The Value of Partial Information in Sequential Medical Scanning Decisions

Ali Kaan Kurbanzade[1], Agni Orfanoudaki[2], Mohammad Zhalechian[1], Charlene Ong[3], Stelios Smirnakis[4]

[1]Kelley School of Business, Indiana University Bloomington
[2]Saïd Business School, University of Oxford
[3]Boston Medical School, Department of Neurology
[4]Harvard Medical School, Department of Neurology

February 2, 2026

Subarachnoid Hemorrhage

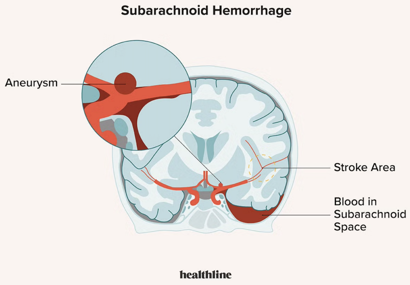

Bleeding (Aneurysm) — Treatment — Clotting (Vasospasm)

nICU

Subarachnoid Hemorrhage

Aneurysm treatment

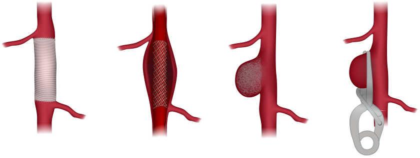

Bleeding (Aneurysm)

Treatment

Clotting (Vasospasm)

nICU

Subarachnoid Hemorrhage

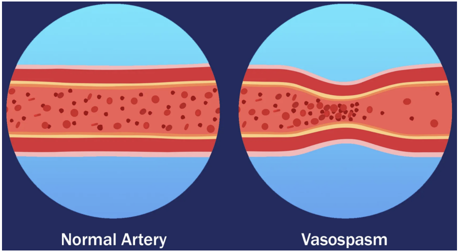

Bleeding (Aneurysm) — Treatment — **Clotting (Vasospasm)**

nICU

Subarachnoid Hemorrhage: Vasospasm Screening

Transcranial Doppler

Accuracy Cost

CT Angiography

Accuracy Cost

Subarachnoid Hemorrhage: Vasospasm Screening

Transcranial Doppler

Accuracy 19–62%

Cost

CT Angiography

Accuracy 93–98%

Cost

Subarachnoid Hemorrhage: Vasospasm Screening

Transcranial Doppler

Accuracy	Cost
19–62%	\$180–680

CT Angiography

Accuracy	Cost
93–98%	\$1,500–10,700

Subarachnoid Hemorrhage: Vasospasm Screening

Transcranial Doppler

Accuracy
19–62%

Cost
\$180–680

CT Angiography

Accuracy
93–98%

Cost
\$1,500–10,700

States: • No, ○ Mild, △ Moderate, ▲ Severe

Example Patient: nICU

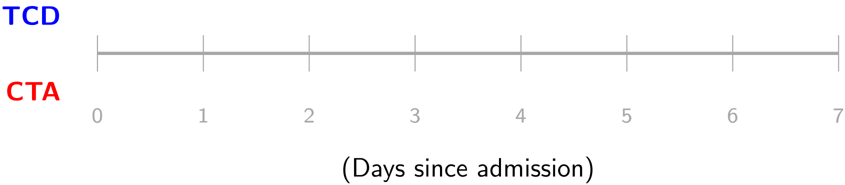

States: • No, ○ Mild, △ Moderate, ▲ Severe

Example Patient: nICU

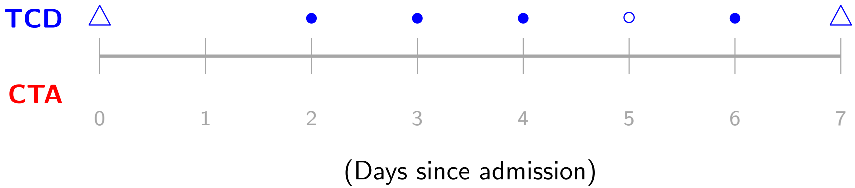

States: • No, ○ Mild, △ Moderate, ▲ Severe

Example Patient: nICU

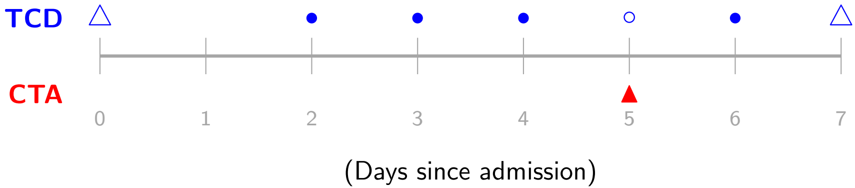

States: • No, ○ Mild, △ Moderate, ▲ Severe

Example Patient: nICU

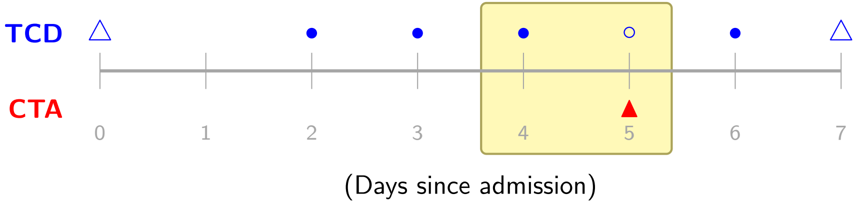

States: • No, ○ Mild, △ Moderate, ▲ Severe

Motivation

- (Wait, Full Scan) $\Rightarrow$ (Wait, Partial Scan, Full Scan).

Motivation

- (Wait, Full Scan) $\Rightarrow$ (Wait, Partial Scan, Full Scan).

- Full-scans are expensive but valuable.

Motivation

- (Wait, Full Scan) $\Rightarrow$ (Wait, Partial Scan, Full Scan).
- Full-scans are expensive but valuable.
- Develop a data-driven policy to decide when to use partial scan.

Motivation

- (Wait, Full Scan) $\Rightarrow$ (Wait, Partial Scan, Full Scan).
- Full-scans are expensive but valuable.
- Develop a data-driven policy to decide when to use partial scan.
- Subarachnoid hemorrhage (Hoh et al., 2023).
  - Requires standardized scanning guidelines.
  - Physicians do not exactly know when to perform a partial scan.

Motivation

- (Wait, Full Scan) $\Rightarrow$ (Wait, Partial Scan, Full Scan).
- Full-scans are expensive but valuable.
- Develop a data-driven policy to decide when to use partial scan.
- Subarachnoid hemorrhage (Hoh et al., 2023).
    - Requires standardized scanning guidelines.
    - Physicians do not exactly know when to perform a partial scan.
- POMDP is very suitable!

Literature Review

- **POMDP Applications in Sequential Medical Scanning:**
  - Chhatwal et al. (2010), Ayer et al. (2012), Zhang et al. (2012), Ayer et al. (2016)
  - Emphasis on binary action sets: Wait or Full Scan.

Literature Review

- **POMDP Applications in Sequential Medical Scanning:**
  - Chhatwal et al. (2010), Ayer et al. (2012), Zhang et al. (2012), Ayer et al. (2016)
  - Emphasis on binary action sets: Wait or Full Scan.
- **Methodological Papers:**
  - Lovejoy (1987), Johnston and Krishnamurthy (2006), Krishnamurthy and Wahlberg (2009), Saghafian (2018)
  - Emphasis on single-threshold approximation algorithms.

Contributions

- First paper to look at SAH modeling using data from a real hospital system ($\sim$ 700 patients).
- Characterizes the optimal policy structure with a partial scan option.
- Study the impact of predictive power of partial-scan option on thresholds.

POMDP Model

- Horizon: Finite.

POMDP Model

- Horizon: Finite.
- State: $\bullet$ No, $\circ$ Mild, $\triangle$ Moderate, $\blacktriangle$ Severe

POMDP Model

- Horizon: Finite.
- State: • No, ○ Mild, △ Moderate, ▲ Severe
- Action: Wait, Partial Scan, Full Scan

POMDP Model

- Horizon: Finite.
- State: $\bullet$ No, $\circ$ Mild, $\triangle$ Moderate, $\blacktriangle$ Severe
- Action: Wait, Partial Scan, Full Scan
- Observation: Discrete probability distribution over states.

POMDP Model

- Horizon: Finite.
- State: $\bullet$ No, $\circ$ Mild, $\triangle$ Moderate, $\blacktriangle$ Severe
- Action: Wait, Partial Scan, Full Scan
- Observation: Discrete probability distribution over states.
- Transition:
  - Full Scan $\Rightarrow$ scan reveals the actual state $\rightarrow$ Belief Update.
  - Wait $\Rightarrow$ carries on the previous signal $\rightarrow$ Belief Update.
  - Partial Scan $\Rightarrow$

POMDP Model

- Horizon: Finite.
- State: • No, ○ Mild, △ Moderate, ▲ Severe
- Action: Wait, Partial Scan, Full Scan
- Observation: Discrete probability distribution over states.
- Transition:
  - Full Scan ⇒ scan reveals the actual state → Belief Update.
  - Wait ⇒ carries on the previous signal → Belief Update.
  - Partial Scan ⇒

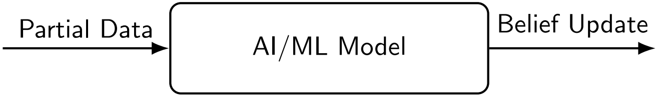

POMDP Model

- Horizon: Finite.
- State: $\bullet$ No, $\circ$ Mild, $\triangle$ Moderate, $\blacktriangle$ Severe
- Action: Wait, Partial Scan, Full Scan
- Observation: Discrete probability distribution over states.
- Transition:
  - Full Scan $\Rightarrow$ scan reveals the actual state $\rightarrow$ Belief Update.
  - Wait $\Rightarrow$ carries on the previous signal $\rightarrow$ Belief Update.
  - Partial Scan $\Rightarrow$

- Belief Update: Bayesian.

POMDP Model

- Horizon: Finite.
- State: • No, ○ Mild, △ Moderate, ▲ Severe
- Action: Wait, Partial Scan, Full Scan
- Observation: Discrete probability distribution over states.
- Transition:
  - Full Scan $\Rightarrow$ scan reveals the actual state $\rightarrow$ Belief Update.
  - Wait $\Rightarrow$ carries on the previous signal $\rightarrow$ Belief Update.
  - Partial Scan $\Rightarrow$

- Belief Update: Bayesian.
- Reward: Depends on action and state.

Theoretical Results

1. How does the partial scan action work under a myopic policy?

Theoretical Results

1. How does the partial scan action work under a myopic policy? $\Rightarrow$ Partial Scan Action is Infeasible in Myopic Policy

Theoretical Results

1. How does the partial scan action work under a myopic policy? $\Rightarrow$ Partial Scan Action is Infeasible in Myopic Policy
2. What happens to threshold values as the predictive power of the partial scan increases?

Theoretical Results

1. How does the partial scan action work under a myopic policy? $\Rightarrow$ Partial Scan Action is Infeasible in Myopic Policy
2. What happens to threshold values as the predictive power of the partial scan increases? $\Rightarrow$ They expand

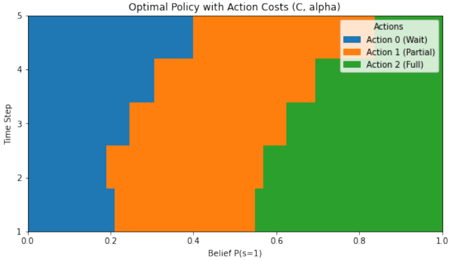

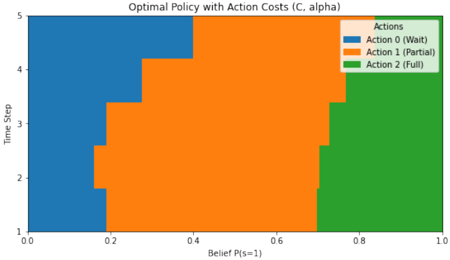

Theoretical Results

3. What is the structure of the non-myopic optimal policy under three actions?

Theoretical Results

3. What is the structure of the non-myopic optimal policy under three actions? $\Rightarrow$ Monotone Two-Threshold Policy

Theoretical Results

3. What is the structure of the non-myopic optimal policy under three actions? $\Rightarrow$ Monotone Two-Threshold Policy

4. How can we approximate the optimal policy?

Theoretical Results

3. What is the structure of the non-myopic optimal policy under three actions? $\Rightarrow$ Monotone Two-Threshold Policy

4. How can we approximate the optimal policy? $\Rightarrow$ Policy Gradient Algorithm

Thank You! Any Questions?

Ali Kaan Kurbanzade
PhD Candidate
Kelley School of Business
Indiana University Bloomington
akurban@iu.edu

References

Ayer, T., O. Alagoz, and N. K. Stout (2012). Or forum—a pomdp approach to personalize mammography screening decisions. *Operations Research 60*(5), 1019–1034.

Ayer, T., O. Alagoz, N. K. Stout, and E. S. Burnside (2016). Heterogeneity in women's adherence and its role in optimal breast cancer screening policies. *Management Science 62*(5), 1339–1362.

Chhatwal, J., O. Alagoz, and E. S. Burnside (2010). Optimal breast biopsy decision-making based on mammographic features and demographic factors. *Operations research 58*(6), 1577–1591.

Hoh, B. L., N. U. Ko, S. Amin-Hanjani, S. H.-Y. Chou, S. Cruz-Flores, N. S. Dangayach, C. P. Derdeyn, R. Du, D. Hänggi, S. W. Hetts, et al. (2023). 2023 guideline for the management of patients with aneurysmal subarachnoid hemorrhage: a guideline from the american heart association/american stroke association. *Stroke 54*(7), e314–e370.

Johnston, L. A. and V. Krishnamurthy (2006). Opportunistic file transfer over a fading channel: A pomdp search theory formulation with optimal threshold policies. *IEEE Transactions on Wireless Communications 5*(2), 394–405.

The Effect of Delivery Unit Congestion on Emergency Caesarean Sections

Presenter: Penelope Stamou, University of Oxford

Co-authors: Agni Orfanoudaki (University of Oxford), Yueyang Zhong (London Business School), Harshita Kajaria-Montag (Indiana University)

Advances in Quantitative Medical Care: The Institute for Mathematical and Statistical Innovation Workshop, February 2-6, 2026

Motivation and Research Question

- Maternity units frequently operate under congestion
- Congestion may alter clinical decision-making, not just outcomes
- Emergency c-sections are risky and costly interventions
- Research question: *Does delivery unit congestion increase emergency c-sections?*

Clinical Setting, Data, and Variables

- **Setting and data:**
  - NHS maternity services
  - Data: NHS Maternity Service Data Set (MSDS)
  - 2021-2025 population-level coverage
- **Key variables:**
  - **Dependent variable:** emergency c-sections
  - **Independent variable:** delivery unit congestion
  - **Control variables:** maternal characteristics, clinical history / comorbidities, obstetric and infant factors, previous obstetric history, temporal factors, hospital characteristics

Model-Free Evidence

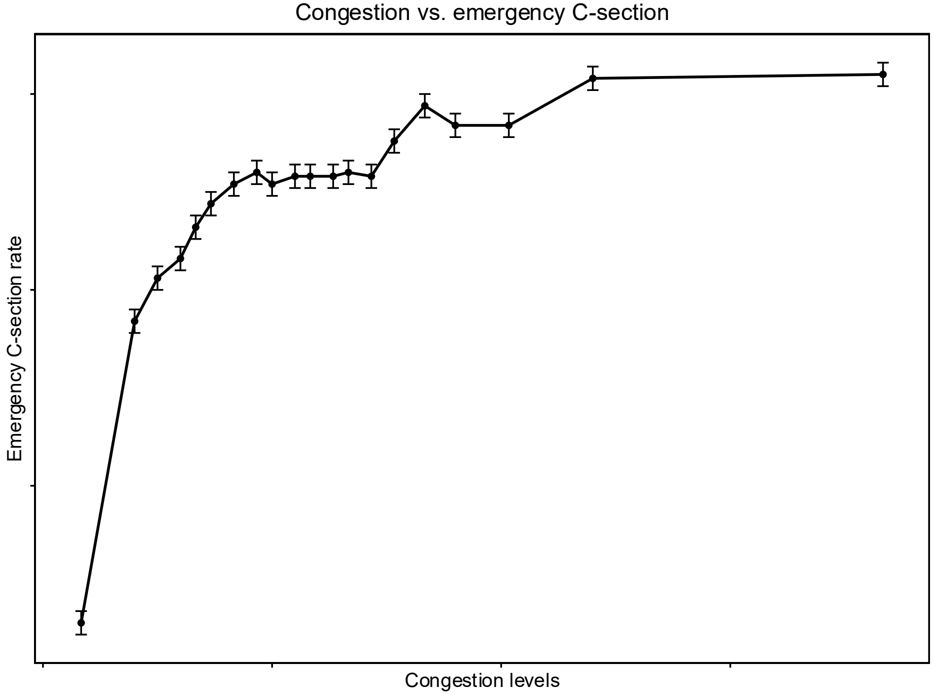

Figure 1: Effect of congestion during one week before labour on emergency c-sections

Method

Hypotheses:

- **Hypothesis 1 (Congestion effect).** Higher pre-labour congestion is associated with a higher probability of emergency caesarean section interventions.
- **Hypothesis 2 (Saturation point).** There will be a saturation point at moderately high congestion levels where the effect of additional congestion on emergency caesarean section interventions diminishes.

Main model estimation: Baseline estimate of the effect of Delivery Unit Congestion (DUC) on the occurrence of emergency c-sections

$$Y_i = \beta_0 + \beta_1 \frac{1}{7} \sum_{k=1}^{7} DUC_{i,t_i-k} + \boldsymbol{X_i \beta} + \tau_{t_i} + \mu_h + \varepsilon_i$$

Where:

- $\boldsymbol{X_i}$: the $1 \times$ n vector of control variables for birth i
- τ_{t_i}: time fixed effects
- μ_h: hospital fixed effects

Main Results

	Emergency c-section
Quintile 0	0.967**
Quintile 1	0.992
Quintile 2 (Reference)	NA
Quintile 3	1.020**
Quintile 4	1.014**

Notes: *$p<.05$, **$p<.01$, ***$p<.001$.

Results: Interpretation and Discussion

- H_1 is supported in our empirical setting, suggesting that under periods of high congestion, clinicians lower their threshold for performing emergency c-sections
- H_2 is supported in our empirical setting, suggesting that there is a point of saturation above which clinicians cannot further increase caesarean deliveries

Robustness Checks

- Robustness checks already carried out:
  - Average congestion up to two weeks before labour ($t_i - 1$ to $t_i - 15$)
  - Average congestion up to three weeks before labour ($t_i - 1$ to $t_i - 22$)
  - Average congestion up to four weeks before labour ($t_i - 1$ to $t_i - 29$)
  - Congestion on the admission day
  - Congestion on the discharge day
  - Average congestion during the length of stay

- The shape and directionality stay the same across checks
- As expected, the effect size diminishes when moving from one week prior delivery to larger time windows

Further Findings

	Emergency c-section	Low APGAR	Stillbirth	Readmission	Neonatal critical incident	NICU	L&D Log LOS (Linear)
Quintile 0	0.967**	0.984	0.985	1.004	0.999	1.024	-0.052**
Quintile 1	0.992	0.998	0.984	1.000	0.953	1.005	-0.012**
Quintile 2 (Reference)	NA	NA	NA	NA	NA	NA	NA
Quintile 3	1.020**	1.017	0.999	1.002	1.019	1.029	-0.003
Quintile 4	1.014**	1.016	1.029	0.991	1.076*	0.963	0.015**

Notes: ***p<.05, **p<.01, ***p<.001.**

We find no evidence that congestion in the delivery unit affects pregnancy outcomes, such as low APGAR score, stillbirth, readmission, neonatal critical incidents, and NICU admission

Future Work

- Strengthen causal identification
- Explore mechanisms (e.g., staffing and operational pressure)

Thank you!

Email: Pinelopi.Stamou@gtc.ox.ac.uk

Assumptions and Limitations

Assumptions:

- Although we control for a rich set of variables, some factors that may jointly influence delivery unit congestion and outcomes are not directly observed in the data

Data limitations:

- Not enough granularity in the MSDS to predict the outcome in case of no emergency c-section intervention
- No direct staffing data and other proxies that could shed light into the underlying mechanisms that drive the results

References

[1] Snowden, J. M., Kozhimannil, K. B., Muoto, I., Caughey, A. B., & McConnell, K. J. (2017). A “busy day” effect on perinatal complications of delivery on weekends: A retrospective cohort study. BMJ Quality & Safety, 26(1), e1.

[2] Benses, Simon (2022), Time to Spare and Too Much Care? Congestion, Treatment Provision and Health Outcomes in the Maternity Unit, Unpublished Manuscript, Statistics Norway, Oslo, Norway.

[3] Freeman, M., Robinson, S., & Scholtes, S. (2020). Gatekeeping under congestion: An empirical study of referral errors in the emergency department (Working Paper No. 2020/09/TOM). INSEAD. Retrieved from https://ssrn.com/abstract=3036999

[4] Miedaner, F., Kuntz, L., Eilermann, K., Roth, B., & Scholtes, S. (2024). Service quality implications of long periods of consecutive working days: An empirical study of neonatal intensive care nursing teams. Manufacturing Service Operations Management, 26(4), 1422–1434. Retrieved from https://doi.org/10.1287/msom.2022.0021 doi: 10.1287/msom.2022.0021

[5] Berry Jaeker, J. A., & Tucker, A. L. (2017). Past the point of speeding up: The negative effects of workload saturation on efficiency and patient severity. Management Science, 63(4), 1042-1062. Retrieved from https://doi.org/10.1287/mnsc.2015.2387 doi: 10.1287/mnsc.2015.2387

Measuring Heterogeneous Effect of Emergency Department Boarding: Towards Efficient and Equitable Inpatient Bed Assignment

Xiaole (Alyssa) Liu
NYU Stern

Jing Dong
Columbia GSB

Yosef Berlyand
Brown Medicine

Martin Copenhaver
Johns Hopkins Medicine

2026.02.02

NEJM Catalyst

Emergency Department Crowding: The Canary in the Health Care System

The solution for this serious threat to ED staff and harm to patients cannot come from a single department, but through engagement of and ongoing commitment by leaders throughout the hospital and, more broadly, by those in the payer and regulatory segments of the health care system as well.

THE WALL STREET JOURNAL.

Emergency-Room Crowding Leads to Higher Mortality, Study Finds

A 10% reduction in an ER's patient volume improves outcomes significantly, the research finds.

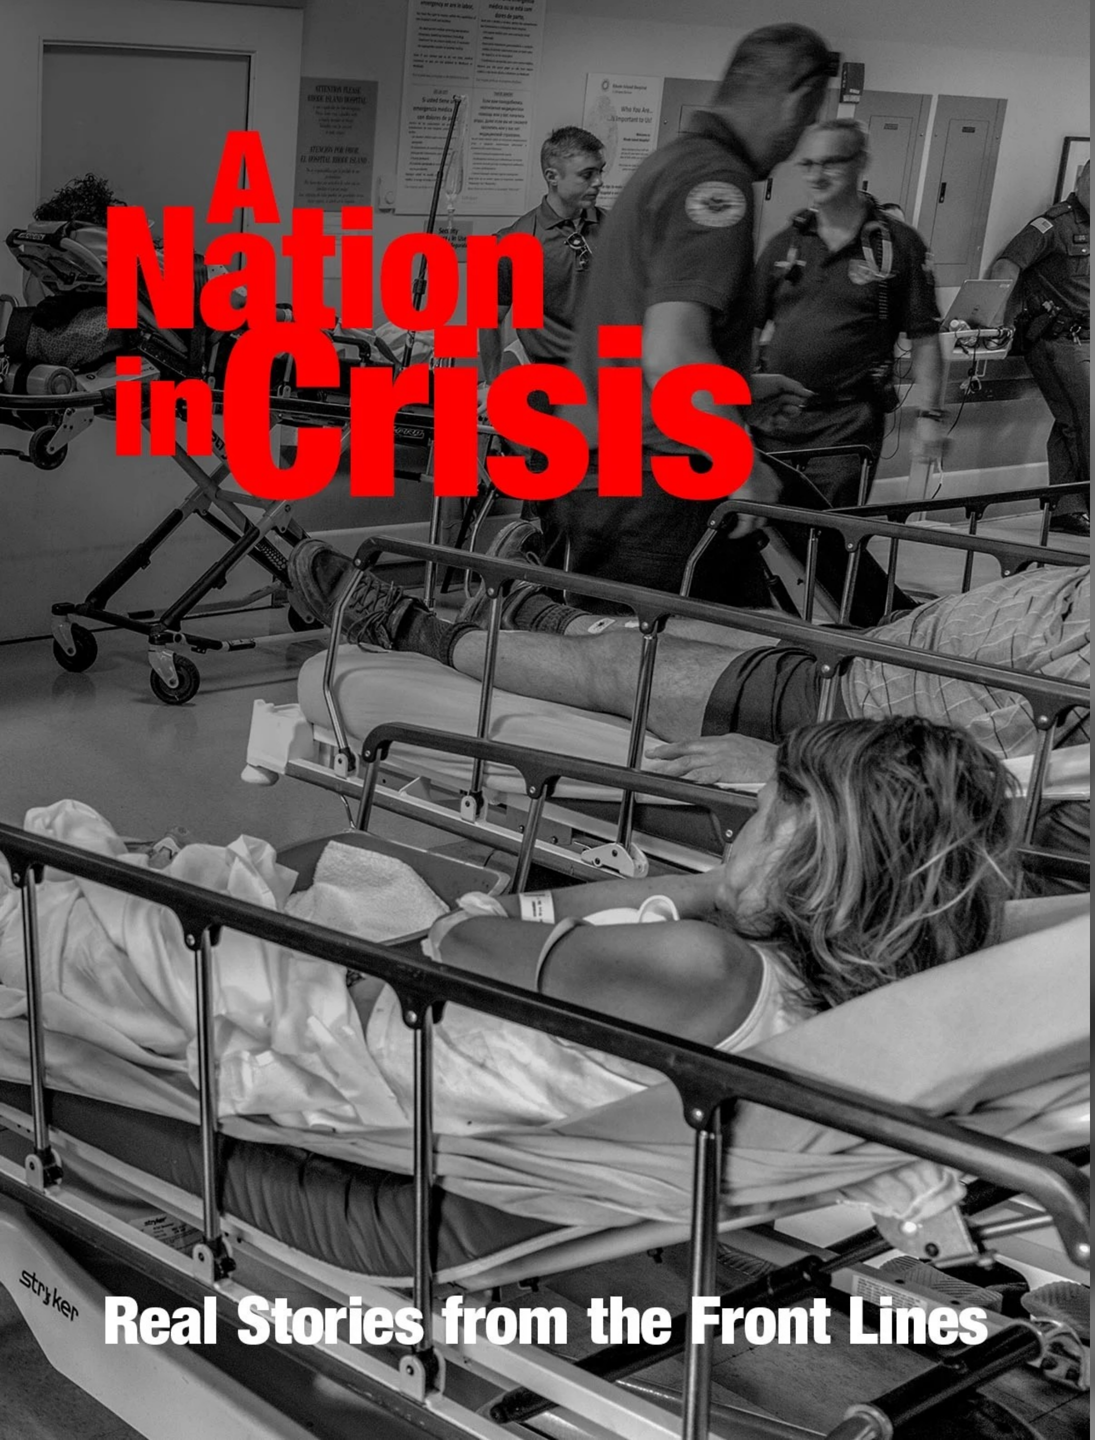

American College of Emergency Physicians®

ACEP Sounds the Alarm on Emergency Department (ED) Boarding as a Public Health Emergency

Boarding (in the ED) refers to holding admitted patients in the ED while awaiting an inpatient bed.

Emergency Department Boarding Crisis

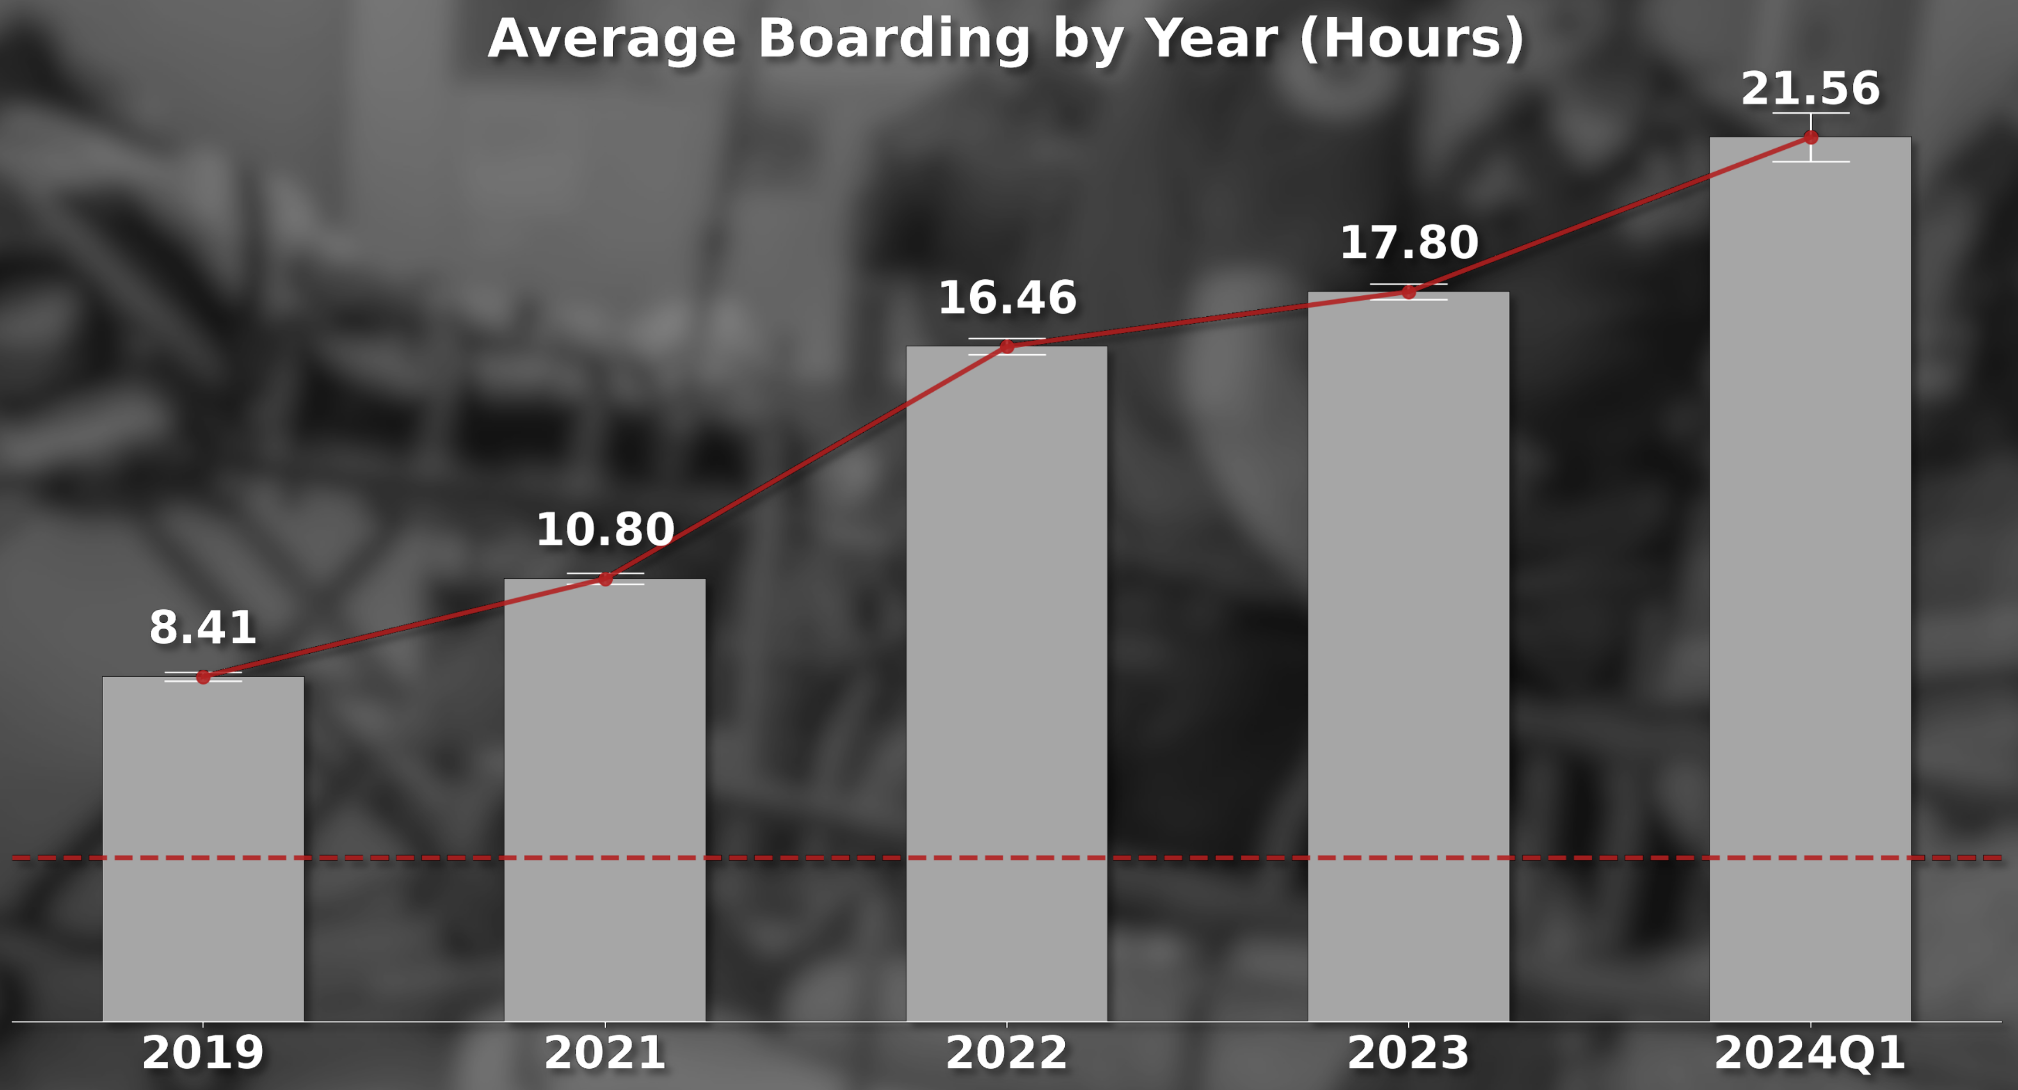

Emergency Department Boarding Crisis

OVERCROWDING EMERGENCY DEPARTMENT

- RISING BURNOUT IN CARE TEAMS (Lokes et al 2023)
- RISING DEATHS AND WORSENING OUTCOMES (Joseph et al 2024)
- INCREASED AVOIDABLE SYSTEM COSTS (Canellas et al 2024)
- PATIENT AND FAMILY STRESS (Berlyand et al 2022)
- RISING WORKPLACE VIOLENCE

