Intelligent Systems (AI-2)

Computer Science cpsc422, Lecture 9

Sep, 28, 2016

An MDP Approach to Multi-Category Patient Scheduling in a Diagnostic Facility

Adapted from: Matthew Dirks

Goal / Motivation

- To develop a mathematical model for multi-category patient scheduling decisions in computed tomography (CT), and to investigate associated trade-offs from economic and operational perspectives.

- Contributions to AI, OR and radiology

Types of patients:

- Emergency Patients (EP)
  - **Critical (CEP)**
  - **Non-critical (NCEP)**
- **Inpatients (IP)**
- Outpatients
  - Scheduled OP
  - **Add-on OP: Semi-urgent (OP)**
- **(Green = Types used in this model)**

Proposed Solution

- Finite-horizon MDP
- Non-stationary arrival probabilities for IPs and EPs
- Performance objective: Max $

MDP Representation

- State
  - $s = (e_{CEP}, w_{OP}, w_{IP}, w_{NCEP})$
  - e_{CEP} CEP arrived
  - w_{type} Number waiting to be scanned
- Action
  - $a = (a_{OP}, a_{IP}, a_{NCEP})$
  - a_{type} Number chosen for next slot

$$a_{\mathrm{OP}} + a_{\mathrm{IP}} + a_{\mathrm{NCEP}} + e_{\mathrm{CEP}} \leq R$$

$$a_j \leq w_j, \quad j = \mathrm{OP}, \mathrm{IP}, \mathrm{NCEP},$$

- State Transition
  - $s' = (d_{CEP}, w_{OP} + d_{OP} - a_{OP}, w_{IP} + d_{IP} - a_{IP}, w_{NCEP} + d_{NCEP} - a_{NCEP})$
  - d Whether a patient type has arrived since the last state

MDP Representation (cont')

- Transition Probabilities

$$P_n(s'|s, a) = p_n(d_{\mathrm{CEP}}) \times p_n(d_{\mathrm{OP}}) \times p_n(d_{\mathrm{IP}}) \times p_n(d_{\mathrm{NCEP}}),$$

