# Robust Policies for Proactive ICU Transfers

Carri W. Chan[1], Vineet Goyal[2], Julien Grand-Clement[2], Gabriel Escobar[3]

[1] Columbia Business School, Columbia University
[2] IEOR Department, Columbia University
[3] Division of Research, Kaiser Permanente

*What is this talk about?*

Uncertainty in transition matrices of Markov Chains.

Geometry of the uncertainty set and tractability.

Application: efficient and robust ICU admission guidelines.

# Background on Intensive Care Unit (ICU).

Critically ill patients are treated in ICU.

Unplanned transfer:

Significant impact on both patients outcomes and operation costs.

- Unplanned transfer to the ICU[1]:
  Stays 10 days longer than average patient.
  Mortality 3 times higher than average patient.

[1] Escobar et al. (2013)
[2] Mullins et al. (2013)
[3] Mildbrandt et al. (2013)

# Background on Intensive Care Unit (ICU).

Critically ill patients are treated in ICU.

Unplanned transfer:

Significant impact on both patients outcomes and operation costs.

- Unplanned transfer to the ICU[1]:
  Stays 10 days longer than average patient.
  Mortality 3 times higher than average patient.
- Cost of ICU patient: up to 2.5 higher than in ward[2].

[1] Escobar et al. (2013)
[2] Mullins et al. (2013)
[3] Mildbrandt et al. (2013)

# Background on Intensive Care Unit (ICU).

Critically ill patients are treated in ICU.

Unplanned transfer:

Significant impact on both patients outcomes and operation costs.

- Unplanned transfer to the ICU[1]:
  Stays 10 days longer than average patient.
  Mortality 3 times higher than average patient.
- Cost of ICU patient: up to 2.5 higher than in ward[2].
- ICU admissions increased by 48.8% from 2002 to 2009[3].

[1] Escobar et al. (2013)
[2] Mullins et al. (2013)
[3] Mildbrandt et al. (2013)

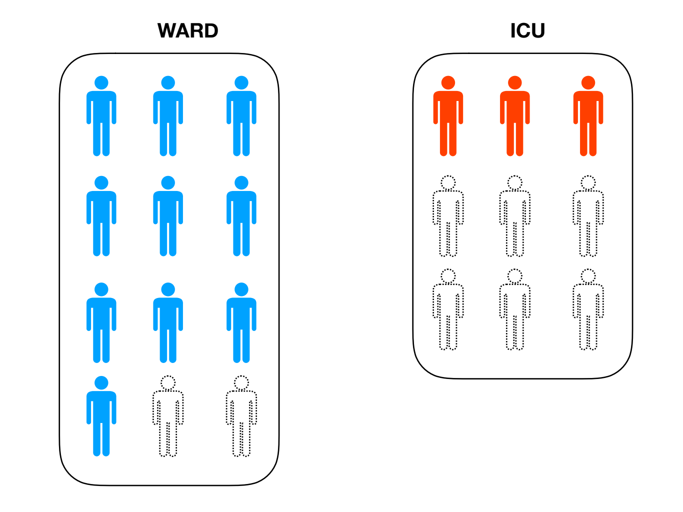


Simplified hospital: ward vs ICU.

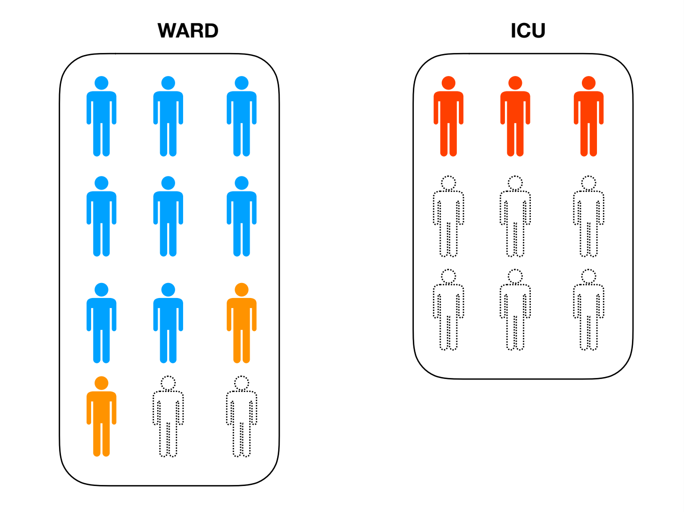


Suboptimal use of ICU resources (congestioned ward, ‘empty’ ICU).

