INTERPRETABLE, ROBUST, RISK-AWARE FRAMEWORK FOR LEARNING SEPSIS TREATMENT STRATEGIES

ANGELA LIN, MIT OPERATIONS RESEARCH CENTER

JOINT WORK WITH:

DR. LIEN LE, TAMPA GENERAL HOSPITAL

DR. DOUGLAS MCCONNELL, TAMPA GENERAL HOSPITAL

PROF. GEORGIA PERAKIS, MIT SLOAN SCHOOL OF MANAGEMENT

PROF. DESSISLAVA PACHAMANOVA, BABSON COLLEGE

SEPSIS IS THE THIRD LEADING CAUSE OF DEATH IN THE U.S.

Sepsis is a **dysregulated** response to **infection.**

SEPSIS IS THE THIRD LEADING CAUSE OF DEATH IN THE U.S.

Sepsis is a **dysregulated** response to **infection.**

Symptoms can include high heart rate, **low blood pressure**, fever, pain and discomfort, shortness of breath, low urine output.

SEPSIS IS THE THIRD LEADING CAUSE OF DEATH IN THE U.S.

Sepsis is a **dysregulated** response to **infection.**

Symptoms can include high heart rate, **low blood pressure**, fever, pain and discomfort, shortness of breath, low urine output.

Treatments include **antibiotics**, **IV fluids**, vasopressors, oxygen. Severe cases can require **ventilator** (lung failure) or **dialysis** (kidney failure).

SEPSIS IS THE THIRD LEADING CAUSE OF DEATH IN THE U.S.

Sepsis is a **dysregulated** response to **infection.**

Symptoms can include high heart rate, **low blood pressure**, fever, pain and discomfort, shortness of breath, low urine output.

Treatments include **antibiotics**, **IV fluids**, vasopressors, oxygen. Severe cases can require **ventilator** (lung failure) or **dialysis** (kidney failure).

Can lead to organ failure, permanent disability

Responsible for **250K deaths** per year in US and **largest portion of hospitalization costs**

TREATING SEPSIS IS A DIFFICULT DECISION-MAKING PROBLEM

- Sepsis is a broad, **heterogeneous** condition
- Each hour treatment is delayed **reduces survival chances** by 7.6%[1]

TREATING SEPSIS IS A DIFFICULT DECISION-MAKING PROBLEM

- Sepsis is a broad, **heterogeneous** condition
- Each hour treatment is delayed **reduces survival chances** by 7.6%[1]

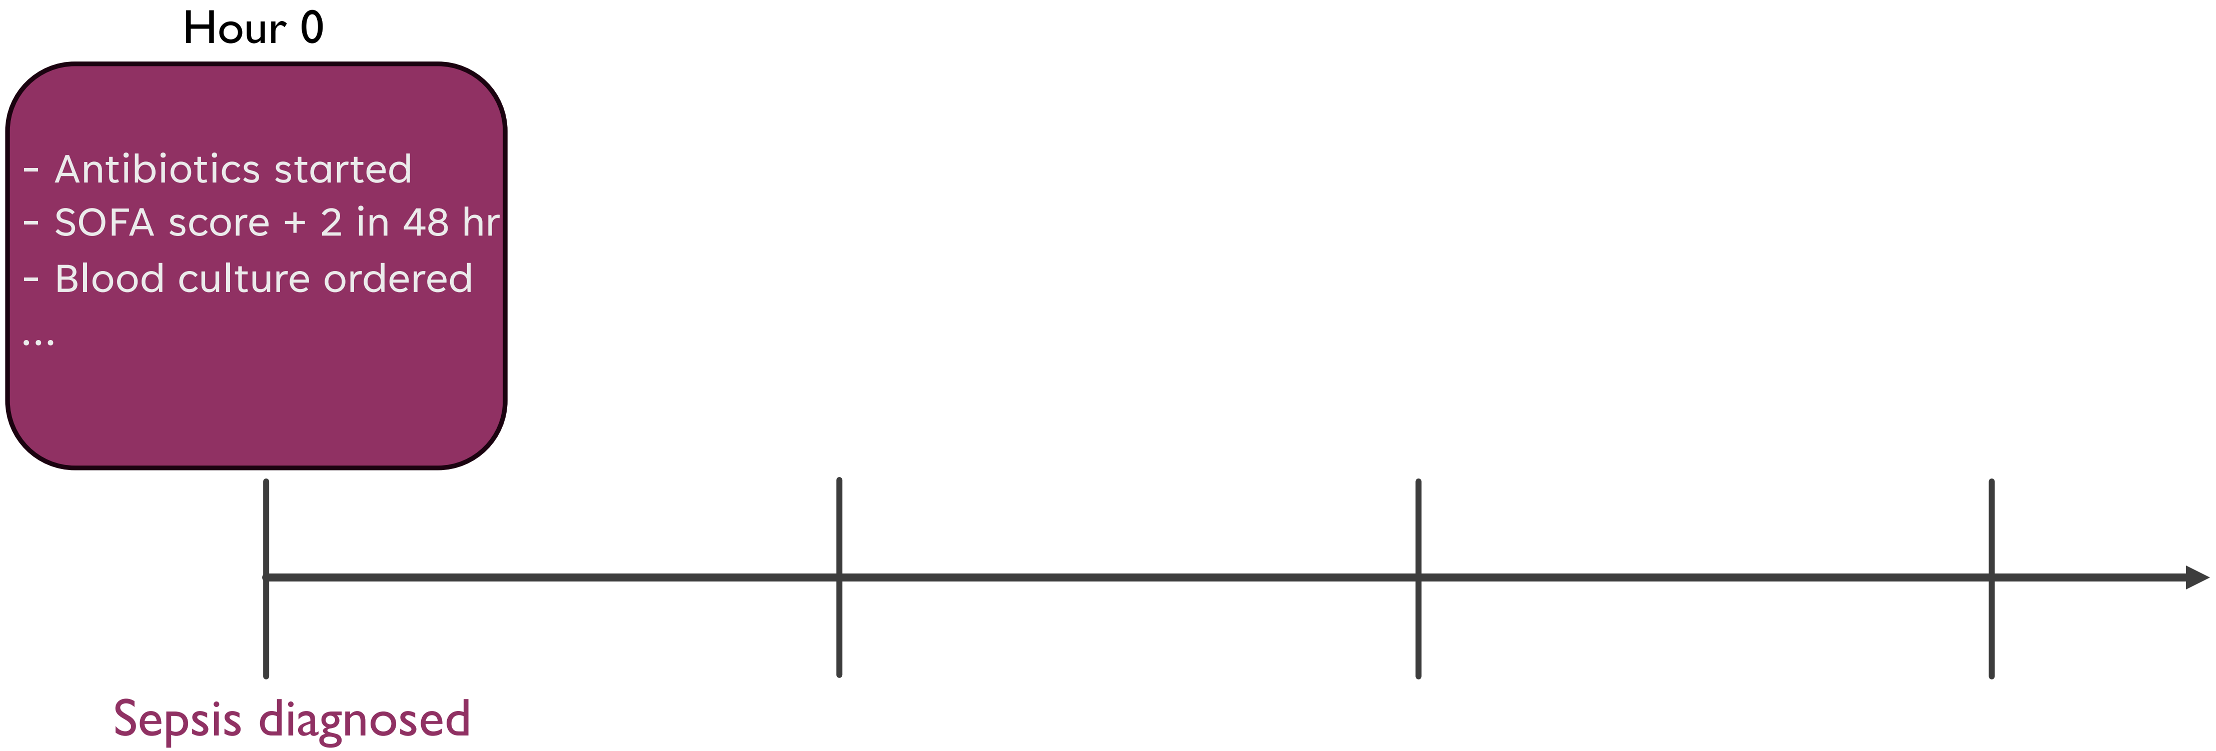

[1]Arulappen et al. (2022)

TREATING SEPSIS IS A DIFFICULT DECISION-MAKING PROBLEM

- Sepsis is a broad, **heterogeneous** condition
- Each hour treatment is delayed **reduces survival chances** by 7.6%[1]

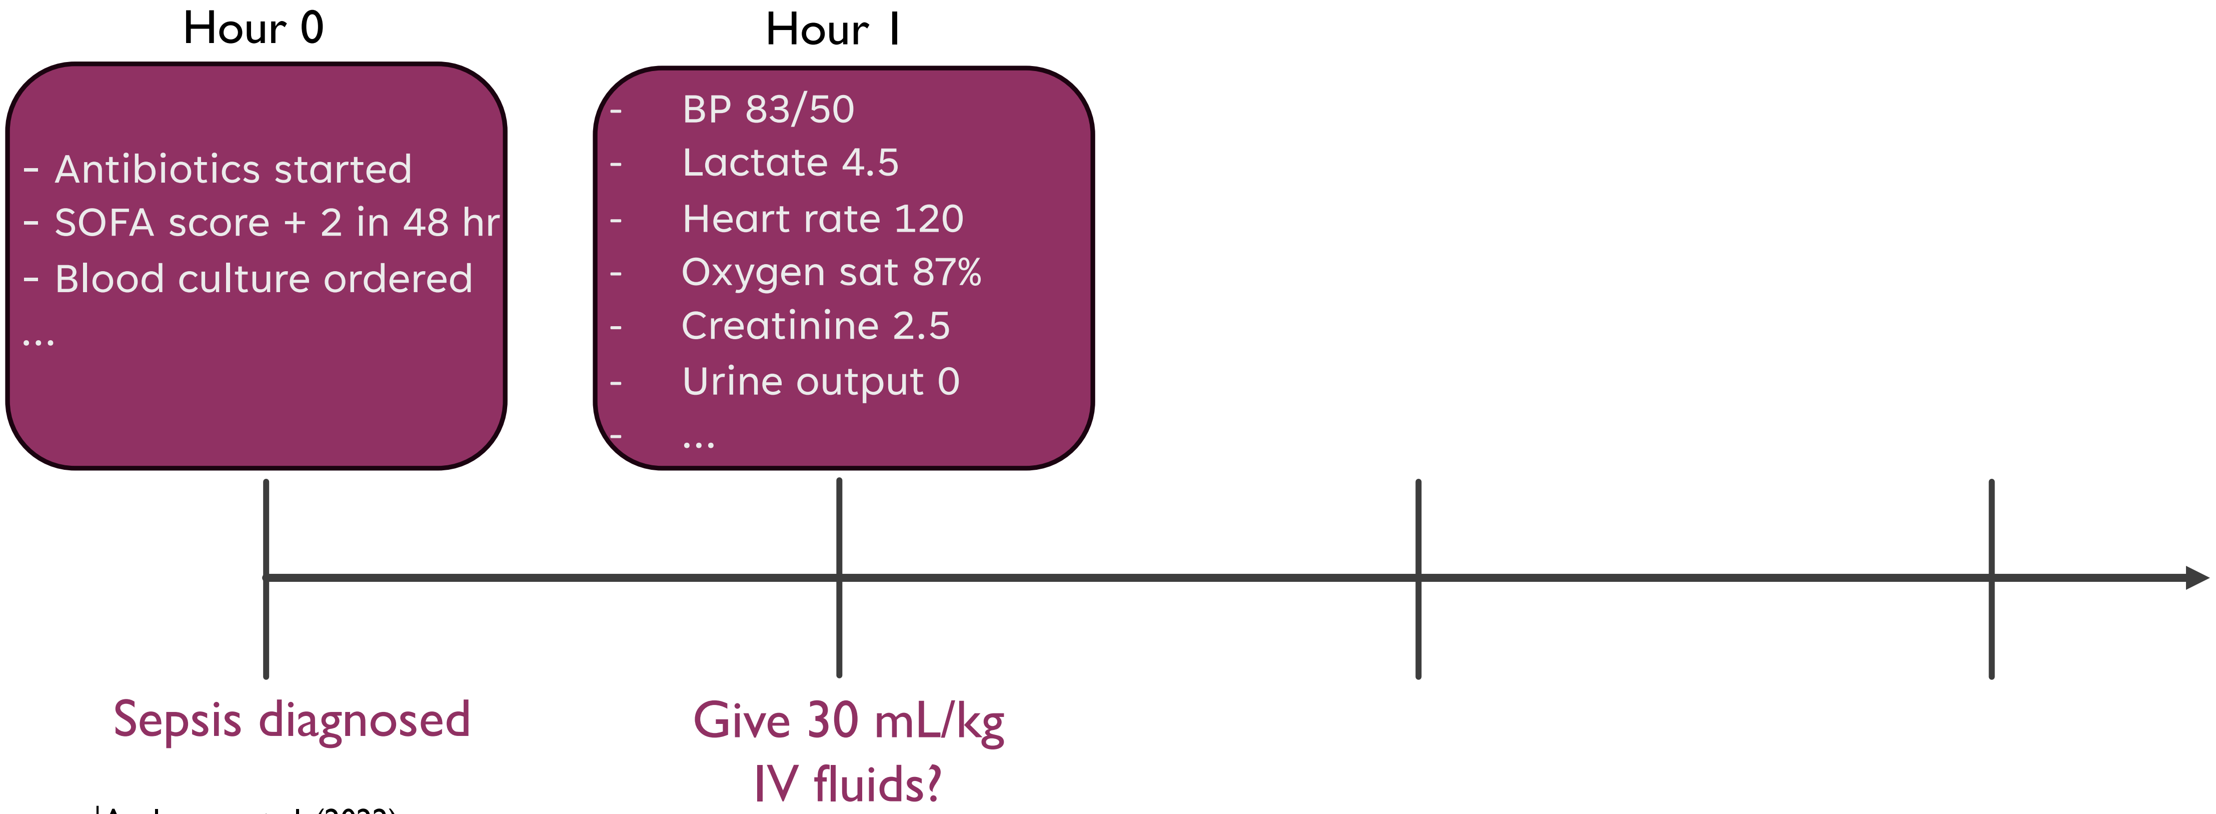

[1]Arulappen et al. (2022)

TREATING SEPSIS IS A DIFFICULT DECISION-MAKING PROBLEM

- Sepsis is a broad, **heterogeneous** condition
- Each hour treatment is delayed **reduces survival chances** by 7.6%[1]

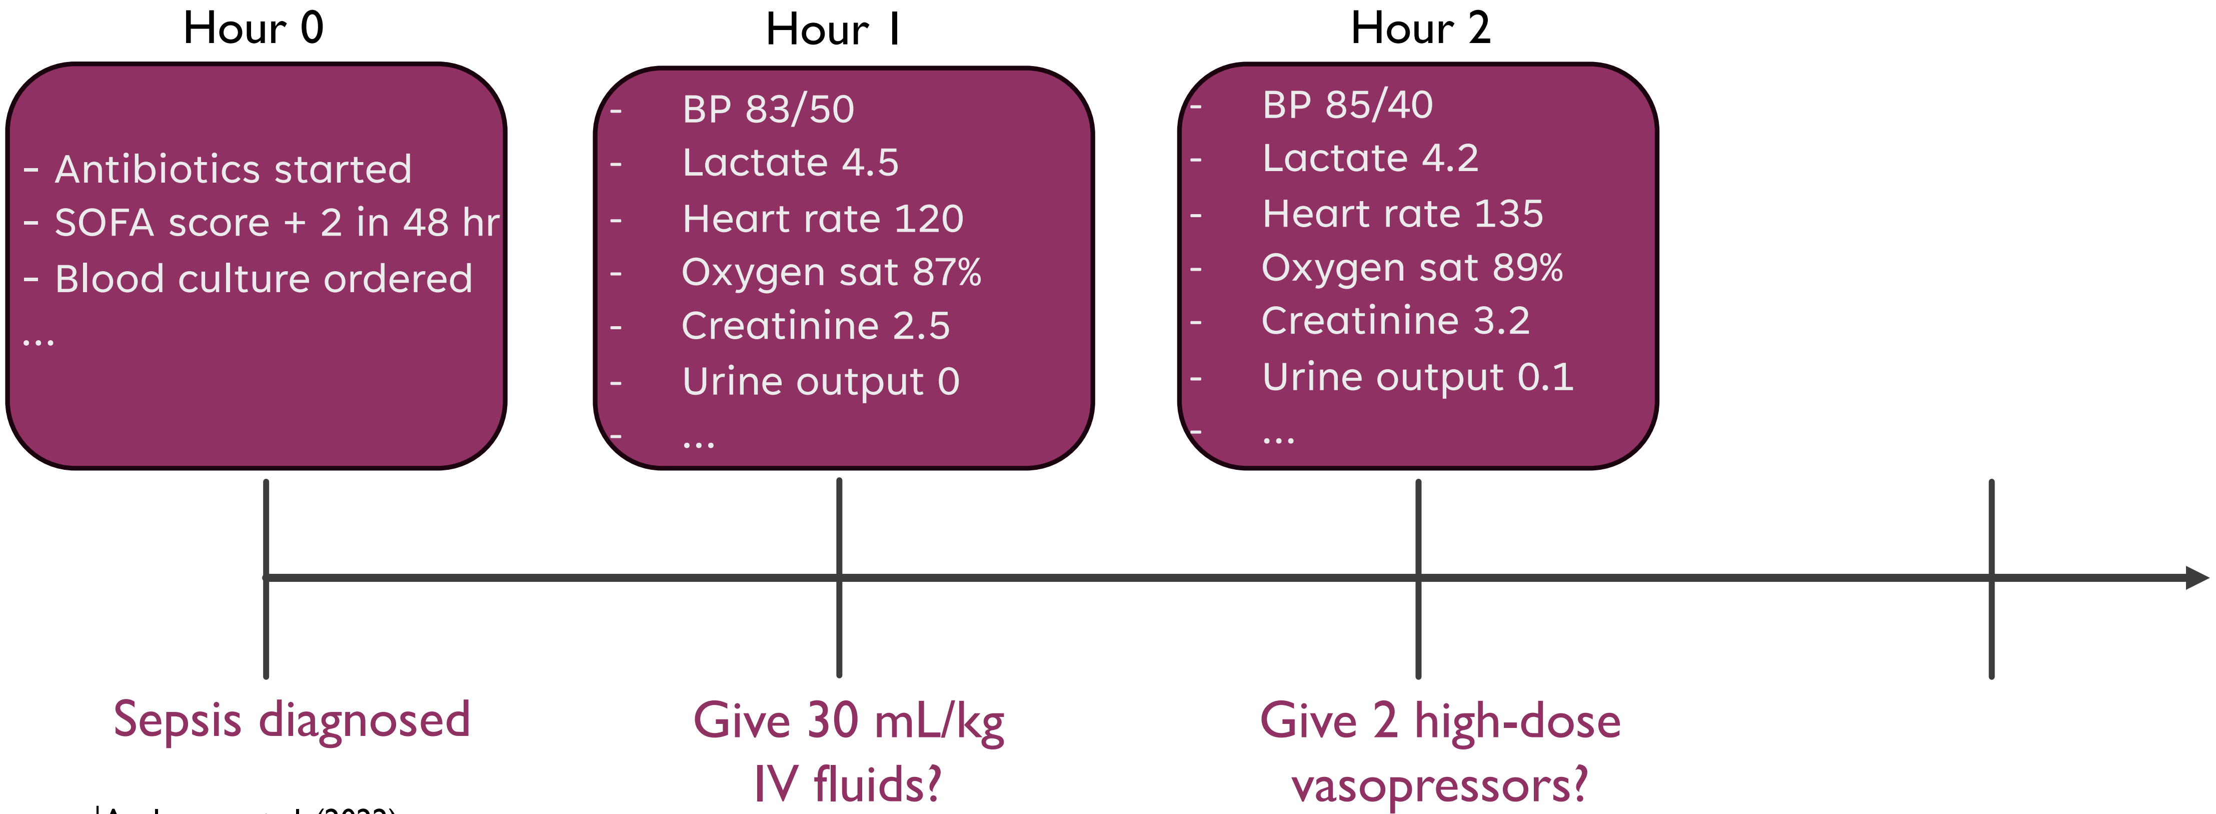

[1]Arulappen et al. (2022)

TREATING SEPSIS IS A DIFFICULT DECISION-MAKING PROBLEM

- Sepsis is a broad, **heterogeneous** condition
- Each hour treatment is delayed **reduces survival chances** by 7.6%[1]

Hour 0

- Antibiotics started
- SOFA score + 2 in 48 hr
- Blood culture ordered

...

Sepsis diagnosed

Hour 1

- BP 83/50
- Lactate 4.5
- Heart rate 120
- Oxygen sat 87%
- Creatinine 2.5
- Urine output 0
- ...

Give 30 mL/kg IV fluids?

Hour 2

- BP 85/40
- Lactate 4.2
- Heart rate 135
- Oxygen sat 89%
- Creatinine 3.2
- Urine output 0.1
- ...

Give 2 high-dose vasopressors?

Hour 3

- BP 75/39
- Lactate 4.7
- Heart rate 150
- Oxygen sat 88%
- Creatinine 3.5
- Urine output 0.1
- ...

Start low-flow oxygen cannula?

[1] Arulappen et al. (2022)

TREATING SEPSIS IS A DIFFICULT DECISION-MAKING PROBLEM

- Sepsis is a broad, **heterogeneous** condition
- Each hour treatment is delayed **reduces survival chances** by 7.6%[1]

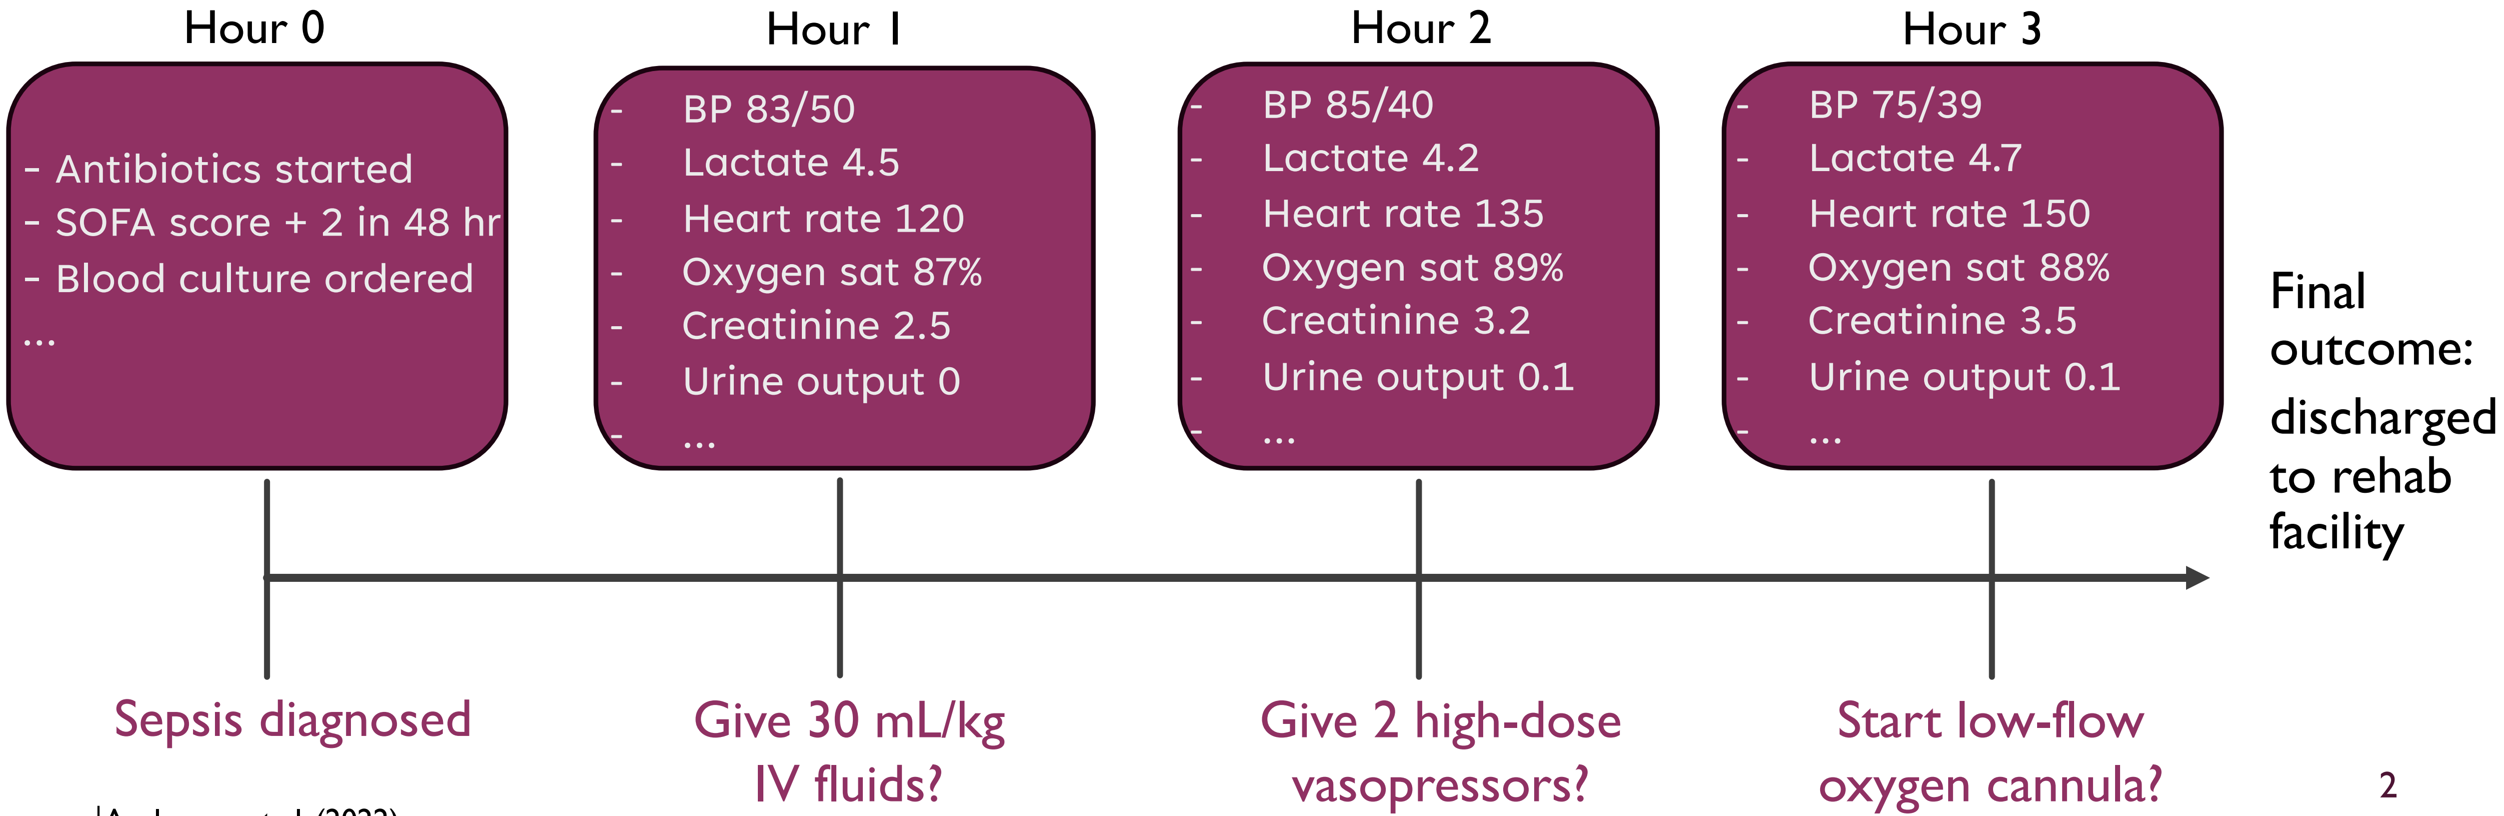

[1] Arulappen et al. (2022)

RESEARCH QUESTIONS AND CHALLENGES

Questions

- What **volume of IV fluids** and vasopressors to give to patients at each stage based on **which characteristics**?
- How to **allocate scarce resources** (ventilators and dialysis machines) for sepsis patients?

