Experimenting under Stochastic Congestion

Shuangning Li

Harvard University

Joint work with

Ramesh Johari

Stefan Wager

Kuang Xu

Li, Shuangning, Ramesh Johari, Stefan Wager, and Kuang Xu. **Experimenting under Stochastic Congestion.** *arXiv preprint arXiv:2302.12093* (2023).

Toy example: ice cream shop

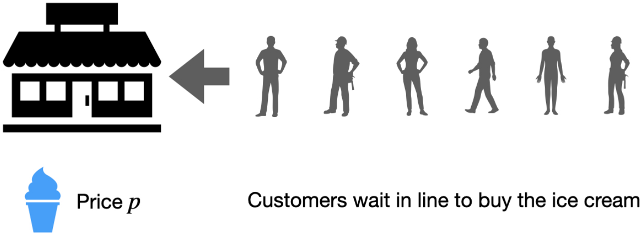

Price p

Customers wait in line to buy the ice cream

Toy example: ice cream shop

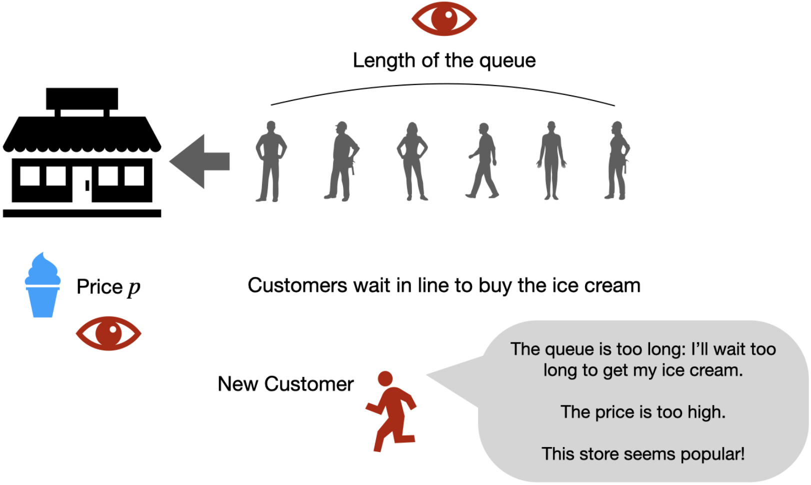

Length of the queue

Price p

Customers wait in line to buy the ice cream

New Customer

The queue is too long: I'll wait too long to get my ice cream.

The price is too high.

This store seems popular!

As the shop owner...

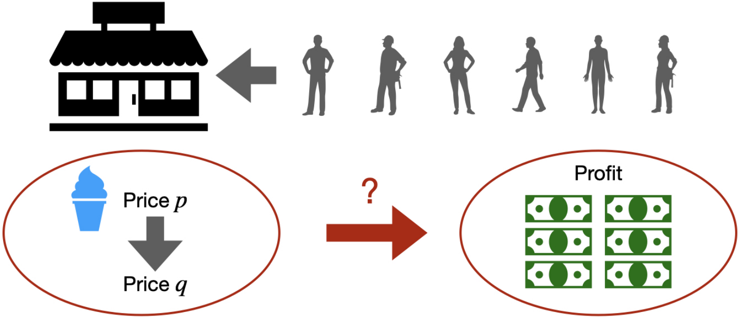

Price p

Price q

?

Profit

As the shop owner...

More customers are attracted to the store.

The line gets longer.

This further impacts customer behaviors.

The key question

How does p impact the (expected) number of arrivals? (1)

The key question

$$\text{How does } p \text{ impact the (expected) number of arrivals?} \tag{1}$$

- We direct our attention to estimating the **derivative** of the (expected) number of arrivals to the system with respect to the system parameter p.
- Why derivative?
  - Small perturbation to the system: protect from risk.
  - Can perform (stochastic) gradient descent to optimize the system.

Two other examples

- **Fast-track in the emergency department.**
  To address the problem of long waits in the emergency department, a hospital implements a fast-track system. Patients with minor illness can go to the "fast-track" area and be treated faster.
  - The parameter p can be one of various quality measures of the fast-track service (e.g., quality of equipment and doctors, efficiency of the service).
  - Question (1) can help the hospital understand how the quality of the fast-track service affects congestion in the emergency department.

Two other examples

- **Flight check-in**. Passengers checking in for their flights have a few options: wait in line to check in in person, check in at the self-service kiosks, or check in online.
  - In order to reduce congestion, the airline sends emails and posts signs encouraging passengers to use online or kiosk check-in. We let the parameter p measure the level of this encouragement.
  - Question (1) studies how much the encouragement actions impact the number of in-person check-ins and alleviate congestion at the airport.