example

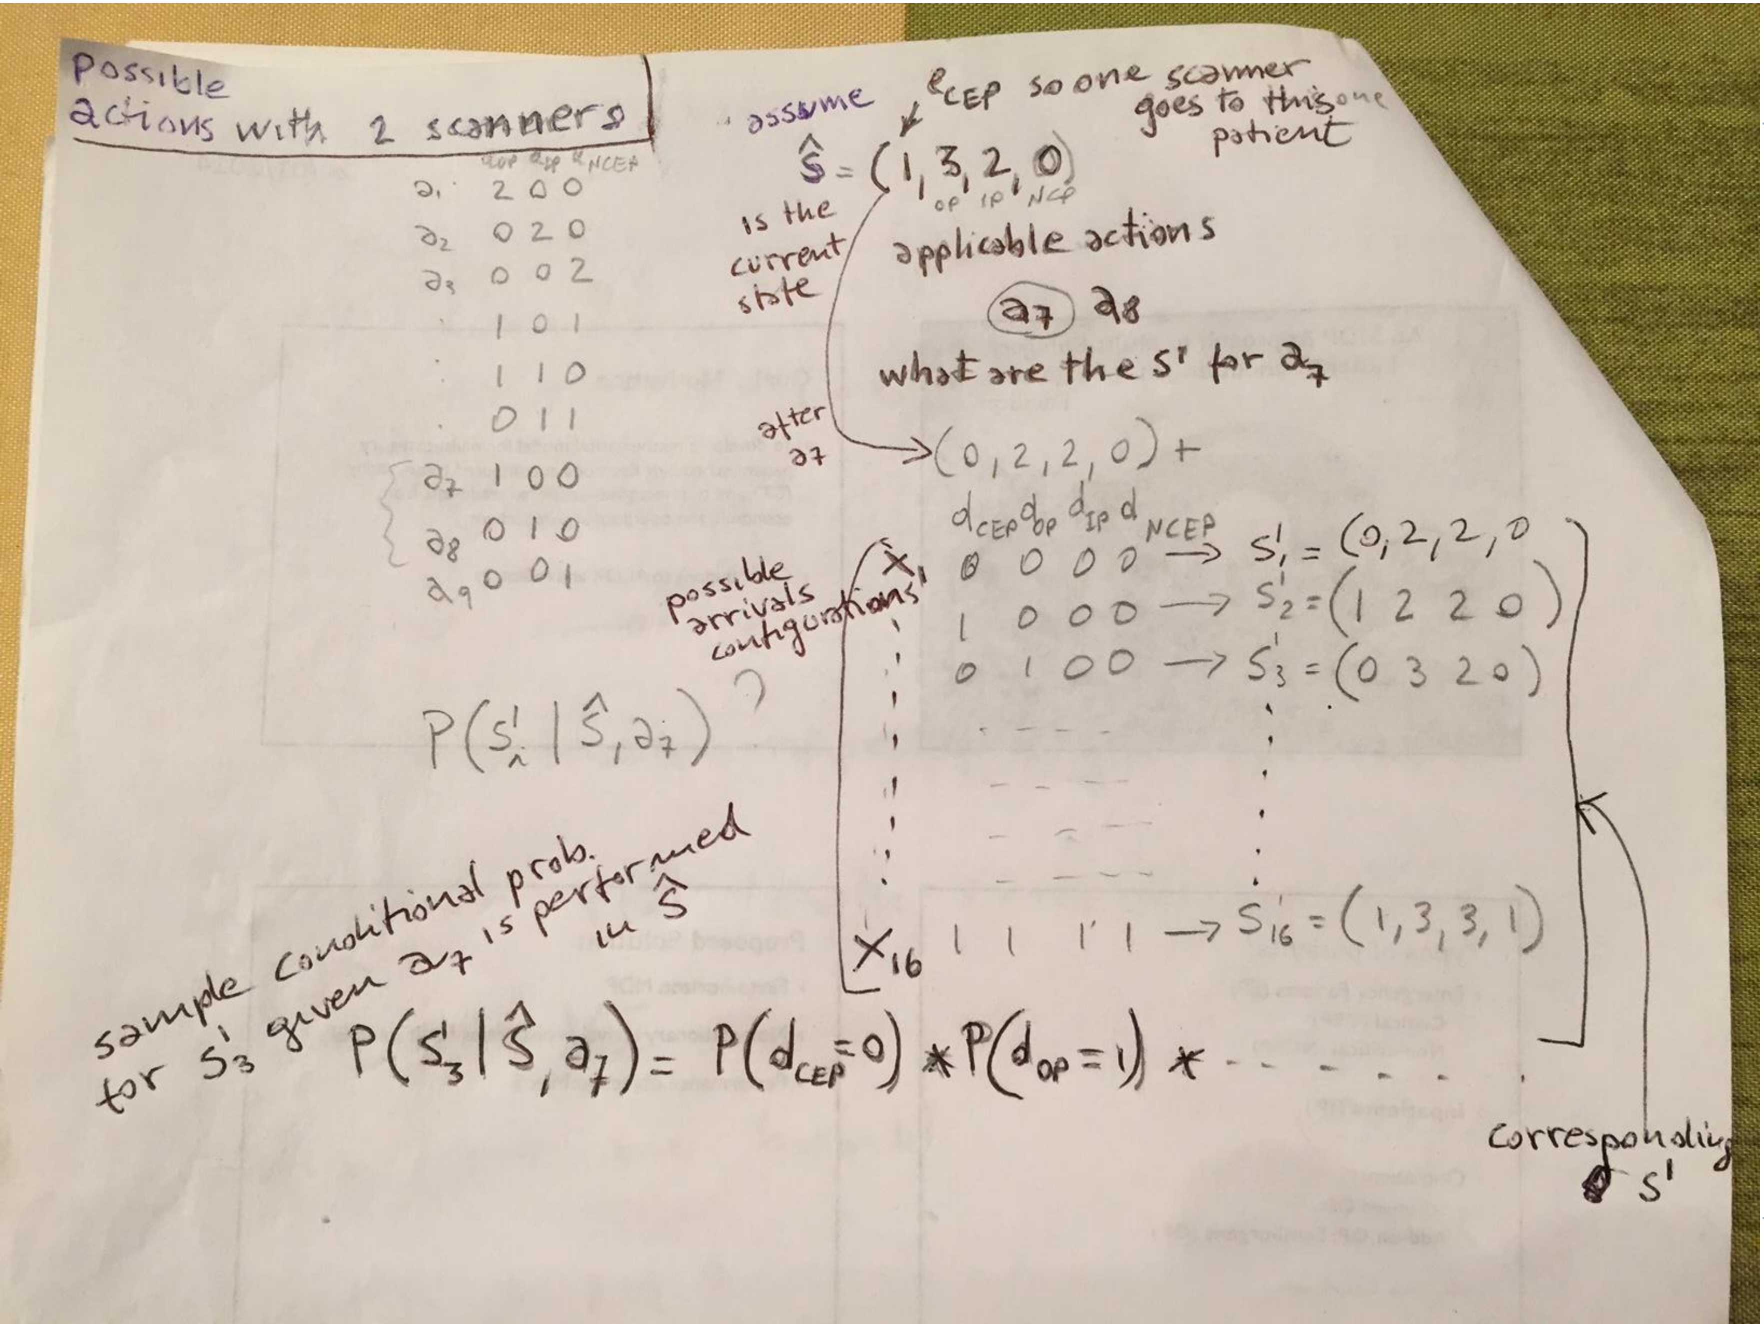
Possible actions with 2 scanners
a1 2 0 0
a2 0 2 0
a3 0 0 2
1 0 1
1 1 0
0 1 1
a7 1 0 0
a8 0 1 0
a9 0 0 1
assume
ŝ = (1, 3, 2, 0)
OP IP NCP
is the current state
CEP so one scanner goes to this one patient
applicable actions
a7 a8
what are the s' for a7
after a7
(0, 2, 2, 0) +
dCEP dOP dIP dNCEP
possible arrivals configurations
X1 0 0 0 0 → s'1 = (0, 2, 2, 0)
1 0 0 0 → s'2 = (1 2 2 0)
0 1 0 0 → s'3 = (0 3 2 0)
X16 1 1 1 1 → s'16 = (1, 3, 3, 1)
P(s'i | ŝ, a7) ?
sample conditional prob. for s'3 given a7 is performed in ŝ
P(s'3 | ŝ, a7) = P(dCEP = 0) * P(dOP = 1) * . . . . . .
corresponding s'

