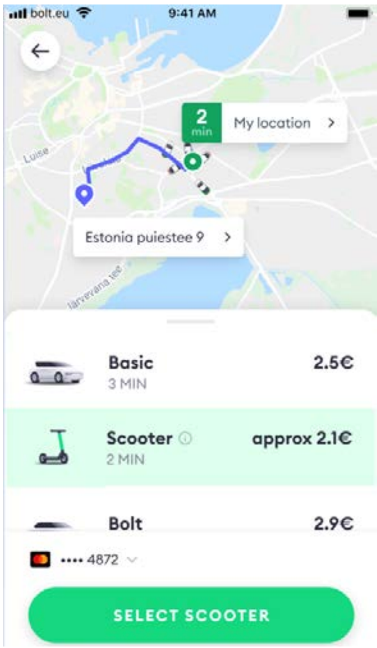

Figure 1.1: The ride hail interface of the Bolt app in case the user was nudged. An "e-scooter" option appears as the second alternative instead of a ride hail option.

Nudging is a term made popular by the book "Nudge: Improving Decisions about Health, Wealth and Happiness" by Thaler and Sunstein, published in 2008. The term 'nudge' refers to a "libertarian paternalistic" approach (Thaler and Sunstein, 2003), or simply put a situation where the freedom of choice is maintained, while a third party attempts to steer people's choices in a certain direction by some non-intrusive intervention. Thaler and Sunstein (2008) define the concept in the following manner:

> *A nudge, as we will use the term, is any aspect of the choice architecture that alters people's behavior in a predictable way without forbidding any options or significantly changing their economic incentives. To count as a mere nudge, the intervention must be easy and cheap to avoid. Nudges are not mandates. Putting fruit at eye level counts as a nudge. Banning junk food does not.*

To count as a "nudge", an intervention should not affect features over which people have explicit preferences (such as the price of ride hail or e-scooters), but rather those features people would typically claim not to care about. The purpose of the nudge experiment conducted by Bolt was to provide users with additional information, by moving information regarding e-scooters to a more convenient position in the app interface.

The experiments were conducted as randomized control trials: users were randomized between a "treatment" group and a "control" group. Users in the "control" group saw no difference in how the app functioned, while users in the "treatment" group were given the nudge displayed in Figure 1.1 provided that their search session met the nudge criteria. By comparing the nudged individuals to those in the control group, we can therefore identify whether nudging affects users' travel behaviours.

1.3 Purpose

The purpose of this report is to shed light on the following questions by empirically analysing data on user behaviour from the nudging experiments:

1. Can we identify behavioural changes in e-scooter use among app users as a result of the nudging experiments?
2. If so, what are the contexts in which nudging has the most substantial impact?
3. To what extent is nudging able to switch users from ride hailing trips to e-scooters?

There are several reasons for why the experiments are novel and the results interesting from a research perspective. First, since e-scooter rental is a relatively new form of micromobility service, there is still little research on travel behaviour in the e-scooter rental market. Considering the recent spike in e-scooter popularity, understanding these travel patterns is vital to ensure that cities can meet the mobility needs of their citizens in the future. Second, the fact that Bolt is a company that promotes a multimodal service by offering both e-scooter rental and ride hailing services offers a unique opportunity to study the interface and substitution patterns between the two markets. Third, experiments are conducted as randomized control trials, in which nudging is distributed randomly among users. The randomization process means that we can give the results a causal interpretation. This is considered the "gold standard" within empirical research, providing the highest quality of evidence (see Section 2.1). Fourth, nudging is a relatively new phenomenon, considered to have several advantages over traditional regulatory measures: users are not forced and there is no additional user cost of not complying with the nudge. Still, behavioral change as a result of nudging is an under-explored topic within the field of transport/mode choice (Franssens, 2021). Finally, the fact that the same experiment is conducted in several European cities means that we can compare differences in behavioral changes across different contexts.

1.4 Structure of the Report

The report is structured as follows. Chapter 2 describes the methods used and what a randomized control trial is. Chapter 3 gives more details regarding the experiment conducted by Bolt and describes the data. Chapter 4 presents the main results. Section 4.1 presents analyses of how users responded to being nudged in a given search session, while Section 4.2 attempts to identify long-term behavioural changes beyond the nudge situation. Finally, Chapter 5 summarises and concludes.

2 Methods

The main purpose of the analyses in this report is to estimate causal effects from several randomized control trials (RCTs) conducted by Bolt during the summer of 2021. All RCTs had the same design, but were conducted in different cities and at different points in time. Users in each city were randomly assigned to either a "treatment" or a "control" group.[7] The treatment consisted of displaying "e-scooter" as an alternative in the ride hailing part of the Bolt app in case certain criteria were met (see Figure 1.1). Throughout this report we will refer to this treatment as "nudging" (see Section 1.2).

This chapter describes how data from these experiments have been analysed to estimate causal effects of nudging on various outcome variables. Section 2.1 describes briefly the theory behind, and virtue of, RCTs. Section 2.2 outlines how results from RCTs can be analysed by means of linear regression. Finally, Section 2.3 characterises discrete choice methods, an alternative to linear regression when the outcome of interest is discrete rather than continuous.

2.1 Randomized control trials

In general, it is difficult to estimate causal effects from data on human behaviour. This difficulty is often formulated as the "fundamental identification problem of causal inference": for individuals that receive a treatment, we do not observe what would have happened in the counterfactual scenario where the treatment was not received (Rubin, 1974). An alternative approach is to compare average outcomes in the population that received treatment to a population that did not receive treatment. However, there are a myriad unobserved factors influencing decision-making processes that can bias this comparison:

- **Selection bias** will occur if individuals can self-select into treatment;
- **Confounding variables** are factors that can influence both the outcome and exposure to treatment, causing a spurious association.

The virtue of RCTs is that selection bias and confounding are eliminated (Schulz and Grimes, 2002). When allocation to treatment is random, unobserved factors that may affect the outcome of interest will mechanically be evenly distributed between the treatment and the control group. Therefore, the effect of unobserved factors will cancel out and differences in expected outcomes between the groups must be caused by the treatment.

For this reason RCTs are considered to be the gold standard when it comes to identifying causal evidence (Meldrum, 2000). This is also the case within transport research: after reviewing 77 transport intervention evaluations, Graham-Rowe et al (2011) concluded that RCTs constitute the highest quality of evidence-based research and should be used whenever possible.

The result of an RCT is most commonly presented in terms of "average treatment effects". This is simply the difference in average outcomes between the treatment and the control

[7] The terms "treatment" and "control" are standard in RCTs and stem from clinical trials that typically compare the effect of drugs, diagnostic procedures or other medical treatments.

group. Moreover, RCTs permit the use of probability theory to express the likelihood that differences in average outcomes are statistically different from zero, rather than being different merely by chance (Schulz and Grimes, 2002).

2.2 Linear regression

In this report, we estimate average treatment effects by means of ordinary least squares by regressing the outcome on a treatment indicator and a constant. This estimates average treatment effects with minimal assumptions, leading to high internal validity (Flannelly et al., 2018): the estimated coefficient on the treatment indicator will be numerically equivalent to the difference in average outcomes (see Athey and Imbens, 2017 for a more thorough discussion on regression analysis in the context of RCTs).

All average treatment effects in this report are displayed graphically with spiked bars indicating 95 percent confidence intervals (CIs), calculated from robust standard errors (White, 1980). Estimates are referred to as significantly different from zero when the value of zero is outside the bounds of the CIs. This is more likely to happen (a) the stronger the true treatment effect is; (b) the less noisy the outcome variable is (i.e. a lower variance); and (c) the more observations that are available in the treatment and control groups.

In Section 4.1.1 we consider the effect of being nudged in a specific search session, and estimate average treatment effects on the probability of choosing one of three different, discrete outcomes: e-scooter, ride hail, or ending the search session without making a trip.

In reality, treatment effects are likely to vary across contexts. In Section 4.1.2 we attempt identify heterogeneity in treatment effects by estimating average treatment effects for sub-populations, where sub-populations are defined according to some dimension of the observed data.

In Section 4.2 we consider the fact that nudging might have influenced behaviour outside the relevant search session where the users were nudged. We do this by aggregating outcomes by user, and estimate average treatment effects for users that were nudged at least once (compared to the subset of the control group that experienced a similar situation at least once). This makes outcomes of interest continuous rather than discrete – instead of estimating changes in the e-scooter probability (the share of observations that are choosing e-scooters), we look at changes in number of e-scooter trips per user over the duration of the experiment.

2.3 Discrete choice models

Some of the outcomes considered in this report are discrete choices; that is, individuals choose one alternative from a given set of options. The discrete choice situation the user is faced with in this context is the outcome of a search session, which will take one of the categorical values “e-scooter”, “ride hail”, or “no trip”. In Section 4.1.3 we use a discrete choice model for the Oslo experiment[8] to estimate how choice probabilities for these outcomes depend on several observed variables, including whether the user received the nudge treatment. When estimating average treatment effects from RCTs, linear regression is generally preferred over discrete choice models, since the choice modelling frameworks rely on both behavioural and functional form assumptions (Athey and Imbens, 2017).

[8] The experiment in Oslo is the only one with sufficient variation in the e-scooter price per hour. For most other cities, Bolt’s price per hour for e-scooters is constant throughout the experiment.

However, discrete choice models are considered more suitable for estimating how other (non-random) attributes, such as prices, influence choices.

Choice modelling is a framework suitable for multivariate analyses, by statistically relating the probability of making a certain choice to observed attributes of either the decision maker, the alternatives or the choice context (McFadden 1978, 1981). Discrete choice models are widely used for analysis in the transport sector, e.g., consumers' choices of car to purchase, and commuters' choice of mode of transportation. When estimating conditional effects of non-randomized attributes, discrete choice models have advantages that linear regression models lack when it comes to internal consistency. First, choice probabilities will always be bounded between zero and one. Second, for each decision-maker the sum of predicted choice probabilities over outcomes will add up to one. Third, the gradient of the choice probability decreases smoothly as the probability approaches zero or one. This is often a more behaviourally realistic assumption than linear regression.

The standard method in the discrete choice literature for modelling these types of situations is the Multinomial Logit model (see, e.g. Ben-Akiva and Lerman, 1985). This method is consistent with random utility theory and based on three main assumptions: (1) the decision-maker chooses the alternative with the highest utility; (2) alternative specific utility functions consist of a deterministic part to be estimated from data and an unobserved random variable; and (3) the unobserved part of the utility follows a particular distribution, meaning that choice probabilities can be calculated from estimated parameters given the data.

3 Experiment design and data

3.1 The Bolt app

Bolt manages ride hail and e-scooter services in various cities. By downloading and signing up to an app, the user can either search for e-scooters nearby or schedule a ride hailing trip by submitting their destination. The app works in the following way:

- When opening the app, the user sees her location on a map. If the user's previous app session was a ride hailing search, she sees the location of any nearby ride hailing vehicles. If the user's previous app session was an e-scooter search, she sees the location of any nearby e-scooters.
- If the user is in the ride hailing part of the app, she can either provide a location to initiate a ride hail search or press an e-scooter icon to switch over to e-scooter searches.
- If the user initiates a ride hailing search by providing a location, she will see the route on the map and get a list of vehicle types to choose from. Three vehicle types are visible immediately, and more vehicle types are shown when scrolling down. The price of the ride and the arrival time of the driver is clearly specified for each alternative.
- If the user searches for e-scooters, she will see available e-scooters (including battery capacity) on the map, as well as the route and walking time to the nearest e-scooter. Since the user has not provided a destination, the price information is displayed in terms of "cost per minute".[9] The map also displays borders around special areas (e.g. areas where it is not allowed to park).

3.2 The nudging experiment

In the summer of 2021, Bolt conducted several experiments consisting of nudging all iOS users of the app from ride hail to e-scooter by providing additional information on e-scooter trips from the ride hailing search menu. The experiments were conducted in 8 different countries and 10 different cities: Krakow in Poland, Brno and Ostrava in the Czech Republic, Lisbon in Portugal, Madrid in Spain, Bordeaux in France, Gothenburg and Stockholm in Sweden, Oslo in Norway and Valletta in Malta.

Users were randomly allocated to "treatment" and "control" groups in each city. The assignment probability varied by experiment, between 20 and 50 percent (see Table 3.1). While users in the control group saw no difference in how the app functioned, users in the treatment group were nudged in all app sessions where certain criteria were met. The criteria were:

- Destination has been specified such that a ride hailing session has started;
- Distance to the nearest e-scooter is no more than 300 meters;

[9] There is also a fixed cost component associated with unlocking an e-scooter, but in most cities this component is set to zero for the time periods we have data.

- Projected distance to the destination no more than 2000 or 3000 meters, depending on the city: In the last set of experiments (Bordeaux, Brno, Madrid, Ostrava, Valetta) the threshold was set to 2000 meters. In remaining cities it was 3000 meters.

This implies that users in the treatment group that did not satisfy the criteria were never nudged. The nudge experiment consisted of replacing the second ride hailing option with an e-scooter option, as visually indicated by Figure 1.1. If a nudge was initiated, the user saw the following information change for the second ride hail alternative:

- Instead of a car icon, the user is shown an e-scooter icon;
- Instead of "waiting time until driver arrives", the e-scooter option displays "walking time to nearest e-scooter".
- Instead of ride hailing price, the user is shown an approximate e-scooter price for the same trip (the word *approximate* being added before the listed price), based on the e-scooter price per minute, the shortest distance between origin and the specified destination and an assumed speed of 10 kilometres per hour.

Each experiment lasted for approximately 4-6 weeks. Table 3.1 displays the location for each experiment, as well as number of users in control and treatment groups. These are all iOS users of the Bolt app with at least one search for either ride hail or e-scooter services during the period. The two last columns display the number of *relevant* users in both the treatment and the control group, defined as the users that satisfied the criteria for the nudge experiment at least once. In Krakow, Lisbon and Stockholm two experiments are conducted right after each other. In the first round of experiments, the assignment probability (i.e. the share of users allocated to the treatment group) for these cities was 20 percent. In Gothenburg and Oslo the assignment probability was 30 percent. In remaining experiments, 50 percent.

Table 3.1: Overview of experiments conducted, sorted by date. Number of total and relevant users in treatment and control groups.

City and experiment	Total number of users		Number of relevant users	
	Control	Treatment	Control	Treatment
Krakow 1	57,727	14,356	15,616	3,931
Lisbon 1	111,063	27,959	25,272	6,424
Stockholm 1	138,090	34,860	17,305	4,410
Gothenburg	50,637	21,851	9,459	4,022
Oslo	30,299	12,888	8,489	3,637
Krakow 2	39,364	39,238	10,116	10,075
Lisbon 2	83,877	84,091	18,180	18,025
Stockholm 2	93,118	93,013	12,072	11,950
Madrid	61,545	61,480	6,558	6,677
Ostrava	4,847	4,709	573	530
Valletta	26,514	26,613	5,001	5,029
Bordeaux	4,560	4,638	178	183
Brno	7,629	7,723	784	782
Total	709,270	433,419	129,603	75,675

Although some users were nudged multiple times, the majority of the users in both the treatment and control group never met the nudge criteria. This can be seen by comparing the number of total users to the number of relevant users. The selection process implied by the nudge criteria is further discussed in Section 3.3.2.

3.3 Data and descriptives

We have obtained data from all users that opened the Bolt app on an iOS device in one of the chosen cities within the time frame of each experiment, i.e. periods of 4-6 weeks. The observational unit in the raw data is "search sessions" including time stamps and fully anonymised user IDs. This means that we can follow the same user over time. A "search session" is the time period from the app is opened until either (1) the user ends a ride, (2) the user kills the app, or (3) about half an hour has passed with the app being inactive. A search session may contain both a ride hail and an e-scooter search, but will always have one of three main outcomes: an e-scooter trip, a ride hail trip or no trip. See Appendix A for a full list of all variables that are available in the raw data.

Section 3.3.1 will go through a few descriptives for the whole dataset, to show differences in user patters across cities. Next, Section 3.3.2 will zoom in on "relevant search sessions", i.e. search sessions that satisfy the criteria for nudging. Note that since only specific ride hail searches are nudged, the relevant search sessions will be initiated by users that are looking to order ride hailing. A consequence of this is that a larger share of the "relevant search sessions" will result in ride hailing trips than e-scooter trips, compared to an average search session.

3.3.1 All search sessions

Table 3.2 displays the number of observations (i.e. search sessions) for each experiment. Experiments in the same city are numbered chronologically (i.e. the experiment "Krakow 1" took place before "Krakow 2").

Table 3.2: User sessions, e-scooter and ride hail trips per experiment. All users (control and treatment group).

Experiment	User sessions (frequency and percent of total)				Sessions per user	Share of rides by e-scooter
	Total	E-scooter rides	Ride hail rides	No rides		
Bordeaux	58,469	7,197	1,611	49,661	6.36	81.7%
		12.3%	2.8%	84.9%		
Brno	109,819	6,401	23,021	80,397	7.15	21.8%
		5.8%	21.0%	73.2%		
Gothenburg	702,203	133,804	101,401	466,998	9.69	56.9%
		19.1%	14.4%	66.5%		
Krakow 1	785,941	97,811	176,802	511,328	10.90	35.6%
		12.5%	22.5%	65.1%		
Krakow 2	755,513	92,175	197,421	465,917	9.61	31.8%
		12.2%	26.1%	61.7%		
Lisbon 1	1,761,582	52,036	701,723	1,007,823	12.67	6.9%
		3.0%	39.8%	57.2%		
Lisbon 2	2,269,242	88,134	819,566	1,361,542	13.51	9.7%
		3.9%	36.1%	60.0%		
Madrid	1,695,366	28,198	261,842	1,405,326	13.78	9.7%
		1.7%	15.4%	82.9%		
Oslo	752,616	277,655	41,661	433,300	17.43	87.0%
		36.9%	5.5%	57.6%		
Ostrava	84,812	5967	17339	61,506	8.88	25.6%
		7.0%	20.4%	72.5%		
Stockholm 1	1,510,823	271,787	271,243	967,793	8.74	50.1%
		18.0%	18.0%	64.1%		
Stockholm 2	1,576,993	314,331	278,207	984,455	8.47	53.1%
		19.9%	17.6%	62.4%		
Valletta	602,191	42,842	193,037	366,312	11.33	18.2%
		7.1%	32.0%	60.8%		
Total	12,665,570	1,418,338	3,084,874	8,162,358	11.08	31.5%
		11.2%	24.4%	64.5%		

The first column displays the total number of search sessions, while columns 2-4 show how these sessions are distributed between the three outcomes “e-scooter ride”, “ride hail ride” and “no ride”, the latter meaning that the user closed the app before a trip took place. Column 5 shows the average number of sessions per user for each experiment (this measure is not directly comparable across experiments, since the number of days each experiment lasted varied). The last column shows the share of rides that are e-scooter rides, ignoring the user sessions where no ride was chosen.

The table clearly reveals large differences across cities. Number of searches per experiment varies between 58 thousand (Bordeaux) and 2.3 million (the second experiment in Lisbon). The popularity of e-scooter relative to ride hail varies from less than 10 percent of trips (Lisbon and Madrid) to more than 80 percent of trips (Bordeaux and Oslo). In Sweden, the split between e-scooter and ride hail trips is close to equal, while remaining cities have roughly 20-30 percent e-scooter trips. Oslo seems to have the most active users by number of sessions, also when taking into consideration the duration of each experiment. Bordeaux and Madrid have the highest share of user sessions not resulting in any ride. This could indicate a lower supply of e-scooters and/or ride hailing drivers. Notably, while e-scooters can be parked freely in most cities, Ostrava and Bordeaux and have mandatory parking spots in the city centre which reduces ridership. Madrid on the other hand is the largest city in terms of area. For the average e-scooter search, the distance to the nearest e-scooter is higher in Madrid and Lisbon compared to other cities, indicating that the supply per square kilometre is lower. However, patterns can also be driven by a number of other differences between cities such as competition, prices and wage levels, user habits, public transport supply, etc.

Figure 3.1 displays the number of e-scooter and ride hail trips per user over the whole duration of each experiment (4-6 weeks). The figure focuses on the control group to unveil the average user pattern in the business-as-usual scenario without interference from the nudge experiment.

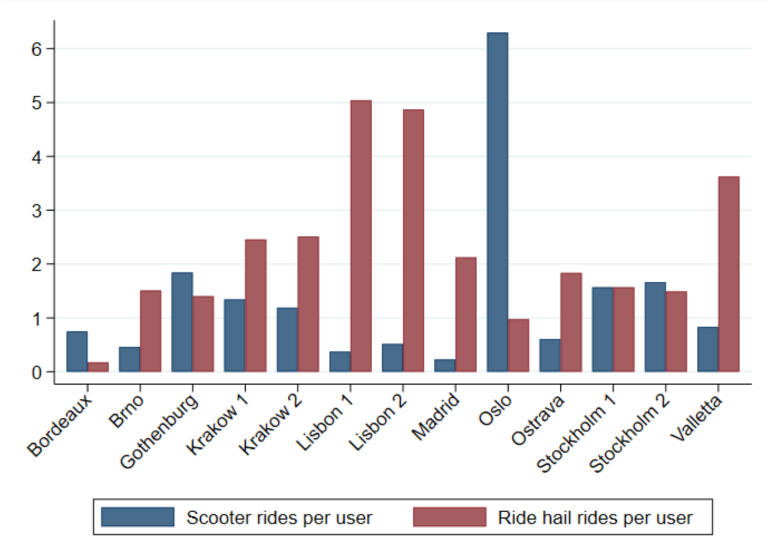

Figure 3.1: Average number of e-scooter and ride hail trips per user in the control group. Whole experiment.

The cities that stand out most in the figure above are Oslo and Lisbon. Oslo has significantly more e-scooter trips per user than any other city, and (except for Bordeaux) the fewest ride hail trips. Lisbon on the other hand has the most ride hail trips and less than 0.5 e-scooter trips per user for each experiment.

To get a better overview of the e-scooter and ride hailing markets in each city and their development over the experiments' duration, Figure 3.2 and Figure 3.3 display e-scooter and ride hail trips per day per city for the first and second rounds of experiments, respectively. The first round of experiments happened within the time period mid-June to mid-August, while the second round of experiments happened in September and October. The histograms display the number of trips per day per city, where red is "ride hail" and blue is "e-scooter". The histogram is purple where ride hail and e-scooter overlaps.

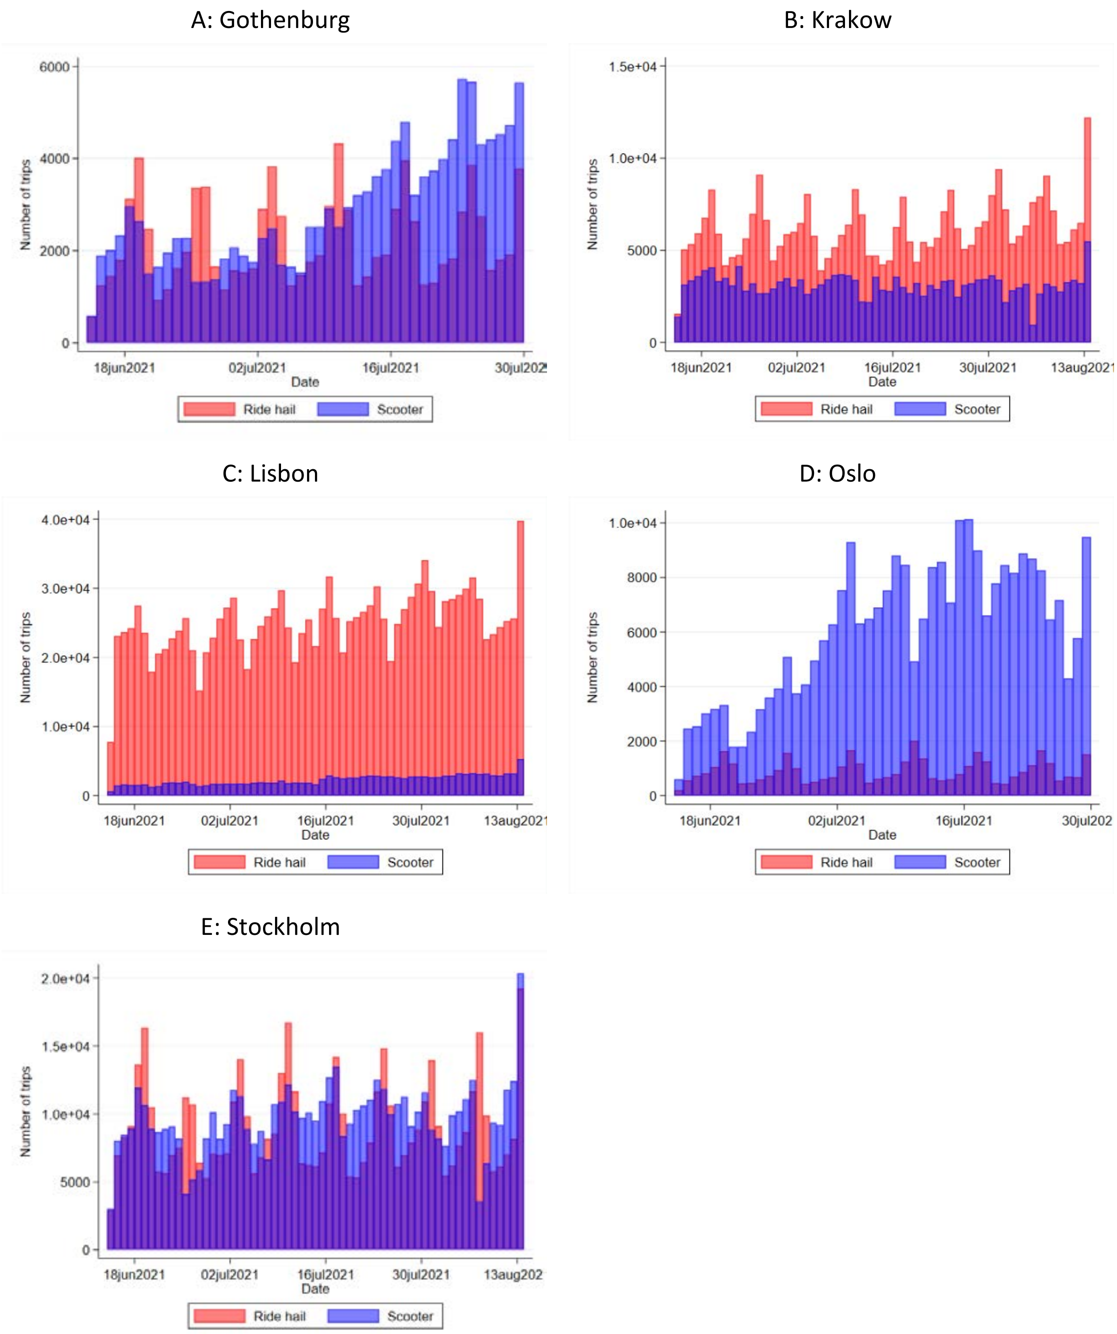

Figure 3.2: Overlapping histograms of ride hail and e-scooter trips mid-June to mid-August.

The saw tooth pattern for ride hail trips reflects that these trips are more popular during weekends. See Appendix A2 for visual representations on how the demand varies over time of day for both weekdays and weekends. The upward sloping trend for e-scooter trips in Gothenburg, Lisbon and Oslo is likely affected by seasonal variation and a general increase in the awareness and popularity of e-scooters throughout the period. For Gothenburg and Oslo, price changes might also have played a role. In the period of June 26th until July 18th Bolt conducted a campaign in Oslo where the price of e-scooters was cut by 90 percent. When the price was increased again, it was set to 50 percent of the initial value. In Gothenburg, the price was cut by 33 percent on July 7th. In remaining cities however, the price has been the same throughout the duration of the experiments.

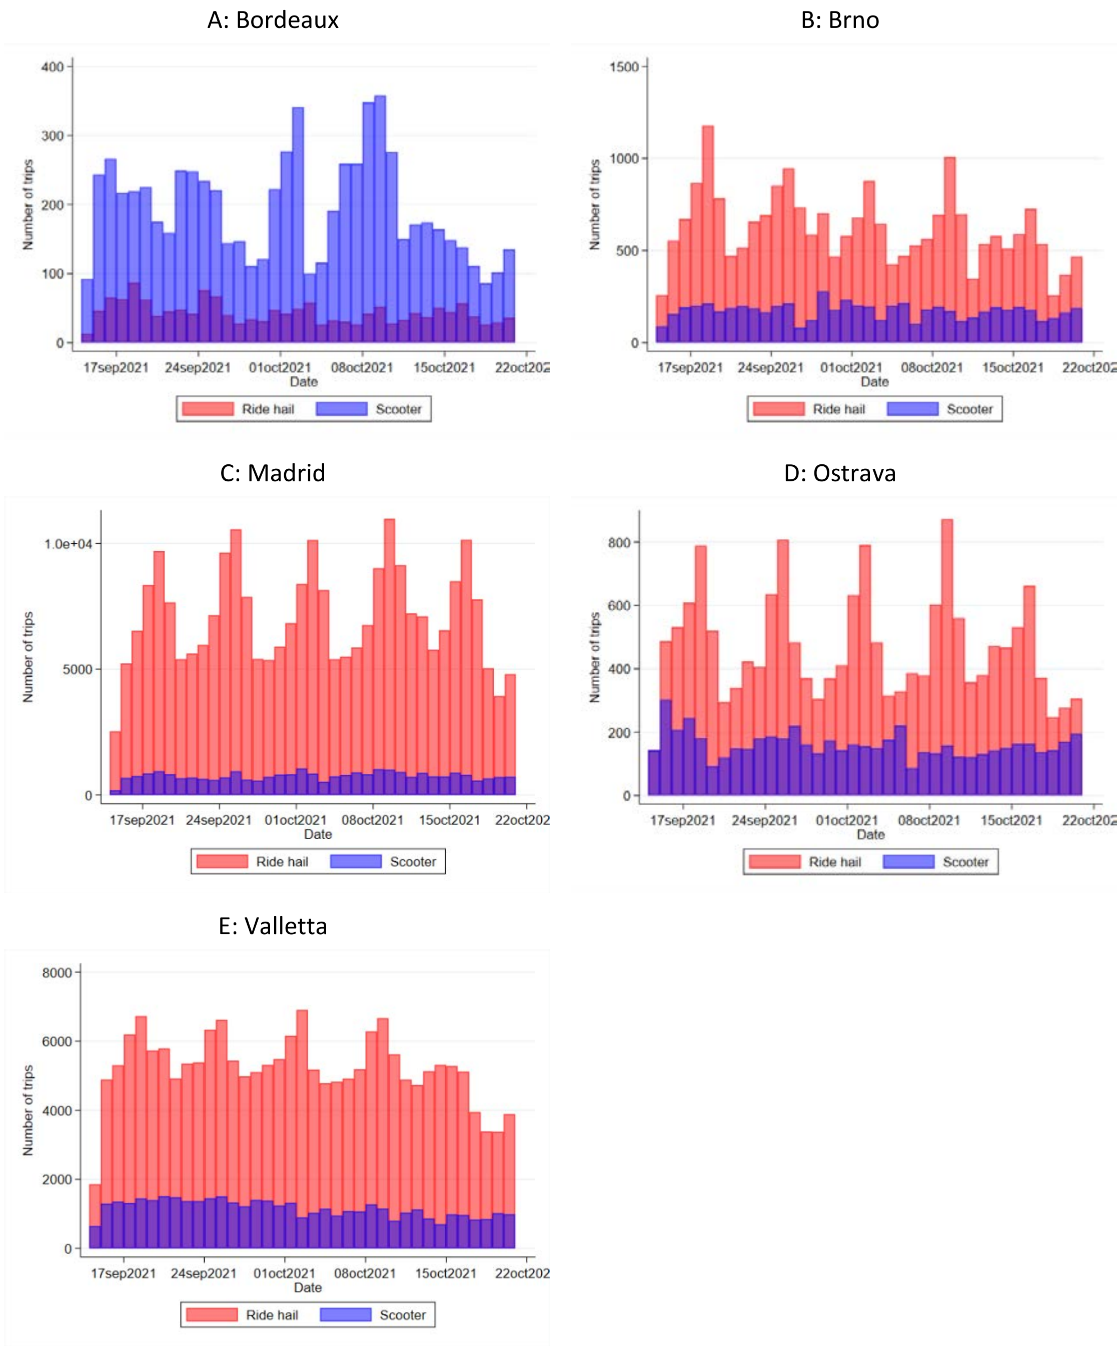

Figure 3.3: Overlapping histograms of ride hail and e-scooter trips mid-September to mid-October.

3.3.2 Relevant search sessions

As indicated by Table 3.2, the total number of search sessions is more than 12.6 million. However, there is only a small fraction of these sessions in which users were eligible for being nudged. Table 3.3 indicates how each of the search criteria affects the number of relevant search sessions. By "relevant" we refer to sessions where the users allocated to the treatment group actually received the nudge treatment. The first column displays *all* sessions; the second shows only sessions that included a ride hail search; the third column limits to sessions where e-scooters were available within 300 meters; and the final column removes user sessions where the trip distance is more than 3 kilometers (2 kilometers in Bordeaux, Brno, Madrid, Ostrava, Valetta).

Table 3.3: Number of search sessions by selection criteria for the nudge experiment. Control and treatment group.

Experiment	All sessions	Only ride hail searches	E-scooter is available within 300 meters	Trip distance within 2/3 km
Bordeaux	58,469	34,816	8,961	610
	100%	59.5%	15.3%	1.0%
Brno	109,819	89,861	34,699	3,124
	100%	81.8%	31.6%	2.8%
Gothenburg	702,203	413,067	203,809	26,386
	100%	58.8%	29.0%	3.8%
Krakow 1	785,939	562,422	403,353	50,718
	100%	71.6%	51.3%	6.5%
Krakow 2	755,513	545,102	390,037	48,147
	100%	72.1%	51.6%	6.4%
Lisbon 1	1,761,582	1,609,475	491,369	112,607
	100%	91.4%	27.9%	6.4%
Lisbon 2	2,269,242	2,044,104	616,579	132,648
	100%	90.1%	27.2%	5.8%
Madrid	1,695,366	1,608,007	475,626	34,136
	100%	94.8%	28.1%	2.0%
Oslo	752,616	226,672	145,293	30,454
	100%	30.1%	19.3%	4.0%
Ostrava	84,812	68,978	35,317	2,912
	100%	81.3%	41.6%	3.4%
Stockholm 1	1,510,822	1,016,852	362,580	41,747
	100%	67.3%	24.0%	2.8%
Stockholm 2	1,576,993	984,540	393,286	47,123
	100%	62.4%	24.9%	3.0%
Valletta	602,191	487,634	285,491	36,749
	100%	81.0%	47.4%	6.1%
Total	12,665,567	9,691,530	3,846,400	567,361
	100%	76.5%	30.4%	4.5%

The result of the selection process is that the number of *relevant* sessions is 567,361 (4.5 percent of the full sample). The various criteria have different effects on sample sizes in different cities. In Oslo, where ride hail is less popular, limiting to sessions including ride hail searches drops the sample to 30 percent. 2/3 of these searches had e-scooters available within 300 meters. In Lisbon however, focusing on ride hail searches keeps 90 percent of the sessions, but only 1/3 of these sessions had e-scooters available nearby. The most restrictive criterion is the trip length, as ride hail trips typically are longer than 2/3 kilometers.

Table 3.4 focuses on the search sessions from the last column in Table 3.3, and shows how outcomes of relevant search sessions are distributed between “e-scooter trip”, “ride hail trip” and “no trip” for the control and treatment group, respectively. “No trip” means that the app is killed during a user session without either an e-scooter or ride hail trip being made. For the remainder of this report, we will refer to these sessions as “relevant search sessions”.

Table 3.4: Number of relevant search sessions by treatment status, outcome and experiment.

Experiment	Control			Treatment		
	E-scooter	Ride hail	No ride	E-scooter	Ride hail	No ride
Bordeaux	49	8	245	51	5	252
	16.23%	2.65%	81.13%	16.56%	1.62%	81.82%
Brno	12	461	1,067	26	418	1,140
	0.78%	29.94%	69.29%	1.64%	26.39%	71.97%
Gothenburg	651	6,576	11,355	423	2,765	4,616
	3.50%	35.39%	61.11%	5.42%	35.43%	59.15%
Krakow 1	462	16,480	23,570	192	4,059	5,955
	1.14%	40.68%	58.18%	1.88%	39.77%	58.35%
Krakow 2	188	11,459	12,475	299	11,296	12,430
	0.78%	47.50%	51.72%	1.24%	47.02%	51.74%
Lisbon 1	147	54,895	35,134	111	13,510	8,810
	0.16%	60.88%	38.96%	0.49%	60.23%	39.28%
Lisbon 2	167	40,121	26,634	444	38,731	26,551
	0.25%	59.95%	39.80%	0.68%	58.93%	40.40%
Madrid	62	3,456	13,362	91	3,406	13,759
	0.37%	20.47%	79.16%	0.53%	19.74%	79.73%
Oslo	1,482	5,567	14,116	923	2,298	6,068
	7.00%	26.30%	66.70%	9.94%	24.74%	65.32%
Ostrava	3	445	1,050	9	435	970
	0.20%	29.71%	70.09%	0.64%	30.76%	68.60%
Stockholm 1	866	13,812	18,603	302	3,500	4,664
	2.60%	41.50%	55.90%	3.57%	41.34%	55.09%
Stockholm 2	435	10,497	12,803	533	9,919	12,936
	1.83%	44.23%	53.94%	2.28%	42.41%	55.31%
Valletta	106	8,912	9,286	117	8,870	9,458
	0.58%	48.69%	50.73%	0.63%	48.09%	51.28%
Total	4,630	172,689	179,700	3,521	99,212	107,609
	1.30%	48.37%	50.33%	1.67%	47.17%	51.16%

Note: The table displays the outcome of all "relevant search sessions", i.e. sessions where the user searched for a ride hail trip that was less than 2/3 kilometers and where an available e-scooter was located less than 300 meters from the user's search location. This is the set of search sessions in which the treated users were nudged.

The table above illustrates that some experiments have very few observations for certain modal choices. This is the case for Bordeaux, Brno and Ostrava in particular, and to some extent for Madrid and Valletta. There are two reasons for this: first, a relatively small share of search sessions is relevant for the nudge experiment in the first place, as shown in Table 3.3. Second, of the relevant search sessions, only a small share of individuals in both the treatment and the control group ends up choosing e-scooter as an option. For example, even if three times as many relevant user sessions in the treatment group in Ostrava ended up choosing e-scooter compared to the control group, the numbers (9 and 3 respectively) are too small for statistical inference. Therefore, Bordeaux, Brno and Ostrava will be excluded from the discussion from this point onwards. To illustrate the competitiveness of e-scooter relative to ride hail in each city, Table 3.5 displays averages of relevant search session characteristics for the control group. See Appendix A3 for how these characteristics are distributed within cities.

Table 3.5: Average characteristics of relevant search sessions, by experiment. Only control group.

	Relevant search sessions per user	Predicted trip distance (km)	Distance to nearest e-scooter (m)	Time until driver arrives (min)	E-scooter price as share of ride hail price
Gothenburg	1.97	2.09	91.4	3.05	0.14
Krakow 1	2.60	2.21	109.7	3.73	0.48
Krakow 2	2.39	2.22	112.6	3.00	0.46
Lisbon 1	3.57	2.15	135.5	2.41	0.21
Lisbon 2	3.68	2.13	118.8	2.47	0.21
Madrid	2.58	1.52	159.8	3.61	0.17
Oslo	2.49	2.08	56.4	3.61	0.05
Stockholm 1	1.92	2.11	57.7	2.56	0.19
Stockholm 2	1.97	2.09	52.0	2.37	0.13
Valletta	3.67	1.45	124.6	2.01	0.34
Total	2.76	2.08	108.5	2.76	0.24

The number of relevant search sessions per user is clearly highest in Lisbon. This is not surprising since relevant search sessions must include ride hail searches, and Lisbon has the most ride hail activity relative to e-scooters. Still, the average number of search sessions per user is relatively similar across cities and experiments. This is also the case for predicted trip distance, the outliers being Madrid and Valetta since relevant search sessions in these cities were required to be less than 2 kilometers. The distance to nearest e-scooter can be seen as an indication of e-scooter availability/density in each city. This variable is significantly lower for Oslo and Stockholm compared to other cities, indicating that e-scooter availability is higher. Similarly, the time until the driver arrives is informative for the availability of ride hailing services. This is highest for Oslo and Madrid and lowest for Valletta, Stockholm and Lisbon, but the variation across cities is smaller than for distance to the nearest e-scooter. The last variable is the e-scooter price as share of the ride hail price. If this value is high, e-scooters are more expensive compared to ride hail trips – a value of one would mean that the prices are equal. Differences are driven by three main factors: the distance of the trip (column 2), the e-scooter and ride hail rates in each city, and the "surge multiplier" for ride hail. This is a factor that increases the prices of ride hail trips in periods of high demand relative to supply, which is typically late at night. When it comes to price, e-scooters are less competitive compared to ride hail in Krakow and Valletta than the other cities. However, the main outlier is Oslo where the average ride hailing price per trip is 20 times higher than the e-scooter price. There are two reasons for this. First, due to higher wage levels, the ride hail price is higher in Oslo than in other cities. Second, Bolt conducted a campaign in Oslo during parts of the experiment period, where e-scooter prices were significantly lower than in other cities.[10]

[10] Interestingly, the low e-scooter price in Oslo did not seem to have a major impact for the main results. Figure B1.8 in Appendix B2 shows that the average probability of choosing e-scooter was about the same for both the treatment and the control group during a time period when the price was five times higher.

4 Results

This chapter presents the results of the ten RCTs conducted in Gothenburg, Krakow, Lisbon, Madrid, Oslo, Stockholm and Valletta. The cities Bordeaux, Brno and Ostrava are excluded from all analyses, since the number of observations is not high enough, as described in Section 3.3.2.

First, we focus on the effect of nudging on user behavior in a relevant search session in Section 4.1. Average treatment effects here are interpreted as the direct effect of nudging on the outcome probabilities in that user session, ignoring spillover effects to subsequent trips. In Section 4.2 we estimate average treatment effects on outcomes aggregated by user IDs, to test whether the treatment group initiated more or fewer searches and trips in total, compared to the control group.

Average treatment effects are estimated by means of simple regression models, as outlined in Section 2.2. All regression models are also repeated with multiple additional control variables, to check whether average treatment effects are different when they are estimated conditional on other factors. As expected, results are nearly identical.[11] Therefore, the more complex regressions are left out of the document.

4.1 Relevant search sessions

In this section we attempt to shed light on how users react to being nudged by restricting the sample to relevant search sessions, as displayed in Table 3.4. The outcomes of interest are the potential outcomes of a search session, which can either be “e-scooter”, “ride hail” or “no choice”. Section 4.1.1 estimates average treatment effects. Section 4.1.2 discusses to what extent this treatment effect varies across dimensions of the observed data by estimating heterogeneous treatment effects. Finally, in Section 4.1.3 we focus on data from Oslo and estimate a discrete choice model for relevant search sessions, where “being nudged” is one of the explanatory variables.

4.1.1 Average treatment effects

E-scooter trips

Figure 4.1 displays, for each experiment, the average number of relevant search sessions in the treatment and control group that resulted in e-scooter trips. The capped spikes (i.e. the black lines) represent 95 percent confidence intervals of the sample means.