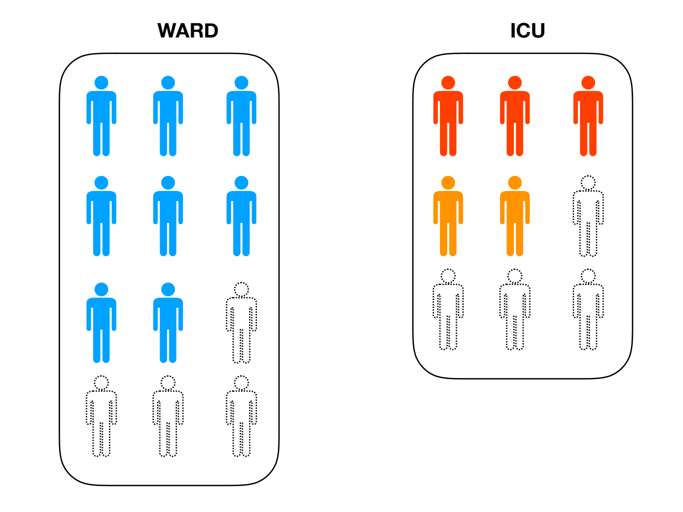


Proactive transfer patients to the ICU: lower mortality.

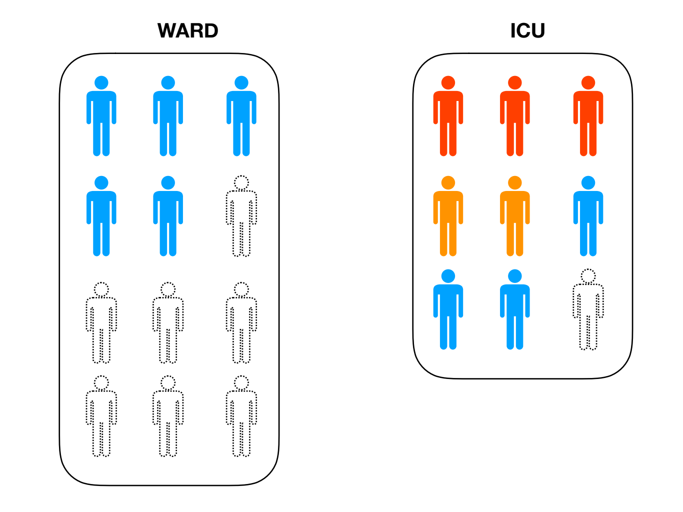


Too many transfers $\Rightarrow$ congestioned ICU: can not face unplanned events.

# Research question and literature.

Research question.

Impact of proactive transfers on hospital performances?

Robustness of predictive models to parameters mispecifications?

# Research question and literature.

Research question.

Impact of proactive transfers on hospital performances?

Robustness of predictive models to parameters mispecifications?

Related works.

*ICU admission.* Shmueli et al. (2004), Peck et al. (2012), Bountourelis et al. (2012), Butcher et al. (2013), Guirgis et al. (2013), Kim et al. (2014), Xu and Chan (2016), Hu et al. (2018).

*Threshold policy.* Altman et al. (2001), Shmueli et al. (2003), Kim et al. (2015), Hu et al. (2018).

*Robust MDP.* Iyengar (2005), Nilim and El Ghaoui (2005), Delage and Mannor (2010), Wiesemann et al. (2013), Goh et al. (2014), Mannor et al. (2016), Goyal and G-C. (2018), Steimle et al. (2018).

# Our results.

1. Markov Decision Process for patient dynamics.
2. Low-rank model for parameter uncertainty.
3. Optimality and Robustness of threshold policies.
4. Numerical study: worst-case analysis of hospital performances.

