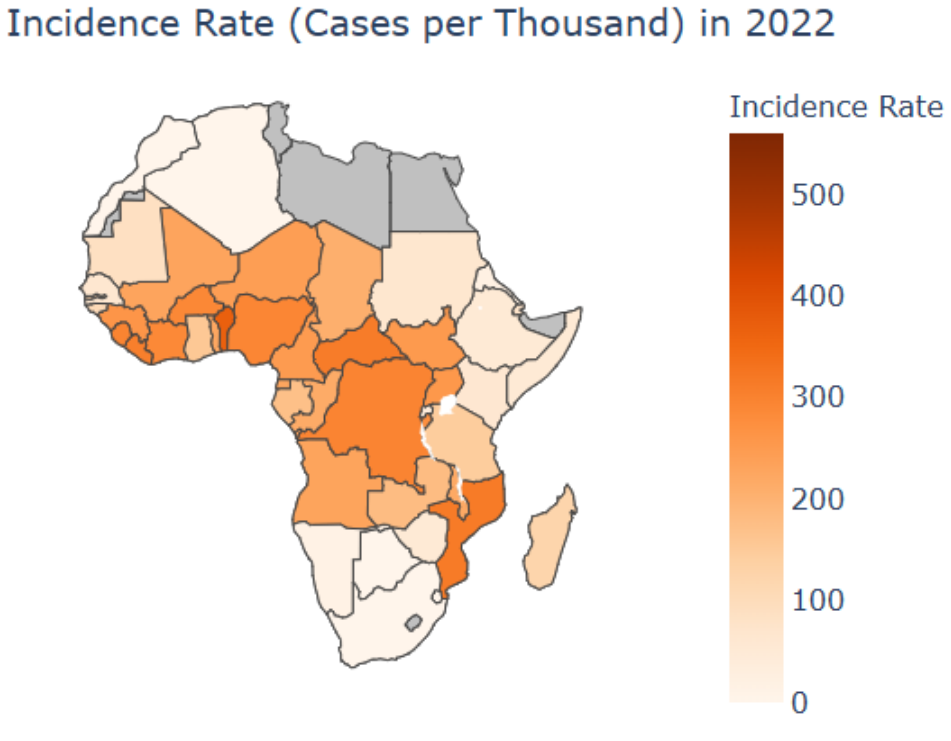
Calibrated Benchmarking and Adaptive Learning for ITN Allocation in Malaria Control

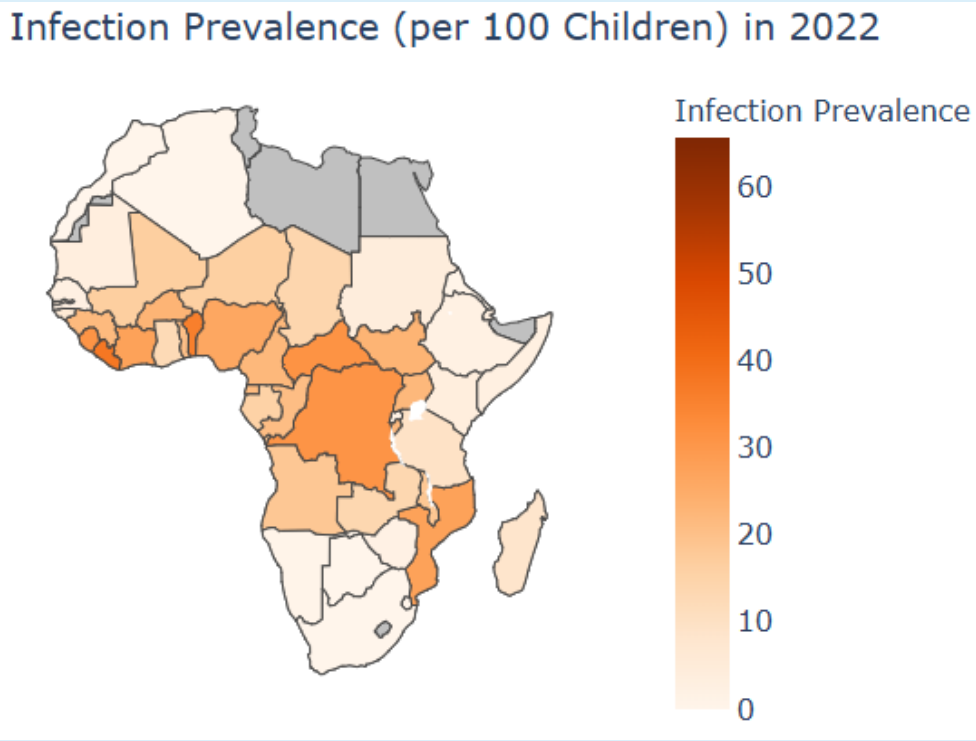
Rajarshi Mandal

PRIMES

HARVARD MEDICAL SCHOOL

Problem Importance

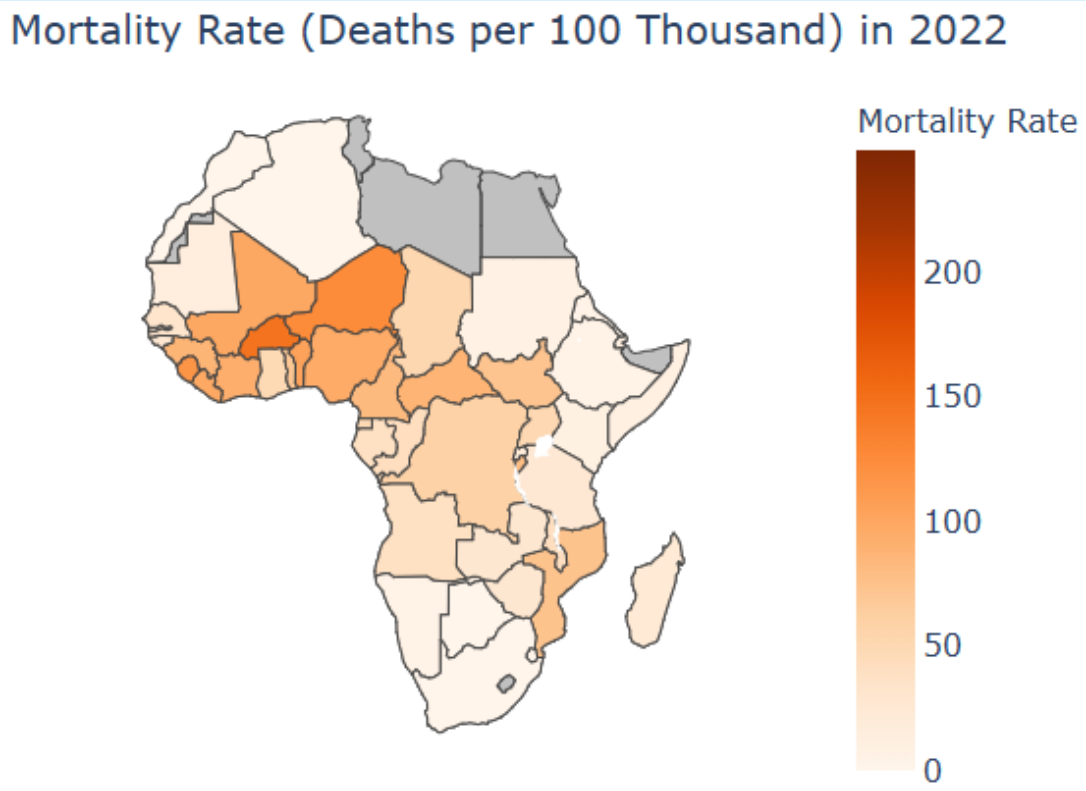




- 282 million cases annually
- 610,000 deaths annually

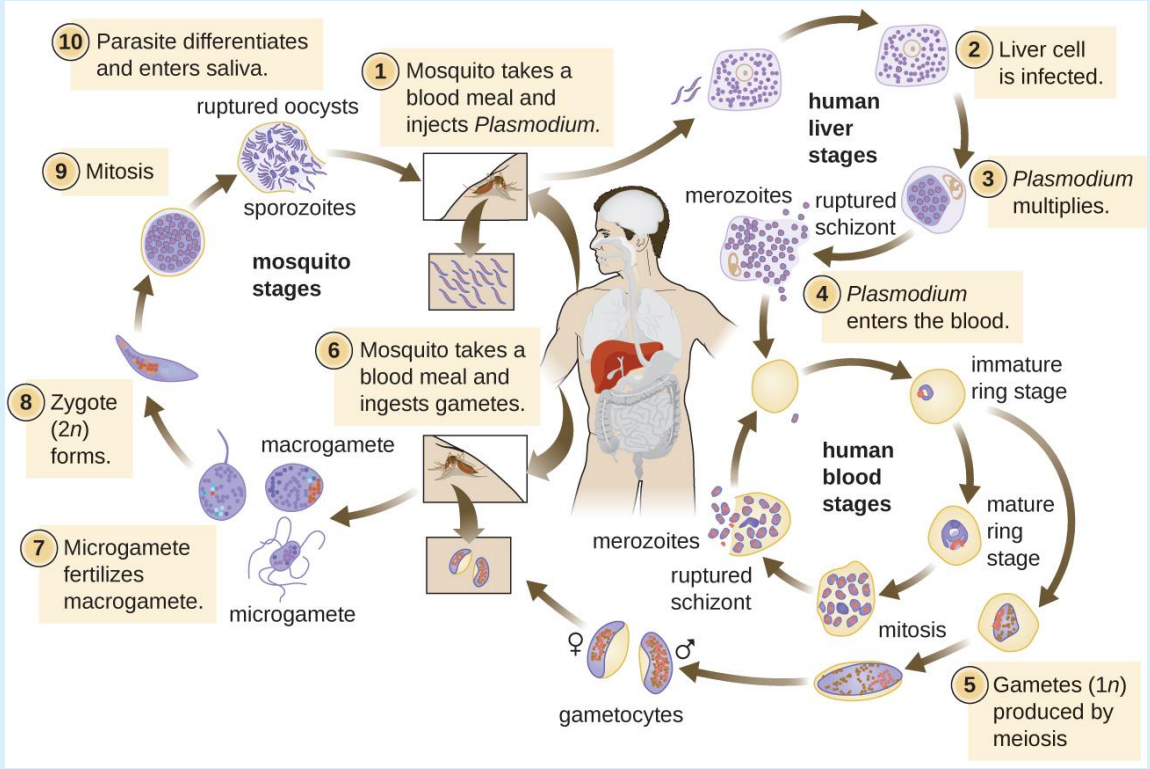
Problem Importance

- 95% of deaths occur in Africa
- Three quarters are children under five
- Stunts childhood development
- Diminishes workforce productivity
- Disproportionately affects rural areas



Problem Importance

- Caused by Plasmodium parasite
- *Anopheles* mosquitoes bite at night
- ITNs are physical and chemical barrier



Problem Importance

- Insecticide-treated nets (ITNs)
- Among the most cost-effective programs
- Trials show ~50% fewer clinical cases
- Distributed in mass campaigns



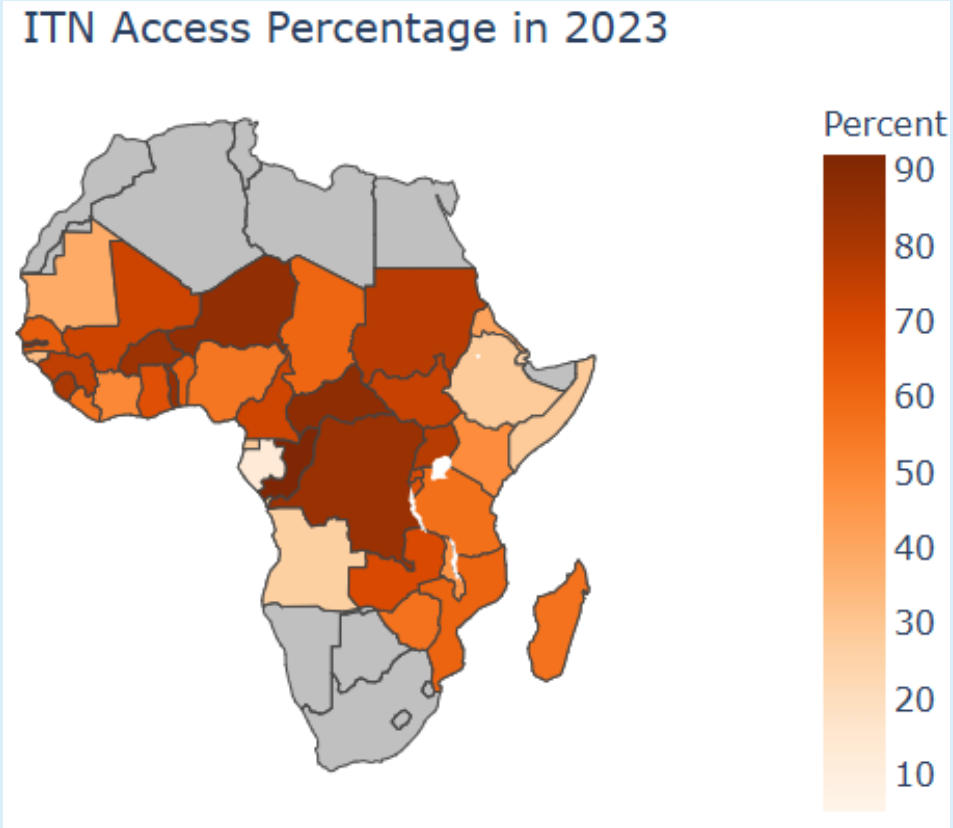
Problem Importance

- Selling new ITNs
- Repurposing ITNs for fishing
- Pyrethroids harm the environment
- Ripple Africa charity’s approach



Problem Importance

- $1 billion of donor nets in 2025 alone
- A fixed supply split across regions
- Allocation changes impact per dollar



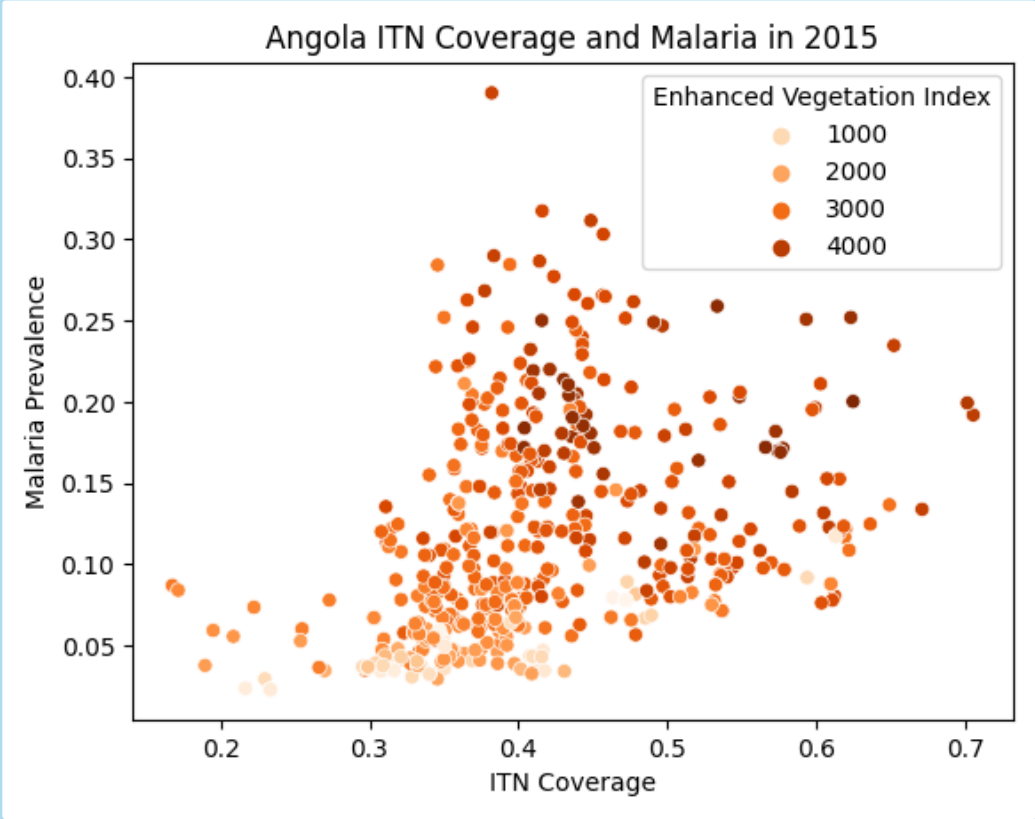
Problem Scope

- Allocate a fixed budget across regions
- Periodic mass campaigns
- Delayed and seasonal transmission
- Net attrition and resistance