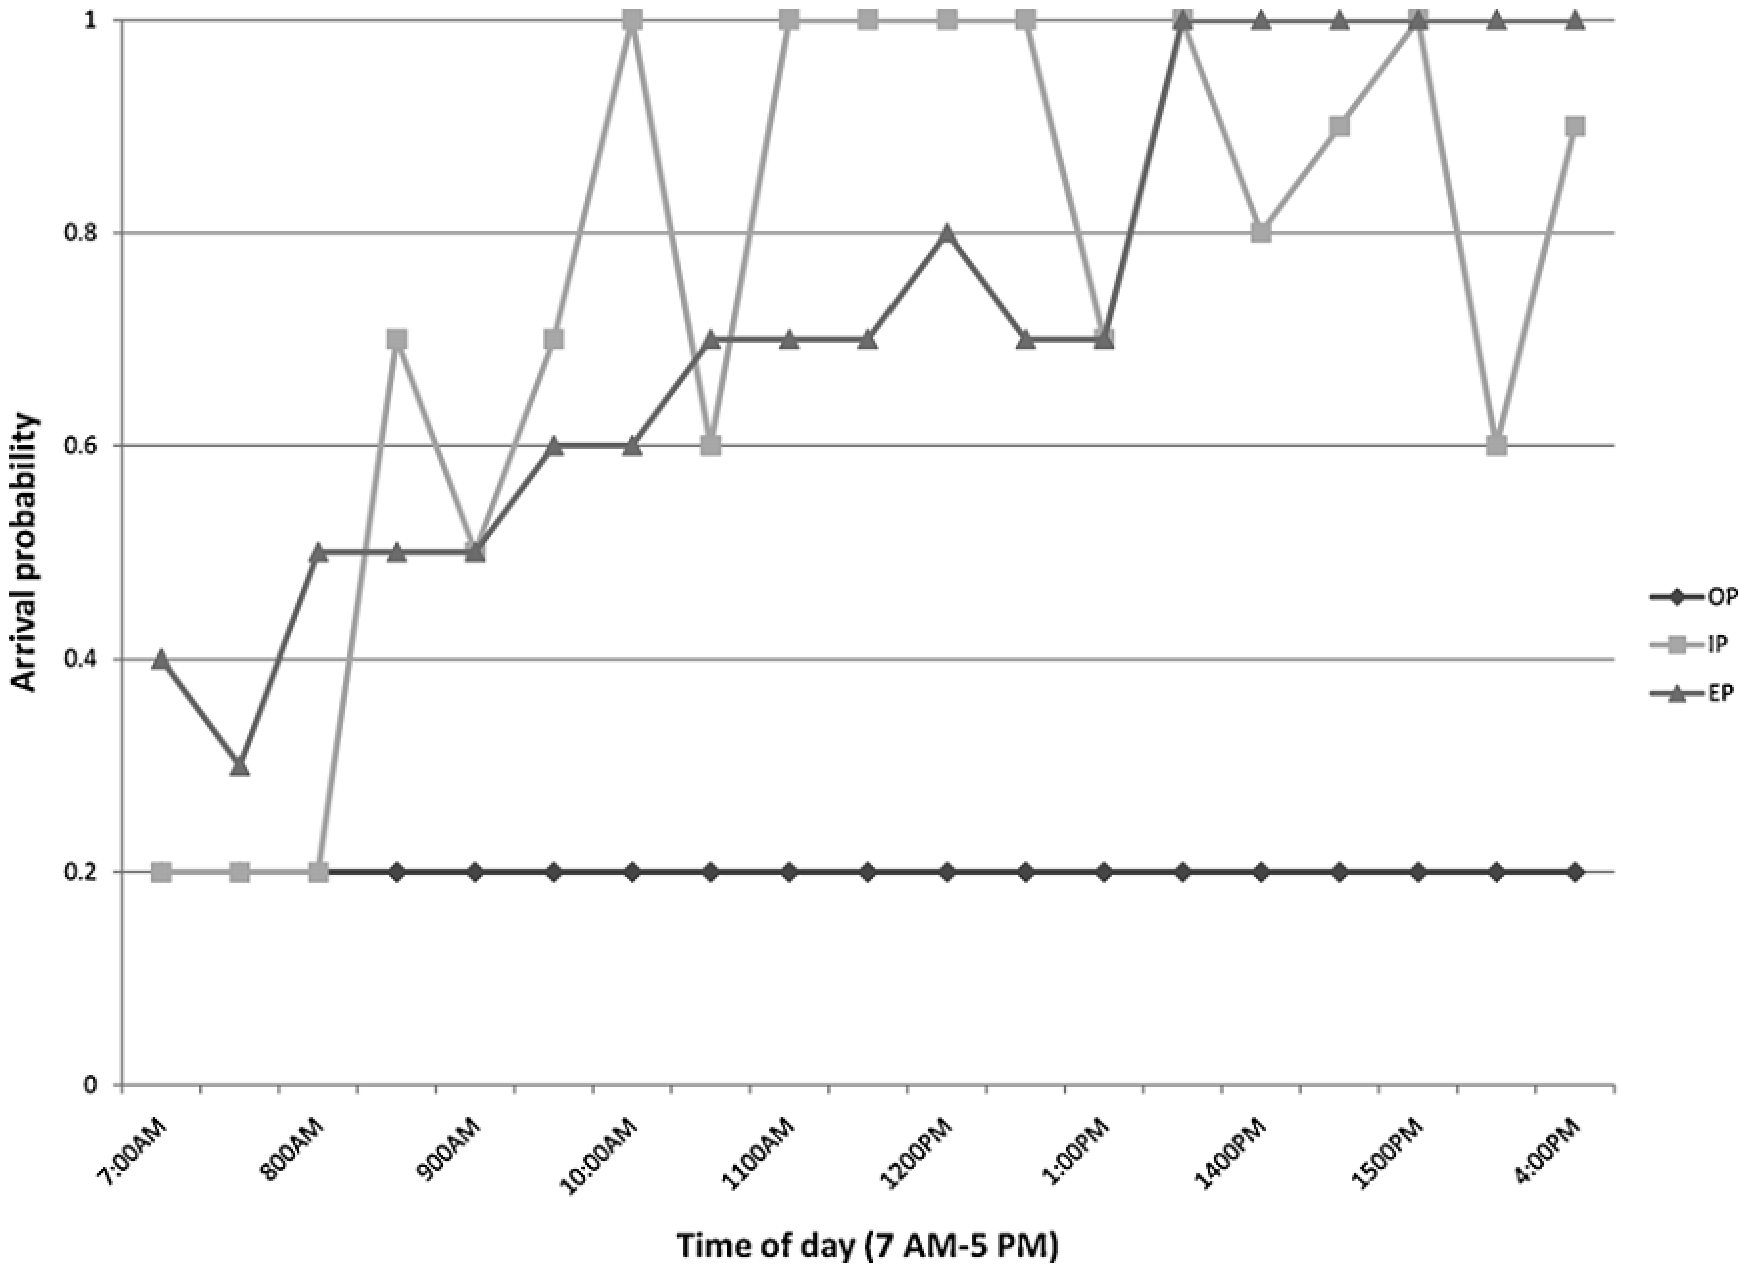

Fig. 1. Arrival probabilities for each patient-type during a work-day. EP includes both CEPs and NCEPs.