Is this problem trivial?

Can we run an A/B test?

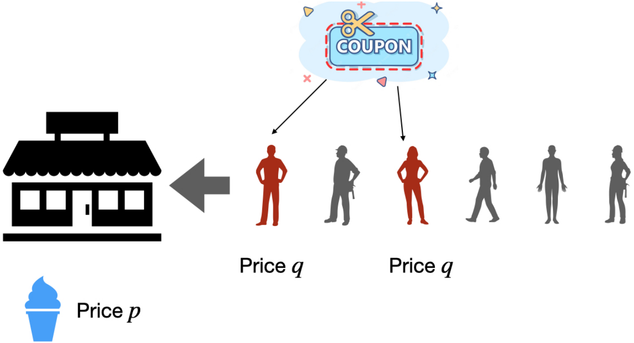

Is this problem trivial?

Can we run an A/B test?

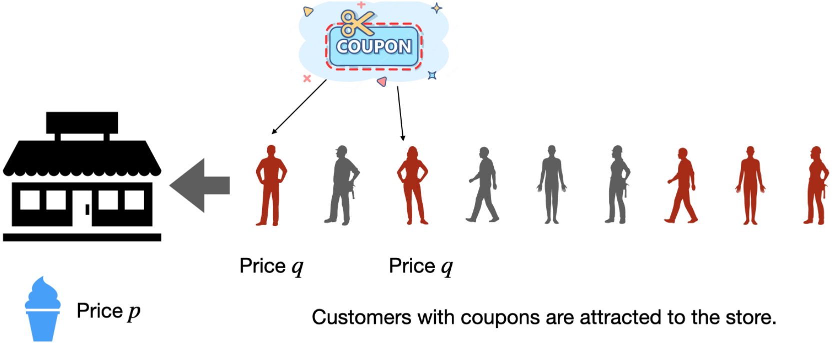

Customers with coupons are attracted to the store.

The line gets longer.

Impacts other customers.

Is this problem trivial?

Can we run an A/B test?

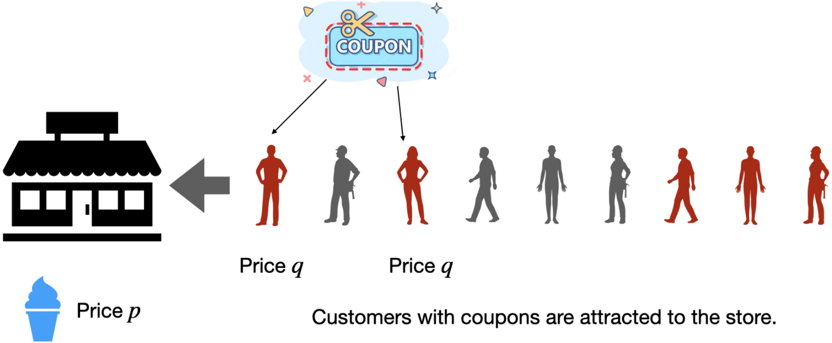

Customers with coupons are attracted to the store.

The line gets longer.

Impacts other customers.

Interference exists!

Setup

- A single-server queueing system that runs in continuous time.
- Service rate μ.
- Arrival rate $\lambda_k(p)$, where k is the current queue length.
- To understand $\lambda_k(p)$, we can think of the system as consisting of a queue with an outside option. We can take $\lambda_k(p) = \lambda_{\text{raw}}\,\text{prob}(k, \mu, p)$

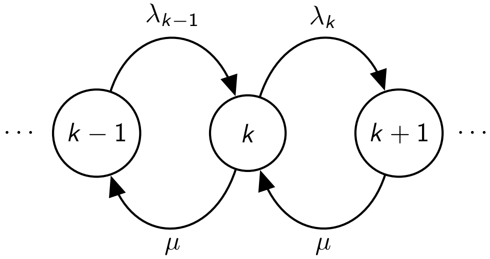

Setup

- The queue length process can be thought of as a continuous-time Markov chain.
- Let $\pi_k(p)$ denote the steady-state probability of being in state k.
- Let $\bar{\lambda}(p)$ be the average arrival rate in steady-state:

$$\bar{\lambda}(p) = \sum_{k=0}^{\infty} \pi_k(p)\lambda_k(p).$$

- We are interested in the derivative $d\bar{\lambda}/dp$.
- We call $d\bar{\lambda}/dp$ the *estimand* of interest.