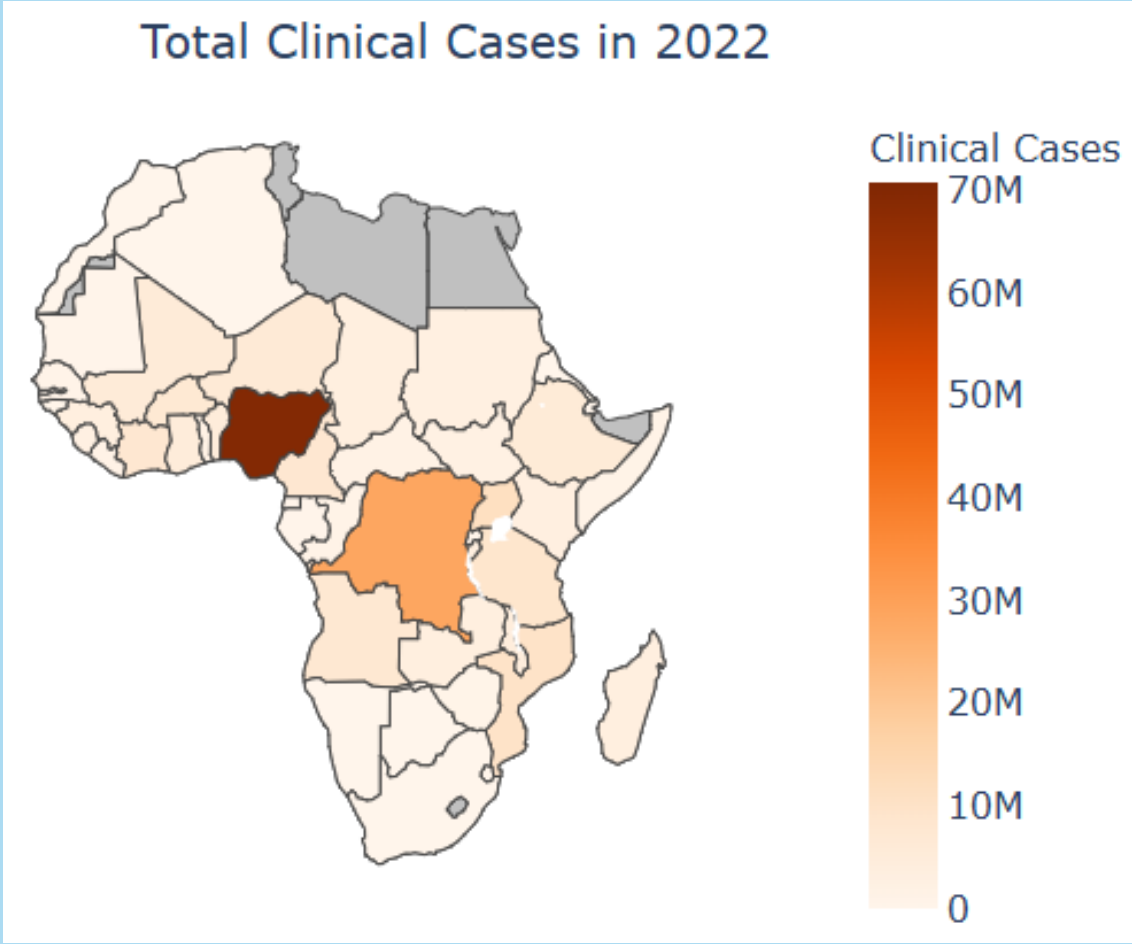
Problem Scope

- Commonly nets split by population
- WHO target of one net per two people
- Prior RL used fictional environments
- Rarely compared to this baseline
- No prior held-out, validation results

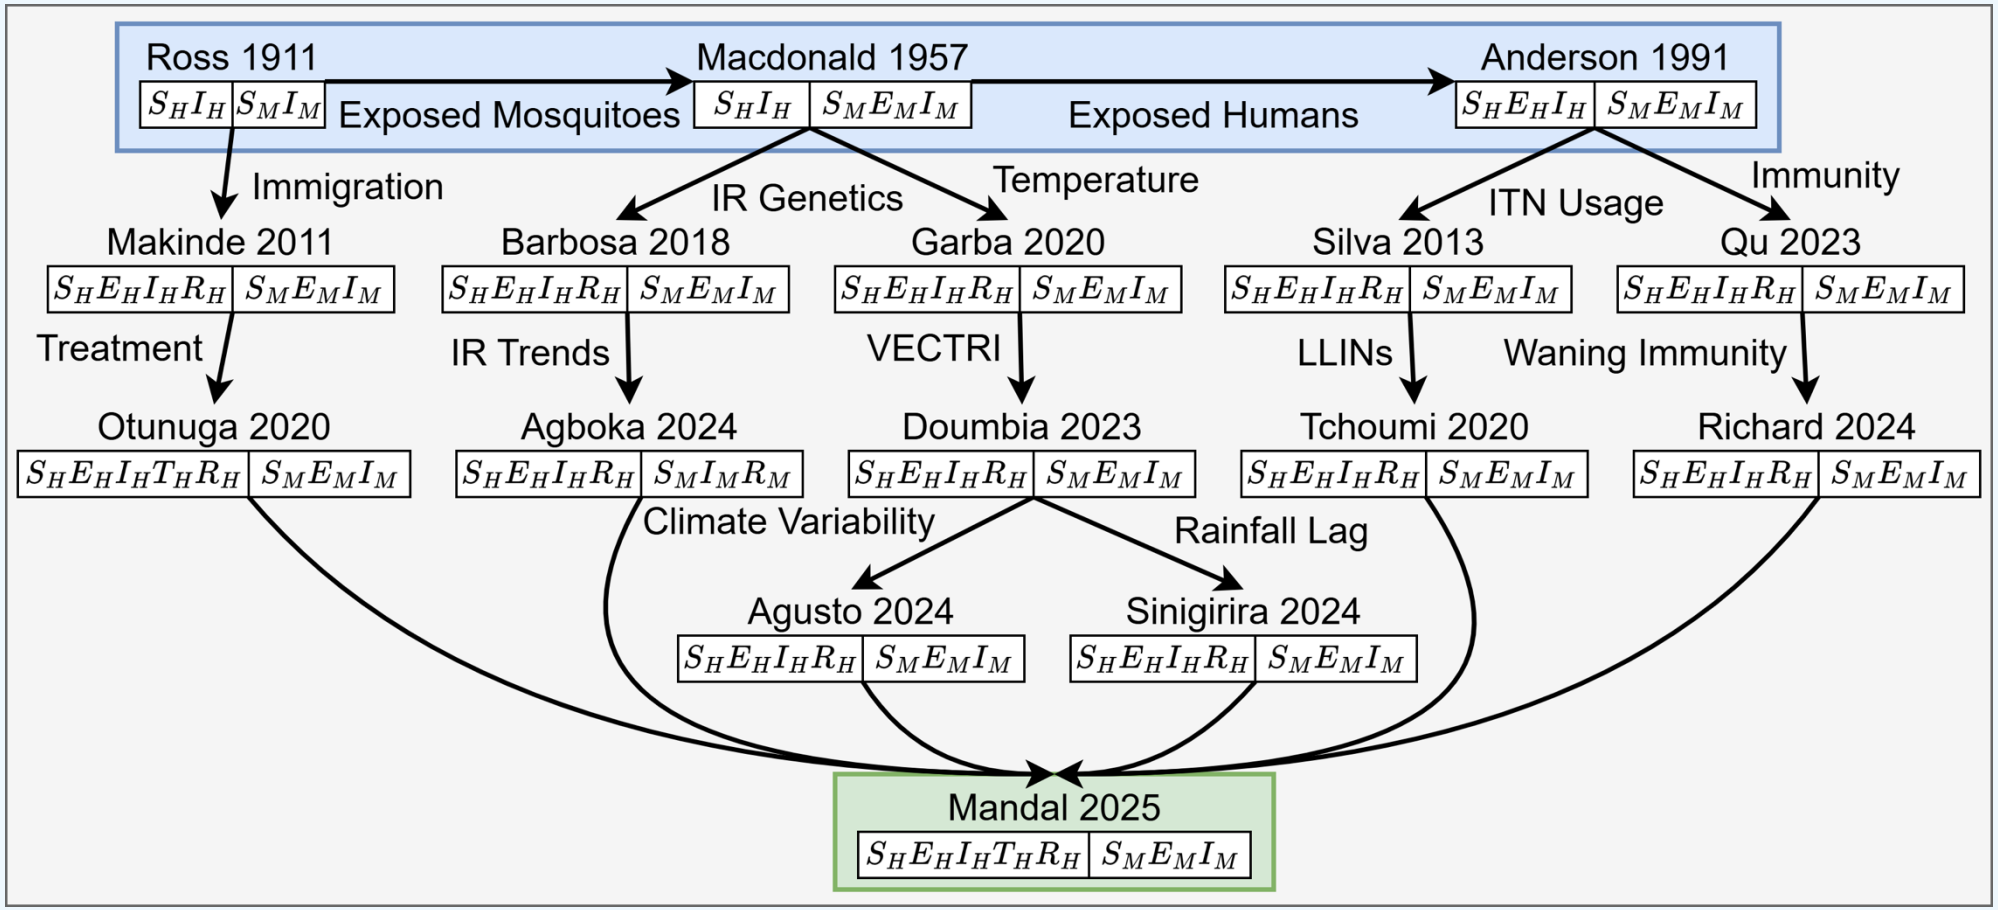


Research Questions

- Can learned allocation beat practice?
- What if transmission is unknown?

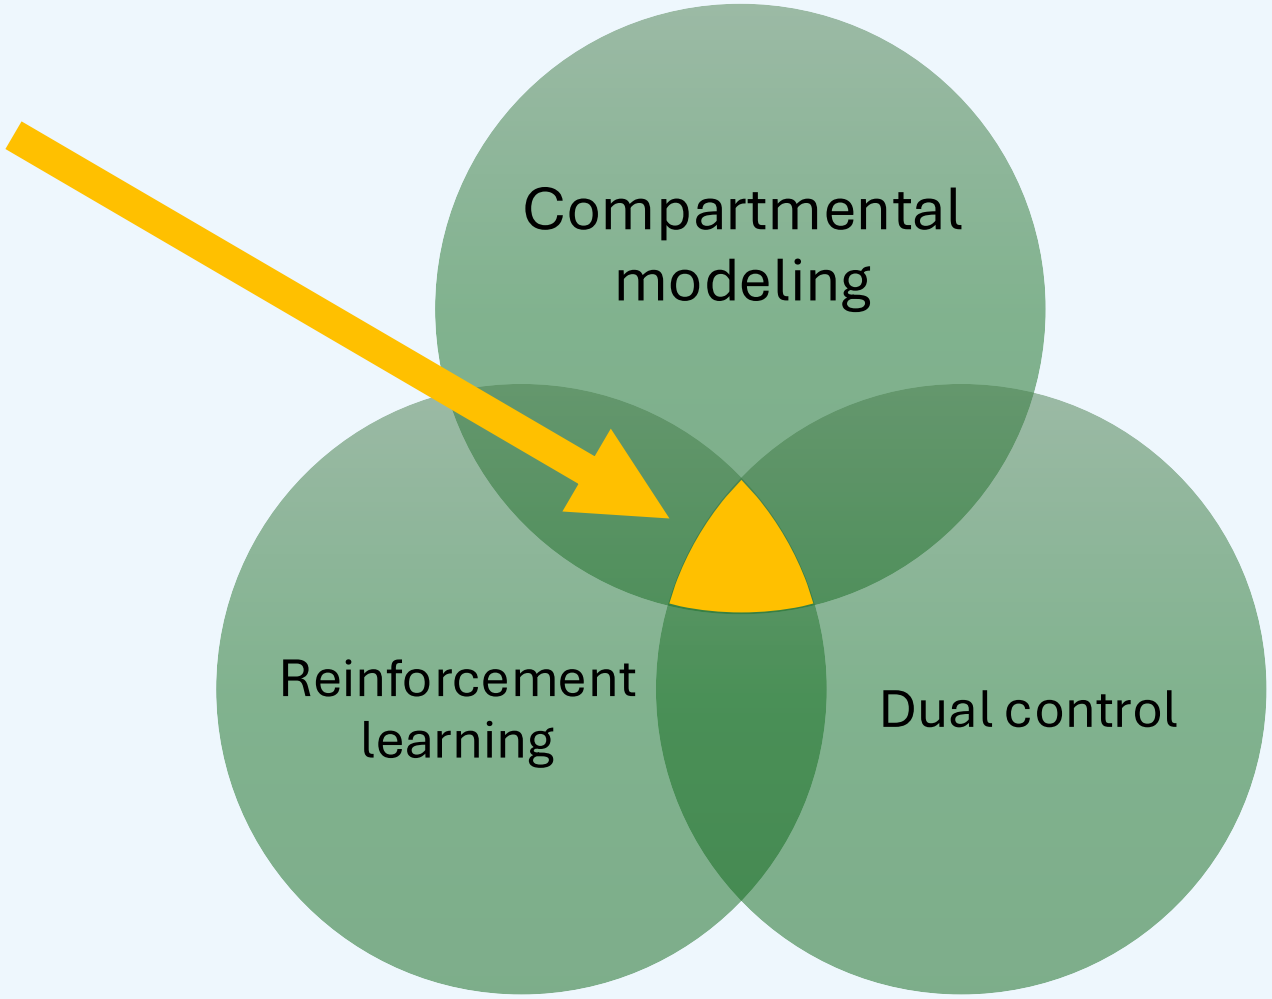


Literature Review



Literature Review

- Illustrative model parameters
- Others stop short of allocation
- RL ignores budget constraints
- No practical baseline
- I unify these with dual control