Emergency Department Boarding Crisis

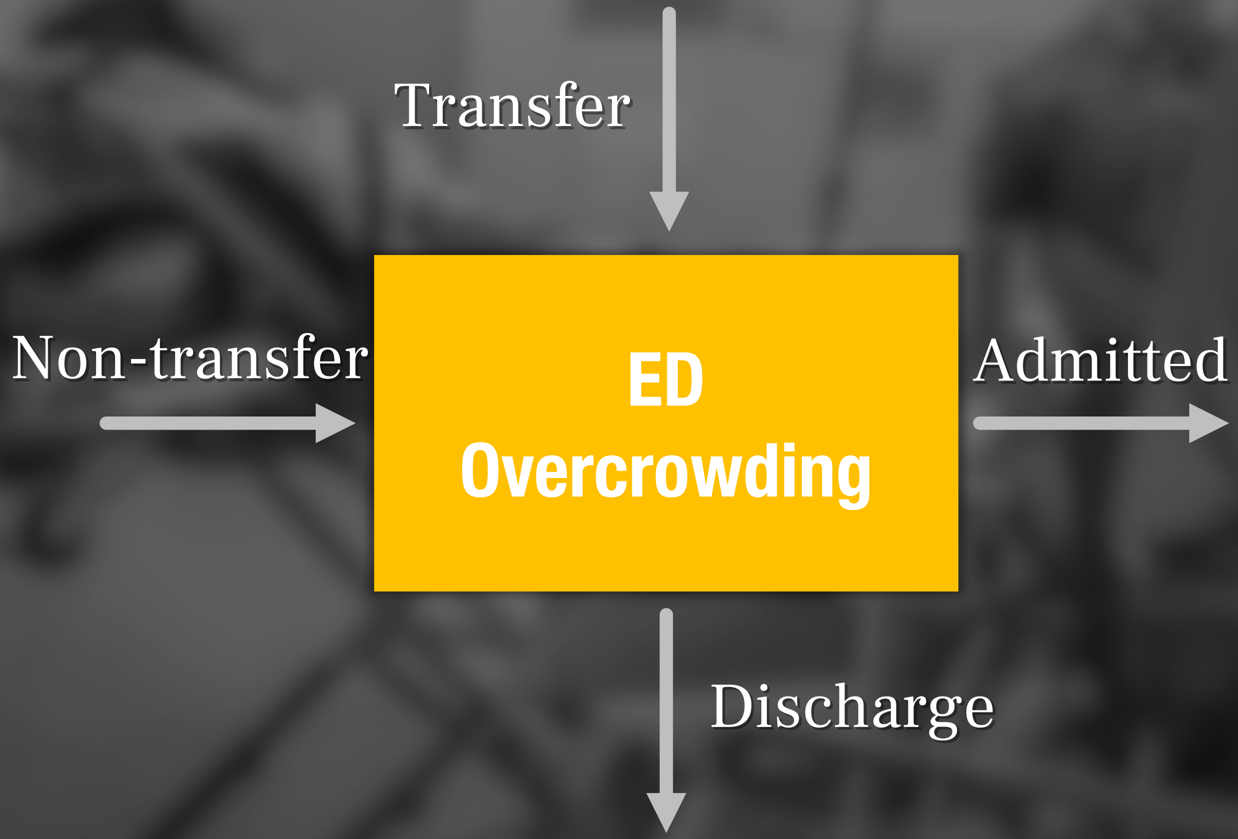

Emergency Department Boarding Crisis

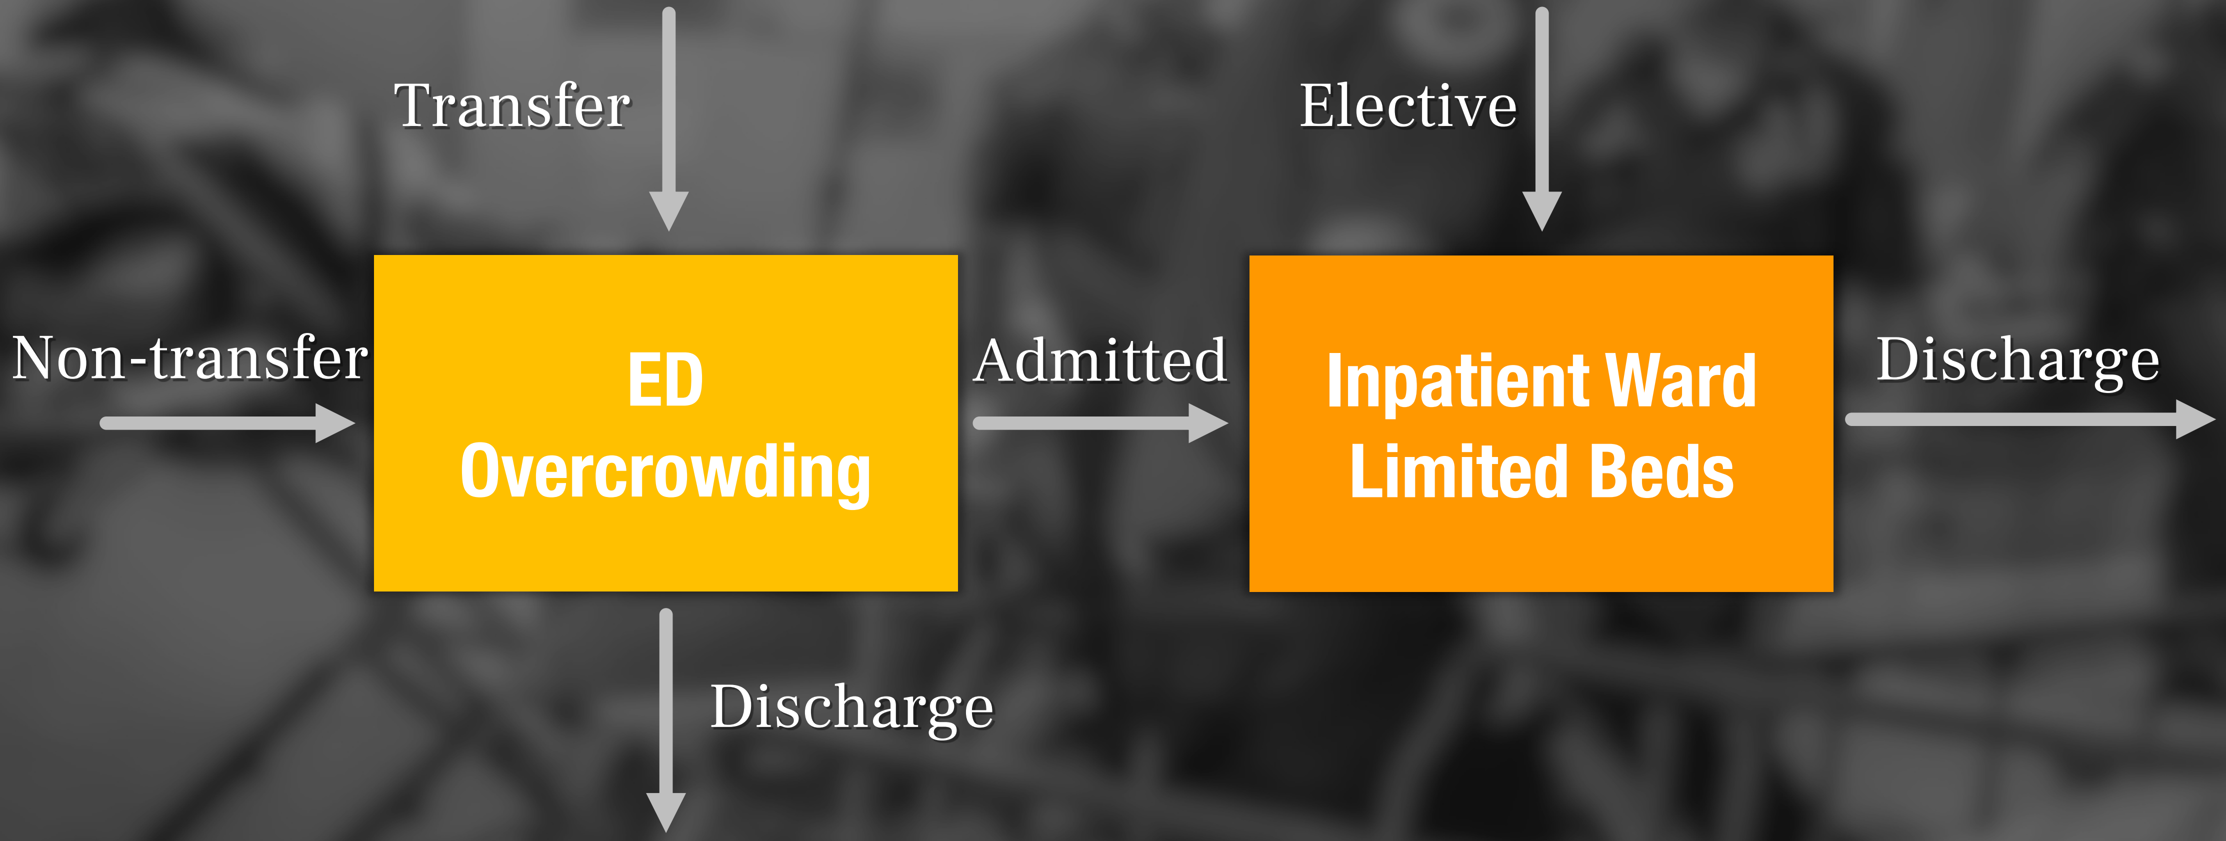

Research Goal

Design bed assignment policies under constrained inpatient capacity

Research Goal

Design bed assignment policies under constrained inpatient capacity

Clinical Collaboration

Data
ED Boarding
Inpatient LOS
Pragmatic Opinion

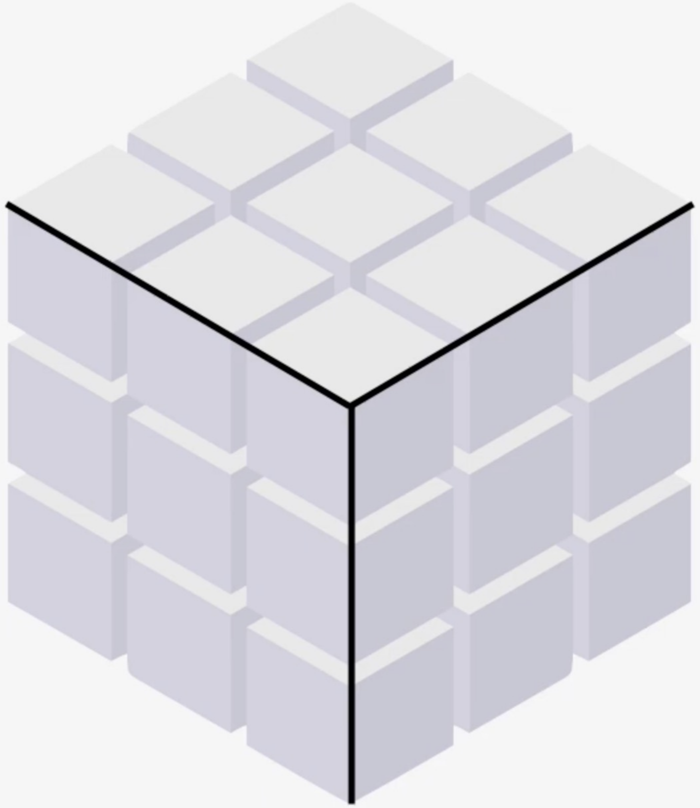

Research Goal

Design bed assignment policies under constrained inpatient capacity

Clinical Collaboration

Data
ED Boarding
Inpatient LOS
Pragmatic Opinion

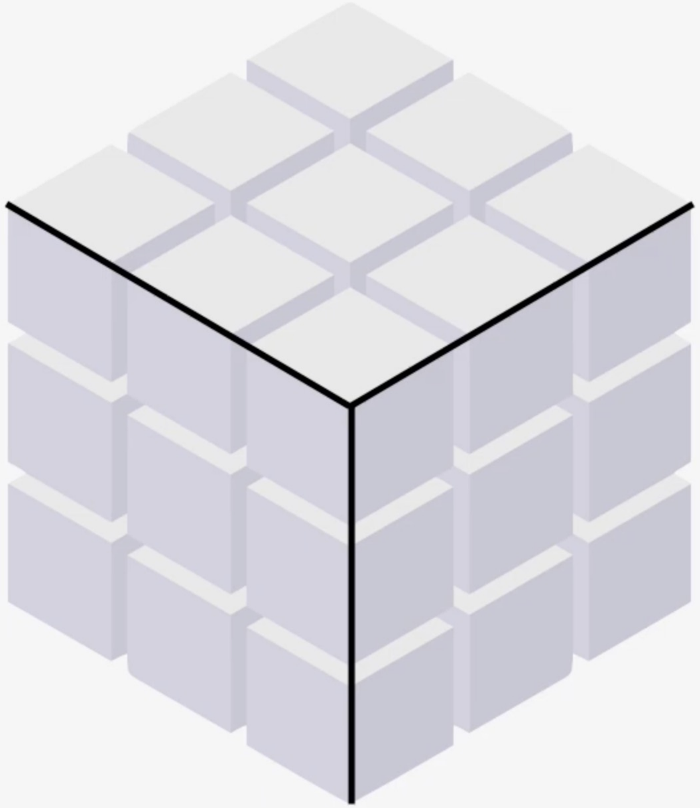

Estimation

Heterogeneous Effect
Econometric Method
Interpretation
Robustness Check

Research Goal

Effect of Boarding on Inpatient Length of Stay (LOS)

Research Goal

Effect of Boarding on Inpatient Length of Stay (LOS)

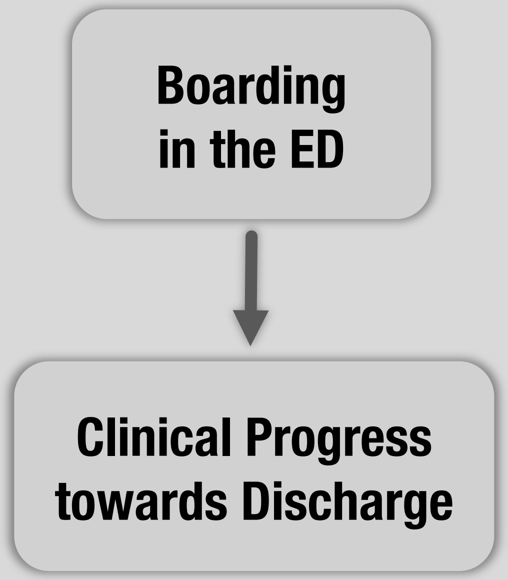

Research Goal

Design bed assignment policies under constrained inpatient capacity

Research Goal

Effect of Boarding on Inpatient Length of Stay (LOS)

Research Goal

Effect of Boarding on Inpatient Length of Stay (LOS)

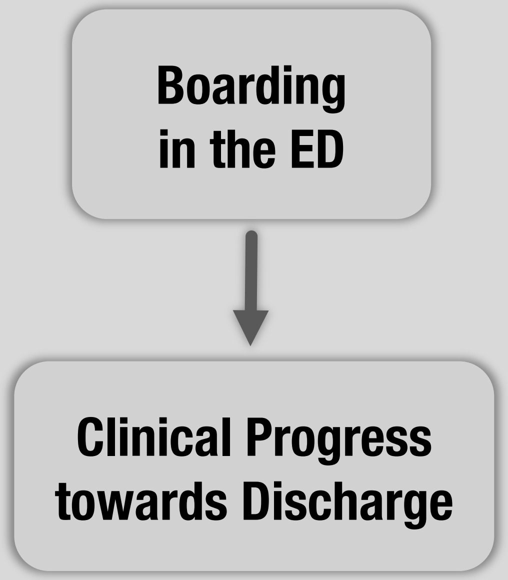

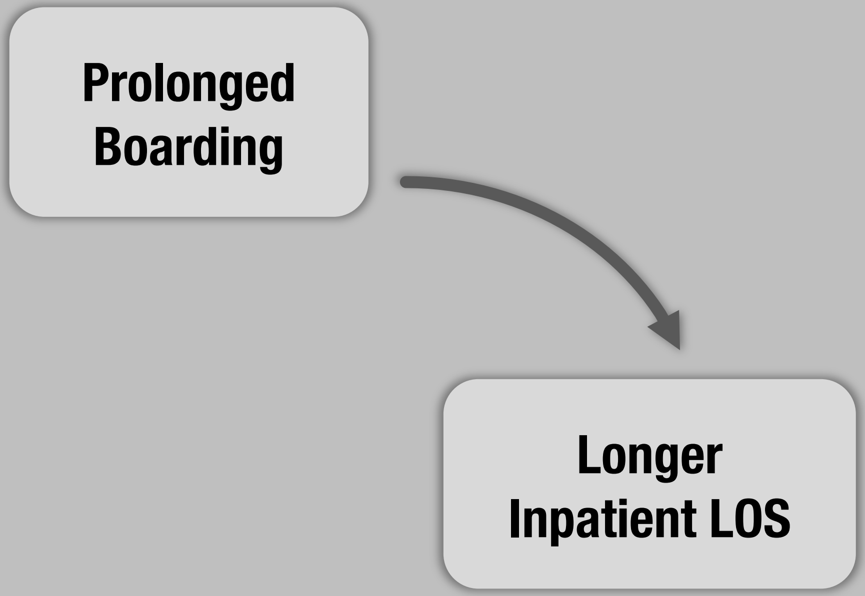

Research Goal

Effect of Boarding on Inpatient Length of Stay (LOS)

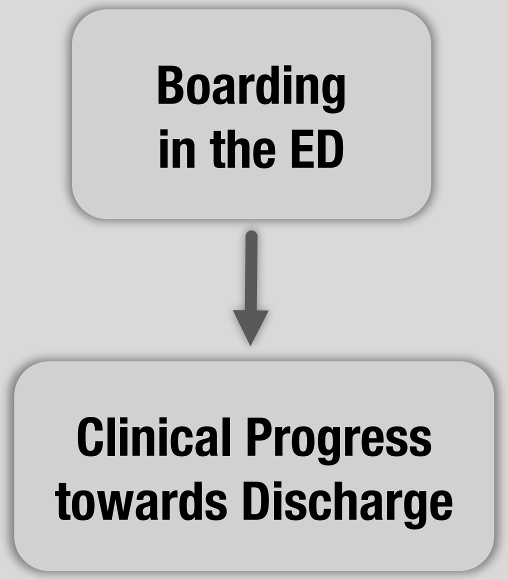

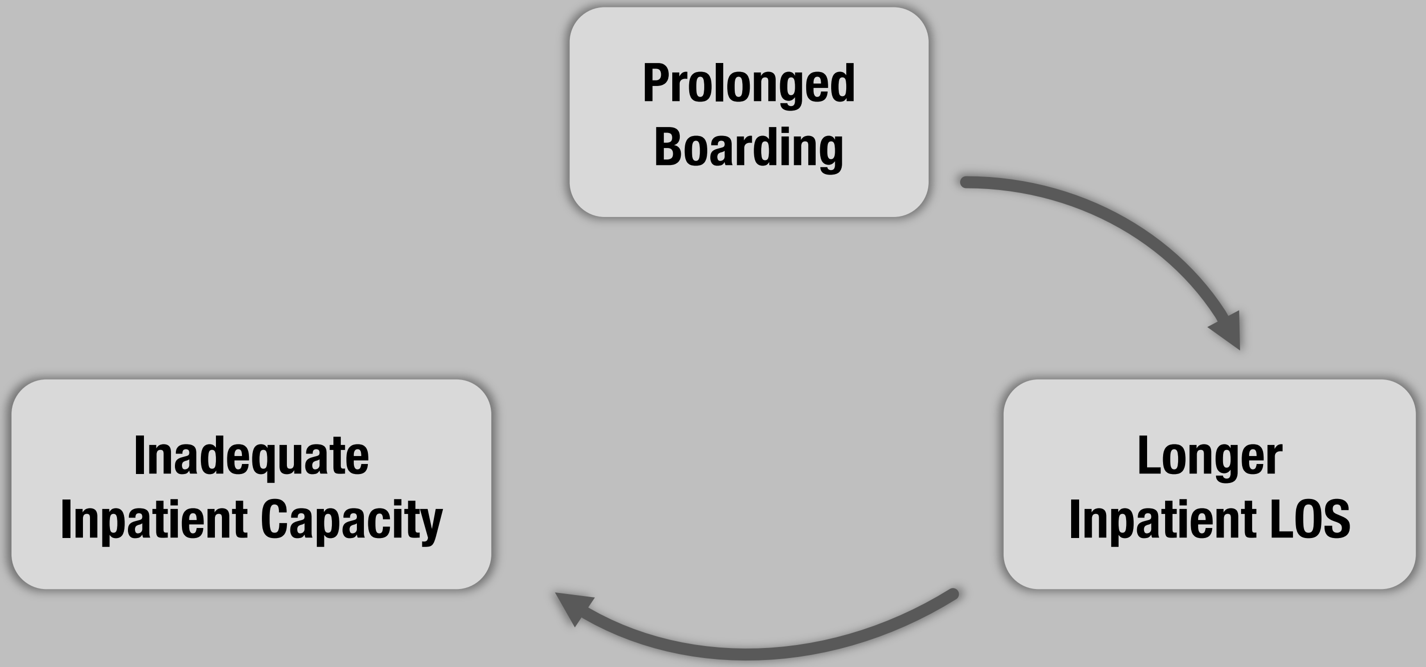

Research Goal

Effect of Boarding on Inpatient Length of Stay (LOS)

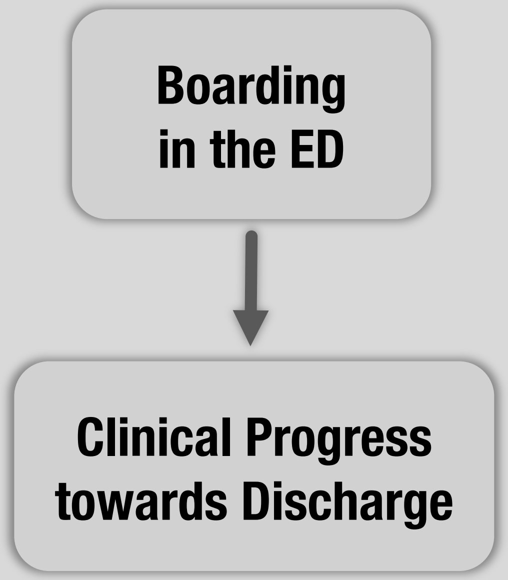

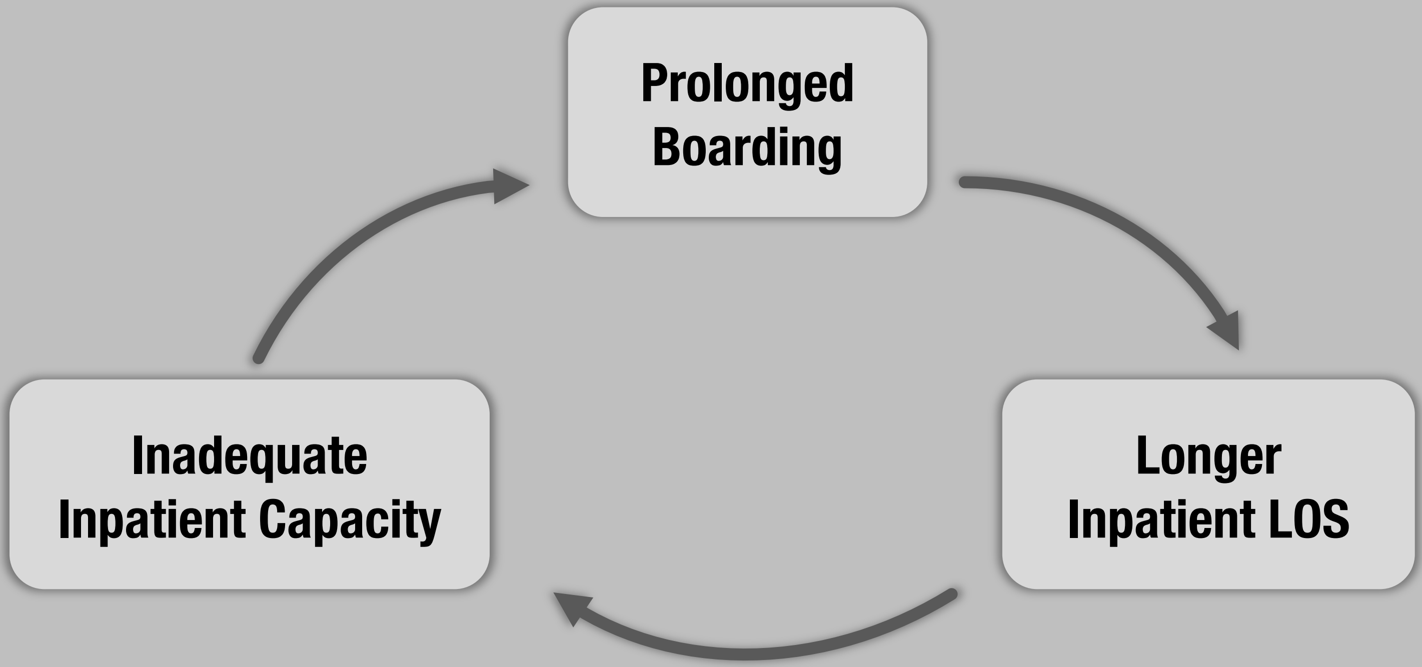

Research Goal

Design bed assignment policies under constrained inpatient capacity

Clinical Collaboration

Data
ED Boarding
Inpatient LOS
Pragmatic Opinion

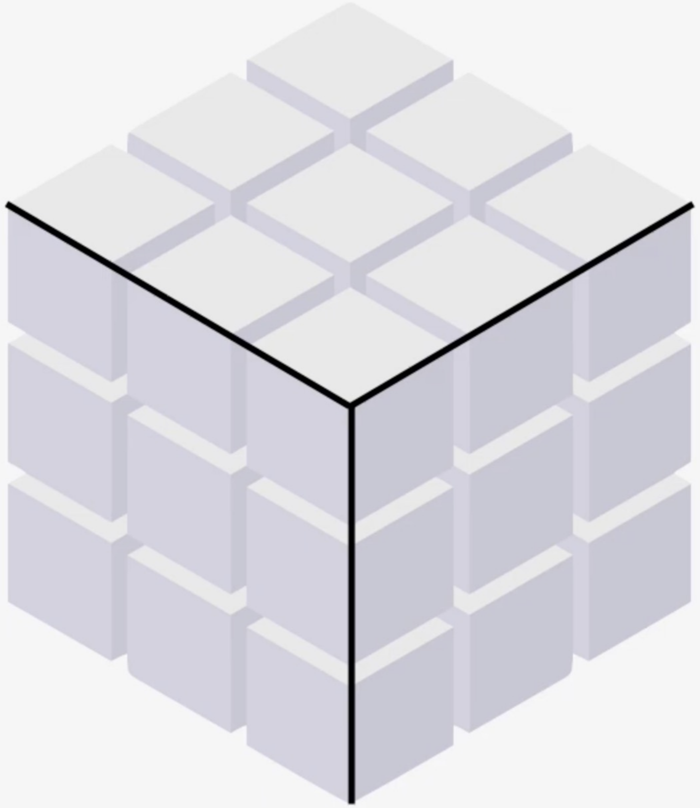

Estimation

Heterogeneous Effect
Econometric Method
Interpretation
Robustness Check

Estimation Model

Measuring heterogenous causal effect of ED boarding on inpatient LOS

$$\boldsymbol{Y_i = \theta(X_i)T_i + g(X_i, W_i) + \varepsilon_i}$$

Estimation Model

Measuring heterogenous causal effect of ED boarding on inpatient LOS

$$\boldsymbol{Y_i = \theta(X_i)T_i + g(X_i, W_i) + \varepsilon_i}$$

Outcome (Y)
Log Inpatient LOS
(ED departure to hospital discharge)

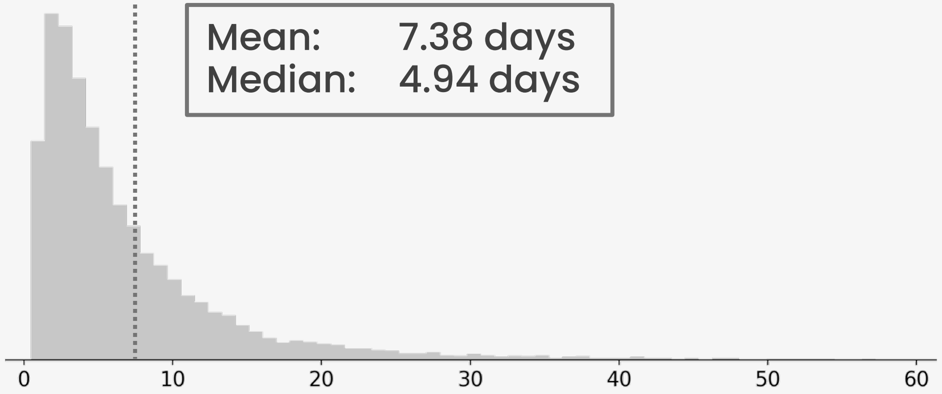

Estimation Model

Measuring heterogenous causal effect of ED boarding on inpatient LOS

$$Y_i = \theta(X_i)T_i + g(X_i, W_i) + \varepsilon_i$$

Outcome (Y)
Log Inpatient LOS
(ED departure to hospital discharge)

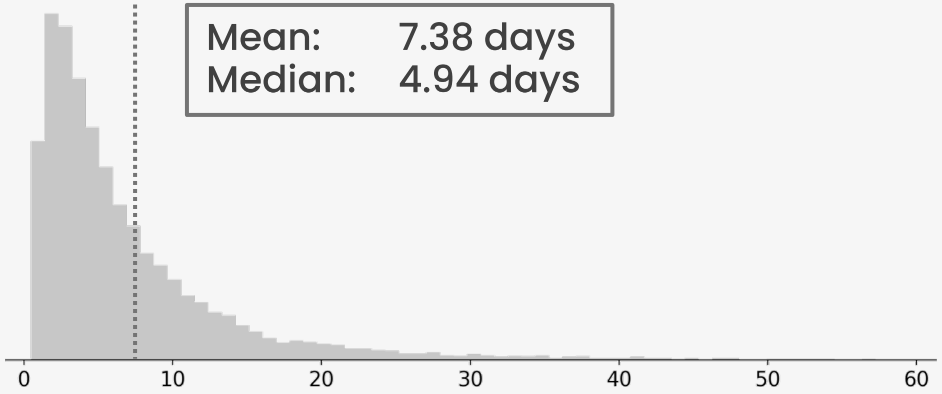

Treatment (T)
Boarding time
(Inpatient bed request to ED departure)

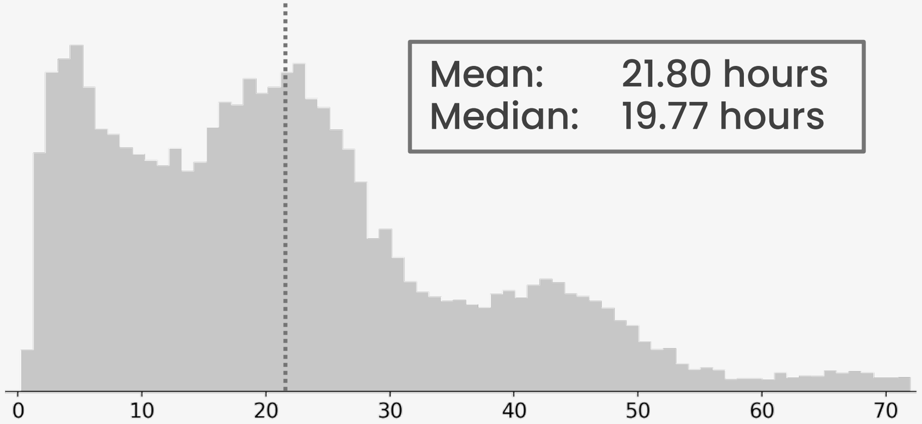

Estimation Model

Measuring heterogenous causal effect of ED boarding on inpatient LOS

$$\boldsymbol{Y_i = \theta(X_i)T_i + g(X_i, W_i) + \varepsilon_i}$$

Control Variables

Features (X)

- Age≥65
- Sex
- Race/Ethnicity
- CCI ≥5 (Charlson Comorbidity Index)
- Comorbidities (diabetes, dementia, hypertension, history of cancer, …)
- Insurance
- Source of admission
- Brought in by EMS
- Long ED duration

Features (W)

Patient visit level: infection status, ED observation, off-service, number of transfers, number of surgical procedures, discharge disposition, etc.