Challenges

- No data-driven, **patient-specific recommendations** for sepsis treatment
- Fixed batch of **observational data** (experimentation not feasible)
- Recommendations need to be **safe** and **explainable** to be implementable

DATA: OBSERVED PATIENT TRAJECTORIES

Patient ID	Demographics	Medical History	Hour	Vital Signs	Lab Values	SOFA Score	Treatments
	age, sex, race, …	Comorbidities, pre-existing conditions, admission reason		heart rate, blood pressure, respiratory rate, urine output…	lactate, white blood cell count, creatinine…	Sequential organ failure assessment (function of six organ systems)	antibiotics, IV fluids, vasopressors, oxygen, dialysis

- 21K sepsis patients + final outcomes (end status at hospital)
- 1.7M rows of data
- Imputed missing values, engineered additional features

Goals: Learn states and transitions from data; explain model to clinicians

MDP MODEL OF SEPSIS TREATMENT

State Space X: vital signs, hour, lab values, engineered features, past treatments

Goals: Learn states and transitions from data; explain model to clinicians

MDP MODEL OF SEPSIS TREATMENT

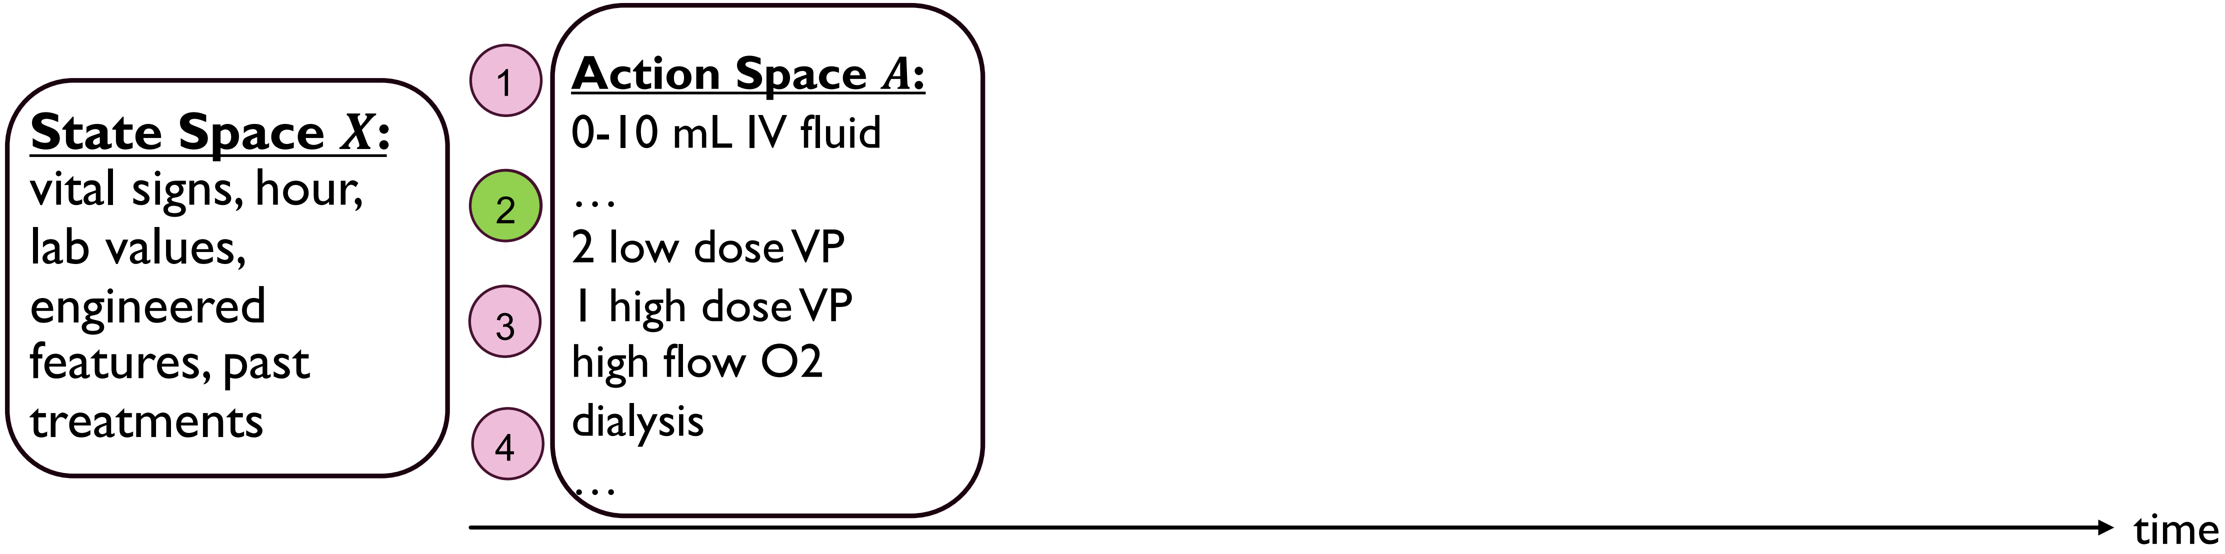

Goals: Learn states and transitions from data; explain model to clinicians

MDP MODEL OF SEPSIS TREATMENT

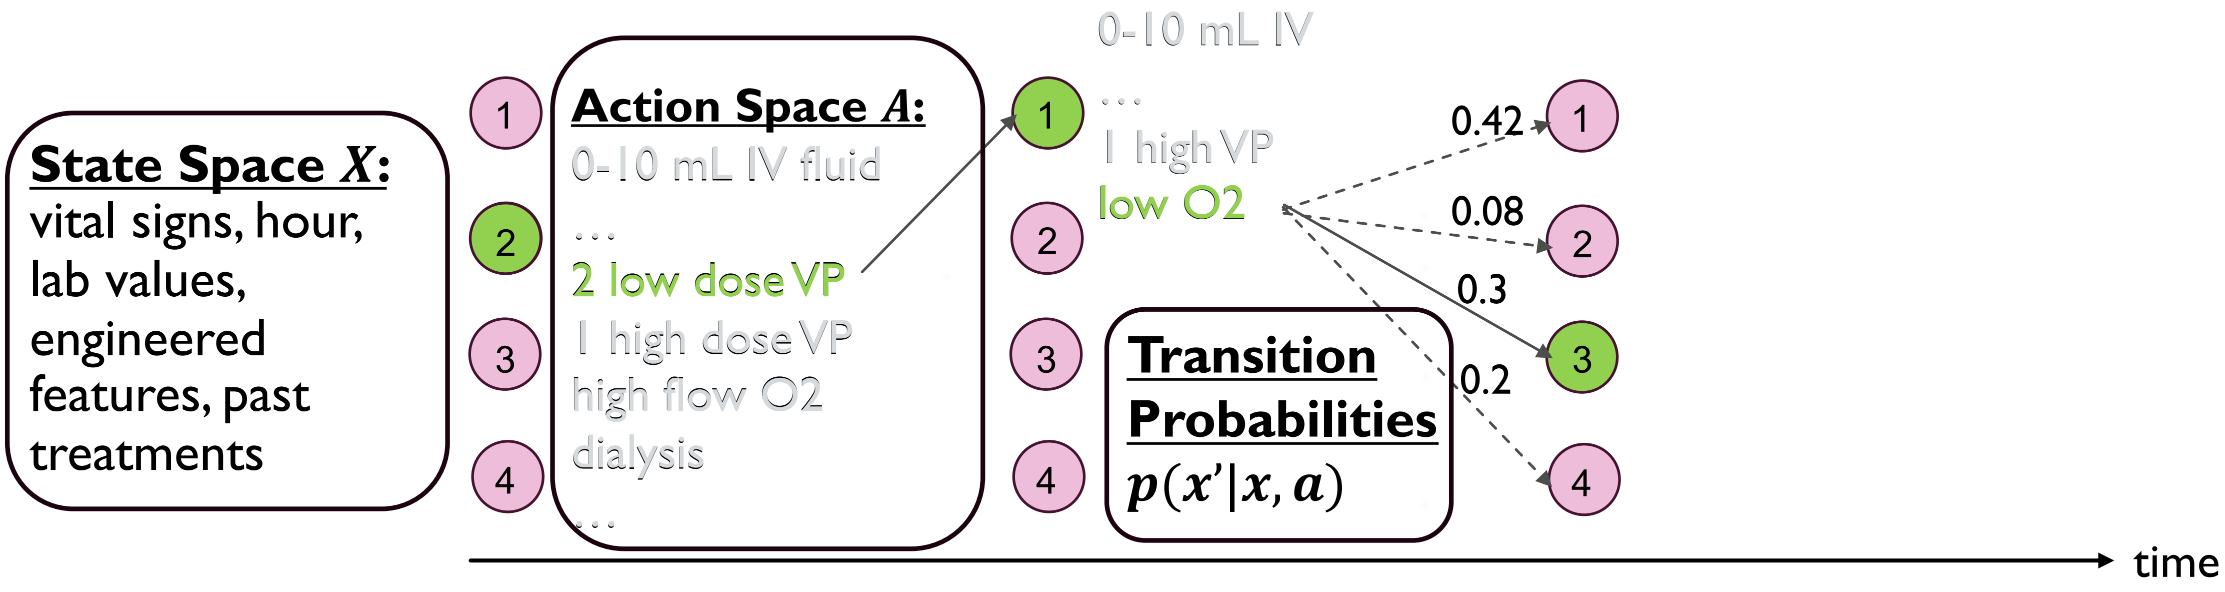

Goals: Learn states and transitions from data; explain model to clinicians

MDP MODEL OF SEPSIS TREATMENT

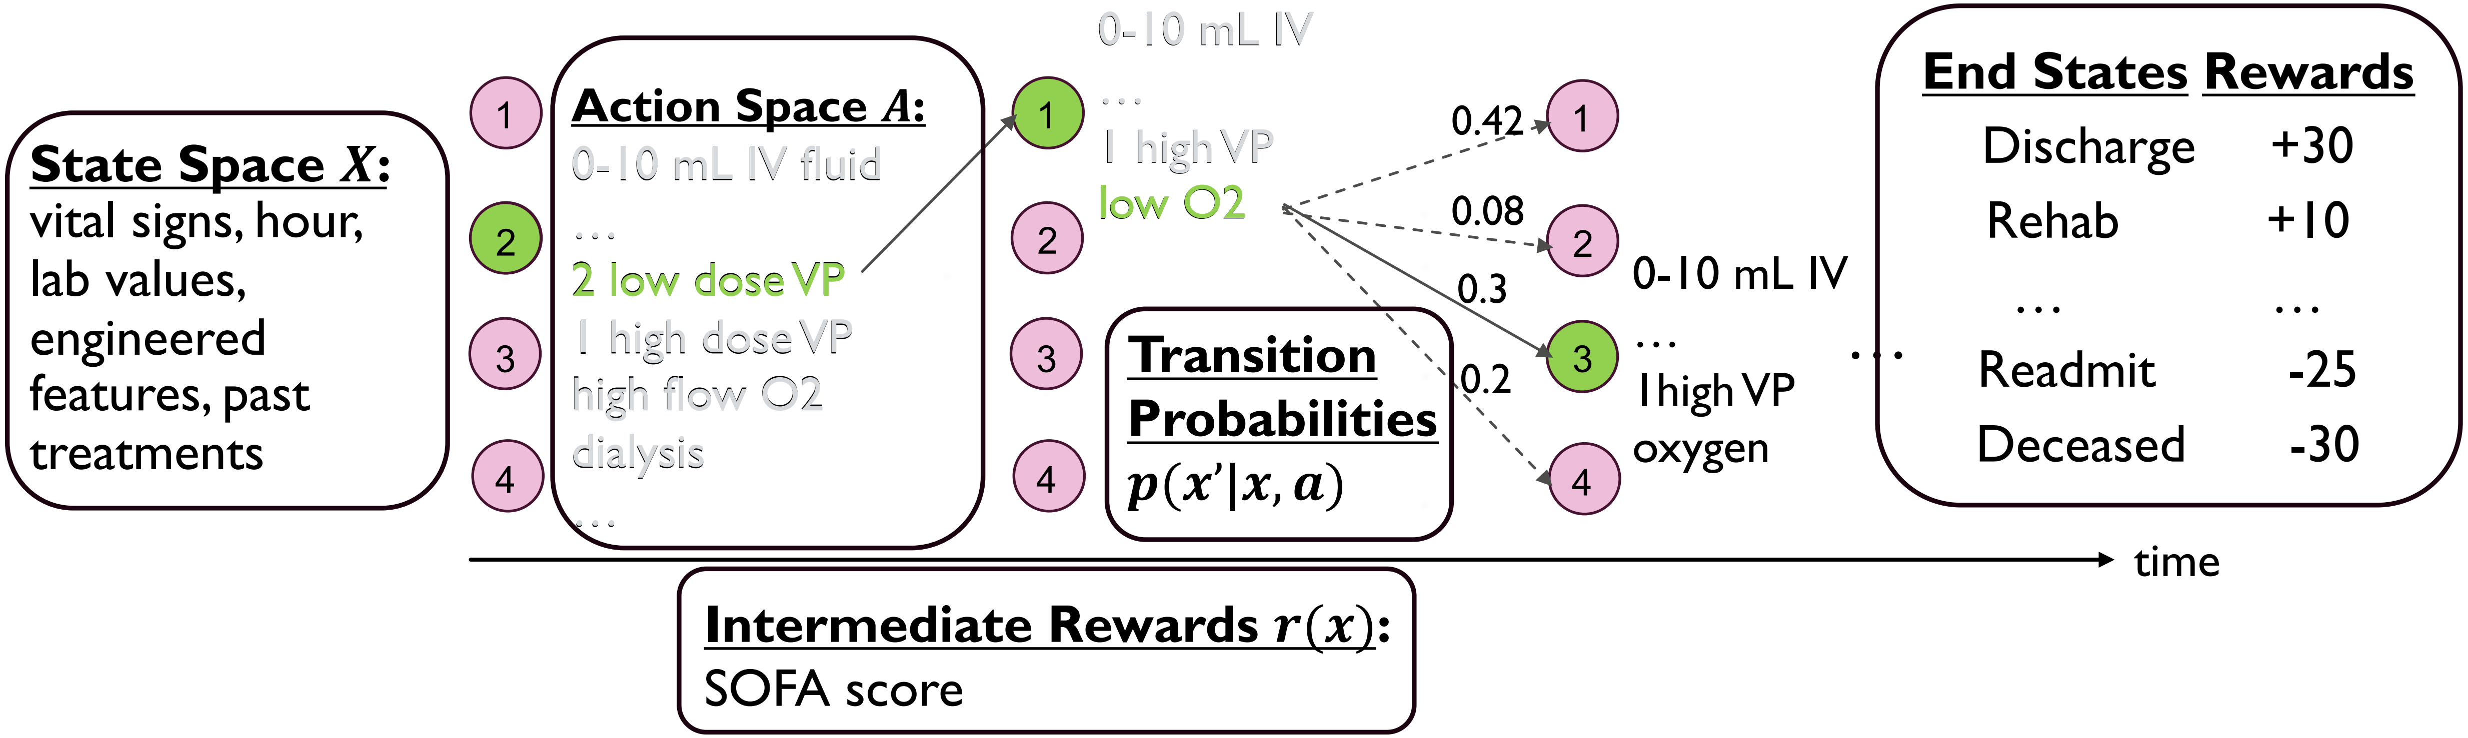

Goals: Learn states and transitions from data; explain model to clinicians

MDP MODEL OF SEPSIS TREATMENT

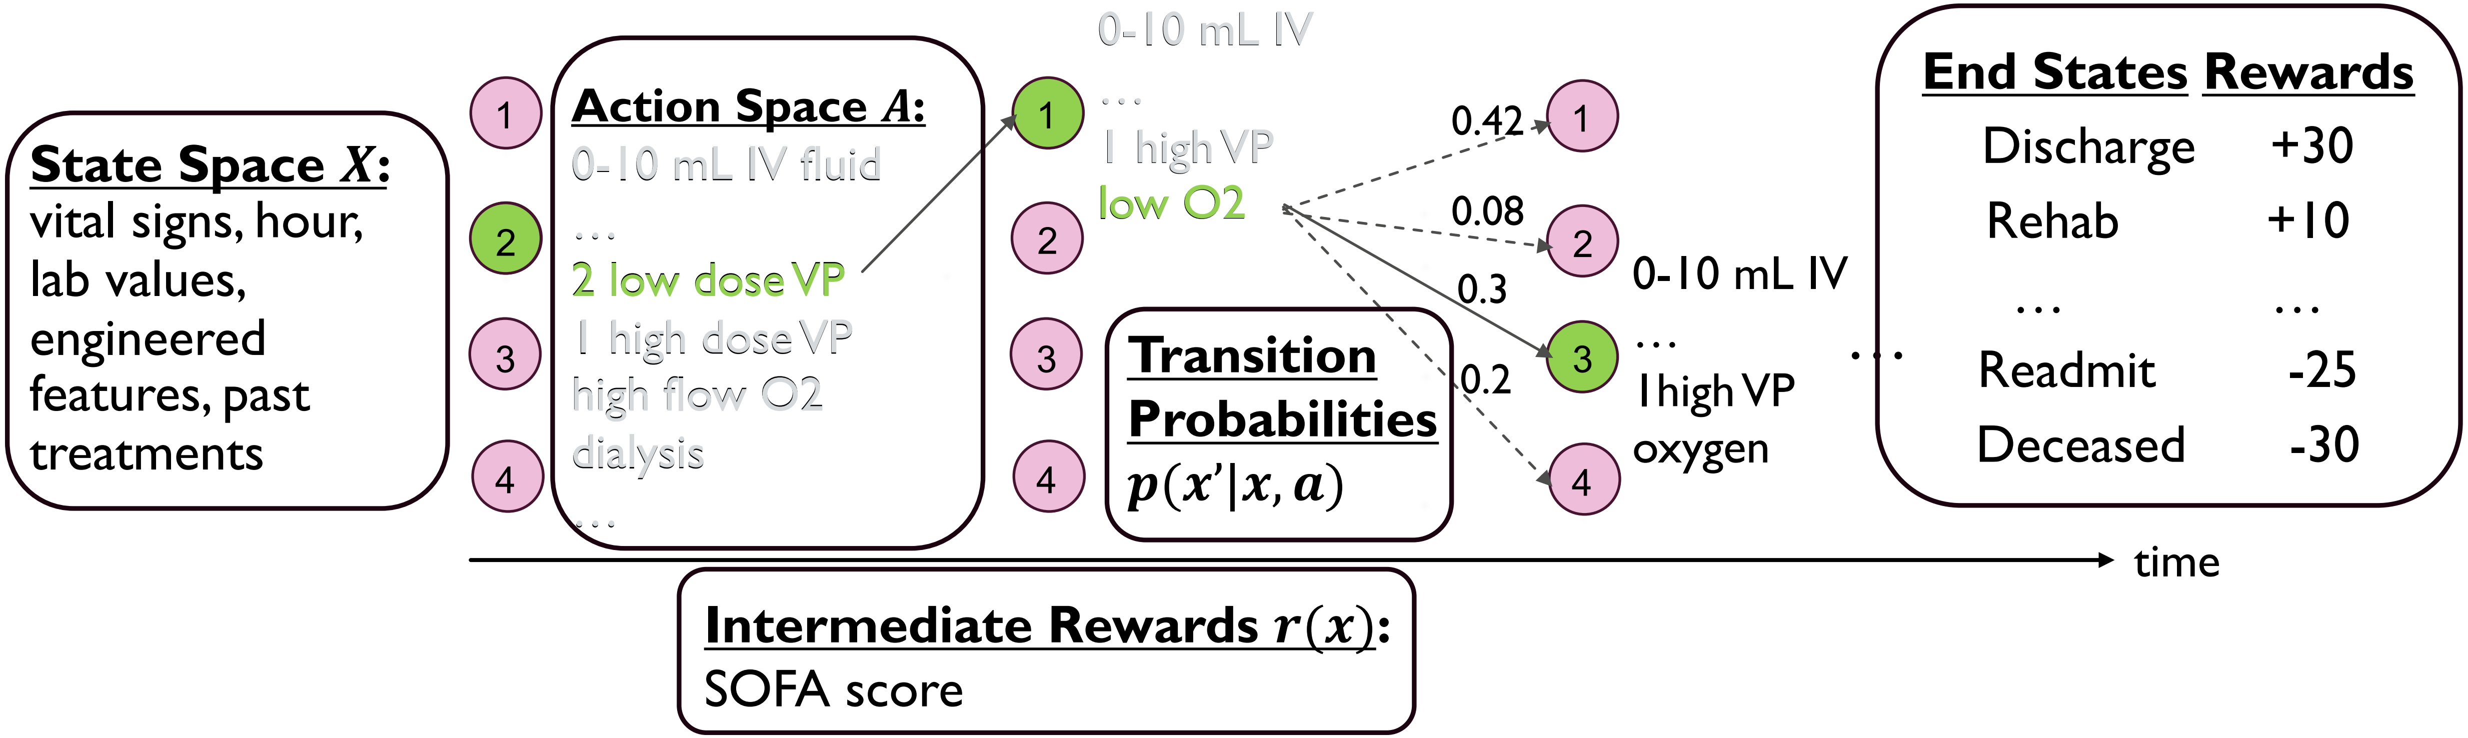

1. Learn states and transition probabilities from a fixed batch of data
2. Solve for the best action for each state (depends on transition probabilities)