Datasets


map


MICS

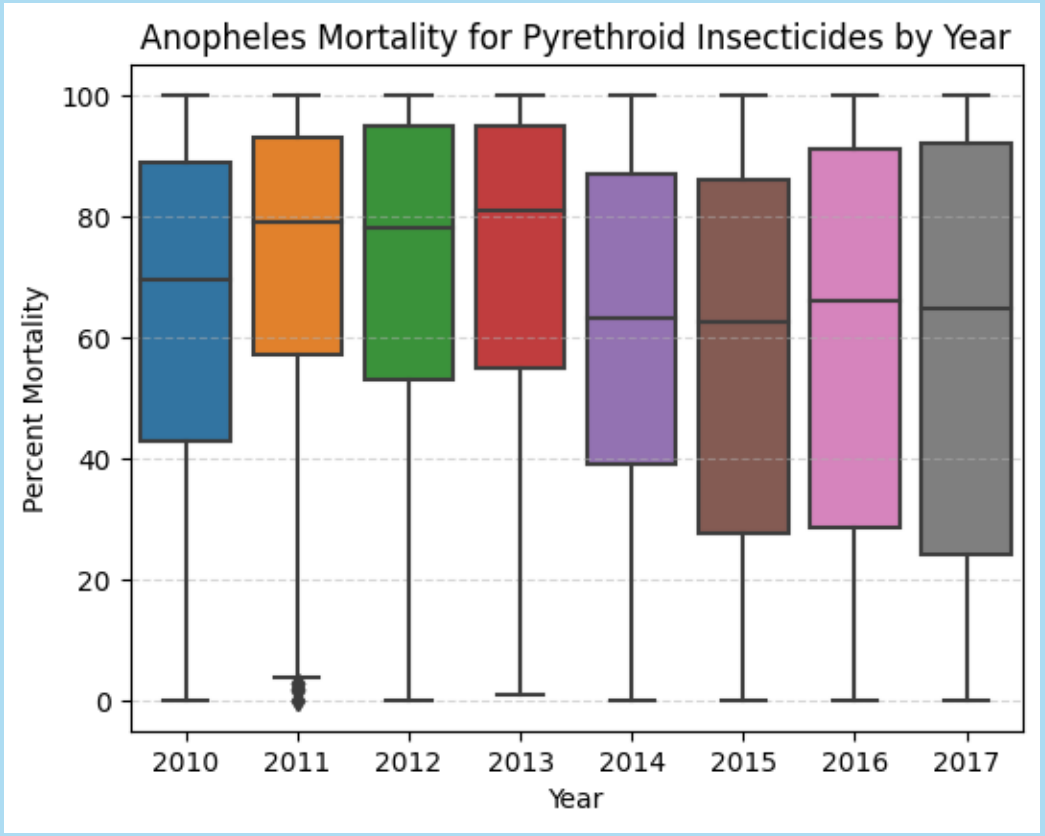

World Health
Organization

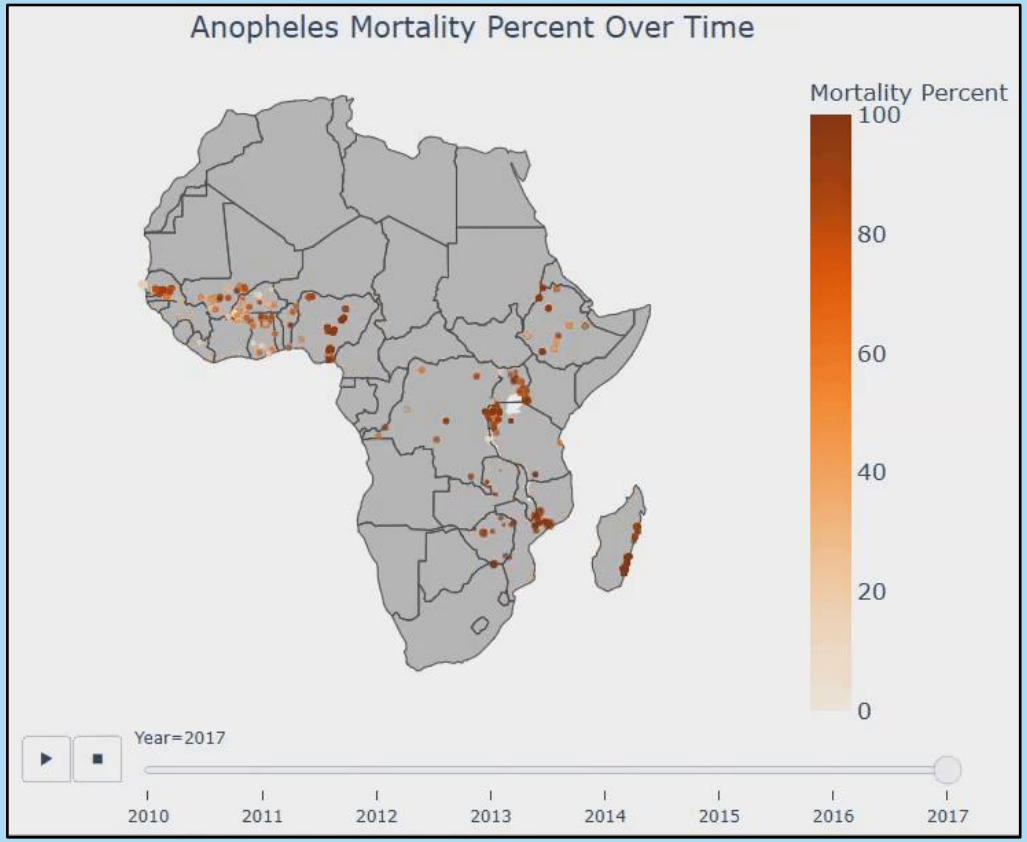

VESTERGAARD
IMPACTING PEOPLE


NASA
EARTHDATA

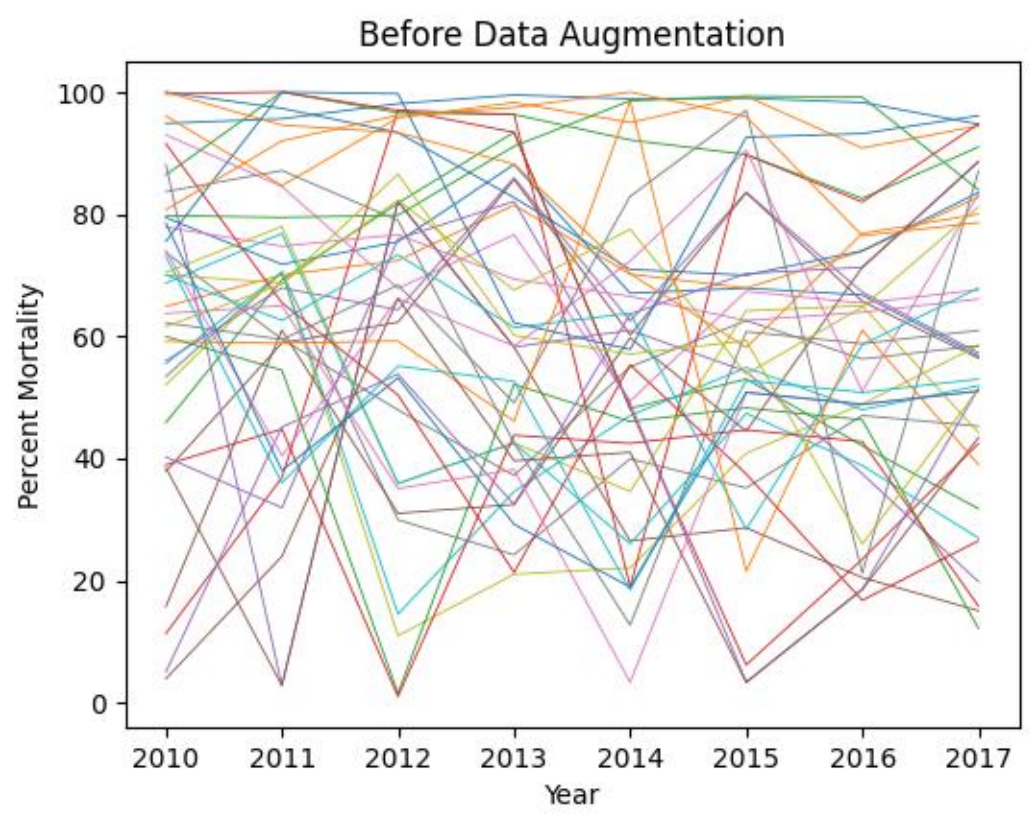

The
DHS Program
Demographic and Health Surveys

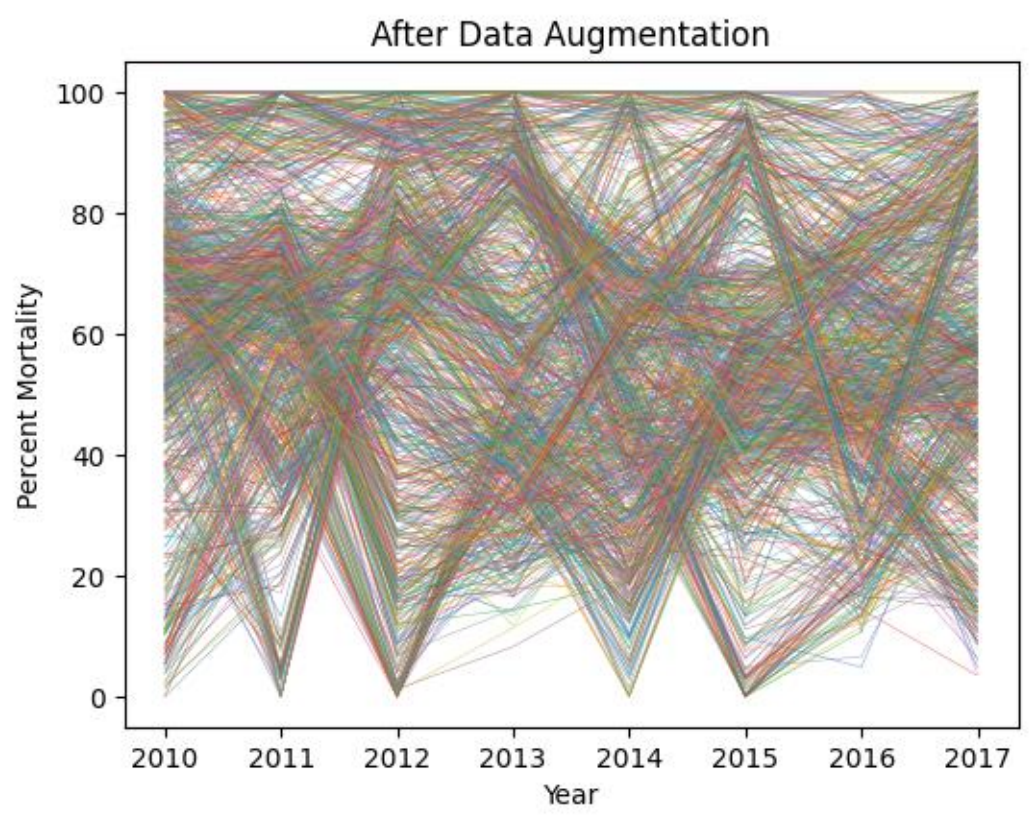
EDA


Anopheles Mortality for Pyrethroid Insecticides by Year
Percent Mortality
100
80
60
40
20
0
2010
2011
2012
2013
2014
2015
2016
2017
Year


Anopheles Mortality Percent Over Time
Mortality Percent
100
80
60
40
20
0
Year=2017
2010
2011
2012
2013
2014
2015
2016
2017