# Single-Patient Markov Decision Process (MDP).

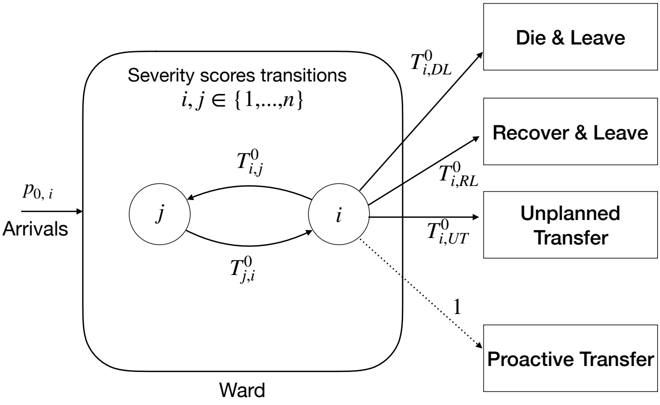


UT/PT: Unplanned/Proactive Transfer. DL/RL: Die/Recover and Leave.

# Single-Patient Markov Decision Process (MDP).

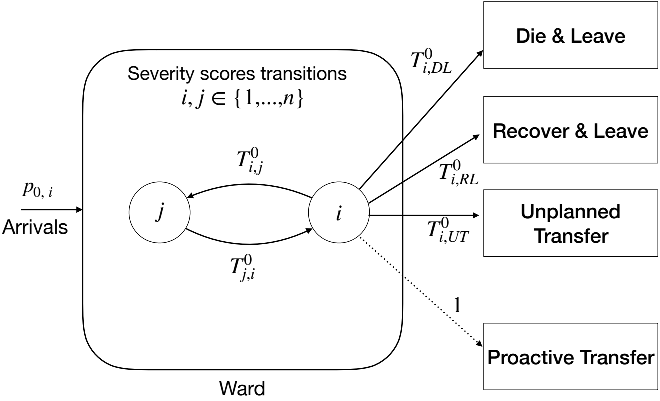


Transition matrix: $\boldsymbol{T}^0$. Transfer policy $\pi : \{$ severity scores $\} \rightarrow [0, 1]$.

# Single-Patient Markov Decision Process (MDP).

**Rewards:** $r_{PT}, r_{UT}, r_{RL}, r_{DL}$.

How to calibrate? Natural assumption:

Ordering by <u>outcome</u> and <u>use</u> of ICU resources.

$$r_{RL} \geq r_{PT} \geq r_{UT} \geq r_{DL}.$$

**Goal:** Find transfer policy $\pi$ to maximize expected reward $R(\pi, \boldsymbol{T}^0)$, for

$$R(\pi, \boldsymbol{T}^0) = E^{\pi, \boldsymbol{T}^0}\left[\sum_{t=0}^{\infty} \lambda^t r_{s_t}\right].$$

# Theoretical guarantees.

Key Assumptions: for all state $i \in \{1, ..., n-1\}$,

$$\text{Reward}\,(\text{ exit the ward} \mid \text{state } i) \geq \text{Reward}\,(\text{ exit the ward} \mid \text{state } i+1)\,. \tag{1}$$

$$\left(\mathbb{P}\,(\text{ exit the ward} \mid \text{state } i)\right)_i \text{ increases fast enough.} \tag{2}$$

**Theorem**

Let $\boldsymbol{T}^0$ a transition matrix that satisfies Assumptions (1)-(2).

There exists a threshold policy $\pi^{\text{nom}}$ that is optimal for $\boldsymbol{T}^0$.

Proof intuition: at every step of Value Iteration, $\pi_s$ is threshold.

Efficient algorithm: Value Iteration or Enumeration of threshold policies.

# Theoretical guarantees.

Key Assumptions: for all state $i \in \{1, ..., n-1\}$,

$$r_{RL} T^0_{i,RL} + r_{DL} T^0_{i,DL} + r_{PT} T^0_{i,PT} \geq r_{RL} T^0_{i+1,RL} + r_{DL} T^0_{i+1,DL} + r_{PT} T^0_{i+1,PT}, \tag{1}$$

$$\frac{1 + \lambda \cdot r_{PT}}{1 + \lambda \cdot r_{RL}} \geq \frac{\left(\sum_{j=1}^{n} T^0_{i+1,j}\right)}{\left(\sum_{j=1}^{n} T^0_{ij}\right)}. \tag{2}$$

**Theorem**

Let $\boldsymbol{T}^0$ a transition matrix that satisfies Assumptions (1)-(2).

There exists a threshold policy $\pi^{\text{nom}}$ that is optimal for $\boldsymbol{T}^0$.

Proof intuition: at every step of Value Iteration, $\pi_s$ is threshold.