CONSTRUCTING STATES TO GET A CONCISE MDP

High level idea: **Iteratively split** feature space X into regions of points (states) S that have **similar rewards** and **respond similarly** when given the same treatments

$$\phi: X \to S$$

CONSTRUCTING STATES TO GET A CONCISE MDP

<u>High level idea:</u> **Iteratively split** feature space X into regions of points (states) S that have **similar rewards** and **respond similarly** when given the same treatments

$$\phi: X \to S$$

CONSTRUCTING STATES TO GET A CONCISE MDP

High level idea: **Iteratively split** feature space X into regions of points (states) S that have **similar rewards** and **respond similarly** when given the same treatments

$$\phi: X \to S$$

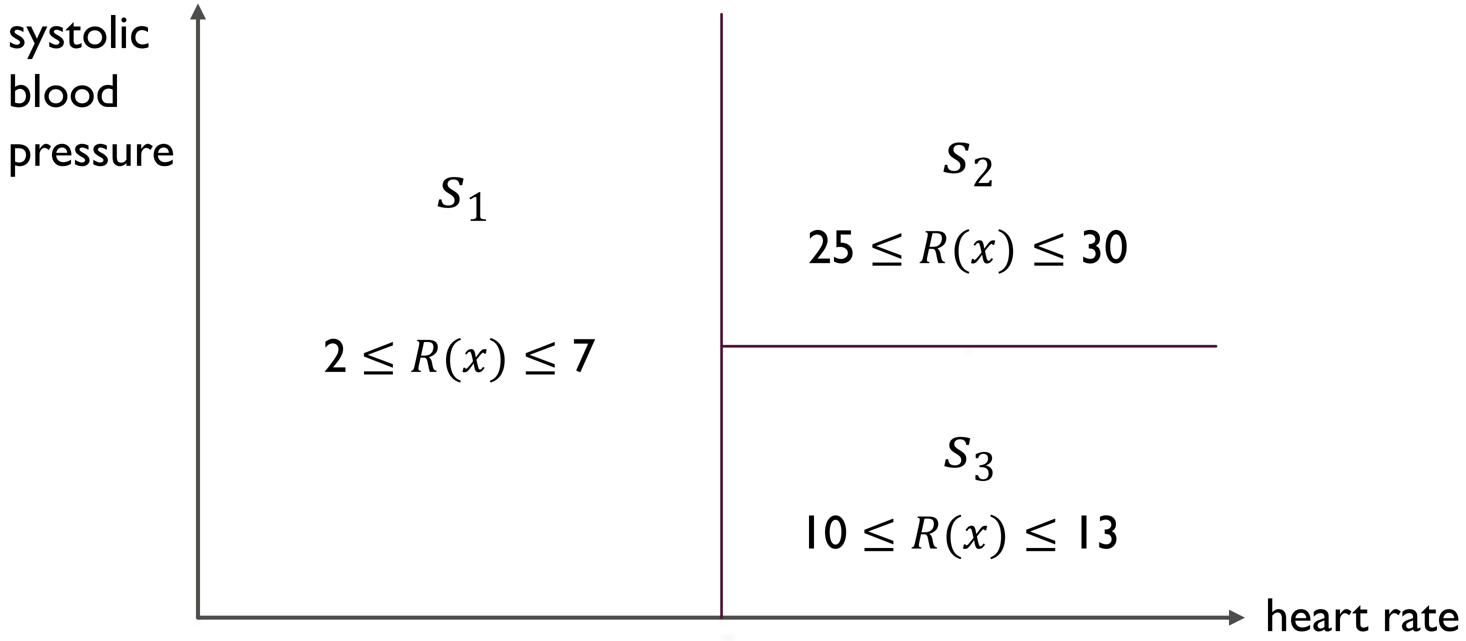

Goals: Learn states and transitions from data; explain model to clinicians

CONSTRUCTING STATES TO GET A CONCISE MDP

High level idea: **Iteratively split** feature space X into regions of points (states) S that have **similar rewards** and **respond similarly** when given the same treatments

$$\phi: X \to S$$

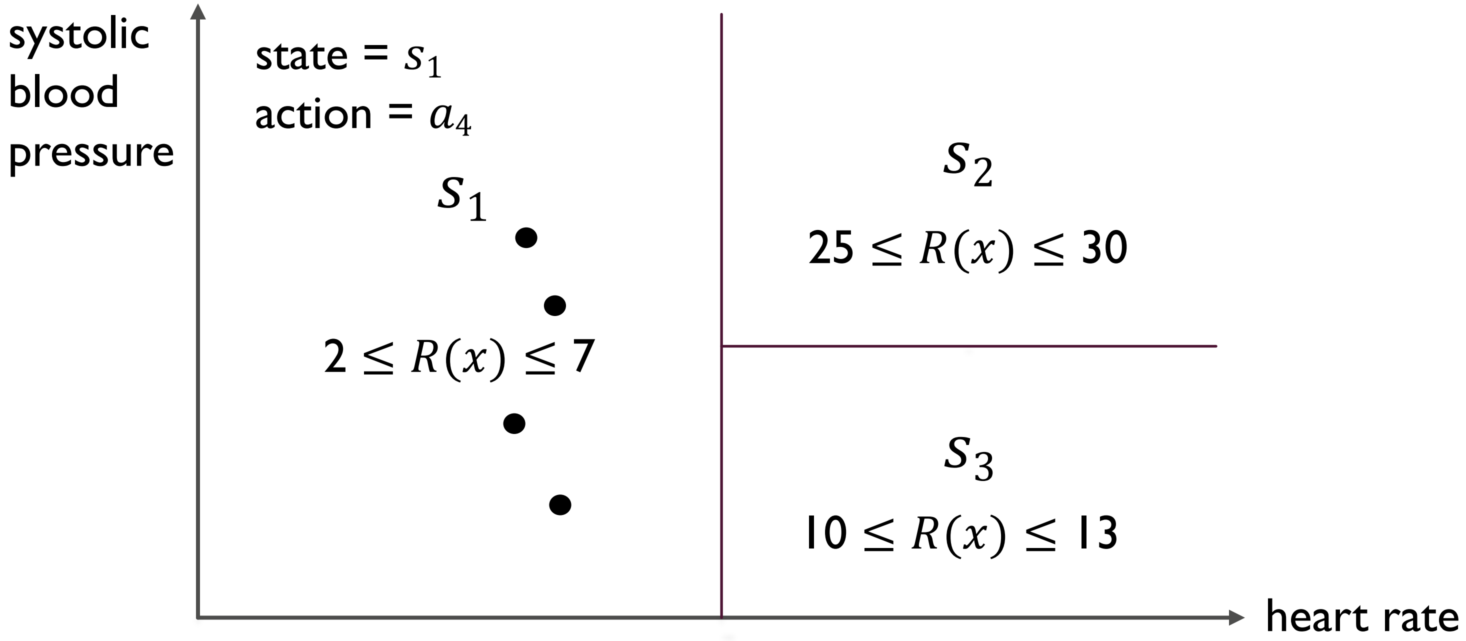

Goals: Learn states and transitions from data; explain model to clinicians

CONSTRUCTING STATES TO GET A CONCISE MDP

High level idea: **Iteratively split** feature space X into regions of points (states) S that have **similar rewards** and **respond similarly** when given the same treatments

$$\phi: X \to S$$

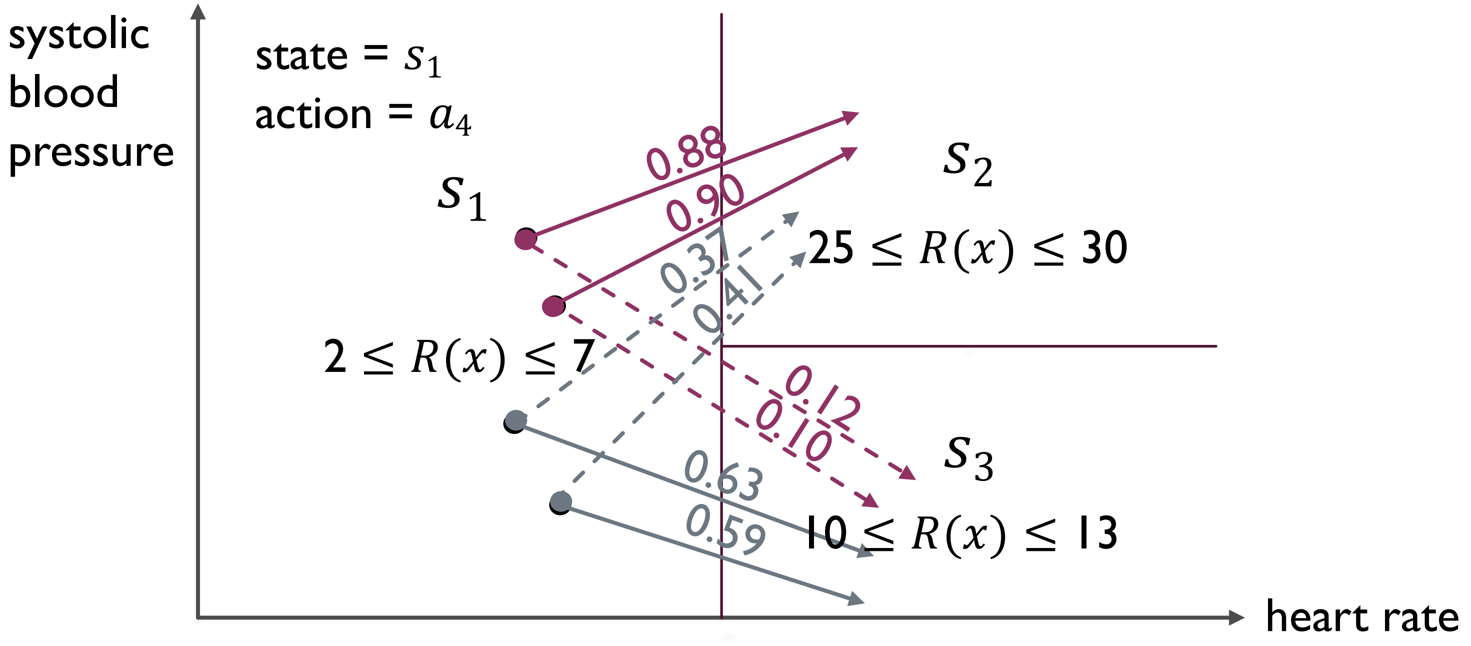

Goals: Learn states and transitions from data; explain model to clinicians

CONSTRUCTING STATES TO GET A CONCISE MDP

High level idea: **Iteratively split** feature space X into regions of points (states) S that have **similar rewards** and **respond similarly** when given the same treatments

$$\phi: X \to S$$

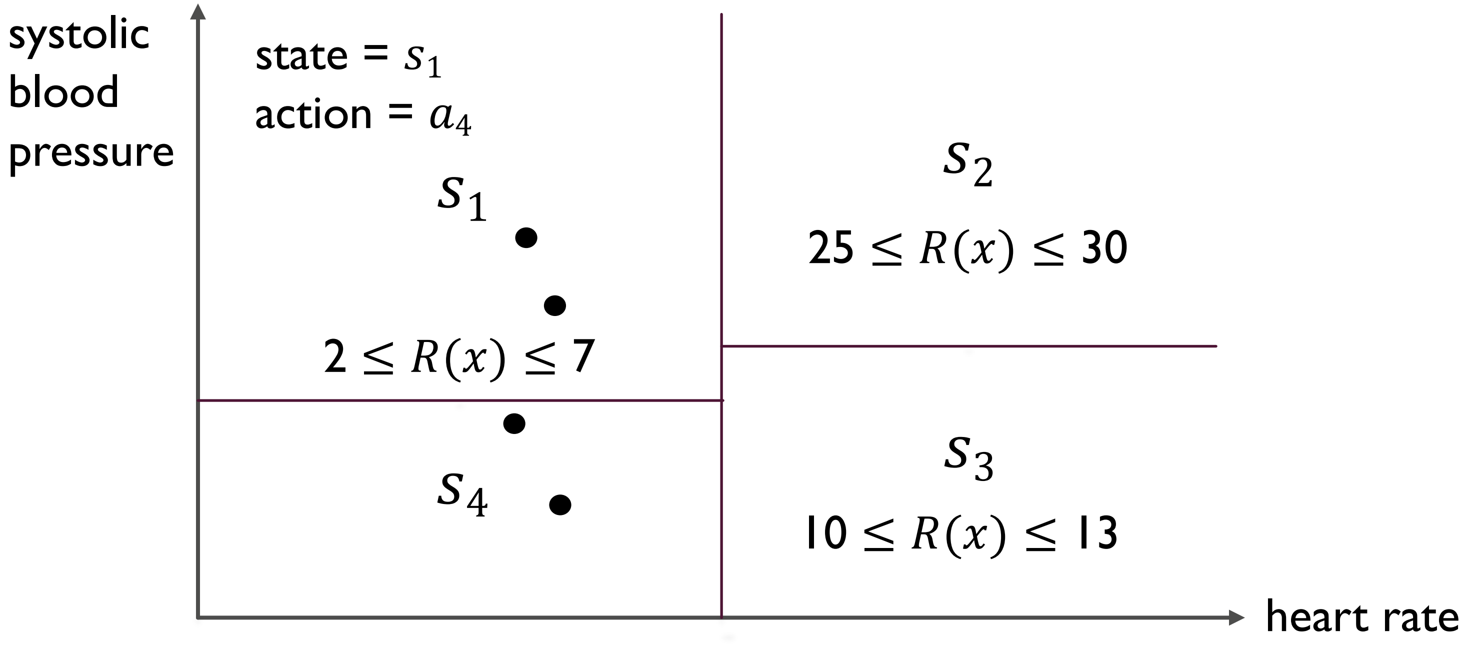

Goals: Learn states and transitions from data; explain model to clinicians

CONSTRUCTING STATES TO GET A CONCISE MDP

High level idea: **Iteratively split** feature space X into regions of points (states) S that have **similar rewards** and **respond similarly** when given the same treatments

$$\phi: X \to S$$

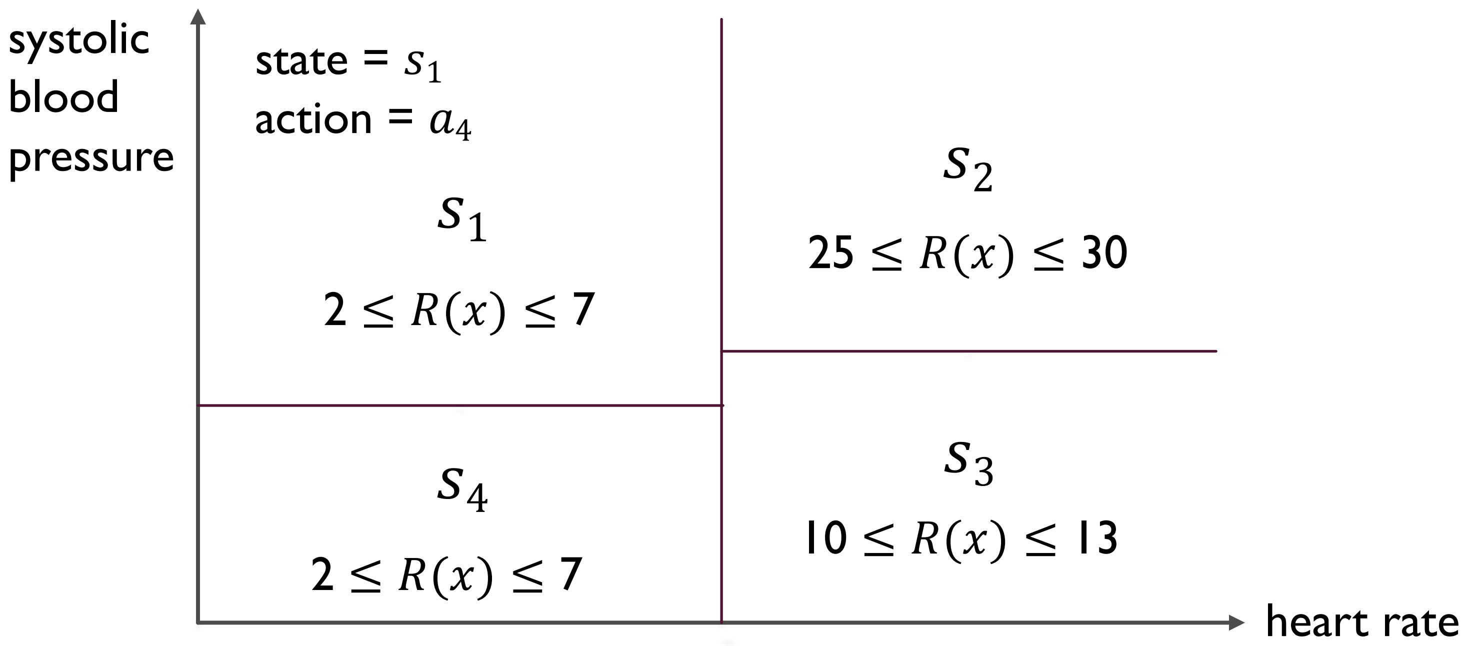

$\hat{r}(s)$ = average reward of data points that fall into s

$$\hat{r}(s) = \frac{\sum_{x \in s} r(x)}{n_s}$$

$\hat{p}(s'|s,a)$ = proportion of points in s taking action a that transition to s'

$$\hat{p}(s'|s,a) = \frac{n_{sas'}}{n_{sa}}$$

INTERPRETING THE FEATURES-TO-STATE MAPPING $\phi: X \to S$

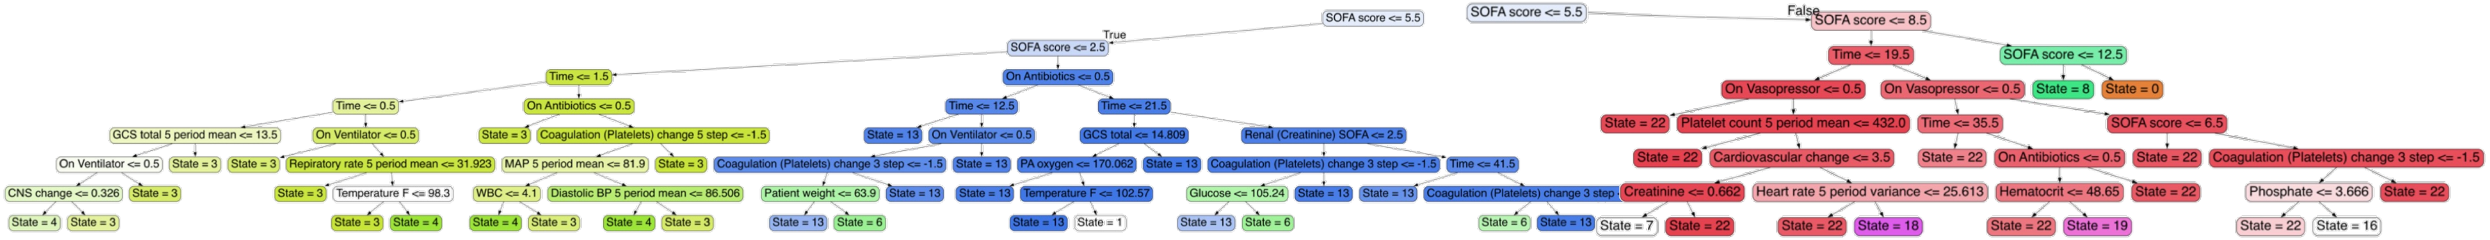

Goals: Learn states and transitions from data; explain model to clinicians

INTERPRETING THE FEATURES-TO-STATE MAPPING $\phi: X \rightarrow S$

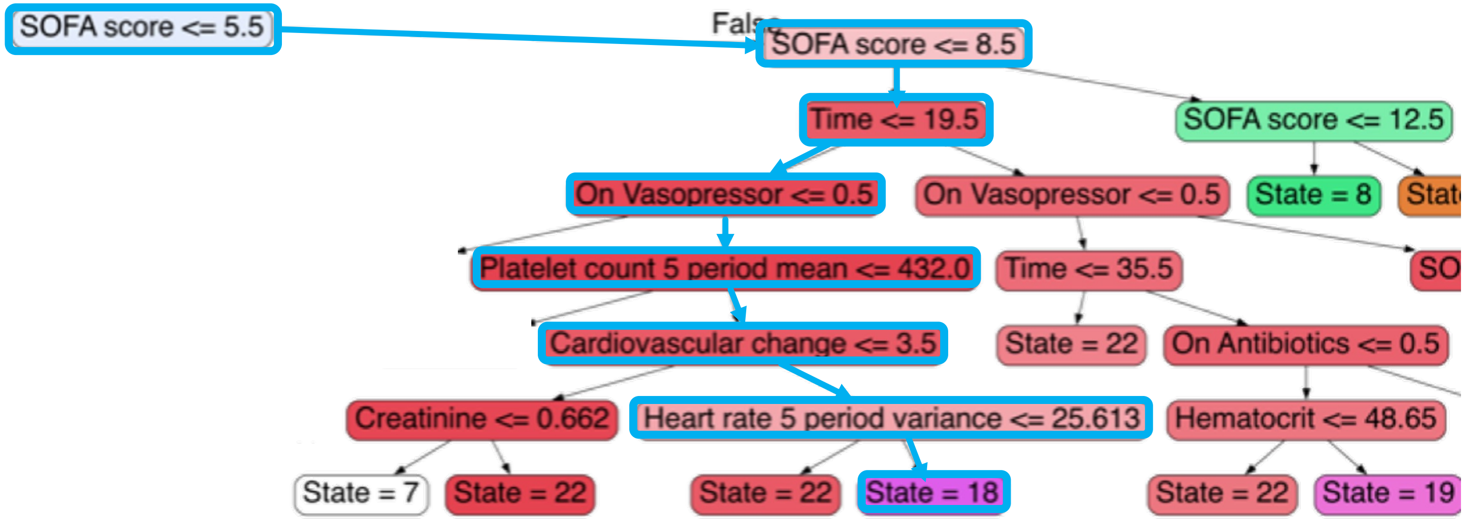

Goals: Learn states and transitions from data; explain model to clinicians

SOLVE LEARNED MDP TO OBTAIN POLICY RECOMMENDATIONS

We now have the states and transition probabilities of the MDP.

Traditionally, **"best actions"** are chosen based on expected value / **average** performance.

SOLVE LEARNED MDP TO OBTAIN POLICY RECOMMENDATIONS

We now have the states and transition probabilities of the MDP.

Traditionally, **"best actions"** are chosen based on expected value / **average** performance.

Since the actions are treatments for critically ill patients, we want:

1. Risk-averse (to bad outcomes) policies
2. Robust (to estimation error) policies
3. Expert-guided decisions

RISK-AVERSE VALUE ITERATION

Traditionally, MDPs are solved for policies that maximize expected (long term) value

Value iteration: $\mathcal{T}V(s) = max_{a\in\mathcal{A}}\{r(s) + \mathbb{E}_{s'}[\gamma V(s')]\} = max_{a\in\mathcal{A}}\{r(s) + \gamma \sum_{s'} p(s'|s,a)V(s')\}$

RISK-AVERSE VALUE ITERATION

Traditionally, MDPs are solved for policies that maximize expected (long term) value

Value iteration: $\mathcal{T}V(s) = max_{a \in \mathcal{A}}\{r(s) + \mathbb{E}_{s'}[\gamma V(s')]\} = max_{a \in \mathcal{A}}\{r(s) + \gamma \sum_{s'} p(s'|s,a)V(s')\}$

When making patient treatment decisions, need to consider **worst-case outcomes**

RISK-AVERSE VALUE ITERATION

Traditionally, MDPs are solved for policies that maximize expected (long term) value

Value iteration: $\mathcal{T}V(s) = max_{a \in \mathcal{A}}\{r(s) + \mathbb{E}_{s'}[\gamma V(s')]\} = max_{a \in \mathcal{A}}\{r(s) + \gamma \sum_{s'} p(s'|s,a)V(s')\}$

When making patient treatment decisions, need to consider **worst-case outcomes**

<u>Conditional Value at Risk (CVaR):</u> $CVaR^{(q)}(Z) = \mathbb{E}[Z | Z \leq VaR^{(q)}(Z)]$

where $VaR^{(q)}(Z)$ is the q[th] quantile of random variable Z

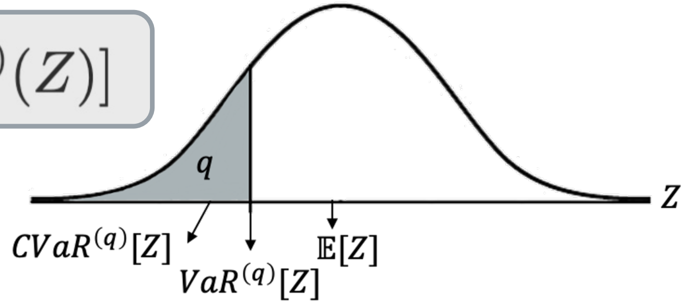

Goals: Risk-averse, robust treatment policies; expert-guided decisions

RISK-AVERSE VALUE ITERATION

Traditionally, MDPs are solved for policies that maximize expected (long term) value

Value iteration: $\mathcal{T}V(s) = max_{a\in\mathcal{A}}\{r(s) + \mathbb{E}_{s'}[\gamma V(s')]\} = max_{a\in\mathcal{A}}\{r(s) + \gamma \sum_{s'} p(s'|s,a)V(s')\}$

When making patient treatment decisions, need to consider **worst-case outcomes**

Conditional Value at Risk (CVaR): $CVaR^{(q)}(Z) = \mathbb{E}[Z|Z \leq VaR^{(q)}(Z)]$

where $VaR^{(q)}(Z)$ is the q^{th} quantile of random variable Z

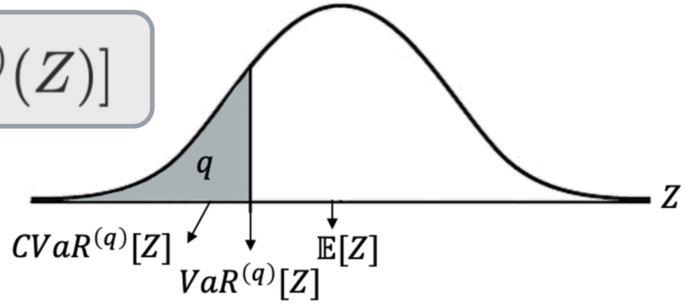

CVaR value iteration: $\mathcal{T}^{(q)}V(s) = max_{a\in\mathcal{A}}\{r(s) + \gamma CVaR^{(q)}_{s'}[V(s')]\}$

$$= max_{a\in\mathcal{A}}\{r(s) + \gamma \sum_{s':V(s')<VaR^{(q)}_{s'}[V(s')]} \frac{1}{q} p(s'|s,a)V(s')\}$$