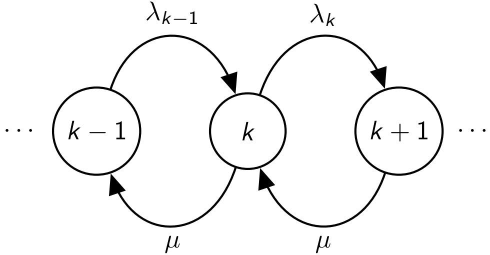

Setup

$$\frac{d}{dp}\bar{\lambda}(p) = \frac{d}{dp}\left(\sum_{k=0}^{\infty}\pi_k(p)\lambda_k(p)\right) = \sum_{k=0}^{\infty}\pi_k'(p)\lambda_k(p) + \sum_{k=0}^{\infty}\pi_k(p)\lambda_k'(p)$$

Indirect effect

Direct effect

- A/B tests can only be used to estimate $\lambda_k'(p)$.
- Can we write $\frac{d}{dp}\bar{\lambda}(p)$ in terms of $\lambda_k'(p)$ only?

Three representations of the estimand

Three representations of the estimand

$$\frac{d\bar{\lambda}(p)}{dp} = -\mu\pi_0'(p) = \mu\pi_0(p)\sum_{k=0}^{K-1}\frac{\lambda_k'(p)}{\lambda_k(p)}\sum_{i=k+1}^{K}\pi_i(p).$$

We need exponential interarrival times.

This holds under more general assumptions.

Finite difference

- The goal of estimation is a derivative with respect to the parameter p.
- Use finite differences.
- Specifically, let $\zeta > 0$ be a small perturbation of the parameter p. We compute finite differences by contrasting the behavior of the queue when the parameter is set to be $p + \zeta$ and $p - \zeta$.

Experimental design

- Interval switchback experiment.
- We divide the time horizon into blocks of equal length l.
- At each time $t = zl$ for some $z \in \mathbb{Z}_+$, the parameter is set to be $p + \zeta$ or $p - \zeta$ with probability $1/2$ independently, and the parameter is kept the same in the time block $[zl, (z+1)l)$ [Bojinov et al., 2022, Hu and Wager, 2022].

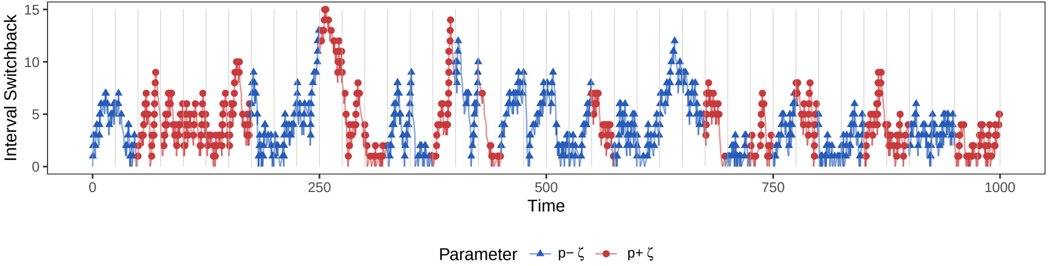

Experimental design

- Regenerative switchback experiment.
- The parameter can only be changed when the queue length is k_r.
- At each visit of state k_r, the parameter is set be $p + \zeta$ or $p - \zeta$ with probability $1/2$ independently [Johari, 2019, Glynn et al., 2022].

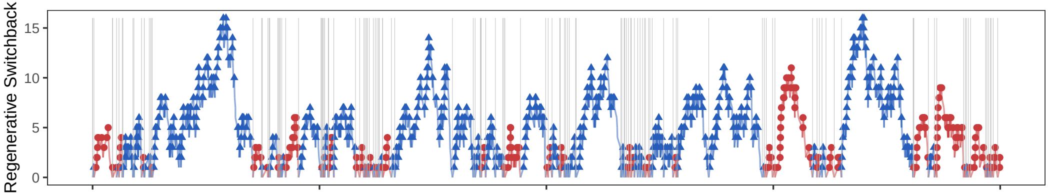

Experimental design

- Local perturbation experiment.
- Instead of changing the parameter for all customers at the same time, we randomize the parameter for each customer. Specifically, let the parameter for customer i be

$$p_i = p + \zeta\varepsilon_i, \qquad \text{where } \varepsilon_i \overset{\text{iid}}{\sim} \{\pm 1\} \text{ uniformly at random.}$$

