Designing and Deploying AI in Public Systems: Evidence from Healthcare and Education

Angel Tsai-Hsuan Chung

2026 Summer Data Driven Discovery Initiative

Outline

- Improving Access to Essential Medicines via Decision-Aware Machine Learning
  - Joint work with Prof. Hamsa Bastani, Prof. Osbert Bastani, and Sierra Leone Government (Jatu Abdulai, Patrick Bayoh, Lawrence Sandi, Francis Smart)
- Effective Personalized AI tutor via LLM-Guided Reinforcement Learning
  - Joint work with Botong Zhang, Ling-Chieh Kung, Hamsa Bastani, and Osbert Bastani
  - In collaboration with Taipei City Government, American Institute in Taiwan, and American Innovation Center

Context in Sierra Leone

- From 2022 report: GDP per capita = $532

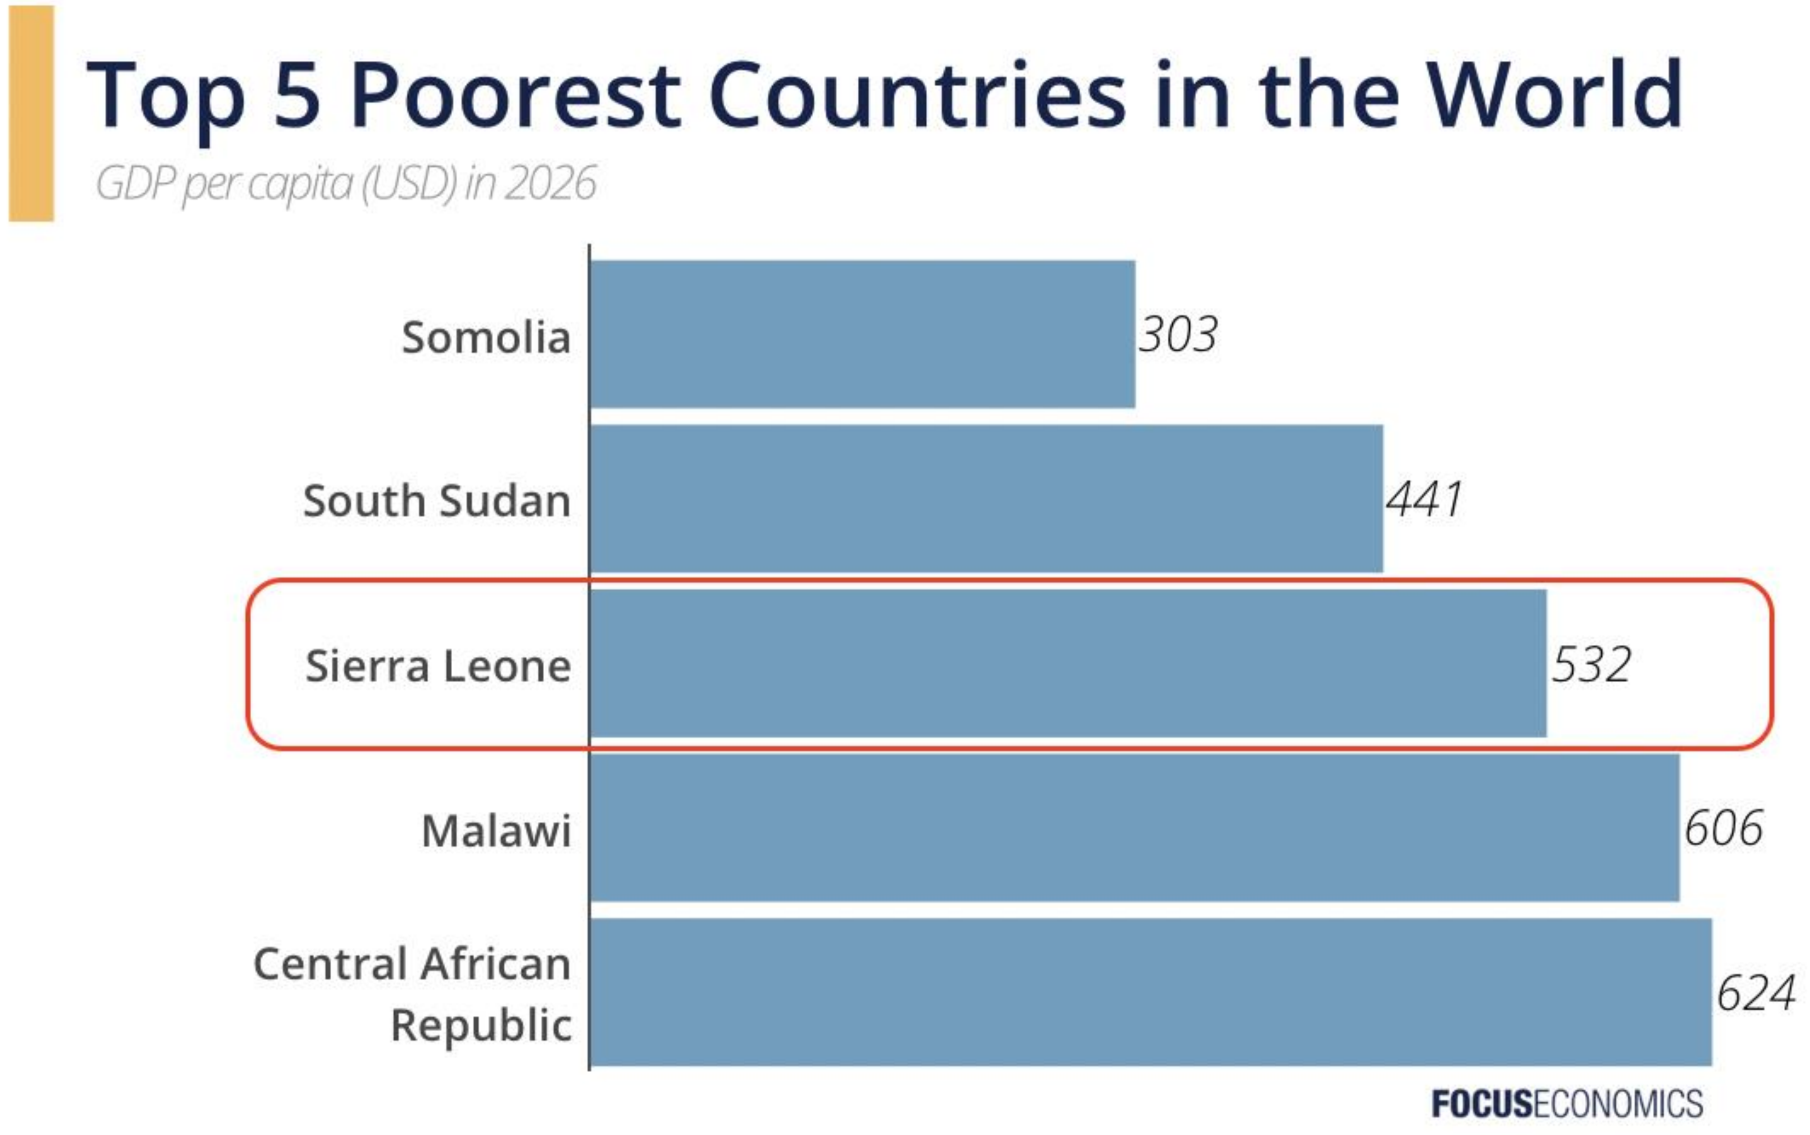

- Child mortality rate is **3.11%** (**0.56%** in the US)
- Maternal mortality: **510** deaths per 100,000 live births (**22.3** per 100,000 in the US)

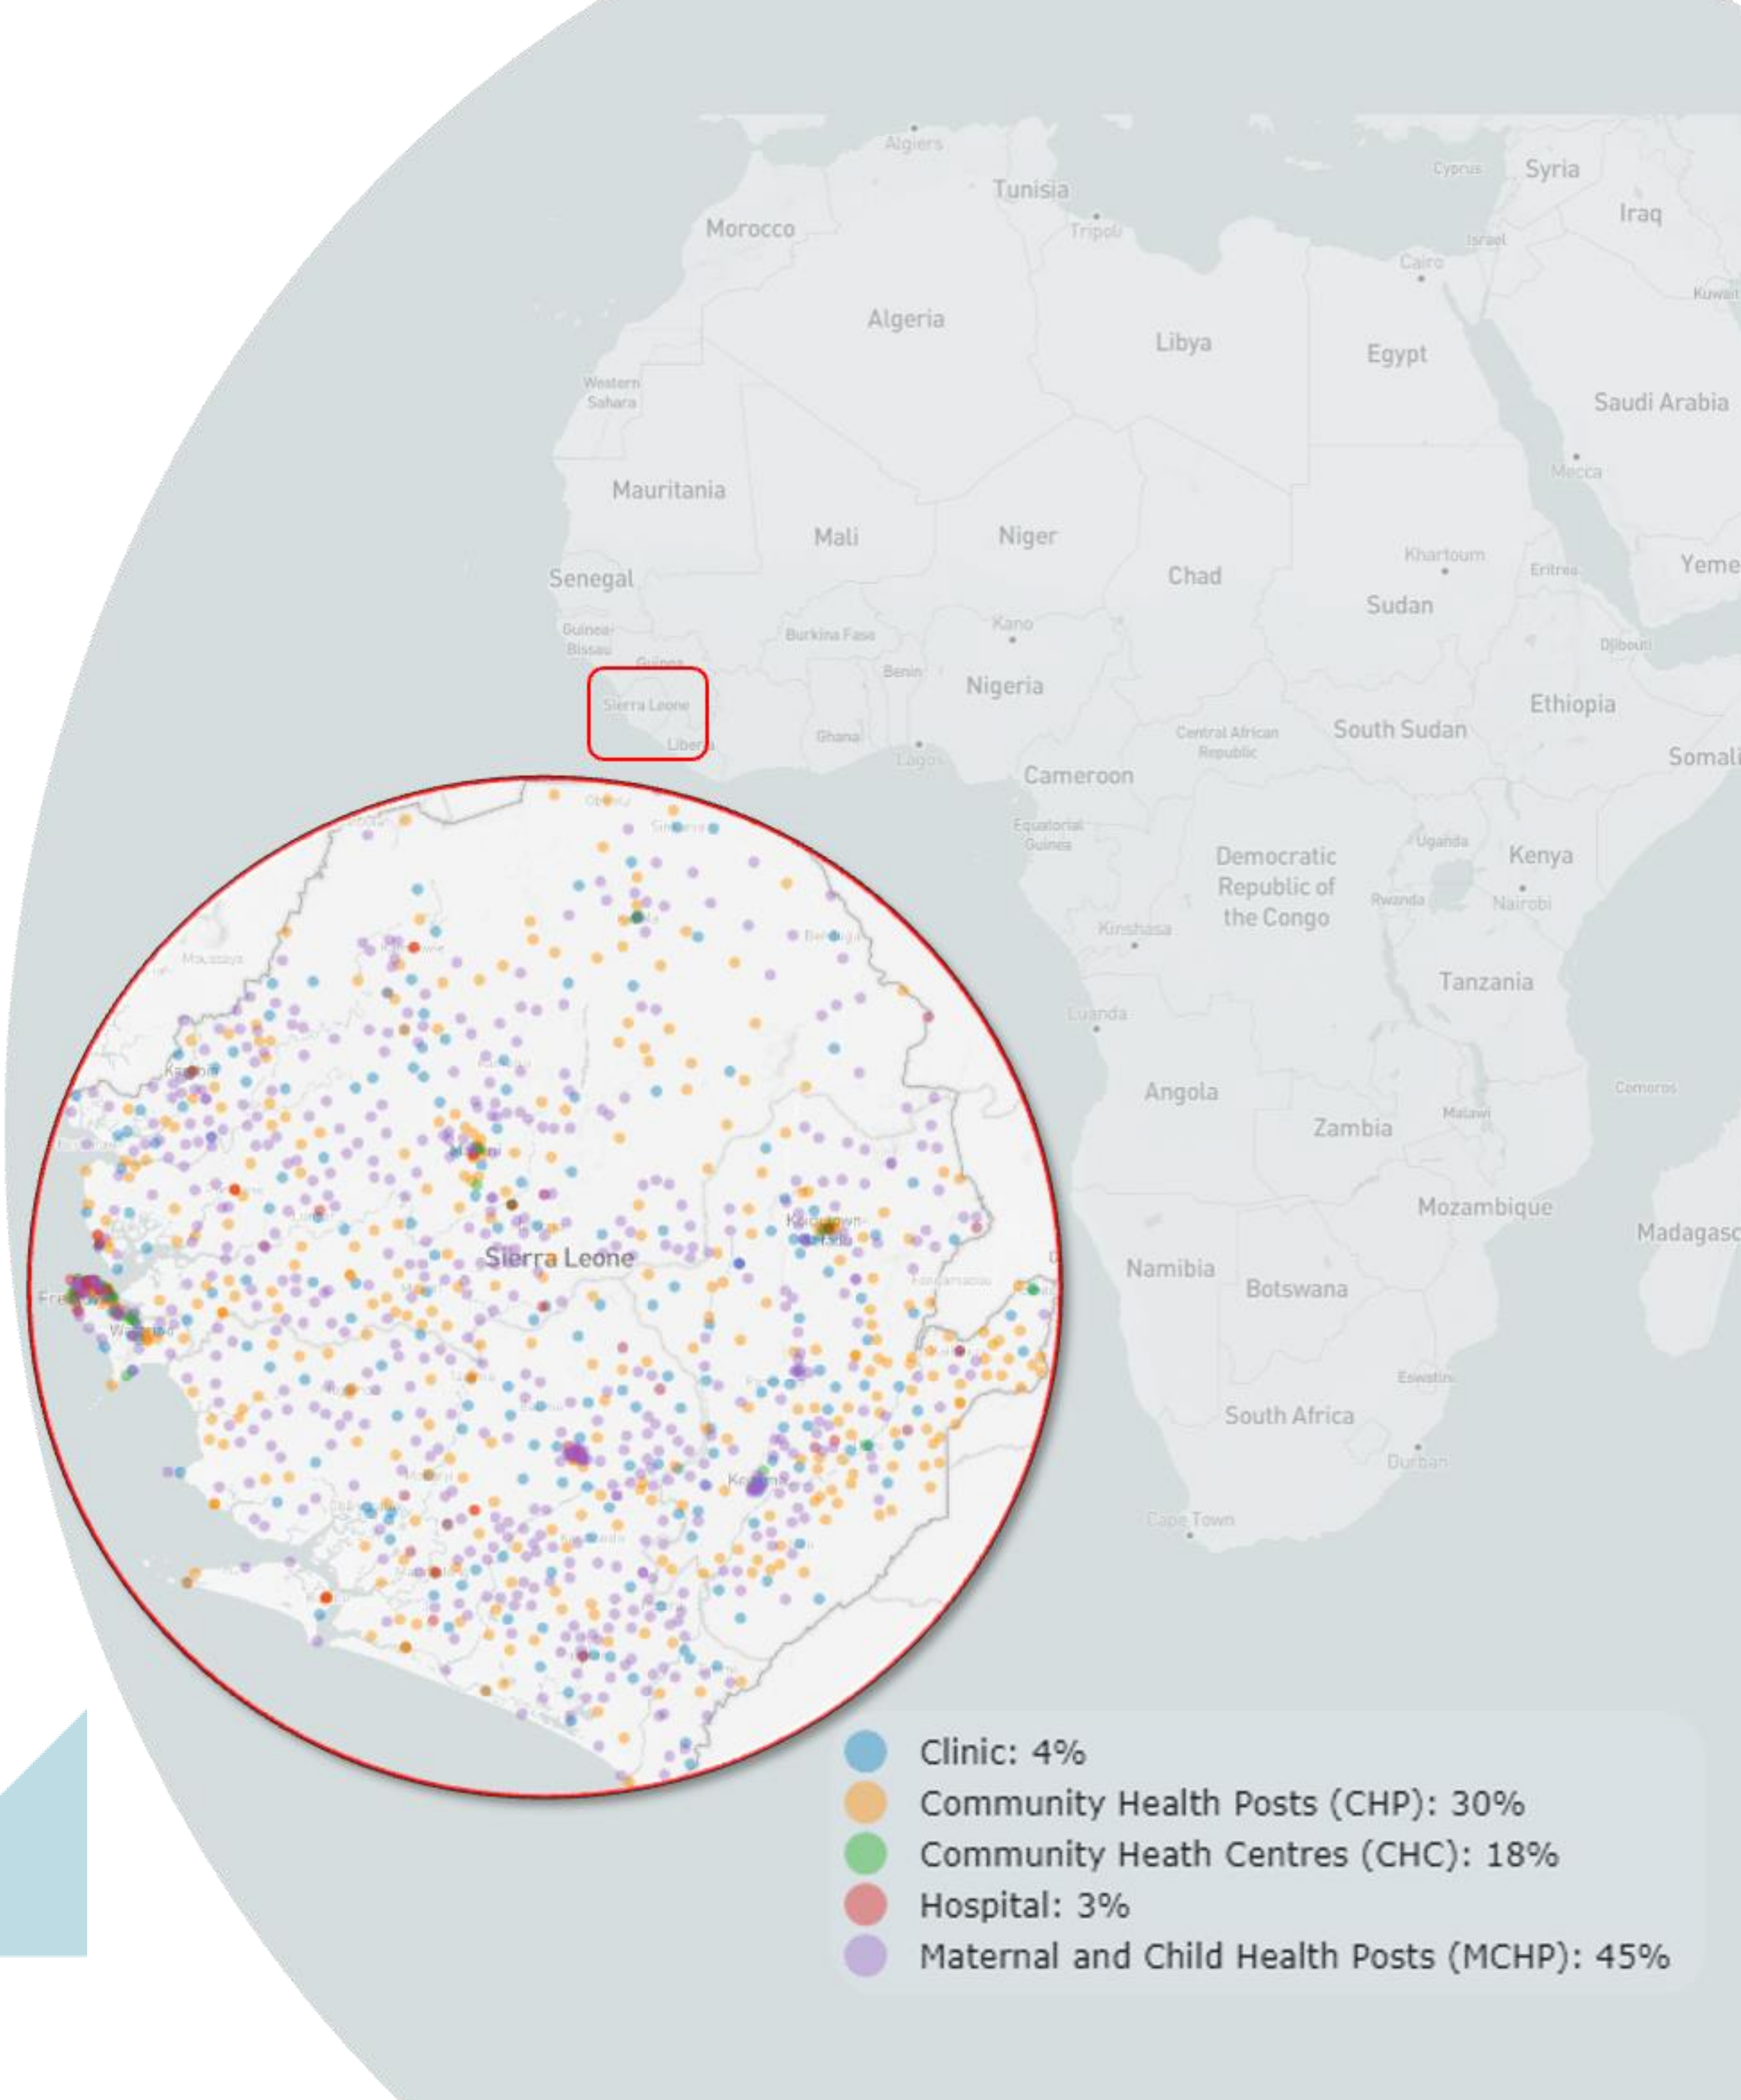

Context in Sierra Leone

- 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

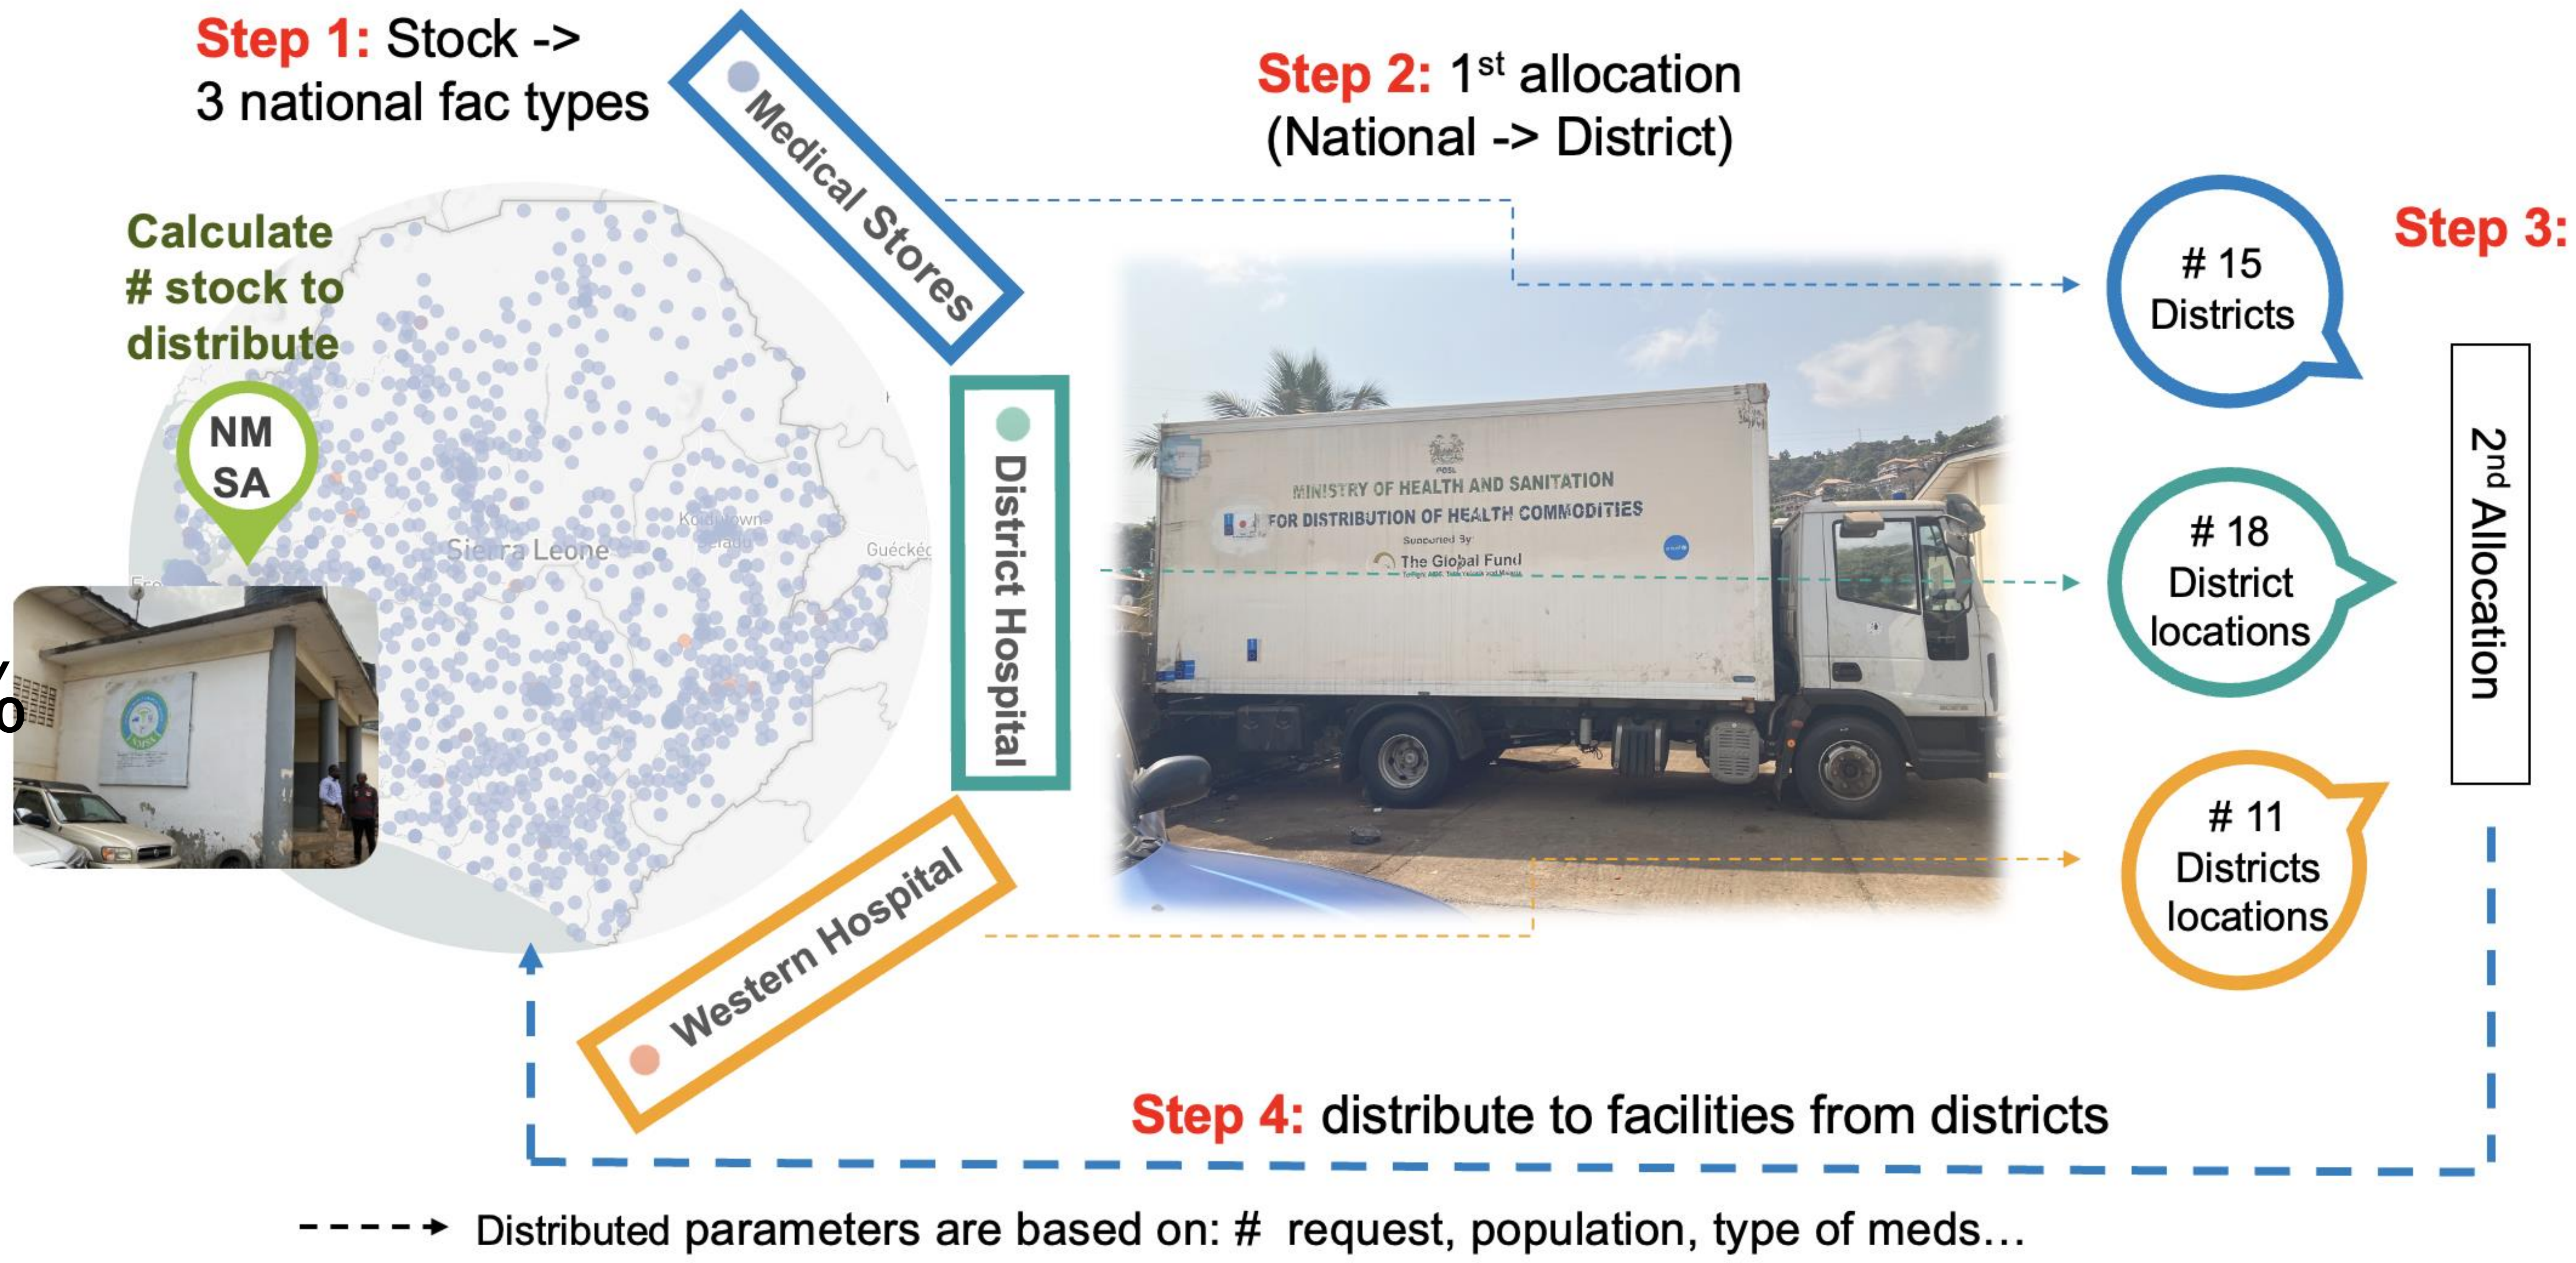

Wharton UNIVERSITY of PENNSYLVANIA

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
10000093	1,384,400	na	83,300	-	67,100	119,000	181,500	65,500	230,000
10000100	1,430,620	na	69,360	-	25,500	77,500	168,000	61,000	87,000
10000095	6,077,450	na	408,000	298,000	234,000	225,000	882,000	360,000	600,000
10000450	839,070	na	27,000	6,020	45,000	20,700	540,800	14,950	33,400
10000703	241,000	na	70,400	-	29,100	1,000	56,400	5,800	-
10000014	27,110	na	1,350	-	4,500	750	-	2,950	5,500
10000684	-	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