Efficient algorithm: Value Iteration or Enumeration of threshold policies.

# Parameter uncertainty.

In practice, $\boldsymbol{T}^0$ is estimated from some data and is uncertain.

From the data, we obtain 95% confidence intervals:

$$T_{i,j}^{\text{true}} \in [\,T_{i,j}^0 - \alpha_i,\, T_{i,j}^0 + \beta_i].$$

Example: $T_{2,1}^0 = 0.3216$, $T_{2,1}^{\text{true}} \in [0.3208, 0.3232]$.

<u>Simulations: deviations of $\approx 10^{-3}$ may increase mortality by 40 % !</u>

# Robust MDP.

**Robust approach:** $\boldsymbol{T}^0 \in \mathcal{U}$, 'uncertainty set': plausible matrices.

$\mathcal{U}$ is computed from the confidence-intervals of $\boldsymbol{T}^0$.

Zero-sum game:

(i) Decision-maker chooses policy $\pi$,
(ii) Nature chooses $\boldsymbol{T} \in \mathcal{U}$ adversarially,
(iii) Decision-maker obtains $R(\pi, \boldsymbol{T})$.

**Our goal**

$$\text{solve } \max_{\pi} \min_{\boldsymbol{T} \in \mathcal{U}} R(\pi, \boldsymbol{T}).$$

**Robust policy:** maximizes worst-case reward over all matrices in $\mathcal{U}$.

$\rightarrow$ The robust policy guarantees a certain level of performance.

# Related works on robust MDPs.

For general $\mathcal{U}$, computing the worst-case reward of a policy $\pi$ is hard:

"$\min_{\boldsymbol{T}\in\mathcal{U}} R(\pi, \boldsymbol{T}) \geq \gamma$ ?" is NP-Hard (Wiesemann et al. (2013)).

Tractable for some special cases:

- $(s, a)$-retangularity (Iyengar (2005), Nilim and El Ghaoui (2005))
- $s$-rectangularity (Wiesemann et al. (2013)).
- $k$-rectangularity (Mannor et al. (2016)).

# Choice of uncertainty set.

*r*-rectangular uncertainty set[4]:

$$\begin{aligned}\mathcal{U} = \{ \boldsymbol{T} \mid \boldsymbol{T} = \boldsymbol{U}\boldsymbol{W}^\top, \\ &\boldsymbol{W} = (\boldsymbol{w}_1, ..., \boldsymbol{w}_r) \in \mathcal{W}^1 \times ... \times \mathcal{W}^r, \\ &\boldsymbol{W}^\top \boldsymbol{e} = \boldsymbol{e}, \boldsymbol{W} \geq \boldsymbol{0}, \\ &\boldsymbol{U} \geq \boldsymbol{0} \text{ is fixed}, \boldsymbol{U}\boldsymbol{e} = 1 \}.\end{aligned}$$

$\rightarrow$ Think about SVD with *uncertain* eigenvectors.

$\rightarrow$ Coupled rows of $\boldsymbol{T}$ = simple transformation of uncoupled $\boldsymbol{w}_1, ..., \boldsymbol{w}_r$.

[4] Goh et al. (2014), Goyal and G-C. (2018)

# Choice of uncertainty set.

*r*-rectangular uncertainty set[5]:

$$\mathcal{U} = \{ \boldsymbol{T} \mid \boldsymbol{T} = \boldsymbol{U}\boldsymbol{W}^\top, \\ \boldsymbol{W} = (\boldsymbol{w}_1, ..., \boldsymbol{w}_r) \in \mathcal{W}^1 \times ... \times \mathcal{W}^r, \\ \boldsymbol{W}^\top \boldsymbol{e} = \boldsymbol{e}, \boldsymbol{W} \geq \boldsymbol{0}, \\ \boldsymbol{U} \geq \boldsymbol{0} \text{ is fixed}, \boldsymbol{U}\boldsymbol{e} = 1 \}.$$

→ Rank $r$ is the number of underlying variables driving the dynamics.

→ We implicetely assume $\boldsymbol{T}^0 = \boldsymbol{U}\boldsymbol{W}_0^\top$ for a feasible $\boldsymbol{W}_0$.

