Nudge: Stochastically Improving Upon FCFS

Isaac Grosof – CMU

Kunhe Yang – Tsinghua University

Ziv Scully – CMU

Mor Harchol-Balter – CMU

M/G/1 Scheduling

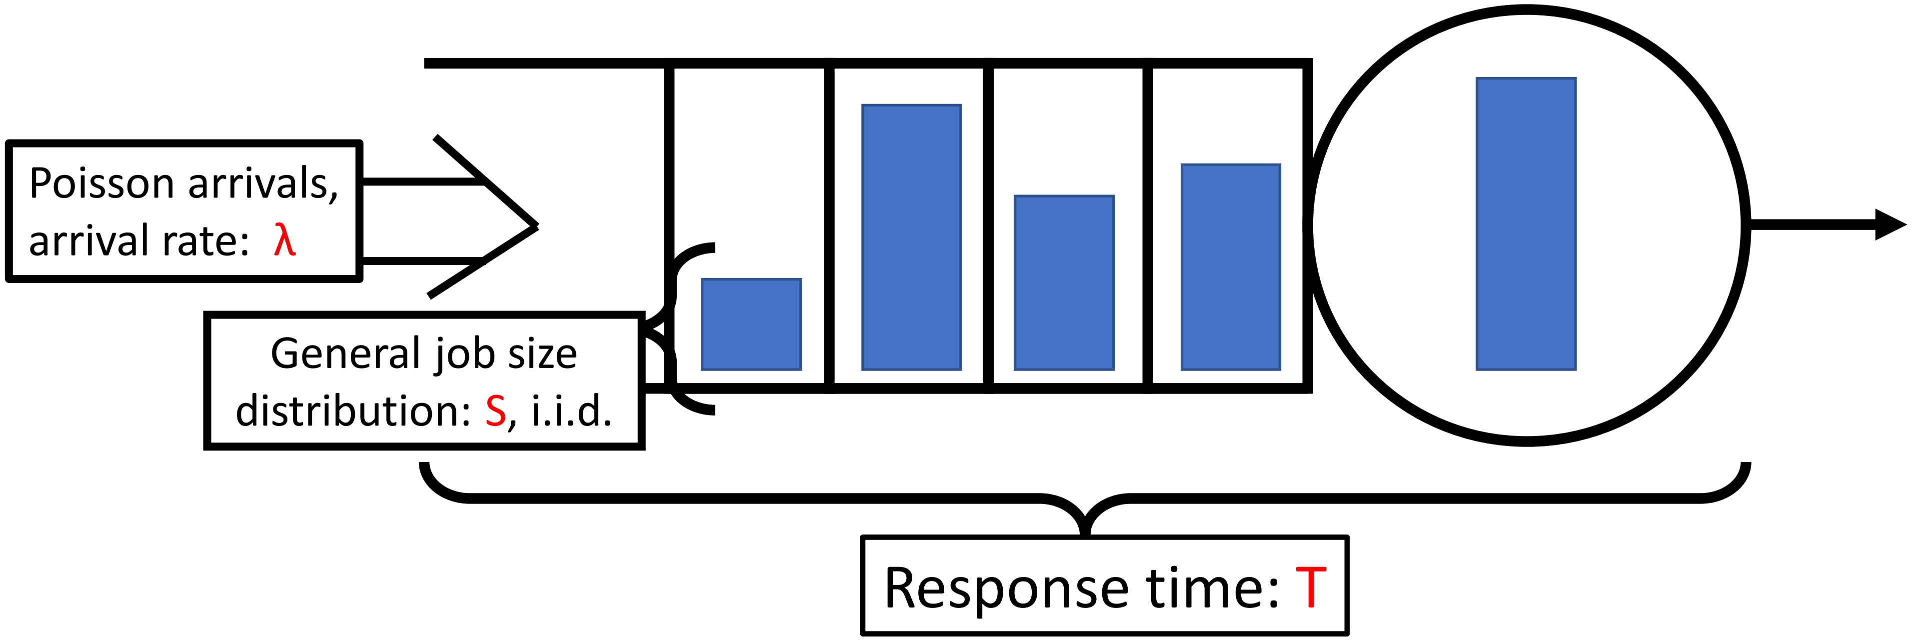

Q: How should we schedule?