Code	DMS - Moyamba	DMS - Port Loko	DMS - Pujehun	DMS - Tonkolili	DMS - Western Area	DMS - Falaba	DMS - Karene
10000093	na	141,600	88,900	56,100	278,800	62,000	10,600
10000100	na	197,100	17,760	178,500	399,000	62,000	87,900
10000095	na	503,000	707,200	287,500	1,069,000	327,000	176,750
10000450	na	57,600	31,000	14,800	23,500	15,500	8,800
10000703	na	24,300	19,200	11,900	-	12,300	10,600
10000014	na	800	7,680	-	1,730	650	1,200
10000684	na	na	na	na	na	na	na
10000520	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:** 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.
- **Nationwide deployment**: 1,058 facilities (all public-owned facilities)
- **Results**: **19%** increase of consumption to medicines for 2 million women & children under 5

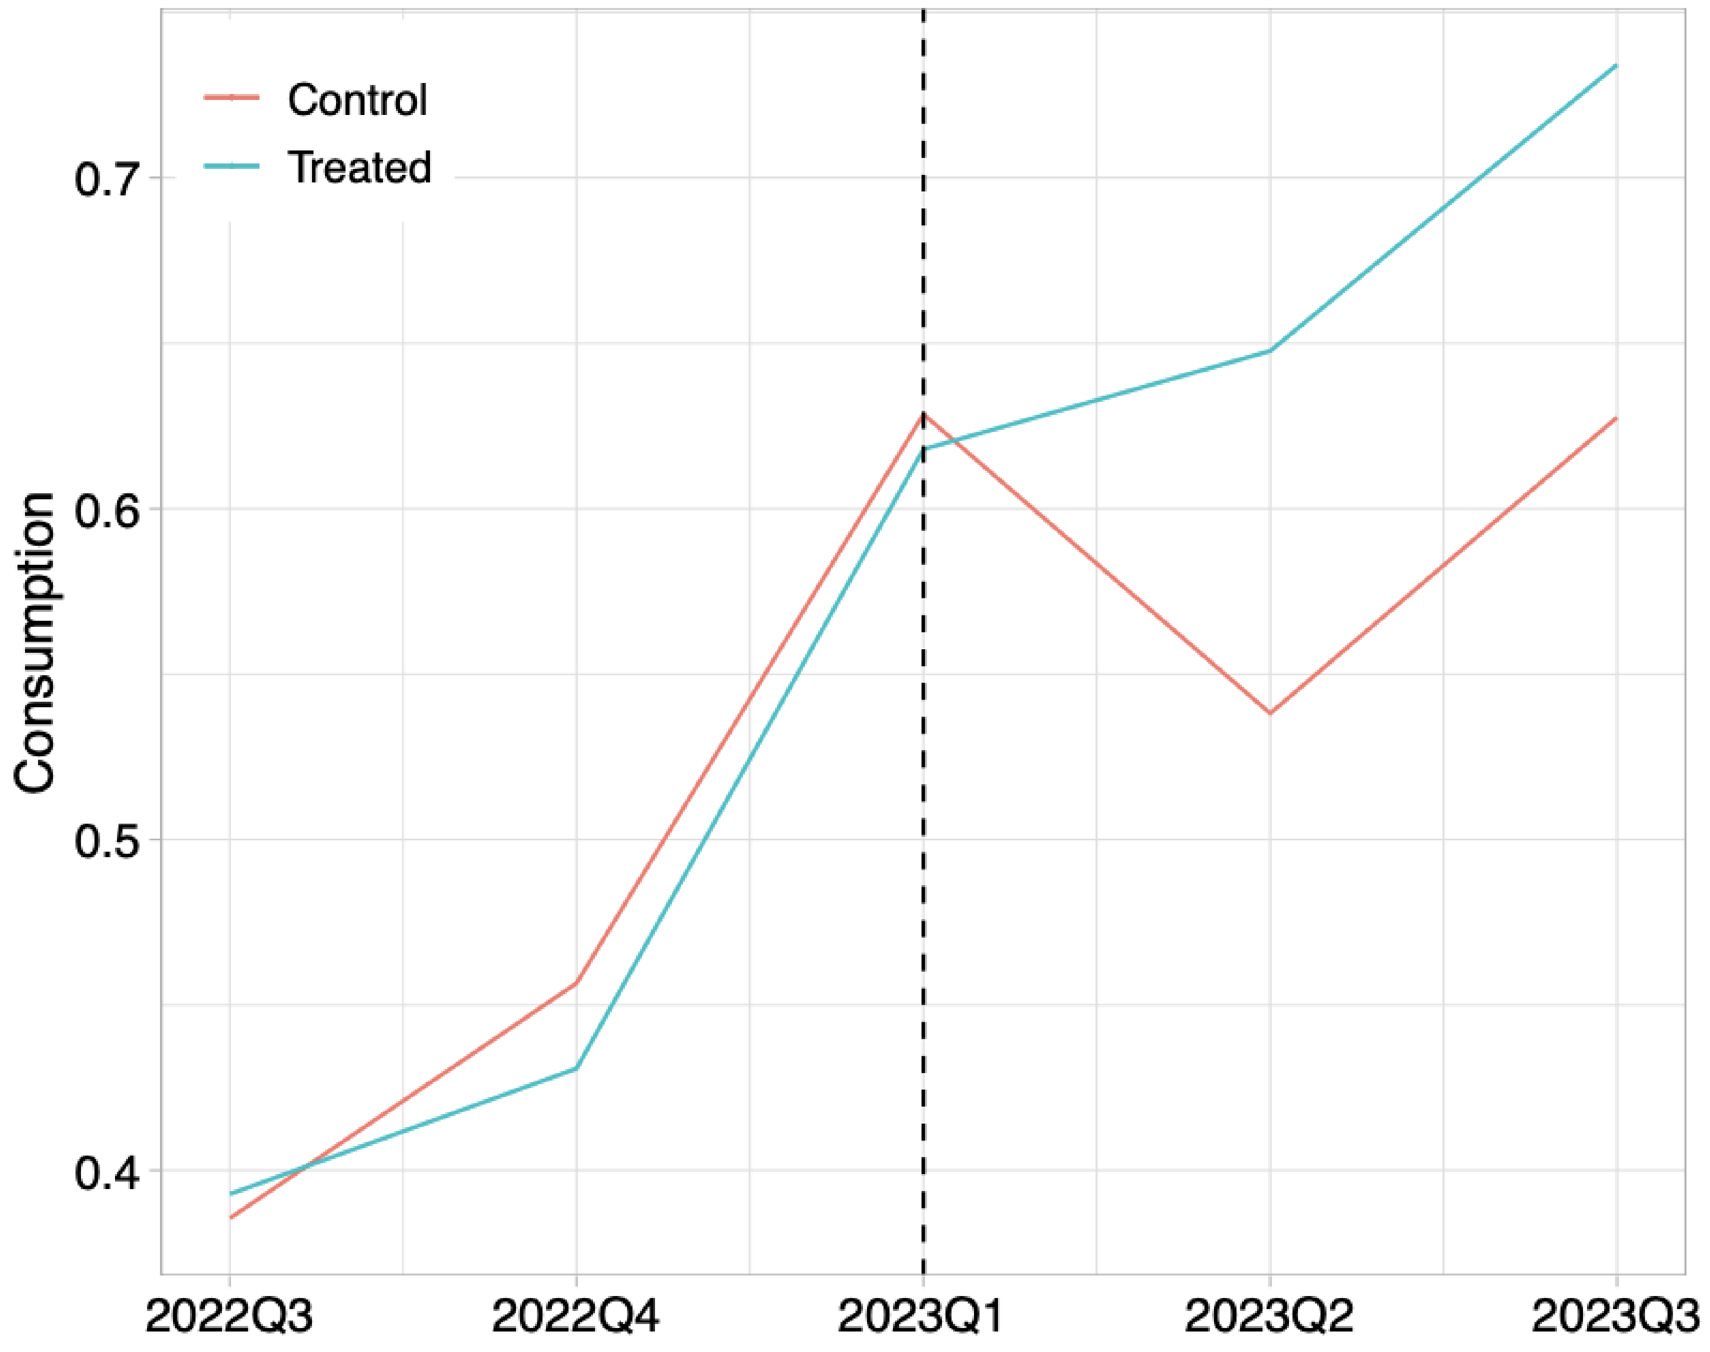

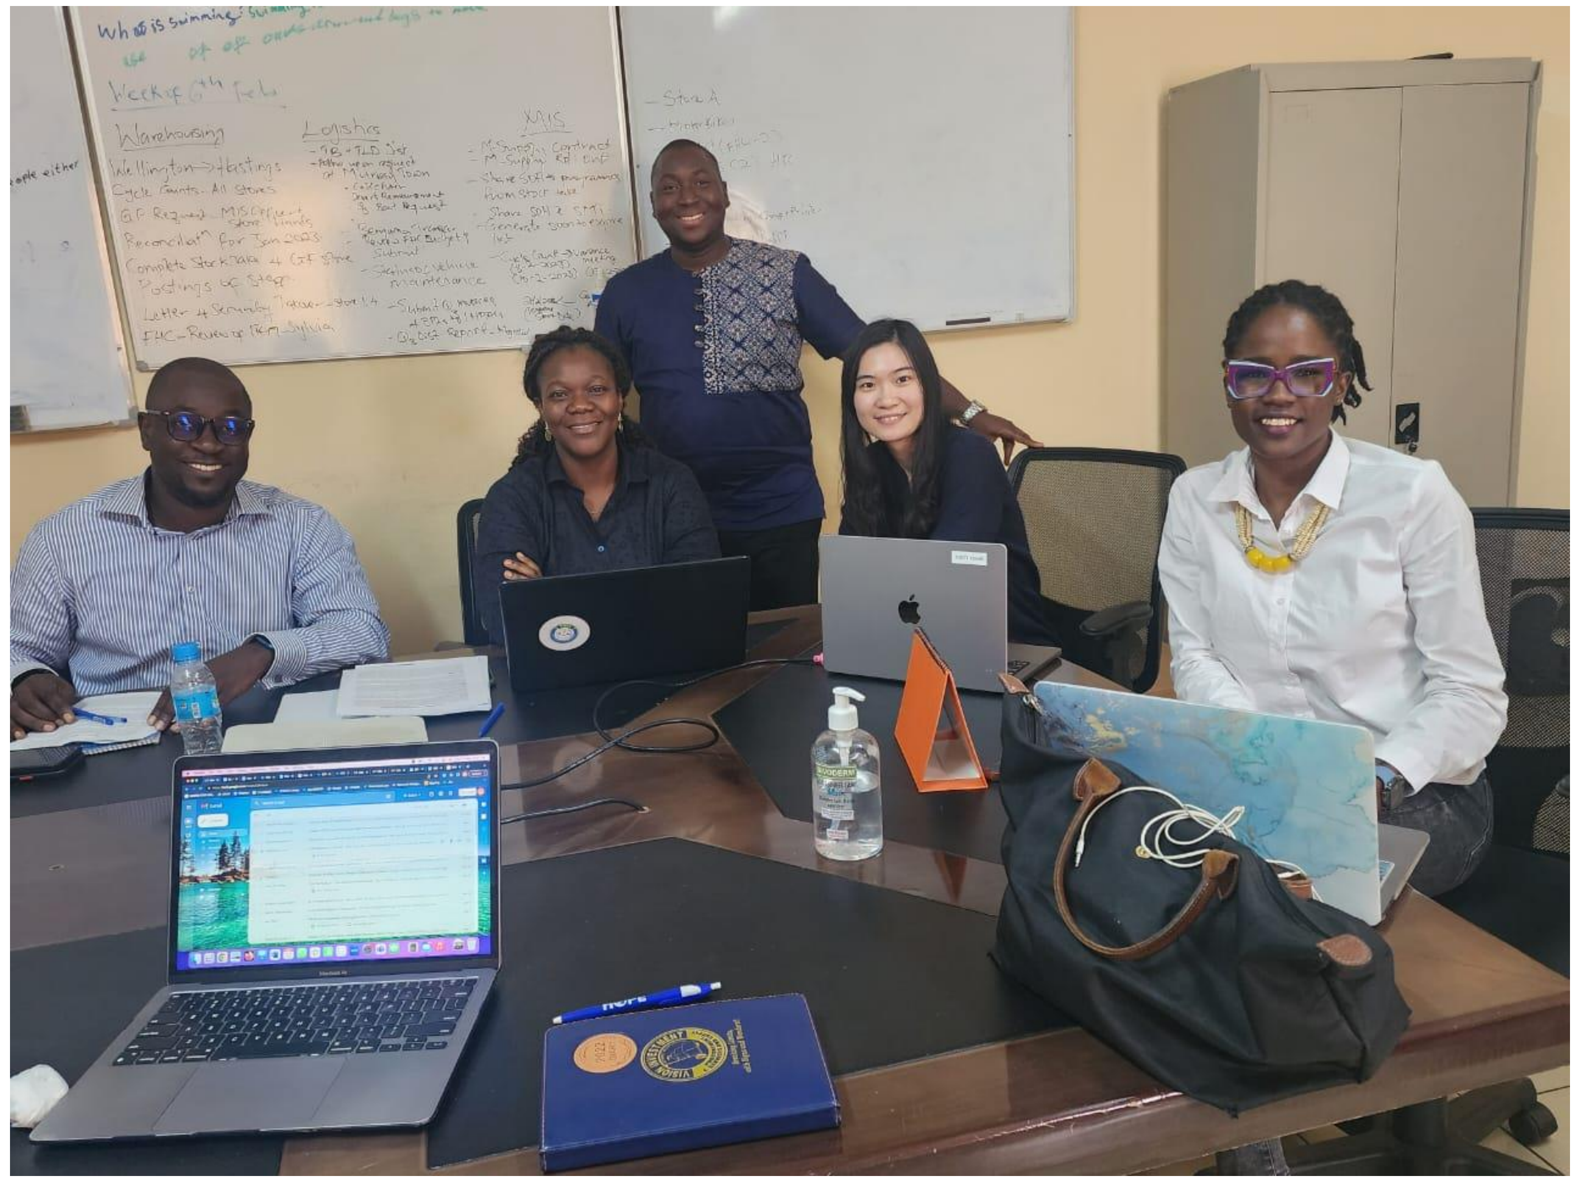

Visit & discuss at Sierra Leone National Medical Supplies Agency's Office

Predict-then-Optimize? (decision-blind)

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

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

Step 2: Optimize based on predictions

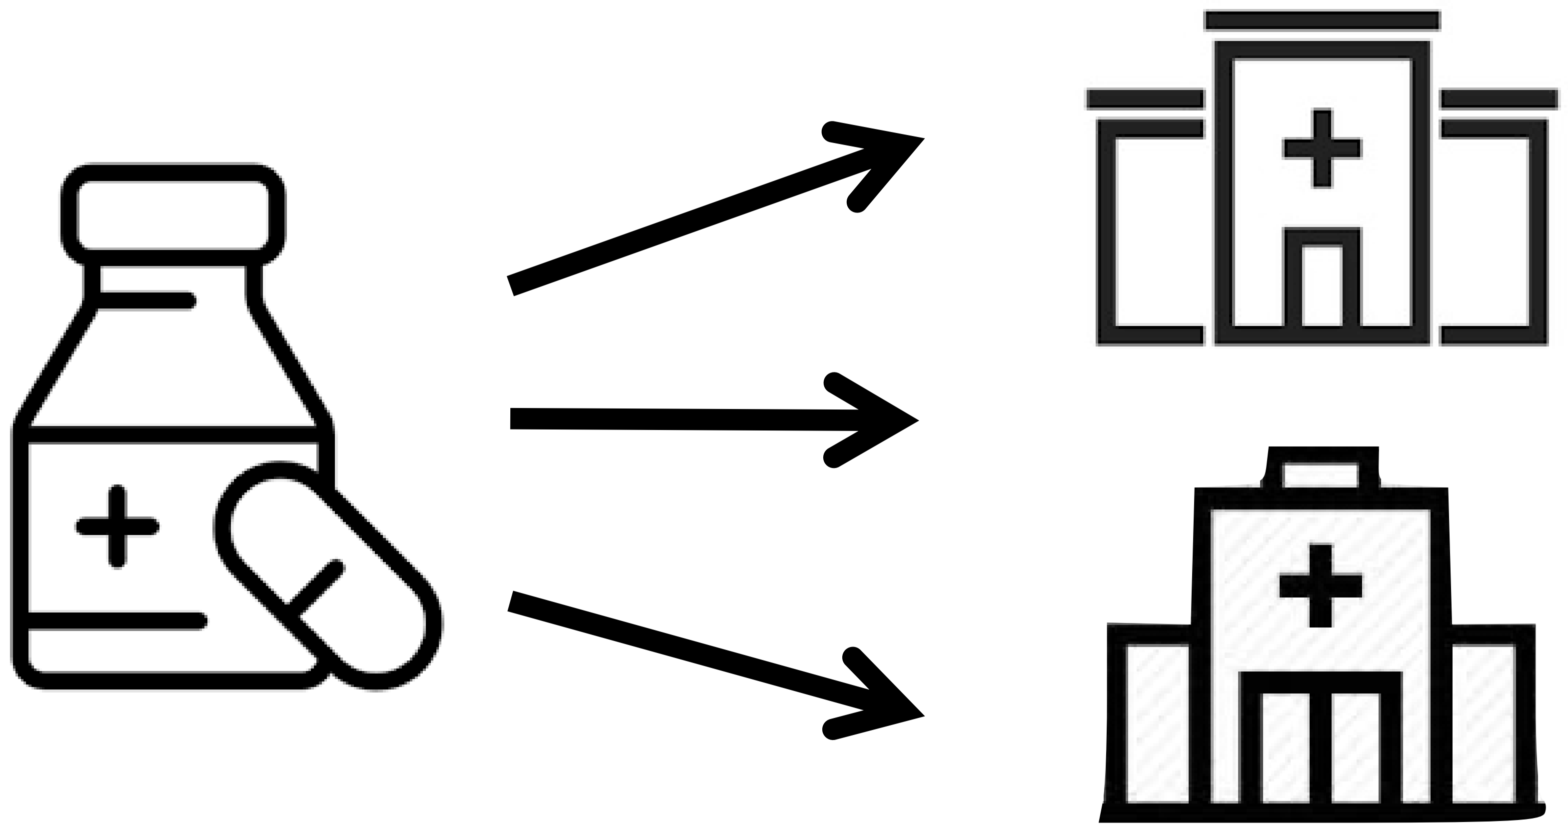

Problem:

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

Why not use decision-blind?

Toy example:

- 2 kinds of facilities: well-stocked & poorly-stocked
- Goal: allocate new inventory to meet the demand
- No point predicting well on well-stocked facilities; only care about poorly-stocked facilities

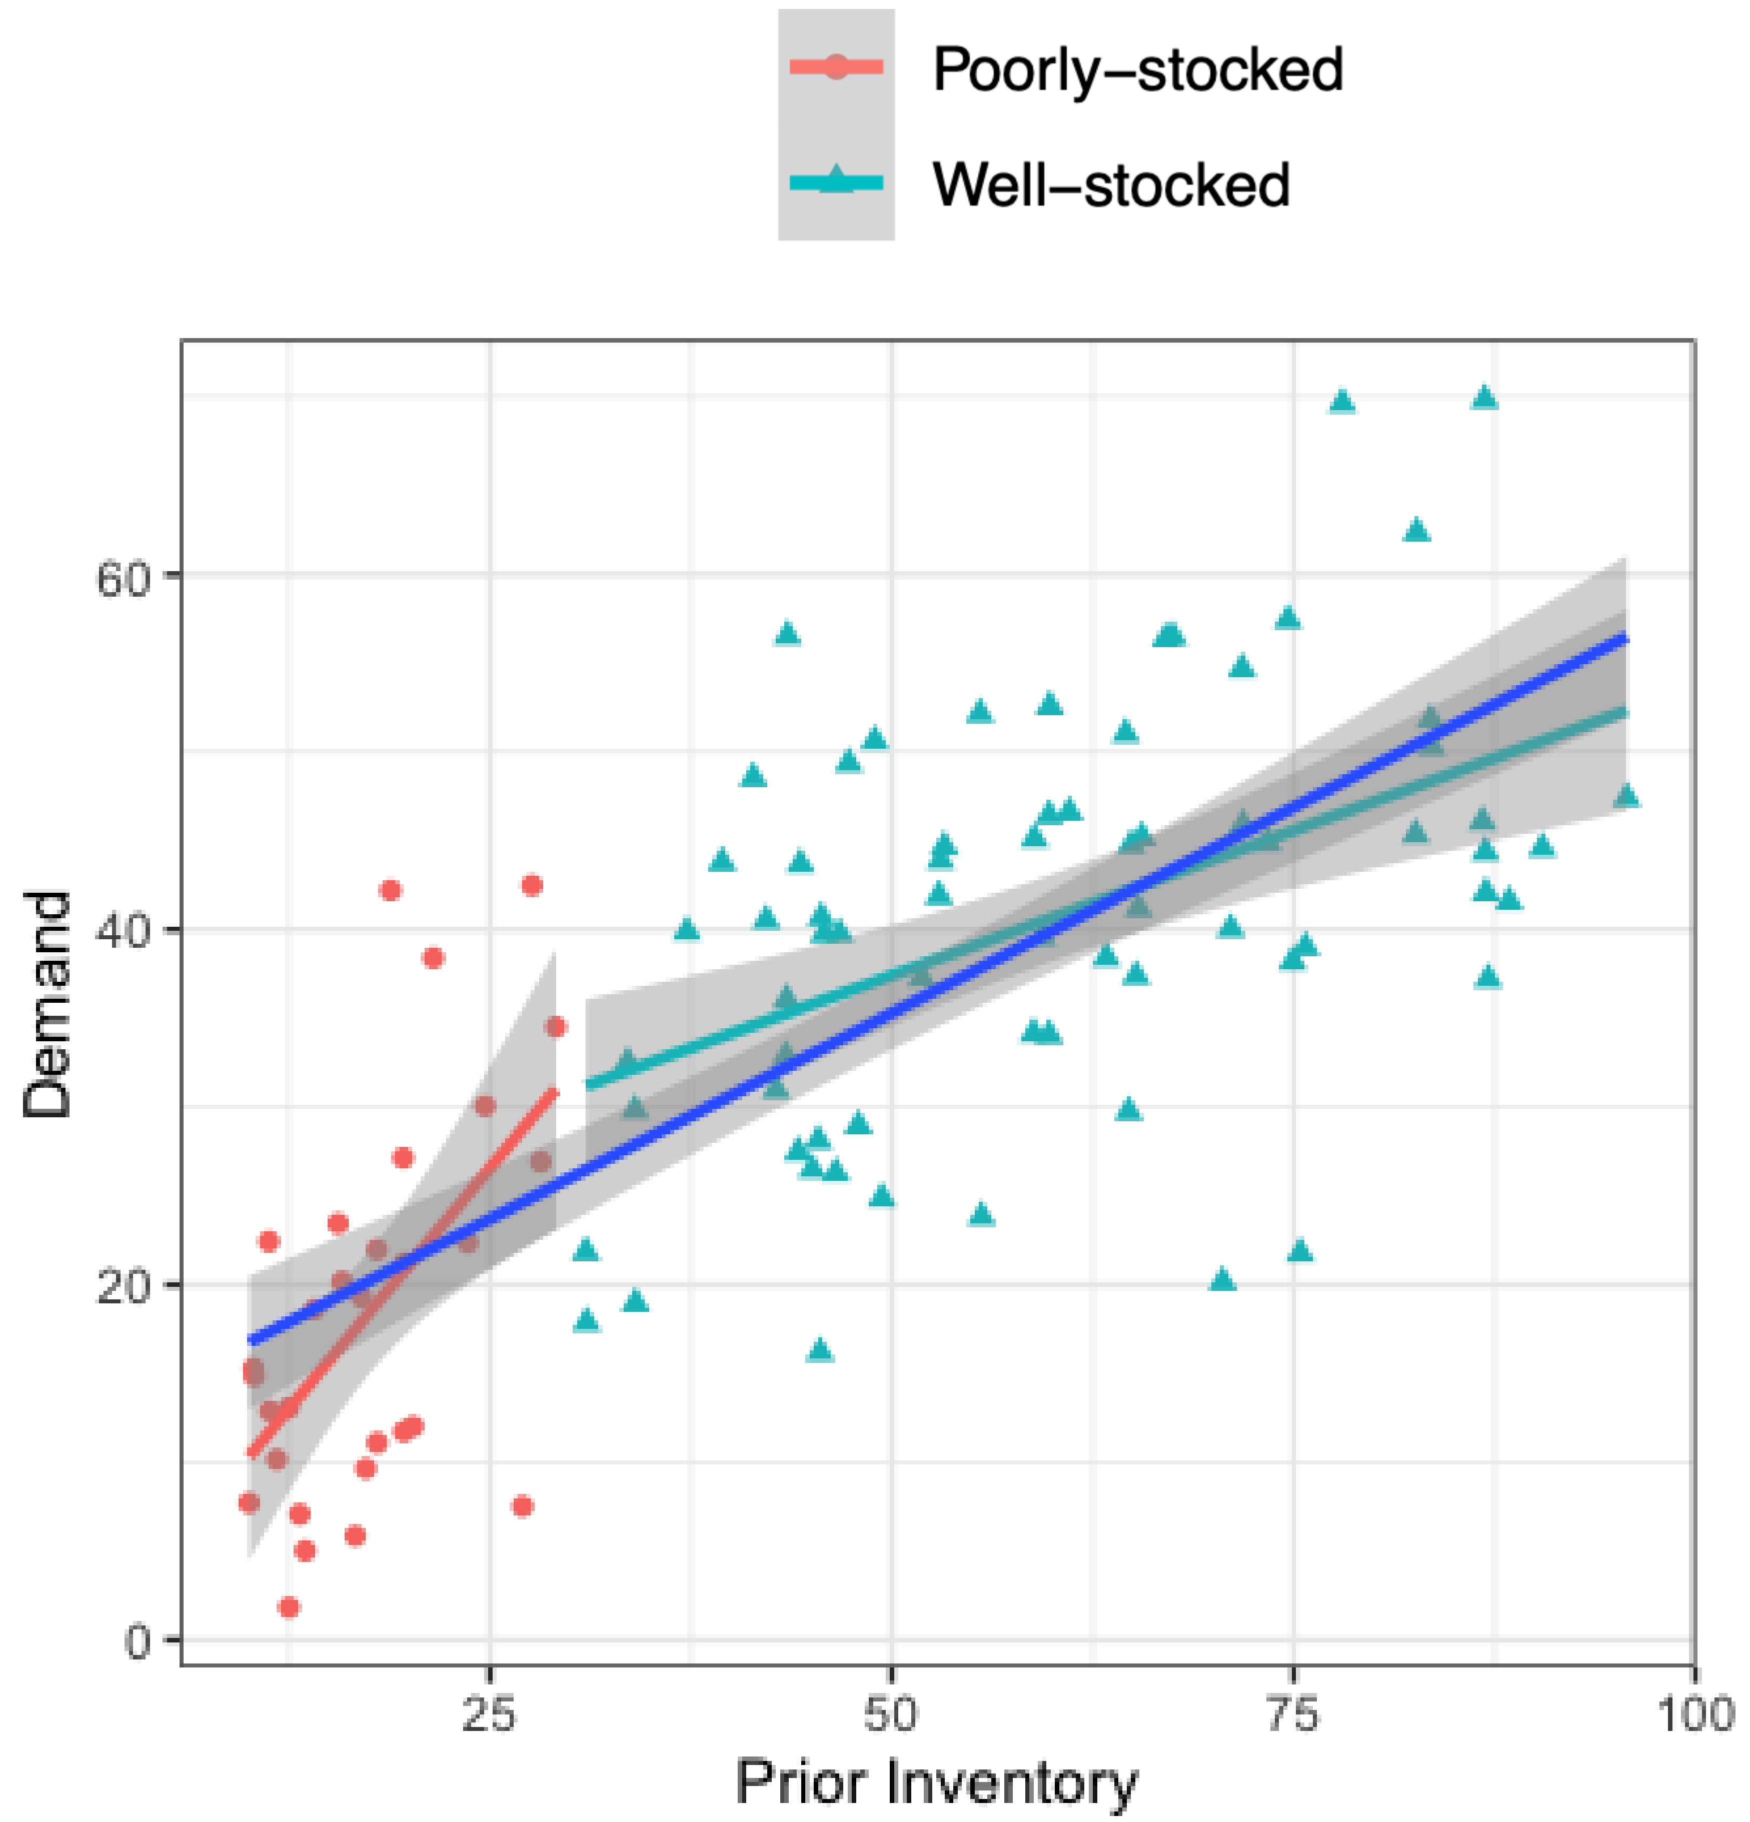

Related work: Decision-Aware learning

- **[Kallus & Mao (2022)]** Specialized approach for random forests + SAA
  - Not general: strategy specific to tree models and one unknown parameter vector per optimization
  - Not computationally tractable: every tree split requires re-solving optimization problem

- **[Wilder et al (2019a/b)]** Backpropagate through decision loss
  - Not general: prediction function must be differentiable
  - Not computationally tractable: need to solve optimization at every gradient step

- **[Elmachtoub & Grigas (2022)]** Approximate decision-aware prediction loss when it is a LP and prediction is linear
  - Not general: strategy specific to linear models and known constraints

Learning and optimization

- Objective:

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

known decision loss
(unmet demand)

optimal
allocation

demand
(unknown)

- 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

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

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?