[11] Since assignment of treatment status is randomized, the treatment status should not be correlated with other potential explanatory variables – this is by design, and one of the virtues of randomized control trials. Hence, the estimated effect of being nudged should not be affected by the inclusion of other control variables in the regressions. The more complex regressions confirm that this is the case, which substantiates that the assignment process of treatment status is satisfactorily random.

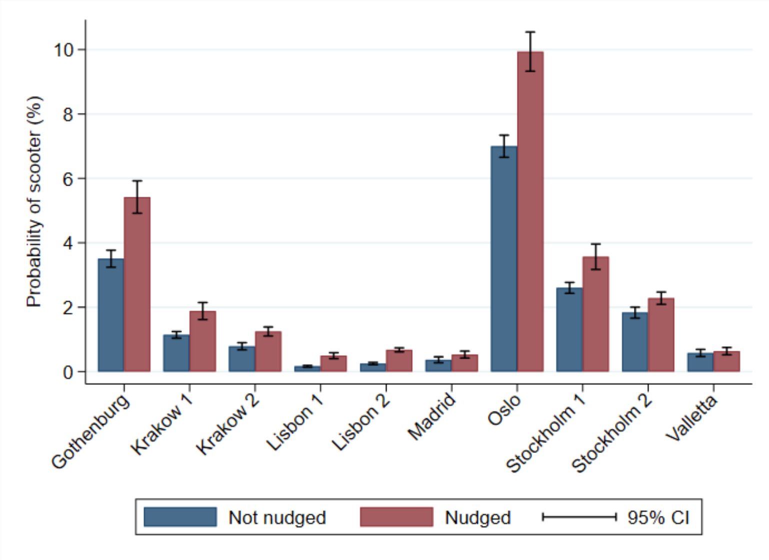

Figure 4.1: The average probability of e-scooter in a relevant choice situation for the treatment and the control group, including 95 percent confidence intervals.

The figure illustrates that the e-scooter probability varies significantly across experiments, from about 0.2 percent in Lisbon to 7 percent in Oslo for the control group (the blue bars). The difference between the blue and the red bar is the average treatment effect, i.e. the effect of being nudged. This effect is greatest in Oslo (a change of almost 3 percentage points) and Gothenburg (almost 2 percentage points), while for other experiments nudging changes the choice probability between 0 and 1 percentage point.

Panel A of Figure 4.2 displays the effect of being nudged in percentage points, i.e. the average difference between the treatment and the control group, including 95 percent confidence intervals. Panel B displays the relative change in the treatment group as a percent of the level in the control group.

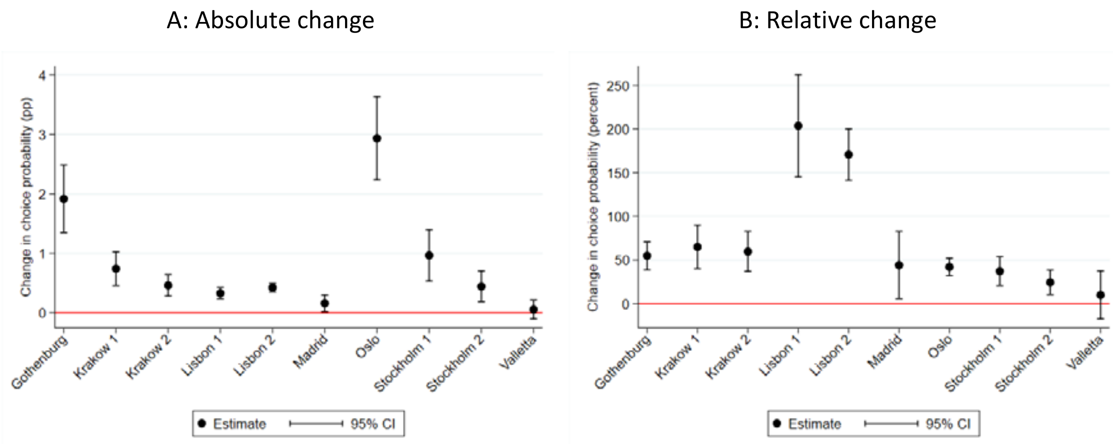

Figure 4.2: Effect of nudging on probability of choosing e-scooter trip. City-specific OLS regressions.

The two panels display the same effect, but in two different metrics. To explain the difference between these metrics more clearly, we will use Oslo as an example. The share of users choosing e-scooters in Oslo is 7 percent in the control group and 10 percent in the group that is nudged, as displayed in Figure 4.1. This means that nudging increases the probability of e-scooter by (10-7=) 3 percentage points, which is what Panel A shows. Another way to display the same effect is relative to a baseline. In this case the natural baseline is the share of users in the control group that chose e-scooters, i.e. 7 percent. Panel B shows that nudging increases e-scooter utilization by (3/7=) 43 percent relative to this baseline. In other words, users that are nudged are 43 percent more likely to choose e-scooters compared to users in the control group.

All estimates are positive and, except for the Valletta experiment, significantly different from zero (i.e. the value of zero is not covered by the confidence intervals). According to Panel B, nudging increases the share of users choosing e-scooters in the remaining cities by 40-60 percent. The exception is the first and the second experiment in Lisbon: Even though the effect is small in percentage points (Panel A) the relative effect of nudging in Panel B is 206 and 172 percent, respectively. This means that for the "Lisbon 1" experiment, being nudged makes selecting an e-scooter more than three times as likely.

There is a clear pattern to how the absolute effect of being nudged varies. To illustrate this, Figure 4.3 plots the estimated effects from Panel A, Figure 4.2 along the y-axis, and the share of search sessions resulting in e-scooters for the control group along the x-axis.

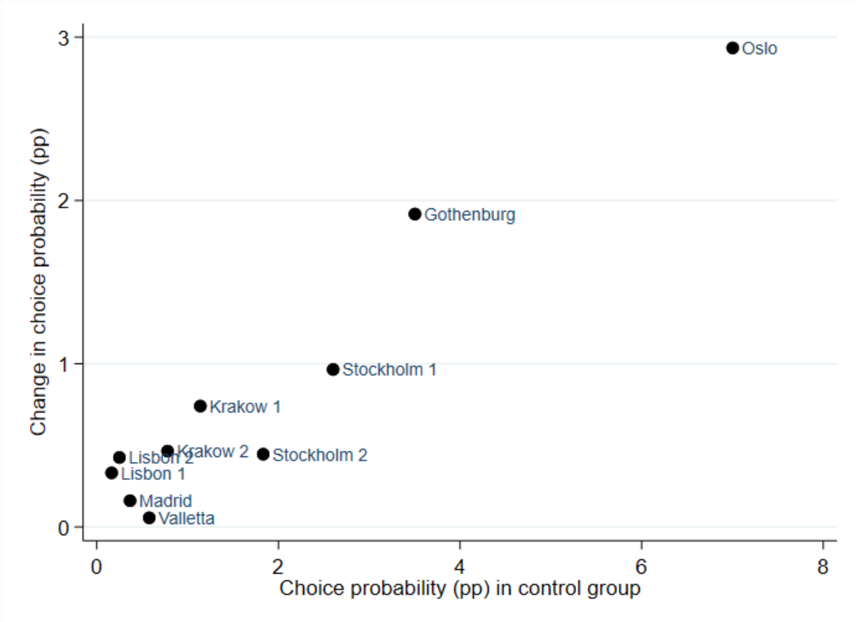

Figure 4.3: Probability of choosing e-scooter. Relationship between the effect of nudging and baseline choice probability.

The figure illustrates that there is a positive and close to linear relationship between the two – in cities where e-scooters are more popular, the average effect of nudging (in percentage points) is larger.

Ride hail trips

The analysis above shows how nudging affects the probability of choosing an e-scooter. However, we would also like to know whether the additional e-scooter trips come in addition to or replace ride hail trips. Figure 4.4 displays how the share of relevant search sessions resulting in ride hail trips vary across experiments depending on whether the user is in the treatment group (nudged) or in the control group.

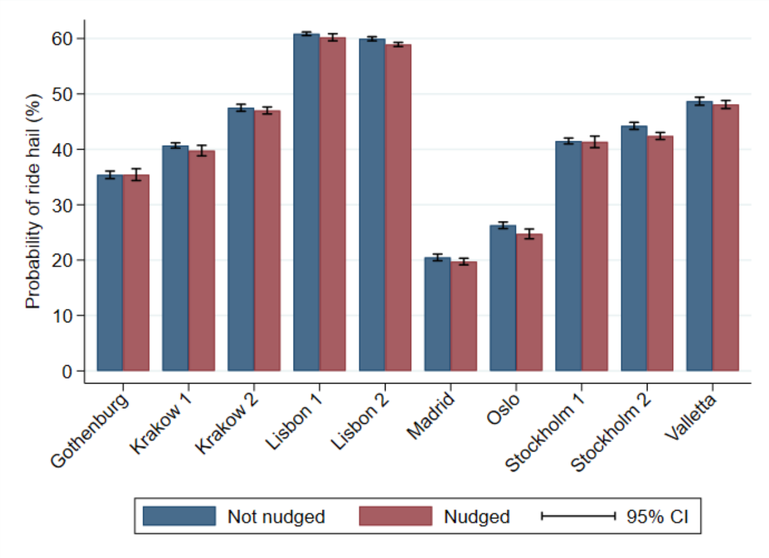

Figure 4.4: The average probability of ride hail in a relevant choice situation for the treatment and the control group, including 95 % confidence intervals.

Again there are large differences across cities, Lisbon being the city in which most of the search sessions result in ride hail trips (about 60 percent). This is less than half in Madrid and Oslo. Figure 4.5 repeats the same analysis as Figure 4.2, but with ride hail as the outcome.

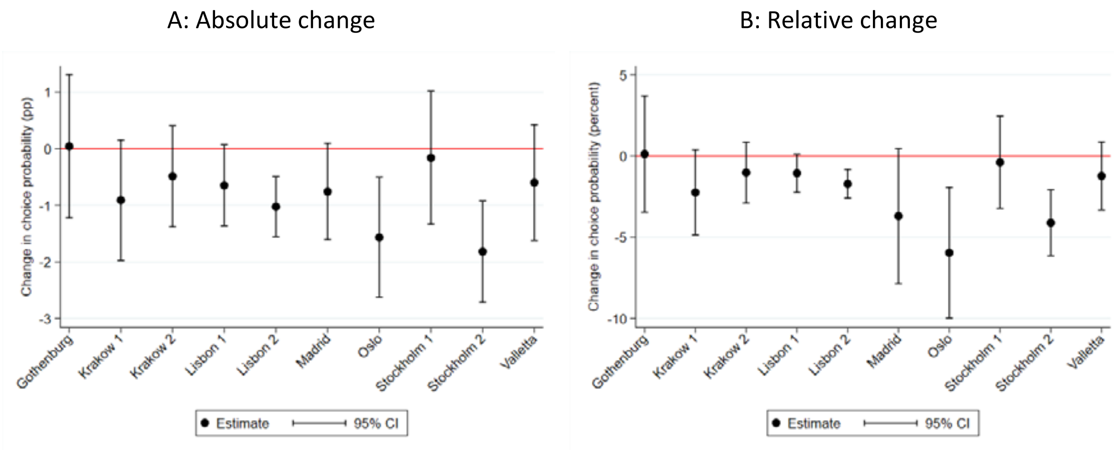

Figure 4.5: Effect of nudging on probability of choosing ride hail. City-specific OLS regressions.

All average treatment effects shown in Figure 4.5 are negative, except for Gothenburg where it is precisely zero. Still, the effect of nudging on ride hail trips is more noisily estimated, resulting in larger confidence intervals – only the effect in Oslo and the second round of experiments in Lisbon and Stockholm are significantly smaller than zero.

For Oslo the point estimate in Panel A is minus 1.6 percentage points, indicating that more than half of the additional e-scooter trips replaced ride hail trips. For Gothenburg on the other hand, the point estimate indicates that there is no change in the number of ride hail trips even though e-scooter trips increased by almost 2 percentage points. In other experiments such as Lisbon, the reduction in ride hail probability is larger than the increase in e-scooter probability. This means that a larger share of the nudged users closed the search session without ordering ride hailing or booking e-scooters, meaning that nudged users had a lower chance of making a trip. This is shown more clearly in the next section.

Although estimates from Panel A are similar in size compared to the e-scooter outcome, the relative changes estimated in Panel B are much lower. This is because the share of users choosing ride hail trips in relevant search sessions is much higher than the share of users booking e-scooters.

No trip

Finally, we present the effect of nudging on the probability of neither choosing an e-scooter nor ride hail.[12] Figure 4.6 shows that the probability of an unsuccessful search session is smallest in Lisbon and largest in Madrid, but due to the scale of the figure it is difficult to spot the exact differences within experiments.

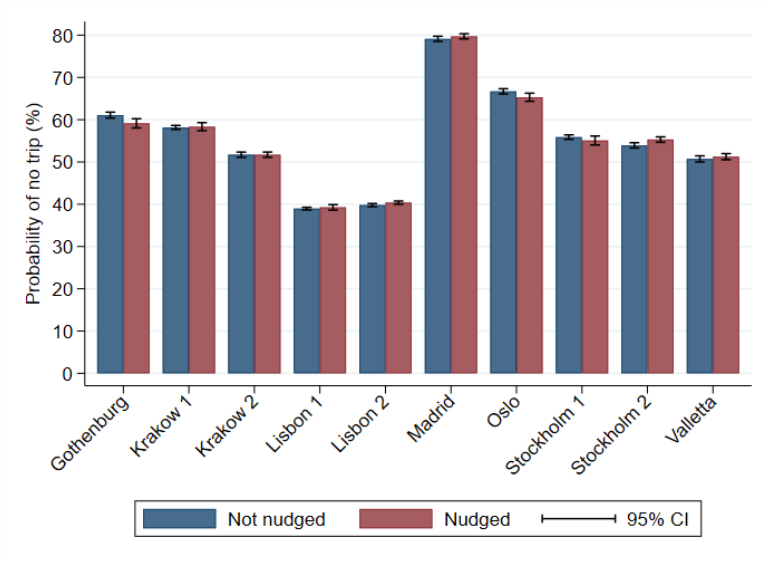

Figure 4.6: The average probability of no trip in a relevant choice situation for the treatment and the control group, including 95 % confidence intervals.

Figure 4.7 displays these differences and their confidence intervals more clearly. When these estimates are negative, it means that the nudging experiment increases the probability of a trip (either e-scooter or ride hail) taking place.

[12] Since the three probabilities for “scooter”, “ride hail” and “no trip” must add up to one, results in this section can in principle be inferred by the two previous sections.

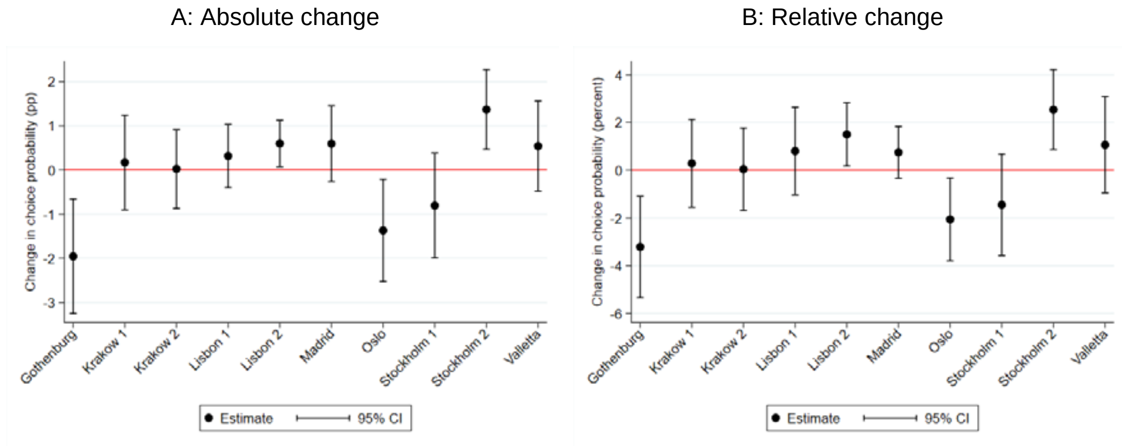

Figure 4.7: Effect of nudging on probability of no trip. City-specific OLS regressions.

The nudging experiment significantly increases the amount of information regarding availability to the user, but slightly reduces the amount of information regarding ride hail (two instead of three ride hail alternatives are visible on the screen without having to scroll down). Whether the increased e-scooter information is more beneficial in order to make a successful trip will vary across users – some will probably never consider an e-scooter trip in the first place, and therefore have no added value of the e-scooter information. As Figure 4.7 indicates, the average effect also varies across different experiments – in some experiments, nudged users are more likely to end their search session in a trip, while in other experiments nudging makes them more likely to close the app.

The two experiments where point estimates are negative and significant (nudging is increasing the chance of a trip) are Oslo and Gothenburg. These are also the experiments with the strongest effects of nudging on the probability of taking an e-scooter (Figure 4.2, Panel A). The two experiments where nudging significantly *reduces* the chance of a successful trip are the second round of experiments in Lisbon and Stockholm. Note that both of these experiments have a low probability of nudging individuals into e-scooters in the first place (less than 0.5 percentage points). There are two likely mechanisms: First, users who would not consider e-scooters in the first place might see the additional e-scooter information as an annoyance. Second, users that are nudged might miss out on the available ride hailing options if they do not scroll down in the app.

Summary

The above comparisons of estimates presented in this section are statistically uncertain, as indicated by the confidence intervals in the figures. In some cases, estimates from different experiments are not significantly different from each other, meaning that comparisons across experiments must be done with caution. Still, some clear patterns emerge:

a. The absolute effect of nudging is larger in cities where e-scooters have a larger market share and are more likely to be chosen in the first place. In Oslo, where the baseline probability of choosing an e-scooter is 7 percent, nudging increases this probability by 3 percentage points (40 percent);
b. It is difficult to know which city-specific factors are driving this result. However, the almost perfect linear relationship from Figure 4.3 indicates that the (observed and/or unobserved) factors driving the effect of nudging are also increasing the share of users that are choosing e-scooters in the control group;
c. In general, the relative effect of nudging is fairly stable, where nudged users are 40-60 percent more likely to choose an e-scooter compared to the control group. In

Lisbon however, the city where e-scooters are least popular compared to ride hailing amongst the control group, the relative effect is much larger. Only about 0.2 percent of relevant search sessions in the control group end up as e-scooter trips, and nudging increases this probability by a factor of three (the treatment effect is about 0.4 percentage points, corresponding to a 200 percent increase). This might indicate that nudging can be effective in spreading information and boost market shares in markets where users are less likely to know about e-scooters in the first place;

d. For experiments where e-scooter shares are high and nudging is effective (Oslo and Gothenburg), users are significantly less likely to end the search session without making a trip. This indicates that nudged users are more satisfied with their in-app options for these cities; and
e. For experiments in which e-scooter shares are low and nudging is less effective (the second round of experiments in Lisbon and Stockholm) users that are nudged are more likely to close the app without making a trip, presumably because a larger share of the users have no value of the additional e-scooter information.

4.1.2 Heterogeneous treatment effects

In Section 4.1.1 we estimated *average* treatment effects of being nudged by comparing (experiment-specific) average outcomes between the treatment and the control group in relevant search sessions. These average effects however are likely to mask differences in behaviors across individuals and situations. Some individuals are more likely to choose e-scooters, while others might never consider it. We have attempted to identify such differences by splitting the sample into sub-groups based on observable variables in the data and estimate treatment effects specific to the sub-groups. We focus on the main outcome, the effect of nudging on the probability of choosing an e-scooter. A comprehensive analysis is presented in Appendix B2, while the main findings are summarized here.

The analysis indicates that there are large differences across sub-groups for when a relevant search session results in an e-scooter trip, for both the treatment and the control group. The probability of choosing an e-scooter is increasing in "waiting time for ride hail driver" and decreasing in "distance to nearest e-scooter" and "trip distance" for most experiments, as expected. The e-scooter probability also varies with e.g. "time of day" and the (self-reported) age of the user, but often in different ways for different experiments.

In terms of e-scooter rates per hour there is little variation within experiments: the rate per hour is constant in most cities, with the main exceptions of Gothenburg and Oslo (see Figure B1.8 in Appendix B2). We do observe a reduction in the probability of choosing e-scooter when the price increases from 60 to 90 SEK in Gothenburg, and when the price increases from 30 to 60 NOK per hour in Oslo (1 SEK ≈ 1 NOK ≈ 0.1 EUR). However, this reduction is similar among both the treated and the non-treated users. Interestingly, Bolt conducted a campaign in Oslo for parts of the experiment when the price per hour of e-scooter rental was only 6 NOK. However, the share of users choosing e-scooters is about the same as when the price is five times higher.

Typically, the probability of choosing an e-scooter changes in similar ways for the treatment and the control group. This means that the effect of being nudged is relatively stable across groups for most of the dimensions considered. A general trend in line with Figure 4.3 however is that the effect of being nudged tends to be larger for sub-groups where the baseline probability of choosing e-scooter is higher. For example, for the Scandinavian cities (Gothenburg, Oslo and Stockholm) the e-scooter probability for the control group is higher in the afternoon, when e-scooters are more popular. This is also when being nudged changes the e-scooter probability the most.

The clearest pattern, true for every experiment, is that the e-scooter probability is higher if there are e-scooters nearby, both for the treatment group and the control group. The effect of nudging is also consistently higher in the proximity of e-scooters, for more or less all experiments. To illustrate this, Figure 4.8 repeats the "distance to e-scooter" results for Oslo and Lisbon. The left panels display the e-scooter probability for quintiles (five equally large groups) of "distance to e-scooter" for both the treatment and the control group. The first quintile has the shortest distance, while the fifth quintile has the longest distance. The right panels display the difference between treated users and the control group, which is the estimate of the causal effect of being nudged for each quintile. To illustrate changes across quintiles, the point estimates are connected by lines.

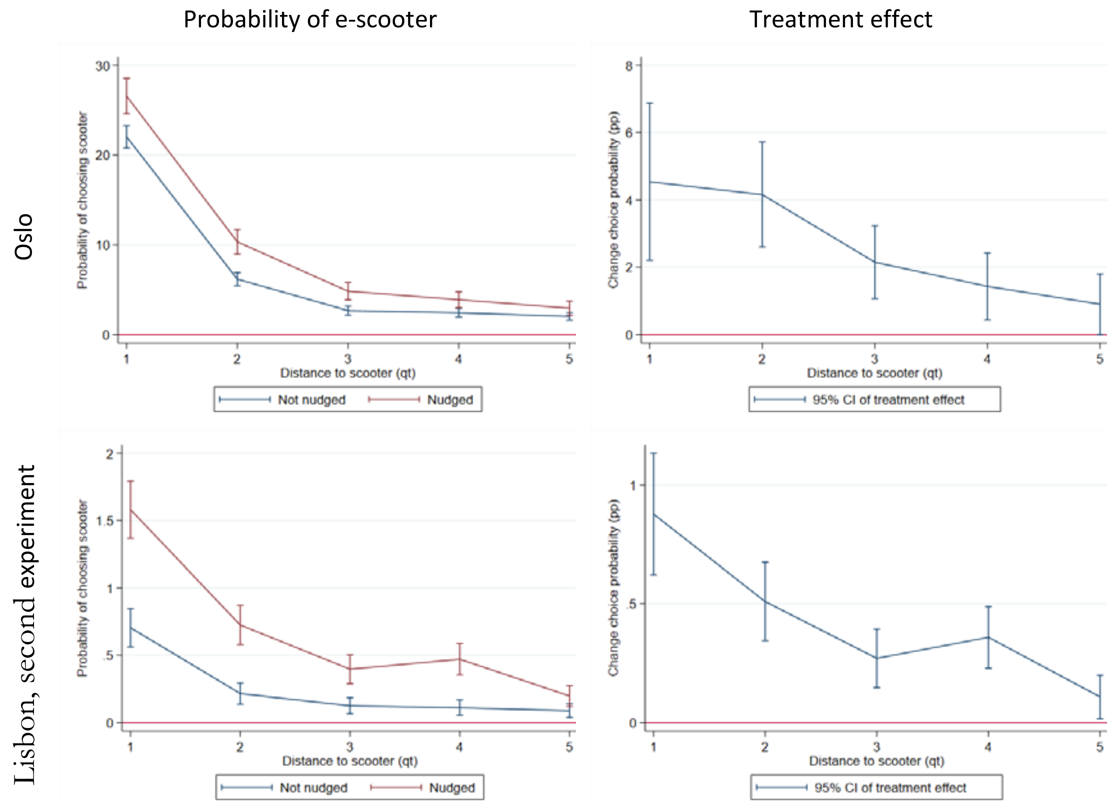

Figure 4.8: Effect of being nudged on probability of e-scooter by quintiles of dist. to nearest e-scooter.

While the average e-scooter probability in Lisbon is 0.68 (0.25) percent for the treatment (control) group (see Table 3.4), it is almost three times as high for the first quintile of "distance to e-scooter" (the left panel of Figure 4.8). For the last quintile it is about three times lower. While the average treatment effect for Lisbon is 0.43 percentage points, the treatment effect for the first quintile is almost twice as high (right panel). A conclusion that can be drawn from this is that for the user to consider e-scooters in the first place, it is vital that the visibility and density of available e-scooters is high enough.

4.1.3 A discrete choice model for Oslo

Motivation

Average treatment effects from randomized experiments, as presented in Sections 4.1.1 and 4.1.2, are most suitably analysed by means of linear regression analysis. However, as described in Section 2.3, linear models are less suitable for understanding how other (non-randomized) attributes impact the choice situation when the outcome variable is discrete rather than continuous. In this section we attempt to estimate conditional effects of different characteristics pertaining to the choice situation (a relevant search session) by means of a Multinomial Logit (MNL) model. The available choices each user is faced with are either “no trip”, “e-scooter trip” or “ride hail trip”. The model will let us compare how being nudged affects the choice probabilities relative to other attributes.

The discrete choice model is only estimated for the Oslo experiment. There are three reasons for this: first, Oslo is an interesting experiment to look further into because previous analyses show that the effect of nudging is most prominent there. Second, according to Table 3.4 Oslo has a reasonable number of relevant search sessions, with a high number of observations for each outcome, allowing us to leverage more variation from the data. Third, Oslo is one of the few cities with variation in Bolt’s e-scooter price per hour – in most other cities this rate is constant within experiments, meaning that we would not be able to isolate a price effect.[13] It should go without saying that results for Oslo are not directly generalizable to other cities, since responses vary significantly across experiments, as shown in Section 4.1.1. Since both the baseline e-scooter probability and the average treatment effect are higher in Oslo than in other cities, we would expect discrete choice model results from other experiments to be weaker/less in favor of e-scooters.

An alternative approach would have been to estimate a full model by pooling datasets from all or several experiments. This is not necessarily advisable because of the large differences in the popularity of e-scooters and ride hail across cities. These differences are likely affected by unobserved factors (e.g. differences in geography, climate, public transport availability, attitudes among users, competition by taxis or other e-scooter and ride hail companies, etc.). If these unobserved factors correlate with explanatory variables, they could potentially bias the results. By focusing on a single city we ensure that nudging is truly randomized and choice situations are comparable across observations, which makes interpretation more straightforward.

Model specification

Alternative specific utility functions are specified as linear in both parameters and variables. We let the “e-scooter trip” alternative depend on the following characteristics: (1) being nudged; (2) the price of the e-scooter trip; (3) distance to nearest e-scooter; and (4) trip distance. The “ride hail trip” alternative depends on (1) being nudged; (2) the price of the ride hail trip; (3) time until the driver arrives; and (4) trip distance. Furthermore, both alternatives depend on several time controls: (5) a linear time trend to take into account that the probability of both e-scooters and ride hail has increased over time, and (6) weekday/weekend specific time of day controls for the time periods 00:00-06:00, 06:00-

[13] Since the predicted price of an e-scooter trip is a function of the price per hour and the distance, the price and the distance of an e-scooter trip will be perfectly collinear unless there is variation in the hourly rate. This means that for most other cities than Oslo we would not be able to separate the effect of a longer trip from the effect of a higher e-scooter price.

12:00, 12:00-18:00 and 18:00-24:00, to allow for differences in demand at these time periods. We let “no trip” be the baseline category, with coefficients normalized to zero. See Appendix C.2 for the mathematical specification of the model.

The model is estimated for all observations where included variables are non-missing, and when the expected trip distance is larger than zero (i.e. the user has specified a destination other than his location in the ride hail menu of the app). This leaves 26,640 observations, 87.5 % of the relevant search sessions in Oslo (see Table 3.4). Note that we have omitted “age” as an explanatory variable, since age is missing for a large fraction (≈60%) of the data.

Results

The estimated coefficients are presented in Table 4.1. See Appendix C.2 for a technical description on how these coefficients relate to the alternative specific utility functions.

Table 4.1: Multinomial logit estimates for the outcome of relevant search sessions, Oslo.

	E-scooter		Ride hail	
	Coeff.	**Std. error**	**Coeff**	**Std. error**
Contextual variables				
Is nudged (0-1)	0.418***	(0.0521)	-0.044	(0.0320)
Price (EUR)	-0.533***	(0.0620)	-0.0594***	(0.00650)
Distance to e-scooter (meters)	-0.03258***	(0.00105)	-	-
Arrival time driver (minutes)	-	-	-0.104***	(0.00725)
Predicted trip distance (km)	-0.300***	(0.0382)	0.287***	(0.0255)
Time categories				
Weekend (0-1)	-0.223**	(0.097)	0.0718	(0.0770)
Weekday 00-06 (0-1)	-1.487***	(0.105)	0.0539	(0.0575)
Weekend 00-06 (0-1)	-2.373***	(0.128)	-0.0334	(0.0754)
Weekday 06-12 (0-1)	-0.207**	(0.101)	0.129	(0.0744)
Weekend 06-12 (0-1)	-0.628***	(0.165)	0.085	(0.105)
Weekday 18-24 (0-1)	-0.475***	(0.0689)	0.248***	(0.05209)
Weekend 18-24 (0-1)	-0.574***	(0.109)	0.153**	(0.07694)
Time trend (days)	0.00489**	(0.00208)	0.00482***	(0.00115)
Alternative specific constant (ASC)	0.638***	(0.105)	-0.607***	(0.0911)
Share of users:	8.6 %		28.8 %	
Number of observations:	26,640			

*Notes: * p<0.10, ** p<0.05, ***p<0.01. Robust standard errors in parentheses. “No trip” is the baseline alternative. For time variables, “weekday 12-18” is the base category.*

The absolute value of coefficients are difficult to interpret directly as they relate to the choice probabilities in a non-linear fashion. Coefficient values describe how much the utility of each alternative changes by a unit increase of the variable in question, meaning that positive coefficients are associated with higher utilities and therefore increased probabilities. Table 4.1 indicates that nudging increases the utility of e-scooters and reduces the utility of ride hail, although the latter effect is not significantly different from zero. All other variables have the expected sign and are statistically significant at the 99% level: E-scooter price, distance to nearest e-scooter and trip distance reduces the utility of e-scooter. E-scooters are also less likely to be chosen during nights and weekends. The utility of ride hail decreases in price and waiting time for the driver, but increases in expected trip distance. Ride hail is significantly more likely to be chosen in evenings (18:00-24:00) conditional on other factors. The time trend is positive for both ride hail and e-scooters, meaning that the probability of "no trip" decreases over time.

Even though the absolute value of coefficients are not directly interpretable, the relationship between coefficients are. The relationship between the nudge coefficient and the price coefficient for e-scooters is -0.78.[14] In other words, the model predicts that the utility of the e-scooter alternative will be the same when a user is nudged as when a non-nudged user gets a discount of 0.78 Euros, everything else equal. This indicates that even the small amount of additional information given to users through nudging has a tangible and economically significant value. Similarly, we calculate that the willingness-to-pay for having an e-scooter one meter closer to the search session is 6 cents, while the willingness-to-pay for the ride hail driver to arrive one minute earlier is 1.75 Euros.[15] To put this in context, the average relevant search session ending in a ride hail trip costs 10.38 Euros and lasts for 6.13 minutes, amounting to 1.69 Euros/minute.

The concept of willingness-to-pay from discrete choice models is frequently used to translate attributes of choice alternatives to monetary values. In the transport literature, the method is commonly used to estimate users' value of travel time savings (see e.g. Hess et al, 2005; Mackie et al., 2001). The fact that users' willingness-to-pay for these attributes appears to be high signifies that users in Oslo care strongly about distance to e-scooters and waiting times compared to price levels. However, the price estimate for e-scooters is relatively uncertain since there is little variation in the hourly rate (it is only observed at three different levels; see Figure B1.8 in Appendix B2). Moreover, estimates are based on a model specification where alternative specific utility functions are assumed to depend linearly on the variables in question. In reality, the effect of one additional Euro (or meter or minute) might vary depending on the base cost (or distance or time) as well. For instance, whether you are close enough to physically see the e-scooter might matter more than the exact number of meters.

Finally, Table 4.2 presents predictions from the choice model. The first row displays average choice probabilities for "no trip", "e-scooter" and "ride hail". A property of the MNL model is that these probabilities are equal to the share of observations that have chosen a particular outcome in the data. The subsequent rows display the average change in choice probabilities for the whole sample when changing each variable by one unit. Units are shown in column 1 while sample averages of the corresponding variable can be found in column 2.

[14] Calculation: $\frac{0.418 \text{ units of utility/nudge}}{-0.533 \text{ units of utility/EUR}} = -0.78 \text{ EUR/nudge}$.

[15] Calculations are done in the same manner as in footnote 14. Value of distance to e-scooter: -0.0326/-0.533 = 0.06. Value of waiting time for driver: -0.104/-0.0594 = 1.75.

Table 4.2: MNL model for Oslo. Baseline choice probabilities and average marginal effects in percentage points.

	Sample average	No trip	E-scooter	Ride hail
Baseline				
Average choice probabilities (%)		62.56***	8.604***	28.83***
		(0.301)	(0.166)	(0.284)
Marginal effects (pp)				
Is nudged (0-1)	30.2 %	-1.263*	3.063***	-1.800***
		(0.657)	(0.379)	(0.614)
Price, e-scooter (1 EUR)	0.53 EUR	2.469***	-3.601***	1.132***
		(0.287)	(0.418)	(0.133)
Price, ride hail (1 EUR)	10.72 EUR	1.055***	0.126***	-1.181***
		(0.115)	(0.014)	(0.129)
Distance to e-scooter (10 m)	55.7 meters	1.508***	-2.200***	0.692***
		(0.050)	(0.072)	(0.024)
Arrival time driver (1 min)	3.63 min	1.845***	0.221***	-2.066***
		(0.127)	(0.016)	(0.142)

*Notes: * p<0.10, ** p<0.05, ***p<0.01. Robust standard errors in parentheses.*

The treatment effect, i.e. the effect of being nudged, is estimated to increase the choice probability of e-scooters by three percentage points. As expected this corresponds to the average treatment effect estimated in Section 4.1.1. According to the predictions, about (-1.80/3.06=) 58 percent of the additional trips are diverted from ride hailing, while the remaining 42 percent would not have booked a trip through the Bolt app in the first place.

The table also indicates how marginal price increases affect choice probabilities (effects of price reductions will have the opposite sign). By reducing the price of e-scooters by 1 EUR (corresponding to 200 percent of the average price) we predict that the market share of e-scooters would increase by 3.6 percentage points. This is slightly higher than the effect of nudging. However, the substitution between e-scooter and ride hail caused by the price change is only (1.13/3.60=) 31 percent. By increasing the price of ride hailing by 1 EUR per trip on the other hand (corresponding to a 9 percent increase compared to the average ride hail price), the market share of ride hailing would decrease by 1.2 percentage points and the market share of e-scooters would increase by 0.13 percentage points: in other words, we predict that only (0.13/1.2=) 10 percent of the users that are deterred away from ride hailing would substitute towards Bolt's e-scooter service.

This shows that substitution from ride hailing to e-scooter is considerably higher in the case when the user is nudged compared to when prices change, because price increases typically push more customers to the "no trip" option. The caveat is that we do not know what the "no trip" users will do instead. They might desist from making the trip, choose to walk or take public transport, or order an e-scooter or ride hailing vehicle from a competing service.

4.2 Aggregate outcomes by user

In Section 4.1 we established that nudging has a significant effect on user behaviour when it comes to the outcome of the search session. However, nudging might have an additional impact if nudged users change their subsequent travel habits as a result. In this section we attempt to shed light on this by aggregating outcomes by users, and comparing these aggregated outcomes across the treatment and the control group.

We only consider users that have complied with the nudge criteria (i.e. experienced a relevant search session) at least once. This ensures that the users in the treatment group actually received some kind of treatment, and that the users in the control group constitute a comparable sample.[16] The number of observations (i.e. relevant users) in the treatment and the control group for each experiment can be seen in the last two columns of Table 3.1.

In this section we focus on three different outcomes: number of search sessions, number of trips and distance travelled for each user, over the duration of each experiment.

4.2.1 Number of search sessions per user

Did the nudged users end up using the Bolt app more or less as a result? Figure 4.10 shows the average number of search sessions per user, for both the treatment and the control group, throughout the duration of the experiment. In Figure 4.11, the average treatment effects are estimated. Even though the outcome variable is defined differently, average treatment effects are calculated in the same manner as in Section 4.1.1. Again, the average difference between the treatment and the control group is displayed in the left panels, while the relative change in the treatment group compared to the control group (percentage increase/decrease) is displayed in the right panels.

Figure 4.9 shows that there is a tendency of more e-scooter searches and fewer ride hail searches in some of the experiments, while Figure 4.10 displays the size of this difference and whether or not it is significantly different from zero.

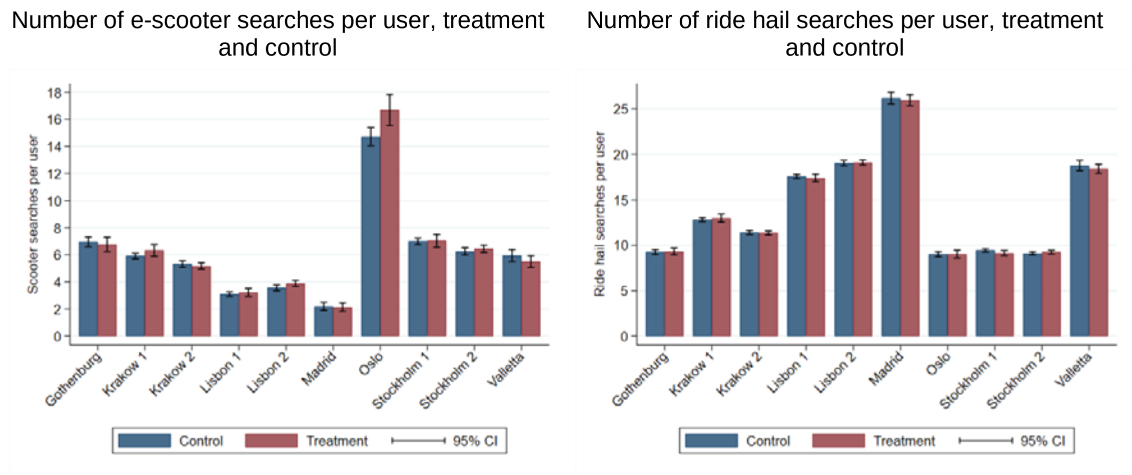

Figure 4.9: Average number of e-scooter and ride hail searches per user, treatment and control.

---

[16] As a check, we have tested whether we can estimate any significant difference in outcomes between members of the treatment and control group that have never been exposed to a search session relevant for nudging. As expected, we find no significant differences.

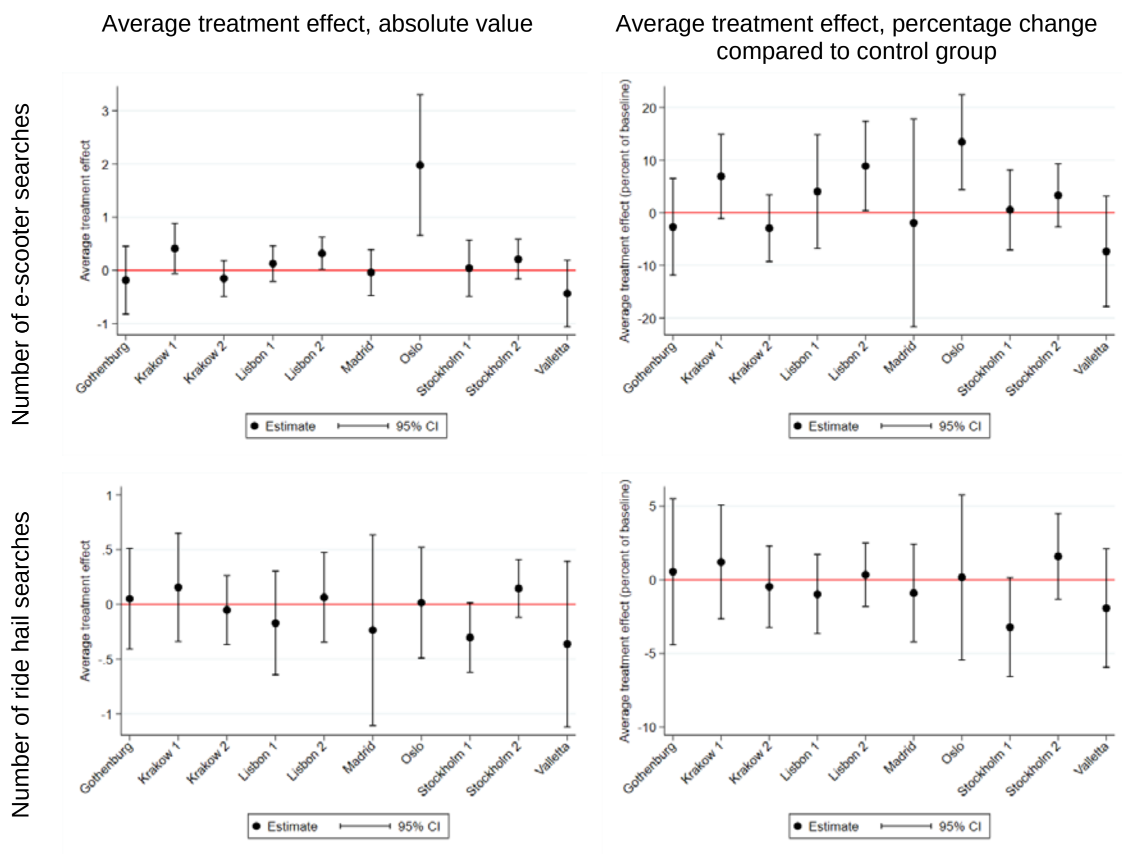

Figure 4.10: Average treatment effect, number of e-scooter and ride hail searches per user.

The change in ride hail searches per user is not significant for any of the cities. However, for Oslo and Lisbon 2 there is a significant increase in the number of e-scooter searches. Treated users initiated an additional 2 e-scooter searches in Oslo, and 0.3 in Lisbon. The right panels reveal that this corresponds to a percentage increase of 13 and 9, respectively.

To predict the number of additional e-scooter searches this amounts to, we can multiply the average treatment effect per user by the number of relevant users in the treatment group from Table 3.1 (3,637 relevant users in Oslo and 18,025 relevant users in Lisbon 2). We then find that the experiment caused more than 7,200 additional e-scooter searches in Oslo and more than 5,400 additional e-scooter searches in Lisbon. This corresponds to about 1% and 0.25% of the total number of searches in each of the cities, respectively (there were 2.27 million search sessions recorded in Lisbon and 753,000 in Oslo; see Table 3.2). While this prediction is uncertain as indicated by the confidence intervals in the figures above, it highlights that nudging has an additional effect beyond increasing the e-scooter probability for the specific search session in which the nudge occurred.