Performance Metrics (over 1 work-day)

- Expected net CT revenue
- Average waiting-time
- Average # patients not scanned by day's end

  - Rewards

$$r(s,a) = \sum_{j \in \{OP, IP, NCEP\}} r_j a_j - \sum_{j \in \{OP, IP, NCEP\}} (w_j - a_j) h_j.$$

revenue for scanning typeJ patient

waiting cost

Terminal reward obtained

$$V_{N+1}(s) = -c_{OP} w_{OP} - c_{IP} w_{IP} - c_{NCEP} w_{NCEP}$$

end of day penalty

- Discount factor? 1

Maximize total expected revenue

- Optimal Policy
  - Solving this gives the policy for each state, n, in the day

$$V_n(s) = \max_{a \in A(s)} \left\{ r(s,a) + \sum_{s'} P_n(s'|s,a) V_{n+1}(s') \right\}$$

- Finite Horizon MDP

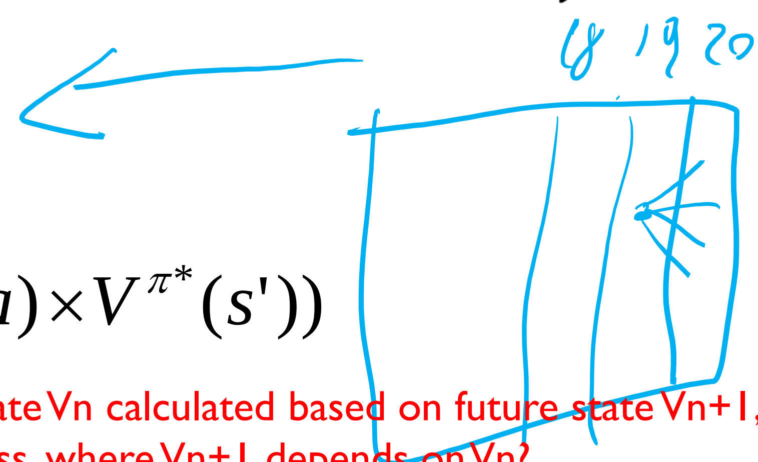

$$V^{\pi^*}(s) = R(s) + \gamma \max_a \sum_{s'} P(s'| s,a) \times V^{\pi^*}(s'))$$

- The recursive equation (3) has value of current state Vn calculated based on future state Vn+1, this contradicts with the equation given during class, where Vn+1 depends on Vn?
- The one in class was Value Iteration (the n index was for the iteration) here we have a finite horizon. We know the Vs at the end so we can compute all the Vs backward. n is an index for the time slice

Evaluation: Comparison of MDP with Heuristic Policies

- 100,000 independent day-long sample paths (one for each scenario)

$\forall n$ sample $p_n(d_x)$

$X = CEP, OP, IP, NCEP$

Result Metric

- Percentage Gap in avg. net revenue =

$$\frac{avg\ net\ revenue\ (optimal\ policy) - avg\ net\ revenue(heuristic\ policy)}{avg\ net\ revenue(optimal\ policy)}\ x\ 100$$

Heuristics

- **FCFS**: First come first serve
- **R-1**: One patient from randomly chosen type is scanned
- **R-2**: One patient randomly chosen from all waiting patients (favors types with more people waiting)
- **O-1**: Priority
  - OP
  - NCEP
  - IP
- **O-2**: Priority:
  - OP
  - IP
  - NCEP

