

Scheduling and dispatching problems in queues:

✓ recently closed,

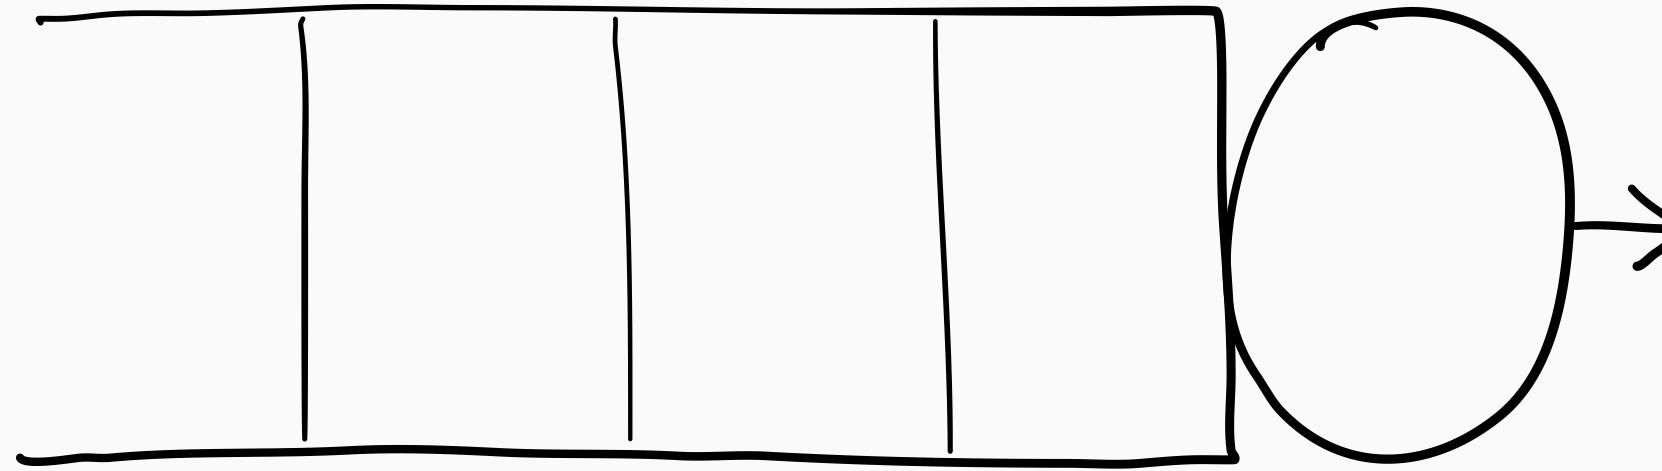
✗ hopelessly open, and

? promisingly ajar

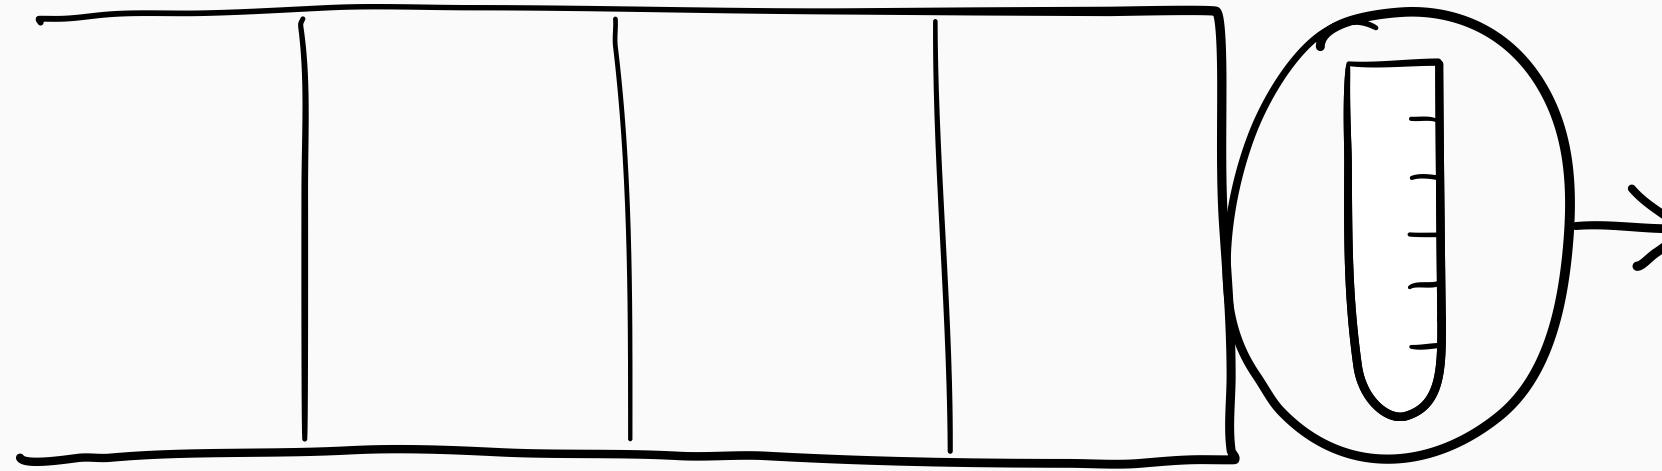
Ziv Scully
Cornell ORIE

The easiest scheduling problem

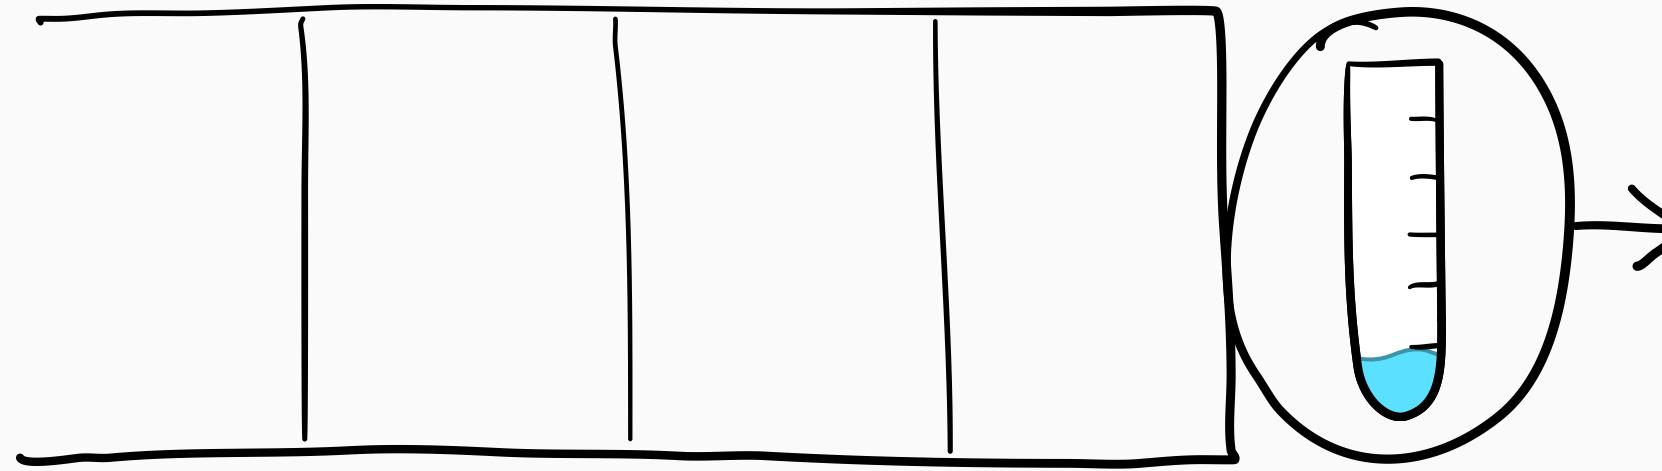
The easiest scheduling problem



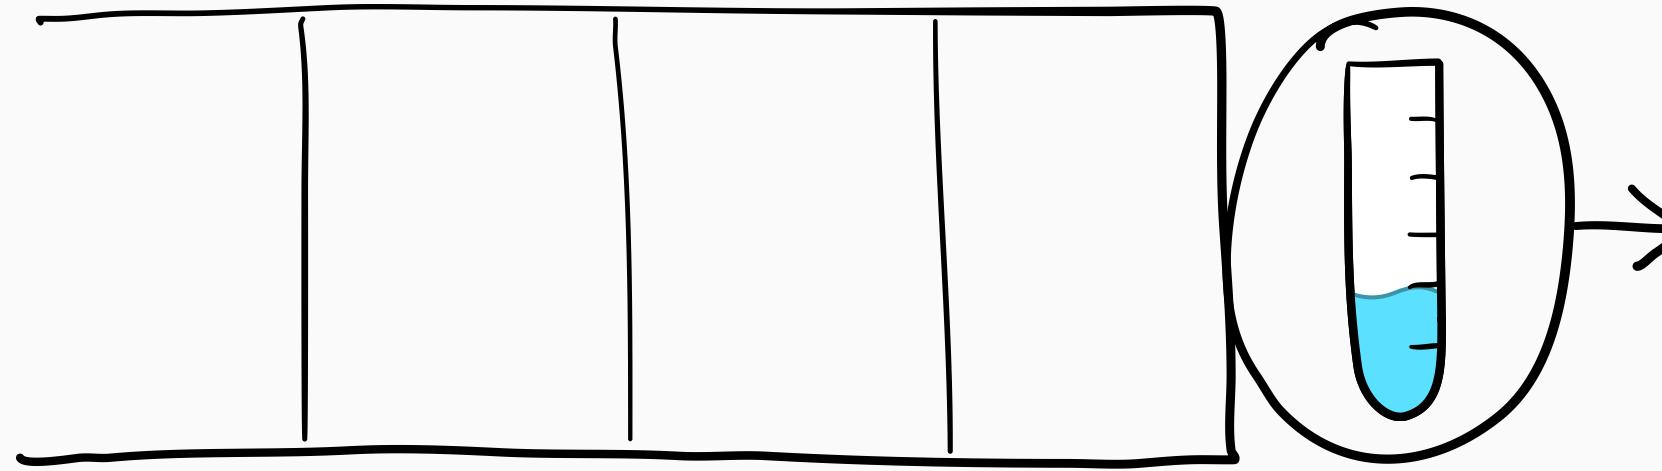
The easiest scheduling problem



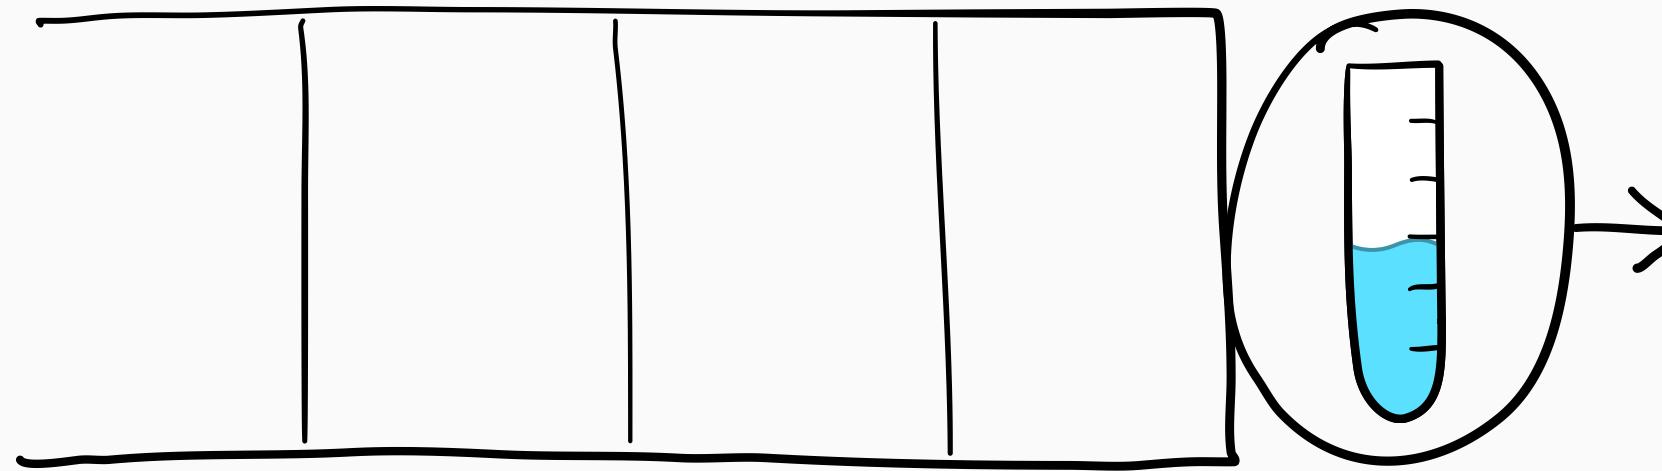
The easiest scheduling problem



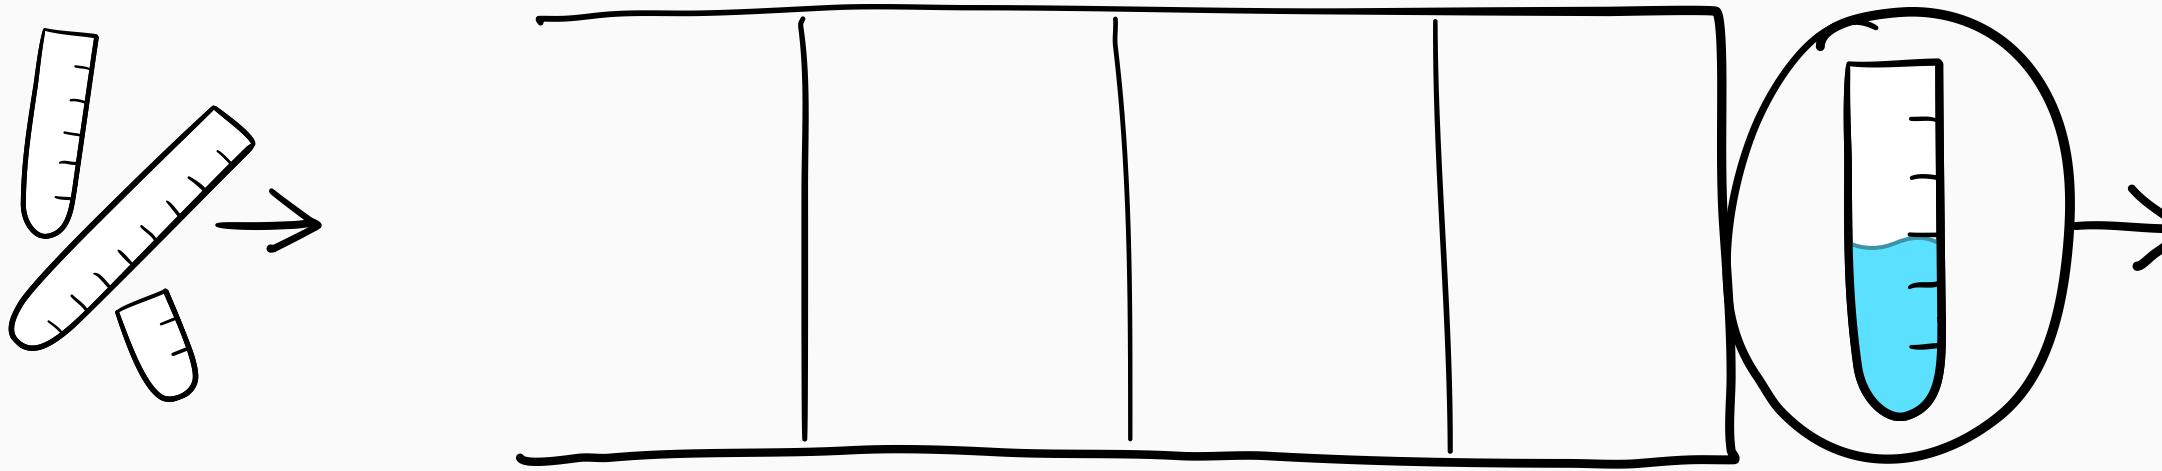
The easiest scheduling problem



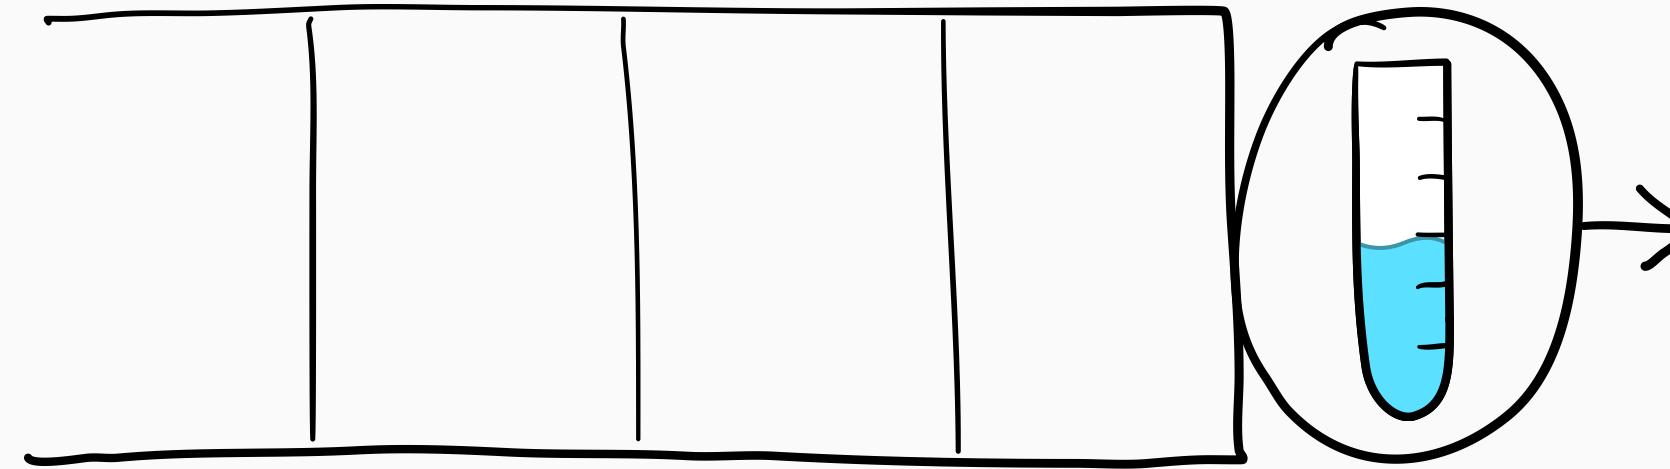
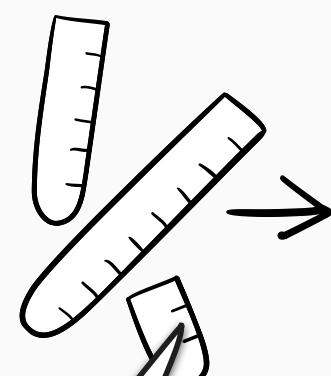
The easiest scheduling problem



The easiest scheduling problem



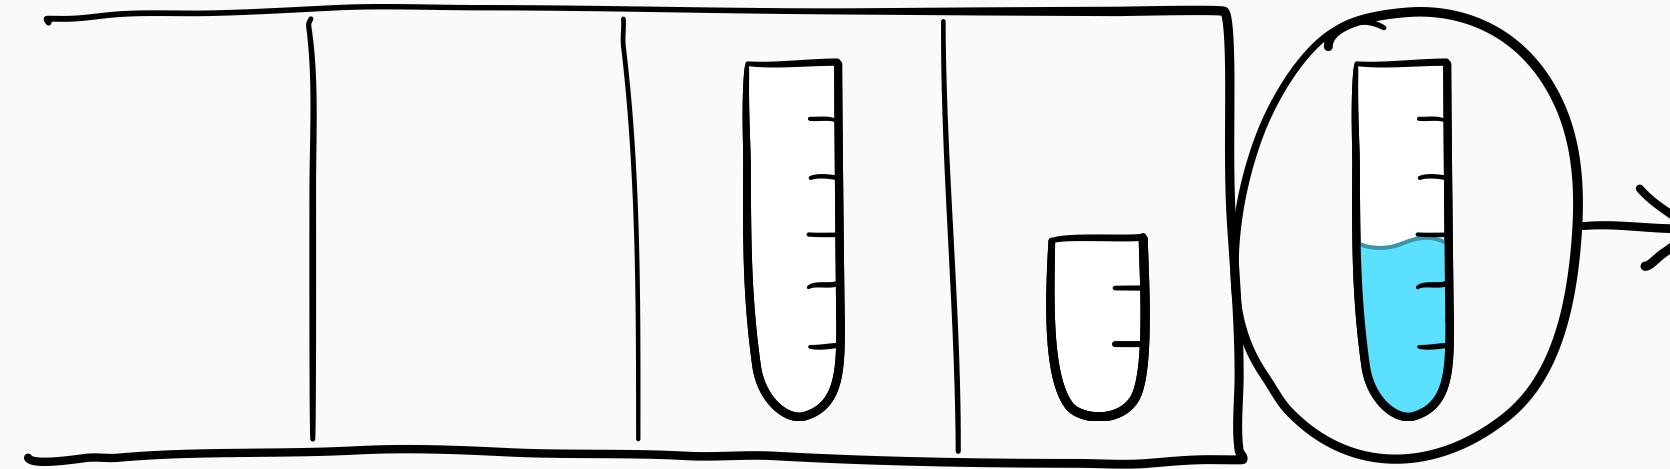
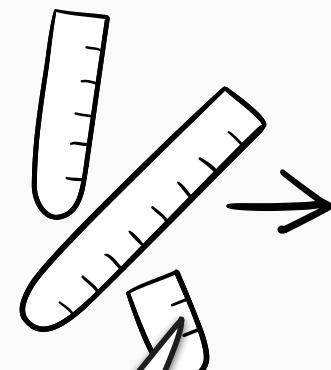
The easiest scheduling problem



M/G arrivals:

- arrival rate λ
- size distribution S

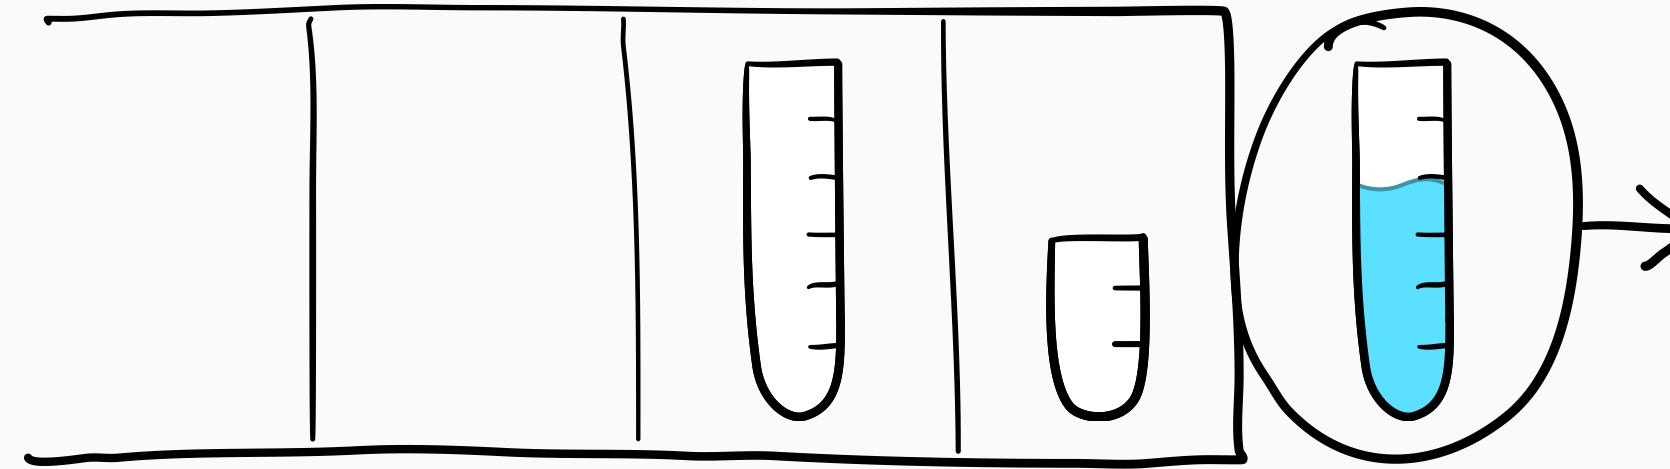
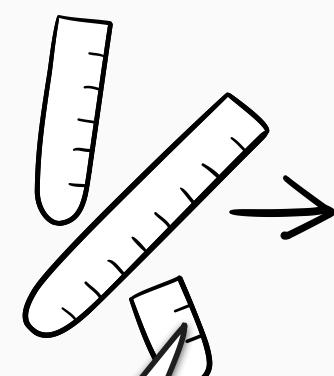
The easiest scheduling problem



M/G arrivals:

- arrival rate λ
- size distribution S

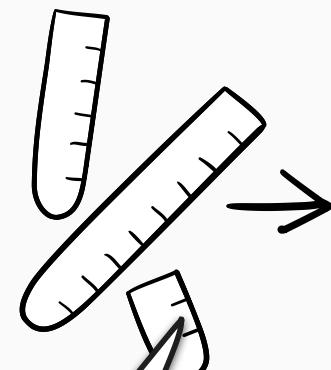
The easiest scheduling problem



M/G arrivals:

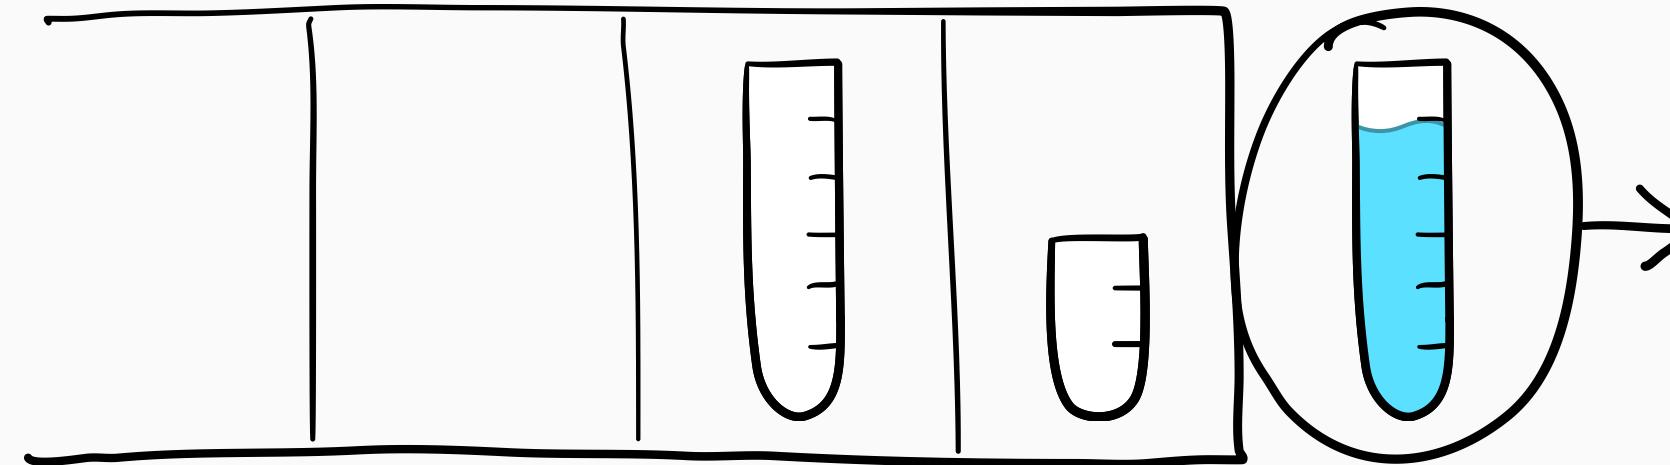
- arrival rate λ
- size distribution S

The easiest scheduling problem

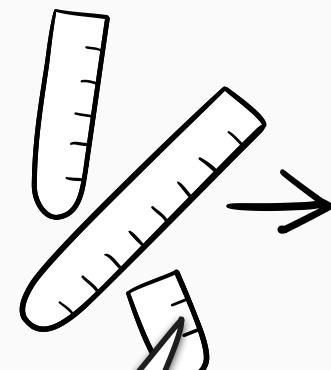


M/G arrivals:

- arrival rate λ
- size distribution S

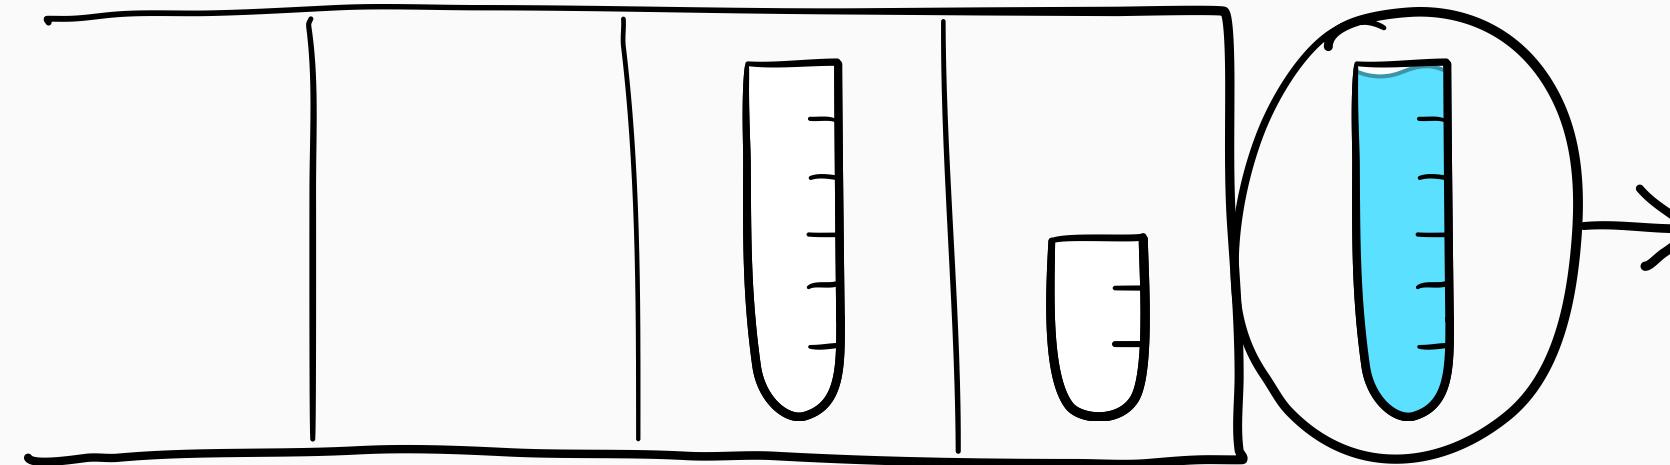


The easiest scheduling problem

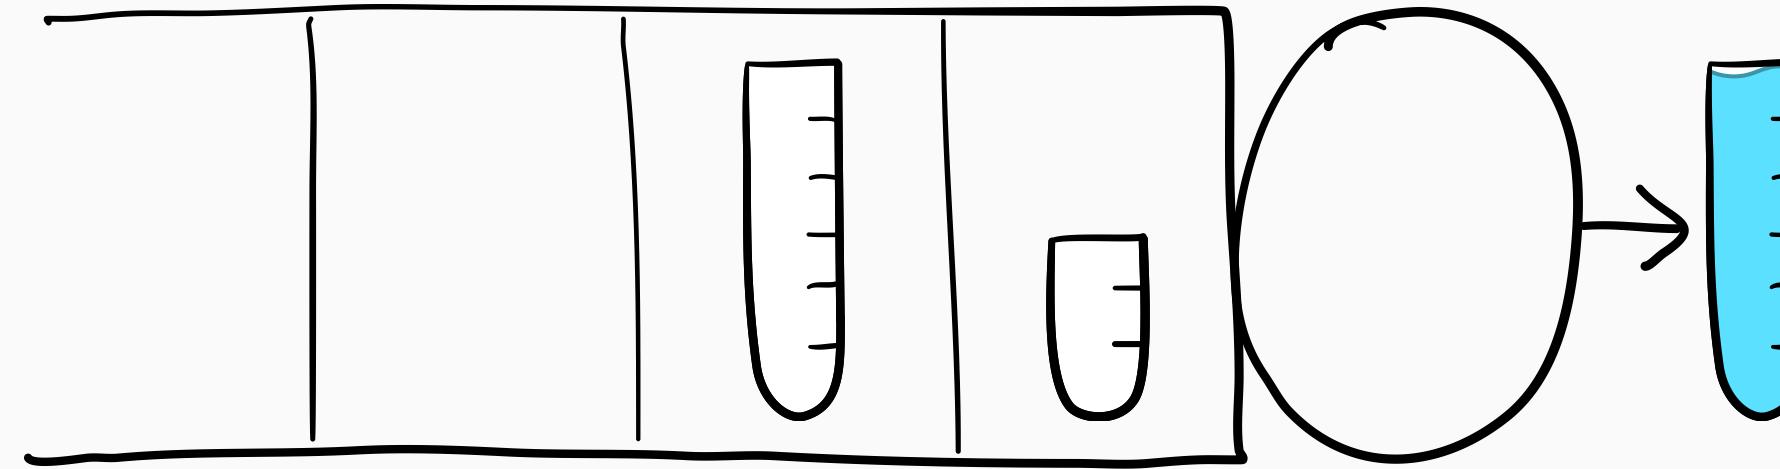
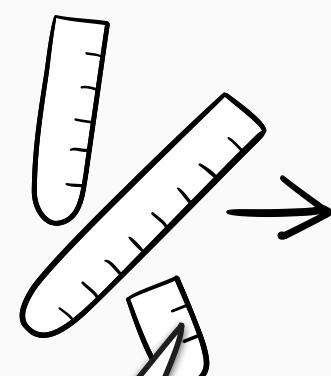


M/G arrivals:

- arrival rate λ
- size distribution S



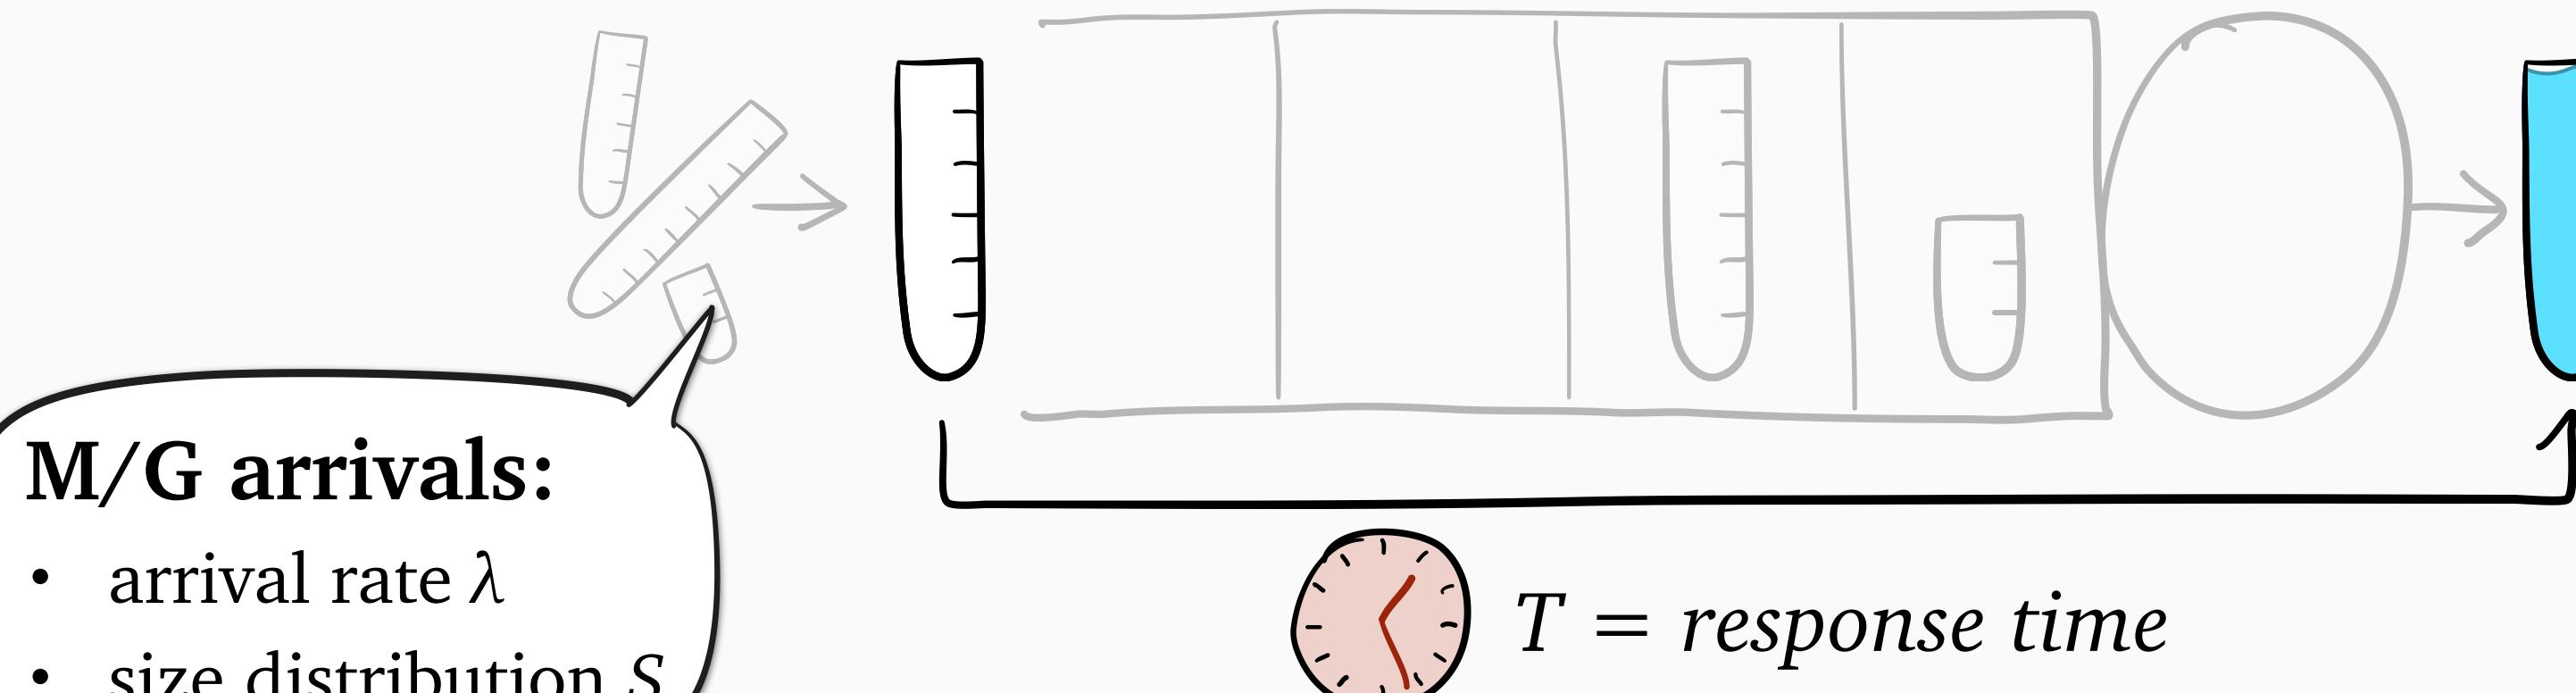
The easiest scheduling problem



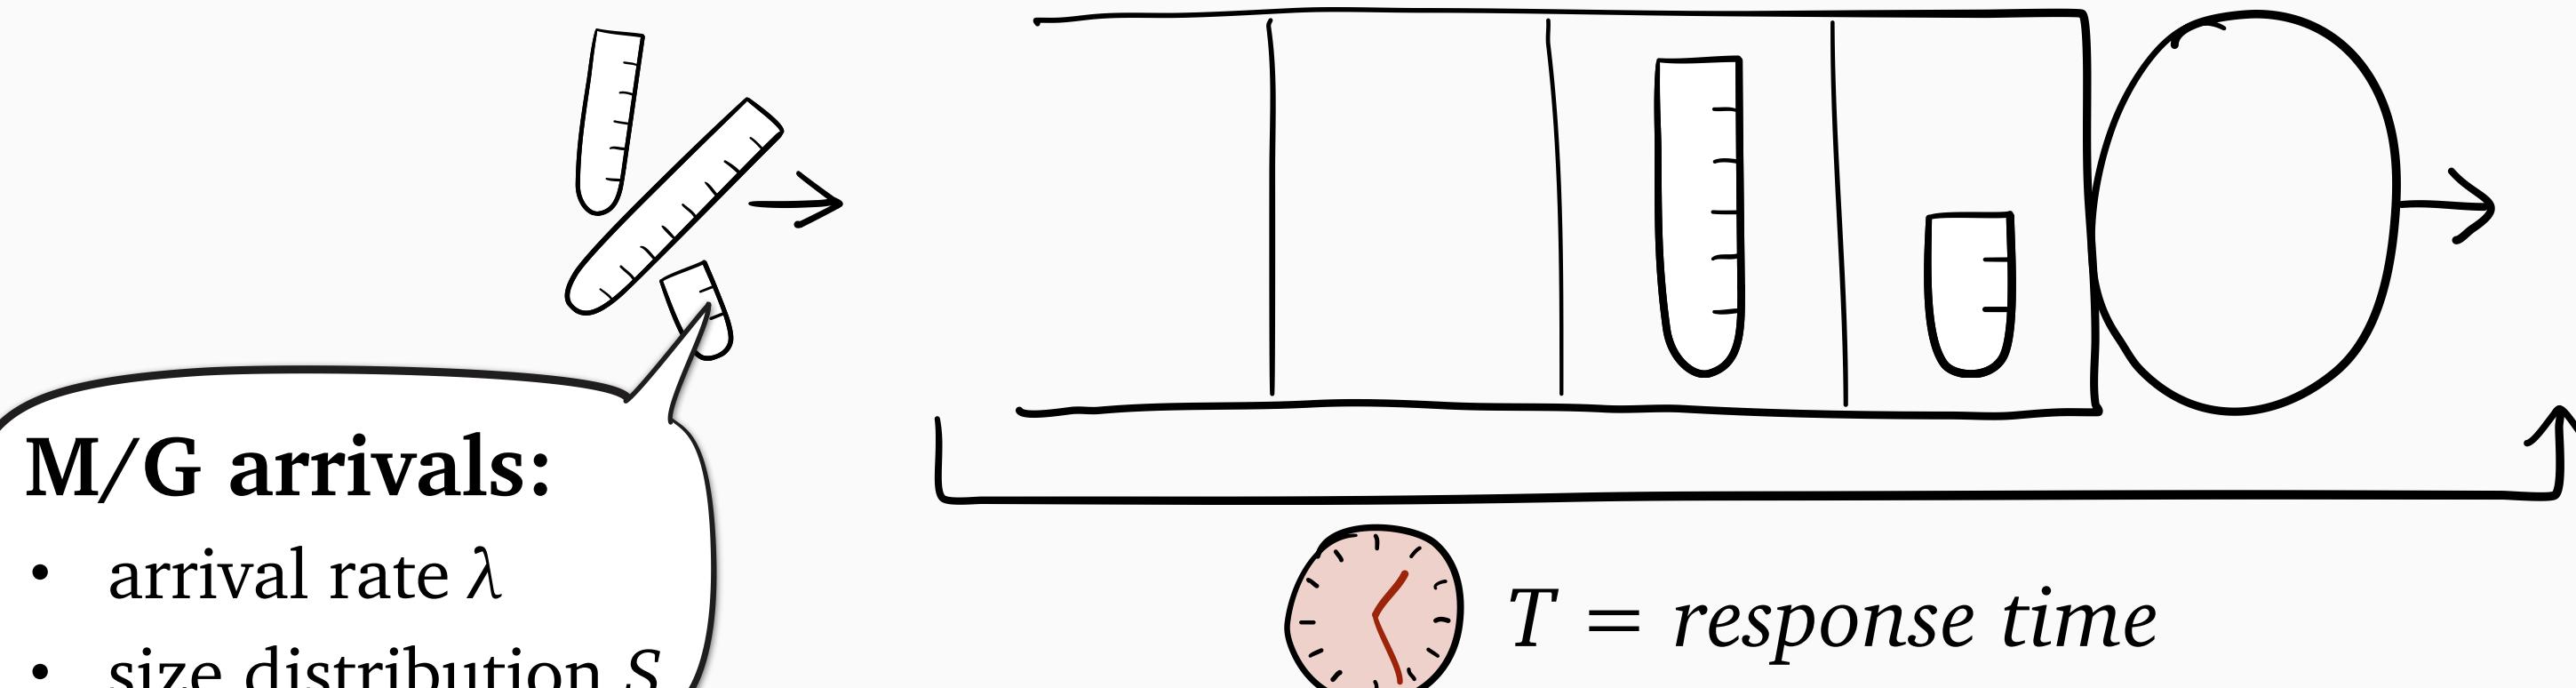
M/G arrivals:

- arrival rate λ
- size distribution S

The easiest scheduling problem



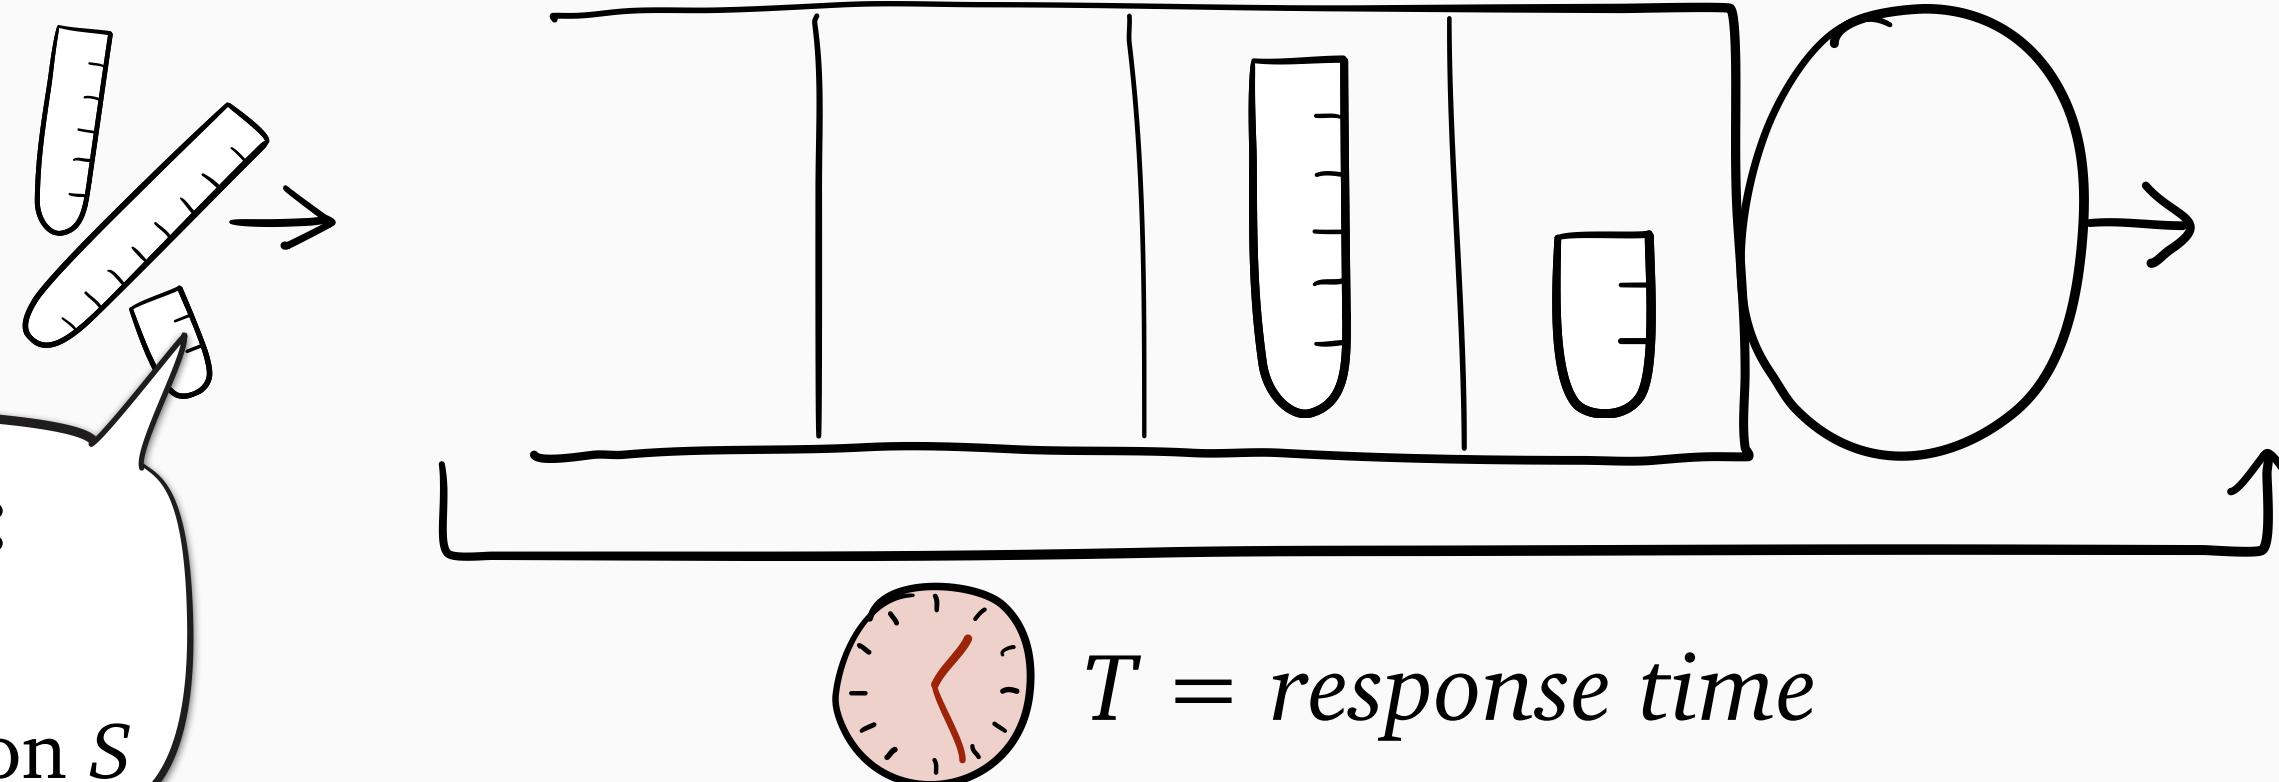
The easiest scheduling problem



The easiest scheduling problem

M/G arrivals:

- arrival rate λ
- size distribution S

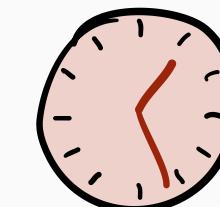
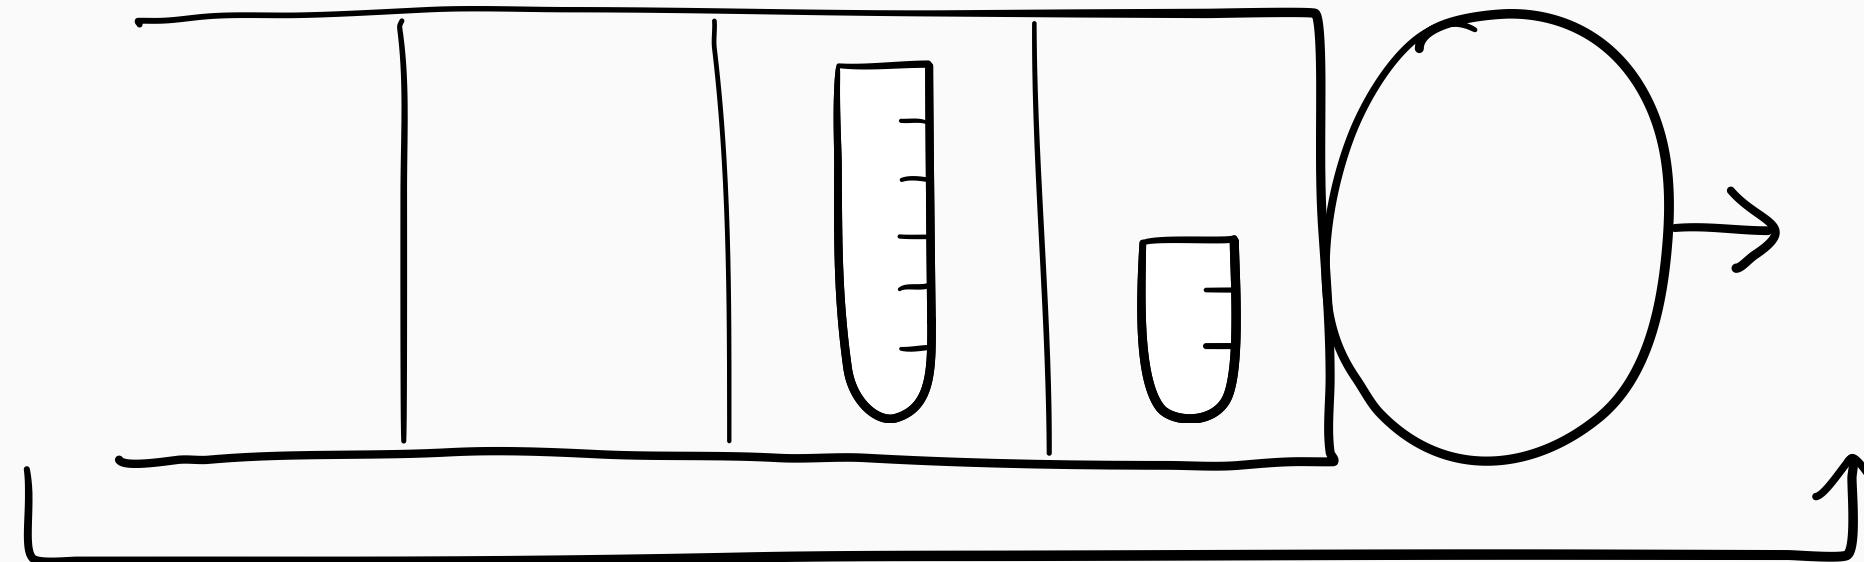
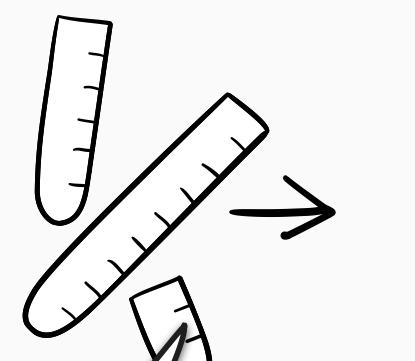


Minimize $E[T]$?

The easiest scheduling problem

M/G arrivals:

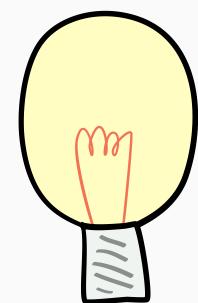
- arrival rate λ
- size distribution S



$T = \text{response time}$



Minimize $E[T]$?

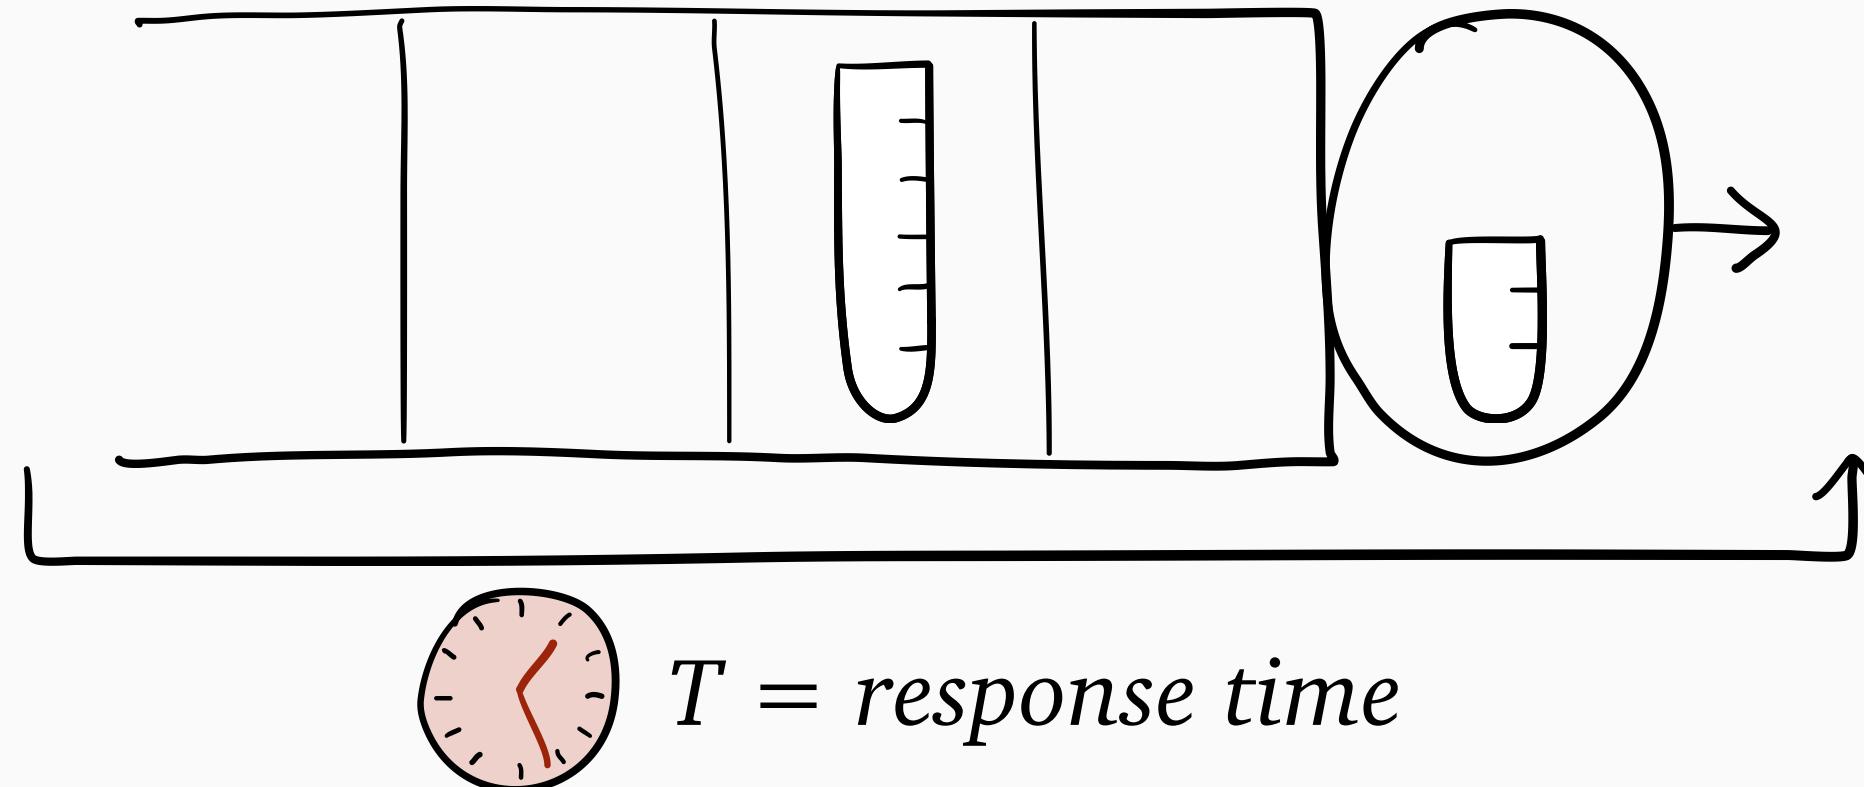
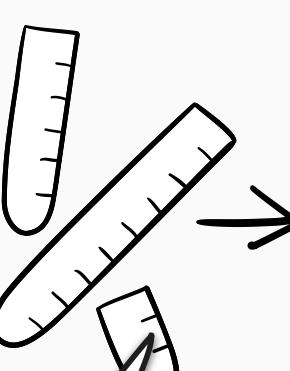


Serve short jobs
before long jobs

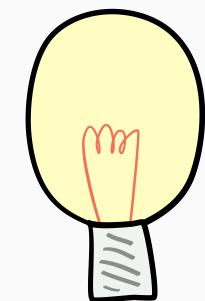
The easiest scheduling problem

M/G arrivals:

- arrival rate λ
- size distribution S



Minimize $E[T]$?

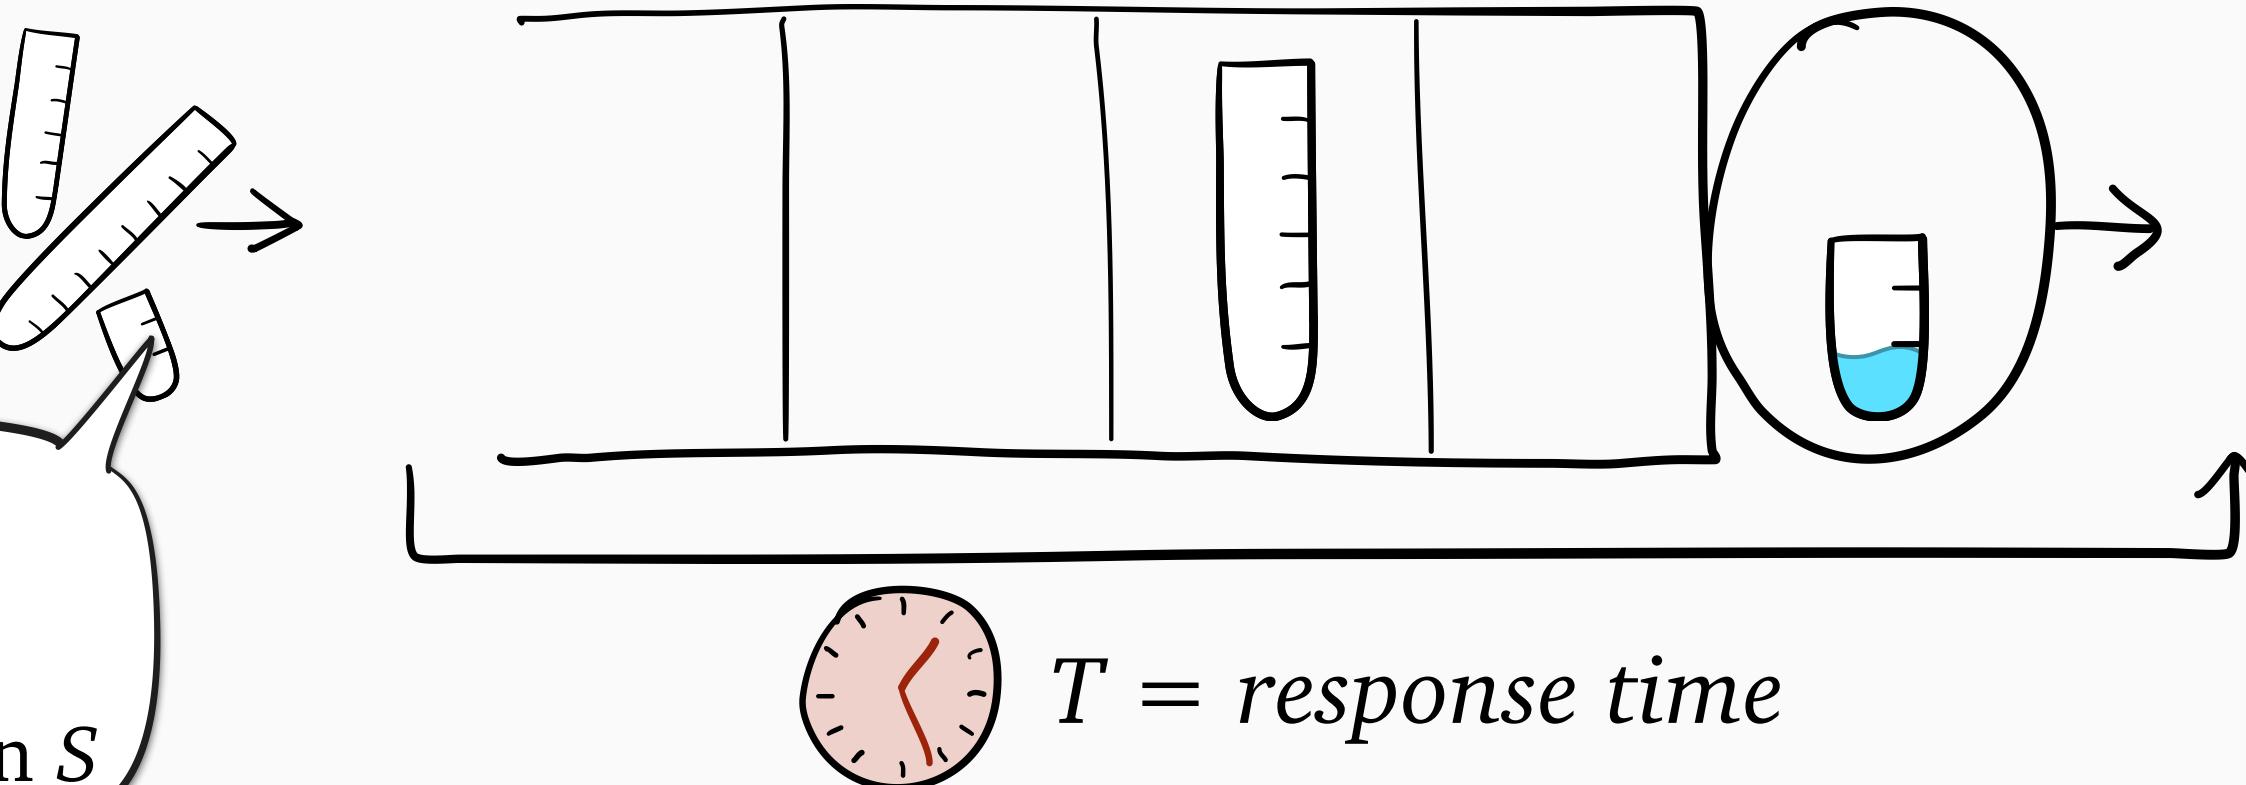


Serve short jobs
before long jobs

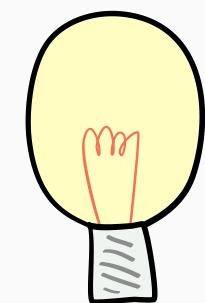
The easiest scheduling problem

M/G arrivals:

- arrival rate λ
- size distribution S



Minimize $E[T]$?