4.2.2 Number of trips per user

The same exercise is carried out in Figure 4.11-Figure 4.12 for the number of e-scooter and ride hailing trips per user.

Again, the figures indicate a tendency that treated users undertake more e-scooter trips and fewer ride hail trips than users in the control group, although effects are not statistically significant. Oslo is the only city in which we observe significantly more e-scooter trips among the treatment group: users that were exposed to nudging conducted in average 1.06 additional e-scooter trips throughout the experiment compared to the control group (left panel), corresponding to a 14 percent increase (right panel).

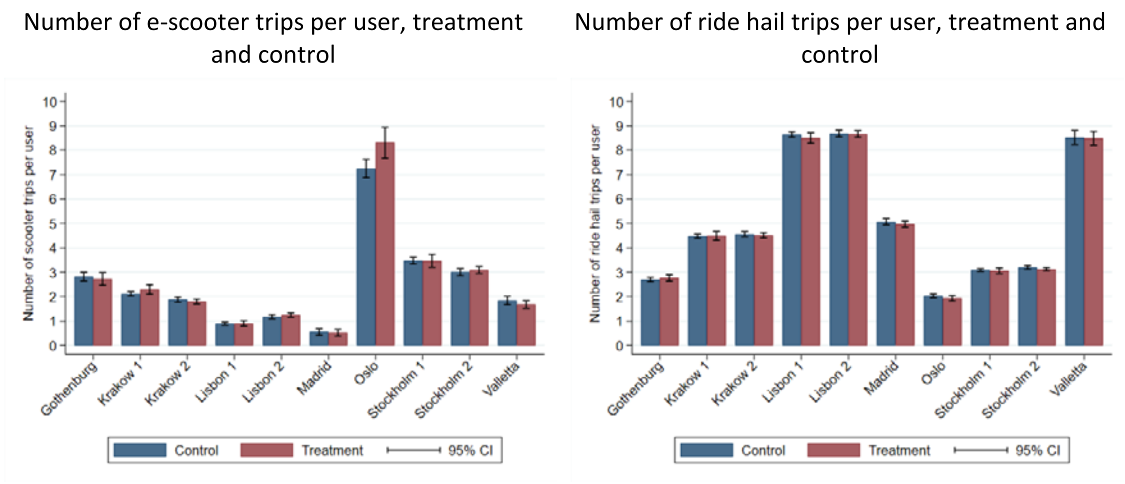

Figure 4.11: Average number of e-scooter and ride hail trips per user, treatment and control.

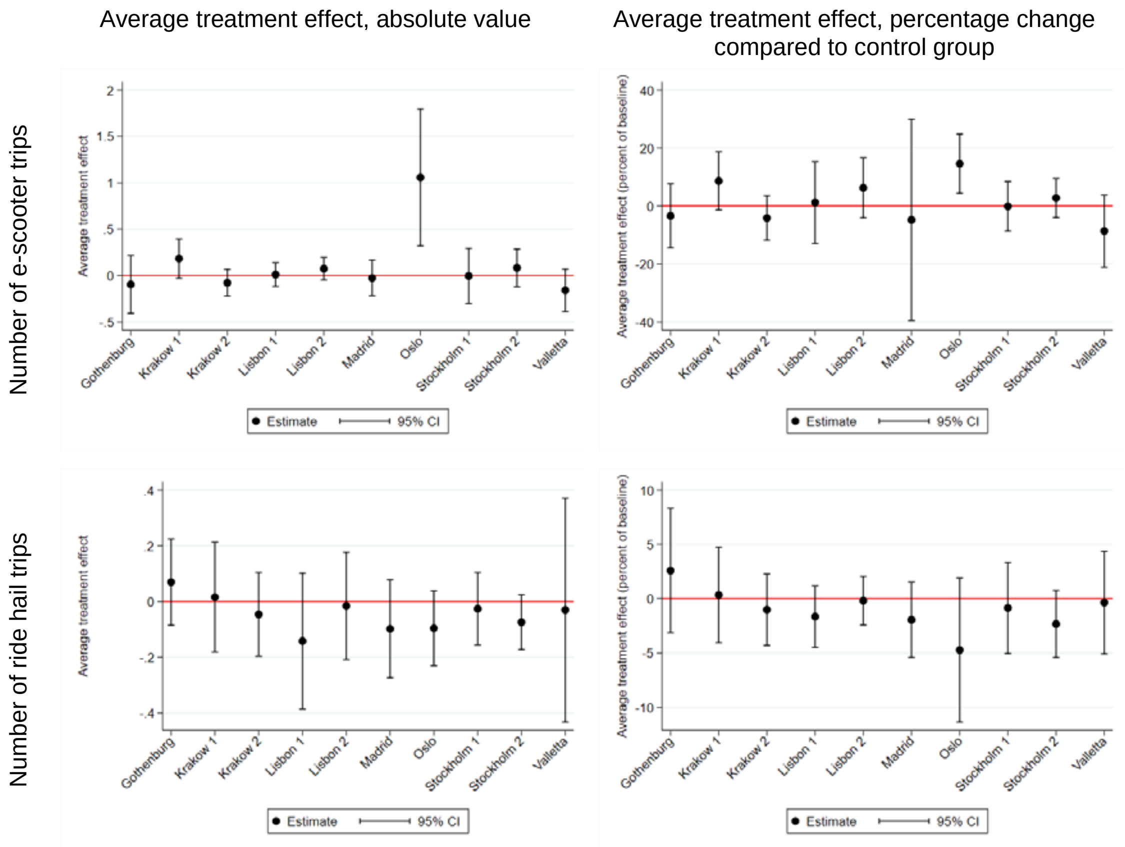

Figure 4.12: Average treatment effect, number of e-scooter and ride hail trips per user.

The effect in Oslo is not only statistically but also economically significant: Multiplied by the number of users that were nudged at least once, this amounts to about 3,800 additional e-scooter trips, or 1.3 percent of the 278,000 e-scooter trips undertaken in total throughout the experiment. Again, the fact that the confidence interval is large (between 0.3 and 1.8), means that the total effect is uncertain. We can do the same calculations for the additional number of e-scooter trips in Lisbon 2, with the caveat that it is not statistically significant (i.e. uncertain). Point estimates from Figure 4.12 reveal that users nudged at least once conducted on average 0.08 additional e-scooter trips (top left panel) over the whole duration of the experiment. This corresponds to an increase in the number of e-scooter trips of 7 percent (top right panel). Multiplied by the number of relevant users, it amounts to 1,400 additional e-scooter trips in Lisbon – about 1.6 percent of the total number of e-scooter trips recorded in total. This is a sizeable share of the e-scooter trips, considering that only about 10 percent of the total number of users in each city were nudged at least once.[17]

We know from Section 4.1.1 that nudging changed the choice probabilities for both ride hail and e-scooters for the relevant search sessions in several experiments. So why are the aggregate results in this section mostly insignificant? Looking at the whole duration of the experiment might be problematic, if the change in behaviour wears off over time (lower signal) or if a long time period amplifies variation in the data that is unrelated to the nudging (more noise). We do several robustness checks of this in Appendix B1.

First, we consider a different outcome: share of users that have conducted at least one e-scooter trip throughout the duration of the experiment. Figure B1.3 reveals that this vary from less than 10 percent in Madrid to more than 50 percent in Oslo, when considering relevant users in the control group. Figure B1.4 shows that the effect indeed is significantly greater than zero in all experiments except Krakow 2, Madrid, Stockholm 1 and Valletta. This illustrates that the experiments caused several users to try e-scooters for the first time. The effect is about 2 percentage points in Lisbon and almost 5 percentage points in Oslo, which constitutes an increase of about 10 percent compared to the control group.

Second, we attempt to reduce noise in the outcome variable by removing the 1 percent of users with the highest number of trips from the treatment and the control group, separately. As indicated by Figure B1.5 this reduces the average number of trips per user in both the treatment and the control group, since the remaining users are less active. As expected, Figure B1.6 shows that confidence intervals are more narrow, and the increase in number of e-scooter trips is significantly different from zero in both Krakow 1 and Lisbon 2 in addition to Oslo. Effects in Krakow and Lisbon are small in absolute value, but the relative effects are similar to Oslo – a 10 percent increase in the number of e-scooter trips compared to the control group.

Finally, in Figure B1.7, we measure average outcomes by user from the first search session she was nudged, for a time span of 7 days. Results confirm that by shortening the time period for analysis the statistical uncertainty decreases: we find a significant increase in both the number of e-scooter searches and e-scooter trips per user in Krakow, Oslo and Lisbon. Furthermore, we find a significant decrease among the treatment group for ride hail searches in Krakow and Stockholm, and a significant decrease in ride hail trips in Krakow and Oslo.

4.2.3 Distance travelled per user

Number of e-scooter trips and ride hailing trips are not necessarily directly comparable, since e-scooter trips tend to be shorter. Therefore, this section compares the treatment and the control group when it comes to distance travelled. Note that this is the actual distance recorded as opposed to the predicted distance based on the destination the user reports in the ride hail menu of the app. Results on e-scooter distance are only significantly different from zero for Oslo, where the experiment caused treated users to travel an additional 1.8 kilometres. However, the point estimates indicates similar relative increases of about 10 percent compared to the control group in Krakow 1 and Lisbon 2 as well. For ride hailing,

[17] In Oslo (Lisbon) there are 3,637 (18,025) relevant users in the treatment group, and 43,187 (167,968) recorded users in total (see Table 3.1).

the effect on distance travelled is largest for users in Oslo (a reduction of about 900 meters over the duration of the experiment), but barely insignificant at the 95 percent level.

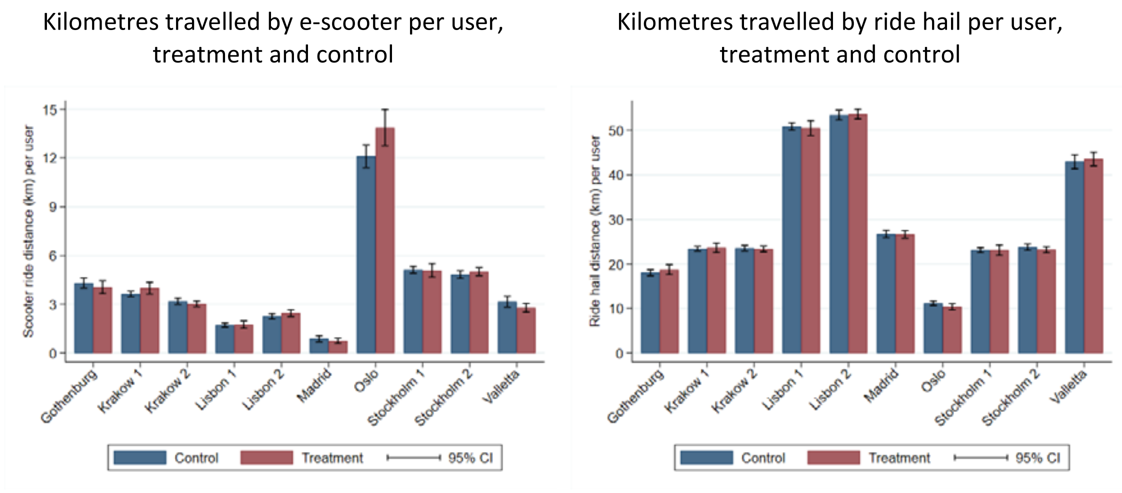

Figure 4.13: Average distance travelled by e-scooter and ride hail per user, treatment and control.

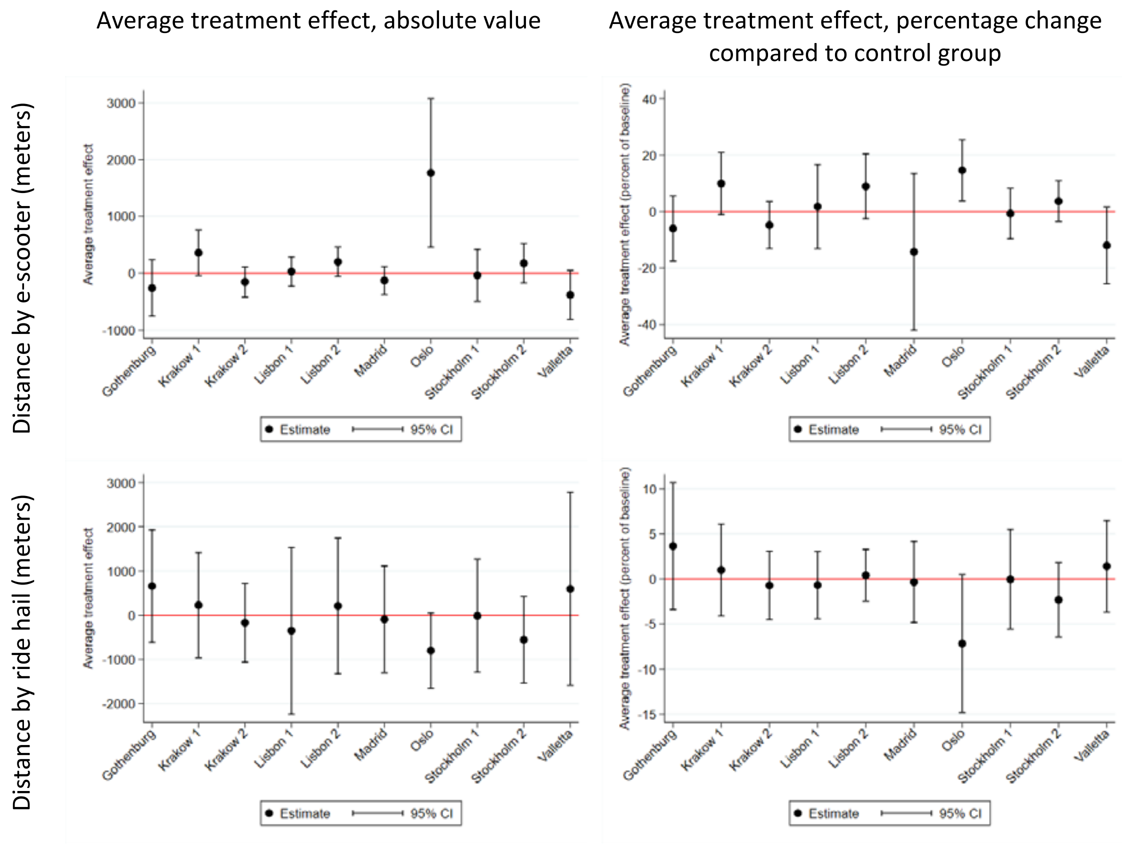

Figure 4.14: Average treatment effect, distance travelled by e-scooter and ride hail per user (meters).

Although several of the results reported in Section 4.2 are imprecisely estimated, it is reassuring to note that both search sessions, number of trips and distance travelled follow the same pattern across experiments, even though outcomes are defined differently.

5 Summary and conclusion

5.1 Summary and discussion

This section will summarize the main findings from the experiments, and discuss potential reasons for why results differ across experiments. The discussion will tend to focus on Oslo and Lisbon for two main reasons. First, Oslo and Lisbon have a large number of observations in each of the outcome categories: i.e. "e-scooter trip", "ride hailing trip" and "no trip", for both the treatment and the control group. This means that the effects of nudging in these experiments are more precisely estimated. Second, Oslo and Lisbon are the two experiments that differ the most in terms of user behaviour: 85 percent of trips booked through the Bolt app in Oslo are e-scooter trips, while more than 90 percent are ride hailing trips in Lisbon.[18] This makes for a more interesting comparison as to how the effect of nudging differs in different contexts.

The main search sessions analysed in this report are the relevant search sessions that satisfied the nudge criteria; i.e., sessions in which the user already has planned to conduct a short ride hailing trip with an e-scooter in close proximity. This is about 4.5 percent of the total number of search sessions, as indicated by Table 3.3. Among these search sessions it is unlikely that the user will end up choosing an e-scooter in the first place, even in cities where e-scooter conditions are favourable and the general e-scooter market share is high.

The share of relevant search sessions that end in e-scooter trips among the control group is 7 percent for Oslo and about 0.2 percent for Lisbon (see Table 3.4); i.e., relevant search sessions are 35 times as likely to end up as e-scooter trips in Oslo compared to Lisbon. For remaining cities, the e-scooter share is somewhere in between. This pattern corresponds well with whether conditions are more favourable for ride hail or e-scooters: for relevant search sessions, we observe that the distance to e-scooters is half as long and the waiting time for ride hail is 50 percent longer in Oslo compared to Lisbon (see Table 3.5). Most strikingly, in Oslo the average price of e-scooters relative to ride hail is four times larger.

The summary of the nudging experiment will be split according to two discussion points. First, whether nudging affects e-scooter use (Section 5.1.1). Second, to what extent the additional e-scooter trips replace ride hailing (Section 5.1.2).

5.1.1 Effect of nudging on e-scooter use

The direct effect of being nudged is found by considering outcomes of relevant search sessions. We find a large and significant increase in the share of these sessions resulting in e-scooter trips for all experiments except Valletta: typically, e-scooters are 40-60 percent more common in sessions that are nudged (right panel of Figure 4.2). The exception is Lisbon, where three times as many (200 percent) of the nudged sessions ended with e-scooter as the outcome. This indicates that the nudge can be a valuable instrument to

[18] See Figure 3.1, which displays number of e-scooter and ride hail trips per user for each experiment. In the Swedish cities Gothenburg and Stockholm, e-scooter and ride hail are about equally popular, while in remaining cities ride hail tends to be more popular than e-scooter.

spread information and boost market shares in markets where users are less likely to know about e-scooters in the first place.

In absolute terms (the share of relevant search sessions that are affected by the nudge) the effect varies notably across experiments. Typically, nudging has a larger impact where e-scooters are more likely to be chosen to begin with. In Oslo, where the share of relevant search sessions resulting in e-scooter trips is 7 percent in the control group, we find that nudging increases this by 3 percentage points, to 10 percent for the treatment group. In Lisbon on the other hand, where only about 0.2 percent of relevant search sessions result in e-scooter trips, the effect of nudging in absolute terms is an increase of about 0.4 percentage points. In the remaining cities, the effect is somewhere in between. It is difficult to know which (observed or unobserved) factors are driving the differences between cities. However, the fact that the effect of nudging on e-scooter use is more or less proportional to the share of users that end their search session with e-scooter in the control groups (see Figure 4.3) indicates that whichever factors that are making e-scooters more popular in general are also increasing the number of users that respond to nudging.

This also seems to hold *within* each city, as indicated by analyses presented in Section 4.1.1 and Appendix B2. Here, we attempt to see whether the effect of nudging in separate experiments varies along some observed dimension of the data. The general trend is that the effect of nudging is stronger in contexts where e-scooters are more popular among the control group in the first place.

The observable variable that affects the probability of e-scooter the most is "distance to nearest e-scooter". For every experiment, we find that the 20 percent of relevant search sessions in the control group where users are closest to e-scooters are two to three (or even more) times as likely to end with an e-scooter trip, compared to the average. Likewise, the effect of being nudged is two to three times higher. Considering that all relevant search sessions are within 300 meters of an available e-scooter in the first place, it signifies that visibility and density of e-scooters is of first-order importance in order for users to consider e-scooters as a viable option to other modes of transport.

The effect of nudging as well as the share of users choosing e-scooter in the control group is also decreasing in the distance of the trip and increasing in the waiting time for ride hail. Furthermore, nudging has a larger impact during time periods in which users are more likely to choose e-scooters – this is typically in the afternoon in the Scandinavian cities, but varies across the different experiments. Even though estimated effects along some of these dimensions are not significantly different in a statistical sense, the fact that the same patterns are observed in several experiments increases our confidence that they reflect patterns in user behavior, and are not differences by mere chance.

When only considering how nudging affects the outcomes of relevant search sessions, the numbers of additional e-scooter trips caused by the experiments are modest. We predict that the experiment caused 283 additional e-scooter trips for the second experiment in Lisbon and 273 additional e-scooter trips in Oslo.[19] However, Section 4.2 reveals that nudging has an effect on the number of e-scooter trips among the treatment group that

[19] Calculation, Lisbon 2: 65,726 relevant search sessions that were nudged times the effect of nudging (0.43 percentage points) is 283 additional e-scooter trips. Calculation, Oslo: 9,289 relevant search sessions that were nudged times the effect of nudging (2.94 percentage points) is 273 additional e-scooter trips. The numbers of relevant search sessions for each experiment are reported in Table 3.4. The effects are largest for Lisbon, since the number of relevant search sessions is higher than in other experiments, and for Oslo, since the effect of nudging is higher than in other city. For other experiments, the effects range from 10 additional e-scooter trips (Valletta) to 150 additional e-scooter trips (Gothenburg).

goes far beyond the *relevant* search sessions, since the nudge increases their subsequent number of e-scooter searches and trips.

Users in Oslo in the treatment group conducted on average 2 additional e-scooter searches and 1.06 additional e-scooter trips over the duration of the experiment compared to the control group (see Figure 4.12). This adds up to more than 7,200 searches and 3,800 trips in total. In Lisbon, 0.3 additional e-scooter searches and 0.08 additional e-scooter trips per user add up to 5,400 searches and 1,400 trips in total. This is a sizeable share of the total number of e-scooter trips in each city (1.4 percent in Oslo and 1.6 percent in Lisbon), considering that only 10 percent of the users were nudged at least once.

5.1.2 Can e-scooters replace car use?

The previous section illustrated that nudged users increased their e-scooter utilization considerably. Along the same lines, we observe a reduction in the probability of relevant search sessions resulting in ride hail trips for virtually all experiments as well, although effects are less precisely estimated.

In Oslo, where effects are most precise, nudging reduces the share of relevant search sessions resulting in ride hail trips by 1.6 percentage points. This constitutes about 55 percent of the additional e-scooter trips. The remaining 45 percent of e-scooter trips caused by nudges were users that otherwise would have closed the app without booking a trip. We don't know what these users did instead – they might have walked, taken public transport, used a private car, ordered a taxi, booked ride hailing or e-scooter from a competing service or refrained from making the trip. However, we can say for certain that for the relevant search sessions, at least 55 percent of the e-scooter trips caused by nudging replaced cars. Other research indicates that on average, e-scooter rides in Oslo have very different substitution patterns: Only 8 percent replace car trips (5 percent replace taxis while 3 percent replace private cars; see Figure 5.1). This shows that in-app information is able to affect users' transport behaviour in a way that significantly reduces car trips to a much larger extent than what has previously been documented in the literature, given the right context.

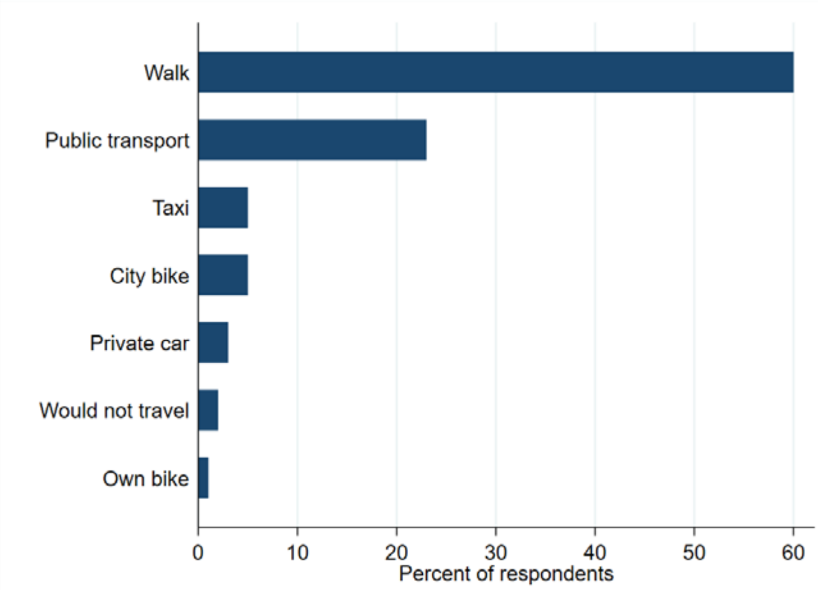

Figure 5.1: "What would you have done on your last e-scooter trip if you couldn't use e-scooter?" Results from a survey among e-scooter users in Oslo. Source: Fearnley et al (2020).

In the second experiment in Lisbon, the picture is different. The reduction in ride hailing trips among nudged users in relevant search sessions is estimated to around 1 percentage point. This is about twice as large as the increase in e-scooter trips, meaning that nudging actually reduced the probability of a successful trip for these sessions: For every additional e-scooter trip, two relevant search sessions are deterred away from ride hailing trips. However, we do not know what the users did instead – they might have booked a car from a competing ride hailing service. This makes it more difficult to draw clear conclusions regarding substitution patterns for experiments in which nudging reduces the total number of trips recorded by the Bolt app.

Comparing results across all experiments, there seems to be a pattern as to where nudging increases/reduces the chance of a successful trip through the Bolt app. In the cities where e-scooters are most common in relevant search sessions among the control group (Oslo, Gothenburg and Stockholm 1) nudging increases the number of completed trips (ride hailing and e-scooter taken together). In remaining experiments however, the share of relevant search sessions resulting in either an e-scooter trip or a ride hail trip was slightly lower among the users that were nudged (although the effect is only significantly different from zero in Lisbon 2 and Stockholm 2). This might be explained by the fact that only a small number of people benefitted from the additional e-scooter information given through the nudge. For the majority of the users who would not consider e-scooters in the first place, the additional e-scooter information might have been considered an annoyance or gotten in the way of the ride hail information sought by the user.

The rest of this section will focus on the results from Oslo. First, to discuss results related to the discrete choice model in Section 4.1.3. Second, because Oslo is the only experiment for which results in Section 4.2 are strong enough that we more confidently can say something about the effect of nudging on ride hail trips outside the relevant search sessions.

The discrete choice model for Oslo indicates that while 58 percent of e-scooter trips caused by nudging is deterred from ride hail, only 31 percent of additional e-scooter trips caused by reductions in e-scooter prices replace ride hailing. When increasing the ride hailing price, only 10 percent of deterred users will substitute towards Bolt's e-scooter service. This shows that substitution from ride hailing to e-scooter in relevant search sessions is considerably higher when users are nudged compared to when prices change. Even though we do not know what the remaining users did instead (they might have booked a trip from a competing company) this indicates that that in-app information can be a powerful tool for changing transport behaviour, that can complement regulatory price-based instruments (e.g. taxes and fees).

Finally, we will discuss whether e-scooter trips replace ride hailing trips also outside relevant search sessions. Effects on ride hailing trips are less precisely estimated, and only sizeable in Oslo. The point estimate indicates that nudged users conducted 0.1 fewer ride hail trips in average over the duration of the experiment, corresponding to a 5 percent decrease. In other experiments the decrease in ride hailing trips among nudged users is in the range of 0-2 percent, although insignificantly different from zero. However ride hailing trips are not directly comparable to e-scooter trips as they tend to be longer. If we instead consider distance, we predict that nudged users in Oslo travelled 1.8 additional kilometres by e-scooter over the duration of the experiment, and 0.9 fewer kilometres by ride hailing. This suggests that about 50 percent of additional e-scooter kilometres caused by the experiment replaced ride hailing kilometres. Although several of the results regarding user behaviour outside relevant search sessions are statistically uncertain, the analyses in Section 4.2 and Appendix B1 indicate that estimates are fairly robust.

5.2 Conclusion

Results from the empirical analyses clearly show that nudging can be a powerful tool in affecting users' transport behaviour. Even though the experiments conducted were of a non-intrusive nature (there was little to no cost for the user associated with ignoring the nudge), several of the behavioural responses analysed are both statistically and economically significant.

While ride hailing outcomes tend to be less precisely estimated, there is a significantly higher share of nudged users that end their search session with an e-scooter trip in all experiments except Valletta. The absolute effect is largest in Oslo, where an additional 3 percentage points (40 percent) of nudged users chose e-scooter compared to the control group. The relative effect is largest in Lisbon, where three times as many (0.4 percentage points) of nudged users chose e-scooter compared to the control group. This demonstrates that the additional information provided to nudged users in many instances enabled them to make better informed travel decisions.

An interesting result is the importance of having access to available e-scooters nearby, both when it comes to the effect of nudging, and the probability of choosing e-scooters for the control group. The proximity to e-scooters explains much of the differences within cities, and is also a likely explanation for differences observed across cities (e-scooter proximity is important in every experiment, and the density of e-scooters is notably higher in Oslo compared to e.g. Lisbon).

The total effect of nudging is modest when only looking at the search sessions in which the nudging took place. We estimate that the direct effect was 283 additional e-scooter trips in Lisbon and 273 in Oslo. This is not because the effect of being nudged is small, but because a small share of search sessions were considered to be relevant for nudging in these experiments. However, users that were nudged at least once were more likely to conduct additional e-scooter searches and e-scooter trips subsequently. We find that about 3,800 additional e-scooter trips were conducted in Oslo and 1,400 in Lisbon over the duration of the experiment, as a result of nudging. This shows that people that are first presented with the additional e-scooter information are more likely to alter their behaviour in the long term as a result. These predictions however are statistically more uncertain than when just considering user sessions relevant for nudging.

We delve more into results for Oslo, the city where nudging had the highest impact. We show that the relative substitution from ride hail to e-scooter is stronger when users are nudged compared to when reducing (increasing) the price of e-scooters (ride hail). In the search sessions relevant for nudging, at least 55 percent of the e-scooter trips caused by the nudge replaced ride hail trips. Results are similar (although less precisely estimated) when looking at distance travelled by nudged users over the whole duration of the experiment. The treatment group travels on average an additional 1.8 kilometres by e-scooters, and 0.9 fewer ride hail kilometres. Again, this indicates that 50 percent of the additional e-scooter kilometres replace driving.

Results from the experiment show that in-app information is able to affect users' transport behaviour in a way that significantly reduces car trips to a much larger extent that what has previously been documented in the literature. This behavioural change is initiated by the nudge experiment, but facilitated from the fact that the app interface is multimodal: hence, the experiment illustrates the effectiveness of "nudging" and in-app information in general, and the benefit of multimodal interfaces for facilitating substitution between the modes.

5.3 Further Research and refinement of nudge experiments

Results indicate that nudging users through multimodal interfaces is a viable way of significantly affecting mode choice, at little to no cost to the user. Hence, exploring different types of nudge experiments in the future can greatly improve our understanding of drivers of micromobility and associated substitution patterns to other modes of transport. The research literature and our analysis suggest that mode substitution varies with predictable patterns, which can be informative for future nudging experiments. Below we will outline a few opportunities for further research.

- US evidence suggests that e-scooters replace car trips to a larger degree in the US than Europe. Hence, focusing on areas with "US like" conditions, i.e. car dependent and sprawled city areas, will likely increase substitution rates.
- The effect of nudging seems to be more or less proportional to the market share in the control group. This indicates that nudging will affect more users if the alternative they are nudged towards is more popular in the first place. This will vary across cities and situations, but can be checked with data.
- However, Lisbon stands out: e-scooters are less popular compared to ride hailing than in any other city, but still three times as many nudged users chose e-scooter compared to the control group. This indicates that in-app information can be an effective way of informing users about alternatives for which there is little utilisation in the first place as well.
- The analysis indicated that in cities where few users chose e-scooters, nudging negatively impacted their chance of making a successful trip. This could either be because users were annoyed by the e-scooter information, or because the additional e-scooter information made them overlook the ride hail information. This could potentially be tested if the additional e-scooter information was displayed to the user in another way (e.g. let nudged users see four alternatives at the same time instead of three).
- Results show that availability of e-scooters within a short distance is critical. This indicates that attempting to nudge users towards e-scooters will have little effect if they have to walk far.
- As found elsewhere in the literature, travel time elements seem to have a larger effect on mode substitution than prices do (Wardman et al., 2018; Fearnley et al., 2017). The fact that price sensitivity appears to be modest, suggests that time saving factors may be more effective in convincing the user to shift transport mode.
- Nudging users between ride hailing and e-scooters can, depending on the context, have benefits when it comes to accident risk reduction, car congestion, and local and global pollutants. An interesting question is whether users will react more strongly to being nudged if they are informed of these benefits (i.e. that their mode choice might have a positive or negative impact for others).
- Regarding e-scooters, accidents represent a major external cost (Fearnley, 2020). E-scooter accident risk is particularly high during weekend nights when driving e-scooters under the influence of alcohol and drugs is widespread. Nudging away from e-scooters towards ride hailing during night time could therefore have important benefits in terms of accident reductions. The fact that ride hailing is more popular during evenings and nights indicates that nudging is likely to have a larger impact at these times.
- Bolt has already implemented a more intrusive version of this in several cities. E-scooter users are required to take a drunk driving test during nights; if they fail the test, ride hailing is suggested. However, the data has not been analysed yet. To

properly evaluate the impact of such measures, app data on user patterns should ideally be matched with accident data. However, this would be difficult in practice, due to privacy considerations and the sensitive nature of these data types.

- The literature suggests that shared e-scooters to a large extent is used as the first or last leg of a multimodal trip, e.g. in combination with public transport (Fearnley et al, 2020). Nudging users by providing them with information regarding how e-scooter trips to a station corresponds with public transport departures could facilitate substitution from cars to e-scooters for longer trips as well.

References

Athey, S., & Imbens, G. W. (2017). The econometrics of randomized experiments. In *Handbook of economic field experiments* (Vol. 1, pp. 73-140). North-Holland.

Axsen, J., Plötz, P., & Wolinetz, M. (2020).Crafting strong, integrated policy mixes for deep CO2 mitigation in road transport. *Nature Climate Change*, 10(9):809:818.

Ben-Akiva, M. E. and S. R. Lerman (1985). *Discrete choice analysis: theory and application to travel demand*, MIT press.

McFadden, D. (1978). *Modelling the choice of residential location*, Institute of Transportation Studies, University of California.

McFadden, D. (1981). *Econometric Models of Probabilistic Choice. Structural Analysis of Discrete Data with Econometric Applications*. C. Manski and D. McFadden. Cambridge, MA, D ed., The MIT Press: MIT Press.

Fearnley, N. (2020). Micromobility—regulatory challenges and opportunities. Book chapter in: Sørensen, C.H., Paulsson, A. (eds.), 2020. *Shaping smart mobility futures. Governance and policy instruments in times of sustainability transitions*. Emerald publishing, https://doi.org/10.1108/978-1-83982-650-420201010

Fearnley, N., Berge, S. H. & Johnsson, E. (2020). Shared e-scooters in Oslo. *TØI report 1748/2020*. Norwegian only (English summary).

Fearnley, N., Johnsson, E., Berge, S.H. (2020). Patterns of E-Scooter Use in Combination with Public Transport, *Transport Findings*, https://doi.org/10.32866/001c.13707

Fearnley, N., Flügel, S., Killi, M., Gregersen, F.A., Wardman, M., Caspersen, E., Toner, J.P., 2017. Triggers of Urban Passenger Mode Shift – State of the Art and Model Evidence, *Transportation Research Procedia*, Vol 26, 2017, pp 62-80, https://doi.org/10.1016/j.trpro.2017.07.009

Flannelly, K. J., Flannelly, L. T., & Jankowski, K. R. (2018). Threats to the internal validity of experimental and quasi-experimental research in healthcare. *Journal of health care chaplaincy*, *24*(3), 107-130.

Franssens, S., Botchway, E., de Swart, W., & Dewitte, S. (2021). Nudging Commuters to Increase Public Transport Use: A Field Experiment in Rotterdam. *Frontiers in psychology*, *12*, 692.

Graham-Rowe, E., Skippon, S., Gardner, B., & Abraham, C. (2011). Can we reduce car use and, if so, how? A review of available evidence. *Transportation Research Part A: Policy and Practice*, *45*(5), 401-418.

Hess, S., Bierlaire, M., & Polak, J. W. (2005). Estimation of value of travel-time savings using mixed logit models. *Transportation Research Part A: Policy and Practice*, *39*(2-3), 221-236.

IEA (2020). *Tracking transport 2020*. IEA, Paris.

Laa, B., & Leth, U. (2020). Survey of E-scooter users in Vienna: Who they are and how they ride. *Journal of transport geography*, *89*, 102874.

Mackie, P. J., Jara-Dıaz, S., & Fowkes, A. S. (2001). The value of travel time savings in evaluation. *Transportation Research Part E: Logistics and Transportation Review*, *37*(2-3), 91-106.

Meldrum, M. L. (2000). A brief history of the randomized controlled trial: From oranges and lemons to the gold standard. *Hematology/oncology clinics of North America*, *14*(4), 745-760.

Reck, D. J., Martin, H., & Axhausen, K. W. (2022). Mode choice, substitution patterns and environmental impacts of shared and personal micro-mobility. *Transportation Research Part D: Transport and Environment*, *102*, 103134.

Rubin, D. B. (1974). Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of educational Psychology*, *66*(5), 688.

Rødseth, K. L., Wangsness, P. B., Veisten, K., Høye, A., Elvik, R., Klæboe, R., Thune-Larsen, H., Fridstrøm, L., Lindstad, E., Rialland, A., Odolinski, K., & Nilsson, J.-E. (2019). Eksterne kostnader ved transport i Norge. *TØI report 1704/2019*. Norwegian only (English summary).

Schulz, K. F., & Grimes, D. A. (2002). Generation of allocation sequences in randomised trials: chance, not choice. *The Lancet*, *359*(9305), 515-519.

Thaler, R. H., & Sunstein, C. R. (2003). Libertarian paternalism. *American Economic Review*, 175-179.

Thaler, R. H., & Sunstein, C. R. (2008). *Nudge: Improving Decisions About Health, Wealth, and Happiness.* Penguin Books.

Van Essen, H., van Wijngaarden, L., Schroten, A., Sutter, D., Bieler, C., Maffii, S., Brambilla, M., Fiorello, D., Fermi, F., Parolin, R & El Beyrouty, K. (2019). *Handbook on the external costs of transport, version 2019* (No. 18.4 K83. 131).

Wang, K., Qian, X., Fitch, D.T., Lee, Y., Malik, J., Circella, G., 2022. What travel modes do shared e-scooters displace? A review of recent research findings, *Transport Reviews*, DOI: 10.1080/01441647.2021.2015639

Wardman, M., Toner, J., Fearnley, N., Flügel, S. & Killi, M. (2018). Review and meta-analysis of inter-modal cross-elasticity evidence. *Transportation Research Part A: Policy and Practice*, https://doi.org/10.1016/j.tra.2018.10.002

White, H. (1980). A heteroskedasticity-consistent covariance matrix estimator and a direct test for heteroskedasticity. *Econometrica: journal of the Econometric Society*, 817-838.

Yanocha, D., & Allan, M. (2019). The Electric Assist: Leveraging E-Bikes and E-Scooters for More Livable Cities. *Institute for Transportation and Development Policy*, New York.

Appendices

Appendix A: Data and additional descriptives

A1: Variables

The raw dataset is at the level of "search sessions", and contains the following variables:

- Indicators for whether the outcome of the search session is an e-scooter trip, a ride hail trip or no trip;
- Anonymised ID variable of the user;
- The self-reported age of approximately 40 percent of the users;
- Indicator for whether the user is in treatment or control group;
- Indicator for when a nudge has taken place, in case the user is in the treatment group and complies with the nudging criteria (see Section 3.2);
- Time stamp for when e-scooter search is initiated, in case of e-scooter search;
- Time stamp for when ride hail search is initiated, in case of ride hail search;
- Distance to nearest e-scooter;
- Cost of renting an e-scooter (per hour and per unlock);

In case a ride hailing search is initiated and the user types in a destination, we also have information on:

- Minimum distance between origin and destination;
- Time until driver arrives;
- Predicted duration of ride hailing trip;
- Predicted price of ride hailing trip;
- Ride hailing surge multiplier (a scaling factor that increases prices of all ride hailing trips e.g. in periods where there is excess demand);

In case a ride hailing or e-scooter trip is taken, we observe:

- A time stamp of when the e-scooter/ride hailing trip is selected;
- The actual price paid for e-scooter/ride hailing trips;
- The actual distance for e-scooter/ride hailing trips;
- The duration of e-scooter trips.

All time stamps are initially in UTC, but changed to local time. All cost variables are in local currency, but recalculated to Euros to be comparable across cities.

A2: Demand and supply for e-scooter and ride hail

The set of figures below display city specific descriptive statistics by time of day. The left panels are weekdays, while the right panels are weekends. The top four panels display the number of searches and trips by ride hail and e-scooter per hour. This gives a good indication of the popularity of the different modes across cities and time of day. The bottom four panels display the average distance to the nearest e-scooter and the average waiting time for drivers in case of ride hail for these time intervals. This can be thought of as an indication of supply, i.e. the density of available e-scooters and ride hail vehicles. Finally, the first and third row consider all ride hail search sessions where the variables in questions are non-missing, while the second and fourth row focus on "relevant search sessions" only. These are the search sessions for which nudging will take place in case the user is in the treatment group.