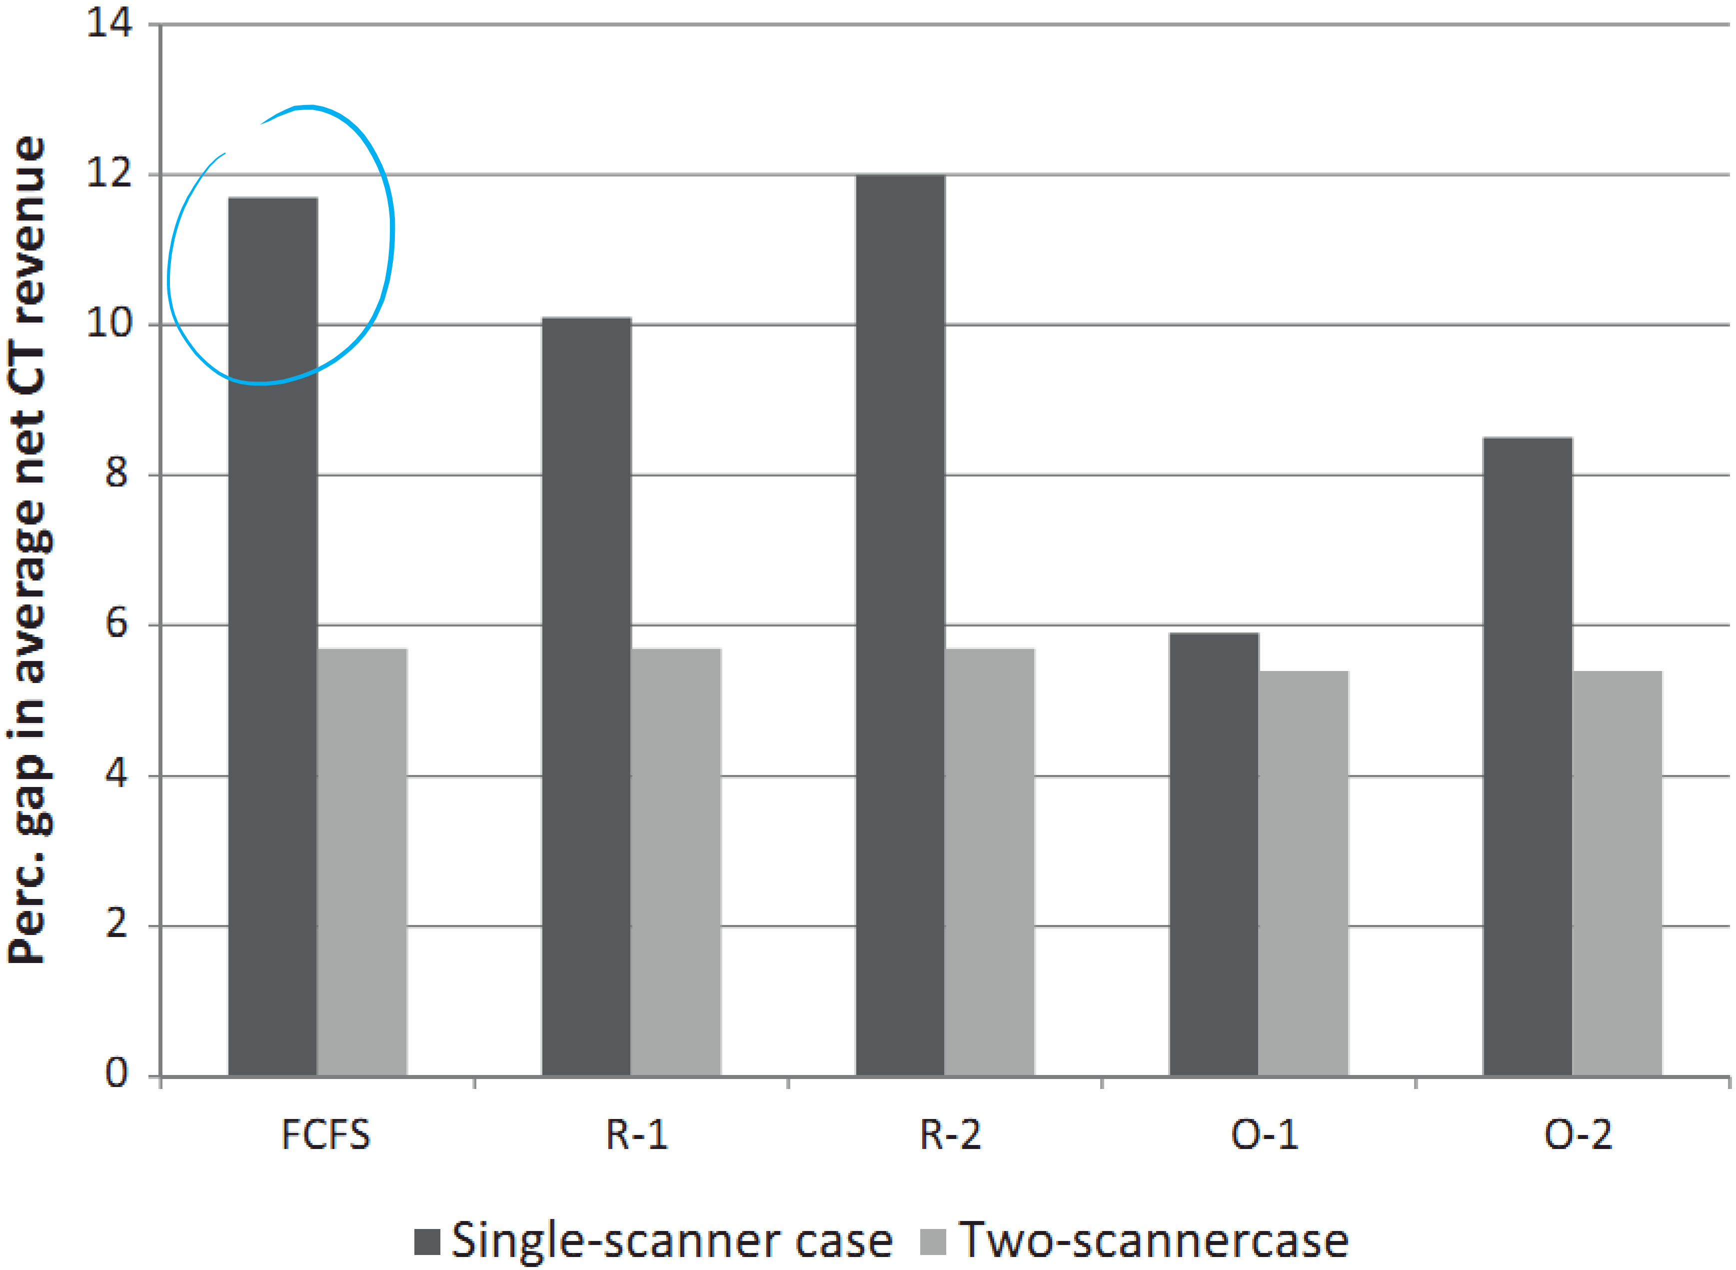

Perc. gap in average net CT revenue
14
12
10
8
6
4
2
0
FCFS
R-1
R-2
O-1
O-2
Single-scanner case
Two-scannercase

priority to OPs

Number of patients not scanned

Table 5
Number of patients not receiving scans by the end of the day under different policies, averaged over all thirty two scenarios.

Different cases	Average number not scanned					
	Optimal policy	FCFS	R-1	R-2	O-1	O-2
OPs						
Single-scanner	3.38	3.50	3.27	3.62	1.73	1.73
Two-scanner	0.72	0.63	0.52	0.64	0	0
IPs						
Single-scanner	10.13	9.97	10.57	9.85	12.01	11.14
Two-scanner	1.19	1.39	1.60	1.37	2.33	1.10
NCEPs						
Single-scanner	1.94	1.99	1.62	1.99	1.71	2.58
Two-scanner	0.51	0.39	0.29	0.41	0.08	1.31

Waiting-time

Table 6
Average waiting-time in minutes of patients before service over all thirty two scenarios.