Baseline: First-Come First-Served

FCFS: Simple, practical, good theoretical properties

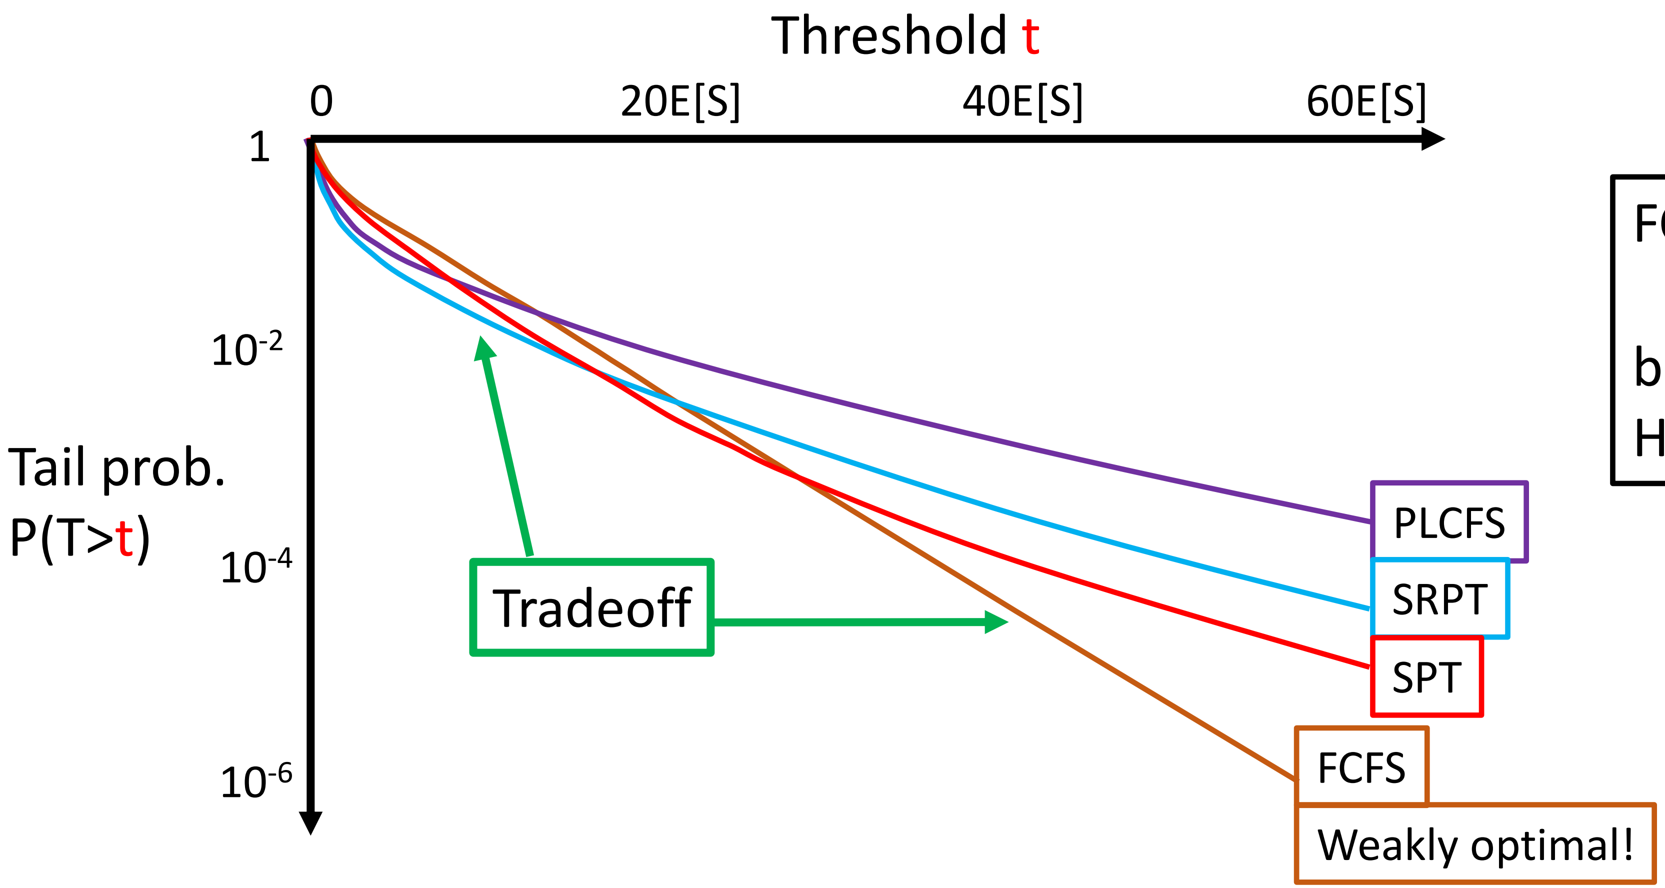

FCFS is weakly optimal:

$$P(T^{FCFS} > t) \sim Ce^{-\theta t},$$

best possible θ.

Holds whenever S is light-tailed.