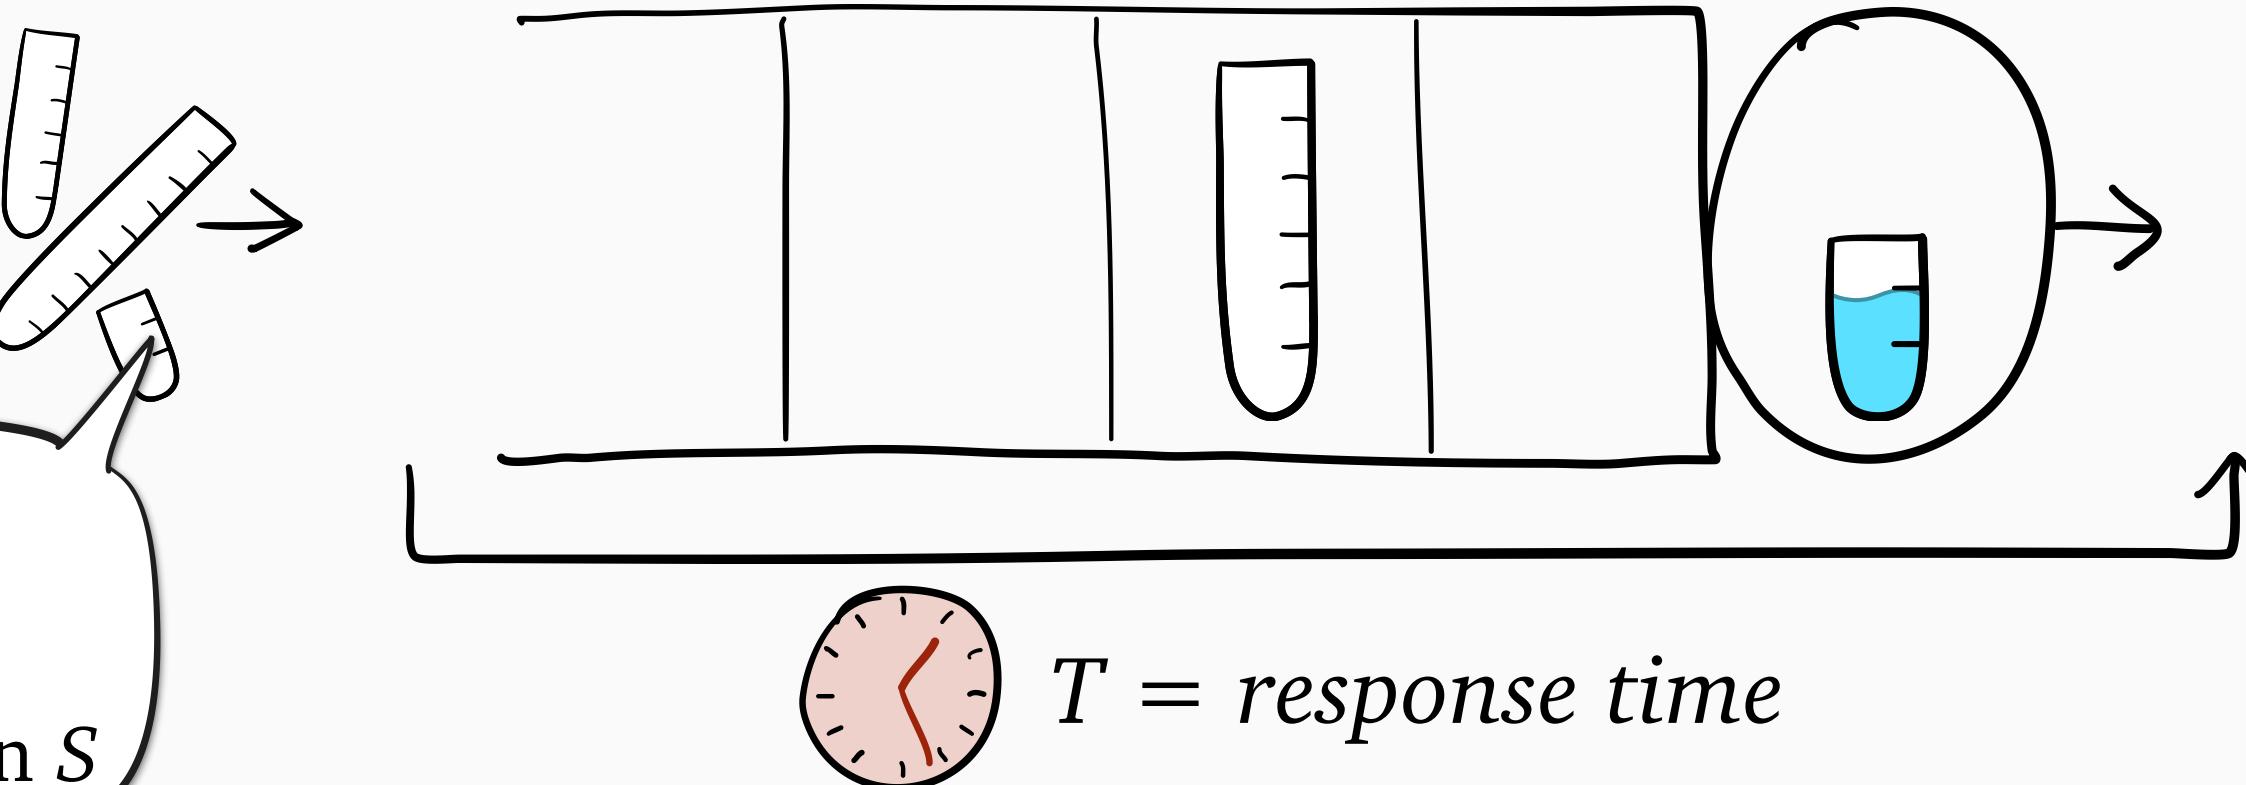


Serve short jobs
before long jobs

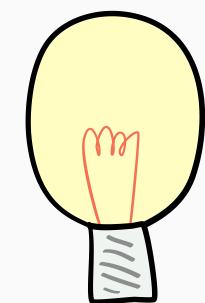
The easiest scheduling problem

M/G arrivals:

- arrival rate λ
- size distribution S

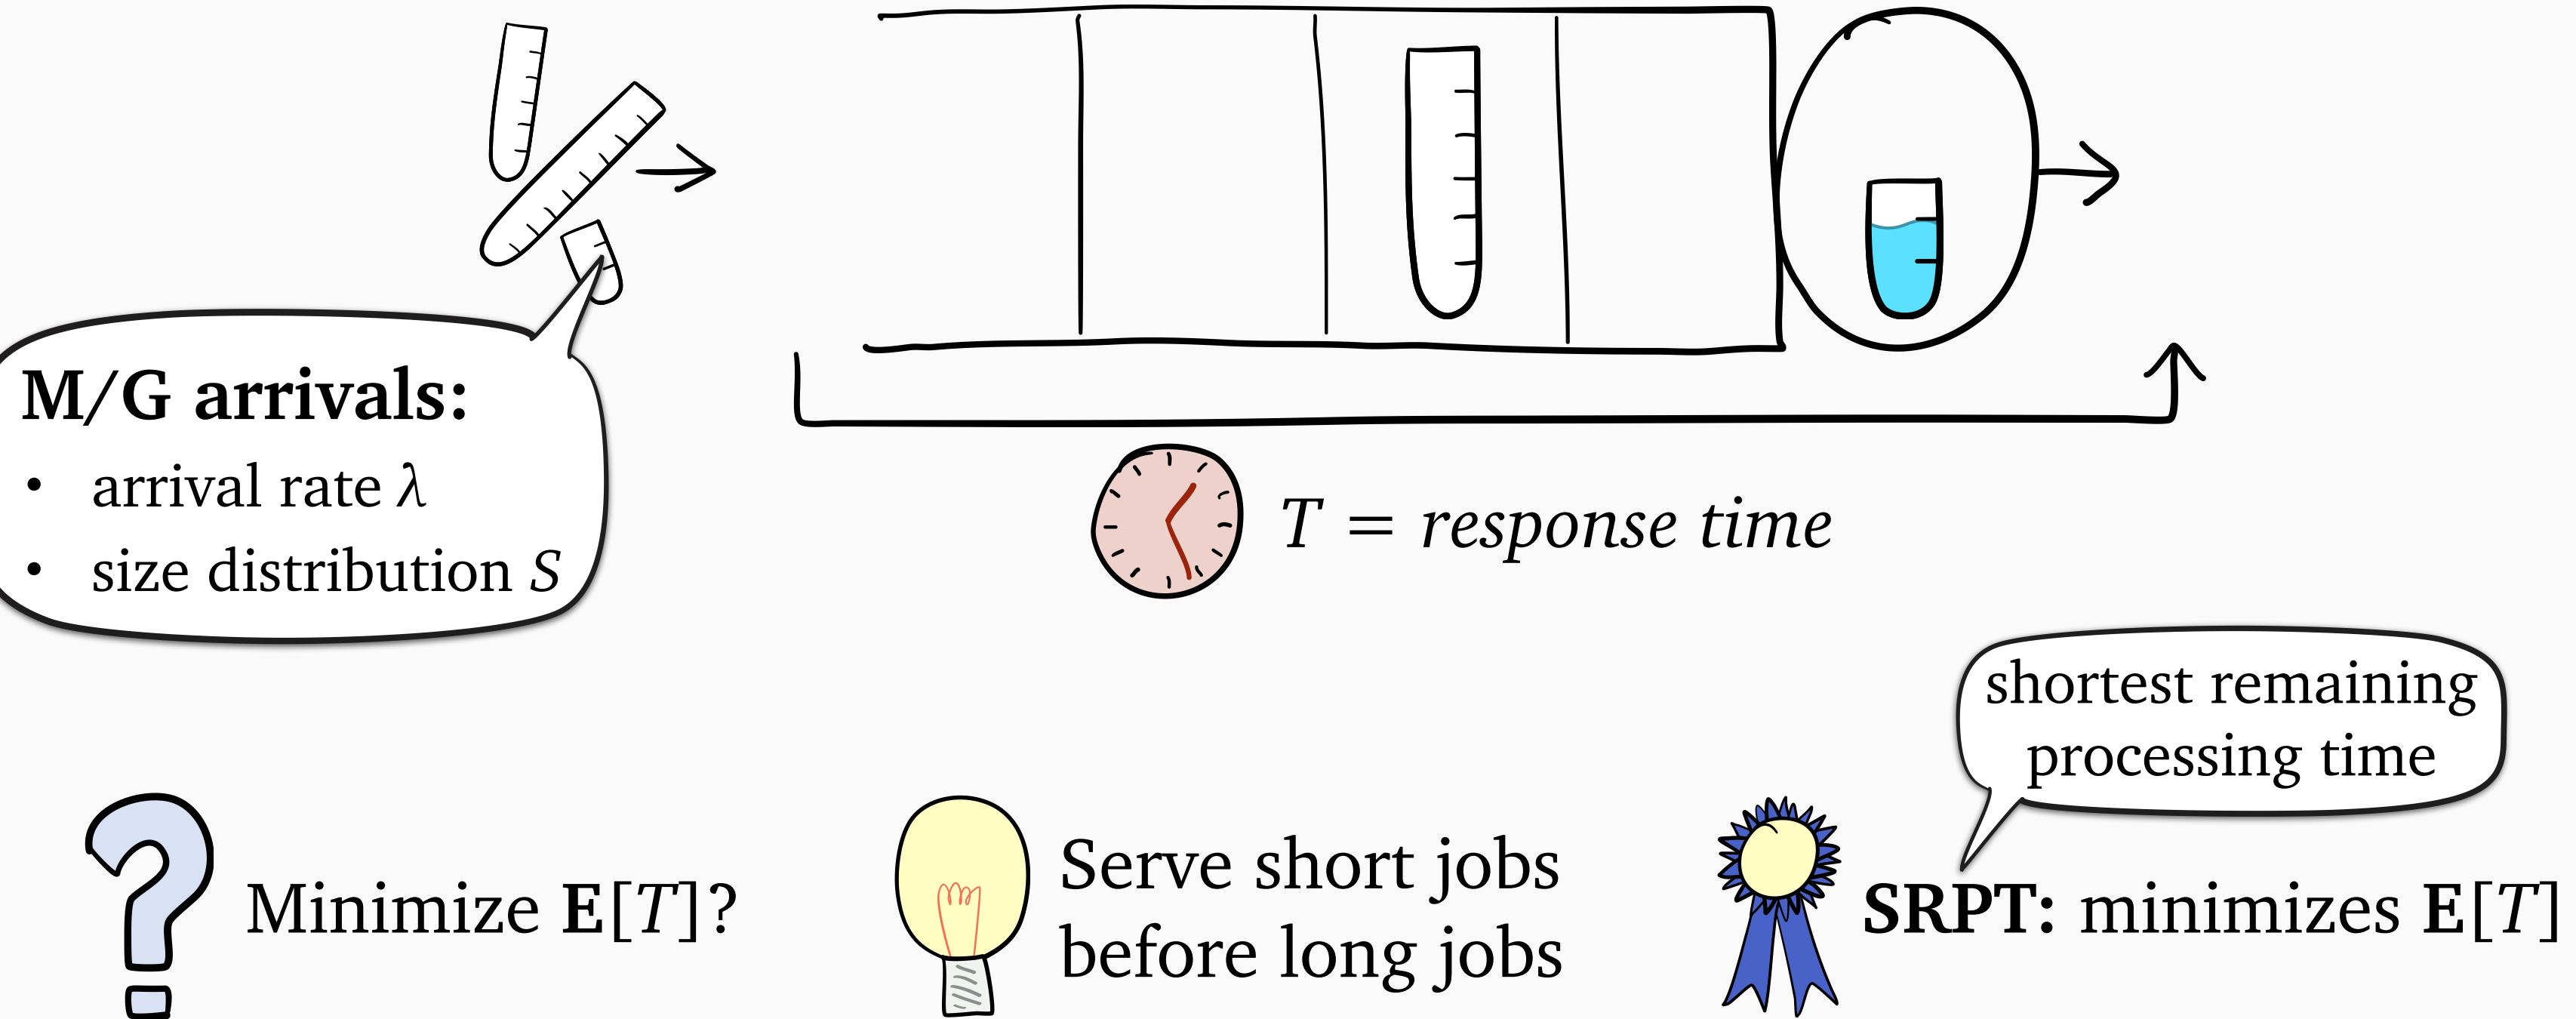


Minimize $E[T]$?



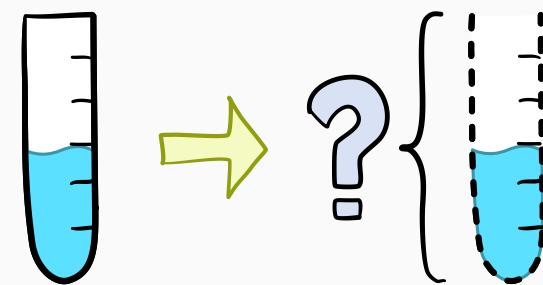
Serve short jobs
before long jobs

The easiest scheduling problem



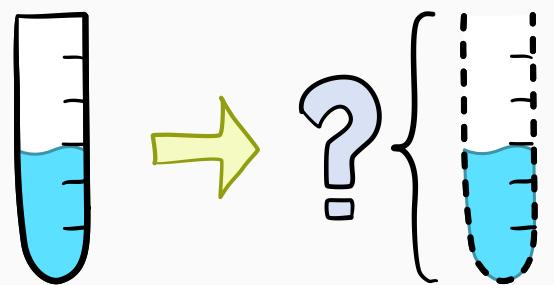
Hard scheduling questions

Hard scheduling questions

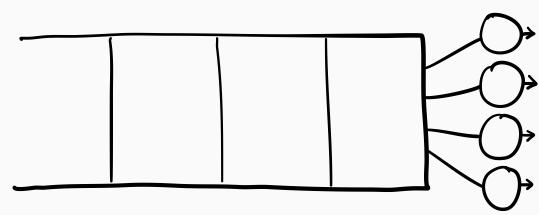


*What if we have
unknown job sizes?*

Hard scheduling questions

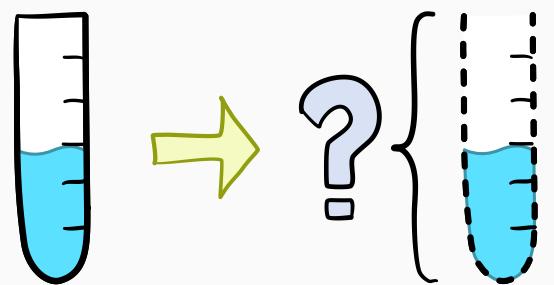


*What if we have
unknown job sizes?*

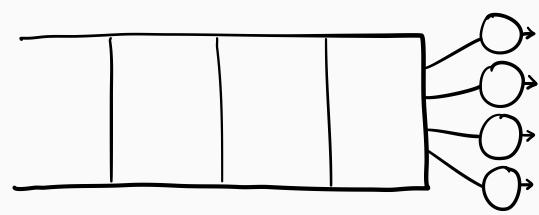


*What if there are
multiple servers?*

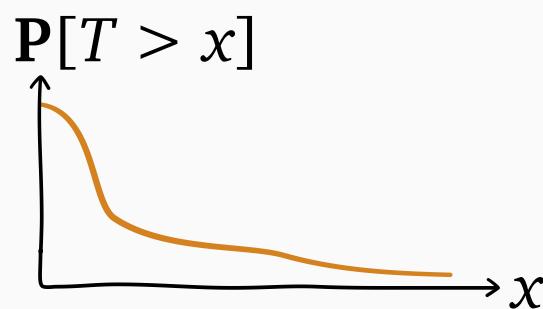
Hard scheduling questions



*What if we have
unknown job sizes?*

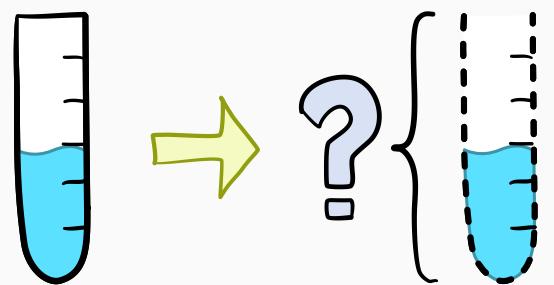


*What if there are
multiple servers?*

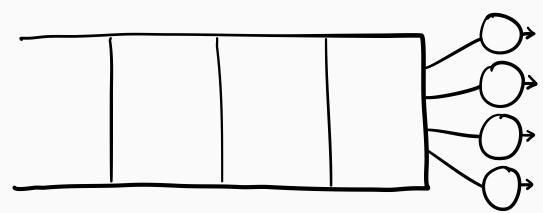


*What if we want to optimize
tails instead of means?*

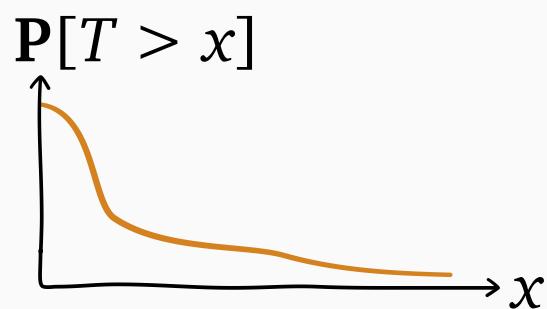
Hard scheduling questions



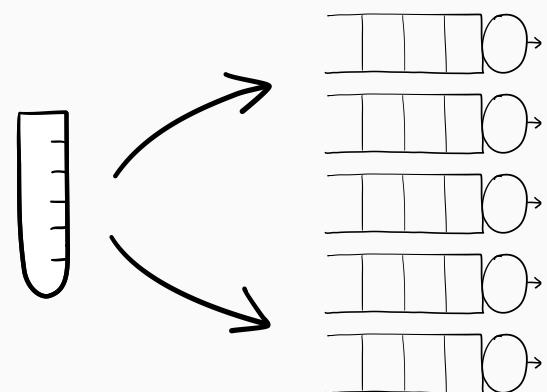
*What if we have
unknown job sizes?*



*What if there are
multiple servers?*

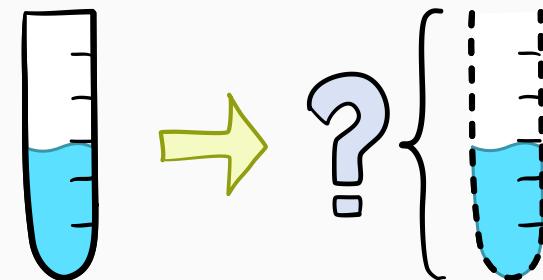


*What if we want to optimize
tails instead of means?*

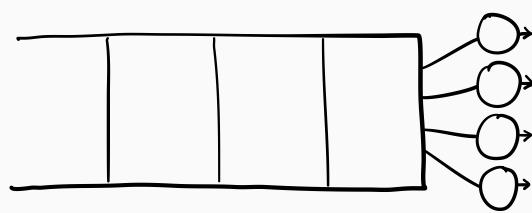


*What if we can only use
dispatching without
fancy scheduling?*

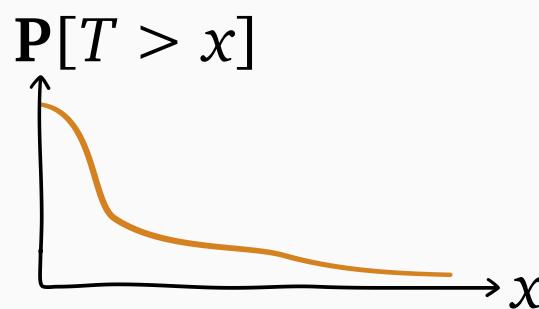
Hard scheduling questions



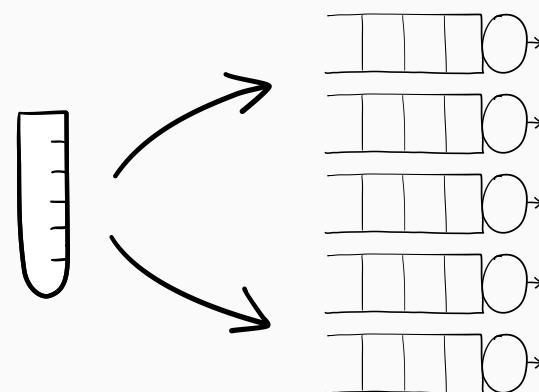
*What if we have
unknown job sizes?*



*What if there are
multiple servers?*



*What if we want to optimize
tails instead of means?*



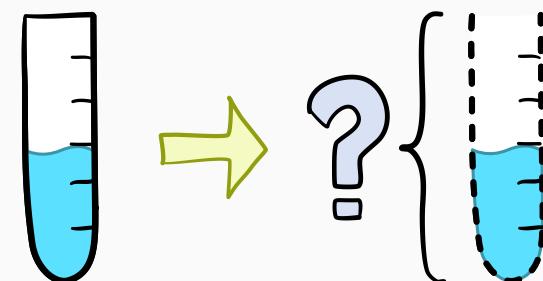
*What if we can only use
dispatching without
fancy scheduling?*



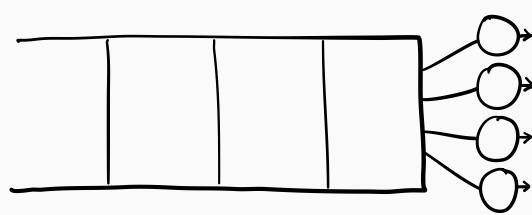
*What are the underlying
theoretical tools?*



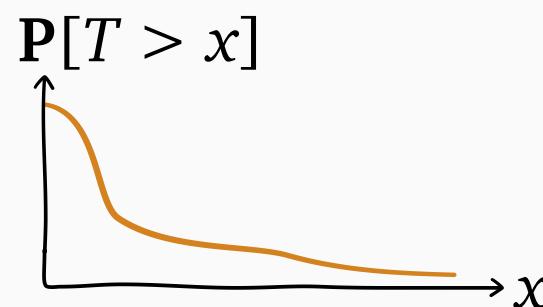
Hard scheduling questions



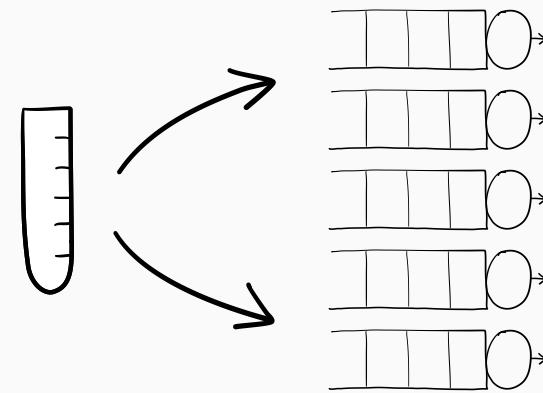
*What if we have
unknown job sizes?*



*What if there are
multiple servers?*



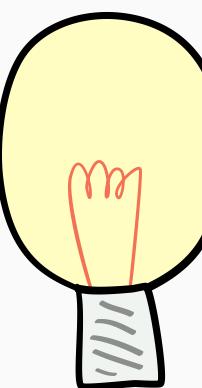
*What if we want to optimize
tails instead of means?*



*What if we can only use
dispatching without
fancy scheduling?*

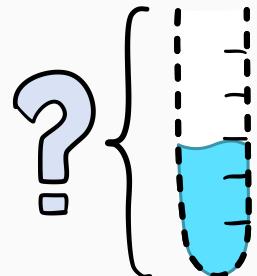
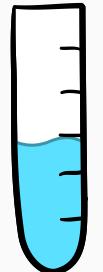


*What are the underlying
theoretical tools?*

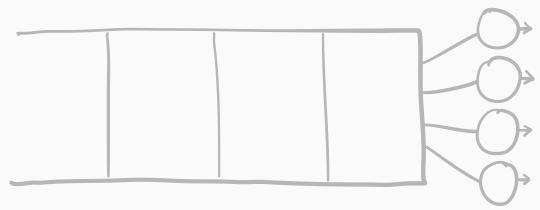


*What are the resulting
practical lessons?*

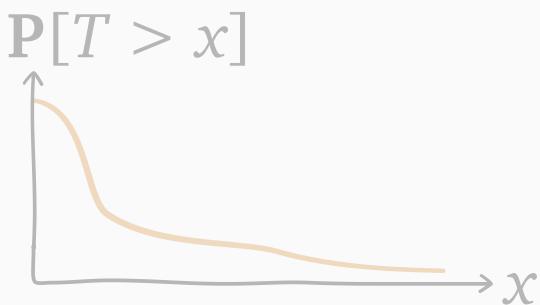
Hard scheduling questions



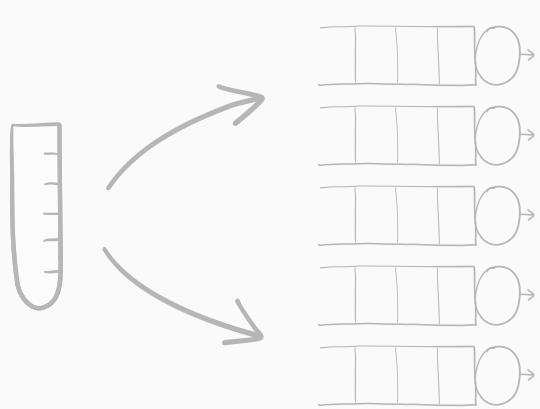
*What if we have
unknown job sizes?*



*What if there are
multiple servers?*



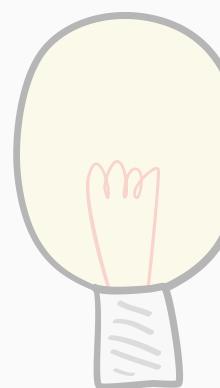
*What if we want to optimize
tails instead of means?*



*What if we can only use
dispatching without
fancy scheduling?*

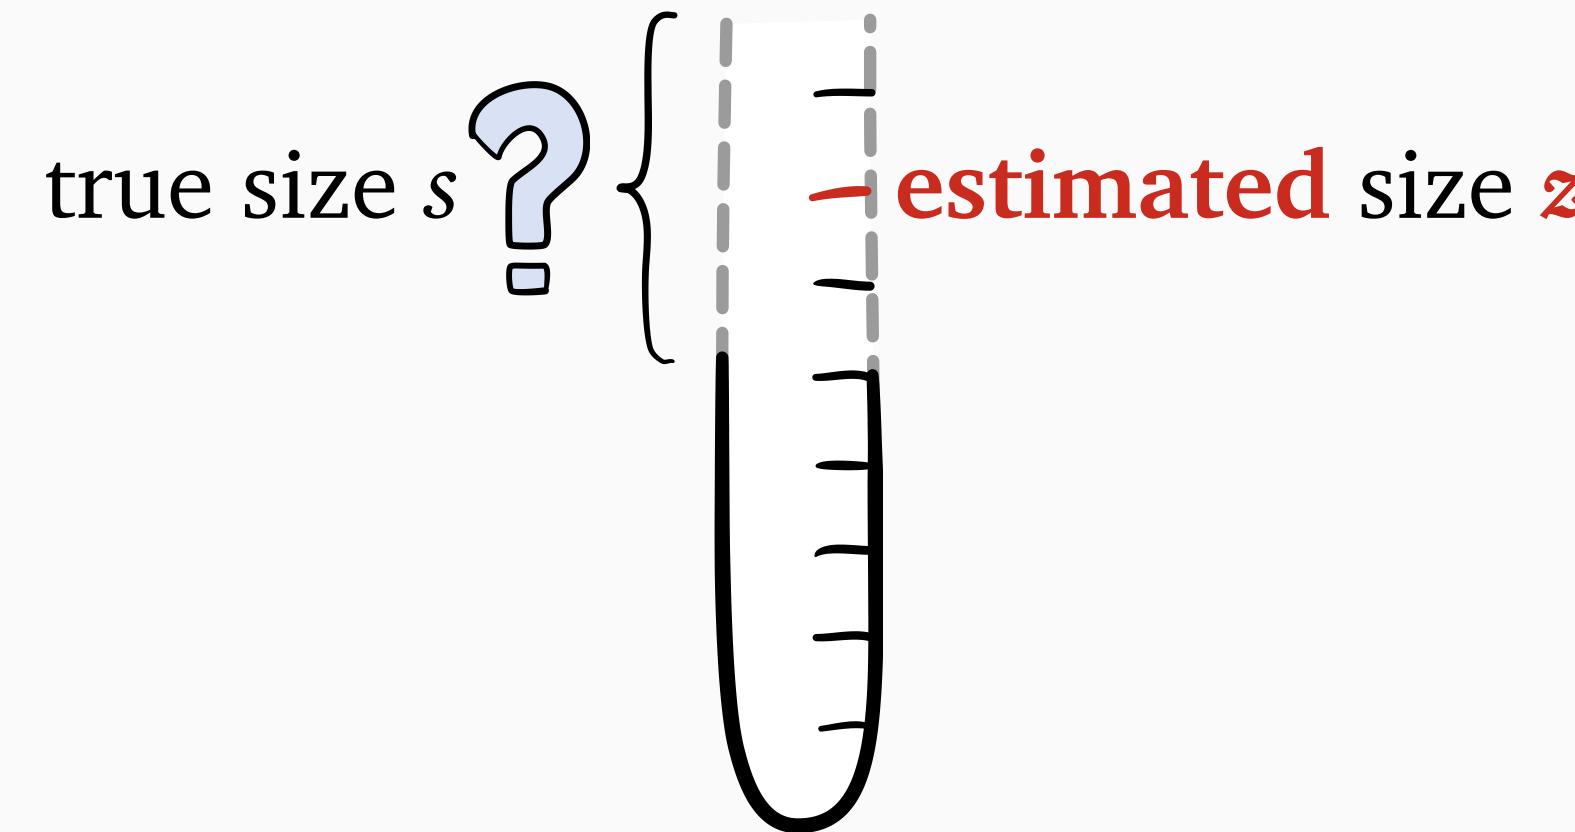


*What are the underlying
theoretical tools?*

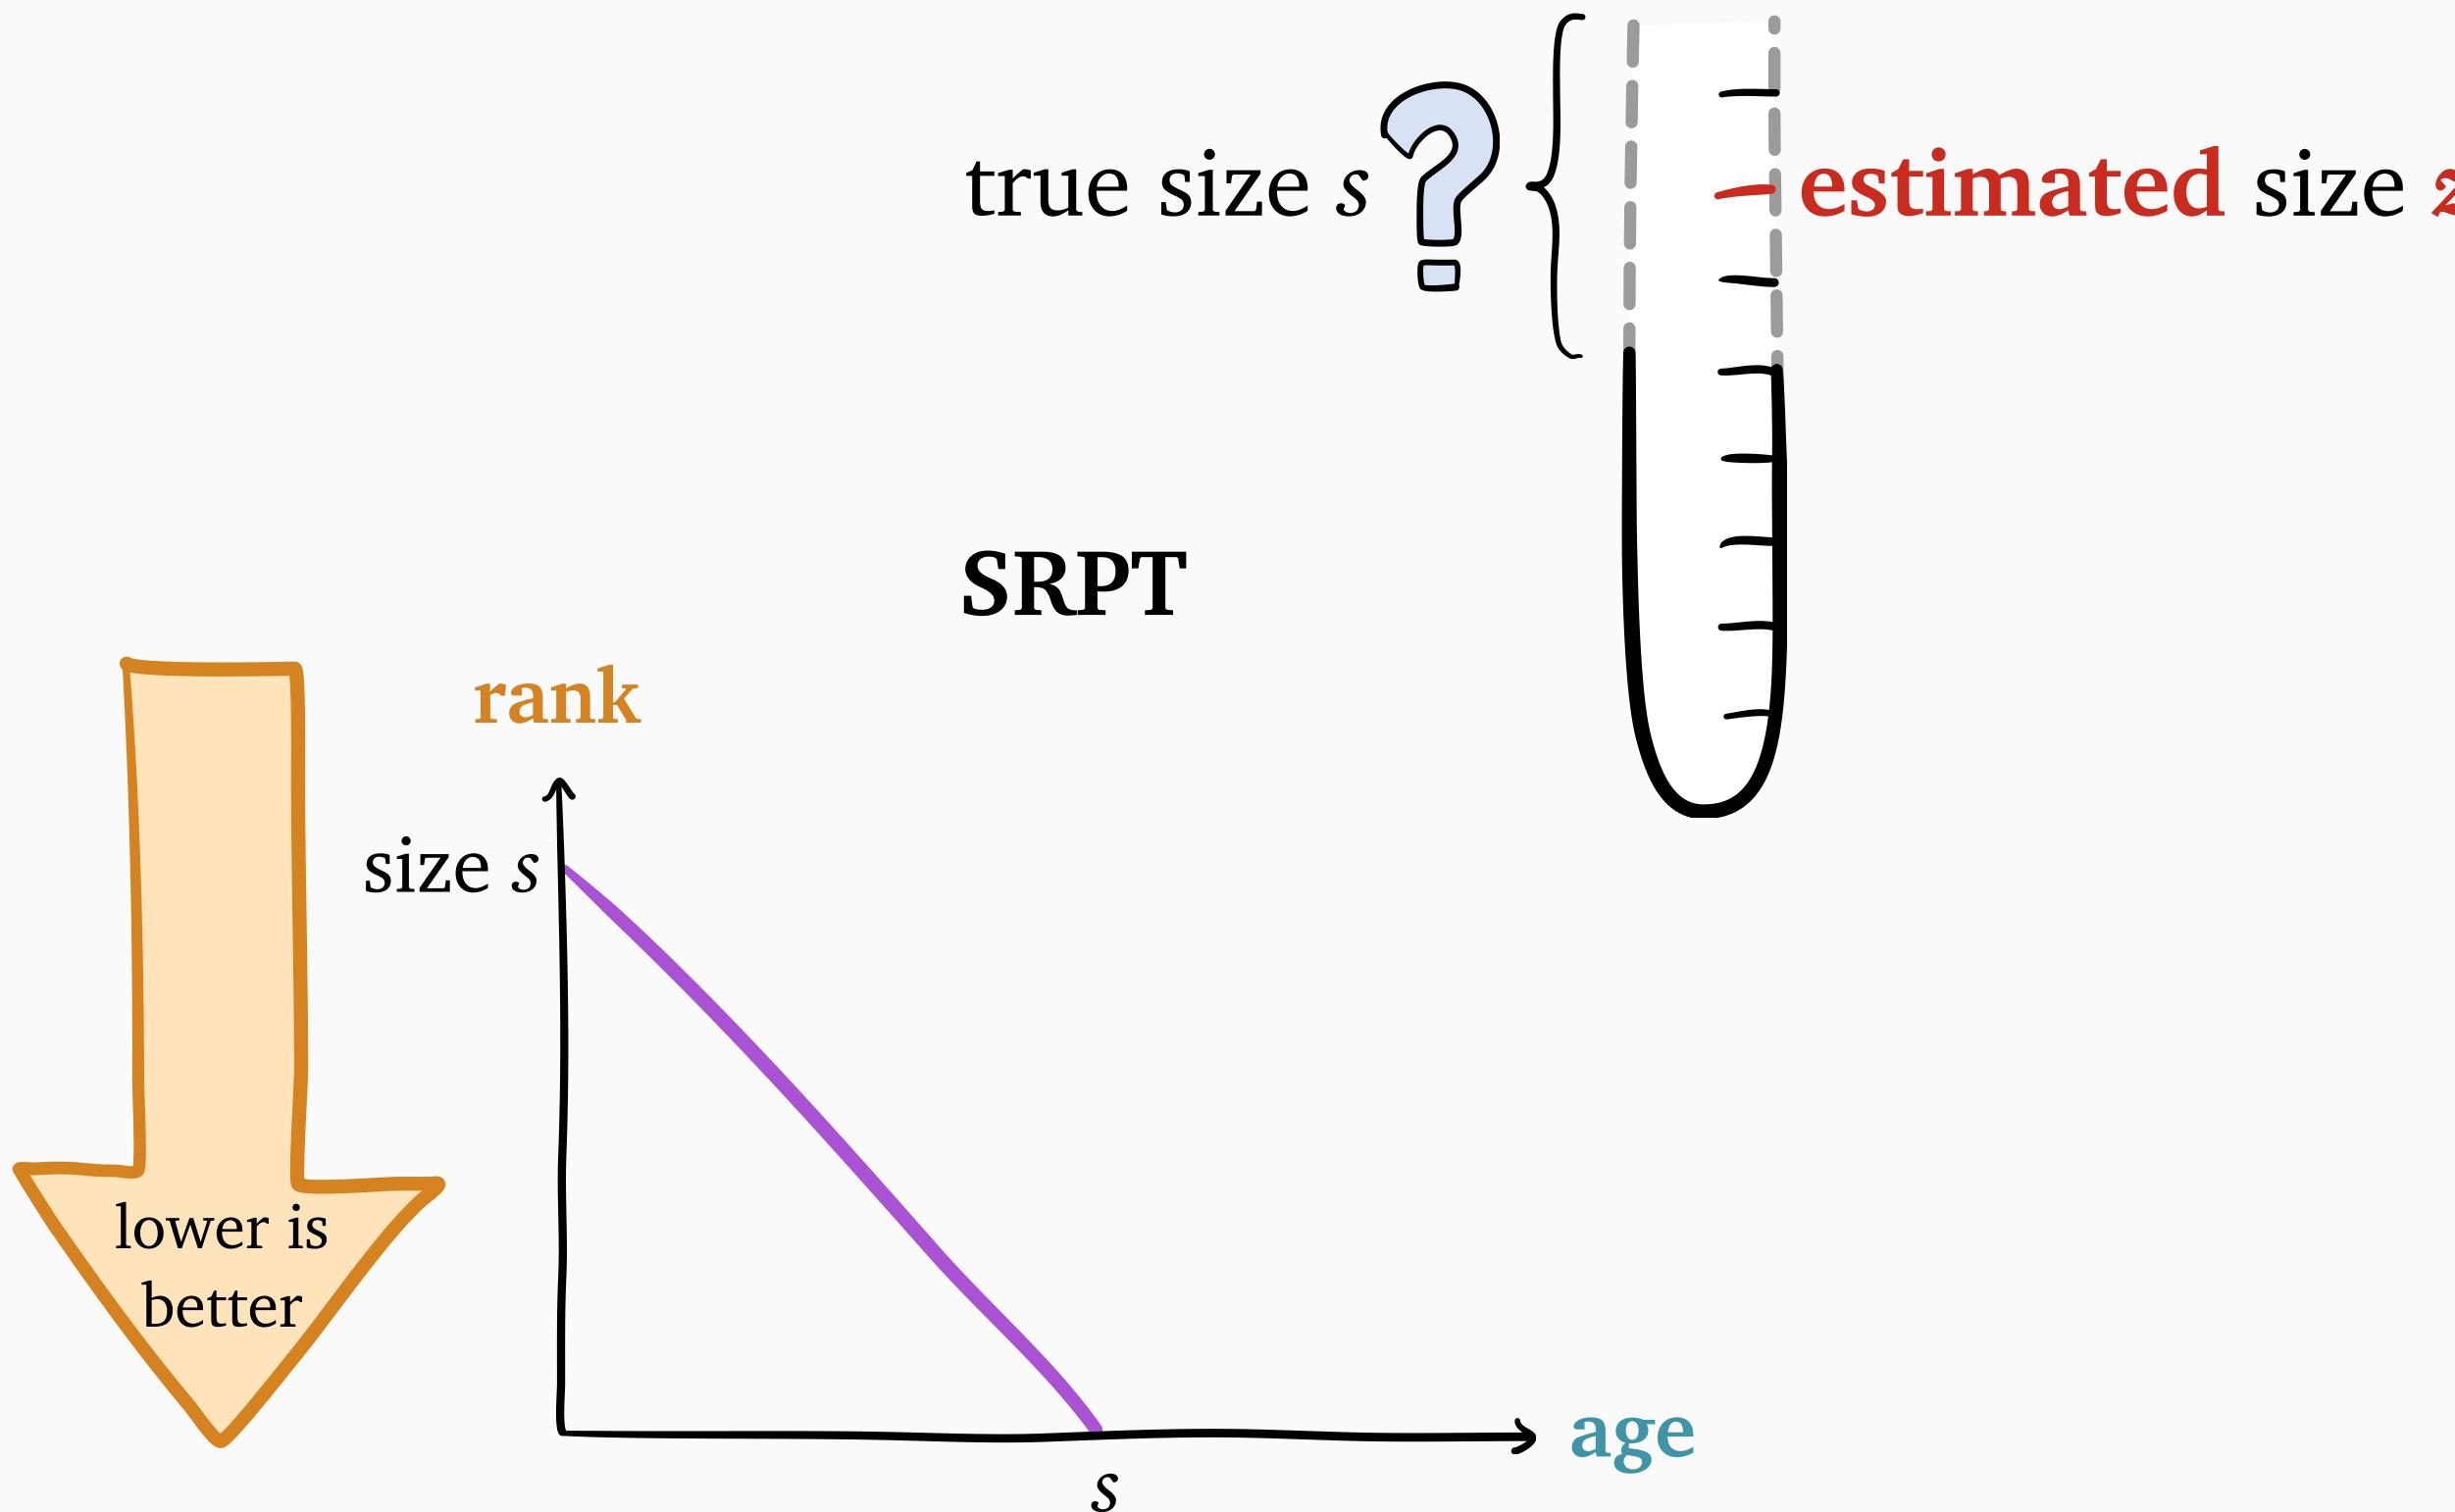


*What are the resulting
practical lessons?*

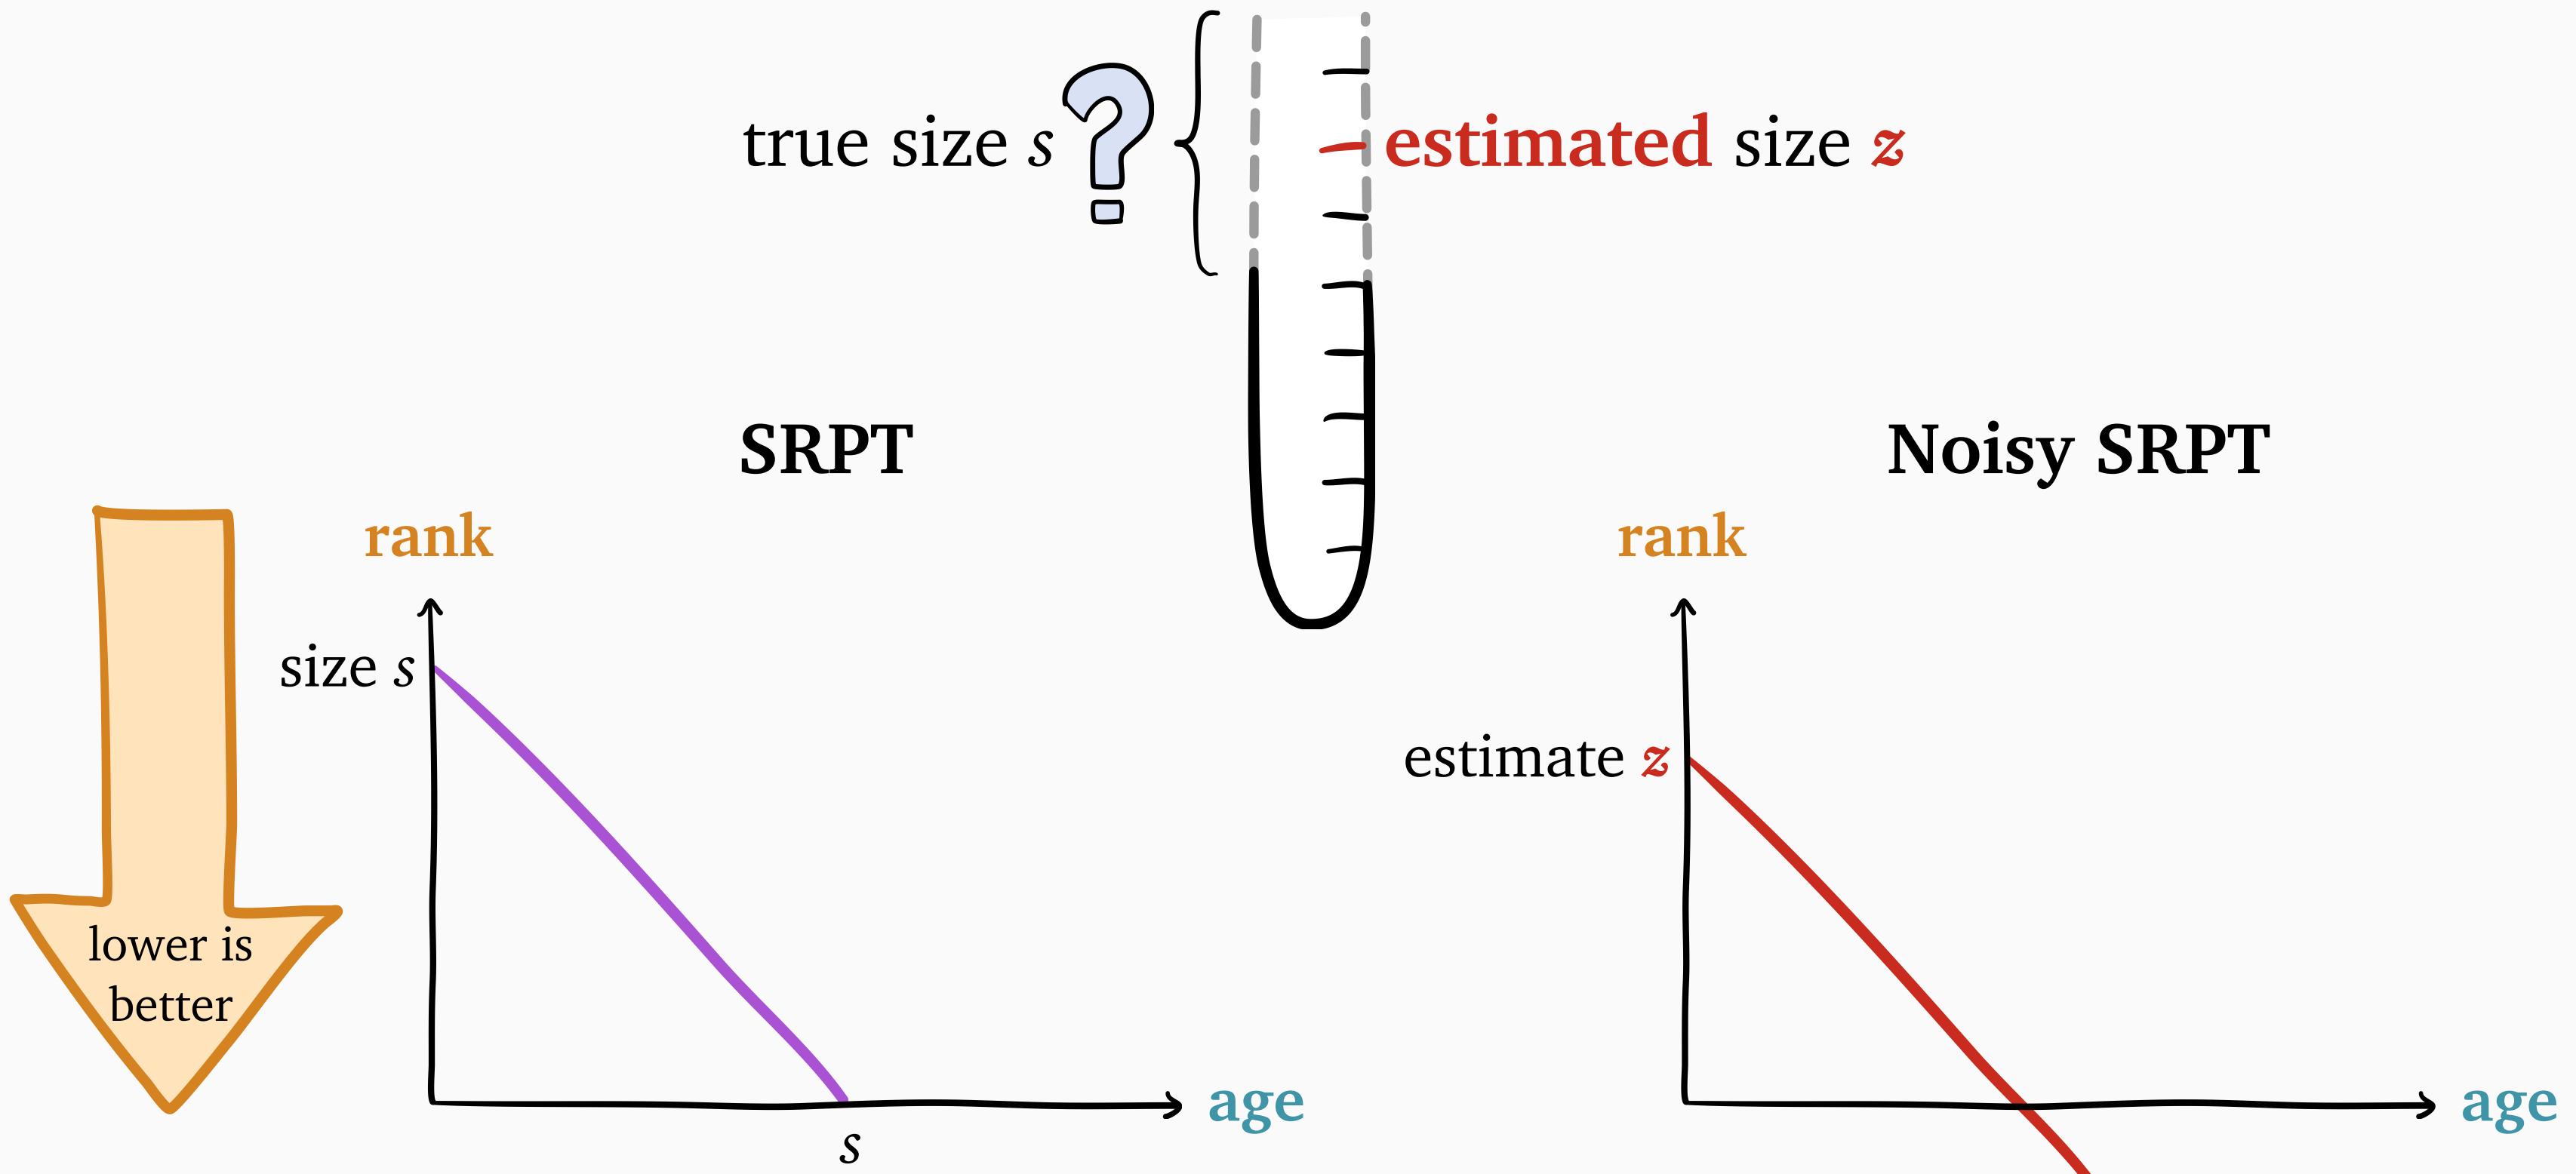
Scheduling with **noisy** size estimates



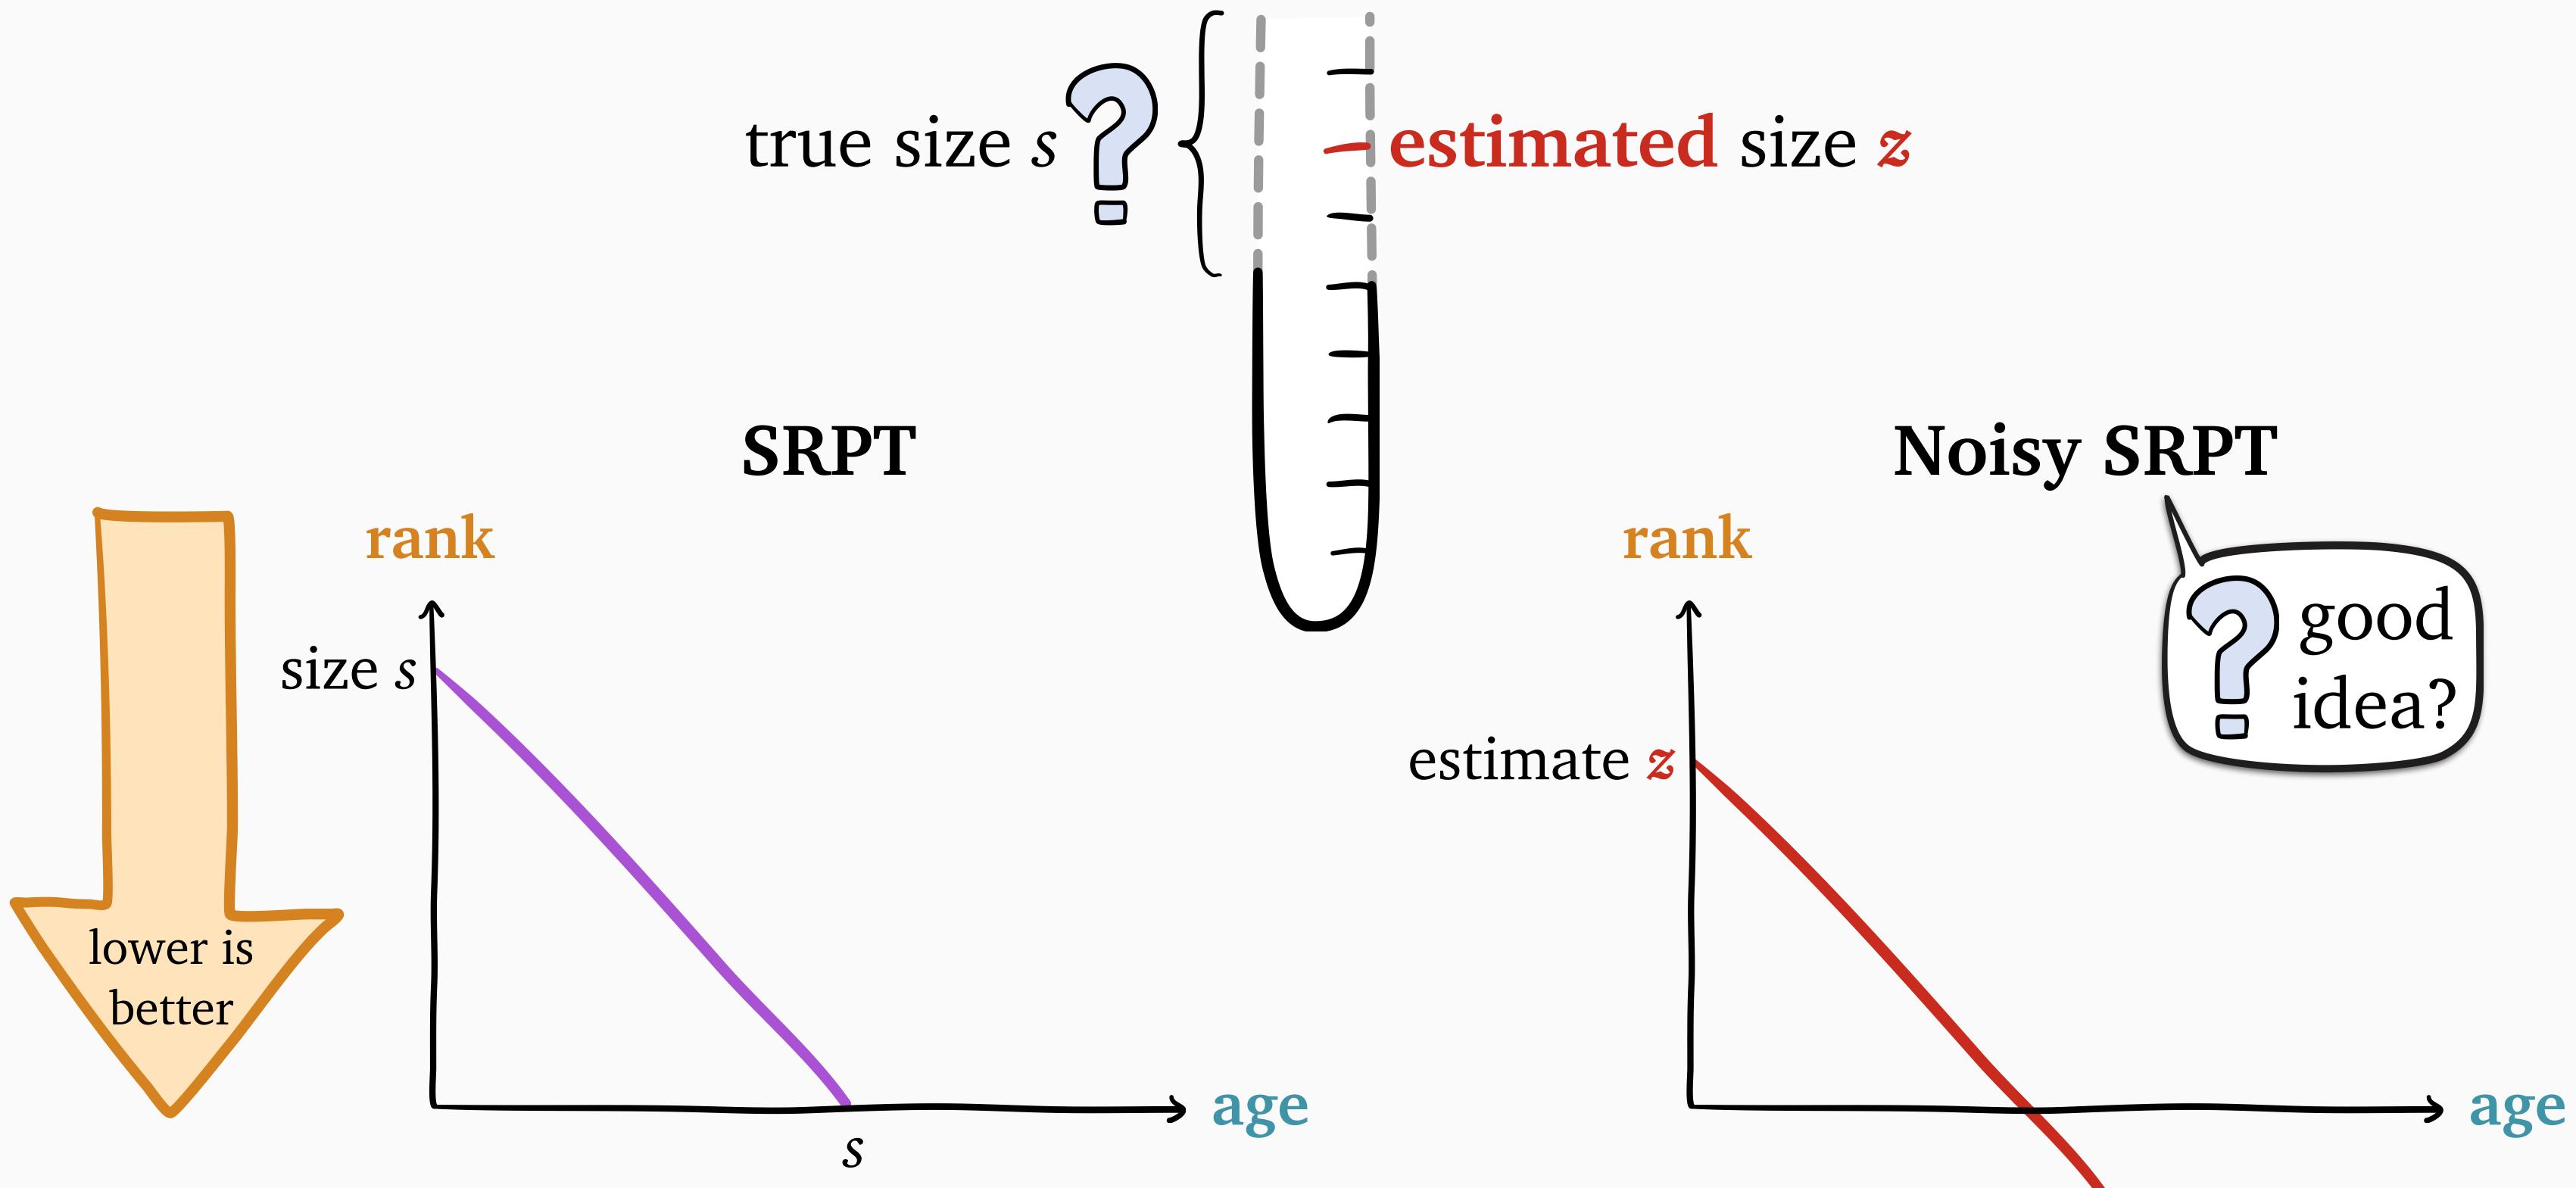
Scheduling with **noisy** size estimates



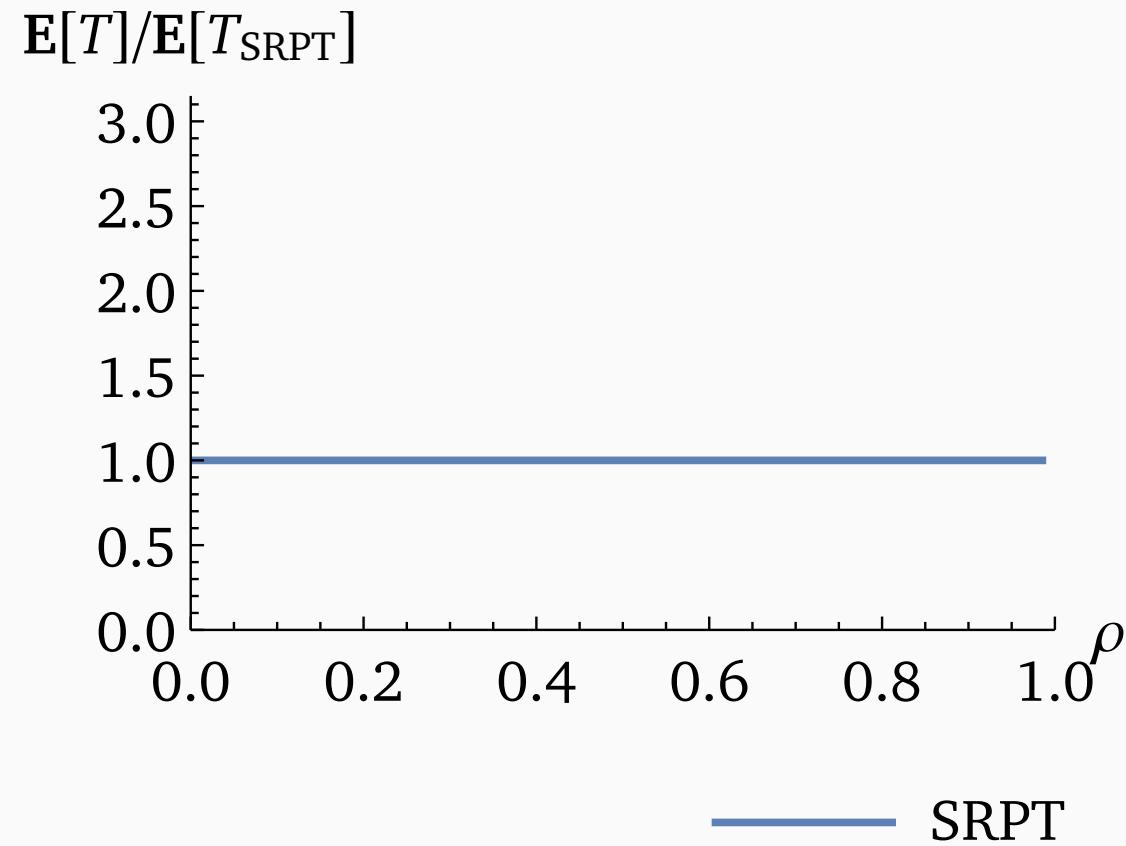
Scheduling with **noisy** size estimates



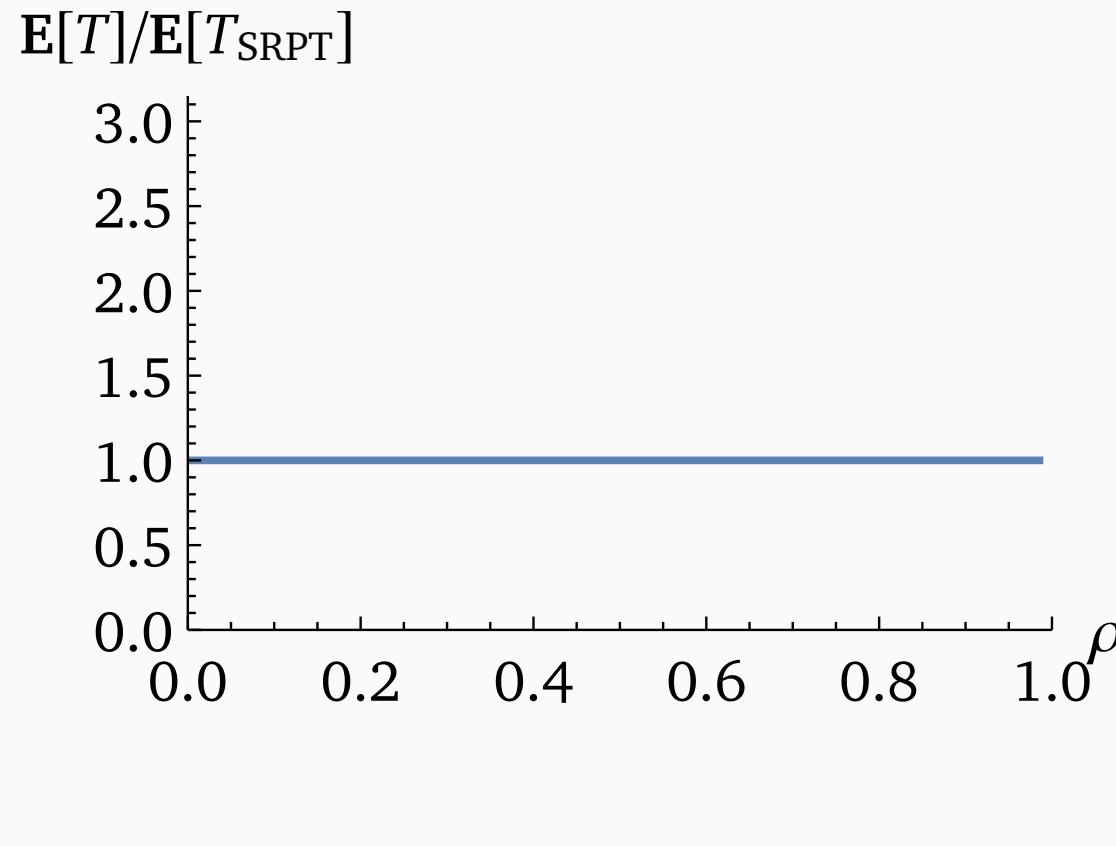
Scheduling with **noisy size estimates**