Inpatient unit level (the inpatient unit the patient is admitted to): average occupancy during the stay, squared average occupancy

Time: year, month, day of the week, time of the day

Estimation Model

Measuring heterogenous causal effect of ED boarding on inpatient LOS

$$\boldsymbol{Y_i = \theta(X_i)T_i + g(X_i, W_i) + \varepsilon_i}$$

Double Machine Learning with Instrumental Variables

(Chernozhukov et al 2018)

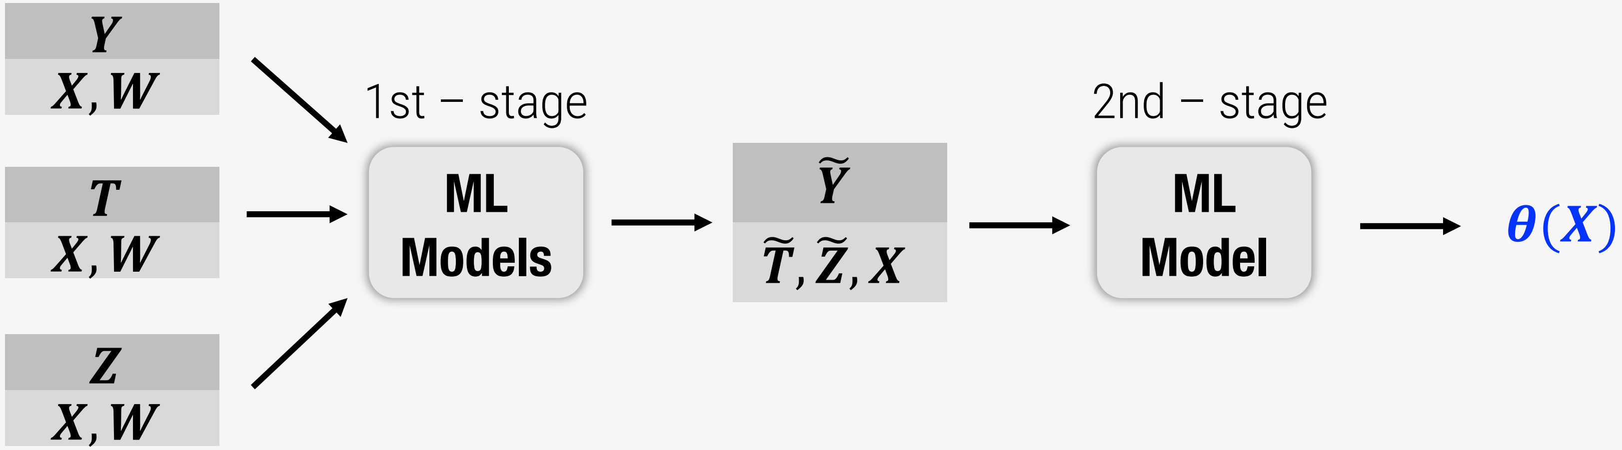

Estimation Result

Measuring heterogenous causal effect of ED boarding on log inpatient LOS $\hat{\boldsymbol{\theta}}(\boldsymbol{X})$

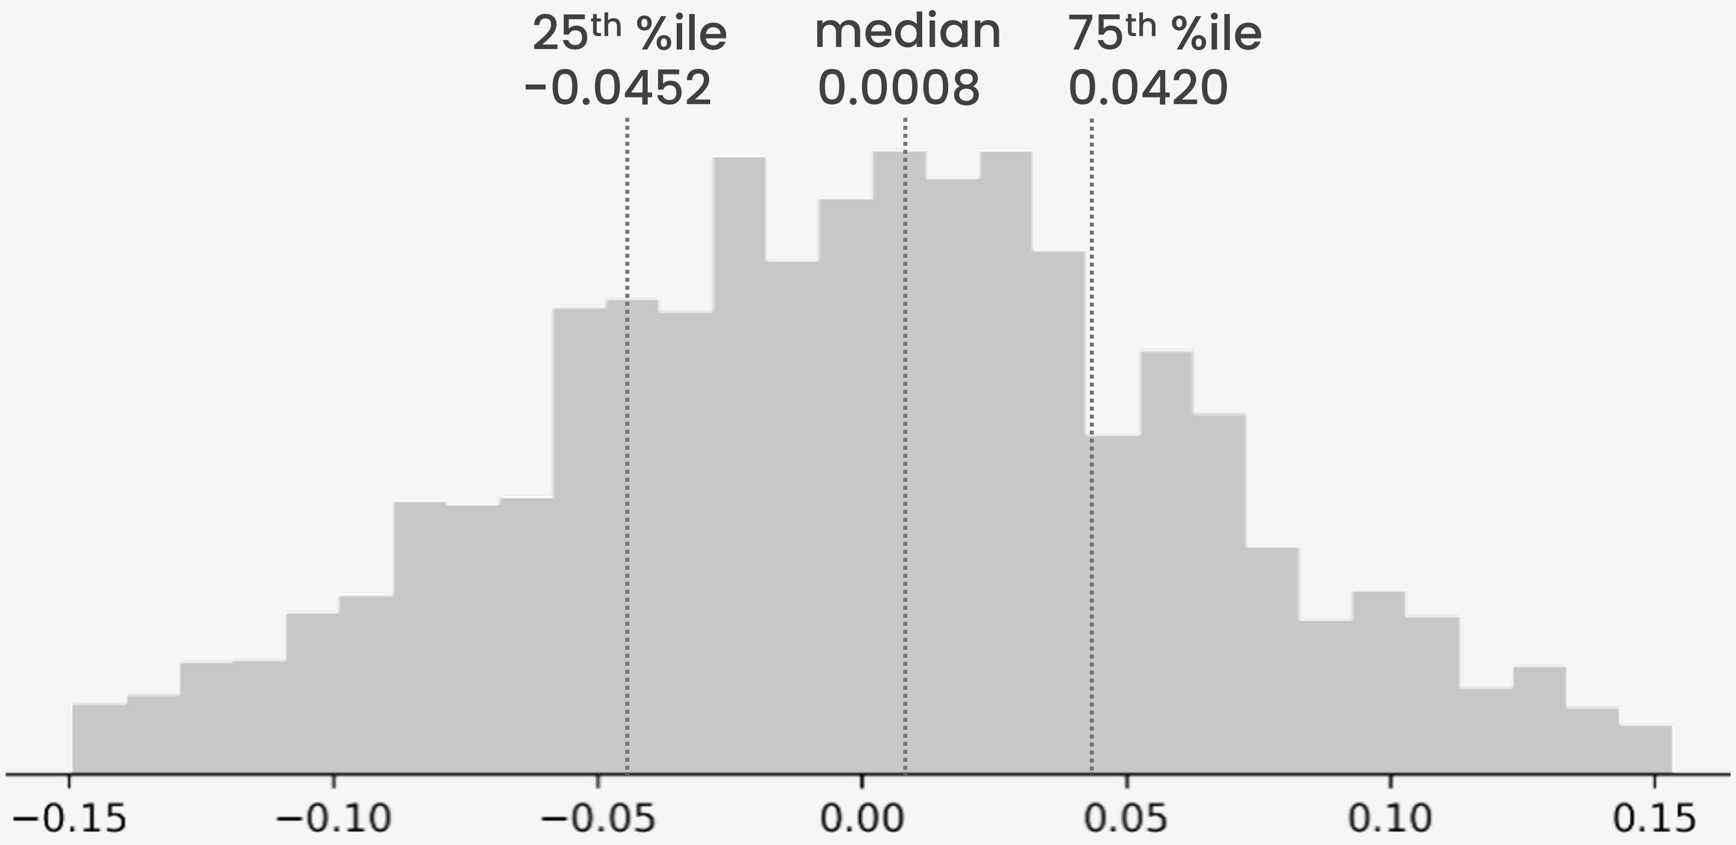

Estimation Interpretation

Measuring heterogenous causal effect of ED boarding on inpatient LOS – Tree Interpreter

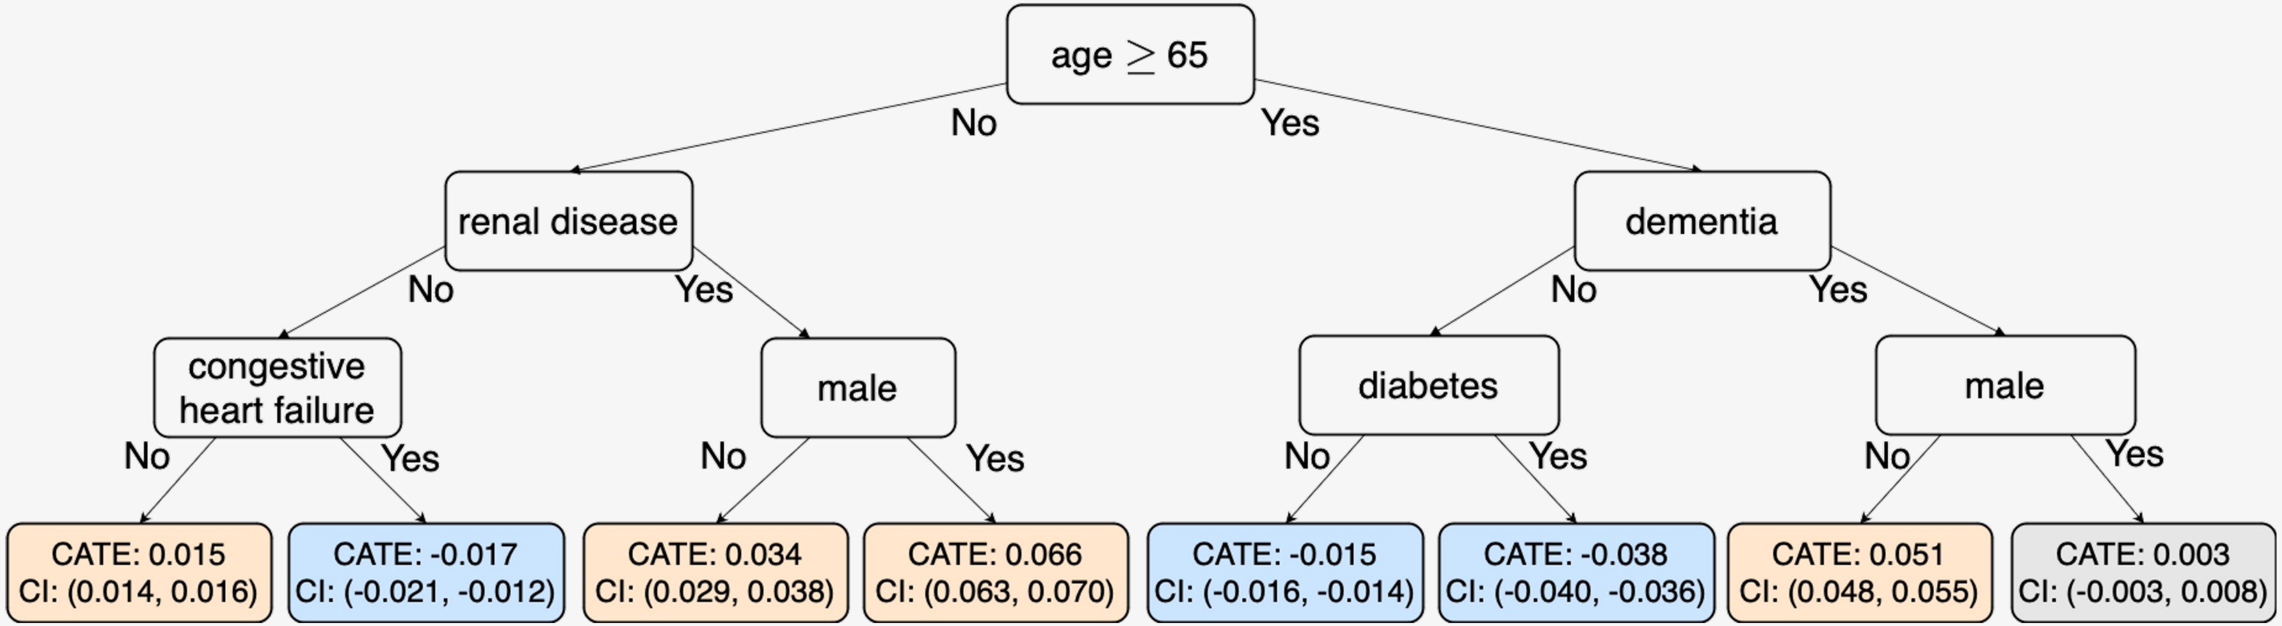

Research Goal

Design bed assignment policies under constrained inpatient capacity

Clinical Collaboration

Data
ED Boarding
Inpatient LOS
Pragmatic Opinion

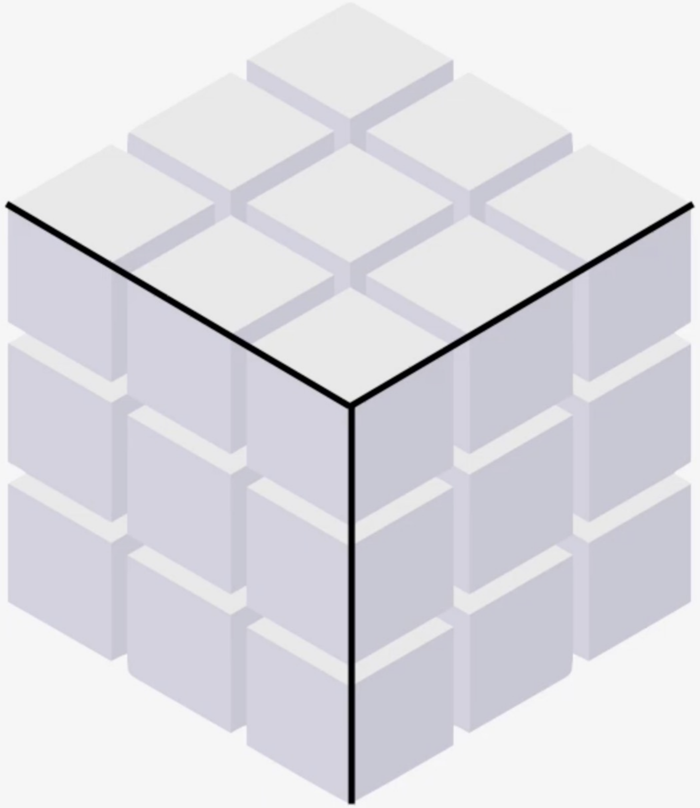

Estimation

Heterogeneous Effect
Econometric Method
Interpretation
Robustness Check

Research Goal

Design bed assignment policies under constrained inpatient capacity

Clinical Collaboration

Data
ED Boarding
Inpatient LOS
Pragmatic Opinion

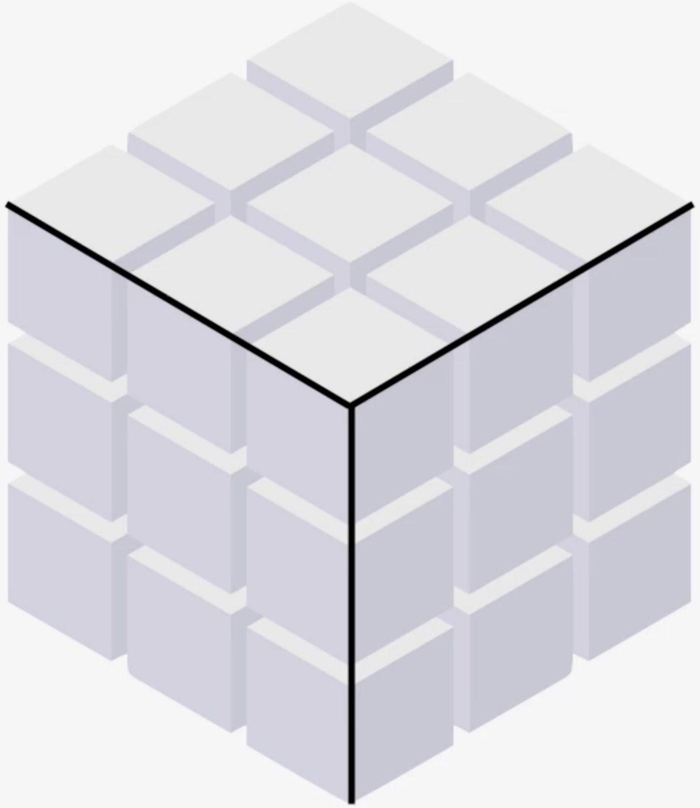

Estimation

Heterogeneous Effect
Econometric Method
Interpretation
Robustness Check

Policy

Inpatient Bed Assignment
Queueing Insight
Trace-Driven Simulation
Efficiency and Fairness

Policy Learning

Develop inpatient bed assignment policies leveraging the heterogeneous estimations

Classic setting:

Static
One-shot resource allocation

(Athey and Wager 2021, Sun et al 2024, Kallus and Zhou 2021, Wang et al 2018)

Our setting:

Dynamic, stochastic system
New bed requests arrive over time
Assign the bed may increase others' wait

Policy Learning

Develop inpatient bed assignment policies leveraging the heterogeneous estimations

Classic setting:

Static
One-shot resource allocation

(Athey and Wager 2021, Sun et al 2024, Kallus and Zhou 2021, Wang et al 2018)

Our setting:

Dynamic, stochastic system
New bed requests arrive over time
Assign the bed may increase others' wait

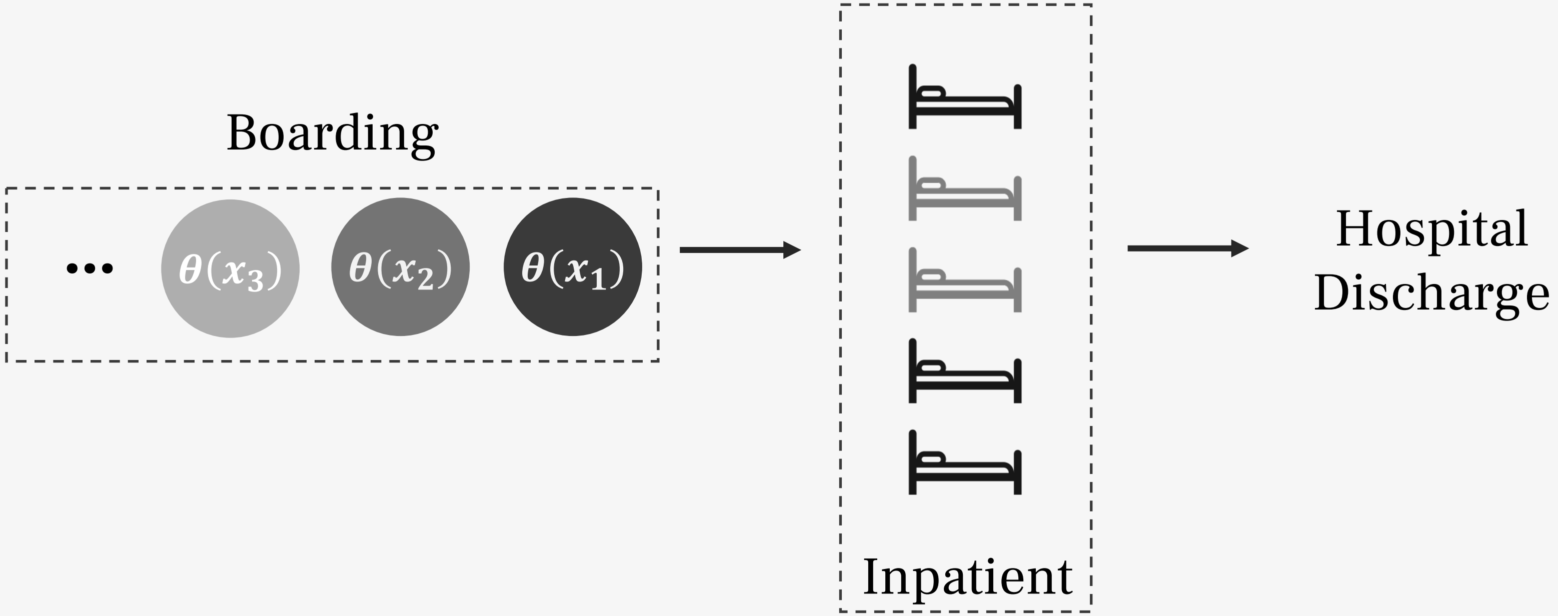

Policy Learning

Develop inpatient bed assignment policies $\boldsymbol{\pi}(\boldsymbol{X})$

HTE$_\mathbf{I}$: prioritize based on individual $\hat{\theta}(X_i)$

Policy Learning

Develop inpatient bed assignment policies $\boldsymbol{\pi}(\boldsymbol{X})$

HTE$_1$: prioritize based on individual $\hat{\theta}(X_i)$

HTE$_2$: $\sup_{\pi} \mathbb{E}\left[\hat{\theta}(X)\pi(X)\right] \quad \mathrm{s.t.}\ \mathbb{E}[\pi(X)] \leq B$

Policy Learning

Develop inpatient bed assignment policies $\boldsymbol{\pi}(\boldsymbol{X})$ while ensuring fairness

HTE$_I$: prioritize based on individual $\hat{\theta}(X_i)$

HTE$_2$: $\sup_{\pi} \mathbb{E}\left[\hat{\theta}(X)\pi(X)\right] \quad \text{s.t.}\ \mathbb{E}[\pi(X)] \leq B$

HTE$_{2,Xsub}$: $\sup_{\pi} \mathbb{E}\left[\hat{\theta}(X)\pi(X_{sub})\right]$

Policy Learning

Develop inpatient bed assignment policies $\pi(X)$ while ensuring fairness

HTE_{I}: prioritize based on individual $\hat{\theta}(X_i)$

HTE_2: $\sup_{\pi} \mathbb{E}\left[\hat{\theta}(X)\pi(X)\right] \quad \text{s.t.}\ \mathbb{E}[\pi(X)] \leq B$

$\text{HTE}_{2,\text{Xsub}}$: $\sup_{\pi} \mathbb{E}\left[\hat{\theta}(X)\pi(X_{sub})\right]$

$\text{HTE}_{2,=}$: $\sup_{\pi} \mathbb{E}\left[\hat{\theta}(X)\pi(X)\right] \quad \text{s.t.}\ \left|\mathbb{E}[\pi(X)|G_i] - \mathbb{E}\left[\pi(X)|G_j\right]\right| \leq \delta$

Policy Learning

Develop inpatient bed assignment policies $\pi(X)$ while ensuring fairness

HTE$_I$: prioritize based on individual $\hat{\theta}(X_i)$

HTE$_2$: $\sup_{\pi} \mathbb{E}\left[\hat{\theta}(X)\pi(X)\right] \quad \text{s.t.}\ \mathbb{E}[\pi(X)] \le B$

HTE$_{2,Xsub}$: $\sup_{\pi} \mathbb{E}\left[\hat{\theta}(X)\pi(X_{sub})\right]$

HTE$_{2,=}$: $\sup_{\pi} \mathbb{E}\left[\hat{\theta}(X)\pi(X)\right] \quad \text{s.t.}\ \left|\mathbb{E}[\pi(X)|G_i] - \mathbb{E}\left[\pi(X)|G_j\right]\right| \le \delta$

High-Fidelity Trace-Driven Simulation

Policy Performance

Develop inpatient bed assignment policies $\pi(X)$ while ensuring fairness

	FCFS	HTE_I	HTE_2	HTE_{2,X_sub}	$HTE_{2,=}$
Mean Boarding	20.29	5.36	4.62	5.02	4.36
Median Inpatient LOS	4.76	3.96	3.98	4.07	4.03

Policy Performance

Develop inpatient bed assignment policies $\pi(X)$ while ensuring fairness

	FCFS	HTE_I	HTE_2	HTE_{2,X_sub}	$HTE_{2,=}$
Mean Boarding	20.29	5.36	4.62	5.02	4.36
Median Inpatient LOS	4.76	3.96	3.98	4.07	4.03

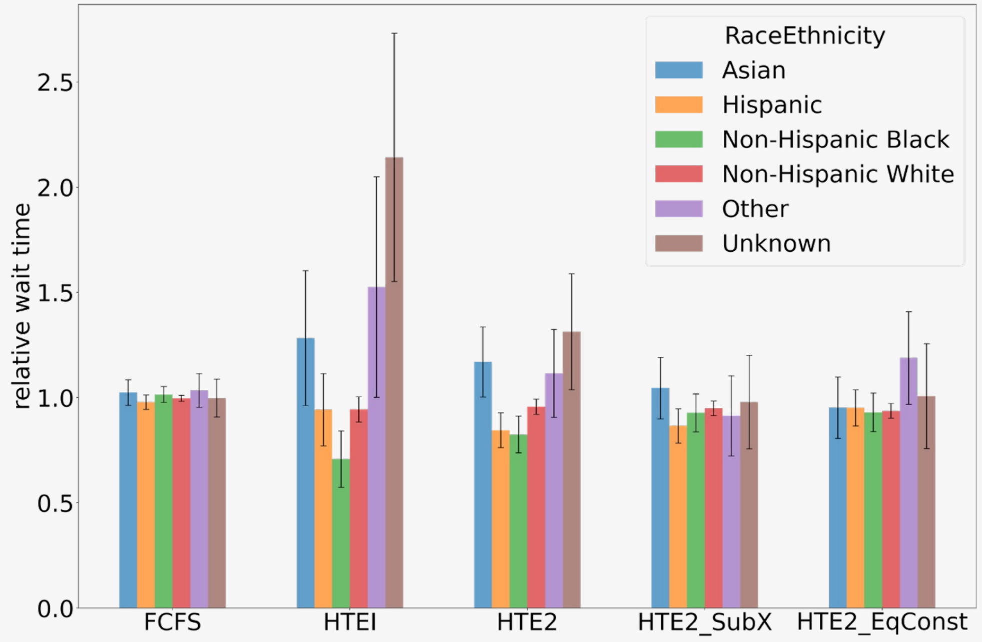

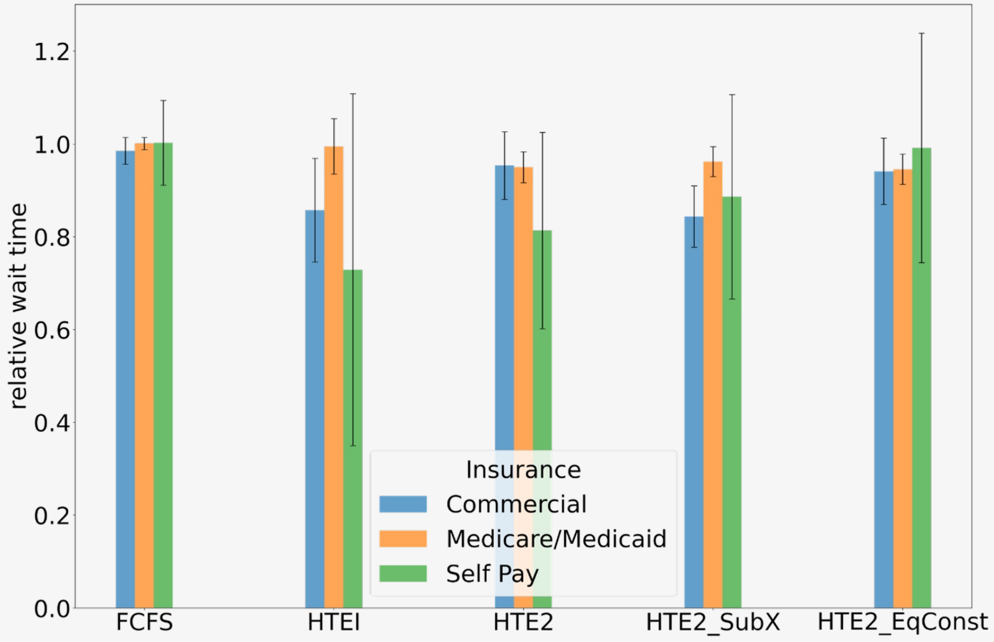

Primary Contribution

Combine heterogeneous treatment effect estimation, queueing theory, and policy learning to design efficient inpatient bed assignment policies that reduce boarding time while promoting fairness.

Thank You

Q&A

2026.02.02

Email xl2500@stern.nyu.edu
URL https://liualyssa.github.io

NYU Stern
Department of Technology, Operations, & Statistics

Policy Performance

Develop inpatient bed assignment policies $\pi(X)$

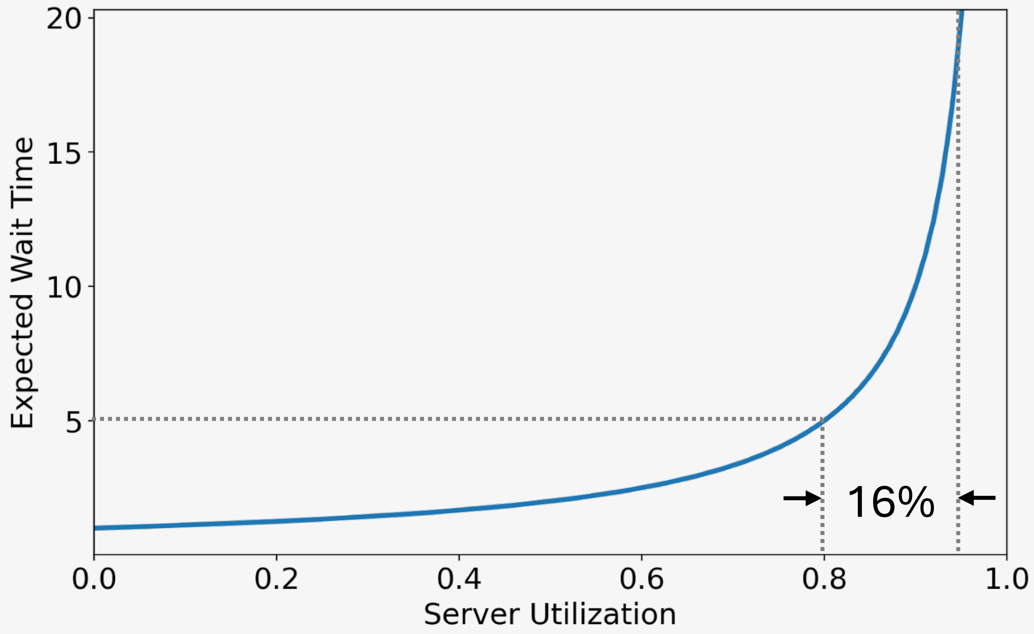

	FCFS	HTE_1	HTE_2
Mean Boarding	20.29	5.36	4.62
Median Inpatient LOS	4.76	3.96	3.98

75% ↓

0.8 days ↓

Policy Performance

Develop inpatient bed assignment policies $\pi(X)$ while ensuring **individual** fairness

$$\widehat{\theta}(X) + \alpha\widetilde{W}$$

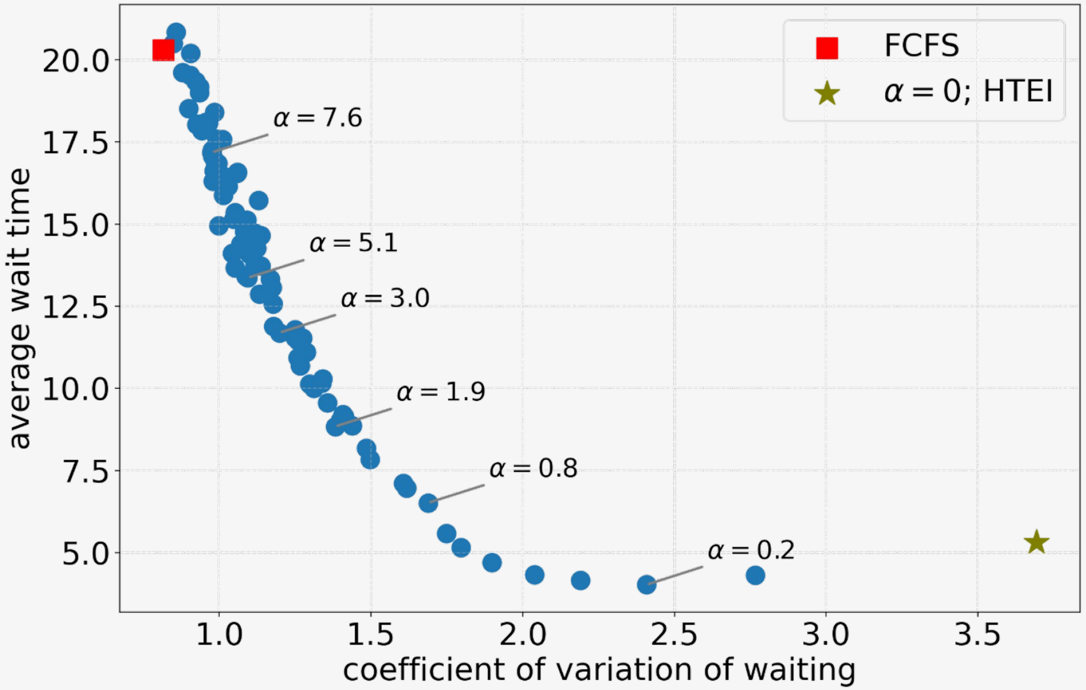

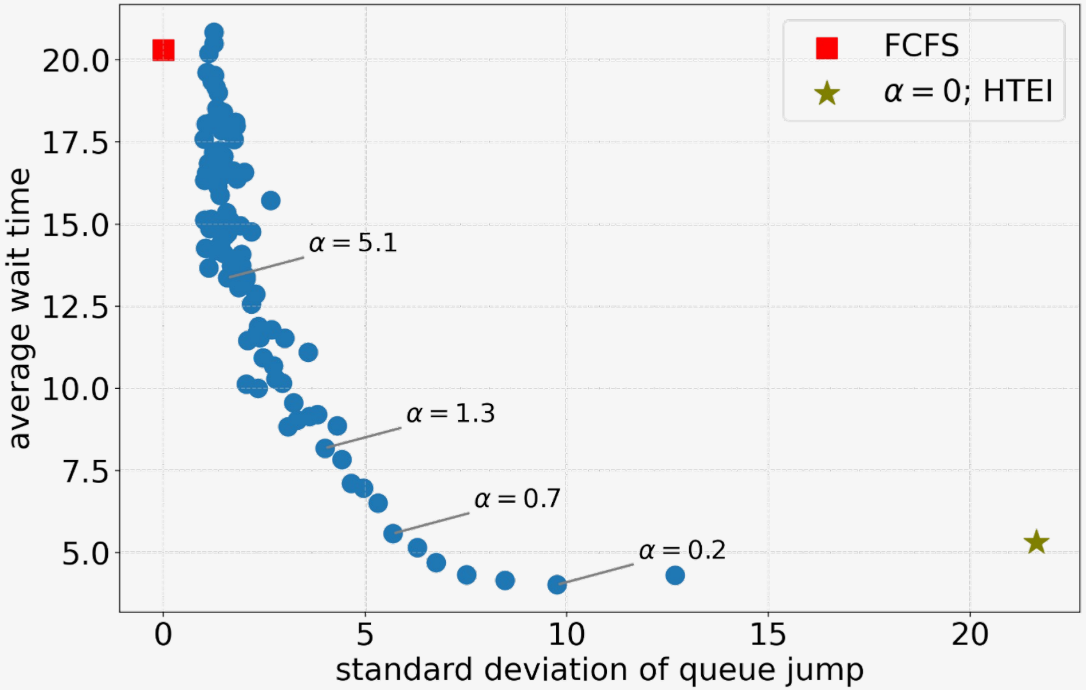

Policy Learning

Interpretability and implementation

$$\textbf{HTE}_{\mathbf{2,Xsub}}: \sup_{\pi} E[\hat{\theta}(X)\pi(X_{sub})]$$

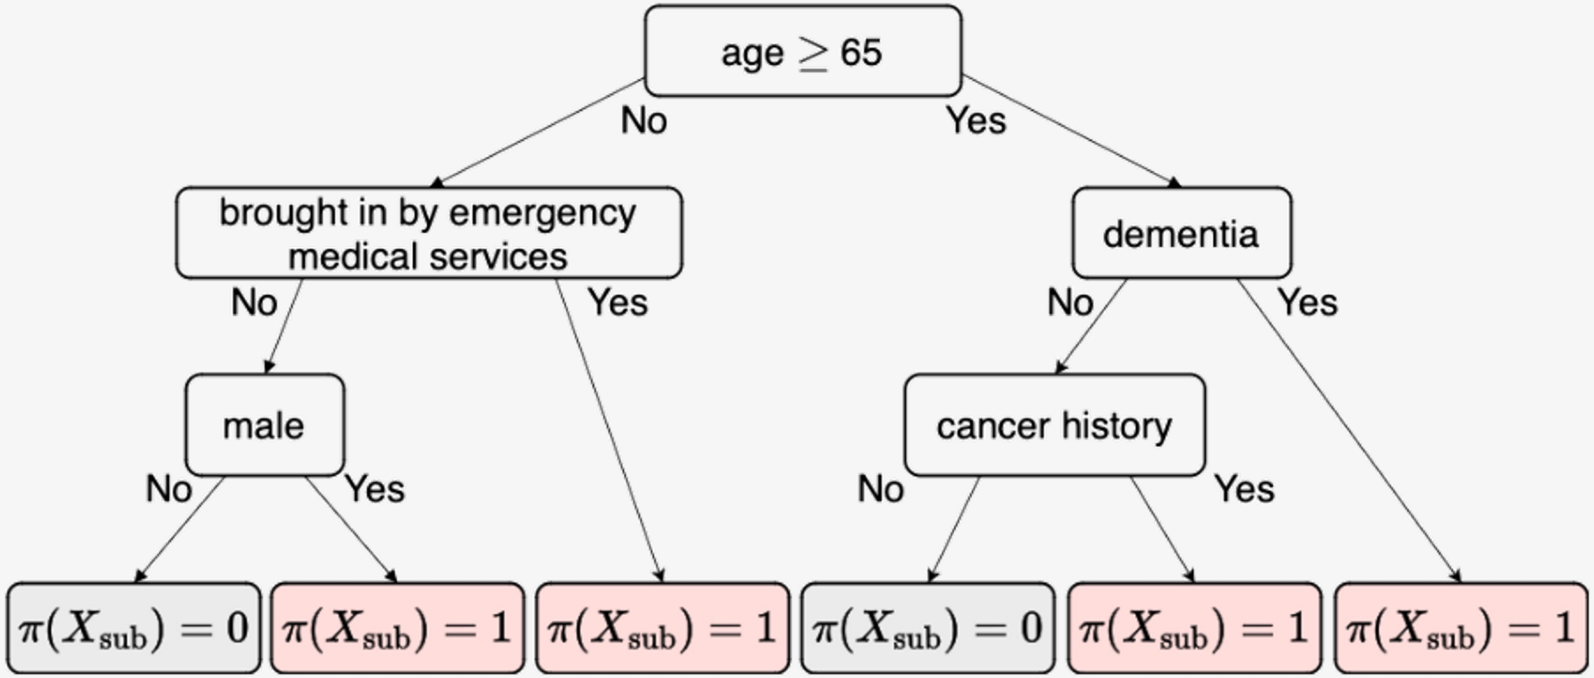

Policy Learning

Develop inpatient bed assignment policies – Queueing insight M/M(W)/n/K

$$\theta\left(\bar{S}\exp\left(\theta\widetilde{W}\right)\right)^{-1}$$

$$\theta\bar{S}^{-1}$$

$$\boldsymbol{\theta}$$

$$\theta\bar{S}\exp\left(\theta\widetilde{W}\right)$$

since $\bar{S}\exp\left(\theta(\widetilde{W}+dt)\right)-\bar{S}\exp\left(\theta\widetilde{W}\right)\approx\theta\bar{S}\exp\left(\theta\widetilde{W}\right)dt$