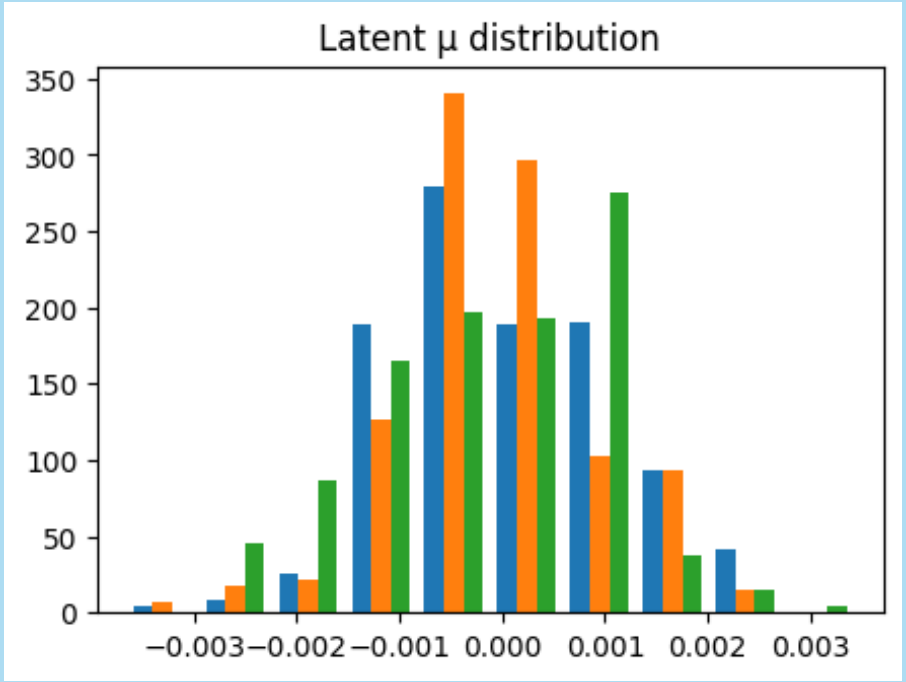
Methodology



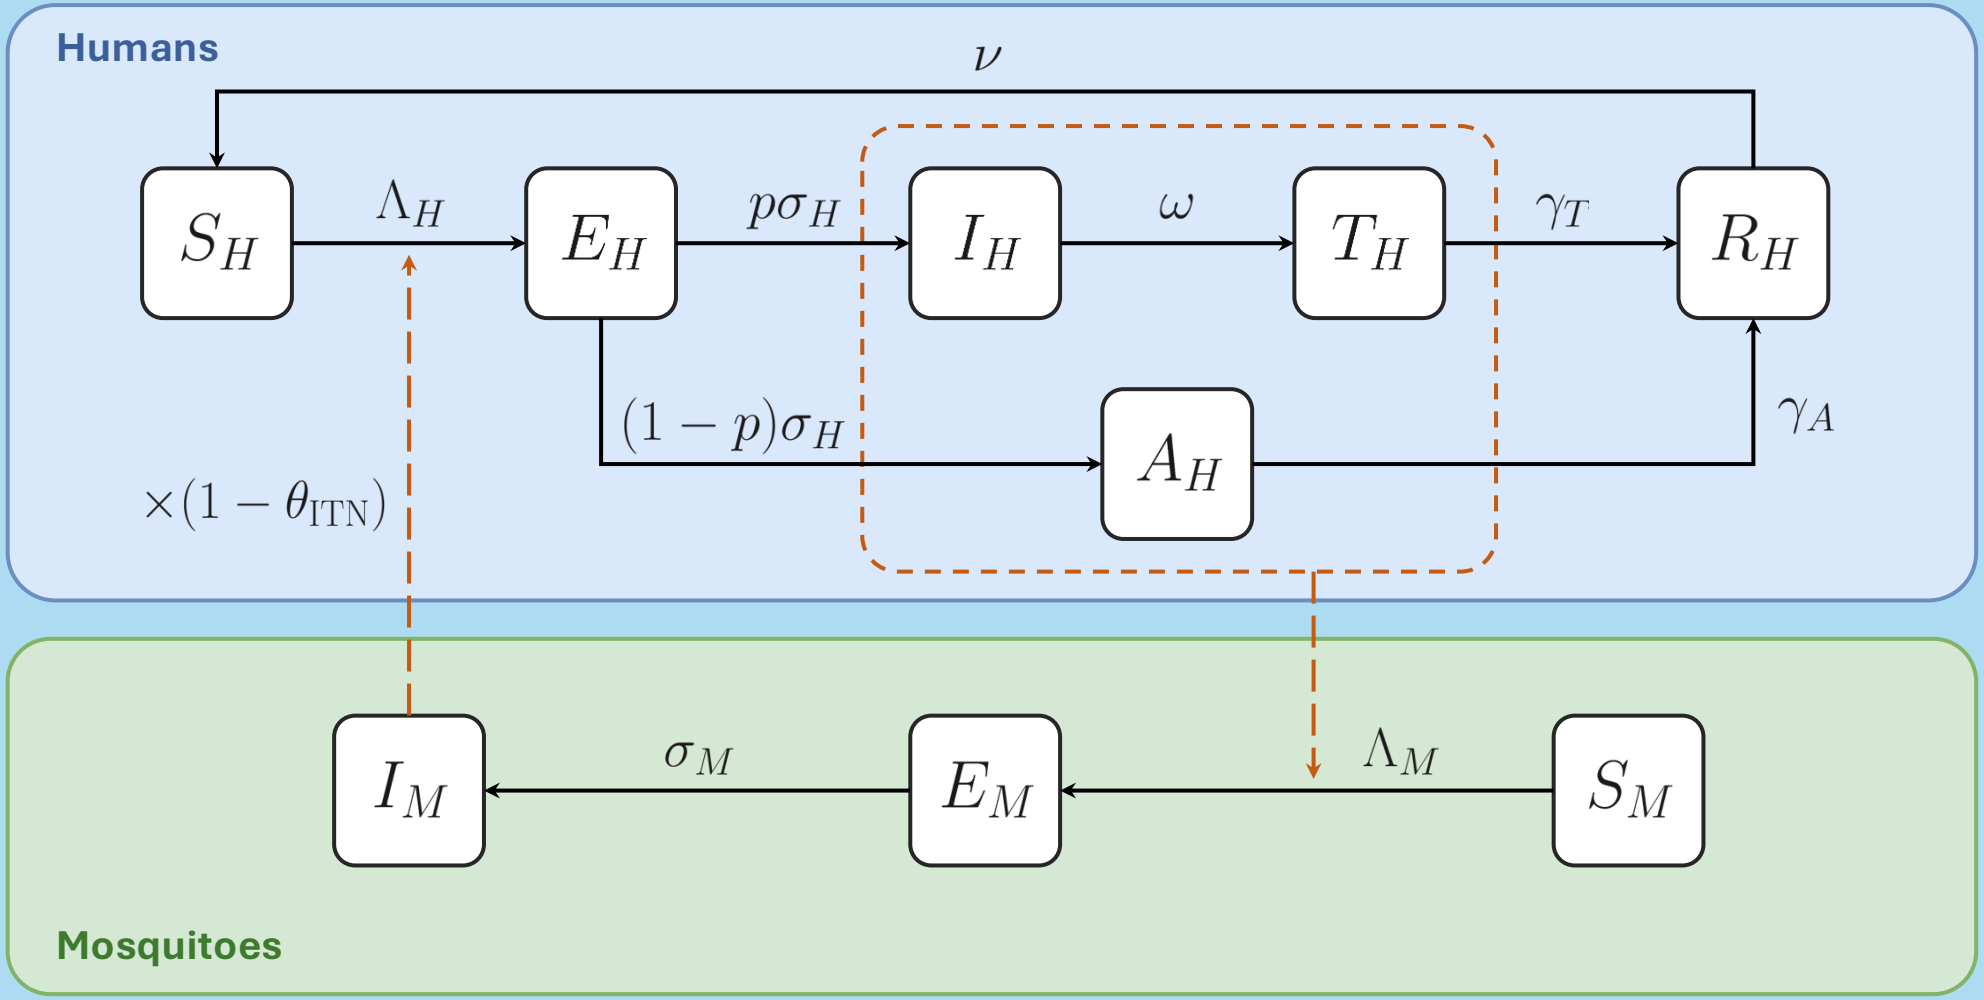


Methodology

- Resistance surveillance is sparse
- Bioassay mortality mapped to R(t)
- Augment by interpolation and jitter
- Autoencoder learns trajectories
- Samples stress test allocation



Transmission Model



Transmission Model

- Force of infection on humans

$$\Lambda_H = a(t)\, c\, \frac{I_M}{N_M}\, (1 - \theta_{\mathrm{ITN}})$$

- Force of infection on mosquitoes

$$\Lambda_M = a(t)\, b\, \frac{I_H + \kappa T_H + \kappa_A A_H}{N_H}$$

- Seasonal biting rate

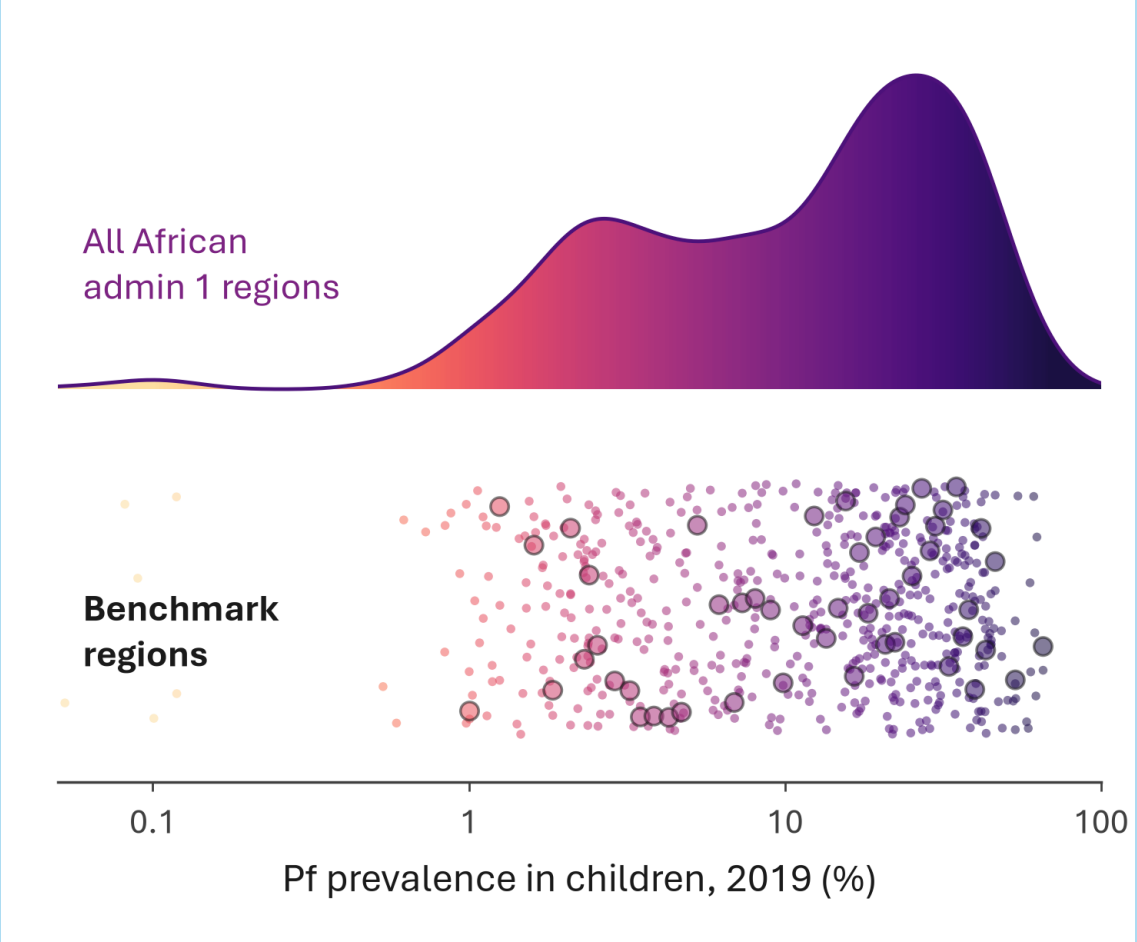
$$a(t) = \bar{a} \left(1 + 0.2 \, \sin \frac{2\pi t}{180} \right)$$

- ITN protection over time

$$\theta_{\mathrm{ITN}}(t) = C\, e^{-\delta \tau} \left(1 - 0.5\, R(t) \right)$$

Asymptomatic Carriage

- Pf prevalence spans 1.8% to over 60%
- Fraction p become symptomatic
- Remaining population carries it silently
- Carriers clear slowly, over ~300 days
- Model fits prevalence and incidence



Problem Formulation

- Constrained MDP over K regions, acting only at campaign days

$$n_t = (n_{1,t}, \ldots, n_{K,t}) \in \mathbb{R}^K_{\geq 0}, \qquad \sum_i n_{i,t} = B_t$$

- Normalized infection burden

$$\mathrm{Inf}(t) = \frac{\sum_i (E_{H,i} + I_{H,i})}{\sum_i N_{H,i}}$$

- Reward is burden over a window of L days

$$r_t = \sum_{s=t}^{t+L-1} \mathrm{Inf}(s), \qquad \max_\pi \ \mathbb{E}_\pi \Big[- \sum_t \gamma^t r_t \Big]$$

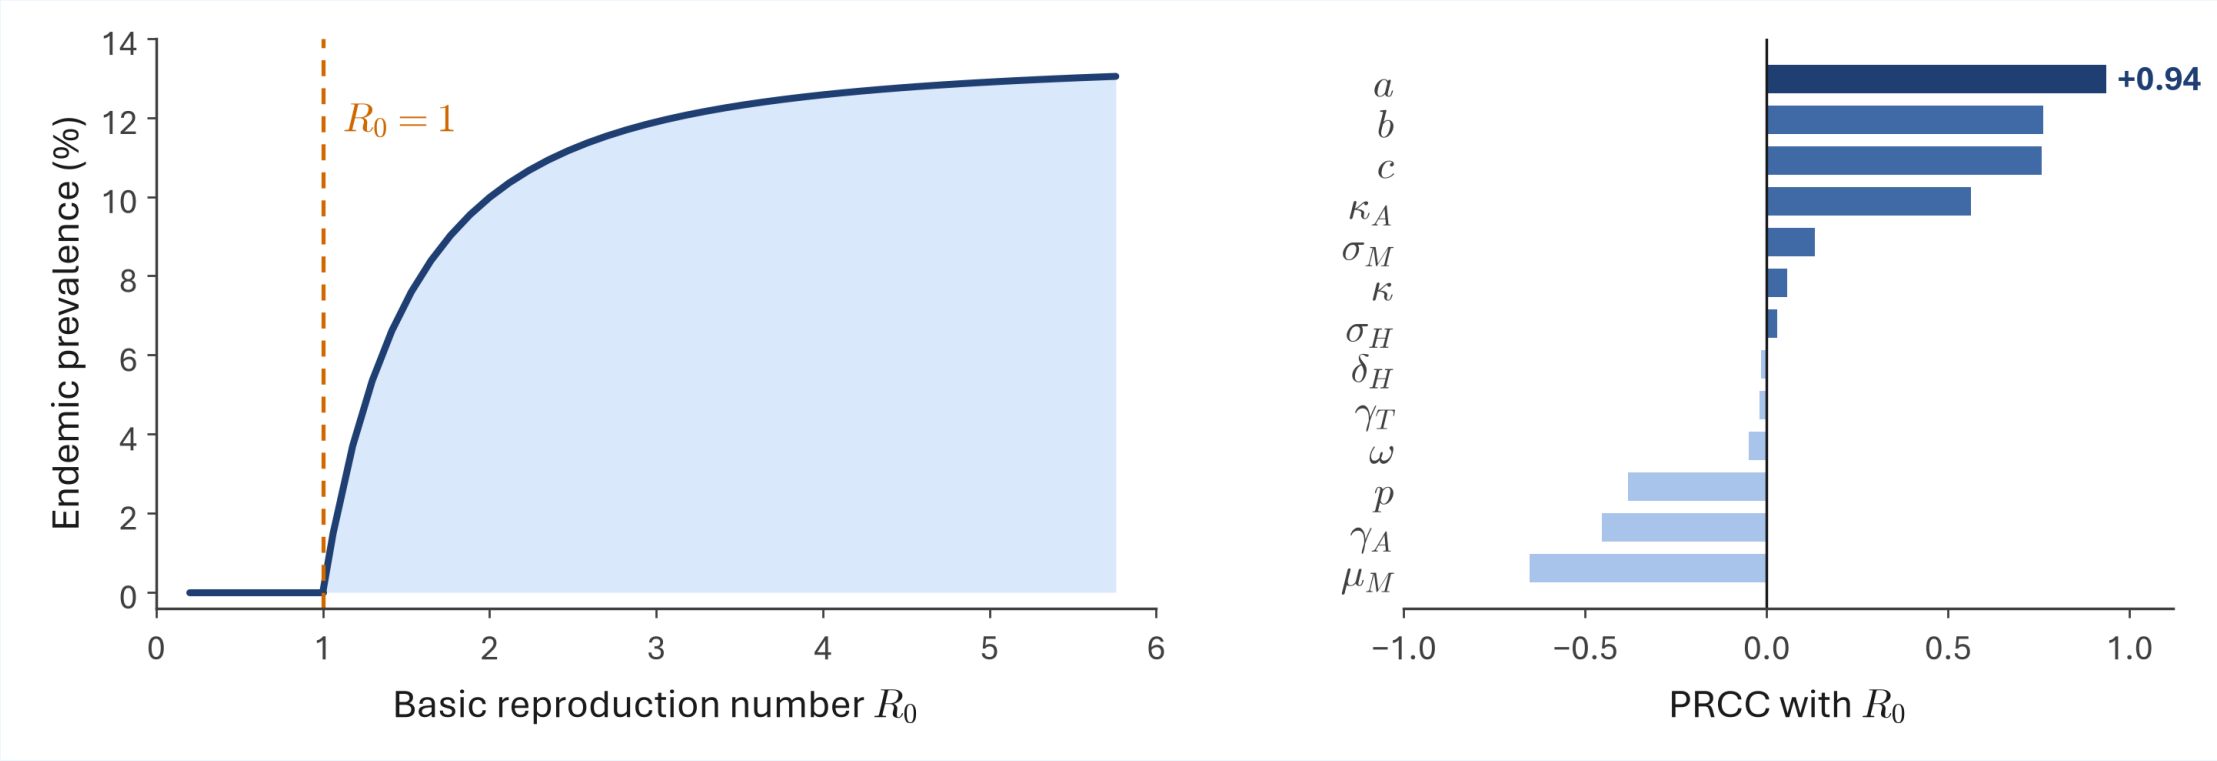
Reproduction Number

- Next-generation matrix gives $R_0 = \rho(FV^{-1})$

$$R_0 = \sqrt{\bar{a}^2 bc \cdot \frac{\sigma_M}{\mu_M(\sigma_M + \mu_M)} \cdot \frac{\sigma_H}{\sigma_H + \mu_H} \cdot \frac{1 + \frac{\kappa\omega}{\mu_H + \delta_T + \gamma_T}}{\omega + \mu_H + \delta_H}}$$