Goals: Risk-averse, robust treatment policies; expert-guided decisions

SIMULATED LONG-TERM VALUES UNDER POLICIES

We simulate the cumulative rewards obtained from following CVaR value iteration policy and traditional value iteration policy
on sepsis dataset

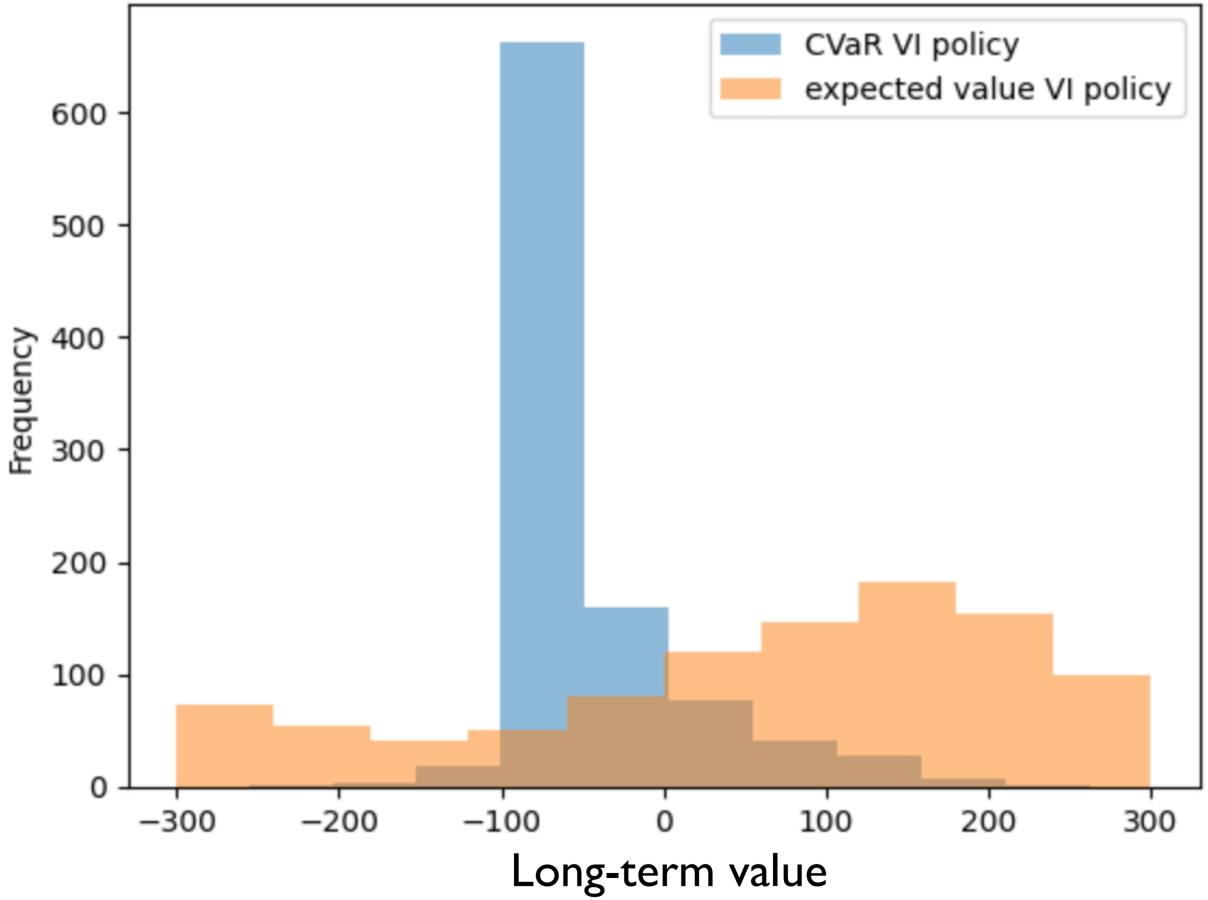

Goals: Risk-averse, robust treatment policies; expert-guided decisions

CVAR VALUE ITERATION IS ROBUST TO ESTIMATION ERROR

Iterated CVaR MDP (solved with CVaR value iteration)

$$\max_{\pi} ICVaR_{\hat{p}}^{(q)}(s)$$

is equivalent to[1] ⟺

Robust MDP (with a particular uncertainty set)

$$\max_{\pi} \min_{p \in U_p} V_p^{\pi}(s)$$

$$\mathcal{U}_p = \{\otimes_{s,a} p(.|sa) | \forall (s,a) \in \mathcal{S} \times \mathcal{A} : 0 \leq p(s'|s,a) \leq \frac{1}{q}\hat{p}(s'|s,a)\ \forall s' \in \mathcal{S}, \sum_{s' \in \mathcal{S}} p(s'|s,a) = 1\}$$

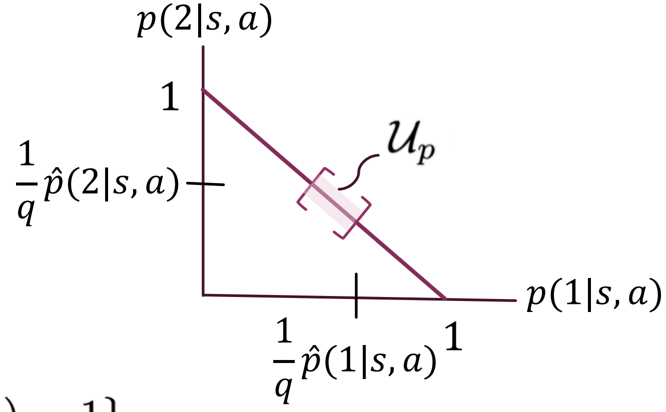

i.e. we get robustness for free when we use CVaR value iteration!

[1]Osagami 2012

PERFORMANCE ON WORST-CASE MIS-ESTIMATION

We evaluated the both robust optimal policy and traditional optimal policy on the **worst-case realization of p** $\in U_p$

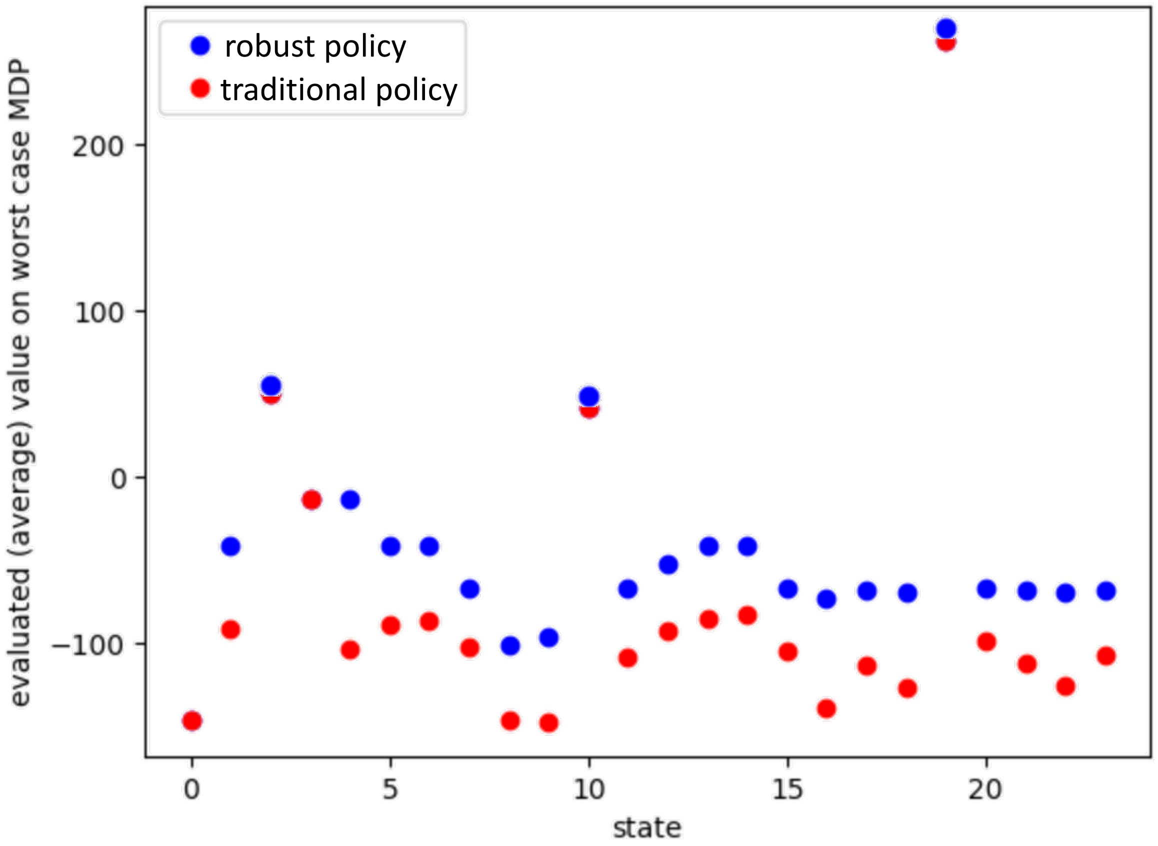

Goals: Risk-averse, robust treatment policies; expert-guided decisions

PILOTING RECOMMENDATIONS FOR SEPSIS TREATMENT

1. Co-design user interface of the recommendation tool

Characteristics defining patient's current state

- SOFA score < 5.5
- SOFA score > 2.5
- Not on antibiotics
- Time in ICU < 21.5 hours
- Glasgow Coma Score < 14
- Temperature > 102.5

state 1

Values of actions for patient's current state

Probabilities of reaching each end state from action*

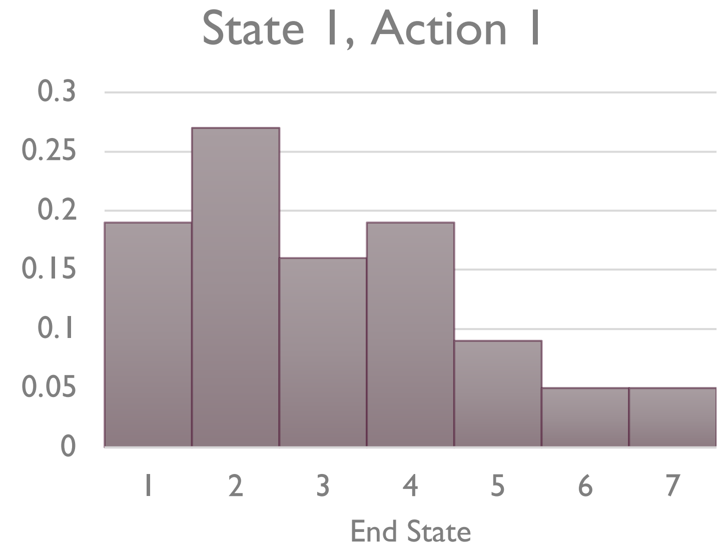

*only holds if optimal policy is followed thereafter

Goals: Risk-averse, robust treatment policies; expert-guided decisions

PILOTING RECOMMENDATIONS FOR SEPSIS TREATMENT

2. In simulations, have clinicians provide feedback on recommendations

Patient characteristics
- BP 83/50
- Lactate 4.5
- Heart rate 120
- Oxygen sat 87%
- Creatinine 2.5
- Urine output 0 …

Clinician inputs their decision

Clinician is shown recommendation

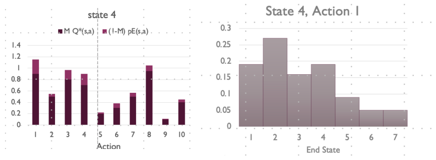

Agree or disagree? Why?

Goals: Risk-averse, robust treatment policies; expert-guided decisions

PILOTING RECOMMENDATIONS FOR SEPSIS TREATMENT

2. In simulations, have clinicians provide feedback on recommendations

Patient characteristics
- BP 83/50
- Lactate 4.5
- Heart rate 120
- Oxygen sat 87%
- Creatinine 2.5
- Urine output 0 …

Clinician inputs their decision

Clinician is shown recommendation

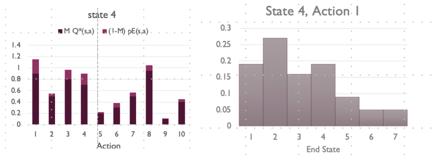

Agree or disagree? Why?

3. While treating live patients, show clinicians the recommendations

Clinician inputs their decision

Clinician is shown recommendation

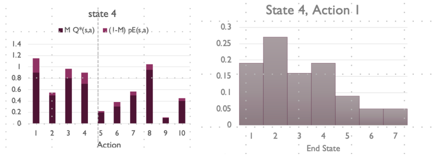

Take or not? Why?

Goals: Risk-averse, robust treatment policies; expert-guided decisions

PATIENTS ARE LINKED BY SHARED RESOURCE CAPACITY

- Assumption so far: each patient's treatment is **independent** of each other
- In reality, patients in a hospital are linked by a **common pool of resources**
- To account for this, model all patients together as a system

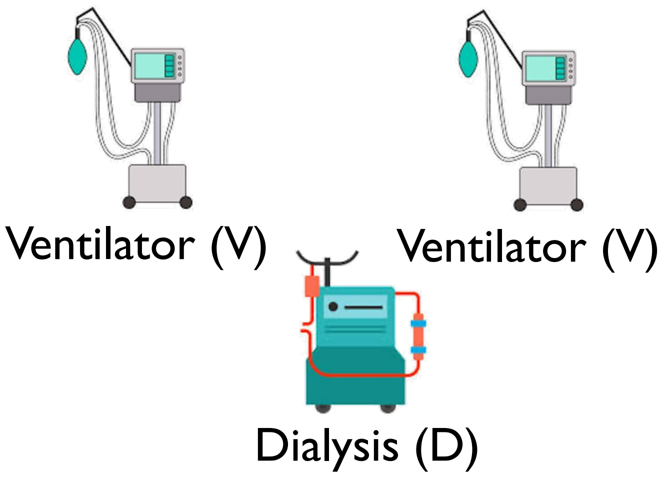

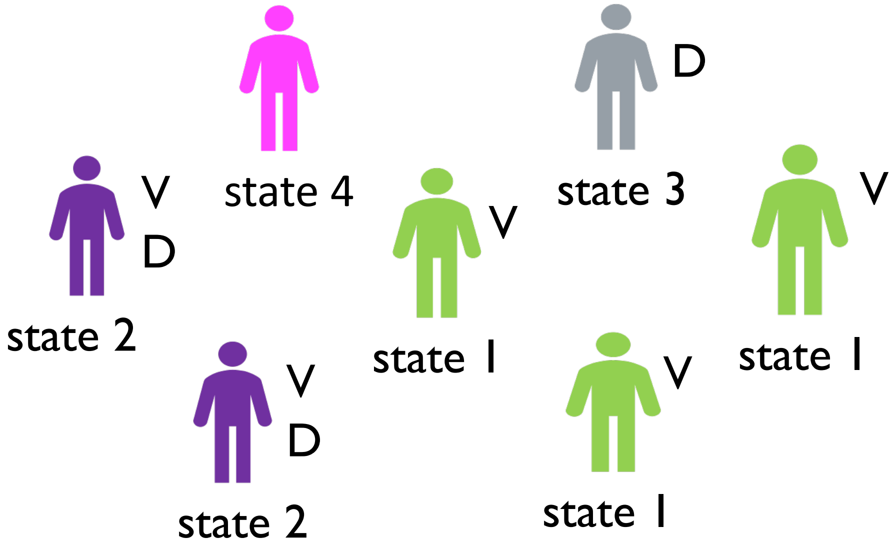

Goals: Capacity-aware MDP solution; forward-looking, tractable formulation for allocating resources

RESOURCE CAPACITY-AWARE MDP

- We can model all patients together as a weakly coupled MDP, **M = (S, A, P, R)** where
  - **S** = $\times_i S_{\boldsymbol{i}}$ 
  - **A**$(\boldsymbol{s})$ = $\{\boldsymbol{a} \in \times_i A_i \colon \sum_{\boldsymbol{i}} \boldsymbol{D_i}(s_i, a_i) \leq \boldsymbol{b}\}$, where $\boldsymbol{D_i}(s_i, a_i)$ is the number of resources used by that state-action combination s_i, a_i 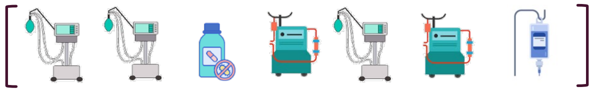
  - **P**$(\boldsymbol{s}'|\boldsymbol{s}, \boldsymbol{a})$ = $\prod_i p_i(s'_i|s_i, a_i)$
  - **R**$(\boldsymbol{s})$ = $\sum_i r_i(s_i)$
- This weakly coupled MDP has an exponentially large state and action space, which makes it intractable to solve for more than a few patients
- Solve Lagrangian relaxation!

Goals: Capacity-aware MDP solution; forward-looking, tractable formulation for allocating resources

ALLOCATING LIMITED RESOURCES IN REAL TIME

$\boldsymbol{x_{sj}}$ = number of patients in state s to assign to resource combination j at the current time

$\boldsymbol{v_{sj}}$ = $\max_{a \text{ in } A(j)} \{Q^*(s,a)\}$ where $A(j)$ is the set of actions that use resource combination j

$$\max_x \sum_{s \in S} \sum_{j \in J} \boldsymbol{v_{sj}} x_{sj}$$

$$\text{s.t.} \sum_{s \in S} \sum_{j \in J} c_{rj} x_{sj} \leq u_r \quad \forall r \in \mathcal{R}$$

$$\sum_{j \in J} x_{sj} = n_s \quad \forall s \in \mathcal{S}$$

$$x_{sj} \in Z^+ \quad \forall s \in \mathcal{S}, j \in \mathcal{J}$$

Maximize total long term value

subject to

Resources used does not exceed resources available

Every patient is assigned to one resource combination

R = {ventilator, dialysis machine}

J = {ventilator and dialysis, ventilator only, dialysis only, no ventilator or dialysis}

Goals: Capacity-aware MDP solution; forward-looking, tractable formulation for allocating resources

SIMULATED DISCHARGES UNDER VARYING RESOURCE LEVELS

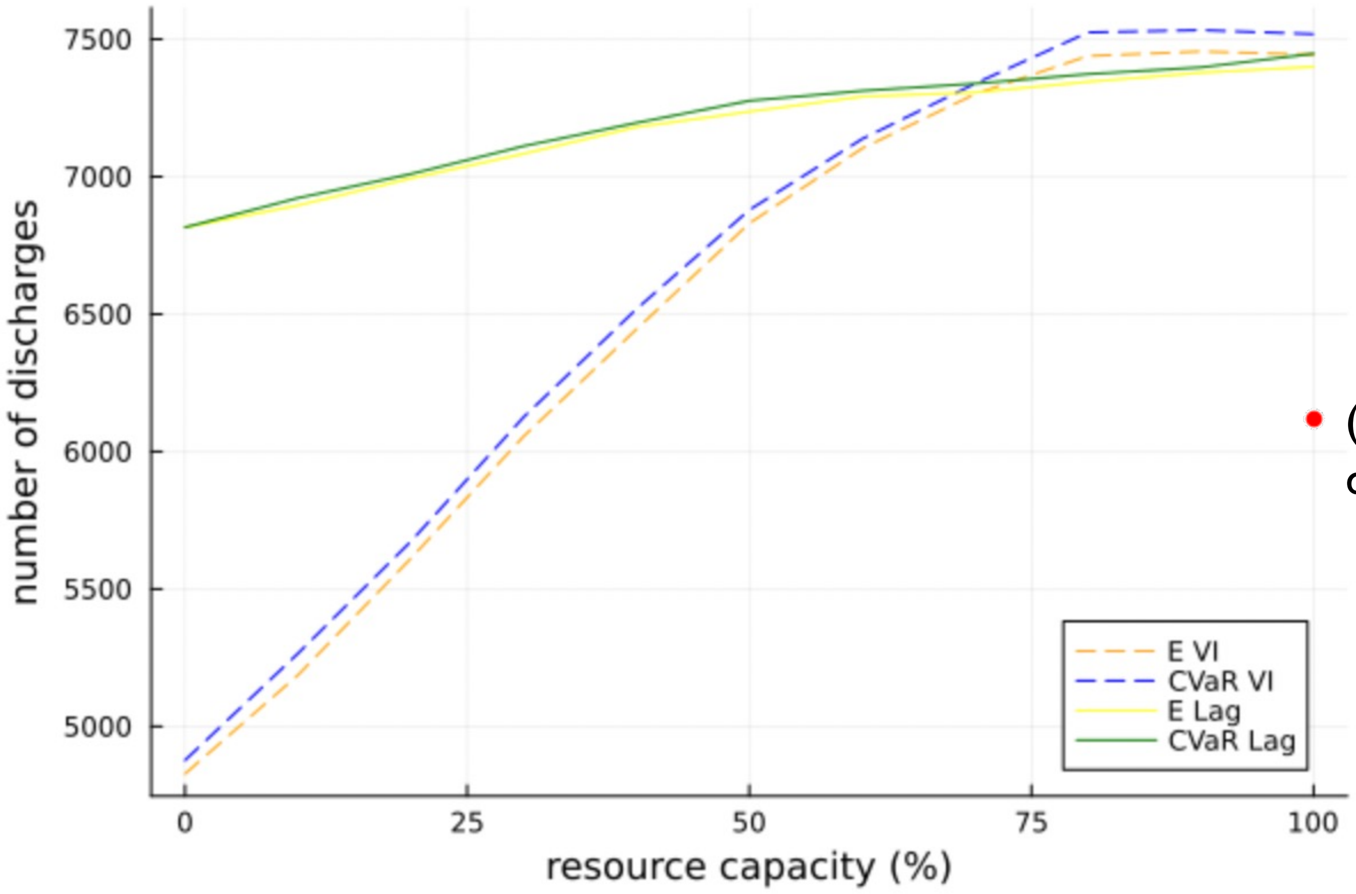

(observed number of discharges)

(observed resource utilization level)

Goals: Capacity-aware MDP solution; forward-looking, tractable formulation for allocating resources

CONCLUSIONS

- Learned the dynamics of **sepsis patient treatment** using an observational dataset by designing a **concise, interpretable MDP model**
- Proposed **risk-aware, robust, capacity-aware MDP solution methods** for obtaining recommended sepsis treatment policies
- Provided a **flow chart** for classifying sepsis patients and **user interface** for recommending treatments
- Designed **ongoing pilot** of decision-support tool
- Formulated tractable optimization for **real-time resource allocation** for sepsis treatment
- Illustrated benefits of recommendations using **simulations** on real patient data

Dynamic Optimization of Workforce Talent

Parshan Pakiman

(University at Buffalo – SUNY)

Dan Adelman (U Chicago)

Adam Mersereau (UNC)

IMSI's Workshop on Advances in Quantitative Medical Care

Motivation

- Our work is motivated by the daily challenges faced by UChicago Medicine's operating room manager in assigning operating rooms.

- Familiarity is a key operational factor in designing effective assignments.

- Empirical evidence suggests that greater familiarity improves patient outcomes and costs (Avgerinos & Gokpinar, 2017; Chen, 2021; Witmer et al., 2023; Tuzcuoglu et al., 2024).

Dynamic Optimization of Familiarity Levels

Broader labor-management concept: workers gain familiarity with jobs by performing them repeatedly, which can lead to improved performance (Wright 1936, Argote and Epple 1990)

- **Service** (Gans and Zhou 2003, Chen et al. 2017, Bavafa and Jónasson 202, ...)
- **Manufacturing** (Huckman et al. 2009, Huckman and Staats 2011, ...)

