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

Matthew Dirks

Overview

- Solution
- Question: Ethics
- Algorithm
- Results & Analysis
- Questions

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)**

What if more than 2 CEPs arrive?

Proposed Solution

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

Performance Metrics (over 1 work-day)

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

Some discussion…

- They assume waiting NCEPs, IPs, and OPs are identical in terms of clinical urgency.
- Focus on $
- Good for the hospital
- What will people think?

Can there be any ethical problems
related to use of such a model in hospitals?

Is revenue a good metric for performance?
Especially if you consider that life and death might
depend on the scheduling results

Algorithm

- 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
- 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

Maximize total expected revenue

- Terminal reward obtained
  - $V_{N+1}(s) = -c_{OP}w_{OP} - c_{IP}w_{IP} - c_{NCEP}w_{NCEP}$
- Optimal Policy
  - Solving this gives the policy for each state, n, in the day

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

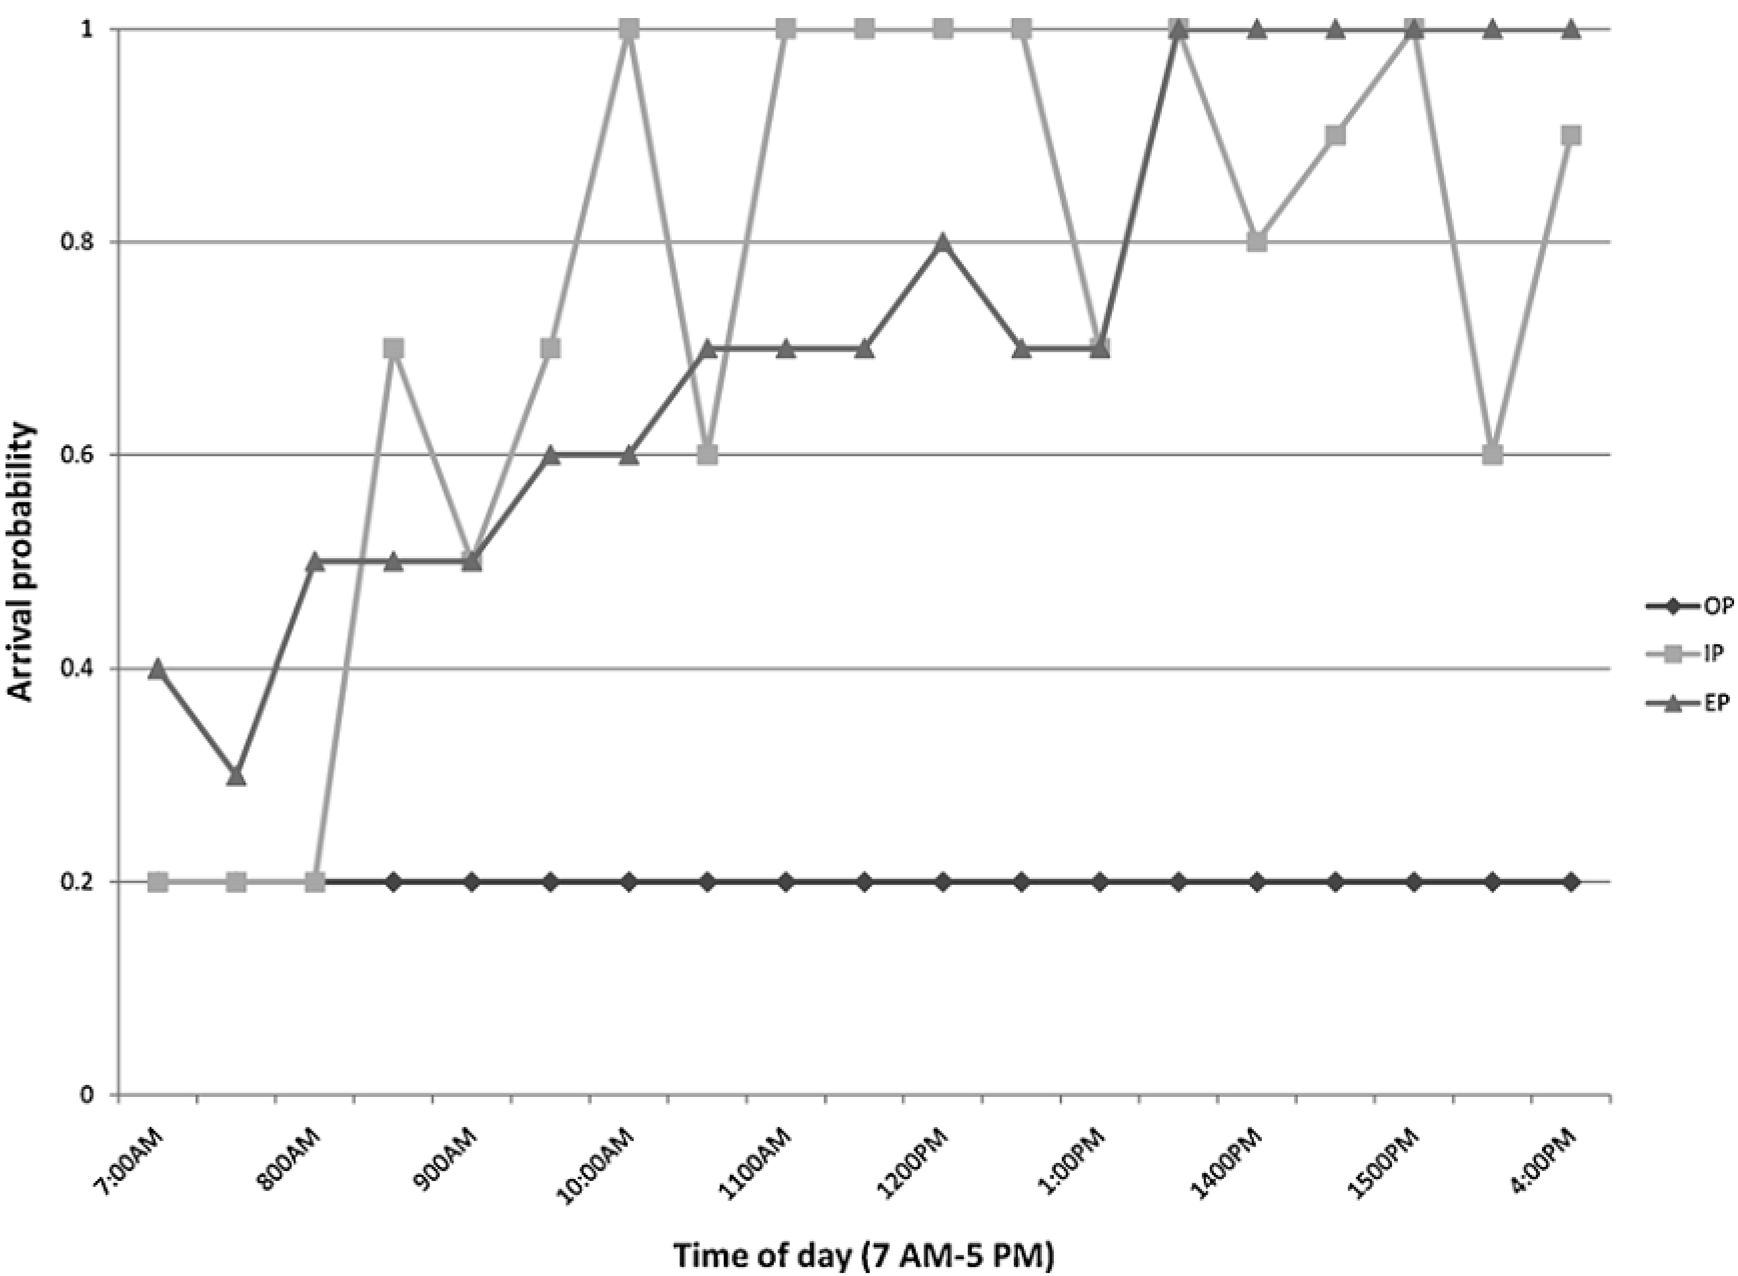

Fig. 1. Arrival probabilities for each patient-type during a work-day. EP includes both CEPs and NCEPs.

Simulation

- 100,000 independent day-long sample paths

Result Metric

- Percentage Gap in avg. net revenue =

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

- Closer you are to 100%, the better.
- 100% means that Heuristic got $0 revenue
- 75% means Heuristic resulted in 4 times less than (25% of) what Optimal got.