[5] Goh et al. (2014), Goyal and G-C. (2018)

# Our results (Goyal and G.-C. (2018)).

Consider an r-rectangular uncertainty set $\mathcal{U}$.

**Robust Value Iteration Algorithm.**

We can efficiently compute an optimal robust policy.

**Strong Duality and Nash Equilibrium.**

$$\max_{\pi} \min_{\boldsymbol{T} \in \mathcal{U}} R(\pi, \boldsymbol{T}) = \min_{\boldsymbol{T} \in \mathcal{U}} \max_{\pi} R(\pi, \boldsymbol{T}).$$

$$\boldsymbol{T}^* \in \arg\min_{\boldsymbol{T} \in \mathcal{U}} R(\pi^*, \boldsymbol{T}),$$

$$\pi^* \in \arg\max_{\pi} R(\pi, \boldsymbol{T}^*).$$

**Interpretability and implementability.**

There exists a <u>deterministic</u> optimal robust policy $\pi^*$.

# Useful tools (Goyal and G.-C. (2018)).

Value vector $\boldsymbol{V}$: for all state $s$, $V_s(\pi, \boldsymbol{T}) = E^{\pi,\boldsymbol{T}} \left[\sum_{t=0}^{\infty} \lambda^t r_{s_t} \mid s_0 = s\right].$

Robust Maximum Principle:

1. Fix a transition matrix $\boldsymbol{T}^0$ and let $\pi^0$ be the optimal policy for $\boldsymbol{T}^0$.

   Then $\boldsymbol{V}(\pi, \boldsymbol{T}^0) \leq \boldsymbol{V}(\pi^0, \boldsymbol{T}^0), \forall\, \pi.$

2. Fix a policy $\pi$ and let $\boldsymbol{T}^{\pi}$ is worst-case matrix for policy $\pi$.

   Then $\boldsymbol{V}(\pi, \boldsymbol{T}^{\pi}) \leq \boldsymbol{V}(\pi, \boldsymbol{T}), \forall\, \boldsymbol{T} \in \mathcal{U}.$

3. Let $\pi^*$ an optimal robust policy.

   Then $\boldsymbol{V}(\pi, \boldsymbol{T}^{\pi}) \leq \boldsymbol{V}(\pi^*, \boldsymbol{T}^*), \forall\, \pi.$

Robust Blackwell Optimality:

Let $\pi^*_{\lambda}$ be the robust policy for the discount factor $\lambda$.

Then $\exists\, \lambda_0 \in [0, 1), \forall\, \lambda \geq \lambda_0, \pi^*_{\lambda} = \pi^*_{\lambda_0}.$

Back to single-patient MDP:

What is the structure of an optimal robust policy?

Numerical study:

How much change in hospital performances when parameters deviate?

# Robust transfer policies.

We make the following assumption:

$$\text{Every matrix } \boldsymbol{T} \text{ in } \mathcal{U} \text{ satisfies Assumptions (1)-(2).} \tag{3}$$

**Theorem**

Let Assumption (3) hold.

There exists an optimal robust policy $\pi^{\text{rob}}$ that is a threshold policy.

Moreover,

$$\text{threshold}(\pi^{\text{rob}}) \leq \text{threshold}(\pi^{\text{nom}}).$$

Intuition: optimal robust policy transfers more patients naive one.

Proof: relies on the robust maximum principle.

Efficient algorithm: Robust Value Iteration or Enumeration.

# Numerical study.

Kaiser Permanente dataset.

- $\approx 300,000$ hospitalizations, 10 severity scores (EDIP2).
- Patients arrive in hospital ward with severity score $i \in \{ 1, ..., 10 \}$.
- Severity score updated every six hours: $i \rightarrow \{ j, UT, RL, DL \}$.

Experimental setup:

1. Construct an hospital model for simulation.
2. Construct uncertainty set $\mathcal{U}$ from the data.
3. Compare nominal and worst-case performances in hospital.