- **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**

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

Wharton University of Pennsylvania

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]

Stochastic optimization

- **Goal:** allocations $\boldsymbol{a}_{\boldsymbol{n}}^*$ across all N facilities that minimizes cost
- **Objective:** sum of **unmet demand** across facilities

$$\ell_n = \boxed{\sum_n \mathbb{E}[\max\{\xi_n - s_n - a_n, 0\}]}$$

  - Current inventory $\boldsymbol{s}_{\boldsymbol{n}}$, demand $\boldsymbol{\xi}_{\boldsymbol{n}}$
- **Constraints:** fixed budget $\boldsymbol{b}$
- **Predictions** ($\hat{y}$): draw random demands $\xi_i^{(k)}$ at each facility based on estimated distribution from historical data
- Right box: our LP formulation with sample average approximation (SAA)

$$\boxed{a^* = \arg\min_{a \in \mathbb{R}^N_{\geq 0}} \sum_{k=1}^{K} \sum_{n=1}^{N} \ell_n^{(k)}}$$

$$\text{subj. to} \quad \sum_{n=1}^{N} a_n \leq b$$

$$\ell^{(k)} \geq \xi^{(k)} - s - a$$

$$\ell^{(k)} \geq 0$$

Wharton
UNIVERSITY of PENNSYLVANIA

Decision-Aware training objective

- Use gradients to obtain approximated predictive model objective:

$$\arg\min_{\theta} \frac{1}{K}\sum_{k=1}^{K}\sum_{n=1}^{N}\left(\mathbb{I}\left(\xi_n^{(k)} \geq s_n + a_n\right) + const\right)\cdot\left|f_\theta(x_n) - \xi_n^{(k)}\right|$$

- 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)

- Recall the toy example: we obtain the right weight

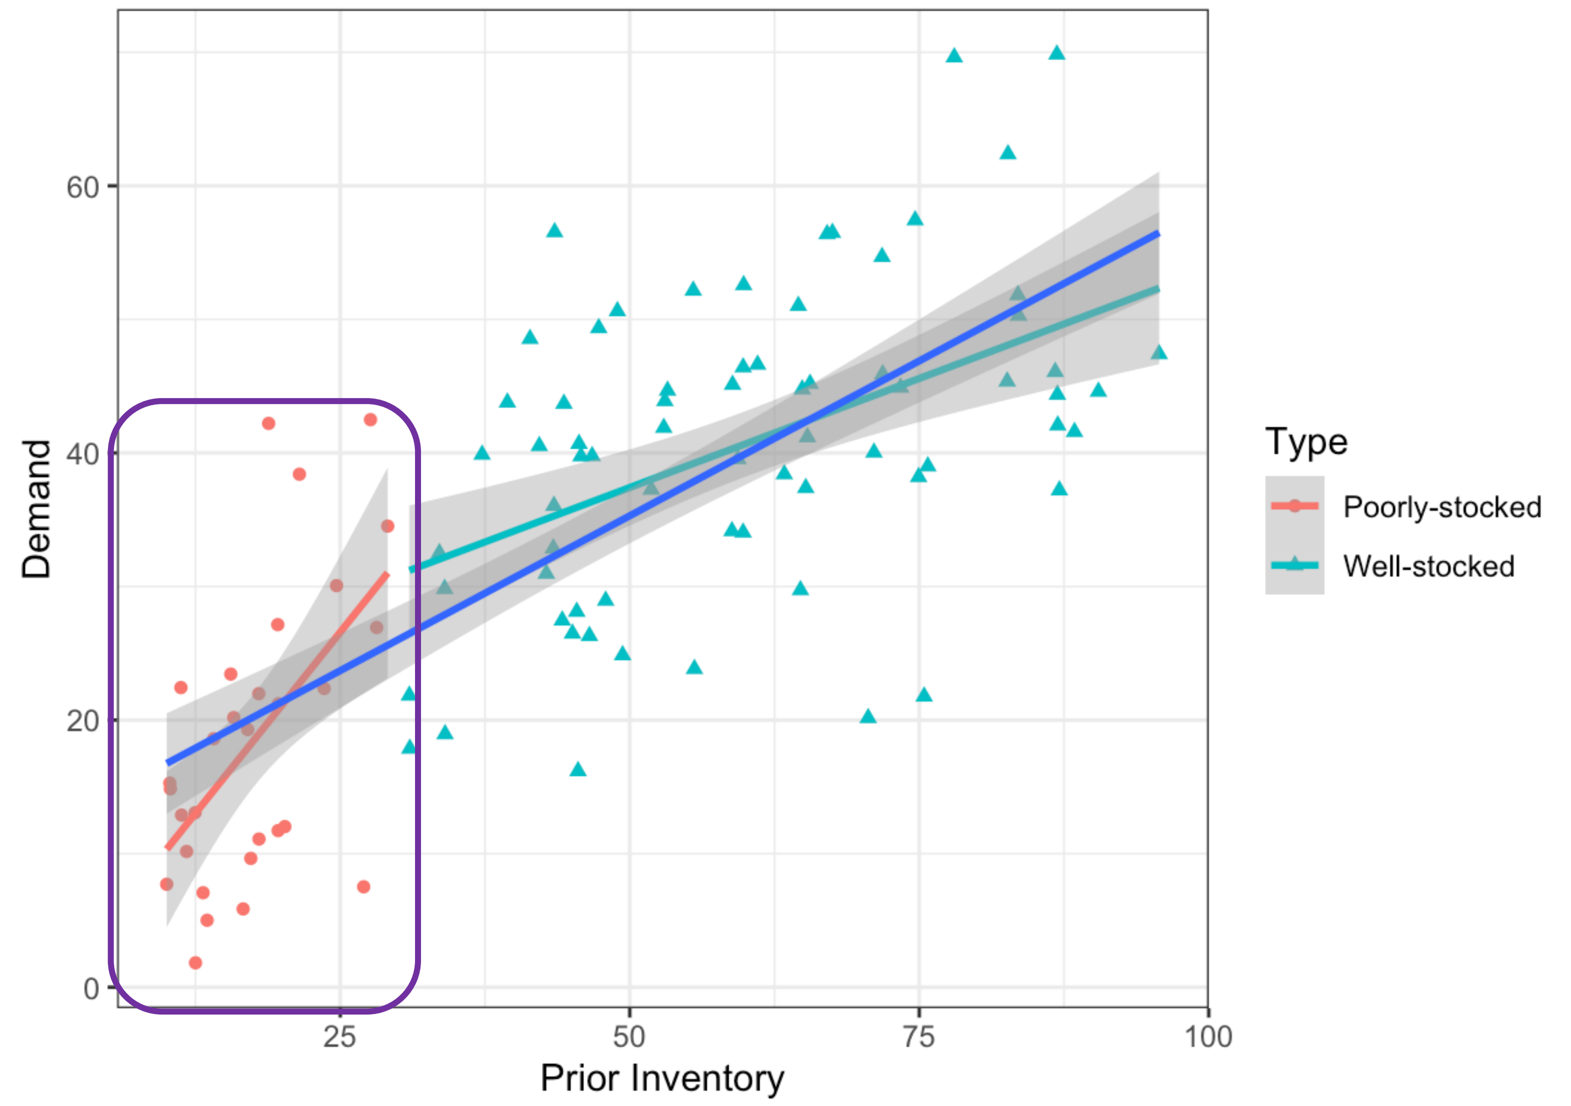

Our Decision-Aware approach

- **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)