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

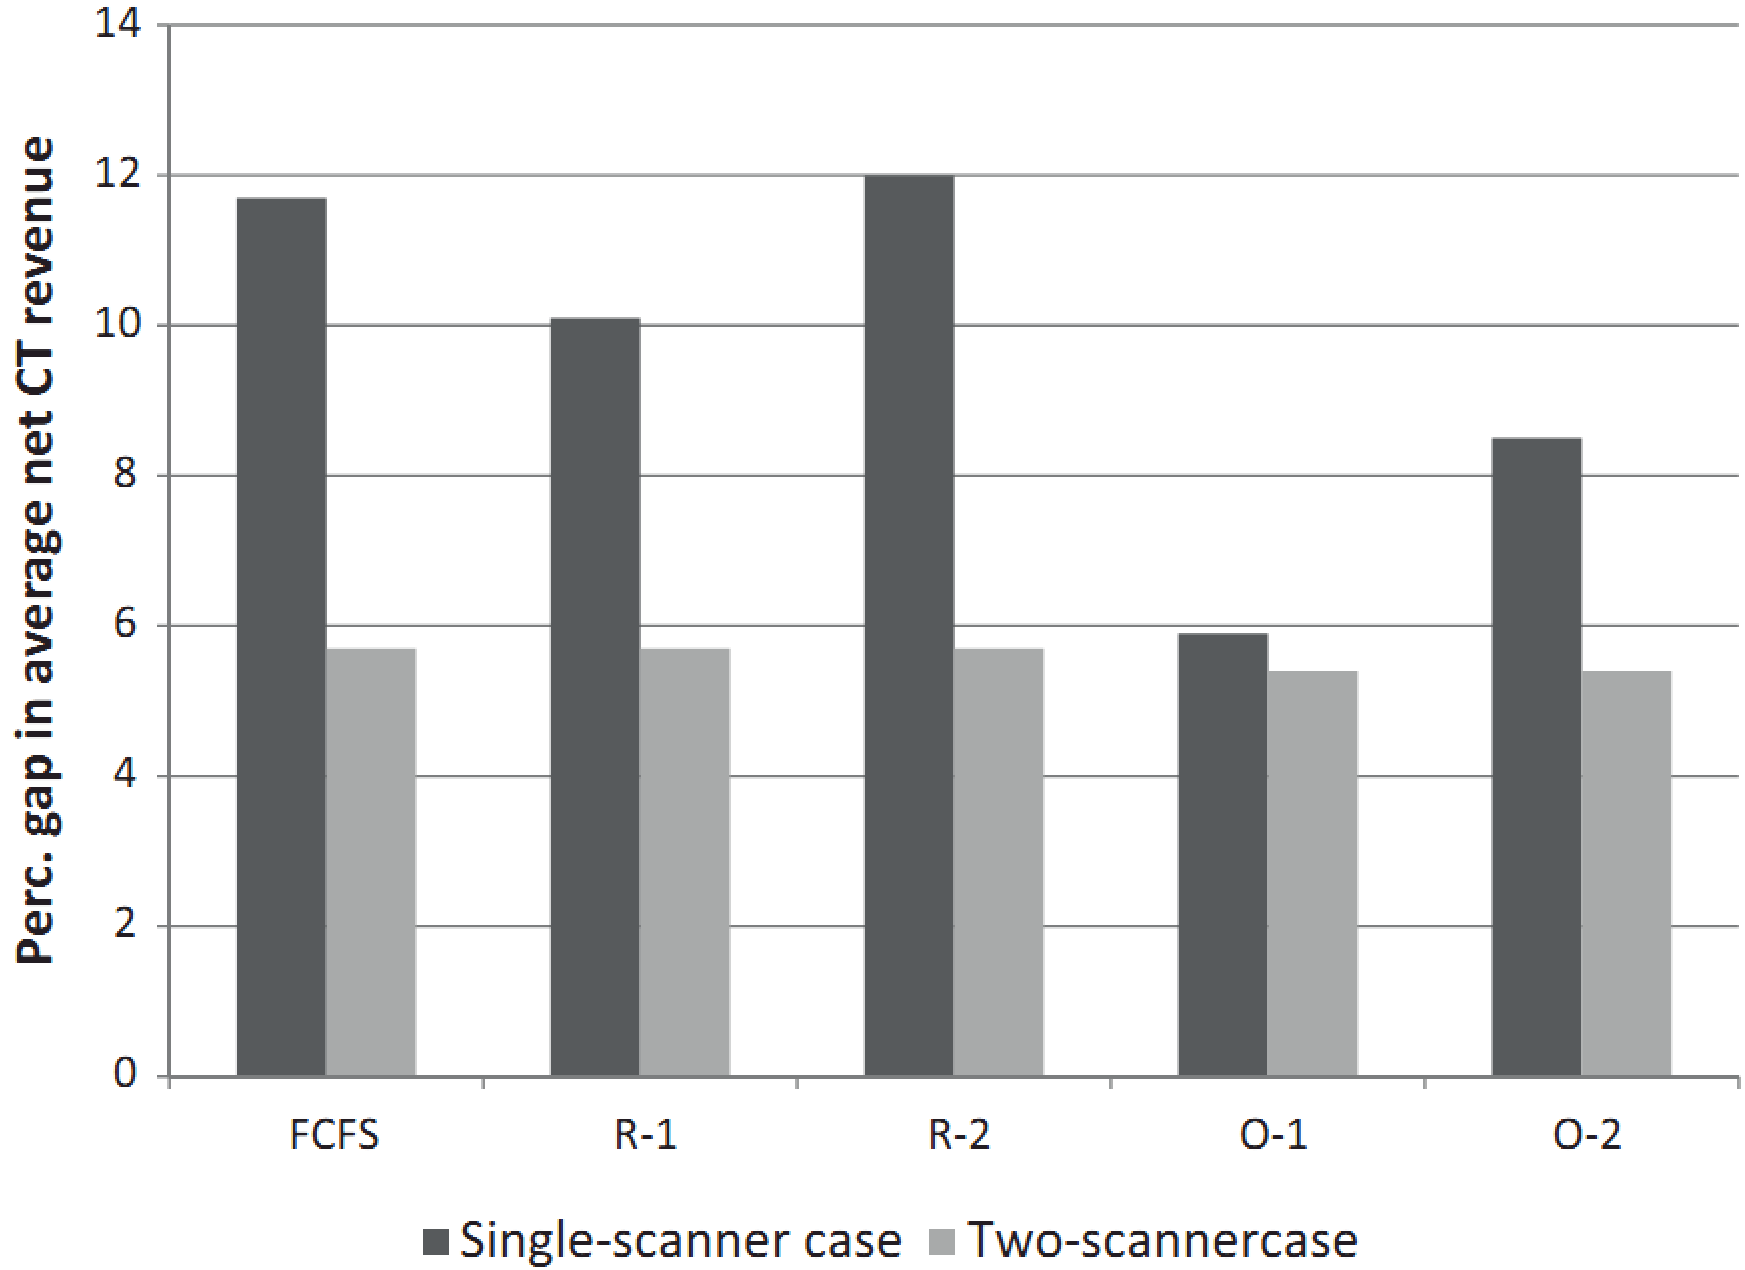
Perc. gap in average net CT revenue
14
12
10
8
6
4
2
0
FCFS
R-1
R-2
O-1
O-2
Single-scanner case
Two-scannercase