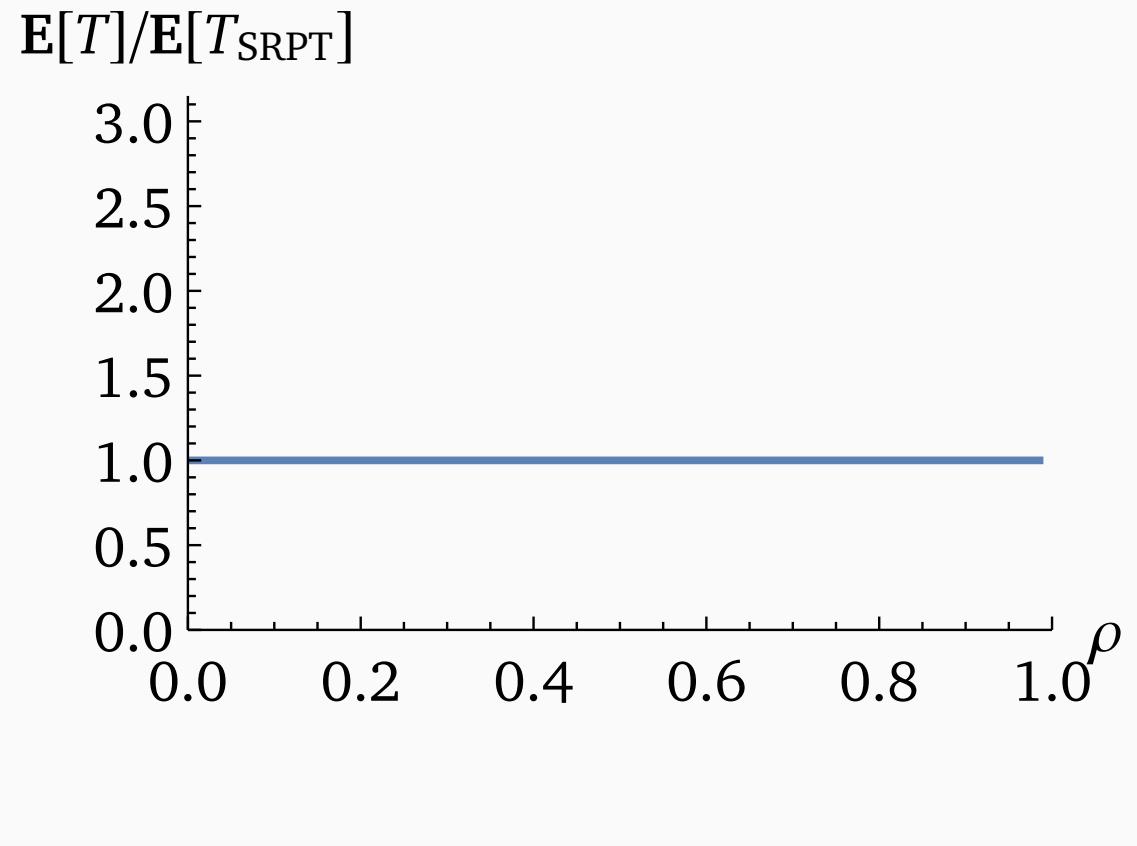
Low noise



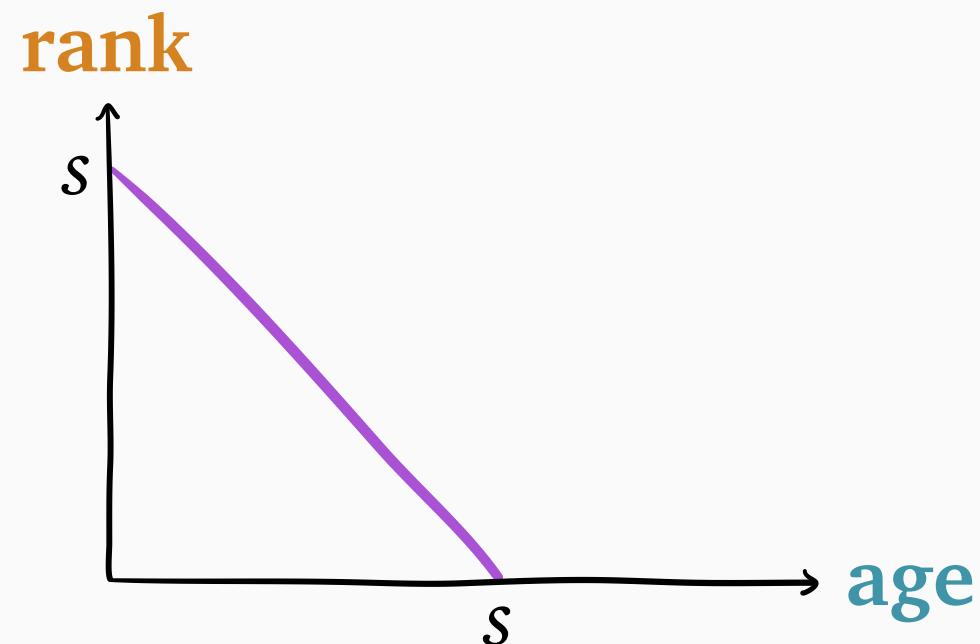
Medium noise



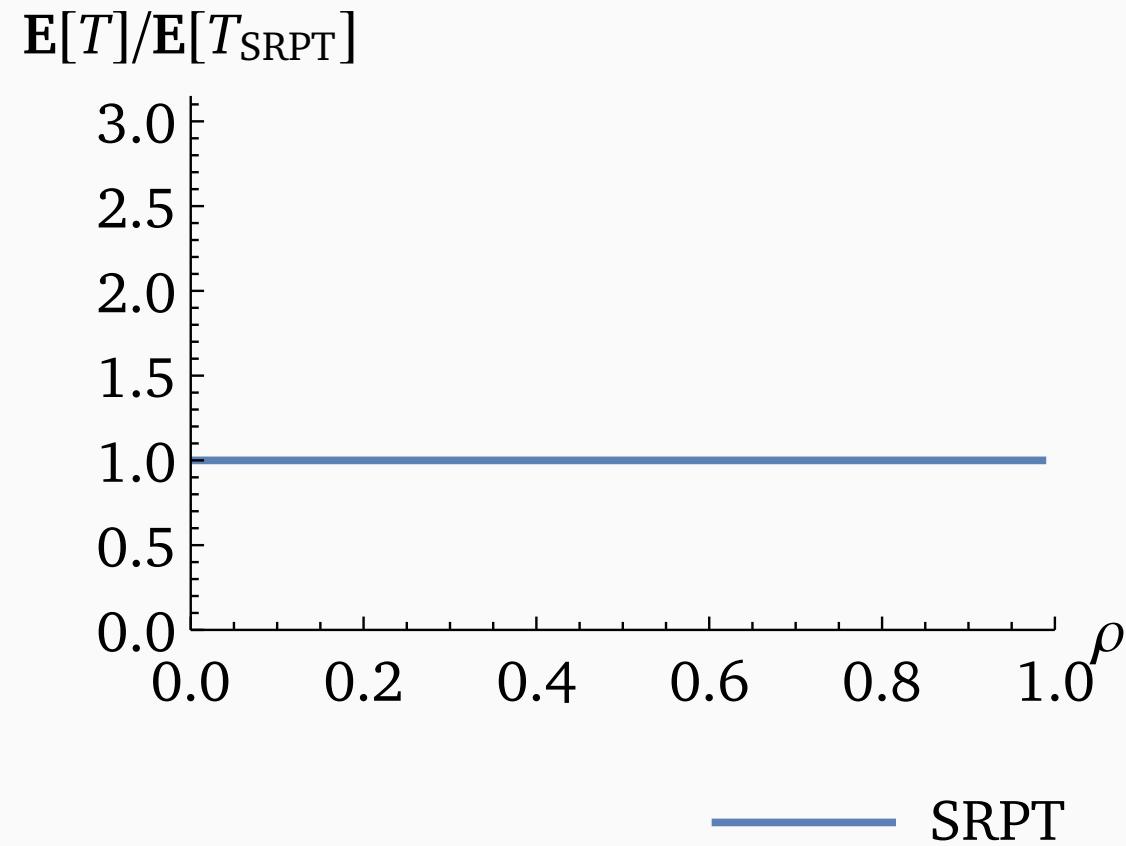
High noise



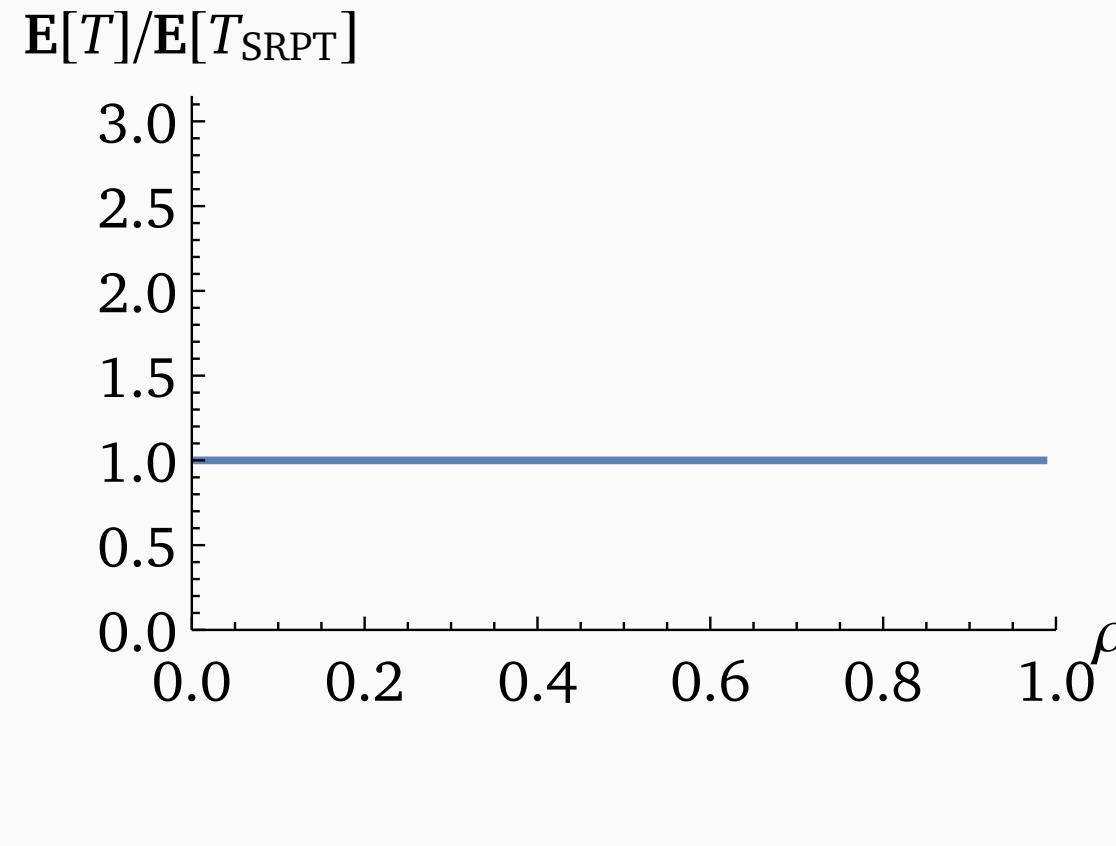
SRPT



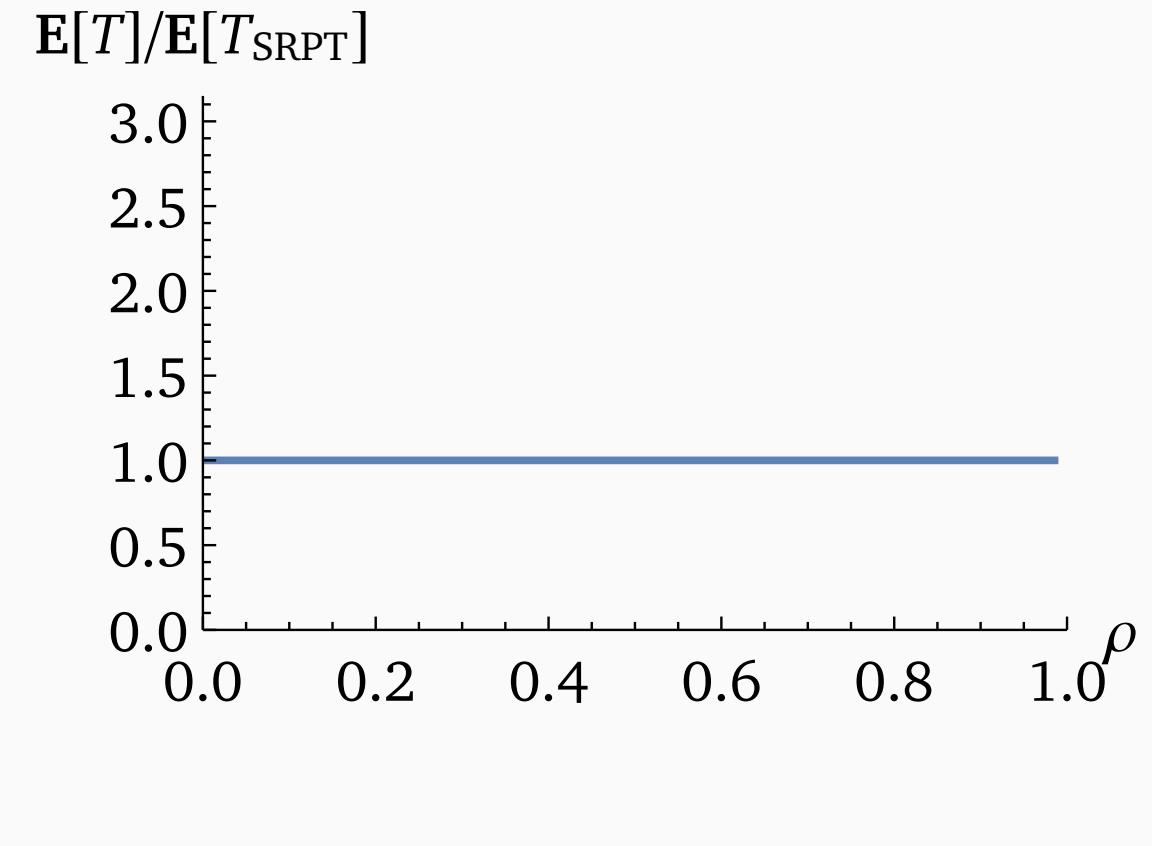
Low noise



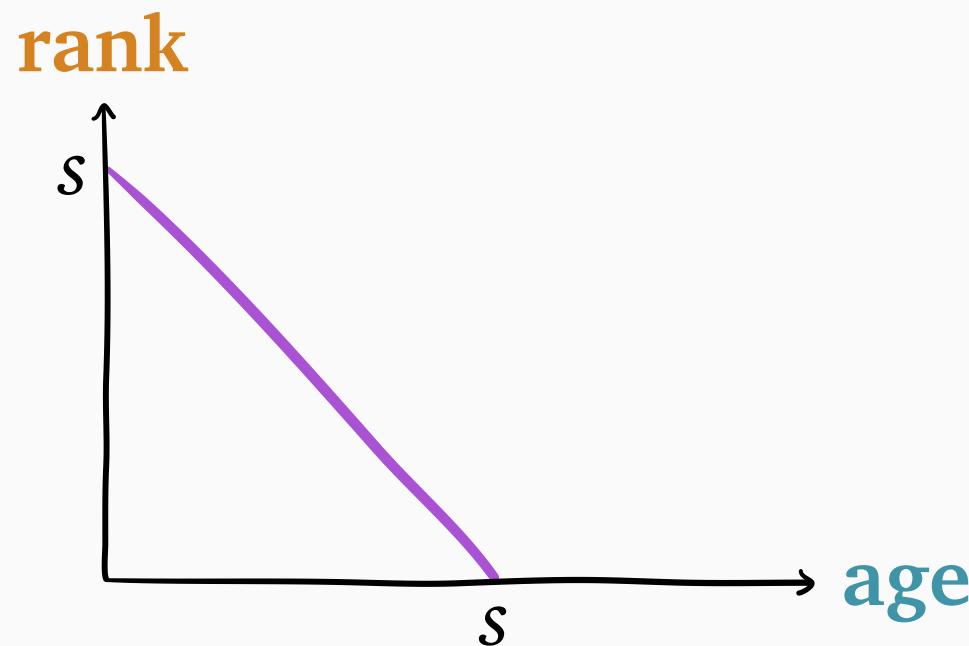
Medium noise



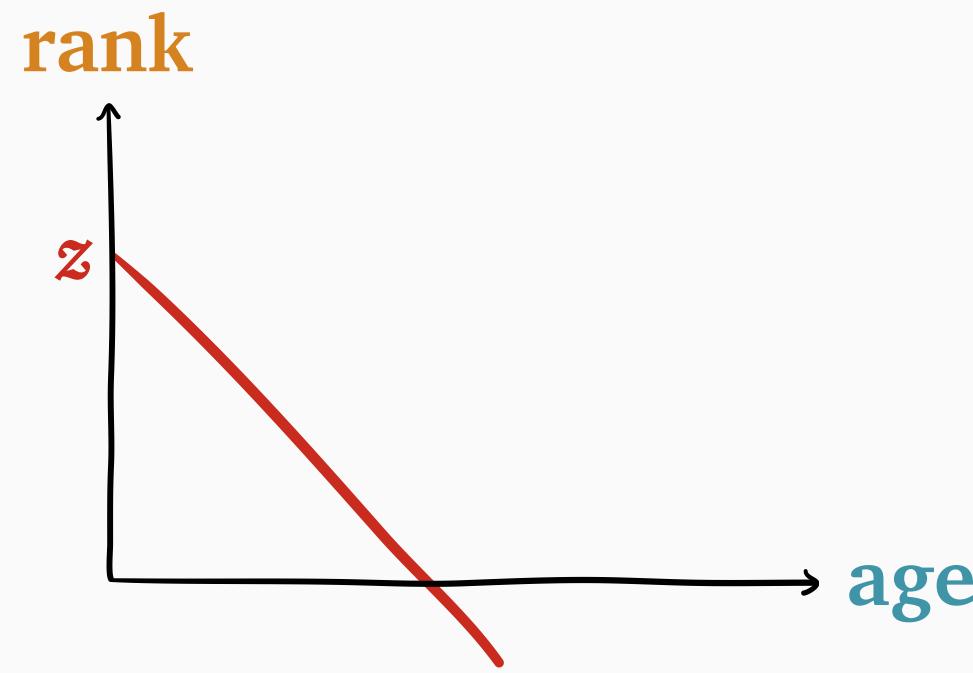
High noise



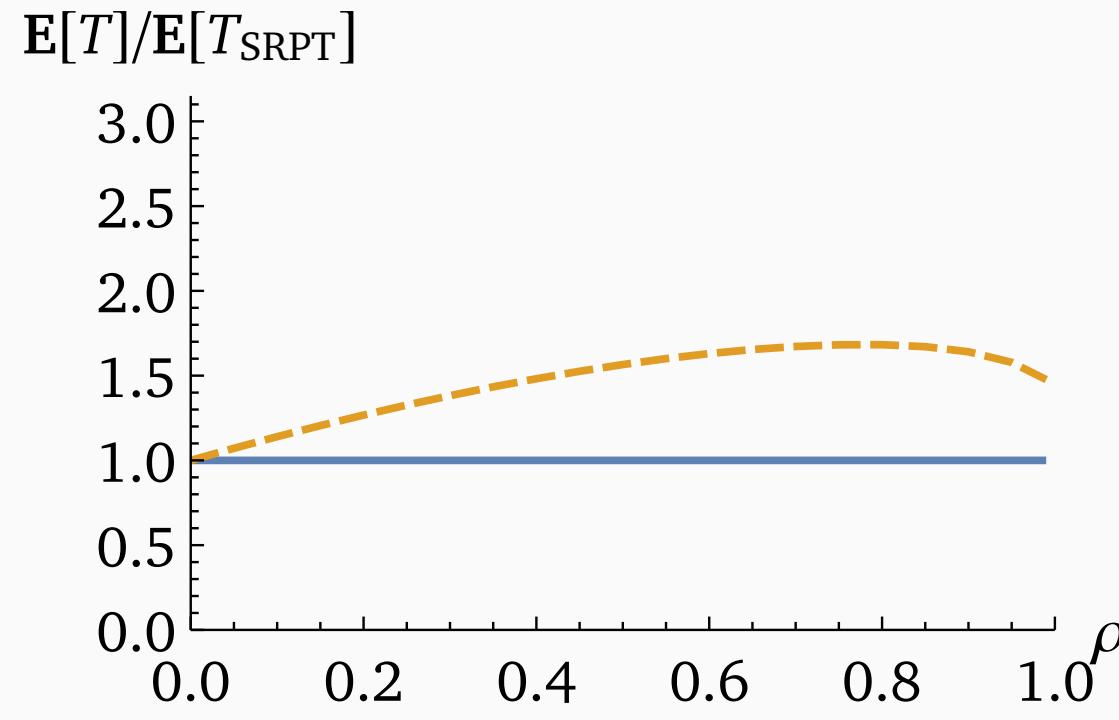
SRPT



Noisy SRPT

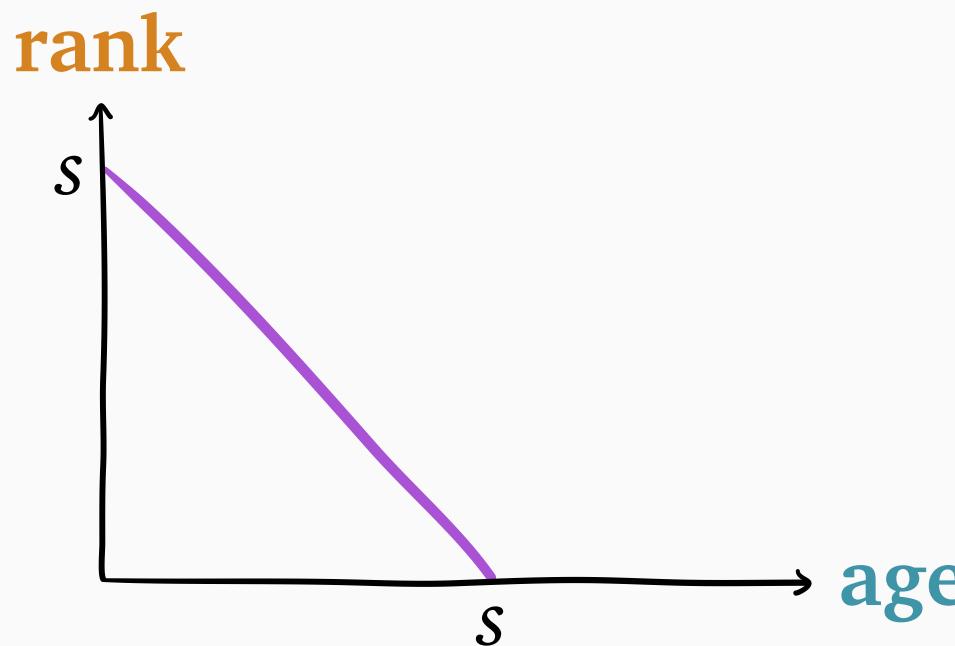


Low noise

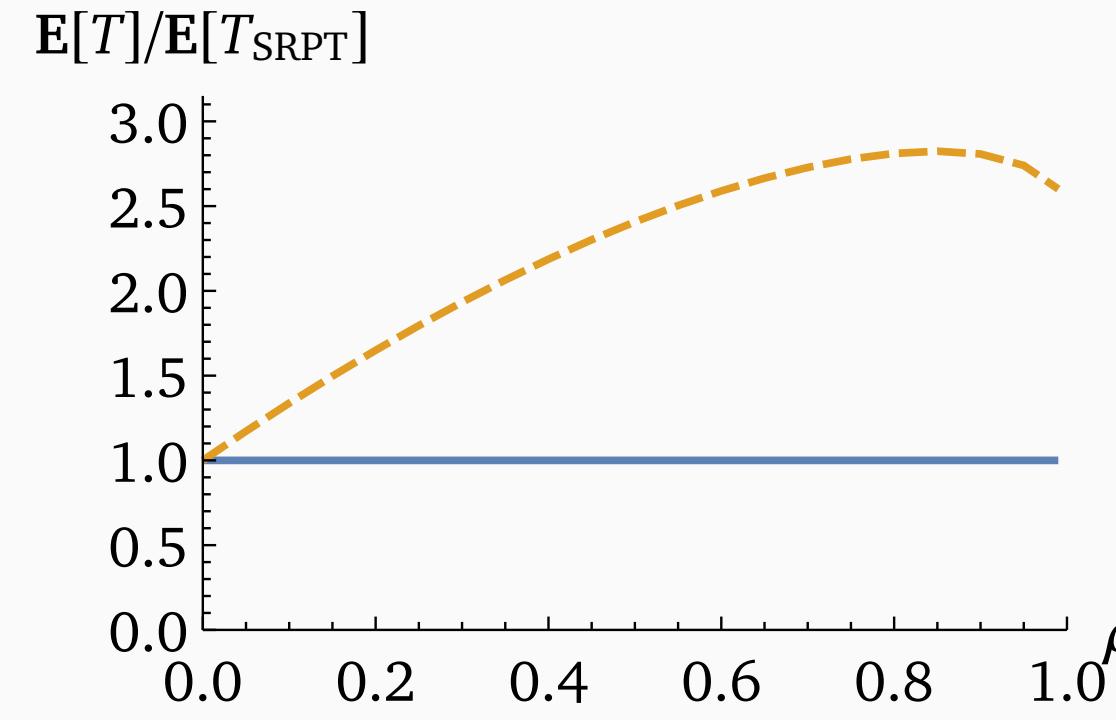


— SRPT
- - - Noisy SRPT

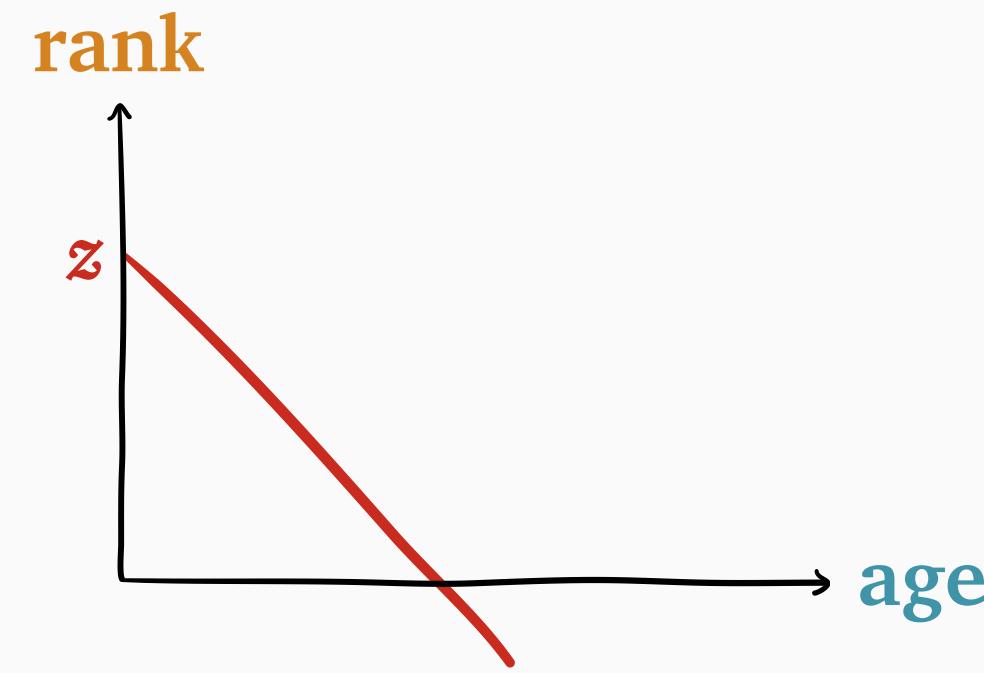
SRPT



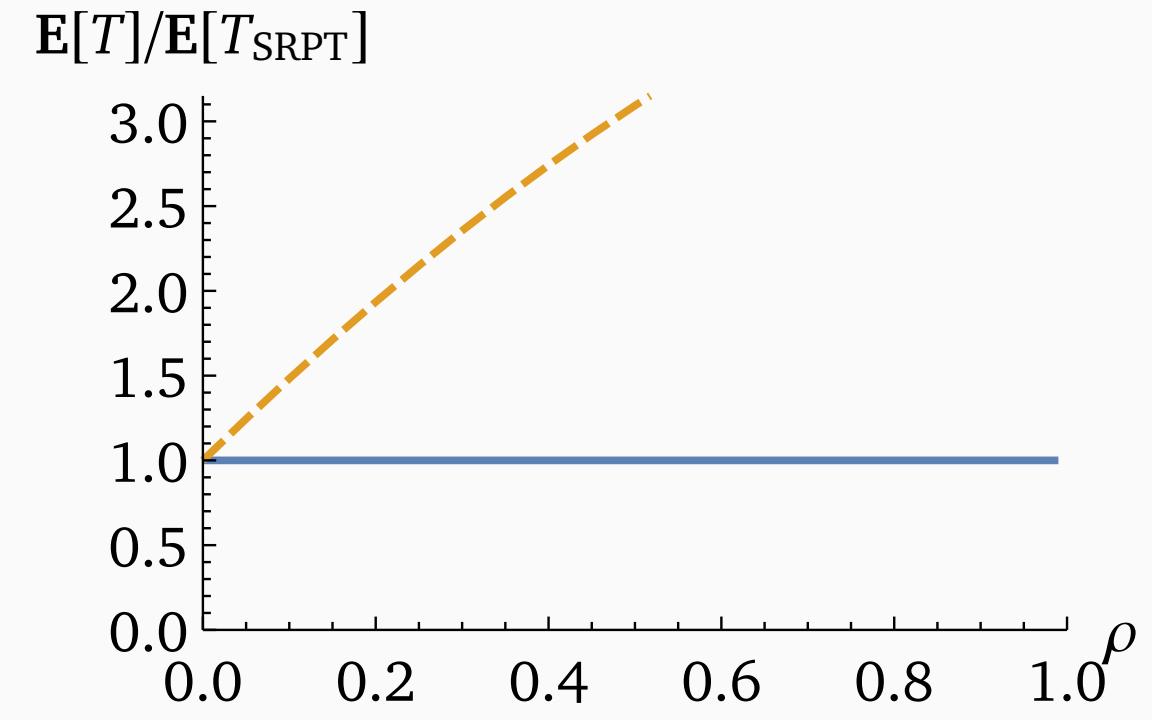
Medium noise



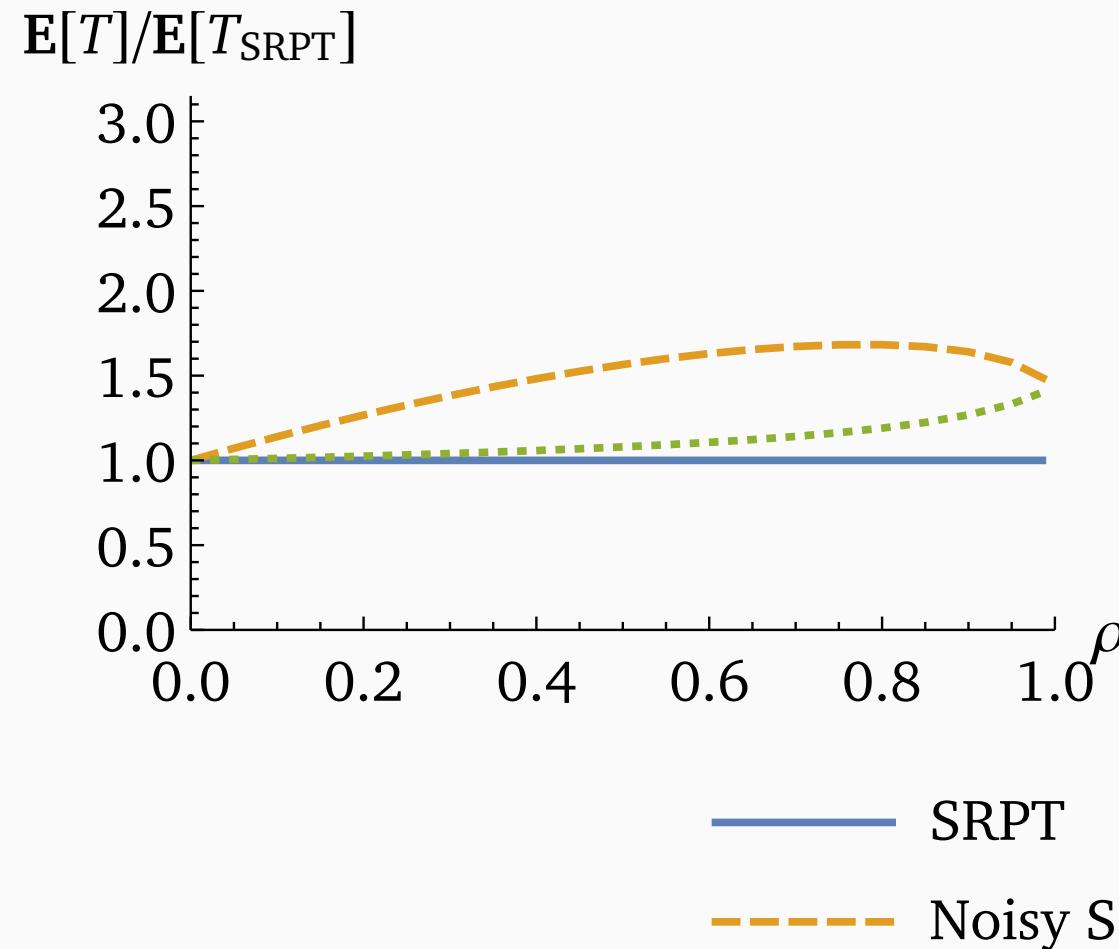
Noisy SRPT



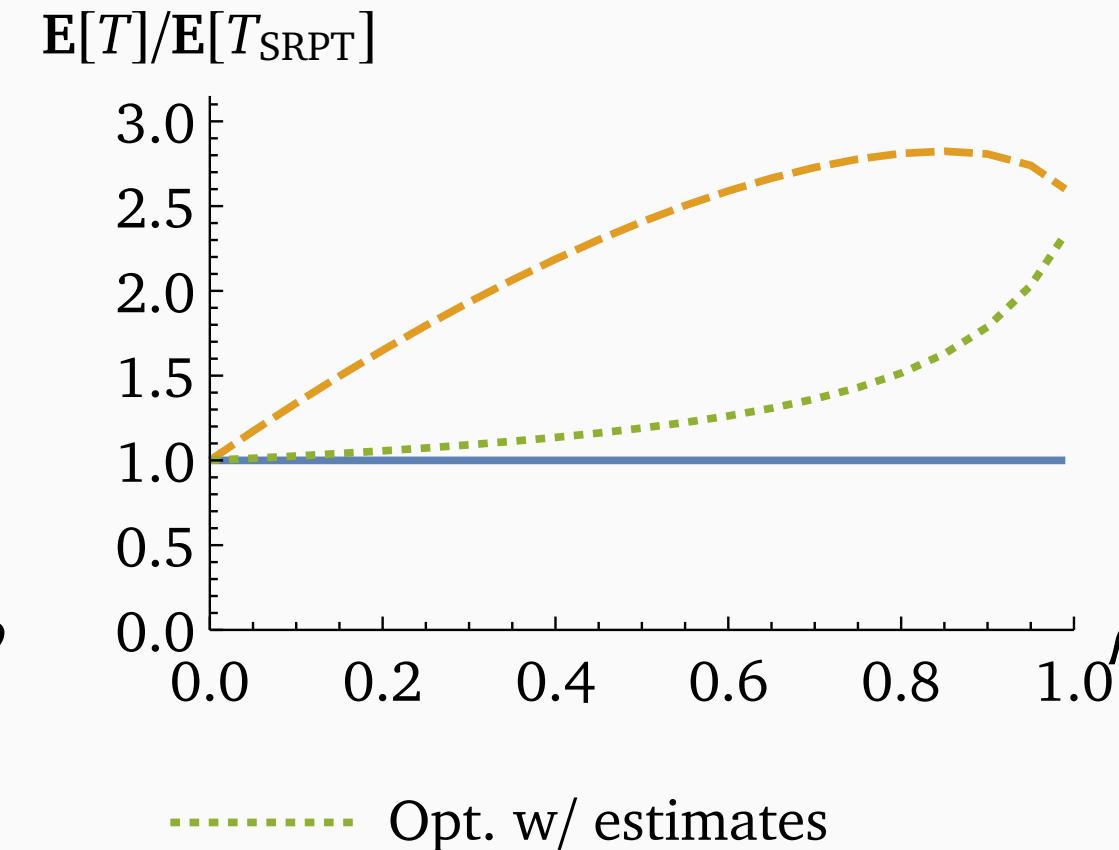
High noise



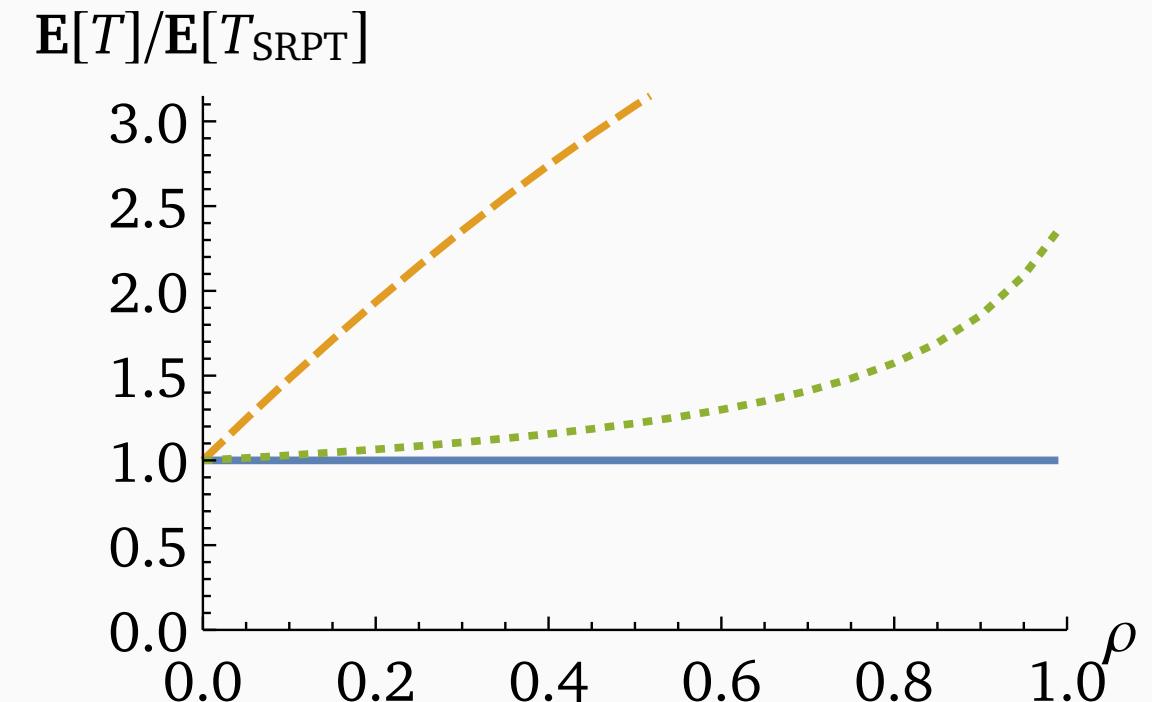
Low noise



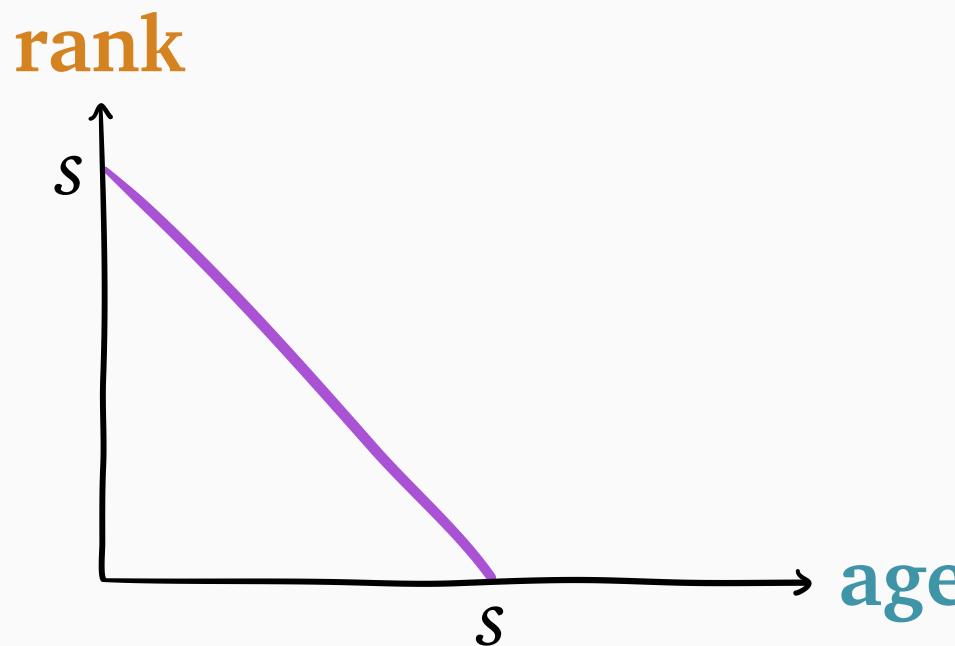
Medium noise



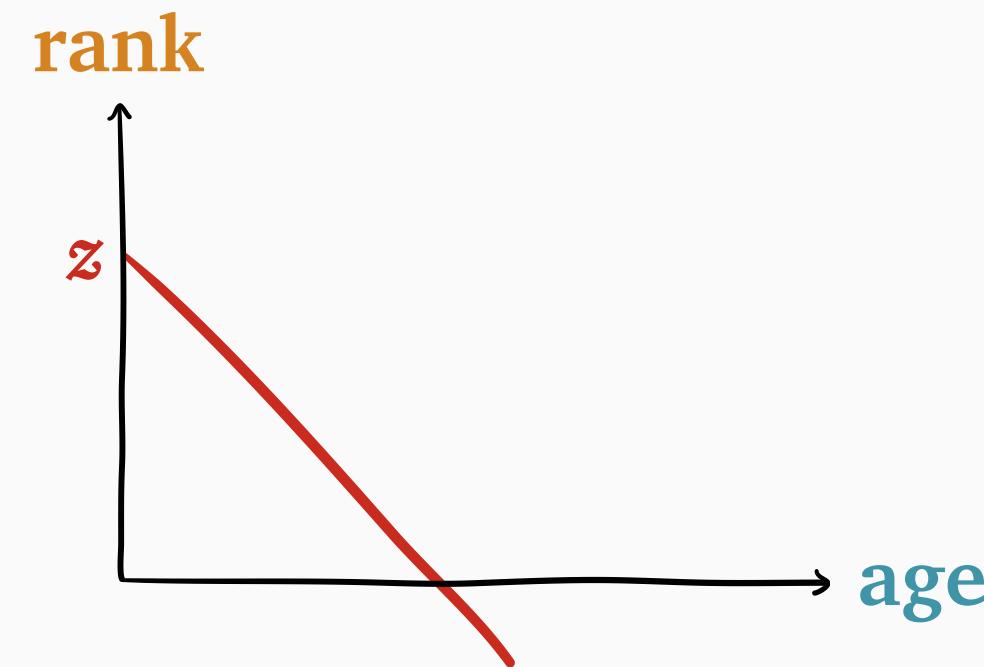
High noise



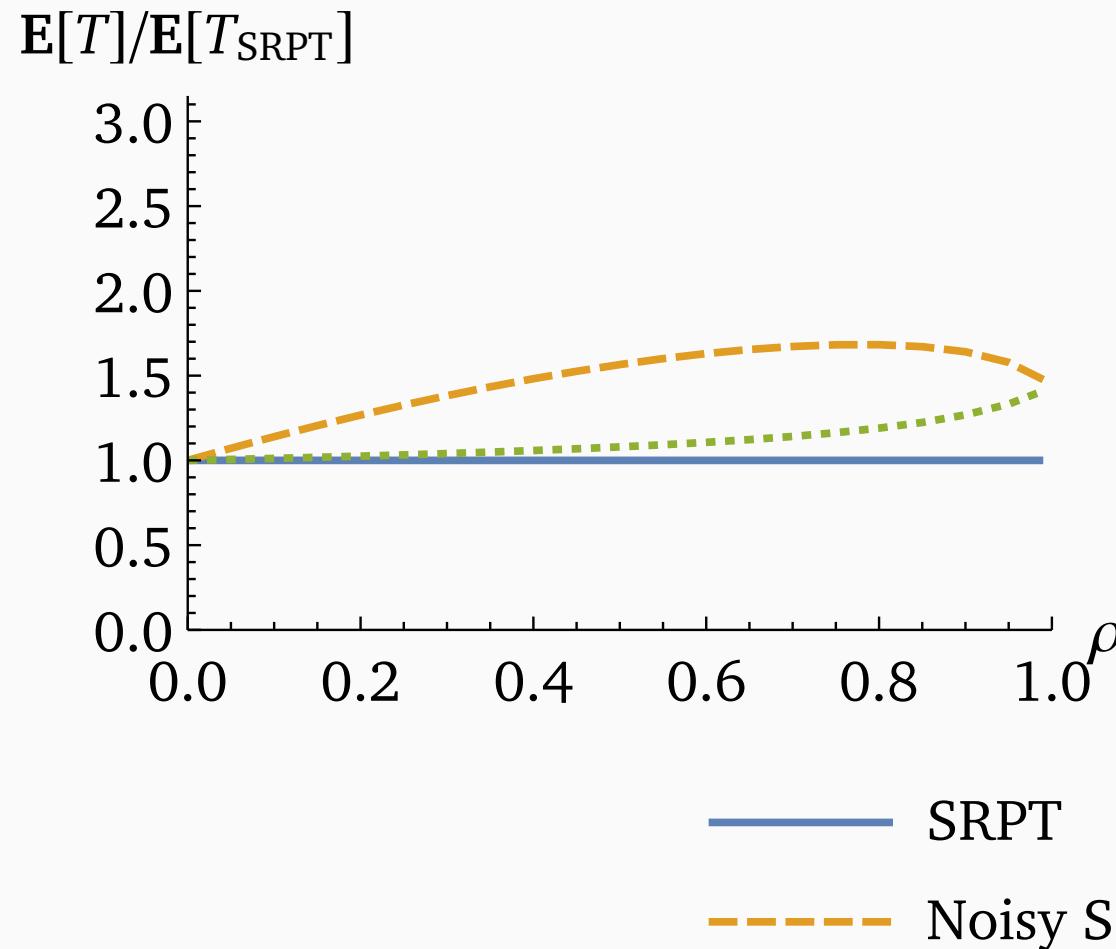
SRPT



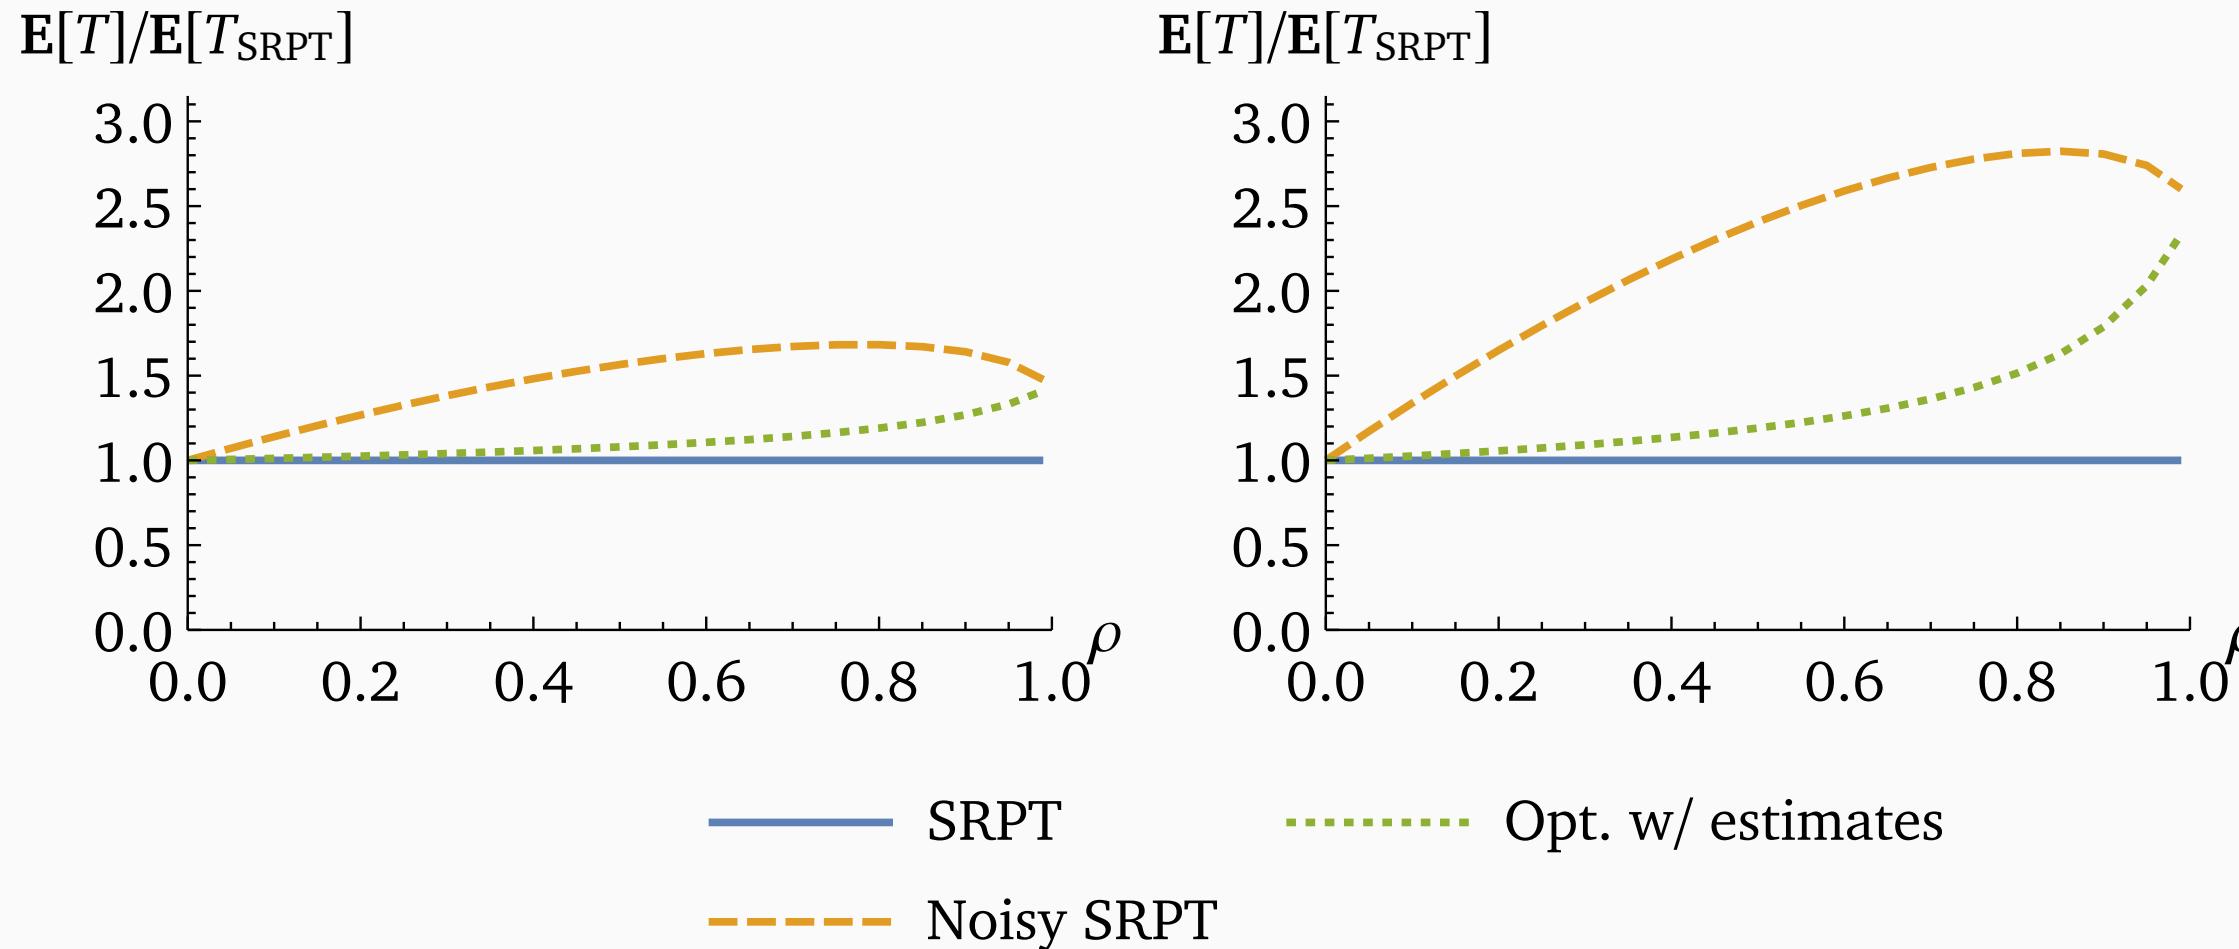
Noisy SRPT



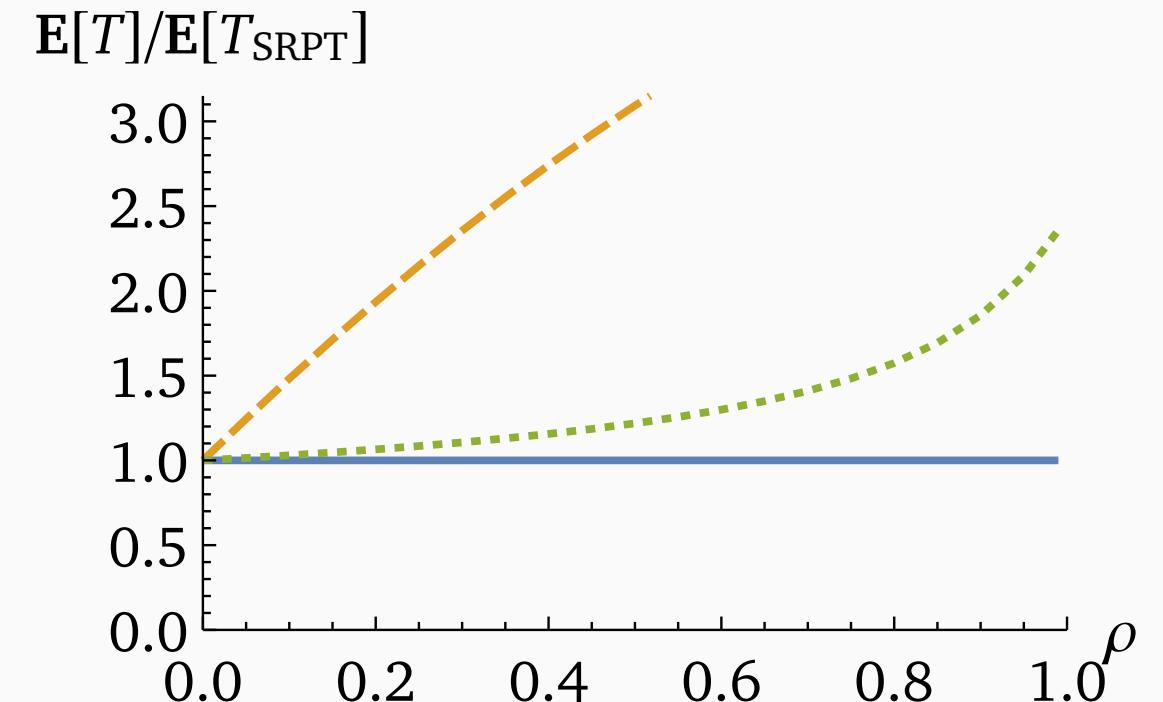
Low noise



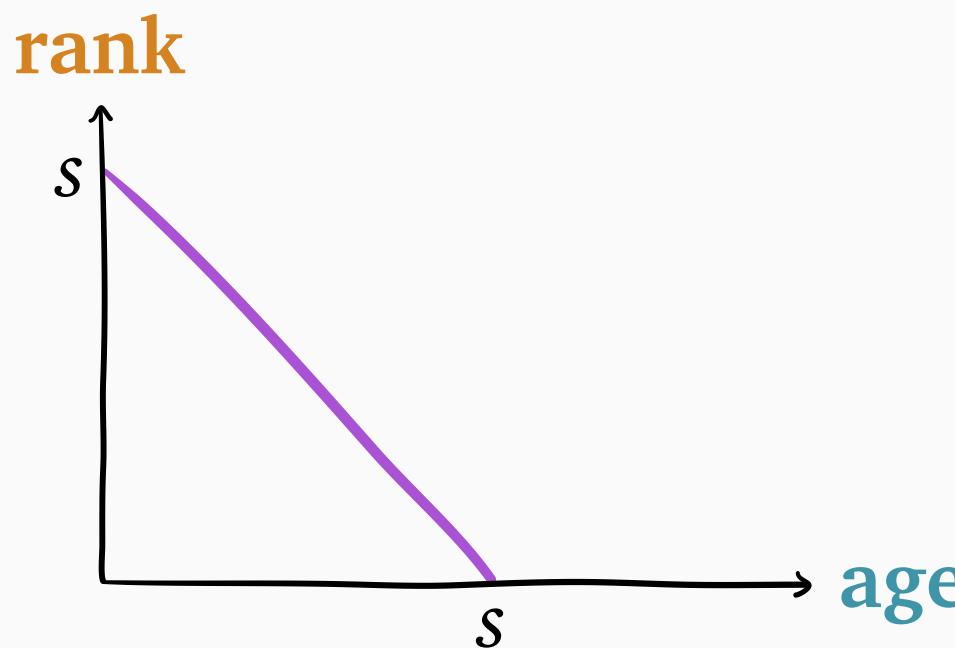
Medium noise



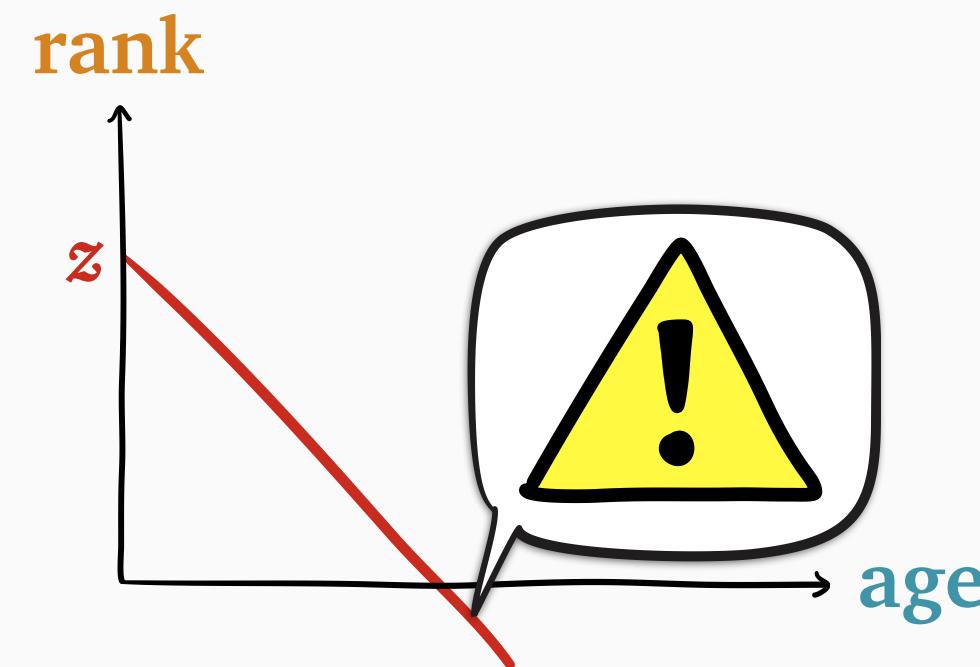
High noise



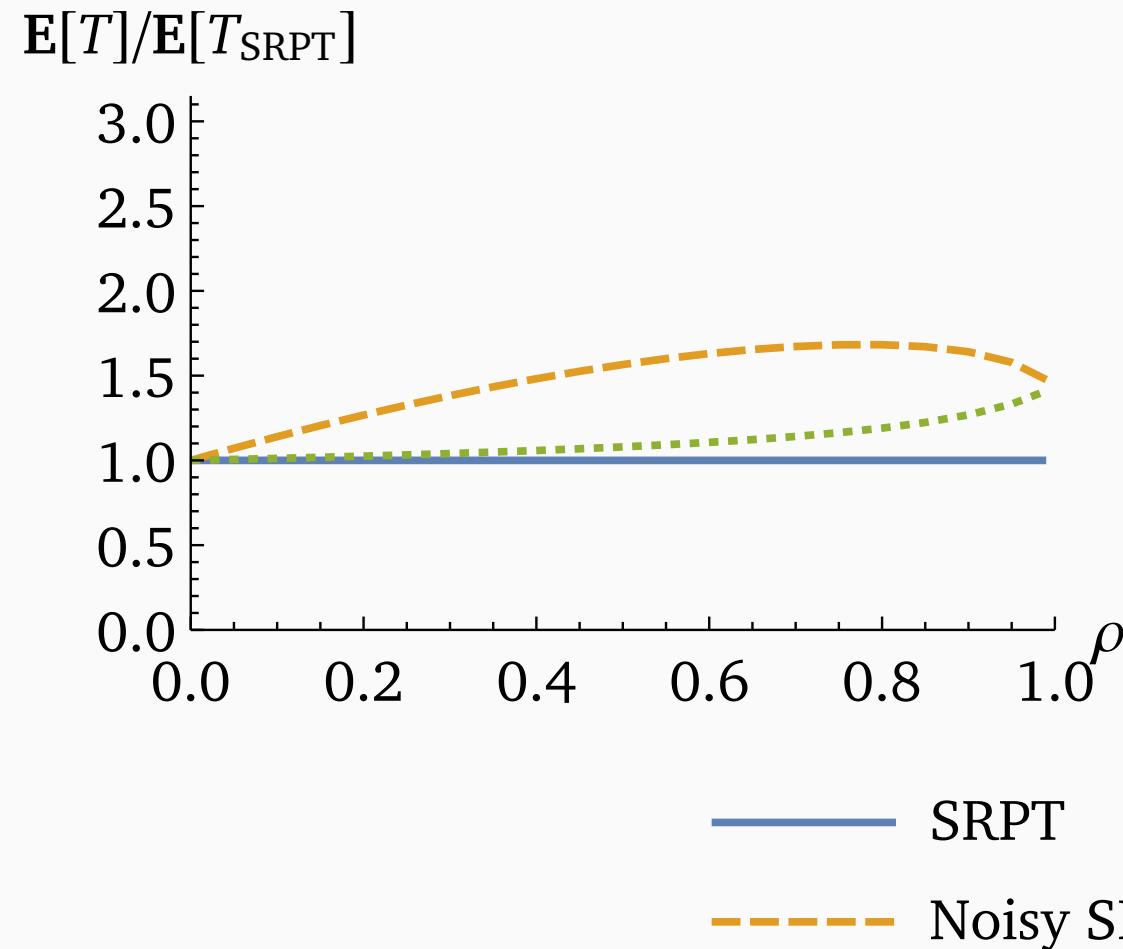
SRPT



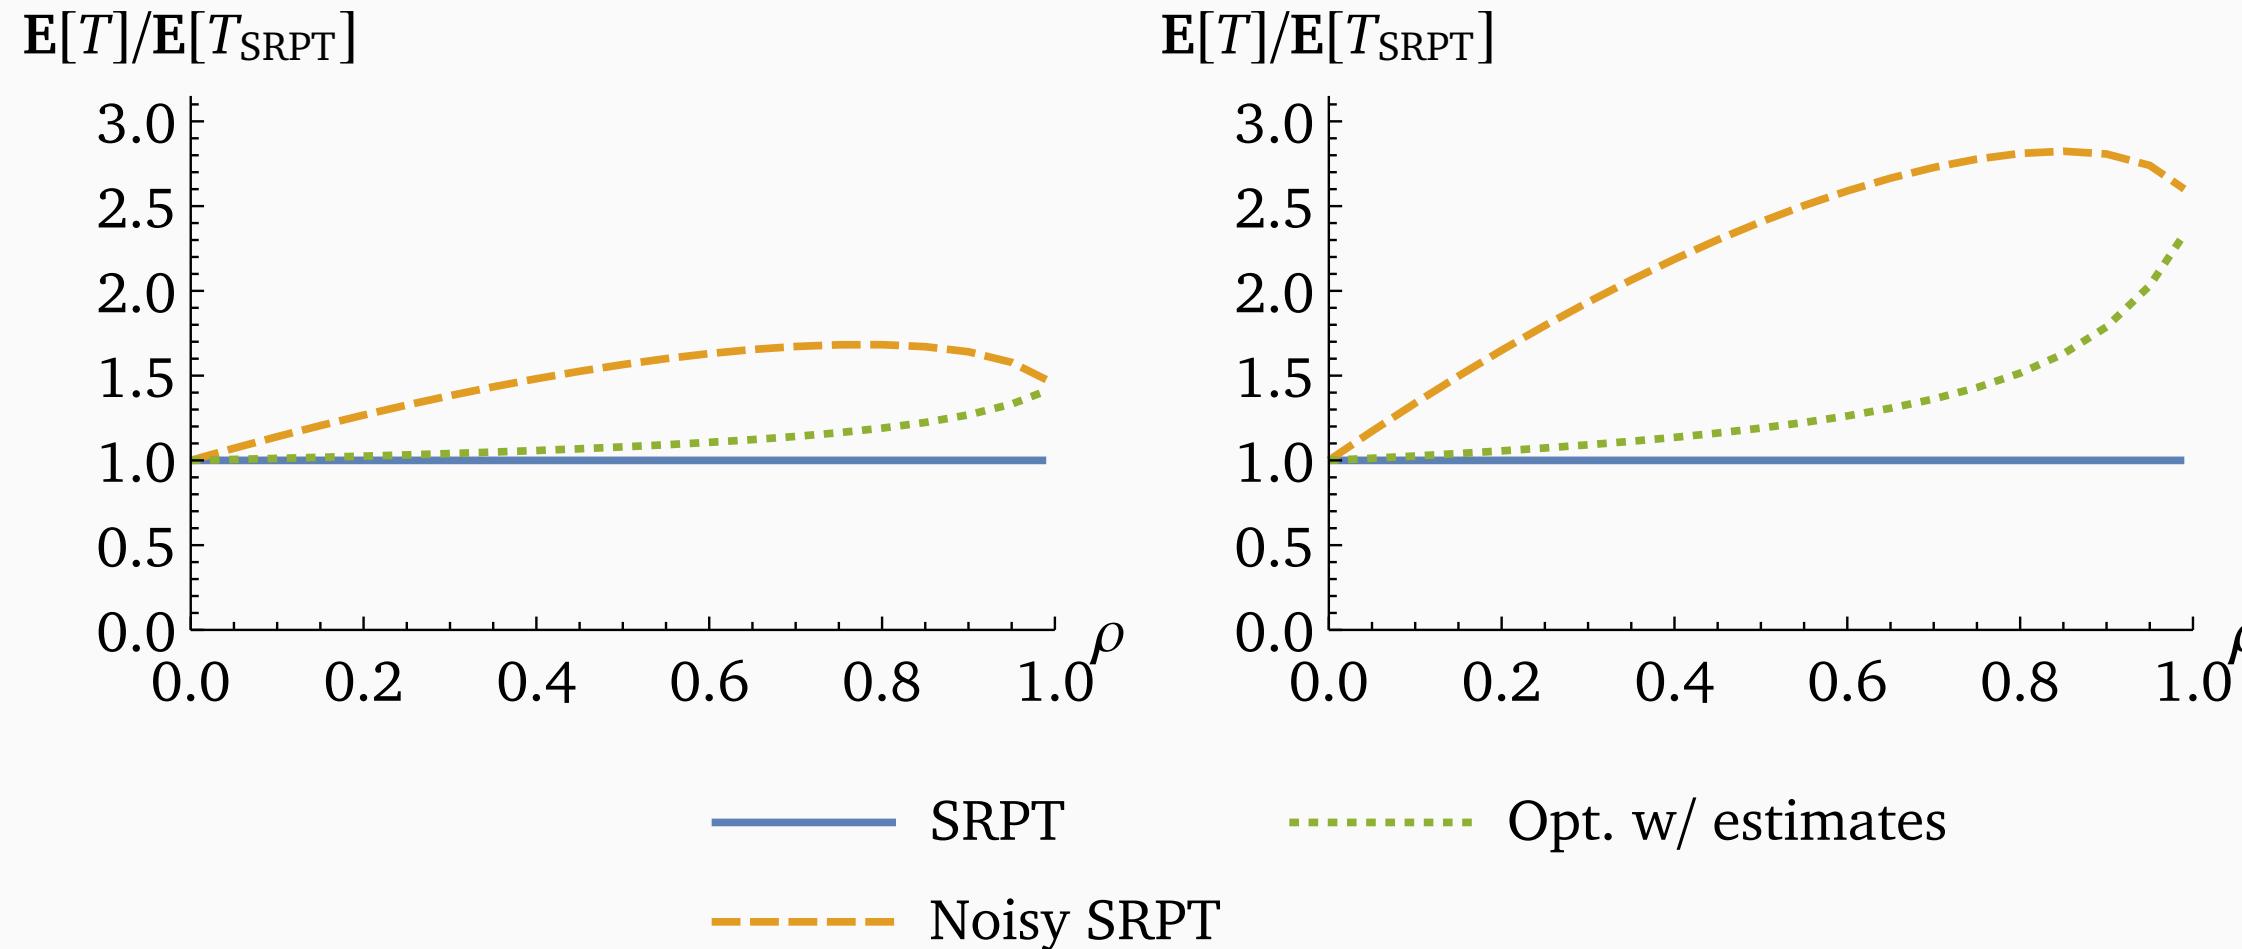
Noisy SRPT



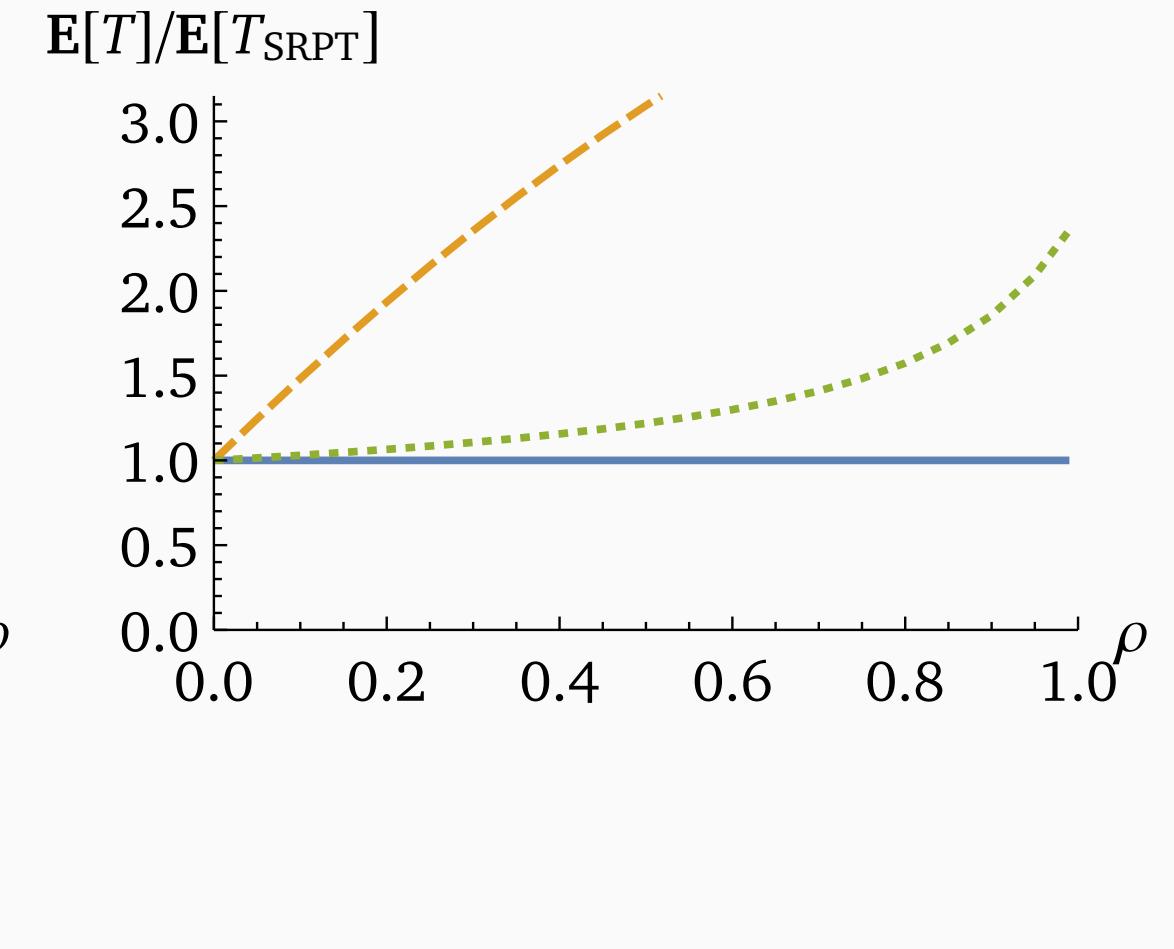
Low noise



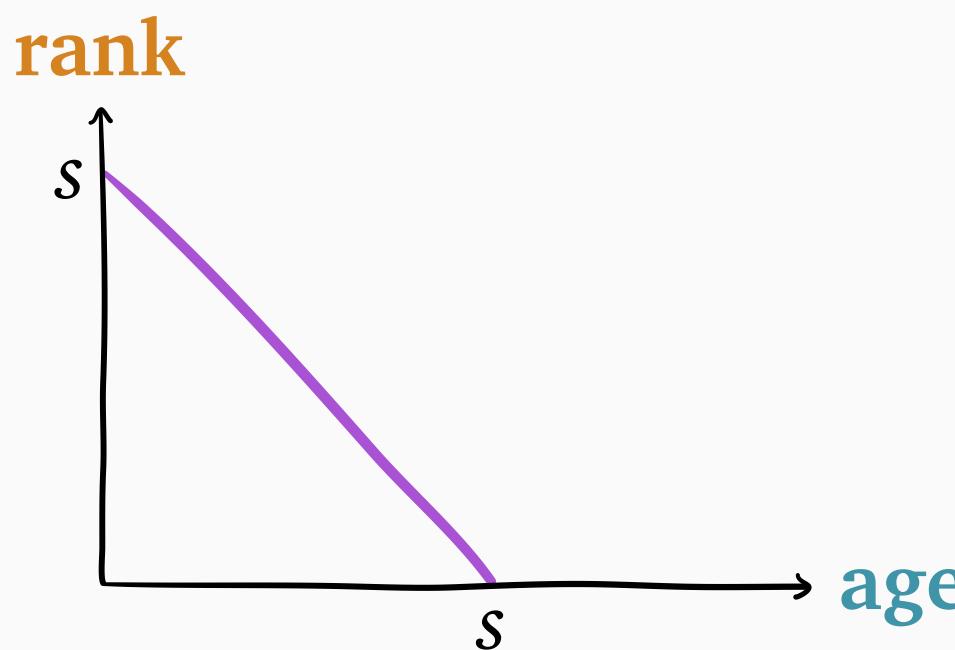
Medium noise



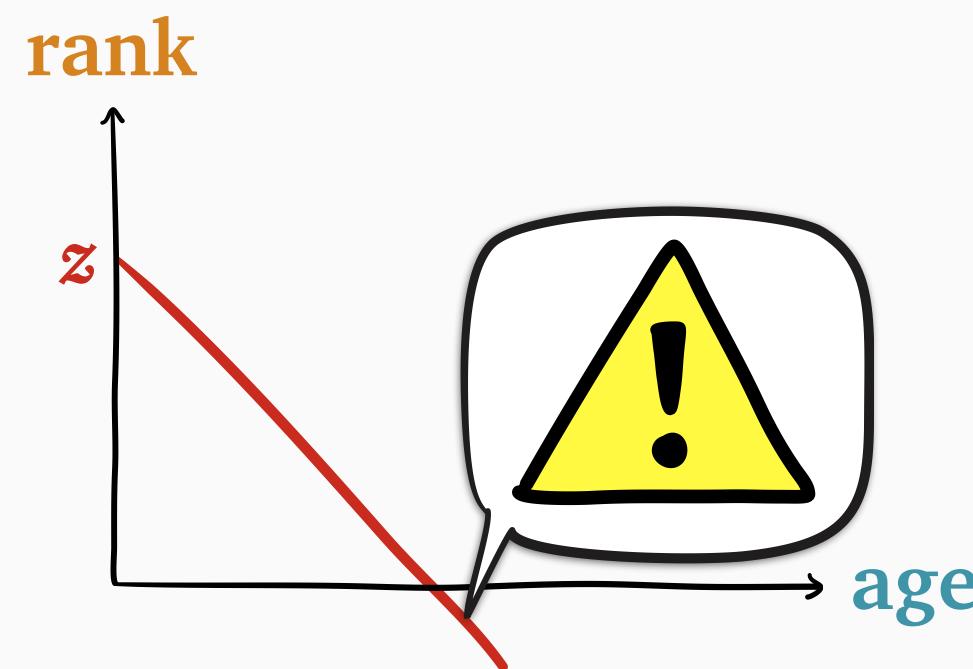
High noise



