Randomized FIFO Mechanisms — Incentive Alignment Using Money & Time

Francisco Castro[2], Hongyao Ma[1], Hamid Nazerzadeh[3] Chiwei Yan[4]

[1]Columbia Business School [2]UCLA Anderson [3]USC Marshall [4]University of Washington

Matching in Two-sided Platforms

(a) Two-sided platforms

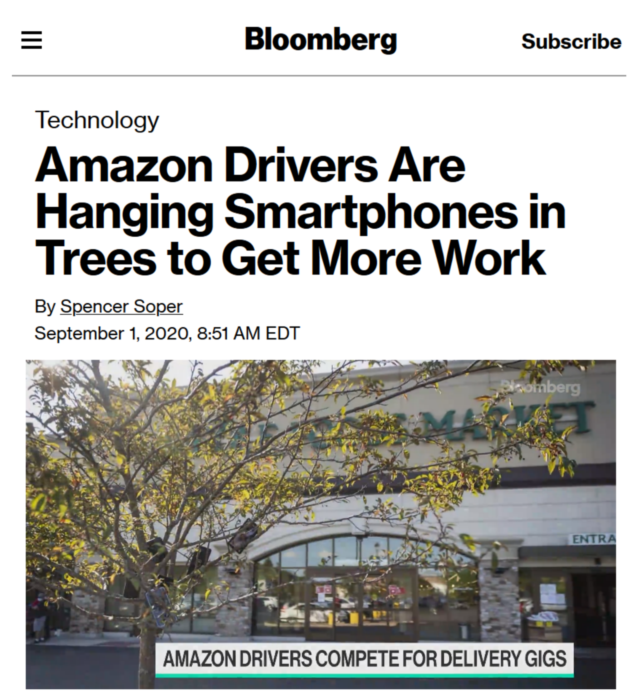

(b) Amazon - assignment to closest drivers

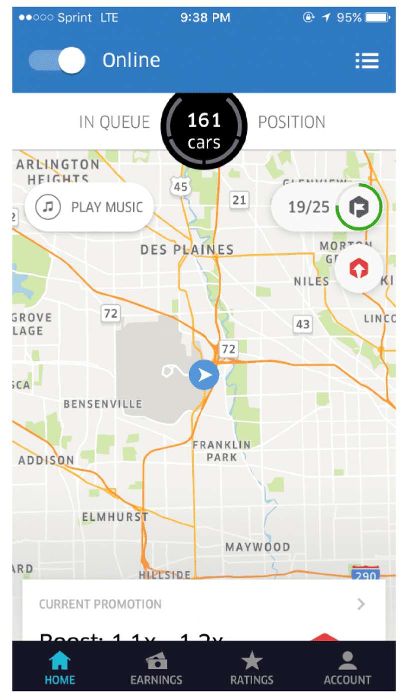

(c) Ridesharing - virtual queues at airports

- Assigning jobs to closest drivers leads to congestion— all drivers try to get closer
- Today's ridesharing platforms (e.g. Uber and Lyft) maintain virtual FIFO (first in first out) queues at airports, for drivers who are waiting in designated areas