Setting: S ~ Hyperexponential with C^2 =3, ρ=0.4

FCFS versus SRPT

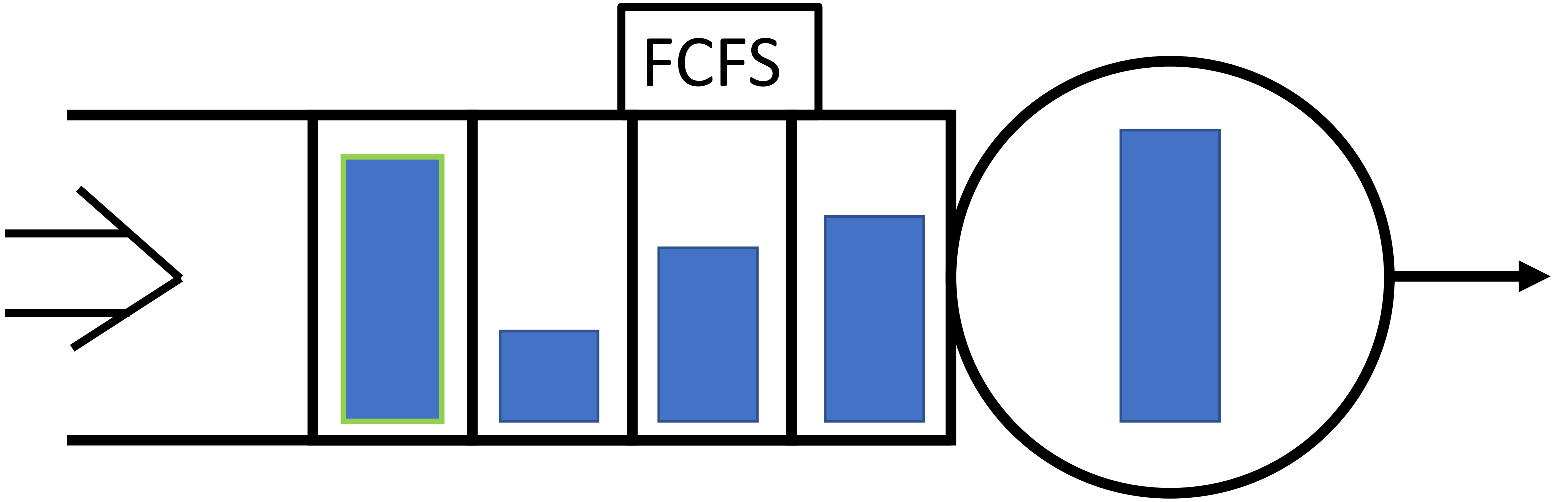

FCFS: Large or small, similar response time

FCFS versus SRPT

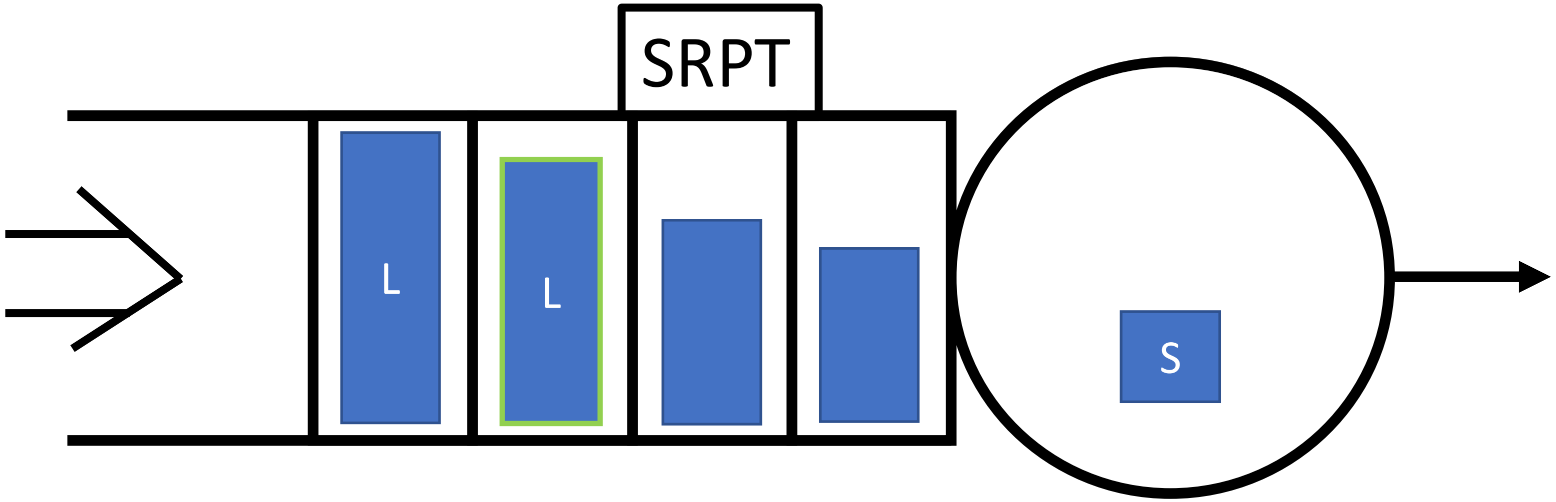

FCFS: Large or small, similar response time

FCFS versus SRPT

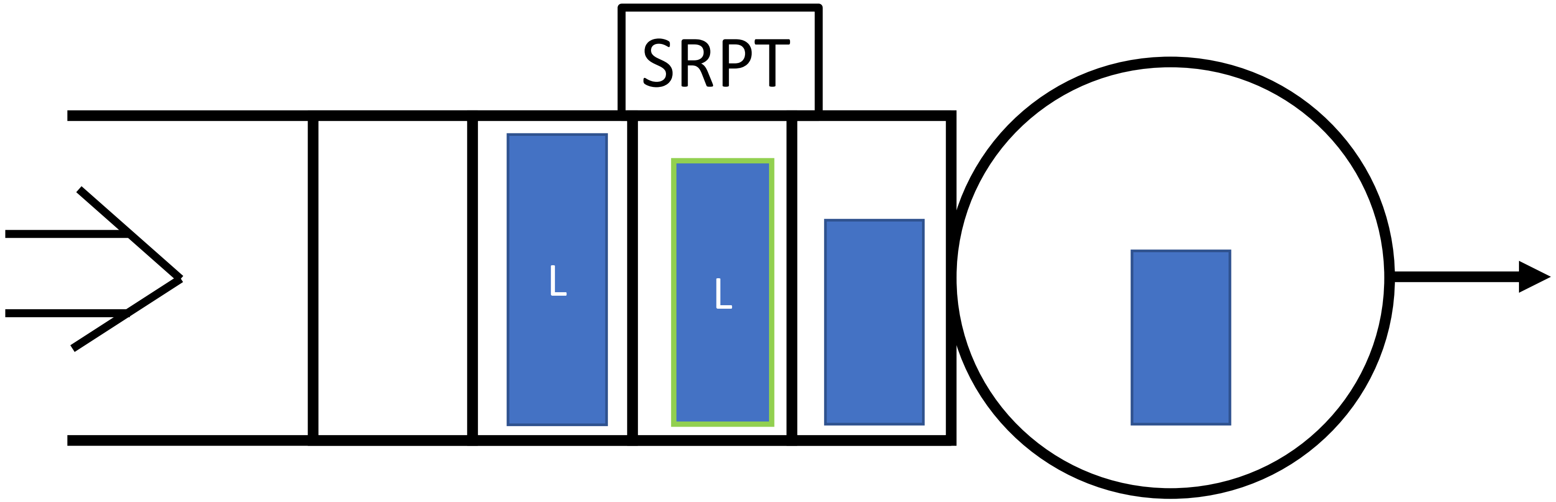

FCFS: Large or small, similar response time

FCFS versus SRPT

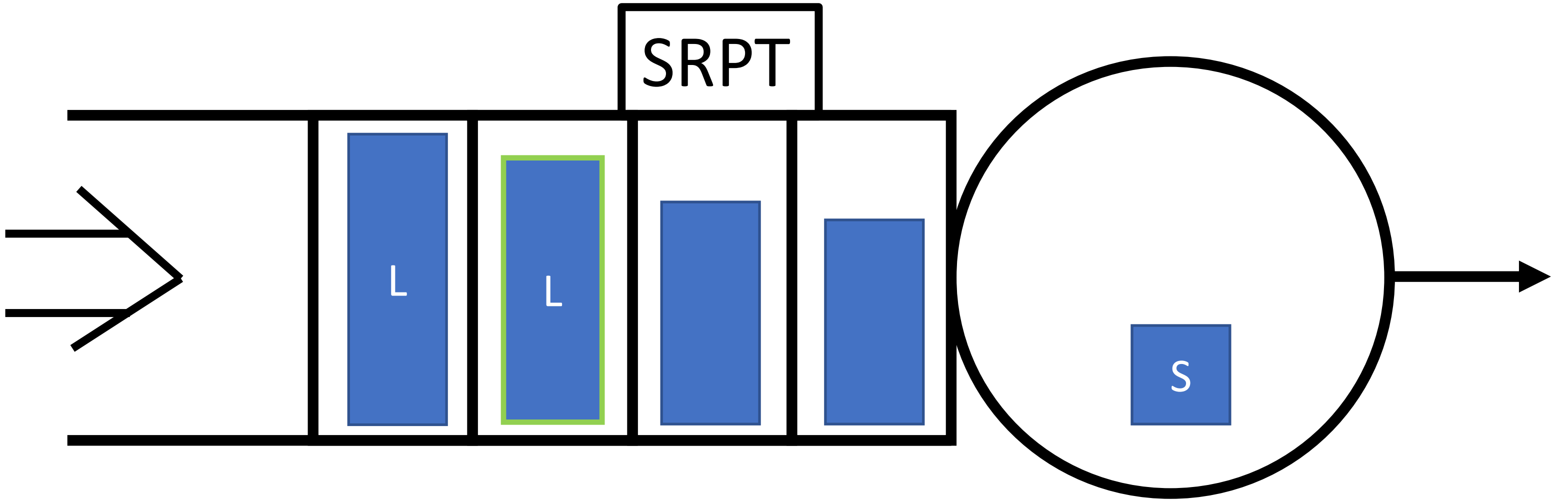

FCFS: Large or small, similar response time

FCFS versus SRPT

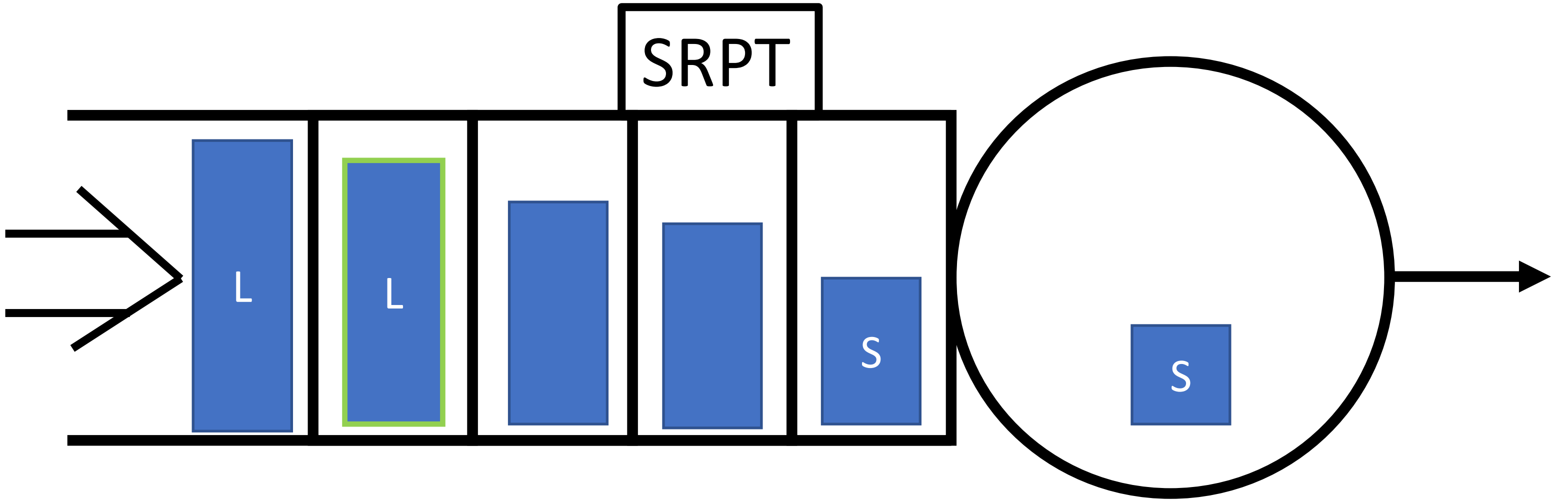

FCFS: Large or small, similar response time

SRPT:

- Helps small jobs, better P(T>t) for small t.
- Delays large jobs, worse P(T>t) for large t.

Fundamental question of tradeoffs

All previous policies have tradeoffs:
Better than FCFS at small t, worse than FCFS at large t.
Is that inevitable?

Is it possible to beat FCFS everywhere? ($\forall t$)

Yes, with Nudge!