Different cases	Average waiting-time					
	Optimal policy	FCFS	R-1	R-2	O-1	O-2
OPs						
Single-scanner	28	80	74	70	45	184
Two-scanner	3	4	3	4	0	0
IPs						
Single-scanner	76	112	95	107	60	245
Two-scanner	4	3	3	3	5	3
NCEPs						
Single-scanner	24	56	56	44	36	3
Two-scanner	12	9	8	10	3	20

Single-scanner

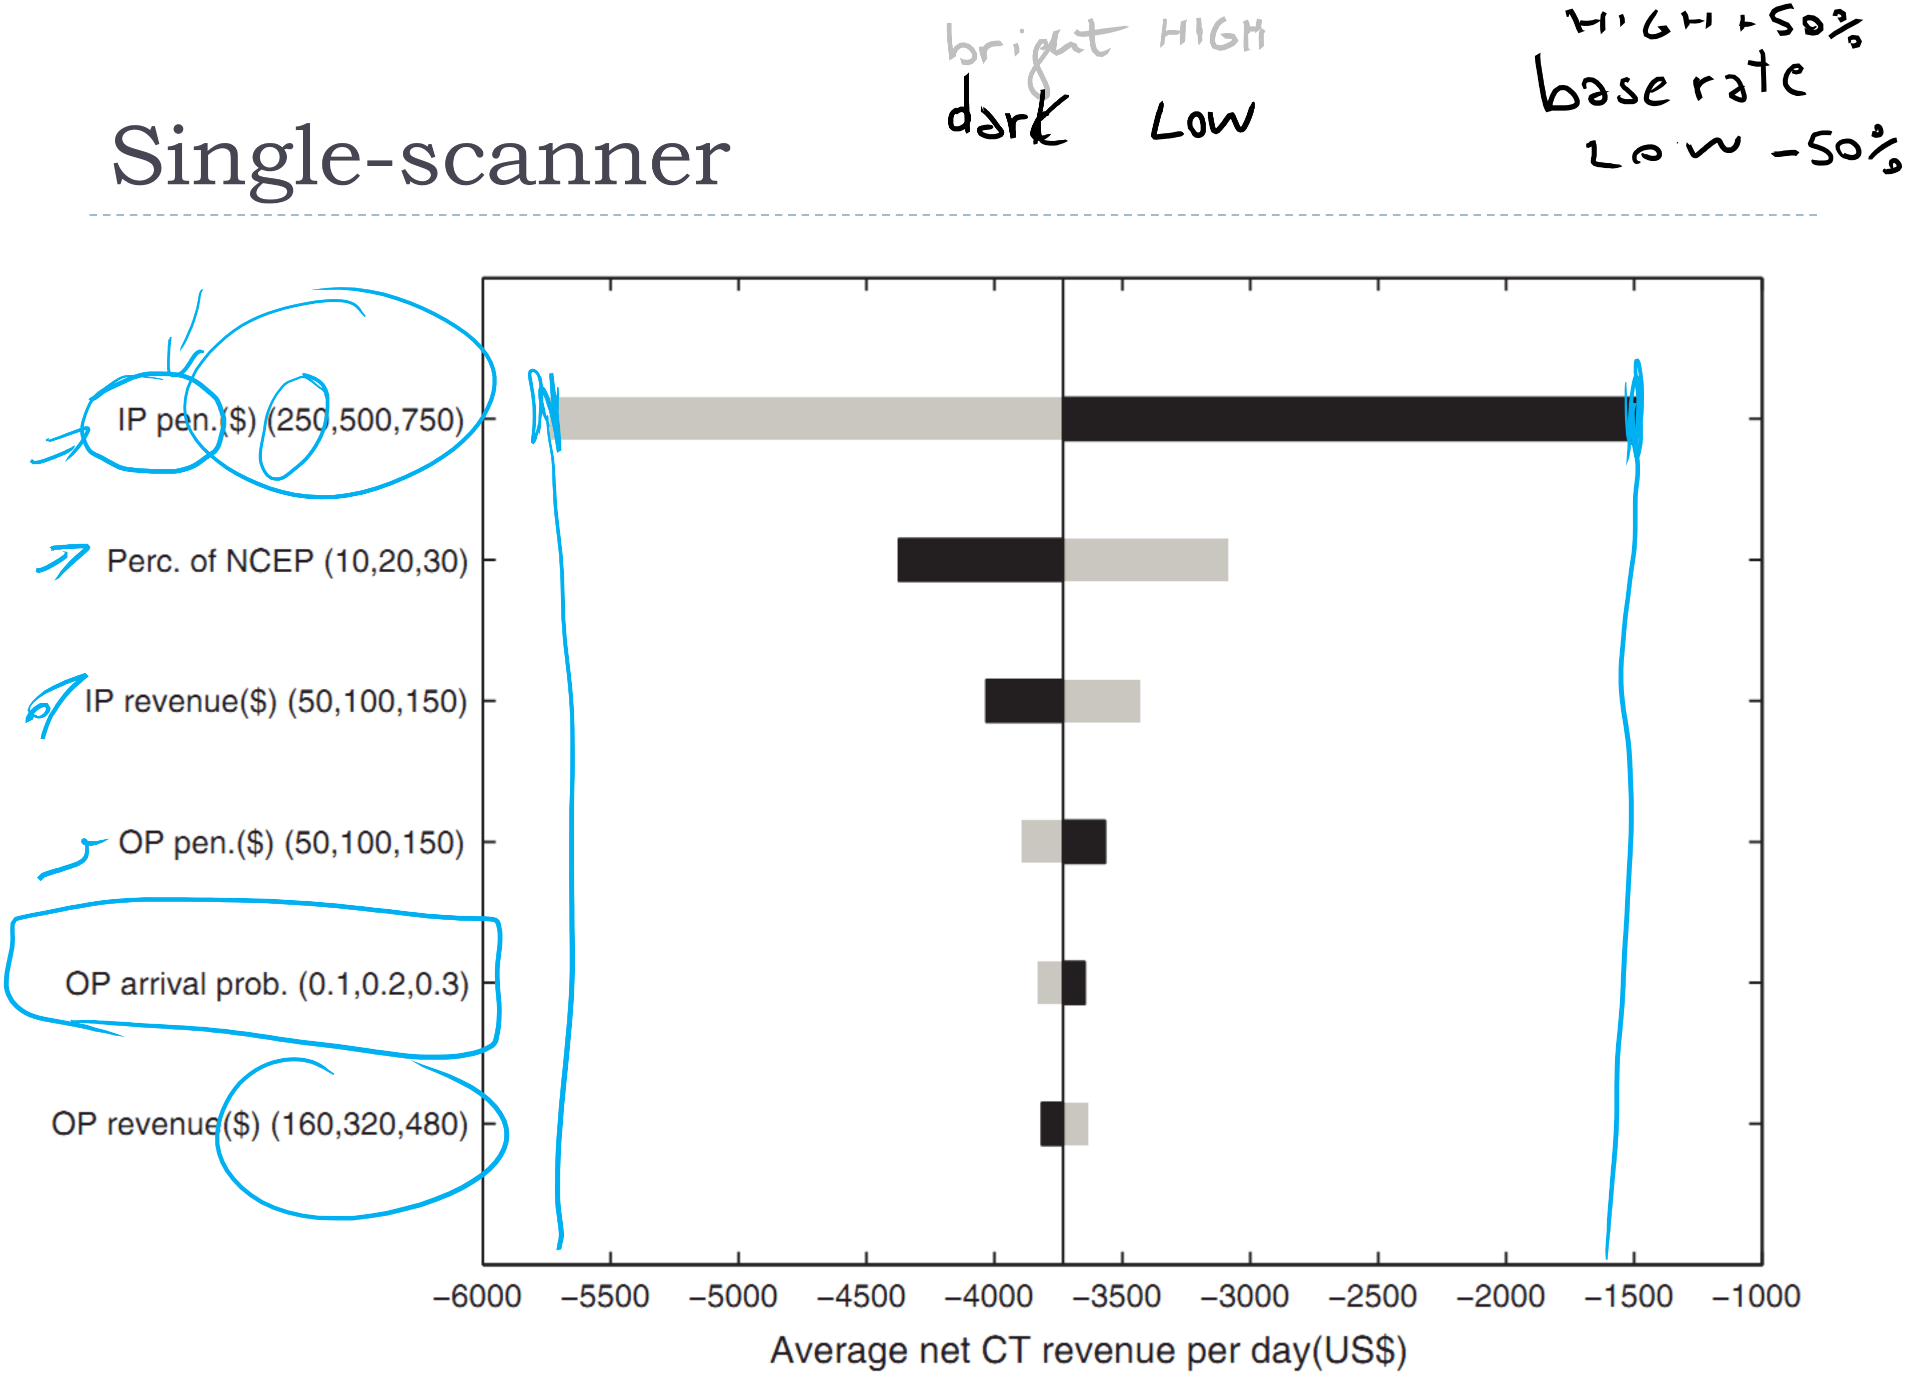

Two-scanner

bright HIGH
dark Low

HIGH +50%
base rate
Low -50%

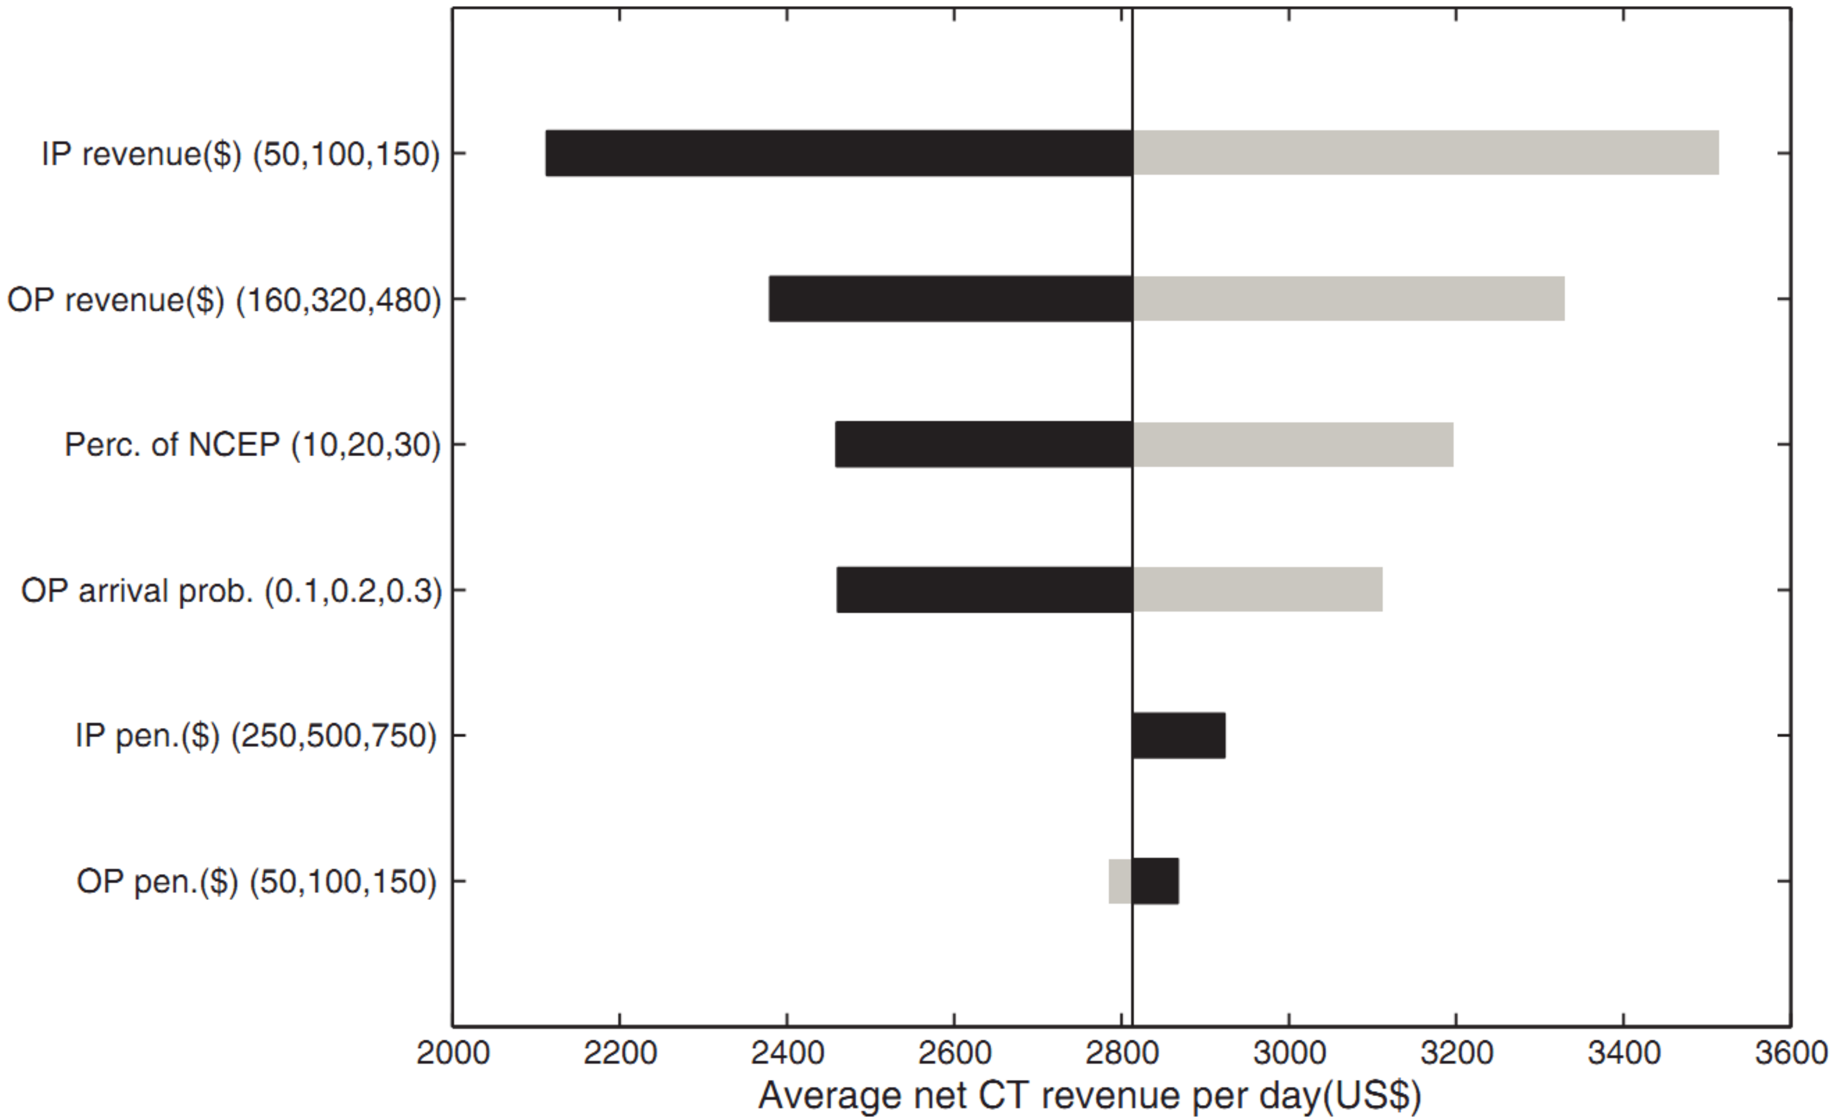

Sample Policy n=12, NCEP=5

CPE = 0

IPs →

OPs ↓

state {0, 0, 3, 5}

6	5	4	4	4	4	4	4	4	4	4	4
3	2	2	2	4	4	4	4	4	4	4	4
1	1	1	2	4	4	4	4	4	4	4	4
1	1	1	2	4	4	4	4	4	4	4	4
1	1	1	2	4	4	4	4	4	4	4	4
1	1	1	2	4	4	4	4	4	4	4	4
1	1	1	2	4	4	4	4	4	4	4	4
1	1	1	2	4	4	4	4	4	4	4	4
1	1	1	2	4	4	4	4	4	4	4	4
1	1	1	2	4	4	4	4	4	4	4	4
1	1	1	2	4	4	4	4	4	4	4	4
1	1	1	2	4	4	4	4	4	4	4	4

i-1 OPs, J-1 IPs

6 → scan two NCEPs

5 → scan one IP and one NCEP

4 → 〃 two IPs

3 → one OP and one NCEP

2 → one OP and one IP

1 → 〃 Two OPs