Single-scanner

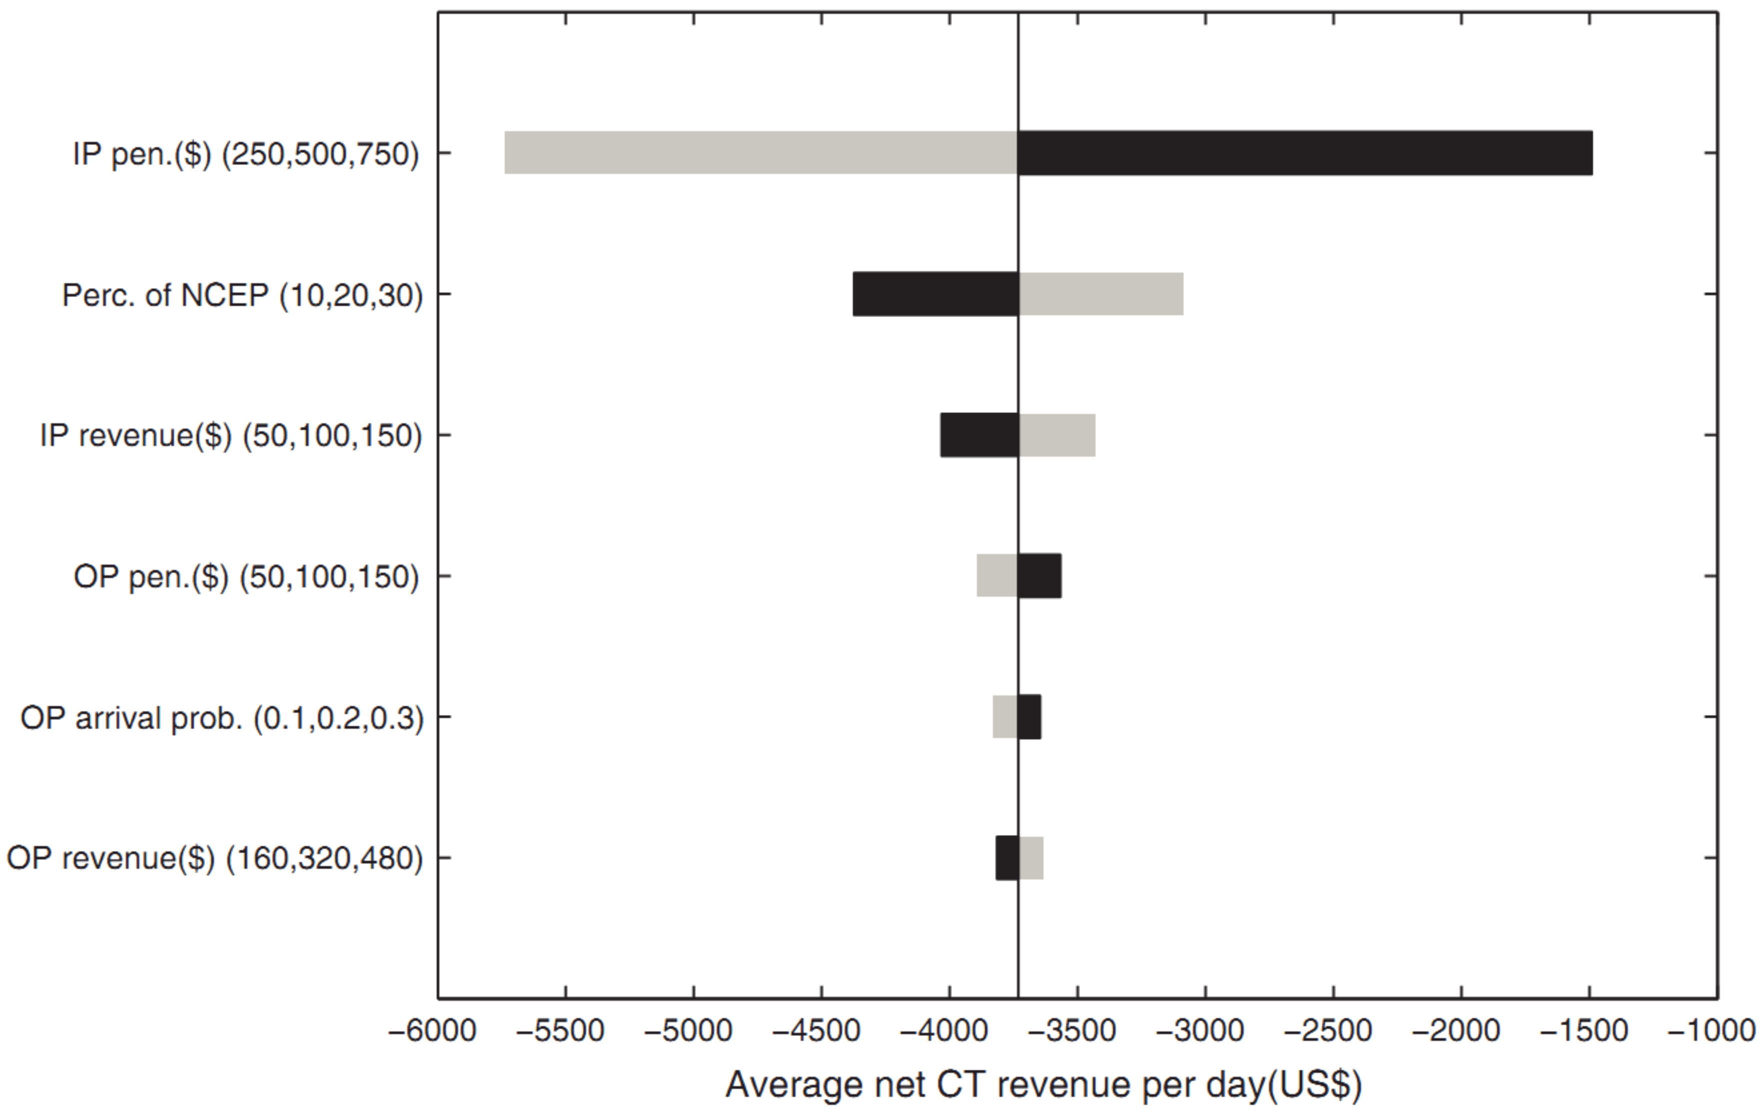

Two-scanner

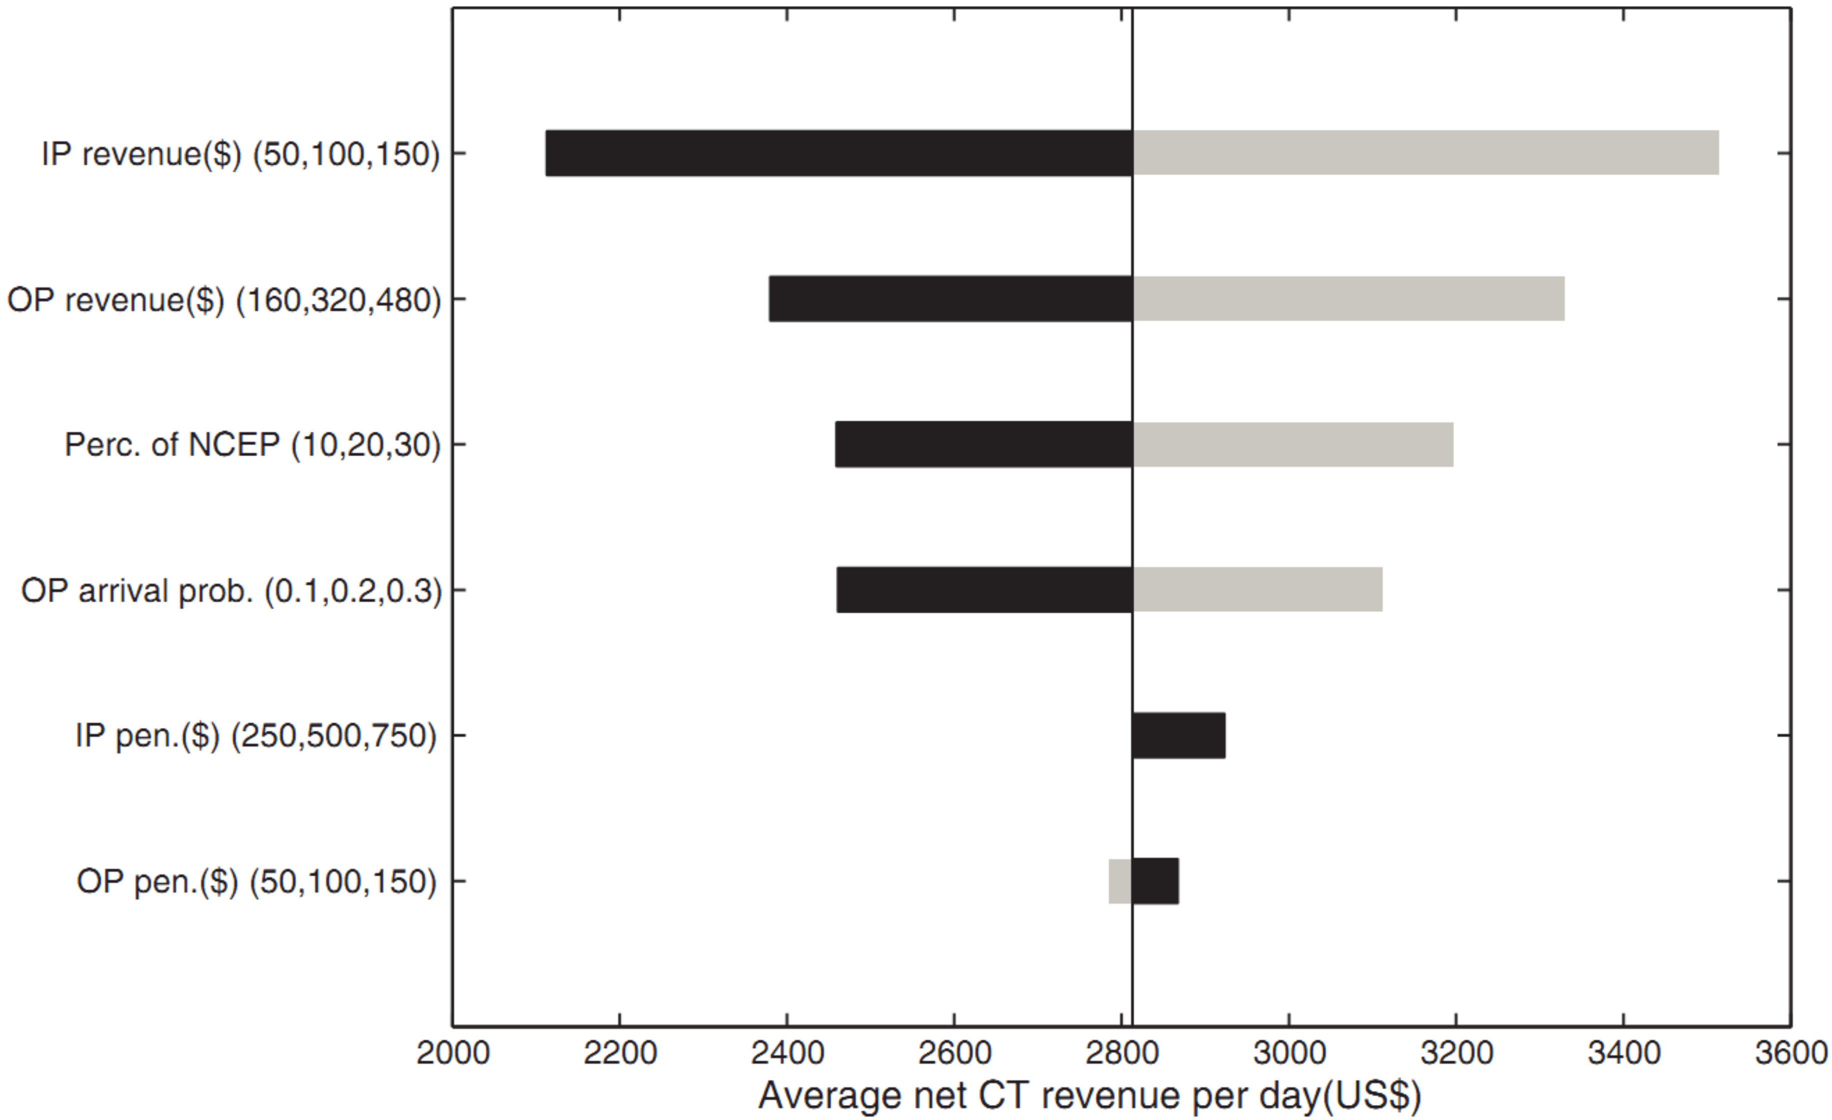
IP revenue($) (50,100,150)
OP revenue($) (160,320,480)
Perc. of NCEP (10,20,30)
OP arrival prob. (0.1,0.2,0.3)
IP pen.($) (250,500,750)
OP pen.($) (50,100,150)
2000
2200
2400
2600
2800
3000
3200
3400
3600
Average net CT revenue per day(US$)

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

Questions

- Why do they try to make the problem so specific?

- This is a scheduling problem with some constraints and an objective function that occurs in many scenarios. Eg. Scheduling multi-category rooms (project-equipped, conference-capable, etc.) and scheduling these rooms given a series of streaming tasks or scheduling different kinds of cars for different requests. How about making a generic framework with some specific parameters that can cater to a variety of problems?

- Their focus is on revenue cost specifically of CT scans. They wanted to test how tweaking the parameters within this specific case could affect revenue. A more general method would not provide specific dollar values about increased/saved revenue.

Questions

- Is it appropriate to use the Markov assumption for this scheduling problem?

- Perhaps we can iteratively improve their schedule after every day, especially if there are some OPs who did not show?

Questions

- The authors acknowledge that simple heuristic policies perform nearly on-par with their relatively complex MDP-derived policies. In practice, how often do medical providers use “smart” policies like the MDP vs. the simpler heuristic policies?

Questions or Comments?