Improving Access to Essential Medicines in Sierra Leone with Decision-Aware Learning

Tsai-Hsuan (Angel) Chung

with Hamsa Bastani, Osbert Bastani, and Sierra Leone Ministry of Health

Context in Sierra Leone

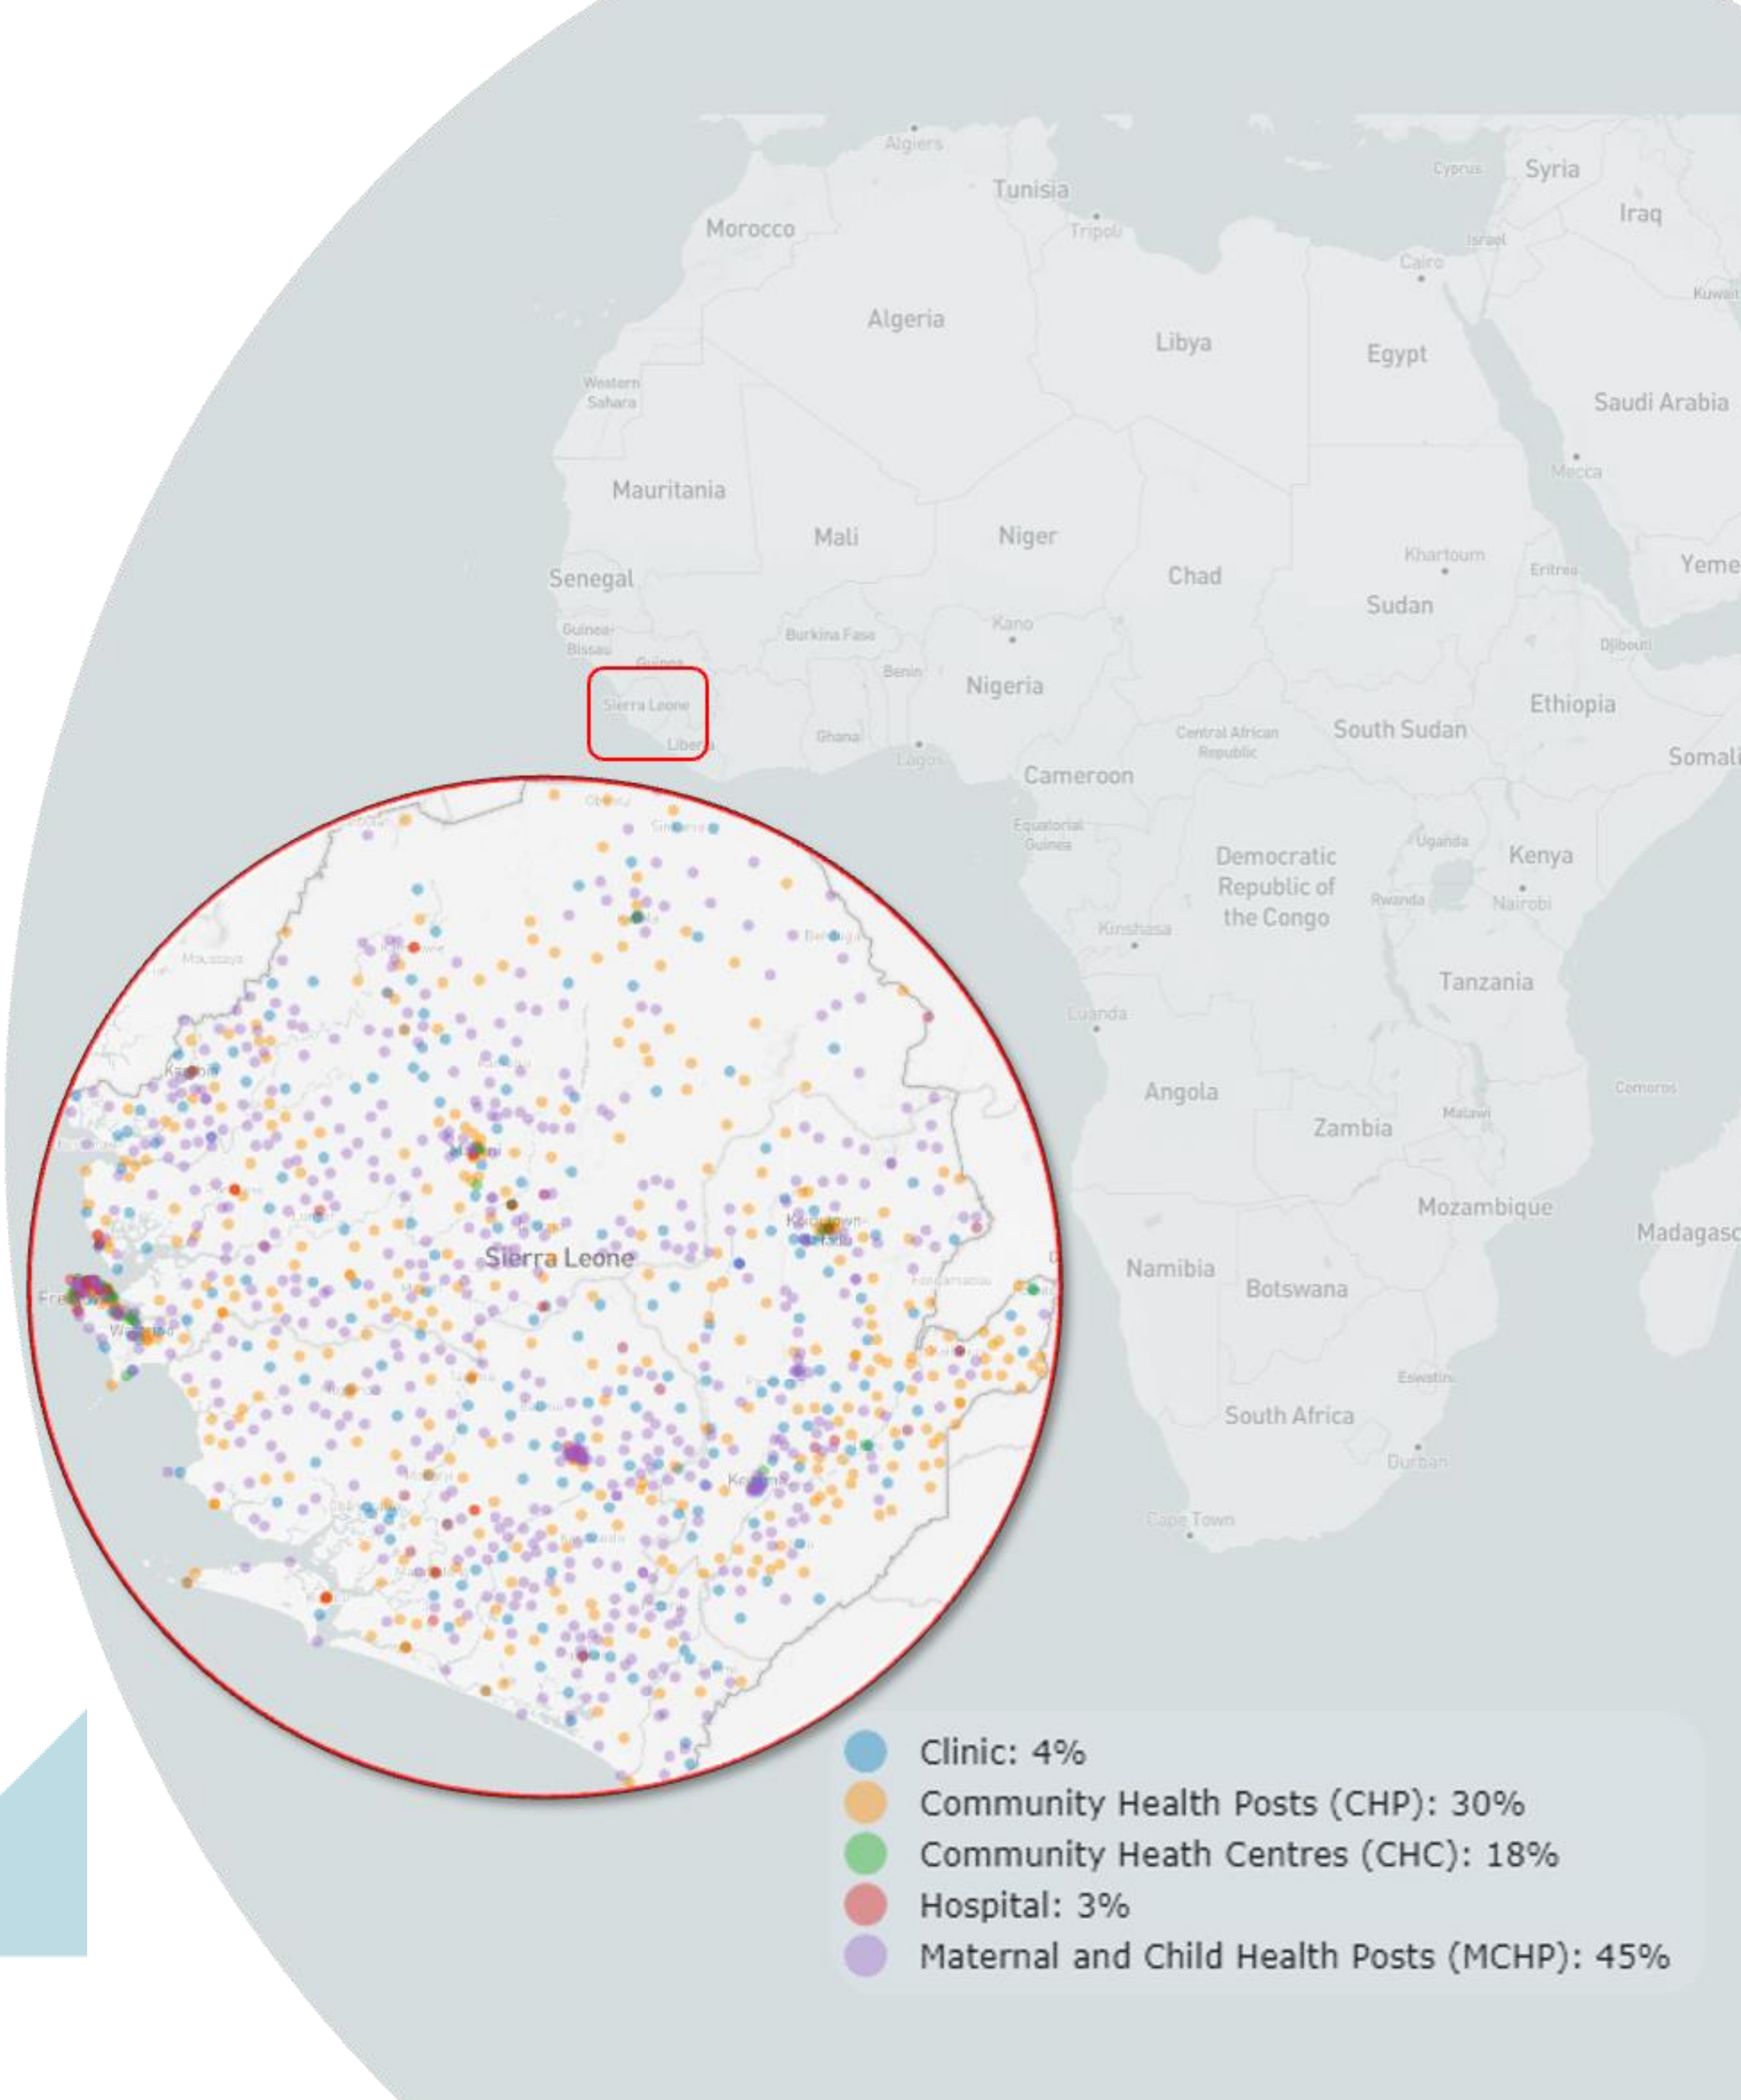

Wharton
University of Pennsylvania

Context in Sierra Leone

- From 2022 report: GDP per capita = $532

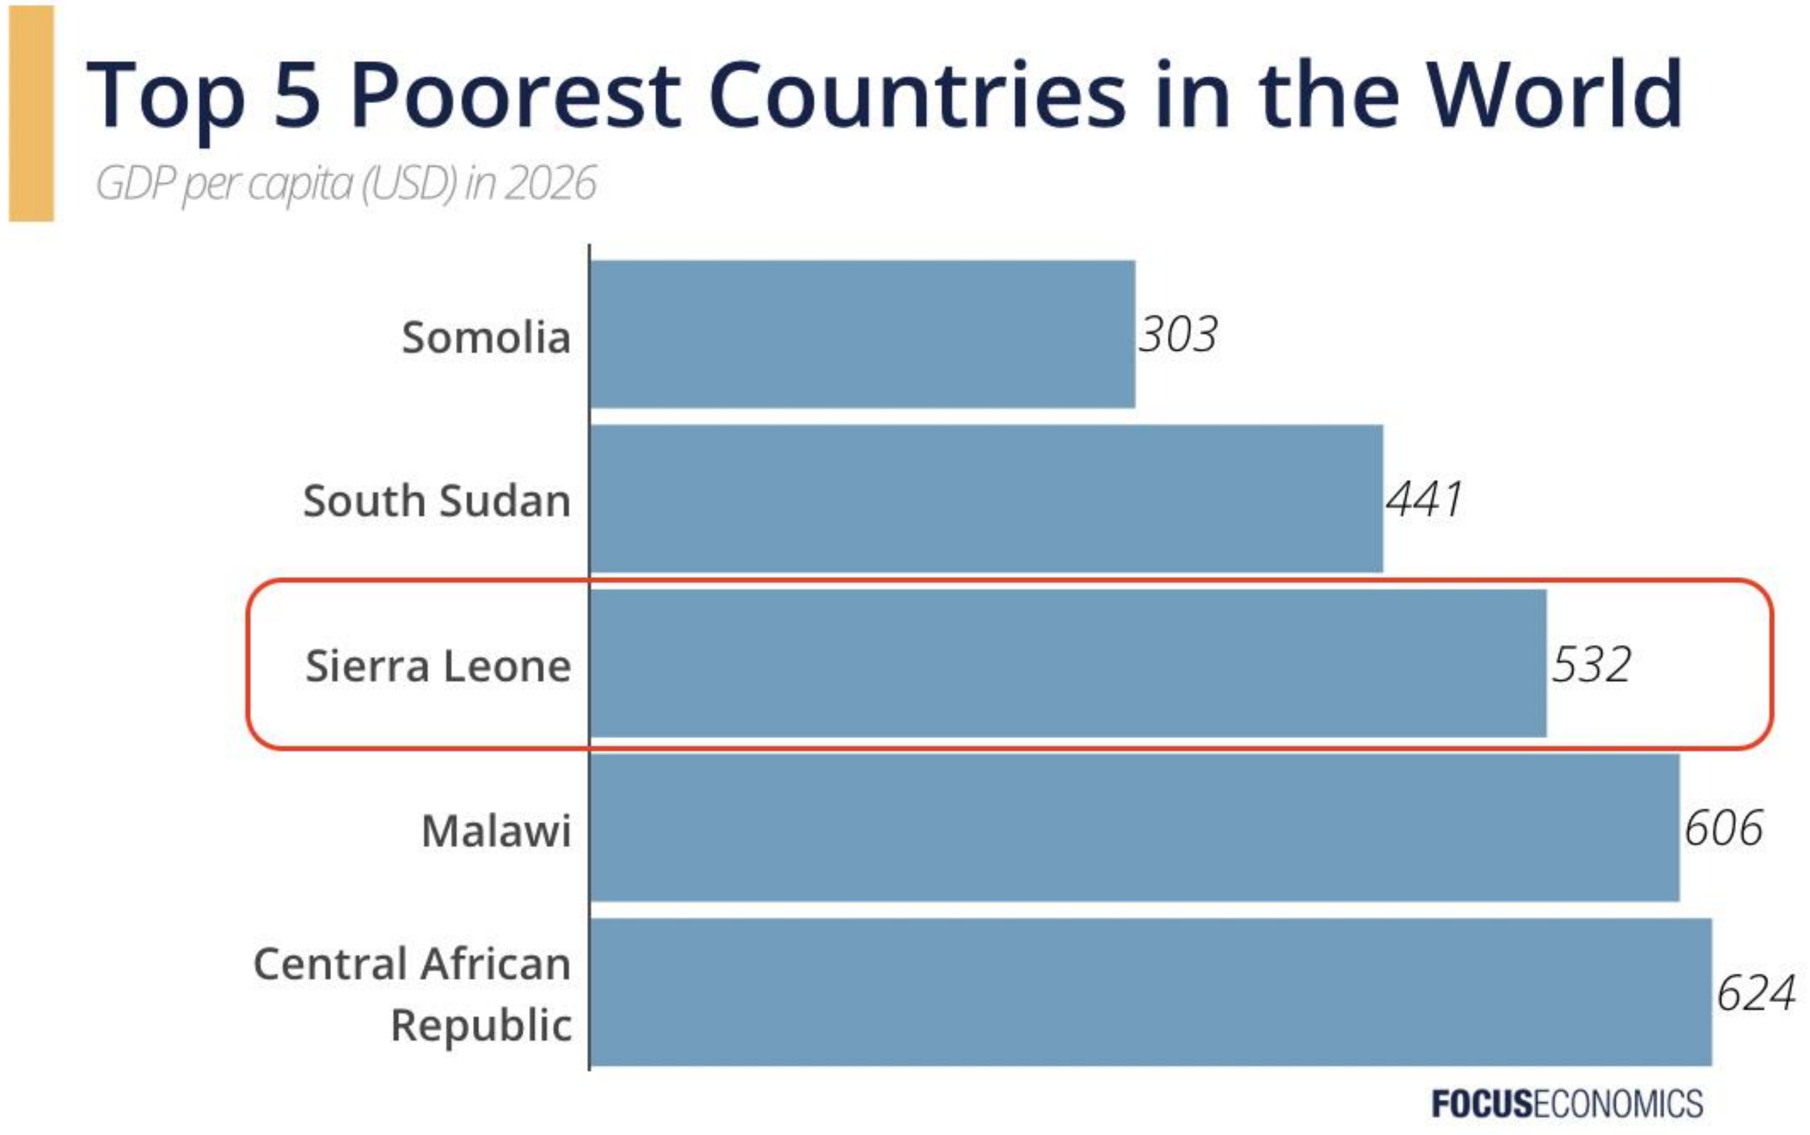

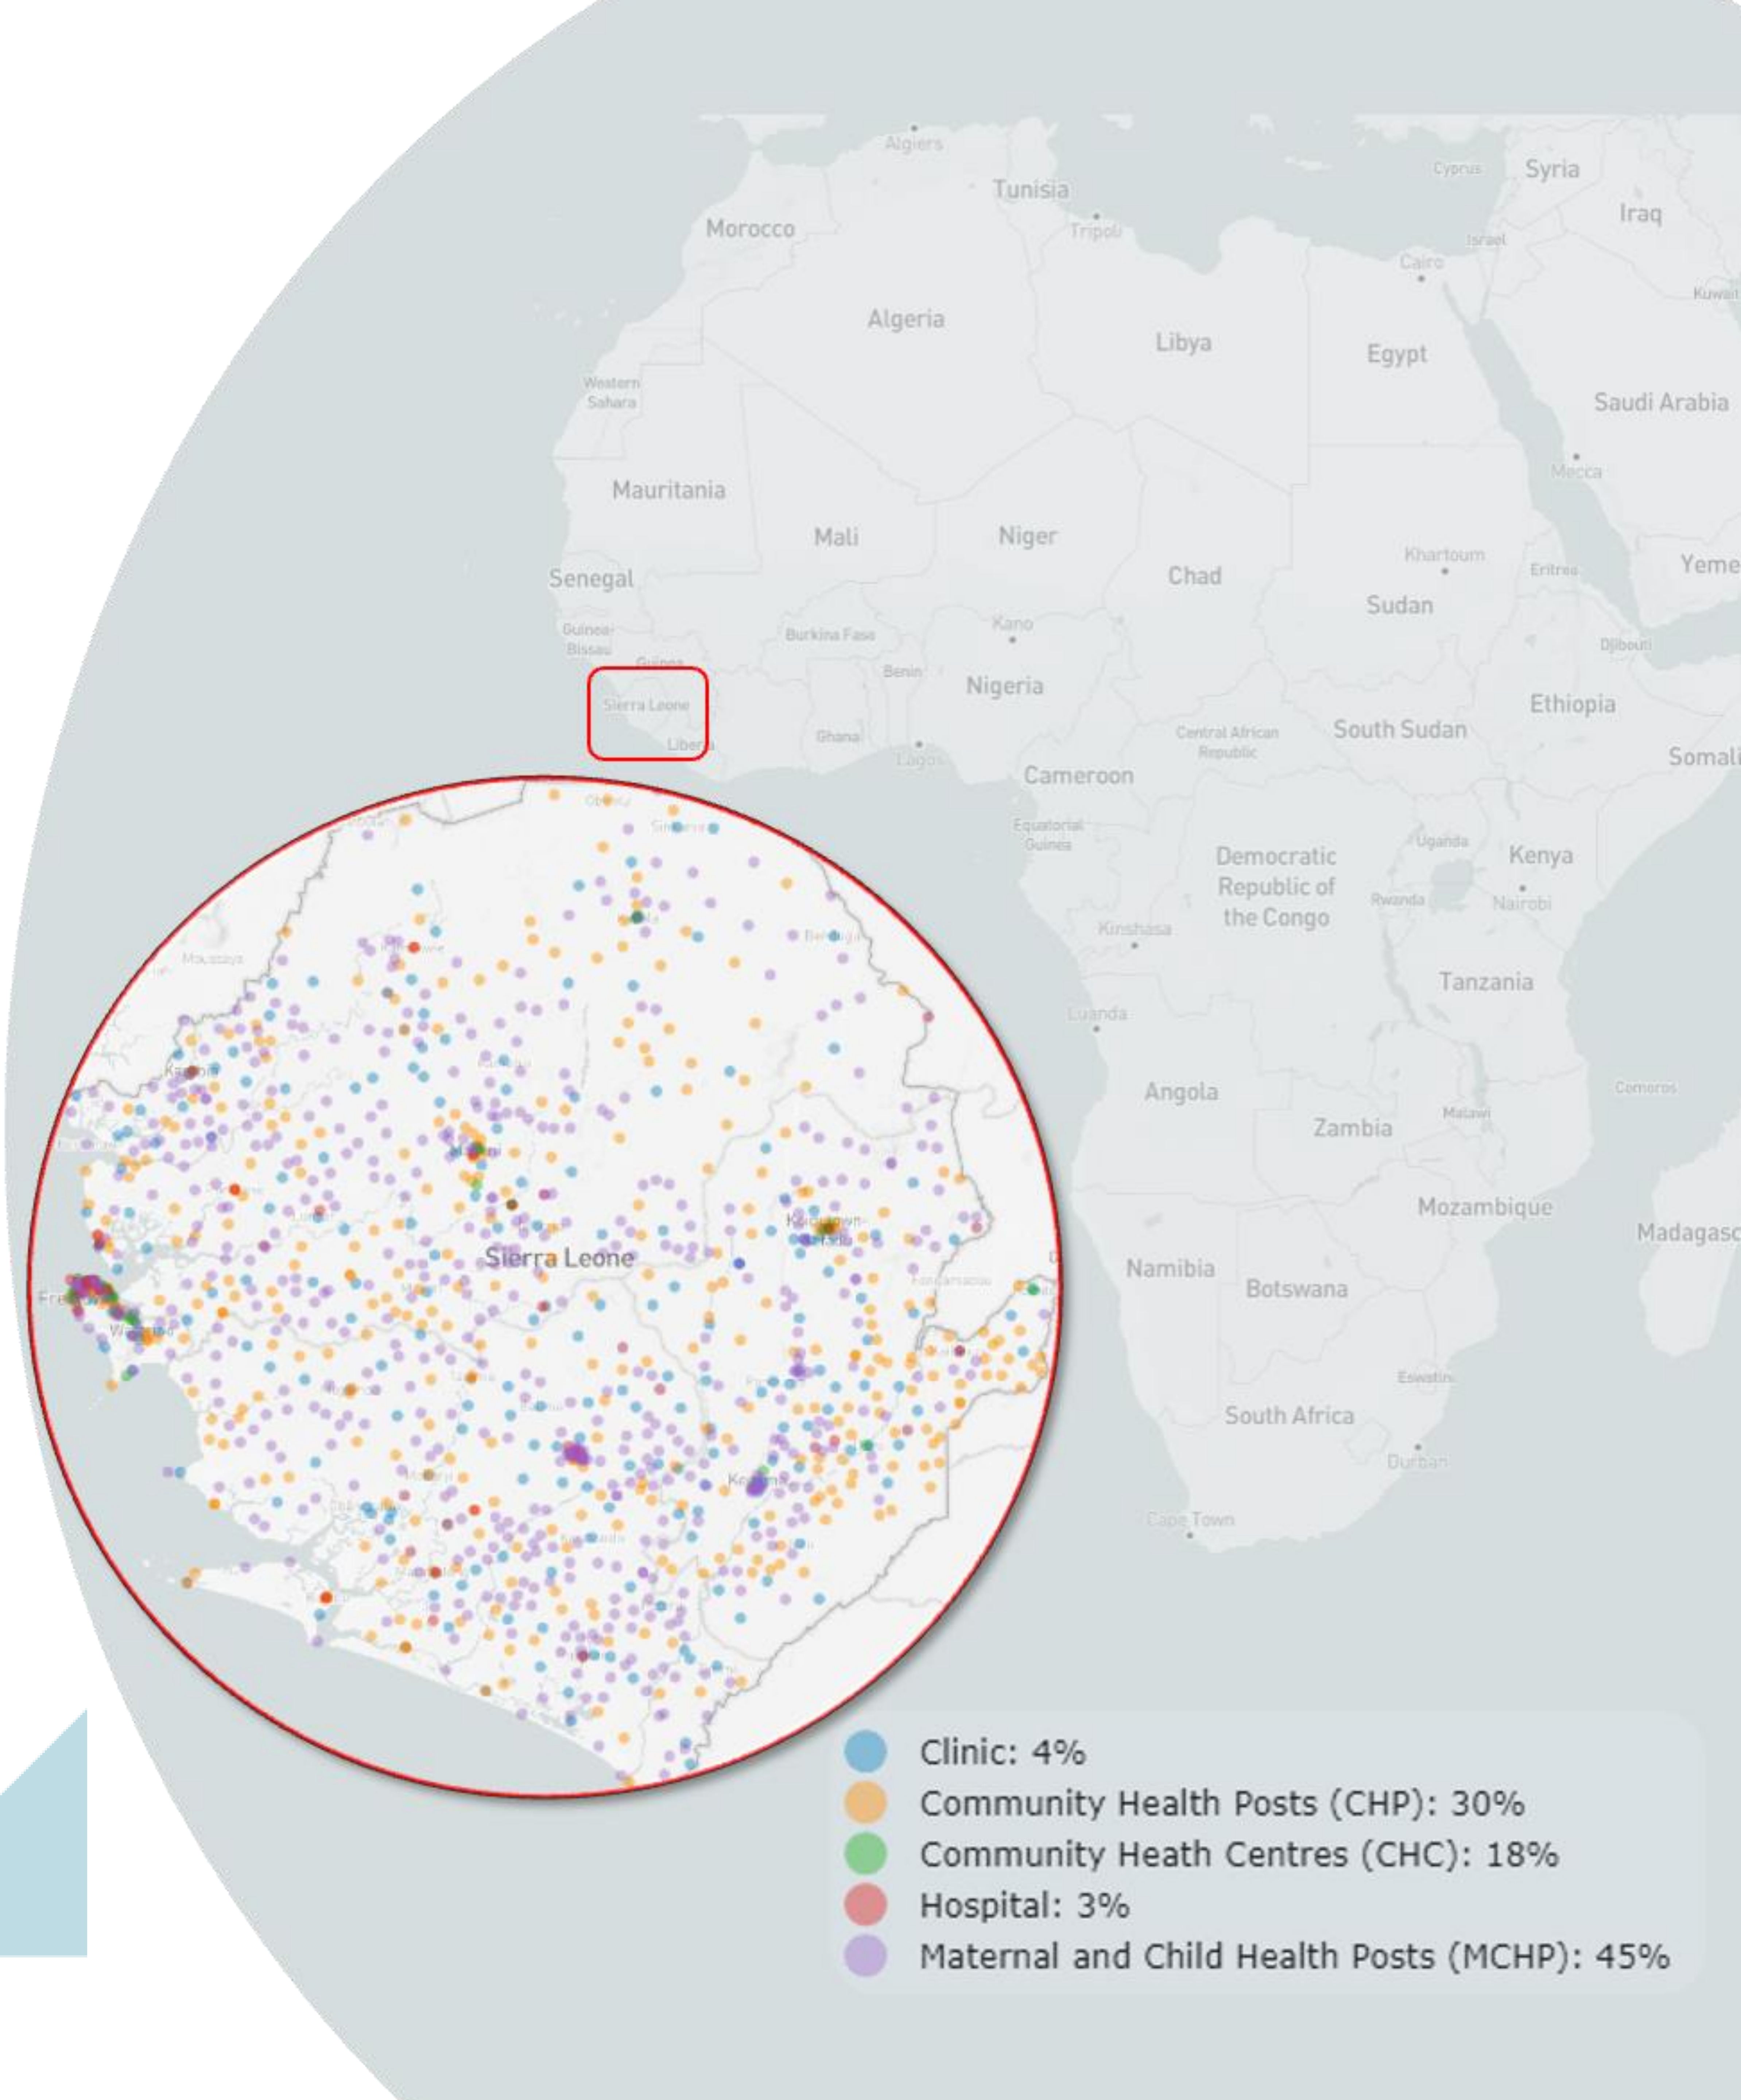

Wharton UNIVERSITY of PENNSYLVANIA

Context in Sierra Leone

- From 2022 report: GDP per capita = $532

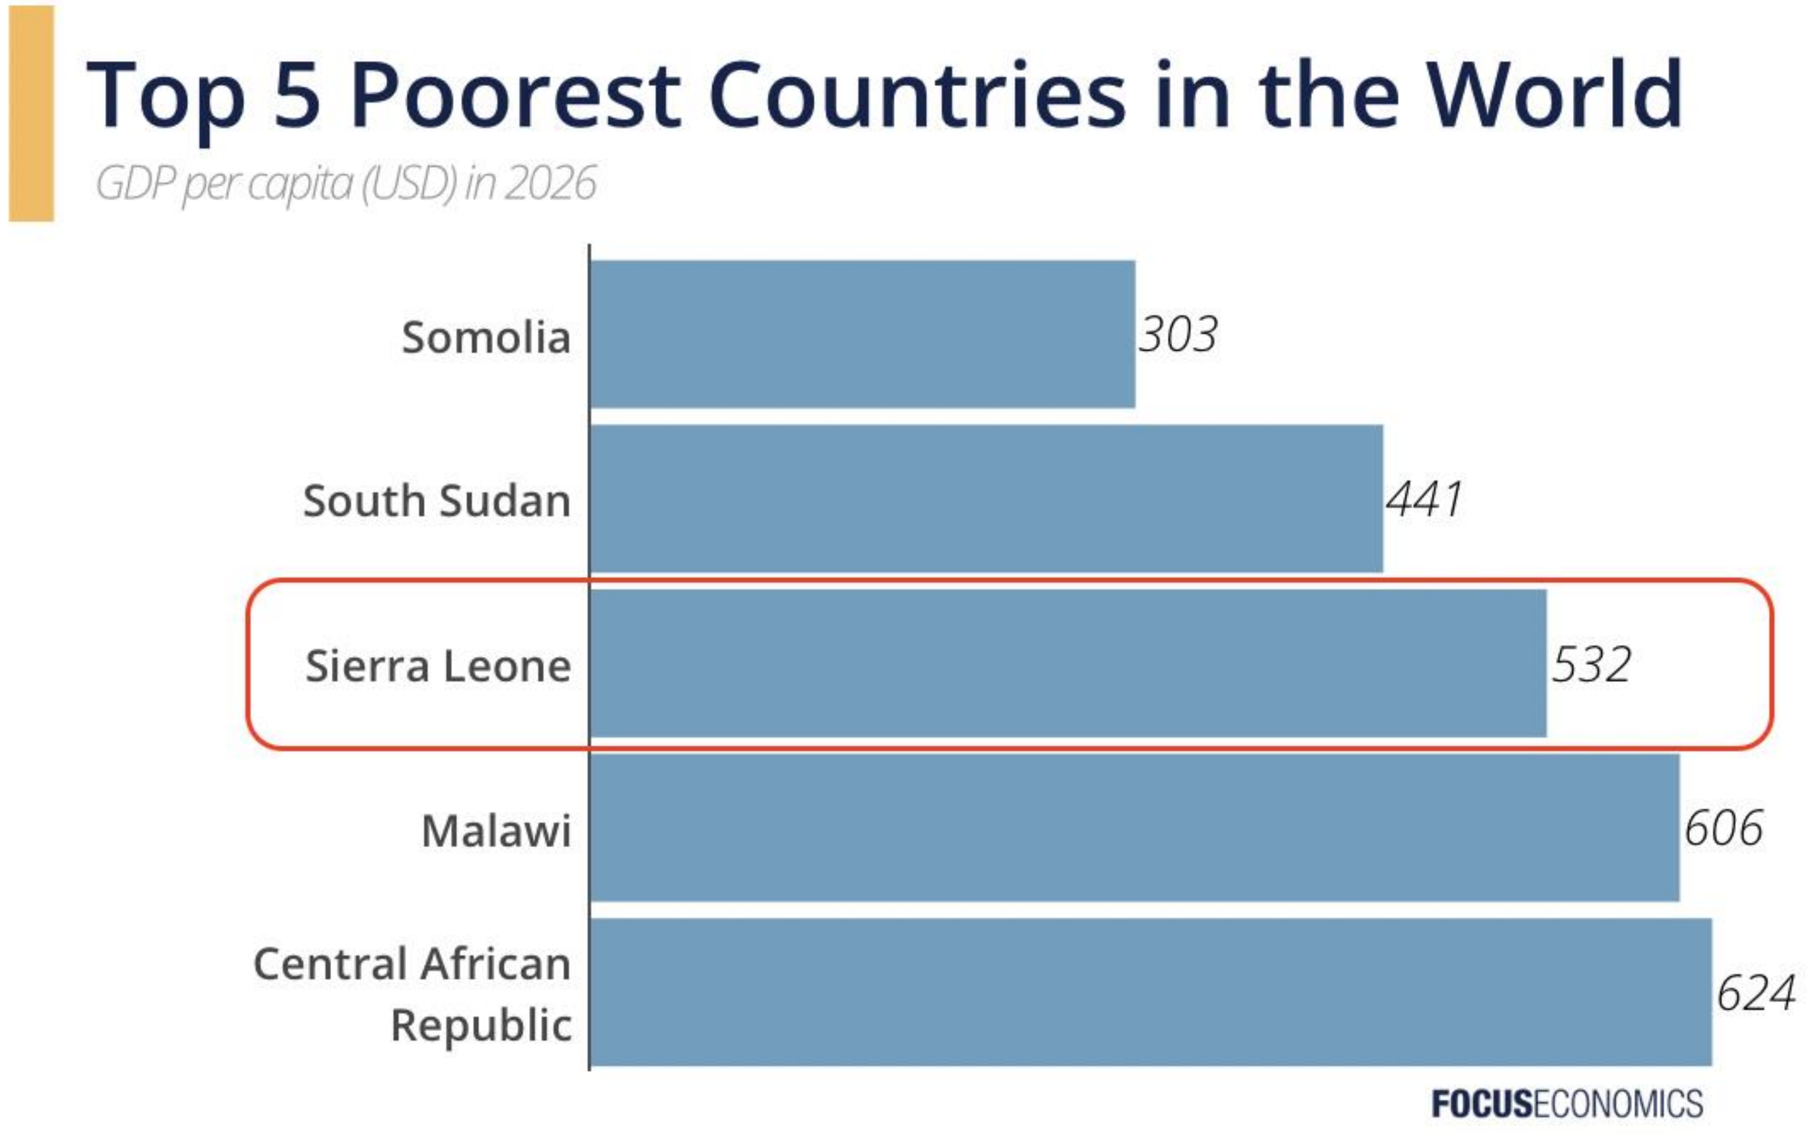

- Child mortality rate is **3.11%** (**0.56%** in the US)
- Maternal mortality: **510** deaths per 100,000 live births (**22.3** per 100,000 in the US)