$\Downarrow$

- **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)…

- **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

Overall System

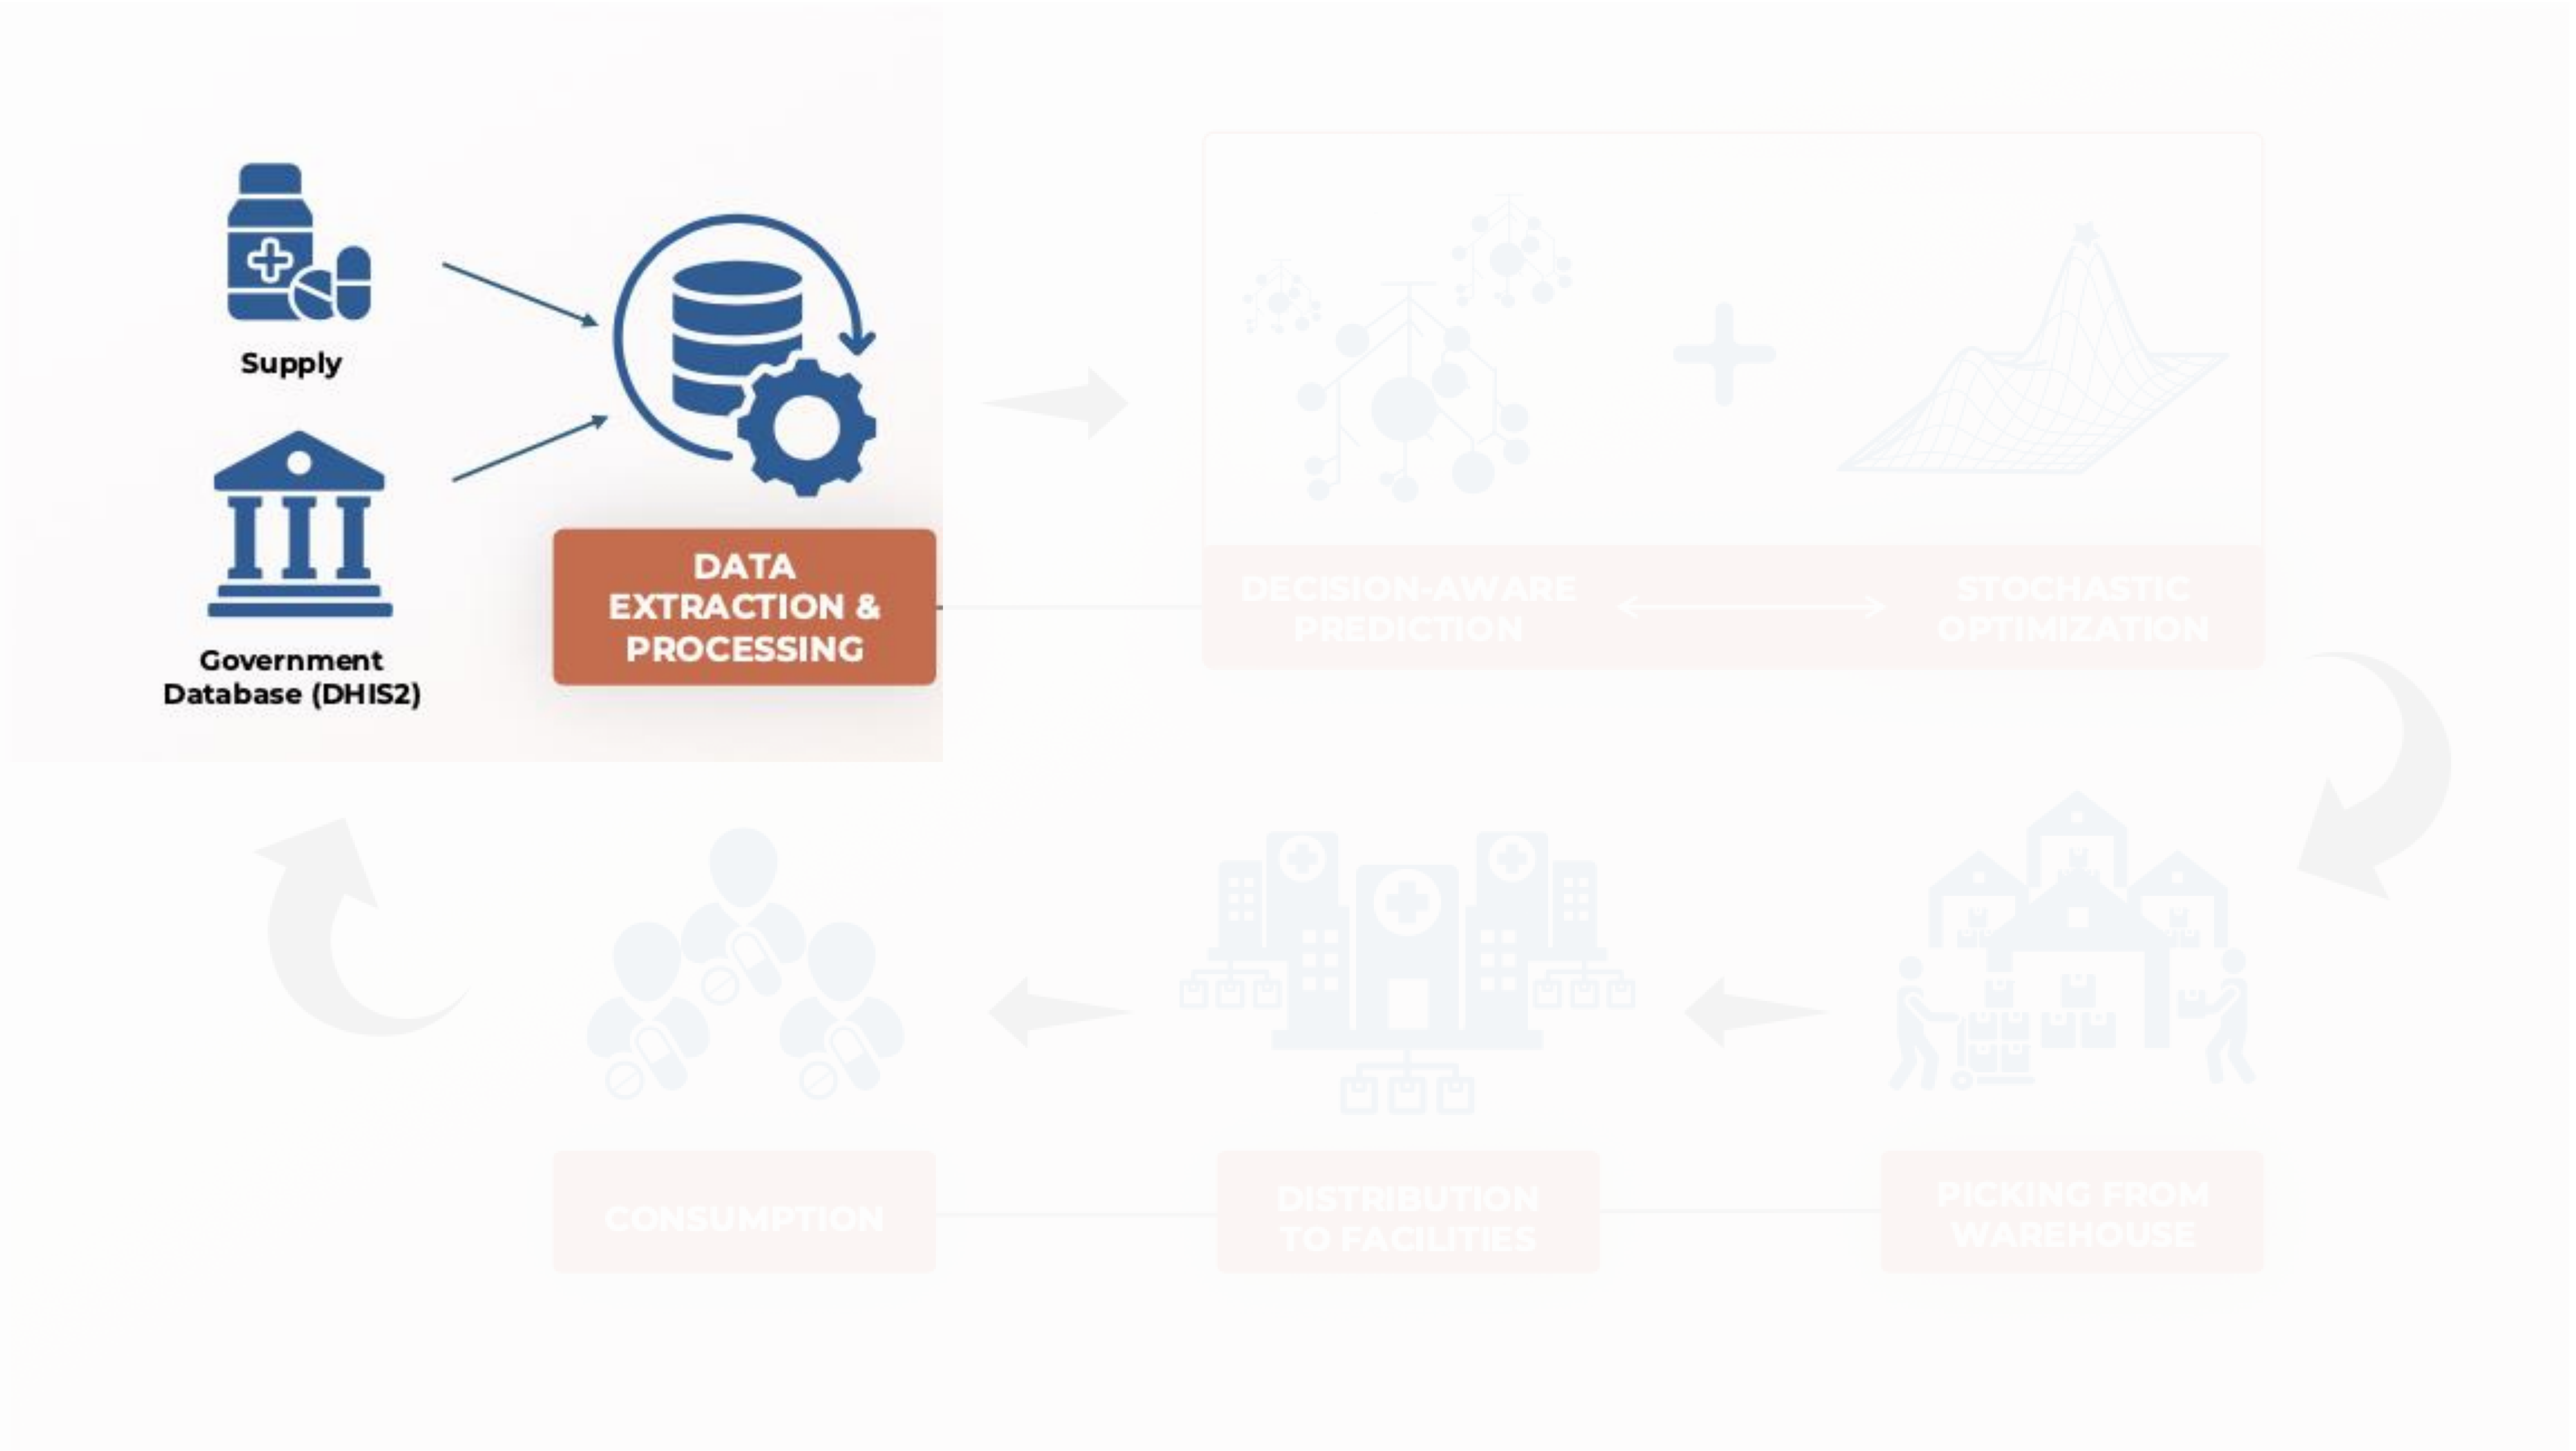

Overall System

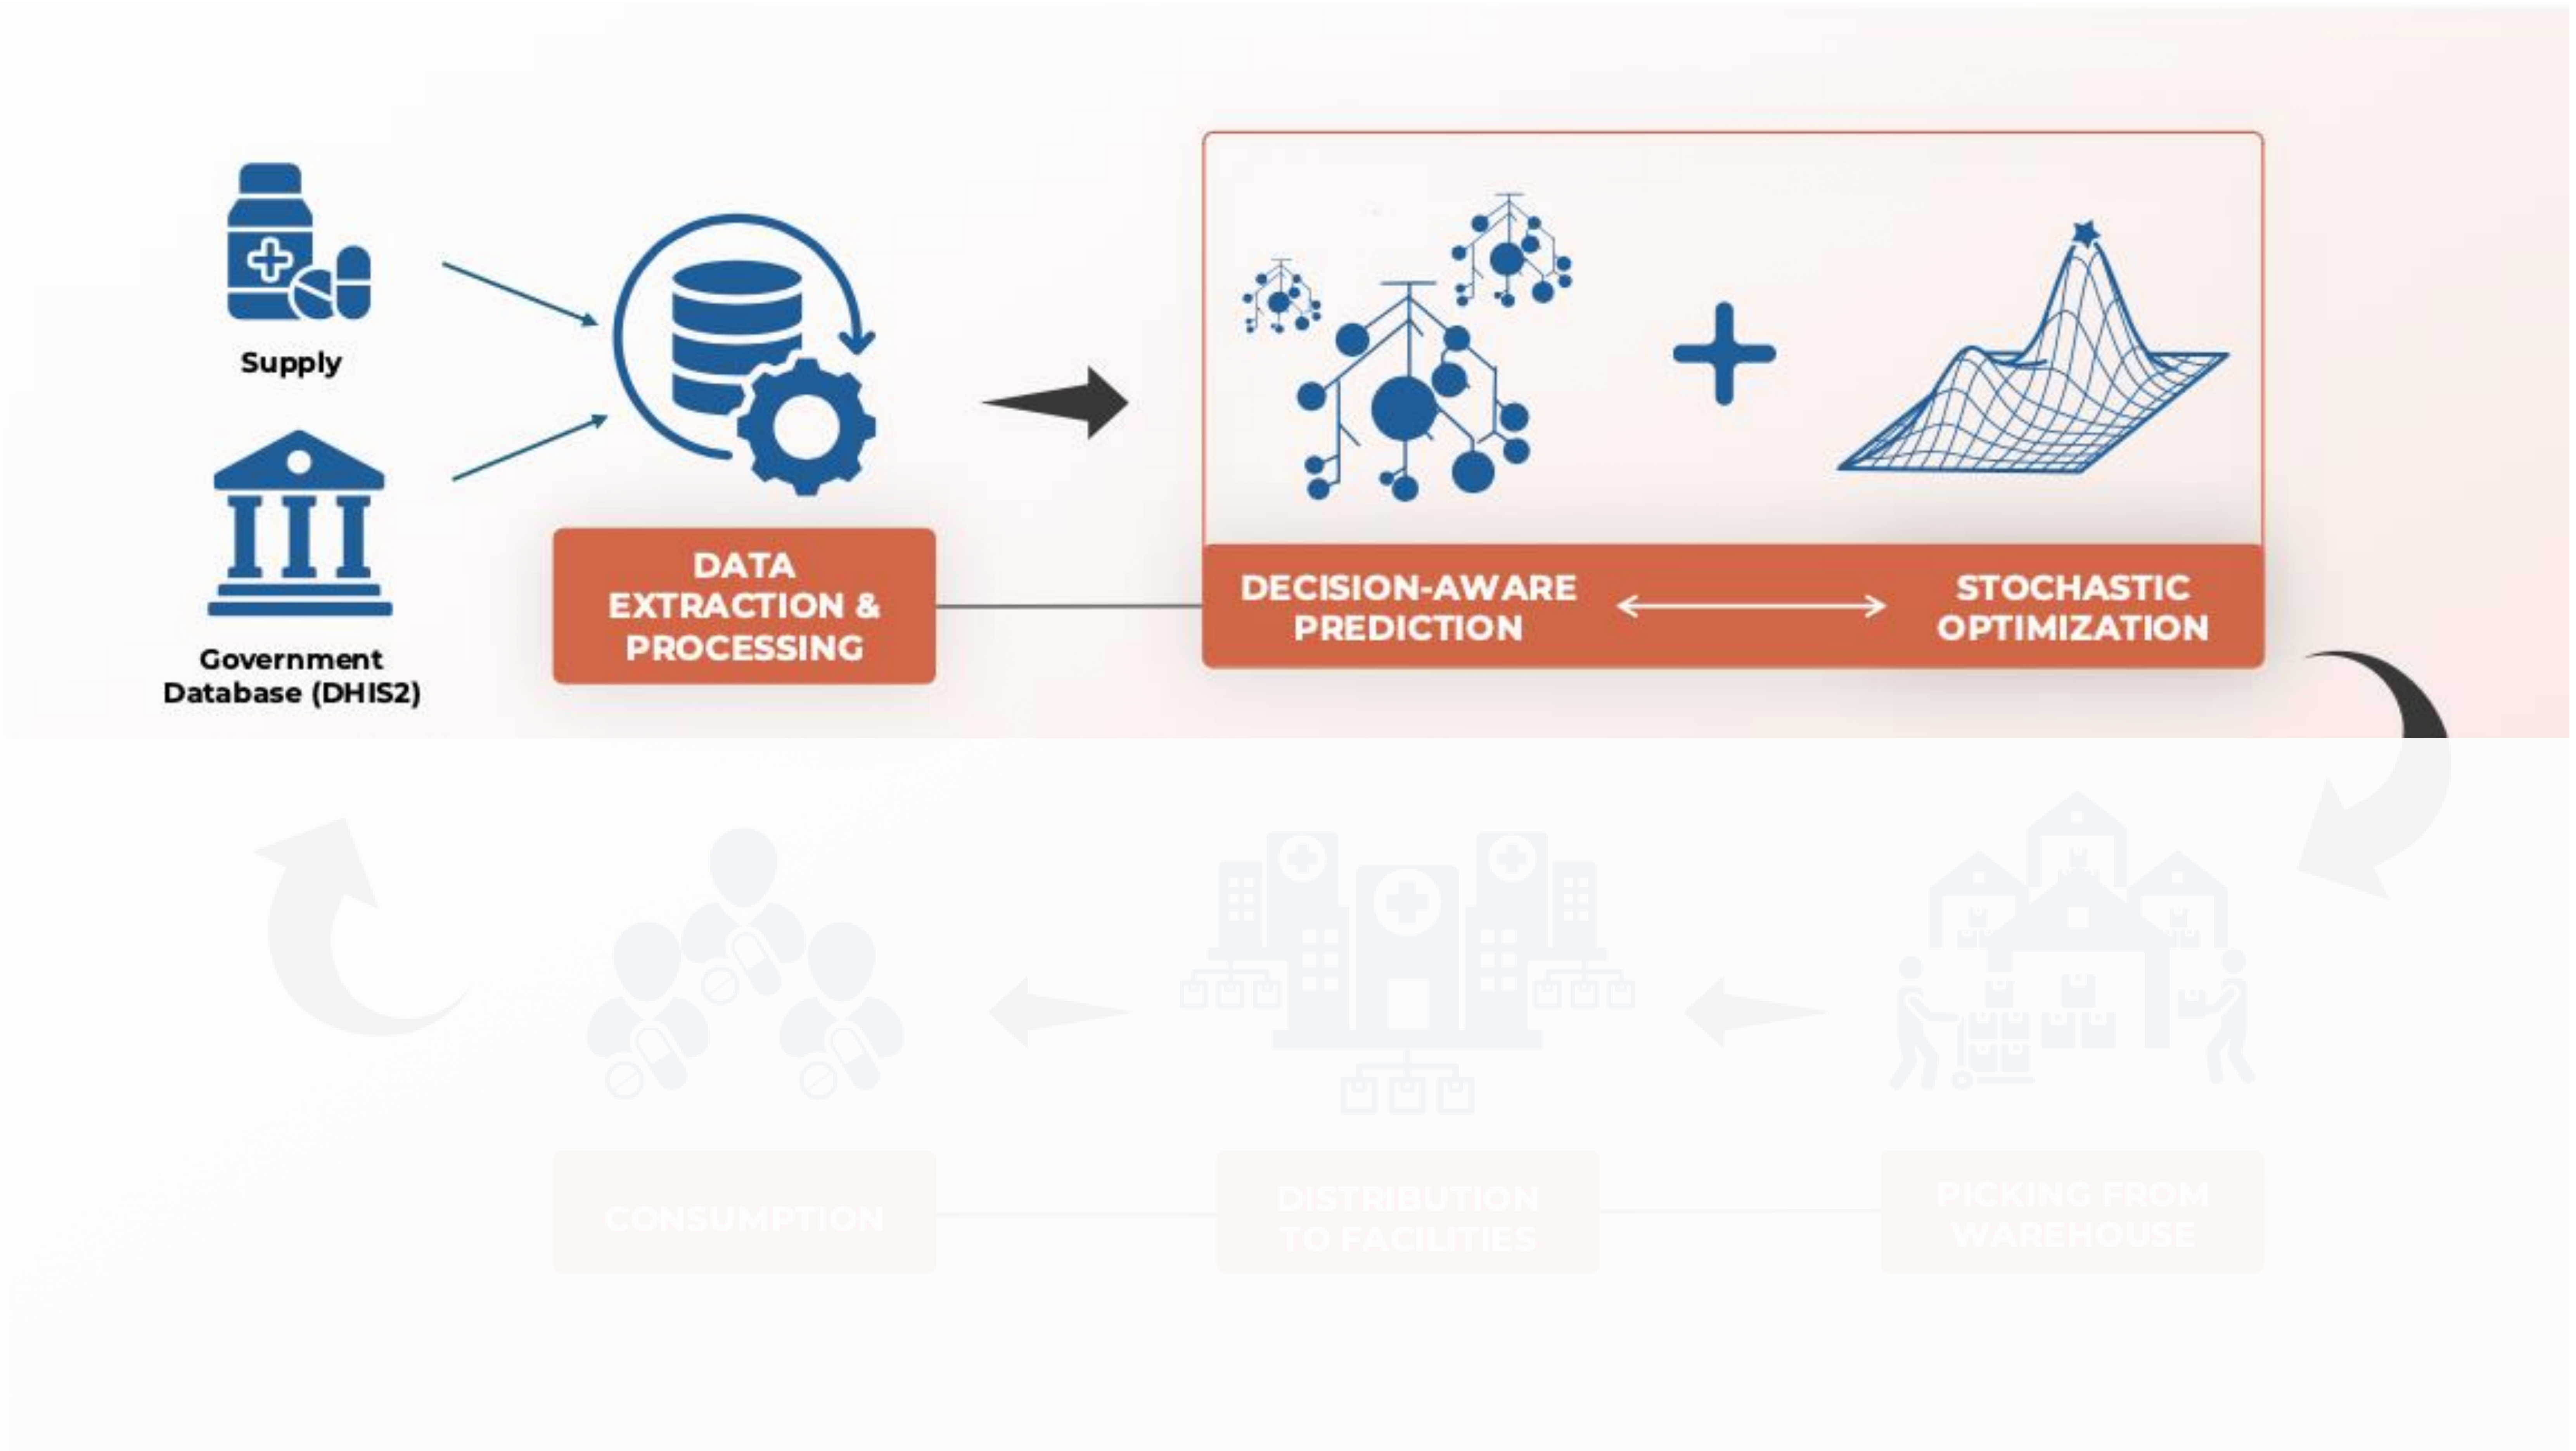

Overall System

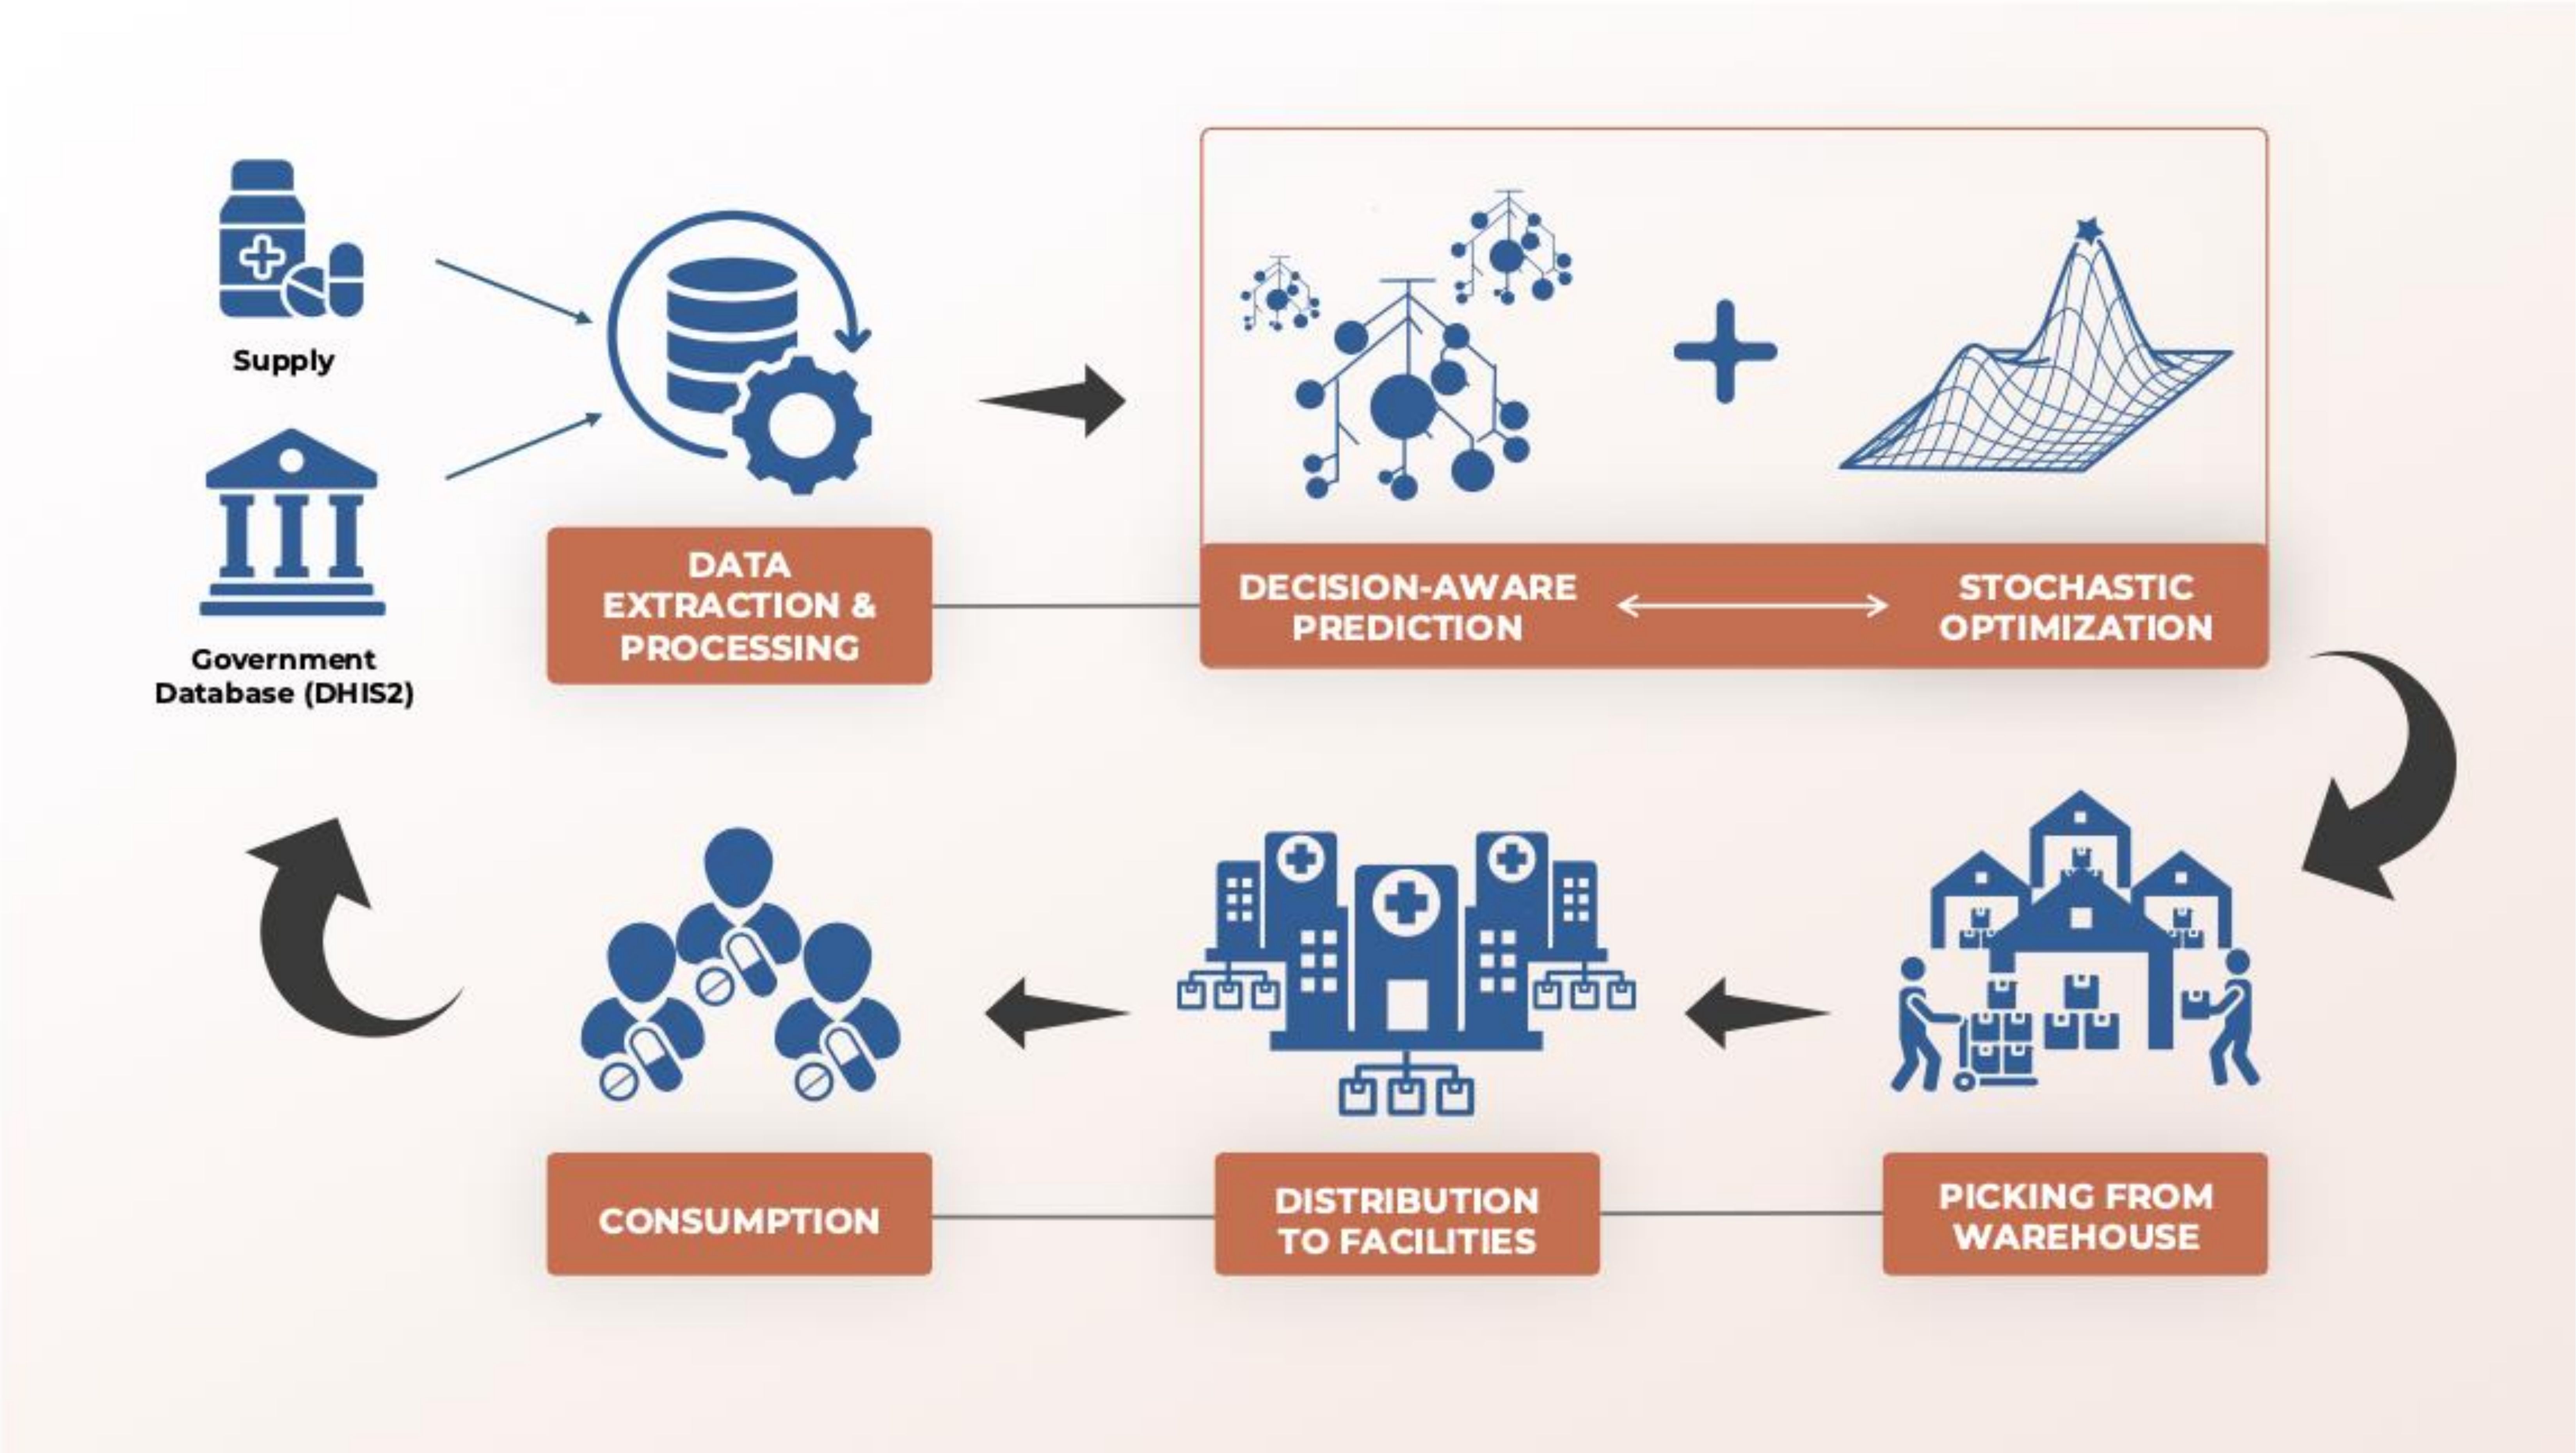

Deployment

- 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

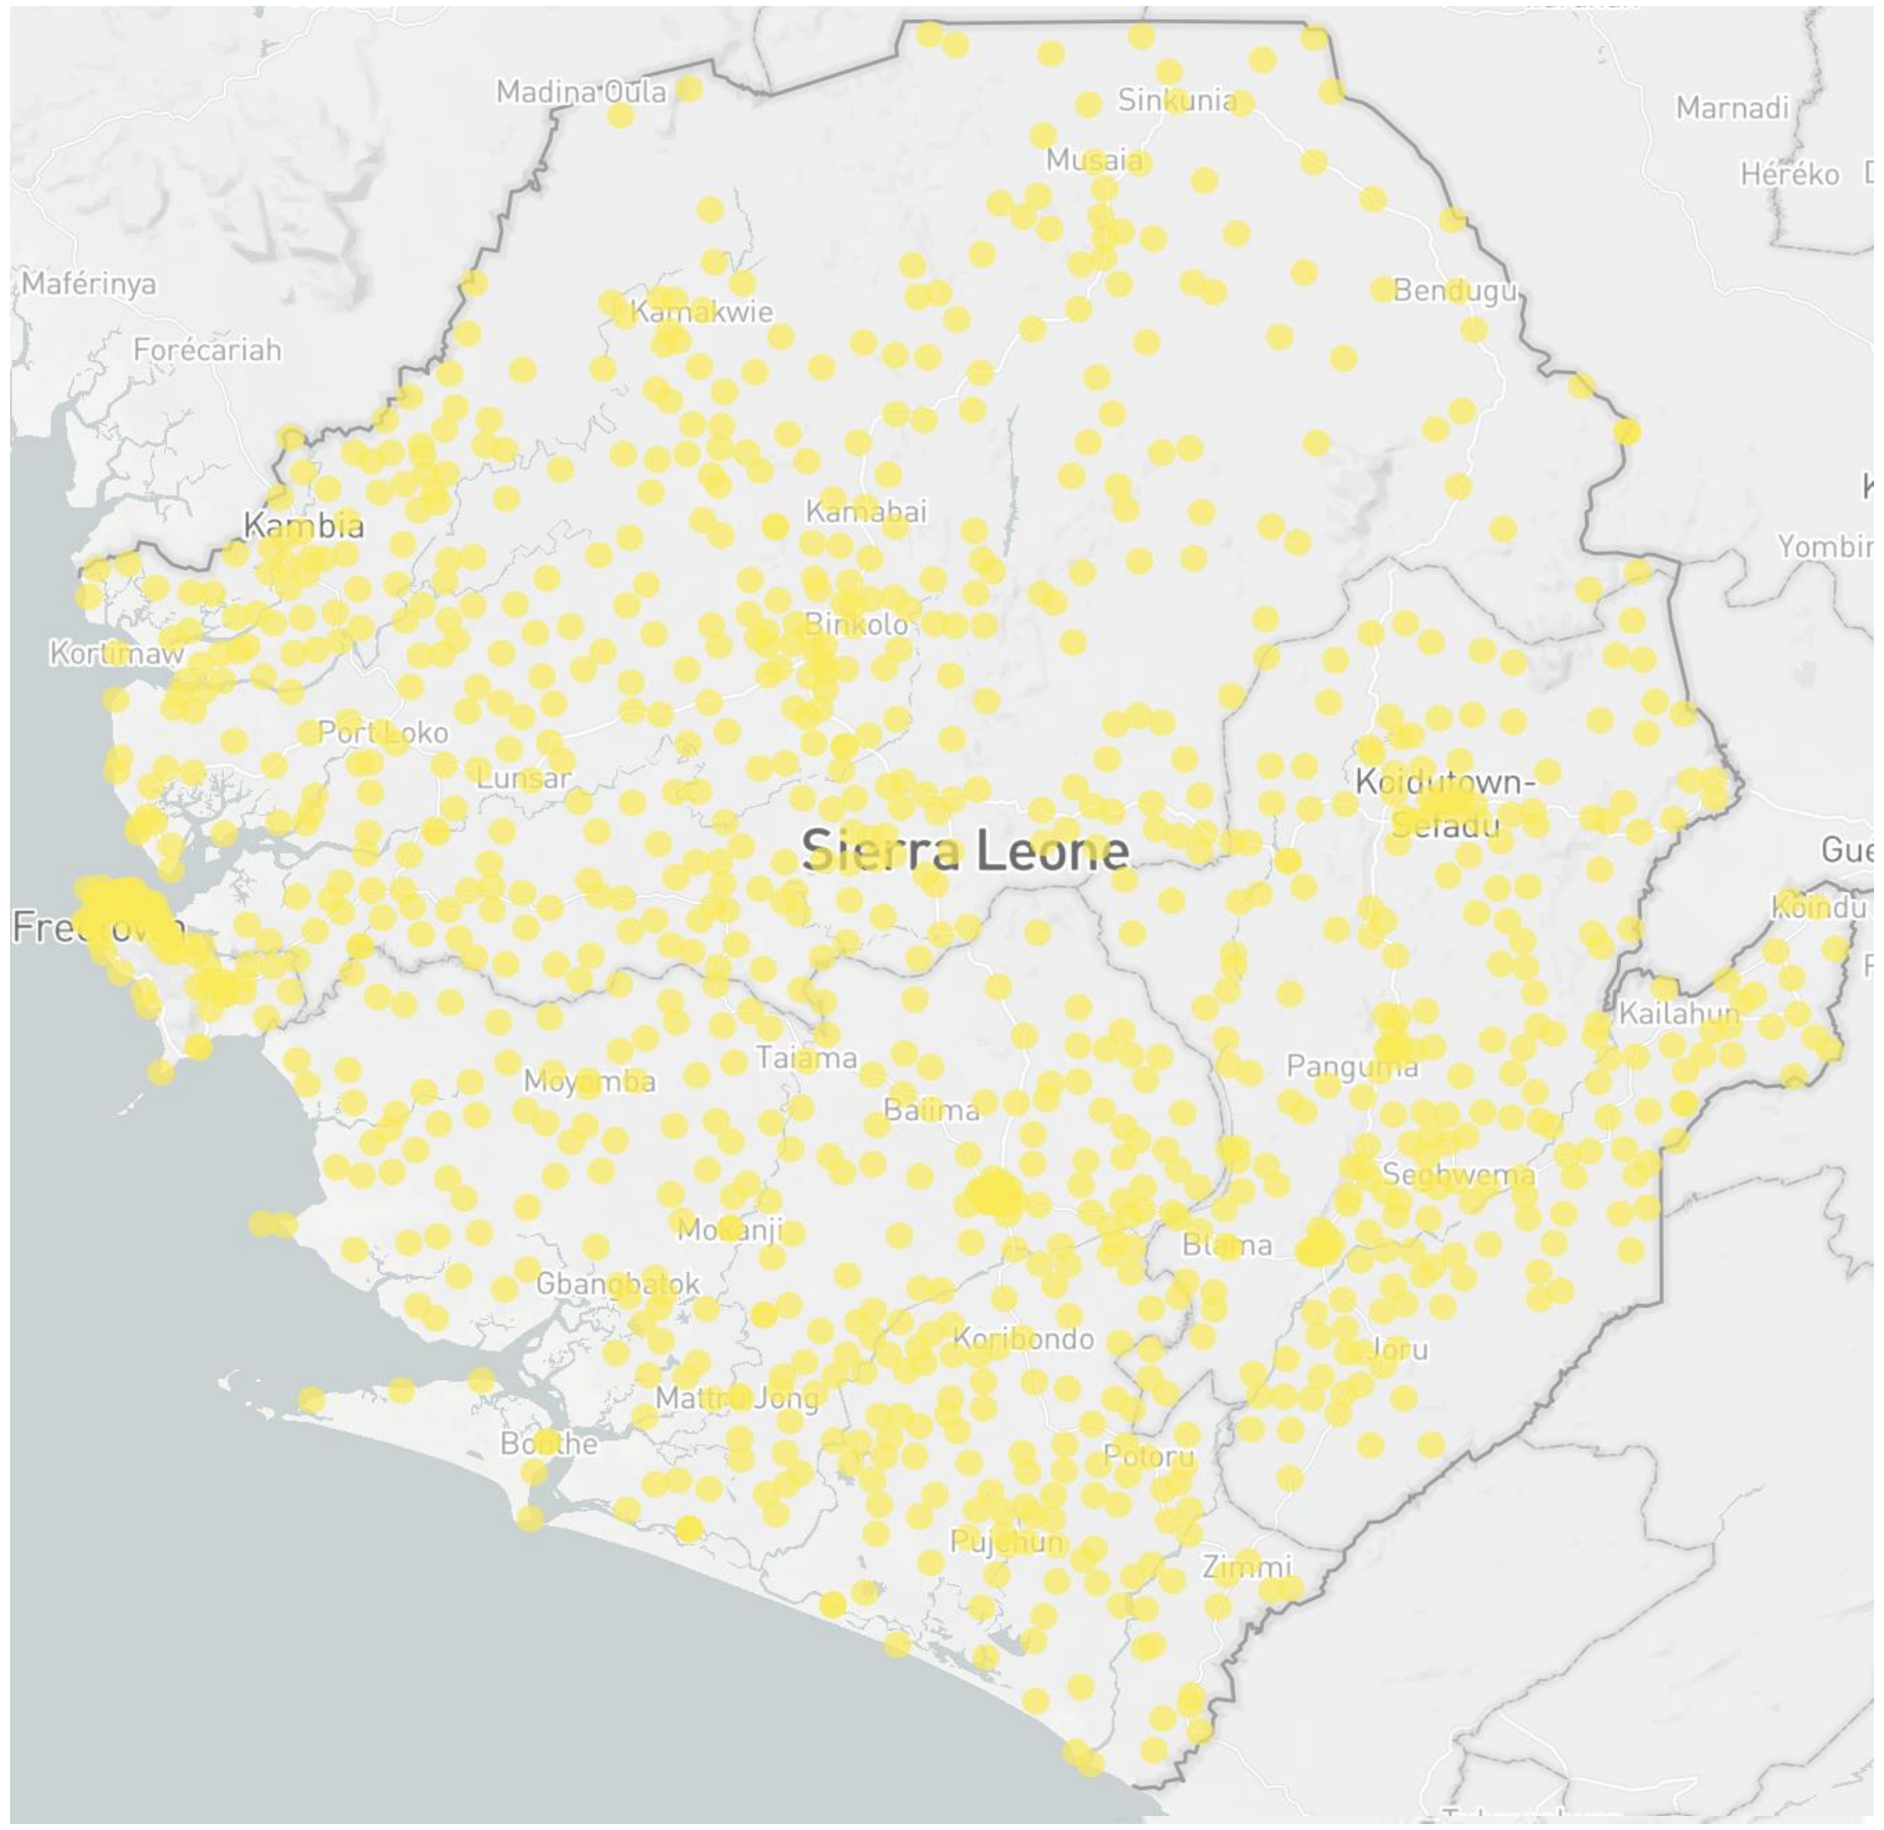

Yellow: treated facilities; Purple: control facilities

Wharton UNIVERSITY of PENNSYLVANIA

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

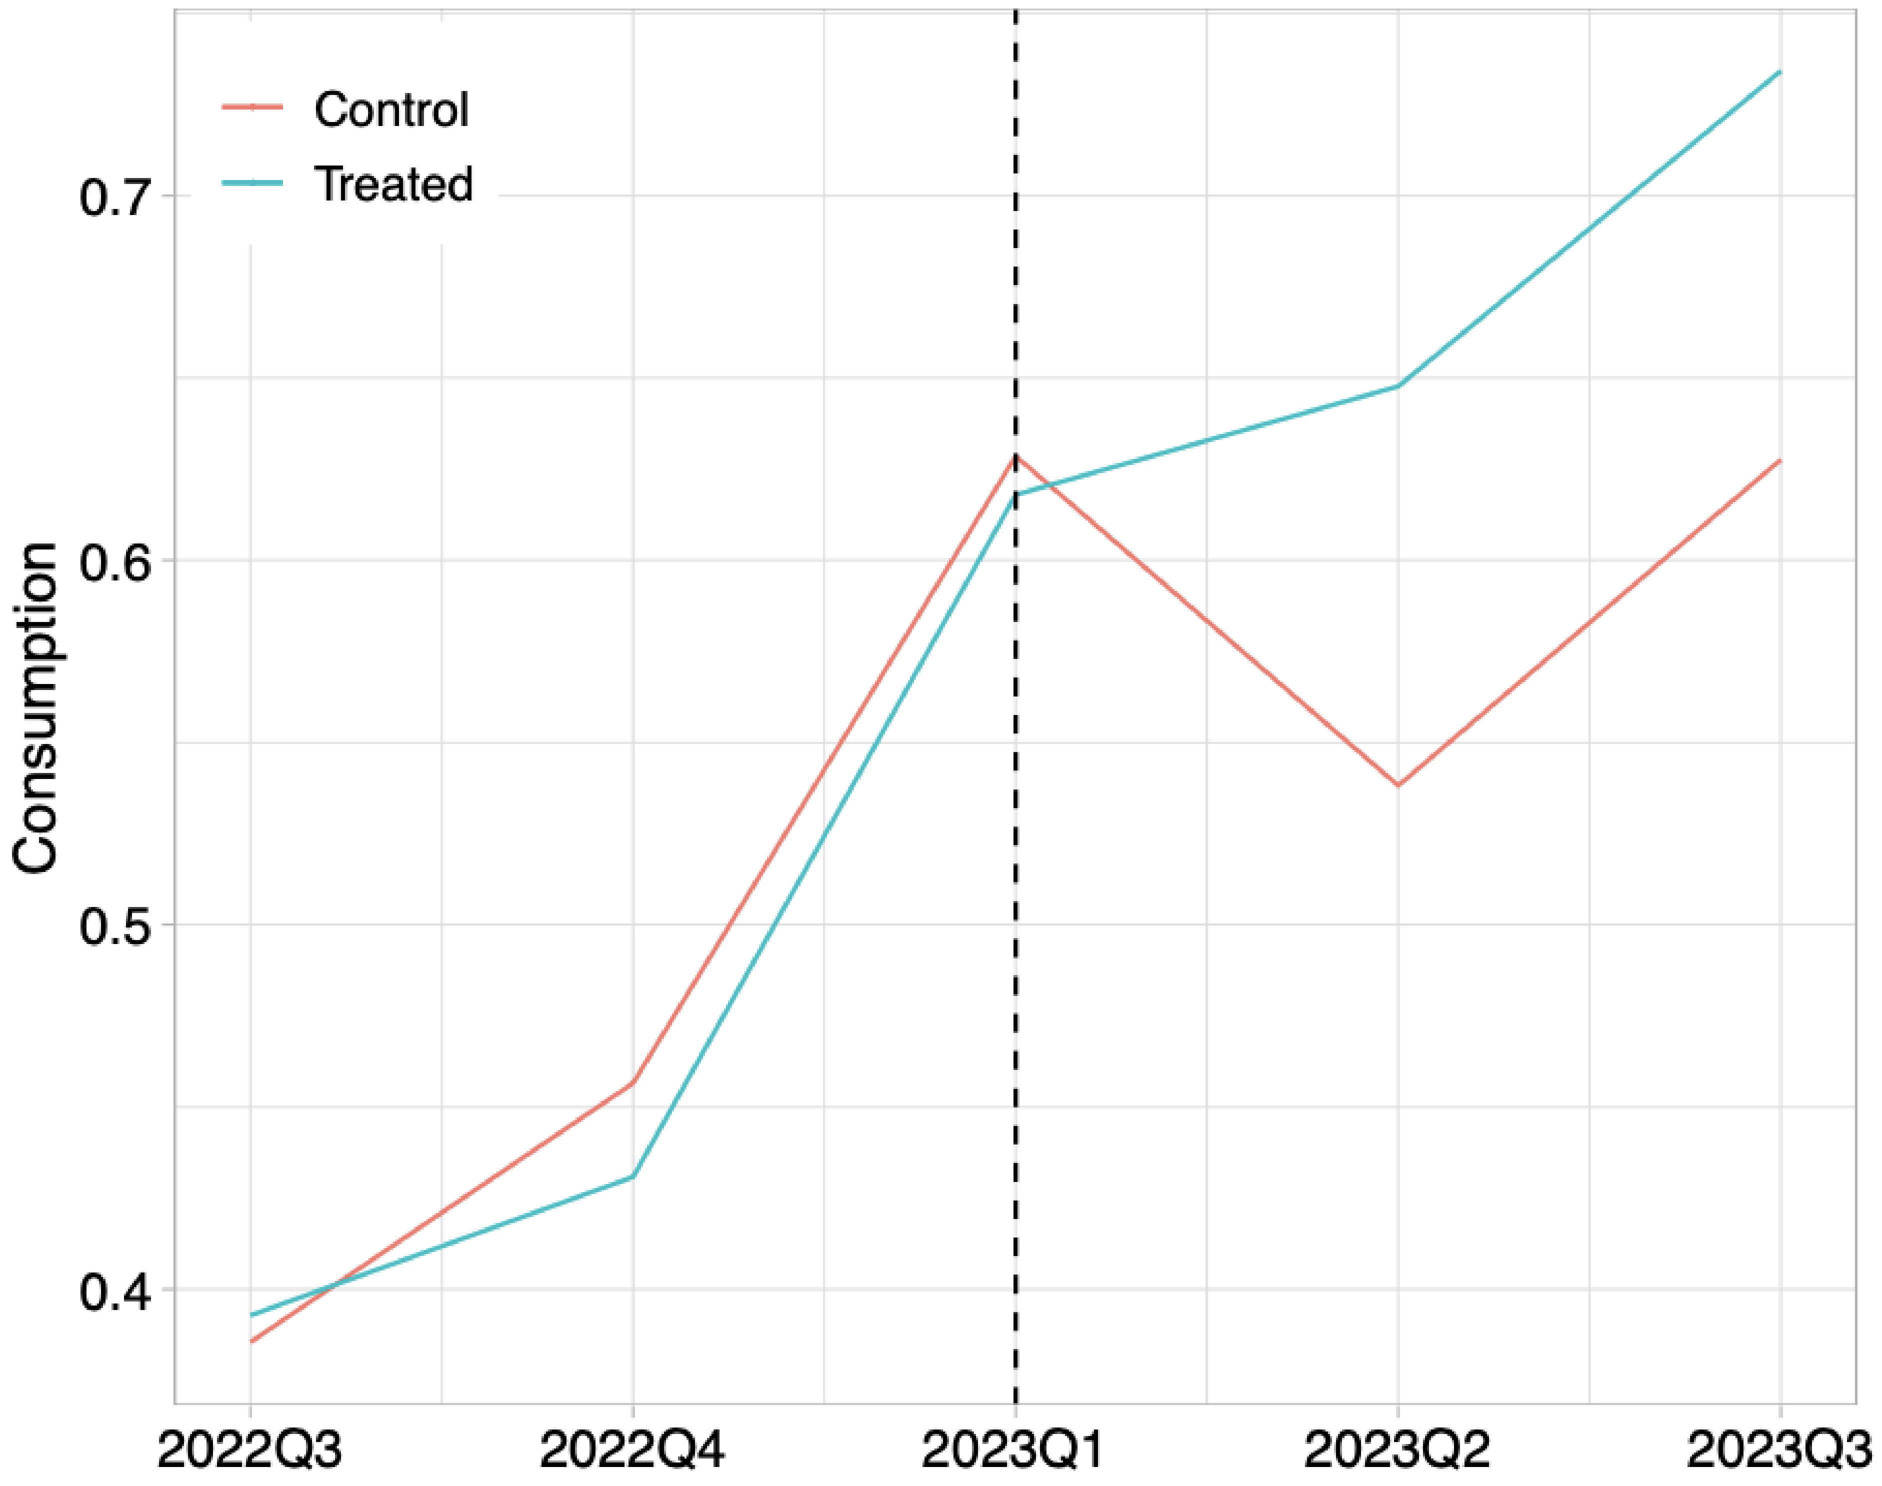

Robustness checks

