Road capacity, congestion and travel behaviour

Tainá Souza Pacheco
Universitat Autònoma de Barcelona

Research question: how do drivers respond to a reduction in road capacity in an urban network?
In 1962, Downs asserted that increasing road capacity reduces congestion. While this principle can be theoretically substantiated, as demonstrated by Anas (2023), empirical testing of this argument faces challenges due to endogeneity issues. Exploiting a natural experiment and a novel dataset, I present short-run evidence of traffic adjustments after a shock in road supply. This paper not only explores congestion but also delves into how drivers respond to a reduction in road capacity.

The natural experiment: an unexpected road blockage in one of the largest cities in the world.
The city of São Paulo (Brazil) has an important ring road, functioning as an urban highway around its central area. The ring road is composed of two segments: an express one with five to seven lanes and a speed of 90 km/h, and a local segment, with two to four lanes and speeds of 60 km/h. In November 2018, a crucial overpass of this road cracked, leading to a four-month road blockage. The overpass blocked all 5 lanes of the express segment, so drivers were left with only two lanes in the affected link.

The data: license plate identified speed cameras records.
I gathered identified speed camera records for the city from 2016 to 2020. With over 700 speed cameras generating more than 10 million daily observations, the dataset provides comprehensive coverage of the city road network.

Methodology: I will employ an event-study and a difference in differences study.
With the identified records, I can see who was using the damaged link before the event, which other links they used and at which time of the day they travelled, and their average speed before, during the blockage and after the re-opening of the overpass. In the event study, I can see what happened to speeds, and how drivers responded to the event: by changing their time of travel, by changing their routes, or by deciding not to travel.

Moreover, I can follow vehicles that were using the damaged link before the event across the city, to understand which is the treated area (speed cameras connected to the damaged link) and the control area (area not connected to the damaged link). I can create a network where each speed camera is a spoke, and the strength of the links connecting them is given by the number of vehicles that pass through the two speed cameras. This allows me to create a diff-in-diff setting, where I investigate what happens to congestion in the affected area versus in the control area.

Expected results: I expect to see a temporary surge in congestion, with a decline in average speeds during the first days following the event. However, this congestion trend should revert to pre-event levels after a short period of time. The question that persists is the mechanism that allows for the reduction in congestion after a certain time. To delve deeper into drivers' behaviour, I will follow vehicles that previously used the damaged road. Vehicles are identified by their license plates. The aim is to investigate how travel behaviour adapted to the shock. Drivers might (i) cease driving, (ii) alter their routes, or (iii) adjust their travel times. Existing literature suggests that the second and third options are more likely to occur. If true, a reduction in road capacity might lead to traffic dispersion and pollution rather than a modal shift towards more sustainable transportation modes.