# Hospital model inspired from Hu et al. (2018).

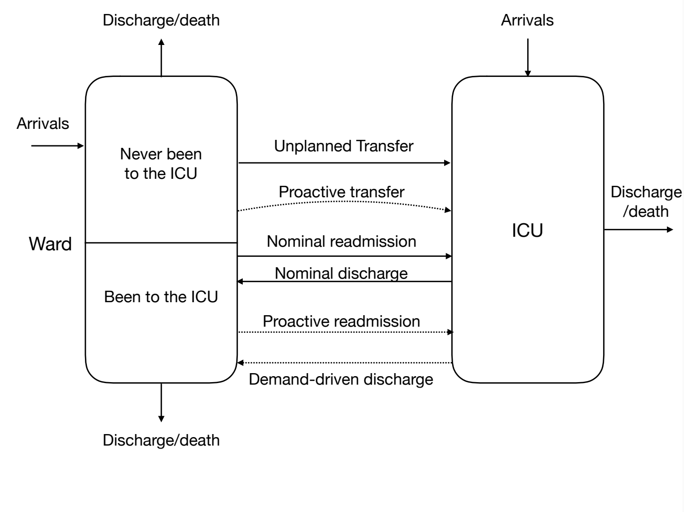


Markov Chain: $T^0_{i,j}$ among ‘severity of illness’ scores $(i, j)$.

# Hospital model inspired from Hu et al. (2018).

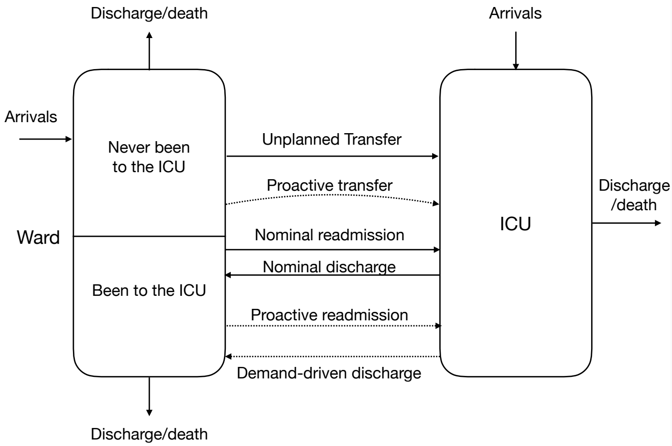


Simulations: compute mortality, length-of-stay, ICU occupancy.

# Parameters estimation.

a) <u>Empirical average:</u> $T^0_{i,j} = \mathbb{P}($ go to score j $|$ being in score i $)$.

From the data, 95% confidence intervals: $T^{\text{true}}_{i,j} \in [T^0_{i,j} - \alpha_i, T^0_{i,j} + \beta_i]$.

Example: $T^0_{2,1} = 0.3216$, $T^{\text{true}}_{2,1} \in [0.3208, 0.3232]$.

b) <u>Hospital worst-case performances:</u> given a policy $\pi$, we compute

$$\boldsymbol{T}^{\pi} \in \arg\min_{\boldsymbol{T} \in \mathcal{U}} R(\pi, \boldsymbol{T}).$$

$\boldsymbol{T}^{\pi}$ is a candidate for worst-case matrix in hospital dynamics.

We compute the hospital performance of $\pi$ for transition matrix $\boldsymbol{T}^{\pi}$.

## How to use r-rectangular $\mathcal{U}$ in simulations?

c) Uncertainty set: $\boldsymbol{T} = \boldsymbol{U}\boldsymbol{W}^{\top}, \boldsymbol{W} \in \mathcal{W}^1 \times ... \times \mathcal{W}^r$.

$\boldsymbol{W}$ is varying, but we have constraints on $\boldsymbol{T}$! (confidence intervals)

(i) $\mathcal{U}_{\text{min}}$: small deviations of rank $r$:

$$\text{maximum deviation from } \boldsymbol{T} \text{ will be } \min\{\, \alpha_i, \beta_i \mid i \in [10] \,\}.$$

(ii) $\mathcal{U}_{\text{std}}$ : empirical std's of rank $r$ matrices in confidence intervals:

$$\text{Sample } \boldsymbol{T}^1, ..., \boldsymbol{T}^N \text{ around } \boldsymbol{T}^0, \text{ compute } \boldsymbol{W}^1, ..., \boldsymbol{W}^N.$$

(iii) $\mathcal{U}_{sa}$ : unrelated parameters deviations:

$$\begin{aligned} \mathcal{U}_{sa} = \{\, \boldsymbol{T} \mid T_{i,j} \in [T^0_{i,j} - \alpha_i, T^0_{i,j} + \beta_i], \forall\, i, j \in [10], \\ \boldsymbol{T}\boldsymbol{e} = \boldsymbol{e}, \boldsymbol{T} \geq 0 \,\}. \end{aligned}$$