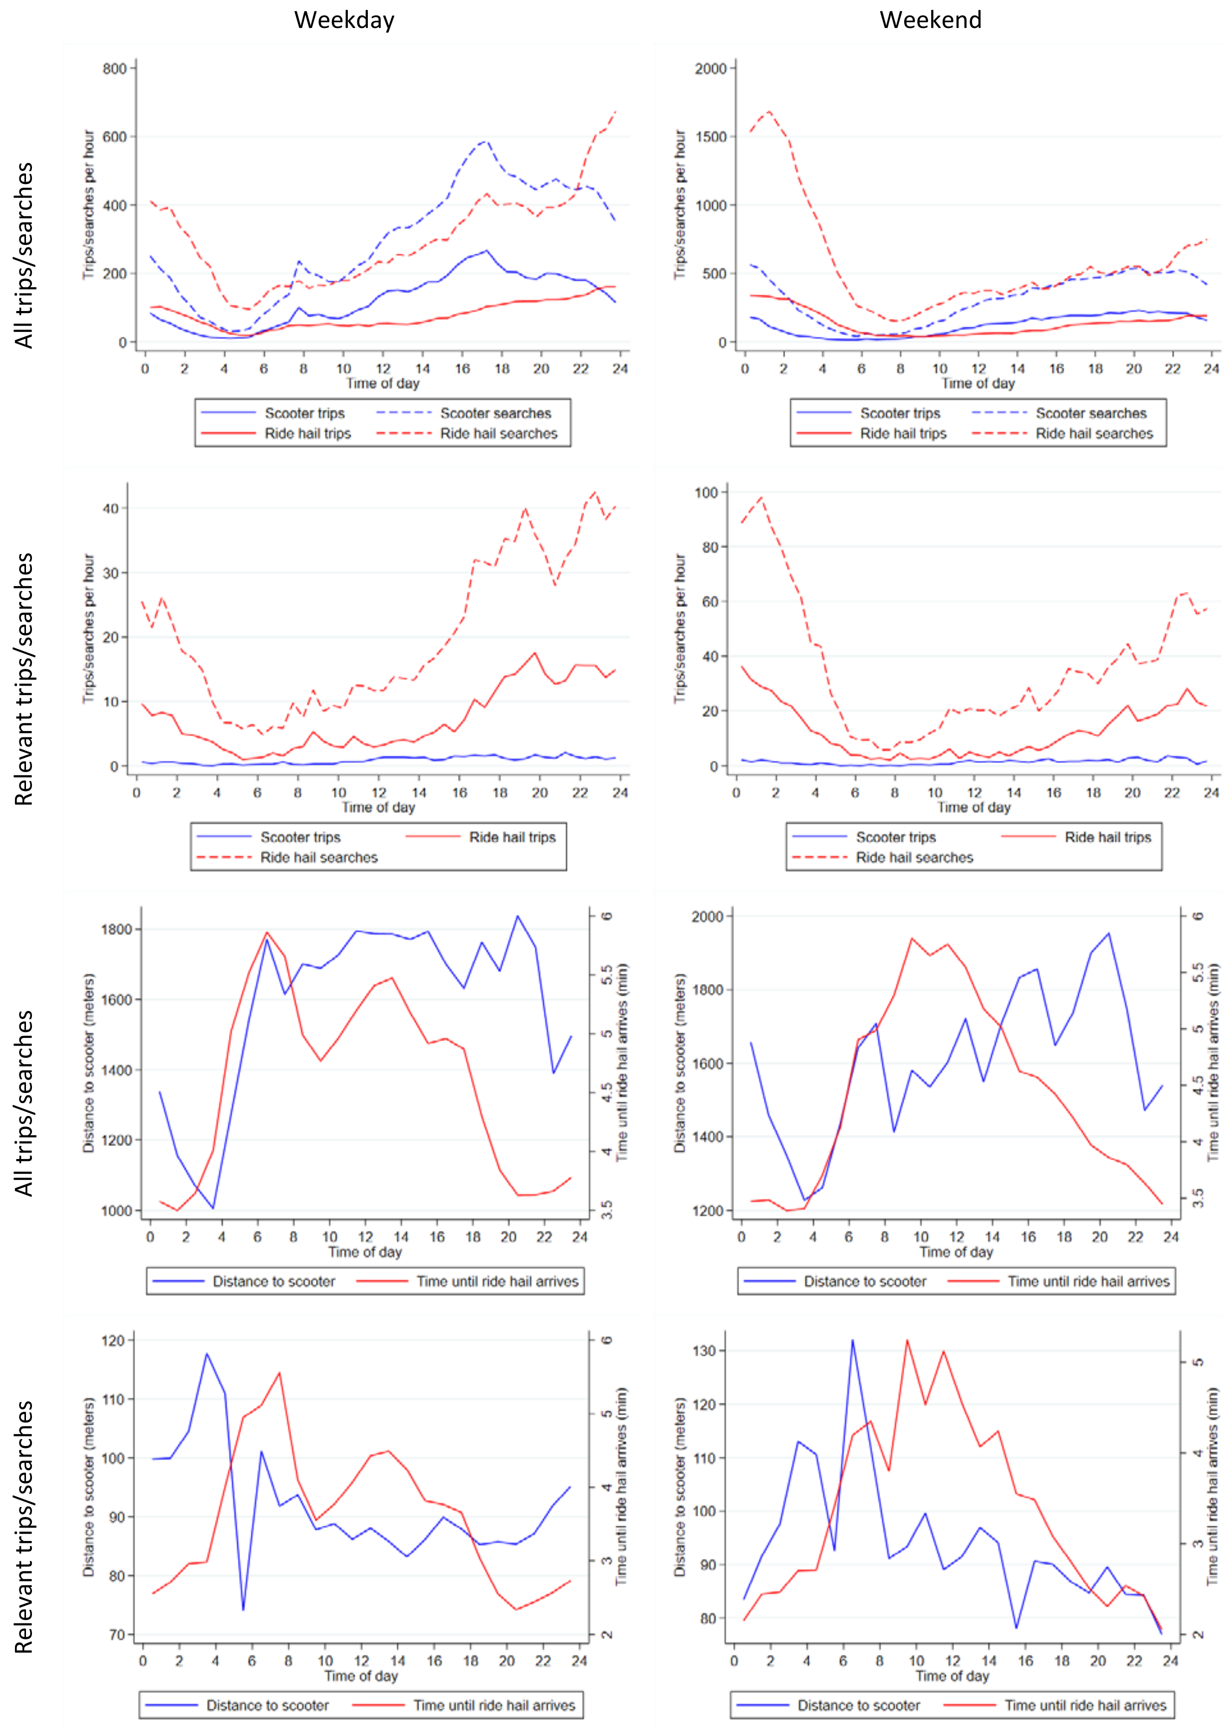

Figure A2.1: Trips and searches by time of day, Gothenburg.

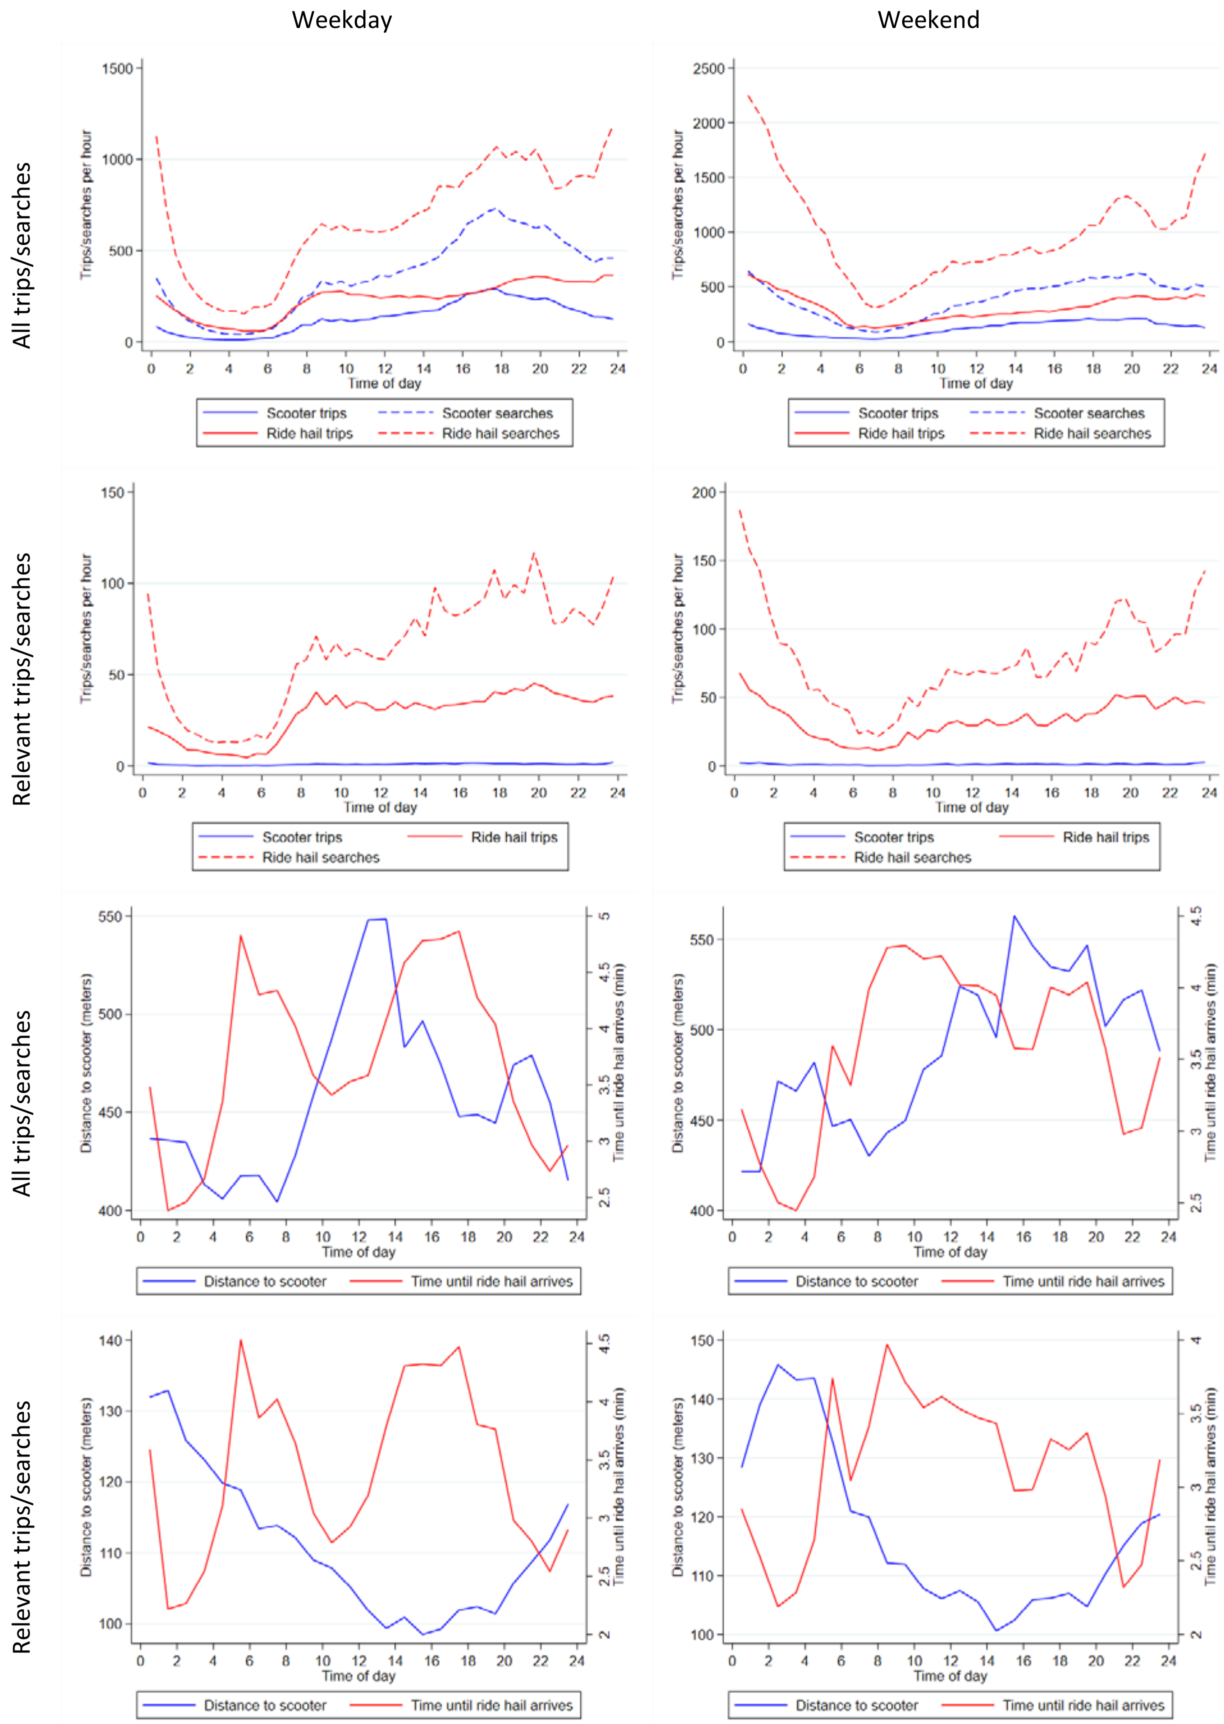

Figure A2.2: Trips and searches by time of day, Krakow.

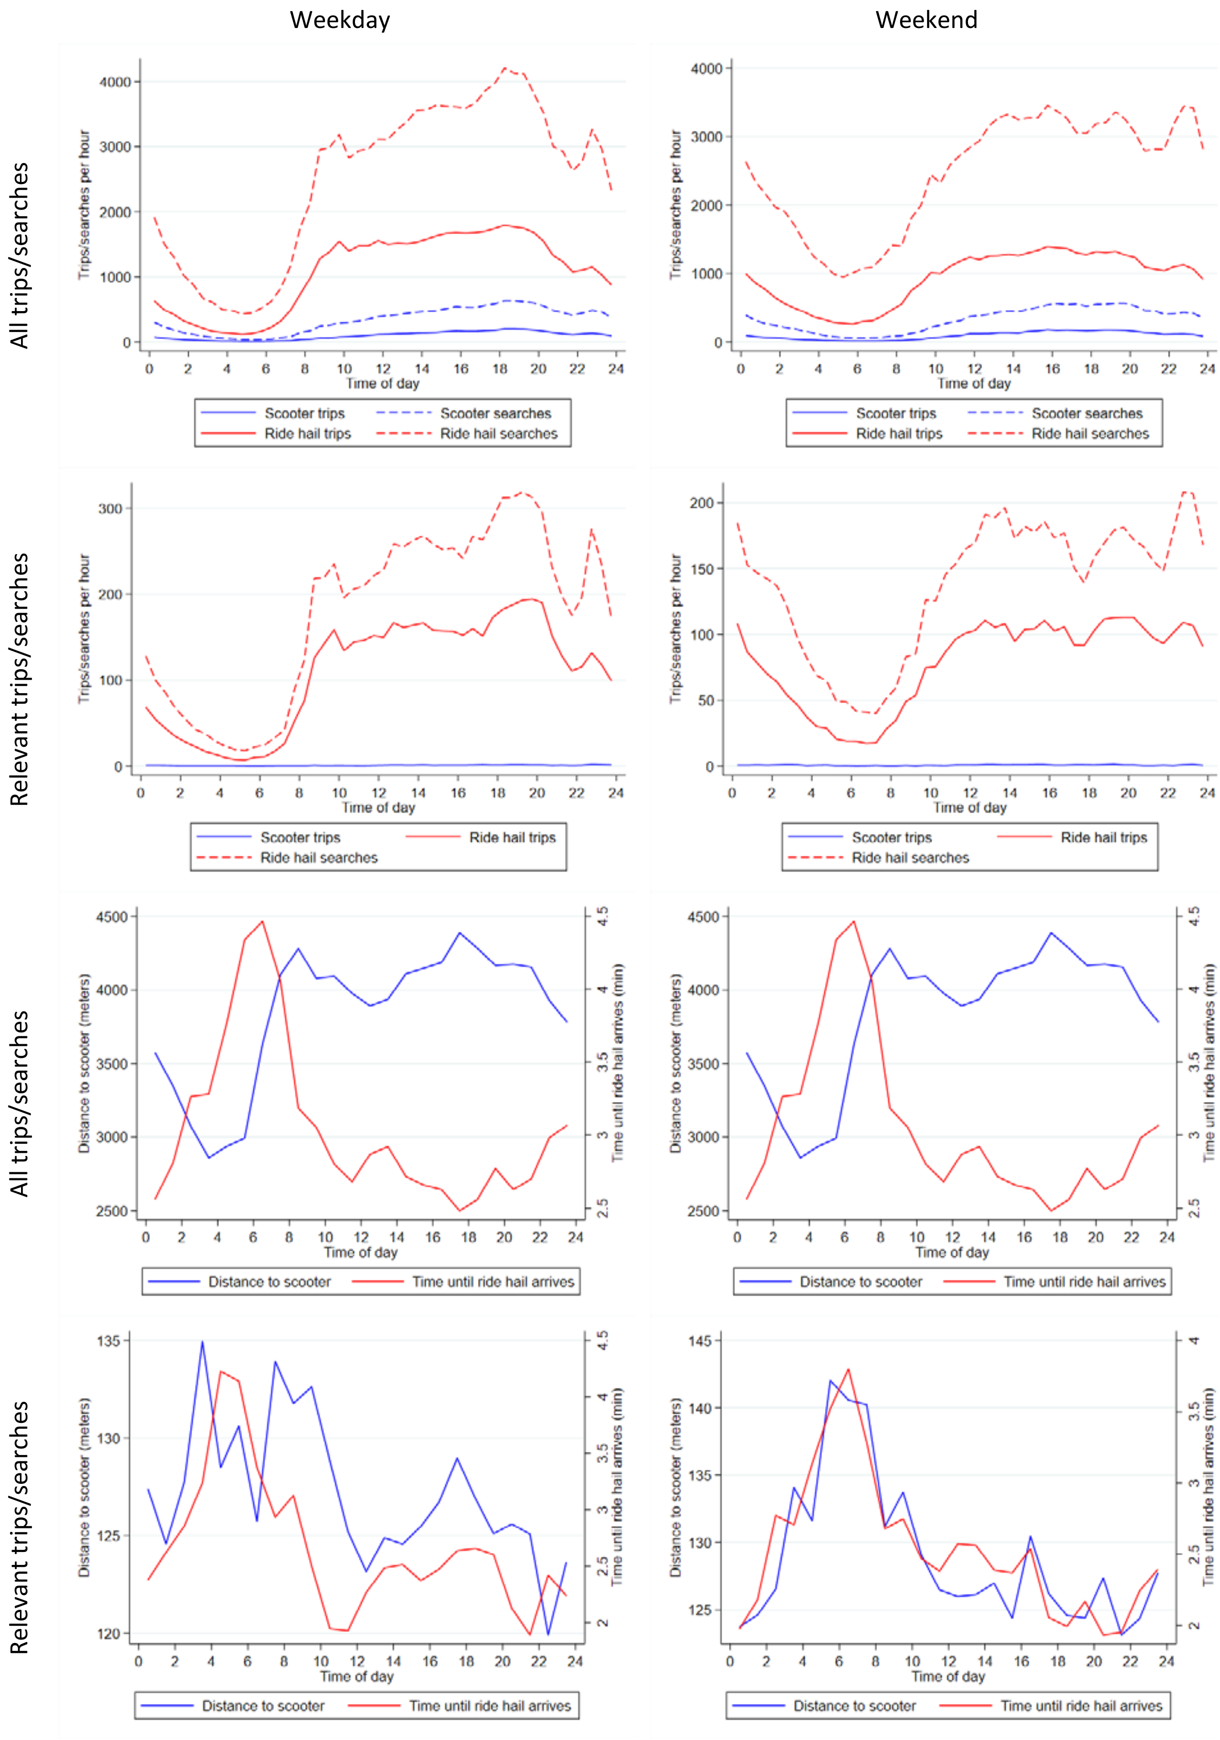

Figure A2.3: Trips and searches by time of day, Lisbon.

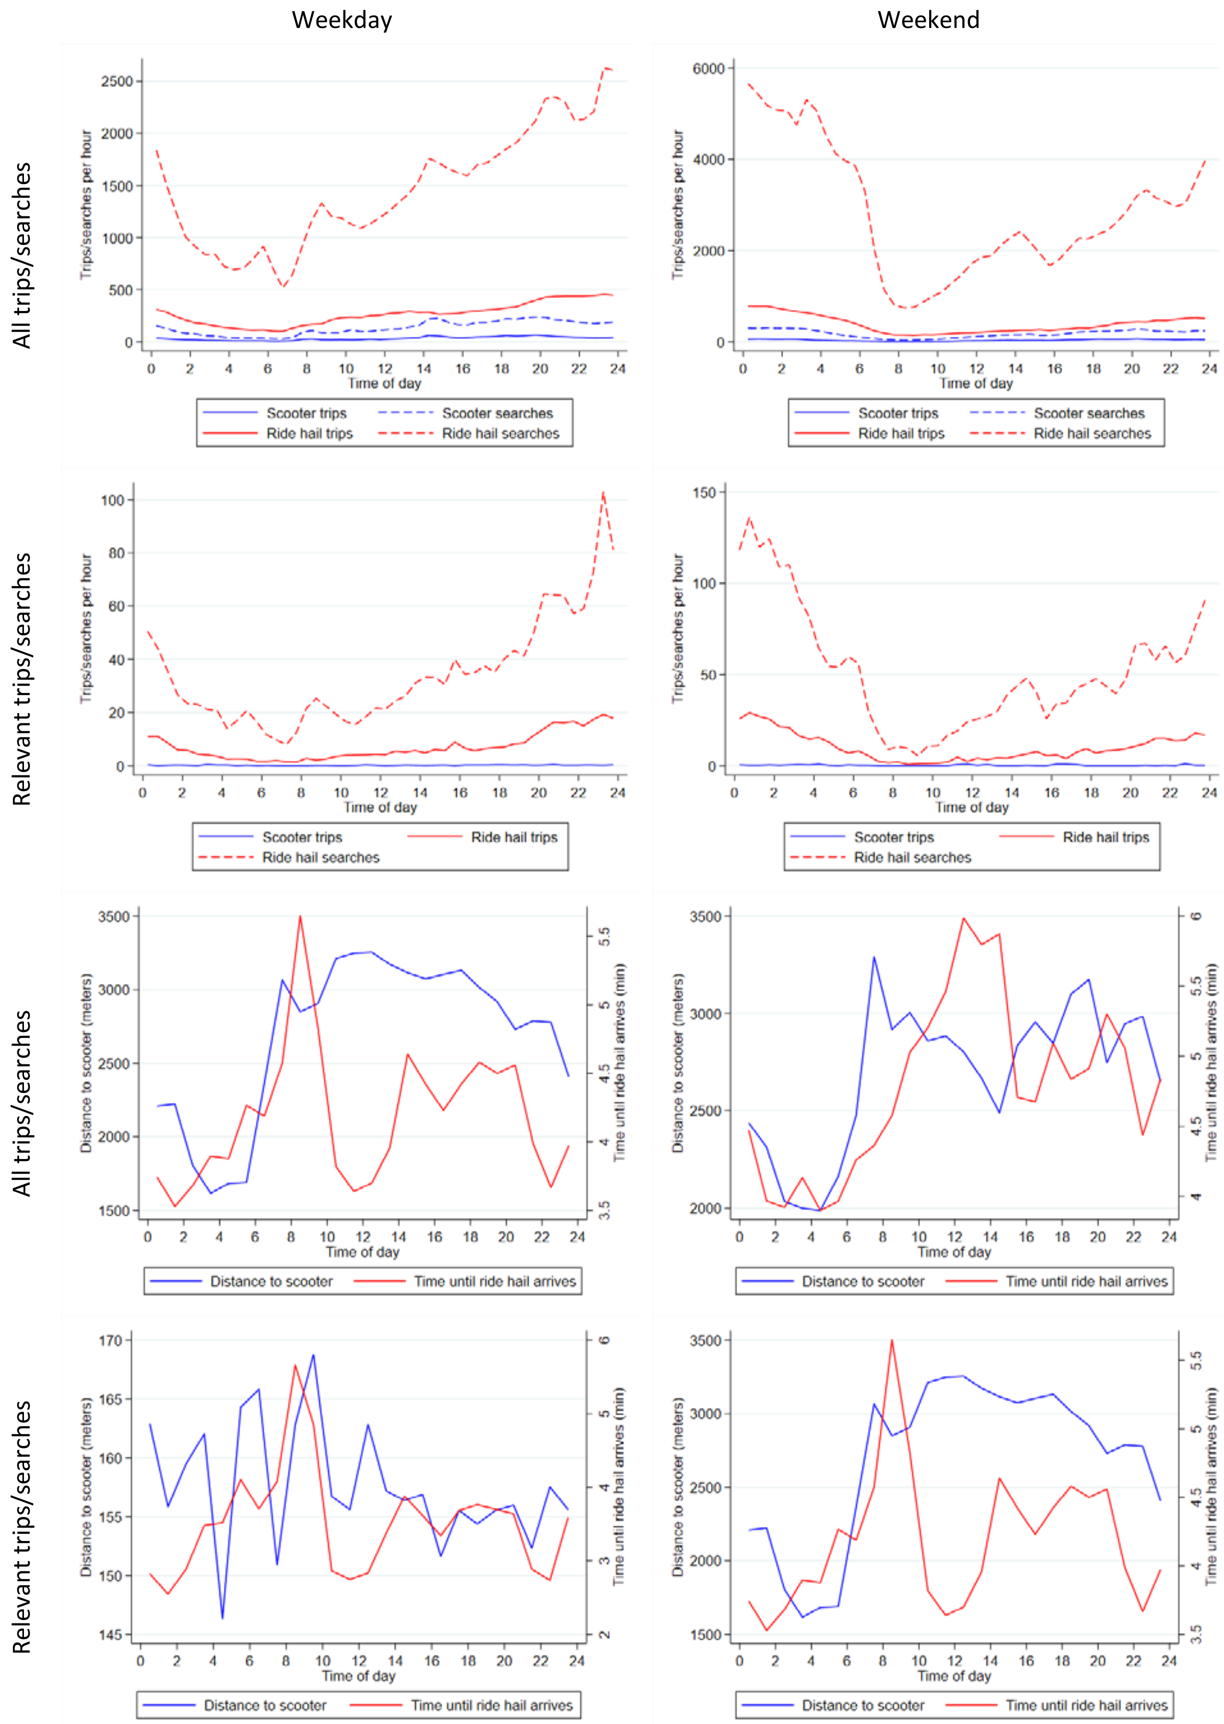

Figure A2.4: Trips and searches by time of day, Madrid.

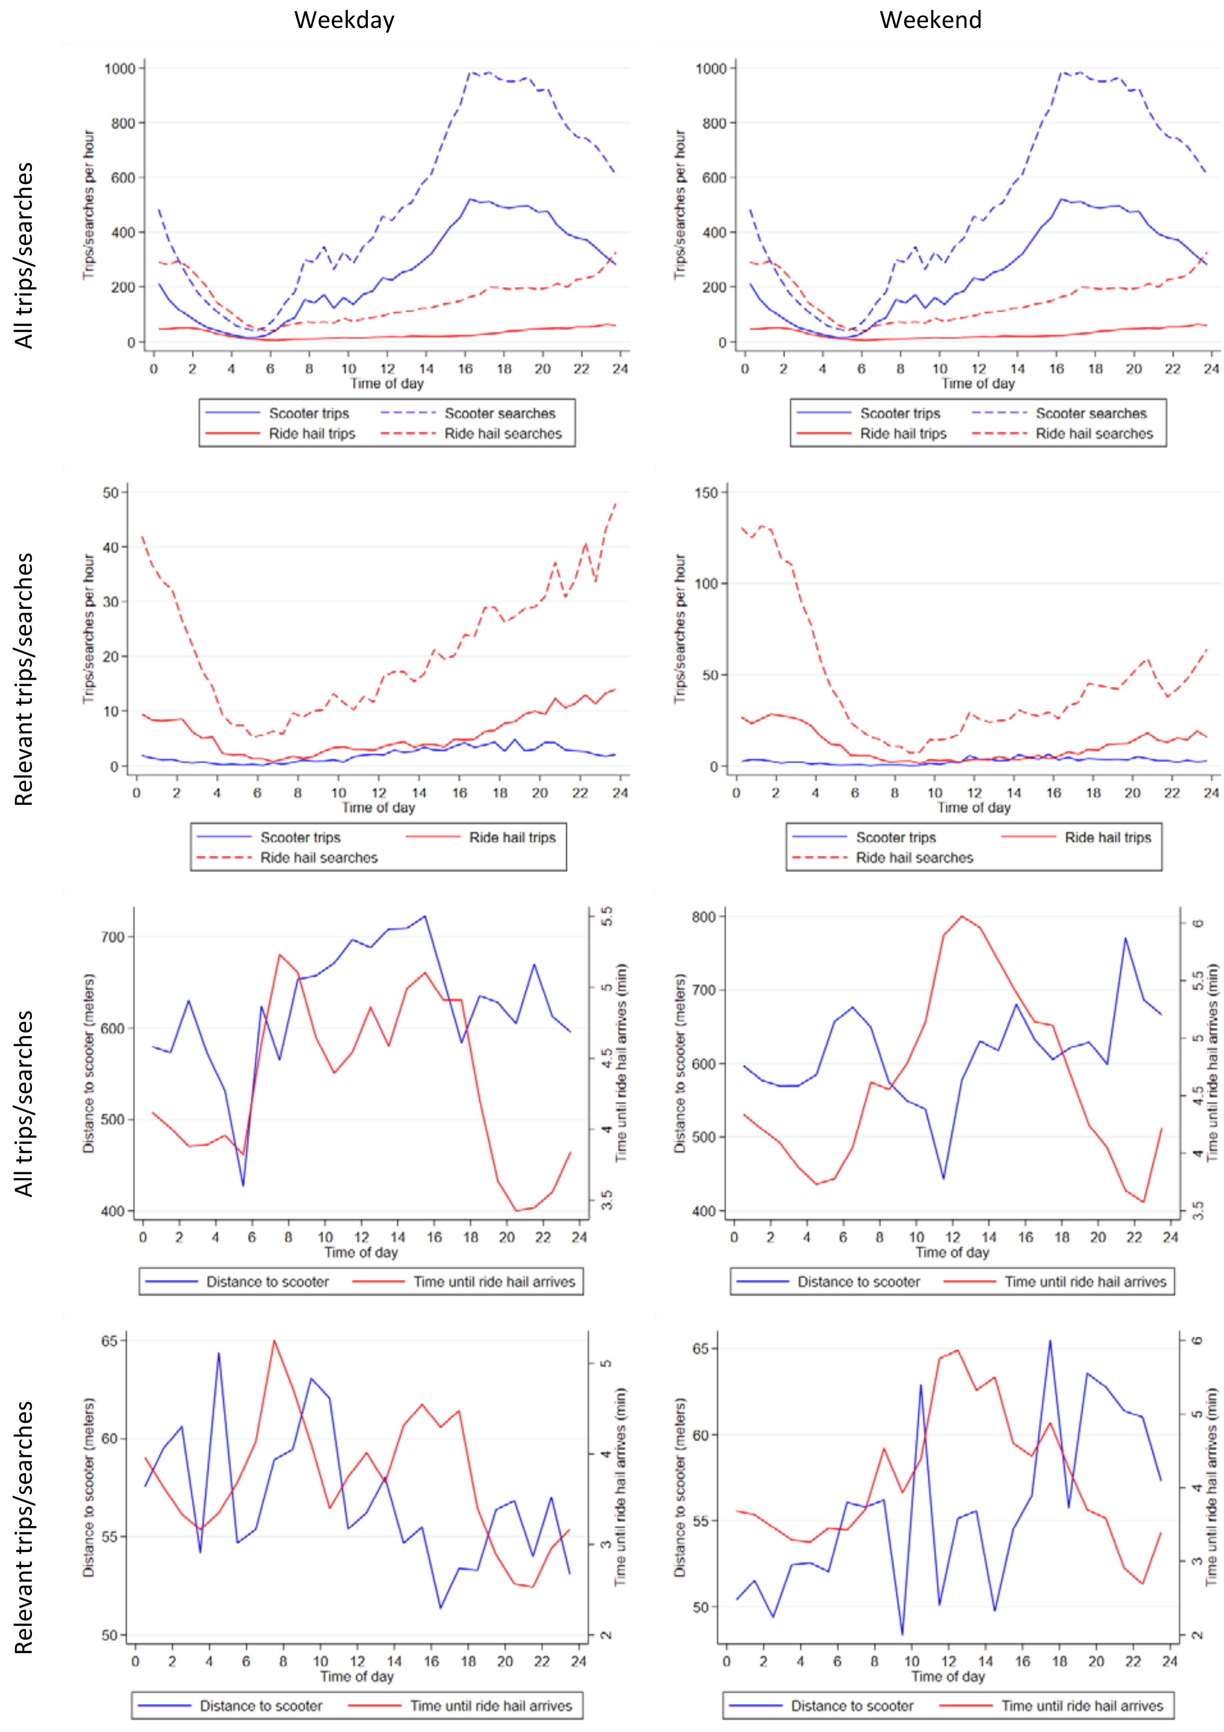

Figure A2.5: Trips and searches by time of day, Oslo.

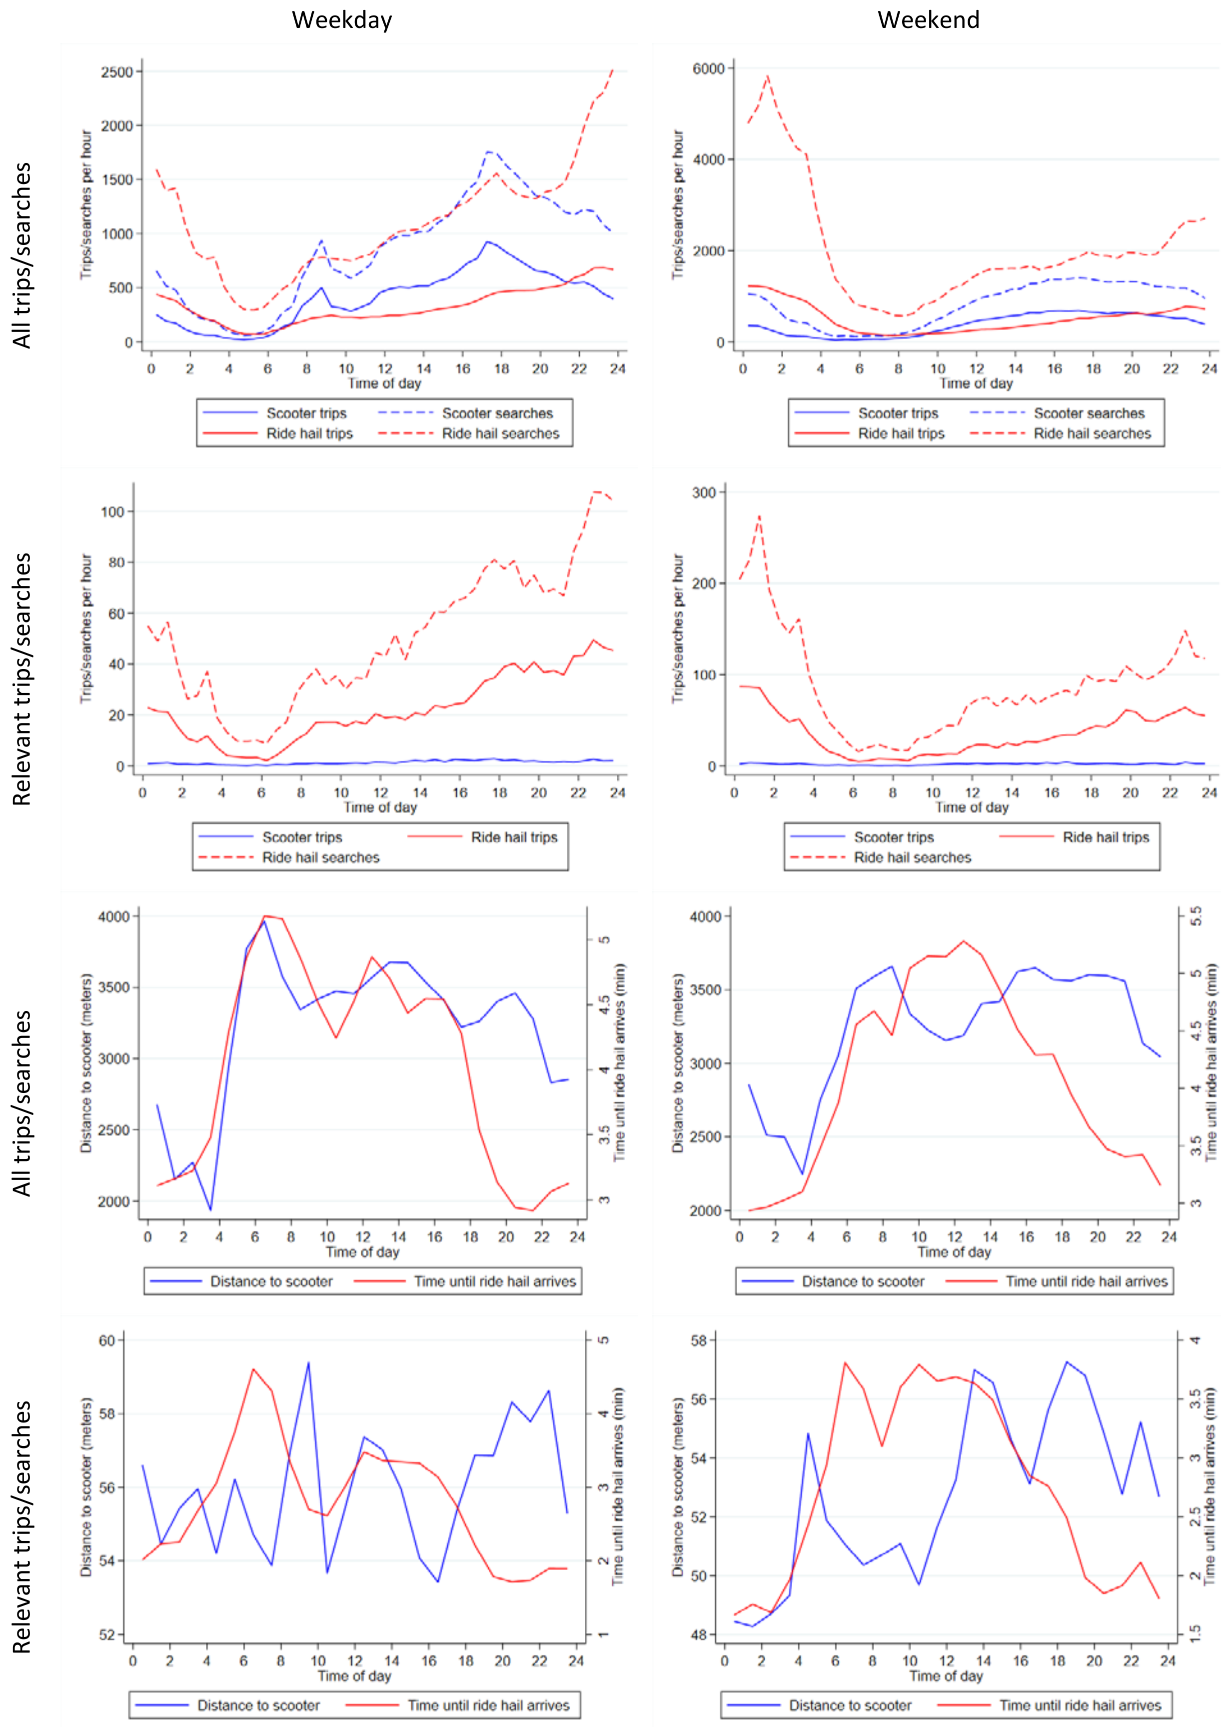

Figure A2.6: Trips and searches by time of day, Stockholm.

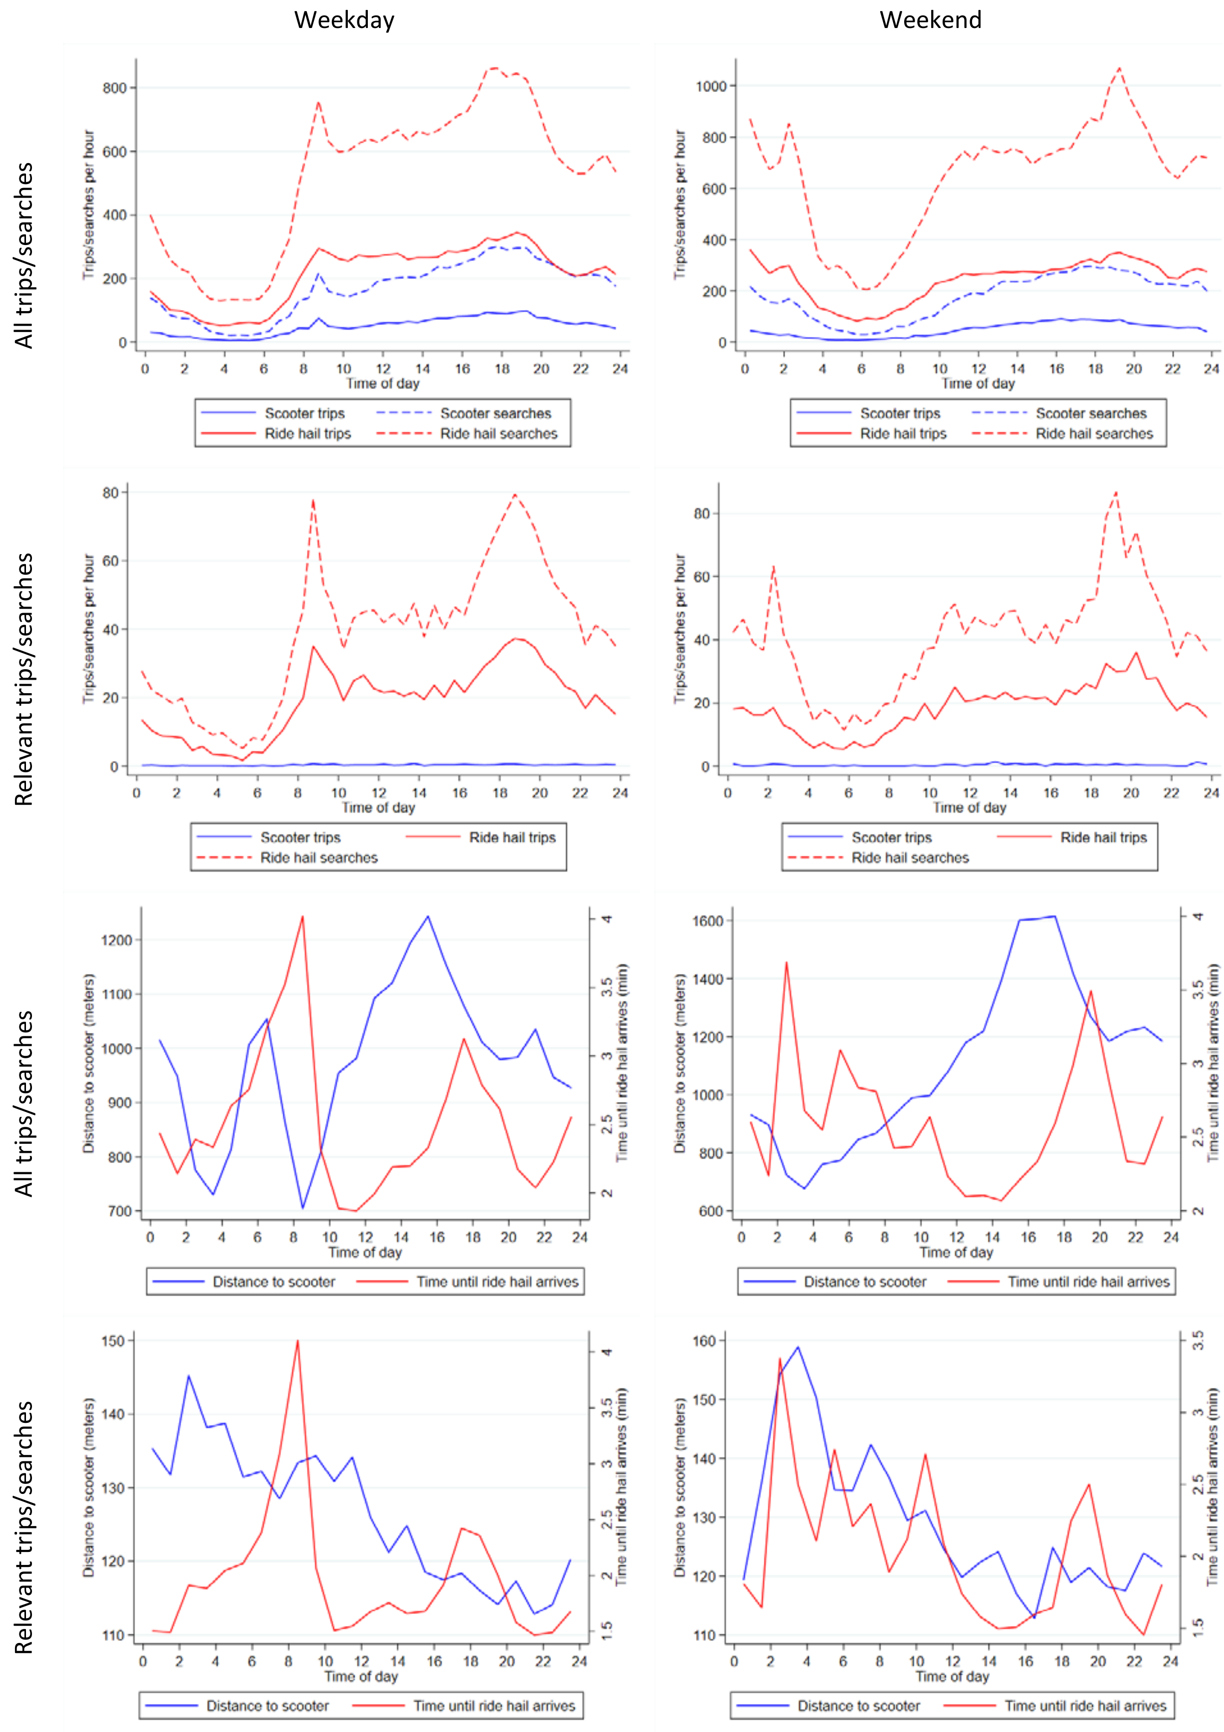

Figure A2.7: Trips and searches by time of day, Valletta.