Nudge: Stochastic Improvement

We introduce a new policy: Nudge

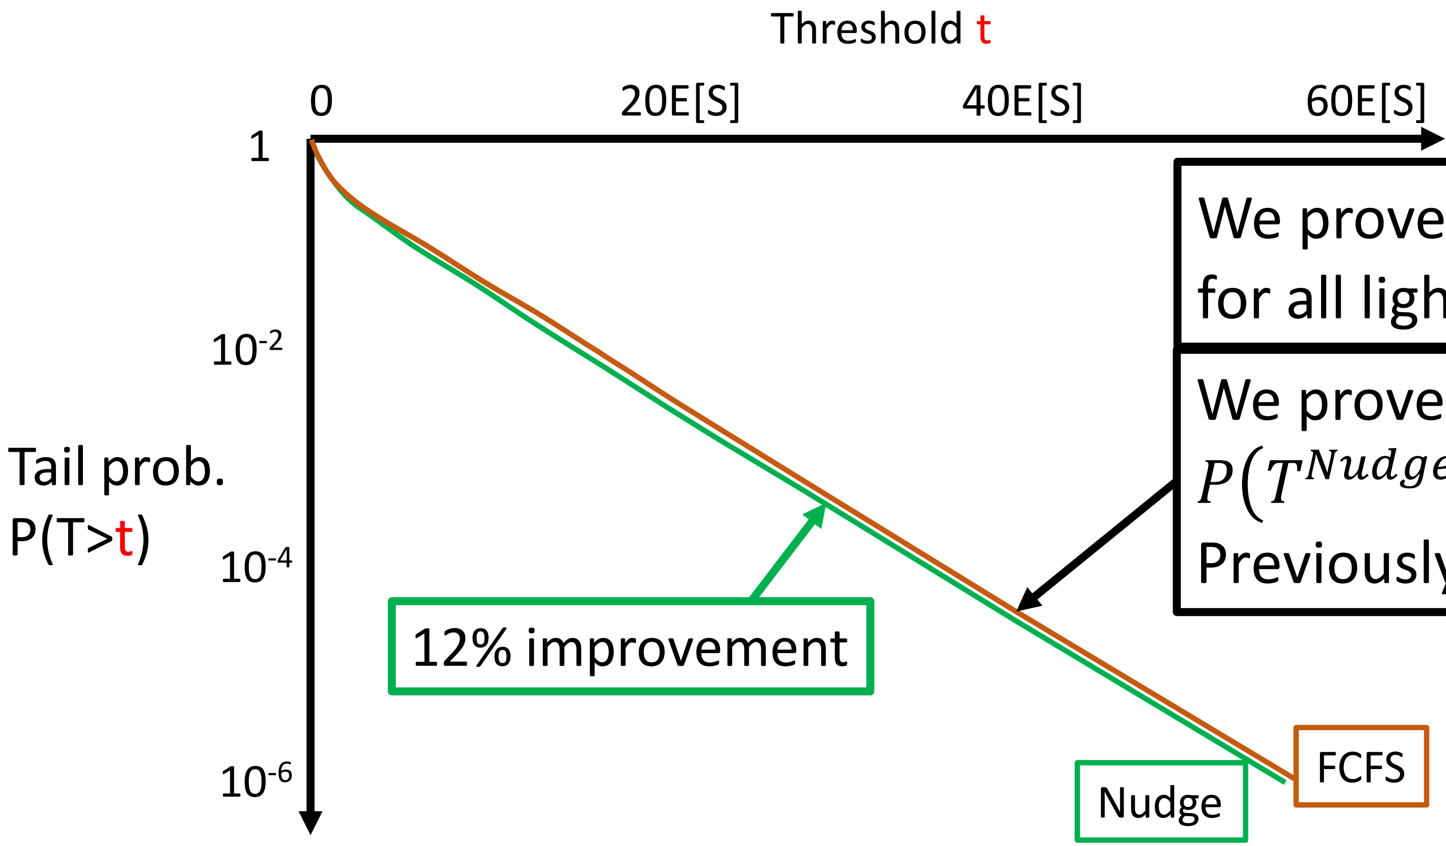

We prove: $P\left(T^{Nudge} > t\right) < P(T^{FCFS} > t)\ \forall t$,
for all light-tailed job size distributions S.

We prove Nudge has better asymptotic tail:
$P\left(T^{Nudge} > t\right) \sim C' e^{-\theta t}$. Same θ, better C′.
Previously conjectured to be impossible [WZ '12]

Setting: S ~ Hyperexponential with $C^2 = 3$, ρ=0.4

Our contribution: Nudge

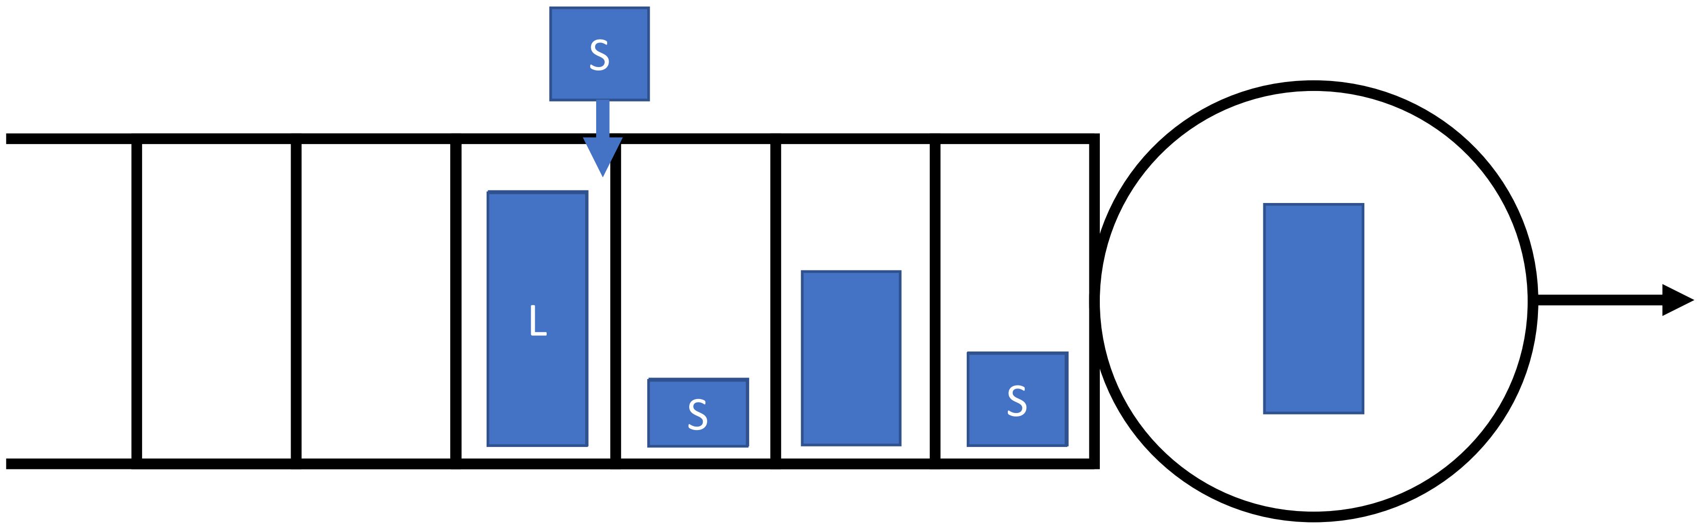

- Default: FCFS
- Classify jobs as small or large by size
- When small arrives, if large is last in queue, small nudges ahead of large