- 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%**

Wharton UNIVERSITY of PENNSYLVANIA

Field work

Travel from airport to capital city

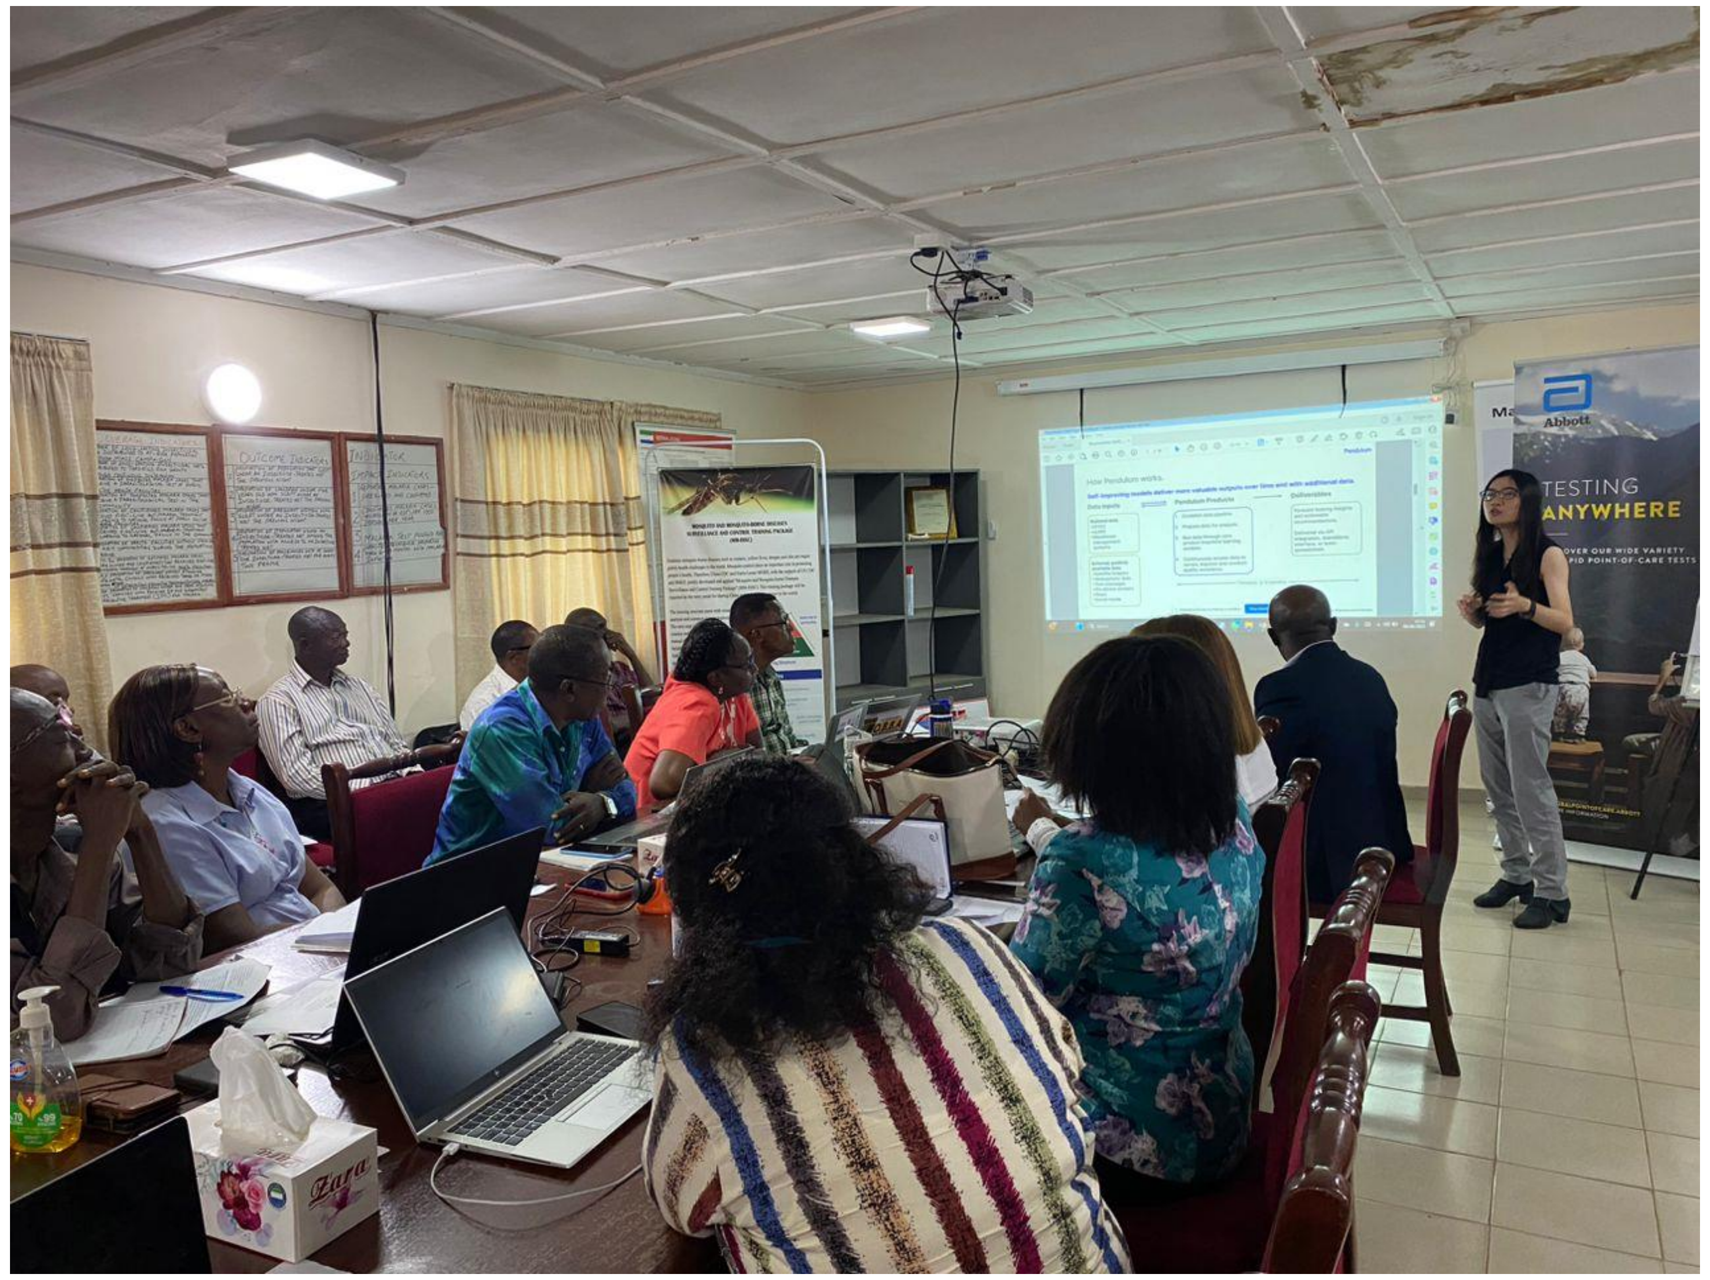

Discuss the work to various stakeholders (WHO, MoHS, US Presidential Malaria Initiative)

Visit the warehouse

Sustainable impact

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

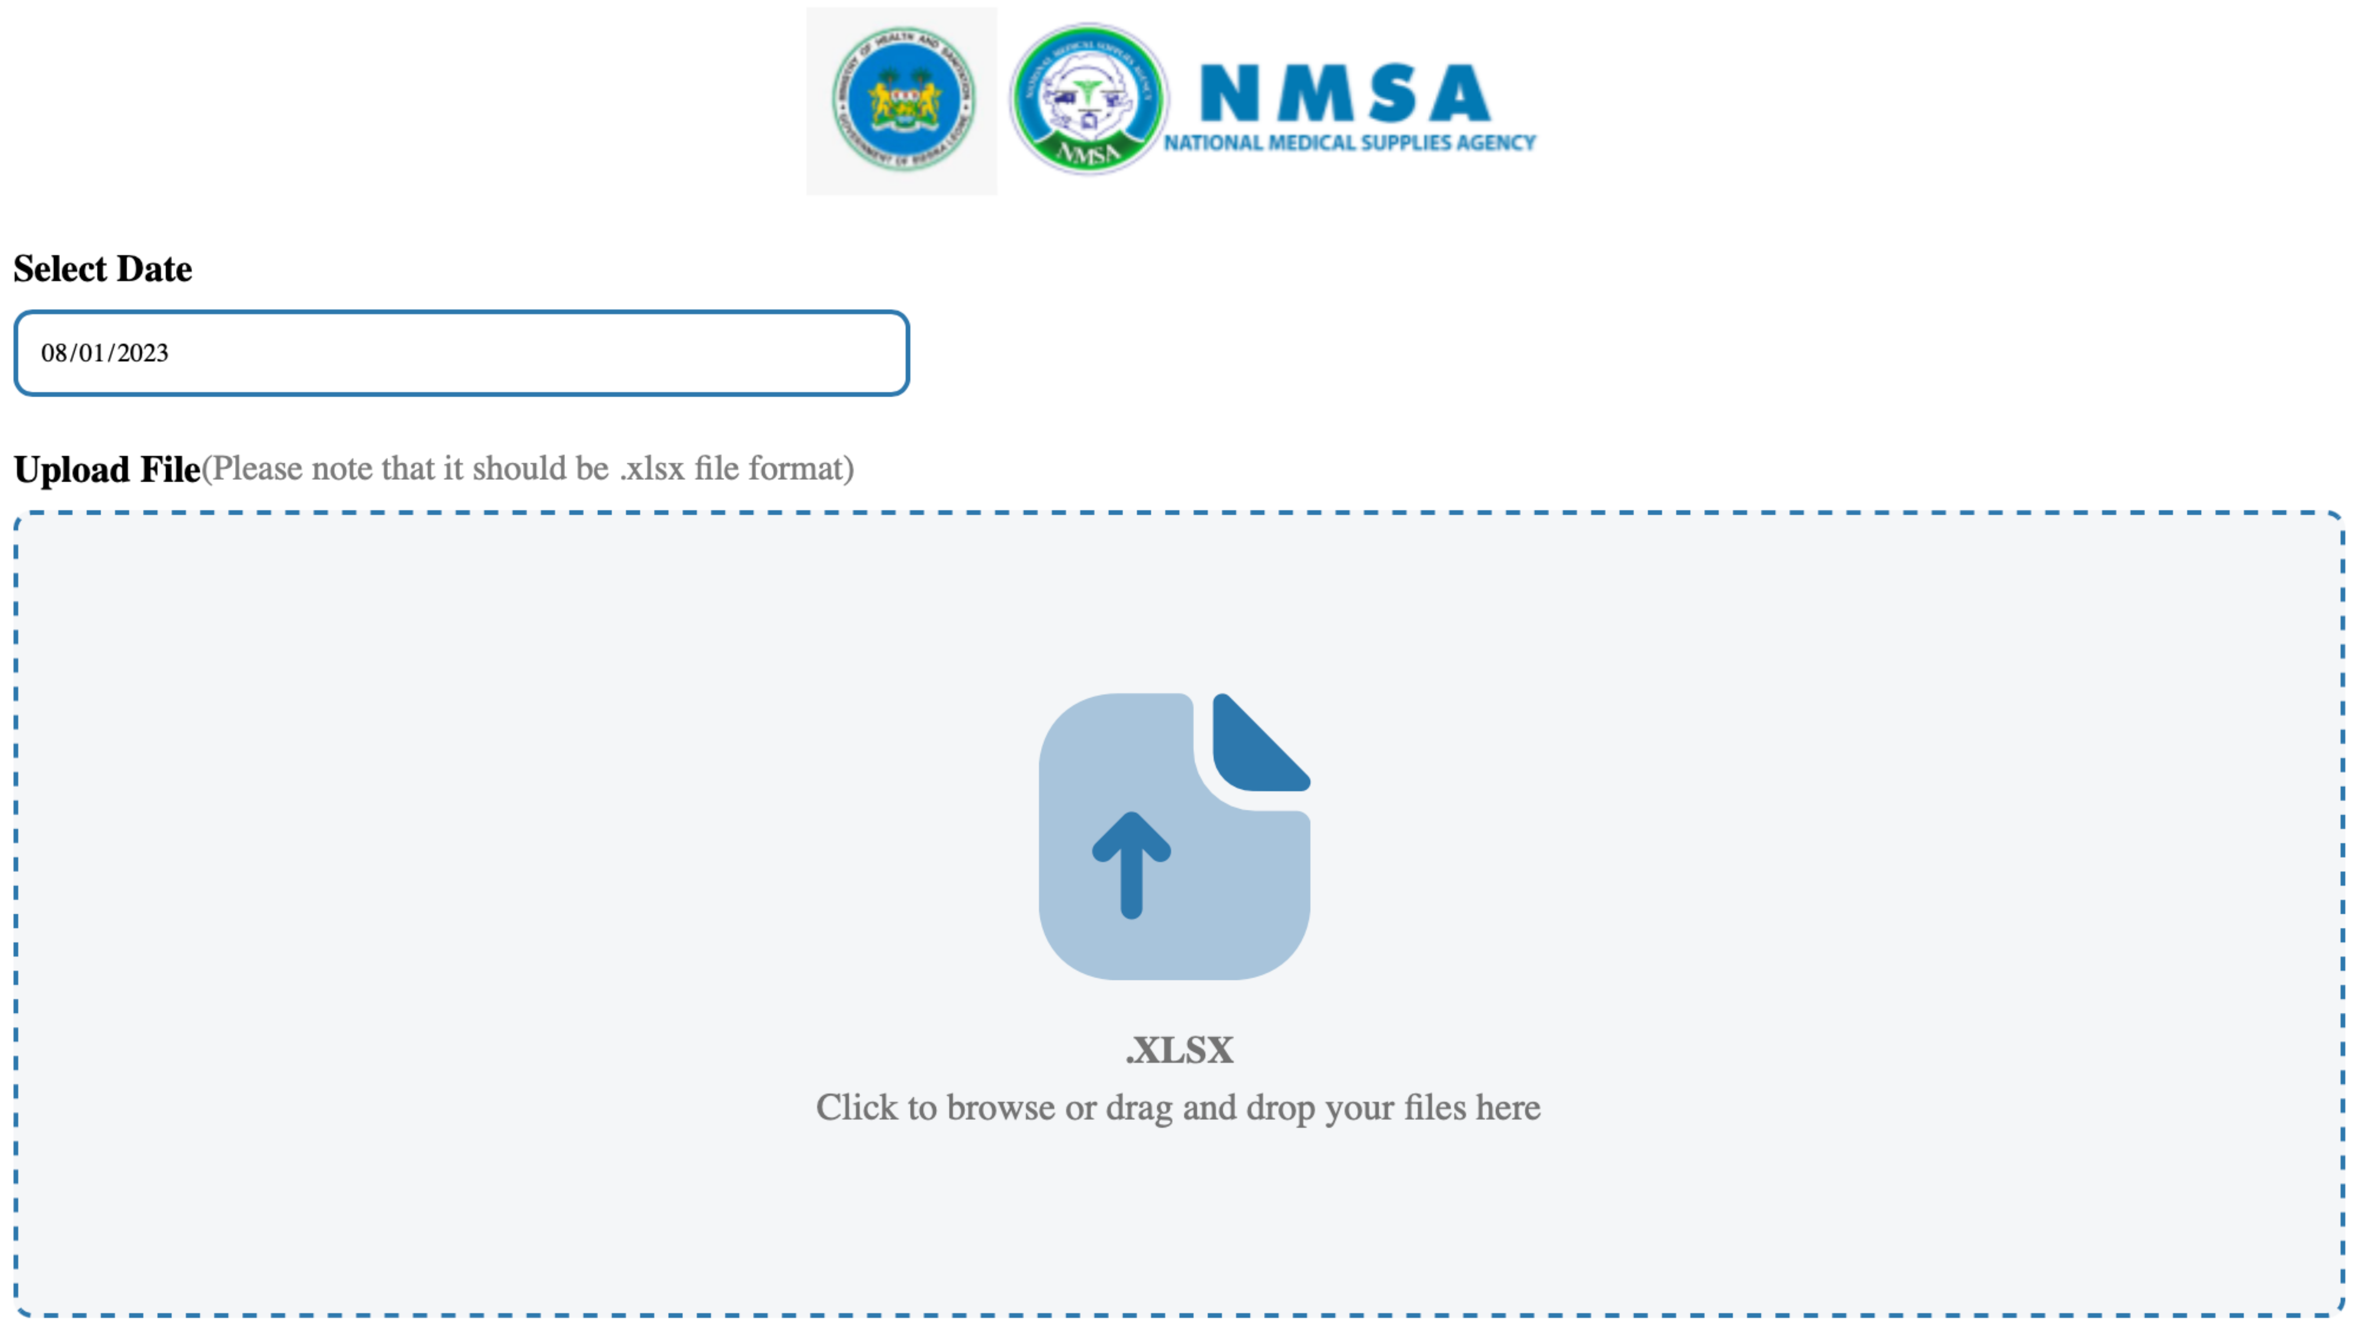

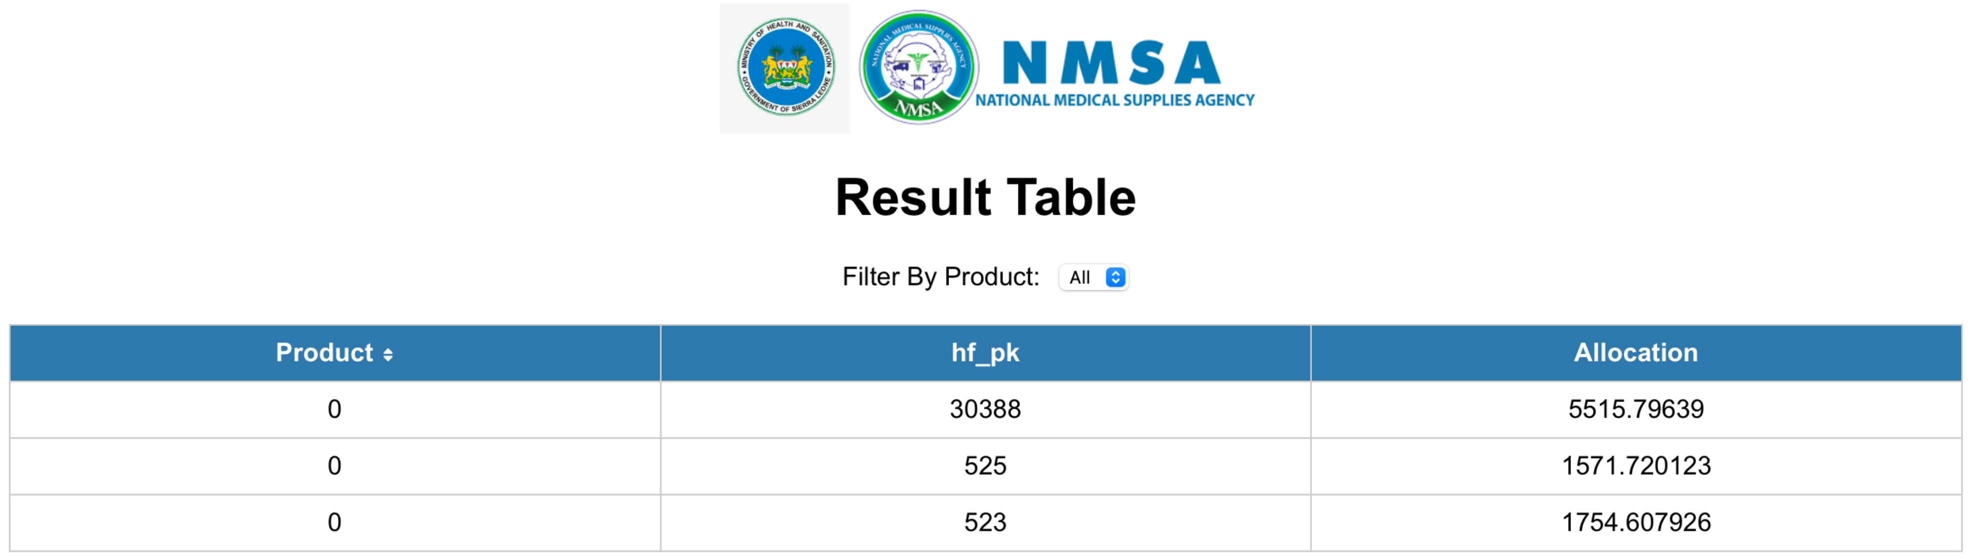

Product	hf_pk	Allocation
0	30388	5515.79639
0	525	1571.720123
0	523	1754.607926

Wharton
UNIVERSITY of PENNSYLVANIA

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.