A3: Distribution of characteristics of relevant search sessions by city

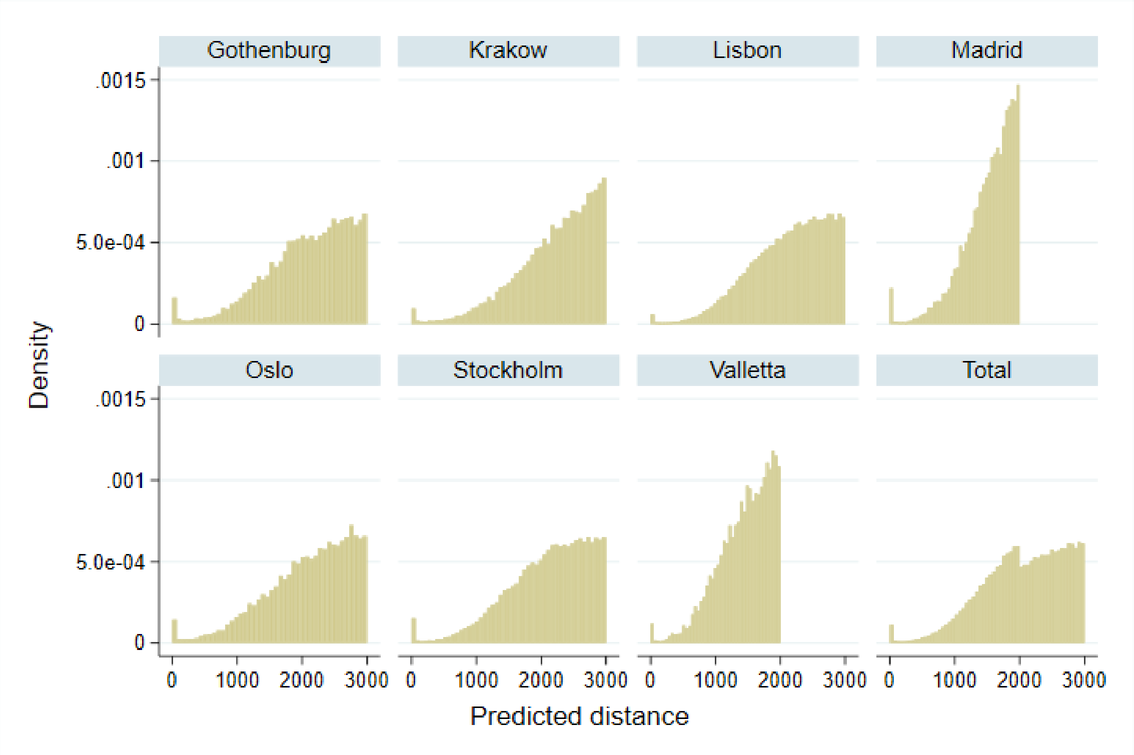

Figure A3.1: Histograms of predicted trip distance (meters) by city. Relevant search sessions only.

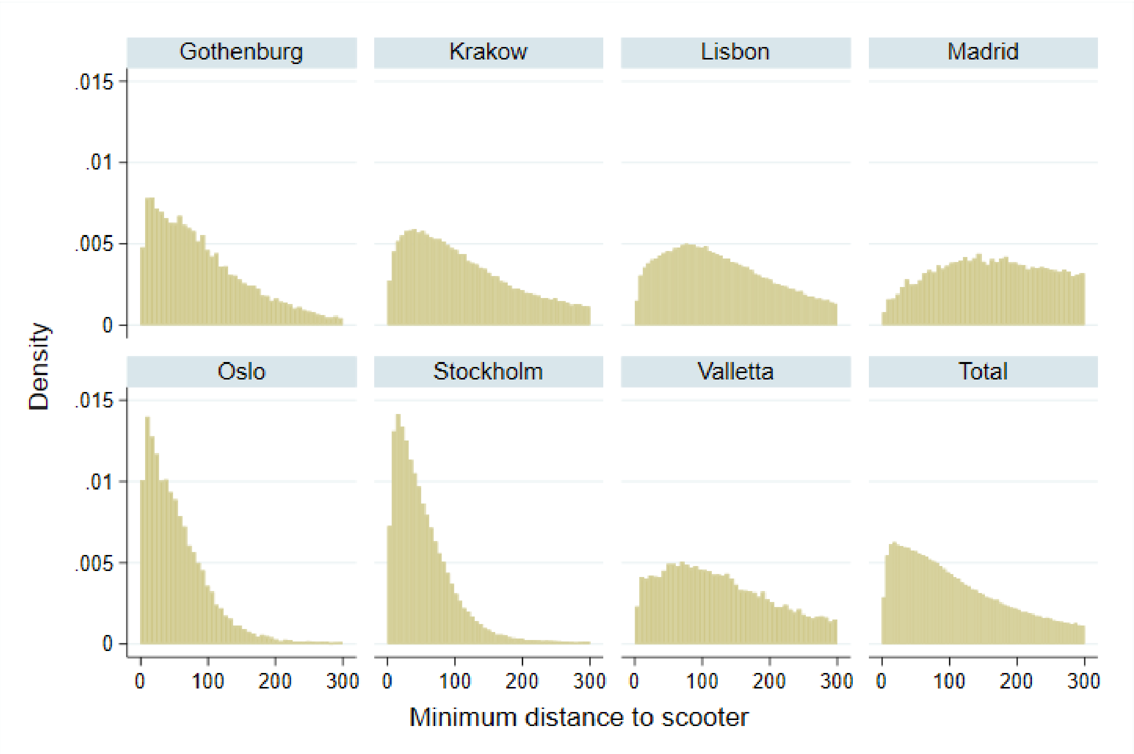

Figure A3.2: Histograms of minimum distance to e-scooter (meters) per city. Relevant search sessions only.

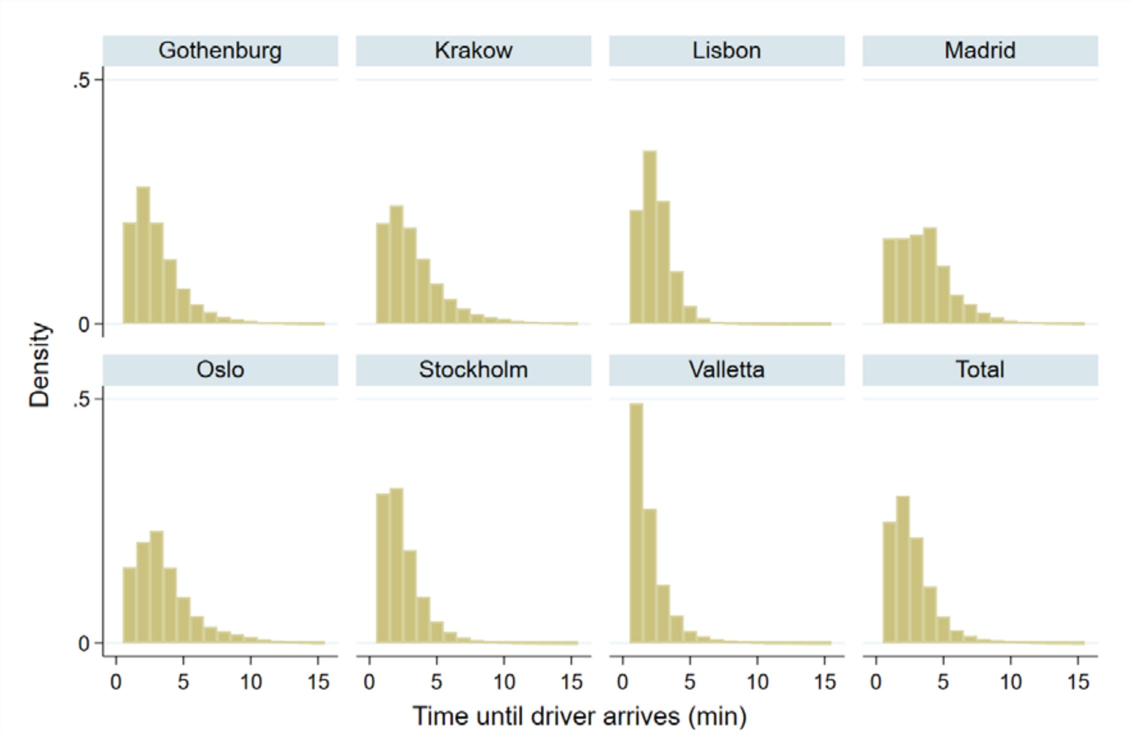

Figure A3.3: Histogram of time until driver arrives (minutes) per city. Relevant search sessions only.

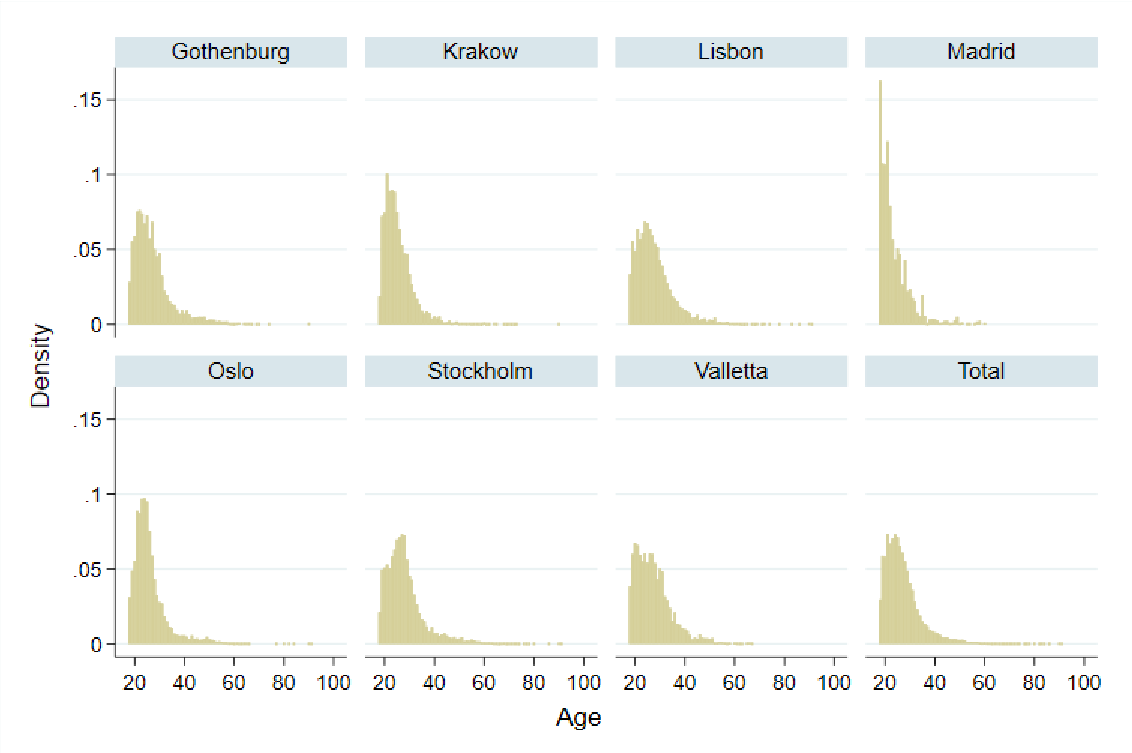

Figure A3.4: Histograms of self-reported age distributions per city. Relevant search sessions only.

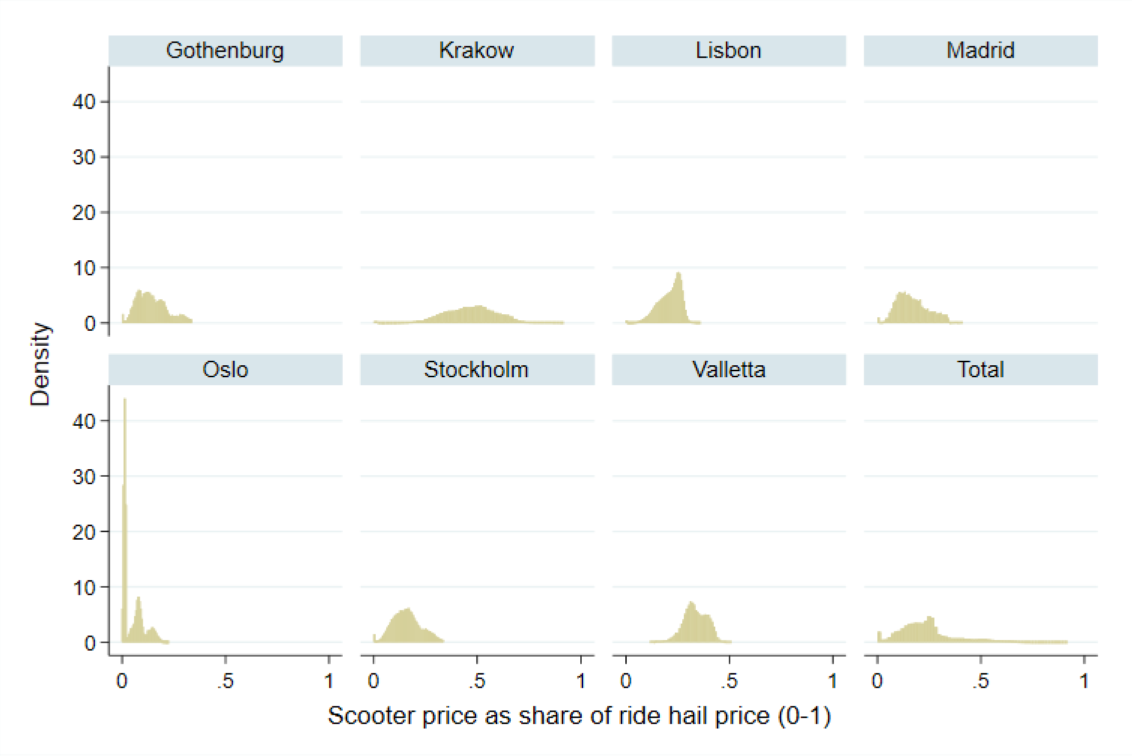

Figure A3.5: Histograms of the relative cost of e-scooters compared to ride hailing per city. Relevant search sessions only.

Appendix B: Additional results

Appendix B1 presents additional results on relevant search sessions and aggregate outcomes, to complement the main text. Detailed results on heterogeneous treatment effects to accompany the discussion in Section 4.1.2 are displayed in Appendix B2.

B1: Additional outcomes, average treatment effects

Relevant search sessions

This section complements results from Section 4.1.1 in the main text. A worry when considering all relevant search sessions is that they are non-randomly distributed between the treatment and the control group, since users that have been nudged once are either more or less likely to initiate another search session. To test this, we repeat the e-scooter analysis, but limit the sample to the first relevant search session experienced by each user. It is comforting to observe that results are still significant, in the same order of magnitudes and follow the same patterns.

We would not expect results to be exactly equal, due to changes in demand/supply throughout the course of the experiment: While results presented in this section are likely to have occurred at the beginning of each experiment, results from Section 4.1.1 will have occurred over the whole experiments' duration. Figure B1.1 should be compared to Figure 4.1, while Figure B1.2 should be compared to Figure 4.2.

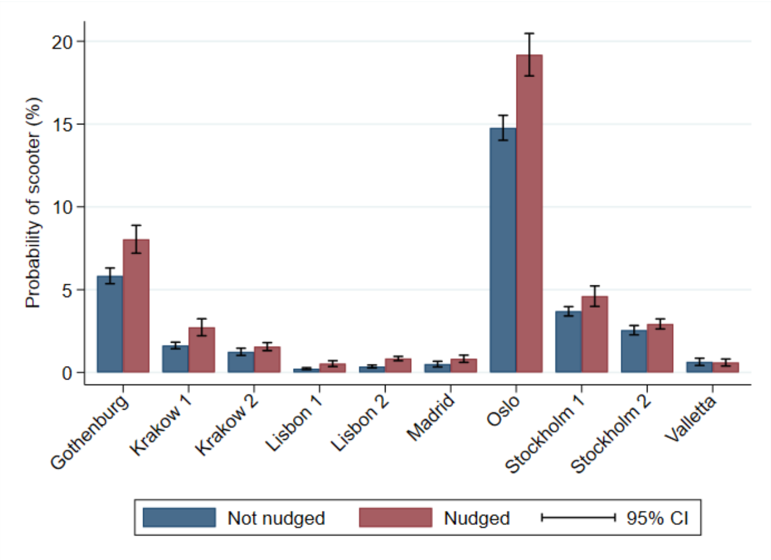

Figure B1.1: Share of users that chose e-scooter in their first relevant search session, treatment and control.

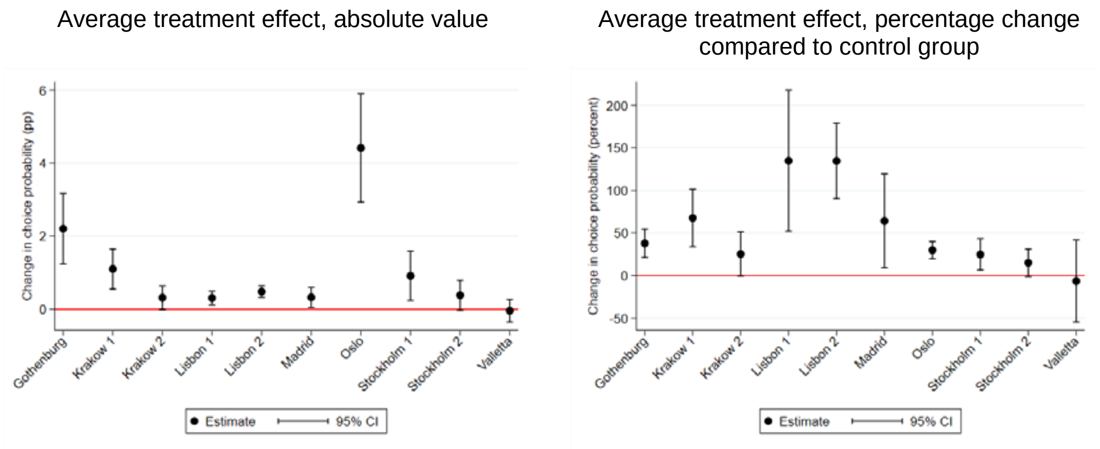

Figure B1.2: Effect of nudging on probability of choosing e-scooter trip. Sample: First relevant search session for each user. City-specific OLS regressions.

Aggregate outcomes by user

Results from Section 4.2 on how nudging affects the total number of trips indicate that effects might be considerable in magnitude, but are imprecisely estimated. This Appendix attempts to shed more light on this. Just as in Section 4.2, the treatment group consists of the full set of users that were nudged at least once. Users in the control group are those who experienced a similar search session at least once, without being nudged. All outcomes are user specific.

A problem with "number of e-scooter trips per user" and "number of ride hail trips per user" is that these outcomes exhibit a lot of variation. A couple of active users conduct hundreds of trips over the duration of the experiments, meaning that it might matter a lot for the results whether these users happen to be allocated in the treatment group or the control group. This variation increases the size of the confidence intervals, making estimated effects more uncertain.

Figure B1.3-Figure B1.4 considers and outcome variable with less variance, i.e. the share of users that have conducted at least one e-scooter trip throughout the duration of the experiments. Figure B1.5-Figure B1.6 consider the same outcomes as in the main text (number of trips) but remove the 1 percent of users with the most trips from the sample, separately for the treatment and the control group.

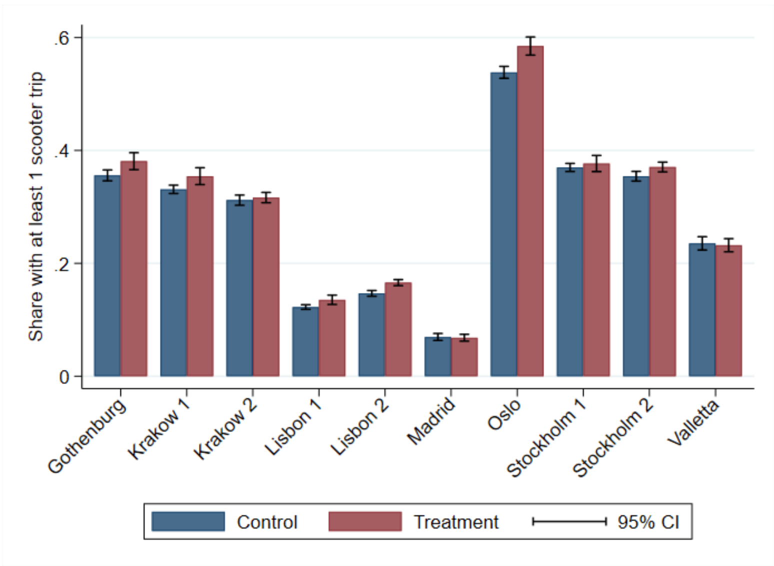

Figure B1.3: Share of users having conducted at least one e-scooter trip throughout the duration of the experiment, treatment and control.

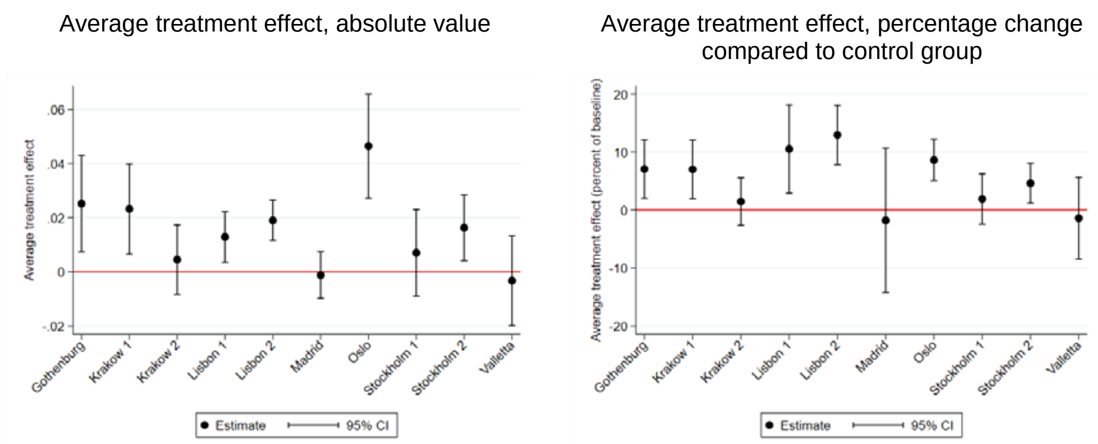

Figure B1.4: Average treatment effect. Outcome: Having conducted at least one e-scooter trip throughout the experiment.

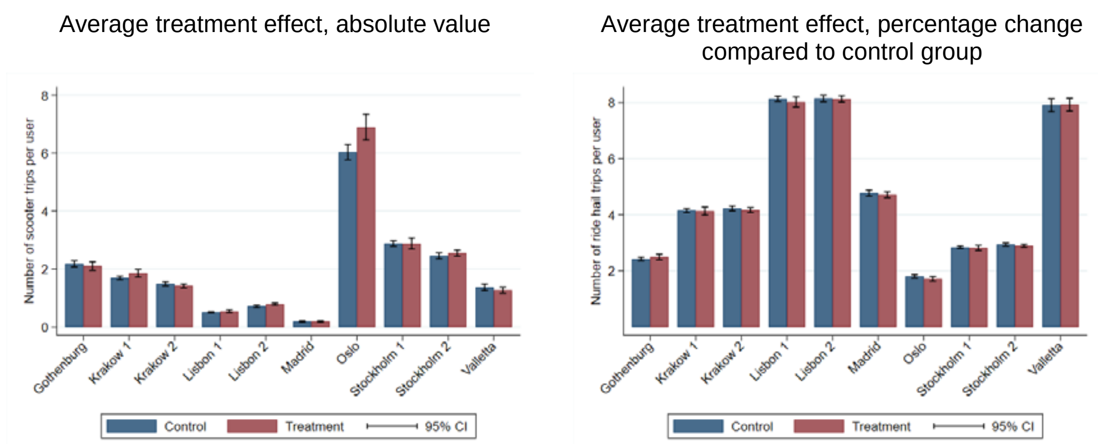

Figure B1.5: Average number of e-scooter and ride hail trips per user, treatment and control. Top 1 percent of most active users removed from both groups.

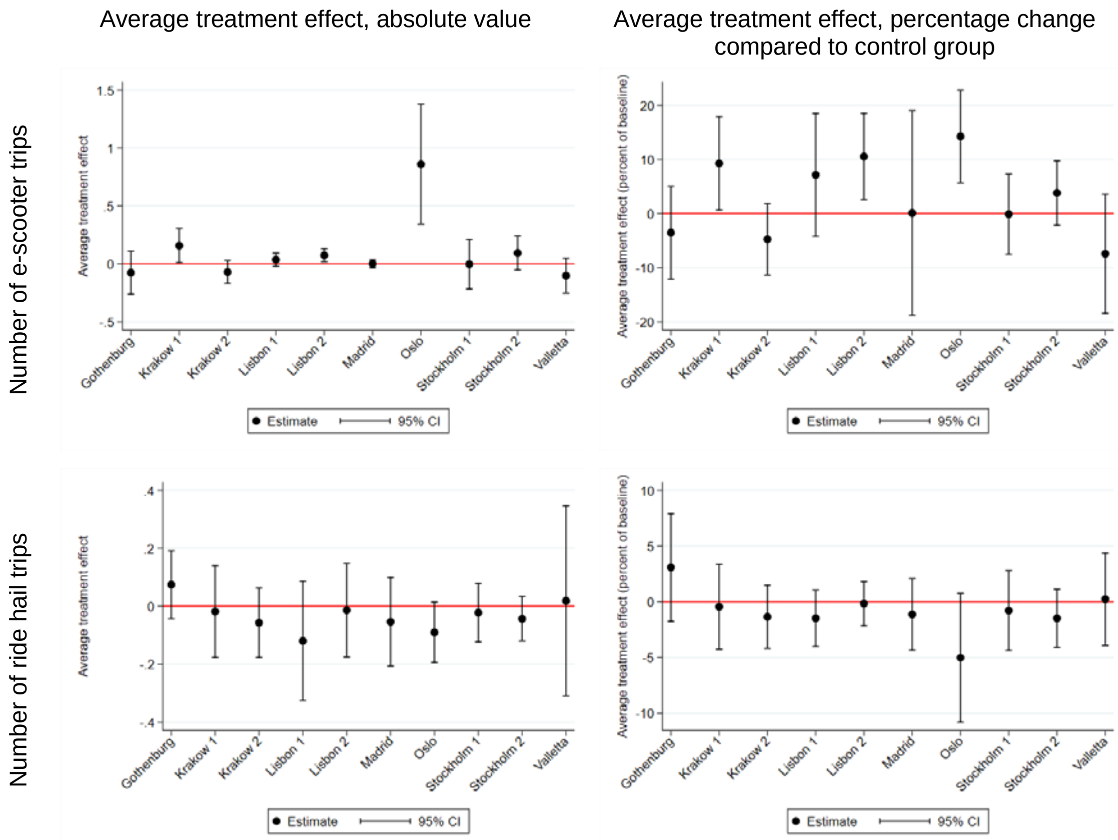

Figure B1.6: Average treatment effect, number of e-scooter and ride hail trips per user. Top 1 percent of most active users removed from both groups.

Finally, Figure B1.7 repeats the same analyses that are presented in the main text, but only considers outcomes app sessions that happened within 7 days since the user conducted her first relevant search session. By contrast, the figures presented in the main text measure outcomes from all search sessions within the duration of the experiment.

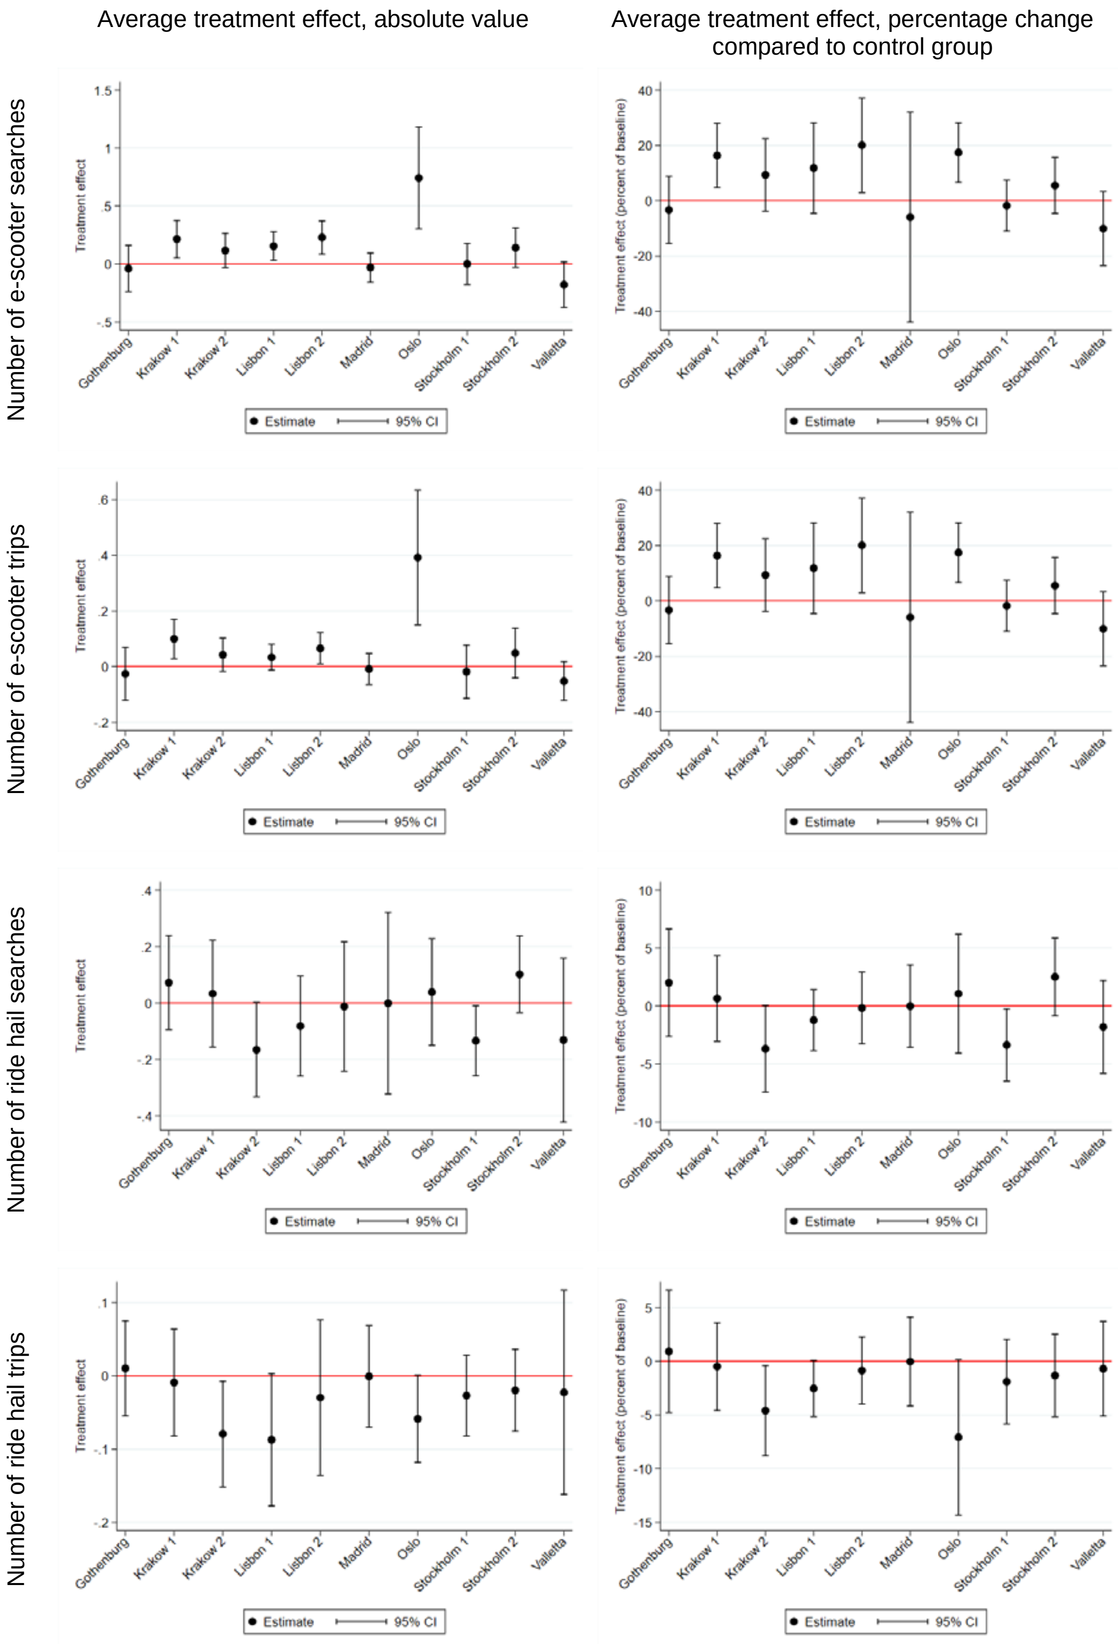

Figure B1.7: User-specific outcomes within one week after first nudge. Left: Average treatment effect in absolute value. Right: Average treatment effect relative to outcome in the control group.

B2: Heterogeneous effects

For each experiment, we have estimated heterogeneous treatment effects along eight different dimensions:

- Time of day
- Distance to nearest e-scooter
- Waiting time for ride hail driver
- Self-reported age of user
- The surge multiplier for ride hail
- Expected trip distance
- Expected price of e-scooter trip
- Expected price of ride hail trip

To examine heterogeneity over time of day, the sample is split into two-hour intervals, and the treatment effect is estimated separately for each interval. For the remaining dimensions, the sample is sorted according to the variable in question (e.g. distance to nearest e-scooter), split in five equally sized bins (i.e. quintiles), and the treatment effect is calculated separately for each bin.

Results are displayed in Figure B1.9-Figure B1.28. The left panels of the figures show changes in the probability of choosing e-scooters for both the treatment and the control group, while the right panels display the difference between them, i.e. the causal effect of being nudged for each group. To illustrate changes across quintiles, the point estimates are connected by lines. Before the figures are displayed, however, we will comment briefly on the main patterns along each dimension (**bold** in the text below).

Although there are clear patterns over **time of day** within some cities, there are differences in these patterns across cities. For the Scandinavian cities (Gothenburg, Oslo and Stockholm), the share of relevant search sessions resulting in e-scooter trips among both the treatment and the control group is higher in the afternoon. For Oslo in particular, this is also the time period where the effect of nudging is highest. Looking at the top rows of the descriptive figures in Appendix A2, this coincides with the time period e-scooter rides are most common in these cities. In Krakow and Lisbon, there is no clear pattern over time of day for when the control group chooses e-scooters. However, the nudging effect is higher in the afternoon for Krakow and in the morning for Lisbon. Looking at the bottom rows of the figures in Appendix A2, these are the time periods when the waiting time for ride hail trips spikes relative to the distance to the nearest e-scooter. This makes sense, given that waiting time for ride hail and distance to e-scooter should proxy these two modes' available supply relatively well.

The clearest pattern, true for every experiment, is that the probability of e-scooter trips is decreasing in **distance to nearest e-scooter** for both the treatment and the control group, and that the effect of nudging is larger if e-scooters are nearby. This is more thoroughly discussed in Section 4.1.2.

When it comes to **waiting time for ride hail drivers**, there are no significant differences in the treatment effects. However, the probability of choosing an e-scooter is higher when the waiting time for ride hail is longer, both in the treatment group and the control group as we would expect.

In the Scandinavian cities, the probability of choosing an e-scooter seems to be u-shaped in the **age distribution** of the user. But there are few age groups with significantly different treatment effects, and the pattern varies across cities. Note that the age variable is self-reported and only available for about 40 percent of the observations. Self-reported age is an in-app requirement to be able to use e-scooters, meaning that these search sessions will mechanically have a higher e-scooter share than the average.

The surge multiplier for ride hail is a scaling factor that increases prices of all ride hailing trips e.g. in periods where there is excess demand. Although we would expect more people to take e-scooters when the price is higher, this typically happens at night, specifically during weekends, when the probability of choosing an e-scooter is low in the first place. For the Scandinavian cities, the trend is therefore that the probability of choosing e-scooters is higher when the surge multiplier is low. Lisbon is the clearest exception to this, where the treatment effect is significantly higher for higher surge multipliers (see top right panel of Figure B1.18). In Lisbon the probability of choosing an e-scooter is more evenly distributed throughout the day to begin with (see top left panel of Figure B1.17), meaning that heterogeneity along the surge multiplier dimension to a smaller extent is driven by other factors that correlate with time of day.

The probability of choosing an e-scooter is also decreasing in **expected trip distance**, for both the treatment and the control group. This variable is derived from the GPS coordinate of the user and the coordinate she specifies as the destination of the ride hail trip. The trend is that the effect of nudging is larger for shorter distances, but the treatment effect is only significantly different for some of the experiments.

The two last dimensions are the **prices of e-scooter and ride hail trips** respectively. Note that since prices are predicted based on expected trip distances, differences in prices are largely driven by differences in distances. For ride hail, the surge multiplier will provide additional price variation. For e-scooters we also observe the hourly and fixed rates for each trip, but they are typically constant within each experiment. This means that for most experiments, heterogeneity along the "trip distance" and "expected e-scooter price" dimensions will be identical, since sorting search sessions along these dimensions will give the same ordering. Two exceptions are Gothenburg and Oslo, where the price actually varied over the experiment. While the subsequent figures (Figure B1.9-Figure B1.28) focus on prices per trip, Figure B1.8 displays how the e-scooter probability varies across hourly price rates for these two cities.

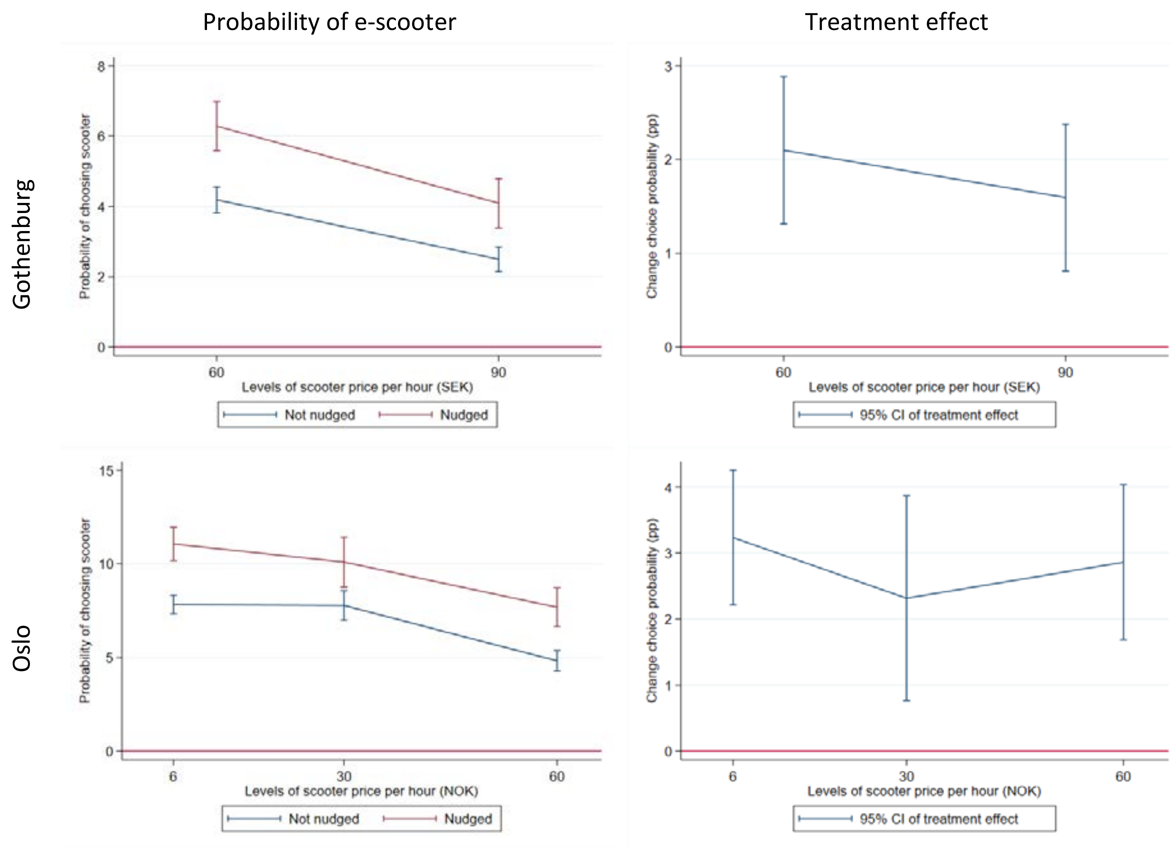

Figure B1.8: Effect of being nudged on probability of e-scooter by levels of e-scooter rate per hour. Gothenburg and Oslo. During the summer of 2021, 1 SEK ≈ 1 NOK ≈0.1 EUR.

During parts of the Oslo experiment, Bolt held a campaign where the price per hour was only 6 NOK (≈0.6 EUR). Remarkably, the probability of choosing an e-scooter for both the treatment and the control group is about the same during the time period when prices were 5 times higher (30 NOK per hour). We do observe a significant reduction in the e-scooter probability for both the treatment and the control group (left panel) when the price changes from 30 to 60 NOK per hour in Oslo, and 60 to 90 SEK per hour in Gothenburg. However, the effect of being nudged does not seem to depend on the e-scooter rate (right panel).

The rest of this section displays the results of the heterogeneity analysis in figures.