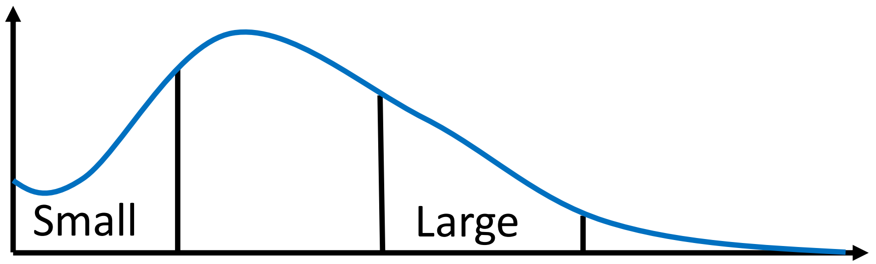

Job size distribution

Our contribution: Nudge

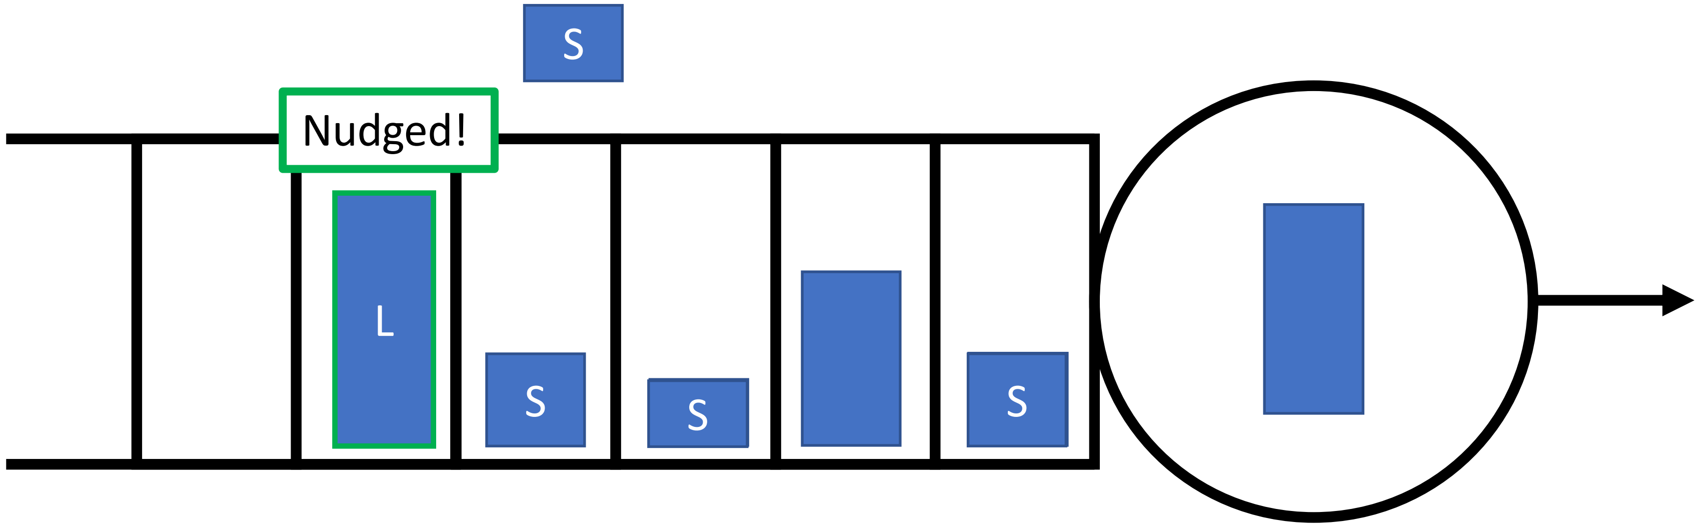

- Default: FCFS
- Classify jobs as small or large by size
- When small arrives, if large is last in queue, small nudges ahead of large
- Large can only be nudged once.

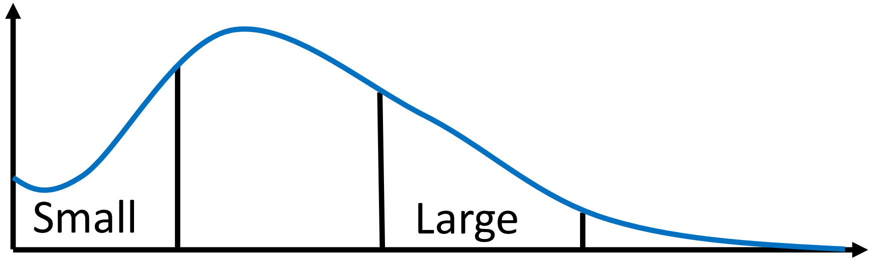

Job size distribution

Our contribution: Nudge

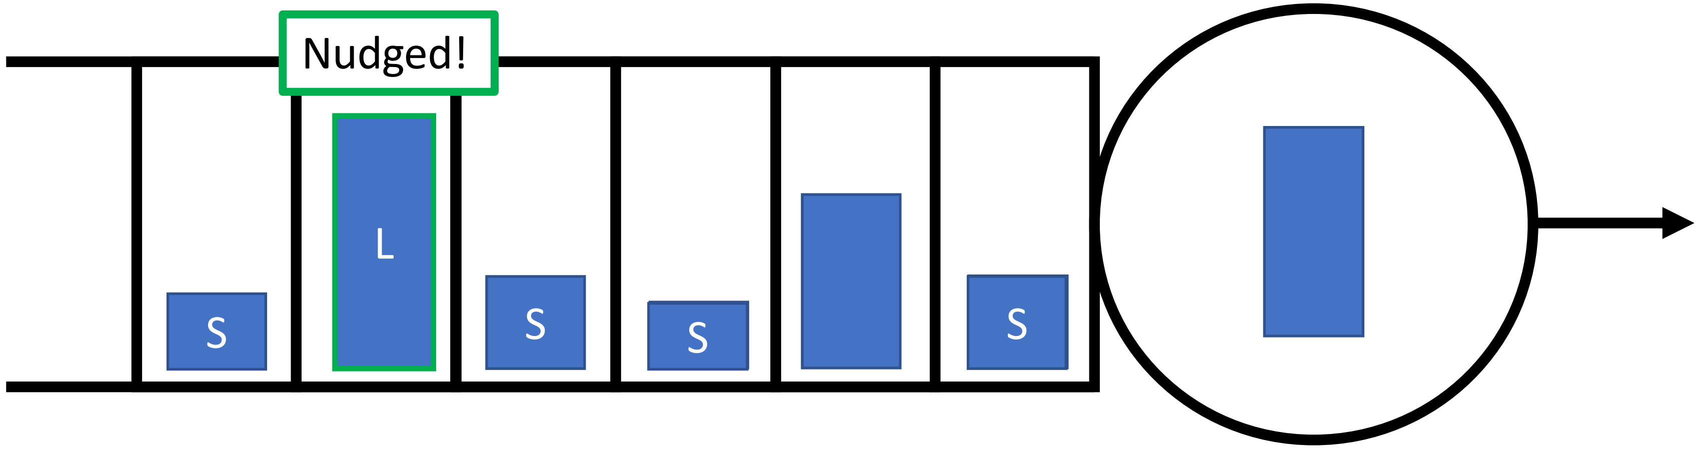

- Default: FCFS
- Classify jobs as small or large by size
- When small arrives, if large is last in queue, small nudges ahead of large
- Large can only be nudged once.