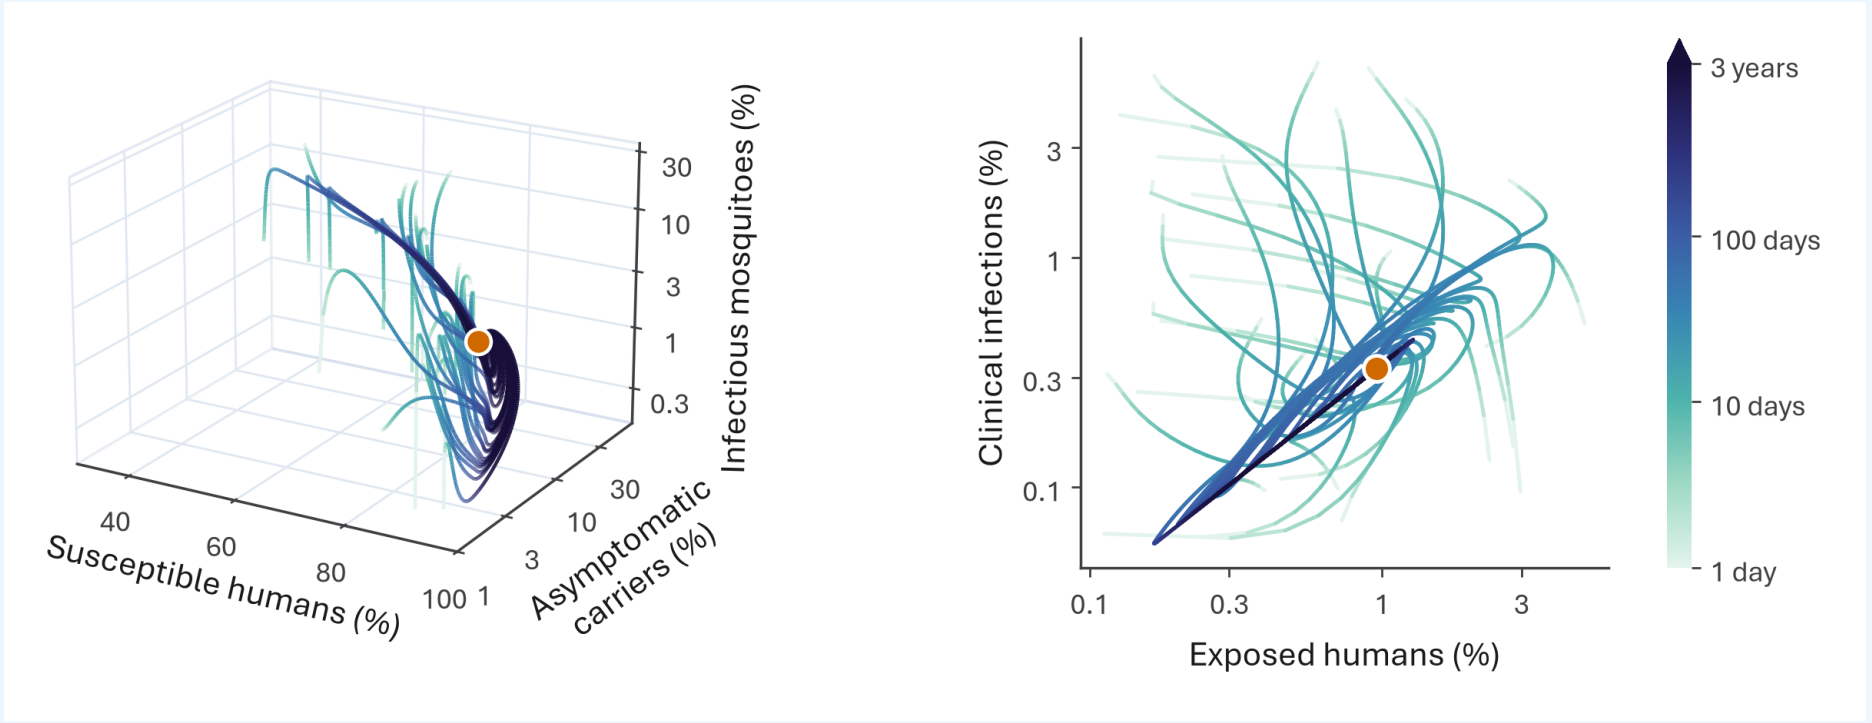
- Asymptomatic reservoir raises R_0 to about **2.40**
- Threshold biting rate falls to $\bar{a} \approx 0.21$
- Closed form agrees with the numerical spectral radius

Bifurcation and Sensitivity



- Forward bifurcation at $R_0 = 1$, so the endemic state is unique
- Biting rate dominates the sensitivity of R_0 (PRCC +0.94)

Phase Portraits

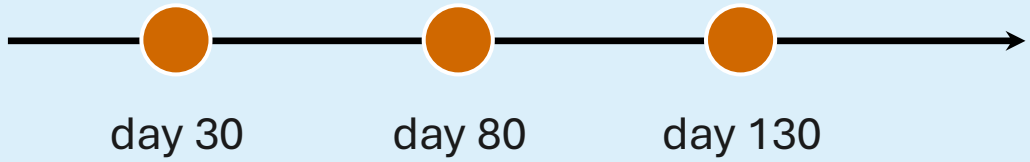


- Random starting states in one calibrated region (Lagos), colored by time
- Every trajectory converges to the same endemic state shown in orange
- Consistent with the forward bifurcation (no bistability)

Benchmark Regions

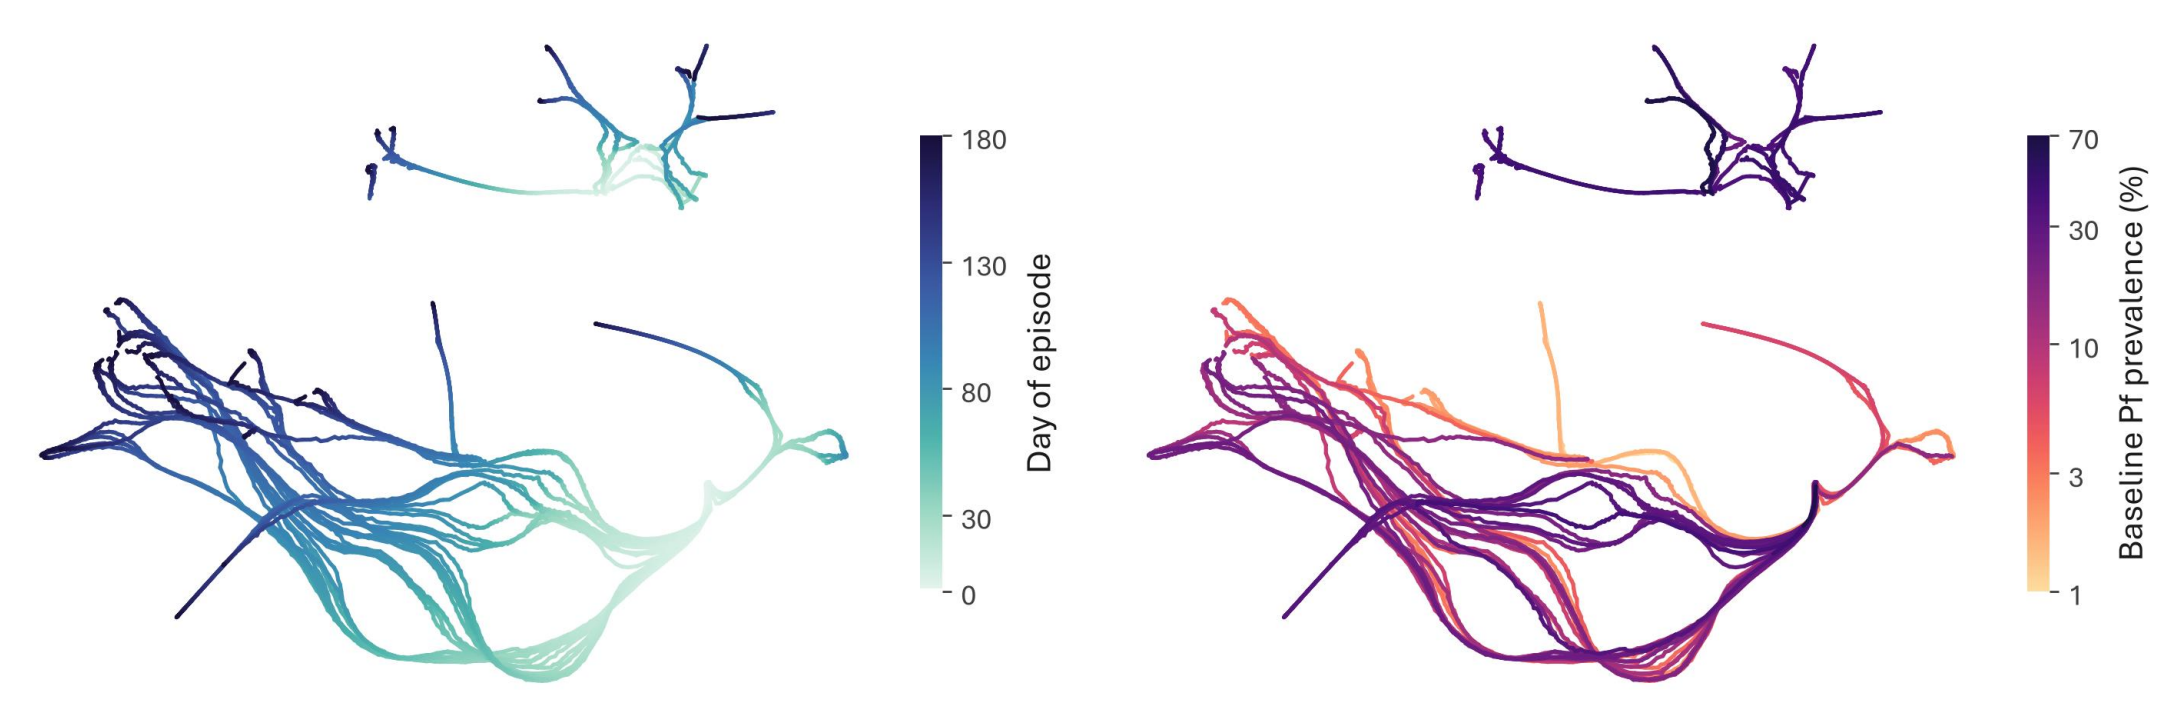
- 50 real admin 1 regions (MAP 2019)
- Pf PR > 1%, population > 50,000
- None are synthetic
- Dar es Salaam 1.8% to Tshuapa 61%

One episode



Three mass campaigns,
one net per four people

Campaign Trajectories



- Each line is one region's daily state under Static, embedded with PaCMAP
- Paths start together and branch through the three campaigns
- Regions respond differently by transmission level, so the split matters

Methods

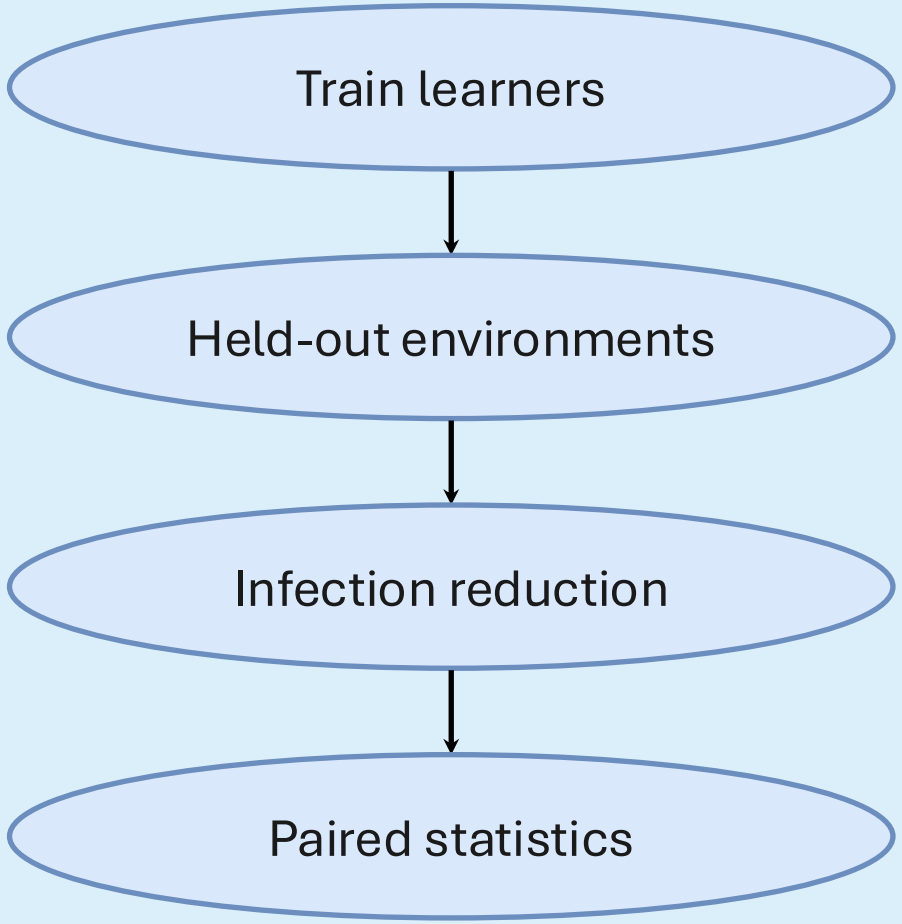
Baselines	Heuristics	Learned	Planners
No ITN	Weight by incidence	PPO	Greedy marginal
Static	Weight by resistance	DQN	MPC with CEM
	Three other rules	Contextual bandit	Oracle upper bound
		QT-Opt with CEM	



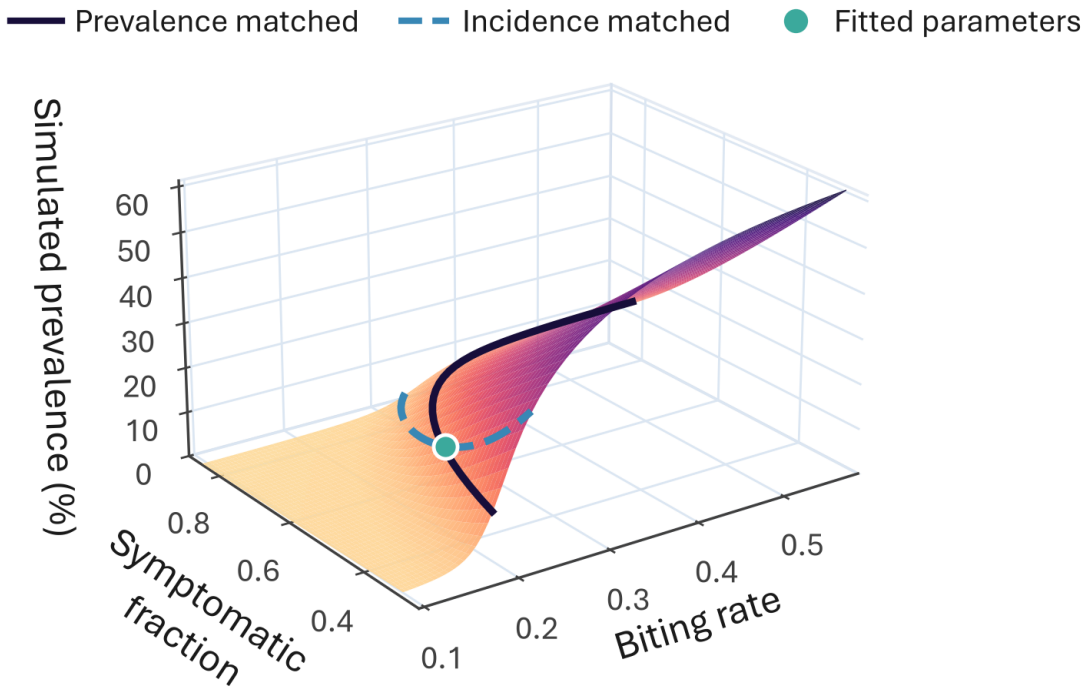
simple complex

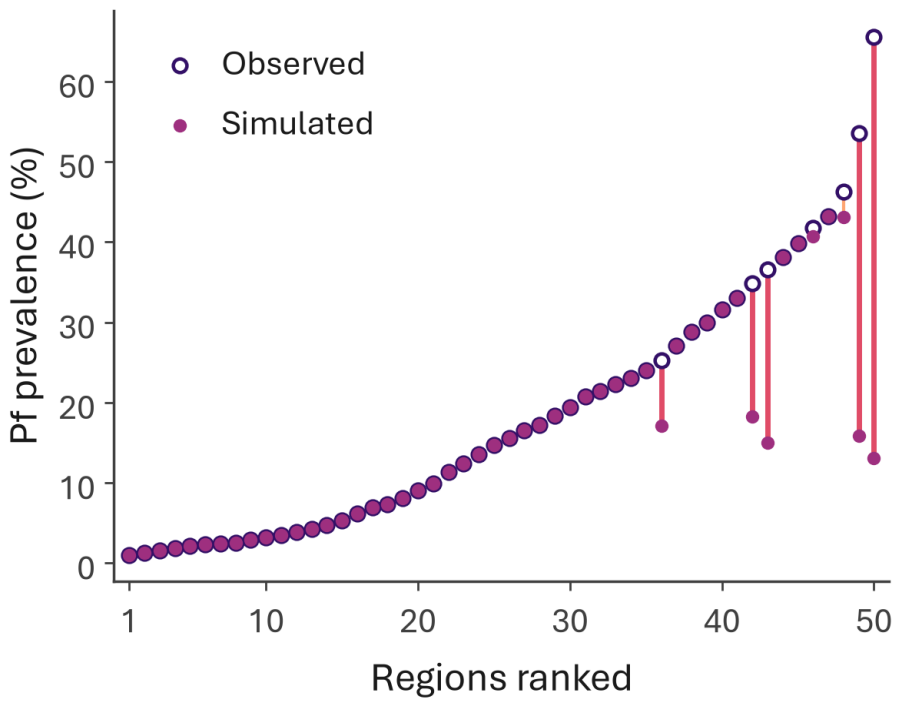
Evaluation Protocol

- Held-out environments only
- Paired comparison with Static
- 95% bootstrap confidence intervals
- Wilcoxon tests, Holm corrected
- Compiled simulator is ~288× faster