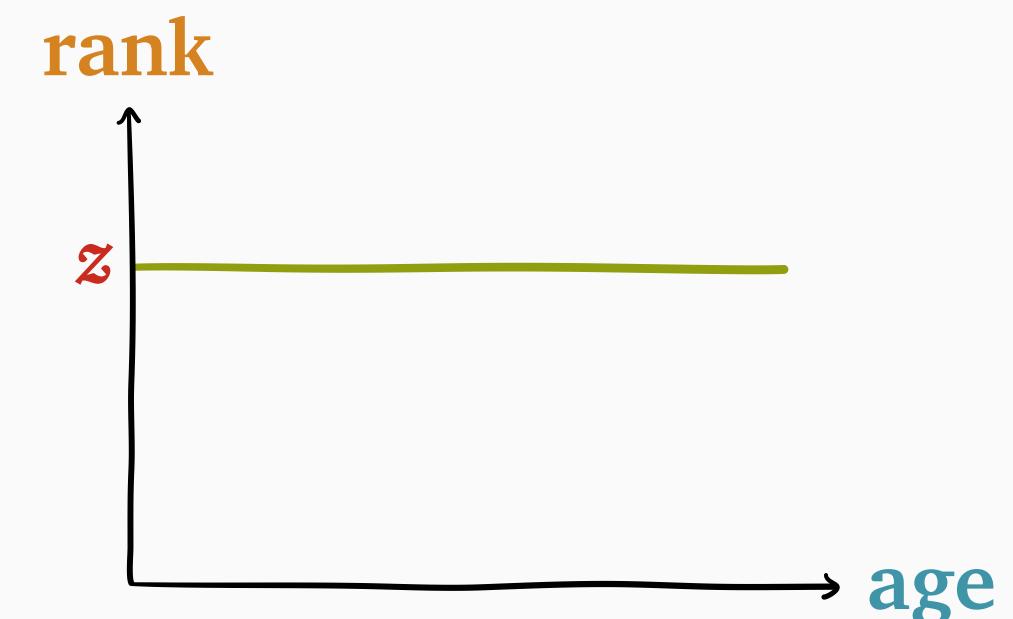
SRPT



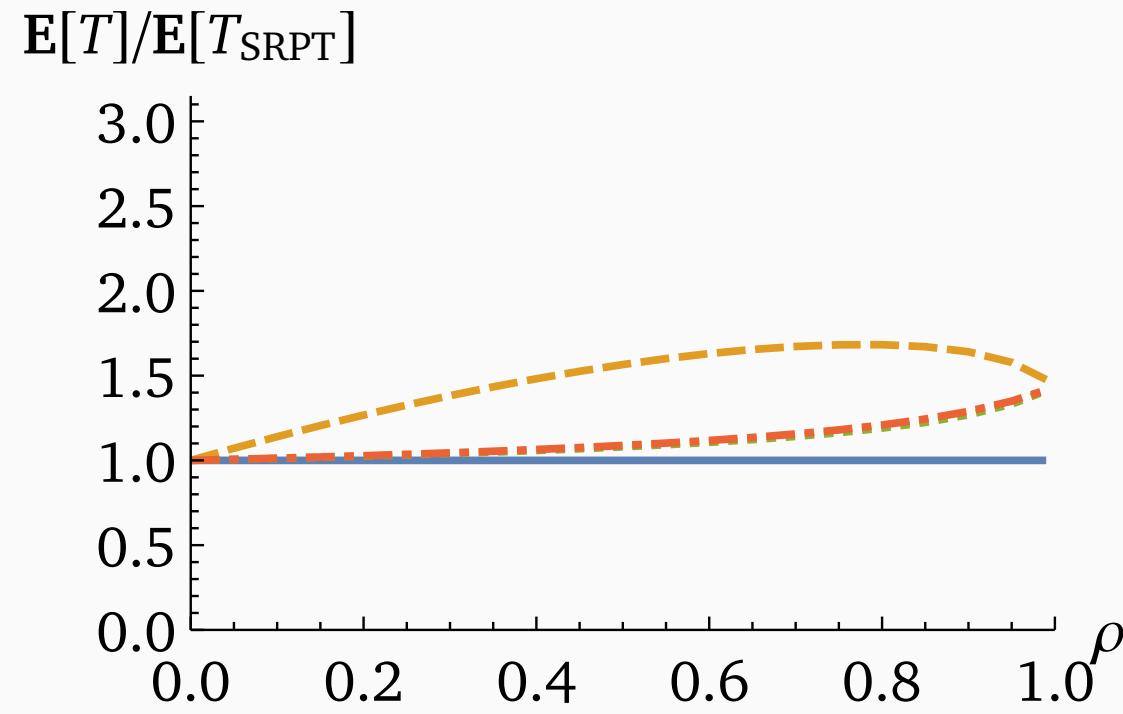
Noisy SRPT



Noisy PSJF

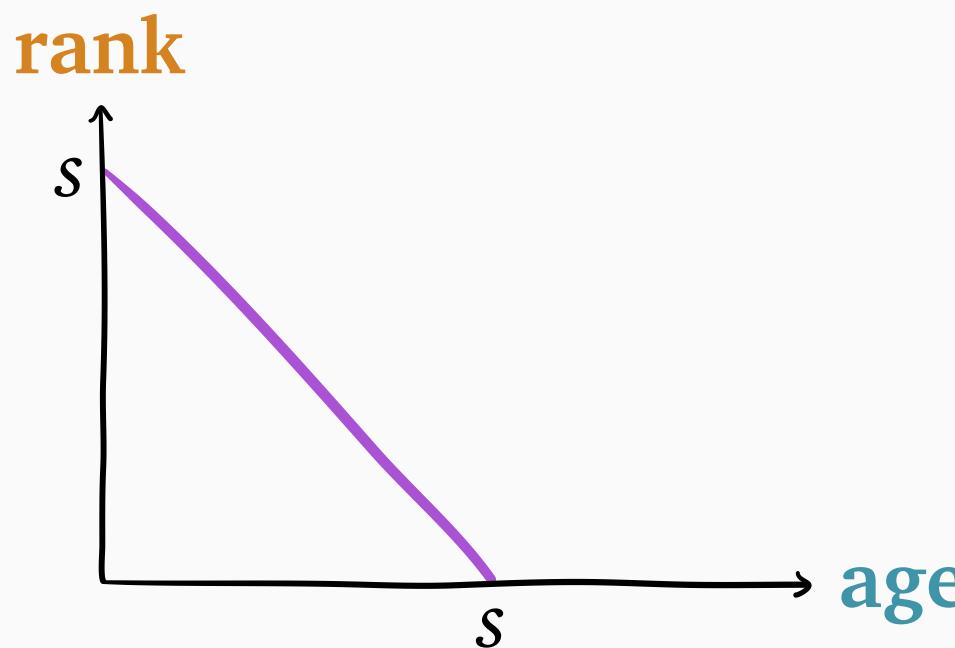


Low noise

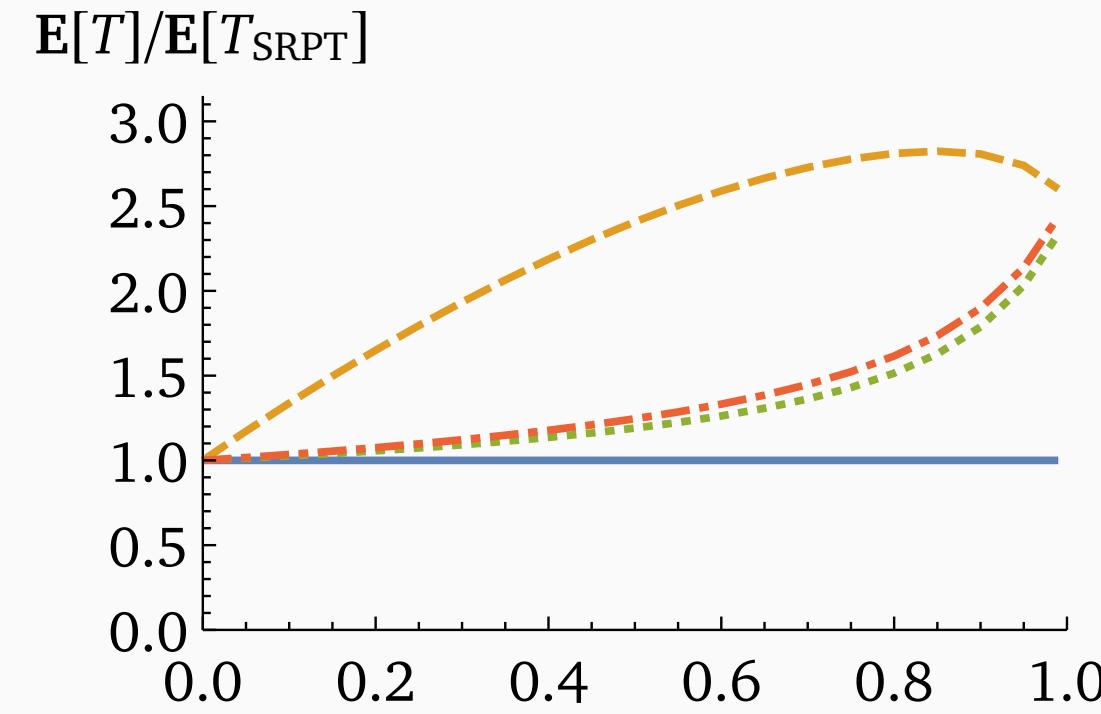


— SRPT
- - - Noisy SRPT

SRPT

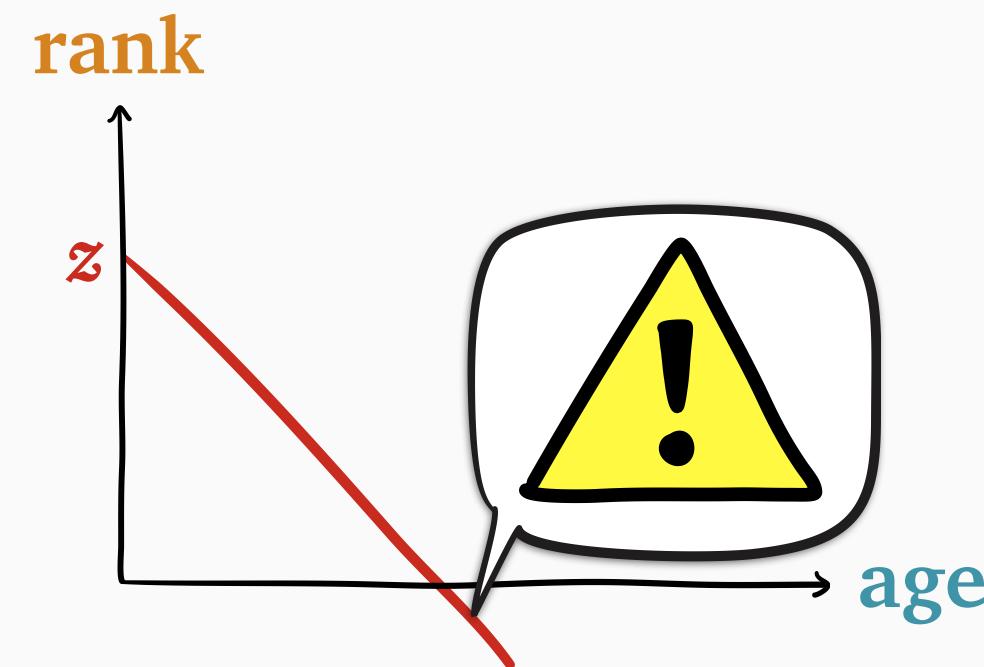


Medium noise

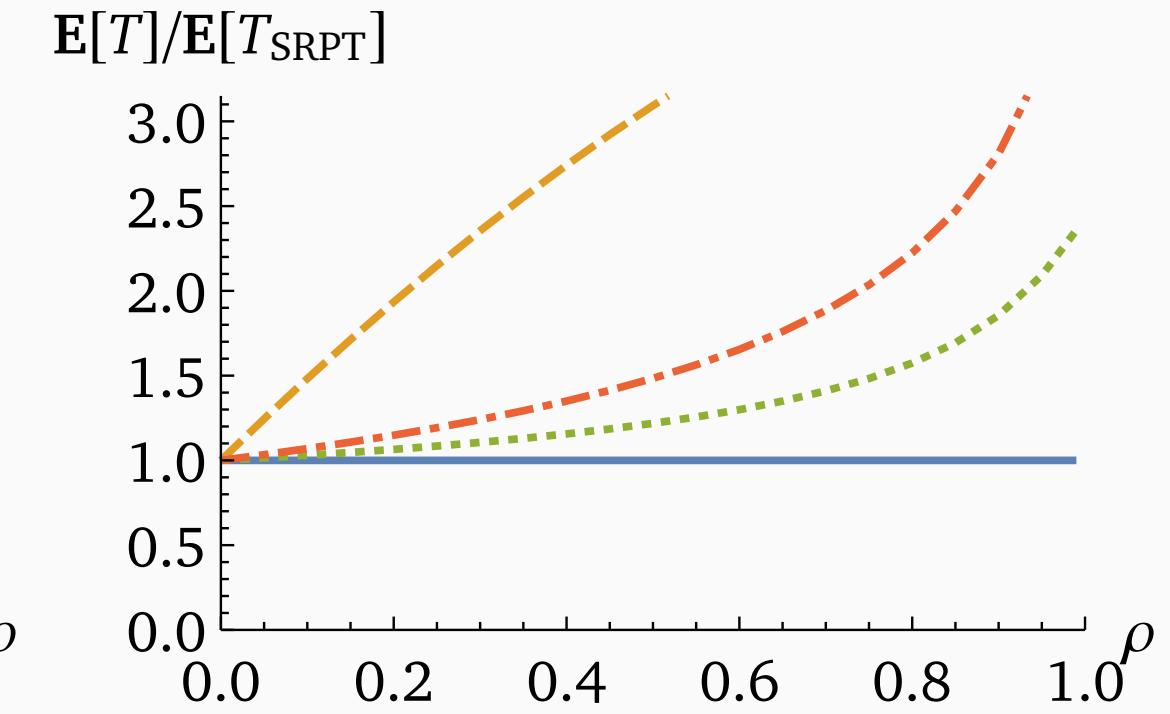


... Opt. w/ estimates
- - - Noisy PSJF

Noisy SRPT

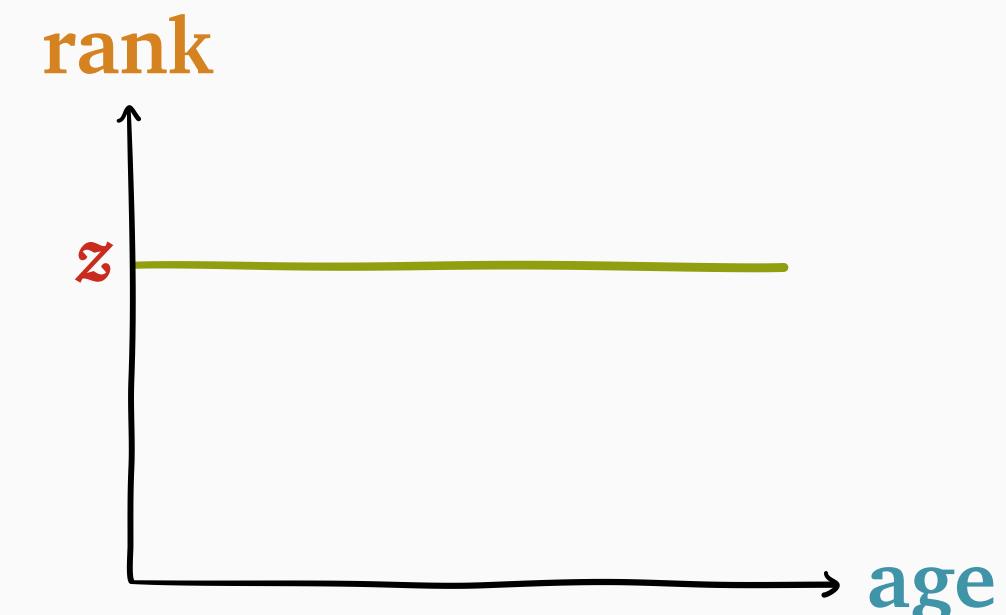


High noise

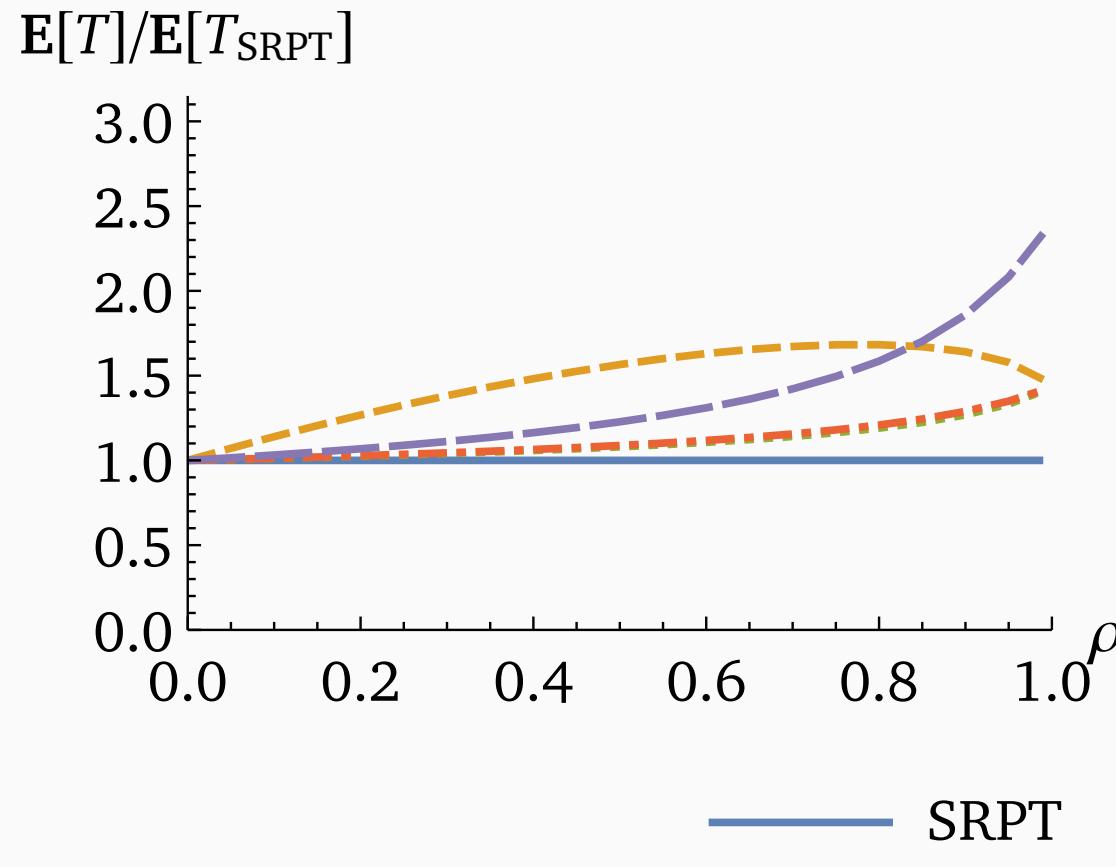


— SRPT
- - - Noisy SRPT

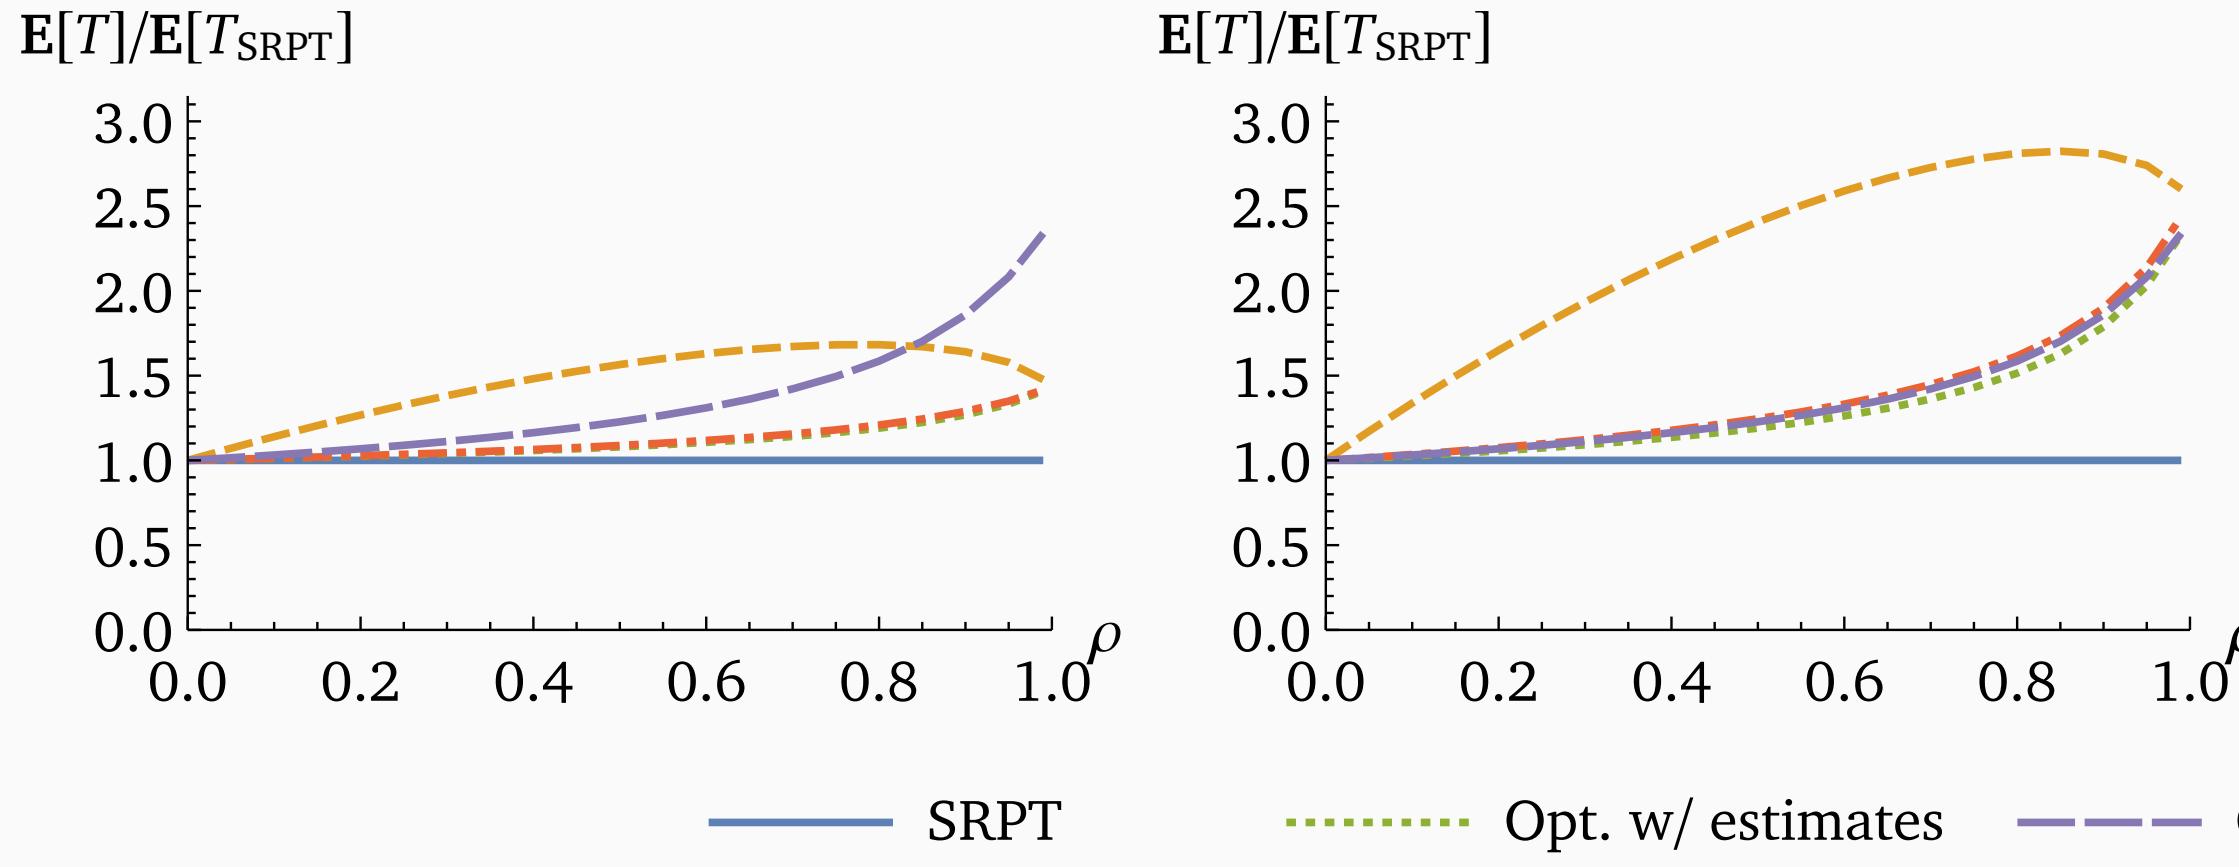
Noisy PSJF



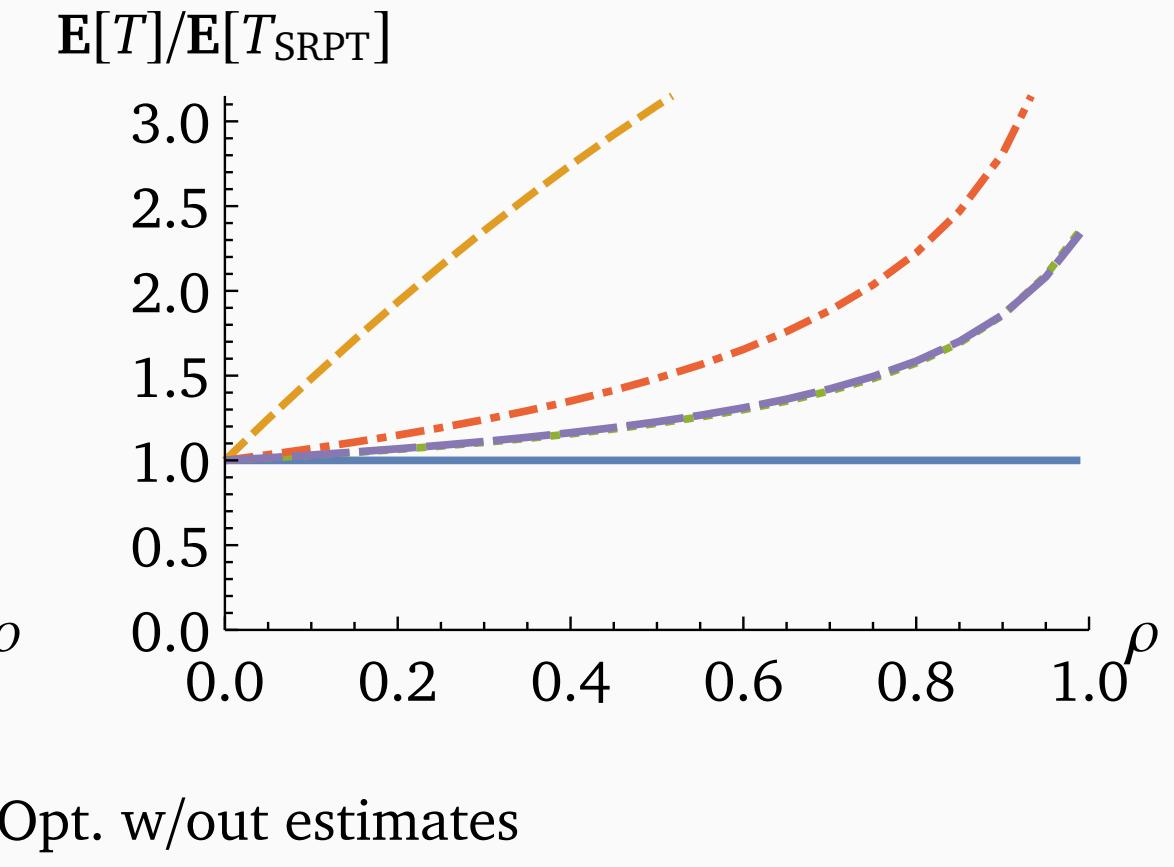
Low noise



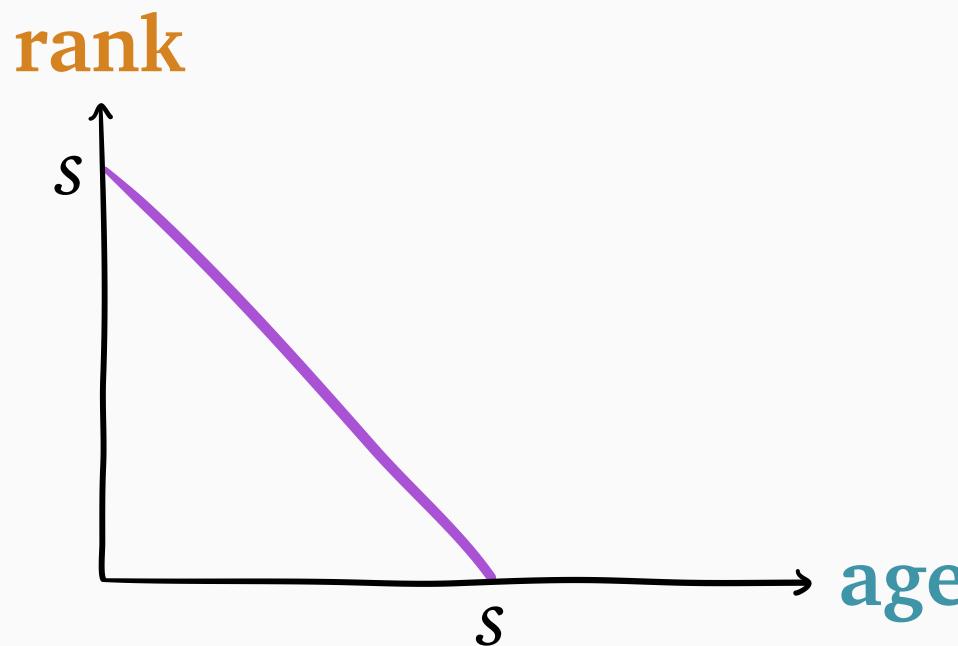
Medium noise



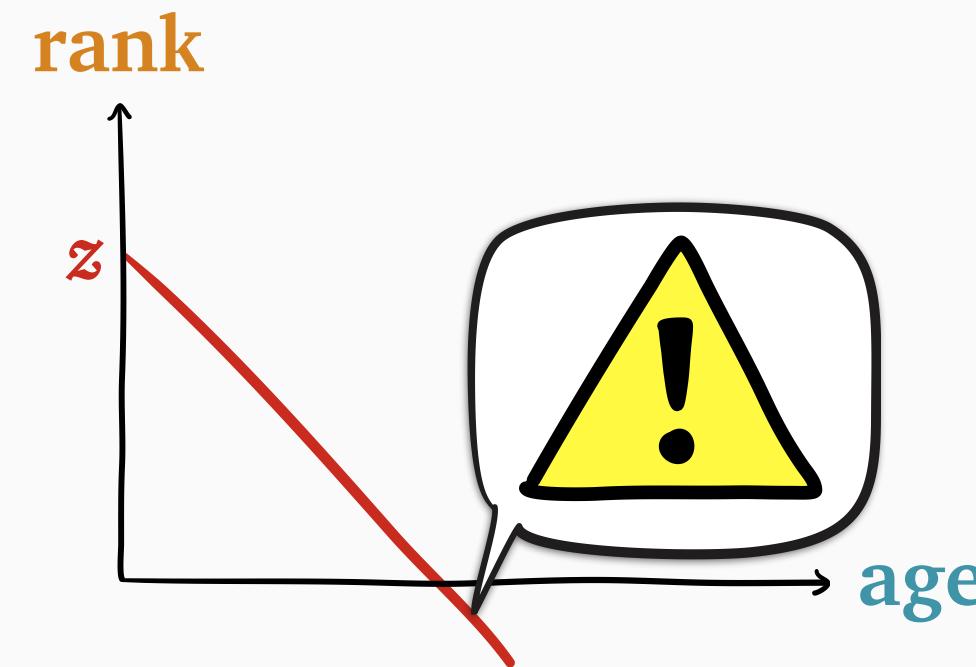
High noise



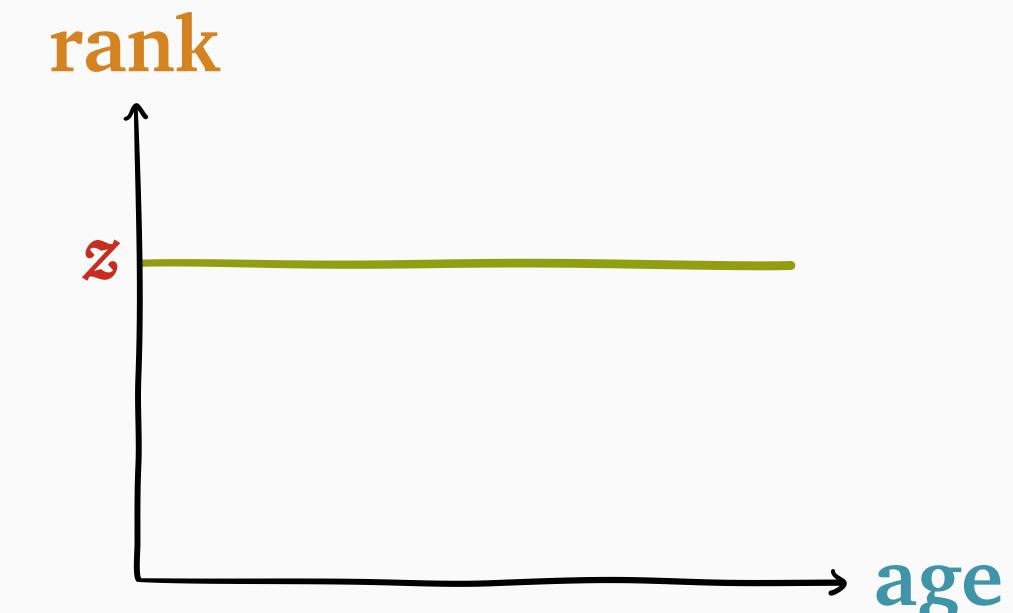
SRPT



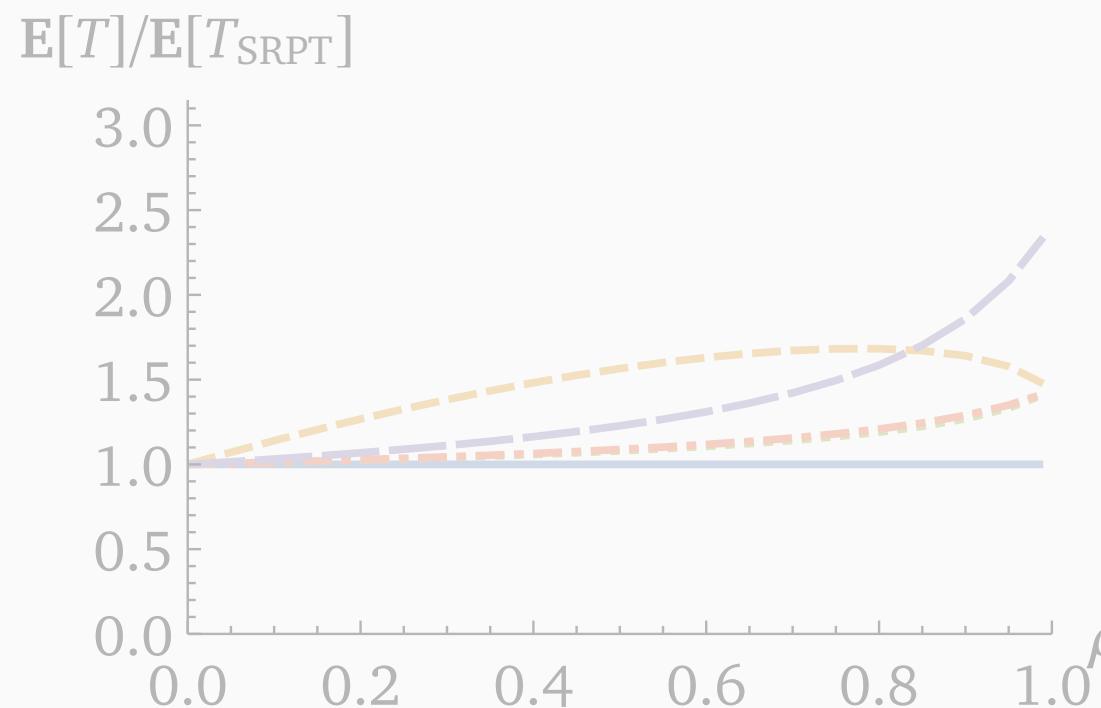
Noisy SRPT



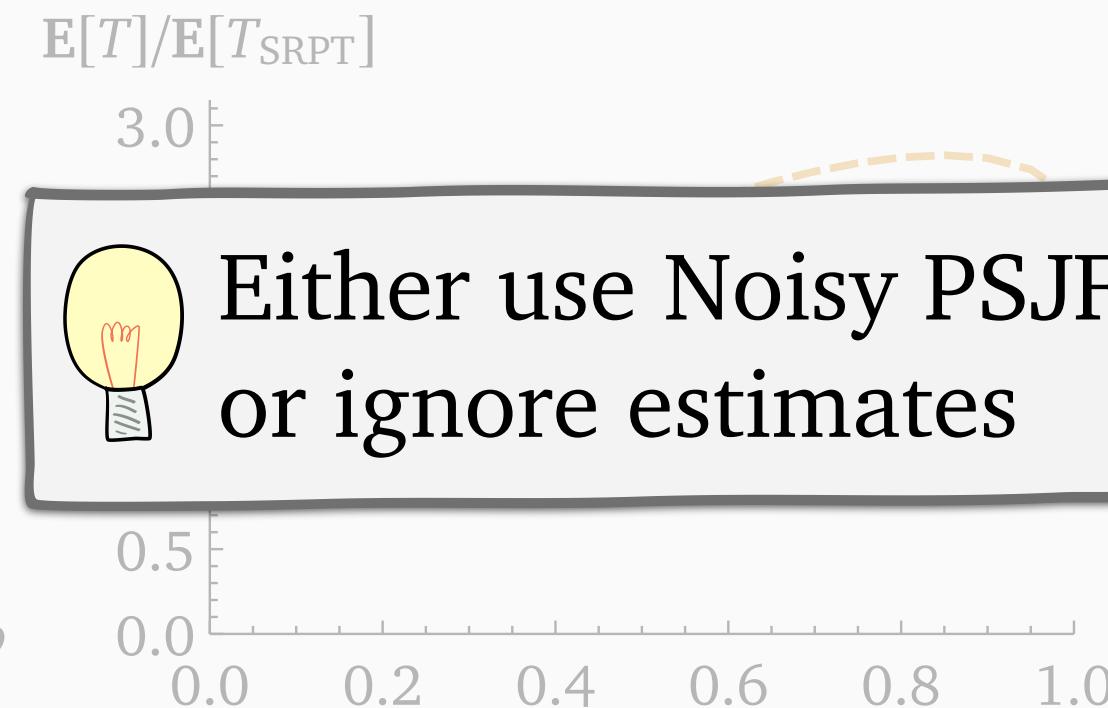
Noisy PSJF



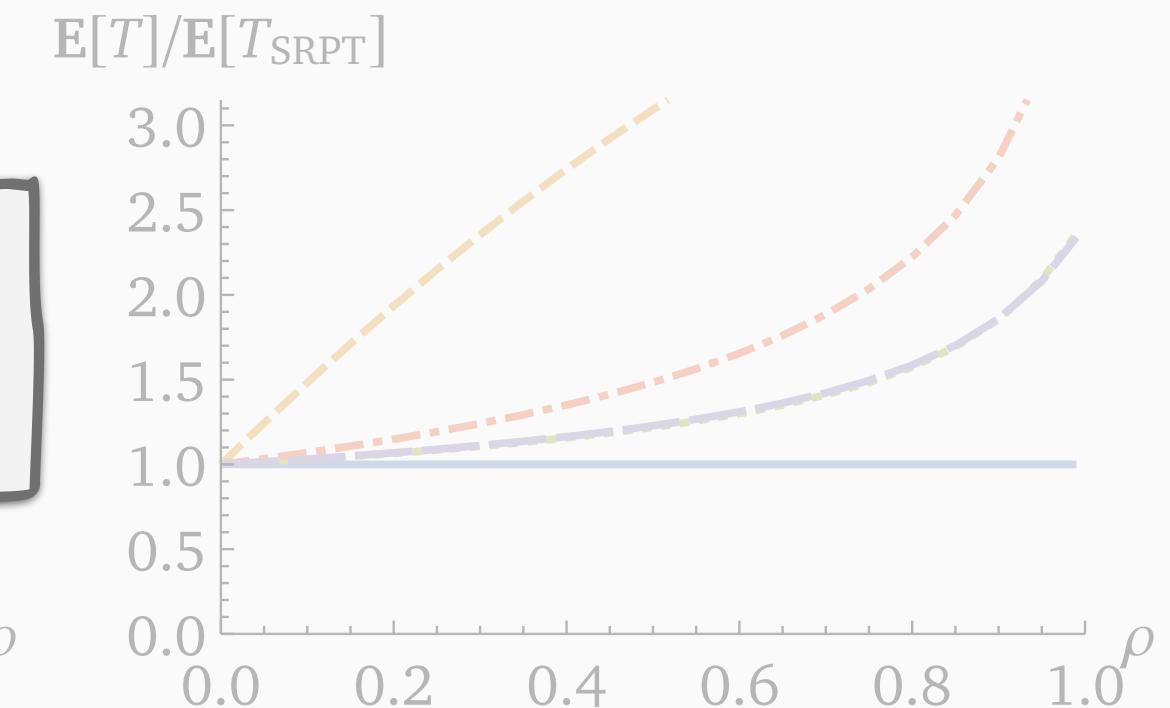
Low noise



Medium noise

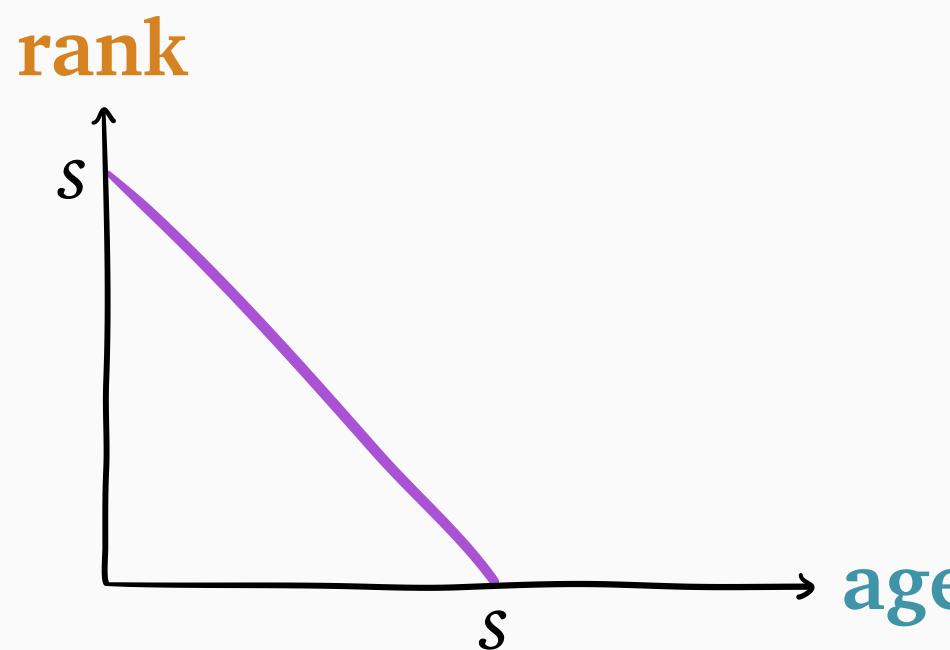


High noise

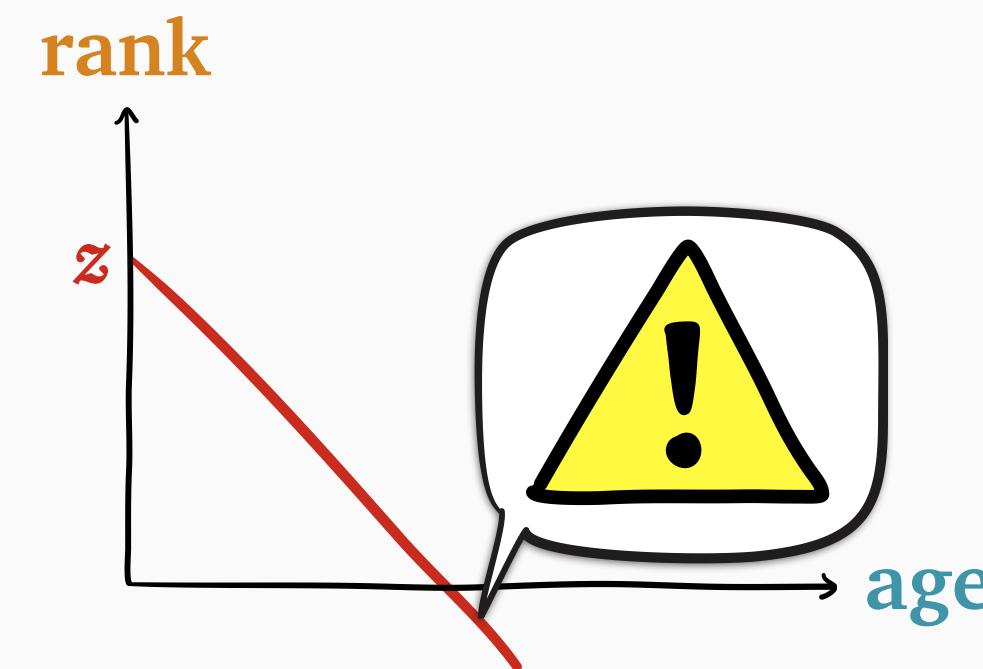


— SRPT
- - - Opt. w/ estimates
- - - Opt. w/out estimates
- - - Noisy SRPT
- - - Noisy PSJF

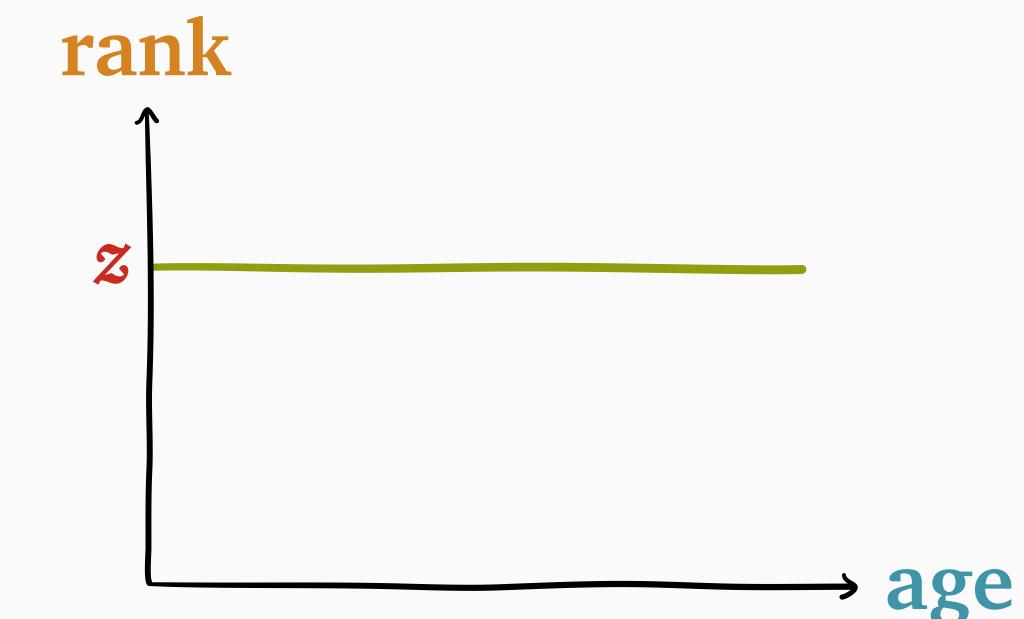
SRPT



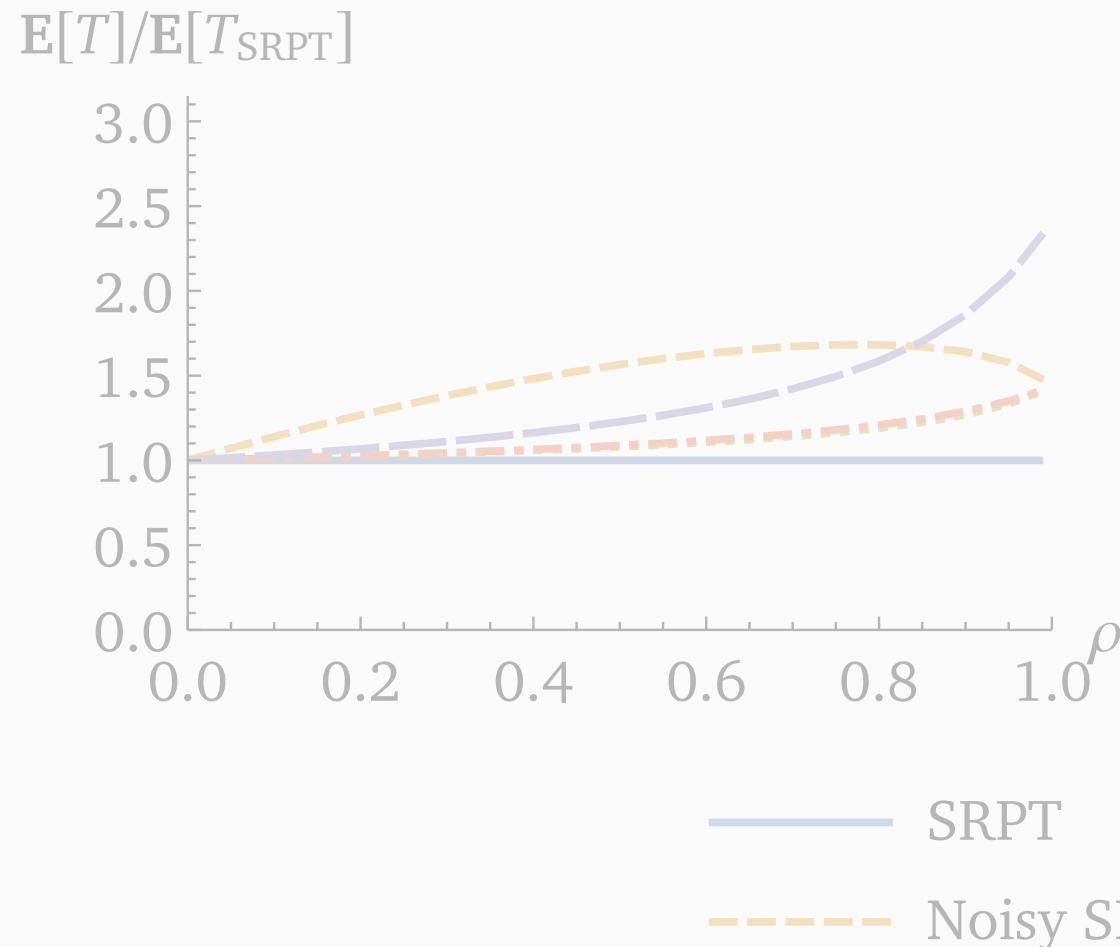
Noisy SRPT



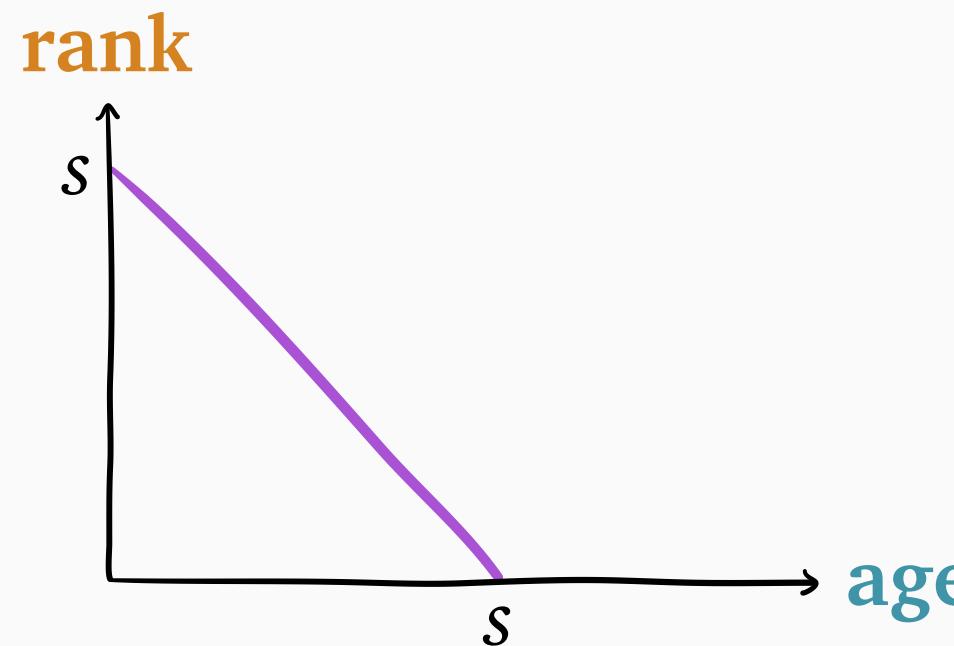
Noisy PSJF



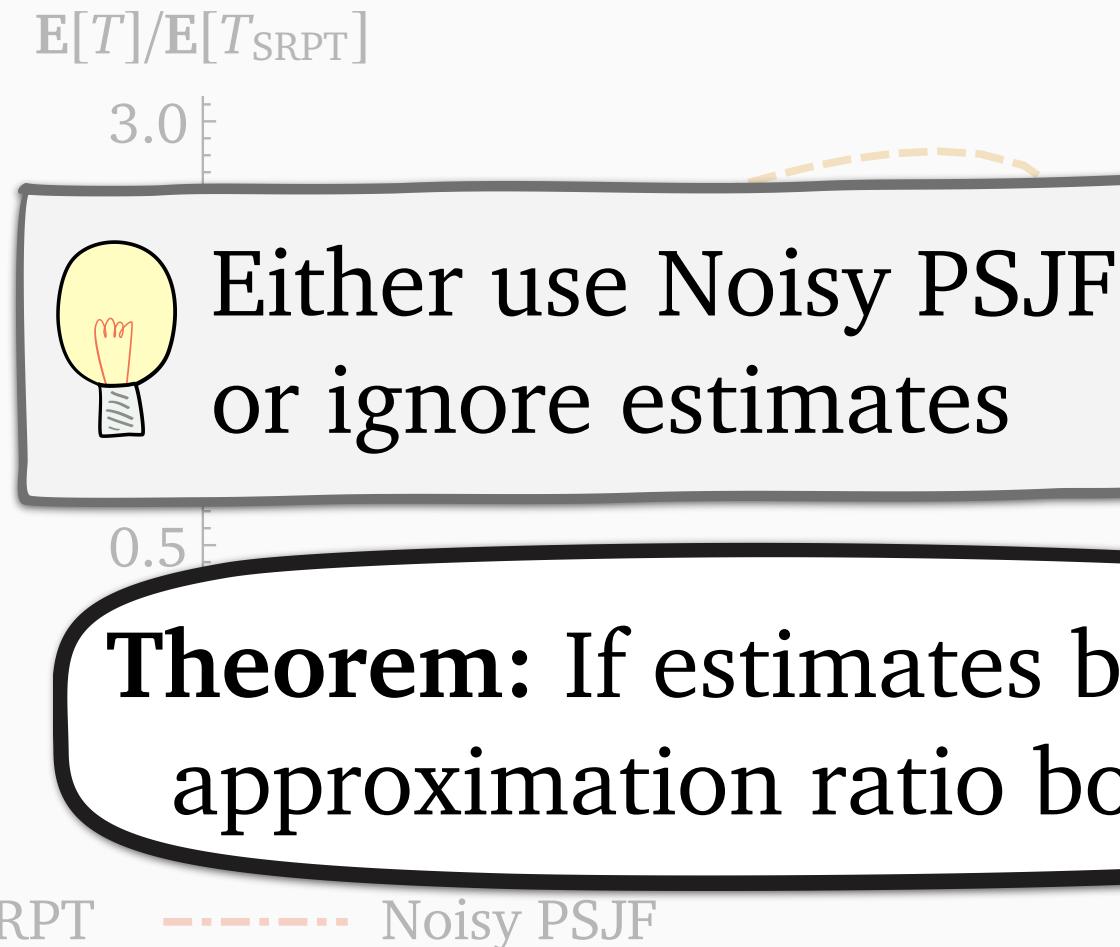
Low noise



SRPT



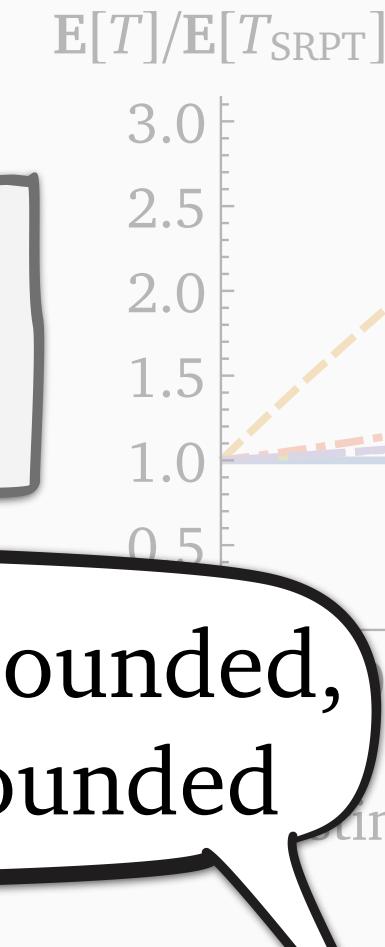
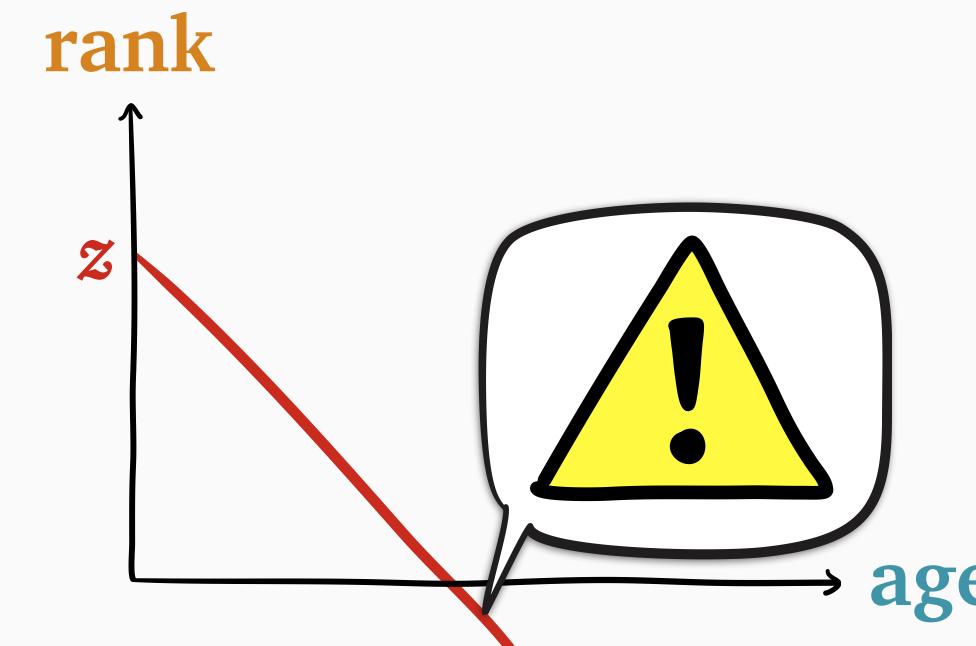
Medium noise