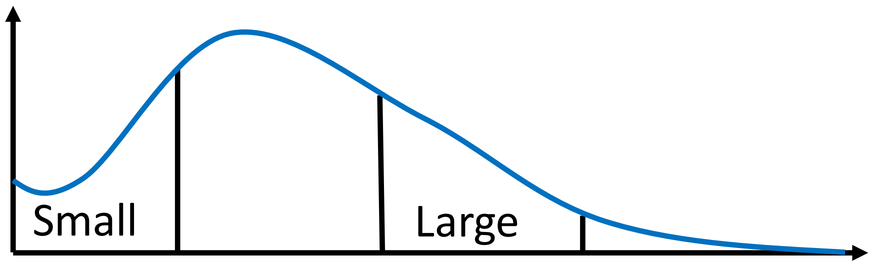

Job size distribution

Nudge intuition: Jack and the Giant

Nudge policy: Jack gets to go ahead of the giant,

Nudge intuition: Jack and the Giant

Nudge policy: Jack gets to go ahead of the giant, but nobody else gets to.

Proof intuition

Want to show:
$$P\left(T^{Nudge} > t\right) < P\left(T^{FCFS} > t\right)\ \forall t$$

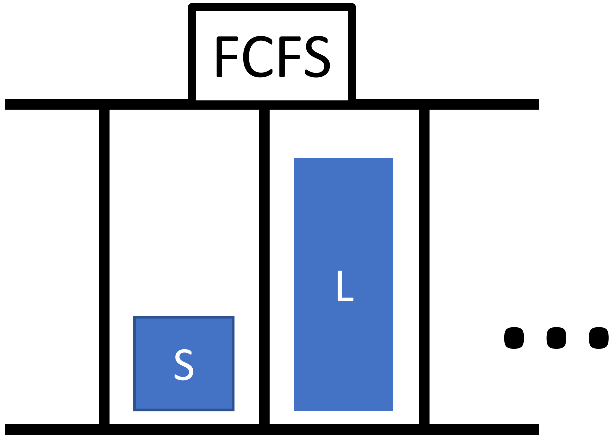

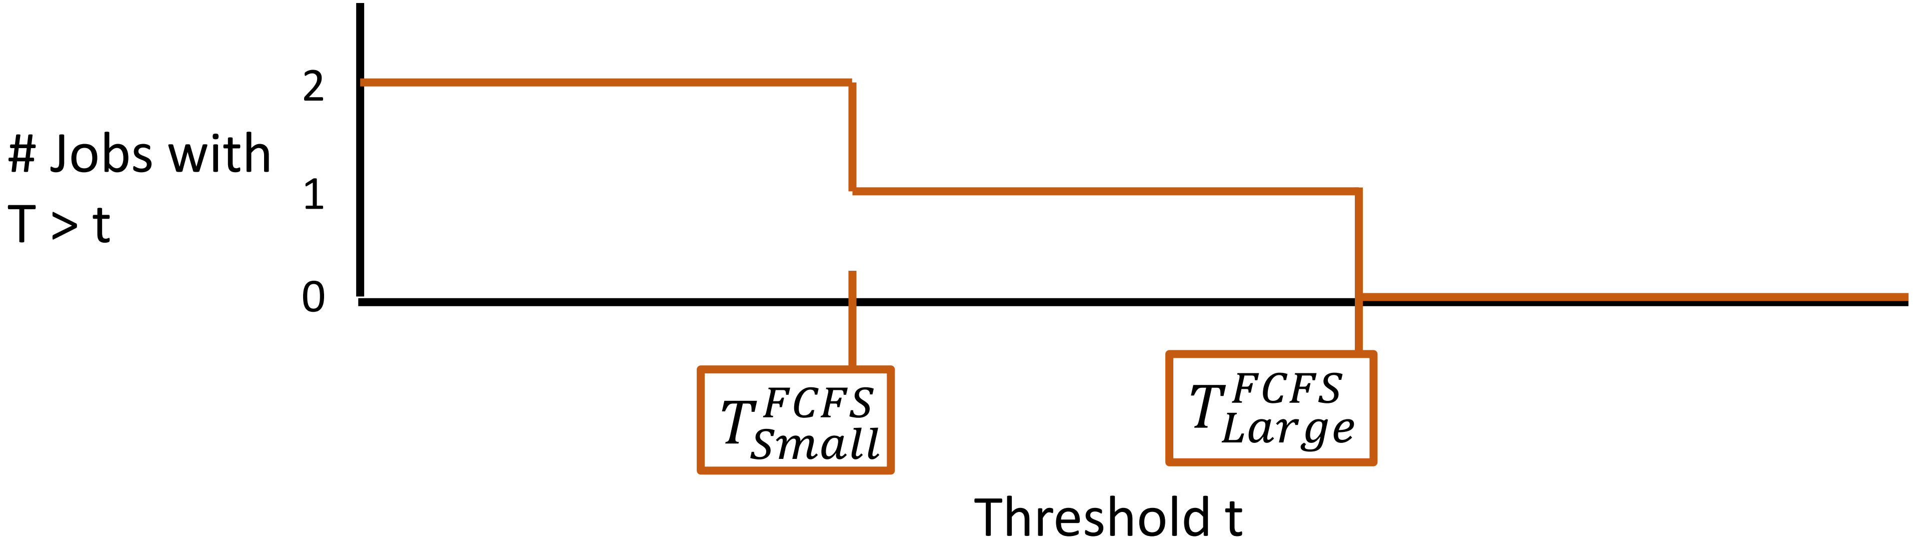

Proof intuition

Want to show:
$$P\left(T^{Nudge} > t\right) < P\left(T^{FCFS} > t\right)\ \forall t$$