Calibration





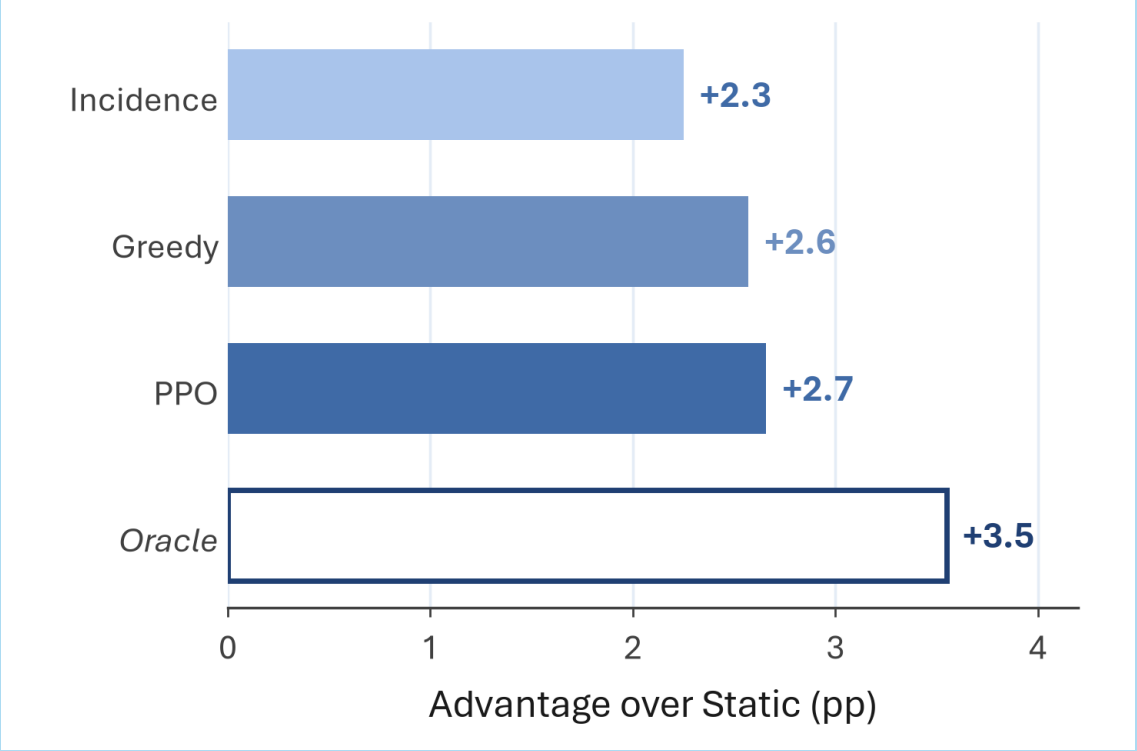
- 45 of 50 regions fit within 3.2 pp prevalence
- Five high-transmission regions are underfit
- Advantage over Static holds in **100%** of bootstrap calibrations

Allocation Comparison

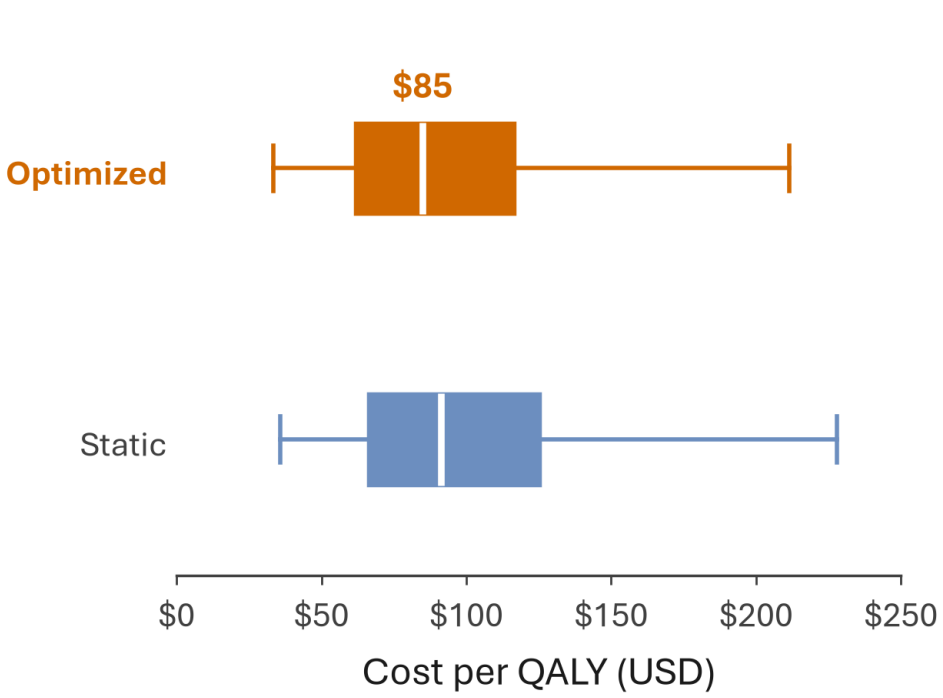
Method	Red. (%)	Adv. vs Static (pp) [95% CI]	Holm p
Static (pop.-proportional)	35.8	—	—
Heuristics			
Incidence-weighted	36.8	+1.0 [+0.9, +1.1]	$< 10^{-14}$
Resistance-aware	37.4	+1.6 [+1.5, +1.6]	$< 10^{-14}$
Learned			
DQN	35.8	0.0 [−0.2, +0.1]	n.s.
Contextual bandit	38.7	+2.9 [+2.5, +3.2]	$< 10^{-13}$
QT-Opt/CEM	38.4	+2.6 [+2.2, +2.9]	$< 10^{-13}$
PPO	40.6	+4.8 [+4.6, +5.0]	$< 10^{-14}$
Planners			
Greedy (marginal)	41.5	+5.7 [+5.5, +6.0]	$< 10^{-14}$
MPC (CEM)	41.3	+5.5 [+5.3, +5.7]	$< 10^{-14}$
Oracle (upper bound)	41.6	+5.8 [+5.6, +6.0]	$< 10^{-14}$

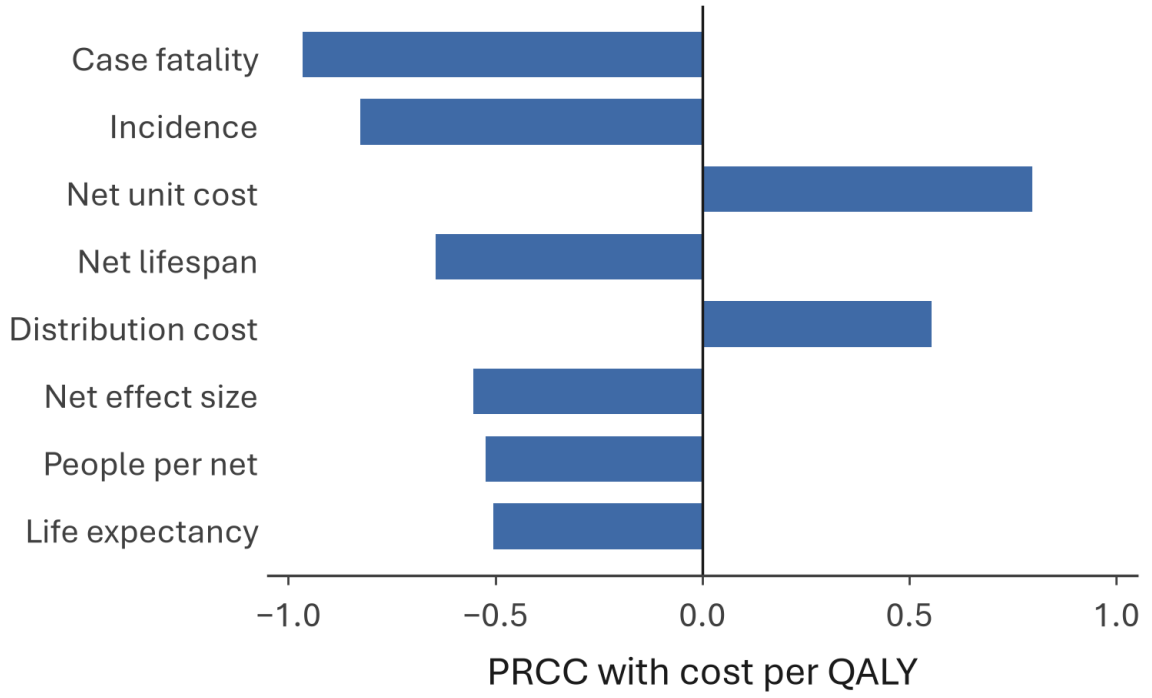
Scaling to 50 Regions

- Static reduces infections by 30.7%
- All methods beat Static in all 40 tests
- Holm-Bonferroni $p < 10^{-10}$
- Greedy leaves 28% of oracle gain



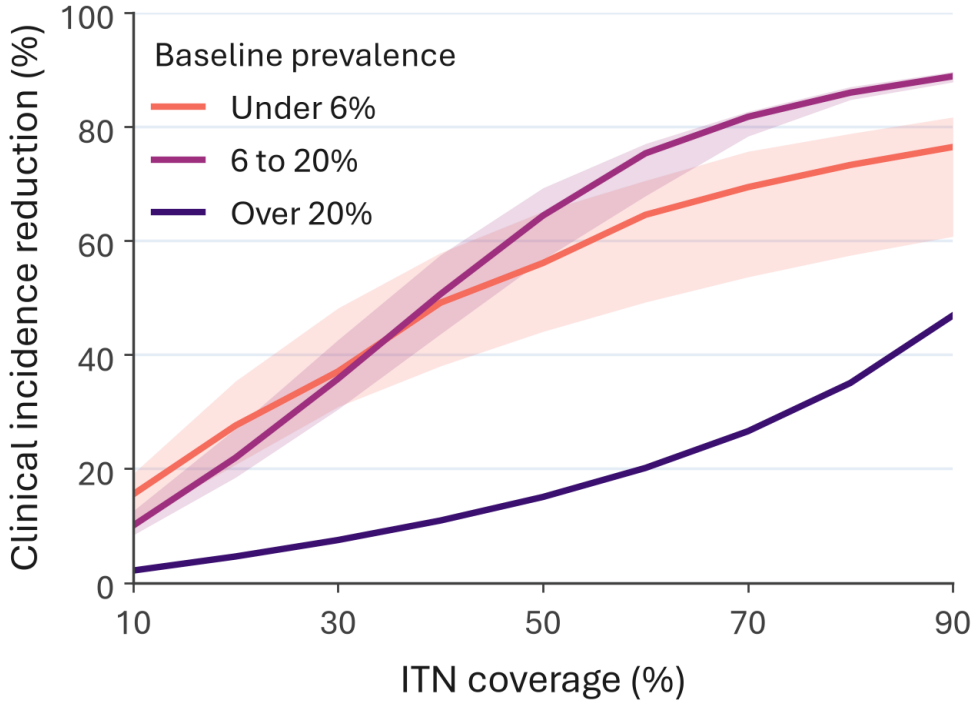
Cost-Effectiveness

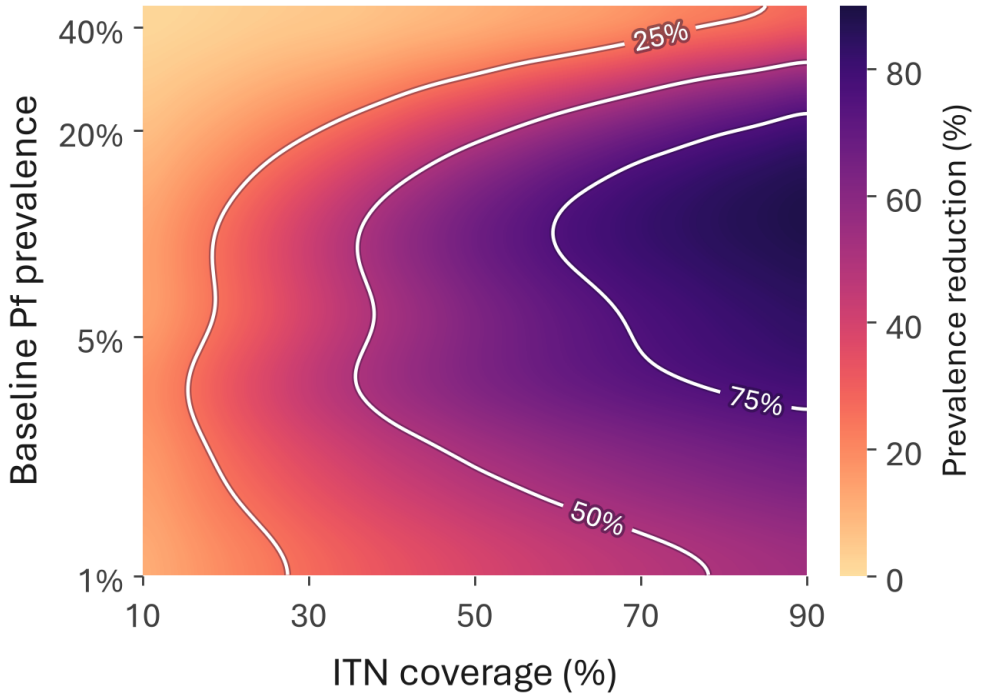




- Optimized allocation costs a median **$85 per QALY** (95% CI $33 to $211)
- The wide range reflects uncertain case fatality, incidence, and net cost
- Those three inputs dominate the sensitivity (|PRCC| ≥ 0.79)