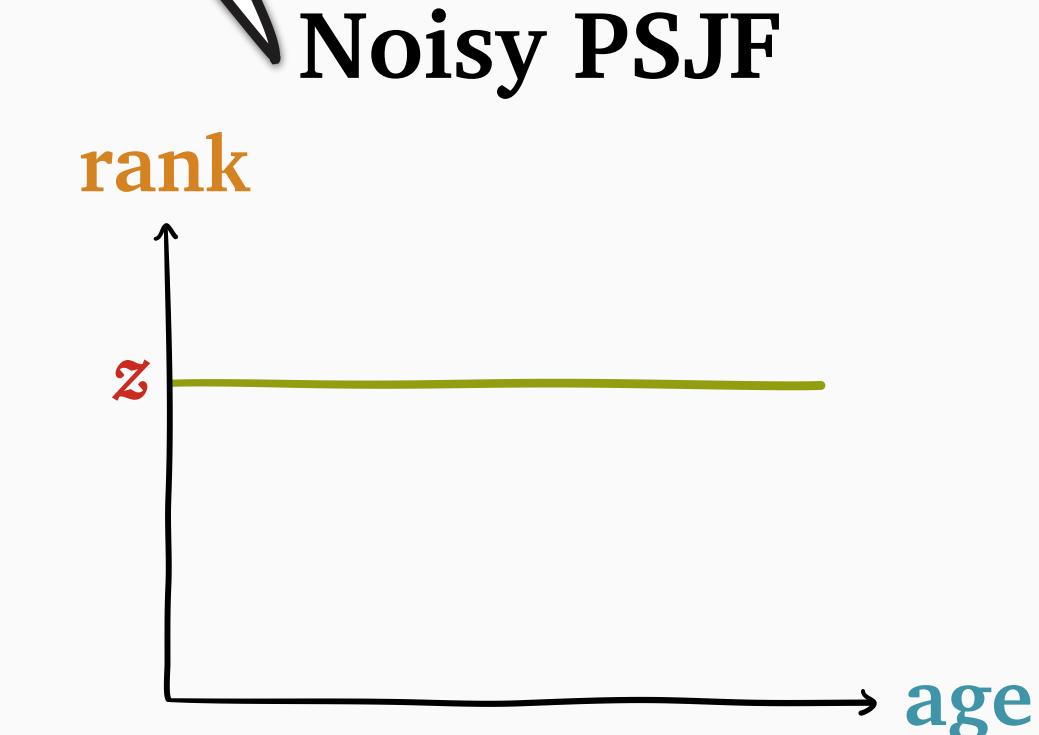
Either use Noisy PSJF
or ignore estimates

Theorem: If estimates bounded,
approximation ratio bounded

Noisy SRPT

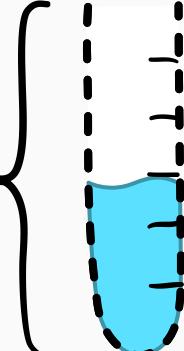


High noise

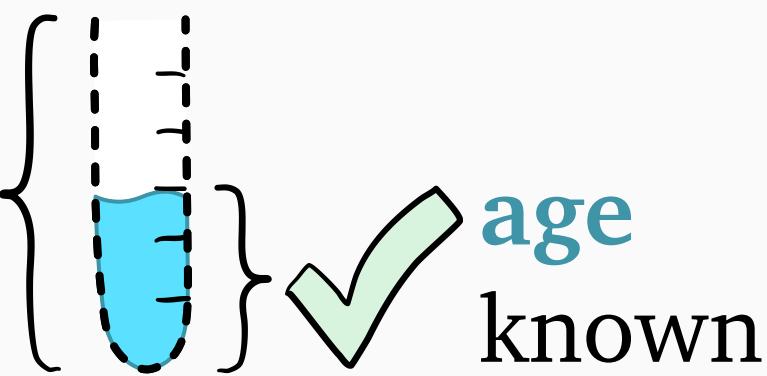


Noisy PSJF

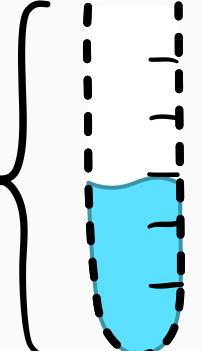
Optimal policy with no estimates

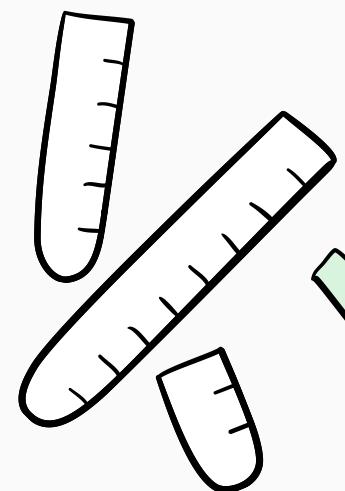
size
unknown ? {  }

Optimal policy with no estimates

size
unknown ?  age
known

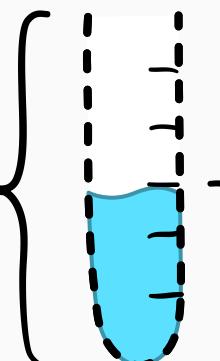
Optimal policy with no estimates

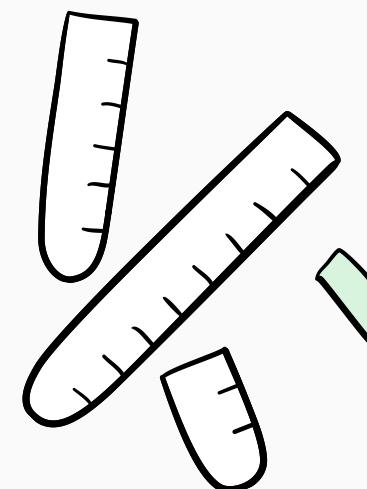
size
unknown ? {  }  age known



 distribution S known

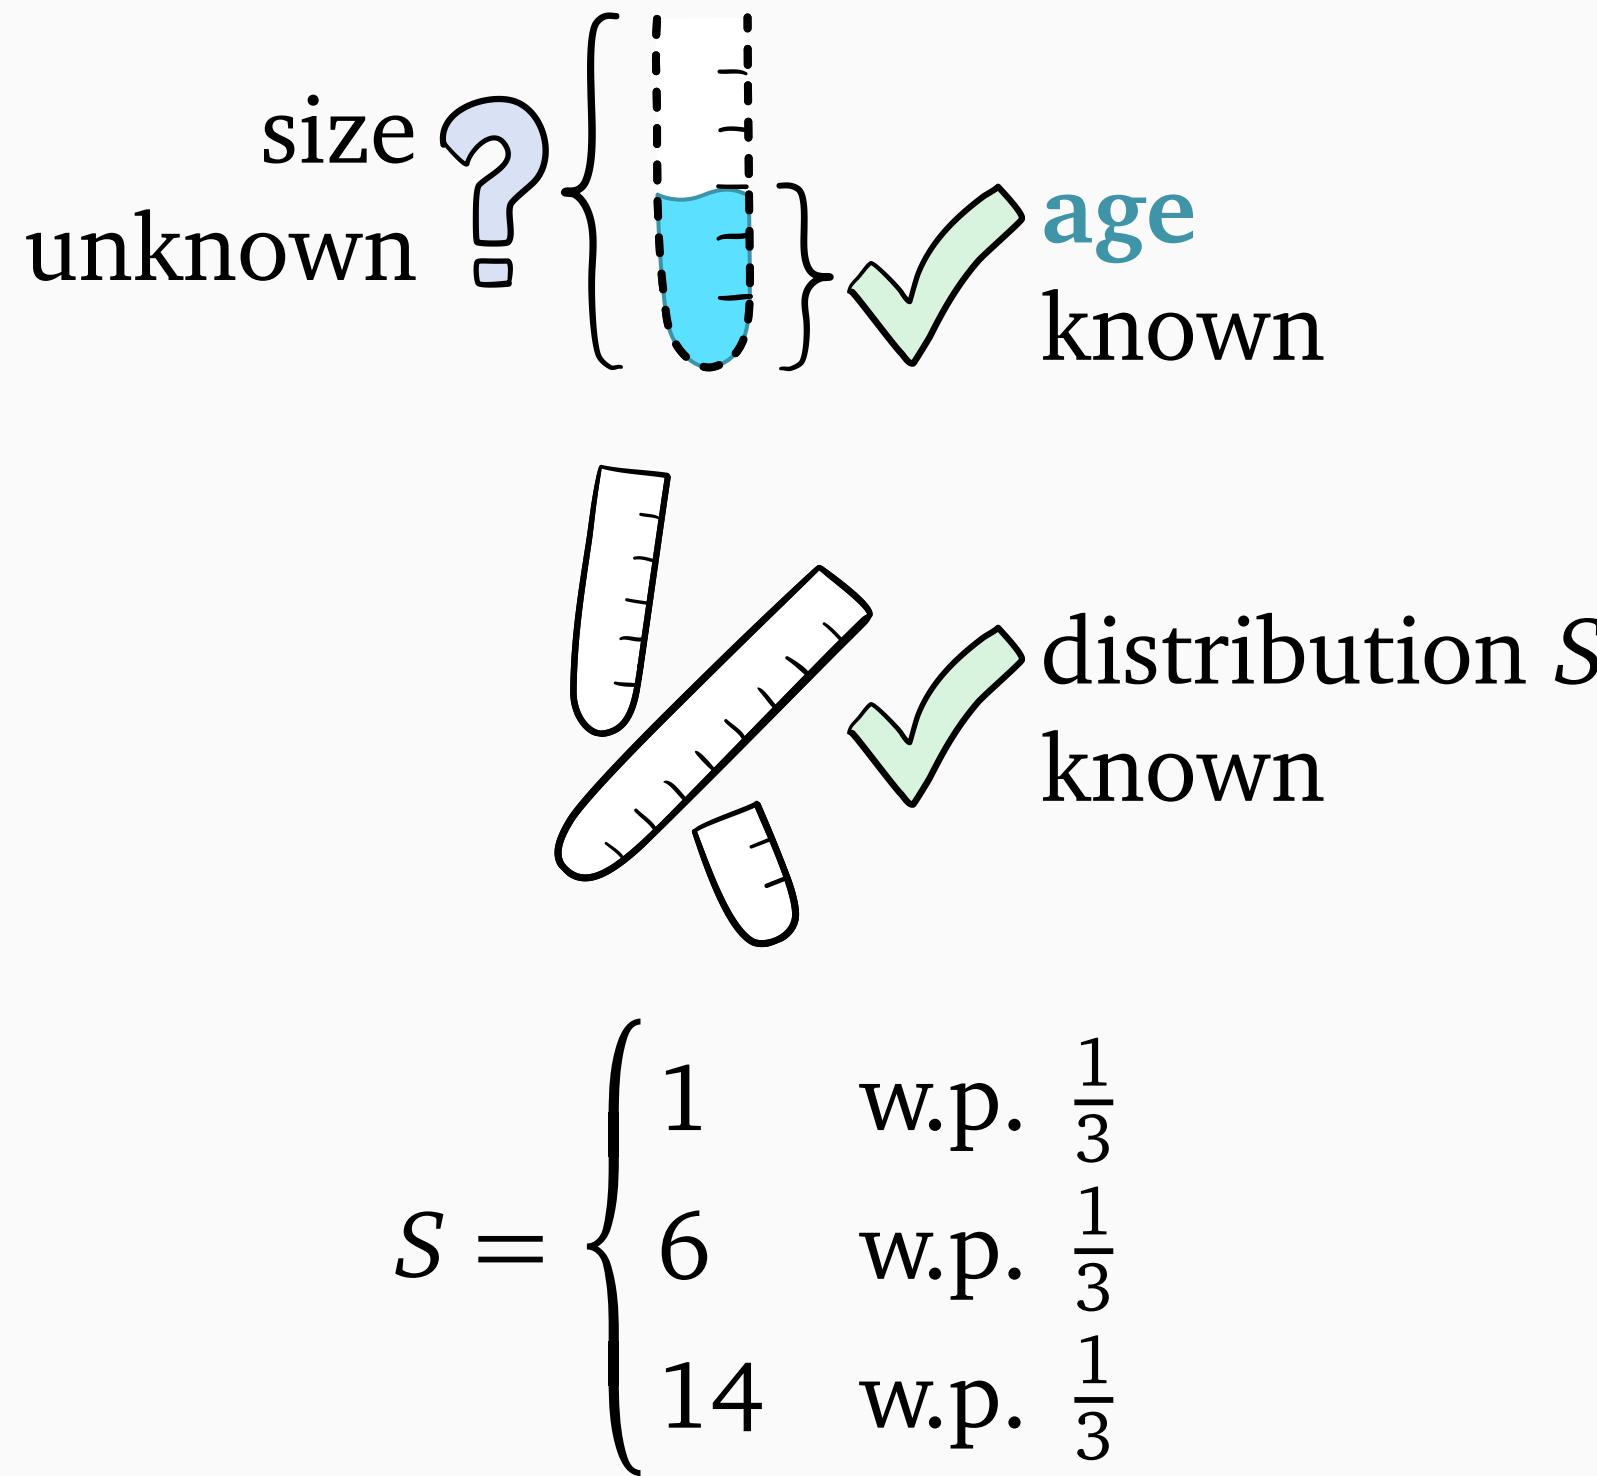
Optimal policy with no estimates

size
unknown ? {  } ✓ age known

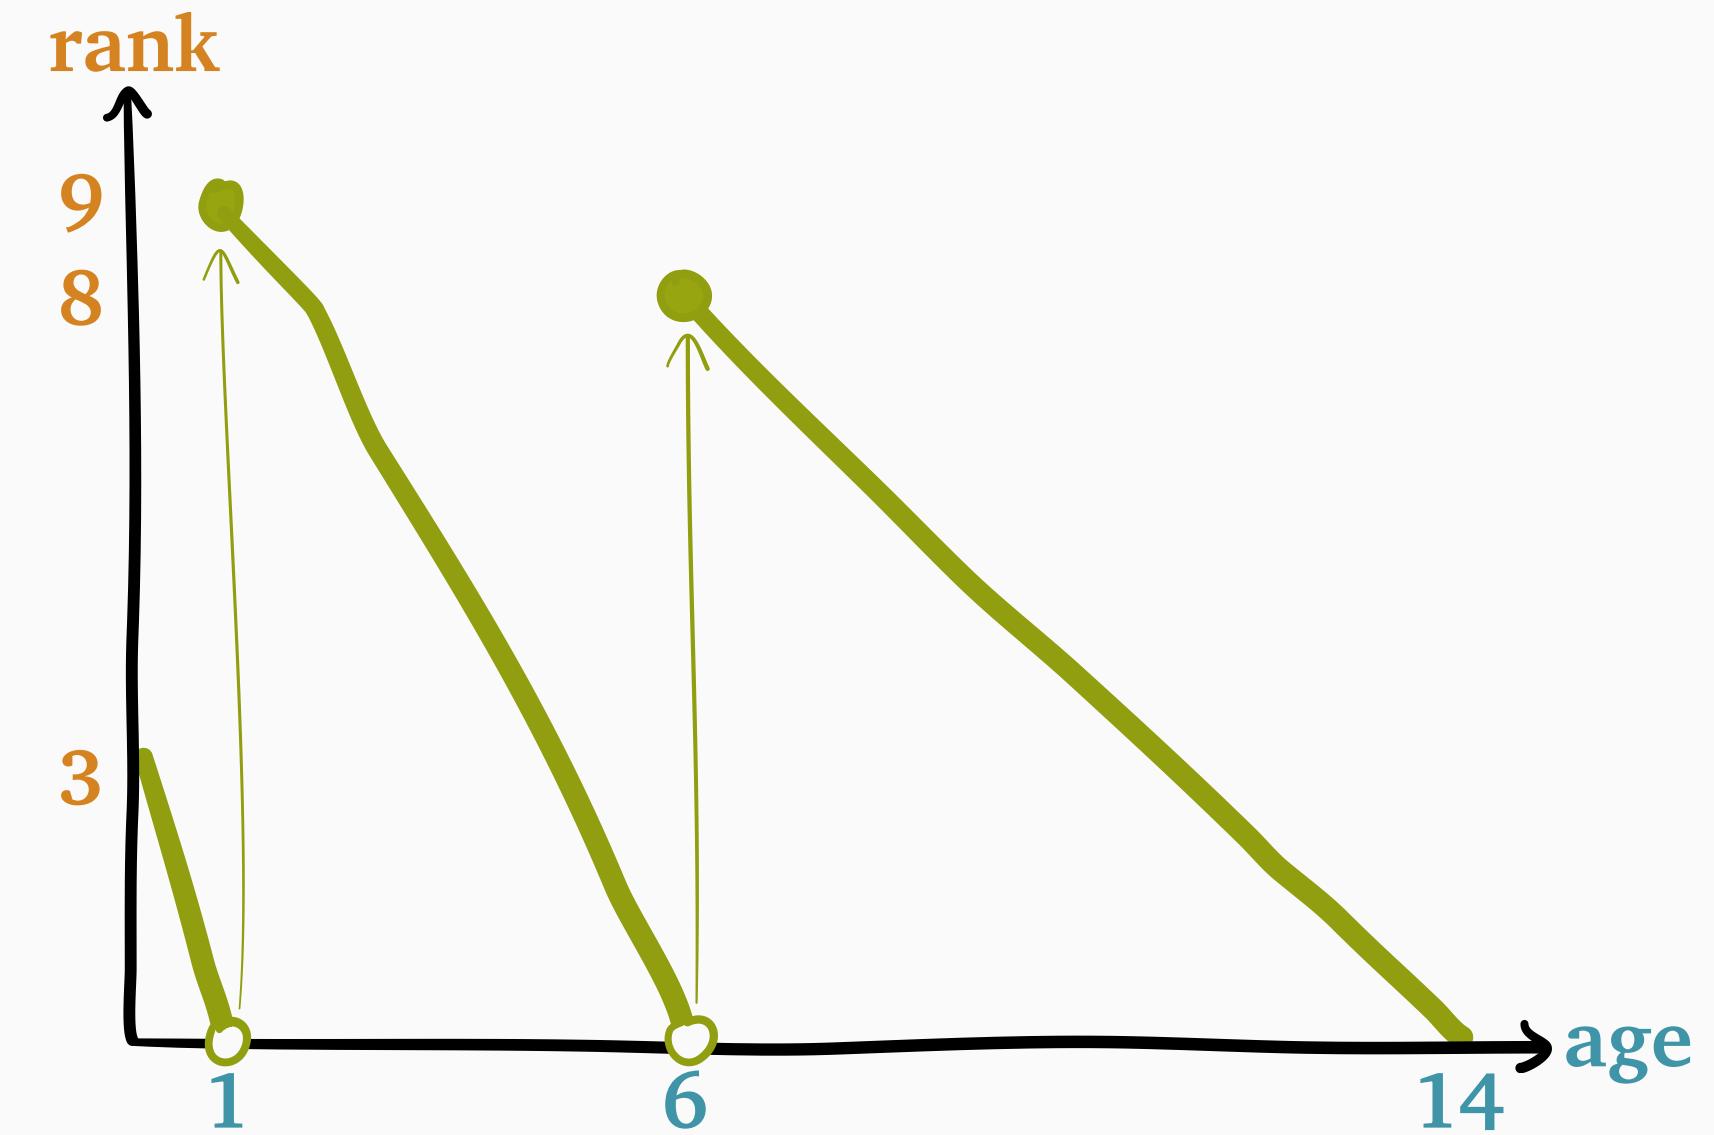
 ✓ distribution S known

$$S = \begin{cases} 1 & \text{w.p. } \frac{1}{3} \\ 6 & \text{w.p. } \frac{1}{3} \\ 14 & \text{w.p. } \frac{1}{3} \end{cases}$$

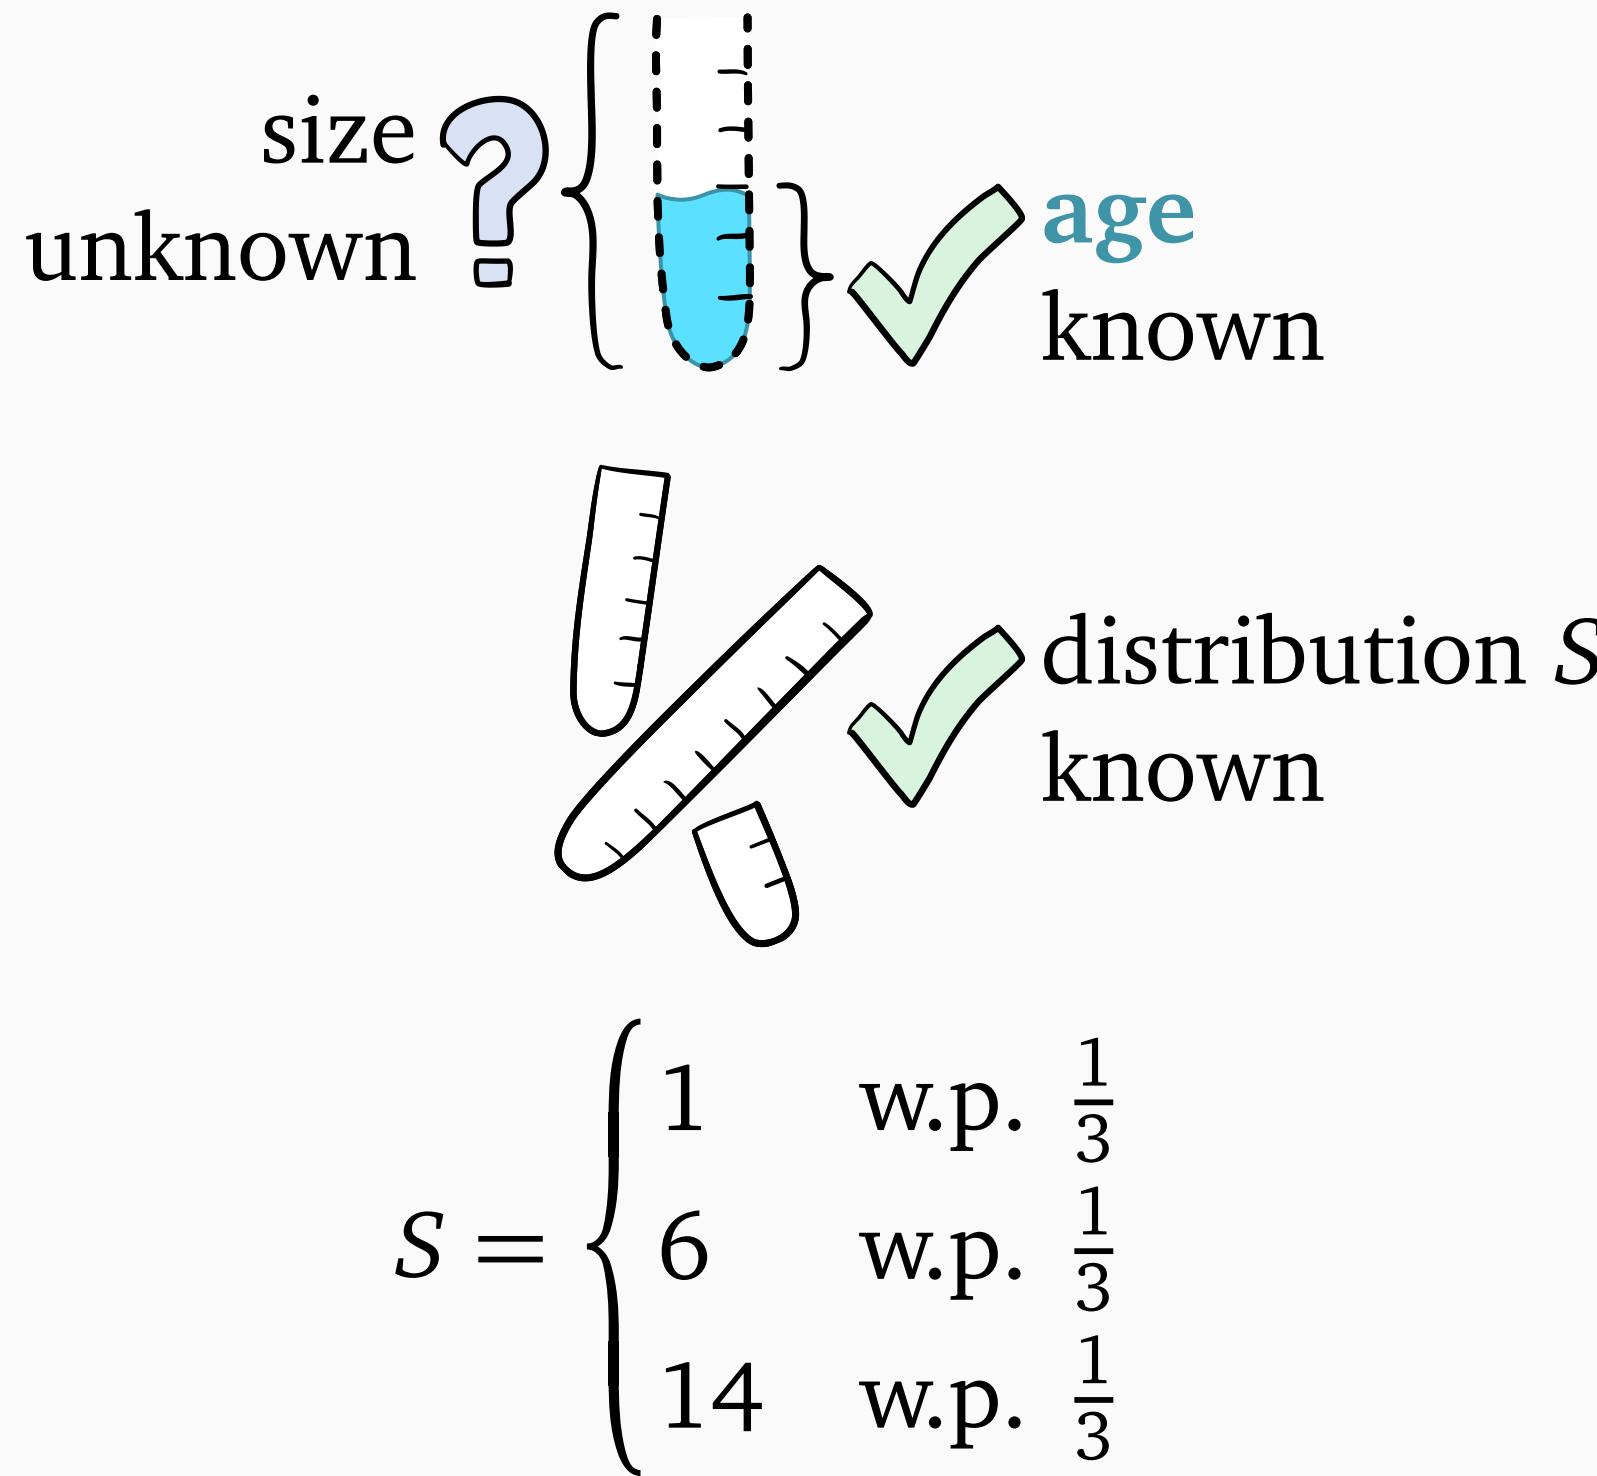
Optimal policy with no estimates



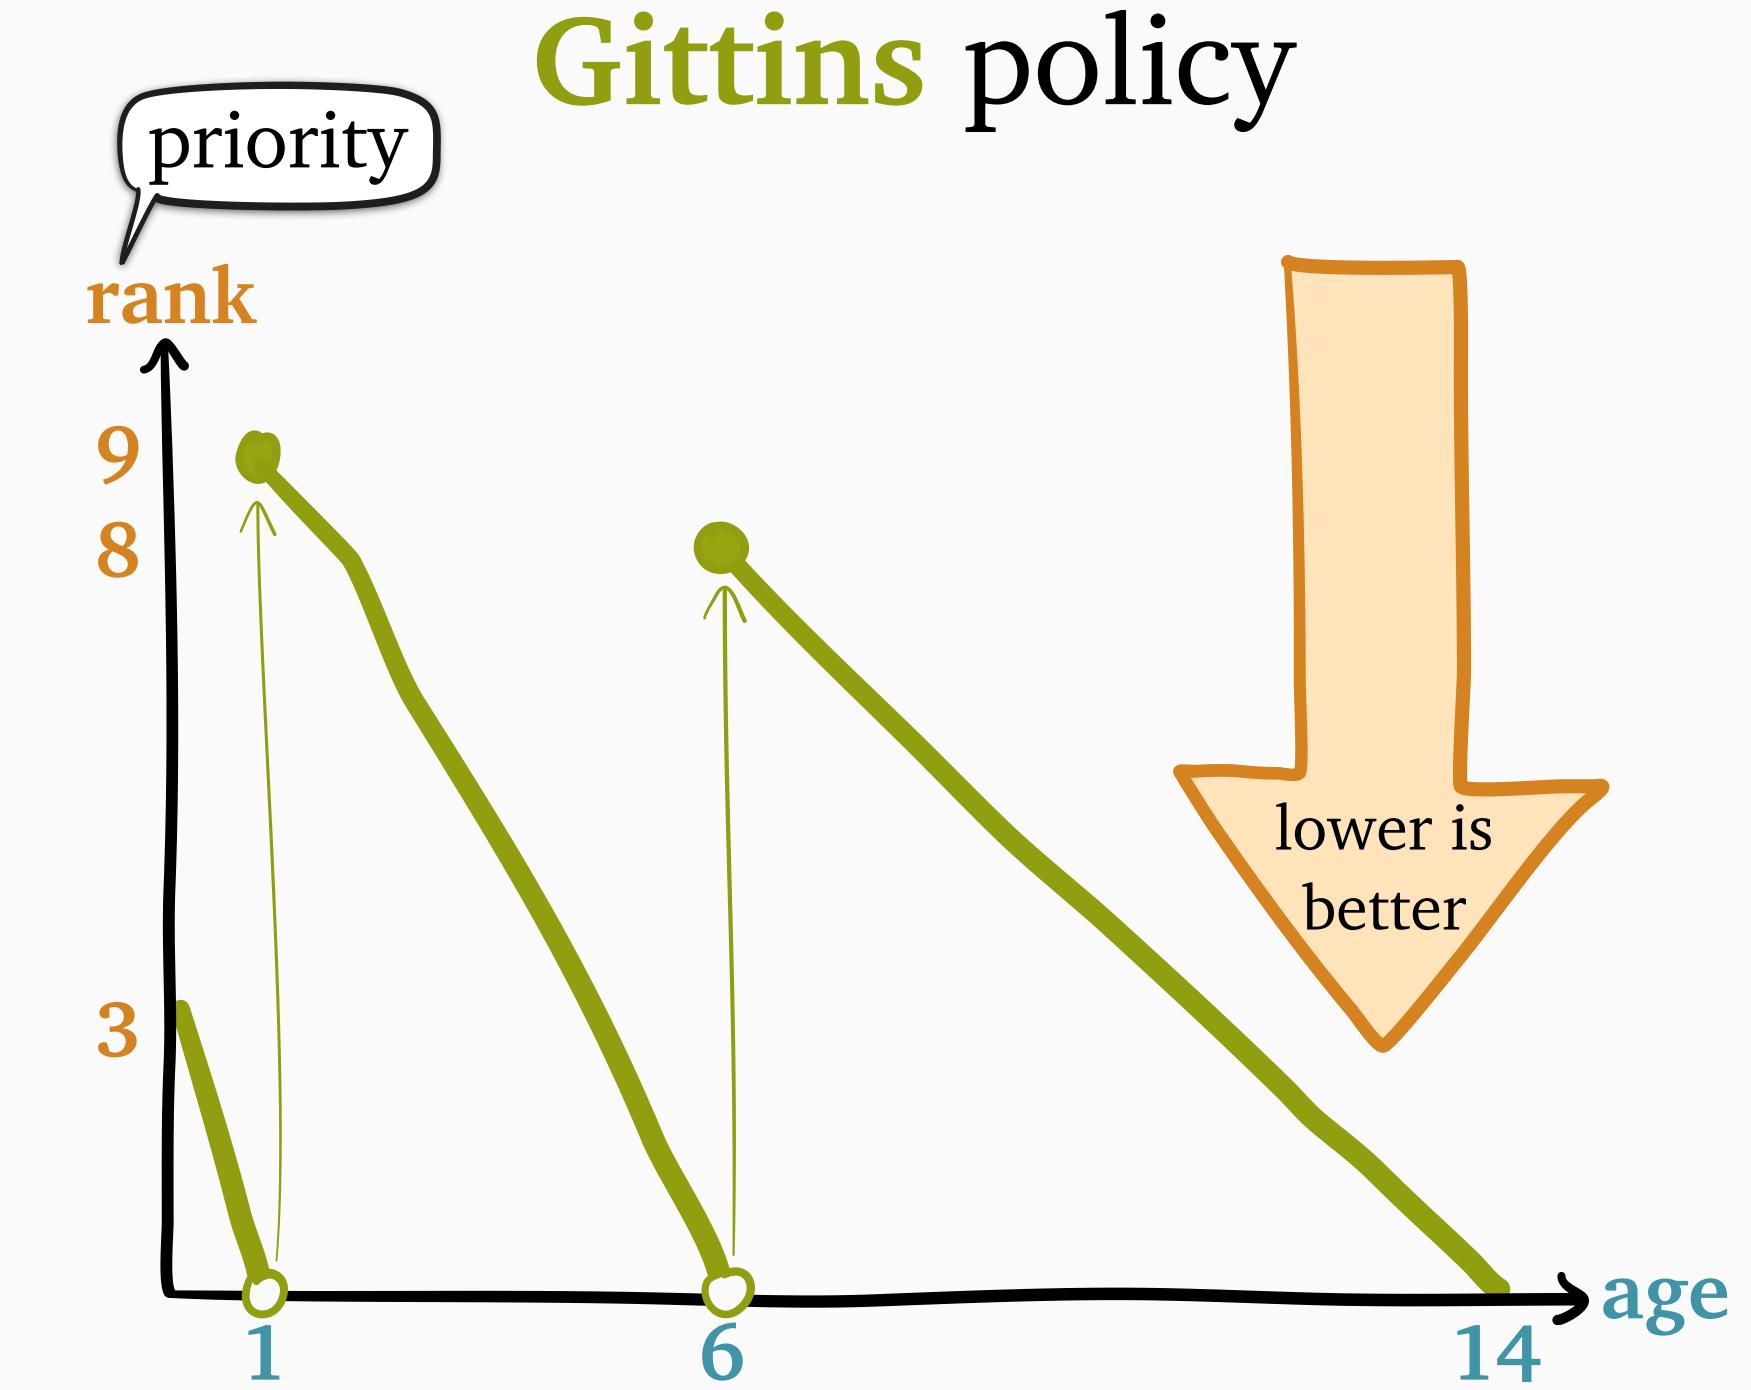
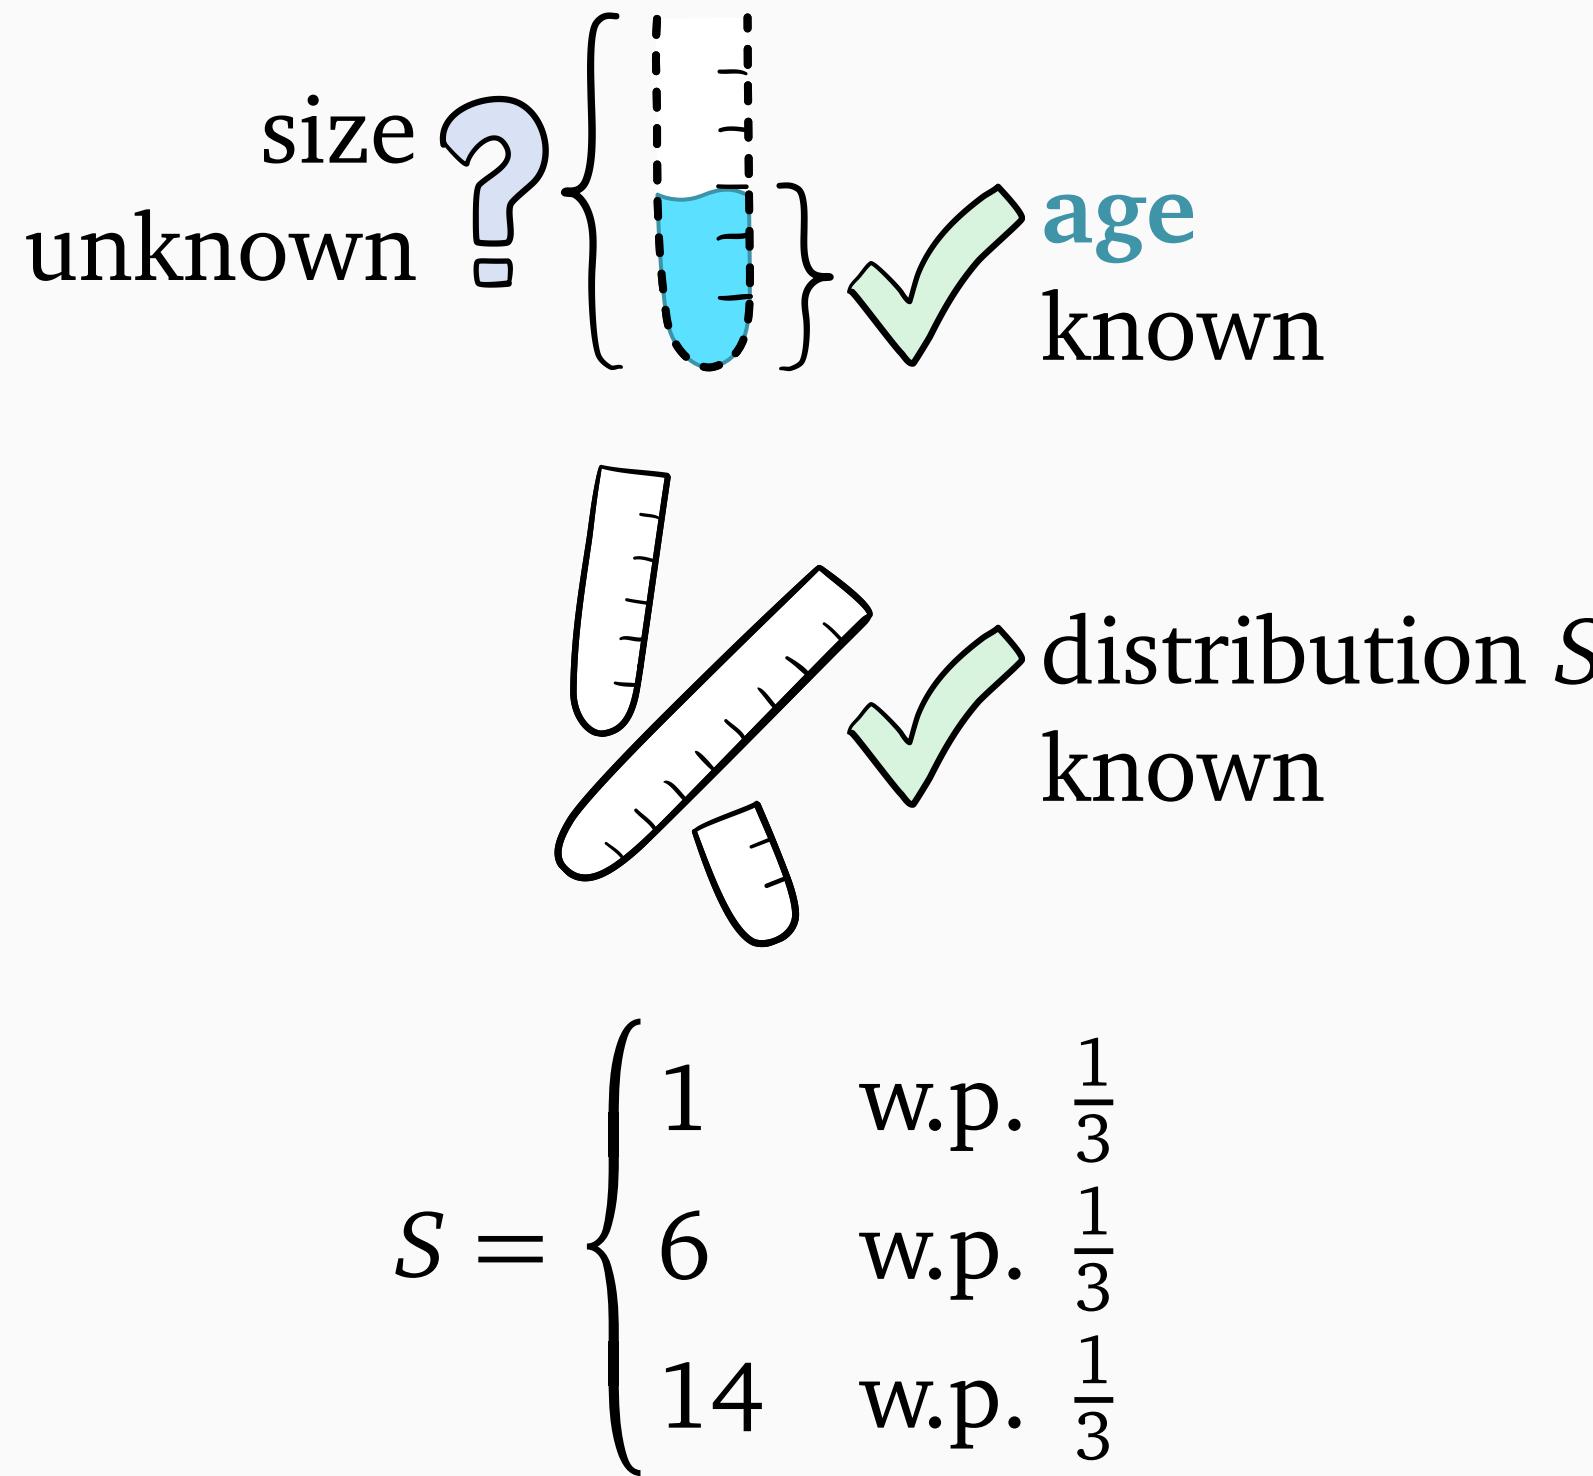
Gittins policy



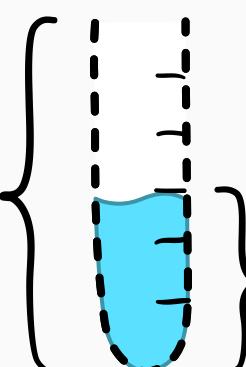
Optimal policy with no estimates

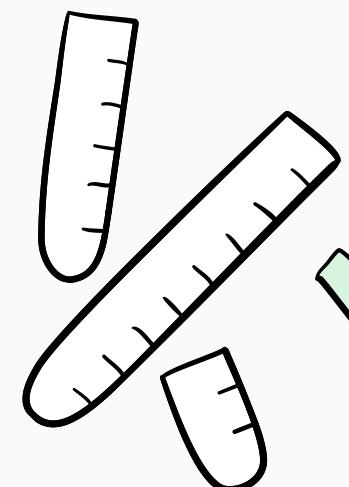


Optimal policy with no estimates



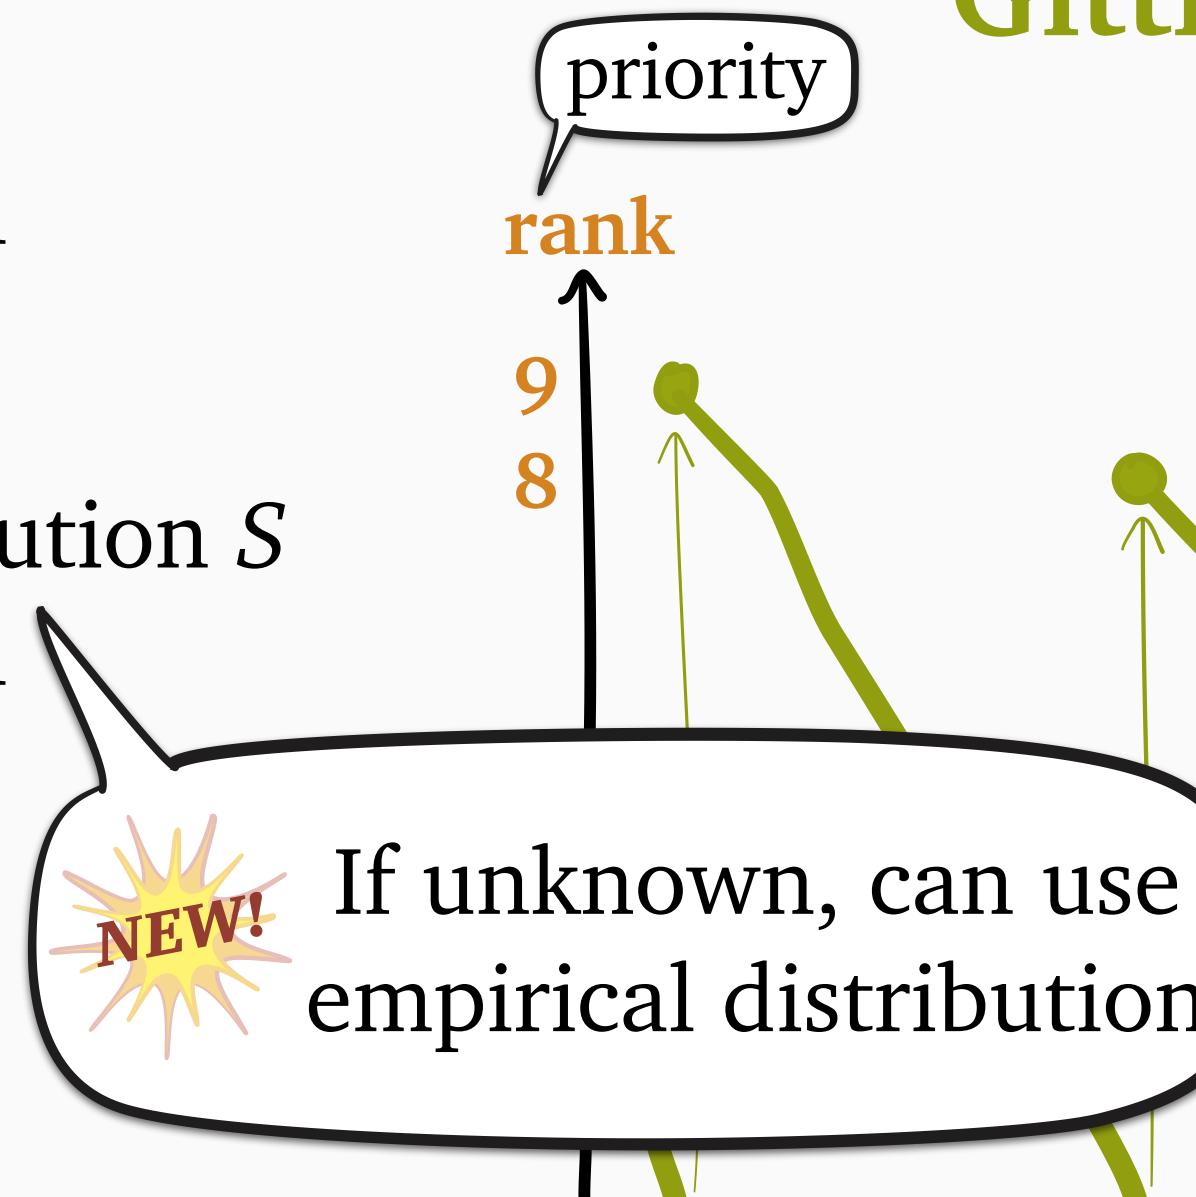
Optimal policy with no estimates

size unknown ? {  } ✓ age known

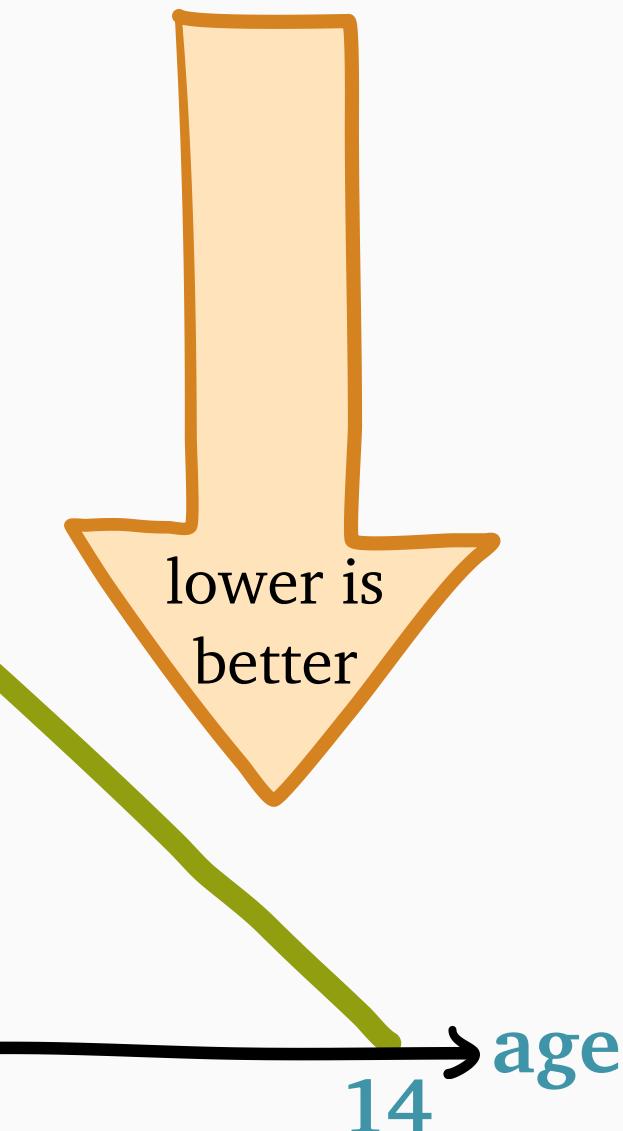


distribution S known

$$S = \begin{cases} 1 & \text{w.p. } \frac{1}{3} \\ 6 & \text{w.p. } \frac{1}{3} \\ 14 & \text{w.p. } \frac{1}{3} \end{cases}$$

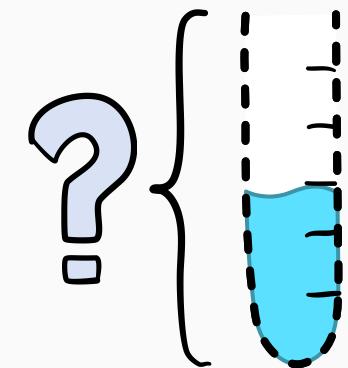


Gittins policy



Gittins construction generalizes

state = age

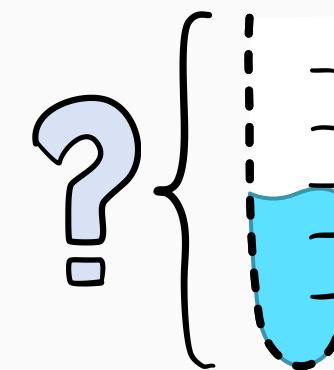


minimize

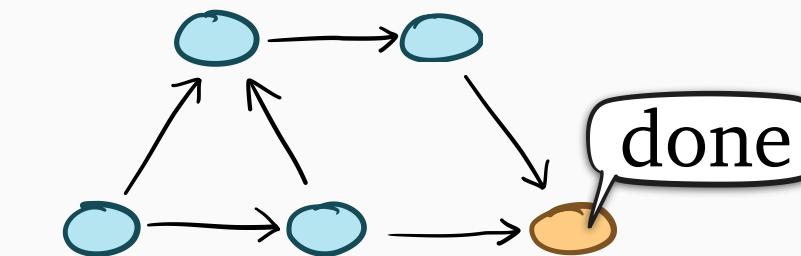
$$\mathbf{E}[N] = \lambda \mathbf{E}[T]$$

Gittins construction generalizes

state = age



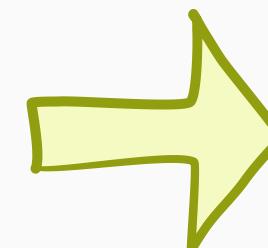
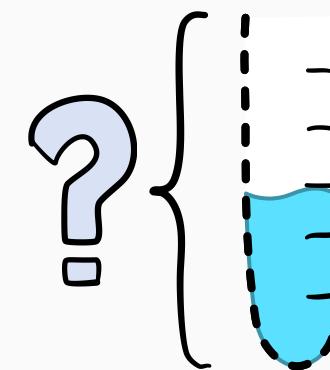
model service as
Markov process



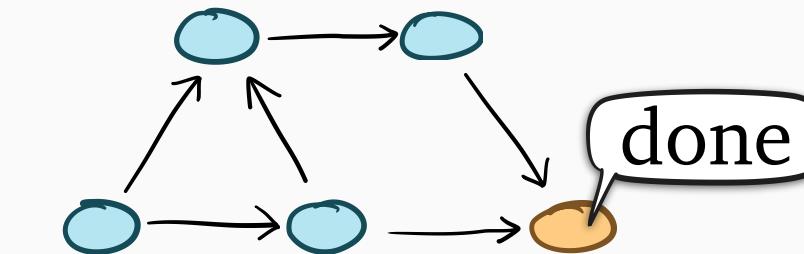
minimize
 $E[N] = \lambda E[T]$

Gittins construction generalizes

state = age



model service as
Markov process



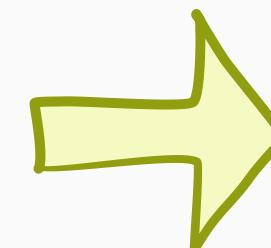
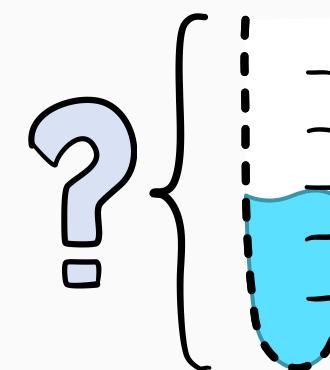
minimize
 $E[N] = \lambda E[T]$



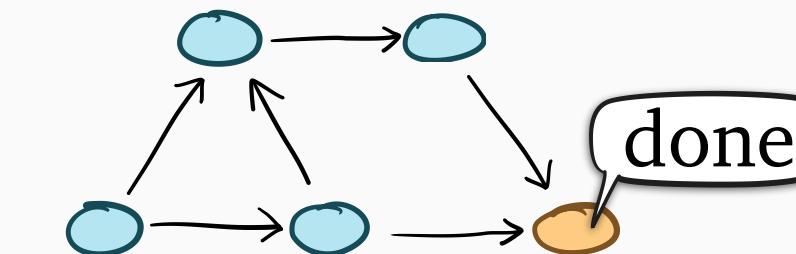
minimize mean sum of
state-based holding costs

Gittins construction generalizes

state = age



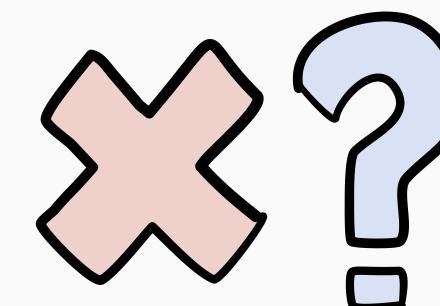
model service as
Markov process



minimize
 $E[N] = \lambda E[T]$

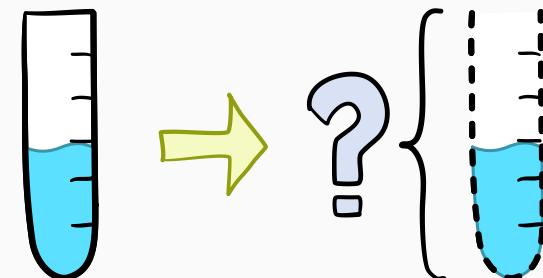


minimize mean sum of
state-based holding costs

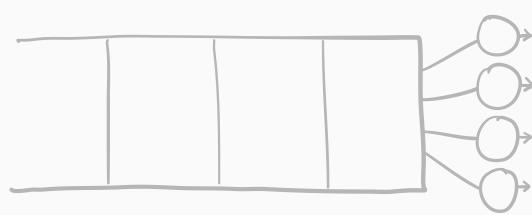


Open: *time-based* holding costs

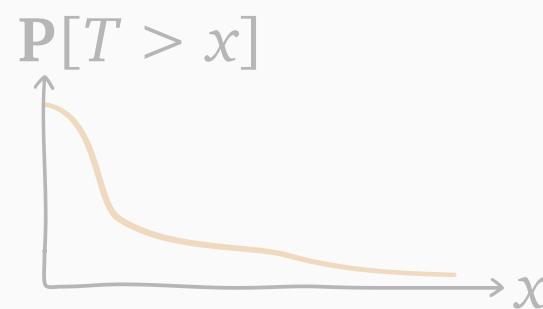
Hard scheduling questions



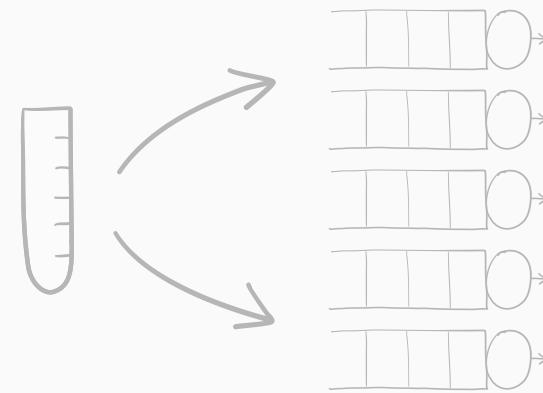
*What if we have
unknown job sizes?*



*What if there are
multiple servers?*



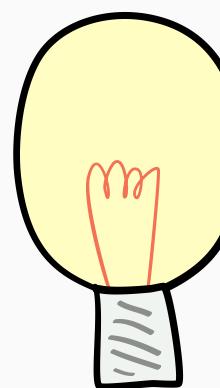
*What if we want to optimize
tails instead of means?*



*What if we can only use
dispatching without
fancy scheduling?*

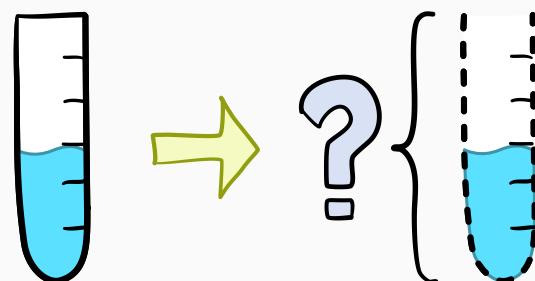


***What are the underlying
theoretical tools?***

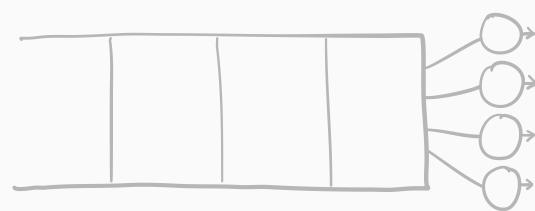


***What are the resulting
practical lessons?***

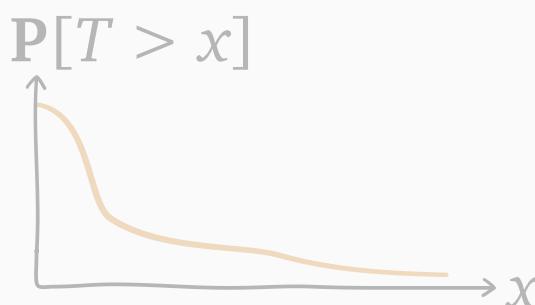
Hard scheduling questions



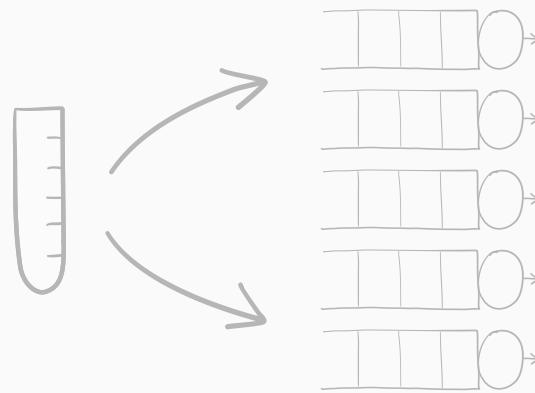
*What if we have
unknown job sizes?*



*What if there are
multiple servers?*



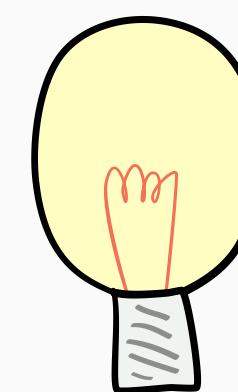
*What if we want to optimize
tails instead of means?*