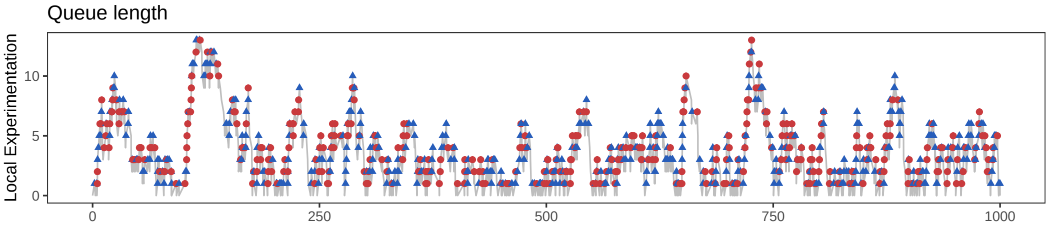

Estimators

Three representations of the estimand

$$\frac{d\bar{\lambda}(p)}{dp} = -\mu\pi_0'(p) = \mu\pi_0(p)\sum_{k=0}^{K-1}\frac{\lambda_k'(p)}{\lambda_k(p)}\sum_{i=k+1}^{K}\pi_i(p).$$

Target representation of estimand	$\frac{d}{dp}\bar{\lambda}(p)$	$-\mu\pi_0'(p)$	$\mu\pi_0(p)\sum_{k=0}^{K-1}\frac{\lambda_k'(p)}{\lambda_k(p)}\sum_{i=k+1}^{K}\pi_i(p)$
Interval switchback	$\hat{\tau}_{\text{IS},\bar{\lambda}}$	$\hat{\tau}_{\text{IS},\pi_0}$	–
Regenerative switchback	$\hat{\tau}_{\text{RS},\bar{\lambda}}$	$\hat{\tau}_{\text{RS},\pi_0}$	–
Local perturbation	–	–	$\hat{\tau}_{\text{LE}}$

Asymptotic Behaviors of the Estimators

Central Limit Theorems

For any method a,

$$\sqrt{T\zeta_T^2}\left(\hat{\tau}_a(T,\zeta_T) - \frac{d}{dp}\bar{\lambda}(p)\right) \Rightarrow \mathcal{N}\left(0, \sigma_a^2(p)\right).$$

- We give explicit formulas for each $\sigma_a^2(p)$.
- We provide variance estimators and get confidence intervals.

Comparison of the estimators

Variance comparison

$$\sigma^2_{\mathrm{LE}}(p) = \frac{1}{2}\sigma^2_{\pi_0}(p),$$

and

$$\sigma^2_{\mathrm{LE}}(p) \leq \sigma^2_{\bar{\lambda}}(p),$$

The equality holds if and only the queue is an $M/M/1$ queue when the parameter is set at p, i.e., $\lambda_k(p) = \lambda_0(p)$ for all $k \geq 0$.

Target representation of estimand	$\frac{d}{dp}\bar{\lambda}(p)$	$-\mu\pi_0'(p)$	$\mu\pi_0(p)\sum_{k=0}^{K-1}\frac{\lambda_k'(p)}{\lambda_k(p)}\sum_{i=k+1}^{K}\pi_i(p)$
Interval switchback	$\hat{\tau}_{\mathrm{IS},\bar{\lambda}}$	$\hat{\tau}_{\mathrm{IS},\pi_0}$	–
Regenerative switchback	$\hat{\tau}_{\mathrm{RS},\bar{\lambda}}$	$\hat{\tau}_{\mathrm{RS},\pi_0}$	–
Local perturbation	–	–	$\hat{\tau}_{\mathrm{LE}}$

This is better!

Why is $\hat{\tau}_{\mathrm{LE}}$ better?

Exploiting the structure of the queueing model →

Target representation of estimand	$\frac{d}{dp}\bar{\lambda}(p)$	$-\mu\pi_0'(p)$	$\mu\pi_0(p)\sum_{k=0}^{K-1}\frac{\lambda_k'(p)}{\lambda_k(p)}\sum_{i=k+1}^{K}\pi_i(p)$
Interval switchback	$\hat{\tau}_{\mathrm{IS},\bar{\lambda}}$	$\hat{\tau}_{\mathrm{IS},\pi_0}$	–
Regenerative switchback	$\hat{\tau}_{\mathrm{RS},\bar{\lambda}}$	$\hat{\tau}_{\mathrm{RS},\pi_0}$	–
Local perturbation	–	–	$\hat{\tau}_{\mathrm{LE}}$

↓ Conditioning on the state of the queue

Thank you!

My Webpage

Paper Link