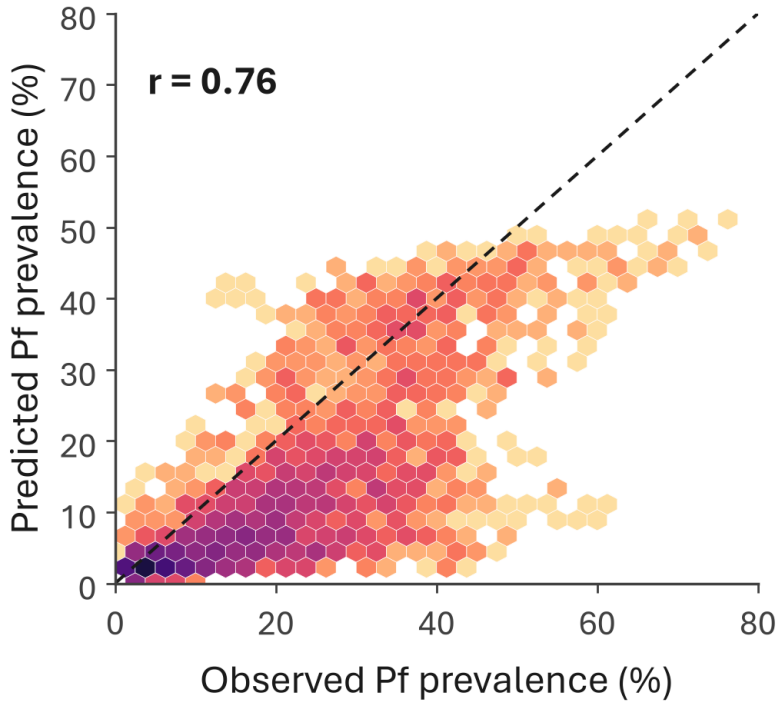
External Validation

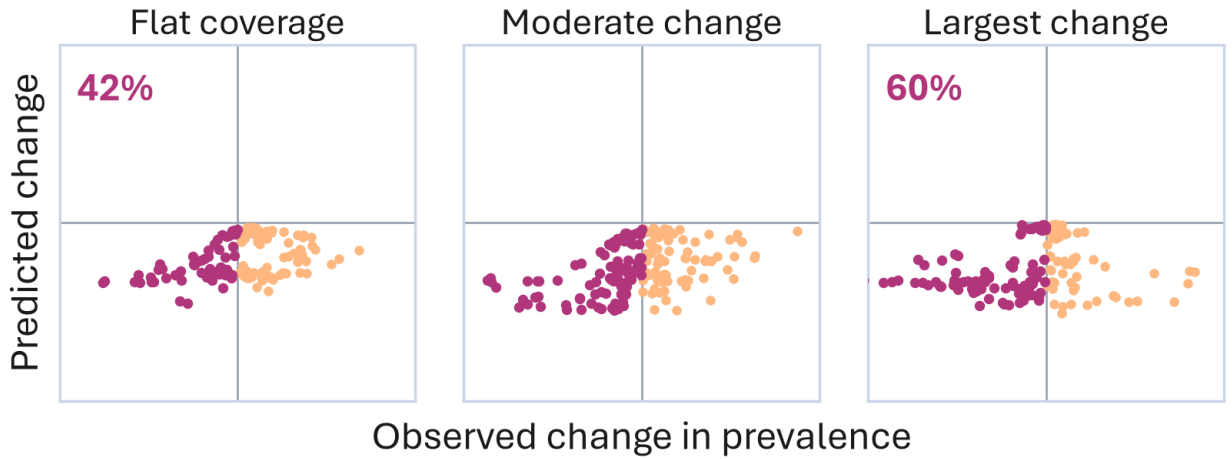




- Never fit to trial outcomes such as Lengeler
- Median 51% fewer cases at 50% coverage, matching trials (~50%)
- Likely overstates ITN effects at high coverage

Out-of-Sample Validation





- Calibrated in 2015, then simulated forward to 2022
- r = **0.76** across 3318 held-out region years
- Magenta dots get the direction right, rising from 42% to 60%

Unknown Transmission

- Greedy suffices if the map is known
- The oracle's edge is knowing the map
- New programs must learn the map
- Each allocation is also an **experiment**

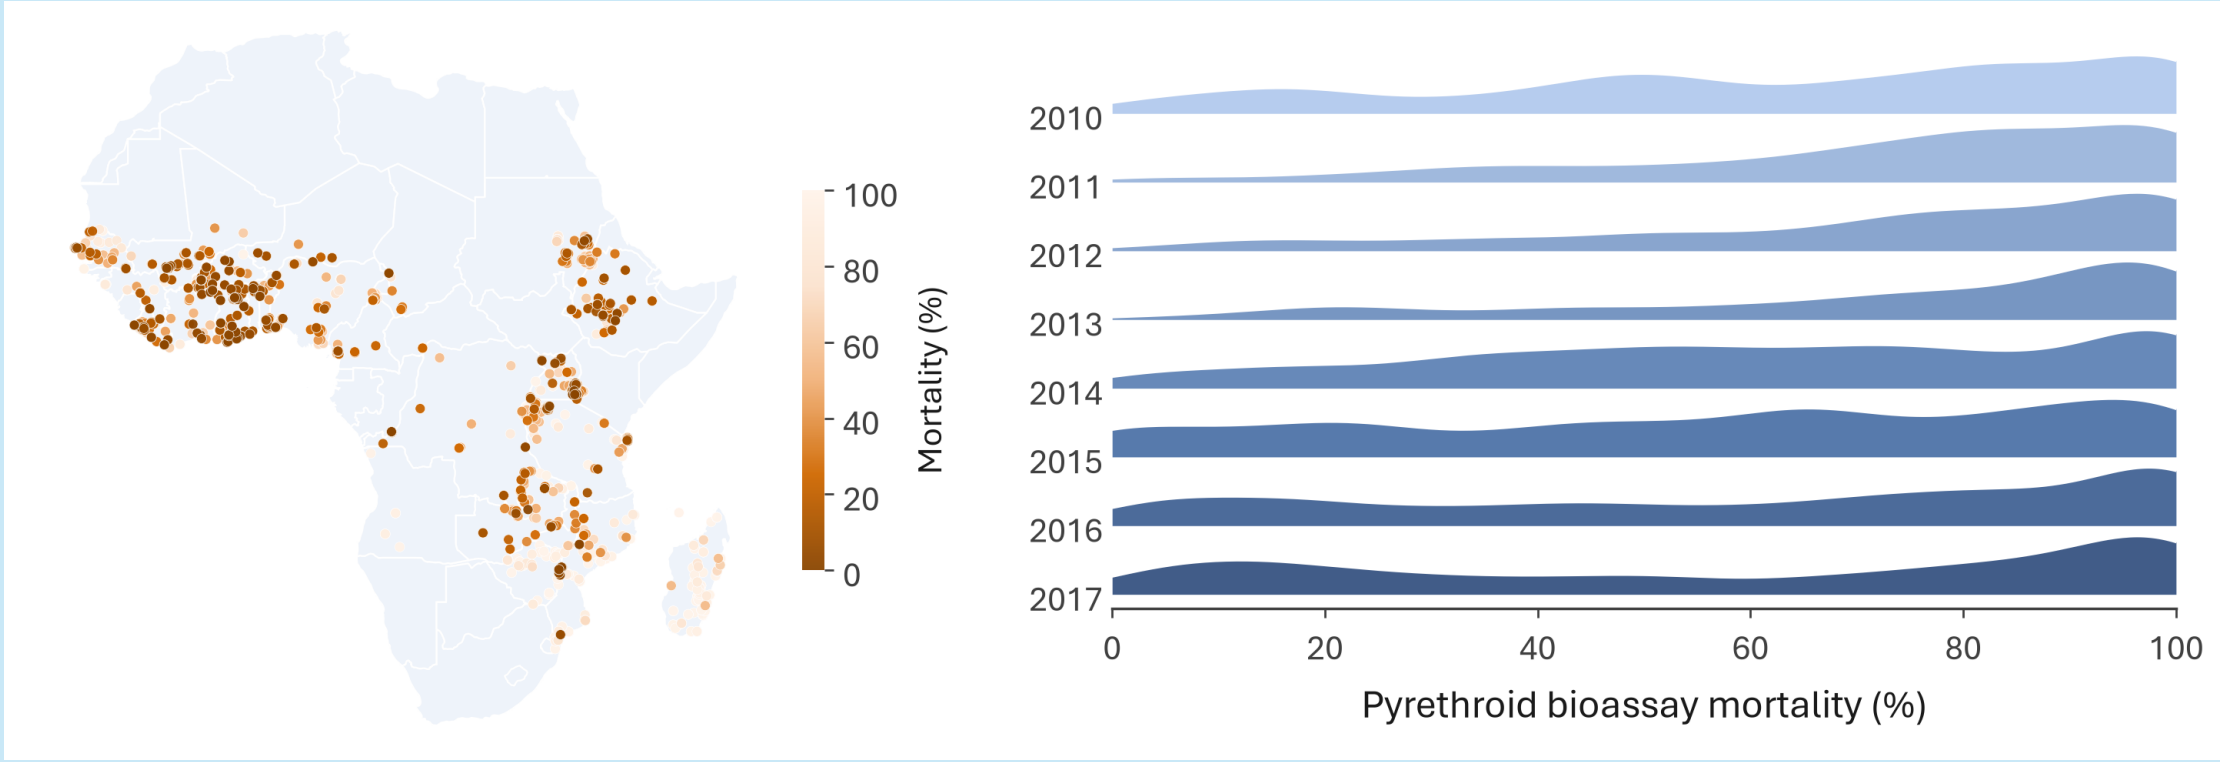


Known



Unknown

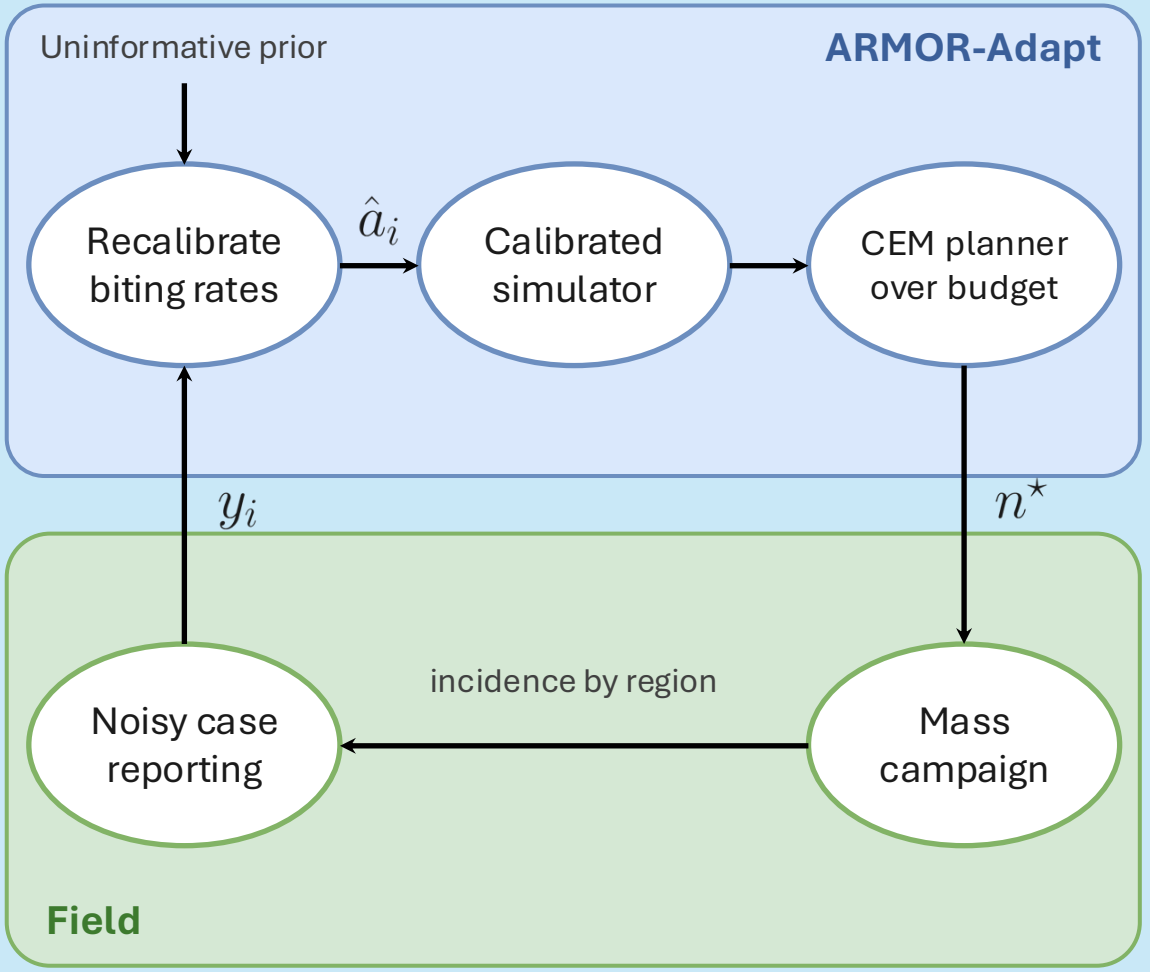
Deployment Dynamics



- Endogenous resistance, human mobility, and net attrition
- Coverage halves in ~2.5 years between campaigns
- Still near-myopic, as greedy stays within **0.7 pp** of the network oracle

Online Learning

- Start from an uninformative prior
- One-shot optimization loses to Static
- ARMOR-Adapt adds **+4.2 pp** over Static
- Ahead of Static in 98% of tests
- Recovers ~77% of the oracle gain