Heterogeneous Earnings & Impatient Riders

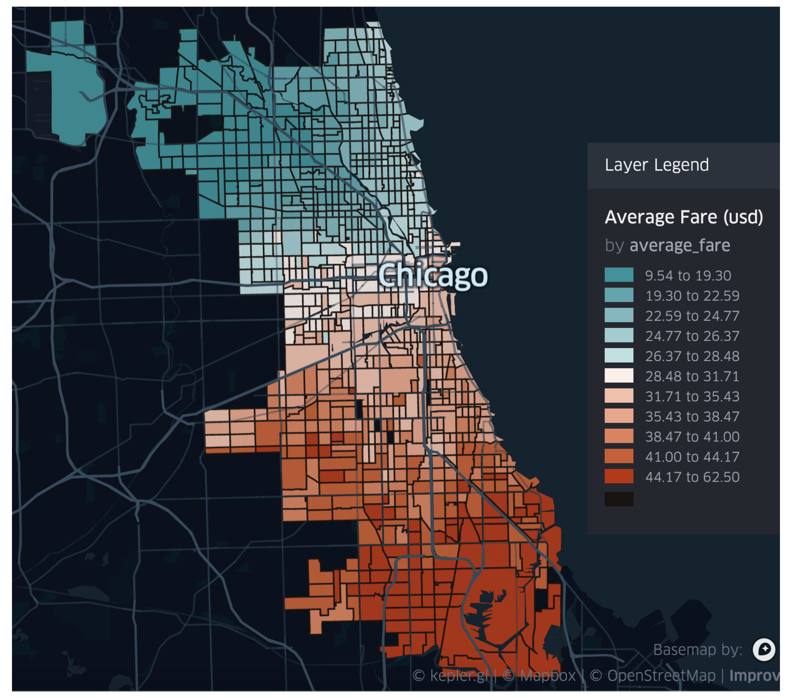

(a) Average fare by destination for trips originating from Chicago O'Hare

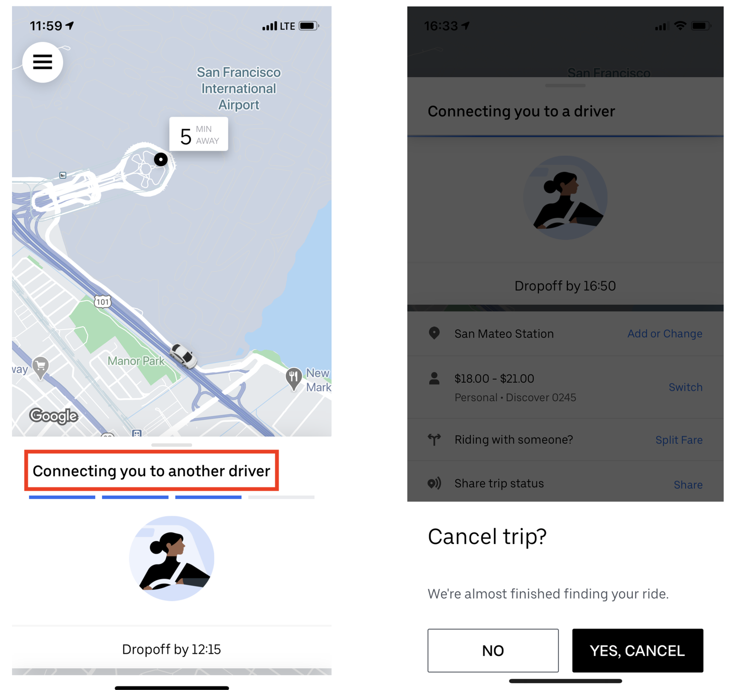

(b) Riders cancel trip requests if getting matched takes too long

Loss of reliability, revenue and trip throughput under FIFO dispatching

- Heterogeneity in earnings by destination: long trips pay substantially more
- Drivers close to the head of the queue are incentivized to cherry-pick based on destinations, leading to repeated declines for lower-earning trips
- Riders have limited patience: repeated declines by drivers → long waiting time for getting matched to a driver → riders canceling trip requests

Pricing alone is not enough for eliminating incentives to cherry-pick

- Difficult to reduce earnings from long trips due to minimum time/distance rates
- Suboptimal to increase fares of short trips to match the earnings from long trips

This work: align incentives using money and *time*, when we do not have the power to tell drivers what to do, or the full flexibility to set prices

Dynamic Dispatching Mechanisms

A simple model

- Continuous time, stationary and non-atomic supply and demand
- Destinations: $\mathcal{L} = \{1, 2, \ldots, L\}$; Arrival rate of riders to destination $i \in \mathcal{L}$: μ_i
- Riders' patience level: P— a rider will cancel trip request after P driver declines
- Arrival rate of drivers: λ; Opportunity cost of driver's time: c
- Net earnings from a trip to location $i \in \mathcal{L}$: w_i. Assume $w_1 > w_2 \cdots > w_L > 0$

Transparency and flexibility

- Drivers know the supply, demand, queue length, their positions in the queue.
- Drivers may freely decline trip dispatches based on trip destination and earnings, or at any time leave the queue, or re-join the queue at the tail.

Goal: optimize platform's net revenue (total driver earnings minus waiting costs) and trip throughput in equilibrium.