*What if we can only use
dispatching without
fancy scheduling?*

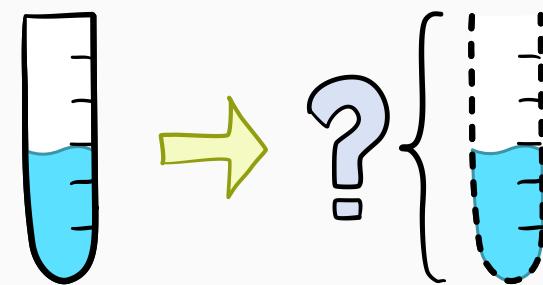


*What are the underlying
theoretical tools?*

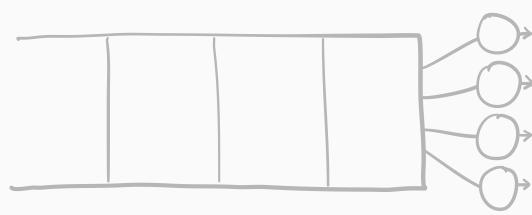


*What are the resulting
practical lessons?*

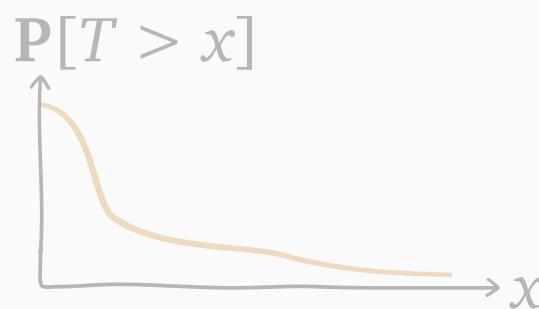
Hard scheduling questions



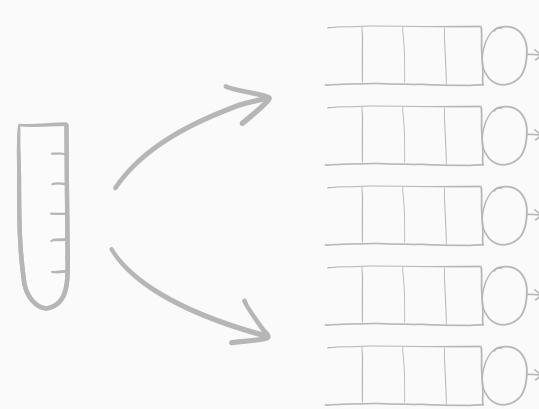
*What if we have
unknown job sizes?*



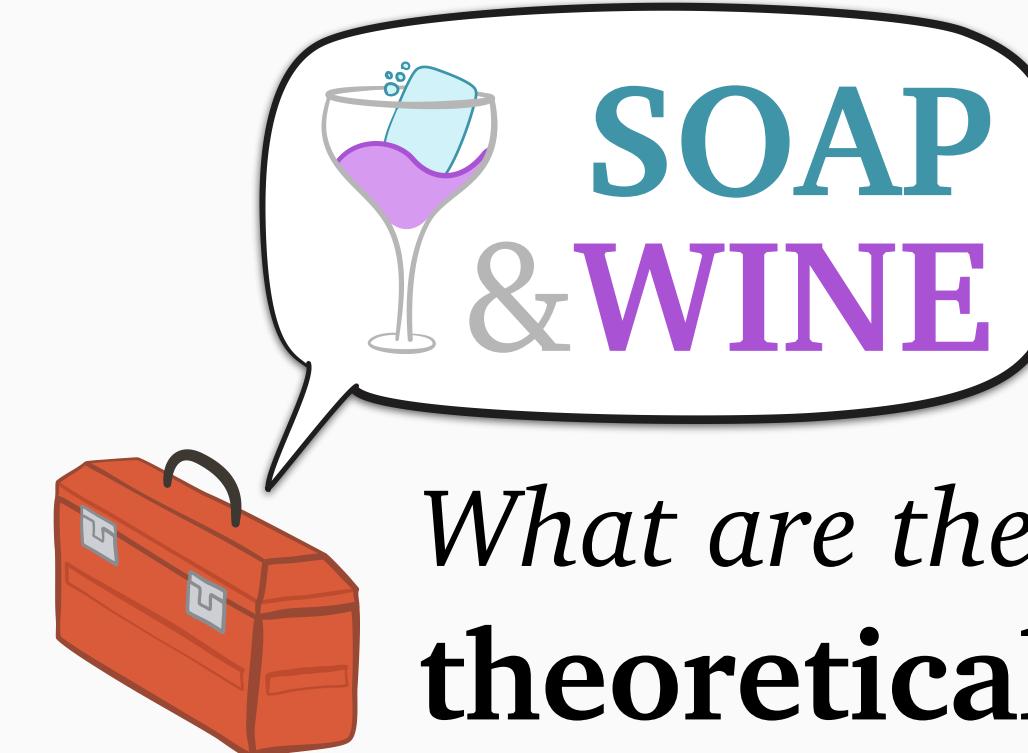
*What if there are
multiple servers?*



*What if we want to optimize
tails instead of means?*

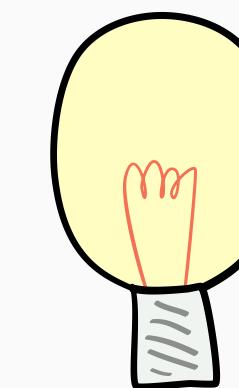


*What if we can only use
dispatching without
fancy scheduling?*



**SOAP
& WINE**

*What are the underlying
theoretical tools?*



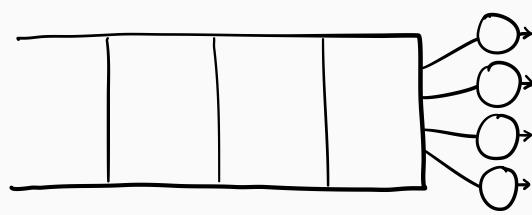
*What are the resulting
practical lessons?*

- PSJF is safer than SRPT
- Ignore terrible estimates
- **Gittins** works with data

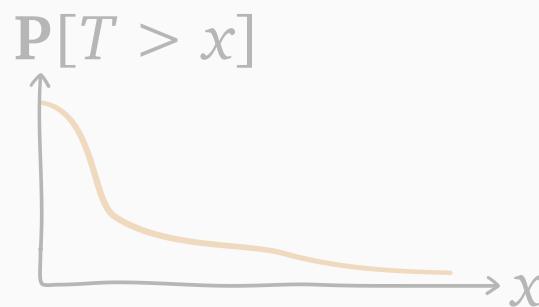
Hard scheduling questions



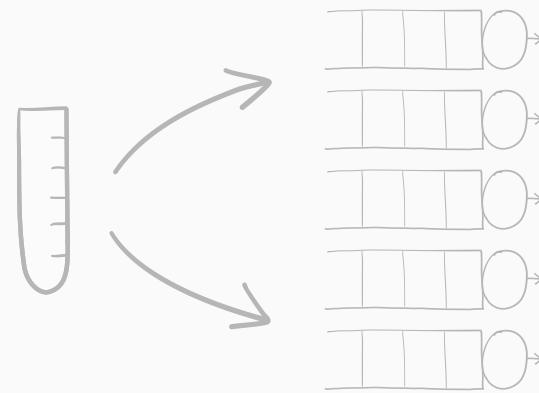
*What if we have
unknown job sizes?*



***What if there are
multiple servers?***



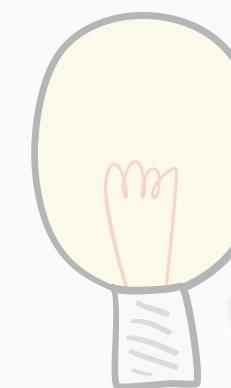
*What if we want to optimize
tails instead of means?*



*What if we can only use
dispatching without
fancy scheduling?*



*What are the underlying
theoretical tools?*



*What are the resulting
practical lessons?*

- PSJF is safer than SRPT
- Ignore terrible estimates
- **Gittins** works with data

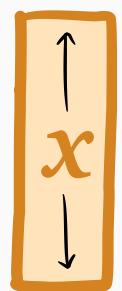
WINE and r -work

$W(r)$ = work relevant to rank r

WINE and r -work

$W(r)$ = work relevant to rank r

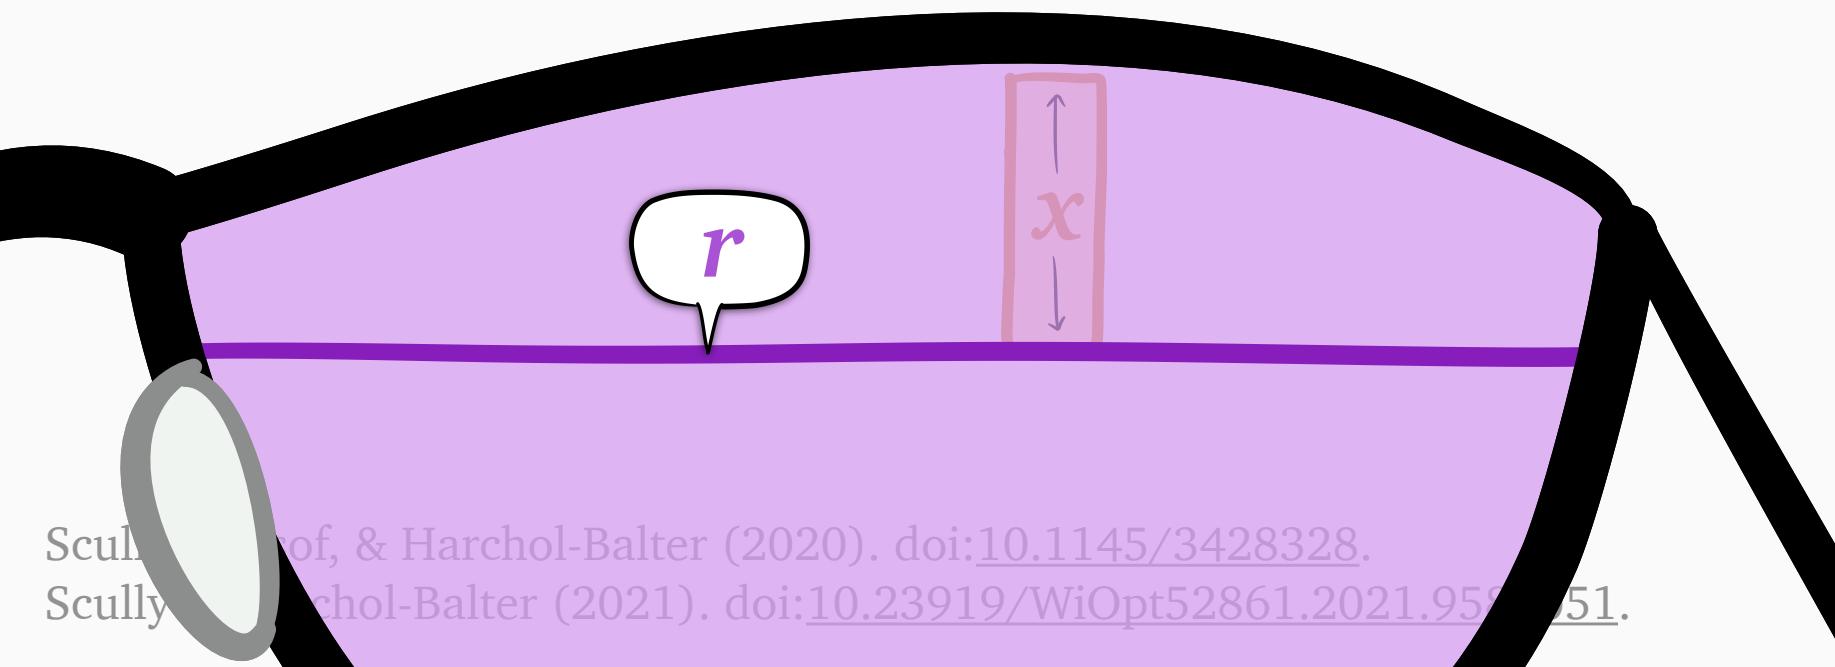
$w_x(r)$ = r -work of *single job* of rem. size x = {



WINE and r -work

$W(r)$ = work relevant to rank r

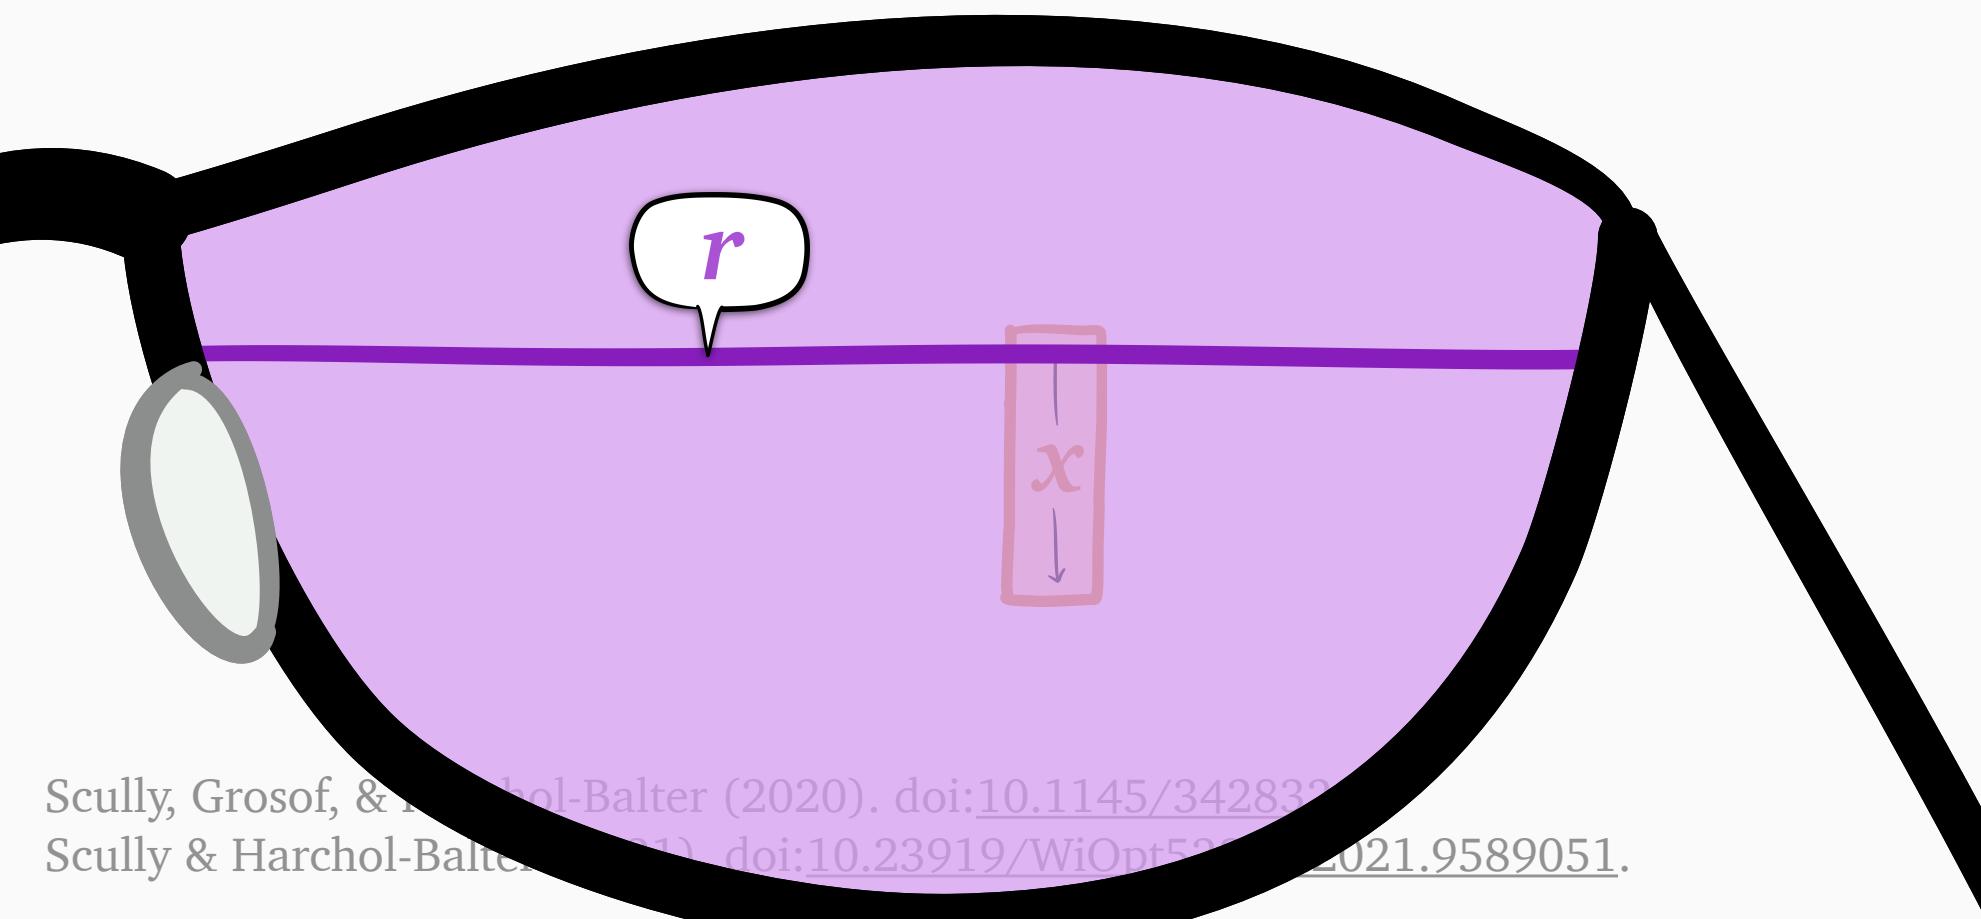
$w_x(r)$ = r -work of *single job* of rem. size x = {



WINE and r -work

$W(r)$ = work relevant to rank r

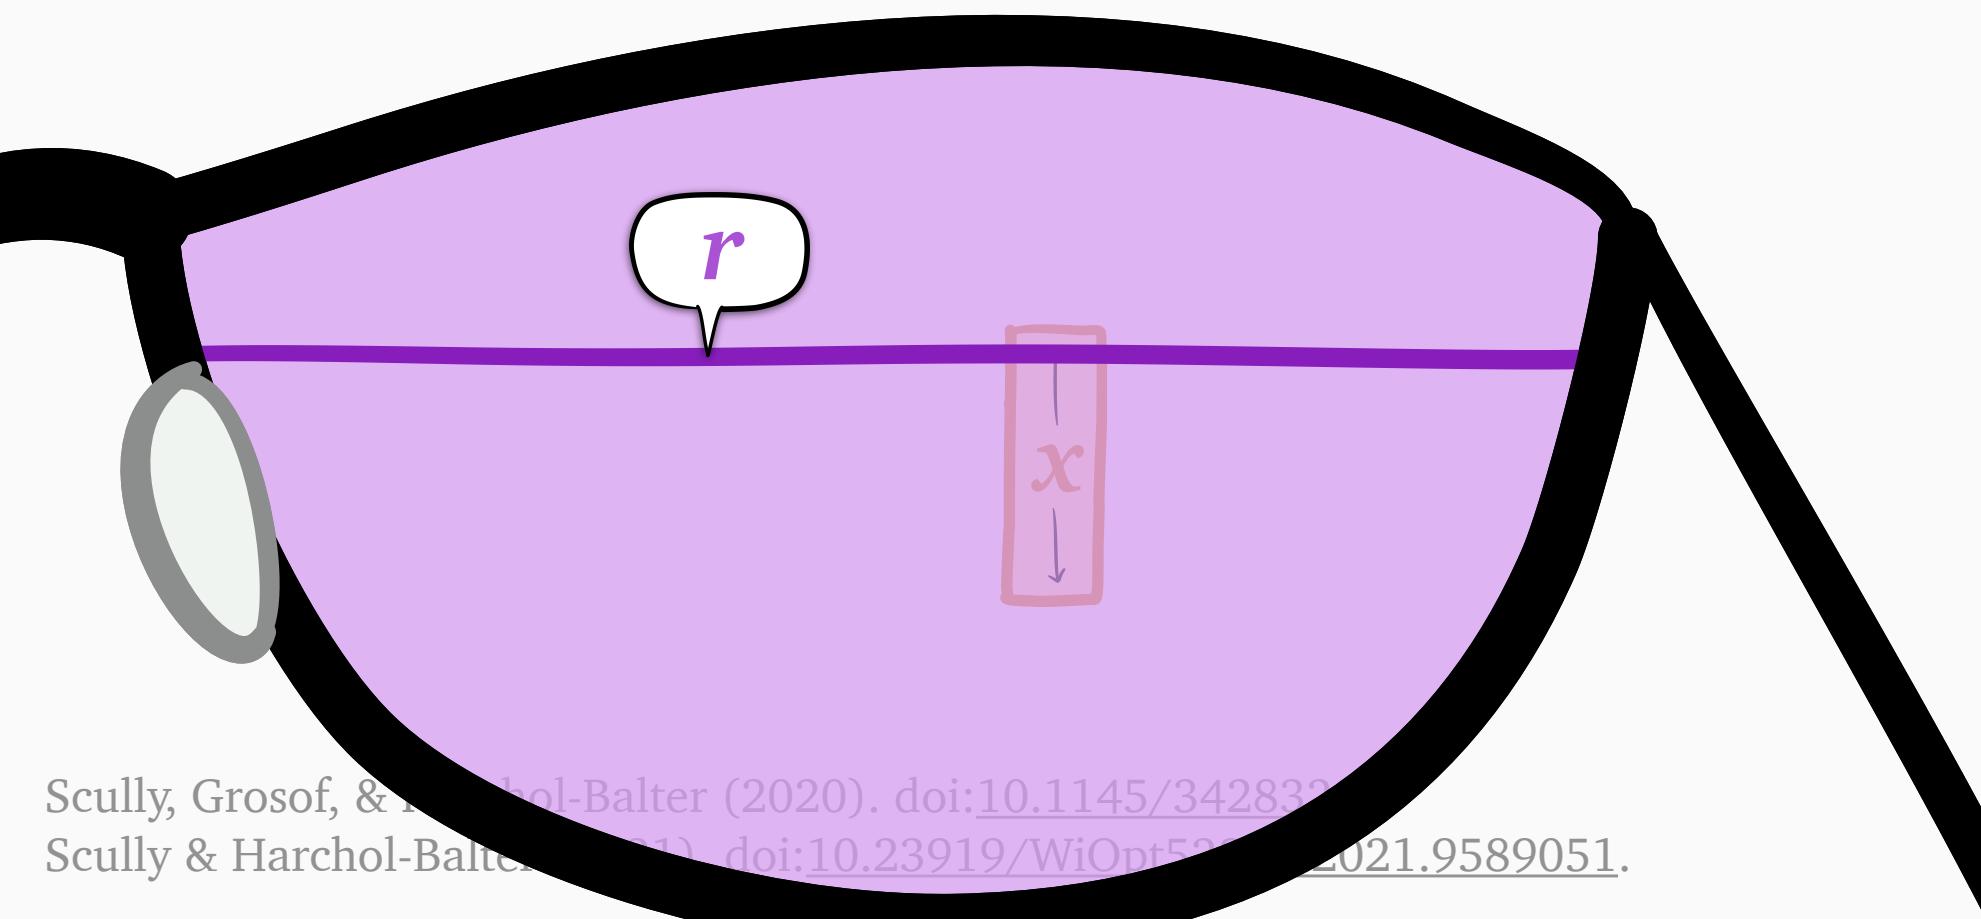
$w_x(r)$ = r -work of *single job* of rem. size x = {



WINE and r -work

$W(r)$ = work relevant to rank r

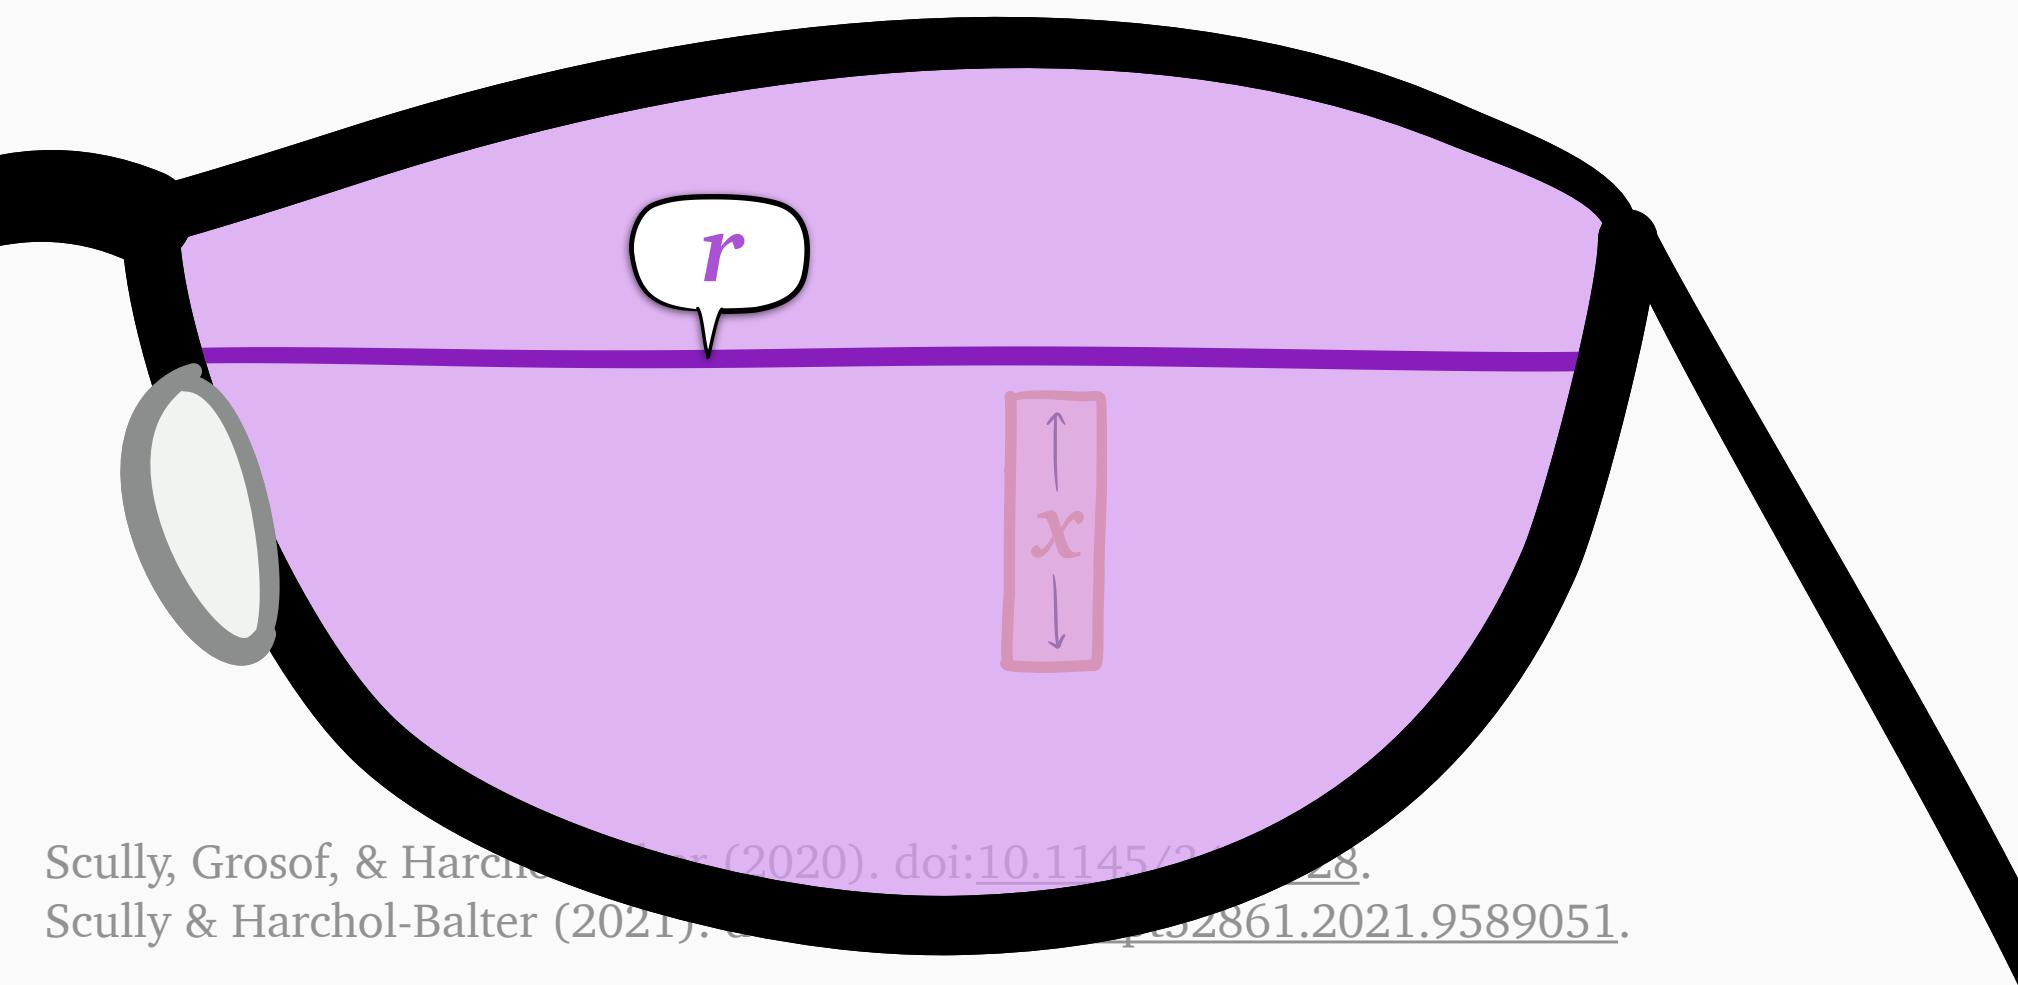
$w_x(r) = r\text{-work of } single job \text{ of rem. size } x = \begin{cases} 0 & \text{if } r < x \\ \dots & \end{cases}$



WINE and r -work

$W(r)$ = work relevant to rank r

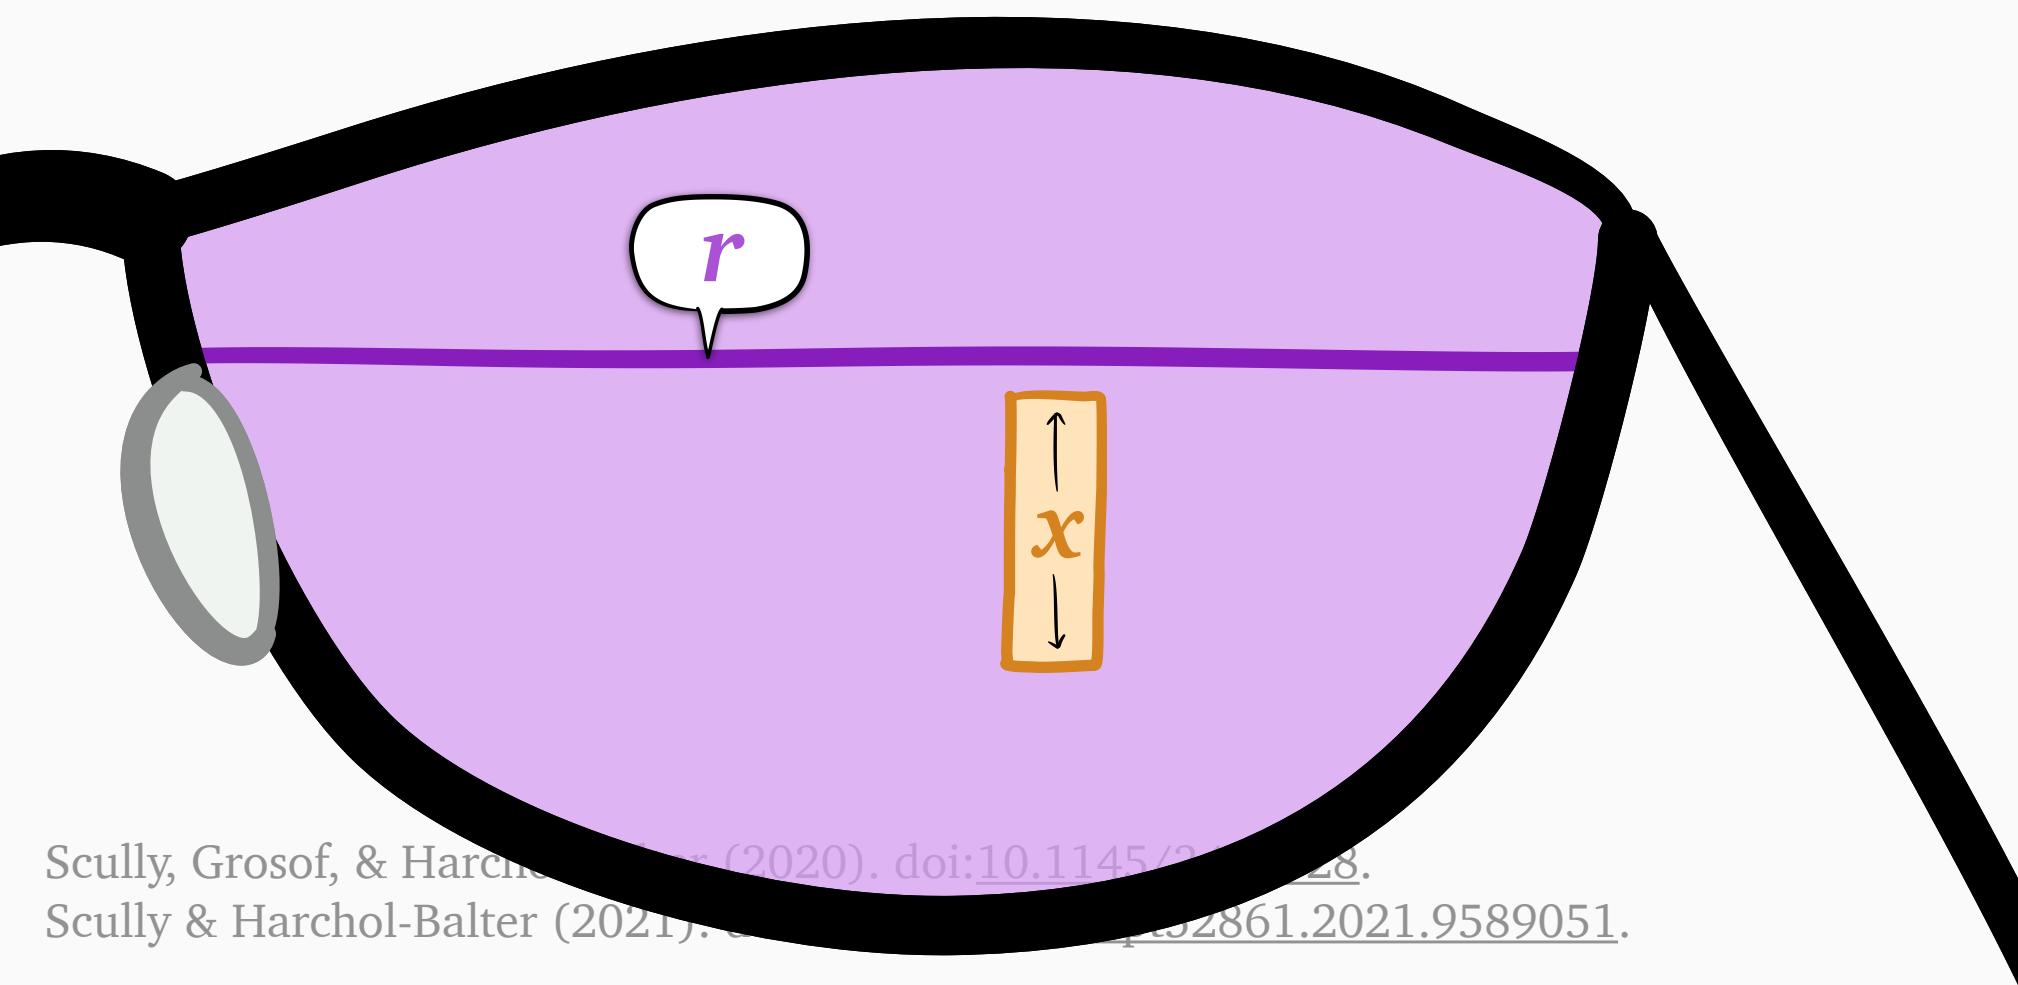
$w_x(r) = r\text{-work of } single job \text{ of rem. size } x = \begin{cases} 0 & \text{if } r < x \\ \end{cases}$



WINE and r -work

$W(r)$ = work relevant to rank r

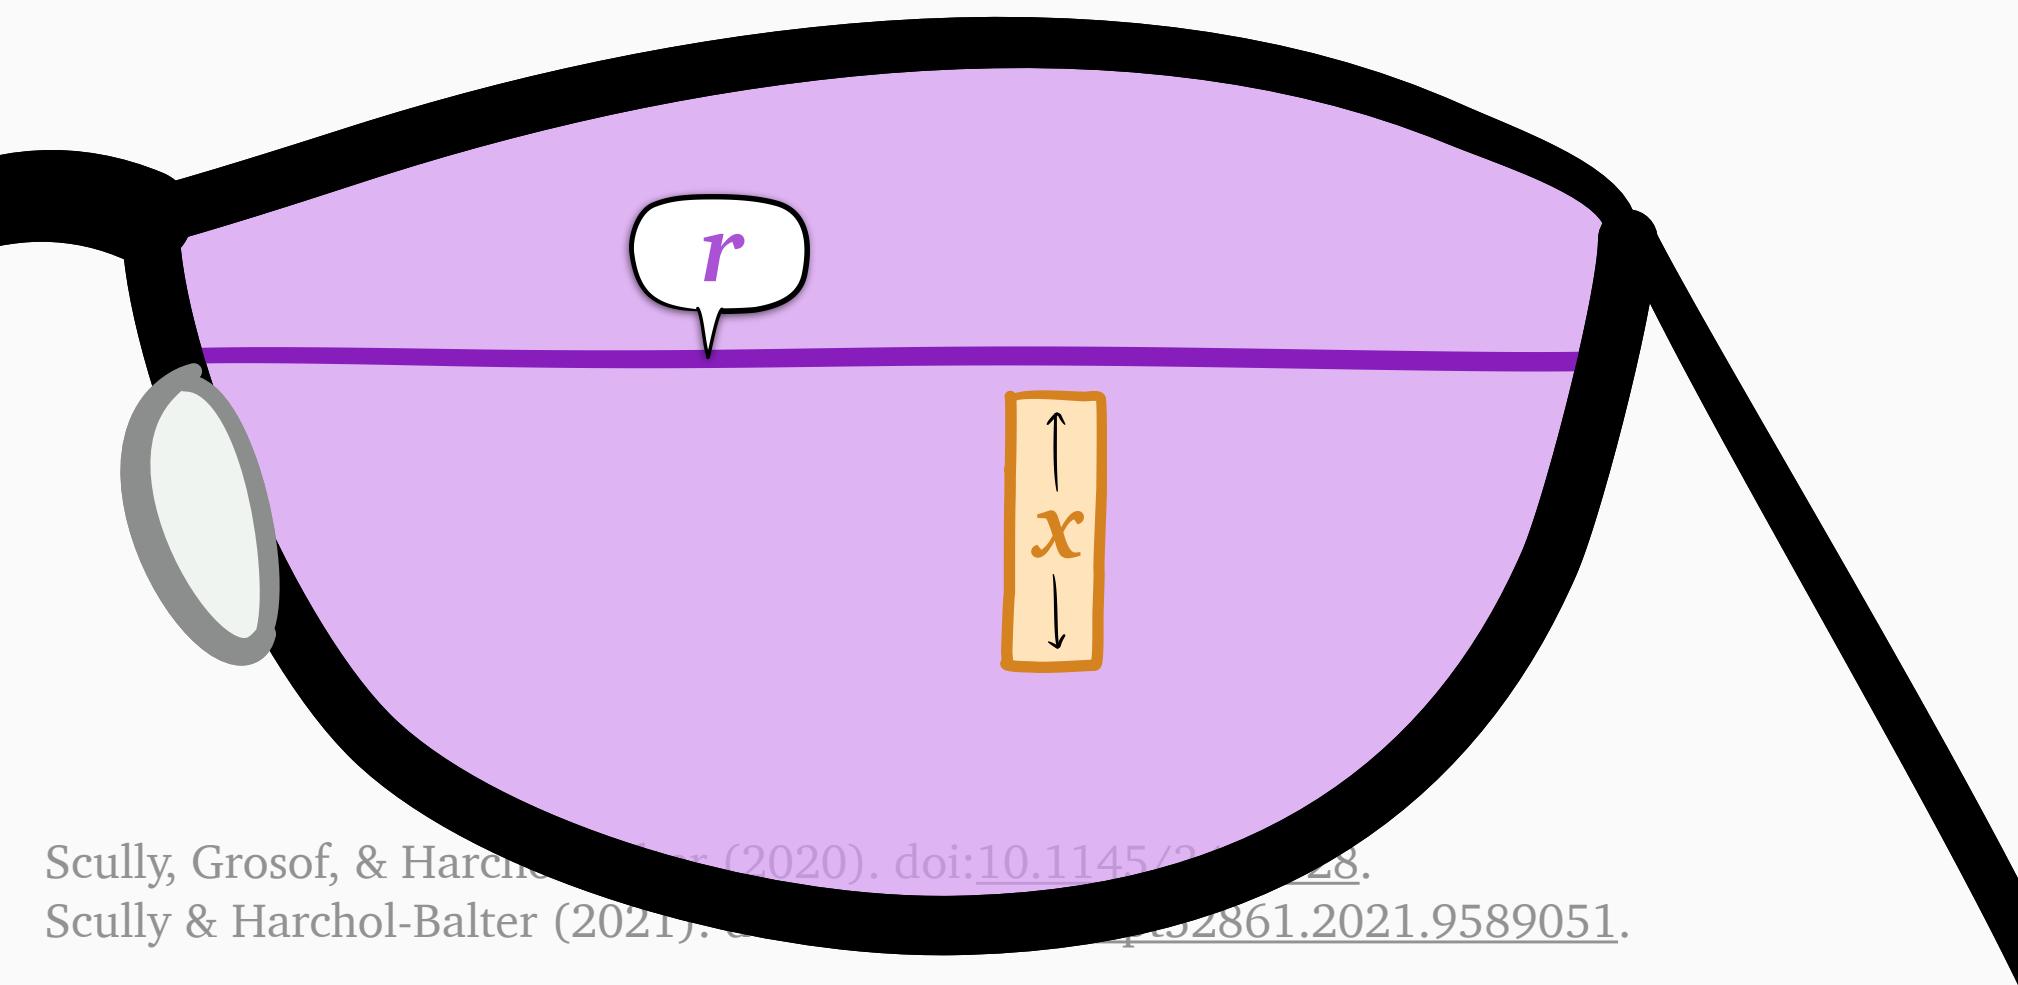
$w_x(r) = r\text{-work of } single job \text{ of rem. size } x = \begin{cases} 0 & \text{if } r < x \\ \end{cases}$



WINE and r -work

$W(r)$ = work relevant to rank r

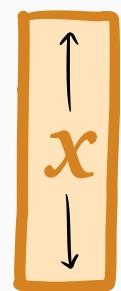
$w_x(r) = r\text{-work of } single job \text{ of rem. size } x = \begin{cases} 0 & \text{if } r < x \\ x & \text{if } r \geq x \end{cases}$



WINE and r -work

$W(r)$ = work relevant to rank r

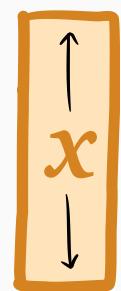
$w_x(r)$ = r -work of *single job* of rem. size x = $\begin{cases} 0 & \text{if } r < x \\ x & \text{if } r \geq x \end{cases}$



WINE and r -work

$W(r)$ = work relevant to **rank r**
= total r -work of all jobs

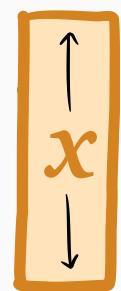
$w_x(r)$ = r -work of *single job* of rem. size x = $\begin{cases} 0 & \text{if } r < x \\ x & \text{if } r \geq x \end{cases}$



WINE and r -work

$W(r)$ = work relevant to rank r
= total r -work of all jobs

$$w_x(r) = r\text{-work of } \textit{single job} \text{ of rem. size } x = \begin{cases} 0 & \text{if } r < x \\ x & \text{if } r \geq x \end{cases}$$



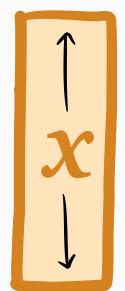
Theorem: in any system,

$$N = \int_0^\infty \frac{W(r)}{r^2} dr$$

WINE and r -work

$W(r)$ = work relevant to **rank r**
= total r -work of all jobs

$$w_x(r) = r\text{-work of } \textit{single job} \text{ of rem. size } x = \begin{cases} 0 & \text{if } r < x \\ x & \text{if } r \geq x \end{cases}$$

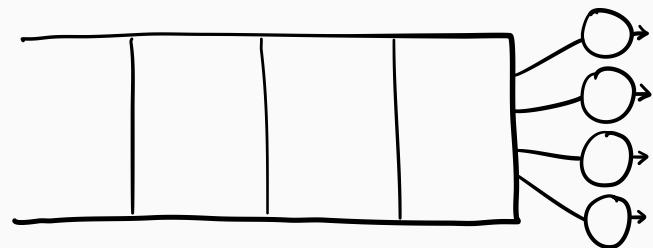


Theorem: in *any* system,

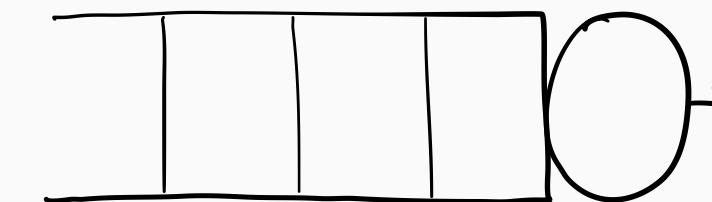
$$N = \int_0^\infty \frac{W(r)}{r^2} dr$$