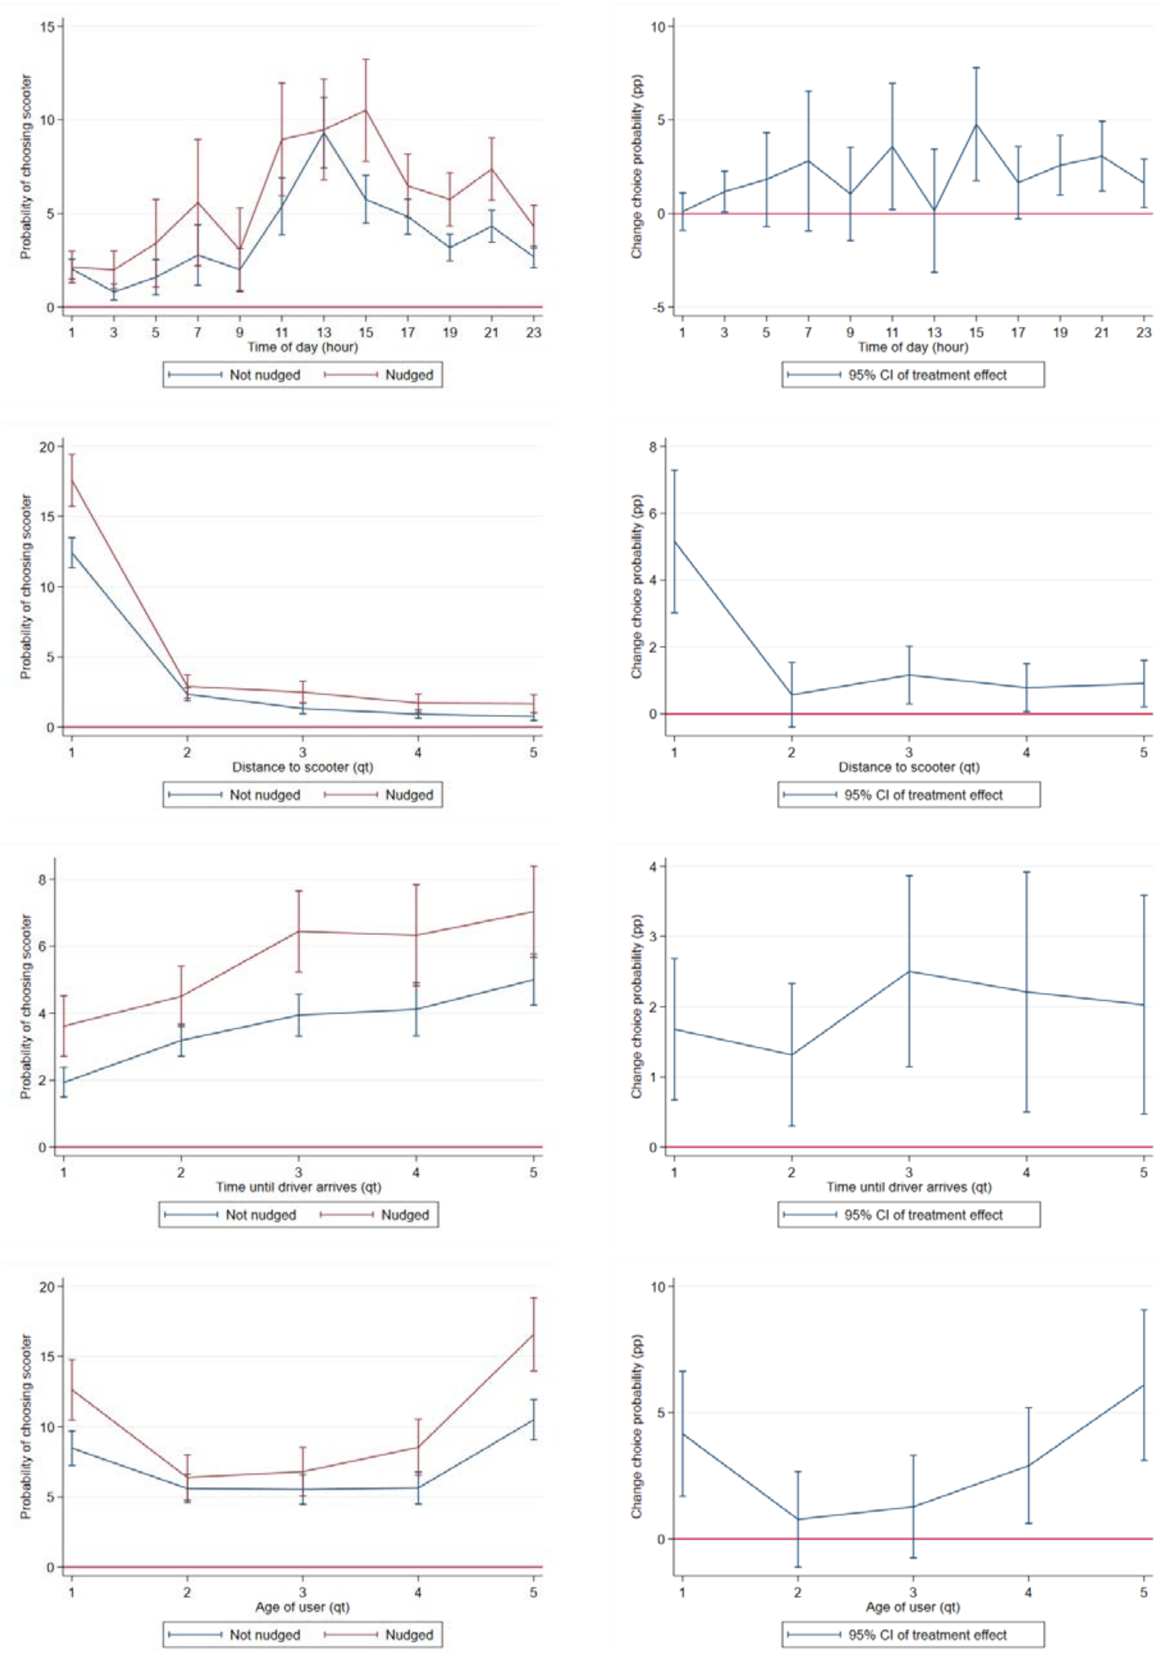

Figure B1.9: Heterogeneous effects of nudging on probability of e-scooter trip. Gothenburg.

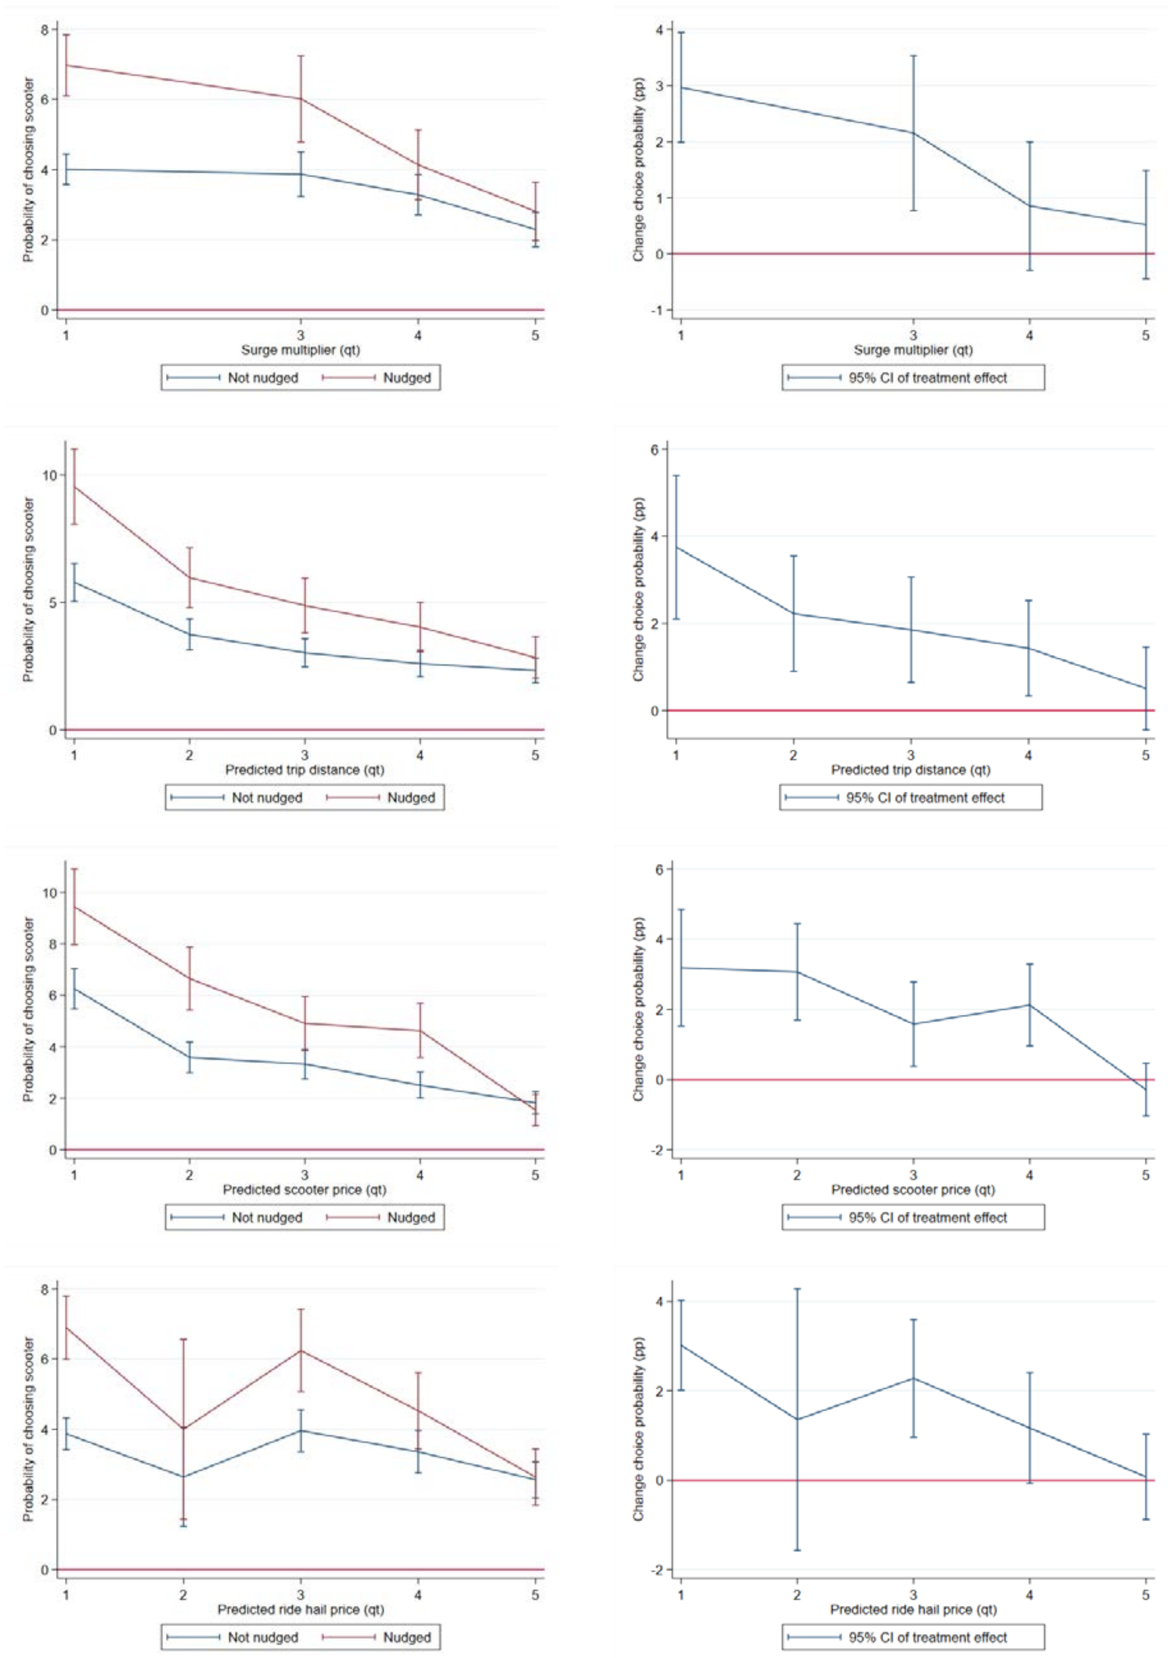

Figure B1:10: Heterogeneous effects of nudging on probability of e-scooter trip. Gothenburg.

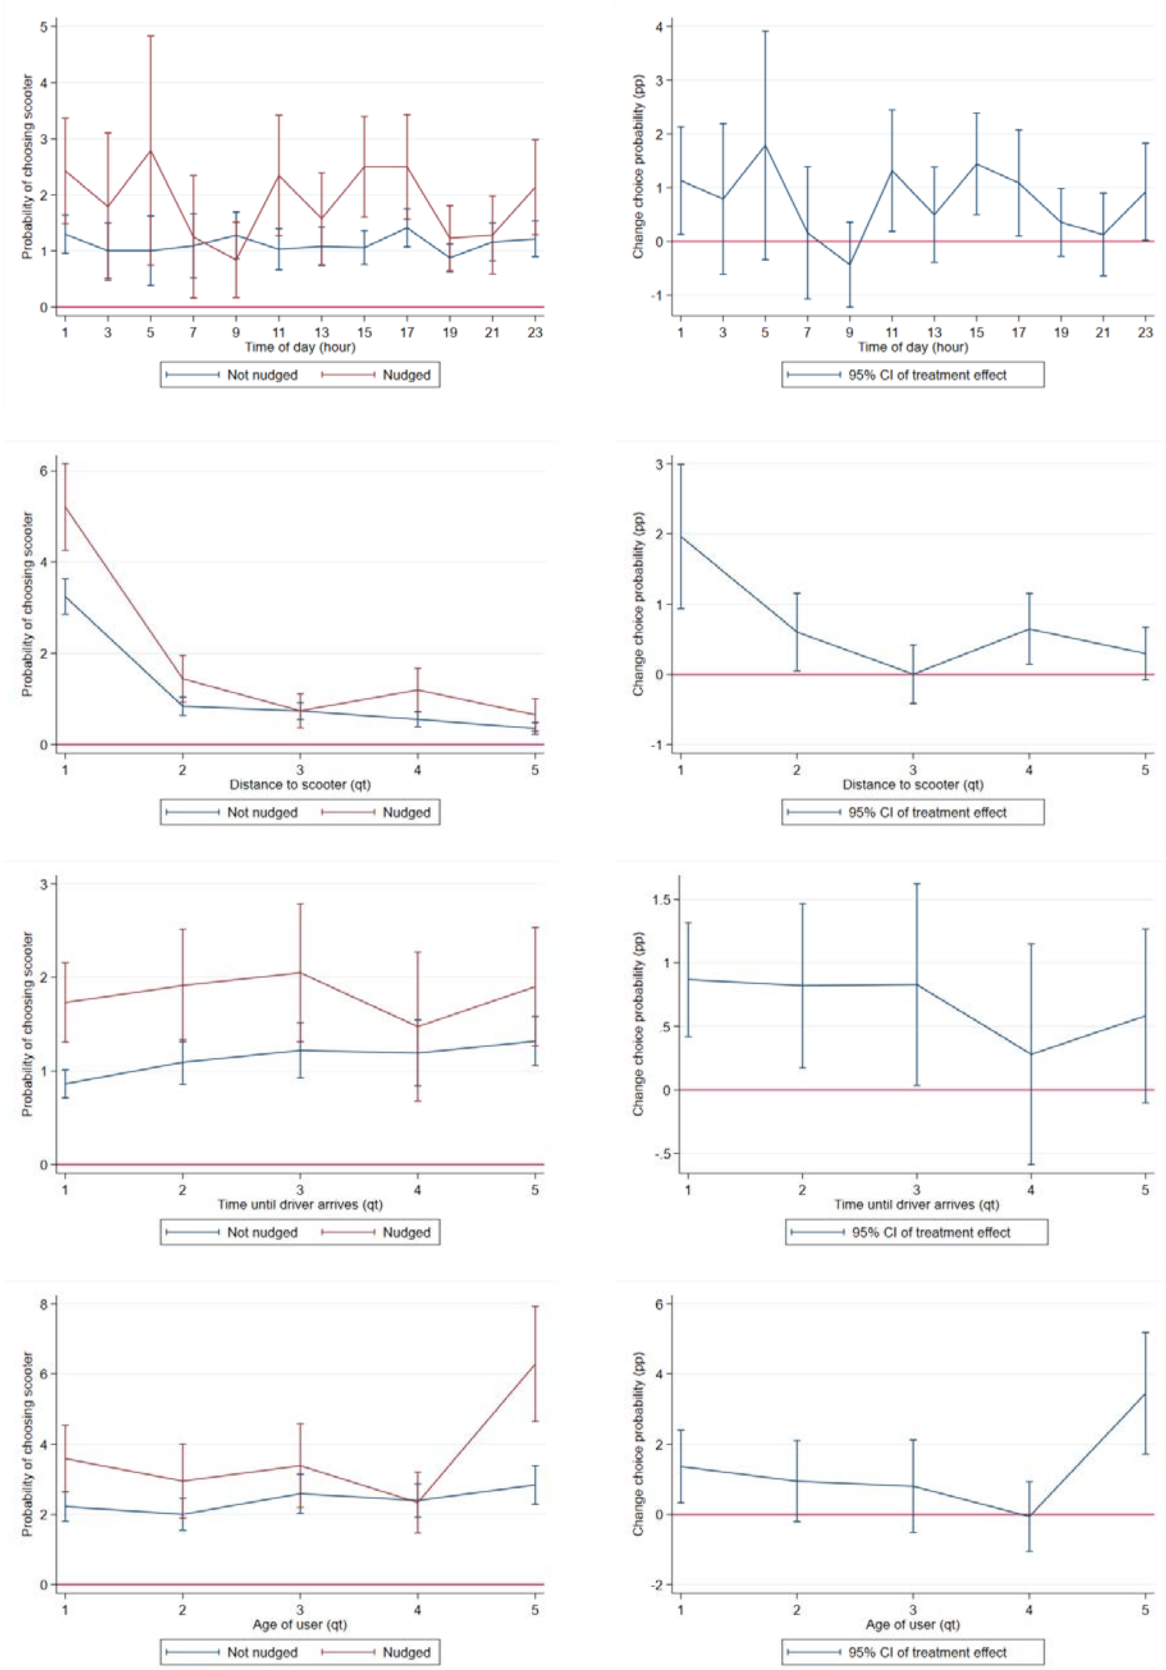

Figure B1.11: Heterogeneous effects of nudging on probability of e-scooter trip. Krakow 1.

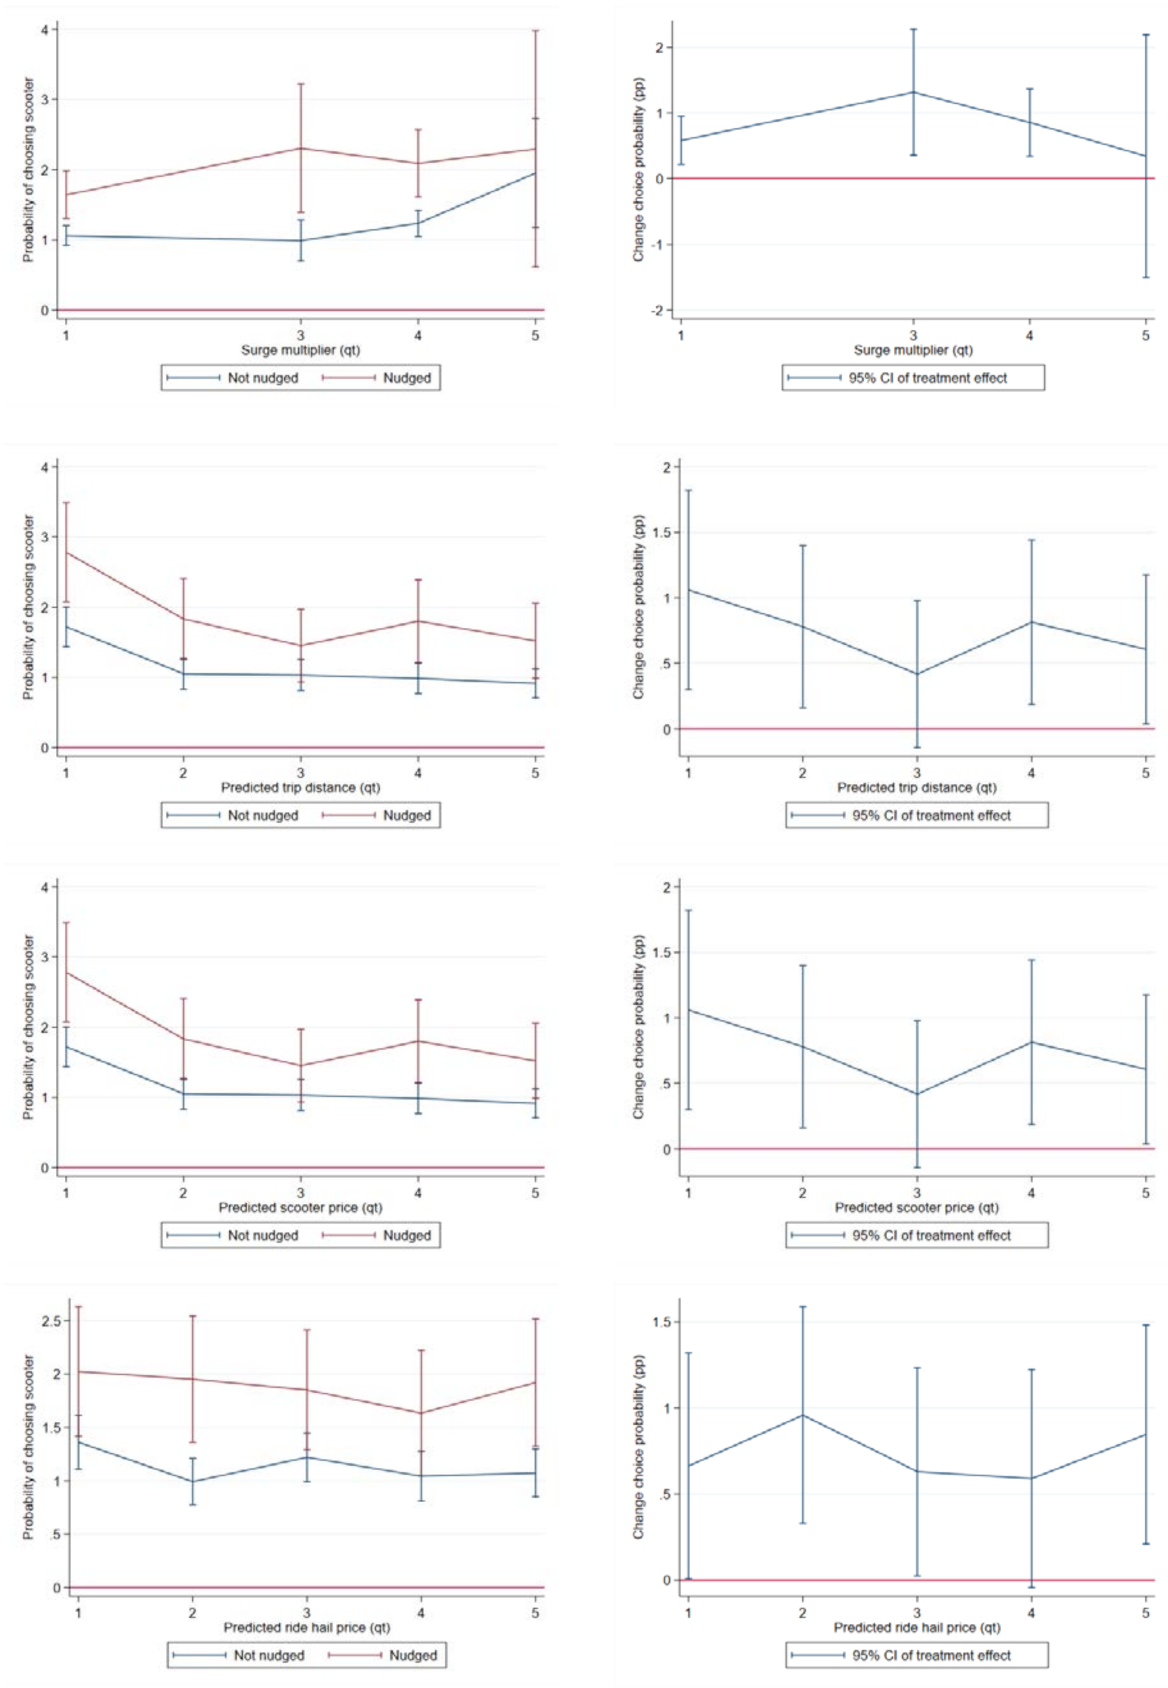

Figure B1.12: Heterogeneous effects of nudging on probability of e-scooter trip. Krakow 1.

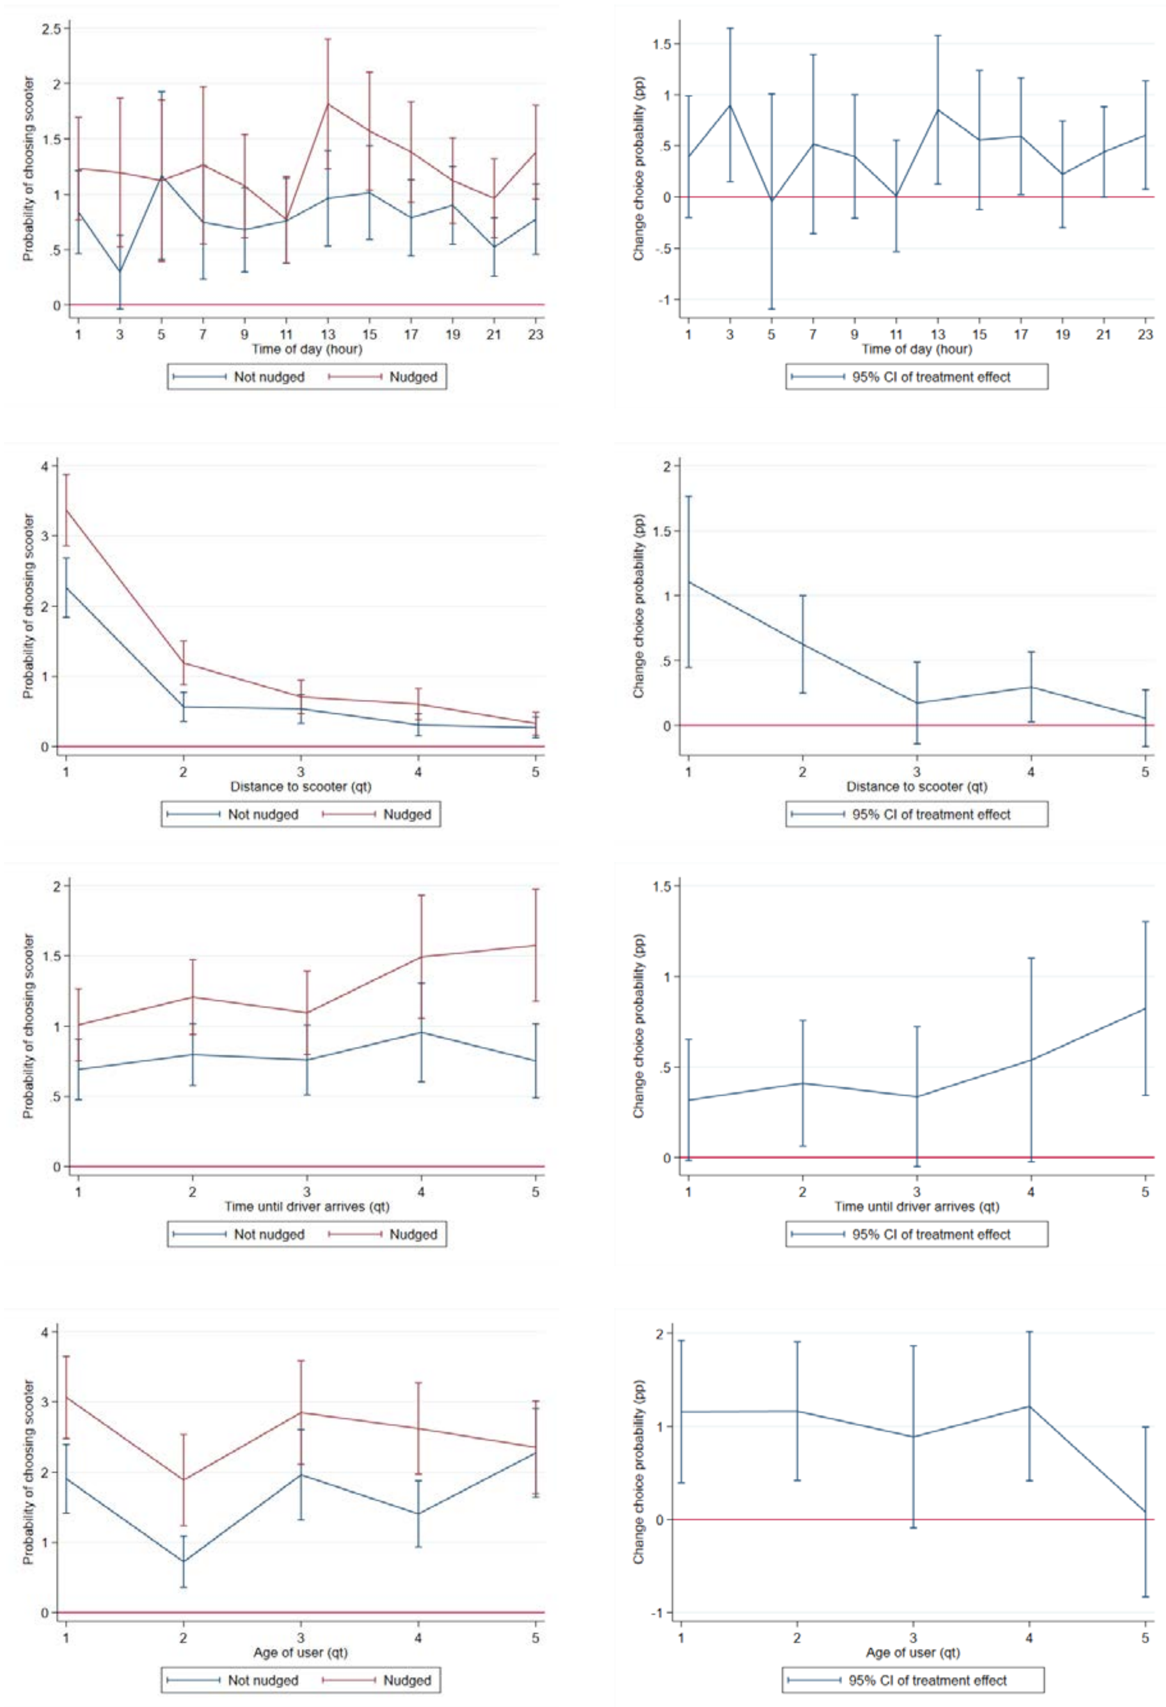

Figure B1.13: Heterogeneous effects of nudging on probability of e-scooter trip. Krakow 2.

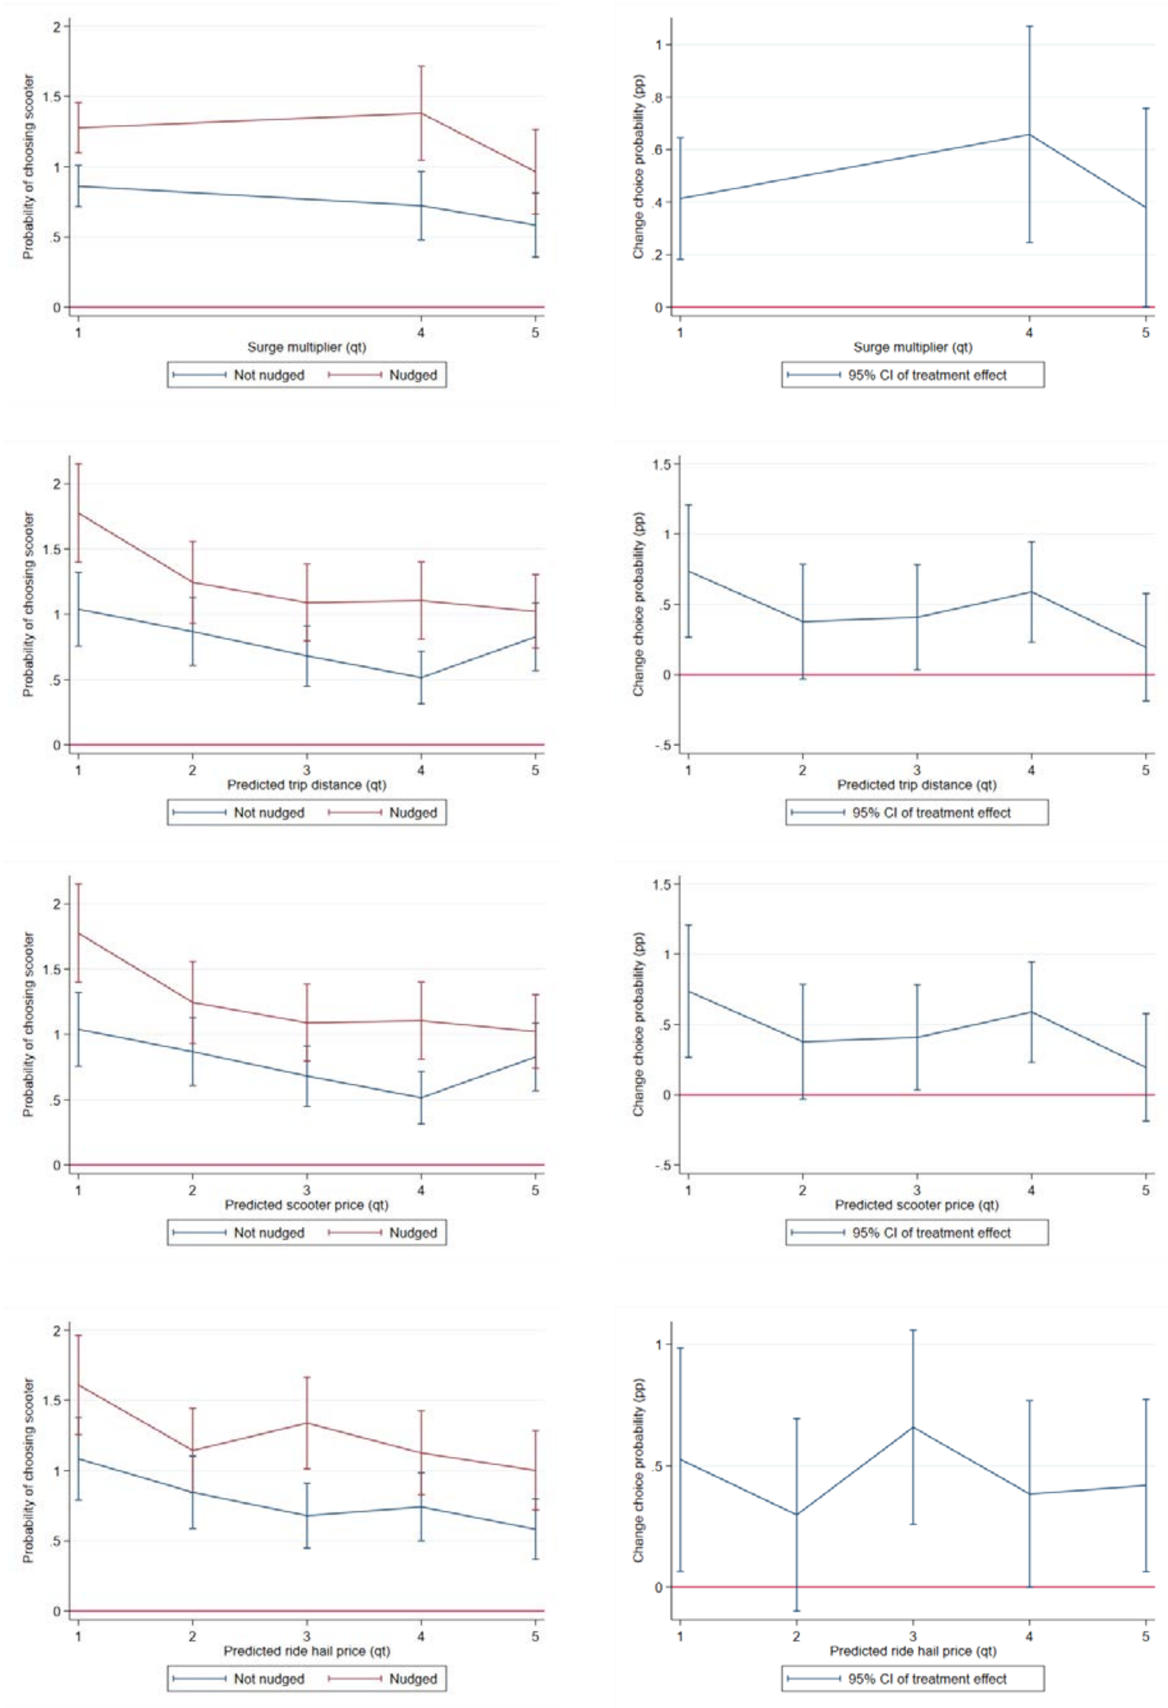

Figure B1.14: Heterogeneous effects of nudging on probability of e-scooter trip. Krakow 2.

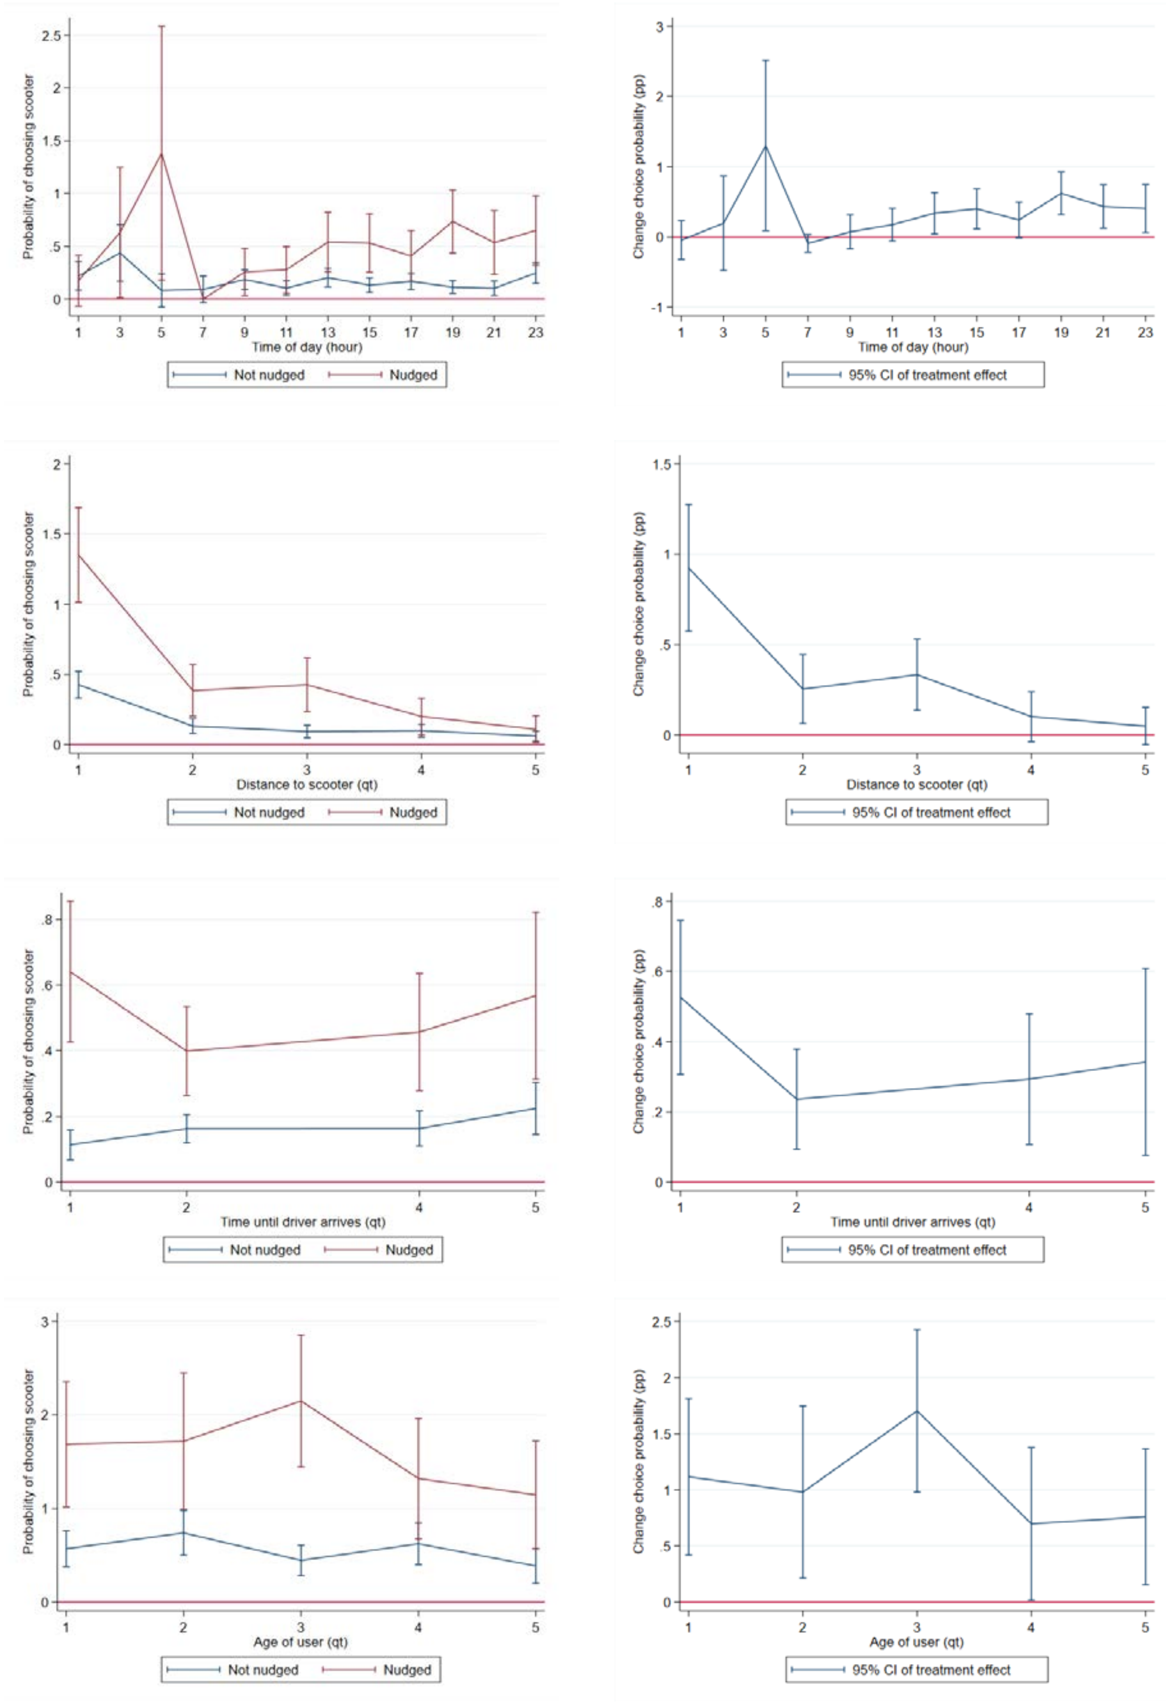

Figure B1.15: Heterogeneous effects of nudging on probability of e-scooter trip. Lisbon 1.