nature

Explore content ⌄ About the journal ⌄ Publish with us ⌄

nature > articles > article

Article | Published: 29 April 2026

Improving access to essential medicines via decision-aware machine learning

Angel Tsai-Hsuan Chung (鍾采璇), Jatu Abdulai, Patrick Bayoh, Lawrence Sandi, Francis Smart, Hamsa Bastani & Osbert Bastani

Nature (2026) | Cite this article

Effective Personalized AI tutor via LLM-Guided Reinforcement Learning

Angel Tsai-Hsuan Chung

Joint work w/ Botong Zhang, Ling-Chieh Kung, Hamsa Bastani, and Osbert Bastani

Skill Development

- Importance of re-skilling with AI [Harvard Business Review 2023; Morandini et al. 2023; Eloundou et al. 2024]

Wharton
UNIVERSITY of PENNSYLVANIA

Skill Development

- Importance of re-skilling with AI [Harvard Business Review 2023; Morandini et al. 2023; Eloundou et al. 2024]
- How can we make learning more efficient?

Skill Development

- Importance of re-skilling with AI [Harvard Business Review 2023; Morandini et al. 2023; Eloundou et al. 2024]
- How can we make learning more efficient?
  – Typically, courses are designed for the “average” learner → inefficient
  – Personalization can help everyone move at their own pace

Wharton
University of Pennsylvania

Can we leverage **rich, real-time signals using LLM** to estimate learner's **knowledge state** to enable better **personalization**?

Overview

- We developed a POMDP-based RL that incorporates real-time richer signals using LLM for better personalization.

- We collaborated with governments to conduct a large-scale field RCT to evaluate it and find significant impact.

DEPARTMENT OF EDUCATION
TAIPEI CITY GOVERNMENT

AIC American innovation center

Wharton UNIVERSITY of PENNSYLVANIA

Learning Platform

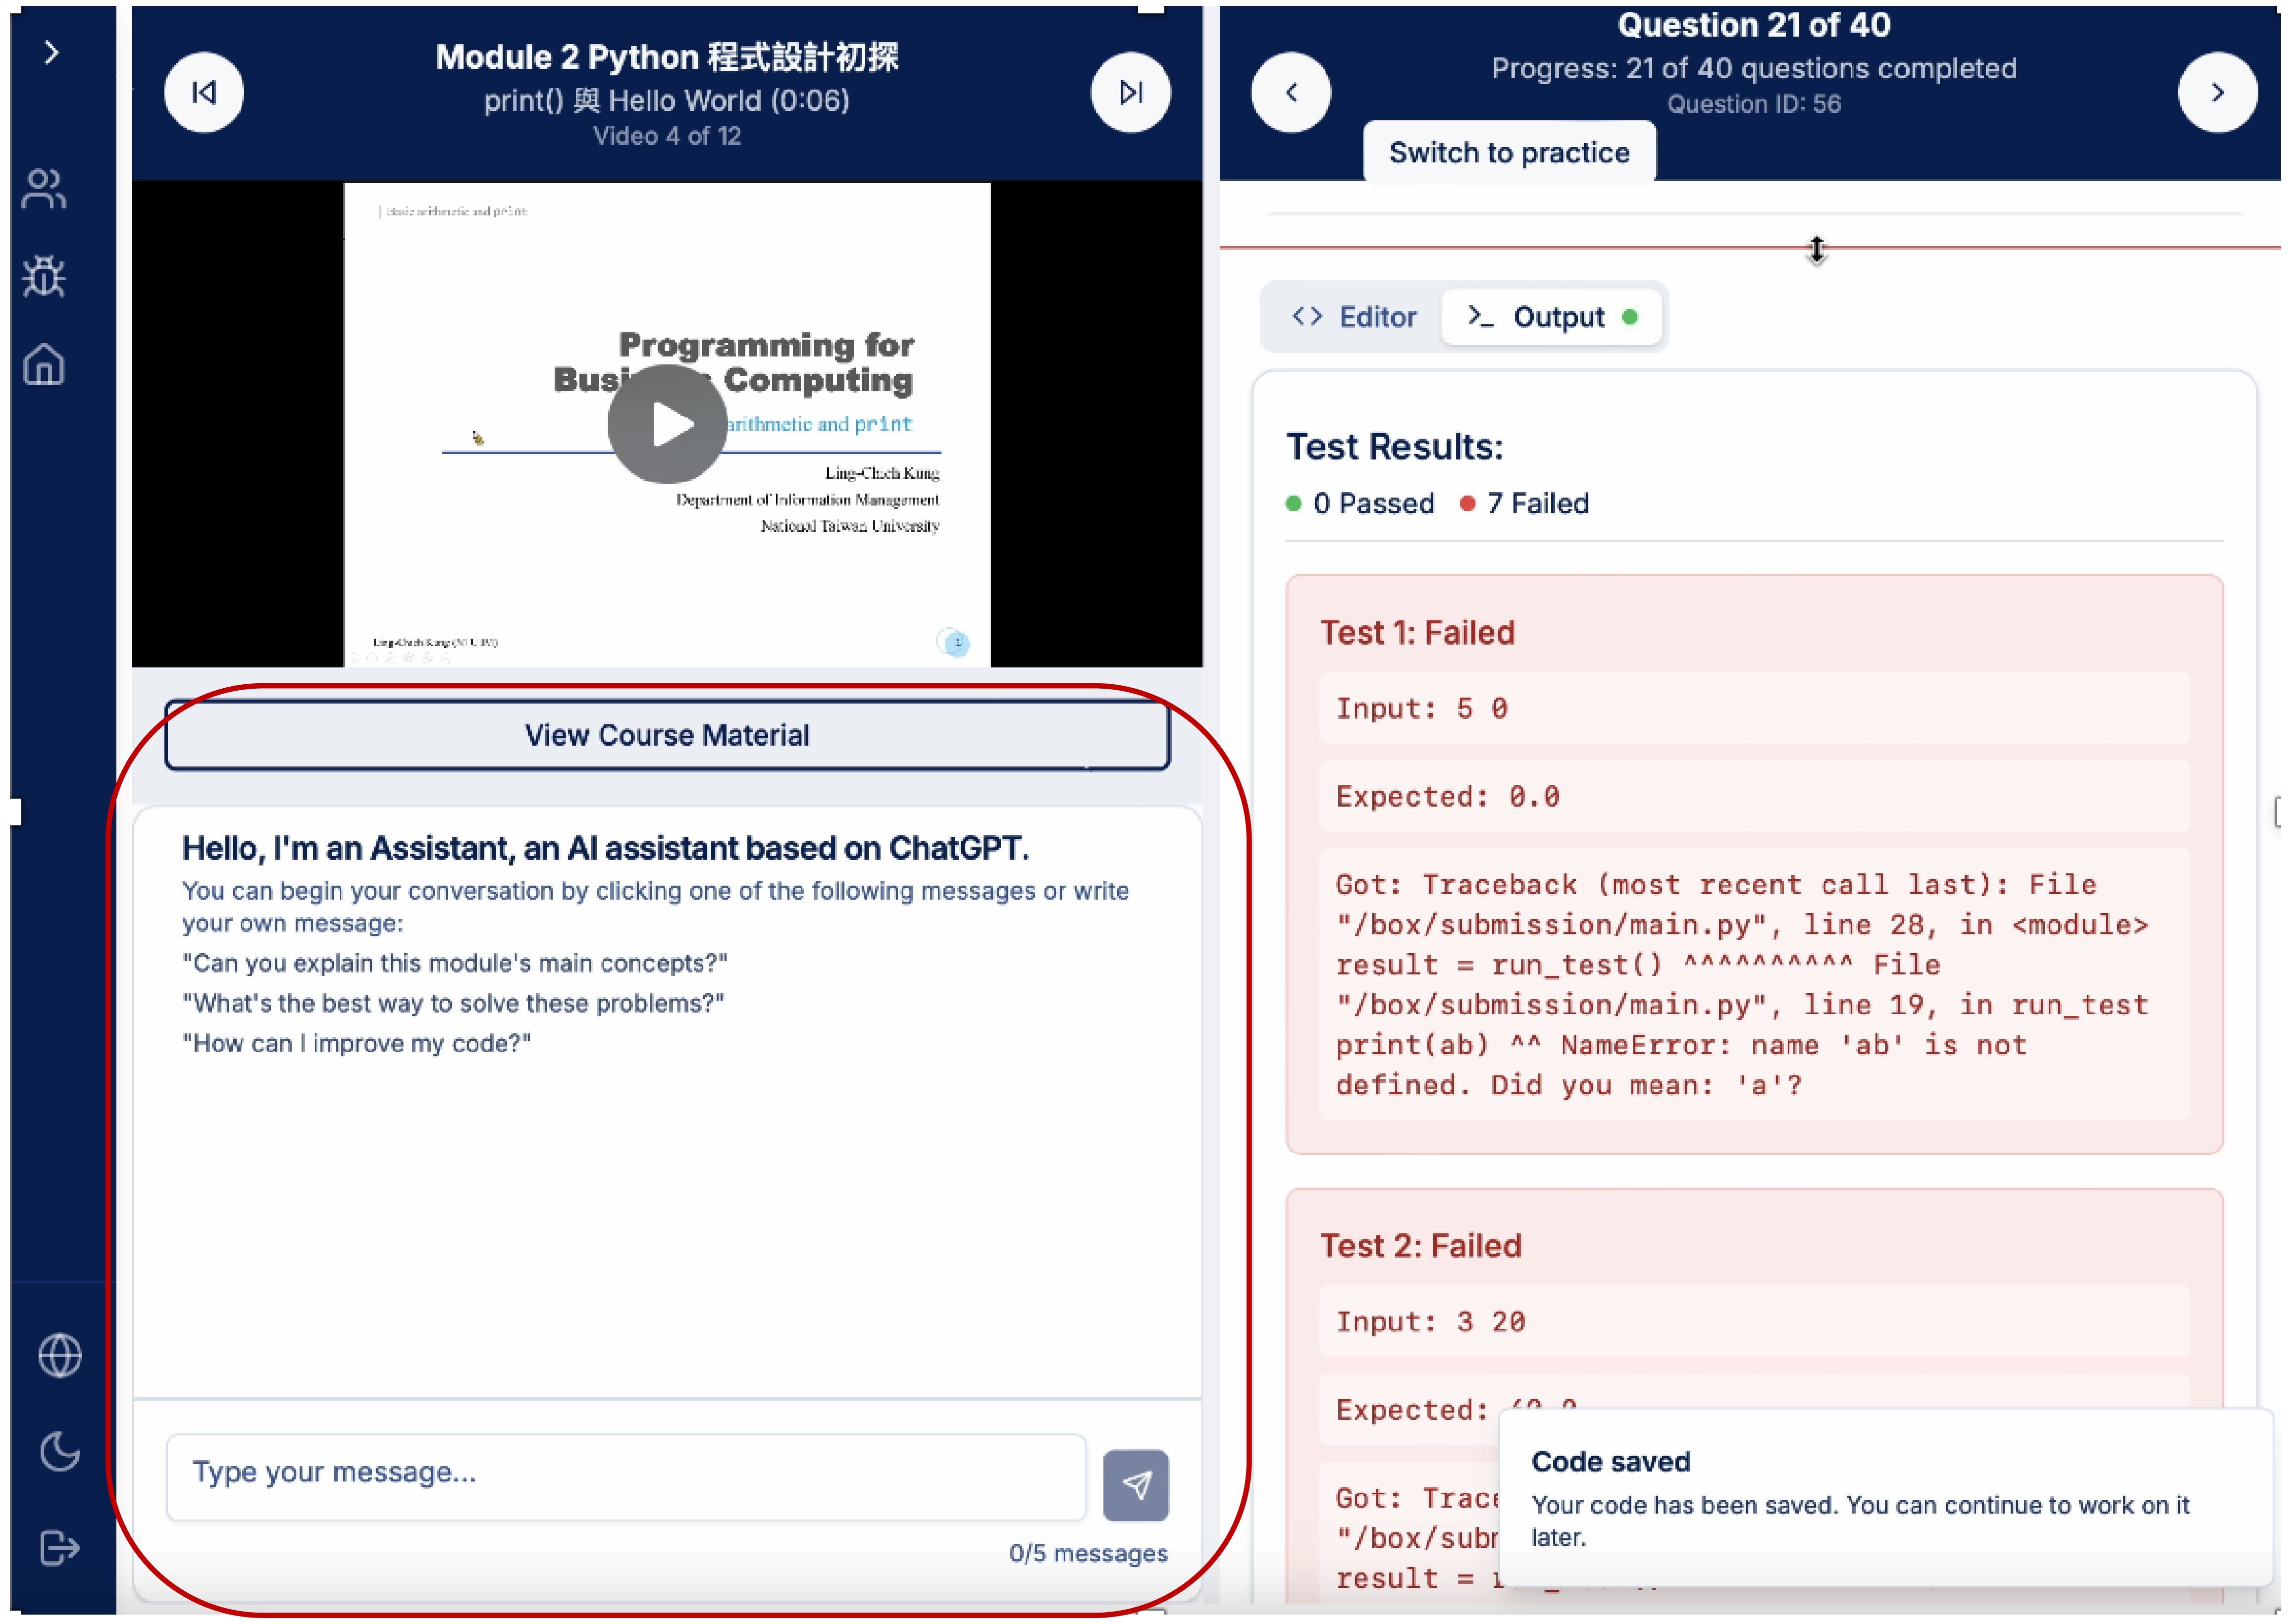

Reinforcement Learning

c_1 initial knowledge
c_2 mastery ceiling
τ learning speed

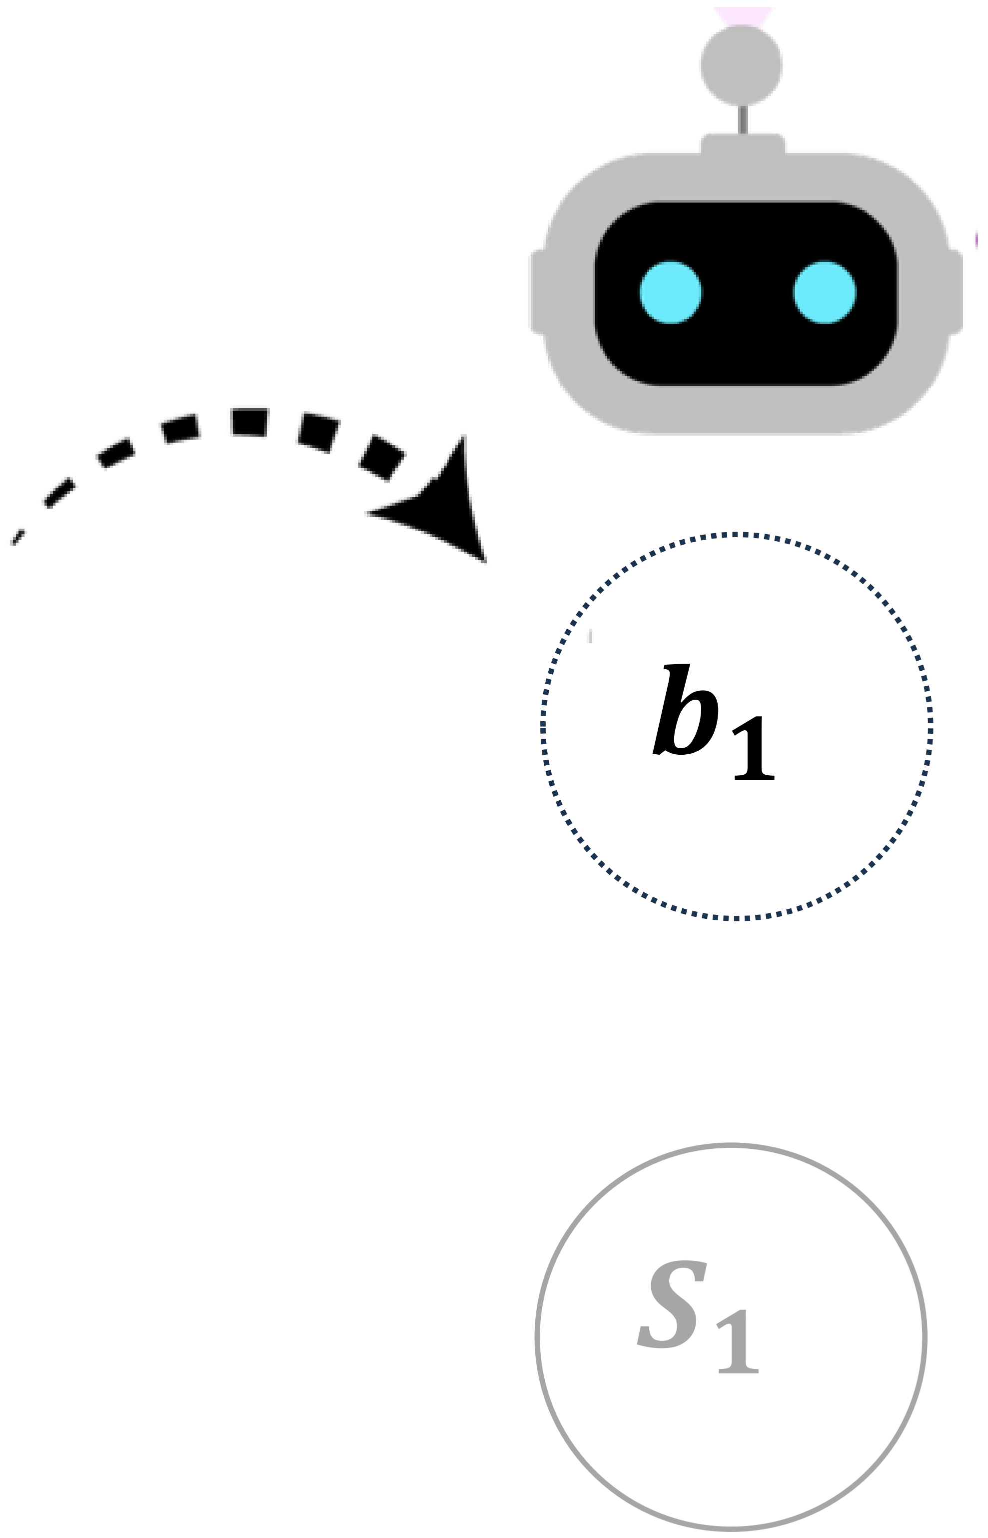

- b: *belief state*
- S: *true state*

Reinforcement Learning

c_1 initial knowledge
c_2 mastery ceiling
τ learning speed

$\mathcal{A}$ = {Easy, Medium, Hard, Very Hard}

$$a_t^* = \pi^*(b_t)$$

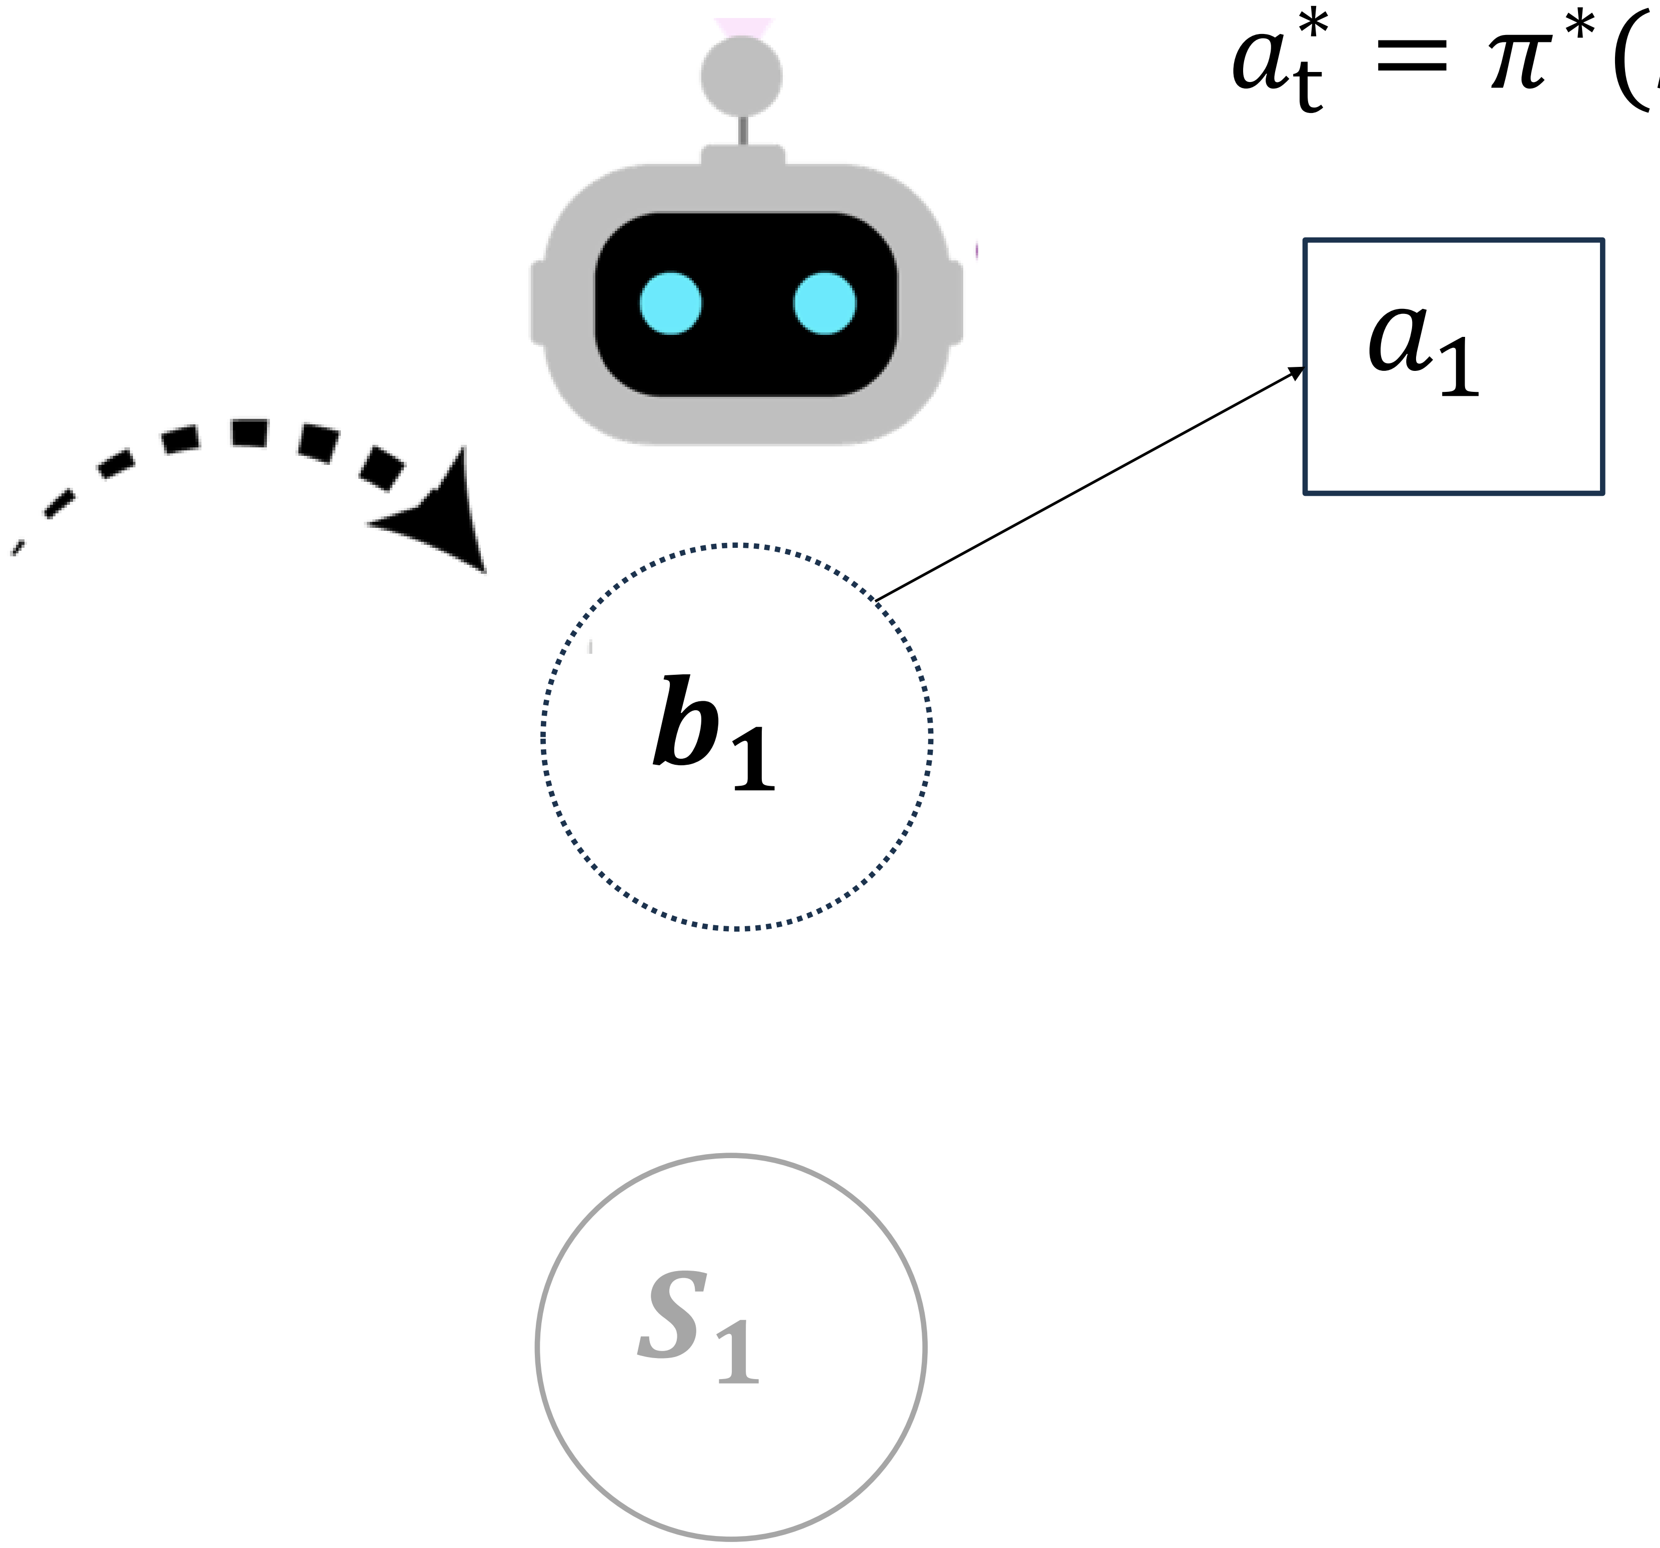

Write a program that reads a positive integer n representing the number of intersections. Assume each intersection has 3 lanes (numbered 1 to 3). For each lane, the number of vehicles is calculated as: vehicles = (intersection number + lane number) * 5....