Dynamic Optimization of Familiarity Levels

Broader labor-management concept: workers gain familiarity with jobs by performing them repeatedly, which can lead to improved performance (Wright 1936, Argote and Epple 1990)

- **Service** (Gans and Zhou 2003, Chen et al. 2017, Bavafa and Jónasson 202, …)
- **Manufacturing** (Huckman et al. 2009, Huckman and Staats 2011, …)

Existing Work	Our Work
Modeling the Right Notion of Familiarity from Data	Dynamically Optimizing Familiarity Levels

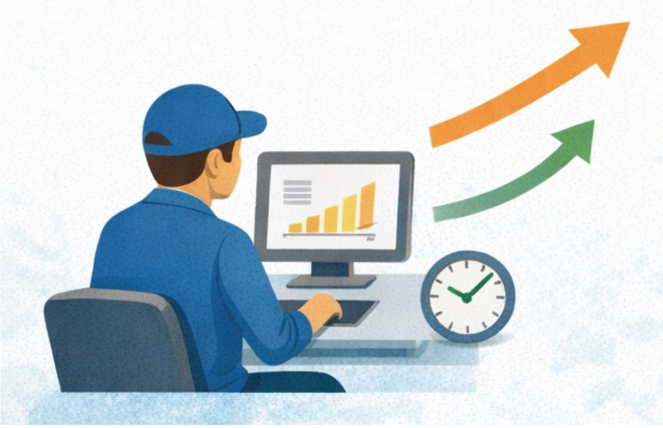

A New Dynamic Assignment Problem

We study a dynamic assignment problem of "jobs" to "workers" under **familiarity** and **availability** dynamics

A New Dynamic Assignment Problem

Familiarity dynamics:

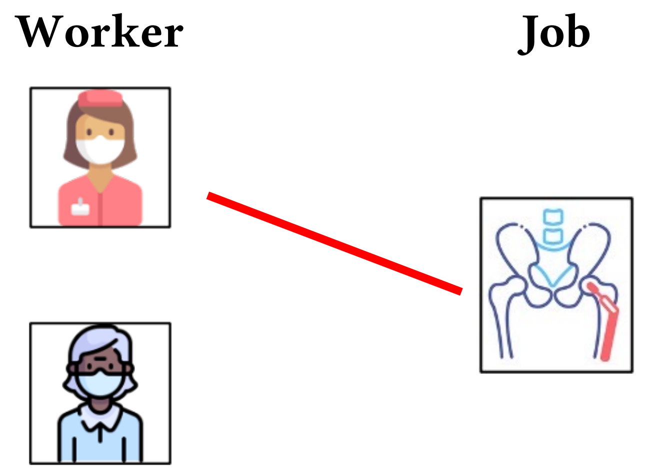

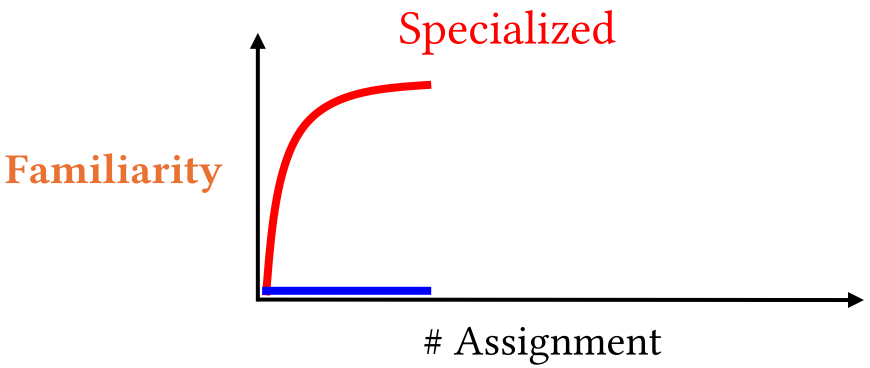

A New Dynamic Assignment Problem

Familiarity dynamics:

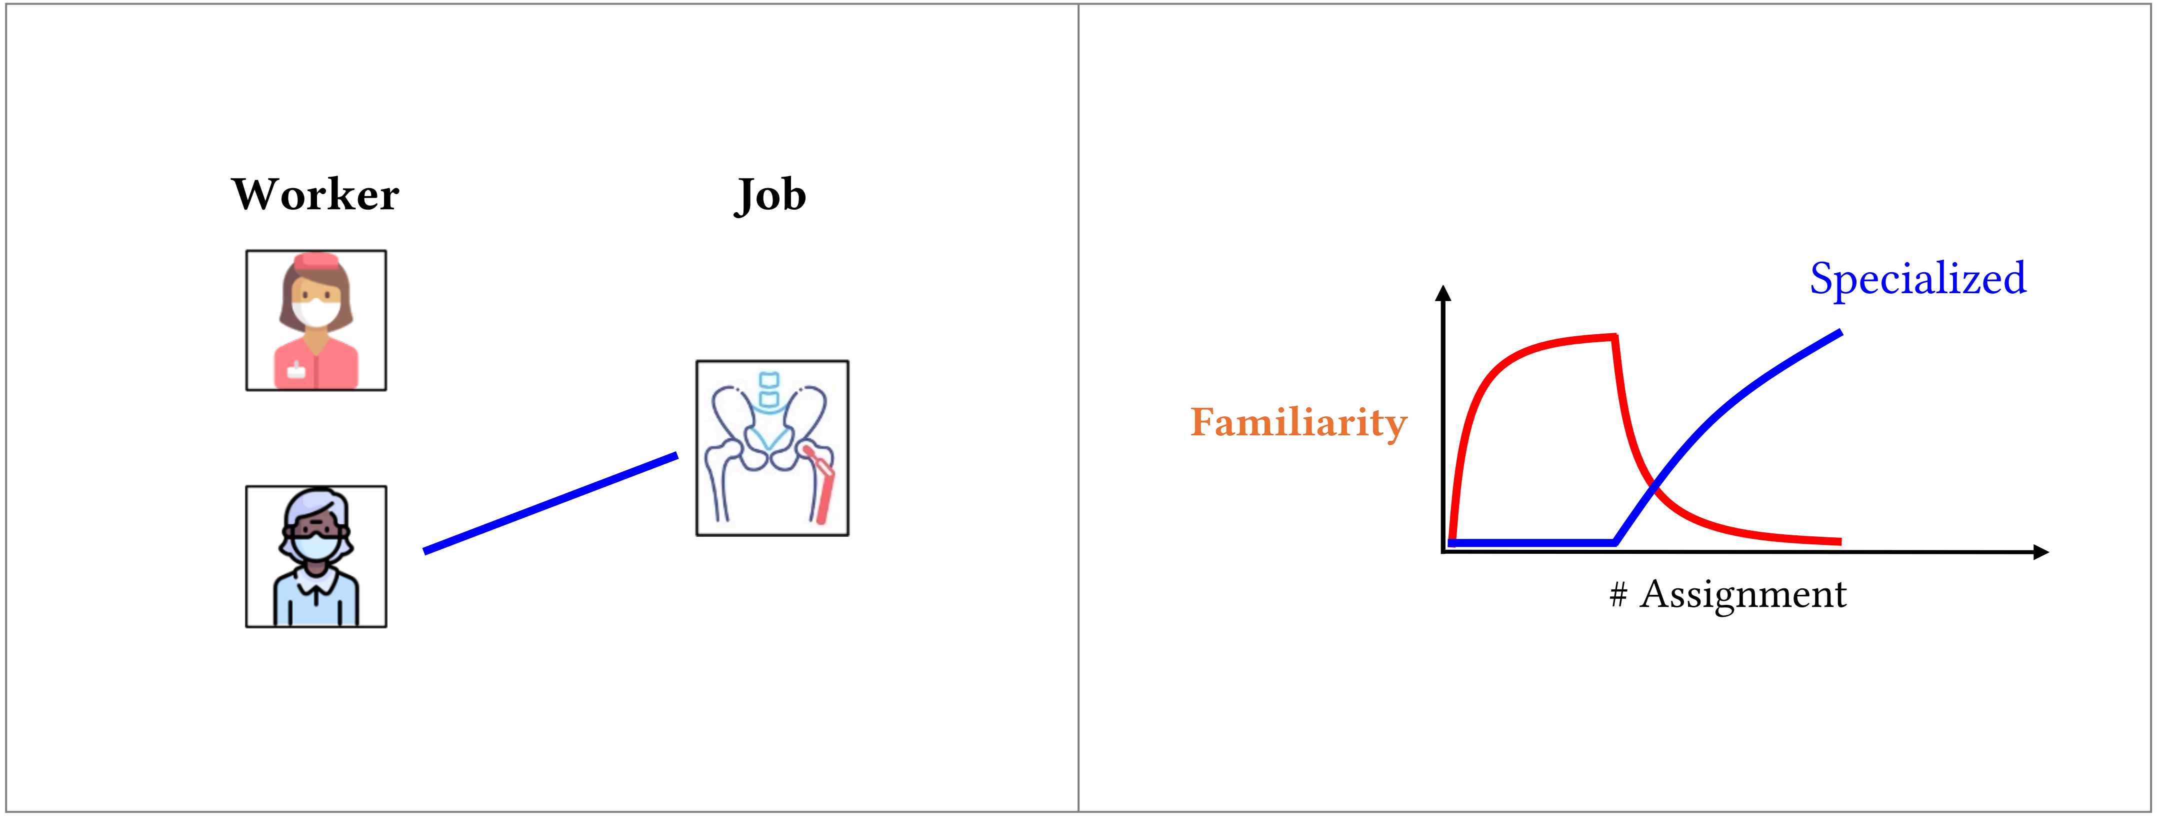

A New Dynamic Assignment Problem

Availability dynamics:

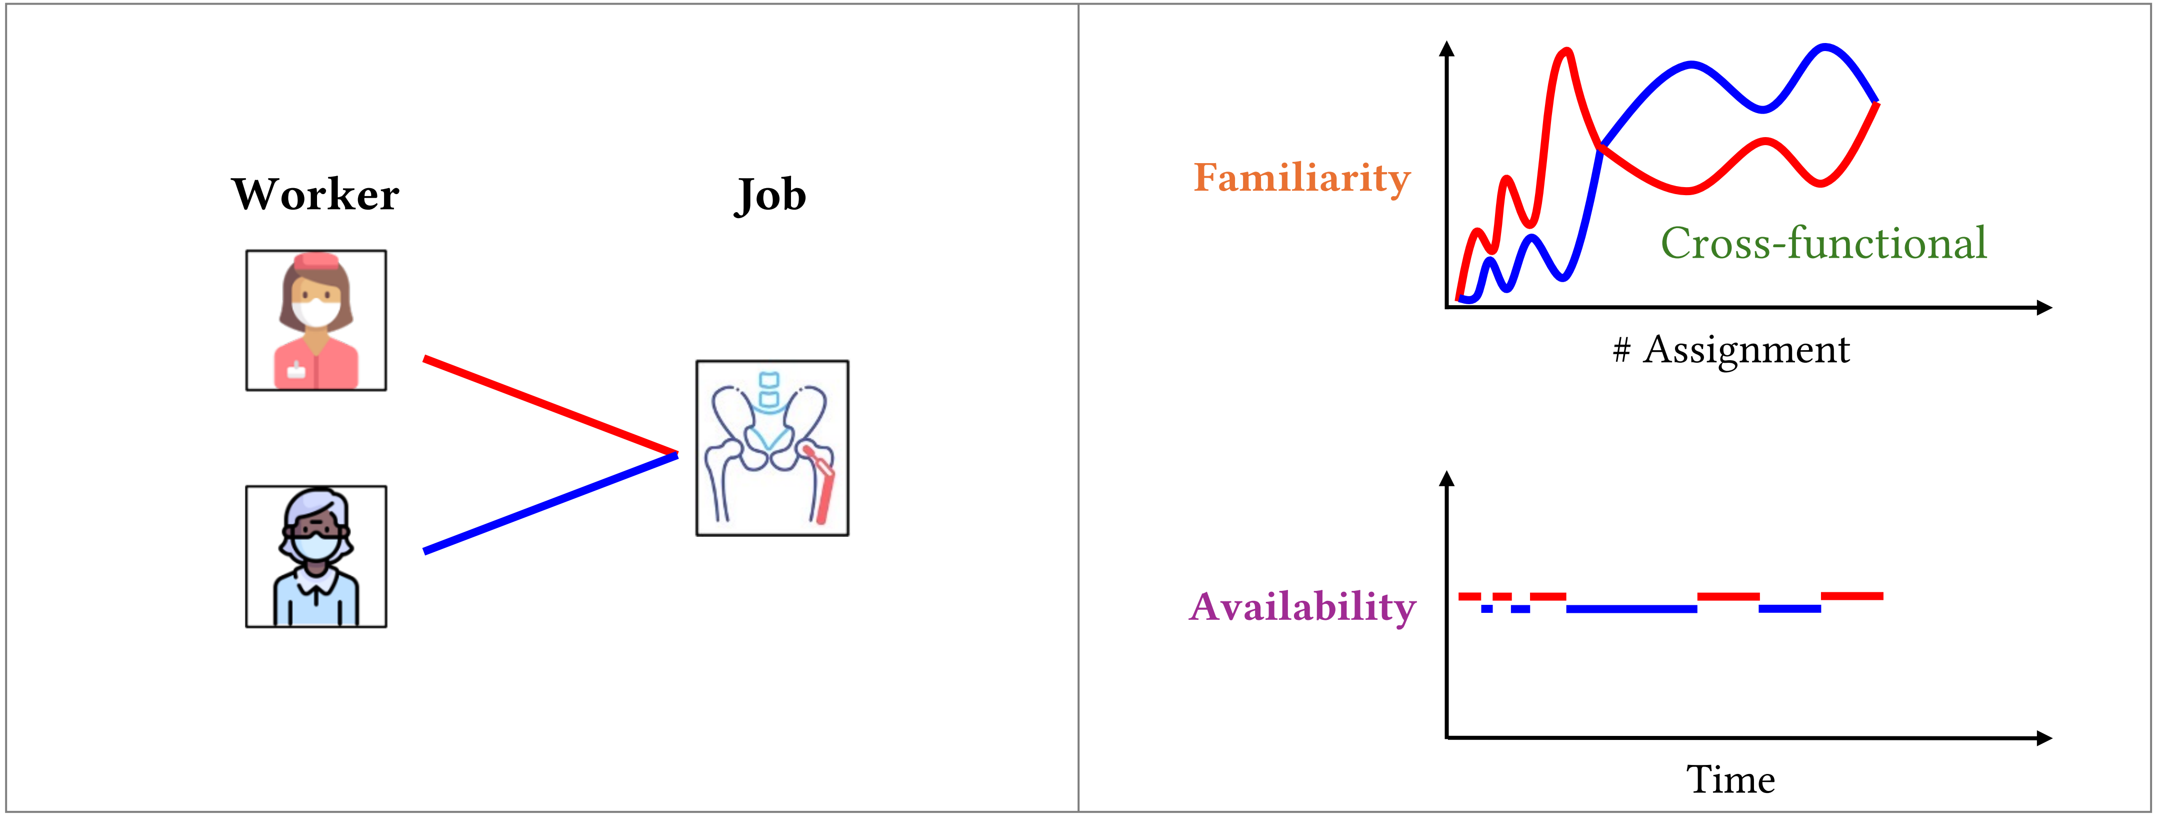

A New Dynamic Assignment Problem

Availability dynamics:

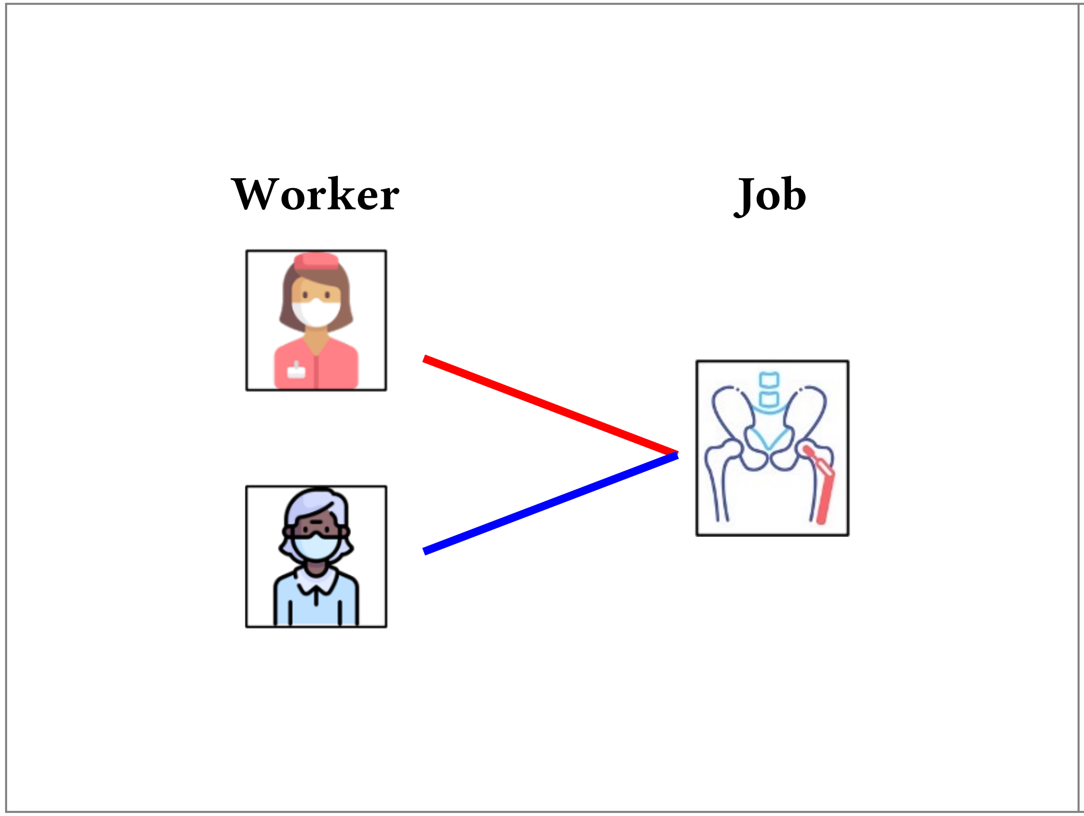

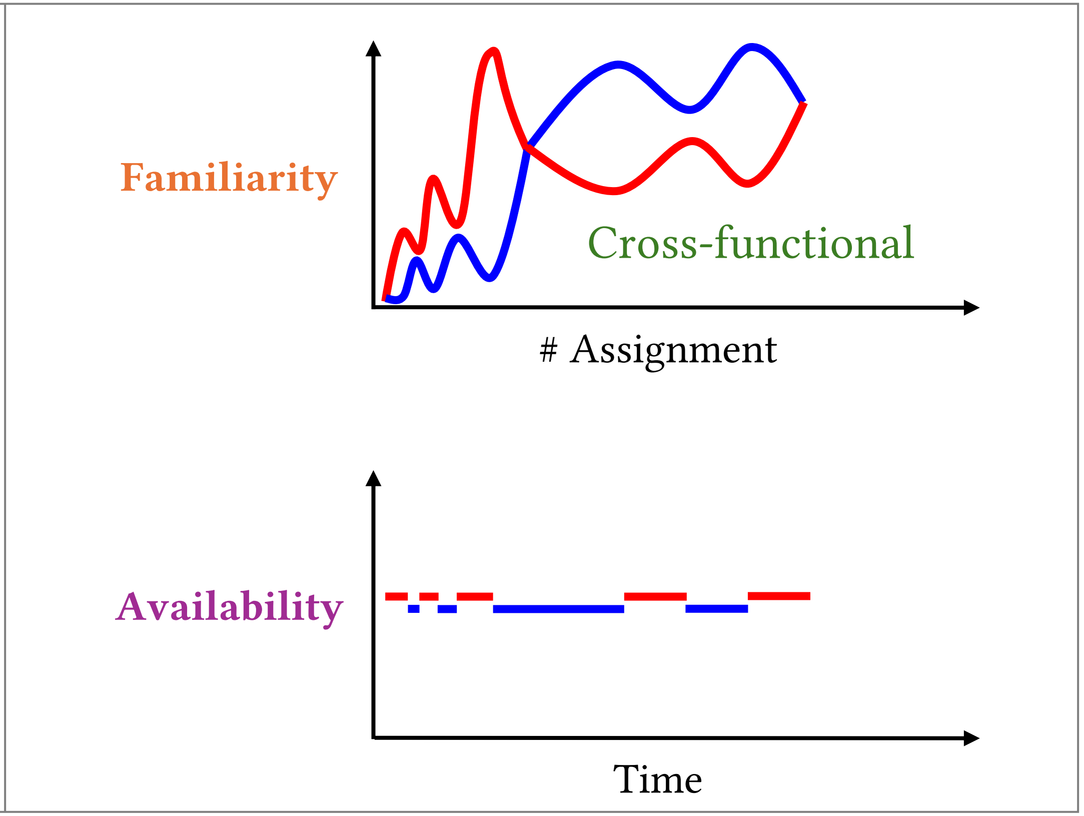

Management trade-off: How should a manager balance the benefits of maintaining specialized workers versus reaping the benefits of cross-functional workers?

Markov Decision Process

- **Exogenous state:** job and worker availability vectors x
- **Endogenous state:** job-worker familiarity matrix F

Markov Decision Process

- **Exogenous state:** job and worker availability vectors x
- **Endogenous state:** job-worker familiarity matrix F
- **Action:** job-worker assignment matrix A
- **Action space:** availability-dependent feasibility constraints $\mathcal{A}(x)$

Markov Decision Process

- **Exogenous state:** job and worker availability vectors x
- **Endogenous state:** job-worker familiarity matrix F
- **Action:** job-worker assignment matrix A
- **Action space:** availability-dependent feasibility constraints $\mathcal{A}(x)$
- **Cost function:** linear cost function $c(F, A)$
- **Assignment policy:** $\pi : \mathcal{X} \times \mathcal{F} \mapsto \mathcal{A}(x)$

$$\min_{\pi:\, \mathcal{X}\times\mathcal{F}\mapsto\mathcal{A}(x)} \ \mathrm{Cost}(\pi) := \ \mathbb{E}^{\pi}\left[\sum_{t=0}^{\infty} \alpha^{t}\, c\big(F^{t,\pi}, A^{\pi}\left(x^{t}, F^{t,\pi}\right)\big)\right]$$

Markov Decision Process

- **Exogenous state:** job and worker availability vectors x
- **Endogenous state:** job-worker familiarity matrix F
- **Action:** job-worker assignment matrix A
- **Action space:** availability-dependent feasibility constraints $\mathcal{A}(x)$
- **Cost function:** linear cost function $c(F, A)$
- **Assignment policy:** $\pi : \mathcal{X} \times \mathcal{F} \mapsto \mathcal{A}(x)$

Our MDP is quite general:

- Availability can evolve according to any exogenous Markov chain
- Learning rates can be heterogeneous across jobs and workers
- Action space can include other operational constraints

Markov Decision Process

- **Exogenous state:** job and worker availability vectors x
- **Endogenous state:** job-worker familiarity matrix F
- **Action:** job-worker assignment matrix A
- **Action space:** availability-dependent feasibility constraints $\mathcal{A}(x)$
- **Cost function:** linear cost function $c(F, A)$
- **Assignment policy:** $\pi : \mathcal{X} \times \mathcal{F} \mapsto \mathcal{A}(x)$

Our MDP is particularly challenging to tackle:

- High-dimensional exogenous and endogenous states
- High-dimensional and state-dependent action space
- High-dimensional expectations

Markov Decision Process

- **Exogenous state:** job and worker availability vectors x
- **Endogenous state:** job-worker familiarity matrix F
- **Action:** job-worker assignment matrix A
- **Action space:** availability-dependent feasibility constraints $\mathcal{A}(x)$
- **Cost function:** linear cost function $c(F, A)$
- **Assignment policy:** $\pi : \mathcal{X} \times \mathcal{F} \mapsto \mathcal{A}(x)$

How does the MDP formulation enable studying the specialization vs cross-functionality trade-off?

Specialization vs Cross-Functionality Trade-off

- Limiting familiarity levels under policy π:

$$\bar{f}_{jk}^{\pi} := \mathbb{E}^{\pi}\left[\lim_{t\to\infty} \boldsymbol{f}_{jk}^{t,\pi}\right]$$

Specialization vs Cross-Functionality Trade-off

- Limiting familiarity levels under policy π:

$$\bar{f}_{jk}^{\pi} := \mathbb{E}^{\pi}\left[\lim_{t\to\infty} \boldsymbol{f}_{jk}^{t,\pi}\right]$$

- An illustrative example

Policy 1

	Worker 1	Worker 2	Worker 3	Worker 4	Worker 5
Job 1					0.9
Job 2			0.8		
Job 3		0.8			

Policy 1 favors specialization

Specialization vs Cross-Functionality Trade-off

- Limiting familiarity levels under policy π:

$$\bar{f}_{jk}^{\pi} := \mathbb{E}^{\pi}\left[\lim_{t\to\infty} \boldsymbol{f}_{jk}^{t,\pi}\right]$$

- An illustrative example

Policy 1

	Worker 1	Worker 2	Worker 3	Worker 4	Worker 5
Job 1					0.9
Job 2			0.8		
Job 3		0.8			

Policy 2

	Worker 1	Worker 2	Worker 3	Worker 4	Worker 5
Job 1		0.3			0.3
Job 2		0.2	0.3		0.3
Job 3			0.2		

Policy 1 favors specialization

Policy 2 favors cross-functionality

Main Contributions

Designed Policy

- Familiarity agnostic (FA) policy
- Lagrangian relaxation (LR) policy

Main Contributions

Designed Policy

- Familiarity agnostic (FA) policy
- Lagrangian relaxation (LR) policy

Uncovered Insight

- FA policy → Favors building cross-functional workers
- LR policy → Favors building specialized workers

Main Contributions

Designed Policy	• Familiarity agnostic (FA) policy • Lagrangian relaxation (LR) policy
Uncovered Insight	• FA policy → Favors building cross-functional workers • LR policy → Favors building specialized workers
Developed Performance Bound	• FA policy → Near-optimal if learning happens slowly • LR policy → Near-optimal if availability minimally impacts assignment feasibility

Main Contributions

Designed Policy	• Familiarity agnostic (FA) policy • Lagrangian relaxation (LR) policy
Uncovered Insight	• FA policy → Favors building cross-functional workers • LR policy → Favors building specialized workers
Developed Performance Bound	• FA policy → Near-optimal if learning happens slowly • LR policy → Near-optimal if availability minimally impacts assignment feasibility
Developed Algorithm	• Math-programming-based method to compute both policies.

