

A major operational challenge we face here at doordash is having the optimal number of drivers out on the road at any given time, such that

- a) Consumer wait times are low
- b) Driver efficiency (deliveries/hour per driver) is high

Since our drivers are paid on a per-delivery basis, these two outcomes typically pull in opposite directions. An overabundance of drivers may result in shorter consumer wait times, but it does so at the expense of driver efficiency. Conversely, a shortage of drivers usually results in higher driver efficiency, but longer consumer wait times.

Complicating driver capacity planning is the general shape of our consumer demand curve. On a typical evening, consumer demand begins increasing steadily before peaking dramatically in mid-evening. After this brief burst, order volume tapers off significantly in the later parts of the night.

In contrast to the volatility of order demand, driver supply tends to be relatively stable over the course of the evening. As a result, we're typically over-staffed in the early and late evening, and understaffed in the middle. Ideally, driver supply and consumer demand would be perfectly in-sync.

We can attack the imbalance from both the demand side as well as the supply side, but for the purpose of this exercise, we'll focus on the supply side.

The Challenge: Design a scheduling tool and/or workflow that helps to allocate drivers in an optimal fashion given the information above.