SRPT in the M/G/ $\textcolor{red}{k}$

SRPT- $\textcolor{red}{k}$



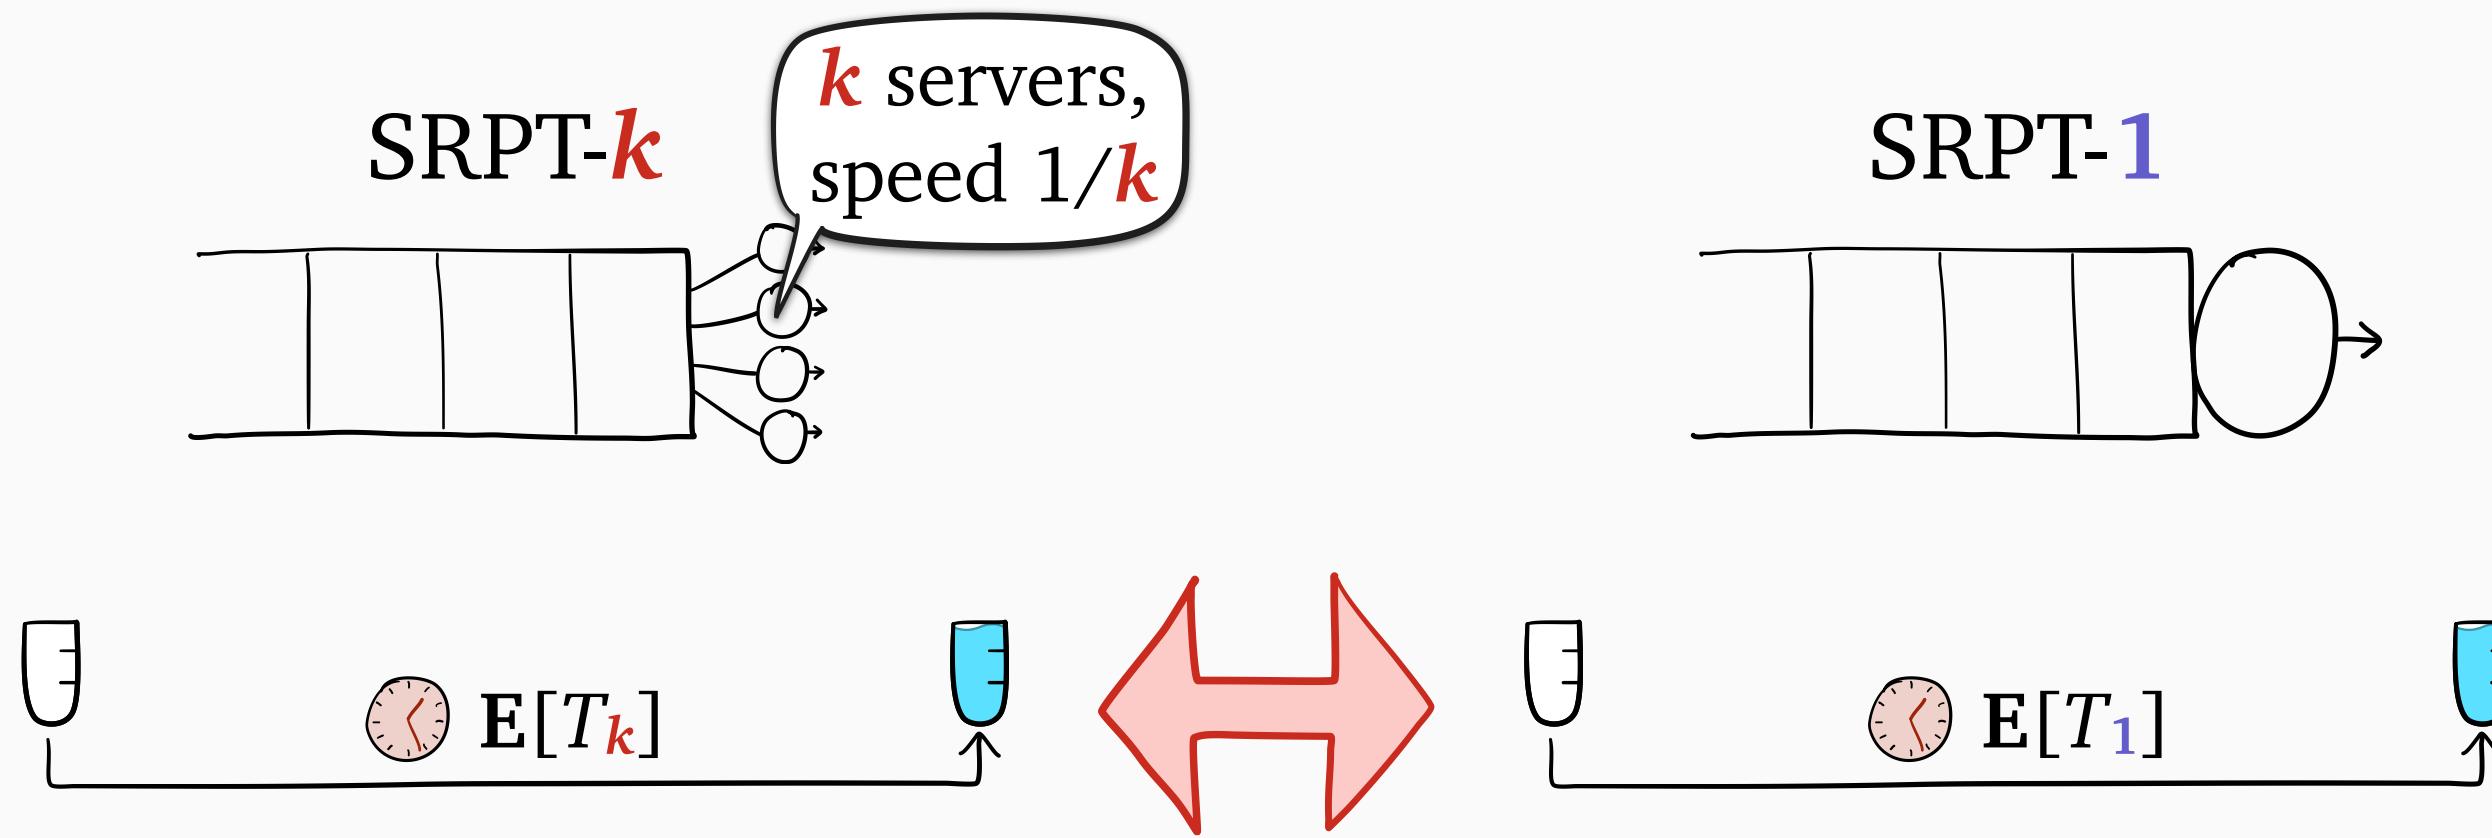
SRPT-1



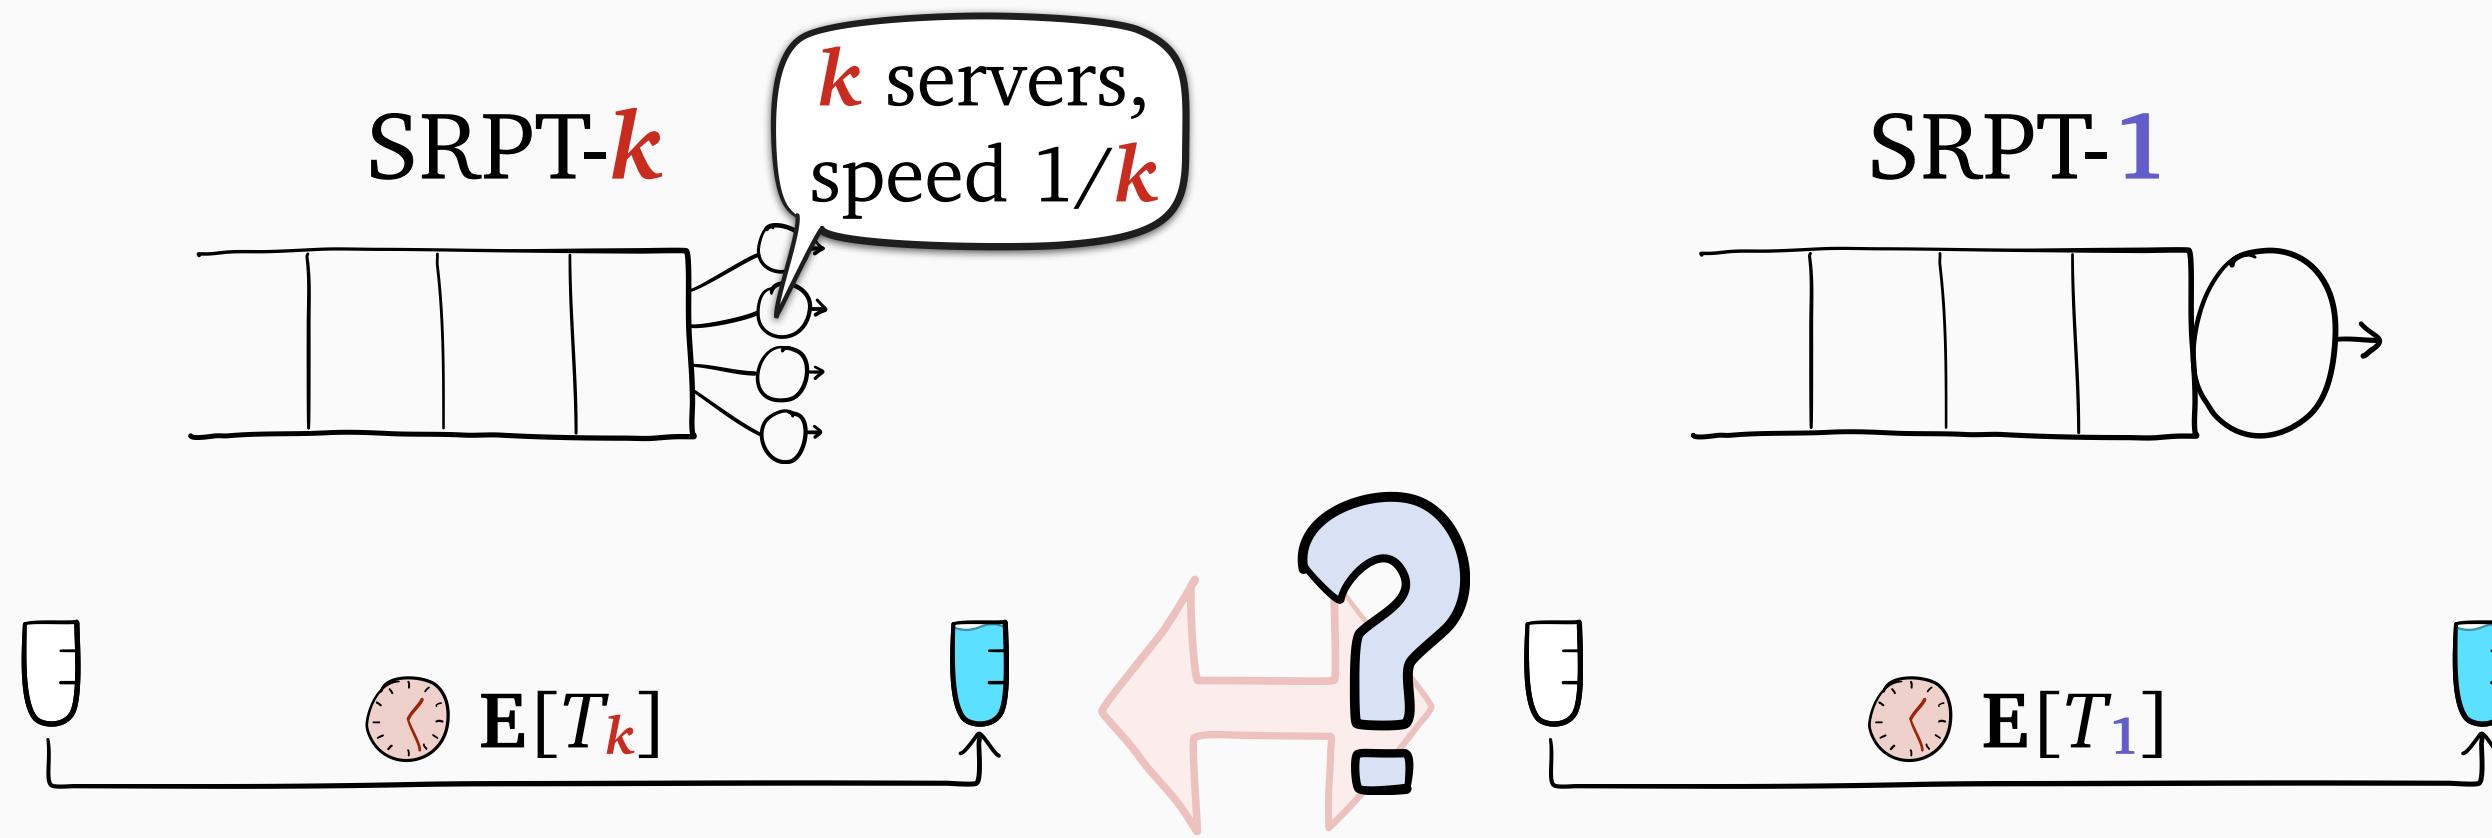
SRPT in the M/G/ k



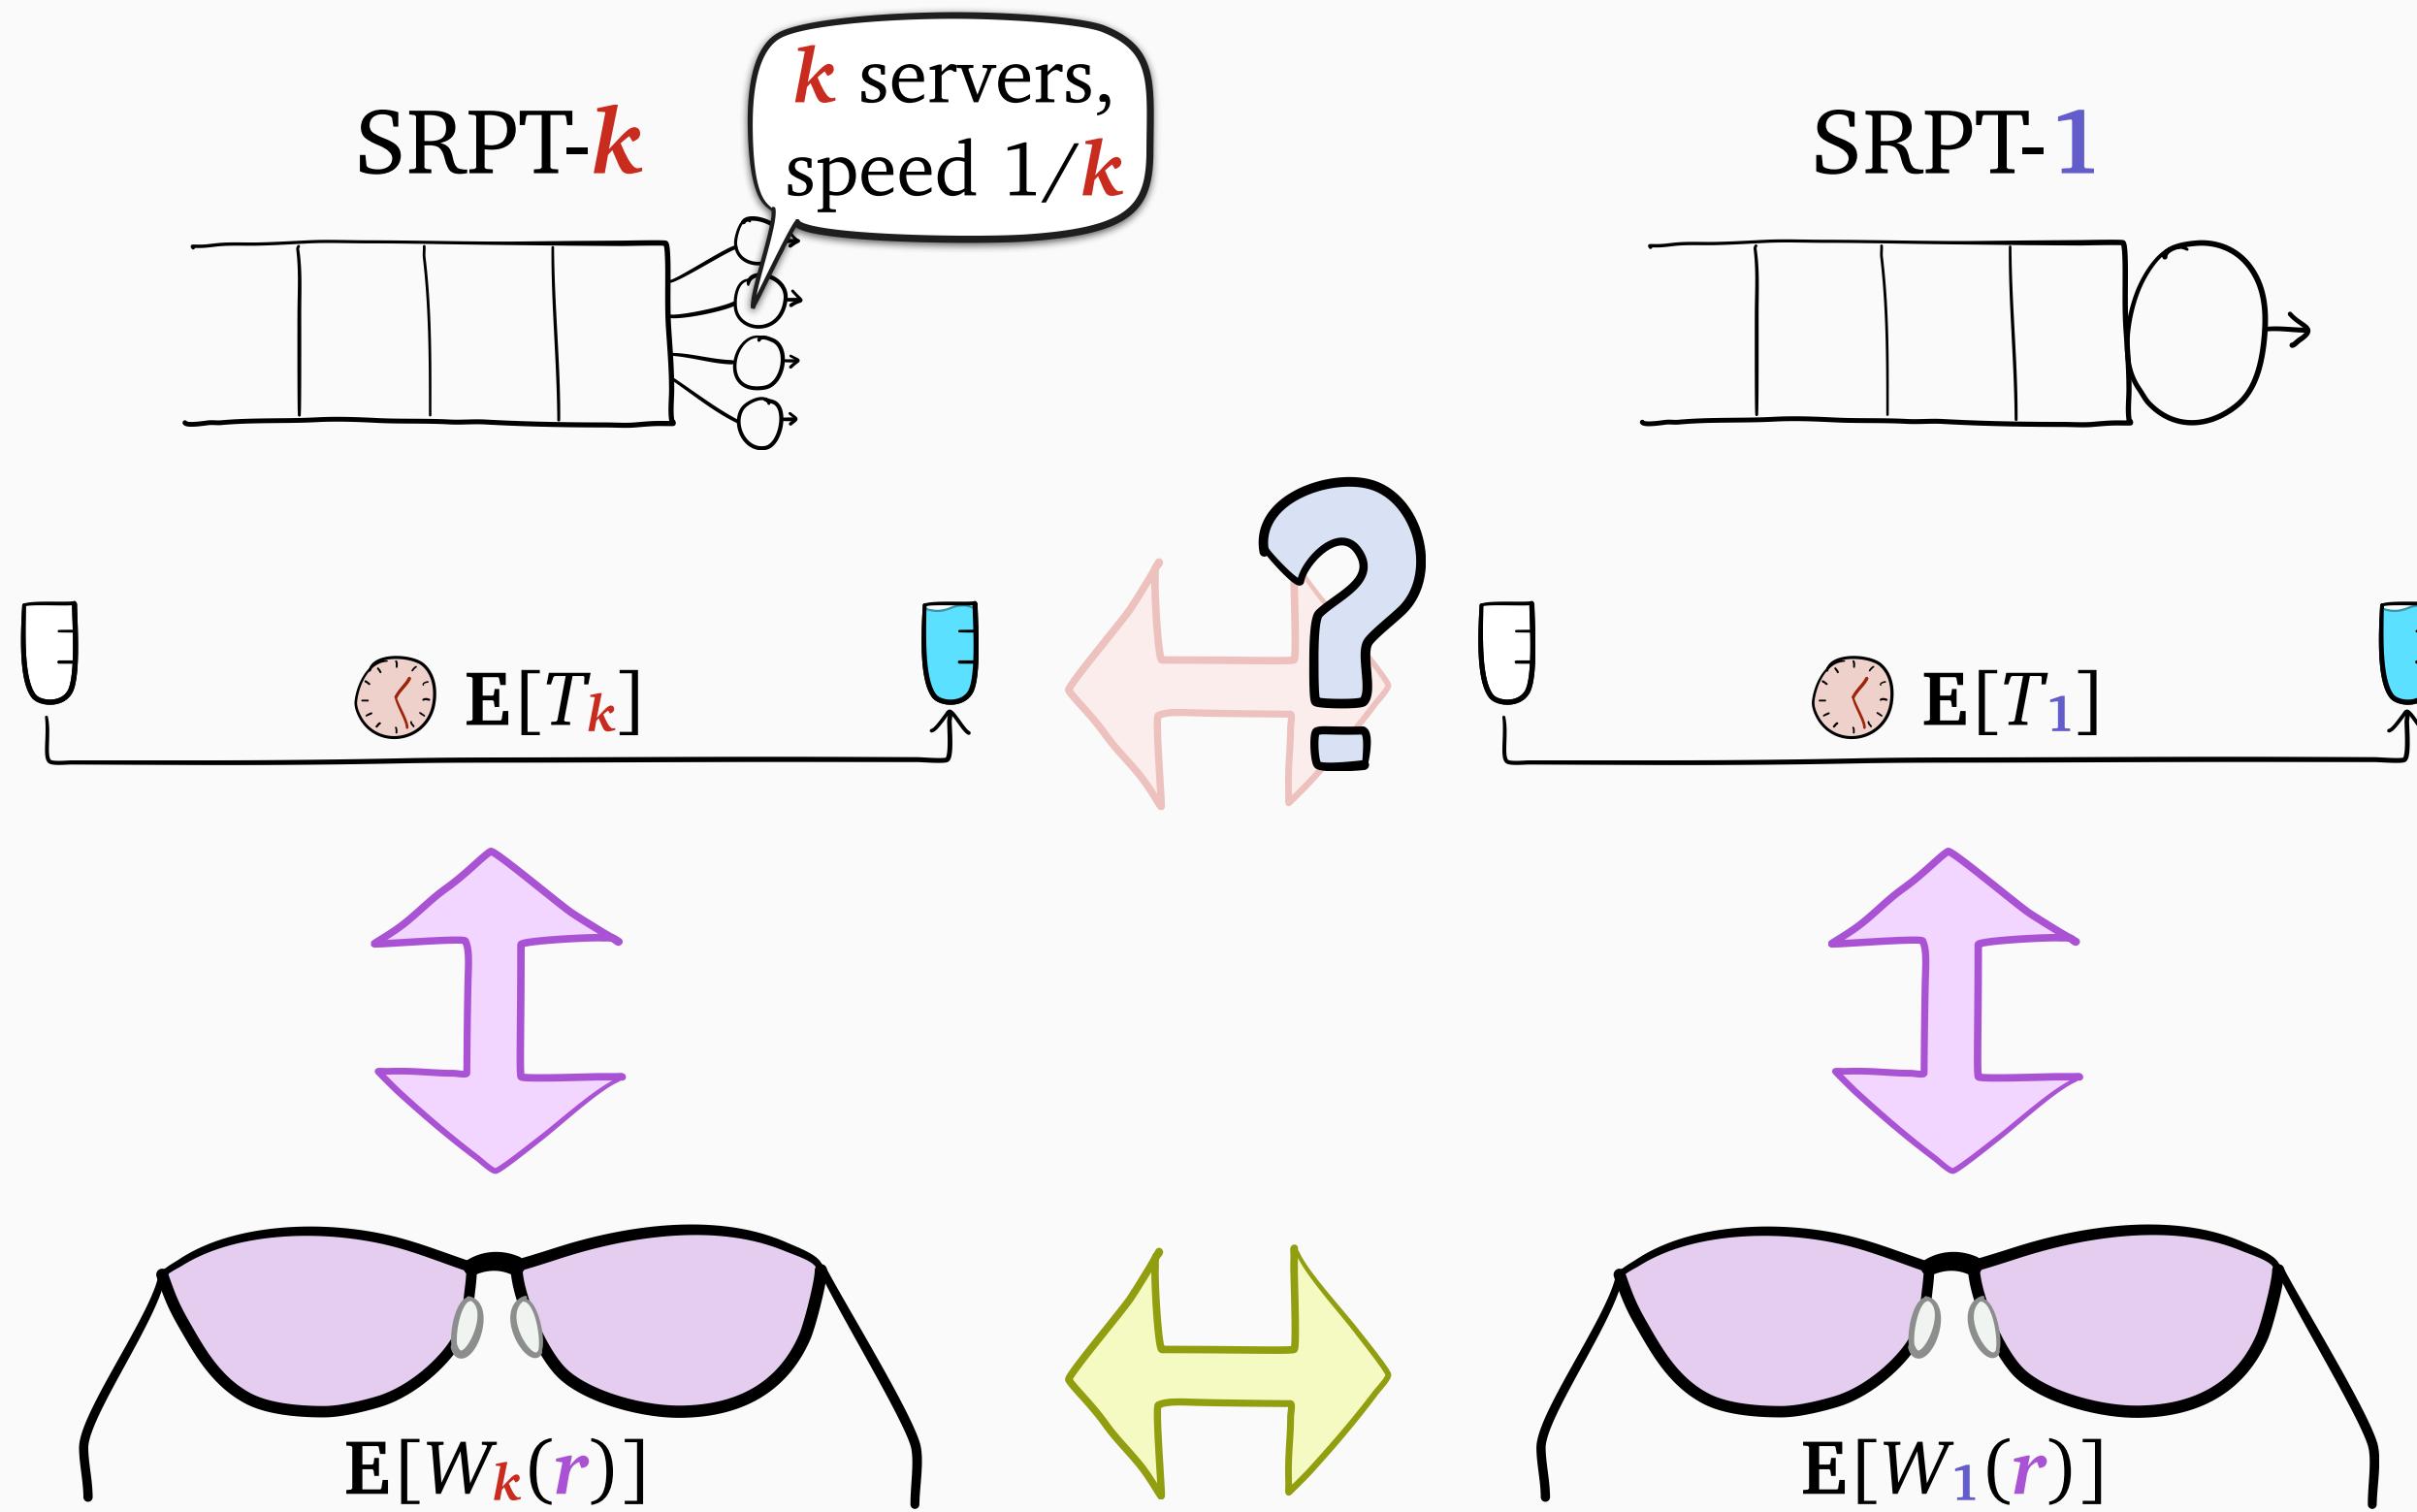
SRPT in the M/G/ k



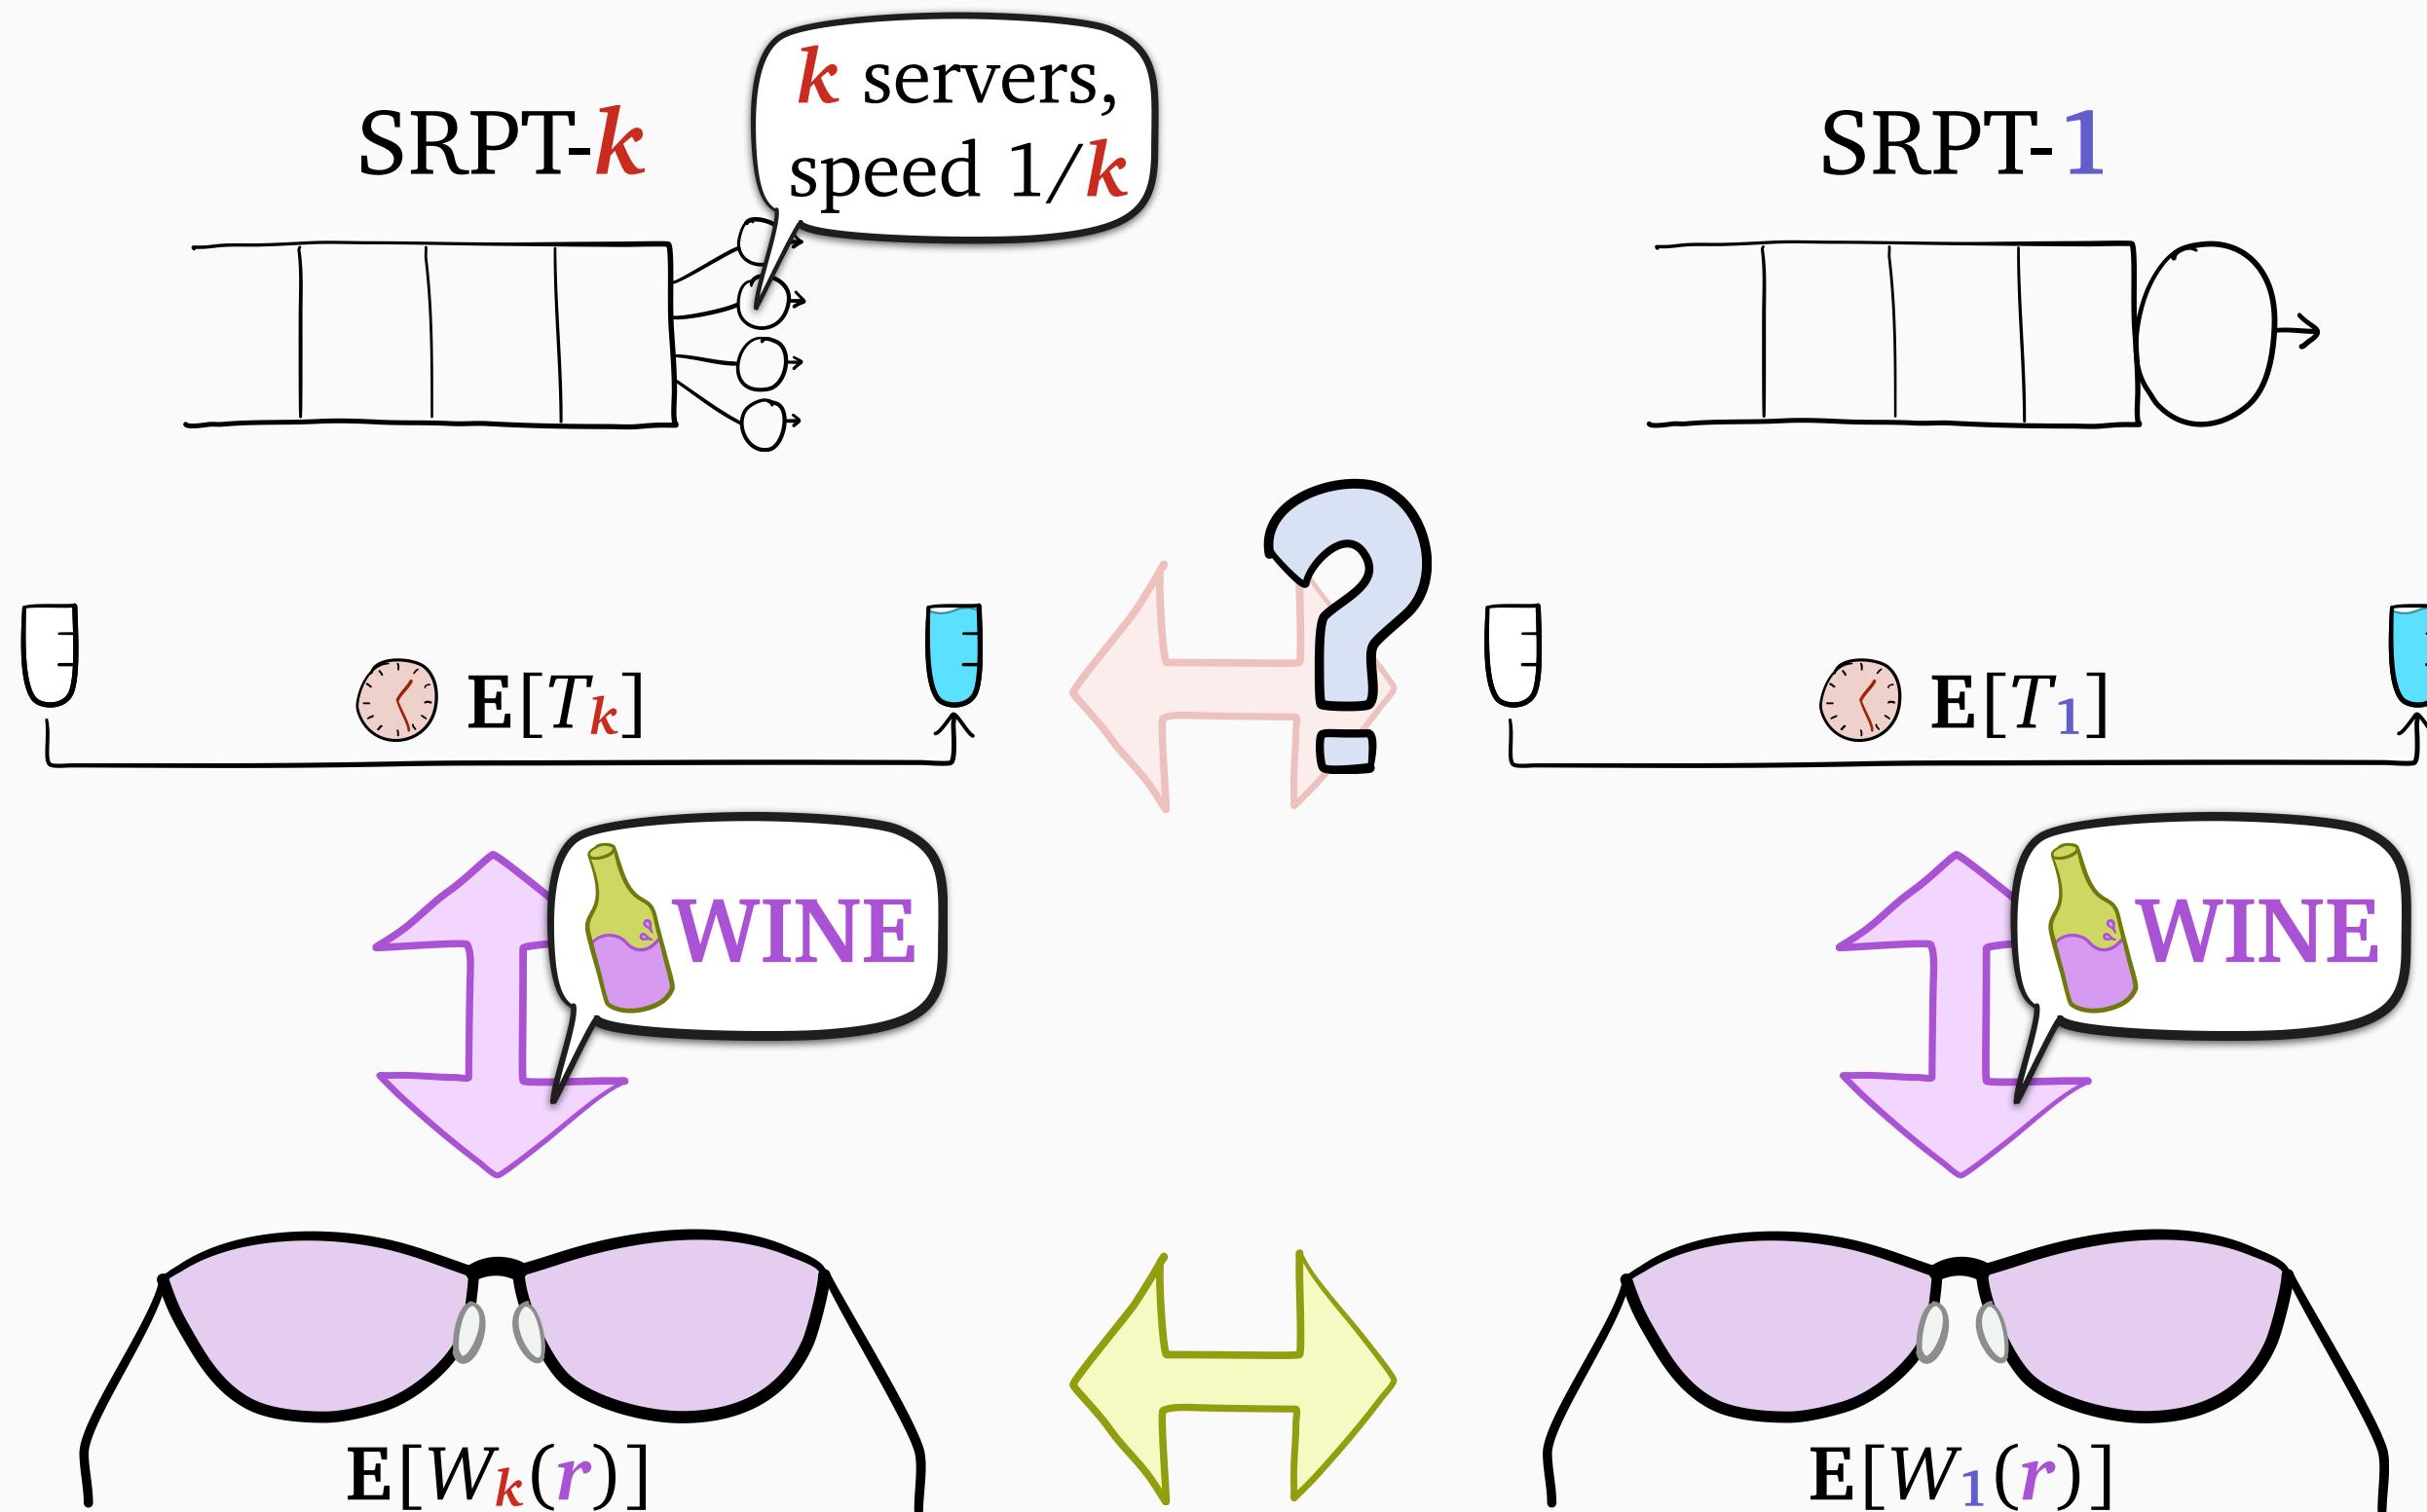
SRPT in the M/G/ k



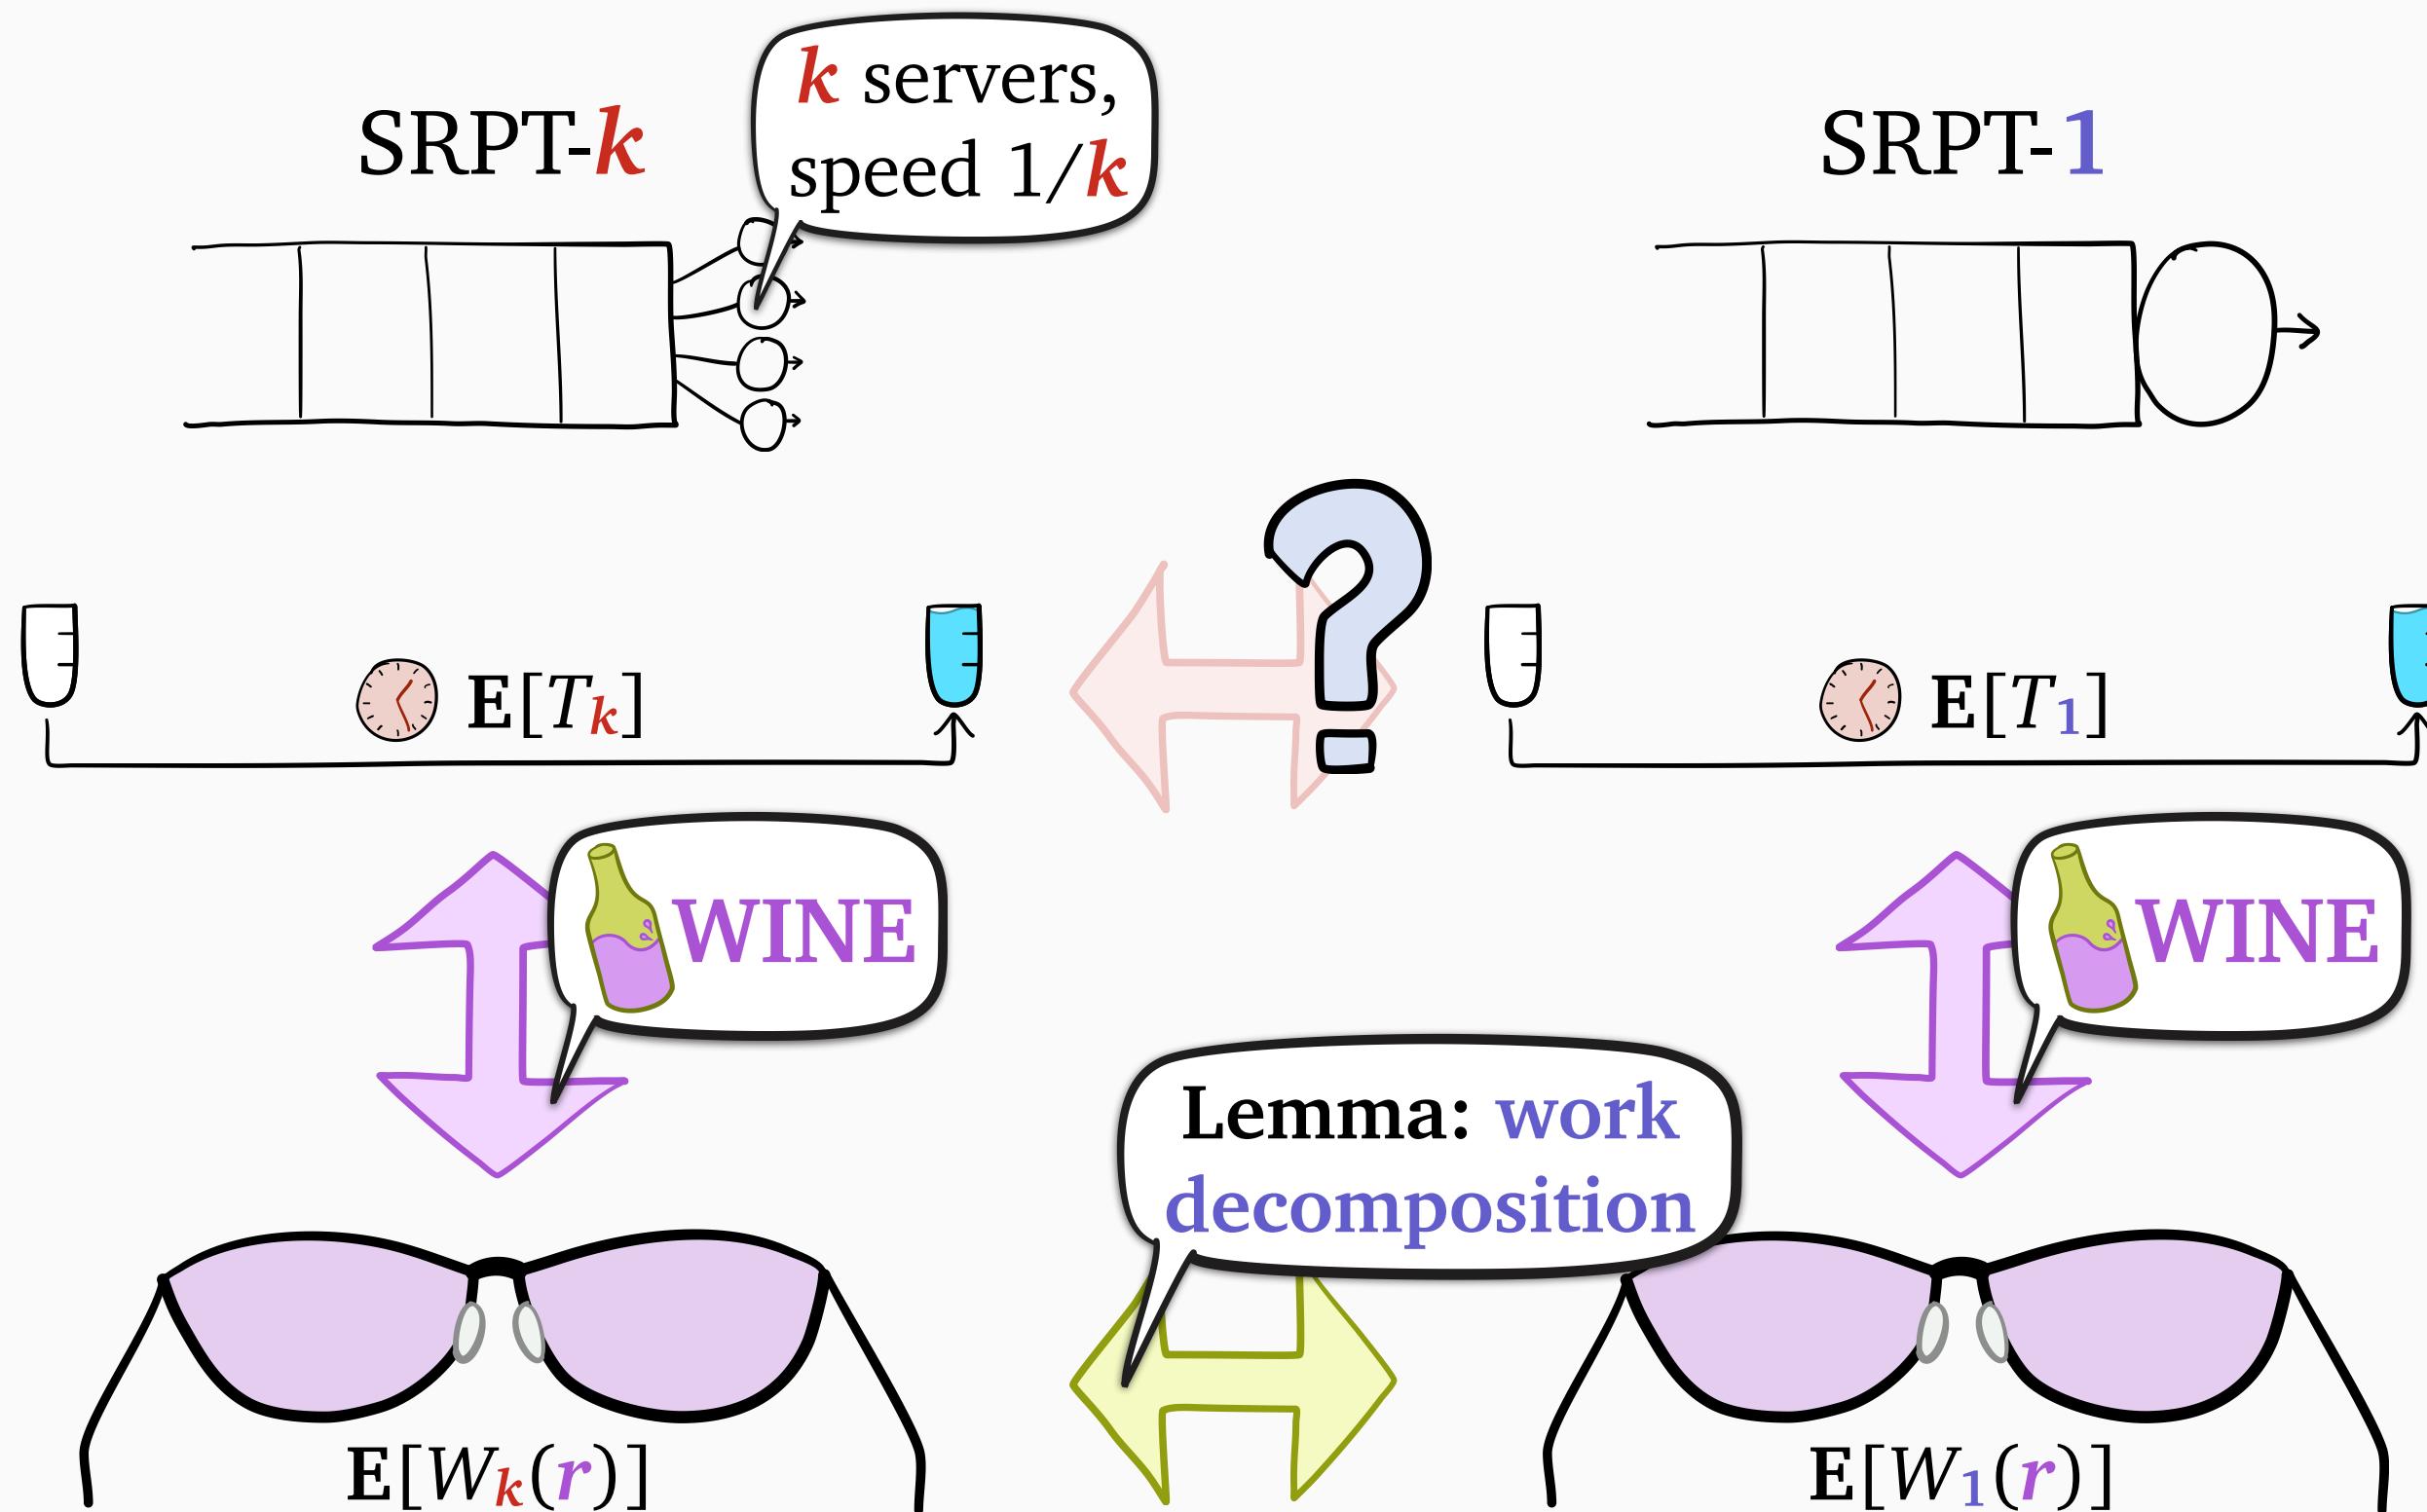
SRPT in the M/G/ k



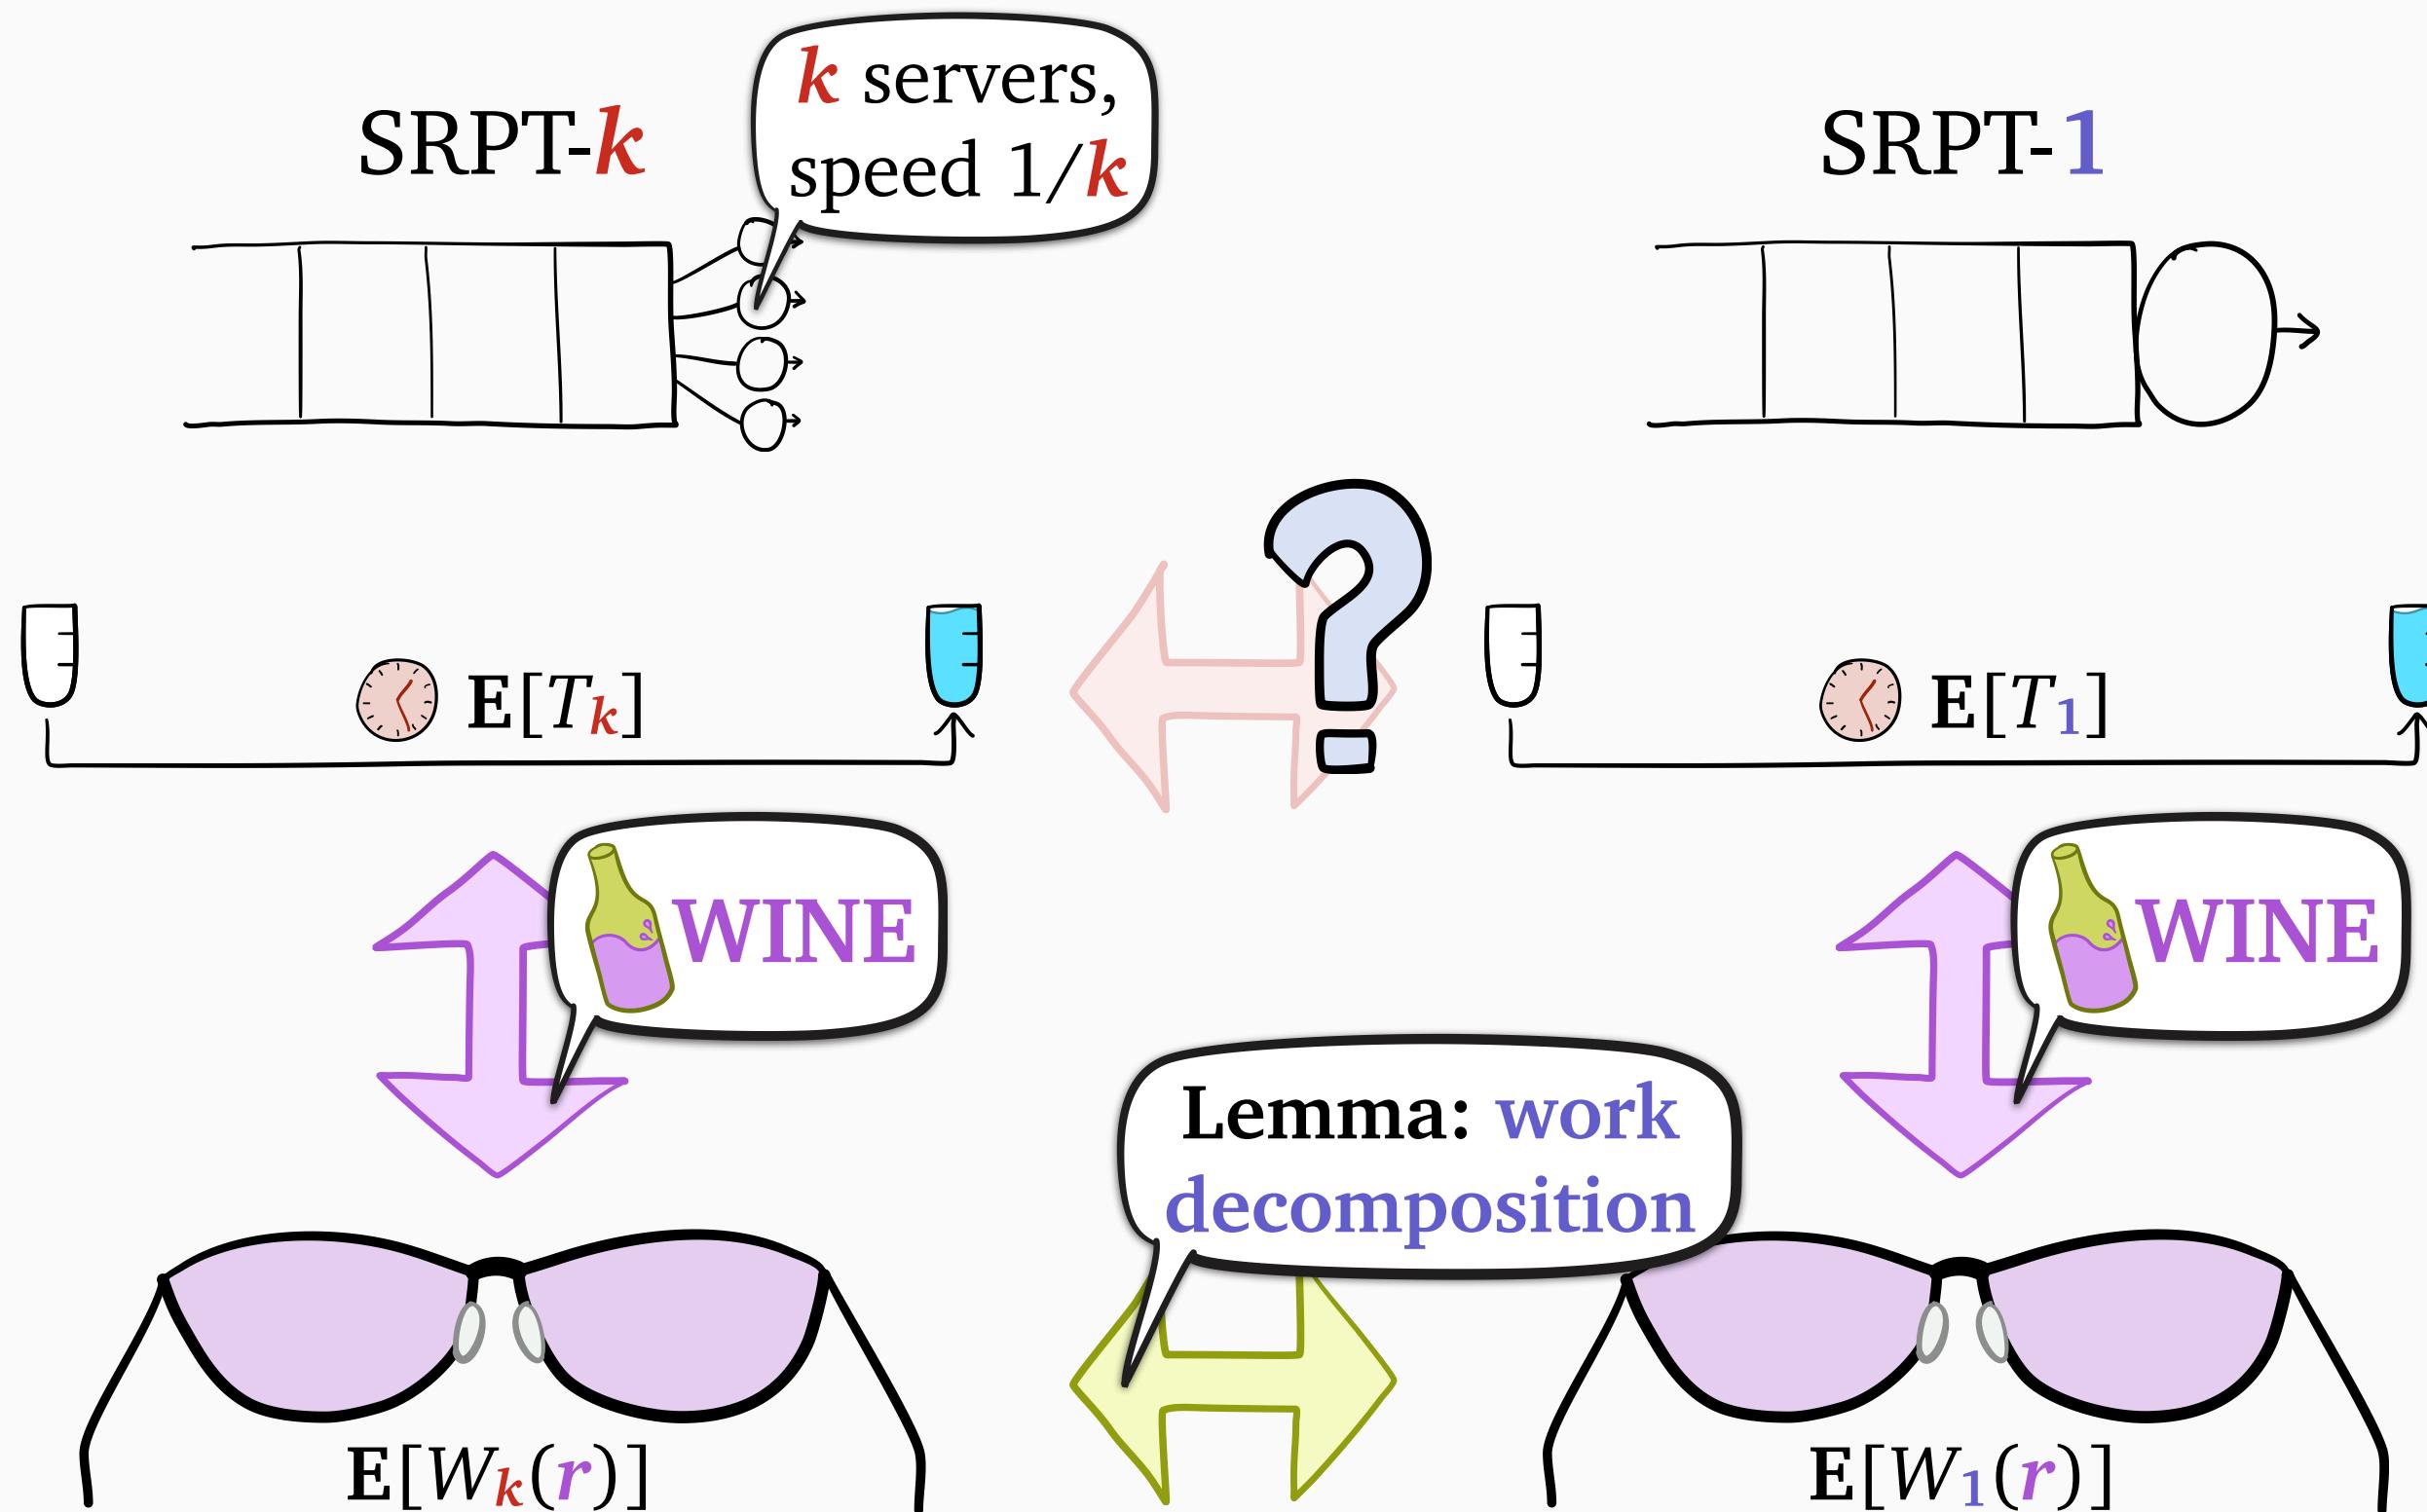
SRPT in the M/G/ k



SRPT in the M/G/ k

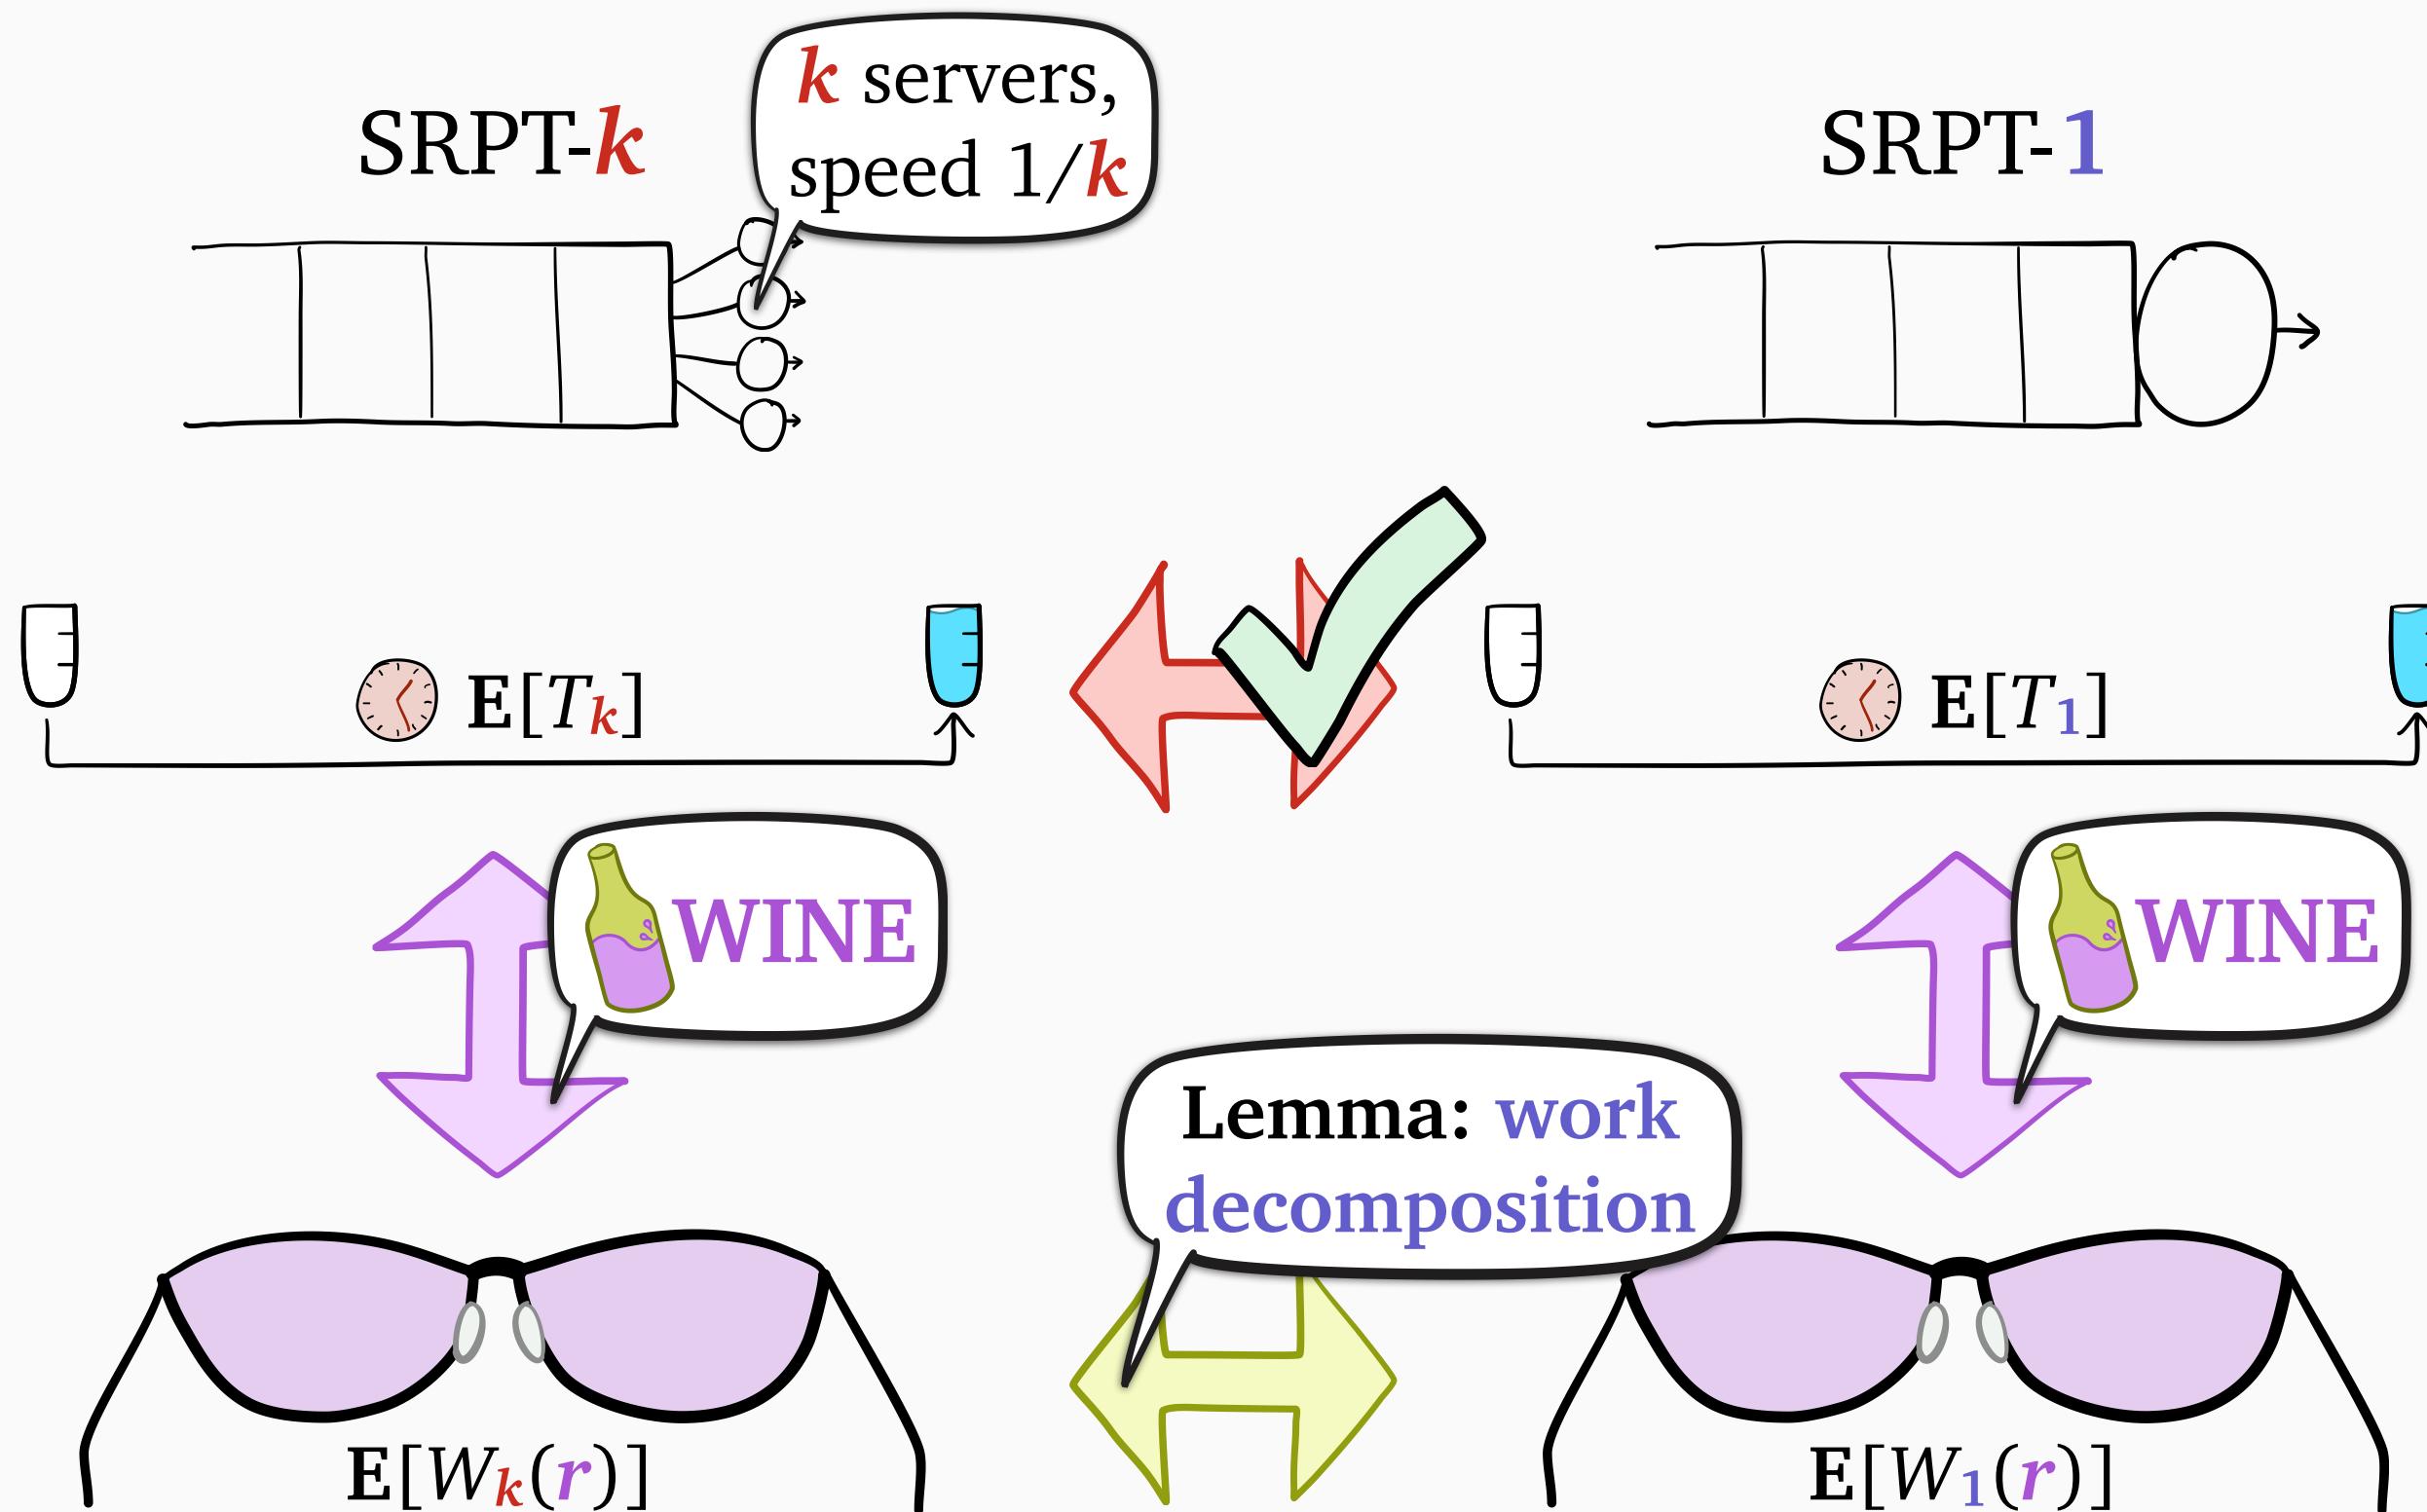


SRPT in the M/G/ k



$$\begin{aligned} \text{Theorem: } & E[S(\log S)^+] < \infty \\ \Downarrow \\ \lim_{\rho \rightarrow 1} \frac{E[T_{SRPT-k}]}{E[T_{SRPT-1}]} &= 1 \end{aligned}$$

SRPT in the M/G/ k



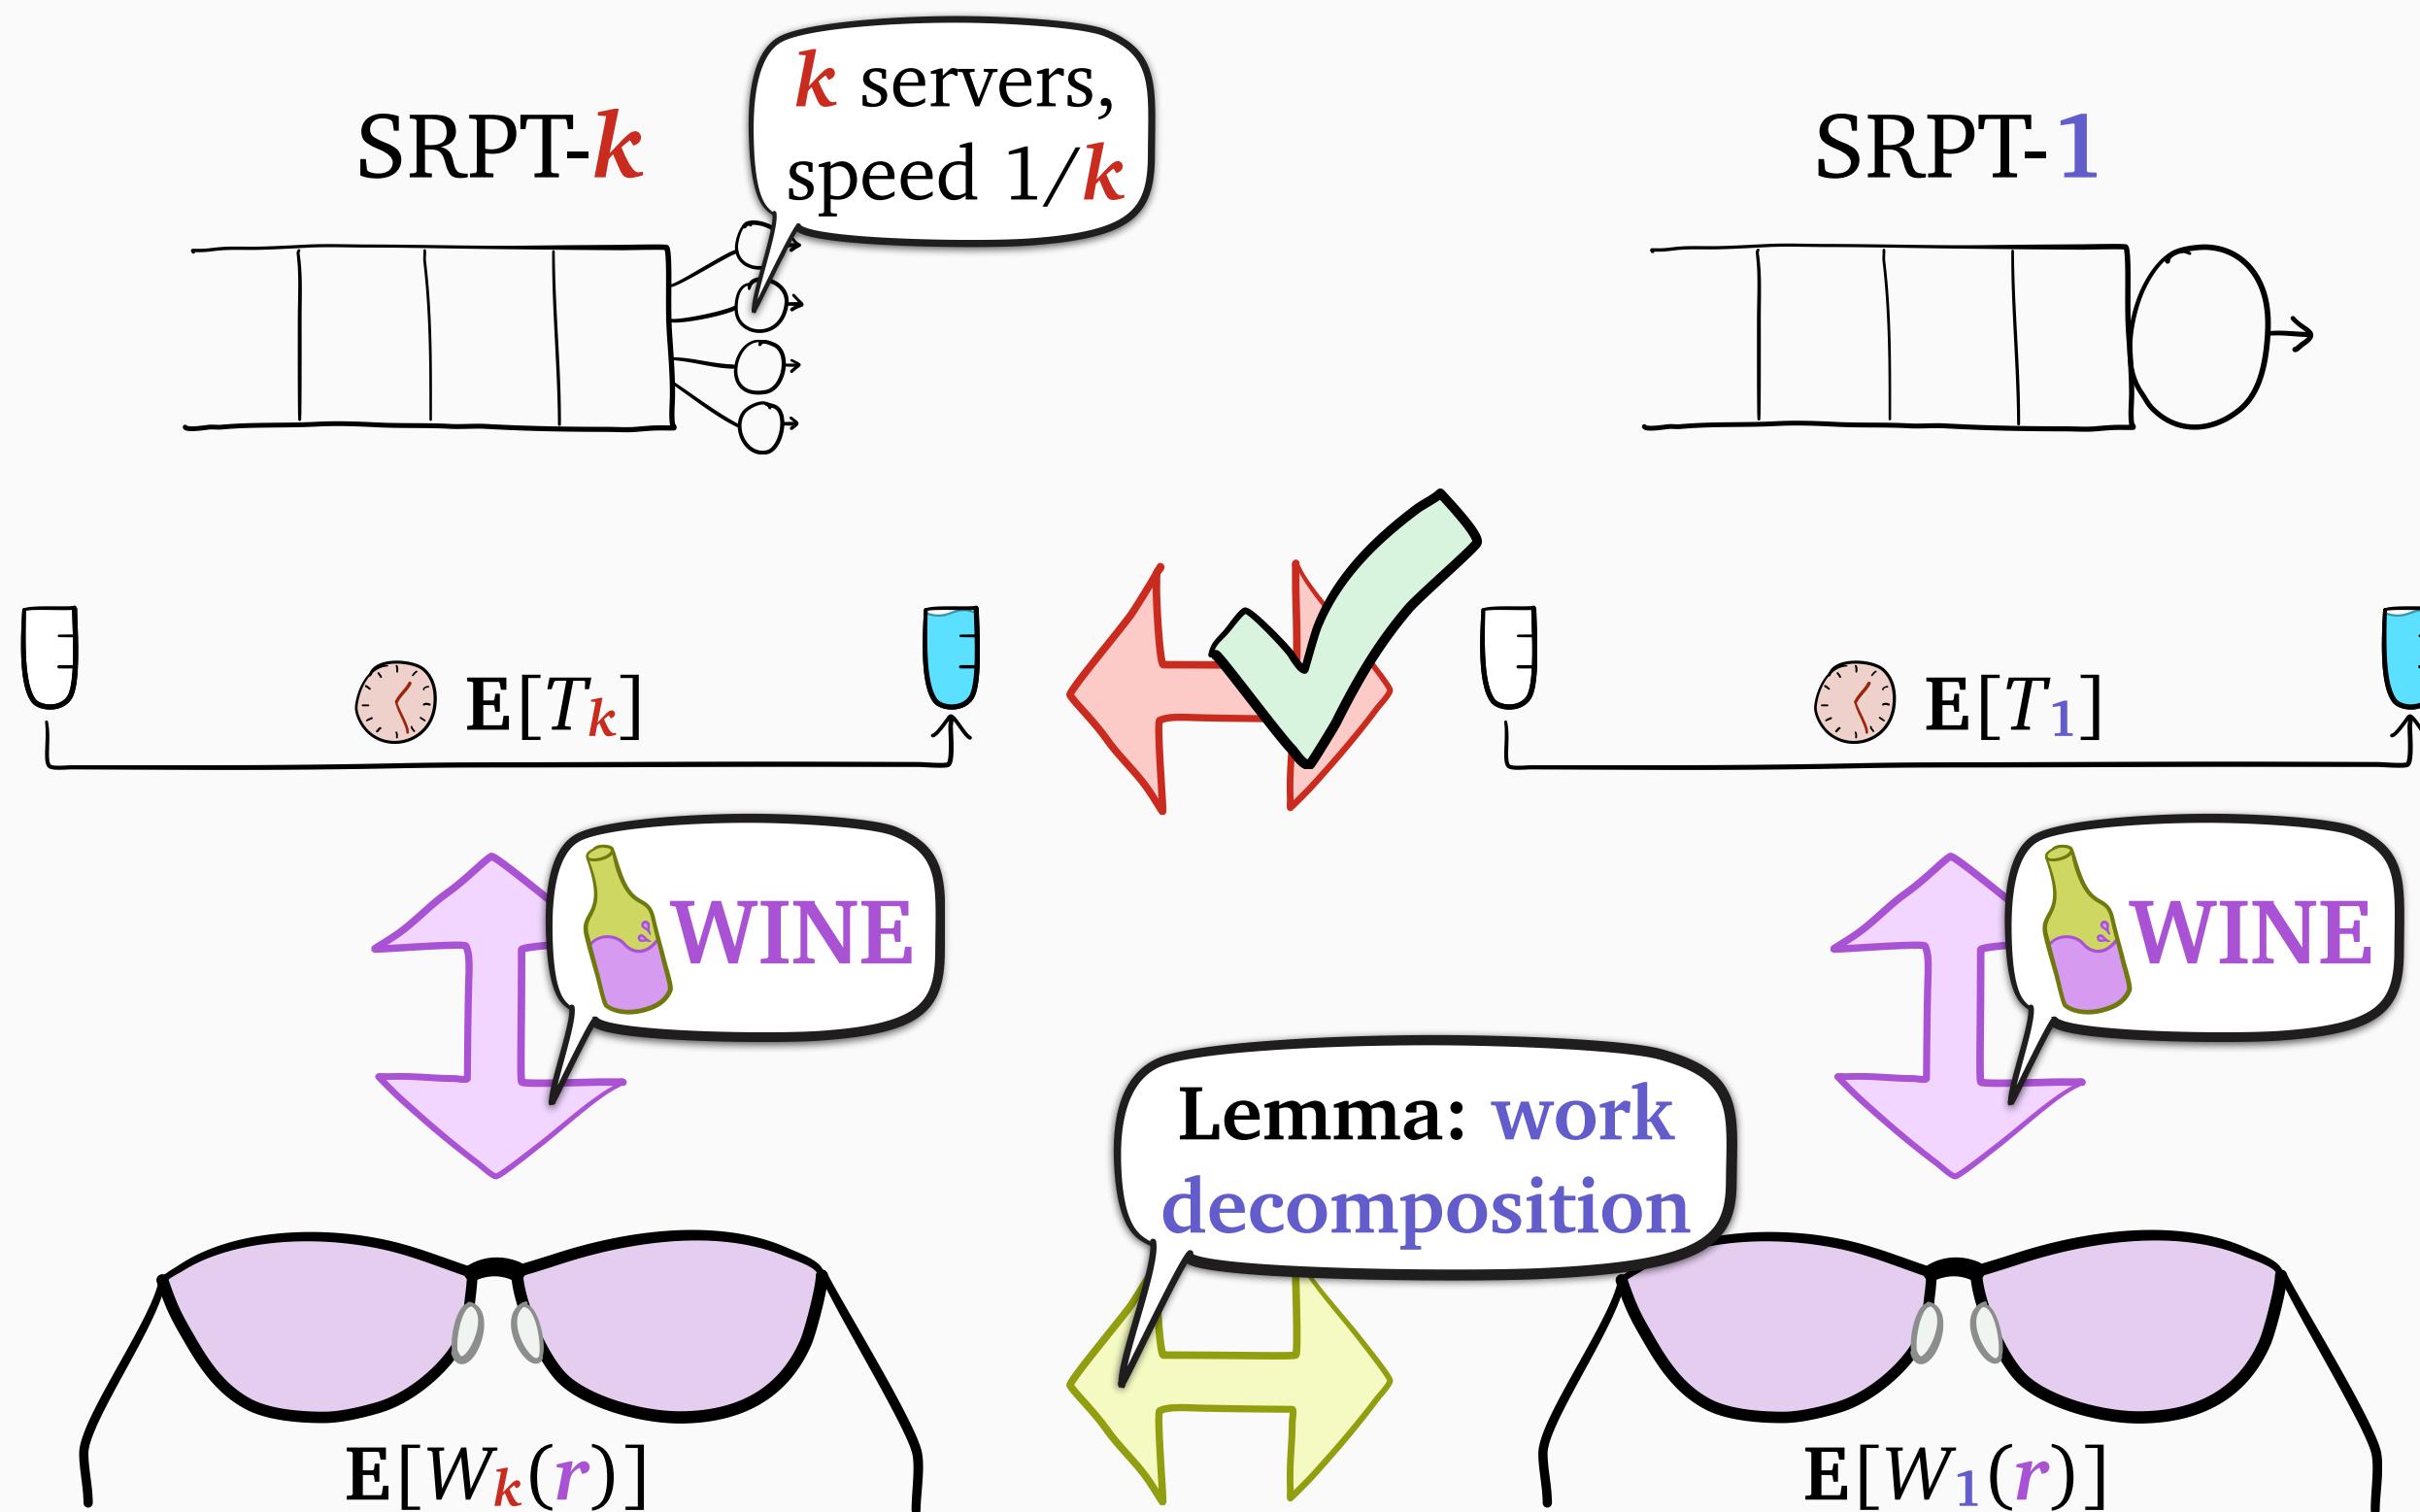
Theorem:

$$E[S(\log S)^+] < \infty$$

$$\Downarrow$$

$$\lim_{\rho \rightarrow 1} \frac{E[T_{SRPT-k}]}{E[T_{SRPT-1}]} = 1$$

SRPT in the M/G/ k



Gittins, too

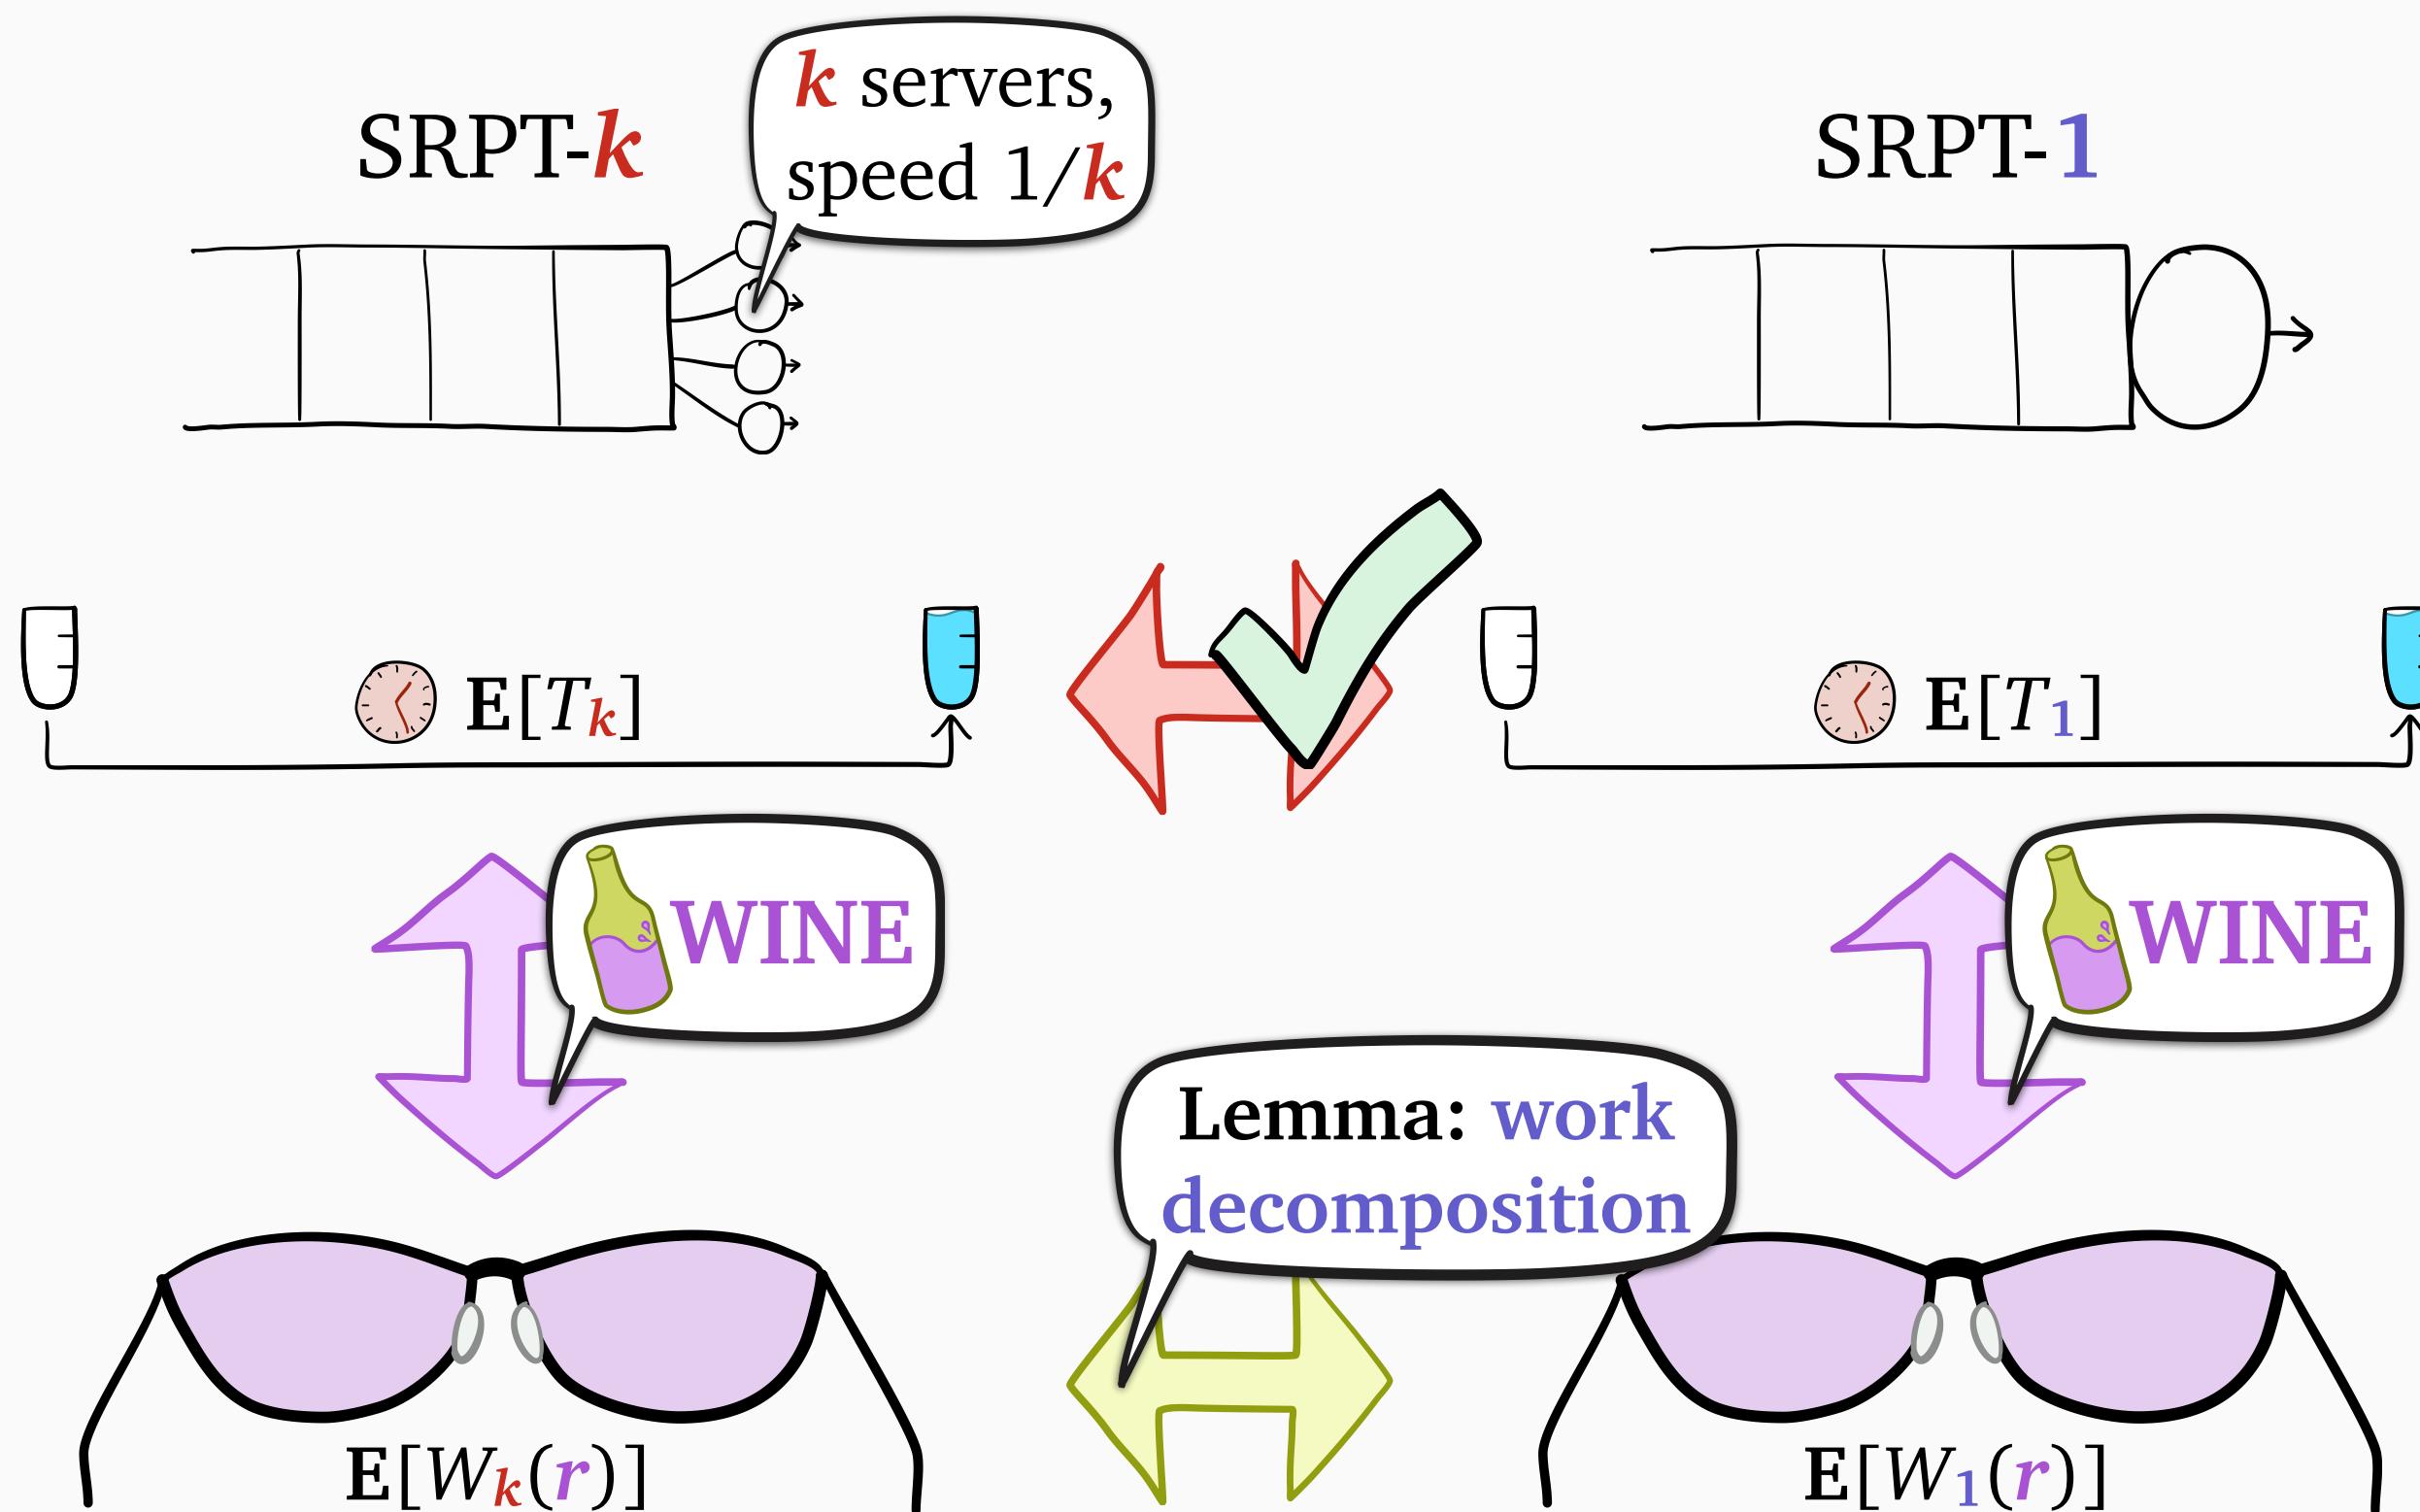
Theorem:

$$E[S(\log S)^+] < \infty$$



$$\lim_{\rho \rightarrow 1} \frac{E[T_{SRPT-k}]}{E[T_{SRPT-1}]} = 1$$

SRPT in the M/G/ k



Gittins, too

Theorem:

$$E[S(\log S)^+] < \infty$$



$$\lim_{\rho \rightarrow 1} \frac{E[T_{SRPT-k}]}{E[T_{SRPT-1}]} = 1$$

Beyond the M/G/**k**

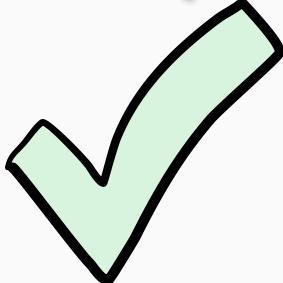
Beyond the M/G/ k



Dispatching with SRPT

Beyond the M/G/ k

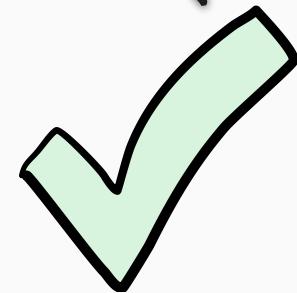


 Dispatching with SRPT

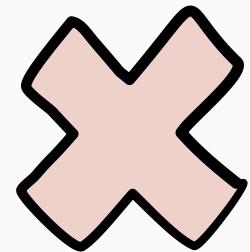
Beyond the M/G/ k



Spread out
small jobs



Dispatching with SRPT



Dispatching with Gittins

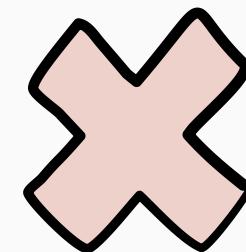
Beyond the M/G/ k



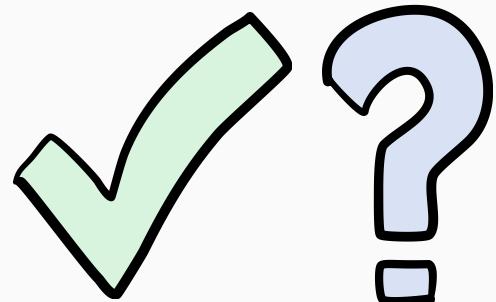
Spread out
small jobs



Dispatching with SRPT



Dispatching with Gittins



Multiserver jobs and other service constraints

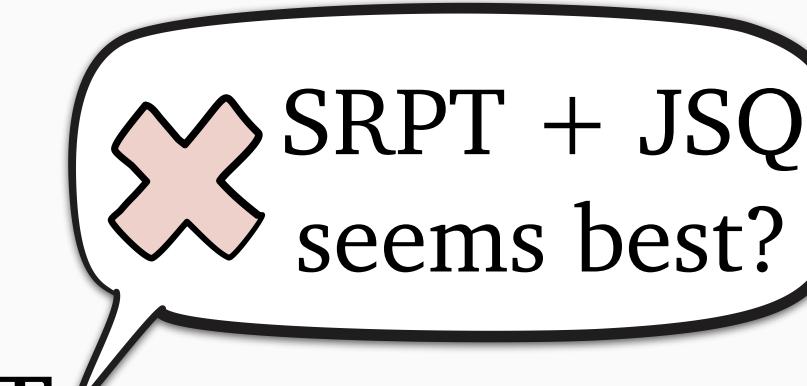
Beyond the M/G/ k



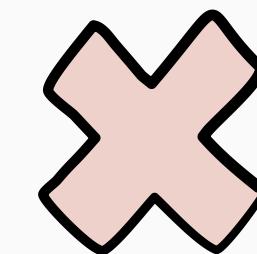
Spread out
small jobs



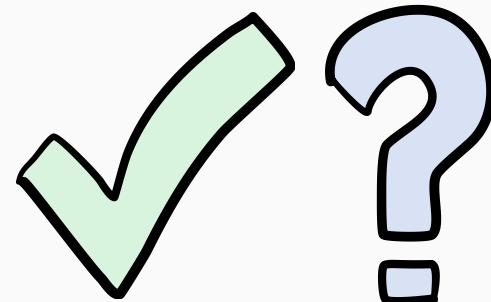
Dispatching with SRPT



SRPT + JSQ
seems best?