action to be performed in state {0, 0, 3, 5}

CPE OP IP NCEP

Question Types from students

- Finite vs. infinite
  - Simplicity. Lots of uncertainty about what can happen overnight
  - Non stationary process – best action depends on time
- Arrival Probabilities
- More scanners
- Modeling more patient types (urgency) / different hospital….. can easily extend the model, Only data from one Hospital (general?)
- Uniform slot length (realistic?)
- the probability distribution of the time for CT scans to be completed rather than to make the assumption that they are all of fixed duration? Finer granularity of the time slots
- Operational Cost of Implementing the policy (take into account): *compute the policy vs. apply the policy*
- Modeling even more uncertainty "Accidents happen randomly without any pattern." "Scanner not working"
- 2 patients at once (need to collect all the prob and consider those in the transition prob)
- P-value
- Why no VI?
- Used in practice ?

- Other models: Is it better to use continuous Markov Chain and queuing theory in analyzing this scheduling problem?
- How would this model handle two CEPs that came in at the same time? Randomly Push one to the next slot ☹
- How does approximate dynamic programming compare to value iteration? *(approximate method, can deal with bigger models but not optimal)*
- Transfer model to other facilities? *Yes…*
- Discount factor 1? *Yes*
- This work failed to take into account human suffering, or the urgency of of scans for in and out patients. Could the reward function to tailored to include such nebulous concepts or is it beyond the capabilities of the model?
- This model is specific to the target hospital
- I think outperforming other MDP-based models can better illustrate the effectiveness of this model's features, so are the choices of comparison methods good in this paper?
- *First step showing that sound probabilistic models can be build and outperform heuristics then you can do the above*