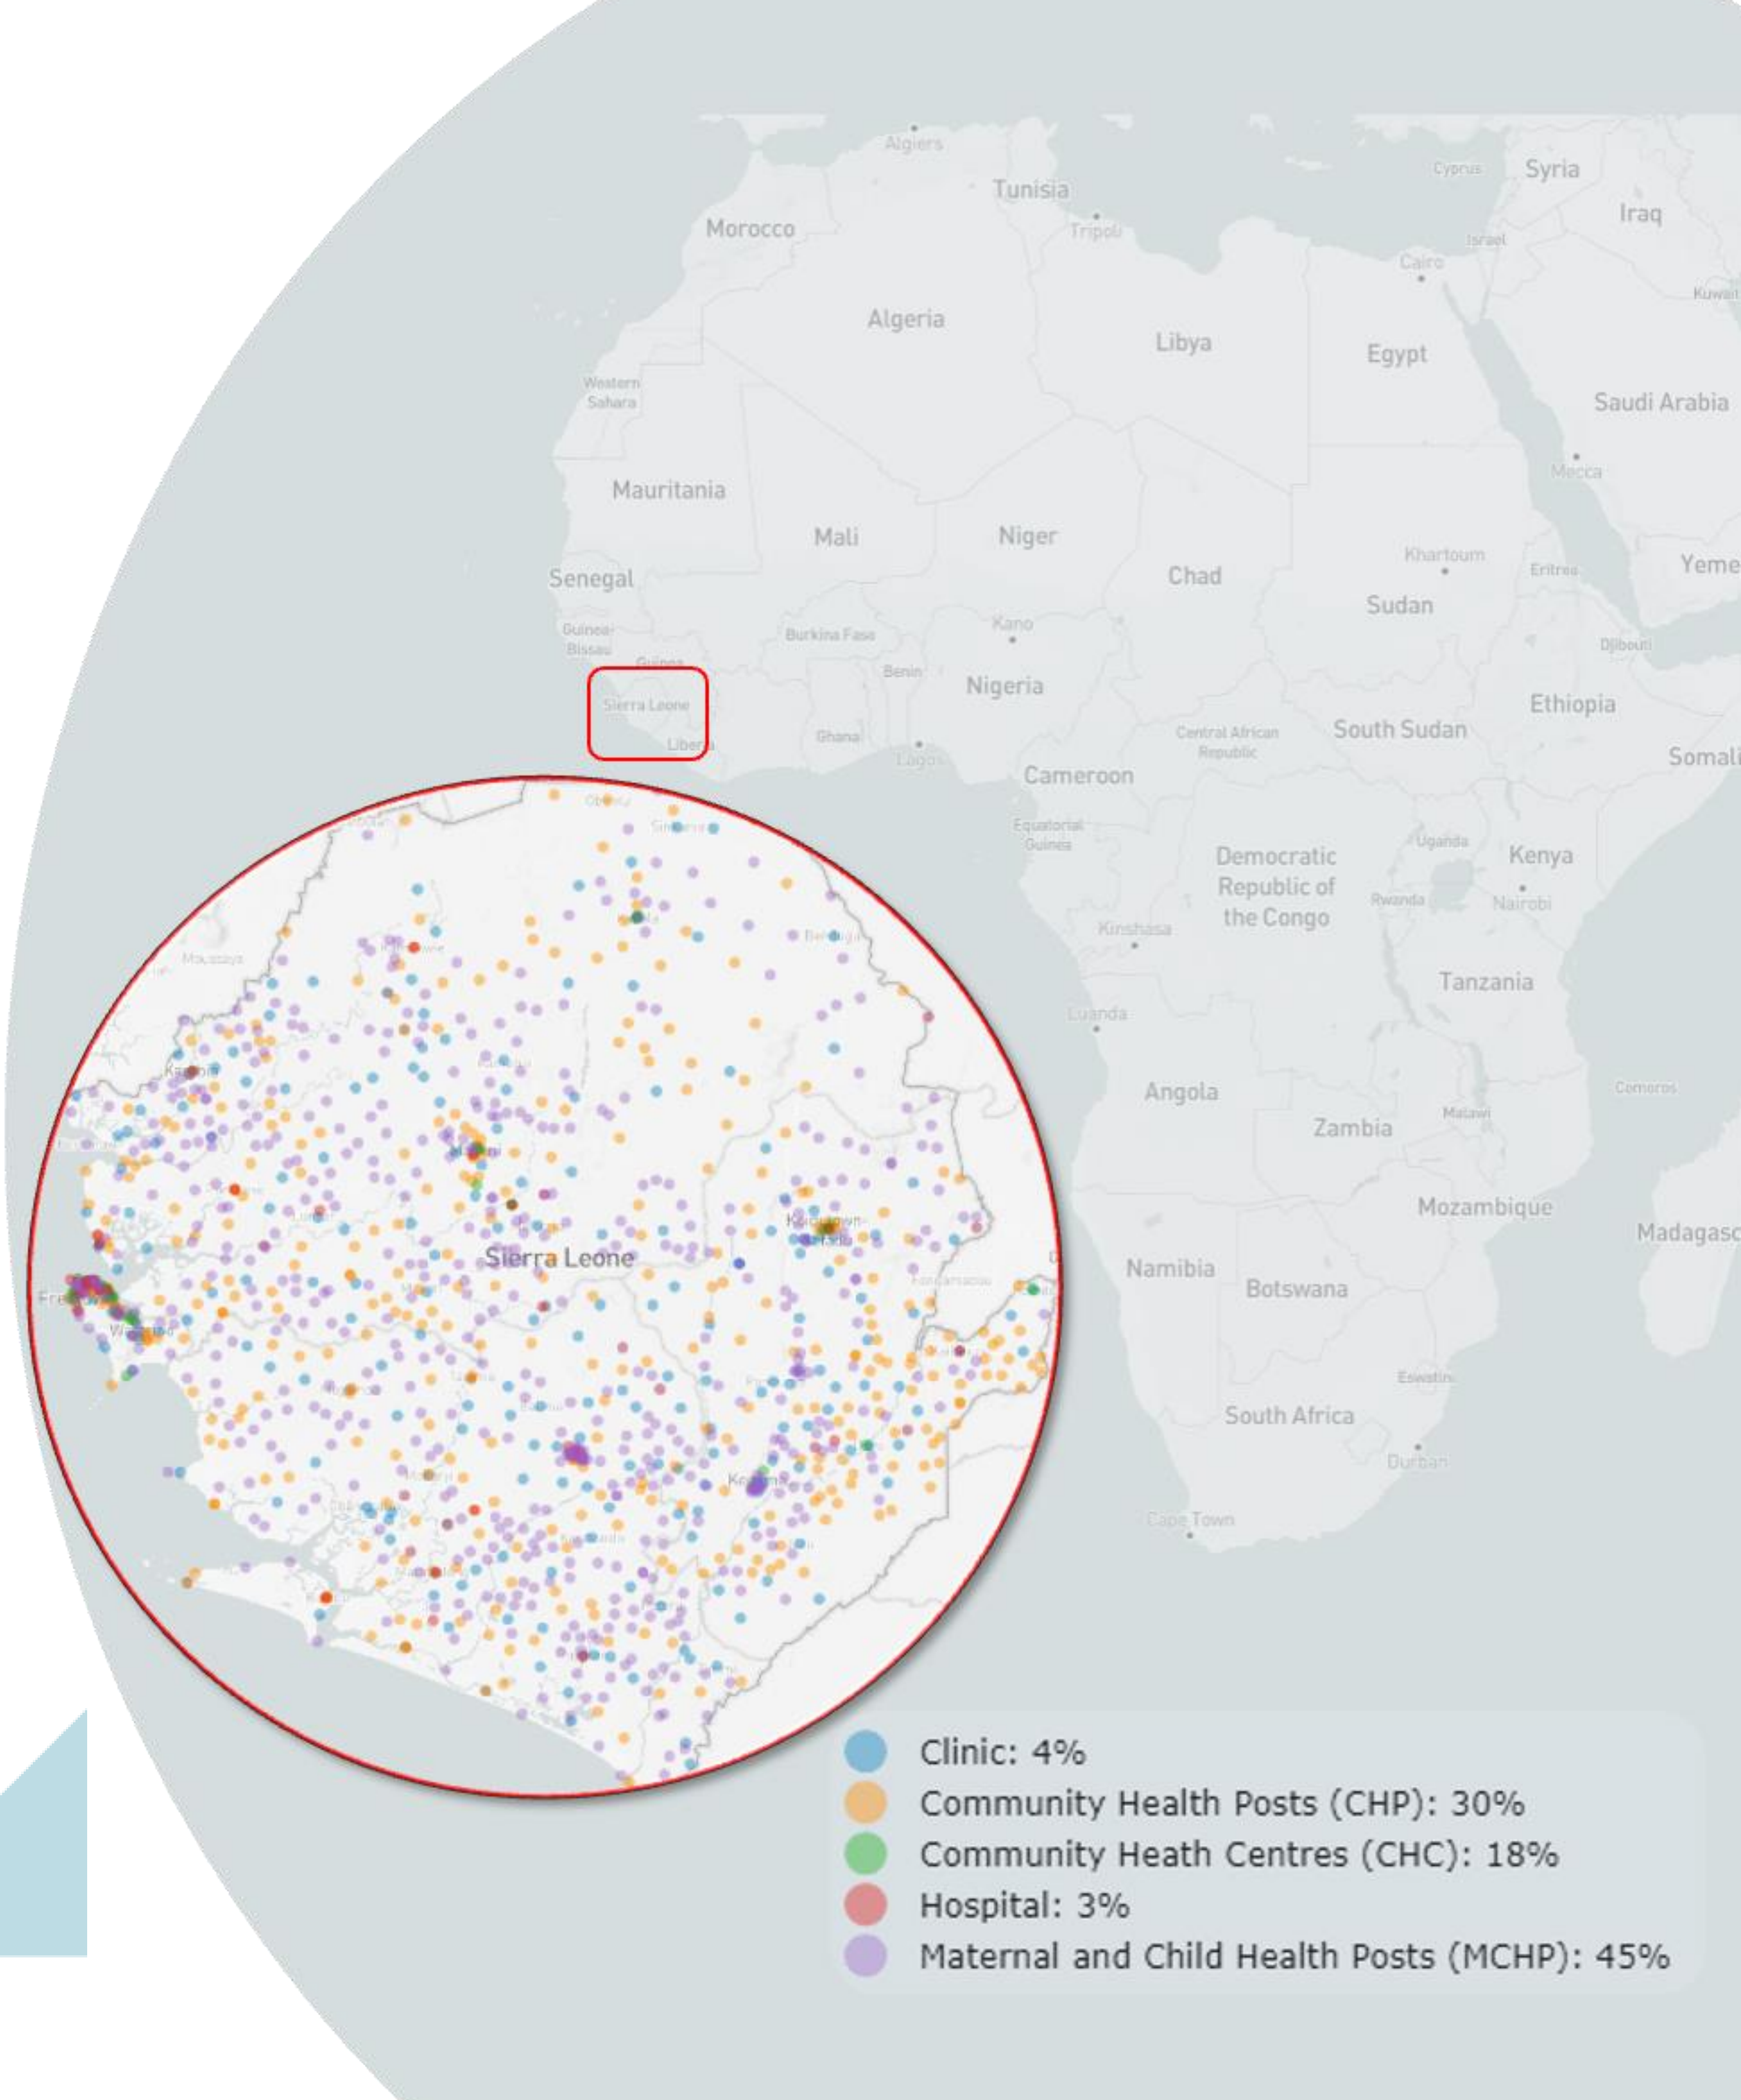

Wharton
UNIVERSITY of PENNSYLVANIA

Context in Sierra Leone

Context in Sierra Leone

- Lead time: ~3 months

Context in Sierra Leone

- Lead time: ~3 months
- Centralized push system

Context in Sierra Leone

- Lead time: ~3 months
- Centralized push system
- Allocate around 70~100 **free** healthcare products

Context in Sierra Leone

- Lead time: ~3 months
- Centralized push system
- Allocate around 70~100 **free** healthcare products
- Unstable & limited supply: 70~80% comes from donation

Context in Sierra Leone

- Lead time: ~3 months
- Centralized push system
- Allocate around 70~100 **free** healthcare products
- Unstable & limited supply: 70~80% comes from donation
- Instability of demand

Wharton
UNIVERSITY of PENNSYLVANIA

Context in Sierra Leone

- Lead time: ~3 months
- Centralized push system
- Allocate around 70~100 **free** healthcare products
- Unstable & limited supply: 70~80% comes from donation
- Instability of demand
- Limited data

Problem & Challenges

Code	Level of Care	Essential	m Type	Item	On request forms	Include in Distribution ?
10000093	ALL	X	FHC	Albendazole 400mg, Tab	Yes	Yes
10000100	Hospital & CHC		FHC	Amoxicillin & Clavulanic Acid (Co-Amoxiclav) 500mg & 125mg, Tab	Yes	Yes
10000095	ALL	X	FHC	Amoxicillin 250mg, Dispersible, Tab	Yes	Yes
10000450	Hospital & CHC	X	FHC	Ampicillin 500mg, Pdr for IM/IV, Inj, Vial	Yes	Yes
10000703	ALL		FHC	Apron, Plastic, Disposable, Pcs	Yes	Yes
10000014	Hospital & CHC		FHC	Atropine Sulphate 1mg/ml, Inj, 1ml, Amp	Yes	No
10000684	Hospital Only		FHC	Bags, Blood, Collecting, 450ml, Pcs	Yes	Yes
10000520	ALL		FHC	Cannula, IV, 18G, Short, Sterile, Disposable, Pcs	Yes	yes

Code	Total Request (Units)	DMS Request Total (Units)	District Hospital Request Total (Units)	WA Hospital Request Total (Units)	DMS Request Total	DMS - Bo	DMS - Bombali	DMS - Bonthe	DMS - Kailahun	DMS - Kambia	DMS - Kenema
10000093	1,690,967	1,384,400	243,600	62,967	1,384,400	na	83,300	-	67,100	119,000	181,500
10000100	1,871,735	1,430,620	166,300	274,815	1,430,620	na	69,360	-	25,500	77,500	168,000
10000095	7,733,420	6,077,450	1,148,436	507,534	6,077,450	na	408,000	298,000	234,000	225,000	882,000
10000450	1,472,982	839,070	403,755	230,157	839,070	na	27,000	6,020	45,000	20,700	540,800
10000703	711,933	241,000	123,633	347,300	241,000	na	70,400	-	29,100	1,000	56,400
10000014	78,678	27,110	40,851	10,717	27,110	na	1,350	-	4,500	750	-
10000684	68,300	-	58,850	9,450	-	na	na	na	na	na	na
10000520	532,114	371,300	121,450	39,364	371,300	na	3,990	17,800	25,600	6,910	7,650

Code	DMS - Koinadugu	DMS - Kono	DMS - Moyamba	DMS - Port Loko	DMS - Pujehun	DMS - Tonkolili	DMS - Western Area	DMS - Falaba	DMS - Karene	District Hospital Request To	Hosp - Bo	Hosp - Bombali Makeni
10000093	65,500	230,000	na	141,600	88,900	56,100	278,800	62,000	10,600	243,600	4,100	9,800
10000100	61,000	87,000	na	197,100	17,760	178,500	399,000	62,000	87,900	166,300	na	na
10000095	360,000	600,000	na	503,000	707,200	287,500	1,069,000	327,000	176,750	1,148,436	45,000	50,000
10000450	14,950	33,400	na	57,600	31,000	14,800	23,500	15,500	8,800	403,755	30,000	27,000
10000703	5,800	-	na	24,300	19,200	11,900	-	12,300	10,600	123,633	5,200	6,800
10000014	2,950	5,500	na	800	7,680	-	1,730	650	1,200	40,851	na	na
10000684	na	na	na	na	na	na	na	na	na	58,850	na	na
10000520	5,400	14,000	na	12,150	192,000	17,100	60,500	4,200	4,000	121,450	na	na

Problem & Challenges

Code	Level of Care	Essential	m Type	Item	On request forms	Include in Distribution ?	Total Request (Units)	DMS Request Total (Units)	District Hospital Request Total (Units)	WA Hospital Request Total (Units)
10000093	ALL	X	FHC	Albendazole 400mg, Tab	Yes	Yes	1,690,967	1,384,400	243,600	62,967
10000100	Hospital & CHC		FHC	Amoxicillin & Clavulanic Acid (Co-Amoxiclav) 500mg & 125mg, Tab	Yes	Yes	1,871,735	1,430,620	166,300	274,815
10000095	ALL	X	FHC	Amoxicillin 250mg, Dispersible, Tab	Yes	Yes	7,733,420	6,077,450	1,148,436	507,534
10000450	Hospital & CHC	X	FHC	Ampicillin 500mg, Pdr for IM/IV, Inj, Vial	Yes	Yes	1,472,982	839,070	403,755	230,157
10000703	ALL		FHC	Apron, Plastic, Disposable, Pcs	Yes	Yes	711,933	241,000	123,633	347,300
10000014	Hospital & CHC		FHC	Atropine Sulphate 1mg/ml, Inj, 1ml, Amp	Yes	No	78,678	27,110	40,851	10,717
10000684	Hospital Only		FHC	Bags, Blood, Collecting, 450ml, Pcs	Yes	Yes	68,300	-	58,850	9,450
10000520	ALL		FHC	Cannula, IV, 18G, Short, Sterile, Disposable, Pcs	Yes	yes	532,114	371,300	121,450	39,364

Code	DMS Request Total	DMS - Bo	DMS - Bombali	DMS - Bonthe	DMS - Kailahun	DMS - Kambia	DMS - Kenema	DMS - Koinadugu	DMS - Kono	DMS - Moyamba	DMS - Port Loko	DMS - Pujehun	DMS - Tonkolili	DMS - Western Area	DMS - Falaba	DMS - Karene
10000093	1,384,400	na	83,300	-	67,100	119,000	181,500	65,500	230,000	na	141,600	88,900	56,100	278,800	62,000	10,600
10000100	1,430,620	na	69,360	-	25,500	77,500	168,000	61,000	87,000	na	197,100	17,760	178,500	399,000	62,000	87,900
10000095	6,077,450	na	408,000	298,000	234,000	225,000	882,000	360,000	600,000	na	503,000	707,200	287,500	1,069,000	327,000	176,750
10000450	839,070	na	27,000	6,020	45,000	20,700	540,800	14,950	33,400	na	57,600	31,000	14,800	23,500	15,500	8,800
10000703	241,000	na	70,400	-	29,100	1,000	56,400	5,800	-	na	24,300	19,200	11,900	-	12,300	10,600
10000014	27,110	na	1,350	-	4,500	750	-	2,950	5,500	na	800	7,680	-	1,730	650	1,200
10000684	-	na	na	na	na	na	na	na	na	na	na	na	na	na	na	na
10000520	371,300	na	3,990	17,800	25,600	6,910	7,650	5,400	14,000	na	12,150	192,000	17,100	60,500	4,200	4,000

Code	District Hospital Request To	Hosp - Bo	Hosp - Bombali Makeni
10000093	243,600	4,100	9,800
10000100	166,300	na	na
10000095	1,148,436	45,000	50,000
10000450	403,755	30,000	27,000
10000703	123,633	5,200	6,800
10000014	40,851	na	na
10000684	58,850	na	na
10000520	121,450	na	na

→ 42% unfulfilled needs

Problem & Challenges

Code	Level of Care	Essential	m Type	Item	On request forms	Include in Distribution?	Total Request (Units)	DMS Request Total (Units)	District Hospital Request Total (Units)	WA Hospital Request Total (Units)
10000093	ALL	X	FHC	Albendazole 400mg, Tab	Yes	Yes	1,690,967	1,384,400	243,600	62,967
10000100	Hospital & CHC		FHC	Amoxicillin & Clavulanic Acid (Co-Amoxiclav) 500mg & 125mg, Tab	Yes	Yes	1,871,735	1,430,620	166,300	274,815
10000095	ALL	X	FHC	Amoxicillin 250mg, Dispersible, Tab	Yes	Yes	7,733,420	6,077,450	1,148,436	507,534
10000450	Hospital & CHC	X	FHC	Ampicillin 500mg, Pdr for IM/IV, Inj, Vial	Yes	Yes	1,472,982	839,070	403,755	230,157
10000703	ALL		FHC	Apron, Plastic, Disposable, Pcs	Yes	Yes	711,933	241,000	123,633	347,300
10000014	Hospital & CHC		FHC	Atropine Sulphate 1mg/ml, Inj, 1ml, Amp	Yes	No	78,678	27,110	40,851	10,717
10000684	Hospital Only		FHC	Bags, Blood, Collecting, 450ml, Pcs	Yes	Yes	68,300	-	58,850	9,450
10000520	ALL		FHC	Cannula, IV, 18G, Short, Sterile, Disposable, Pcs	Yes	yes	532,114	371,300	121,450	39,364

Code	DMS Request Total	DMS - Bo	DMS - Bombali	DMS - Bonthe	DMS - Kailahun	DMS - Kambia	DMS - Kenema	DMS - Koinadugu	DMS - Kono	DMS - Moyamba
10000093	1,384,400	na	83,300	-	67,100	119,000	181,500	65,500	230,000	na
10000100	1,430,620	na	69,360	-	25,500	77,500	168,000	61,000	87,000	na
10000095	6,077,450	na	408,000	298,000	234,000	225,000	882,000	360,000	600,000	na
10000450	839,070	na	27,000	6,020	45,000	20,700	540,800	14,950	33,400	na
10000703	241,000	na	70,400	-	29,100	1,000	56,400	5,800	-	na
10000014	27,110	na	1,350	-	4,500	750	-	2,950	5,500	na
10000684	-	na	na	na	na	na	na	na	na	na
10000520	371,300	na	3,990	17,800	25,600	6,910	7,650	5,400	14,000	na

Code	DMS - Port Loko	DMS - Pujehun	DMS - Tonkolili	DMS - Western Area	DMS - Falaba	DMS - Karene	District Hospital Request To	Hosp - Bo	Hosp - Bombali Makeni
10000093	141,600	88,900	56,100	278,800	62,000	10,600	243,600	4,100	9,800
10000100	197,100	17,760	178,500	399,000	62,000	87,900	166,300	na	na
10000095	503,000	707,200	287,500	1,069,000	327,000	176,750	1,148,436	45,000	50,000
10000450	57,600	31,000	14,800	23,500	15,500	8,800	403,755	30,000	27,000
10000703	24,300	19,200	11,900	-	12,300	10,600	123,633	5,200	6,800
10000014	800	7,680	-	1,730	650	1,200	40,851	na	na
10000684	na	na	na	na	na	na	58,850	na	na
10000520	12,150	192,000	17,100	60,500	4,200	4,000	121,450	na	na

→ 42% unfulfilled needs

- **Problem:** one supplier multilocation problem [Clark and Scarf, 1960; Federgruen and Zipkin, 1984; Chen and Zheng, 1994; Cachon, 1999…]

Wharton
UNIVERSITY of PENNSYLVANIA

Problem & Challenges

Code	Level of Care	Essential	m Type	Item	On request forms	Include in Distribution ?
10000093	ALL	X	FHC	Albendazole 400mg, Tab	Yes	Yes
10000100	Hospital & CHC		FHC	Amoxicillin & Clavulanic Acid (Co-Amoxiclav) 500mg & 125mg, Tab	Yes	Yes
10000095	ALL	X	FHC	Amoxicillin 250mg, Dispersible, Tab	Yes	Yes
10000450	Hospital & CHC	X	FHC	Ampicillin 500mg, Pdr for IM/IV, Inj, Vial	Yes	Yes
10000703	ALL		FHC	Apron, Plastic, Disposable, Pcs	Yes	Yes
10000014	Hospital & CHC		FHC	Atropine Sulphate 1mg/ml, Inj, 1ml, Amp	Yes	No
10000684	Hospital Only		FHC	Bags, Blood, Collecting, 450ml, Pcs	Yes	Yes
10000520	ALL		FHC	Cannula, IV, 18G, Short, Sterile, Disposable, Pcs	Yes	yes

Code	Total Request (Units)	DMS Request Total (Units)	District Hospital Request Total (Units)	WA Hospital Request Total (Units)
10000093	1,690,967	1,384,400	243,600	62,967
10000100	1,871,735	1,430,620	166,300	274,815
10000095	7,733,420	6,077,450	1,148,436	507,534
10000450	1,472,982	839,070	403,755	230,157
10000703	711,933	241,000	123,633	347,300
10000014	78,678	27,110	40,851	10,717
10000684	68,300	-	58,850	9,450
10000520	532,114	371,300	121,450	39,364

Code	DMS Request Total	DMS - Bo	DMS - Bombali	DMS - Bonthe	DMS - Kailahun	DMS - Kambia	DMS - Kenema	DMS - Koinadugu	DMS - Kono	DMS - Moyamba
10000093	1,384,400	na	83,300	-	67,100	119,000	181,500	65,500	230,000	na
10000100	1,430,620	na	69,360	-	25,500	77,500	168,000	61,000	87,000	na
10000095	6,077,450	na	408,000	298,000	234,000	225,000	882,000	360,000	600,000	na
10000450	839,070	na	27,000	6,020	45,000	20,700	540,800	14,950	33,400	na
10000703	241,000	na	70,400	-	29,100	1,000	56,400	5,800	-	na
10000014	27,110	na	1,350	-	4,500	750	-	2,950	5,500	na
10000684	-	na	na	na	na	na	na	na	na	na
10000520	371,300	na	3,990	17,800	25,600	6,910	7,650	5,400	14,000	na

Code	DMS - Port Loko	DMS - Pujehun	DMS - Tonkolili	DMS - Western Area	DMS - Falaba	DMS - Karene
10000093	141,600	88,900	56,100	278,800	62,000	10,600
10000100	197,100	17,760	178,500	399,000	62,000	87,900
10000095	503,000	707,200	287,500	1,069,000	327,000	176,750
10000450	57,600	31,000	14,800	23,500	15,500	8,800
10000703	24,300	19,200	11,900	-	12,300	10,600
10000014	800	7,680	-	1,730	650	1,200
10000684	na	na	na	na	na	na
10000520	12,150	192,000	17,100	60,500	4,200	4,000

Code	District Hospital Request To	Hosp - Bo	Hosp - Bombali Makeni
10000093	243,600	4,100	9,800
10000100	166,300	na	na
10000095	1,148,436	45,000	50,000
10000450	403,755	30,000	27,000
10000703	123,633	5,200	6,800
10000014	40,851	na	na
10000684	58,850	na	na
10000520	121,450	na	na

→ 42% unfulfilled needs

- **Problem:** one supplier multilocation problem [Clark and Scarf, 1960; Federgruen and Zipkin, 1984; Chen and Zheng, 1994; Cachon, 1999…]
- **Challenges**
  - Demand prediction with **small data** and **unstable demand**

Wharton
University of Pennsylvania

Problem & Challenges

Code	Level of Care	Essential	m Type	Item	On request forms	Include in Distribution ?	Total Request (Units)	DMS Request Total (Units)	District Hospital Request Total (Units)	WA Hospital Request Total (Units)
10000093	ALL	X	FHC	Albendazole 400mg, Tab	Yes	Yes	1,690,967	1,384,400	243,600	62,967
10000100	Hospital & CHC		FHC	Amoxicillin & Clavulanic Acid (Co-Amoxiclav) 500mg & 125mg, Tab	Yes	Yes	1,871,735	1,430,620	166,300	274,815
10000095	ALL	X	FHC	Amoxicillin 250mg, Dispersible, Tab	Yes	Yes	7,733,420	6,077,450	1,148,436	507,534
10000450	Hospital & CHC	X	FHC	Ampicillin 500mg, Pdr for IM/IV, Inj, Vial	Yes	Yes	1,472,982	839,070	403,755	230,157
10000703	ALL		FHC	Apron, Plastic, Disposable, Pcs	Yes	Yes	711,933	241,000	123,633	347,300
10000014	Hospital & CHC		FHC	Atropine Sulphate 1mg/ml, Inj, 1ml, Amp	Yes	No	78,678	27,110	40,851	10,717
10000684	Hospital Only		FHC	Bags, Blood, Collecting, 450ml, Pcs	Yes	Yes	68,300	-	58,850	9,450
10000520	ALL		FHC	Cannula, IV, 18G, Short, Sterile, Disposable, Pcs	Yes	yes	532,114	371,300	121,450	39,364

Code	DMS Request Total	DMS - Bo	DMS - Bombali	DMS - Bonthe	DMS - Kailahun	DMS - Kambia	DMS - Kenema	DMS - Koinadugu	DMS - Kono	DMS - Moyamba
10000093	1,384,400	na	83,300	-	67,100	119,000	181,500	65,500	230,000	na
10000100	1,430,620	na	69,360	-	25,500	77,500	168,000	61,000	87,000	na
10000095	6,077,450	na	408,000	298,000	234,000	225,000	882,000	360,000	600,000	na
10000450	839,070	na	27,000	6,020	45,000	20,700	540,800	14,950	33,400	na
10000703	241,000	na	70,400	-	29,100	1,000	56,400	5,800	-	na
10000014	27,110	na	1,350	-	4,500	750	-	2,950	5,500	na
10000684	-	na	na	na	na	na	na	na	na	na
10000520	371,300	na	3,990	17,800	25,600	6,910	7,650	5,400	14,000	na

Code	DMS - Port Loko	DMS - Pujehun	DMS - Tonkolili	DMS - Western Area	DMS - Falaba	DMS - Karene	District Hospital Request To	Hosp - Bo	Hosp - Bombali Makeni
10000093	141,600	88,900	56,100	278,800	62,000	10,600	243,600	4,100	9,800
10000100	197,100	17,760	178,500	399,000	62,000	87,900	166,300	na	na
10000095	503,000	707,200	287,500	1,069,000	327,000	176,750	1,148,436	45,000	50,000
10000450	57,600	31,000	14,800	23,500	15,500	8,800	403,755	30,000	27,000
10000703	24,300	19,200	11,900	-	12,300	10,600	123,633	5,200	6,800
10000014	800	7,680	-	1,730	650	1,200	40,851	na	na
10000684	na	na	na	na	na	na	58,850	na	na
10000520	12,150	192,000	17,100	60,500	4,200	4,000	121,450	na	na

→ 42% unfulfilled needs

- **Problem:** one supplier multilocation problem [Clark and Scarf, 1960; Federgruen and Zipkin, 1984; Chen and Zheng, 1994; Cachon, 1999…]
- **Challenges**
  - Demand prediction with **small data** and **unstable demand**
  - Mis-alignment of prediction and optimization objective

High level preview

- **Solution:** a scalable decision-aware learning framework.