Equilibrium Outcome Under Strict FIFO

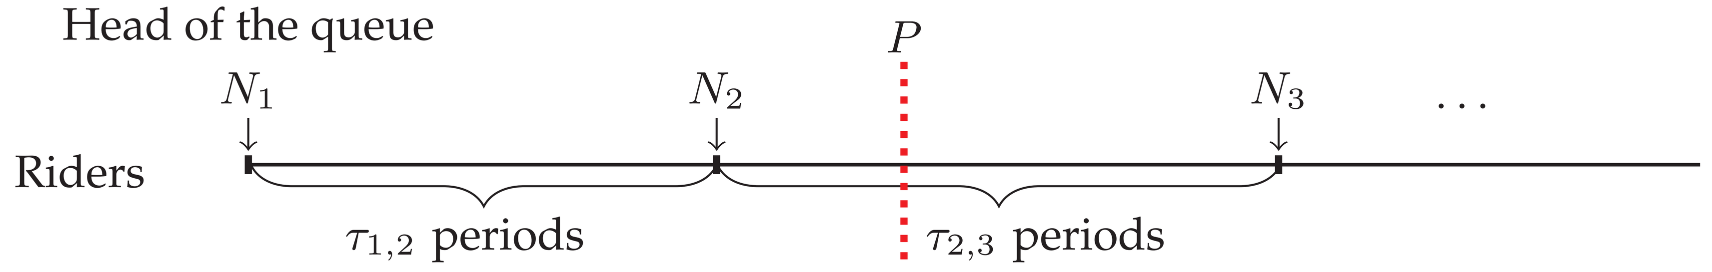

- Driver at the head of the queue: accept only trips to location 1 (i.e. highest earning trips). First position in the queue willing to accept location 1 trips: $N_1 = 0$.
- In comparison to location 2, a driver is willing to wait for an additional $\tau_{1,2}$ periods for a trip to location 1. We know $w_1 - \tau_{1,2}c = w_2 \Rightarrow \tau_{1,2} = (w_1 - w_2)/c$.
- Little's Law ⇒ first position willing to accept location 2 trips $N_2 = \tau_{1,2}\mu_1$. Can similarly find the first position N_i where driver is willing to go to location $i \geq 3$.
- With rider patience level P, a location 3 trip (offered to drivers starting from the head of the queue under strict FIFO) is canceled by the rider after P declines.
- All trips to location i with $N_i > P$ are *unfulfilled*— poor revenue and throughput.

The Direct FIFO Mechanism

Direct FIFO. Dispatch location i trips starting from the N_i^{th} position in the queue.

Theorem.

- It is a subgame-perfect equilibrium (SPE) for drivers to accept all dispatches from direct FIFO. The equilibrium outcome is ex-post envy-free.
- The mechanism achieves in SPE the *second best*, i.e. the highest achievable revenue and trip throughput when drivers are strategic.

Discussion.

- The option to skip the rest of the line incentivizes drivers further from the head of the queue to accept lower earning trips
- Drivers with no experience or information about supply/demand may still optimize earnings by simply following the mechanism's dispatches

The Randomized FIFO Mechanism

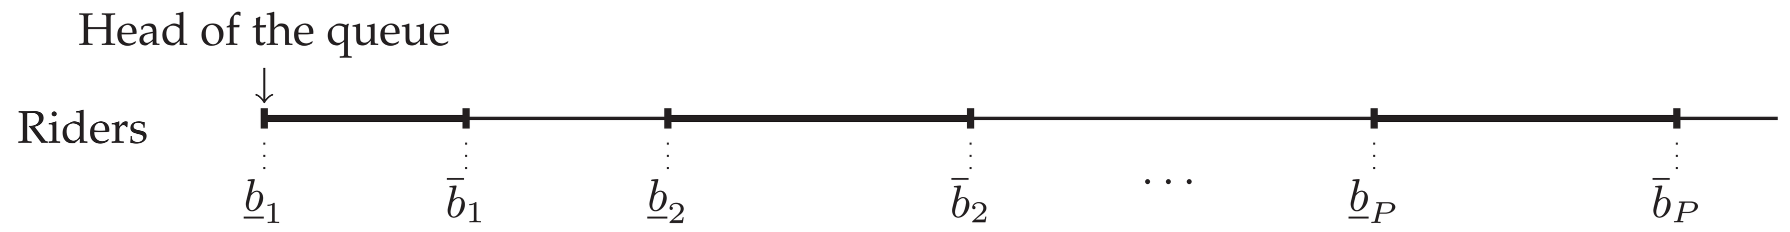

A randomized FIFO mechanism is specified by P "bins". A trip is first dispatched to a driver in the first bin $[\underline{b}_1, \bar{b}_1]$ uniformly at random. If declined for $k^{\text{th}} - 1$ times, then for the k^{th} time a trip request is dispatched, select a random driver from $[\underline{b}_k, \bar{b}_k]$.