Related Work

Dynamic Assignment / Matching × **Modeling Familiarity / Experience**

Dynamic Assignment / Matching: Caldentey et al. 2009, Bertsimas et al. 2013, Golrezaei et al. 2014, Gurvich and Ward 2015, Ma and Simchi-Levi 2020, Johari et al. 2021, Afèche et al. 2022, Aouad and Sarıtac 2022, Delong et al. 2024, Bansak et al. 2024, …

Modeling Familiarity / Experience: Gans and Zhou 2003, Reagans et al. 2005, Avgerinos and Gokpinar 2017, Ramdas et al. 2018, Zhang et al. 2024, Tuzcuoglu et al. 2024, …

- **Quasi-open-loop policies** (Adelman and Mancini 2016, Bai et al. 2023)
- **Inventory-agnostic policies** (Bai et al. 2022)
- **Fast-slow moving MDPs** (Wang and Jiang 2023)
- **Weakly-coupled MDPs** (Hawkins 2003, Adelman and Mersereau 2008, Brown and Smith 2020, Torrico and Toriello 2022, Nadarajah and Cire 2024, …)
- **Math-programming-based ADP** (De Farias and Van Roy 2003, Desai et al. 2012, Zhang and Adelman 2009, Tong and Topaloglu 2013, Nadarajah et al. 2015, Pakiman et al. 2024, …)

Familiarity-Agnostic Policy

Construction

$$\pi : \mathcal{X} \times \mathcal{F} \mapsto \mathcal{A}(x) \longrightarrow \hat{\pi} : \mathcal{X} \mapsto \mathcal{A}(x)$$

Familiarity-Agnostic Policy

Construction

$$\pi : \mathcal{X} \times \mathcal{F} \mapsto \mathcal{A}(x) \longrightarrow \hat{\pi} : \mathcal{X} \mapsto \mathcal{A}(x)$$

Insights (Proposition 1)

$\min_{\hat{\pi}} \mathrm{Cost}(\hat{\pi}) \equiv$ Optimizing familiarity levels and assignment probabilities in the steady state

Enables the development of a cross-functional workforce

Familiarity-Agnostic Policy

Construction

$$\pi : \mathcal{X} \times \mathcal{F} \mapsto \mathcal{A}(x) \longrightarrow \hat{\pi} : \mathcal{X} \mapsto \mathcal{A}(x)$$

Insights (Proposition 1)

$\min_{\hat{\pi}} \text{Cost}(\hat{\pi})$ $\equiv$ Optimizing familiarity levels and assignment probabilities in the steady state

Enables the development of a cross-functional workforce

Computation (Theorem 2)

$$\min_{\hat{\pi}} \text{Cost}(\hat{\pi}) \equiv \text{ a stochastic quadratic program}$$

Familiarity-Agnostic Policy

Construction

$$\pi : \mathcal{X} \times \mathcal{F} \mapsto \mathcal{A}(x) \longrightarrow \hat{\pi} : \mathcal{X} \mapsto \mathcal{A}(x)$$

Insights (Proposition 1)

$\min_{\hat{\pi}} \text{Cost}(\hat{\pi}) \equiv$ Optimizing familiarity levels and assignment probabilities in the steady state

Enables the development of a cross-functional workforce

Computation (Theorem 2)

$$\min_{\hat{\pi}} \text{Cost}(\hat{\pi}) \equiv \text{a stochastic quadratic program}$$

Performance Bound (Theorem 1)

FA policy is near-optimal if learning and forgetting happen slowly

Lagrangian Relaxation Policy

Construction

$$\mathcal{A}(x) := \left\{ A \in \{0,1\}^{J \times K} \;\middle|\; J + K \text{combinatorial constraints} \right\} \quad \lambda = (\lambda^{\mathrm{W}}, \lambda^{\mathrm{J}})$$

Lagrangian Relaxation Policy

Construction

$$\mathcal{A}(x) := \left\{A \in \{0,1\}^{J\times K} \;\middle|\; J + K \text{combinatorial constraints}\right\} \quad \lambda = (\lambda^{\mathrm{W}}, \lambda^{\mathrm{J}})$$

Insights (Theorem 3)

$$\bar{\pi}^{\lambda}_{jk}(f_{jk}) = \begin{cases} \text{Assign} & \text{if } f_{jk} \geq 1 - \frac{(1-\alpha(1-\ell_{jk}))}{(1-\alpha)c^1_{jk}}\left[\lambda^{\mathrm{K}}_k + \lambda^{\mathrm{J}}_j + c^1_{jk} - c^0_{jk}\right] \\ \text{Not Assign} & \text{otherwise} \end{cases}$$

The LR approach enables the development of specialized workers

Lagrangian Relaxation Policy

School of Management

Construction

$$\mathcal{A}(x) := \{A \in \{0,1\}^{J\times K} \mid J + K \text{combinatorial constraints}\} \quad \lambda = (\lambda^{\mathrm{W}}, \lambda^{\mathrm{J}})$$

Insights (Theorem 3)

$$\bar{\pi}^{\lambda}_{jk}(f_{jk}) = \begin{cases} \text{Assign} & \text{if } f_{jk} \geq 1 - \frac{(1-\alpha(1-\ell_{jk}))}{(1-\alpha)c^1_{jk}}\left[\lambda^{\mathrm{K}}_k + \lambda^{\mathrm{J}}_j + c^1_{jk} - c^0_{jk}\right] \\ \text{Not Assign} & \text{otherwise} \end{cases}$$

The LR approach enables the development of specialized workers

Computation (Theorem 3)

The LR policy can be obtained using binary programs.

Lagrangian Relaxation Policy

Construction

$$\mathcal{A}(x) := \left\{A \in \{0,1\}^{J\times K} \;\middle|\; J + K \text{combinatorial constraints}\right\} \qquad \lambda = (\lambda^{\mathrm{W}}, \lambda^{\mathrm{J}})$$

Insights (Theorem 3)

$$\bar{\pi}^{\lambda}_{jk}(f_{jk}) = \begin{cases} \text{Assign} & \text{if } f_{jk} \geq 1 - \frac{(1-\alpha(1-\ell_{jk}))}{(1-\alpha)c^1_{jk}}\left[\lambda^{\mathrm{K}}_k + \lambda^{\mathrm{J}}_j + c^1_{jk} - c^0_{jk}\right] \\ \text{Not Assign} & \text{otherwise} \end{cases}$$

The LR approach enables the development of specialized workers

Computation (Theorem 3)

The LR policy can be obtained using binary programs.

Performance Bound (Proposition 3)

LR policy is near optimal if assignment feasibility constraints change minimally in x

Numerical Experiments

- Total number of instances: 120
- Numbers of jobs and workers: $(J, K) \in \{(2,5), (2,8), (3,5), (3,6), (4,5)\}$
- Worker-job availability: deterministic, low-uncertainty, high-uncertainty
- Learning rates: zero, half, one
- Proposed policies: FA, LR, and ALP policies

Policy Comparison

Policy gap = *% (Cost of Policy - Cost of Best Proposed Policy) / (Cost of Best Proposed Policy)*

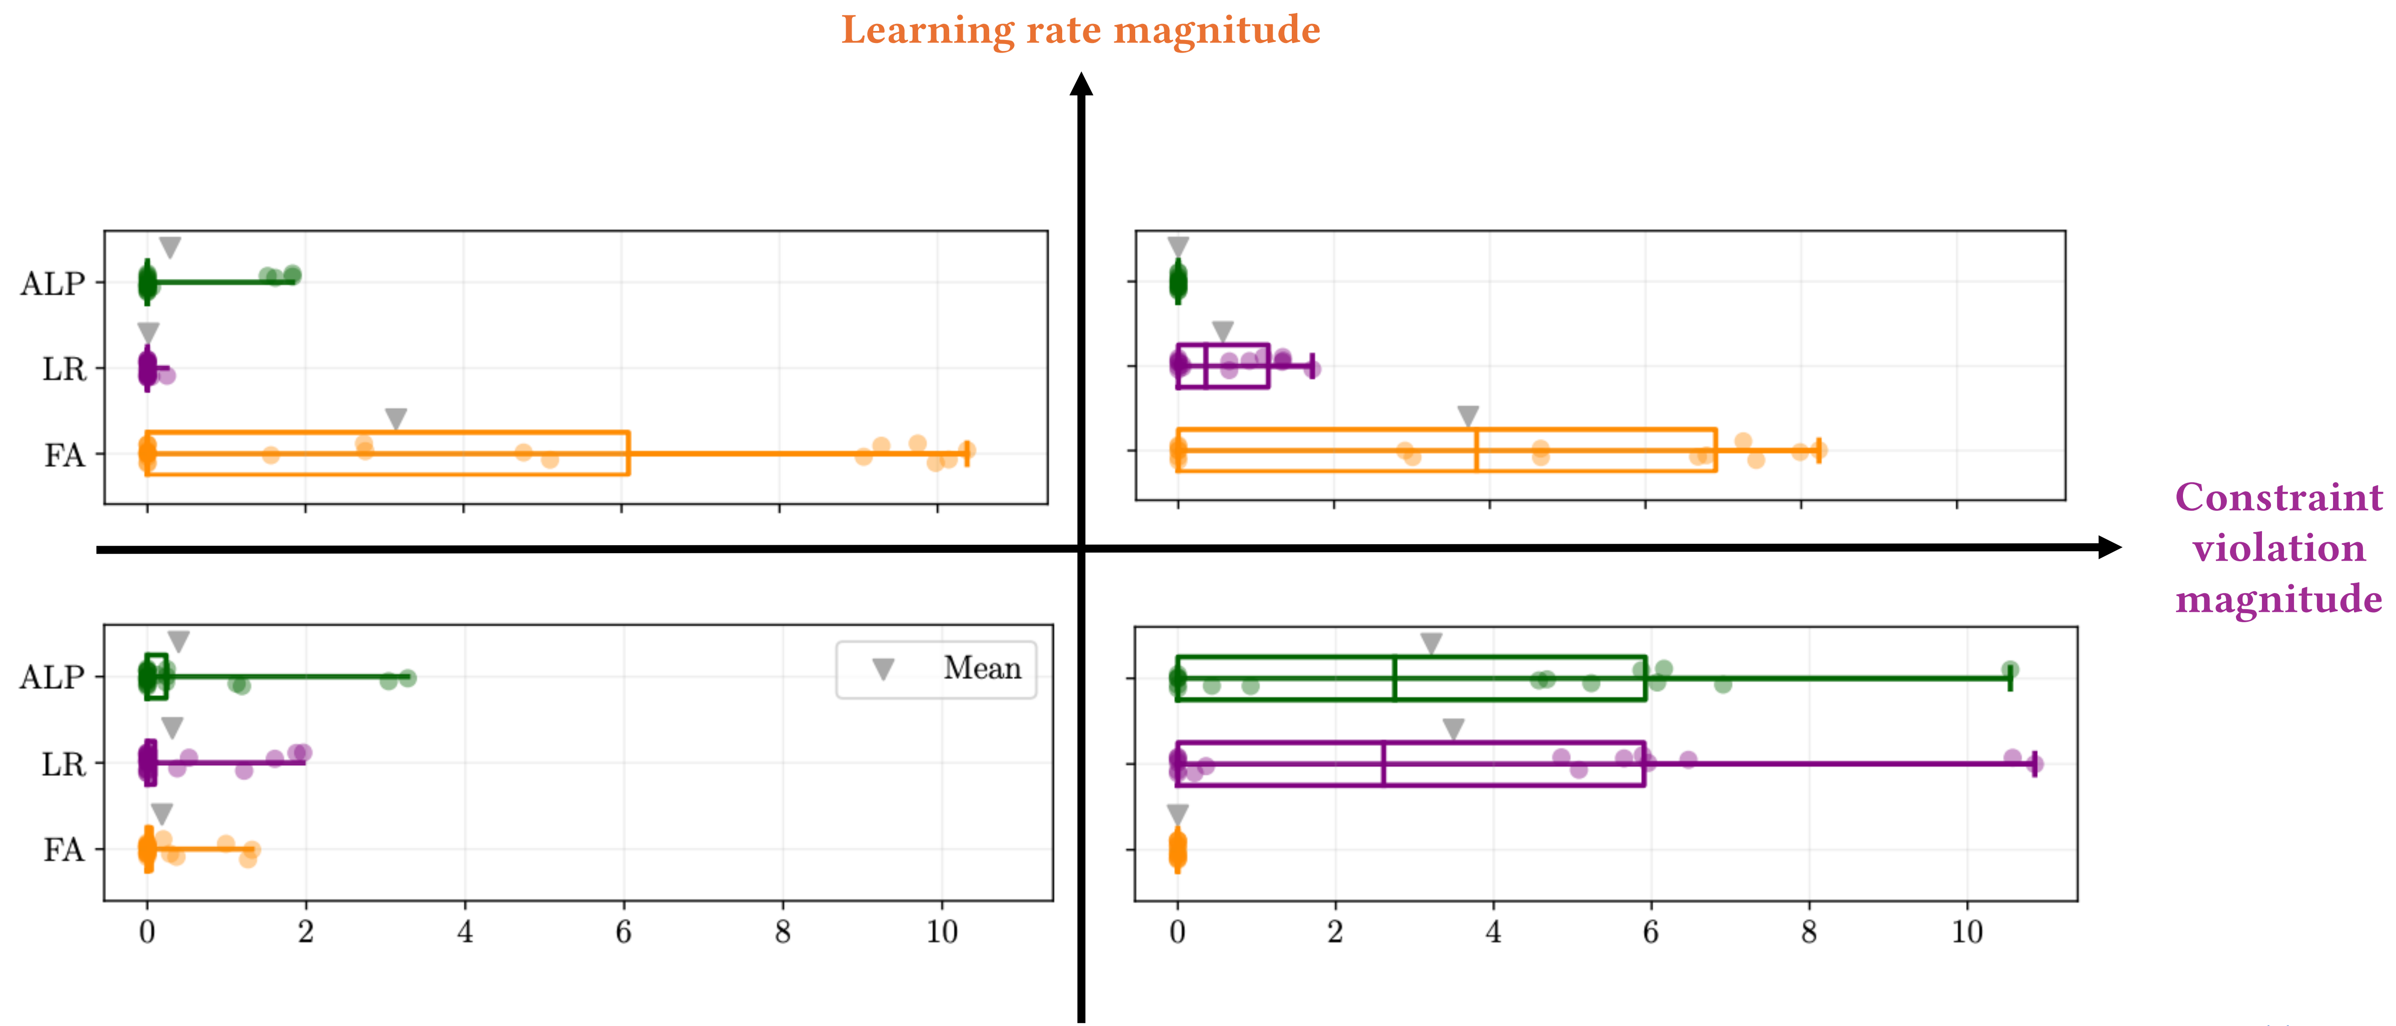

Specialization vs Cross-Functionality

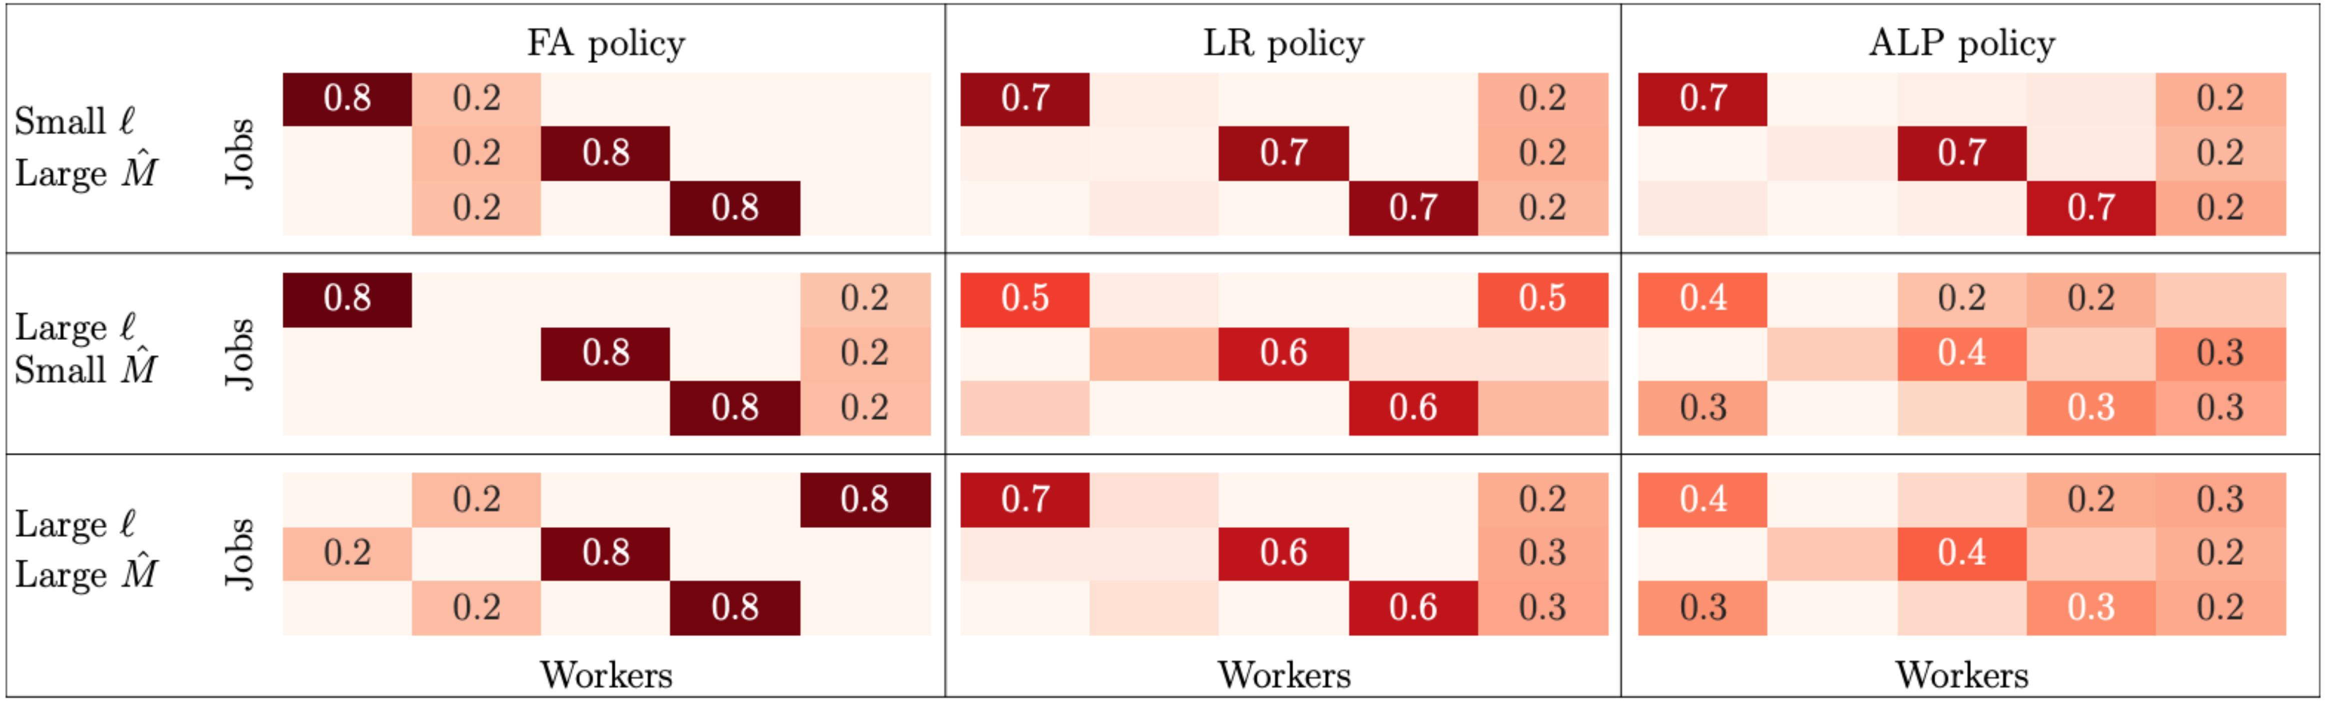

Key Findings

- **Low impact of availability:**
  - LR policy outperforms FA policy and is near-optimal
  - LR policy delivers strong portfolios of specialized workers

Key Findings

School of Management

- **Low impact of availability:**
  - LR policy outperforms FA policy and is near-optimal
  - LR policy delivers strong portfolios of specialized workers

- **High impact of availability** & **low learning/forgetting rates:**
  - FA policy outperforms LR policy and is near-optimal
  - FA policy delivers strong portfolios of cross-functional workers

Key Findings

- **Low impact of availability:**
  - LR policy outperforms FA policy and is near-optimal
  - LR policy delivers strong portfolios of specialized workers

- **High impact of availability** & **low learning/forgetting rates:**
  - FA policy outperforms LR policy and is near-optimal
  - FA policy delivers strong portfolios of cross-functional workers

- **Note:** We have developed an approximate linear programming (ALP) policy that performs well on instances with high availability impact and high learning/forgetting rates.

Broader Takeaway

Our research enables the study of complex labor–management dynamics and provides a framework for dynamically optimizing talent across the workforce.

Q&A

Sequential diagnostic testing with patient refusal dynamics

A decision-making framework with correlation + mixed test outcomes + refusal dynamics

Shakiba Rahnama • Joint work with Negar Soheili & Ludwig Dierks

Information Decision Sciences — University of Illinois Chicago

Diagnosis is like "20 Questions" in a clinic

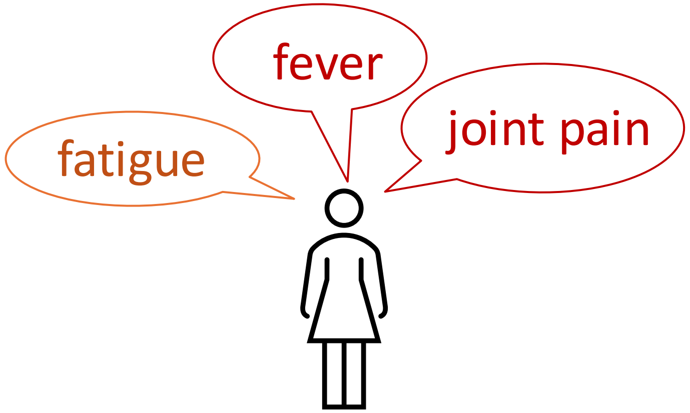

Symptoms

A patient comes to the clinic with symptoms that could match several diseases and they might even have more than one at the same time.

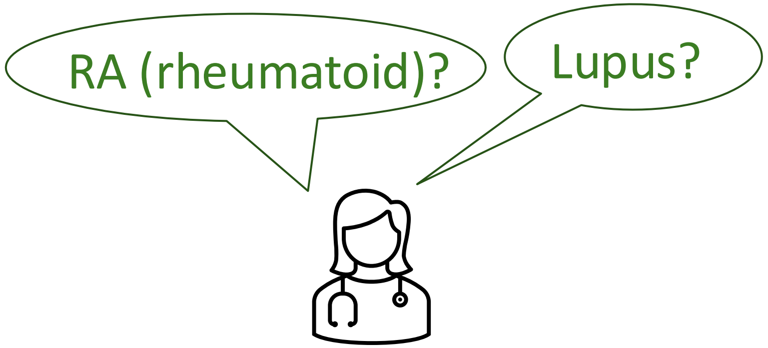

Diagnosis

The provider tries to identify the underlying disease based on the patient's symptoms.

Possible tests?

- ESR
- CRP
- Anti-CCP

Tests can be correlated

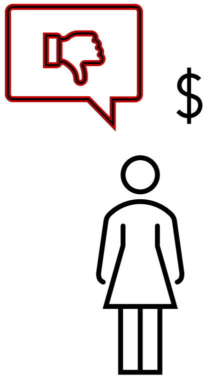

Patient can refuse doing the test

Goal: learn enough before the patient opts out

Why is sequential diagnosis hard in the real world?

Multiple diseases

- Multiple diseases can co-exist, so it's not "disease vs no disease," it's among many combinations.
- The number of possible disease patterns grows quickly (about 2^M for M diseases).

Correlated results

- Tests are correlated because a hidden patient factor (severity/inflammation) shifts many results together.
- Ignoring correlation overestimates certainty.

Patient Refusal

- Patients may refuse tests as the burden accumulates.
- The plan must anticipate dropout, not assume completion.

We need a sequential decision model that can update beliefs quickly, while handling **multi-disease uncertainty, correlated tests, and patient refusal.**

A gap in the literature

World A: OR / decision models

Model diagnosis testing as a sequential decision problem (MDP/POMDP): update beliefs after each test and decide “test vs stop”.

Steimle and Denton (2015), Alagoz et al. (2015), Zhang et al., (2022).

- Outcomes are simplified for tractability.
- Tests are assumed to be independent.
- Patient dropout is typically not modeled.

World B: diagnostic statistics

Study how to combine multiple tests (cutoffs, series/parallel rules) and evaluate the accuracy of a fixed pathway, often allowing dependence.

Leeflang et al. (2008), Fanshawe et al. (2024), and Böttcher & Felder (2025).

- Pathway is usually not optimized sequentially.
- Often assumes a single outcome type/scale.
- Refusal is typically not part of the decision.

We model sequential testing while accounting for patient dropout, correlated tests, and multiple diseases.