High level preview

- **Solution:** a scalable decision-aware learning framework.
- **Nationwide deployment**: 1,058 facilities (all public-owned facilities)

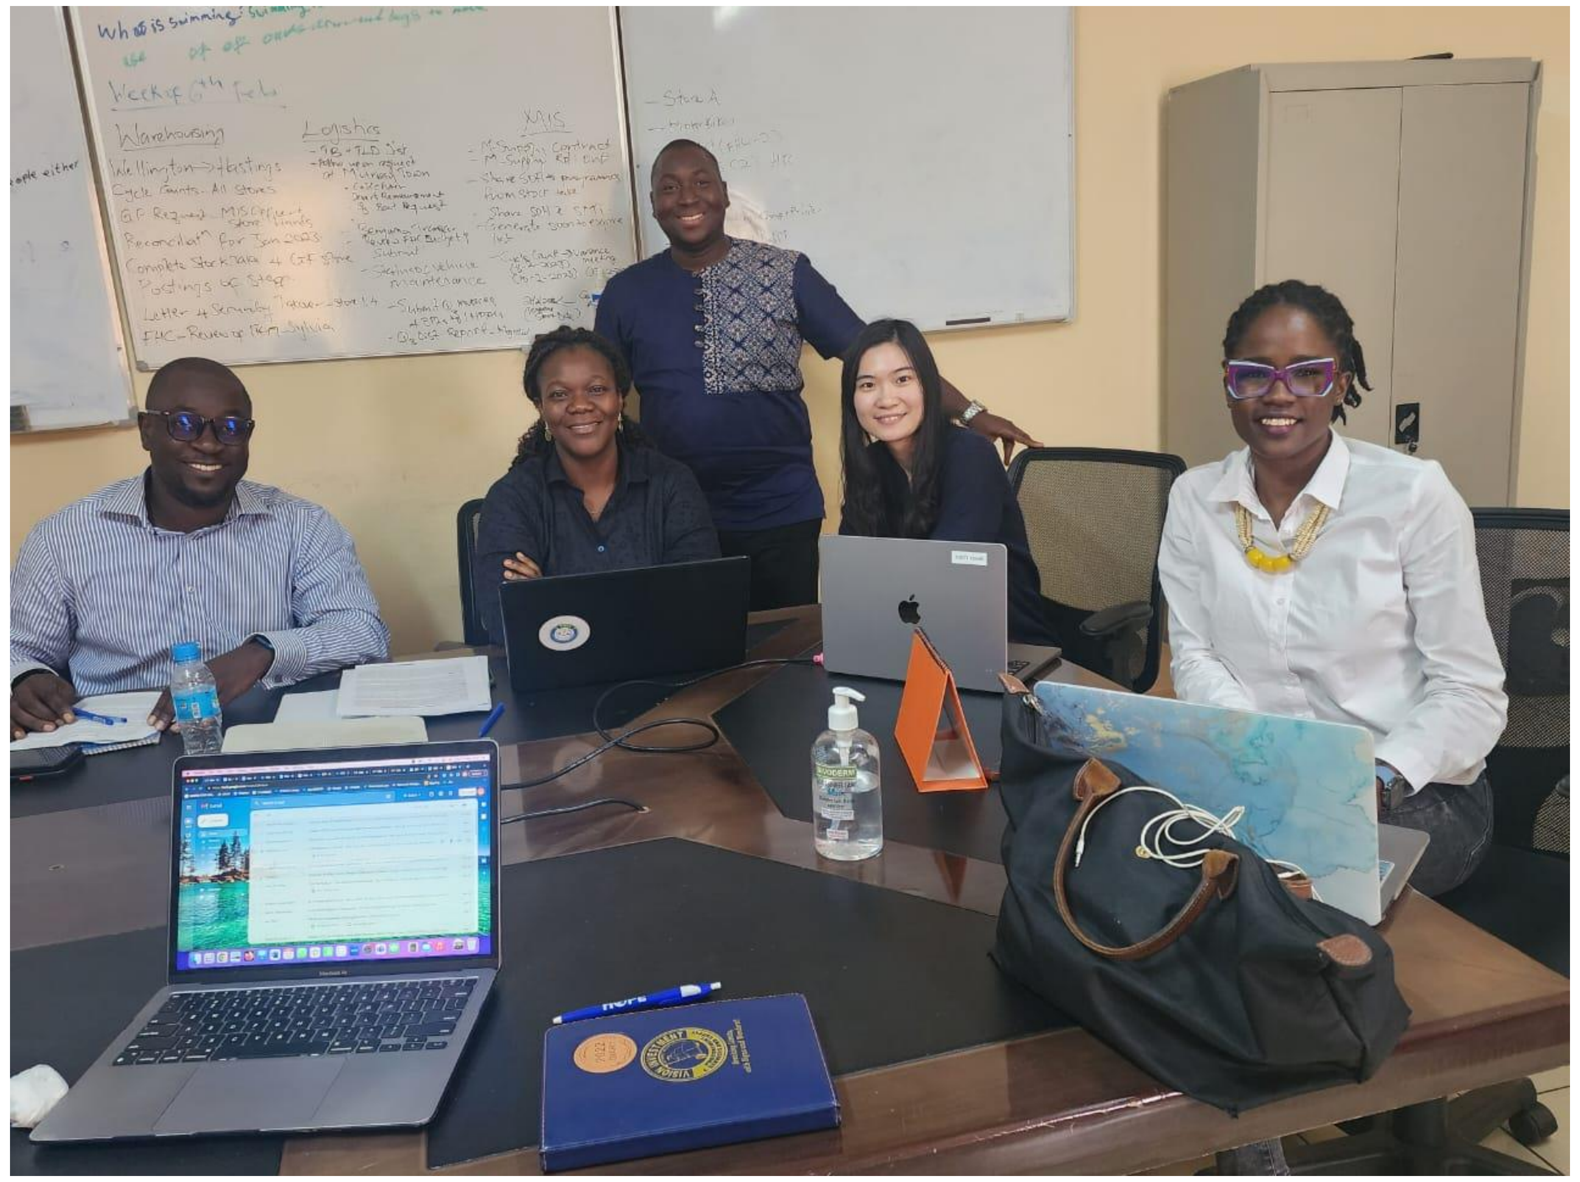

Visit & discuss at Sierra Leone National Medical Supplies Agency's Office

High level preview

- **Solution:** a scalable decision-aware learning framework.
- **Nationwide deployment**: 1,058 facilities (all public-owned facilities)
- **Results**: **19%** increase of consumption to medicines for 2 million women & children under 5

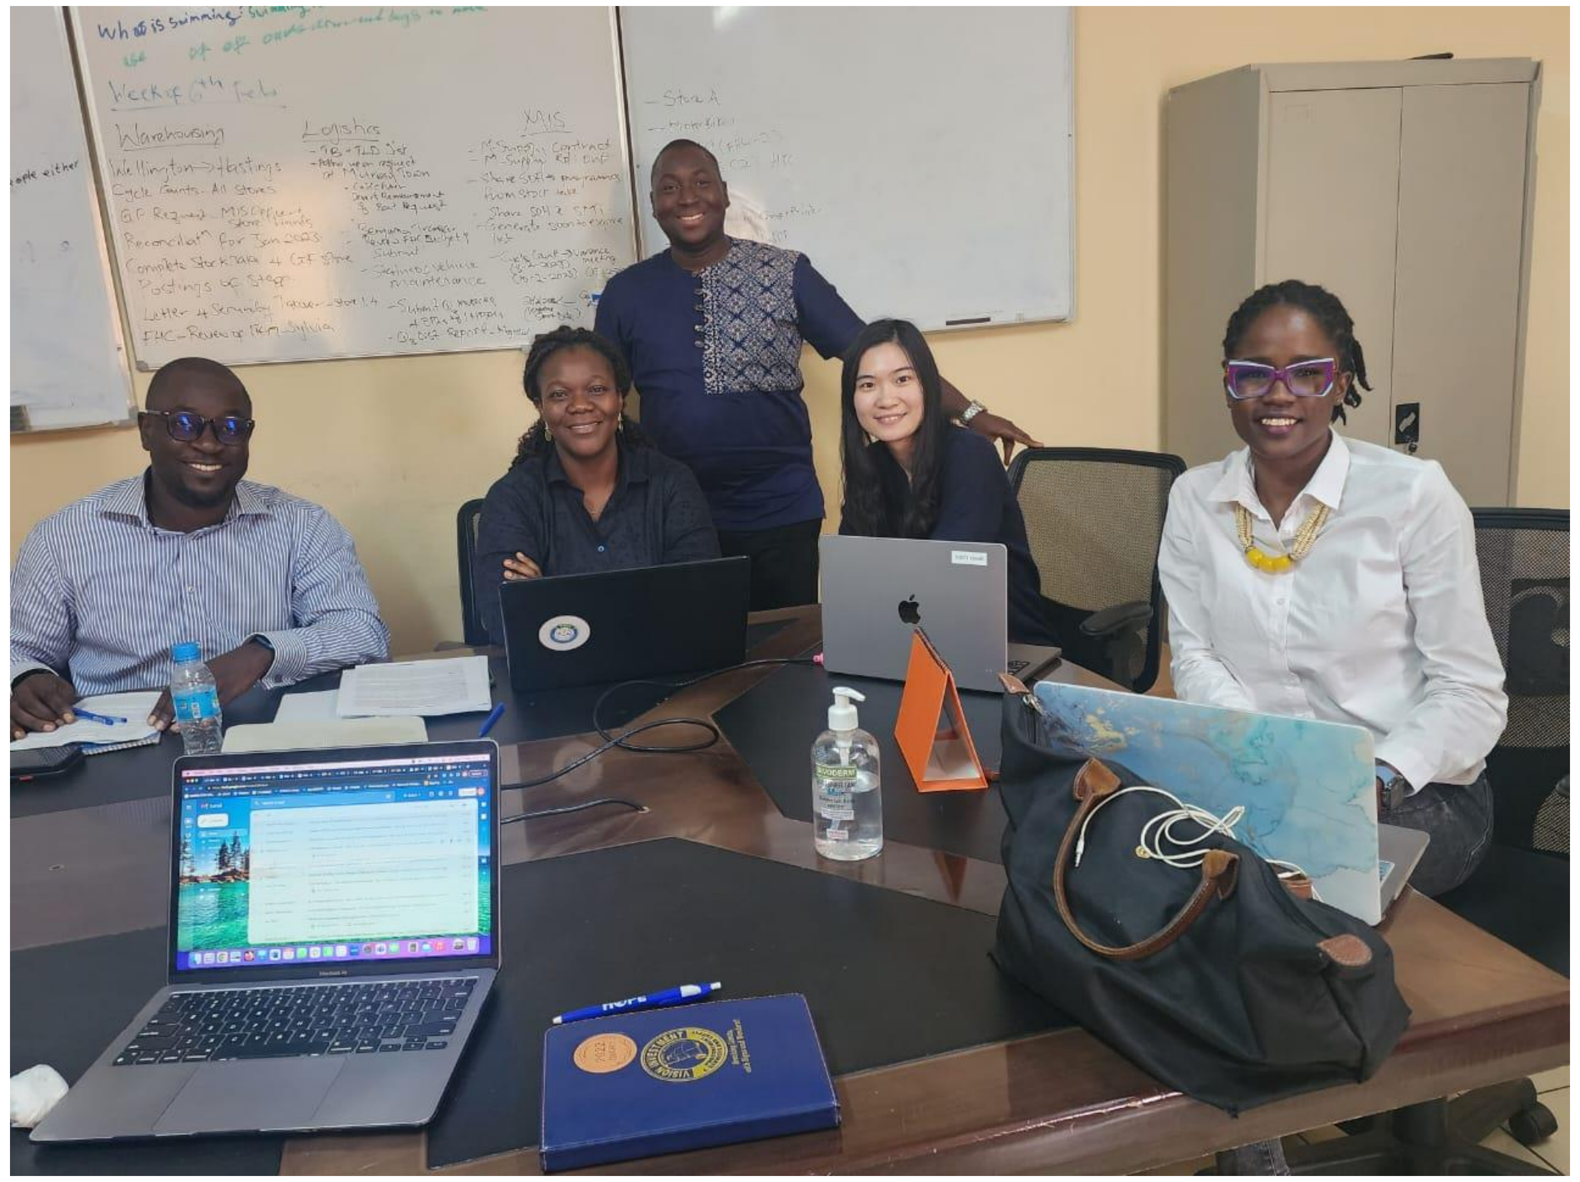

Visit & discuss at Sierra Leone National Medical Supplies Agency's Office

Predict-then-Optimize? (decision-blind)

- Combine ML + OR: from predictive to prescriptive analytics [Bertsimas and Kallus, 2019]

Predict-then-Optimize? (decision-blind)

- Combine ML + OR: from predictive to prescriptive analytics [Bertsimas and Kallus, 2019]

Step 1: Train demand prediction model for every product across all facilities for every quarters

Predict-then-Optimize? (decision-blind)

- Combine ML + OR: from predictive to prescriptive analytics [Bertsimas and Kallus, 2019]

Step 1: Train demand prediction model for every product across all facilities for every quarters

Step 2: Optimize based on predictions

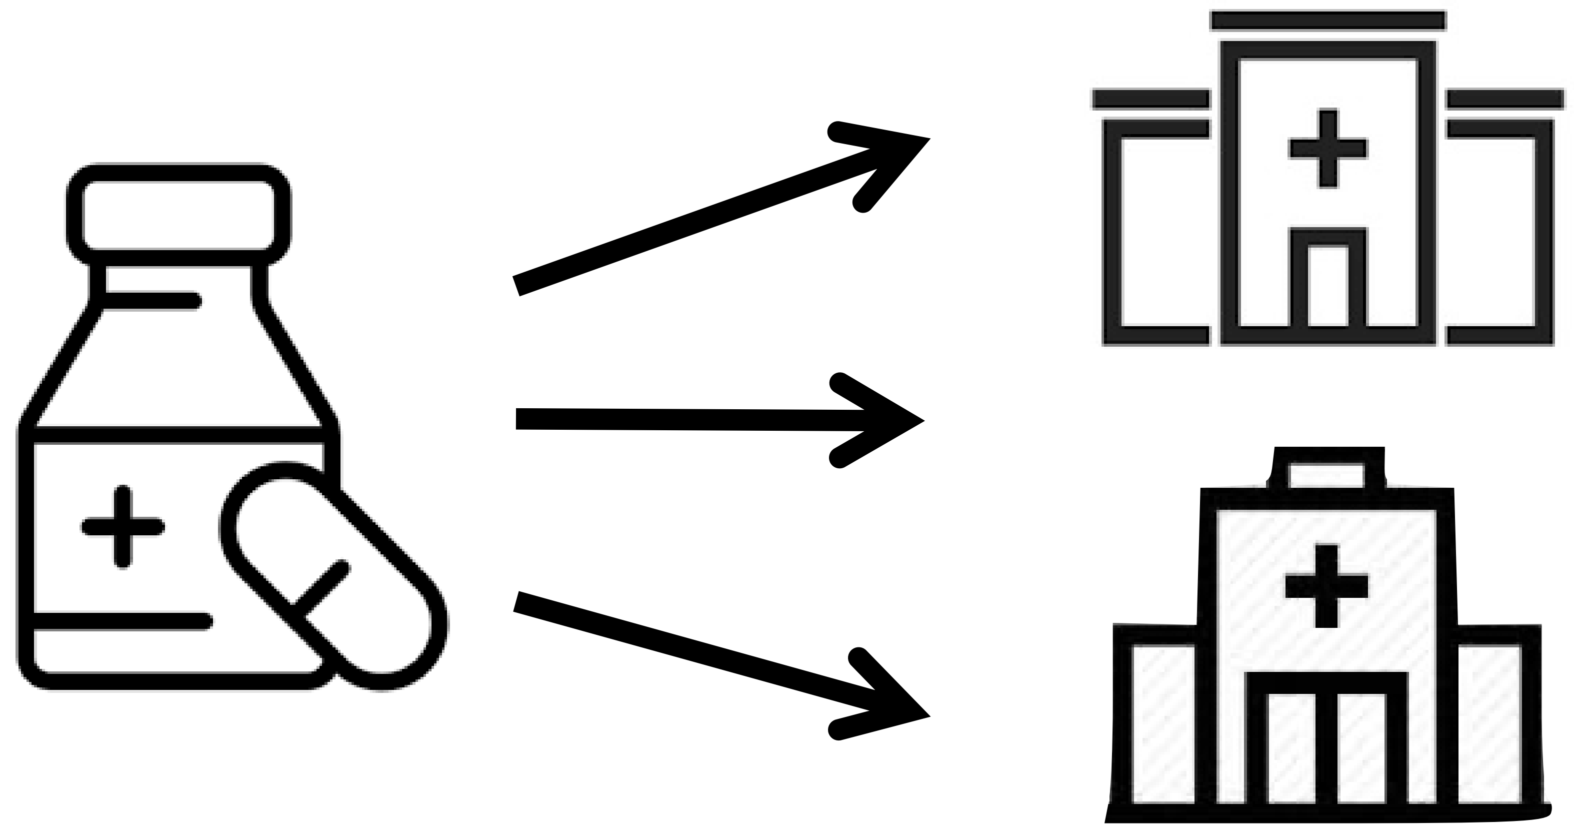

Predict-then-Optimize? (decision-blind)

- Combine ML + OR: from predictive to prescriptive analytics [Bertsimas and Kallus, 2019]

Step 1: Train demand prediction model for every product across all facilities for every quarters

Step 2: Optimize based on predictions

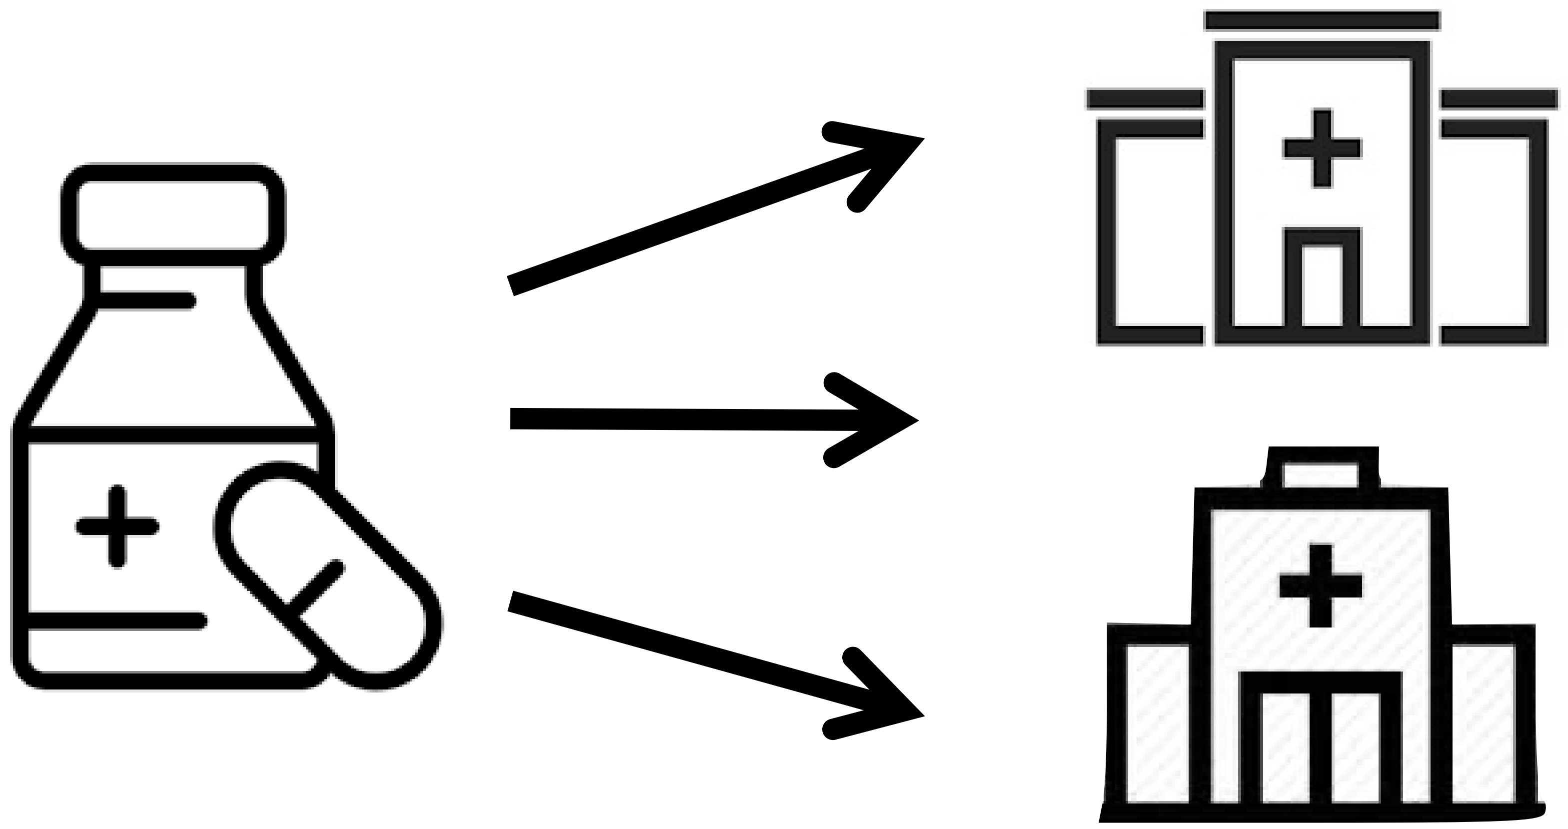

Problem:

- The prediction model did not account for the downstream optimization problem (primary interest)

Learning and optimization

- Objective:

known decision loss
(unmet demand)

$$\underbrace{a^*}_{\text{optimal allocation}}(\underbrace{y^*}_{\text{demand}}) = \operatorname*{argmin}_a \ell(a; y^*)$$

Wharton UNIVERSITY of PENNSYLVANIA

Learning and optimization

- Objective:

known decision loss
(unmet demand)

$$a^*(y^*) = \operatorname*{argmin}_a \ell(a; y^*)$$

optimal allocation

demand (unknown)

- Problem: y^* are unknown

Learning and optimization

- Objective:

known decision loss (unmet demand)

$$a^*(y^*) = \underset{a}{\operatorname{argmin}}\, \ell(a; y^*)$$

optimal allocation

demand (unknown)

- Problem: y^* are unknown

Step 1 Predict

$$\hat{y} = f_{\hat{\theta}}(x)$$

Historical data

Learning and optimization

- Objective:

$$\underbrace{a^*}_{\text{optimal allocation}}(\underbrace{y^*}_{\text{demand (unknown)}}) = \underset{a}{\operatorname{argmin}}\, \underbrace{\ell(a; y^*)}_{\text{known decision loss (unmet demand)}}$$

- Problem: y^* are unknown

Step 1 Predict

$$\hat{y} = f_{\hat{\theta}}(\underbrace{x}_{\text{Historical data}})$$

Step 2 Optimize

$$a^*(\hat{y}) = \underset{a}{\operatorname{argmin}}\, \ell(a; \hat{y})$$

Wharton UNIVERSITY of PENNSYLVANIA

Learning and optimization

- Objective:

$$\underbrace{a^*}_{\text{optimal allocation}}(\underbrace{y^*}_{\text{demand (unknown)}}) = \underset{a}{\operatorname{argmin}}\, \ell(a; y^*)$$

known decision loss
(unmet demand)

- Problem: y^* are unknown

Step 1 Predict

$$\hat{y} = f_{\hat{\theta}}(x)$$

Historical data

$$\hat{\theta} = \arg\min_{\theta} \sum_i \tilde{\ell}(f_\theta(x_i); y_i^*)$$

Prediction loss

Step 2 Optimize

$$a^*(\hat{y}) = \underset{a}{\operatorname{argmin}}\, \ell(a; \hat{y})$$

Learning and optimization

Step 1 Predict

$\hat{y} = f_{\hat{\theta}}(x)$

Historical data

$\hat{\theta} = \arg\min_{\theta} \sum_{i} \tilde{\ell}(f_{\theta}(x_i); y_i^*)$

Prediction loss

Step 2 Optimize

known decision loss
(unmet demand)

$a^*(\hat{y}) = \operatorname{argmin}_{a} \ell(a; \hat{y})$

Key question:
What prediction loss $\tilde{\ell}(\hat{y}; y^*)$ to use in training?

Learning and optimization

Step 1 Predict

$$\hat{y} = f_{\hat{\theta}}(x)$$

Historical data

$$\hat{\theta} = \arg\min_{\theta} \sum_{i} \tilde{\ell}(f_{\theta}(x_i); y_i^*)$$

Prediction loss

Step 2 Optimize

known decision loss (unmet demand)

$$a^*(\hat{y}) = \underset{a}{\operatorname{argmin}}\, \ell(a; \hat{y})$$

Key question:
What prediction loss $\tilde{\ell}(\hat{y}; y^*)$ to use in training?

- **Common approach:** Use standard loss such as mean-squared error (MSE)

$$\tilde{\ell}_{\mathrm{MSE}}(\hat{y}; y^*) = (\hat{y} - y^*)^2$$

Learning and optimization

Step 1 Predict

$$\hat{y} = f_{\hat{\theta}}(x)$$

Historical data

$$\hat{\theta} = \arg\min_{\theta} \sum_i \tilde{\ell}(f_\theta(x_i); y_i^*)$$

Prediction loss

Step 2 Optimize

known decision loss (unmet demand)

$$a^*(\hat{y}) = \operatorname*{argmin}_a \ell(a; \hat{y})$$

Key question:
What prediction loss $\tilde{\ell}(\hat{y}; y^*)$ to use in training?

- **Common approach:** Use standard loss such as mean-squared error (MSE)

$$\tilde{\ell}_{\mathrm{MSE}}(\hat{y}; y^*) = (\hat{y} - y^*)^2$$

- **Problem:** MSE loss does not account for downstream objective!

Learning and optimization

Step 1 Predict

$\hat{y} = f_{\hat{\theta}}(x)$

Historical data

$\hat{\theta} = \arg\min_{\theta} \sum_i \tilde{\ell}(f_\theta(x_i); y_i^*)$

Prediction loss

Step 2 Optimize

known decision loss
(unmet demand)

$a^*(\hat{y}) = \operatorname{argmin}_a \ell(a; \hat{y})$

Key question:
What prediction loss $\tilde{\ell}(\hat{y}; y^*)$ to use in training?

- **Common approach:** Use standard loss such as mean-squared error (MSE)

$$\tilde{\ell}_{\mathrm{MSE}}(\hat{y}; y^*) = (\hat{y} - y^*)^2$$

- **Problem:** MSE loss does not account for downstream objective!
- **Ideal: Use downstream decision loss**

$$\tilde{\ell}(\hat{y}; y^*) := \ell(a^*(\hat{y}); y^*) = \ell(\arg\min_a \ell(a; \hat{y}); y^*)$$

Our strategy – Taylor expansion to derive weights

- Taylor expand y around $\hat{y}$:

$\hat{y} = f_{\hat{\theta}}(x)$

$$\ell(a^*(\hat{y}); y^*) \approx \ell(a^*(\hat{y}); y^*) + \underbrace{\nabla_a \ell(a^*(\hat{y}); y^*)^\top \nabla_y a^*(\hat{y})}_{\text{importance of improving } \hat{y} \text{ on improving the loss}} (y - \hat{y})$$

- Can compute gradient through OPT objective and OPT decision
  - Can be computed numerically efficiently for general class of convex programs [Agrawal-Amos-Barratt-Boyd-Diamond-Kolter, 2018]

Decision-Aware training objective

- Use gradients to obtain approximated predictive model objective:

$$\arg\min_{\theta} \frac{1}{K} \sum_{k=1}^{K} \sum_{n=1}^{N} \Big(\mathbb{I}\big(\underbrace{\xi_n^{(k)}}_{\text{demand}} \geq \overset{\text{inventory}}{s_n} + \overset{\text{allocation}}{a_n} \big) + const \Big) \cdot \left| f_\theta(x_n) - \xi_n^{(k)} \right|$$

- Up-weight training examples with unmet demand: instead of treating all observations equally, the model focus on facilities which are more relevant to our decision loss (minimizing unmet demand)

Our Decision-Aware approach

- **Step 1:** Train demand prediction model $f_{\hat{\theta}}(x)$

$$\Downarrow$$

- **Step 2:** run optimization to obtain the weight

Our Decision-Aware approach

- **Step 1:** Train demand prediction model $f_{\hat{\theta}}(x)$

$\Downarrow$

- **Step 2:** run optimization to obtain the weight

$\Downarrow$

- **Step 3:** re-train model with weighted training data (up-weight unmet demand)

Our Decision-Aware approach

- **Step 1:** Train demand prediction model $f_{\hat{\theta}}(x)$

⇩

- **Step 2:** run optimization to obtain the weight

⇩

- **Step 3:** re-train model with weighted training data (up-weight unmet demand)

⇩

- **Step 4:** run optimization problem to get decision-aware result

Our Decision-Aware approach

- **Step 1:** Train demand prediction model $f_{\hat{\theta}}(x)$

$$\Downarrow$$

- **Step 2:** run optimization to obtain the weight

$$\Downarrow$$

- **Step 3:** re-train model with weighted training data (up-weight unmet demand)

$$\Downarrow$$

- **Step 4:** run optimization problem to get decision-aware result

→ ***Computationally tractable & light touch***

Data

- **Dhis2** forms (District Health Information Software) from Feb 2020 to Jan 2023
  - Consumption, # received, opening/closing balance, stockout (binary)…

Data

- **Dhis2** forms (District Health Information Software) from Feb 2020 to Jan 2023
  - Consumption, # received, opening/closing balance, stockout (binary)…

- **Feature engineering**
  - Lagged consumptions for each product in that facility for last 10 months
  - Month, year fixed effects
  - Rolling average of last 2/3/4/5/6/8/10 months + variance of last 3/6 months
  - Facility region & type

Data

- **Dhis2** forms (District Health Information Software) from Feb 2020 to Jan 2023
  - Consumption, # received, opening/closing balance, stockout (binary)…

- **Feature engineering**
  - Lagged consumptions for each product in that facility for last 10 months
  - Month, year fixed effects
  - Rolling average of last 2/3/4/5/6/8/10 months + variance of last 3/6 months
  - Facility region & type

- Use **multi-task learning** to train the prediction model using random forest (RF)
  - RF outperform historical average, ARIMA, NN, LASSO regression

MSE comparison | Multi-task learning

Wharton
University *of* Pennsylvania

Deployment

- 2023 Q2: 5 out of 16 districts
  - Picked randomly by the government
    - 36 treated products
    - 24 control products