Theorem. Randomized FIFO achieves the second best in Nash equilibrium.

Discussion.

- Drivers who have waited the longest in the queue have the highest priority for trips to any destination— fair, and robust to drivers' idiosyncratic preferences
- Randomization increases the waiting times for the next dispatch (vs the driver at the head of the queue under strict FIFO), raising the costs of cherry-picking
- Drivers who have waited longer in the queue (i.e. in earlier bins) will accept higher earning trips → small variance/uncertainty in drivers' net payoffs

Simulation Results

Data from the City of Chicago

- Ridesharing trips originating from Chicago O'Hare, Nov. 2018 - Mar. 2020
- A total of around 800 destinations (census tracts in Chicago)

The *first best*. Drivers are dispatched upon arrival to locations in dec. order of w_i.

Varying arrival rate of drivers λ

Total rider arrival rate: 12 per min; Assuming rider patience P = 12

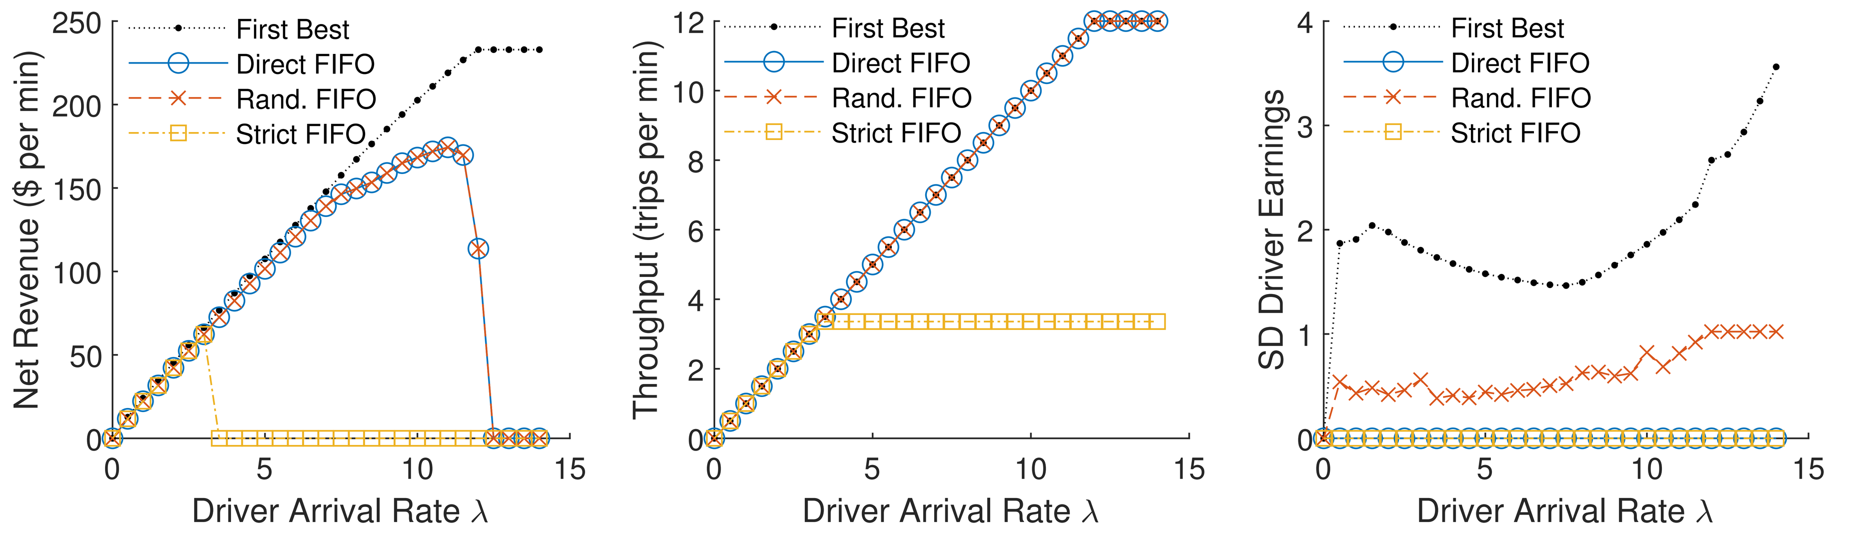

Varying rider patience level P

Fixing rider arrival rate at 12 per min, and driver arrival rate at 10 per min