Dependence model: one latent physiology drives many tests

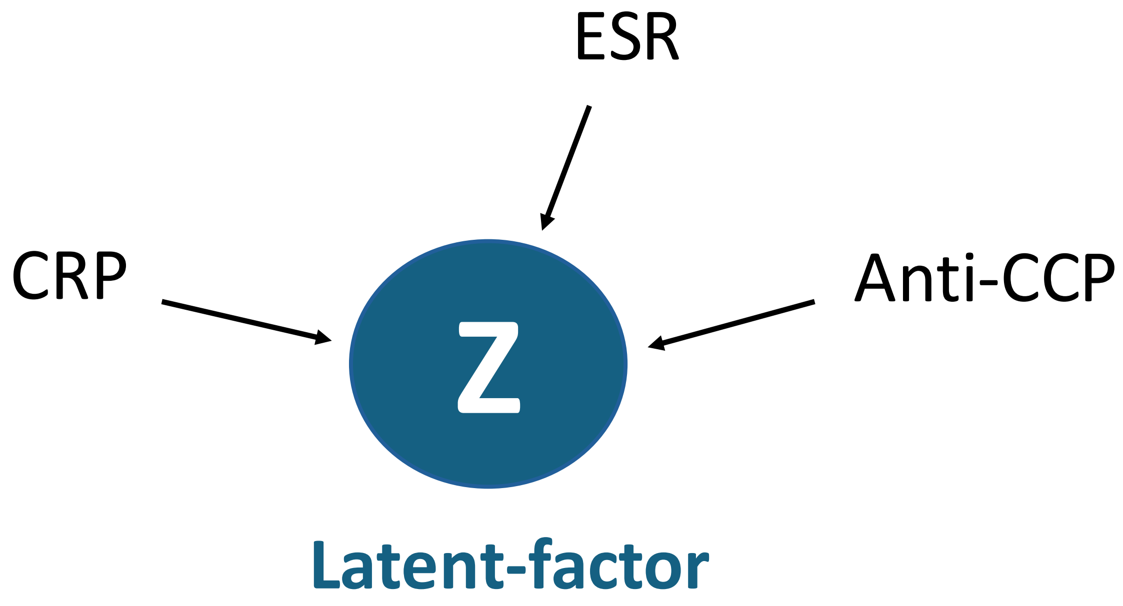

- We capture within-patient correlation using a single latent factor Z.
- Given the true disease pattern and Z, the remaining test noise is independent - dependence comes from unobserved Z.

Model: Sequential diagnosis as a POMDP

(S, A, O, T, C, γ)- decide "test vs stop" under partial information

States

D = disease pattern → Un-observed

Z = latent factor → Un-observed

C_k = cumulative cost → Observed

Belief

Belief $b_k(d)$ over disease patterns

$Z \mid \text{History}: \mathcal{N}(\mu_k, \sigma_k^2)$

Action

Offer a test t_k

Or stop and diagnose

Observation

Patient: refuse with prob $P_{t_k}(C_k) \in [0,1]$

If accept → result (binary or numeric)

Belief update and tractability

- Update disease-pattern weights $b_k(d)$ by **Bayes' rule** using disease pattern likelihood of observation $\ell_{d,k}(x_k)$:

$$b_{k+1}(d) = \frac{b_k(d), \ell_{d,k}(x_k)}{\sum_{d'} b_k(d'), \ell_{d',k}(x_k)}$$

- Update a Gaussian summary for the latent correlation factor Z, i.e., $Z \approx \mathcal{N}(\mu_k, \sigma_k^2)$.

Issue

Correlation makes exact inference blow up: After a new observation, the true posterior couples $\boldsymbol{D}$ and $\boldsymbol{Z}$, creating a mixture over all 2^M disease patterns.

Our fix

Step 1: Update Z within disease patterns that impact test outcomes:

- Numeric result: Z stays Gaussian (**Kalman-style** update)
- Binary result: Z is non-Gaussian → approximate with a Gaussian (**Laplace**)

Step 2: Collapse the mixture **back to one** Gaussian

Match mean and variance to keep $Z \mid \text{History} \approx \mathcal{N}(\mu_k, \sigma_k^2)$.

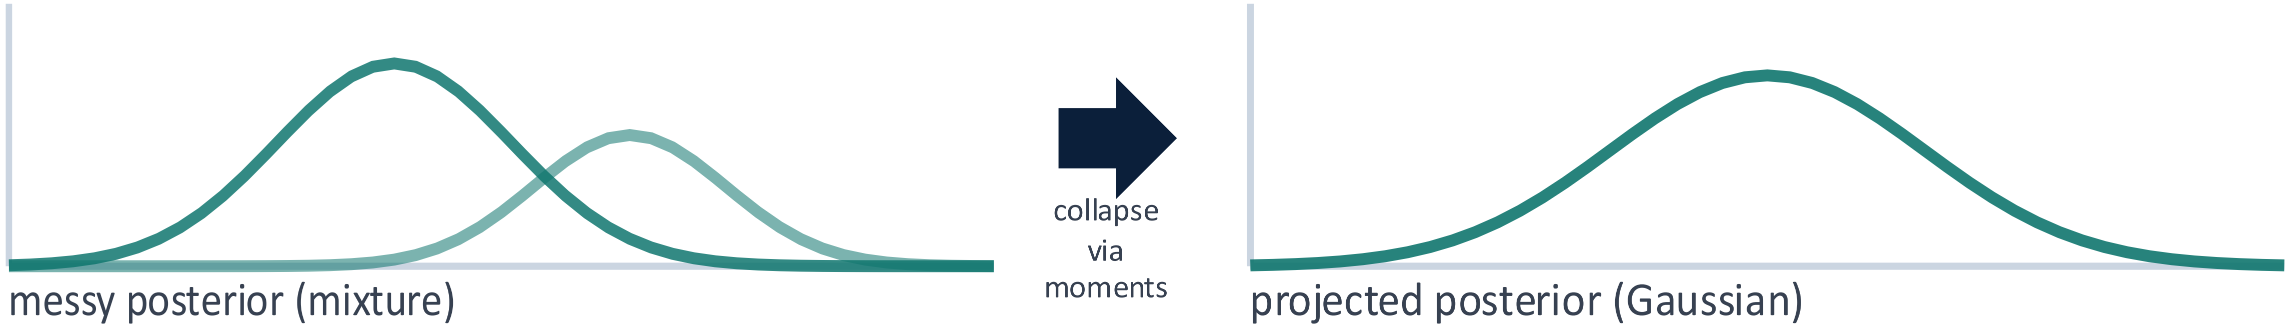

Decision rule: stop now vs offer another test

Stop: misdiagnosis risk

- $\boldsymbol{c_i}^-, \boldsymbol{c_i}^+ > 0$: Cost of a false negative/false positive.

$$Q^{stop}(B_k) := \sum_{i=1}^{M} \min \{ \underbrace{\mathrm{p_{i,k}} \boldsymbol{c_i}^-}_{\text{Expected cost of FN}},\ \underbrace{(1 - \mathrm{p_{i,k}}) \boldsymbol{c_i}^+}_{\text{Expected cost of FP}} \}$$

Where $\mathrm{p_{i,k}} = \Pr(D_i = 1 \mid H_k)$ is the current probability that disease i is present, given all test results so far.

Test: expected cost-to-go

- $\boldsymbol{c_{t_k}} \geq 0$: monetary cost of test t_k.

$$Q_k^{t_k}(\mathrm{B}_k, C_k) = P_{t_k} . Q^{stop}(B_k) + (1 - P_{t_k}) c_{t_k}$$
$$+ (1 - P_{t_k}) . (E \ \ [V_{k+1}(B_{k+1}, C_{k+1}) \mid \mathrm{B}_k, a_k = t_k])$$

The expectation is over the predicted test outcome under belief B_k=$(b_k(d), (\mu_k, \sigma_k^2))$; if the test is accepted, the cumulative cost updates as $\boldsymbol{C_{k+1}} = \boldsymbol{C_k} + \boldsymbol{c_{t_k}}$

$$V_k(B_k, C_k) = min \ \{Q^{stop}(B_k), \min_{t_k \in \mathcal{T}_k} Q_k^{t_k}(B_k, C_k)\},$$

Analytical insight from a baseline

Simple model: One disease, tests are binary (positive or negative), and independent

Goal: Finding the structure of optimal policy

$$\Delta(b, C) > 0 \iff Q^{stop}(b) - E[Q^{stop}(b^o)] > c_j + \frac{p_{t_k}}{1 - p_{t_k}}(bc^- - Q^{stop}(b))$$

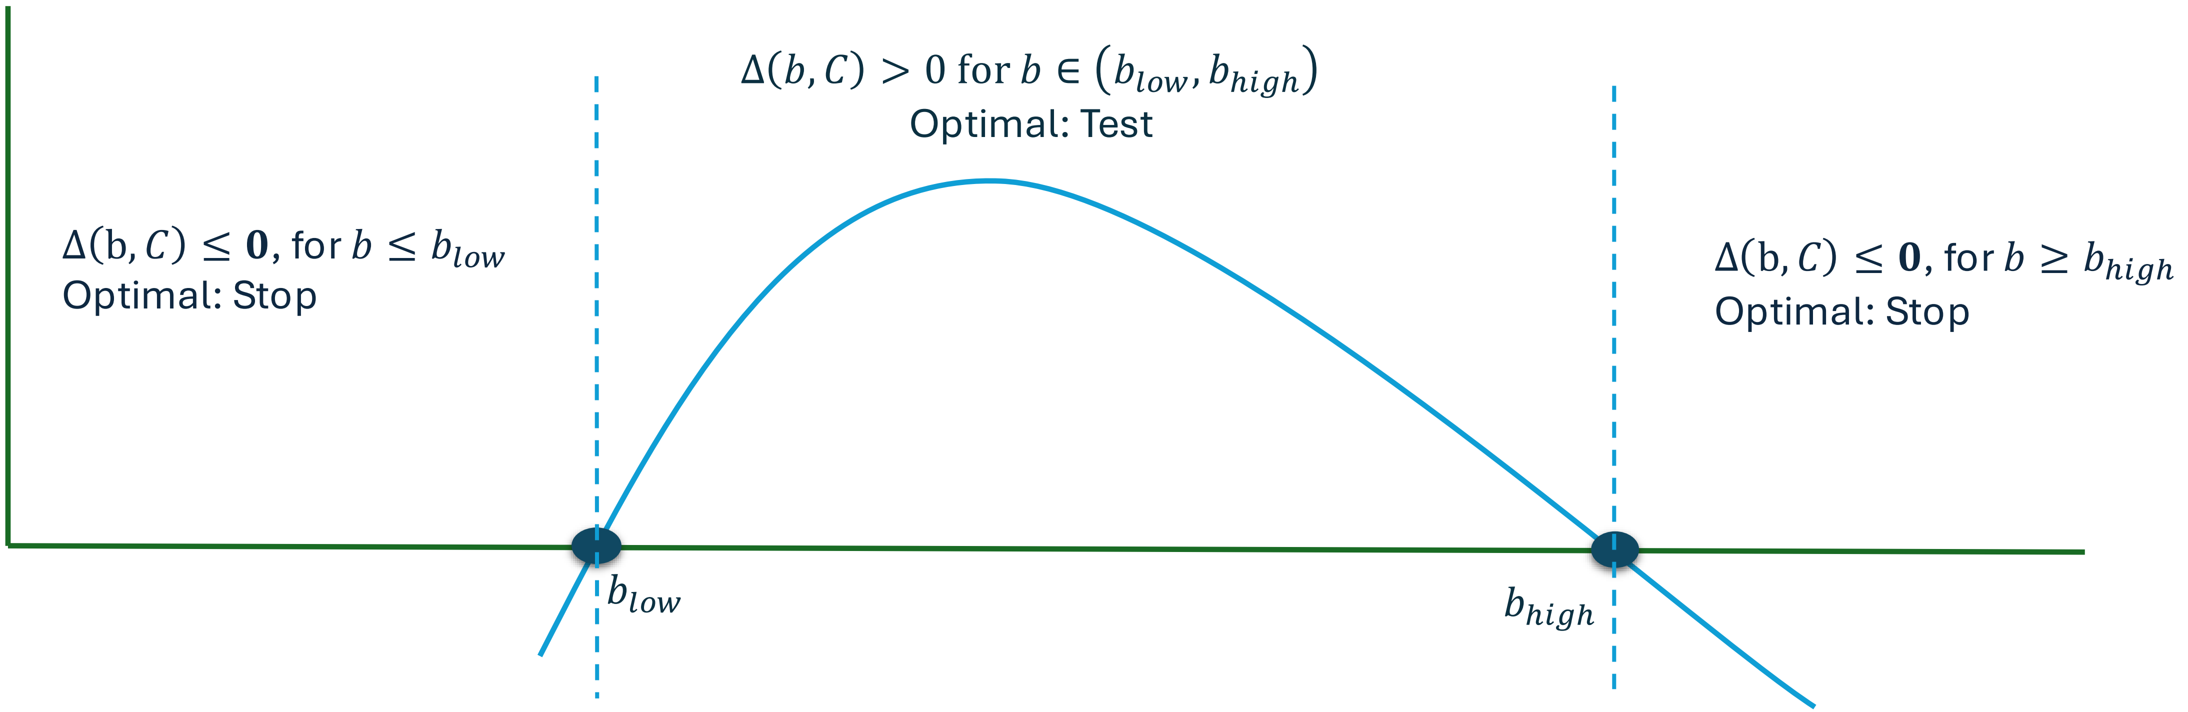

Analytical insight from a baseline

$$Q^{stop}(b) - E[Q^{stop}(b^o)] > c_j + \frac{p_{t_k}}{1 - p_{t_k}}(bc^- - Q^{stop}(b))$$

Information gain = how much the test reduces expected misdiagnosis cost on average

Information gain

Cost of stopping now

Expected cost of stopping after seeing the test result

Cost of the test

Effective Cost

The expected cost from being forced to stop early because of refusal

- When $p_j = 0$, then test iff information gain worth its price (> c_j)
- When $p_j > 0$, then test iff information gain worth its price and the risk of losing the chance to learn
- Testing region shrinks as refusal risk increases
- Some tests disappear entirely (even if informative)
- Information is front loaded: Test earlier, stop sooner

Thank you!

srahna3@uic.edu

Q&A

NYU Langone Health

CSTAR

Center for Surgical and Transplant Applied Research

"EBAY" FOR HIGH-KDPI KIDNEYS:
SIMULTANEOUS OFFERS TO DECREASE NONUSE

Meghan E. Meredith[1], Bonnie E. Lonze[1], Dorry L. Segev[1], Nicholas L. Wood[2], Allan B. Massie[1], Sommer E. Gentry[1]

[1]Department of Surgery, New York University Grossman School of Medicine, NYU Langone Health, New York, NY, [2]Hennepin Healthcare Research Institute, Minneapolis, MN

DECEASED DONOR KIDNEY TRANSPLANTATION

NYU Langone Health

Kidney transplantation is the gold standard treatment for end-stage kidney failure

- Kidneys eliminate toxins in blood and send waste out as urine
- Kidney failure can be caused by chronic conditions or acute injury, and leads to waste build up in the body
- There are two treatments for kidney failure:
  1. Dialysis: machine or fluid drip used to eliminate toxins
  2. Kidney transplant: surgical implantation of a donor kidney

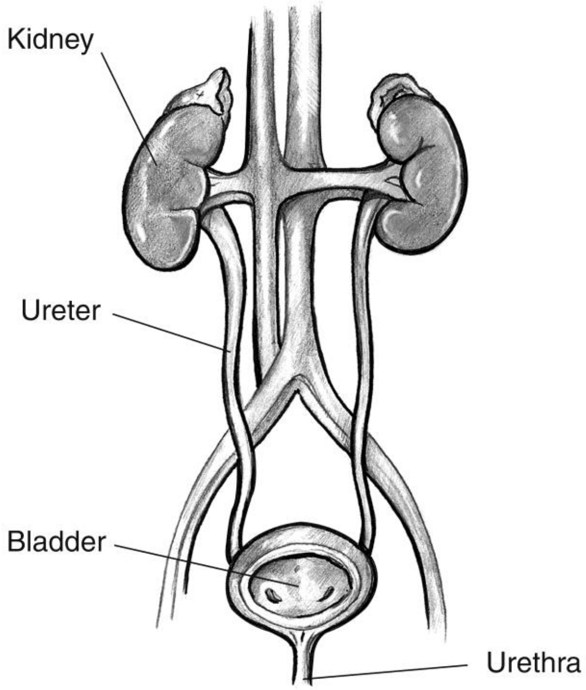

National Institute of Diabetes and Digestive and Kidney Diseases, National Institutes of Health.

NYU Langone Health

Donor kidneys for transplantation are a scarce resource

In the United States in 2023,

- 35,500,000 people with kidney disease
- 815,000 people with kidney failure
- 555,000 people on dialysis
- 142,000 people on the waiting list for a transplant (including 47,000 people added in 2023)
- 28,000 kidney transplants performed

National Institute of Diabetes and Digestive and Kidney Diseases, National Institutes of Health.
Lentine KL, Smith JM, Lyden GR, Miller JM, Booker SE, Dolan TG, Temple KR, Weiss S, Handarova D, Israni AK, Snyder JJ. OPTN/SRTR 2023 Annual Data Report: Kidney. Am J Transplant. 2025 Feb;25(2S1):S22-S137. doi: 10.1016/j.ajt.2025.01.020. PMID: 39947805; PMCID: PMC12414513.

Donor kidneys for transplantation are a scarce resource

In the United States in 2023,

35,500,000 people with kidney disease

815,000 people with kidney failure

555,000 people on dialysis

142,000 people on the waiting list for a transplant
(including 47,000 people added in 2023)

28,000 kidney transplants performed

4,000 people died on the waitlist

8,600 (28%) recovered deceased donor kidneys were not used

National Institute of Diabetes and Digestive and Kidney Diseases, National Institutes of Health.
Lentine KL, Smith JM, Lyden GR, Miller JM, Booker SE, Dolan TG, Temple KR, Weiss S, Handarova D, Israni AK, Snyder JJ. OPTN/SRTR 2023 Annual Data Report: Kidney. Am J Transplant. 2025 Feb;25(2S1):S22-S137. doi: 10.1016/j.ajt.2025.01.020. PMID: 39947805; PMCID: PMC12414513.

Kidneys are often not used because they are not placed efficiently enough

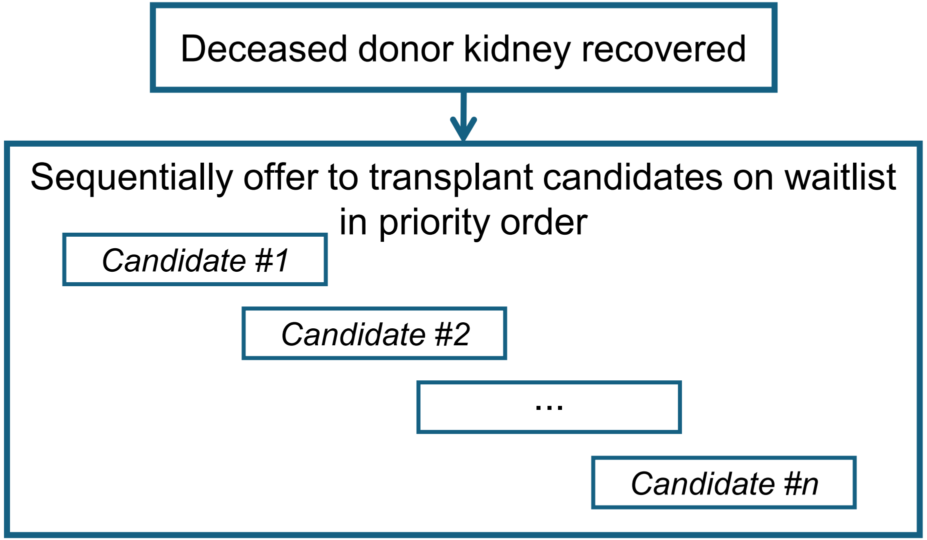

The priority order is based on:

- donor-candidate compatibility
- candidate wait time
- geography
- expected survival benefit

OPTN Policies Effective as of Dec 11, 2025

Kidneys are often not used because they are not placed efficiently enough

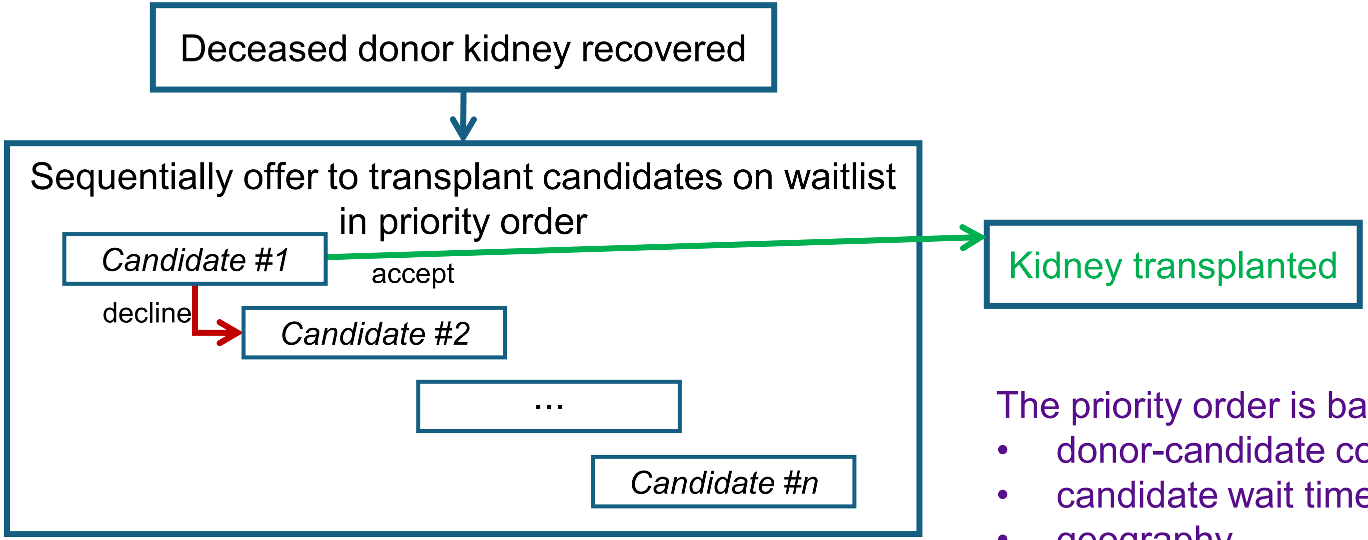

The priority order is based on:

- donor-candidate compatibility
- candidate wait time
- geography
- expected survival benefit