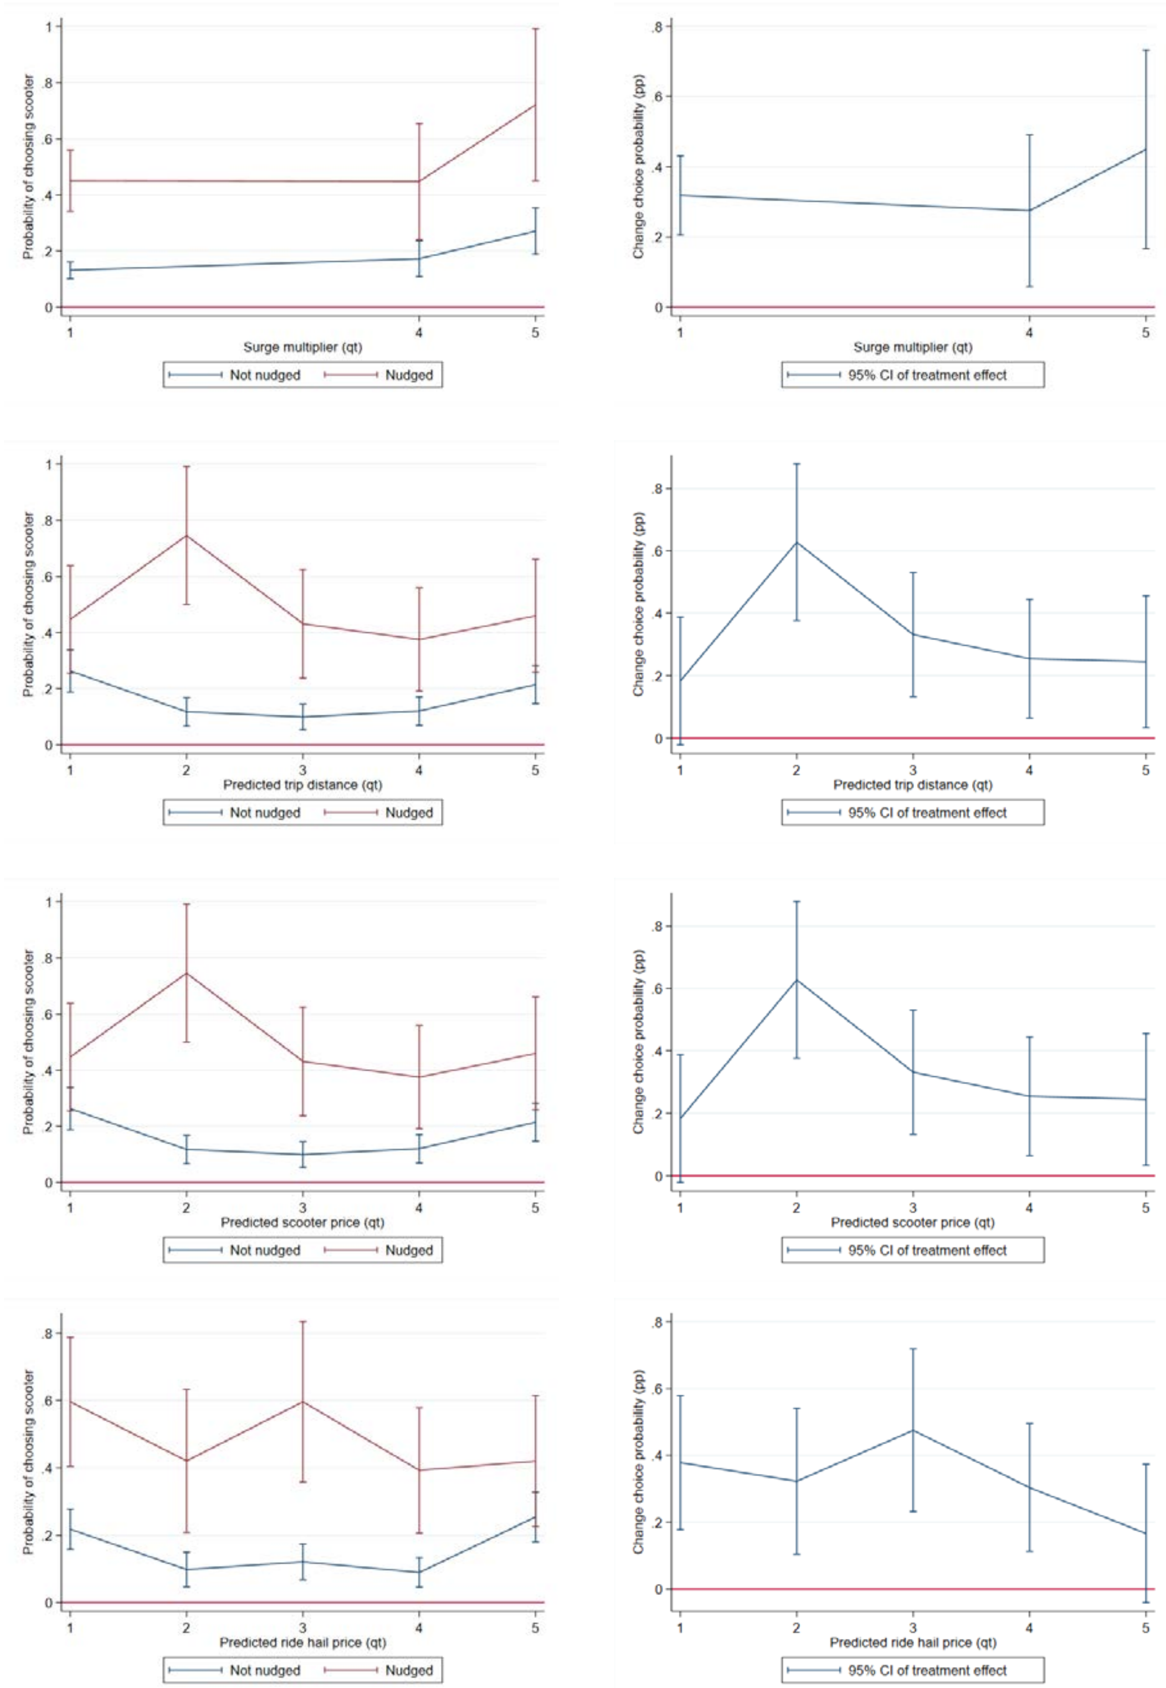

Figure B1.16: Heterogeneous effects of nudging on probability of e-scooter trip. Lisbon 1.

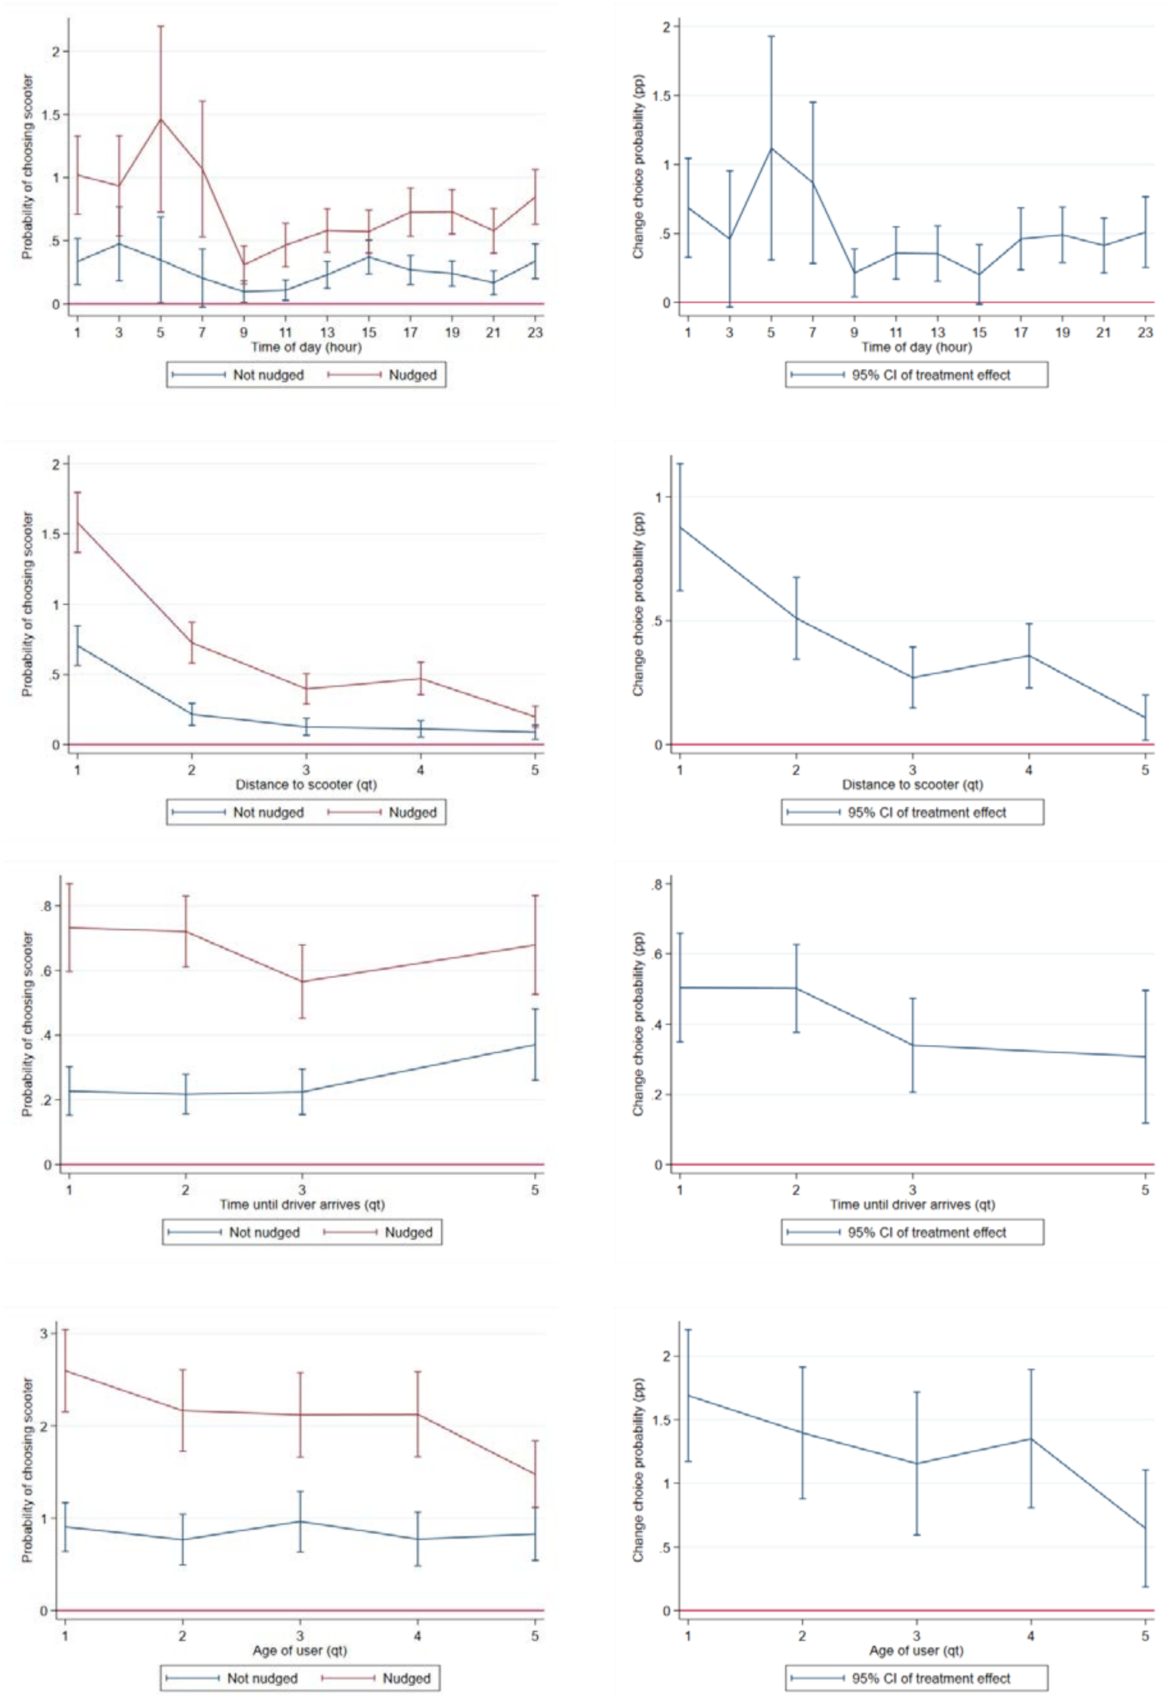

Figure B1.17: Heterogeneous effects of nudging on probability of e-scooter trip. Lisbon 2.

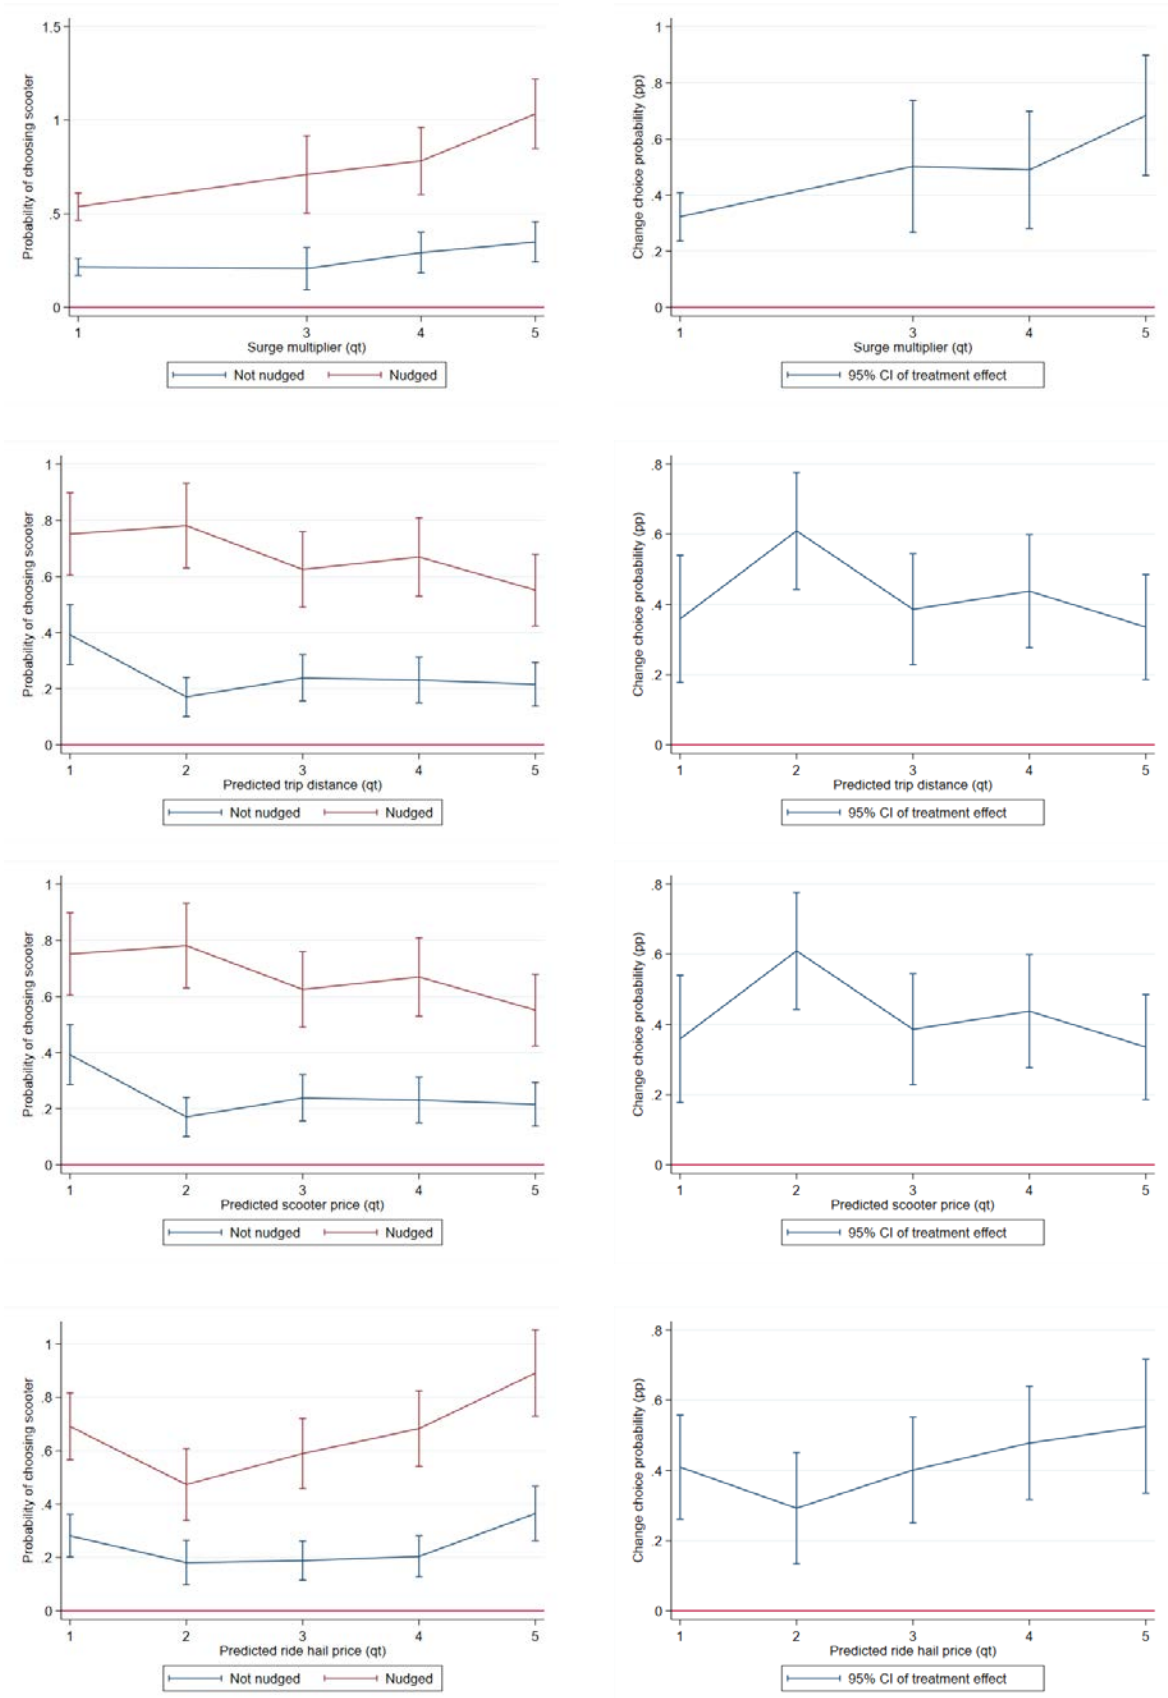

Figure B1.18: Heterogeneous effects of nudging on probability of e-scooter trip. Lisbon 2.

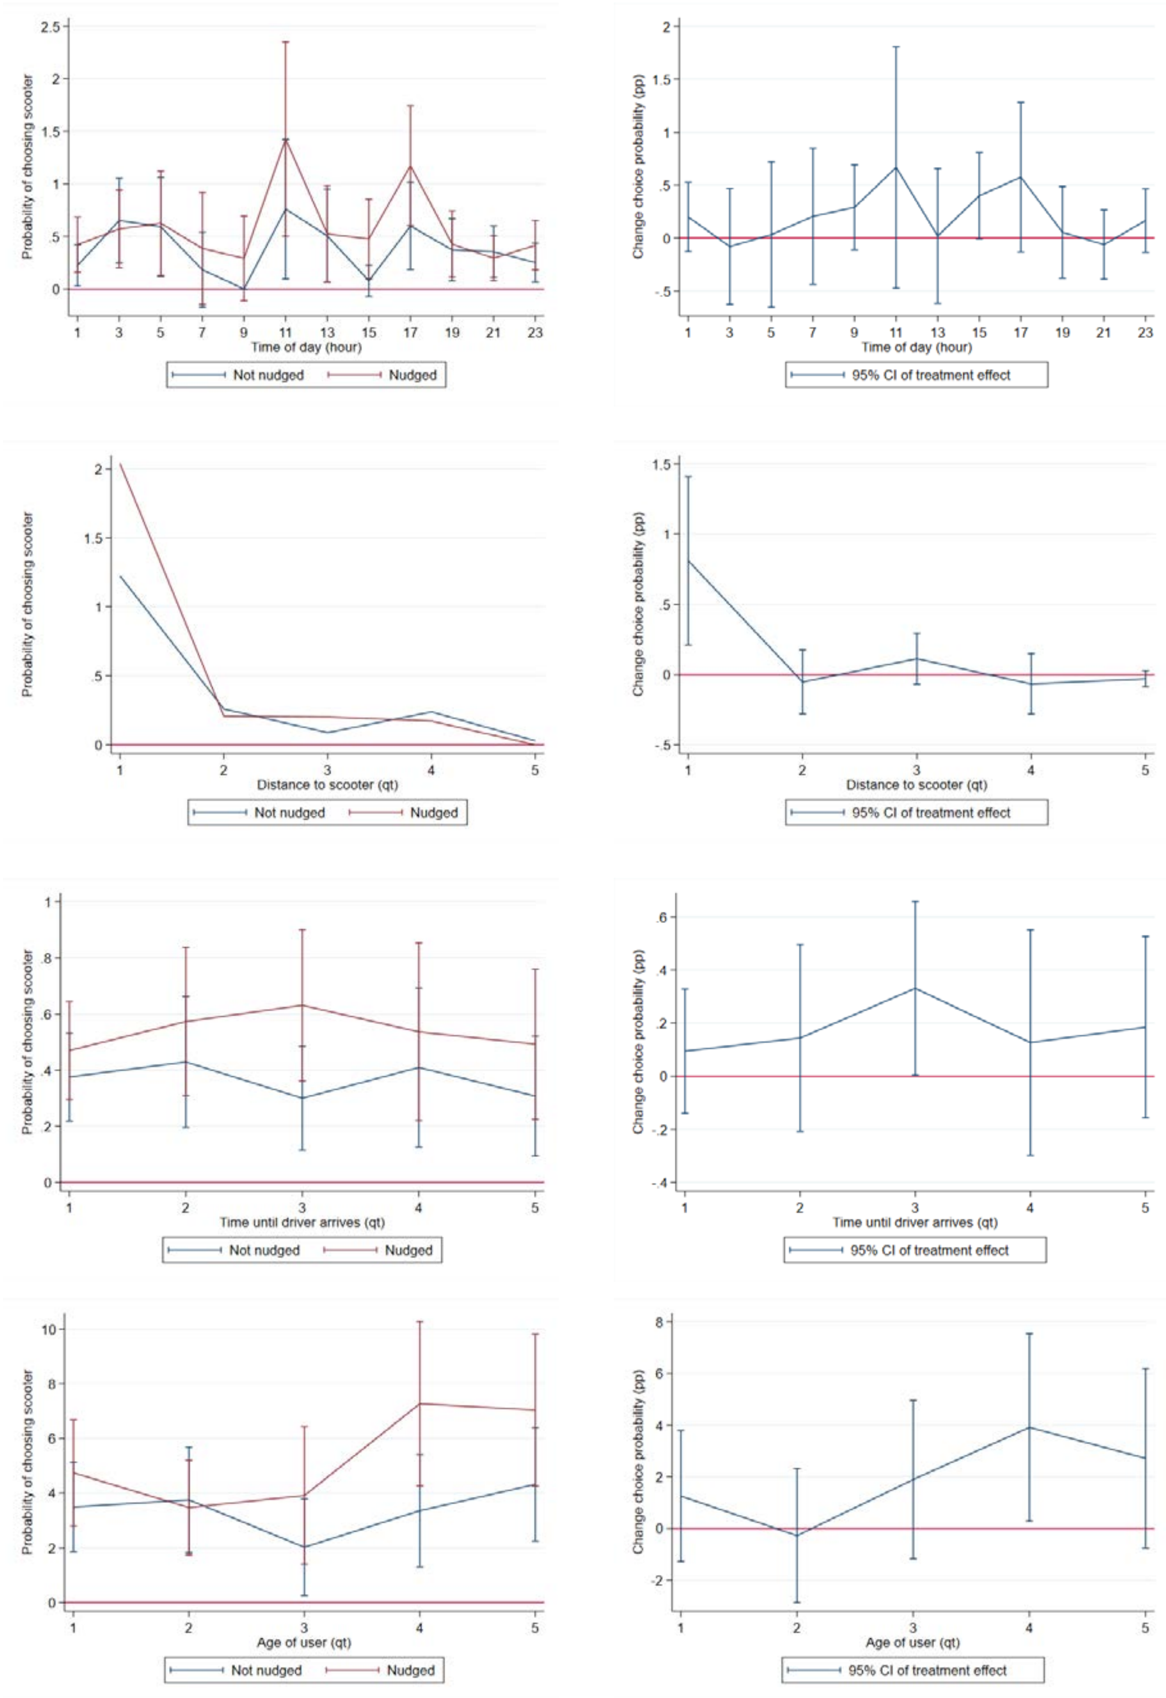

Figure B1.19: Heterogeneous effects of nudging on probability of e-scooter trip. Madrid.

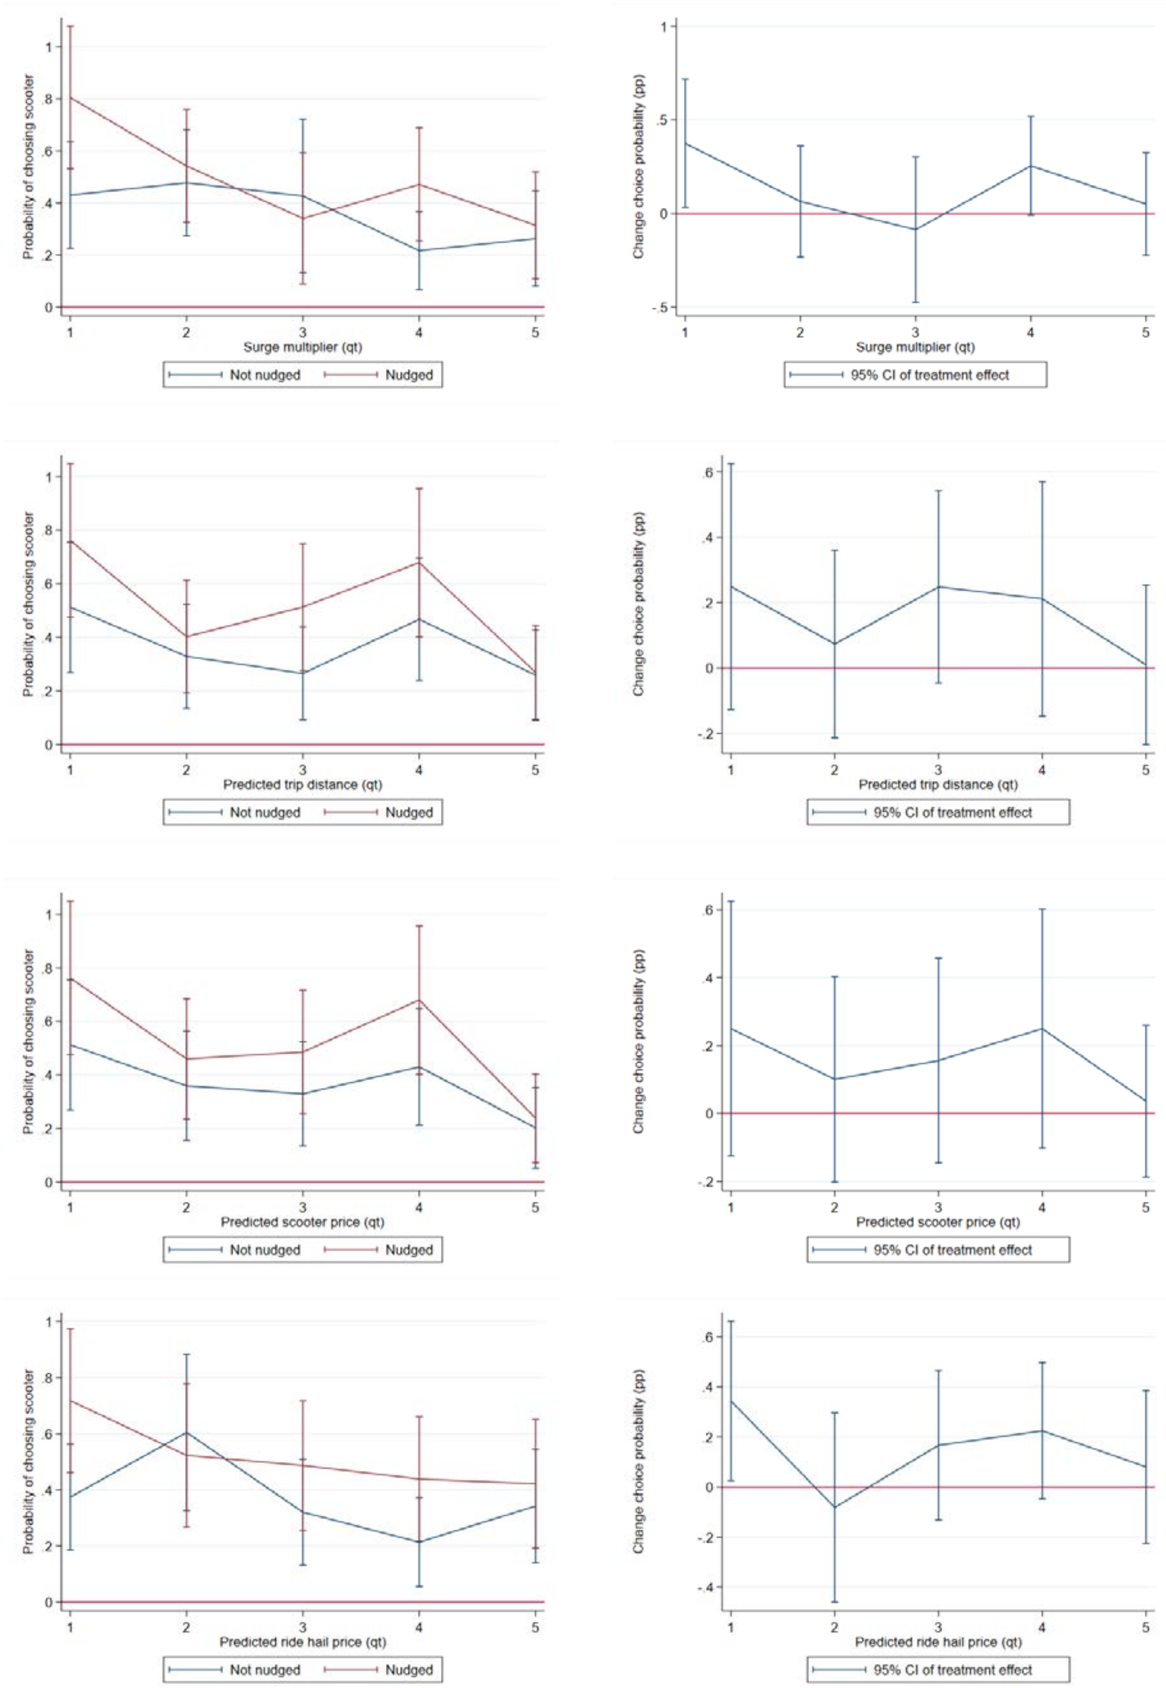

Figure B1.20: Heterogeneous effects of nudging on probability of e-scooter trip. Madrid.

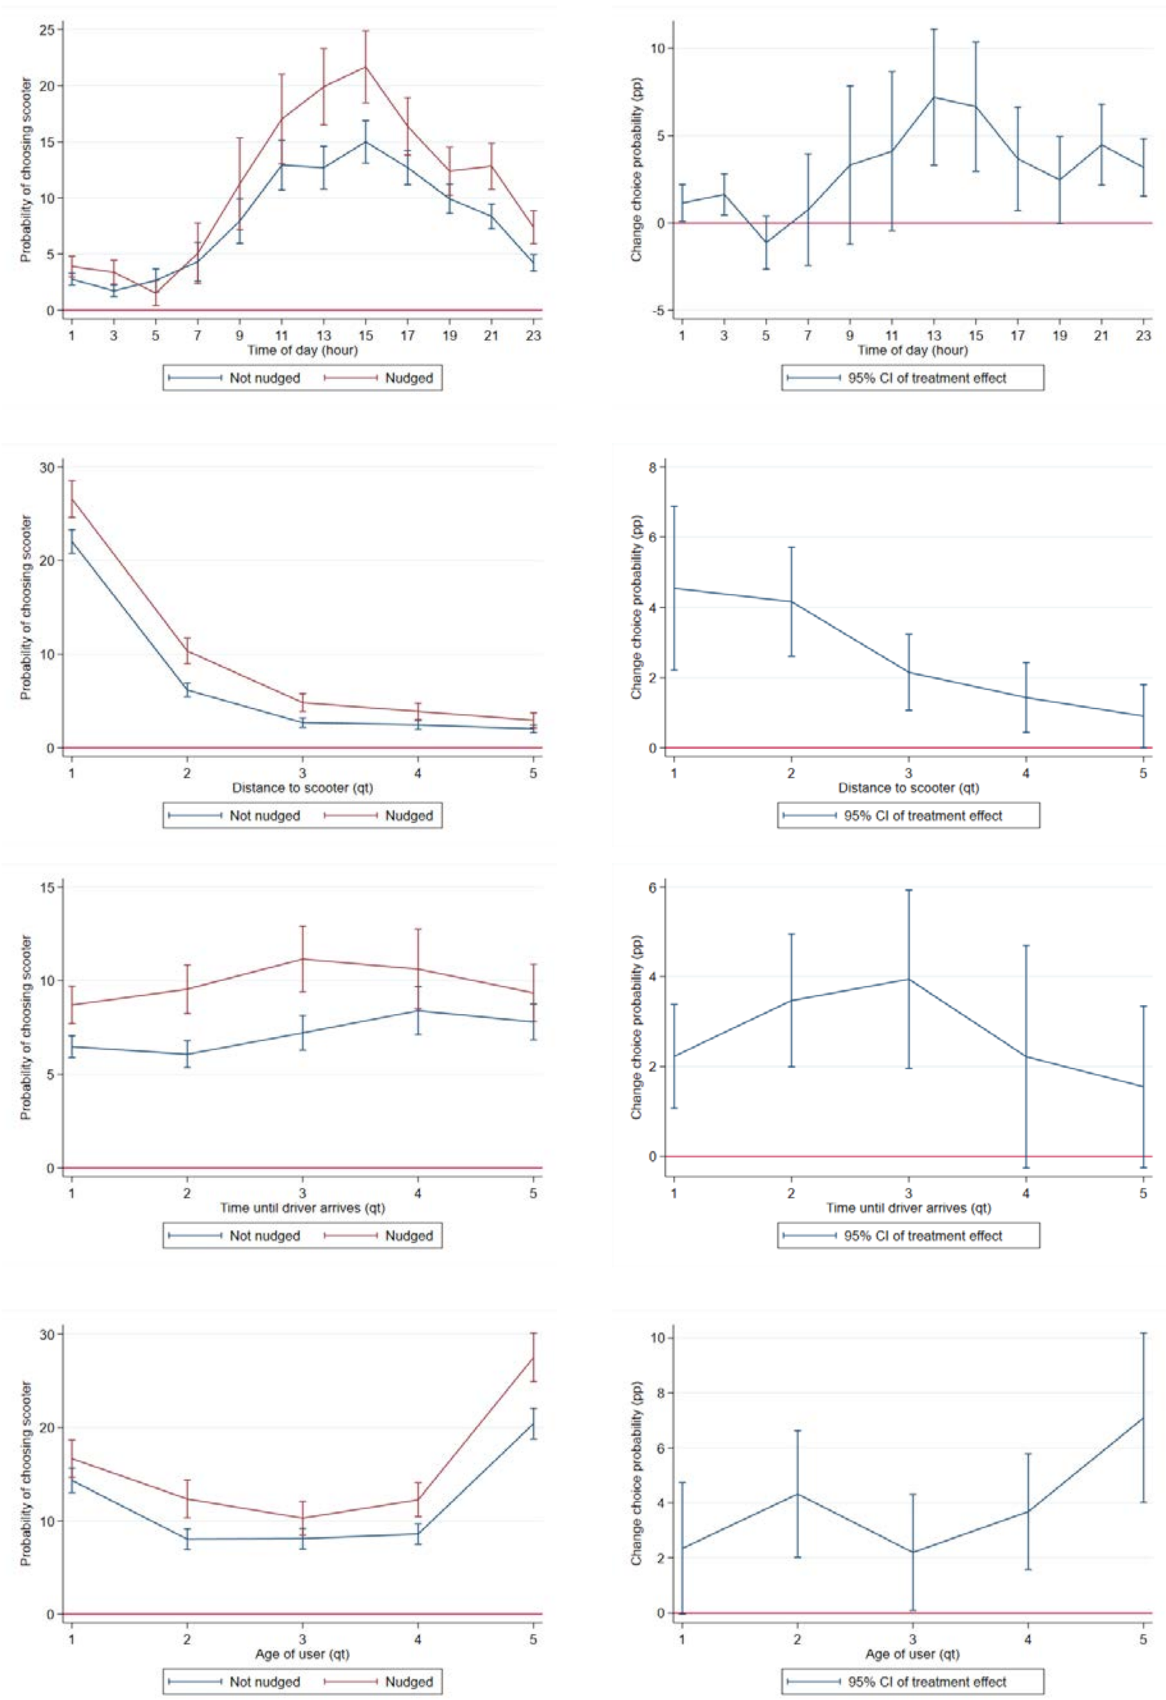

Figure B1.21: Heterogeneous effects of nudging on probability of e-scooter trip. Oslo.

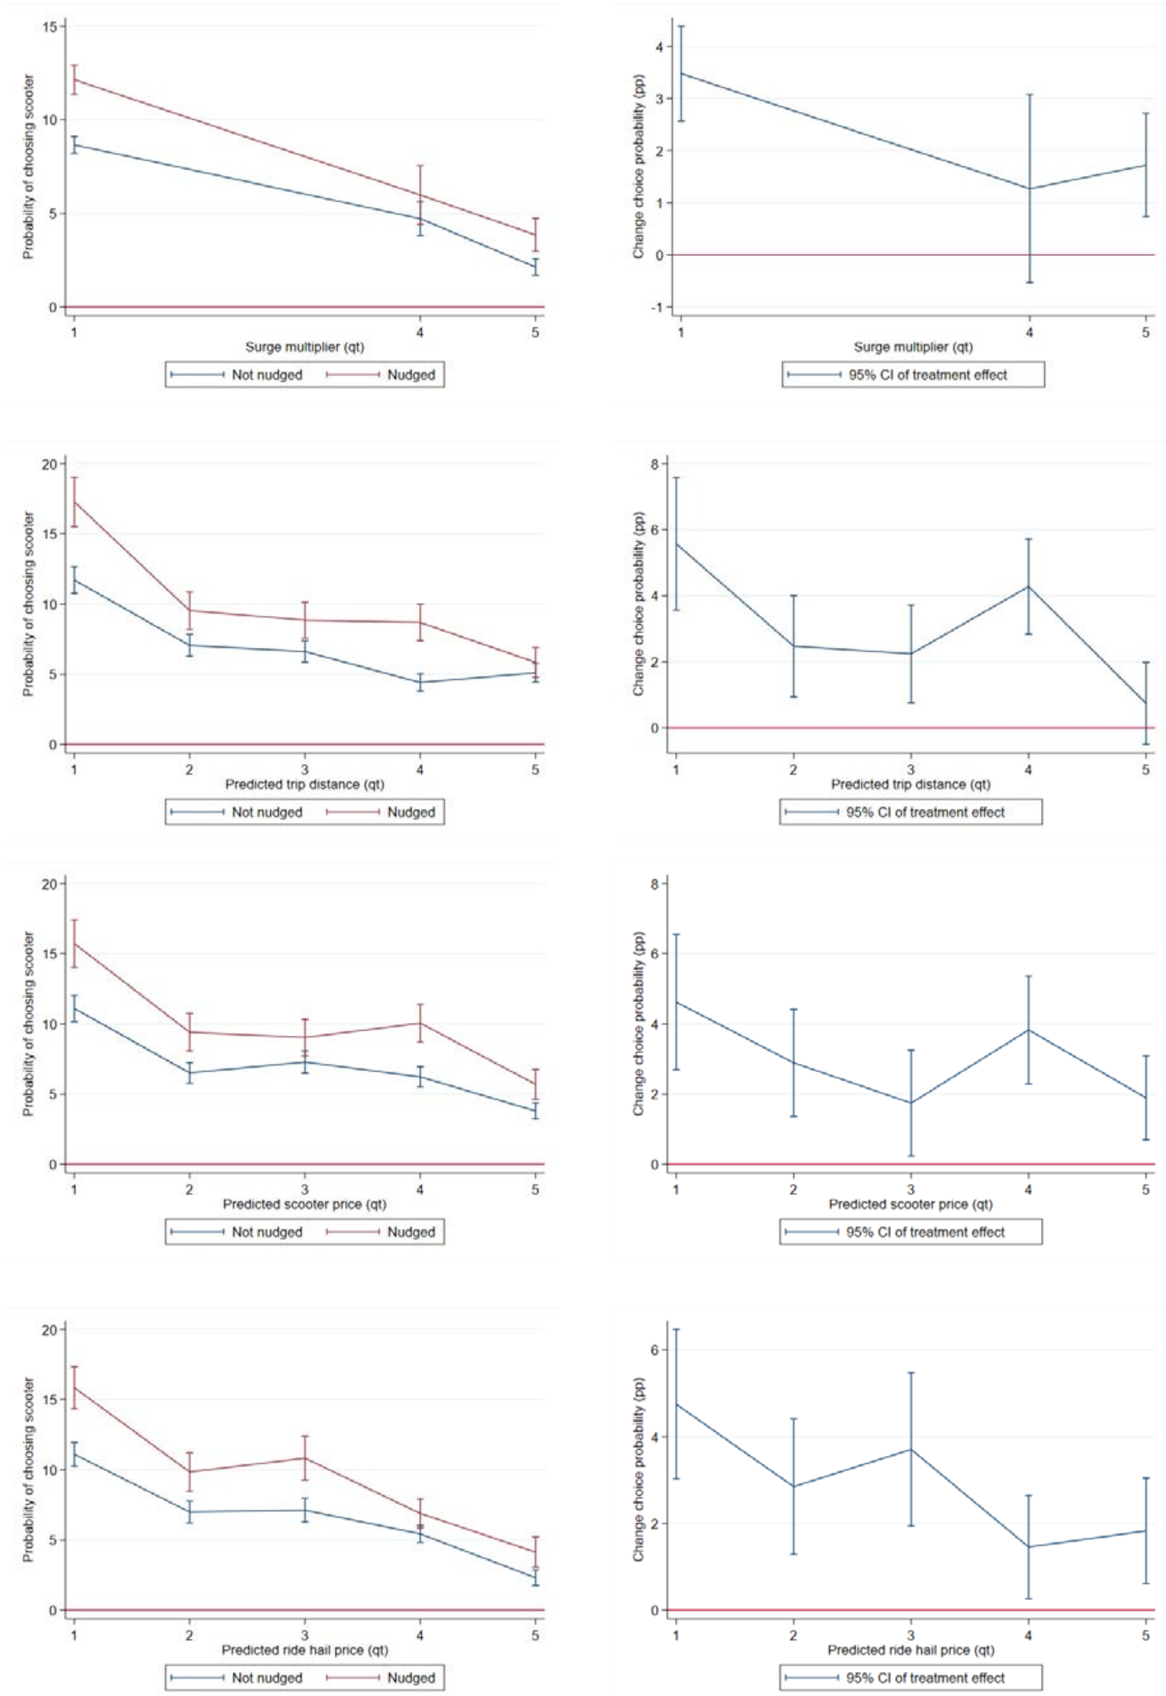

Figure B1.22: Heterogeneous effects of nudging on probability of e-scooter trip. Oslo.