# Numerical results 1: nominal performances.

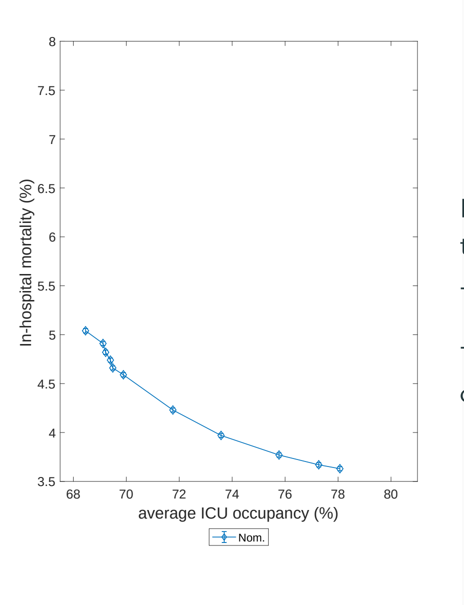


Hospital performances of threshold policies.

Transition matrix: $\boldsymbol{T}^0$.

Trade-off Mortality vs ICU occupancy.

# Numerical results 2: random samples.

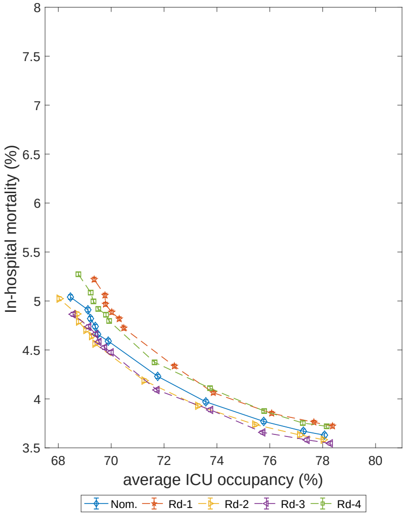


25 random samples around $\boldsymbol{T}^0$:

Mortality may increase by 6.1%.
ICU occupancy may increase by 2.6%.

$\Rightarrow$ Model looks stable.

# Numerical results 3: worst-case matrices.

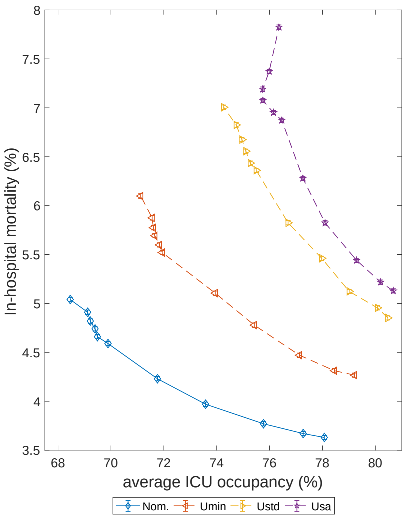


$\mathcal{U}_{\min}, \mathcal{U}_{\text{std}}$: rank-r deviations.

$\mathcal{U}_{sa}$: full-rank deviations.

1. Worst-case approach: Mortality may increase by 40%. ICU occupancy may increase by 12%.
2. Different insight for $\mathcal{U}_{sa}$: no trade-off Mortality/ICU occupancy (for high thresholds).

# Conclusion.

**Our contributions:**

For proactive transfer policies,

1. Guarantees for optimality/robustness of threshold policies.
2. Model of *r-rectangular* uncertainty from robust MDP.
3. Simulations for worst-case deviations.

**Main take-aways.**

When using predictive models:

1. Small parameters deviations may lead to worse performances.
2. Different insights on the model:

   Nominal parameters & random samples vs. worst-case analysis.
3. Important: choice of uncertainty sets: $(s, a)$-rec.,$s$-rec.,$r$-rec., etc.

**Main take-aways.**

When using predictive models:

1. Small parameters deviations may lead to worse performances.
2. Different insights on the model:

   Nominal parameters & random samples vs. worst-case analysis.
3. Important: choice of uncertainty sets: $(s, a)$-rec.,$s$-rec.,$r$-rec., etc.

Thanks! Any questions?

jg3728@columbia.edu