OPTN Policies Effective as of Dec 11, 2025

Kidneys are often not used because they are not placed efficiently enough

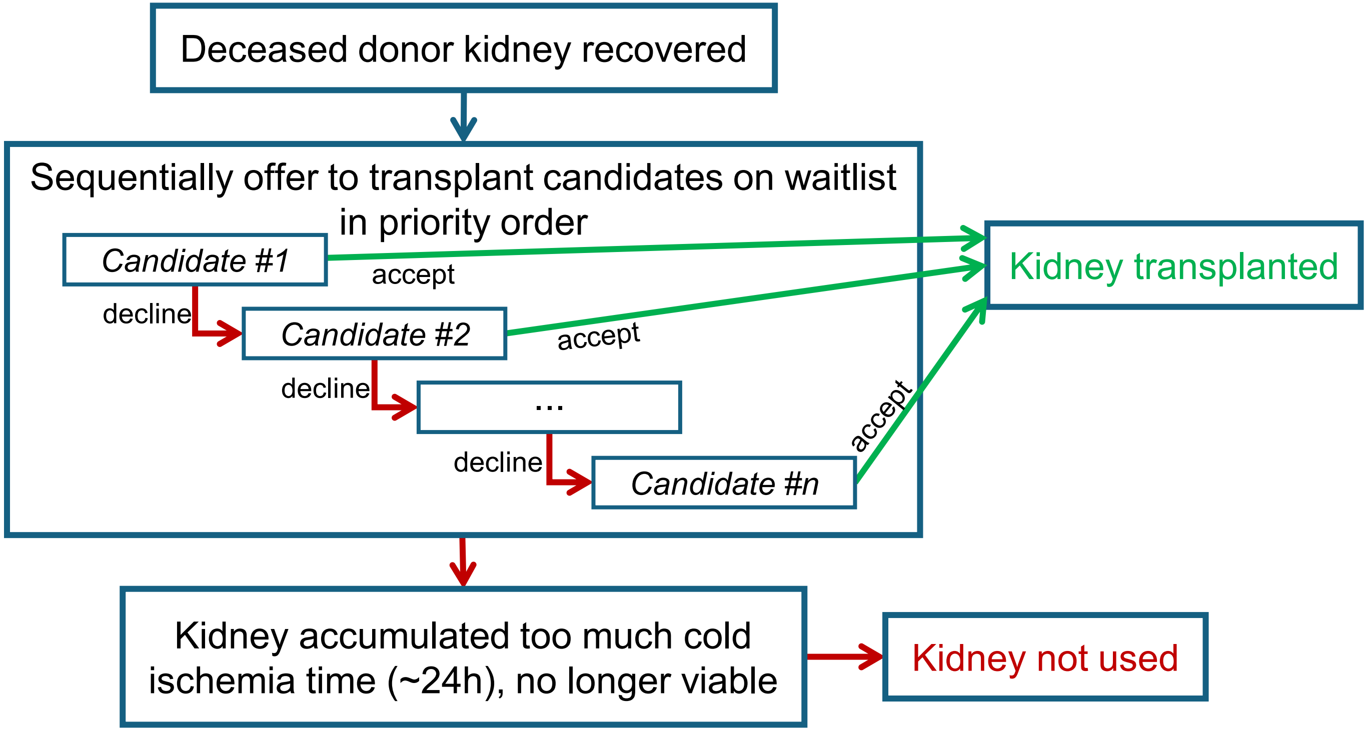

OPTN Policies Effective as of Dec 11, 2025

NYU Langone Health

Kidneys are often not used because they are not placed efficiently enough

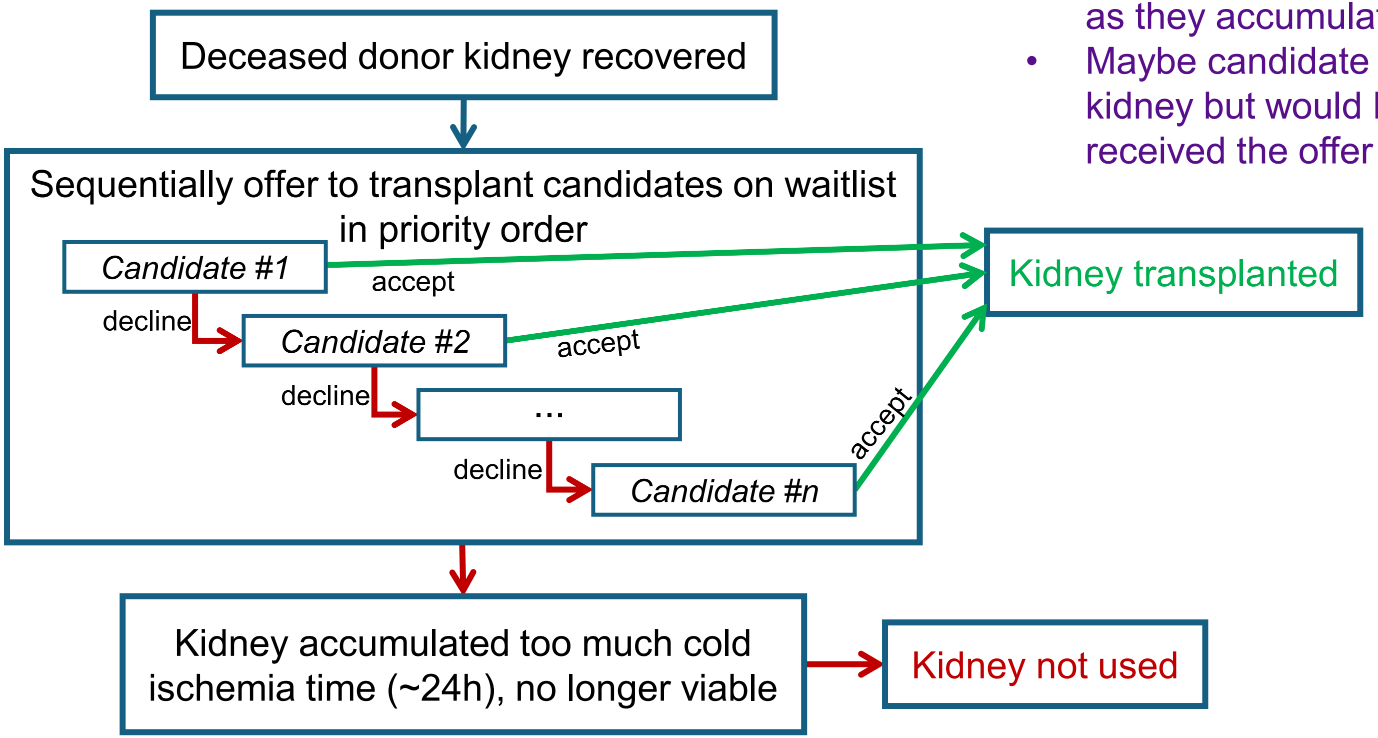

- Kidneys are less likely to be accepted as they accumulate cold ischemia time
- Maybe candidate n declined the kidney but would have accepted if they received the offer sooner?

OPTN Policies Effective as of Dec 11, 2025

AIM: PROPOSE AN ALLOCATION POLICY TO EFFICIENTLY PLACE HARD-TO-PLACE KIDNEYS TO DECREASE NONUSE

Aim: Design an allocation policy to more efficiently place hard-to-place kidneys to decrease nonuse

1. Hard-to-place: Kidneys with a high Kidney Donor Profile Index (KDPI) (> 75%)
   - KDPI is a numerical score (0-100%) characterizing a kidney's risk of transplant failure compared to other recovered kidneys
   - More than half of KDPI > 75% kidneys go unused
2. Alternative allocation policy for high-KDPI (hard-to-place) kidneys: Simultaneous (rather than sequential) offers within geographic thresholds

Simultaneous offers for high-KDPI kidneys

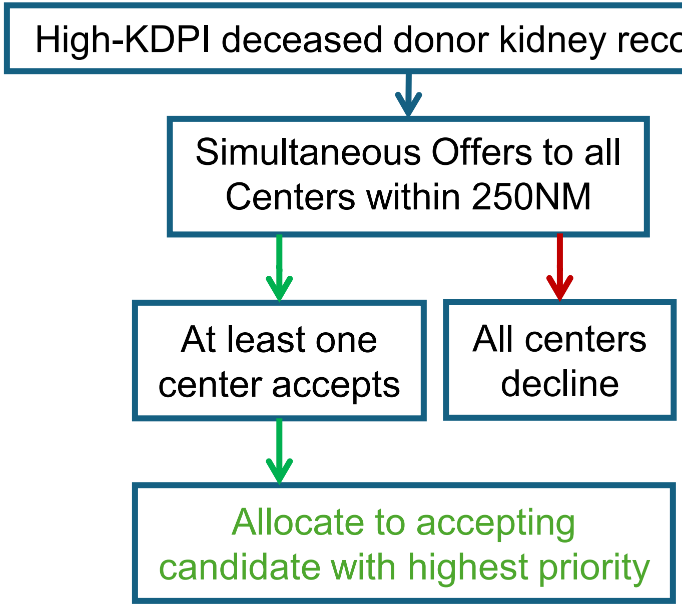

Centers are given 1-2 hours to respond

NYU Langone Health

Simultaneous offers for high-KDPI kidneys

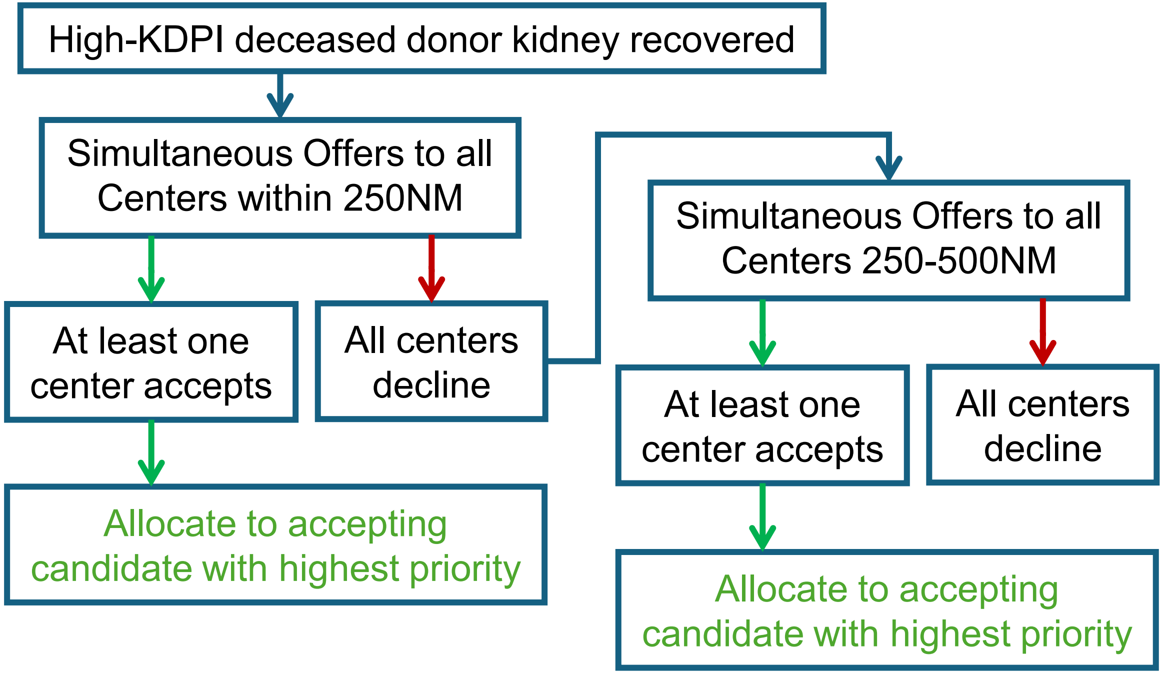

Simultaneous offers for high-KDPI kidneys

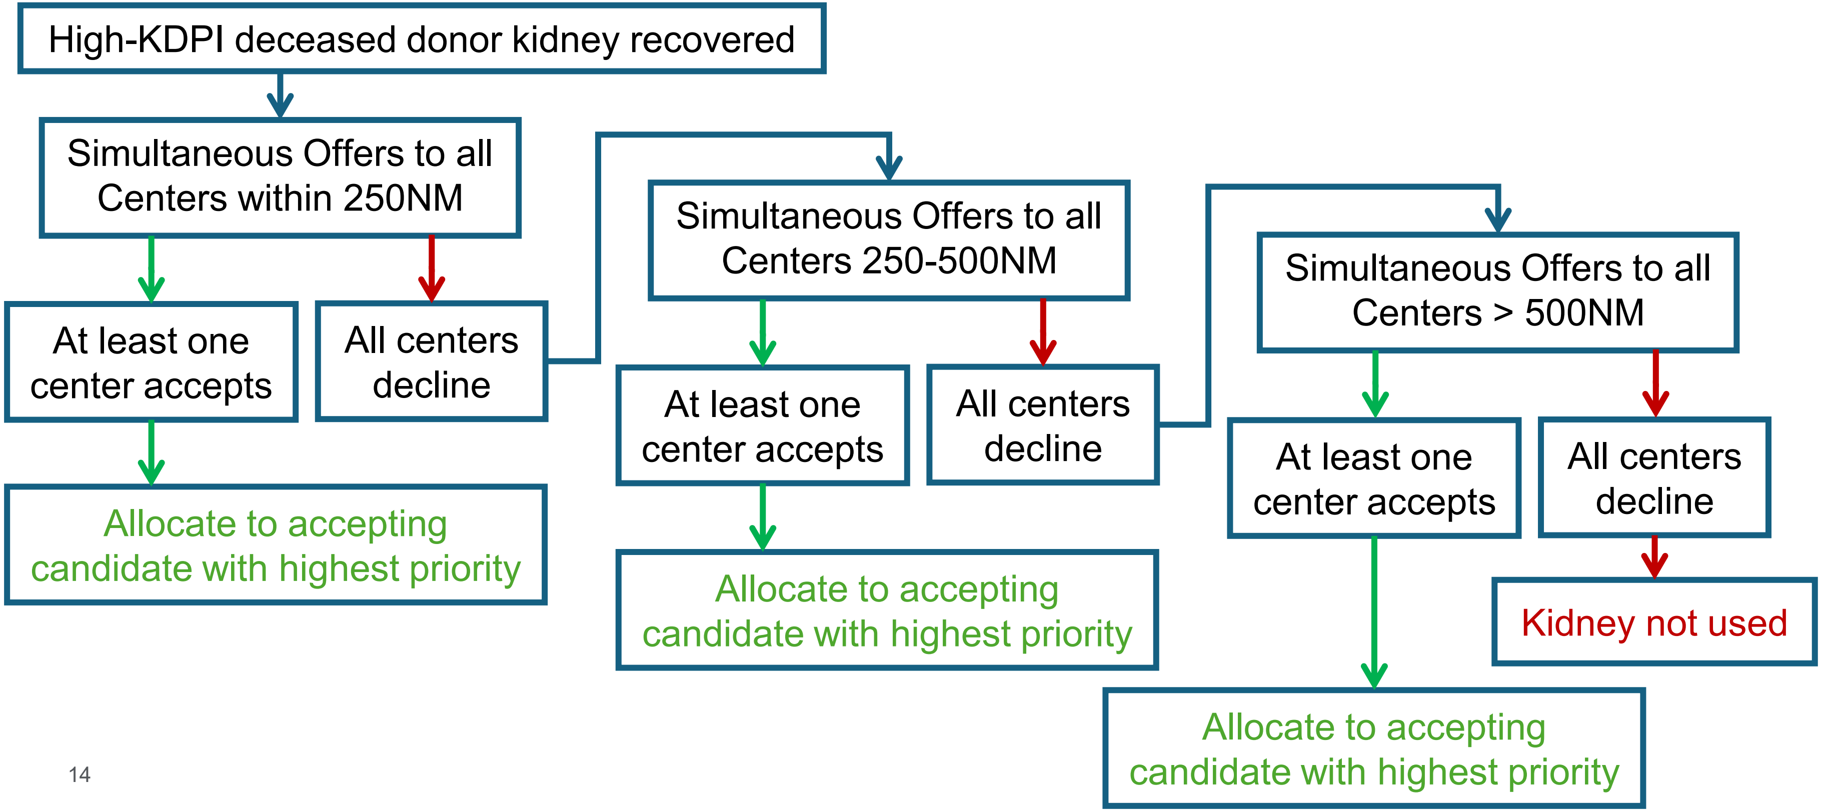

We test our proposed simultaneous offers policy with the Kidney-Pancreas Organ Allocation Simulator (OASim)

- Developed in 2024 by the Scientific Registry of Transplant Recipients (SRTR)
  - SRTR's mission: provide analysis to aid the DHHS in their oversight of the national organ transplantation system
- **Goal of OASim (2024):** simulate historic metrics of allocation policies (e.g., logistical burden of evaluating offers) AND organ utilization
  - Trained and tested on data from all recovered kidneys, regardless of utilization

Default logic of the Kidney-Pancreas Organ Allocation Simulator (OASim)

Utilization logic

- Assume that a kidney is not used if it is not accepted by the 1,000th offer

Accept/decline model

- Logistic regression to predict a candidate's likelihood of accepting a particular donor kidney
- Features: Donor characteristics, candidate characteristics, candidate-donor mismatch, geography, center-level, and allocation-related metrics

To represent our proposed simultaneous offers policy, we adjust the default OASim logic

Adjusted utilization logic

- Assume that a high-KDPI kidney is not used after being declined by all candidates at all centers

Adjusted accept/decline model

- Adjust allocation-related metrics to reflect that simultaneous offers do not have the same cold ischemia time accrued by time of offer or the information gained by previous candidates decline decisions

NYU Langone Health

Accept/decline model allocation-related metric: Center number

A candidate's center number is the number of unique transplant centers (i.e., hospitals) to which the donor kidney had been previously offered, including their own center

Candidate	Priority Ordering	Center	Center Number
A	1	TNVU	1
B	2	PAPT	2
C	3	TNVU	2
D	4	NJBI	3
E	5	CTYN	4
F	6	MAMG	5
G	7	CTYN	5
H	8	NJBI	5

Accept/decline model allocation-related metric: Center number

A candidate’s center number is the number of unique transplant centers (i.e., hospitals) to which the donor kidney had been previously offered, including their own center

Candidate	Priority Ordering	Center	Center Number
A	1	TNVU	1
B	2	PAPT	2
C	3	TNVU	2
D	4	NJBI	3
E	5	CTYN	4
F	6	MAMG	5
G	7	CTYN	5
H	8	NJBI	5

In the accept/decline logistic regression model:

- Center number coefficient: -0.10
- Center number odds ratio: 0.90

Interpretation: If a candidate’s center number increases by 1, they have a 10% decrease in the odds of accepting the offer

NYU Langone Health

Accept/decline model allocation-related metric: Center number

A candidate's center number is the number of unique transplant centers (i.e., hospitals) to which the donor kidney had been previously offered, including their own center

Candidate	Priority Ordering	Center	Center Number
A	1	TNVU	1
B	2	PAPT	2
C	3	TNVU	2
D	4	NJBI	3
E	5	CTYN	4
F	6	MAMG	5
G	7	CTYN	5
H	8	NJBI	5

In the accept/decline logistic regression model:

- Center number coefficient: -0.10
- Center number odds ratio: 0.90

Interpretation: If a candidate's center number increases by 1, they have a 10% decrease in the odds of accepting the offer

Center number represents a transplant candidate's location on the match run, which is a proxy for time since organ recovery and information gained by previous offer's declined

NYU Langone Health

Adjustments to allocation-related metrics: Center number

Allocation-related metric	Value in accept/decline models			
	OASim default	Adjustment for simultaneous offers policy (values for accepting/declining KDPI > 75% kidney)		
		Candidates at centers ≤ 250NM	Candidates at centers 250-500NM	Candidates at centers > 500NM
Center number	# unique centers from previous offers	1	1 + Total # unique centers ≤ 250NM	1 + Total # unique centers ≤ 500NM

NYU Langone Health

Adjustments to allocation-related metrics: Center number

Allocation-related metric	OASim default	Adjustment for simultaneous offers policy		
		Candidates at centers ≤ 250NM	Candidates at centers 250-500NM	Candidates at centers > 500NM
Center number	# unique centers from previous offers	1	1 + Total # unique centers ≤ 250NM	1 + Total # unique centers ≤ 500NM

Example:

Default OASim

Candidate	Priority Ordering	Distance	Center	Center Number
A	1	600NM	TNVU	1
B	2	300NM	PAPT	2
C	3	600NM	TNVU	2
D	4	50NM	NJBI	3
E	5	100NM	CTYN	4
F	6	400NM	MAMG	5
G	7	100NM	CTYN	5
H	8	50NM	NJBI	5

Adjusted for Simultaneous Offers

Candidate	Priority Ordering	Distance	Center	Center Number
D	4	50NM	NJBI	1
E	5	100NM	CTYN	1
G	7	100NM	CTYN	1
H	8	50NM	NJBI	1
B	2	300NM	PAPT	3
F	6	400NM	MAMG	3
A	1	600NM	TNVU	5
C	3	600NM	TNVU	5

OASim Data

- Four datasets of two-months between 2020-2023
- Each dataset contains:
  - ~90,000 transplant candidates
  - 3 samples of ~2,000 donors

We run 60 total simulations for both the existing allocation policy and our proposed simultaneous offers policy

RESULTS

Our proposed simultaneous offers policy decreases nonuse from 56.0% to 20.9% among high-KDPI kidneys

	Existing Policy	Simultaneous Offers Policy
Nonuse	percentage	
All kidneys	21.6%	12.5%
KDPI > 75% kidneys	56.0%	20.9%
KDPI ≤ 75% kidneys	9.7%	9.6%

NYU Langone Health

Our proposed simultaneous offers policy increases the logistical burden of evaluating offers

	Existing Policy	Simultaneous Offers Policy
Logistical burden	mean	
# kidneys offered / center / mo	249.3	331.7
# KDPI > 75% kidneys offered / center / mo	136.7	207.6
# KDPI ≤ 75% kidneys offered / center / mo	128.8	128.81

NYU Langone Health

CONCLUSION

NYU Langone Health

Conclusions

- Deceased donor kidneys are a scarce and precious resource
- Thousands of kidneys are not used every year, while thousands of transplant candidates die on the waitlist
- Simultaneous offers for high-KDPI kidneys within geographic thresholds is a promising allocation policy that:
  - Reduces kidney nonuse from 21.6% to 12.5%
  - But increases the logistical burden of evaluating more kidney offers from 249.3 to 331.7 kidneys offered / center / month

NYU Langone Health

THANK YOU

NYU Langone Health

Dorry Segev, MD, PhD
ERGOT Founder and Director
C-STAR Director

NYU Langone MEDICAL CENTER

JOHNS HOPKINS MEDICINE

Faculty

NYU

***Sommer Gentry, PhD**
Co-Director, C-STAR
Director, Computational Core
***Macey Levan, JD, PhD**
Director, C-STAR (Qualitative Core)
***Allan Massie, PhD**
Director, C-STAR (Quantitative Core)
Associate Professor of Surgery, Epidemiology
***Mara McAdams-DeMarco, PhD, MS**
Associate Professor of Surgery and Population Health
Associate Vice Chair for Research
***Sunjae Bae, KMD, PhD**
Assistant Professor of Surgery
***Byoungjun Kim, PhD**
Assistant Professor of Surgery
***Bonnie Lonze, MD, PhD**
Assistant Professor
***Darren Stewart, MS**
Associate Director of Registry Studies
Robert Montgomery, MD, DPhil
Professor of Surgery
Brendan Parent, JD
Medical Ethics
Laura Kimberly, PhD
Plastic Surgery
Natasha Gupta, MD
Urology
Morgan Grams, MD, PhD
Division Director of Precision Medicine

JHU

***Christine Durand, MD**
Associate Professor of Medicine
Director, JH Transplant Research Center
***William Werbel, MD**
Assistant Professor of Medicine
***Betsy King, MD, PhD**
Director, JH Transplant Research Center
Assistant Professor of Surgery

JHU

***Tanjala Purnell, PhD, MPH**
Dir. of Community and Stakeholder Engagement
Assistant Professor of Epidemiology and Surgery
***Andrew Cameron, MD, PhD**
Professor of Surgery
***Nadia Chu, PhD, MPH**
Assistant Professor of Surgery and Epidemiology
***Douglas Mogul, MD, PhD**
Assistant Professor of Pediatrics Hepatology and Nutrition
***Abimereki Muzaale, MD, MPH**
Assistant Professor of Surgery
Victor Chen, MD
Transplant Hepatology
***Daniel Warren, PhD**
Director, JH Transplant Research Center
Associate Professor, Islet Transplantation
***Robin Avery, MD**
Professor of Medicine
Anna Beavis, MD
Gynecology
Gerald Brandacher, MD
Plastics and Reconstructive Surgery
Daniel Brennan, MD
Transplant Nephrology
Errol Bush, MD
Lung Transplant Program

**Core Faculty*

JHU

Olga Charnaya, MD
Pediatric Nephrology
Josef Coresh, MD, PhD
Epidemiology
Niraj Desai, MD
Transplant Surgery
Shane Ottman, MD
Transplant Surgery

California

Fawaz Al Ammary, MD, PhD
Medical Director, Living Donation

Wisconsin

Jacqueline Garonzik-Wang, MD, PhD
Associate Professor of Surgery

Thank You

Trainees

Post Doc /Fellows
Alvin Thomas, MSPH
Maarten Coemans, PhD
Aly Strauss, MD MIE
Jamilah Perkins, MD MHS
Sharon Weeks, MD

Surgery Residents
Jennifer Alejo, MD
Jacqueline Kim, MD
Amrusha Musunuru, MD
Jessica Ruck, MD
Victoria Bendersky, MD
Amy Chang, MD
Mackenzie Eagleson, MD
Kieran Nijhar, MD
Andrew Hallett, MD
Jonathan Mitchell, MBBS

Medical Students
Shivani Bisen
Mary Grace Bowring, MPH
Sarah Frey
Caroline Qin

Research Scientist
Karen Vanterpool, PhD, MPH
Nidhi Ghildayal, PhD
Samantha Klitenic, JD

Research Program Manager
Carolyn Sidoti, BS

Research Project Manager
Natalia Martinez-López, MPA

Data Analysts
Tanveen Ishaque, MBBS ScM
Jennifer Motter, MHS
Po-Yu Chiang, MD MPH
Yiting Li, MPH
Kadiatou Diallo, MPH
Katya Kaplow, MPH
Yi Liu, MS
Gayathri Menon, MHS
Akanksha Nalatwad, MPH
Yusi Chen, MHS
Jingyao Hong, MHS
Laura Bowles Zeiser, ScM
Suhani Patel, MPH
Michael Mankowski
Hannah Sung, PhD
Sarah Rasmussen

Research Program Coordinators
Max Downey, BS
Jasmine Akhtar, BS
Aura Abedon
Mayan Teles
Alexa Jefferis
Julia López

Administrative Coordinators
Johanna Paguay
Morgan Johnson
Quinten Stearns

Interns
Devyn Jones
Denise Nemeth
Evelien Quint

Research Assistants/Associates
• Lauranna Maresca • Helen Crosby • Swati Sharma • Barbara Krawczuk • Julia Panchenko • Eric Donohue • Alain Phung • Vedant Jain • Whitney Langlee • Alex Alex • Nicole Fortune Hernandez • Rivka Abedon • Snigdha Panda • Angela Ramirez • Wendy Huang • Ross Greenberg • Letitia (Kenzie) Thomas • Roni Shanoada • Janetta Brundage • Ellie Kim • Mimi Mensah • Molly Ma • Sofia Romano • Sachin Mehta • Linda Paredes • Bahati Kuffer • Jiashu Xue•

NYU Langone Health