- b: *belief state*
- S: *true state*
- a: *action*

Wharton
UNIVERSITY *of* PENNSYLVANIA

Reinforcement Learning

c_1 initial knowledge
c_2 mastery ceiling
τ learning speed

- b: *belief state*
- S: *true state*
- a: *action*
- O: *observation*

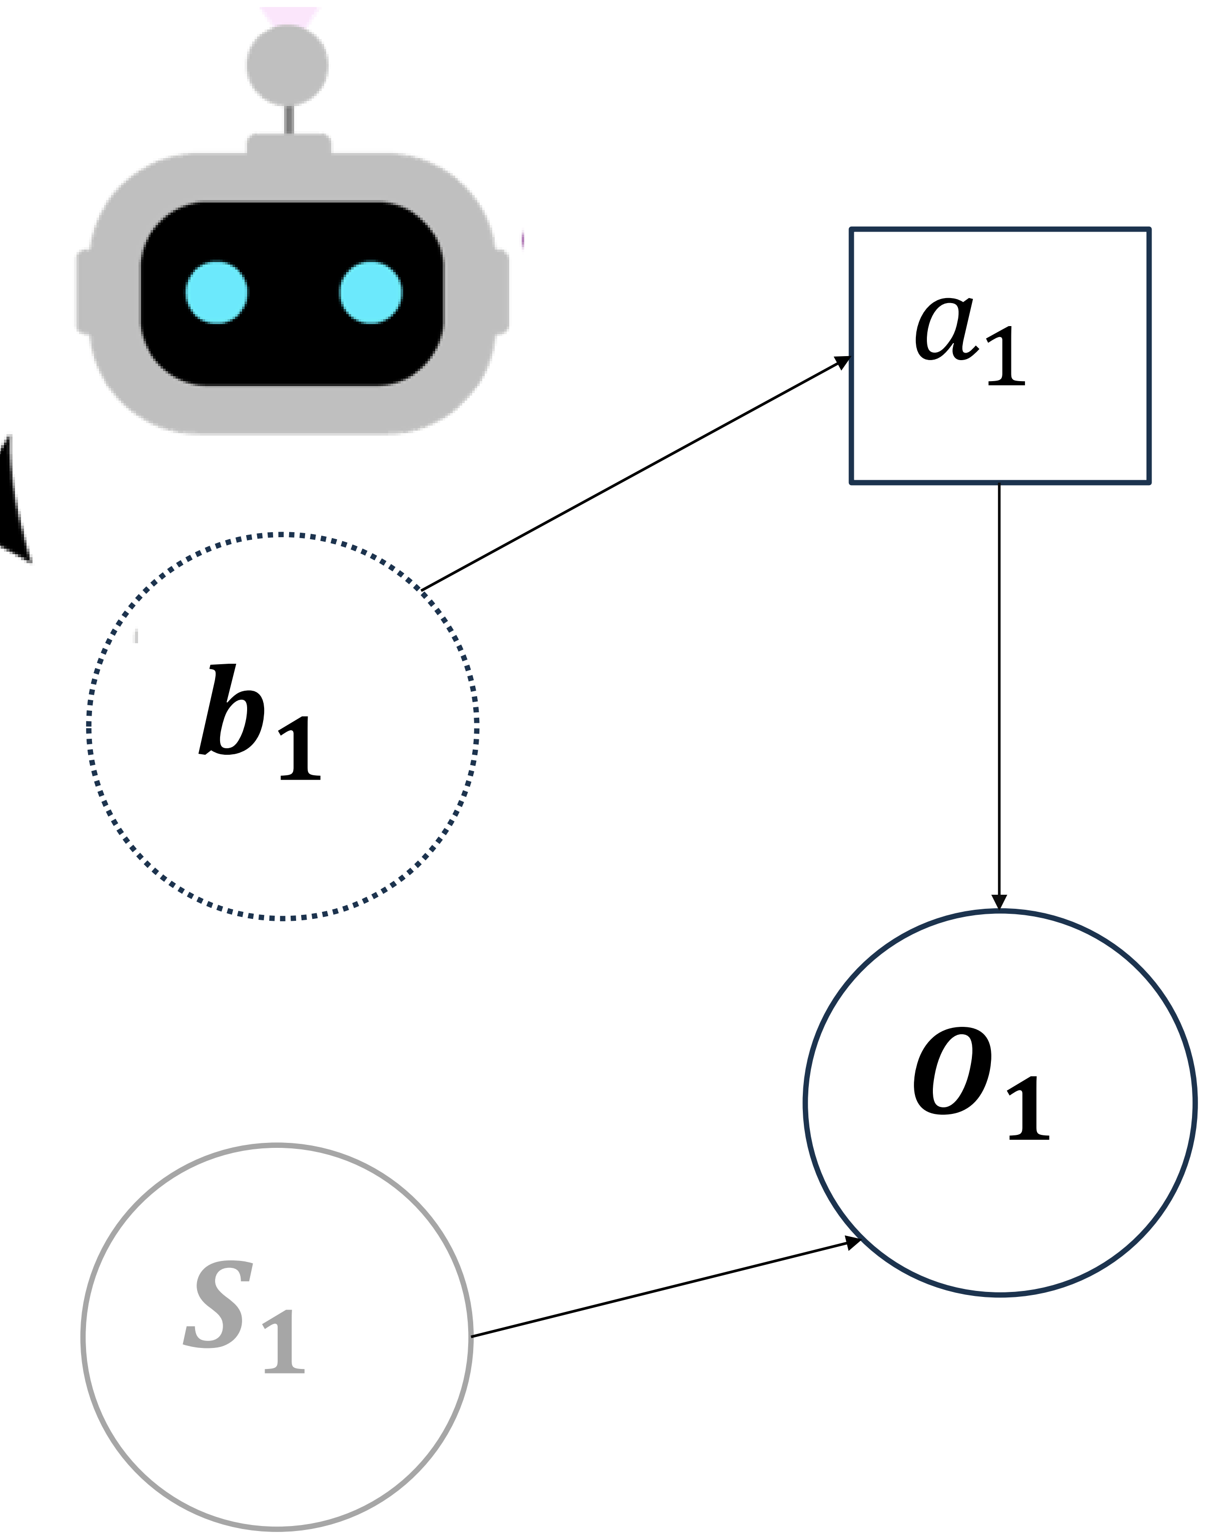

Write a program that reads a positive i nteger n representing the number of i ntersections. Assume each intersectio n has 3 lanes (numbered 1 to 3). For e ach lane, the number of vehicles is cal culated as: vehicles = (intersection nu mber + lane number) * 5....

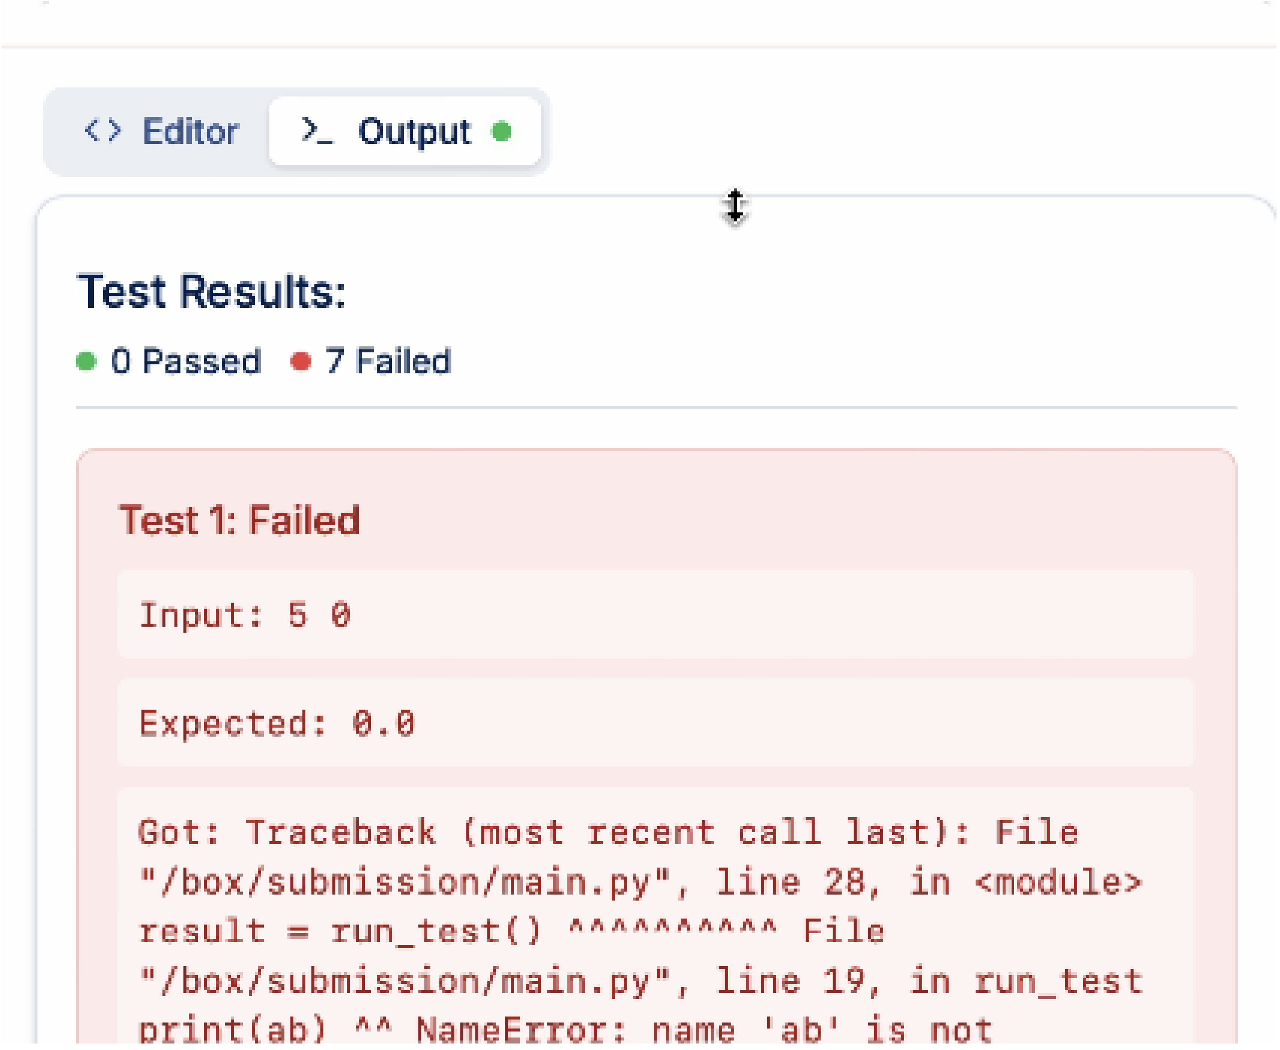

Reinforcement Learning

c_1 initial knowledge
c_2 mastery ceiling
τ learning speed

- b: *belief state*
- S: *true state*
- a: *action*
- O: *observation*

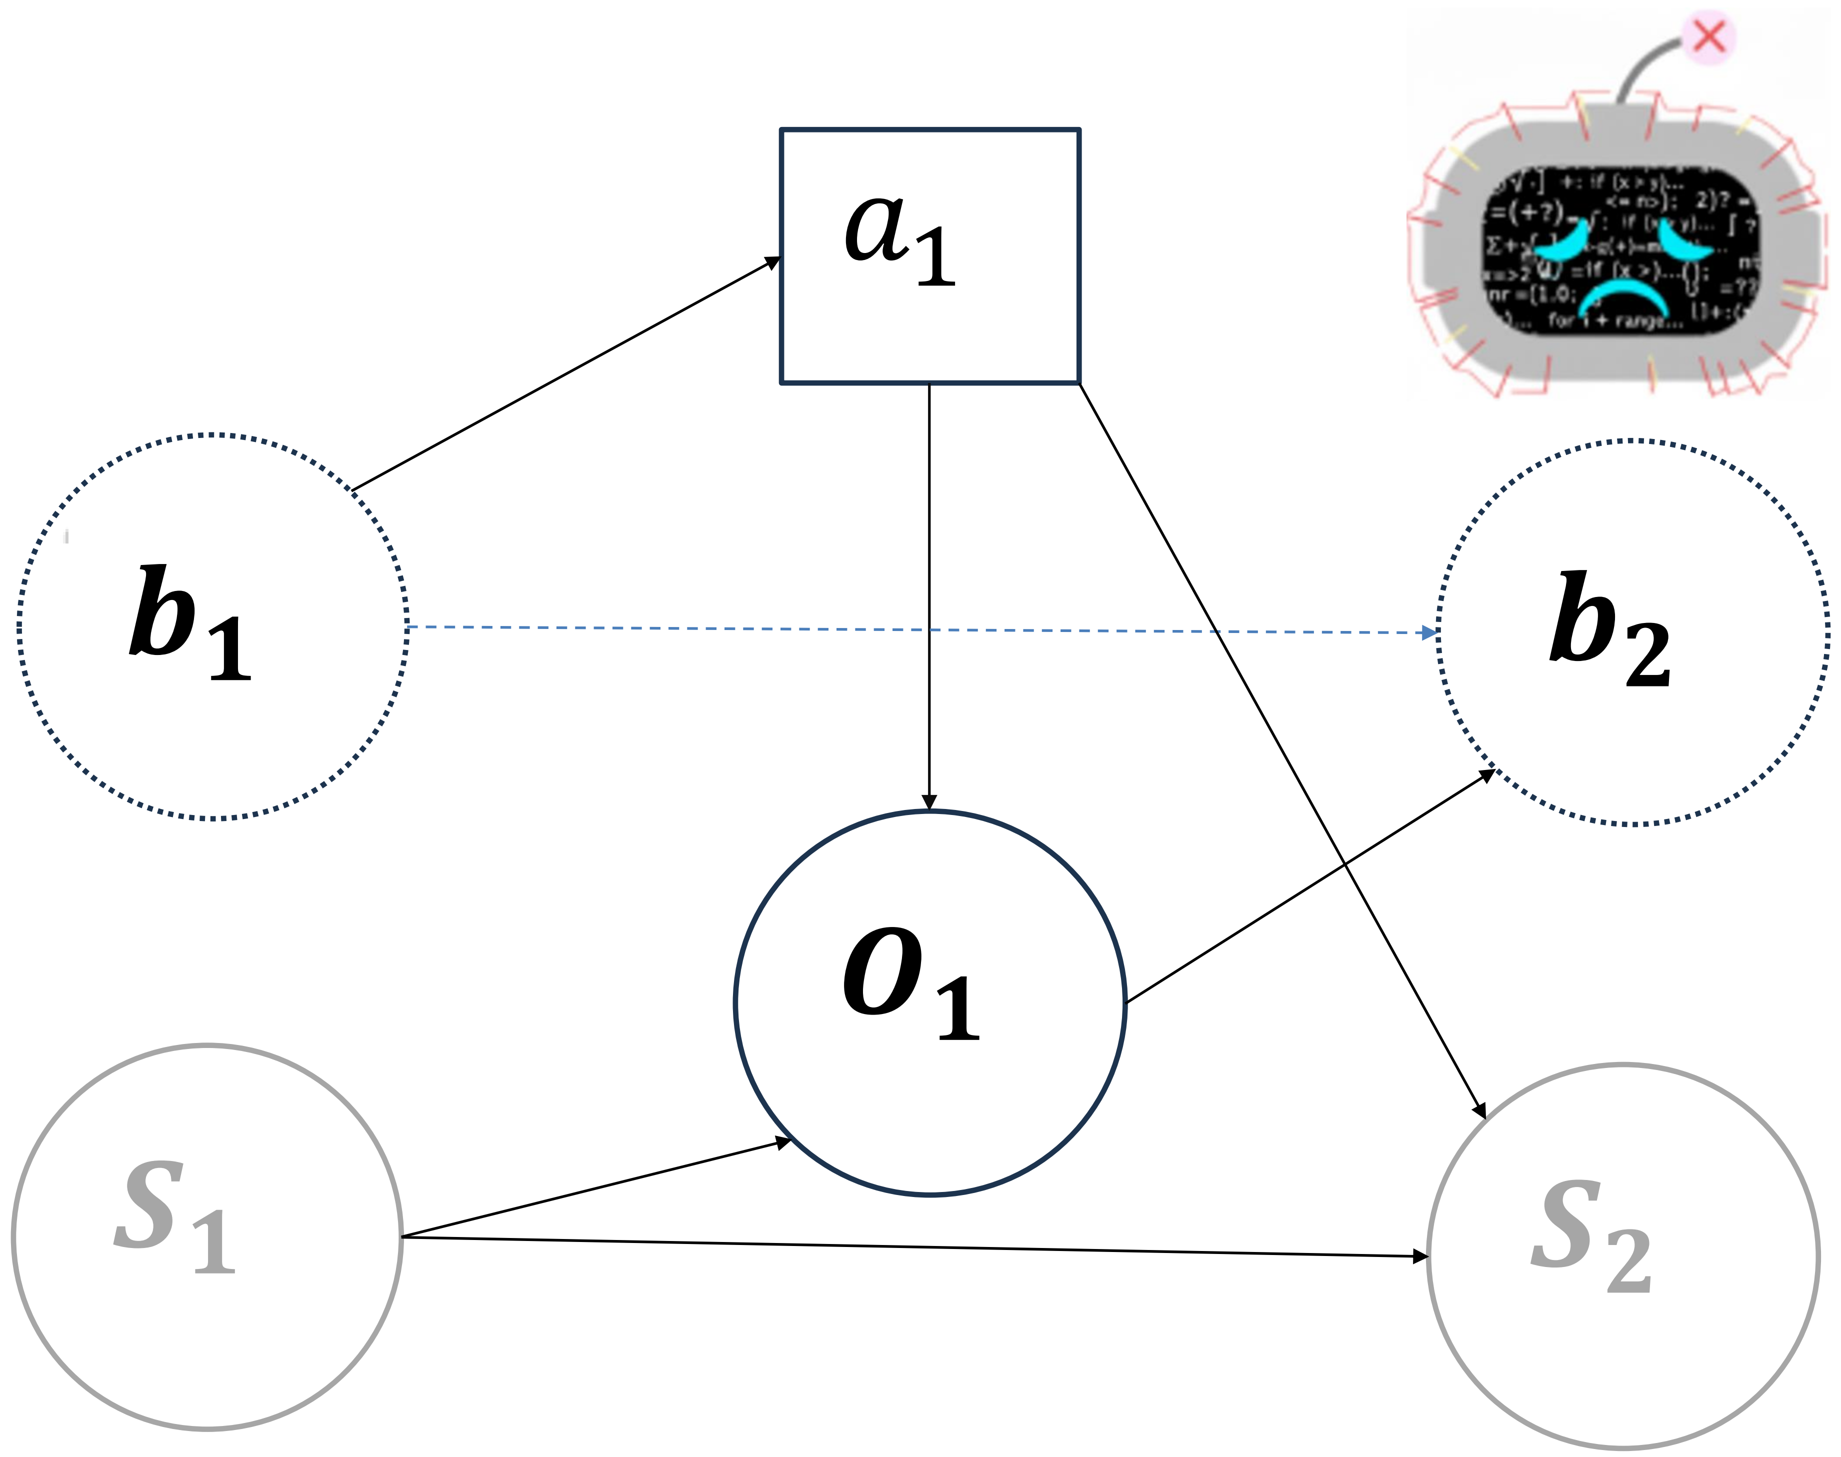

Experimental Context & Timeline

- Government certification program
  - ~1,000 students across 10 high schools in Taipei

2024 July — Pilot Study

2024 Fall — Recruiting
- Two briefings
(1) Teacher
(2) Student & parent

2025 Jan — Project Launched

2025 Spring — Field RCT 5 months (Feb-June)

In-person
Paper-based Exam
(No digital device allowed)

Wharton
UNIVERSITY of PENNSYLVANIA

Randomized Control Trial (RCT) Design

- **Treatment assignment**
  - RL + AI tutor
  - Fixed (expert designed) + AI tutor