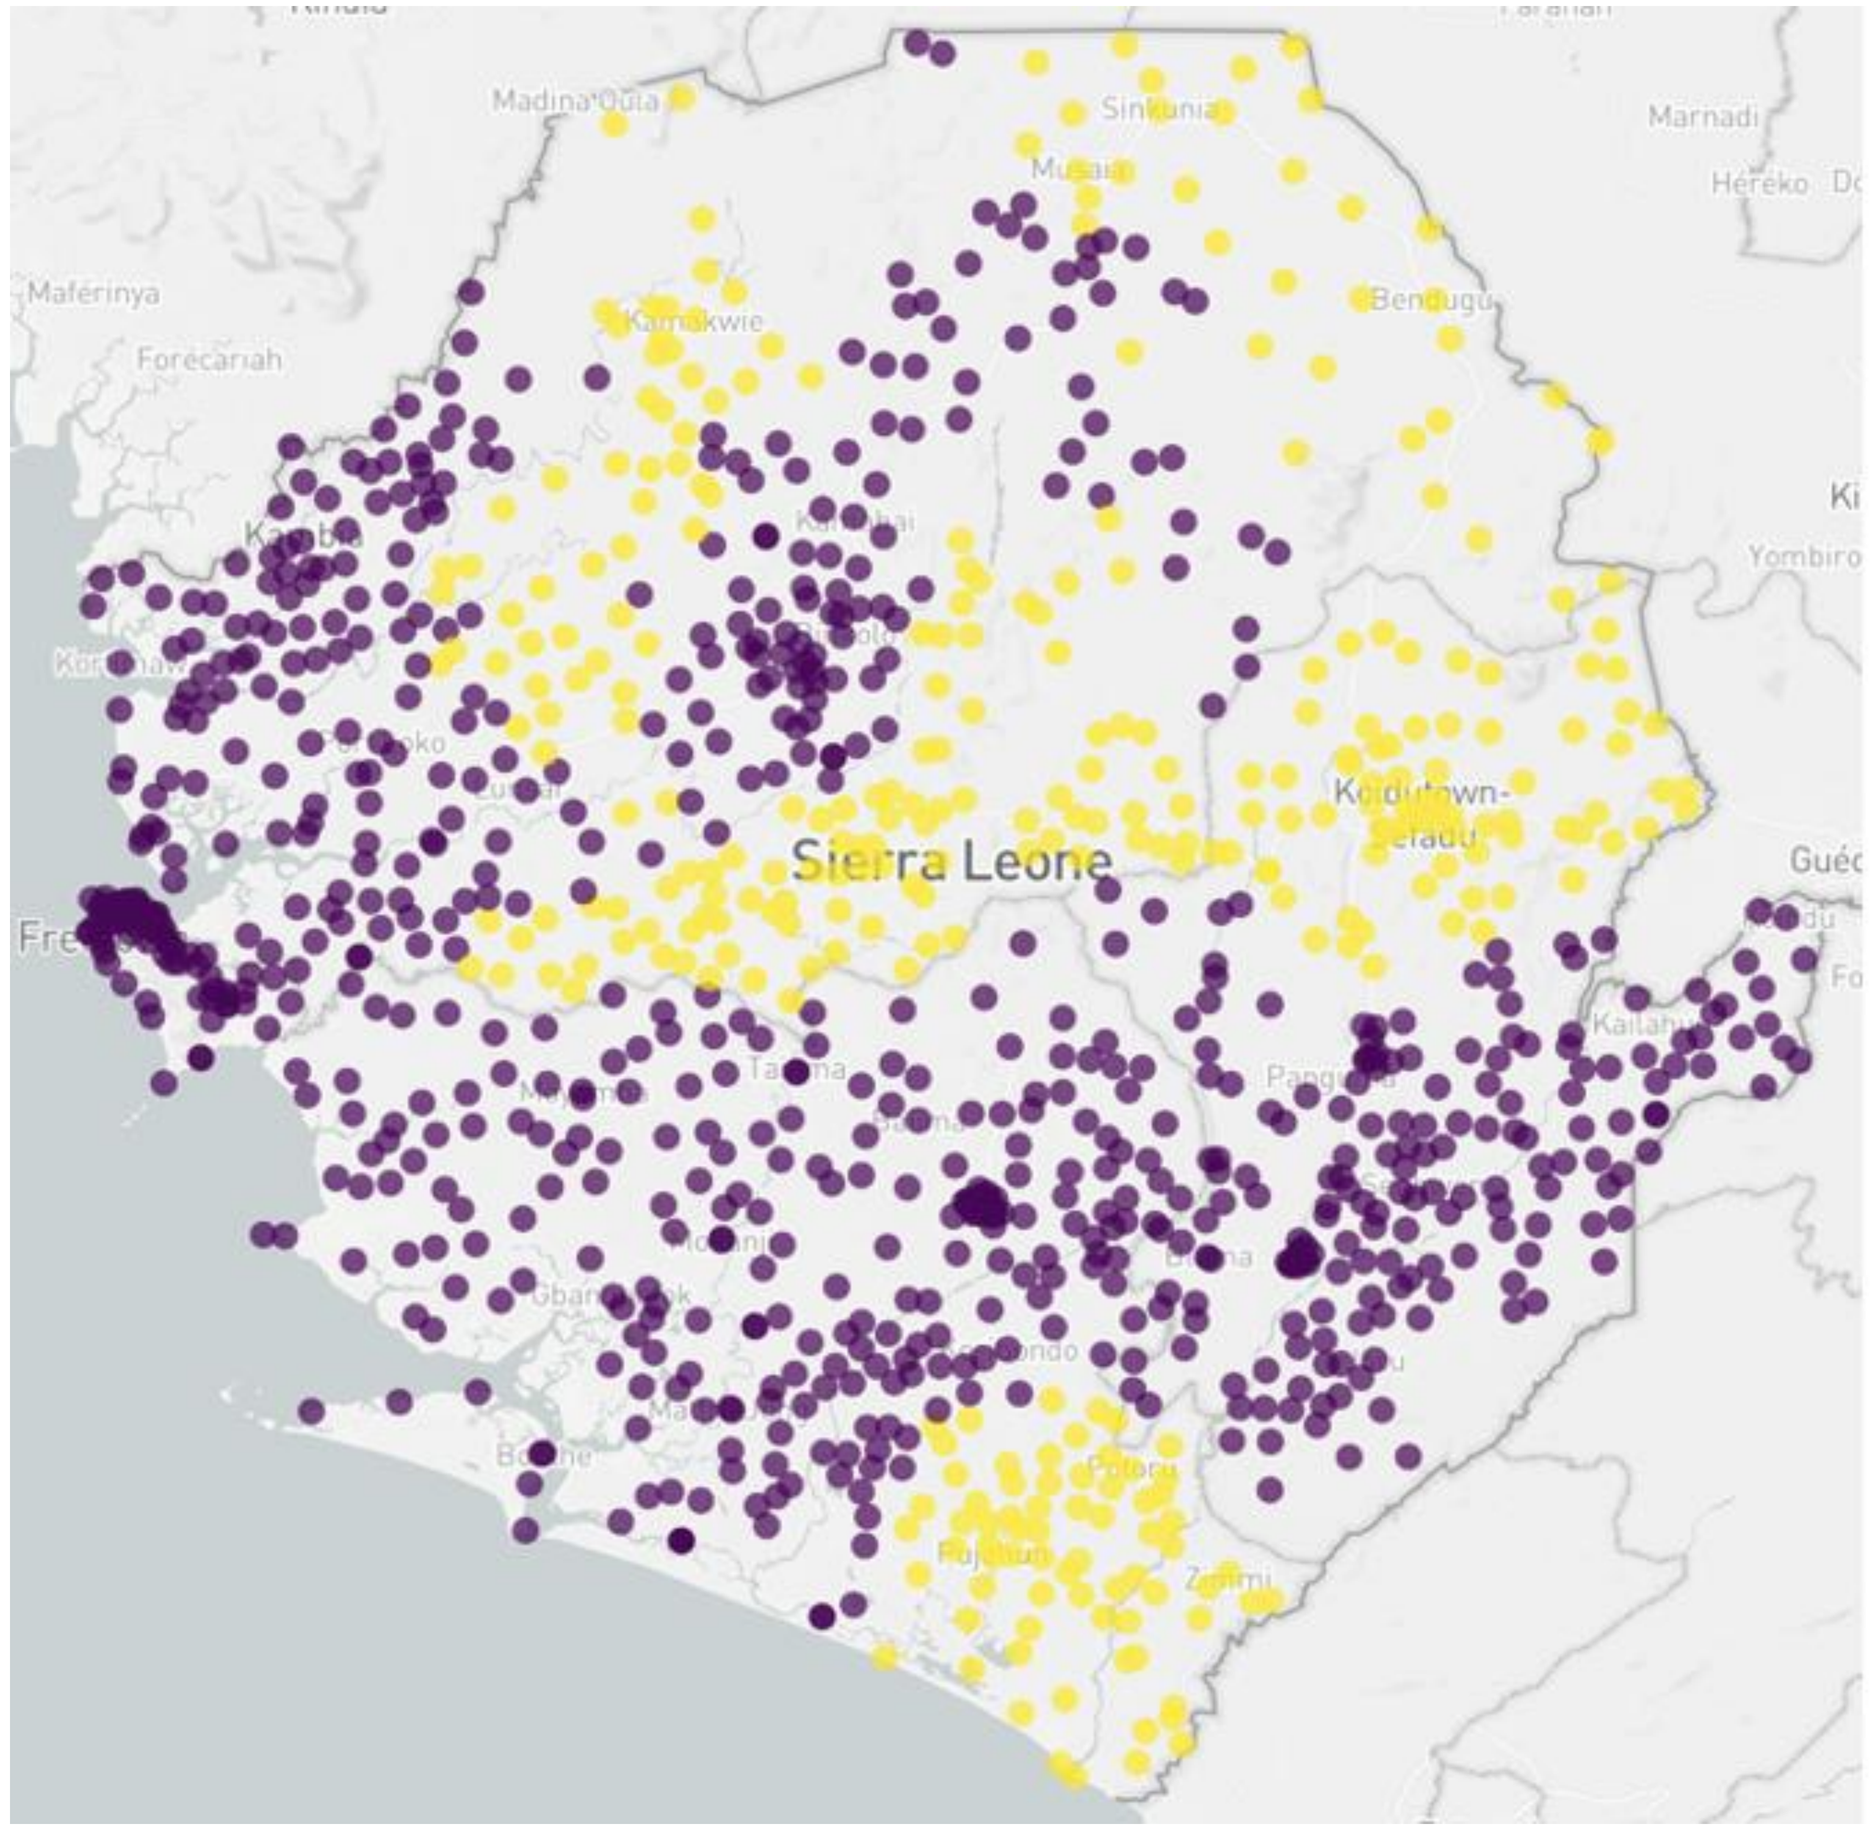

Yellow: treated facilities; Purple: control facilities

Wharton UNIVERSITY of PENNSYLVANIA

Deployment

- 2023 Q2: 5 out of 16 districts
  - Picked randomly by the government
    - 36 treated products
    - 24 control products
- 89~100% compliance

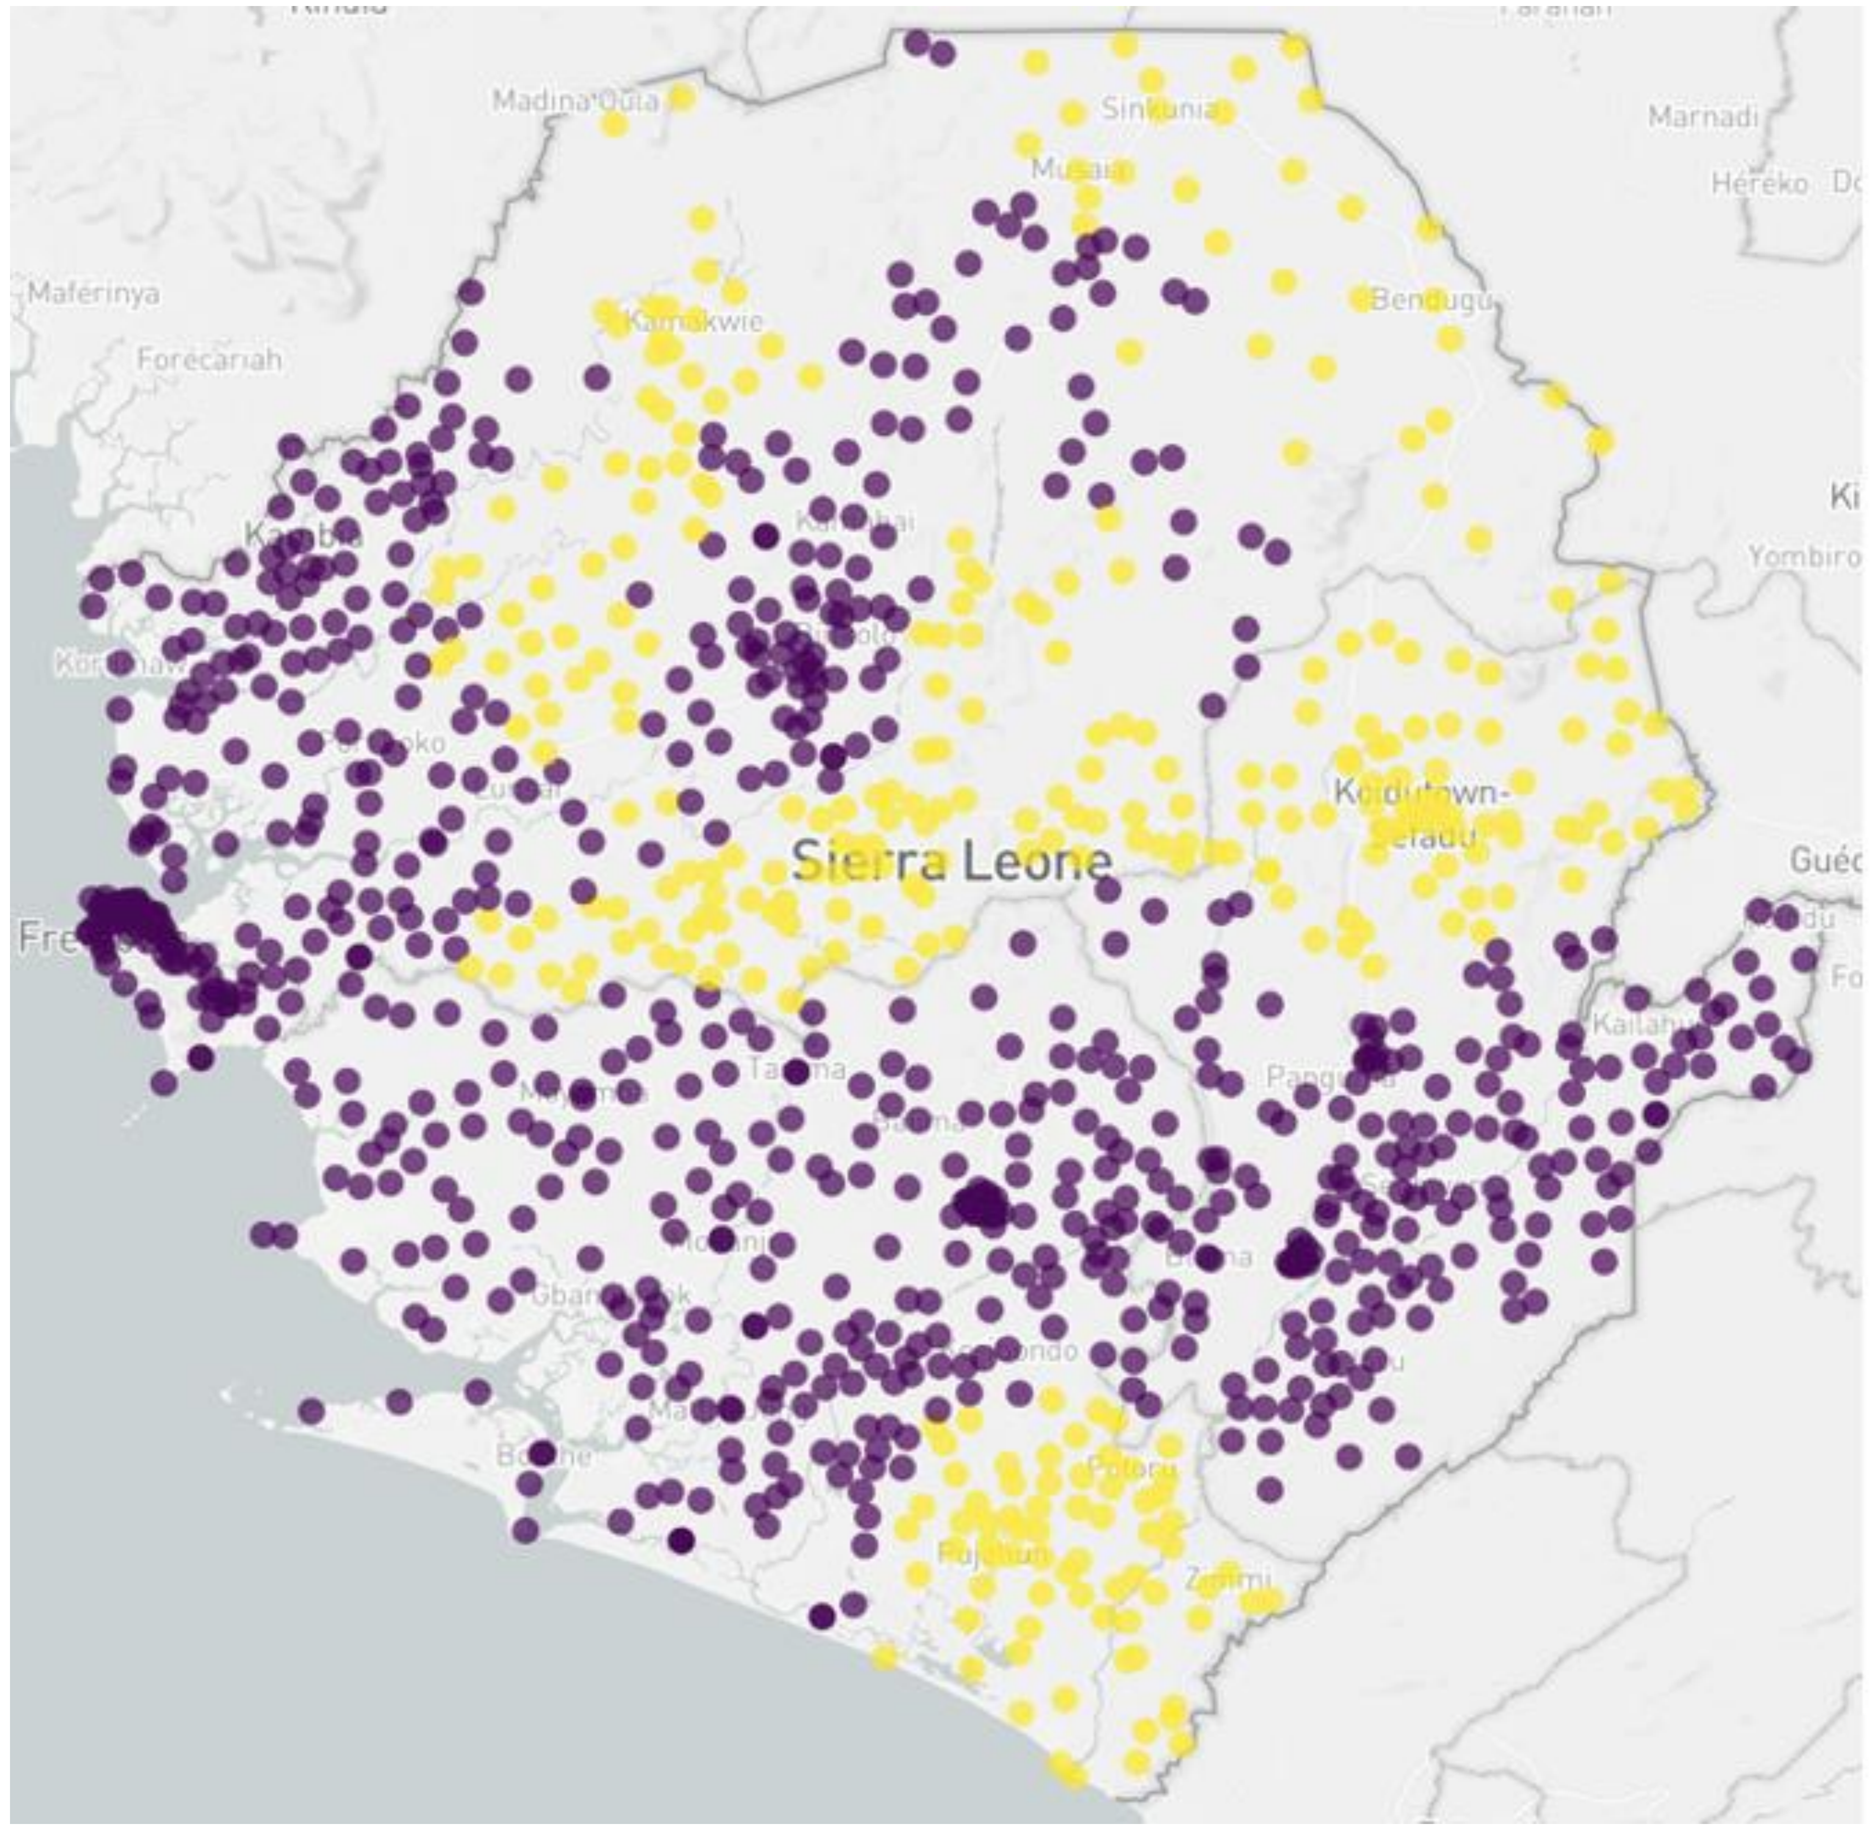

Yellow: treated facilities; Purple: control facilities

Deployment

- 2023 Q2: 5 out of 16 districts
  - Picked randomly by the government
    - 36 treated products
    - 24 control products
- 89~100% compliance
- Expand to all 16 districts in 2023 Q3

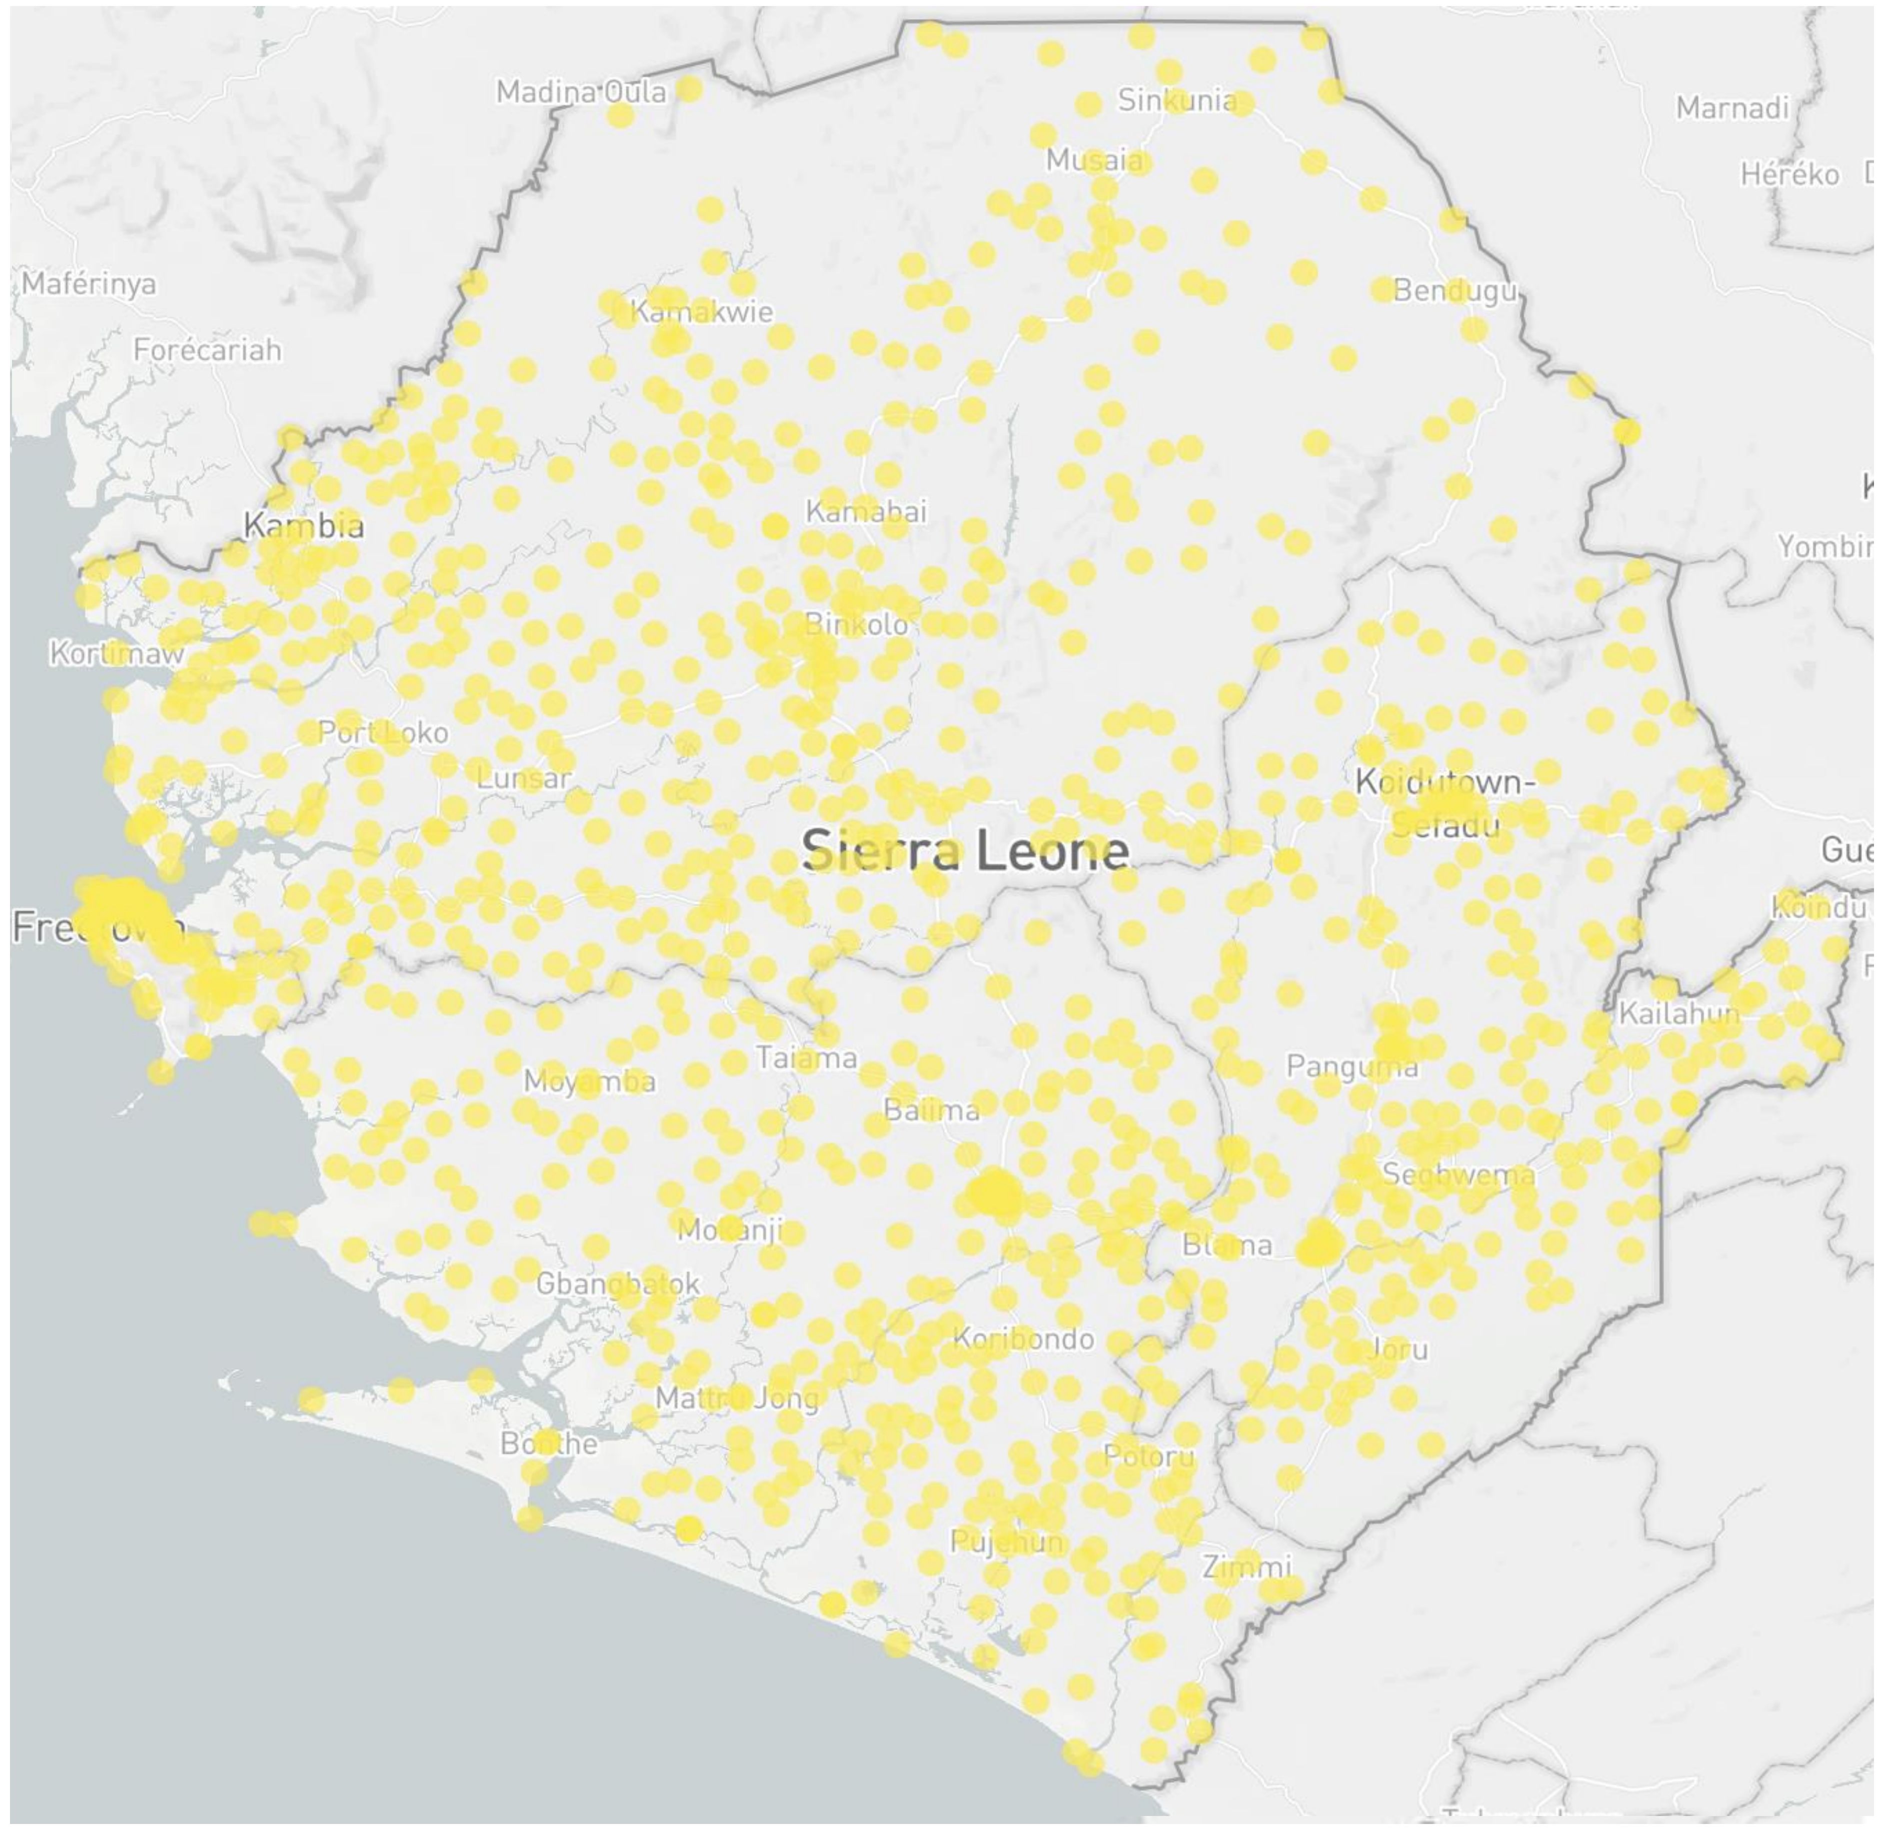

Yellow: treated facilities; Purple: control facilities

Evaluation: Q2 all treated districts

- Dependent variable: Consumption
- Synthetic Difference-in-differences (Arkhangelsky, Athey, Hirshberg, Imbens, and Wager, 2021)
- On average, consumption significantly increased around 19% compared with controls

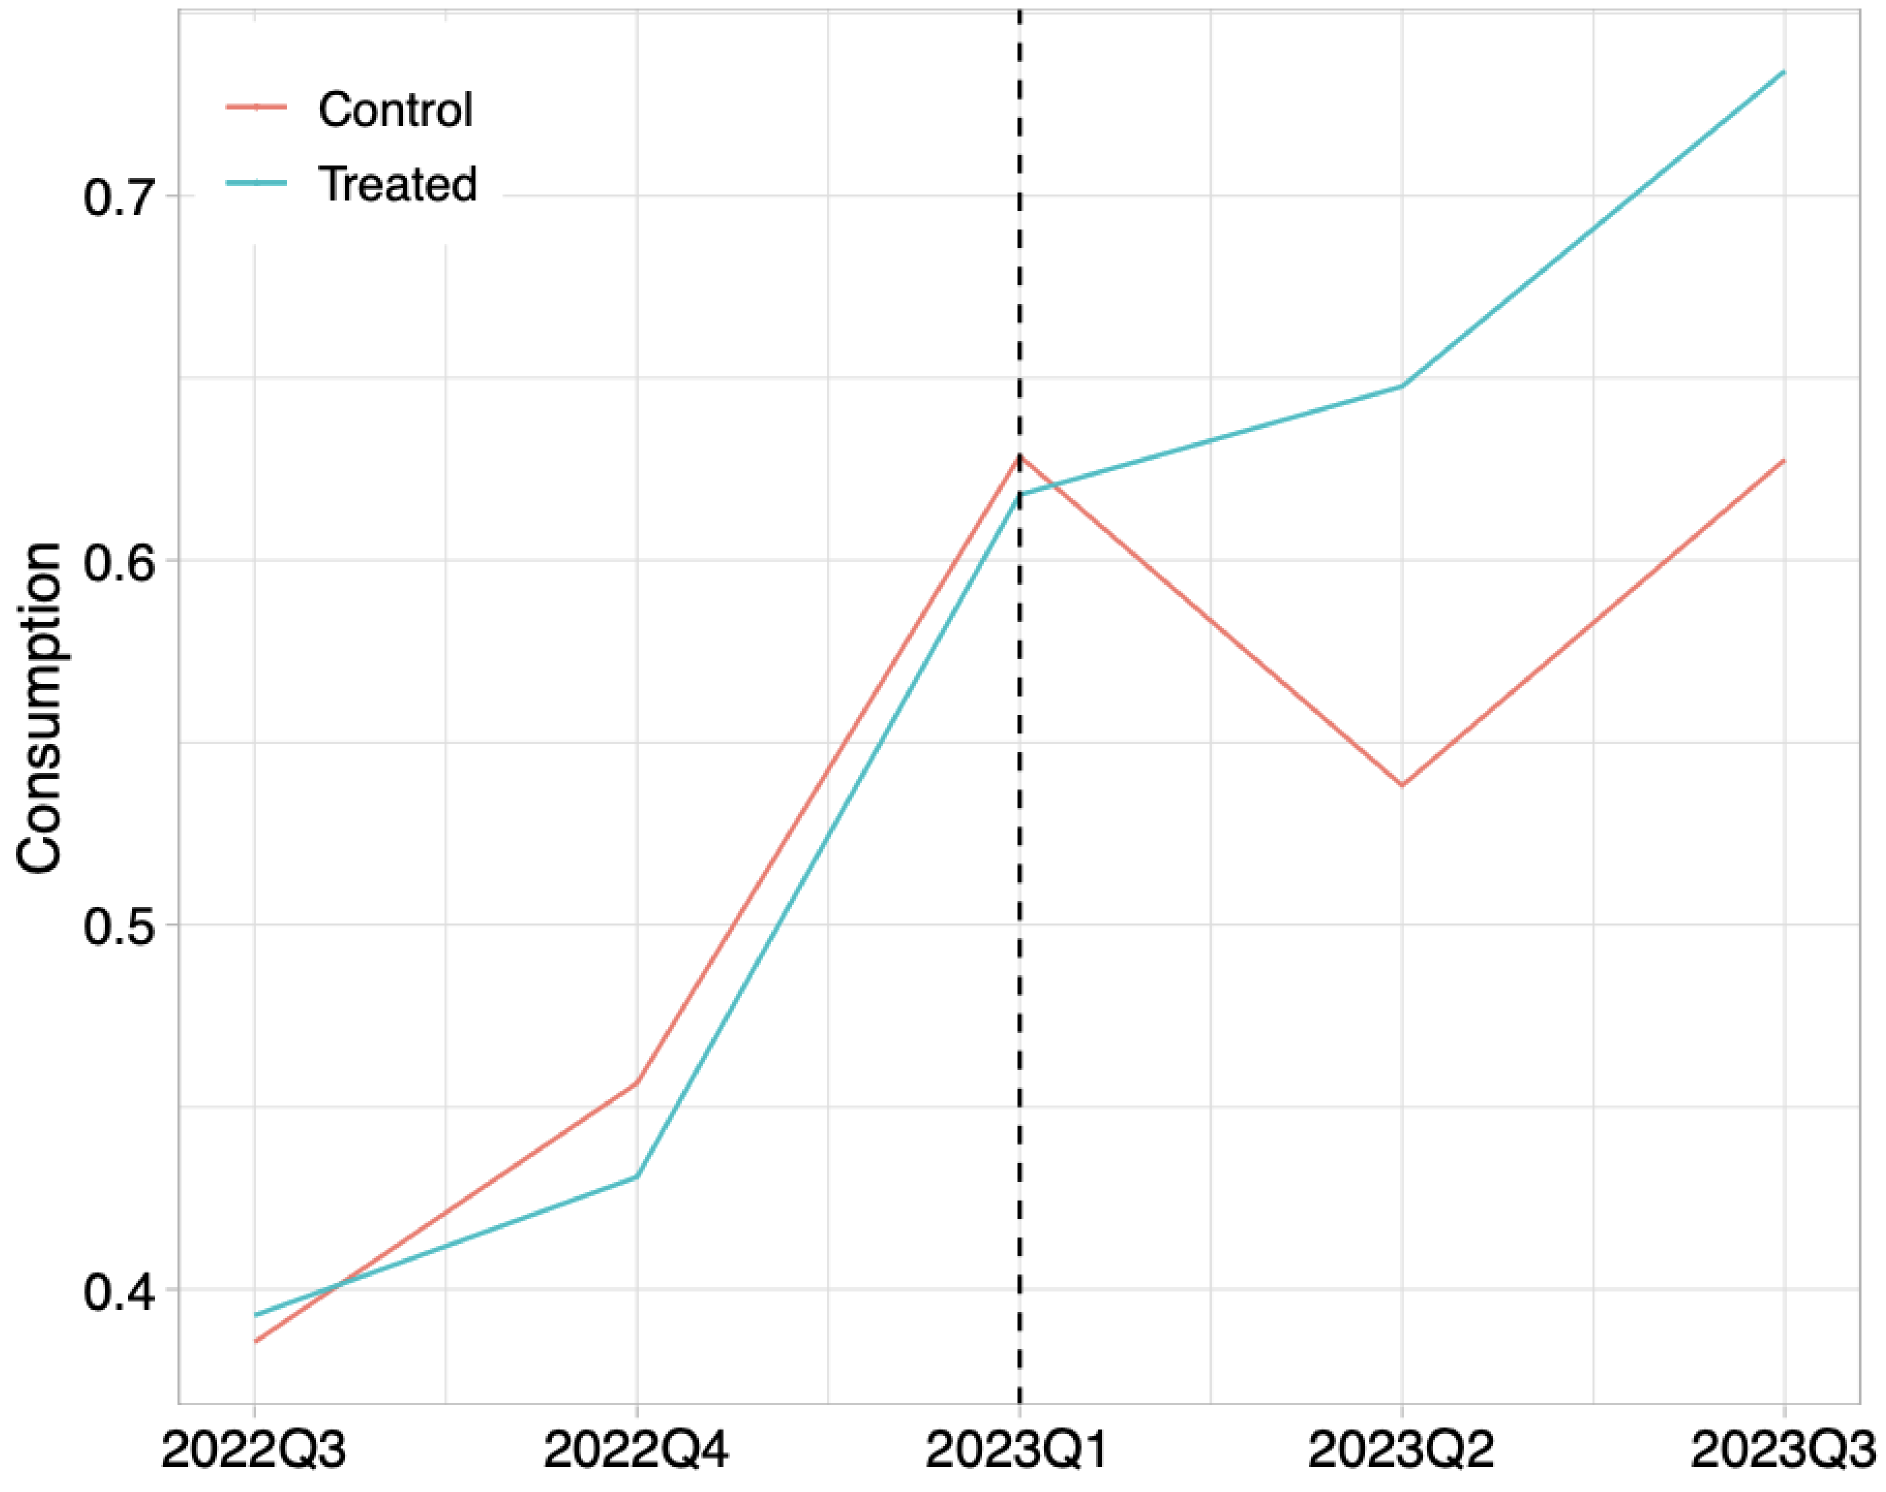

Robustness checks

- All robustness analysis shows **consistent significant impact.**
  - Intention to Treat (ITT)
  - Alternative ITT: use treated and control products
  - Local Average Treatment Effect (LATE): average treatment effect on highly compliant districts