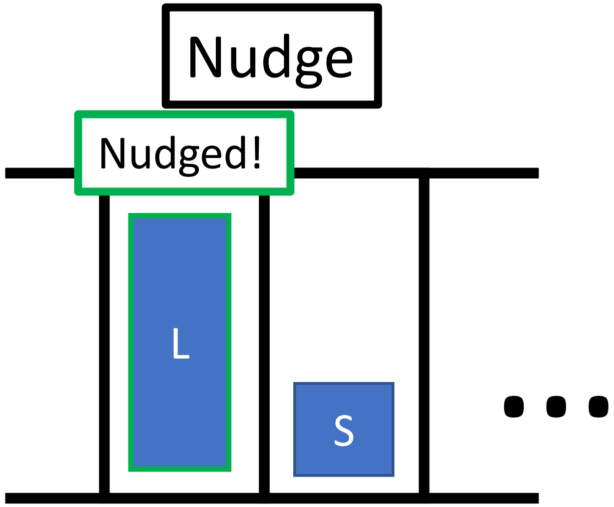

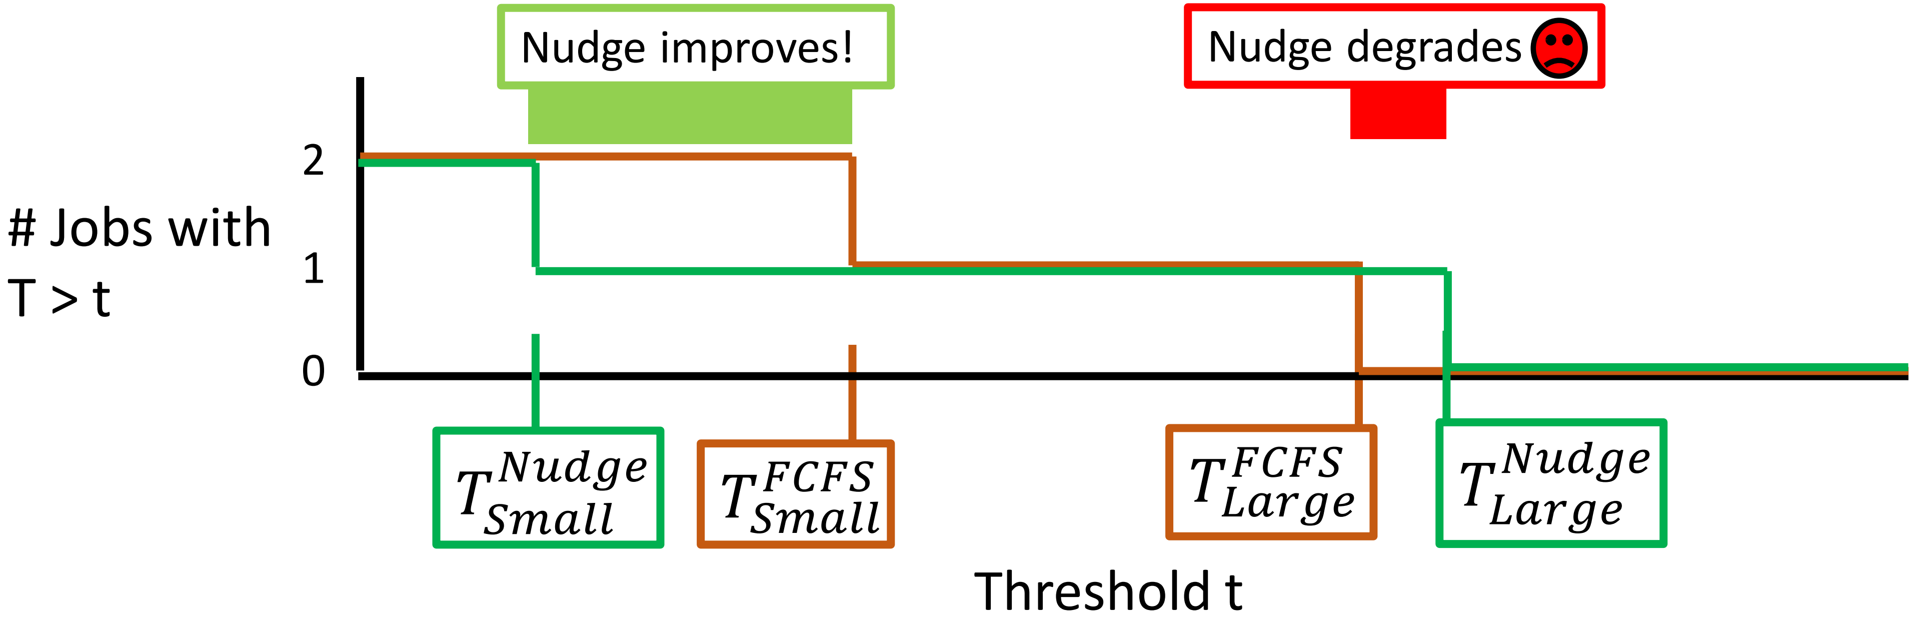

Proof intuition: One t, many nudges

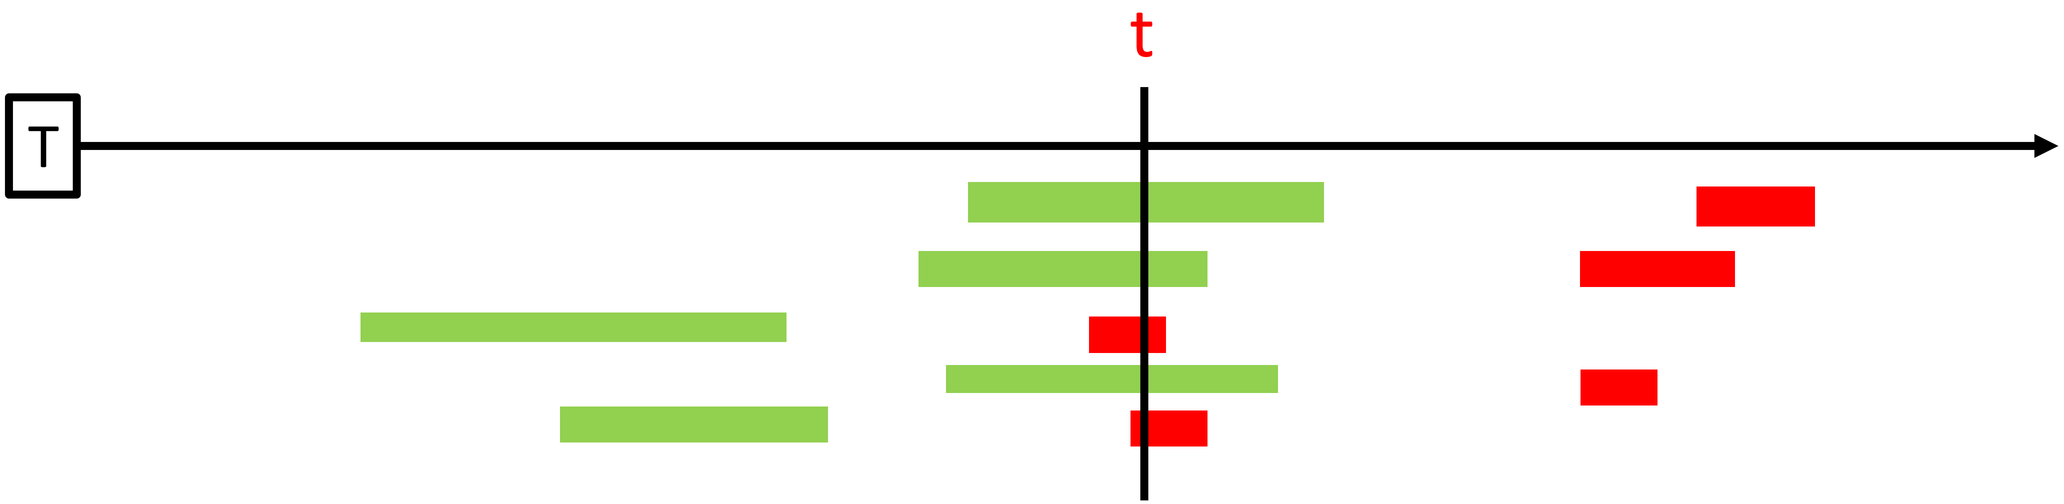

Want to show: Rate of improve exceeds degrade.

- Key idea 1: Rates of improve and degrade determined by FCFS tail and pdf.
- Key idea 2: Bound FCFS tail and pdf relative to limiting exponential.
- Bounds imply more improve than degrade, given correct small and large cutoffs.

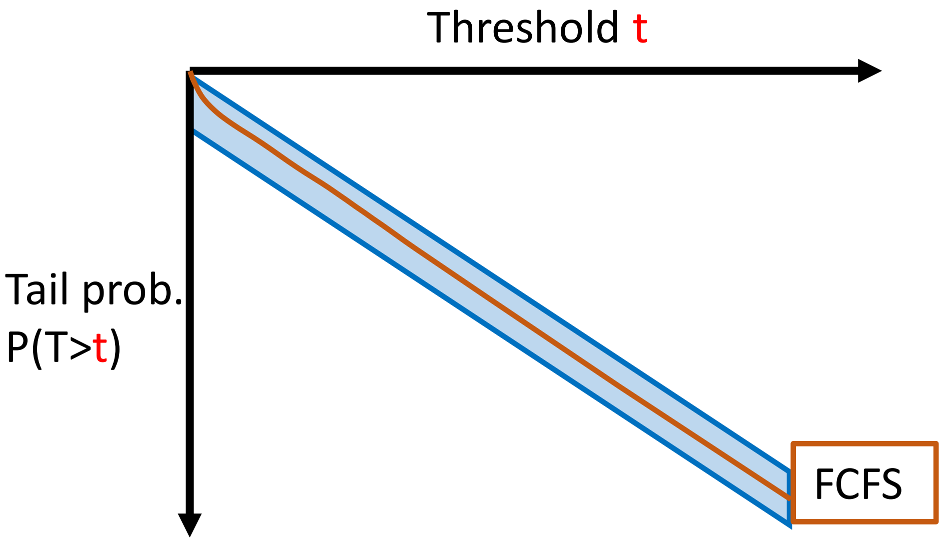

Empirical results

Tail improvement ratio:

$$1 - \frac{P\left(T^{Nudge} > t\right)}{P\left(T^{FCFS} > t\right)}$$

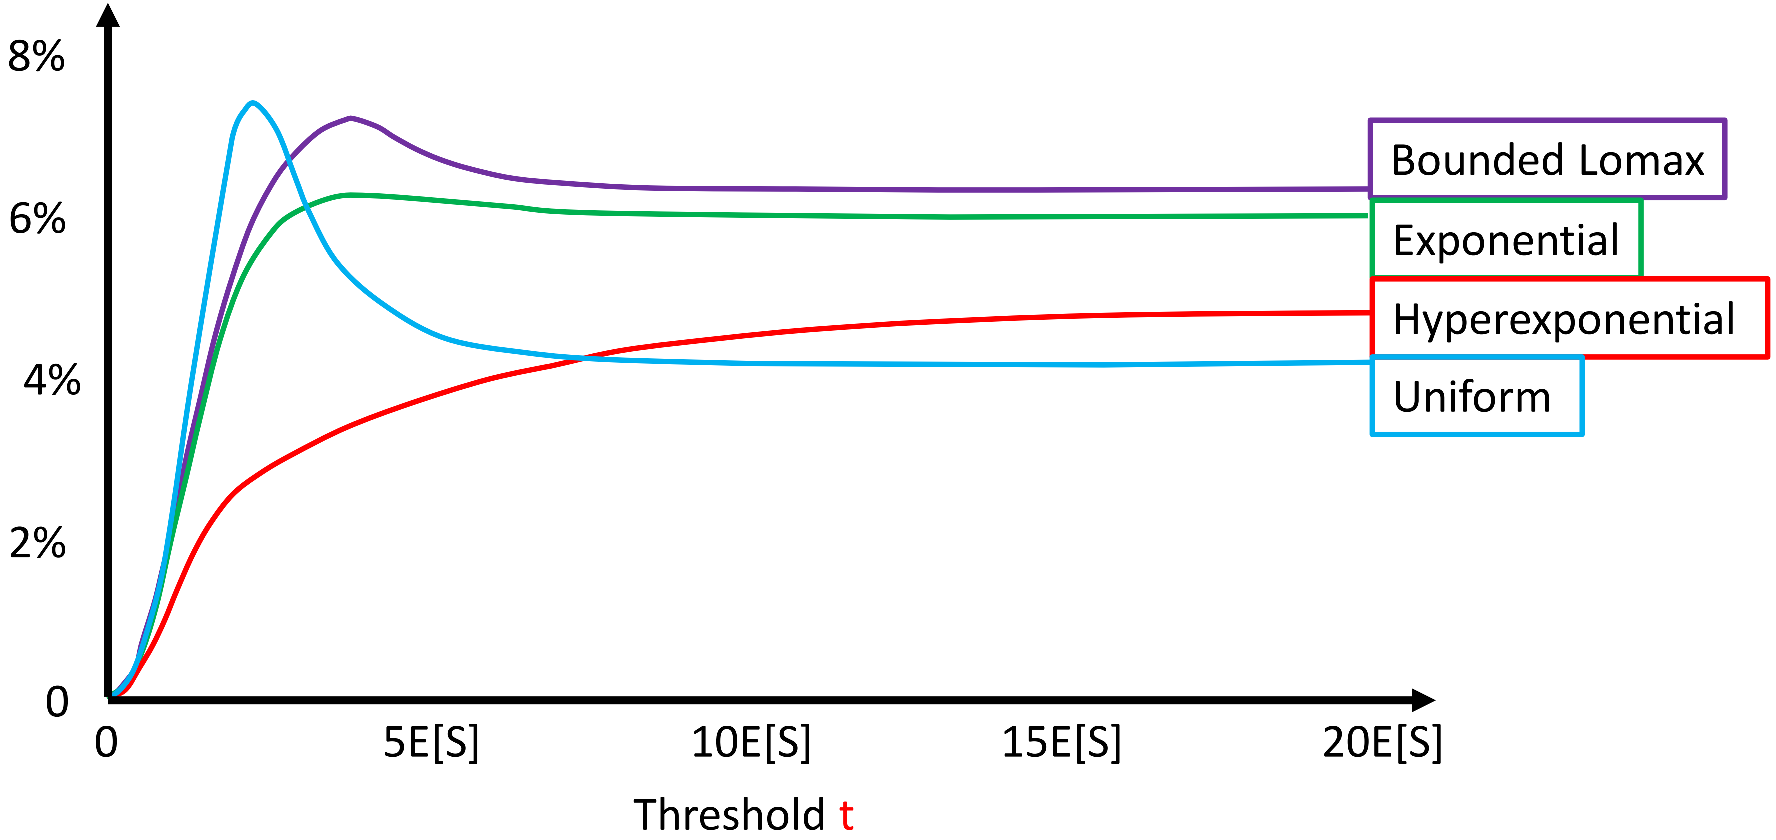

Load ρ=0.8

Future directions

- What about 2+ swaps per job? Constant # is important.
- What range of size cutoffs work?
  - Single threshold, all jobs either large or small?
- Beyond FCFS, what other policies can be improved everywhere?

Current cutoffs:

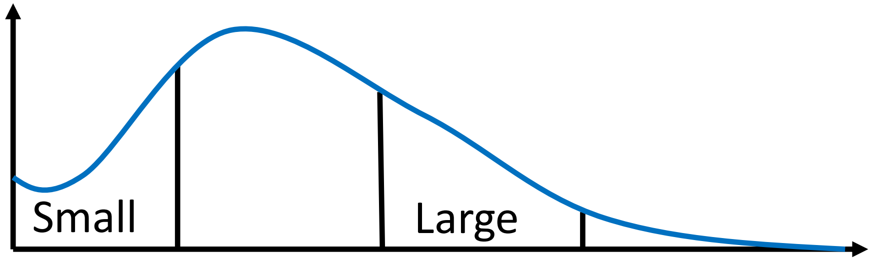

Job size distribution

Future result?

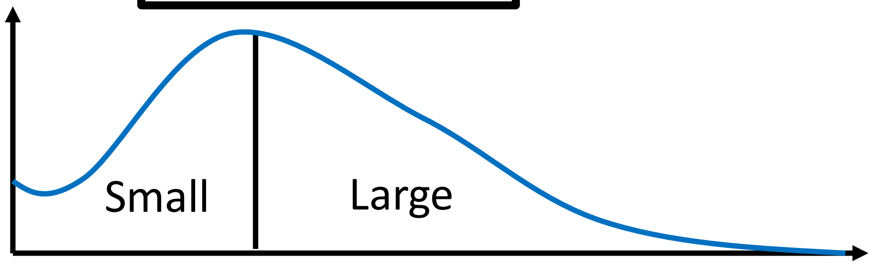

Job size distribution

Conclusion

Introduce policy called Nudge.

First policy to achieve stochastic improvement over FCFS, for any light-tailed job size distribution.

First policy to achieve multiplicative asymptotic improvement over FCFS.