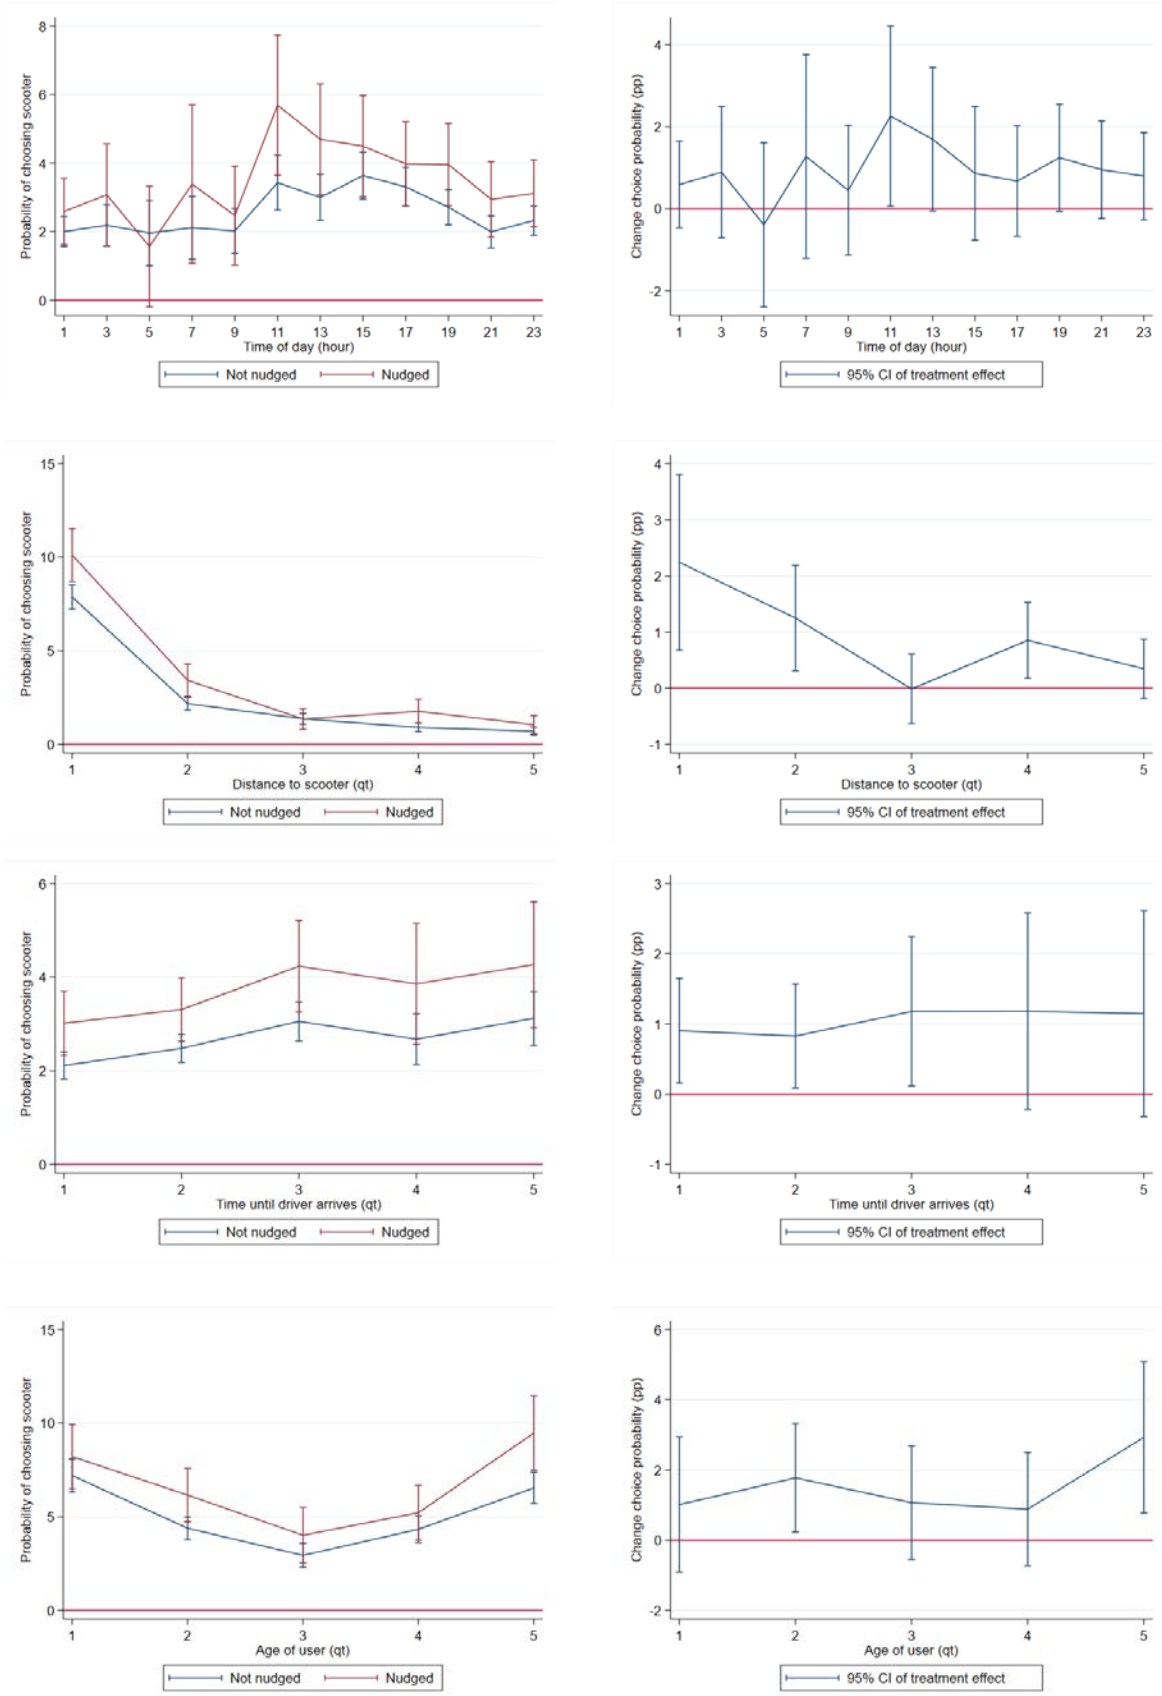

Figure B1.23: Heterogeneous effects of nudging on probability of e-scooter trip. Stockholm 1.

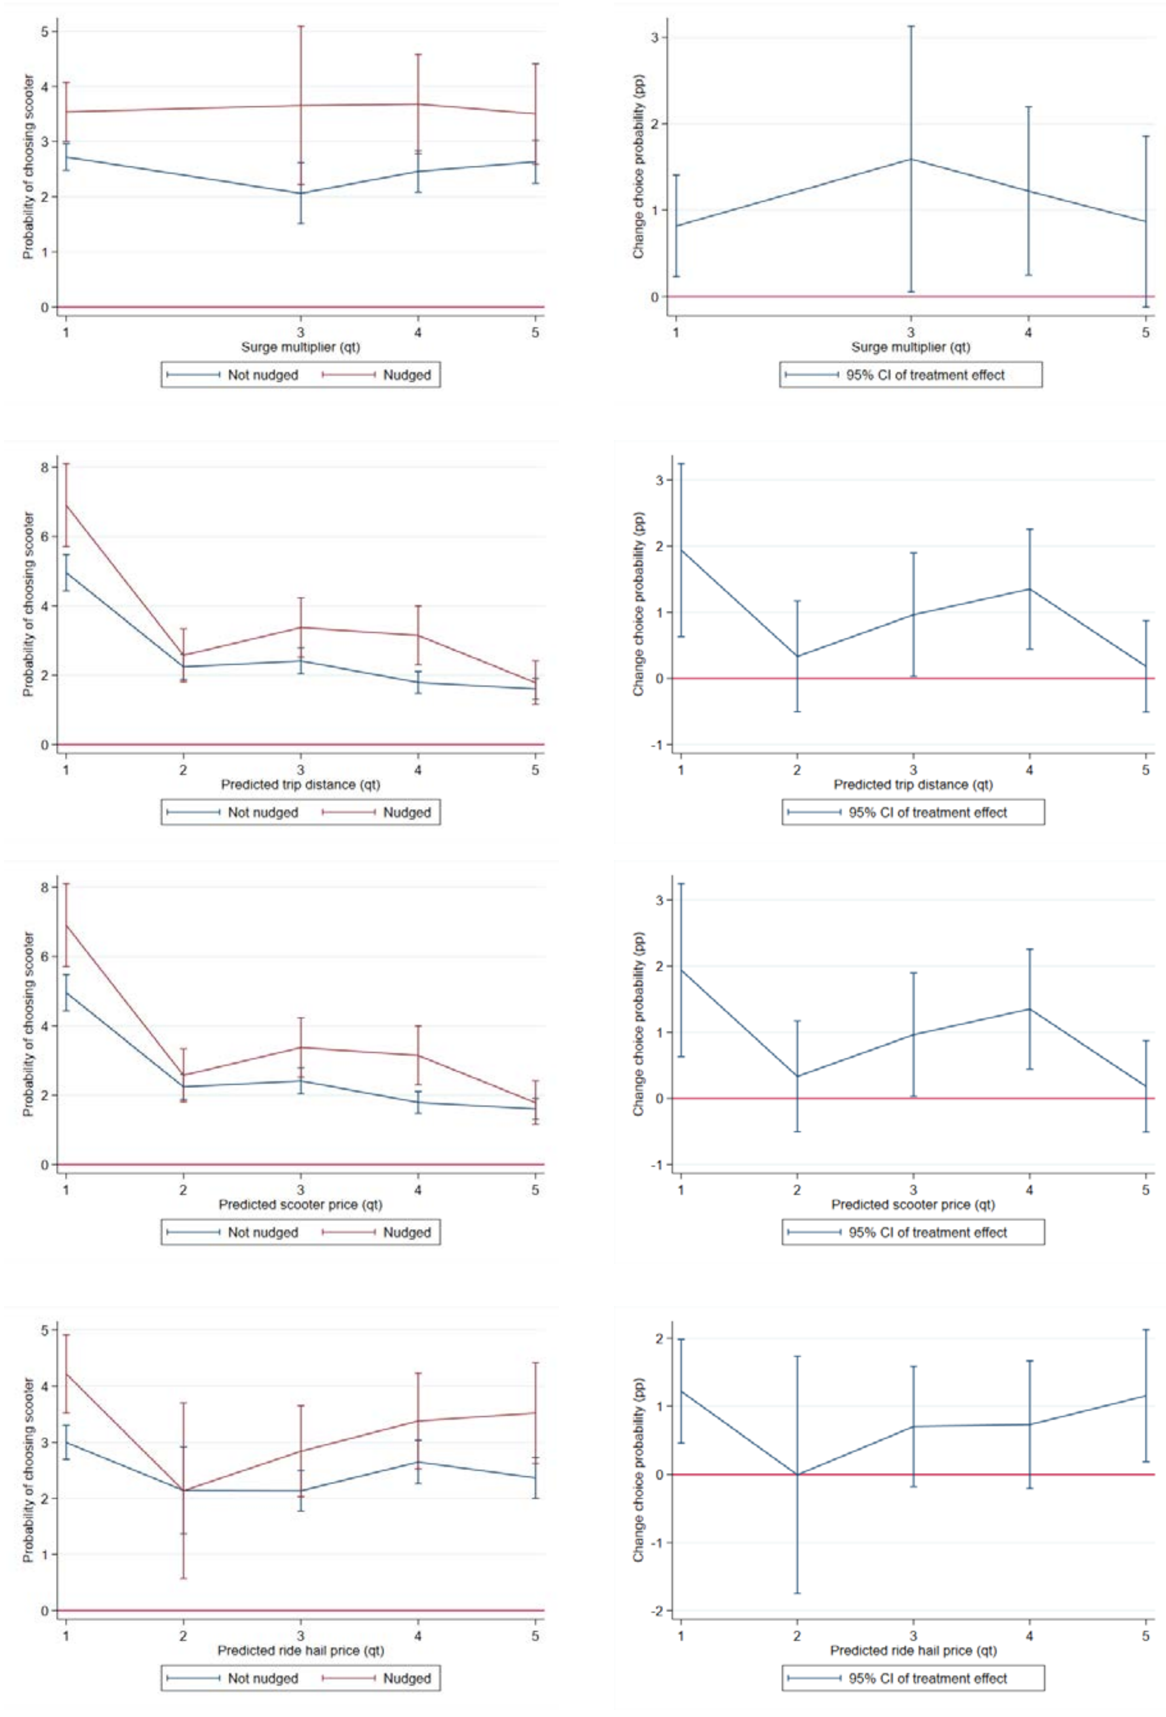

Figure B1.24: Heterogeneous effects of nudging on probability of e-scooter trip. Stockholm 1.

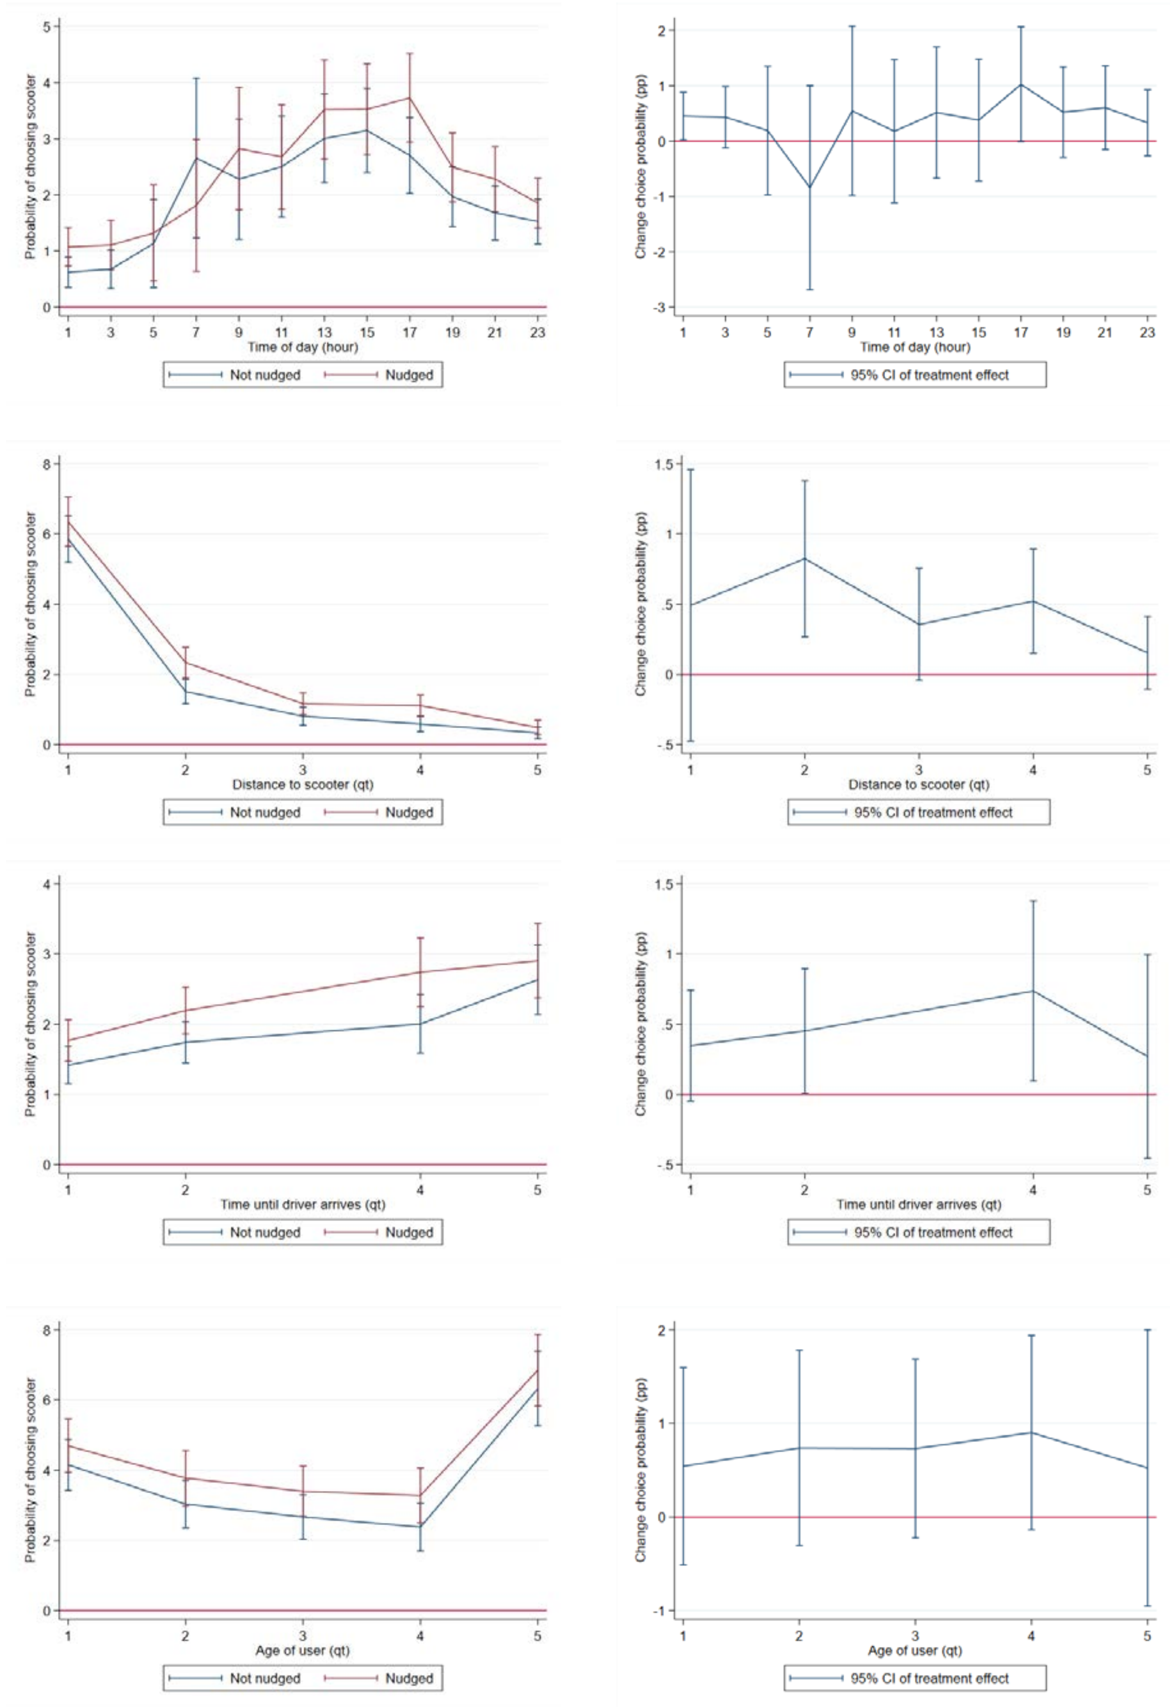

Figure B1.25: Heterogeneous effects of nudging on probability of e-scooter trip. Stockholm 2.

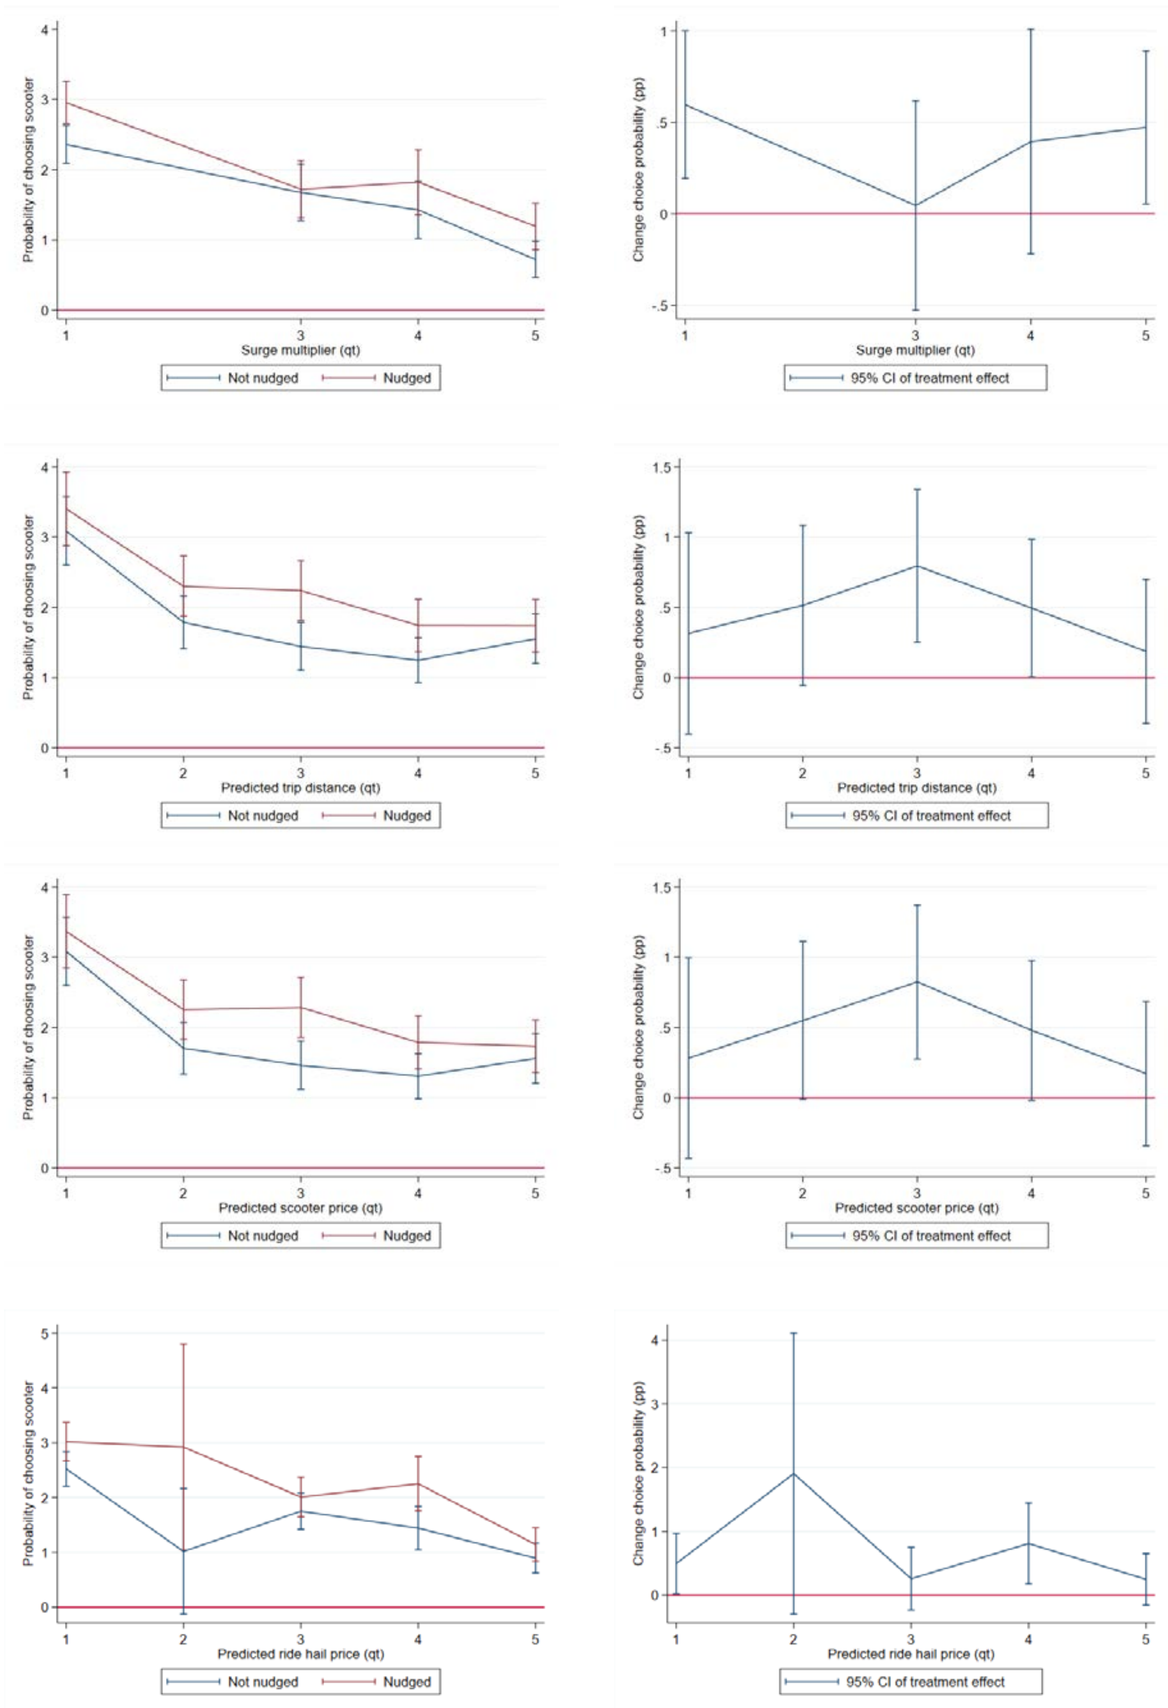

Figure B1.26: Heterogeneous effects of nudging on probability of e-scooter trip. Stockholm 2.

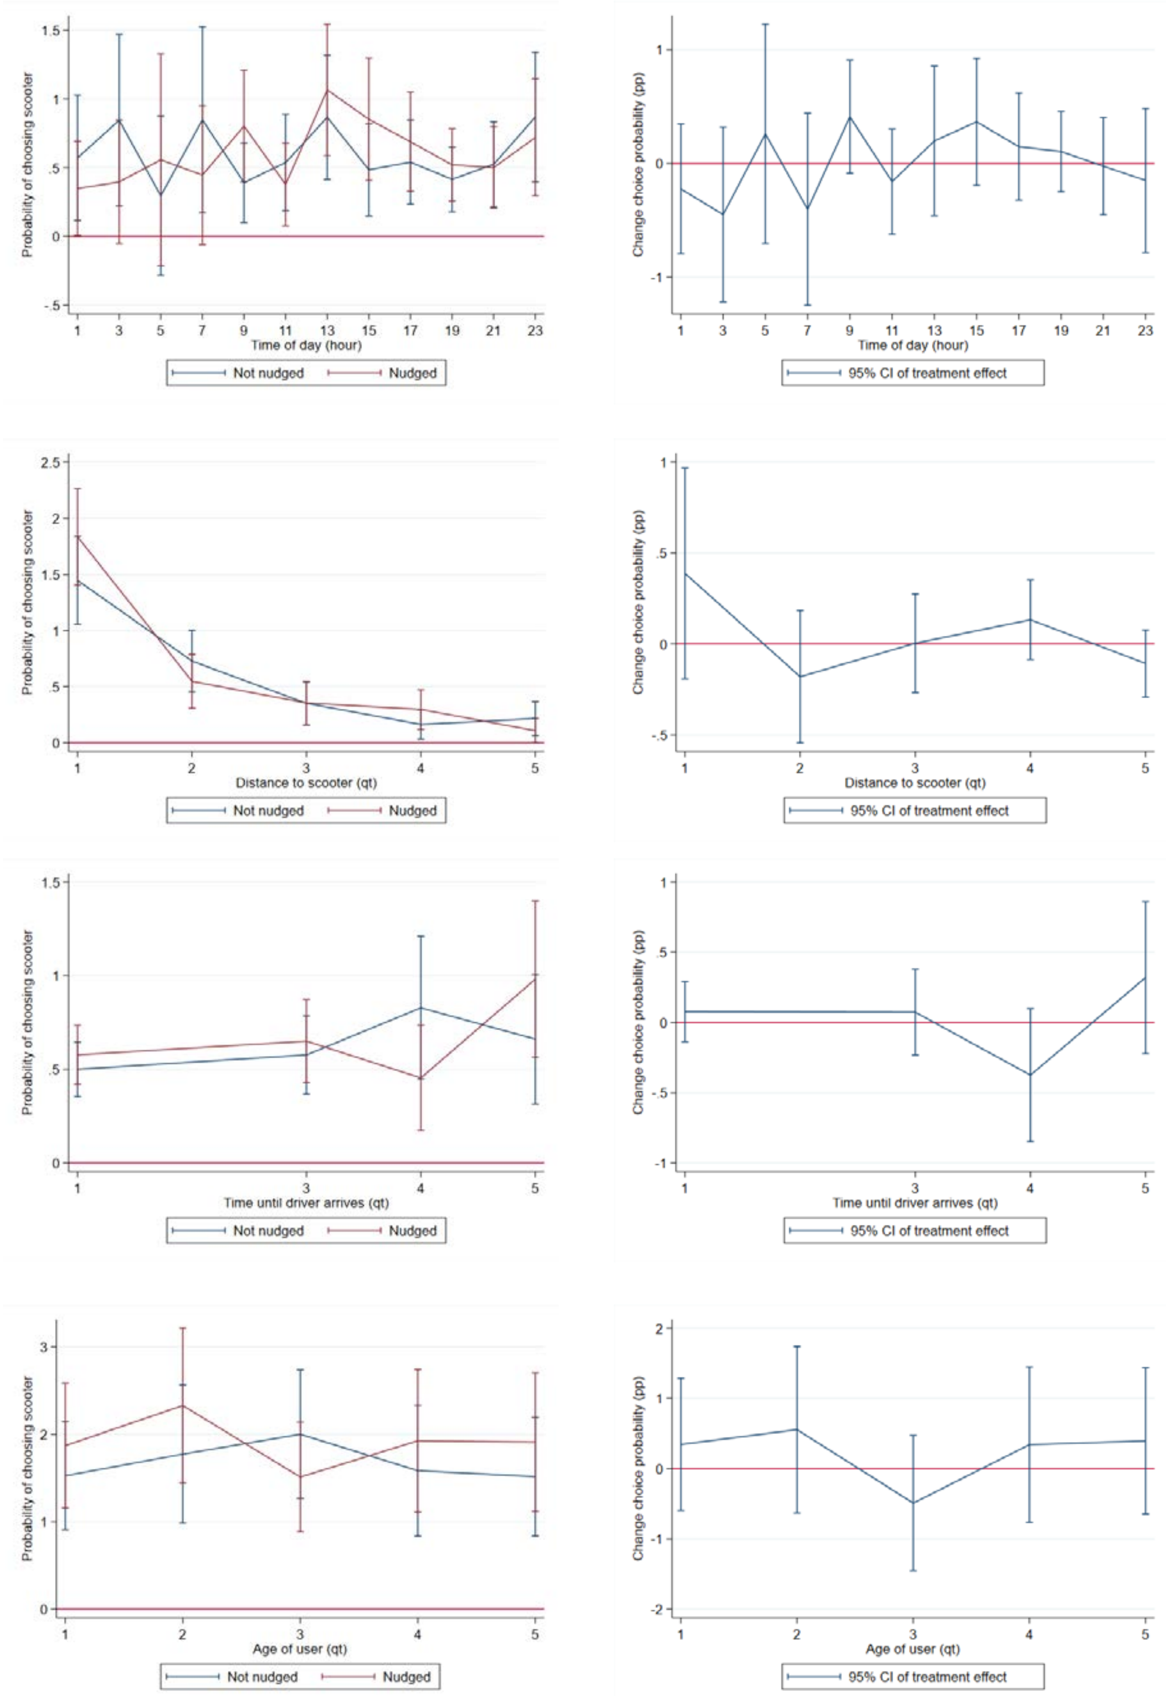

Figure B1.27: Heterogeneous effects of nudging on probability of e-scooter trip. Valletta.

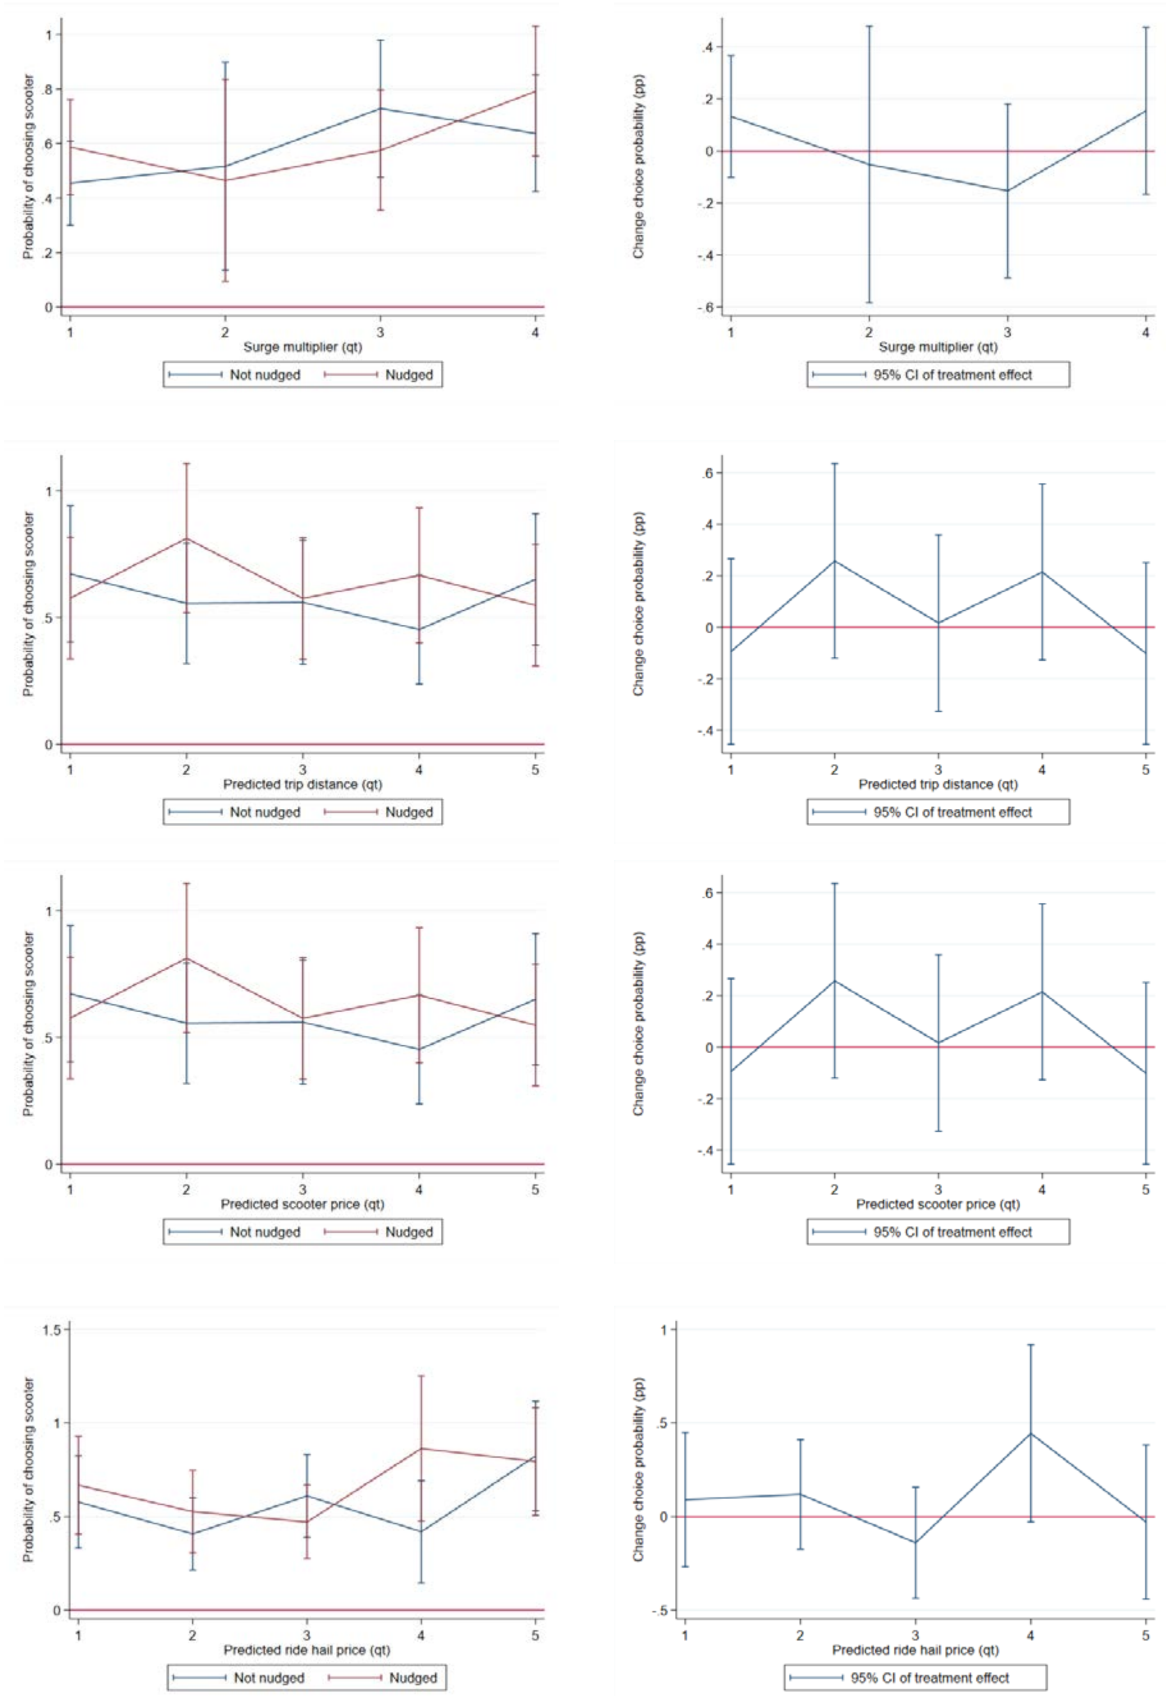

Figure B1.28: Heterogeneous effects of nudging on probability of e-scooter trip. Valletta.

Appendix C: Model specifications

C1: Average treatment effects by linear regressions

The average treatment effect can be estimated by means of ordinary least squares (OLS), by regressing an observed outcome on a treatment indicator and a constant. Let Y_i denote the observed outcome for individual i and $X_i \in \{0,1\}$ denote the binary treatment, such that a value of 0 refers to being in the control group while a value of 1 refers to being in the group that is nudged. Then the regression equation for individual i can be written as

$$Y_i = \alpha + \tau X_i + \varepsilon_i,$$

where α and τ are parameters to be estimated and ε_i is the error term. OLS regression implies finding parameter values that minimize the sum of squared errors (ε_i) from the equation above:

$$(\hat{\alpha}, \hat{\tau}) = \arg \min_{\alpha,\tau} \sum_{i=1}^{N} (Y_i - \alpha - \tau X_i)^2 .$$

It can be shown that the algebraic solution for the estimated treatment effect, $\hat{\tau}$, is

$$\hat{\tau} = \bar{Y}_t - \bar{Y}_c,$$

Where $\bar{Y}_t$ is the average outcome in the treatment group (for individuals where $X_i = 1$), while $\bar{Y}_c$ is the average outcome in the control group (for individuals where $X_i = 0$). See Athey and Imbens (2017) for an account of the detailed calculations, including standard errors.

C2: Discrete choice model specification

In the multinomial logit (MNL) modelling framework, it is assumed that individuals choose the alternative, d, with the highest utility, from a given set of options, $d \in \mathcal{D}$. The discrete choice situation the individual is faced with in this context is the outcome of a search session, which can take the following values: $\mathcal{D} = \{\text{no trip},\ \text{scooter},\ \text{ride hail}\}$.

For a given choice situation i, the alternative specific utility functions are assumed to take the following form: $U_{id} = V_{id} + \varepsilon_{id}$. They are additively composed of a deterministic part consisting of observed variables and parameters to be estimated, V_{id}, and a random error component capturing unobserved factors, ε_{id}. For the model estimated in Section 4.1.3, the following functional form specifications are chosen for the deterministic utility functions V_{id}:

$$\begin{aligned}
V_{i,\text{no trip}} &= 0 \\
V_{i,\text{scooter}} &= \alpha_1 + \tau_1 X_i + \beta_{11}\text{price} + \beta_{12}\text{trip distance} + \beta_{13}\text{distance to scooter} + \gamma_1' Z_i \\
V_{i,\text{ride hail}} &= \alpha_2 + \tau_2 X_i + \beta_{21}\text{price} + \beta_{22}\text{trip distance} + \beta_{23}\text{arrival time driver} + \gamma_2' Z_i
\end{aligned}$$

The deterministic utility of the "no trip" alternative is normalized to zero. X_i is the treatment indicator, equal to one if individual i is nudged. Z_i is a collection of time controls as displayed in Table 4.1. This includes a time trend (number of days since the first day of the experiment) as well as dummy variables for weekends and weekdays interacted with six hour intervals, where the baseline category is "weekdays 12:00-18:00". The remaining variables are price of the trip (in Euros), expected distance of the trip (in kilometers), distance to the nearest e-scooter (in meters) and arrival time of the driver (in minutes). Greek letters are parameters to be estimated.

By assuming that the error terms ε_{id} are drawn independently from an Extreme Value Type I distribution, the choice probabilities (i.e. the probability that U_{id} is higher than the utility of the two other alternatives) will have the following closed form solutions:

$$P_{id} = \frac{e^{V_d}}{\sum_{j \in \mathcal{D}} e^{V_j}},$$

Where P_{id} is individual i's probability of choosing alternative d, and e is the exponential function. This means that for a given set of parameter values and observed variables, the context specific choice probabilities for i can be calculated as:

$$P_{i,\text{no trip}} = \frac{1}{1 + e^{V_{i,\text{scooter}}} + e^{V_{i,\text{ride hail}}}},$$

$$P_{i,\text{scooter}} = \frac{e^{V_{i,\text{scooter}}}}{1 + e^{V_{i,\text{scooter}}} + e^{V_{i,\text{ride hail}}}},$$

$$P_{i,\text{ride hail}} = \frac{e^{V_{i,\text{ride hail}}}}{1 + e^{V_{i,\text{scooter}}} + e^{V_{i,\text{ride hail}}}}.$$

Table 4.2 in the main text displays average effects on choice probabilities of marginally increasing a set of variables. These are calculated as the average (over individuals) change in the choice probabilities presented above by increasing the variable in question by one unit.

Institute of Transport Economics (TØI)
Norwegian Centre for Transport Research

Established in 1964, the Institute of Transport Economics is an interdisciplinary, applied research centre with approximately 90 professionals. Its mission is to develop and disseminate transportation knowledge that has scientific quality and practical application.

A private, non-profit foundation, TØI receives basic funding from the Research Council of Norway. However, the greater part of its revenue is generated through contract research. An important part of its activity is international research cooperation, mostly in the form of projects under the Framework Programmes of the European Commission.

TØI participates in the Oslo Centre for Interdisciplinary Environmental and Social Research (CIENS) located near the University of Oslo. See www.ciens.no

TØI covers all modes of transport and virtually all topics in transportation, including road safety, public transport, climate change and the environment, travel behaviour, tourism, land use and urban planning, decision-making processes, freight and travel demand, as well as general transport economics.

Claiming copyright to its products, TØI acts independently of its clients in matters of scientific approach, professional judgment and evaluation. TØI reports are generally downloadable for free at www.toi.no.

Visiting and postal address:
Institute of Transport Economics
Gaustadalléen 21
NO-0349 Oslo

+ 47 22 57 38 00
toi@toi.no
www.toi.no