Method	Analysis	95% CI	P-value	**Increase %**
Instrument Variable (Angrist and Imbens, 1995)	LATE	(0.15, 0.37)	<0.05	**26%**
SynthDiD	ITT	(0.08, 0.26)	<0.00	**19%**
	Alternative ITT	(0.05, 0.14)	<0.00	**18%**
Differences-in-differences	ITT	(0.09,0.16)	<0.01	**21%**

Field work

Travel from airport to capital city

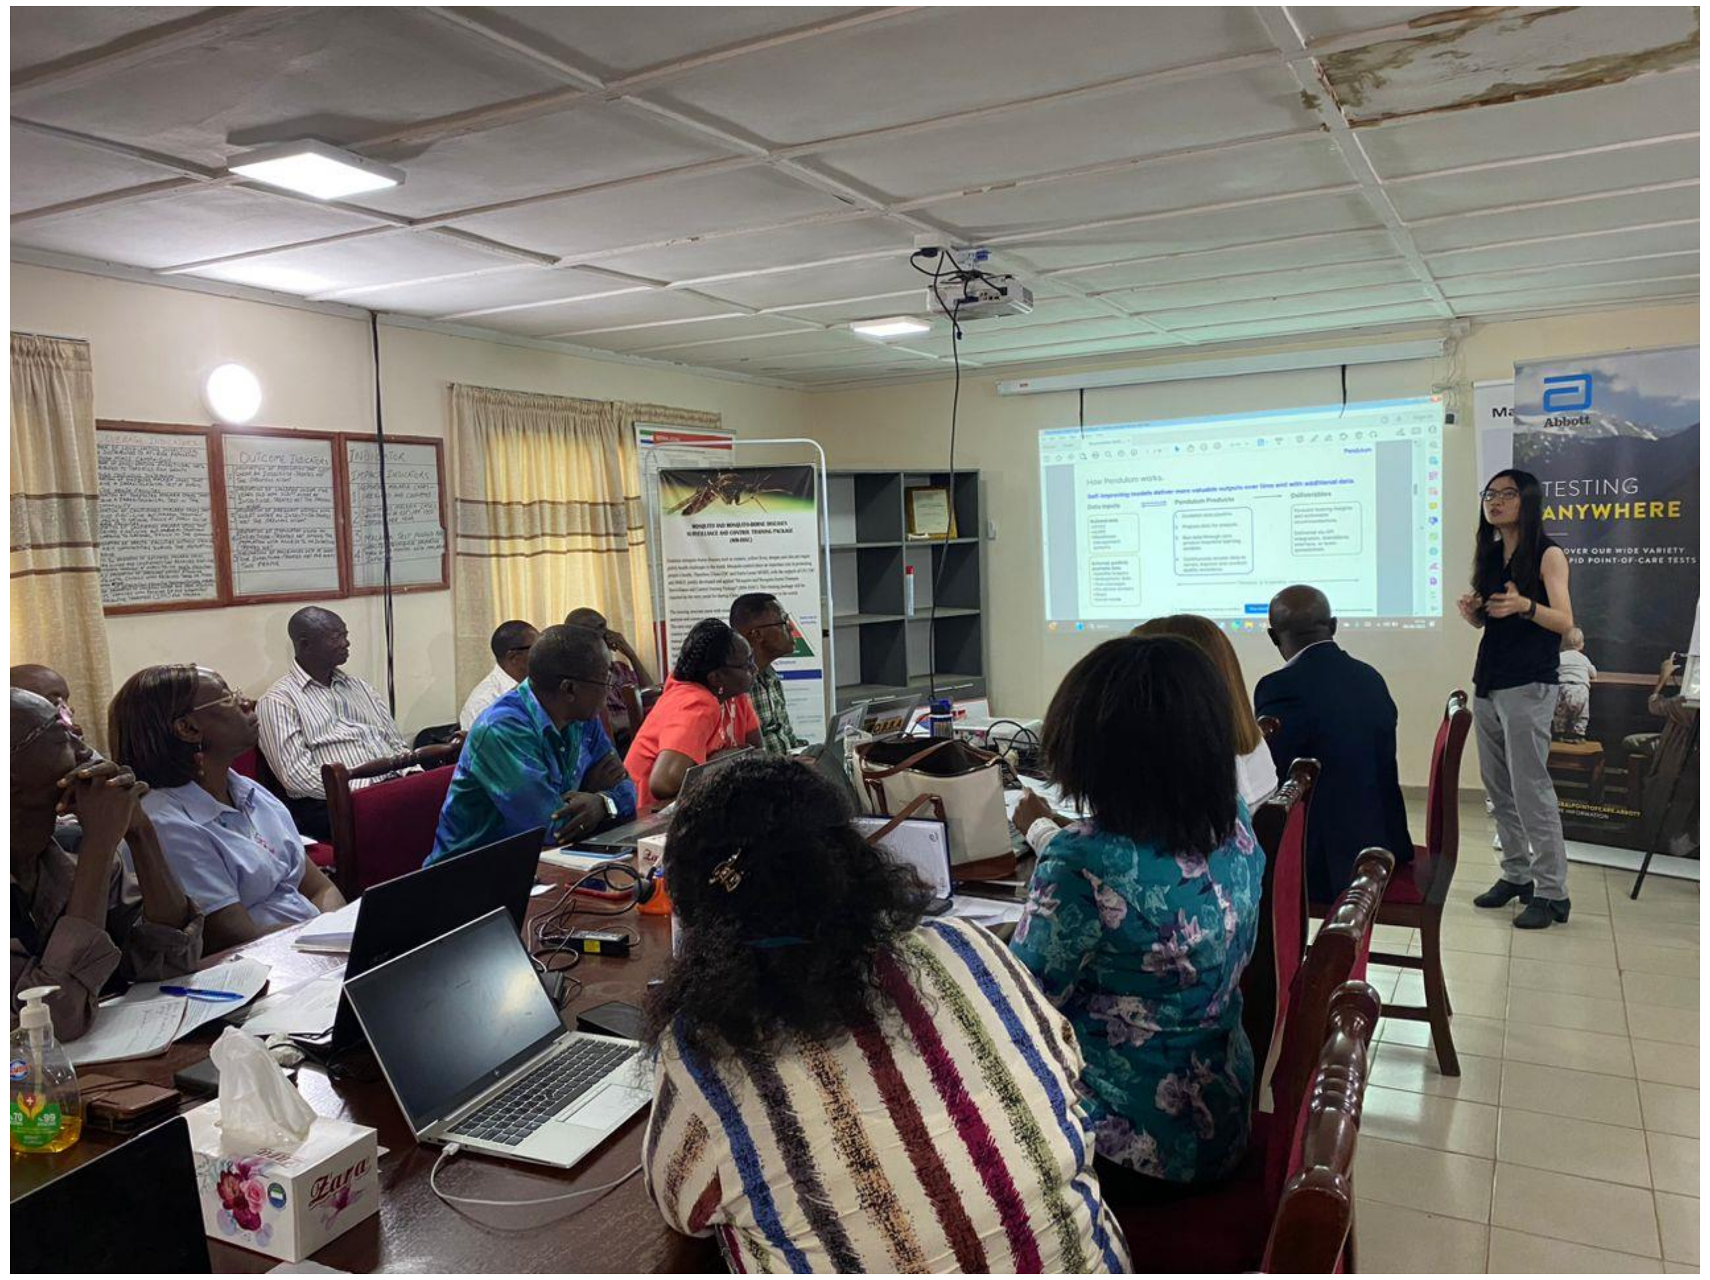

Discuss the work to various stakeholders (WHO, MoHS, US Presidential Malaria Initiative)

Visit the warehouse

Sustainable impact

- Empower the government officials and local staff to use our tool easily in the future.

Select Date

08/01/2023

Upload File(Please note that it should be .xlsx file format)

.XLSX

Click to browse or drag and drop your files here

Result Table

Filter By Product: All

Product	hf_pk	Allocation
0	30388	5515.79639
0	525	1571.720123
0	523	1754.607926

Conclusion

- We propose a novel **decision-aware learning** framework that could address the challenges in a **resource-constrained** setting.
- We **successfully deployed** on the ground at **national level** and shows a significant real impact → **19%** increase of access to medicines for 2 million of women and children under five.

Thank You!

angelchg@wharton.upenn.edu

MINISTRY OF HEALTH AND SANITATION GOVERNMENT OF SIERRA LEONE

Wharton UNIVERSITY of PENNSYLVANIA

https://angel-chung.github.io

Distributional impact & fairness

- Significantly large impact on facilities with historical stockouts.

Dep. var.: Consumption	By Facility Type		By Location				
			0.25 miles		1 mile		Under-Served
	Large Facility	Small Facility	Rural	Urban	Rural	Urban	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Treatment Effect	0.368* (0.183)	0.052 (0.037)	0.107** (0.043)	0.291 (0.395)	0.090* (0.044)	0.334 (0.315)	0.182** (0.069)
Observations	1,075	4,215	4,865	425	4,430	860	4,055
Improvement %	36%	10% (Insig)	19%	24% (Insig)	16%	32% (Insig)	32%

*Note: *p < 0.05, **p < 0.01, ***p < 0.001. SE in parentheses.*

All robustness checks results

Table 7: **Robustness Analysis Results.**

Model	**Coefficient**	**Std. Error**	**Observations**	**Improvement %**
Dependent Variable = Consumption				
SynthDiD	0.116*	(0.046)	5,290	19%
DiD	0.128**	(0.046)	5,290	21%
Matching (5km)	0.154**	(0.066)	1,125	25%
Matching (10km)	0.166**	(0.064)	1,815	30%
Matching (15km)	0.121**	(0.054)	2,415	21%
Imputation (low rank)	0.037**	(0.014)	5,455	15%
Imputation (avg consump)	0.067**	(0.031)	5,455	21%
Imputation (pop)	0.076**	(0.028)	5,455	27%
No Missingness Imbalance	0.132*	(0.058)	5,285	18%
Substitution (Method 2)	0.132*	(0.048)	5,290	20%
Alt. Control	0.095***	(0.022)	10,520	18%
Substitution (Method 1)	0.119*	(0.048)	5,290	18%
LATE (Continuous)	0.115**	(0.038)	5,290	37%
Dependent Variable = Stockout				
SynthDiD	−0.280	(0.179)	5,290	−4.6% (Insig)
Dependent Variable = Missingness				
SynthDiD	−0.007	(0.006)	5,290	−1.4% (Insig)

Notes: $^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$. Standard errors in parentheses. The balanced dataset includes 1,058 facilities across five quarters from 2022 Q3. The Improvement % column reports the relative change in consumption compared to the counterfactual mean for each specification.

Balanced table

Table 13: **Balance Table: District Characteristics by Treatment Status**

Variable	Control Mean	Control (SD)	Treatment Mean	Treatment (SD)	Difference p-value	Difference FDR Adj. p-value
Days of service delivery	0.46	(0.17)	0.38	(0.17)	0.35	0.71
Hours of service delivery	1.40	(0.56)	1.26	(0.55)	0.14	0.71
Facilities where women give birth	6.44	(2.54)	5.23	(2.47)	0.19	0.71
Availability of basic emergency obstetric and neonatal care	0.43	(0.38)	0.19	(0.18)	0.35	0.71
Availability of priority drugs	3.62	(1.39)	3.12	(1.58)	0.78	0.88
Availability of all tracer drugs	1.84	(1.41)	2.40	(1.78)	0.24	0.71
Availability of vaccines	6.32	(2.31)	5.31	(2.46)	0.68	0.88
Vaccines storage: refrigerators 2–8 °C	1.72	(2.72)	1.06	(2.14)	0.87	0.88
Availability of communication equipment	3.65	(2.56)	2.56	(1.50)	0.41	0.71
Access to various forms of communication	3.65	(2.56)	2.56	(1.50)	0.41	0.71
Total proportion of facilities carrying out safe health care waste disposal	5.00	(2.10)	3.86	(1.66)	0.68	0.88
Availability of basic equipment	1.99	(0.67)	2.00	(0.91)	0.31	0.71
Availability of Standard Treatment Guidelines	3.63	(1.97)	2.53	(2.56)	0.44	0.71
Outpatient caseload (median per facility)	0.63	(0.52)	0.51	(0.21)	0.88	0.88
Facilities with community health workers	5.44	(2.54)	4.84	(2.82)	0.78	0.88
Average number of community health workers	0.69	(0.24)	0.64	(0.33)	0.57	0.85
Average health workers per facility	0.47	(0.50)	0.41	(0.38)	0.81	0.88
Population Share	6.95	(3.13)	4.72	(2.45)	0.15	0.71
Human Development Index (HDI)	0.03	(0.01)	0.03	(0.01)	0.18	0.71
Health Index (health dimension of HDI based on life expectancy at birth)	0.04	(0.01)	0.05	(0.01)	0.37	0.71
Income Index (income dimension of HDI based on Gross National Income per capita)	0.03	(0.01)	0.03	(0.01)	0.06	0.71

Notes: Means and standard deviations are shown as Mean (SD). p-values are from two-sample Welch's t-tests. Multiple testing adjustments control the false discovery rate (FDR) using the Benjamini–Hochberg (BH) procedure.

Evaluation: compliance

- Treatment compliance: the absolute difference between actual allocation and our deployed suggestion
  - 89~100% compliance

District	Average absolute difference across products and facilities
Tonkolili	0.000
Falaba	0.028
Karene	0.039
Kono	0.073
Pujehun	0.109

Compliance of treated districts

Baselines comparison

Method	– Improvement %
Our Framework	0%
Decision-Blind Ablation	5%
Population Based Census	27%
Distribution Modeling	82%
Global Health (3mth Avg)	88%
StochOptForest	92%
Existing Excel Tool	98%

Average % Improvement of Our Framework vs. Baselines Across Q2 Products

Multi-task learning

- Learn across facilities & products [Bastani (2021), Bastani, Simchi-Levi & Zhu (2022), Xu & Bastani (2022)]
- Use random forest

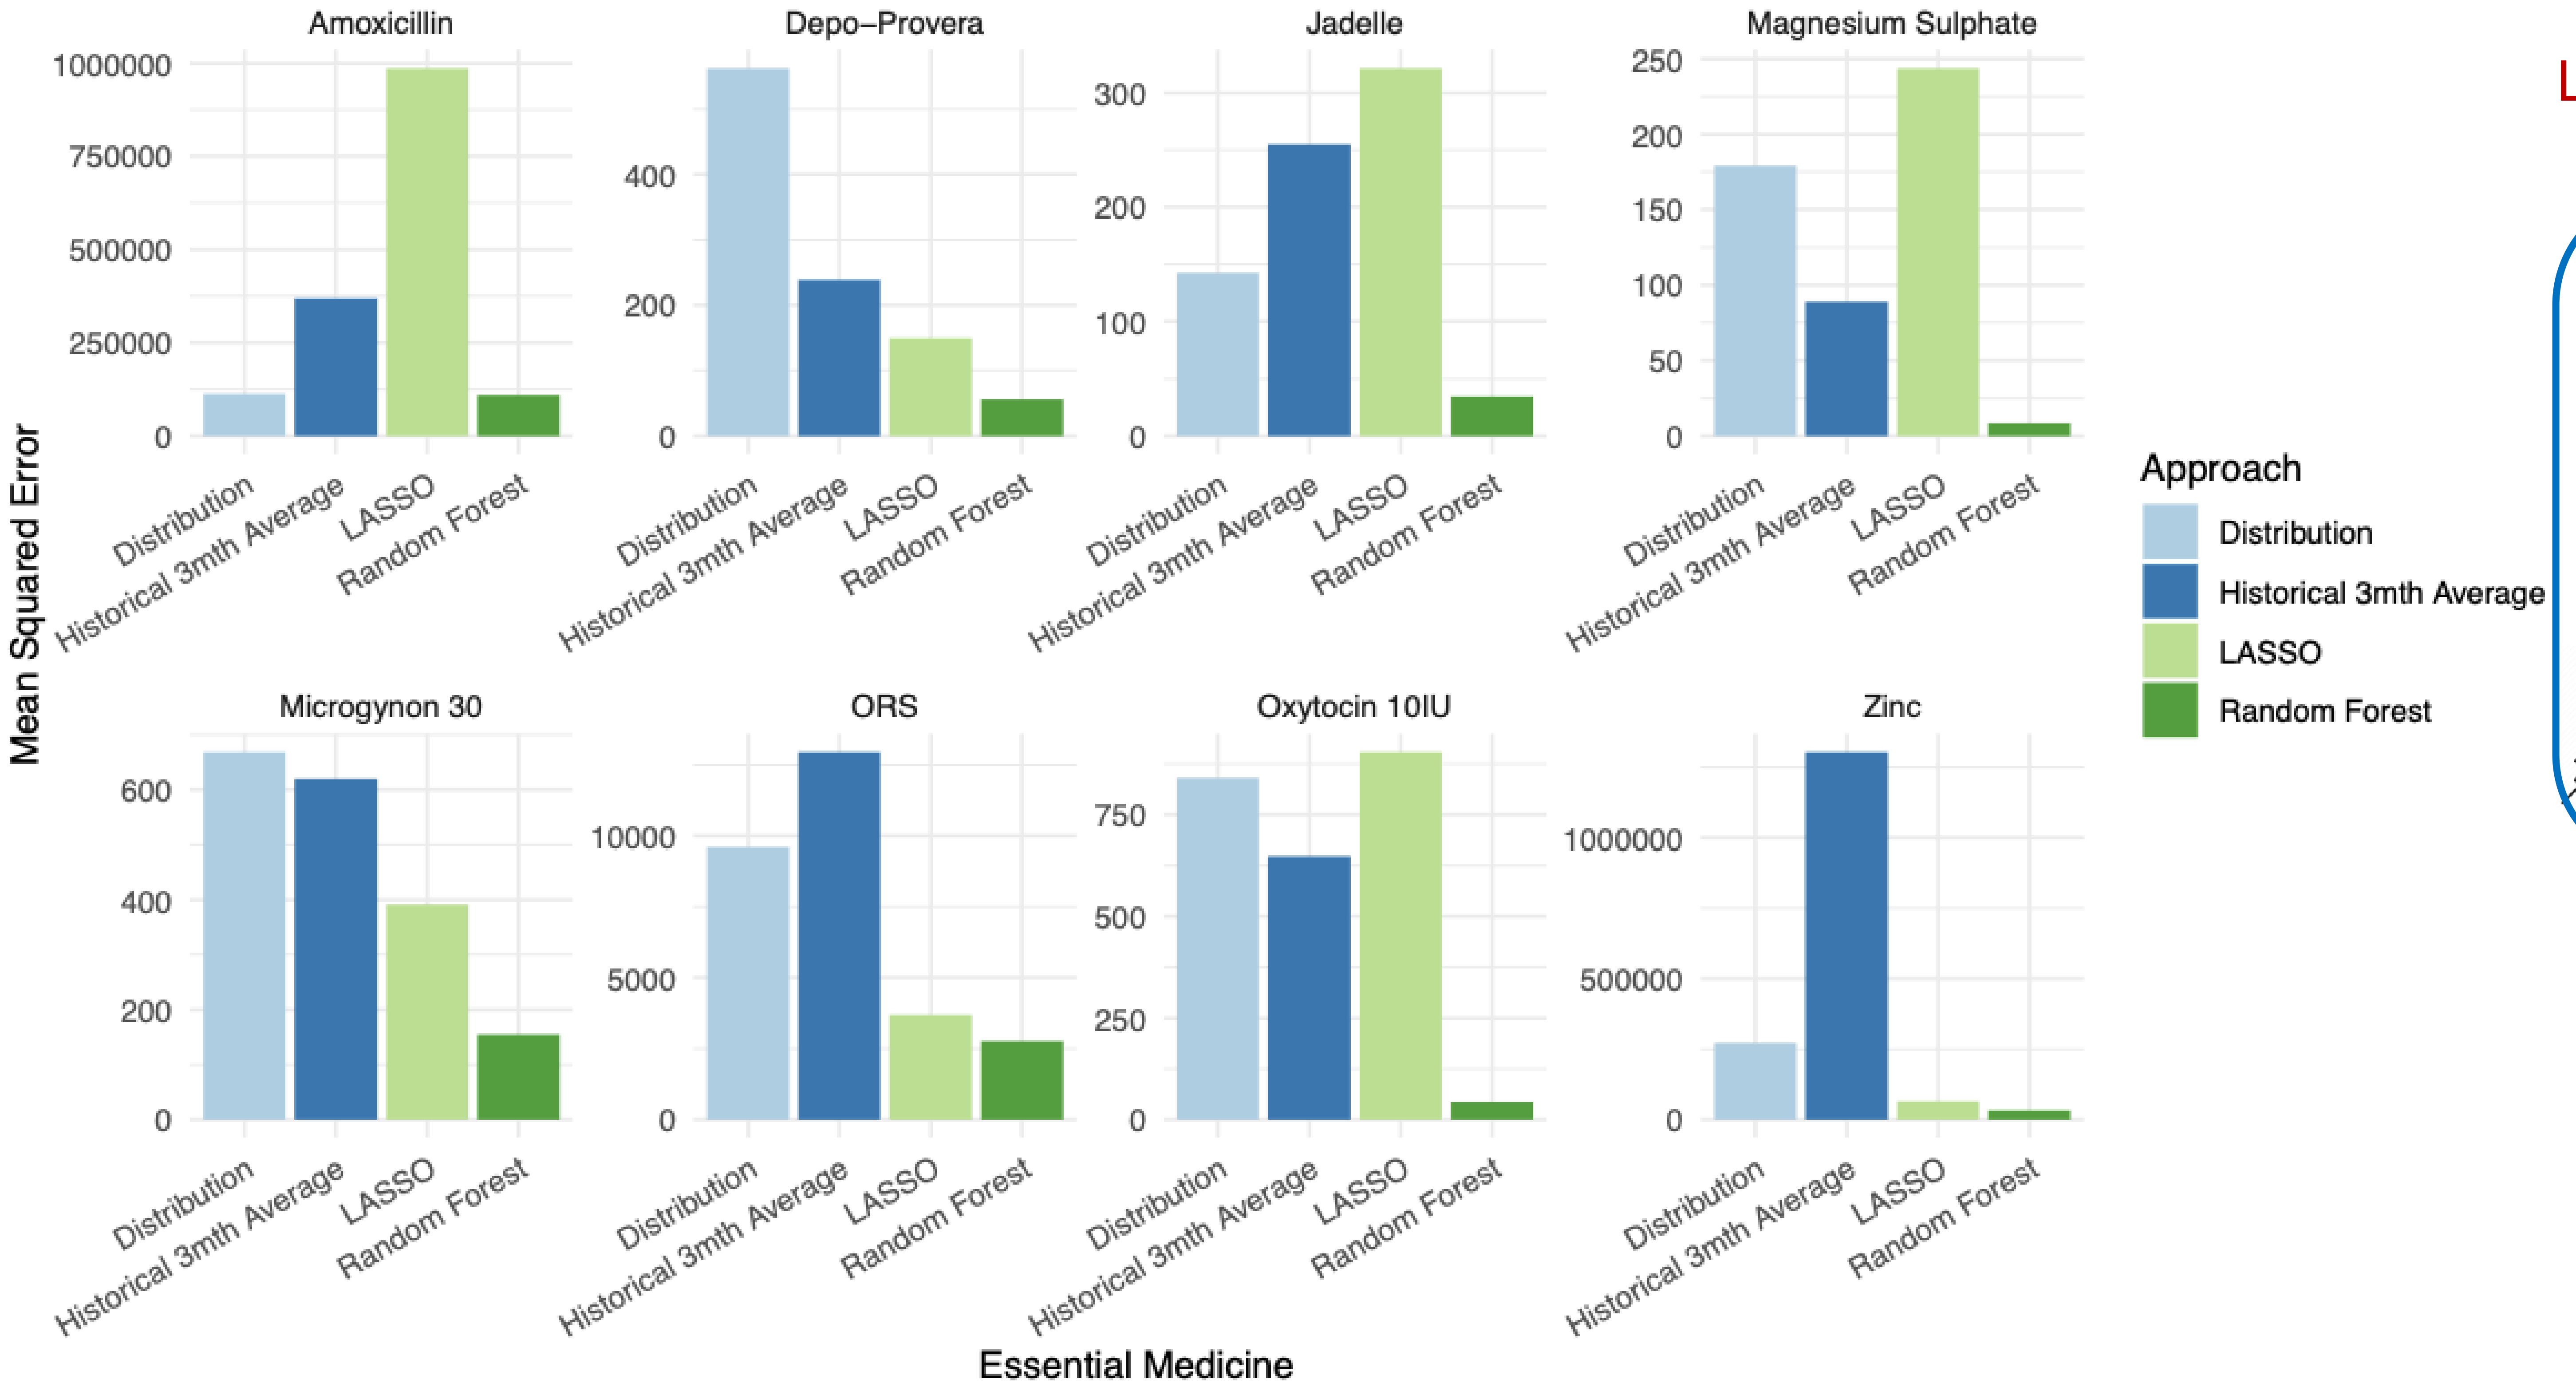

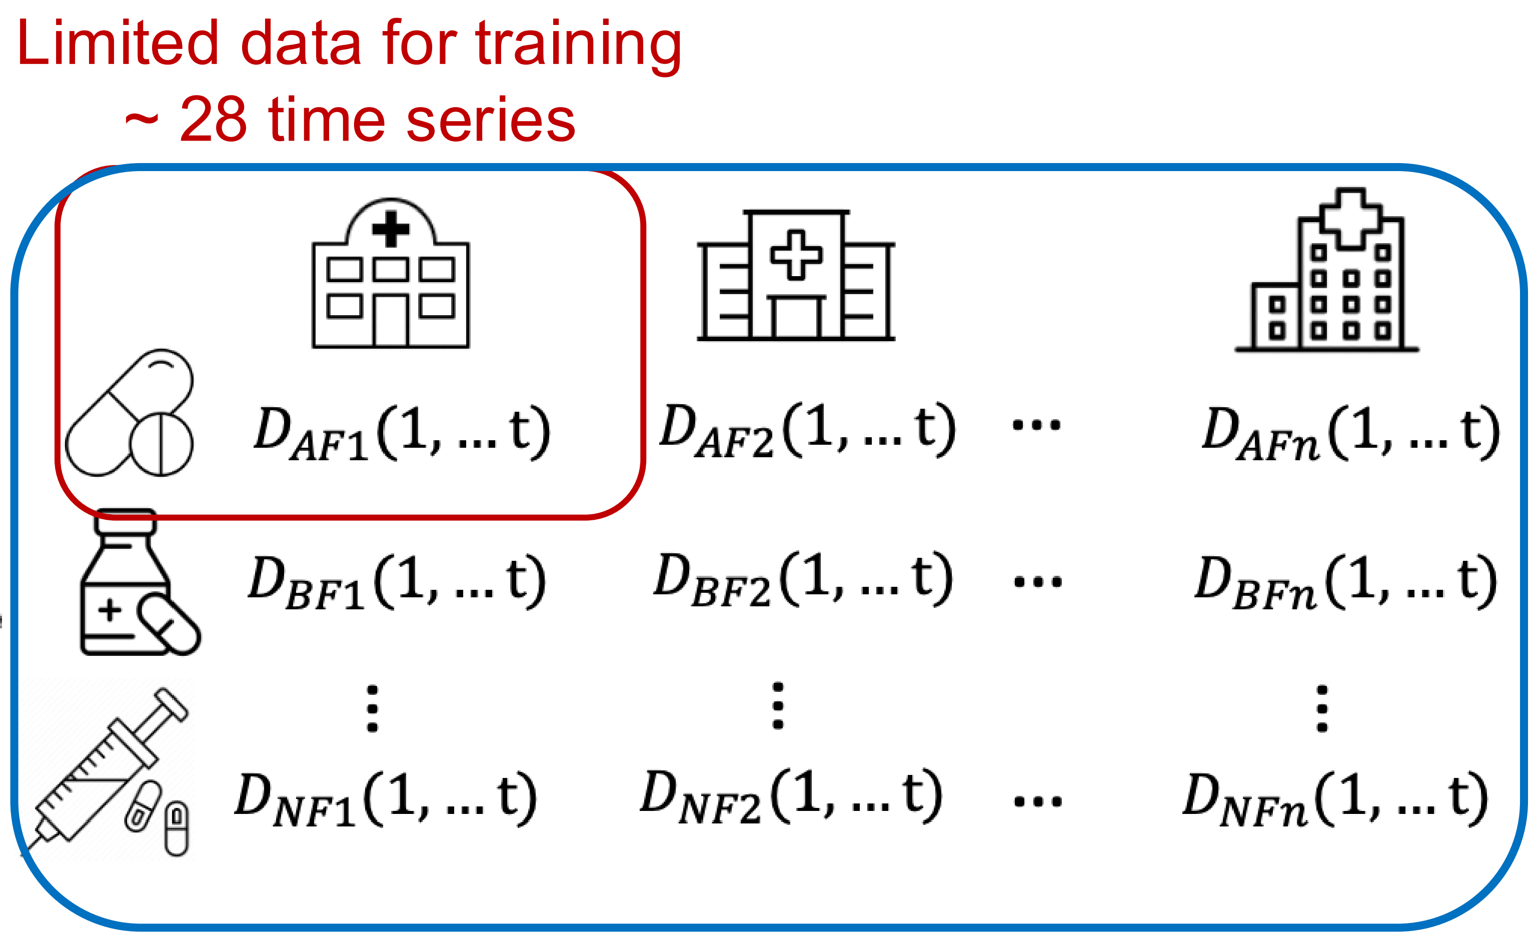

Addressing potential data equity issue

- Impute missing values [Agarwal et.al, 2021; Athey et.al, 2020] -> doesn't work

- **Catalytic prior** [Huang et.al, 2020]
  - regularize the complex model towards the simple model

- **Simple model**:
  - Use at-risk population (Census, UN)
  - Ignore temporal & seasonal variation

- **Complex model**:
  - Random forest with dhis2 data
  - Issue of missing not at random (MNAR)

Missing Completely at Random (MCAR)

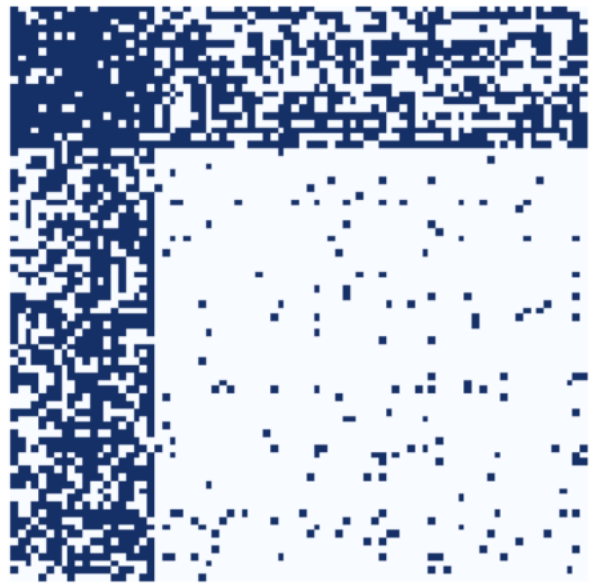

Missing Not at Random (MNAR)

- We combine the model and train the random forest on this augmented dataset with the synthetic data.

Empirical strategy

- Difference-in-differences:

$$\left(\hat{\tau}^{\text{did}},\, \hat{\mu},\, \hat{\alpha},\, \hat{\beta}\right) = \underset{\alpha,\beta,\mu,\tau}{\arg\min} \left\{ \sum_{i=1}^{N} \sum_{t=1}^{T} \left(Y_{it} - \mu - \alpha_i - \beta_t - W_{it}\tau \right)^2 \right\}$$

- SynthDiD:

$$\left(\hat{\tau}^{\text{sdid}},\, \hat{\mu},\, \hat{\alpha},\, \hat{\beta}\right) = \underset{\tau,\mu,\alpha,\beta}{\arg\min} \left\{ \sum_{i=1}^{N} \sum_{t=1}^{T} \left(Y_{it} - \mu - \alpha_i - \beta_t - W_{it}\tau \right)^2 \hat{\omega}_i^{\text{sdid}} \hat{\lambda}_t^{\text{sdid}} \right\}$$

  - Unit weights: let average outcome for the treated units is approximately parallel to the weighted average for control units.
  - Time weights: let average post-treatment outcome for each of the control units differ by a constant from the weighted average of the pre-treatment outcomes for the same control units.
  - Our Y= consumption
    - Minimize unmet demand is equivalent to maximize consumption when demand is censored.