Dual Control



Allocate nets
under a fixed budget



Infections fall
learning payoff

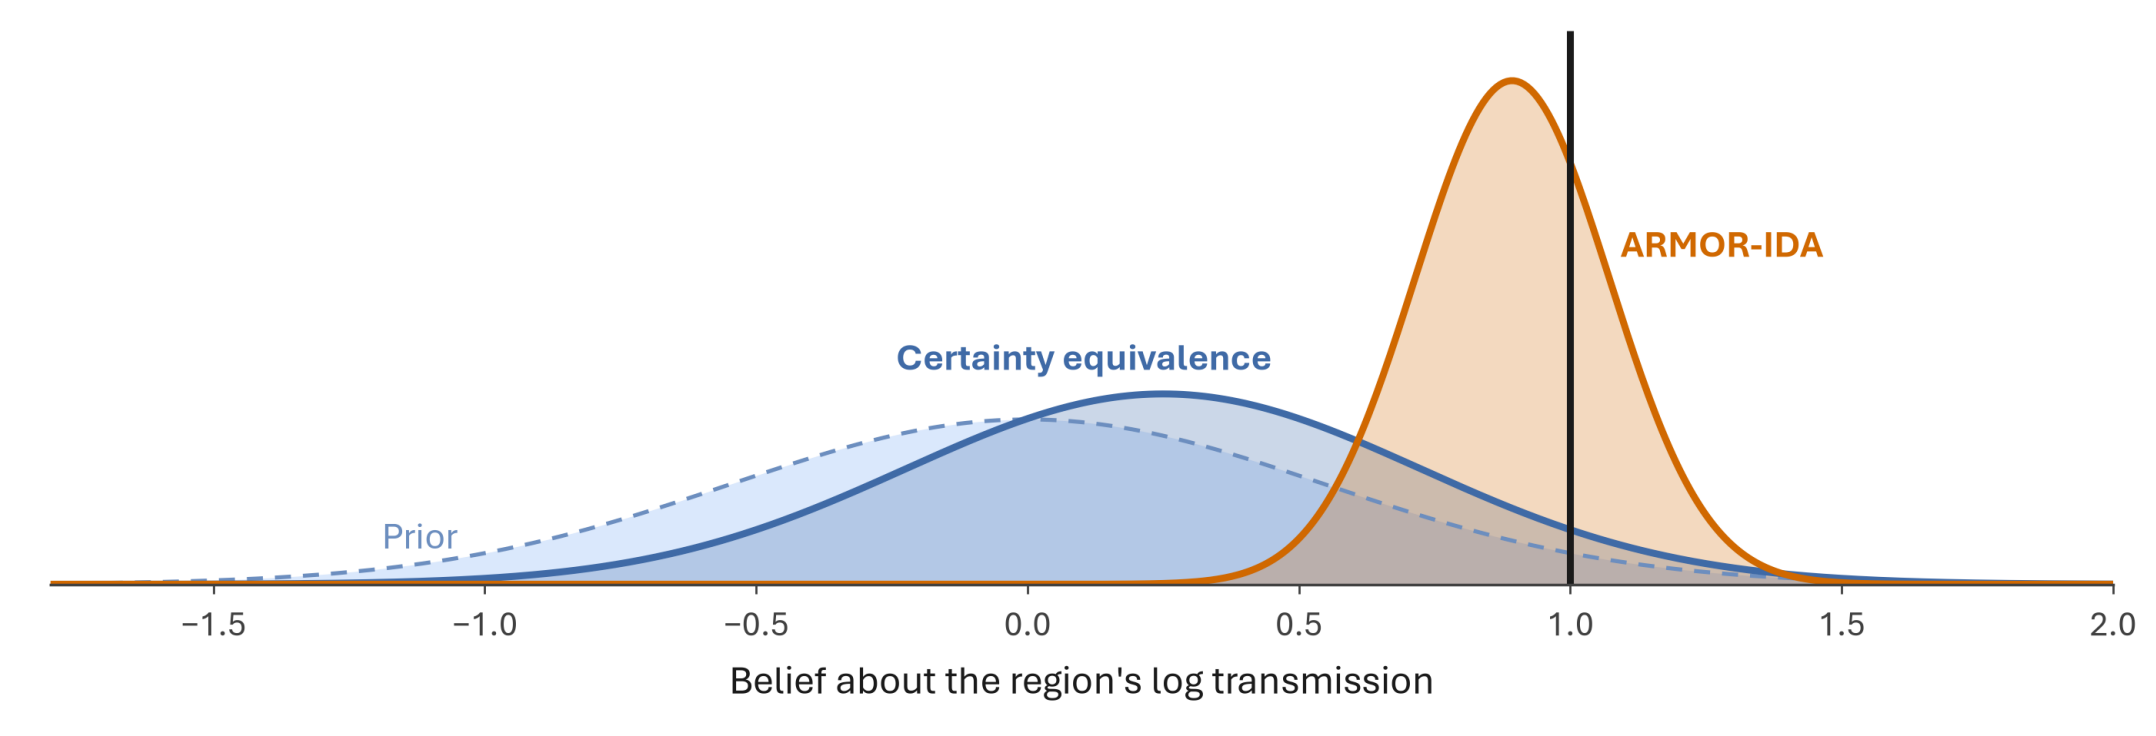
Reporting sharpens
where nets arrive





Beliefs update
the experiment

Certainty Equivalence Fails



- Idealized case from the Proposition, one underestimated region
- Certainty equivalence sends no nets so its belief barely moves
- ARMOR-IDA funds the region and learns its true transmission

Proposition

Assume a region is funded only above an estimate threshold, and that precision rises with allocation from a floor κ_0/σ^2

Certainty Equivalence

If $\kappa_0 = 0$, it converges to a suboptimal allocation with positive probability, because an underfunded region stays unobserved and its error is never corrected

Information-Directed

For large enough λ, it funds any high-burden region that certainty equivalence leaves unfunded, keeping its precision above the floor

ARMOR-IDA

$$n^\star = \arg\min_{n \in \Delta_B} \; \mathbb{E}\big[\text{burden}(n; \hat{\beta})\big] - \lambda \sum_i w_i(n)\big(v_i - v_i^{\text{post}}(n_i)\big)$$

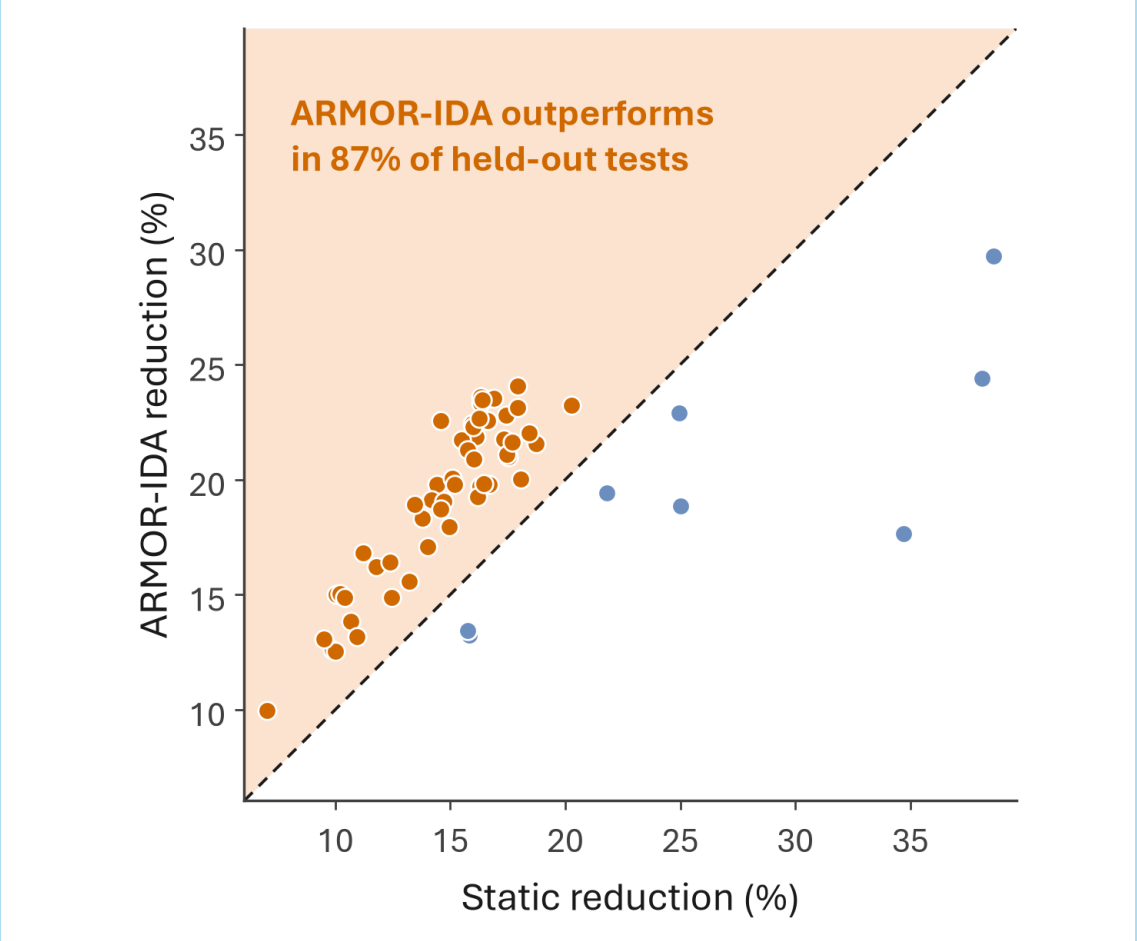
- Expected burden minus a bonus
- Bonus weighted by burden w_i
- $\lambda = 0$ is certainty equivalence
- Explores where decisions change

Algorithm 1 ARMOR-IDA: information-directed adaptive allocation

prior $\mu_i \leftarrow \bar{\mu}$, $v_i \leftarrow v_0$ for all regions i
for campaign $j = 1, \ldots, J$ **do**
  set $\hat{\beta}_i \leftarrow e^{\mu_i}$
  $n^\star \leftarrow$ minimize Eq. (2) over Δ_B (CEM)
  deploy $n^\star$; observe noisy incidence y_i, precision $\propto \kappa_0 + \kappa_1 C_i$
  fit $\tilde{\beta}_i$ to y_i; Bayes-update (μ_i, v_i) with $\log \tilde{\beta}_i$
end for

Noisy Surveillance

- 60 held-out tests, 8 campaigns each
- ARMOR-IDA reaches **19.4%** reduction
- Certainty equivalence reaches 17.7%
- Adds +3.0 pp, ahead of Static in 87%
- Best worst case of any adaptive rule



Results and Discussion

- The exploration rule drives the gain
- $\lambda = 0$ (certainty equivalence) is worst
- Gains rise steadily over $\lambda \in [2, 32]$
- We report $\lambda = 8$, below the best

Method	Mean red. (%)	CVaR (%)
Static (population-proportional)	16.4	10.3
Greedy on observed incidence	17.3	11.6
Certainty-equiv. MPC (=ARMOR-Adapt)	17.7	11.2
Thompson sampling	17.7	11.7
UCB / optimistic	15.6	8.2
ARMOR-IDA (information-directed)	**19.4**	**13.5**
Oracle (known transmission)	22.3	15.1

Deployment Guidance



Greedy or heuristic allocation

Transmission map known



Does allocation impact surveillance?

no

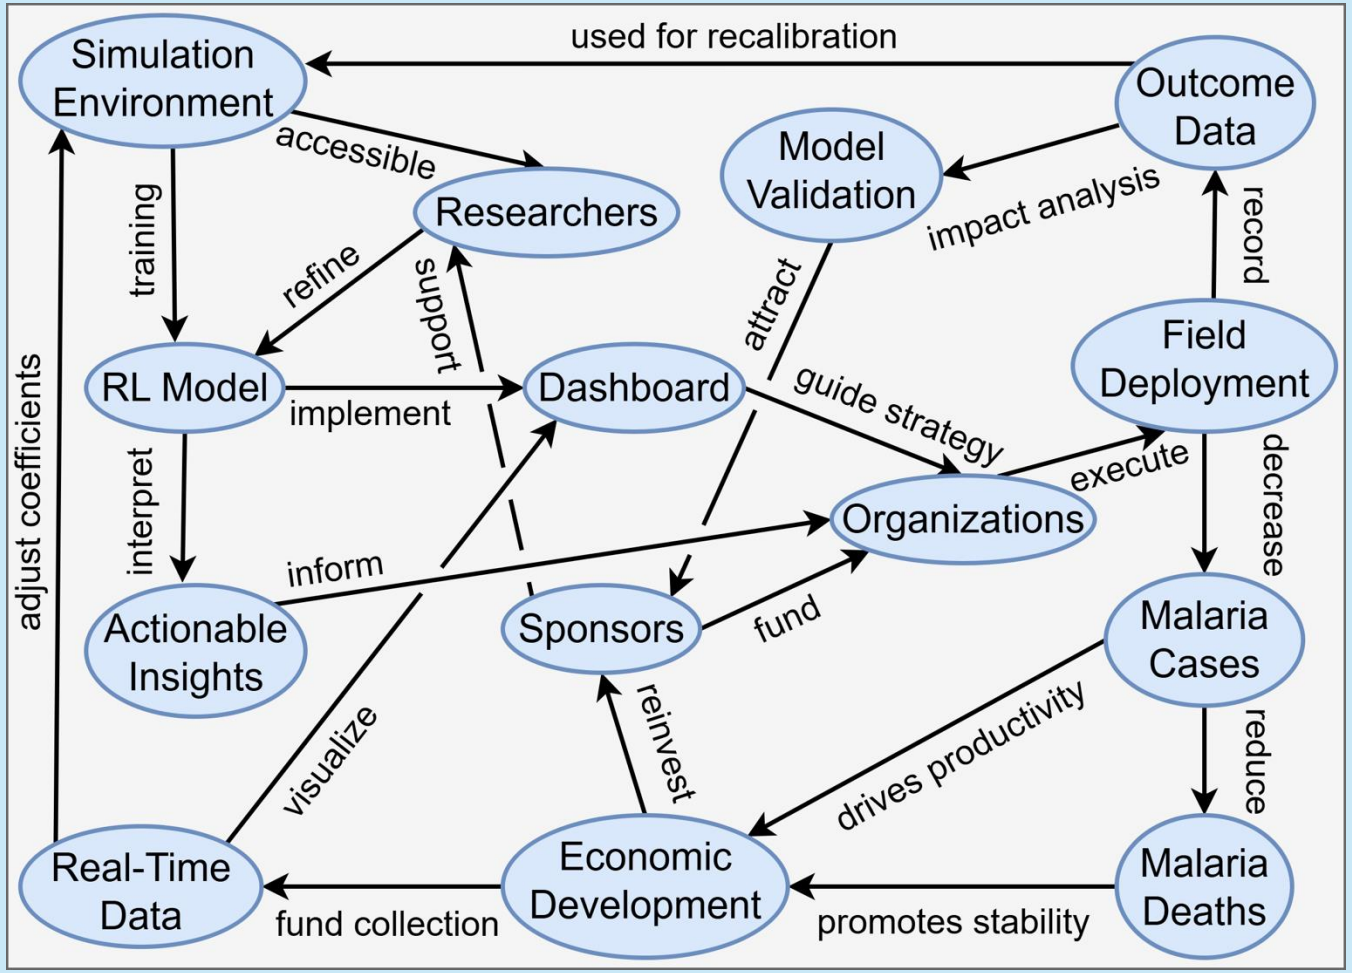
yes

ARMOR-Adapt

ARMOR-IDA

Transmission map unknown

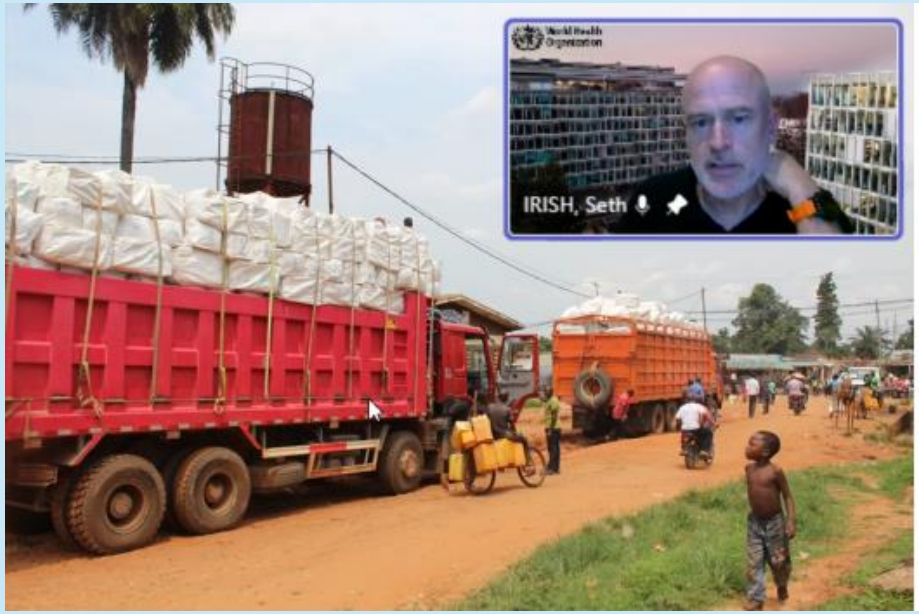
Theory of Change



The Global Fund, PMI, AMF, WHO, and Malaria Consortium

Limitations and Uncertainties

- Validated on prevalence, not on controlled ITN campaigns
- Model likely overstates ITN effects at high coverage
- Assumes regional incidence is reported after each campaign





Conclusion

- **Calibrated** ITN allocation benchmark on real subnational data
- Optimized allocation beats population-proportional by up to **~6 pp**
- Known mapping problem is near-myopic, so greedy nearly matches the oracle
- Leverage appears when transmission must be learned online
- **ARMOR-IDA** beats certainty equivalence, Thompson sampling, and UCB

Funding





Special thanks to Dr. Seth Irish and Dr. Clara Champagne
for sharing their valuable expertise during interviews.

Acknowledgements

We would like to thank Prof. Gil Alterovitz and Dr. Ning Xie for their support throughout this project along with the entire MIT PRIMES community.