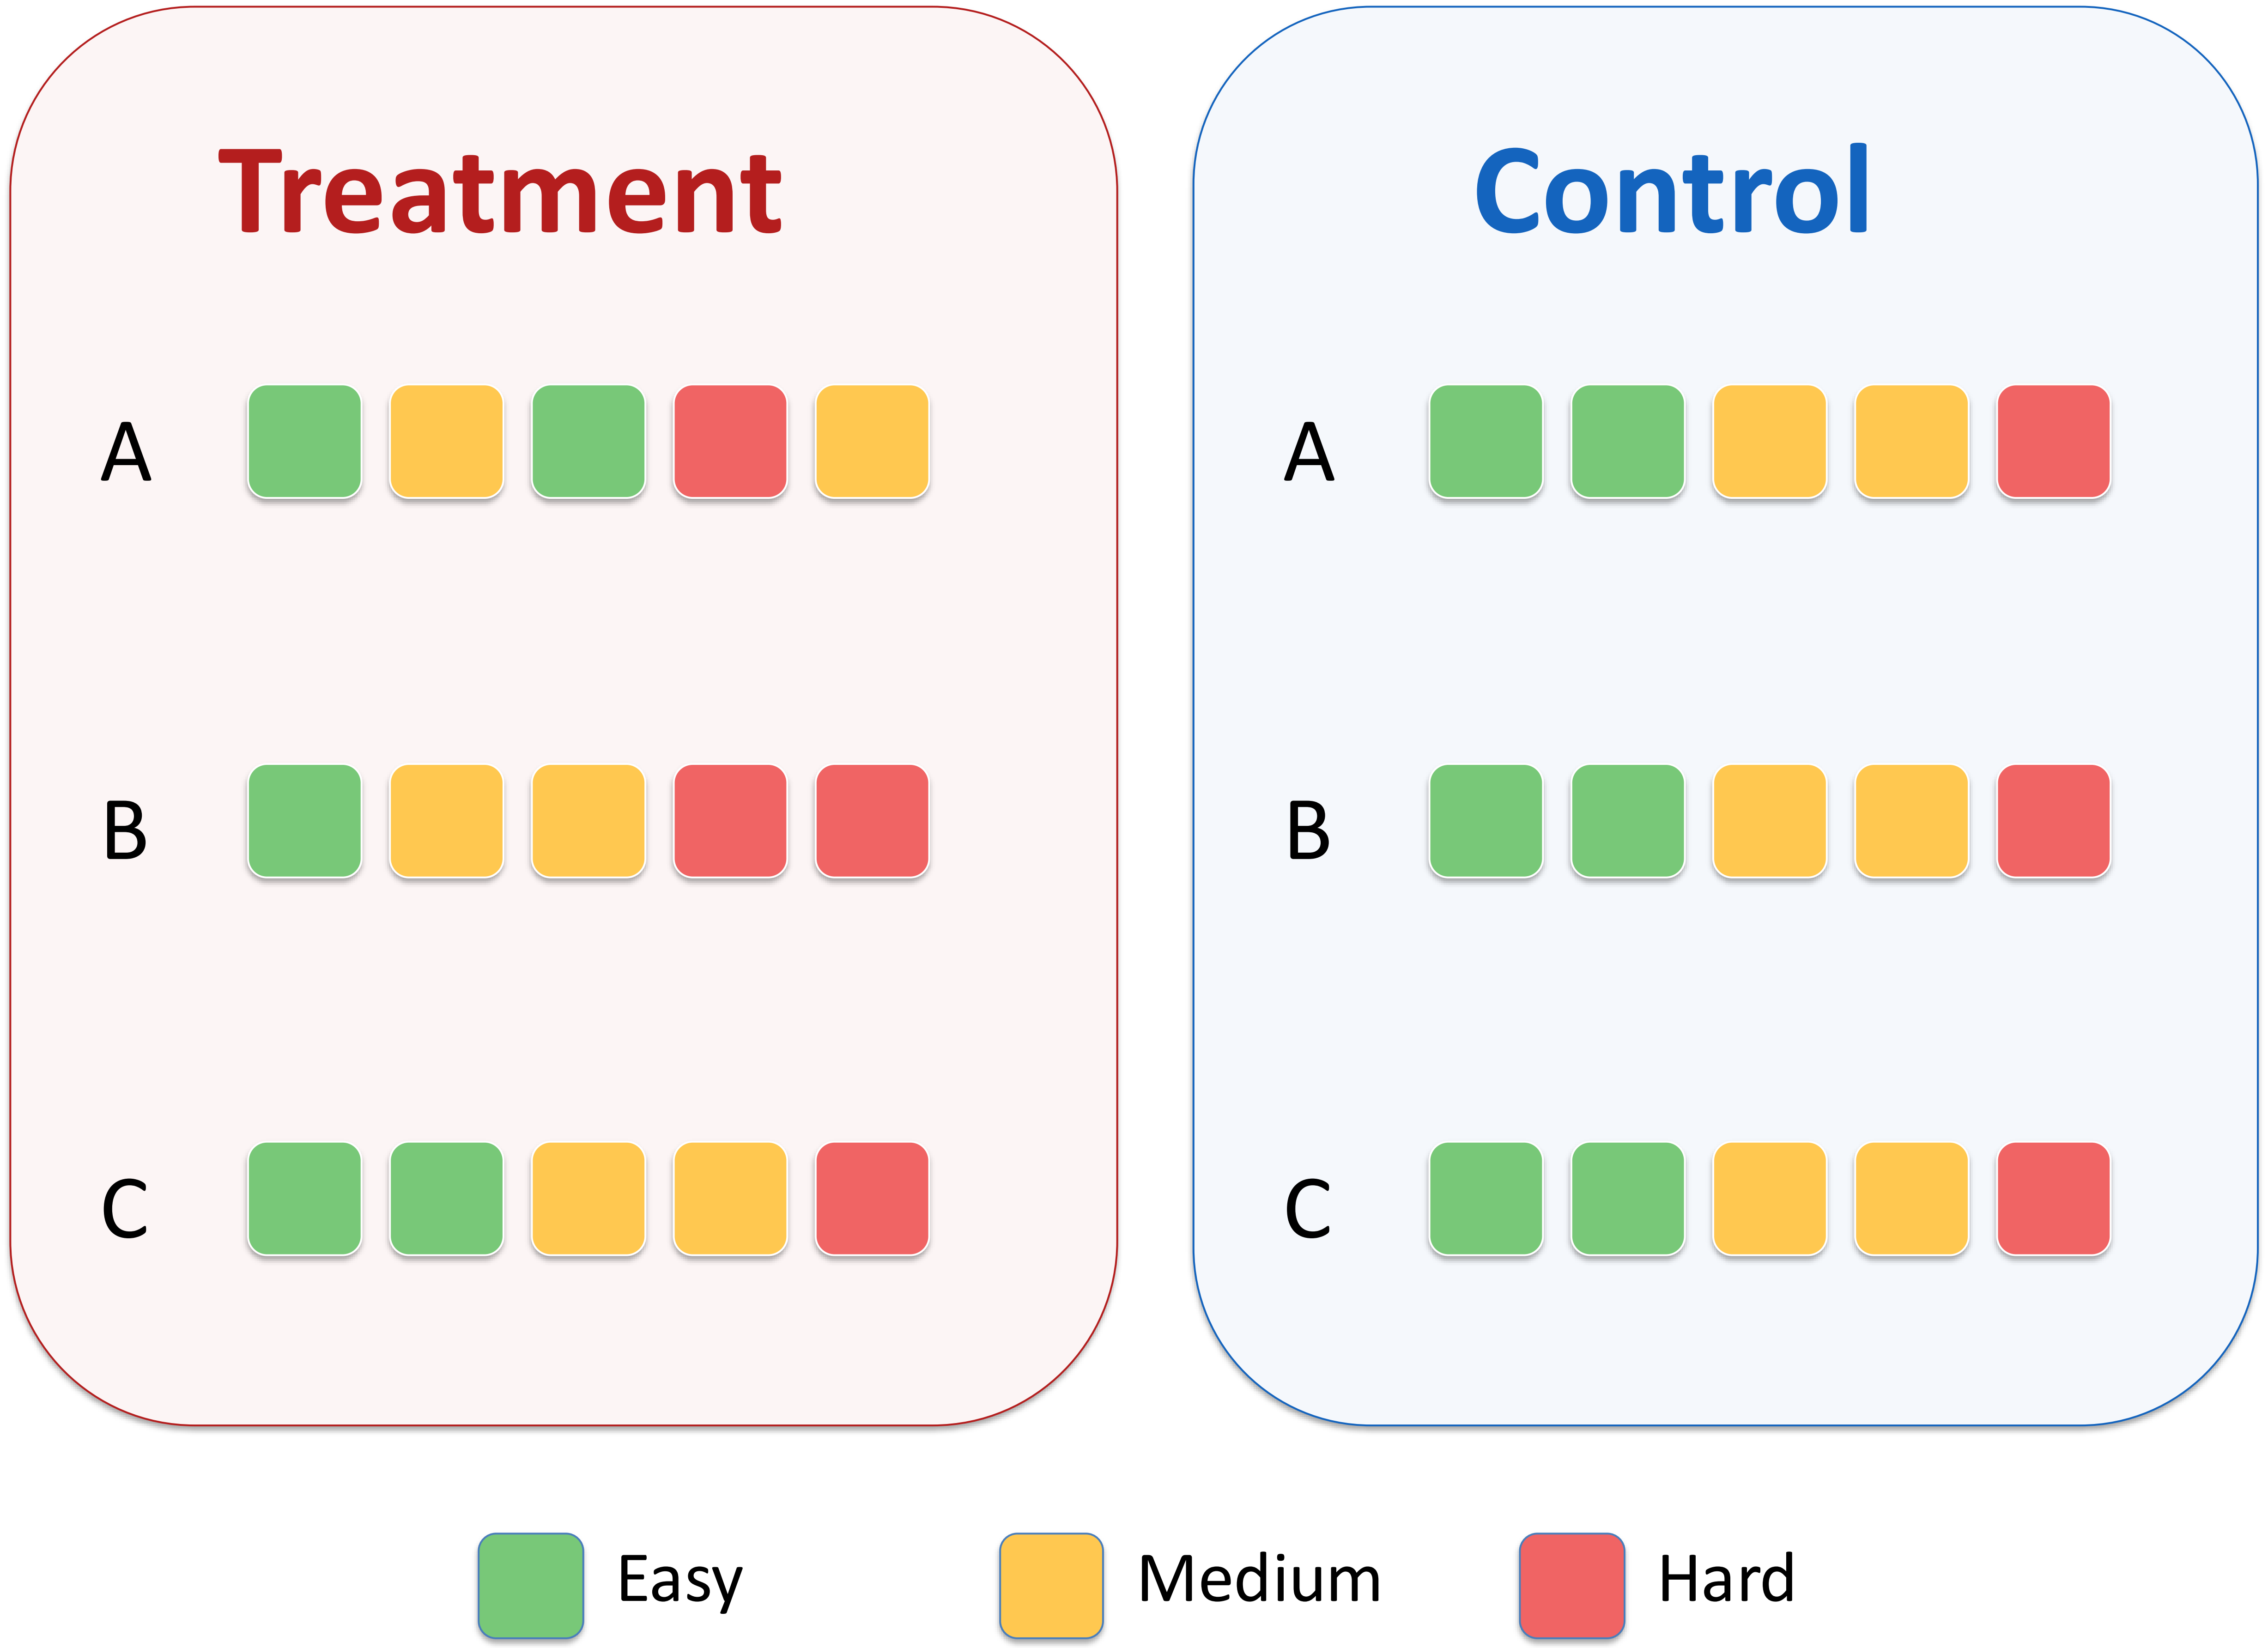

Wharton
University of Pennsylvania

Result: Performance

- RL group achieved a statistically significant higher score by 0.15 standard deviation.
  - Equivalent to 6-9 school months worth of learning [Ritchie & Tucker-Drob 2018; Evans & Yuan 2019]

	Dependent variable: Standardized score		
	(1) OLS	(2) HC1 robust	(3) School-clustered
(Intercept)	−1.273**	−1.273***	−1.273*
	(0.316)	(0.271)	(0.455)
RL	0.150*	0.150*	0.150*
	(0.067)	(0.067)	(0.069)
Python ability (ordinal)	0.125**	0.125**	0.125
	(0.042)	(0.043)	(0.088)
Male	0.224**	0.224**	0.224
	(0.078)	(0.075)	(0.208)
Has programming experience	0.246**	0.246**	0.246
	(0.085)	(0.081)	(0.164)
Parent education (ordinal)	0.018	0.018	0.018
	(0.041)	(0.040)	(0.029)
Academic performance (ordinal)	0.079***	0.079***	0.079***
	(0.021)	(0.022)	(0.020)
School fixed effects	YES	YES	YES
Grade level fixed effects	YES	YES	YES
Motivation controls	YES	YES	YES
R^2	0.253	0.253	0.253
Adjusted R^2	0.226	0.226	0.226
Observations	716	716	716

Wharton
University *of* Pennsylvania

Result: Heterogeneity

- RL gains are larger for beginners. Aligned with existing findings [Doroudi, Aleven, and Brunskill 2019; Lyu et al 2024; Zhang et al 2024].

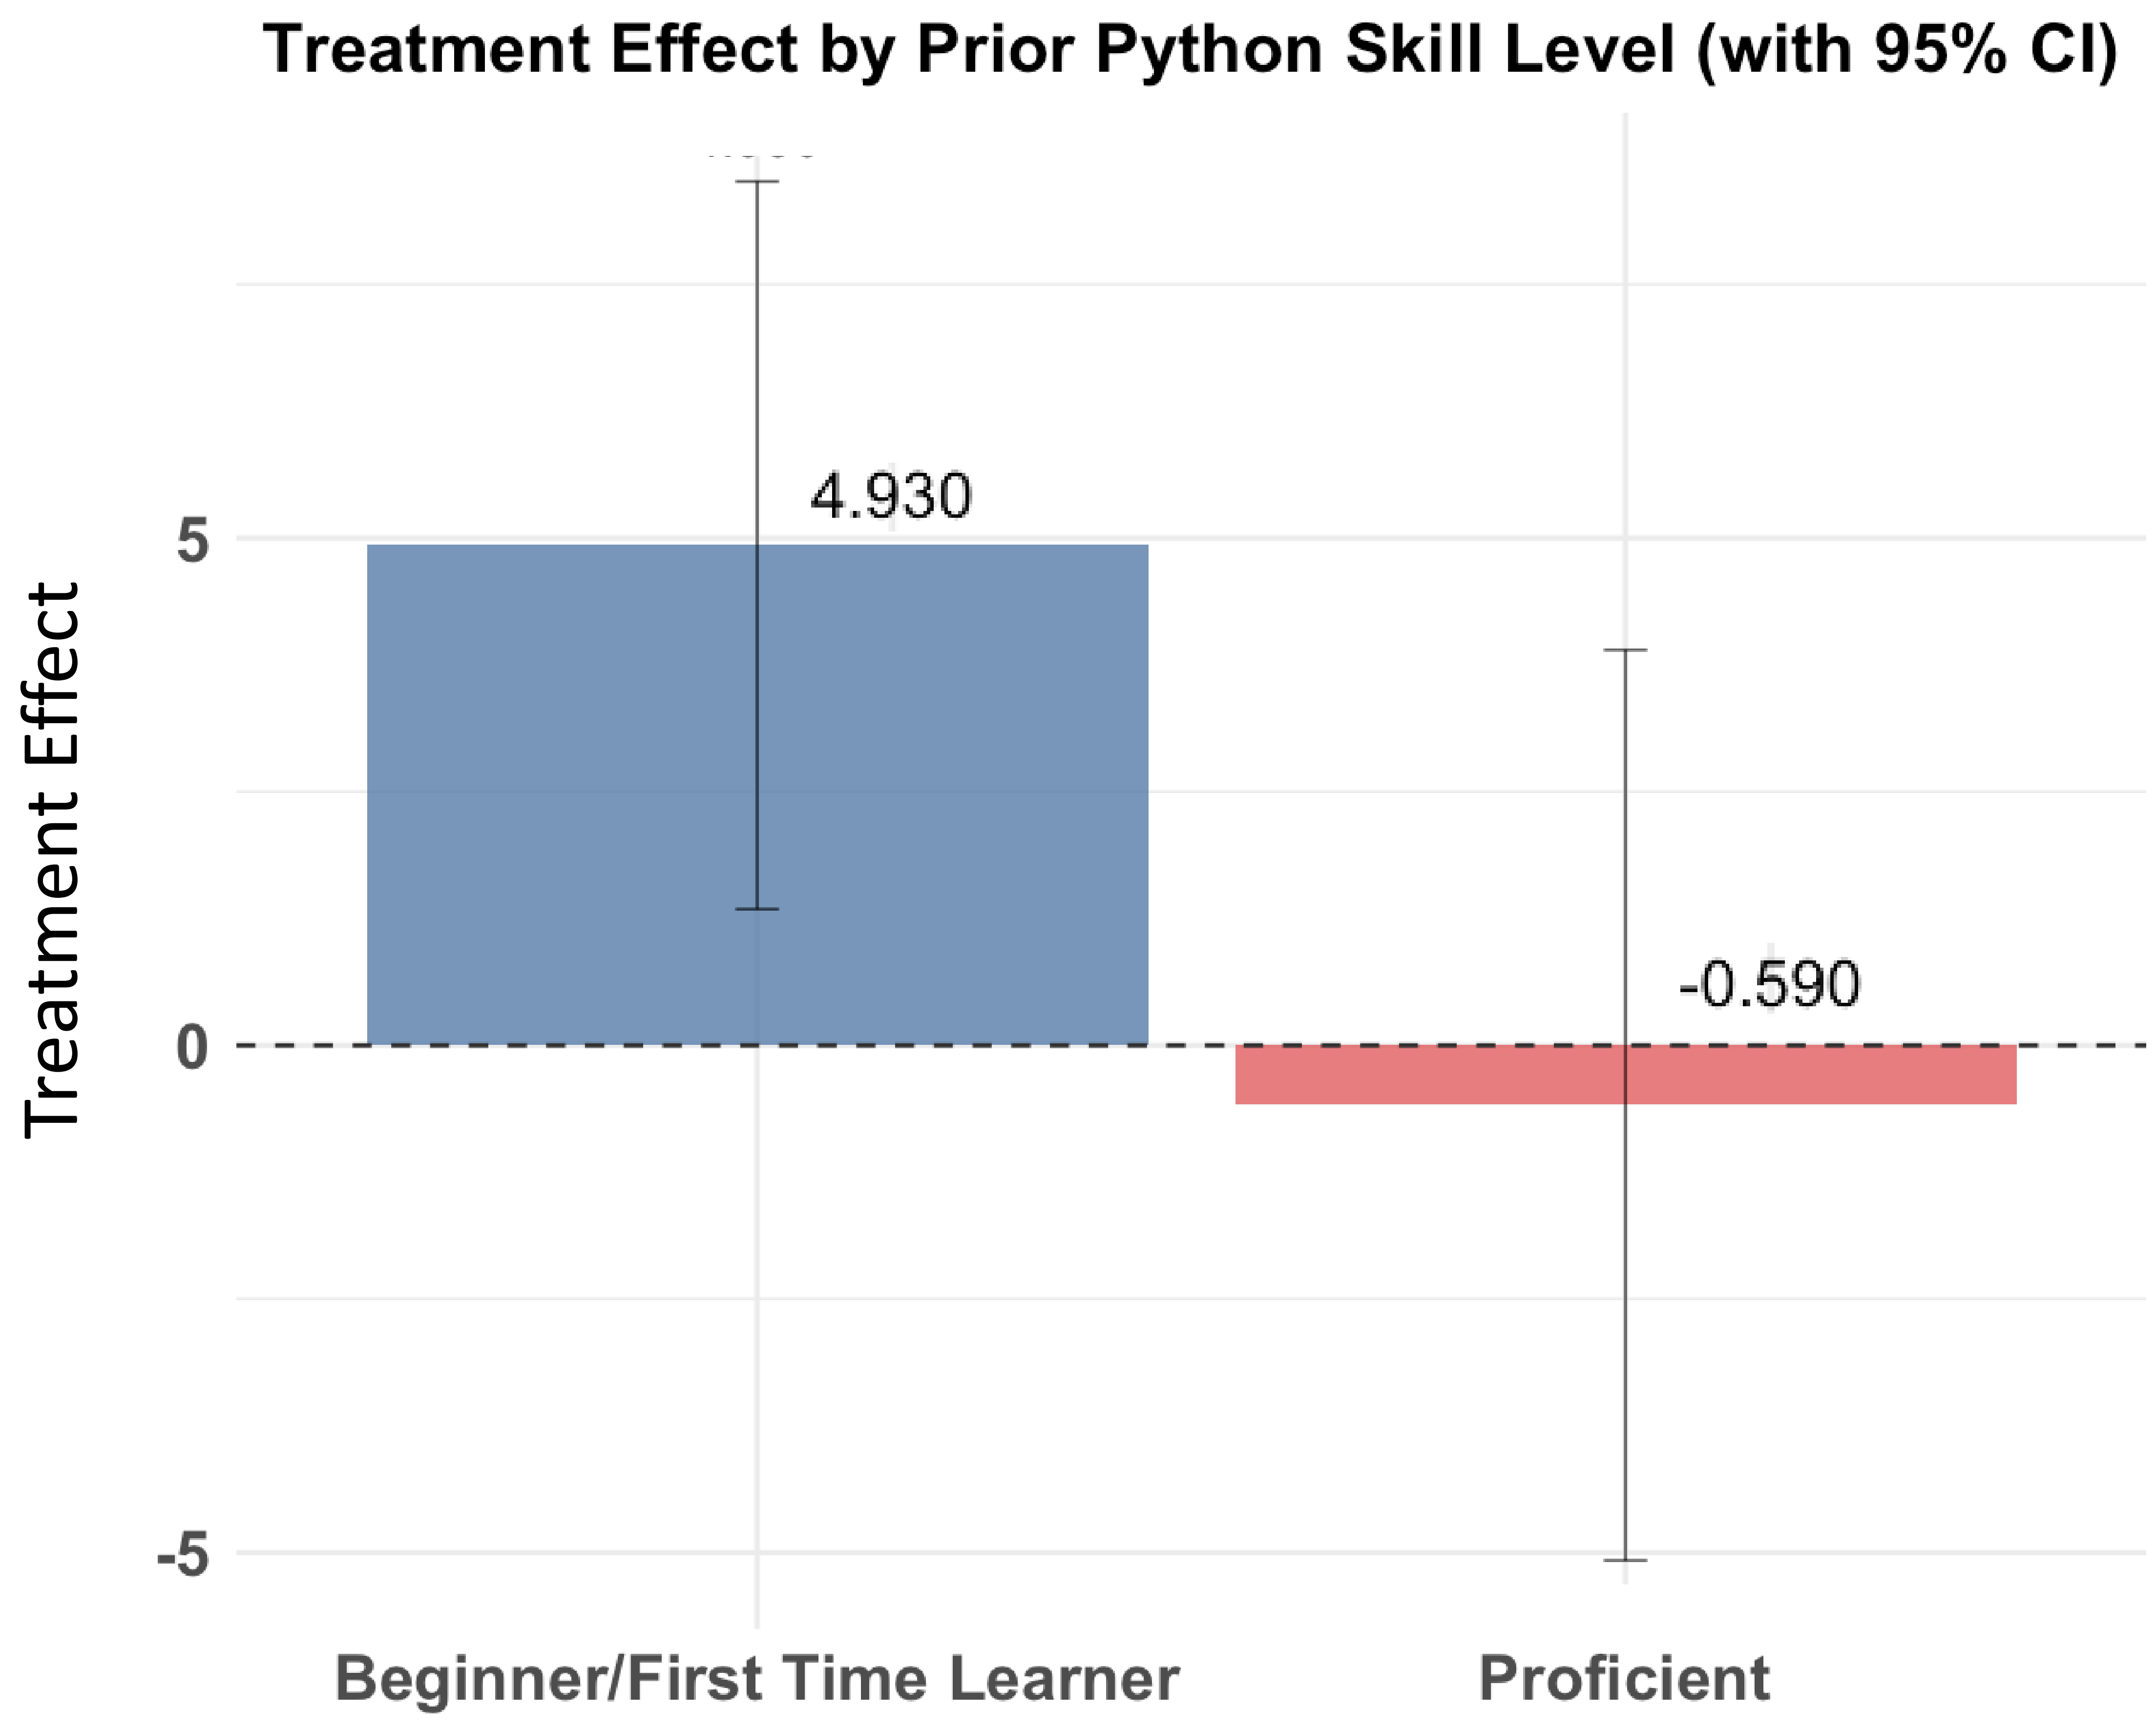

Potential Mechanism: Engagement

- Mediation analysis [Tingley et al. 2014] suggests that RL group spent more time and more attempts.
- They also have statistically significantly higher chat quality score – use AI to learn instead of simply seeking for answer.

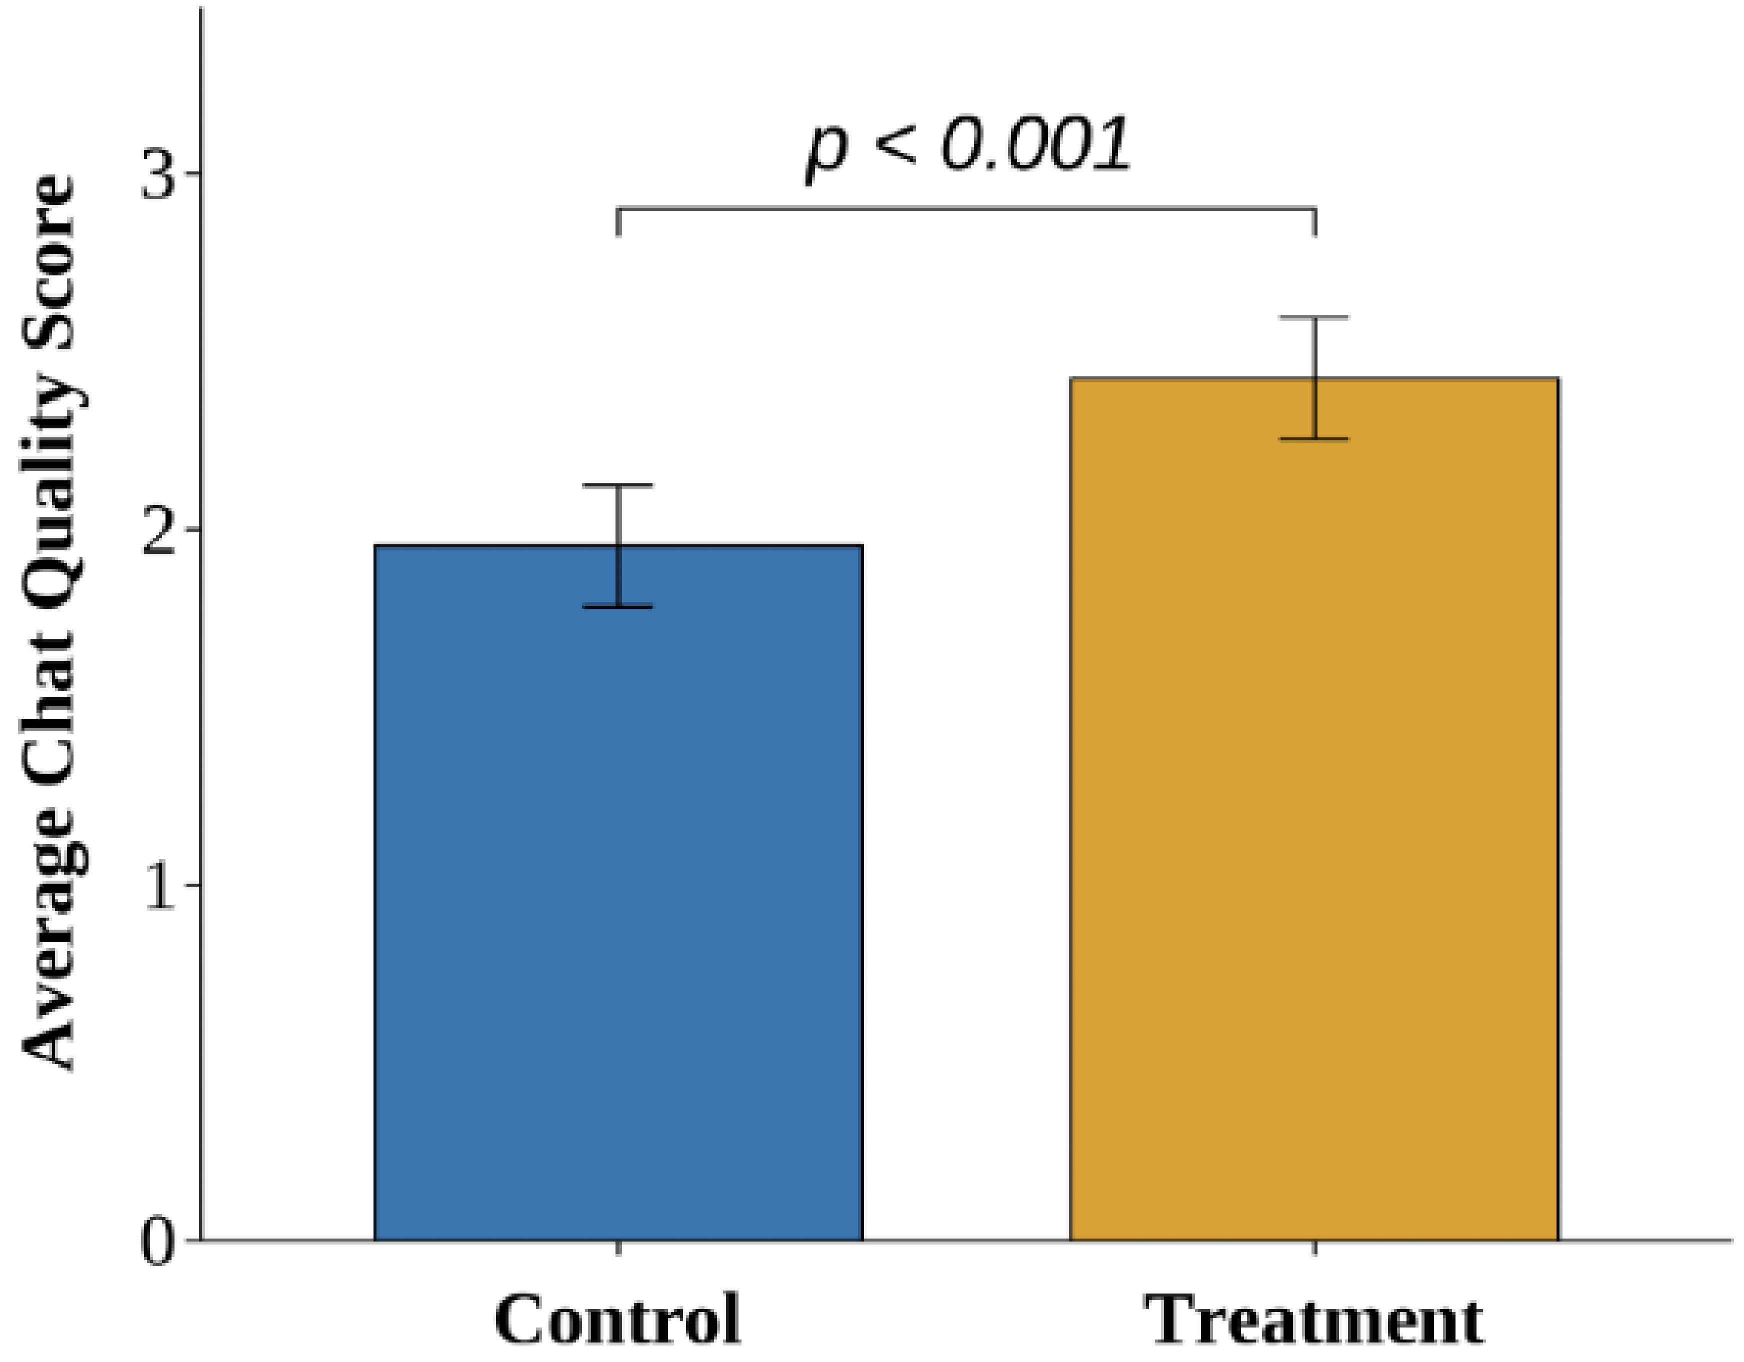

Thank You!

angelchg@wharton.upenn.edu

https://angel-chung.github.io

Penn Engineering
UNIVERSITY of PENNSYLVANIA

William and Phyllis
MACK INSTITUTE
for INNOVATION MANAGEMENT

PennClub

National Taiwan University

DEPARTMENT OF EDUCATION
TAIPEI CITY GOVERNMENT

AIC American innovation center

AMERICAN
INSTITUTE IN
TAIWAN
美國在台協會

Wharton
UNIVERSITY of PENNSYLVANIA