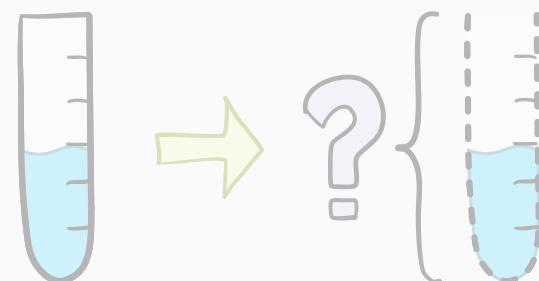


Dispatching with Gittins

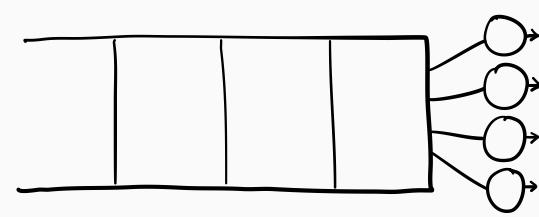


Multiserver jobs and other service constraints

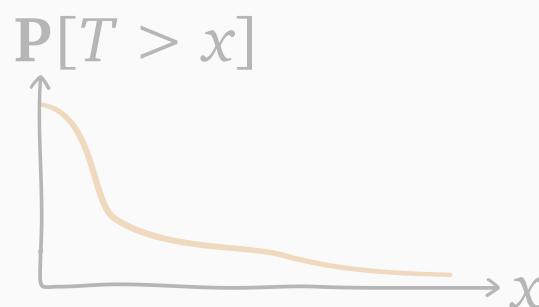
Hard scheduling questions



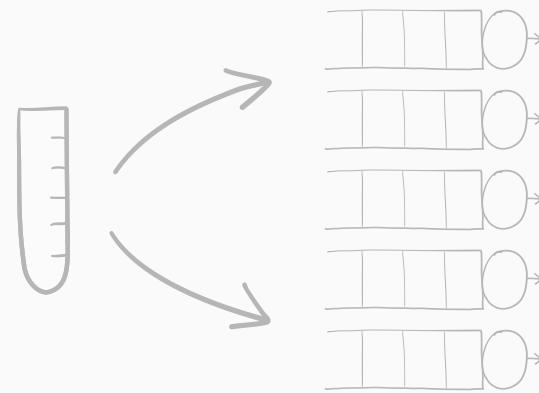
*What if we have
unknown job sizes?*



**What if there are
multiple servers?**



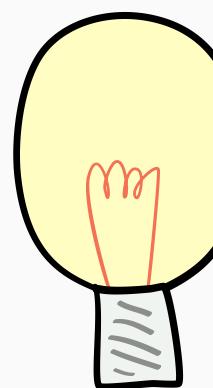
*What if we want to optimize
tails instead of means?*



*What if we can only use
dispatching without
fancy scheduling?*

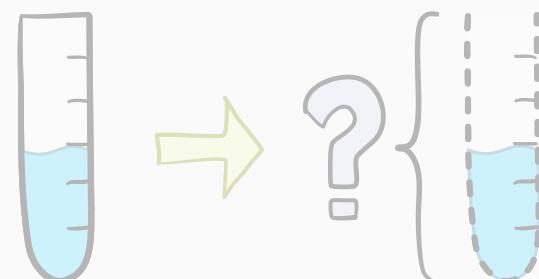


**What are the underlying
theoretical tools?**

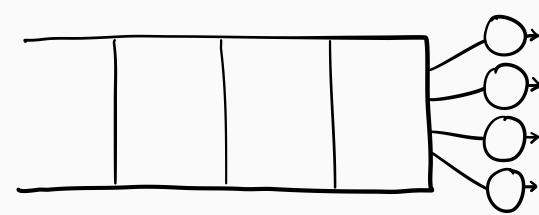


**What are the resulting
practical lessons?**

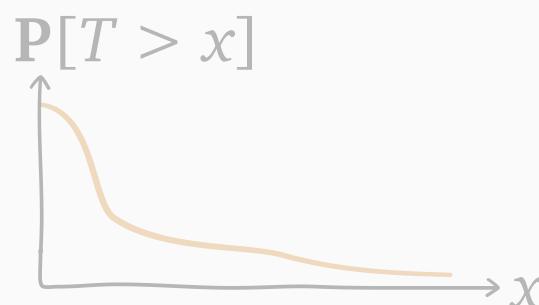
Hard scheduling questions



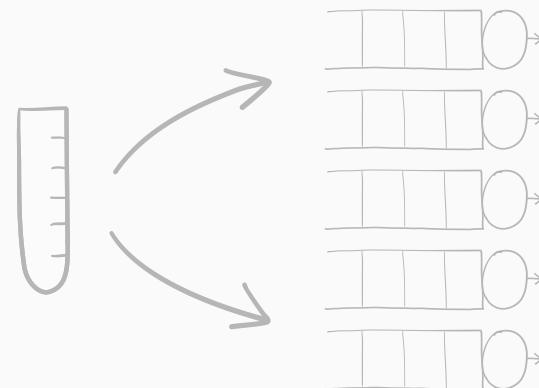
*What if we have
unknown job sizes?*



**What if there are
multiple servers?**



*What if we want to optimize
tails instead of means?*

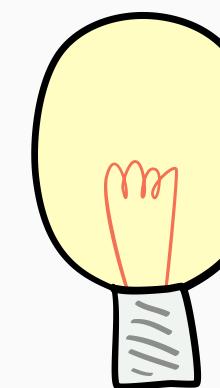


*What if we can only use
dispatching without
fancy scheduling?*



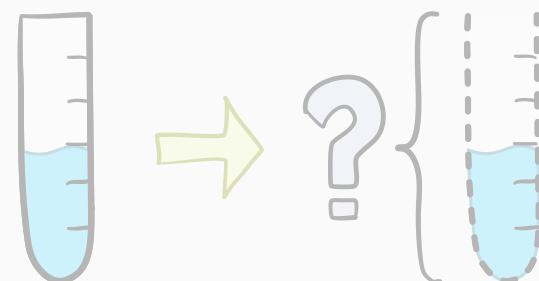
**WINE and work
decomposition**

**What are the underlying
theoretical tools?**

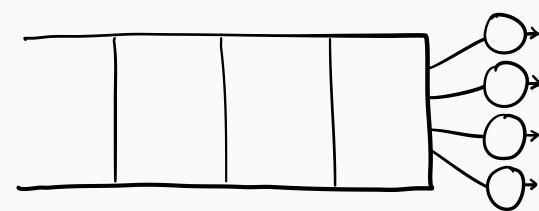


**What are the resulting
practical lessons?**

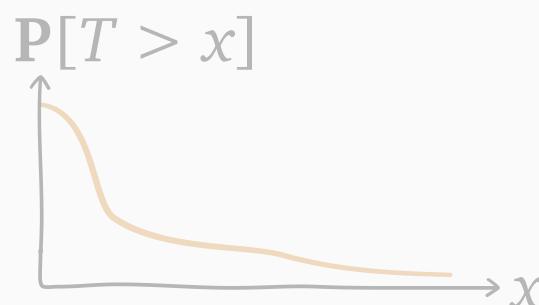
Hard scheduling questions



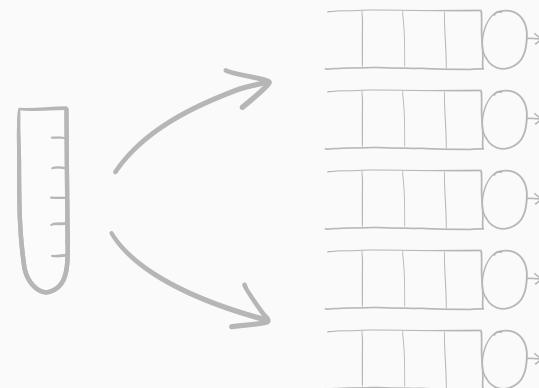
What if we have unknown job sizes?



What if there are multiple servers?



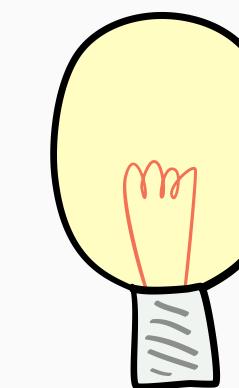
What if we want to optimize tails instead of means?



What if we can only use dispatching without fancy scheduling?



What are the underlying theoretical tools?

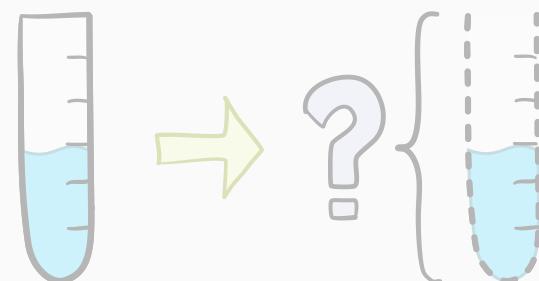


What are the resulting practical lessons?

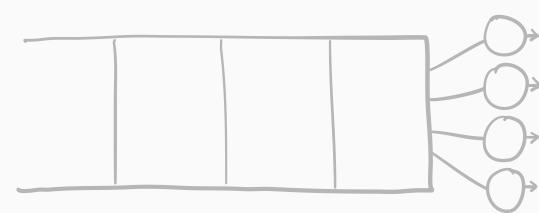
- Prefer central queueing
- When dispatching, spread out the small jobs

WINE and work decomposition

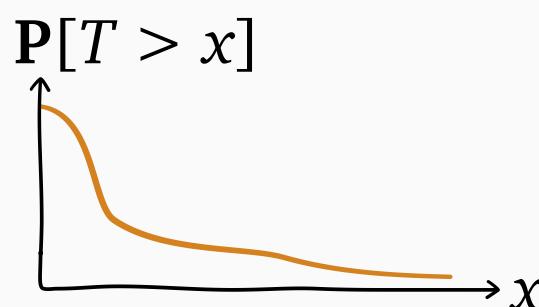
Hard scheduling questions



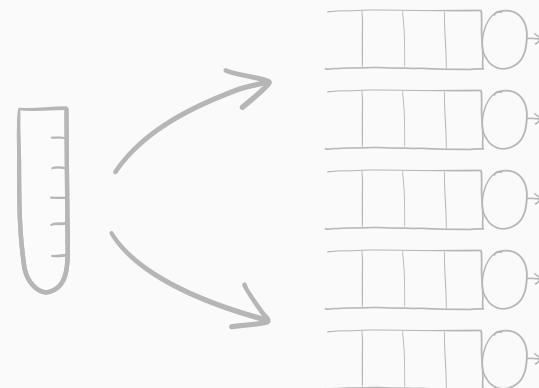
What if we have unknown job sizes?



What if there are multiple servers?



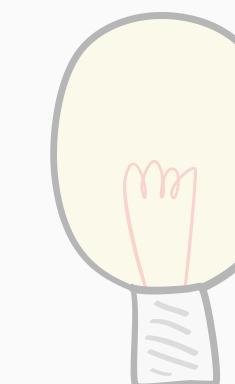
What if we want to optimize tails instead of means?



What if we can only use dispatching without fancy scheduling?



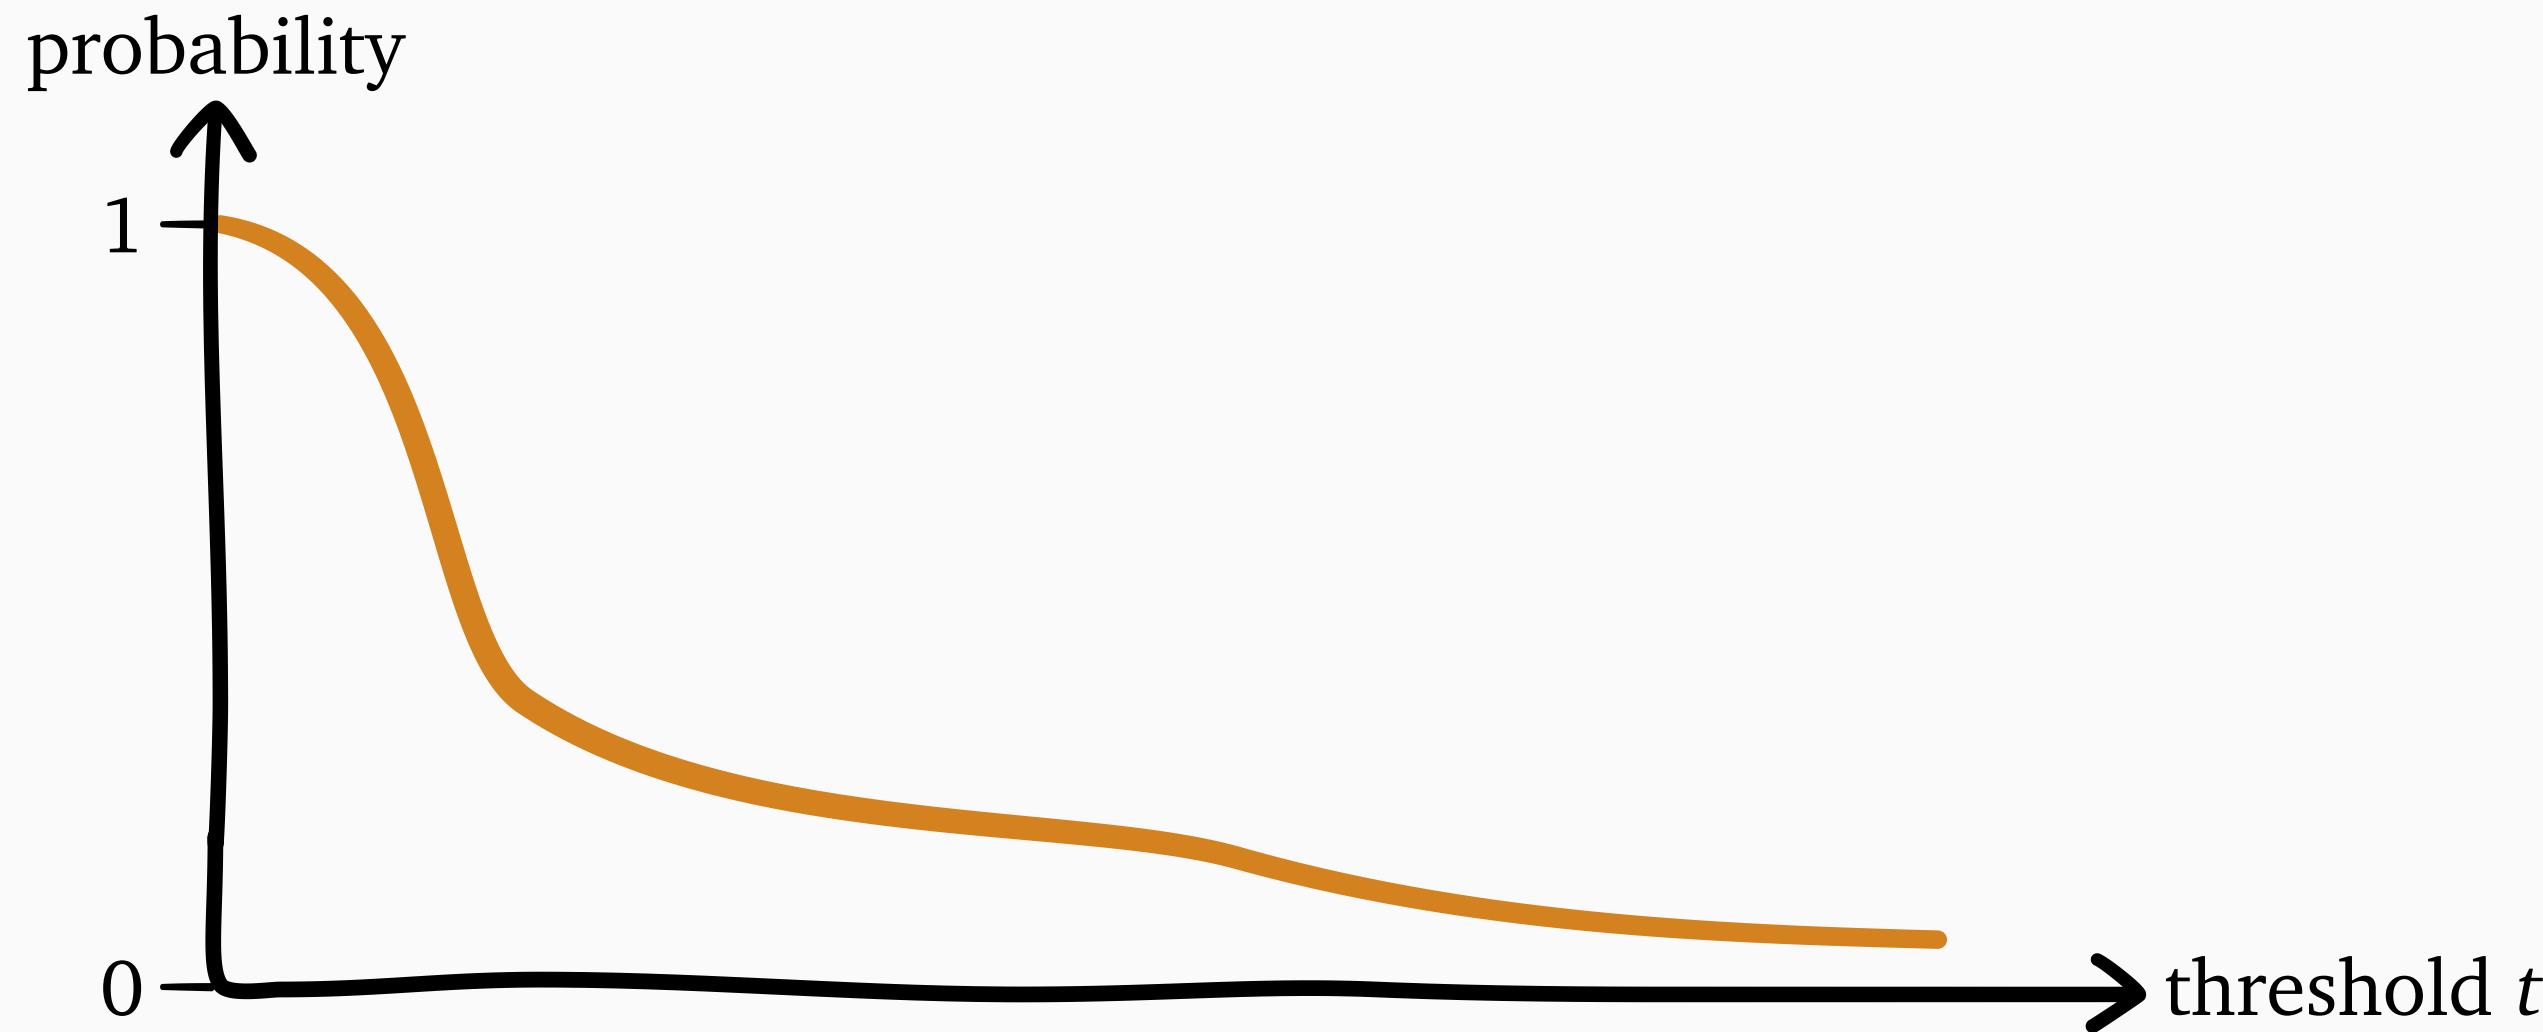
What are the underlying theoretical tools?



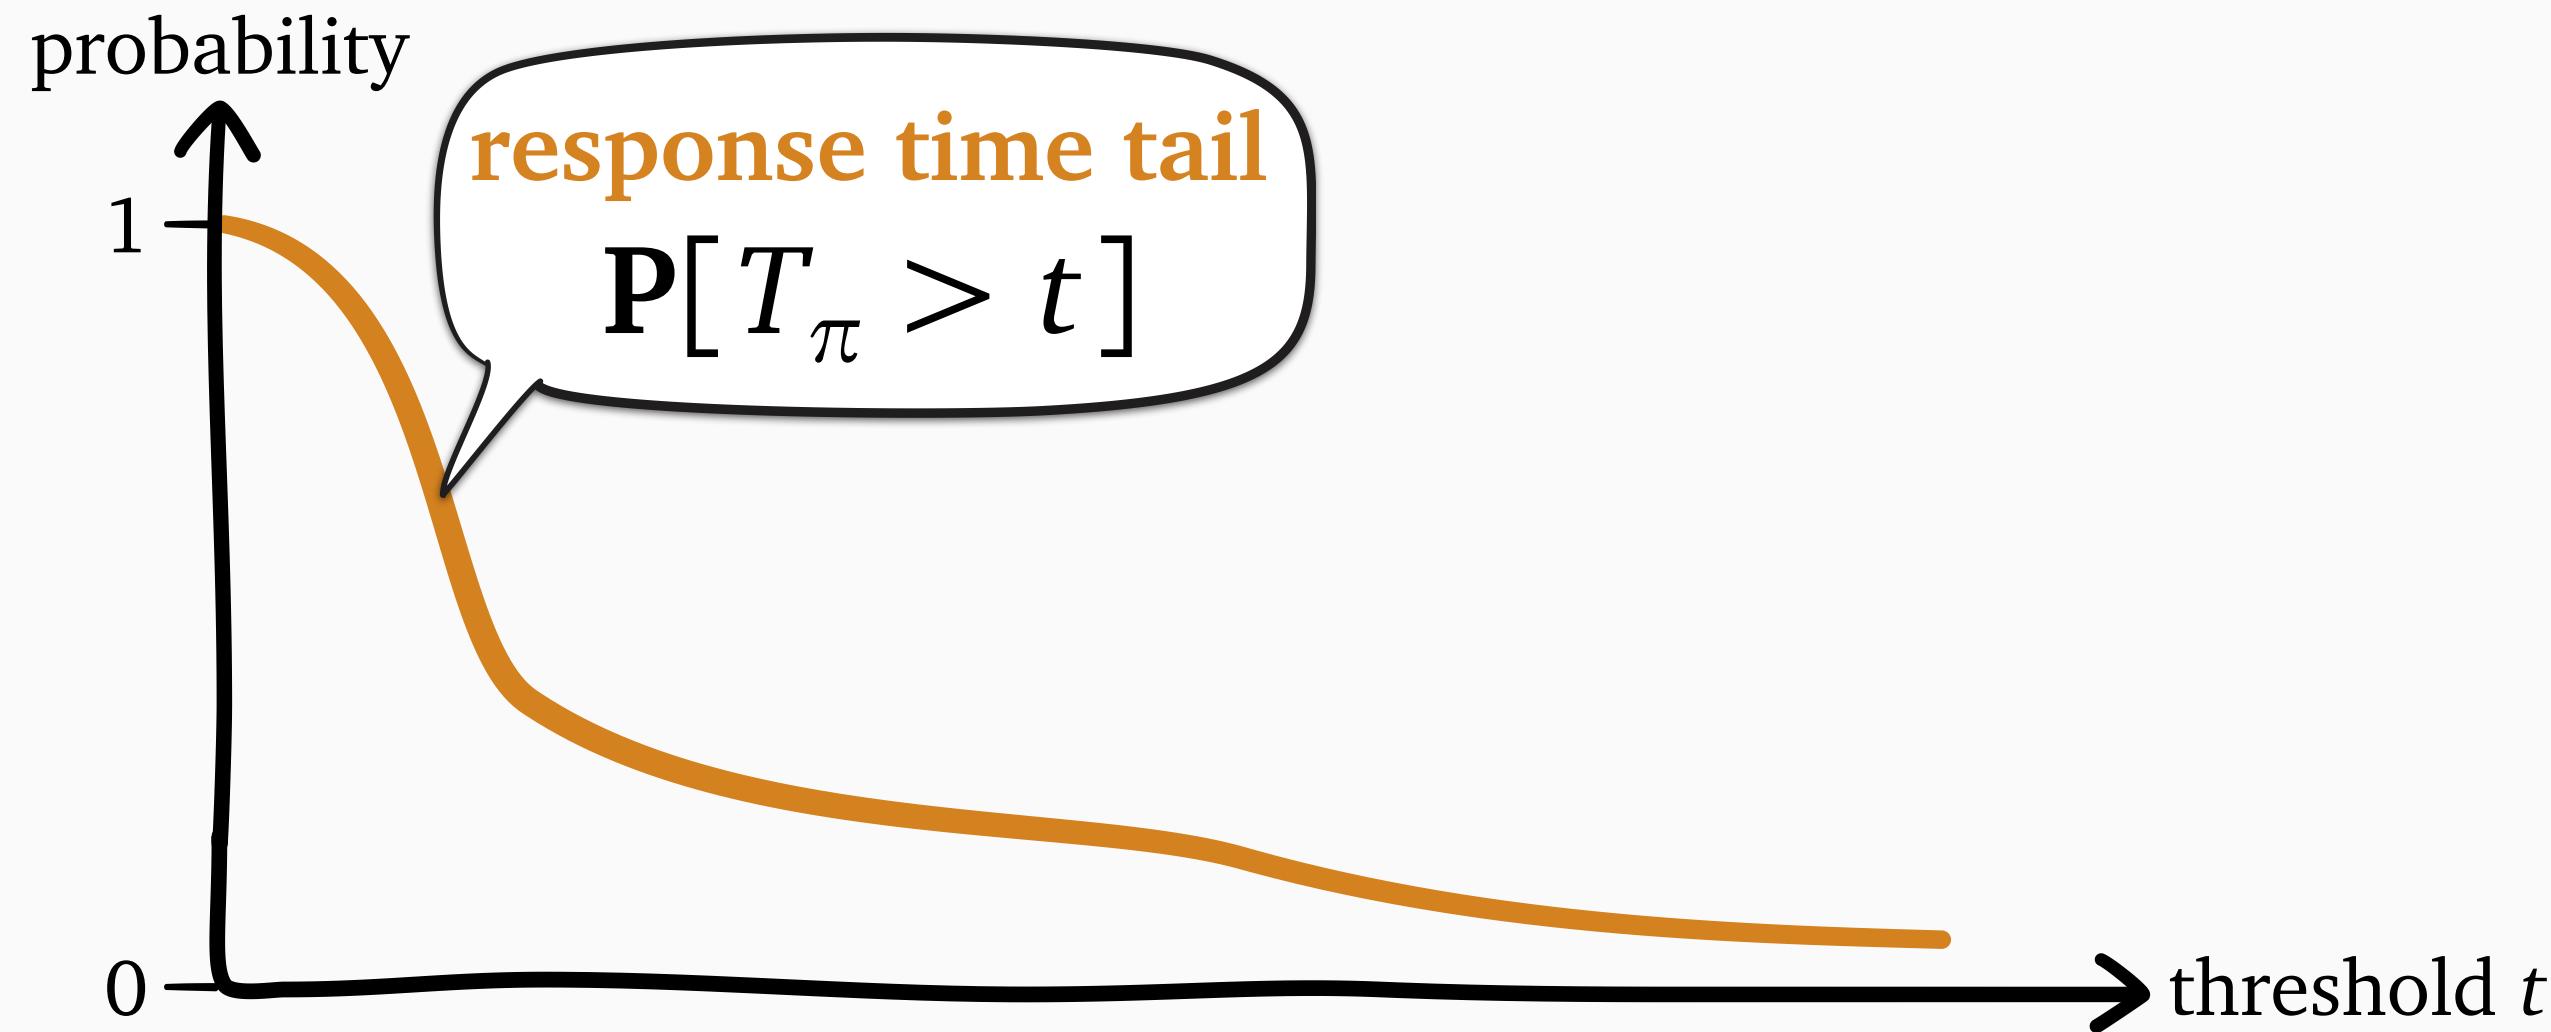
What are the resulting practical lessons?

- Prefer central queueing
- When dispatching, spread out the small jobs

Asymptotic response time tail

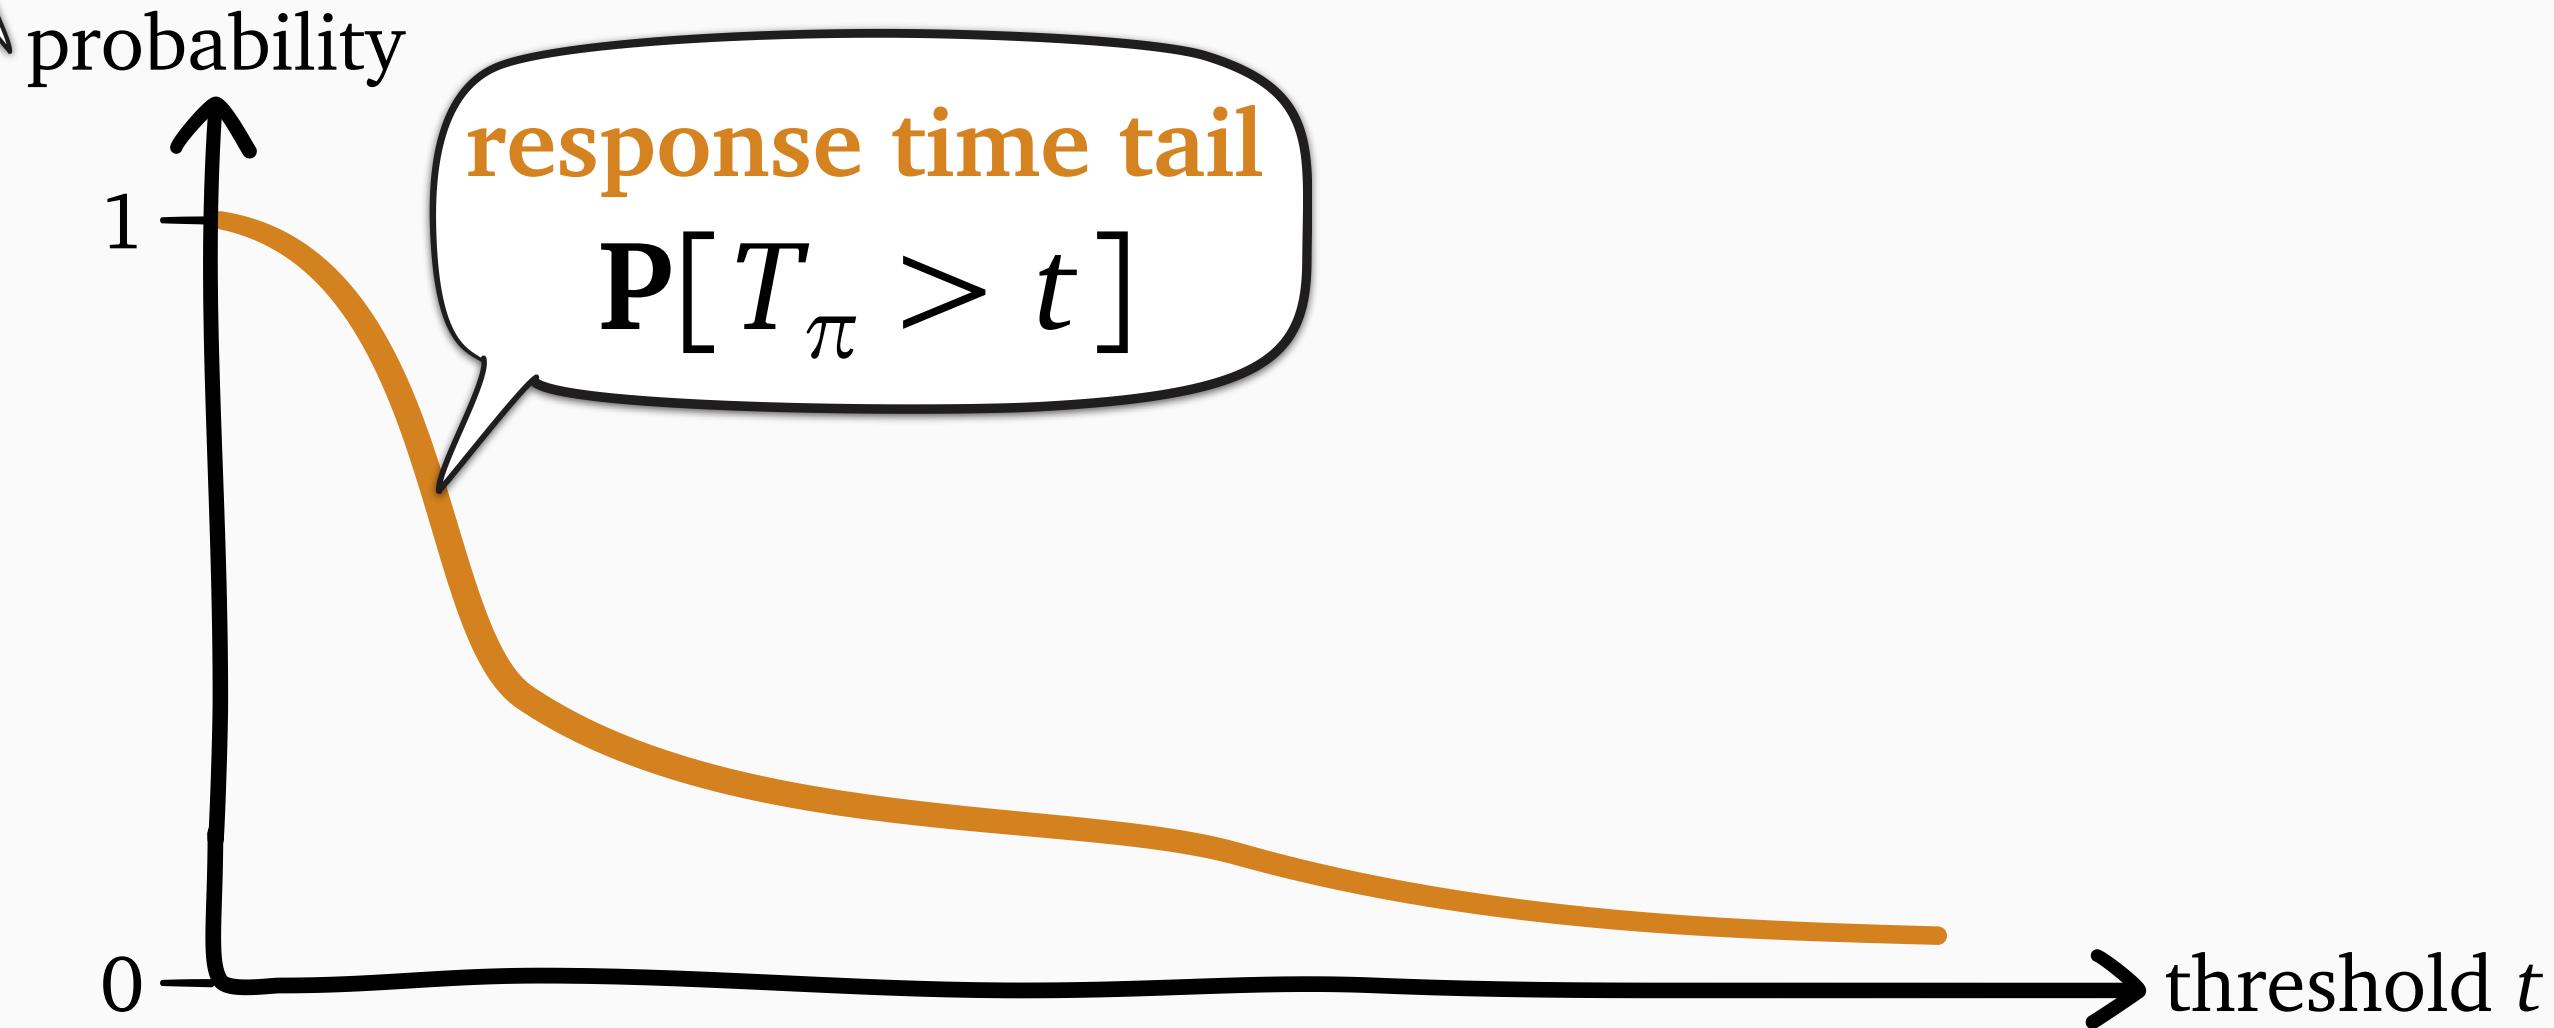


Asymptotic response time tail



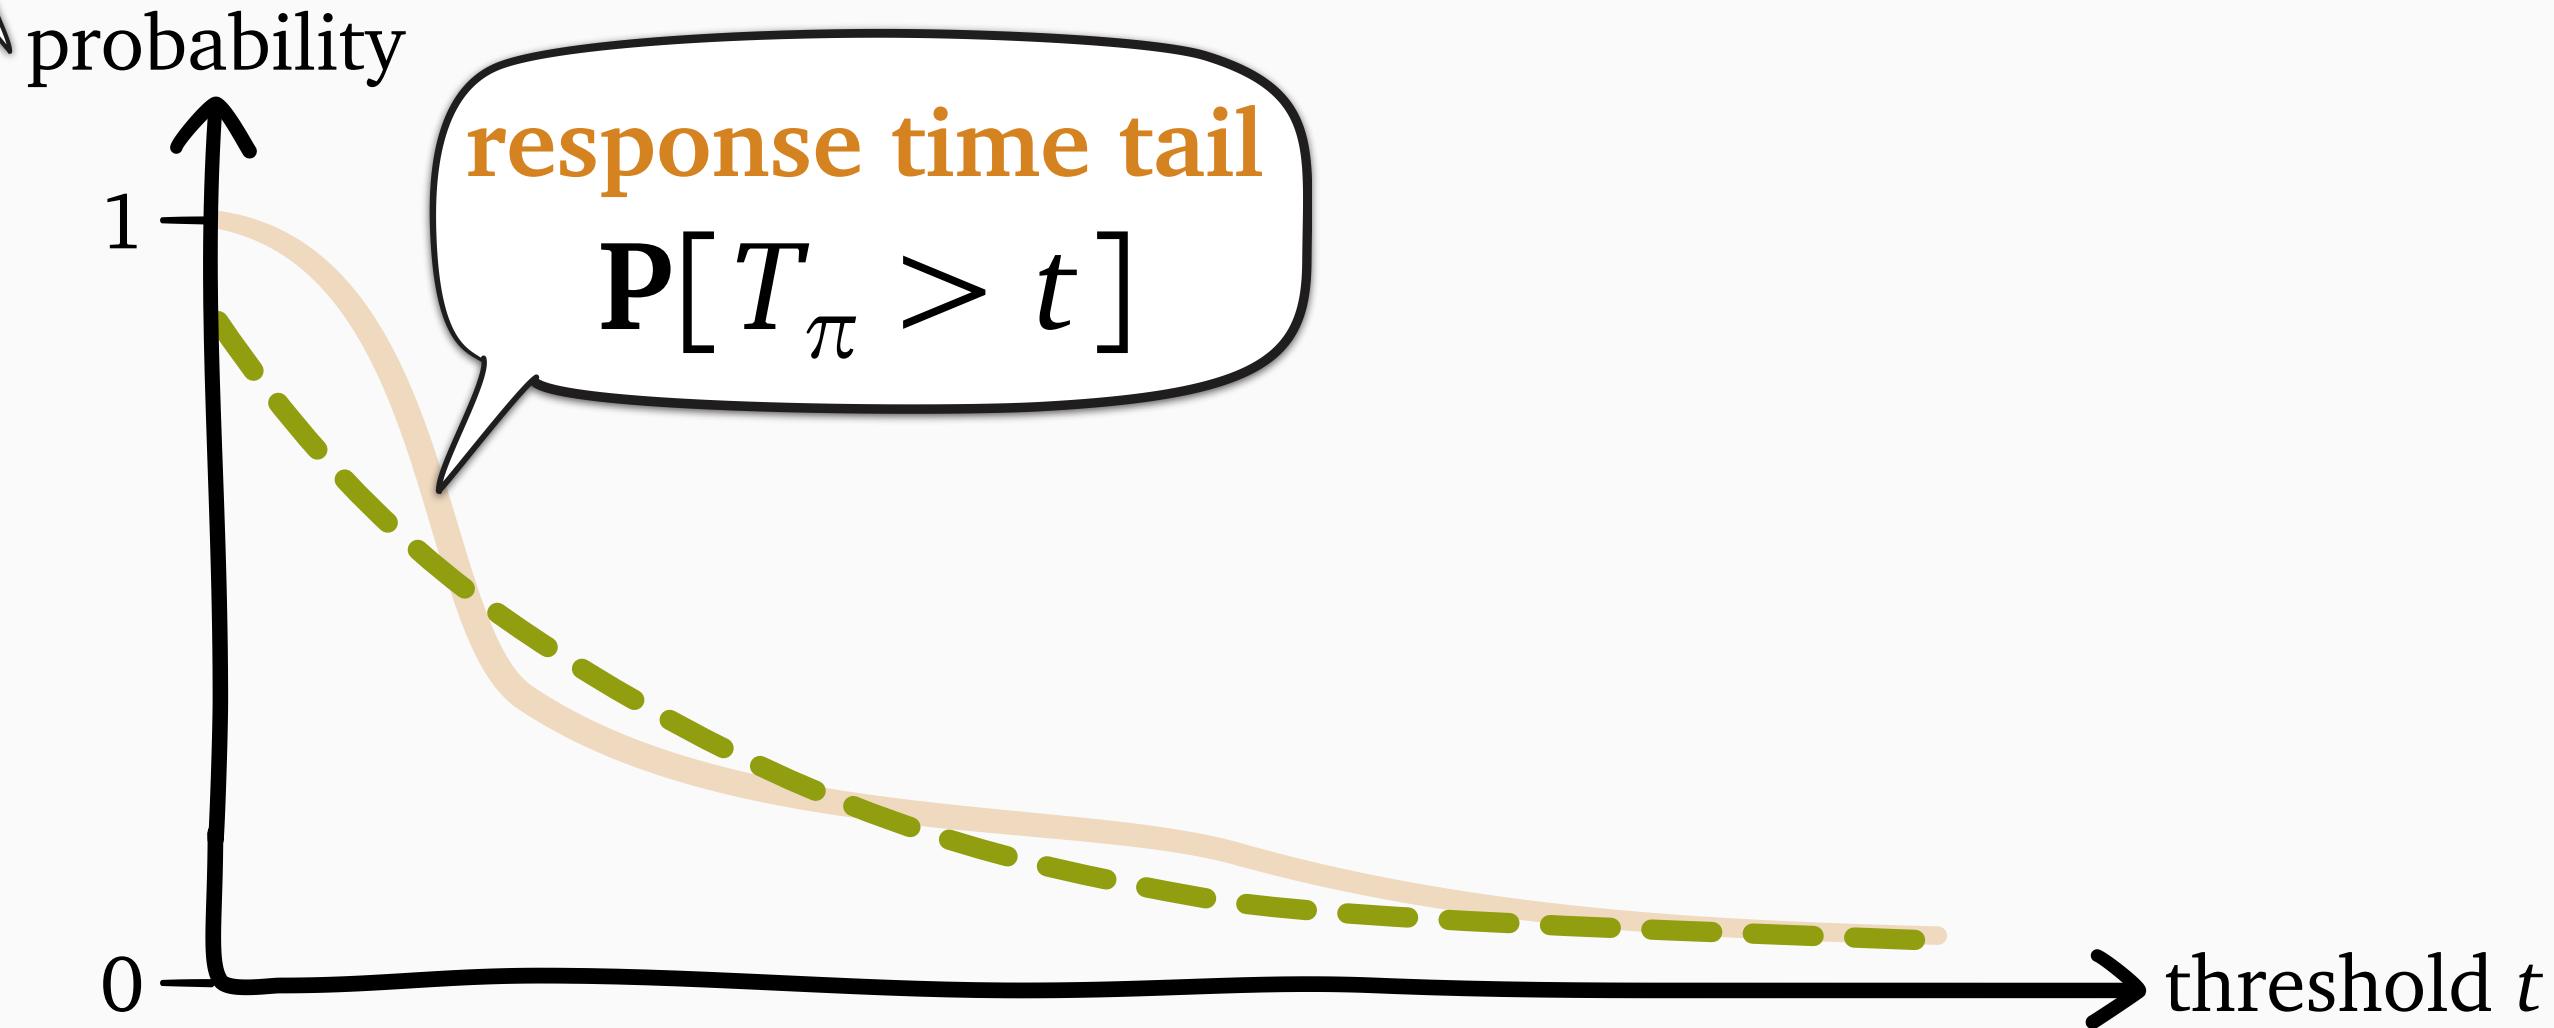
depends on
policy π

Asymptotic response time tail



depends on
policy π

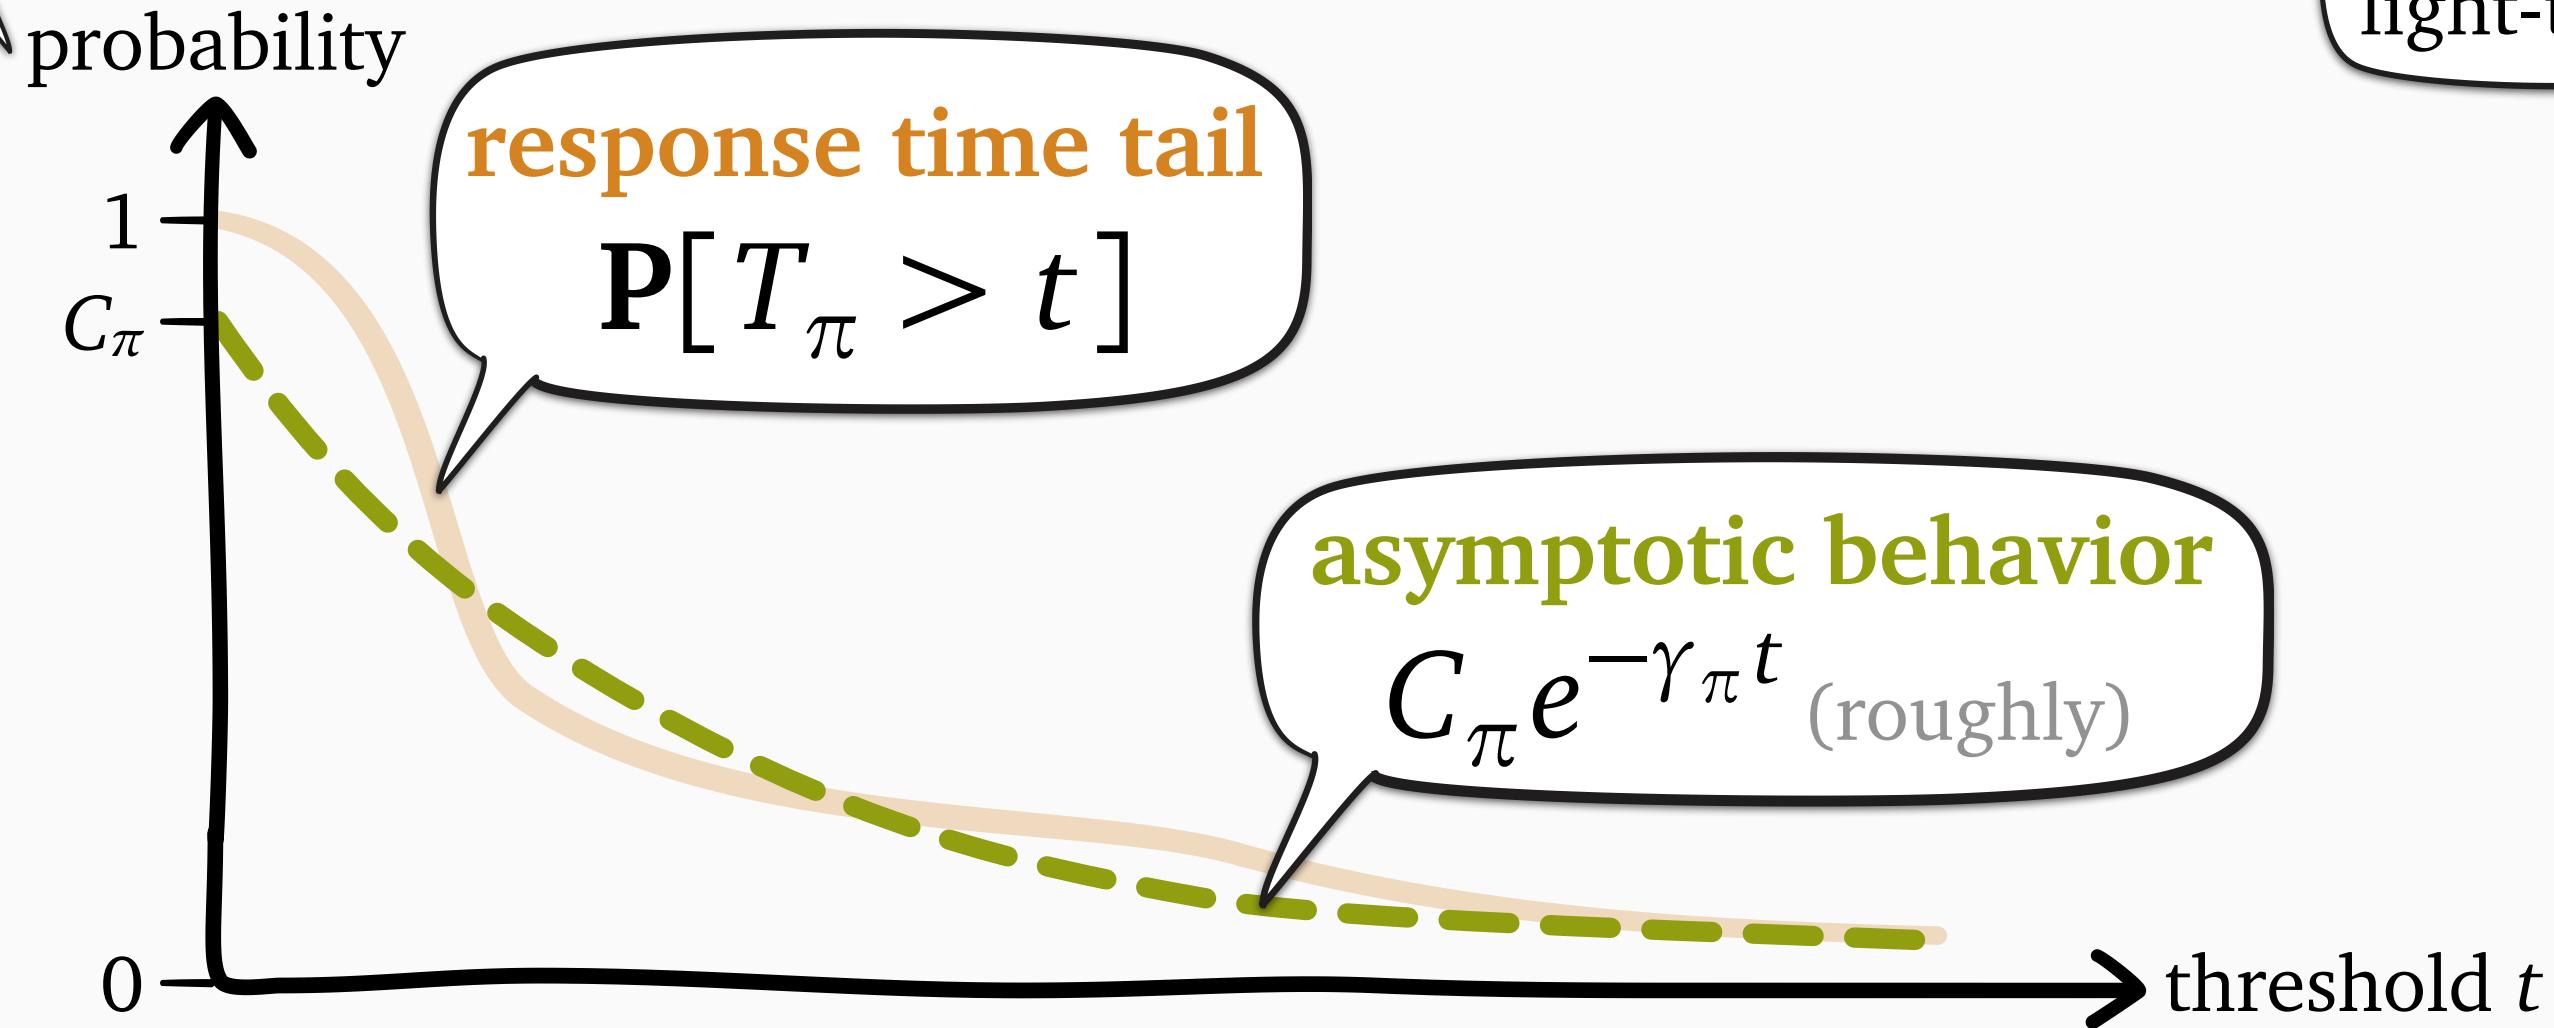
Asymptotic response time tail



depends on
policy π

Asymptotic response time tail

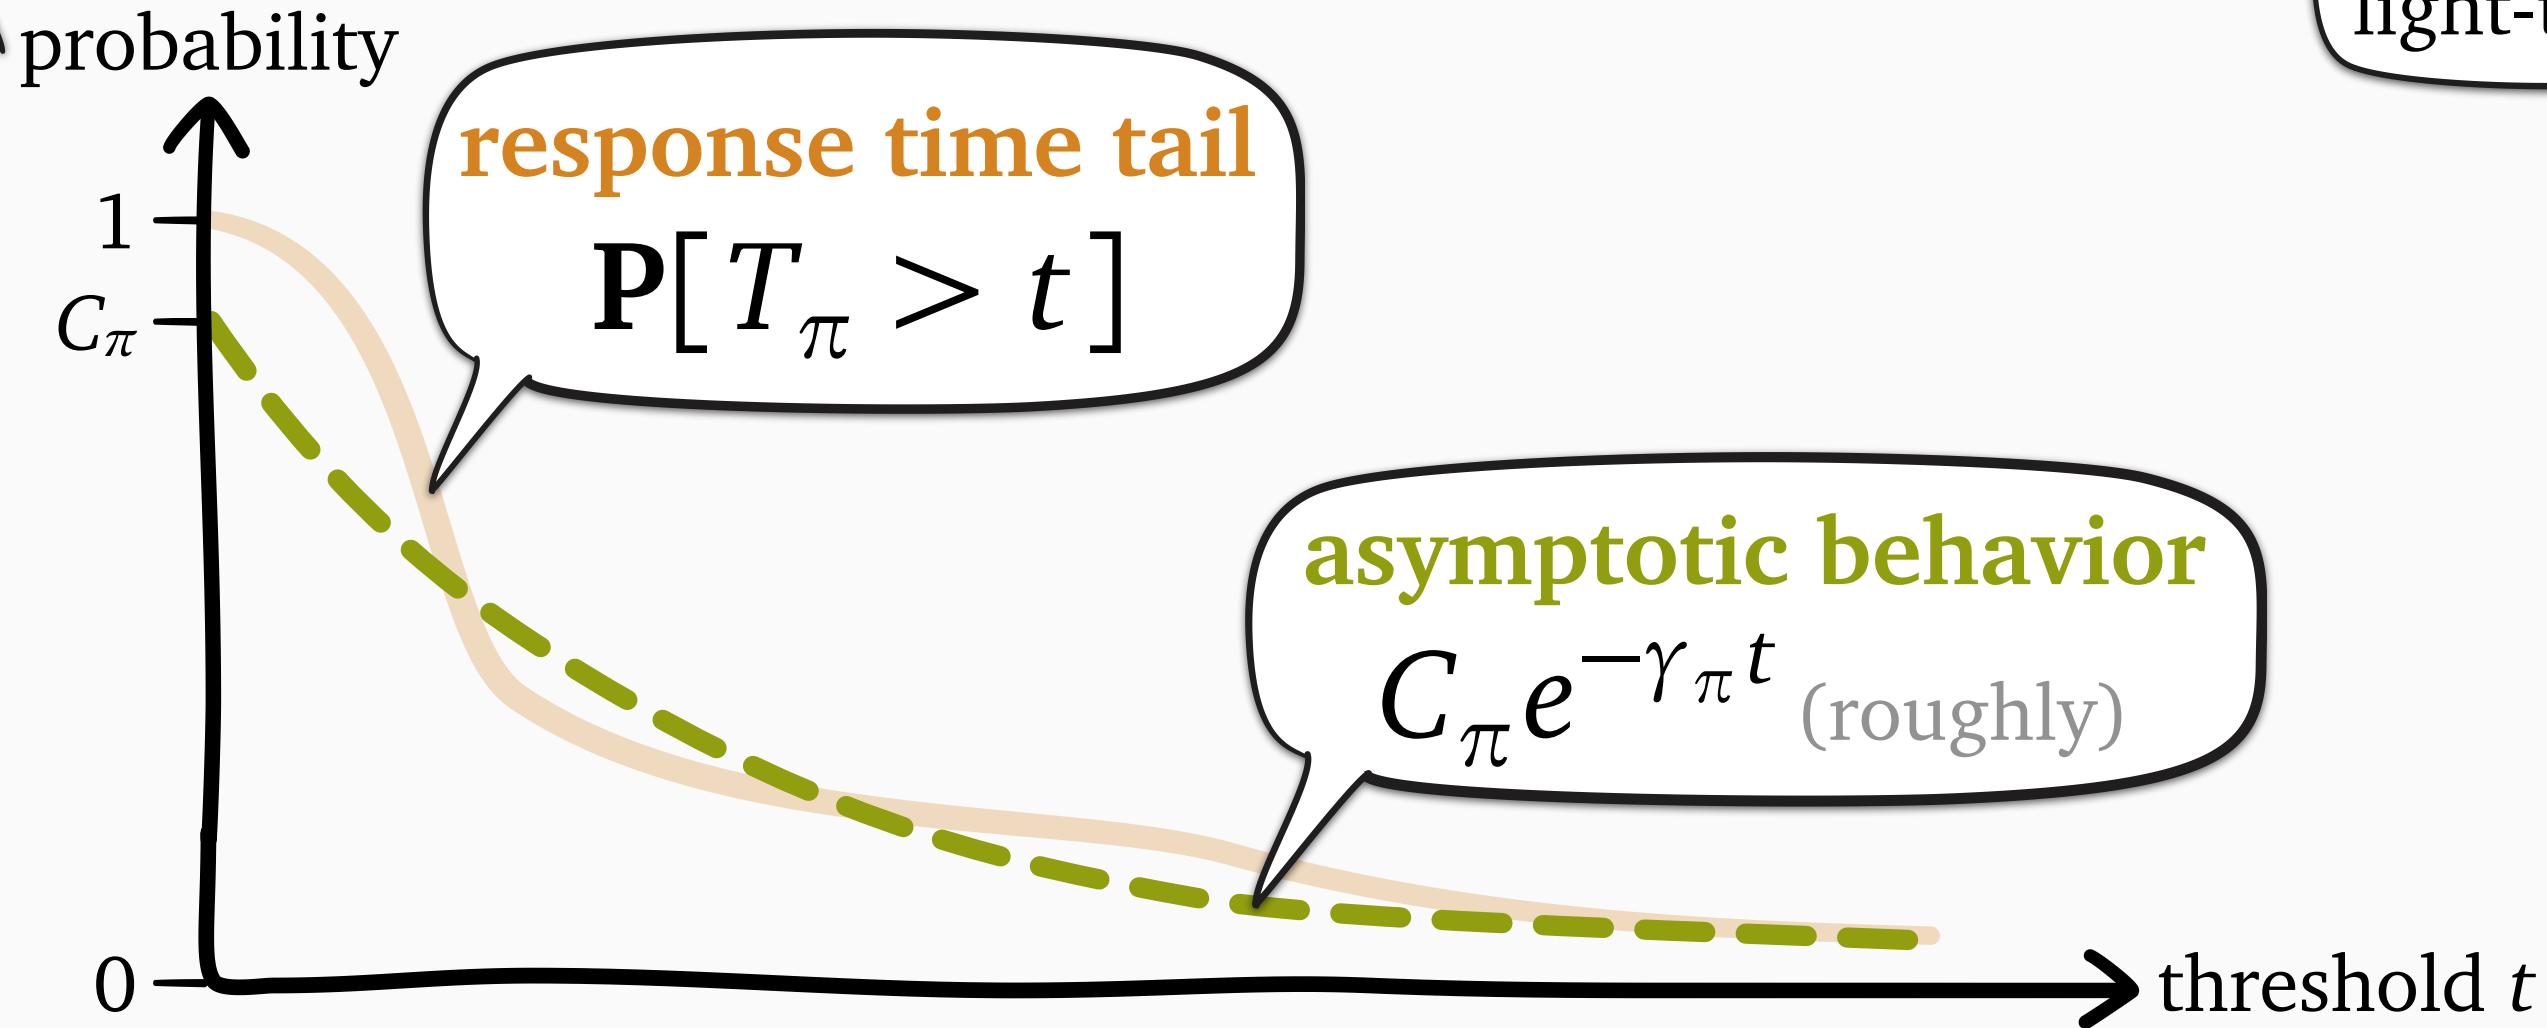
when S is
light-tailed



depends on
policy π

Asymptotic response time tail

when S is
light-tailed



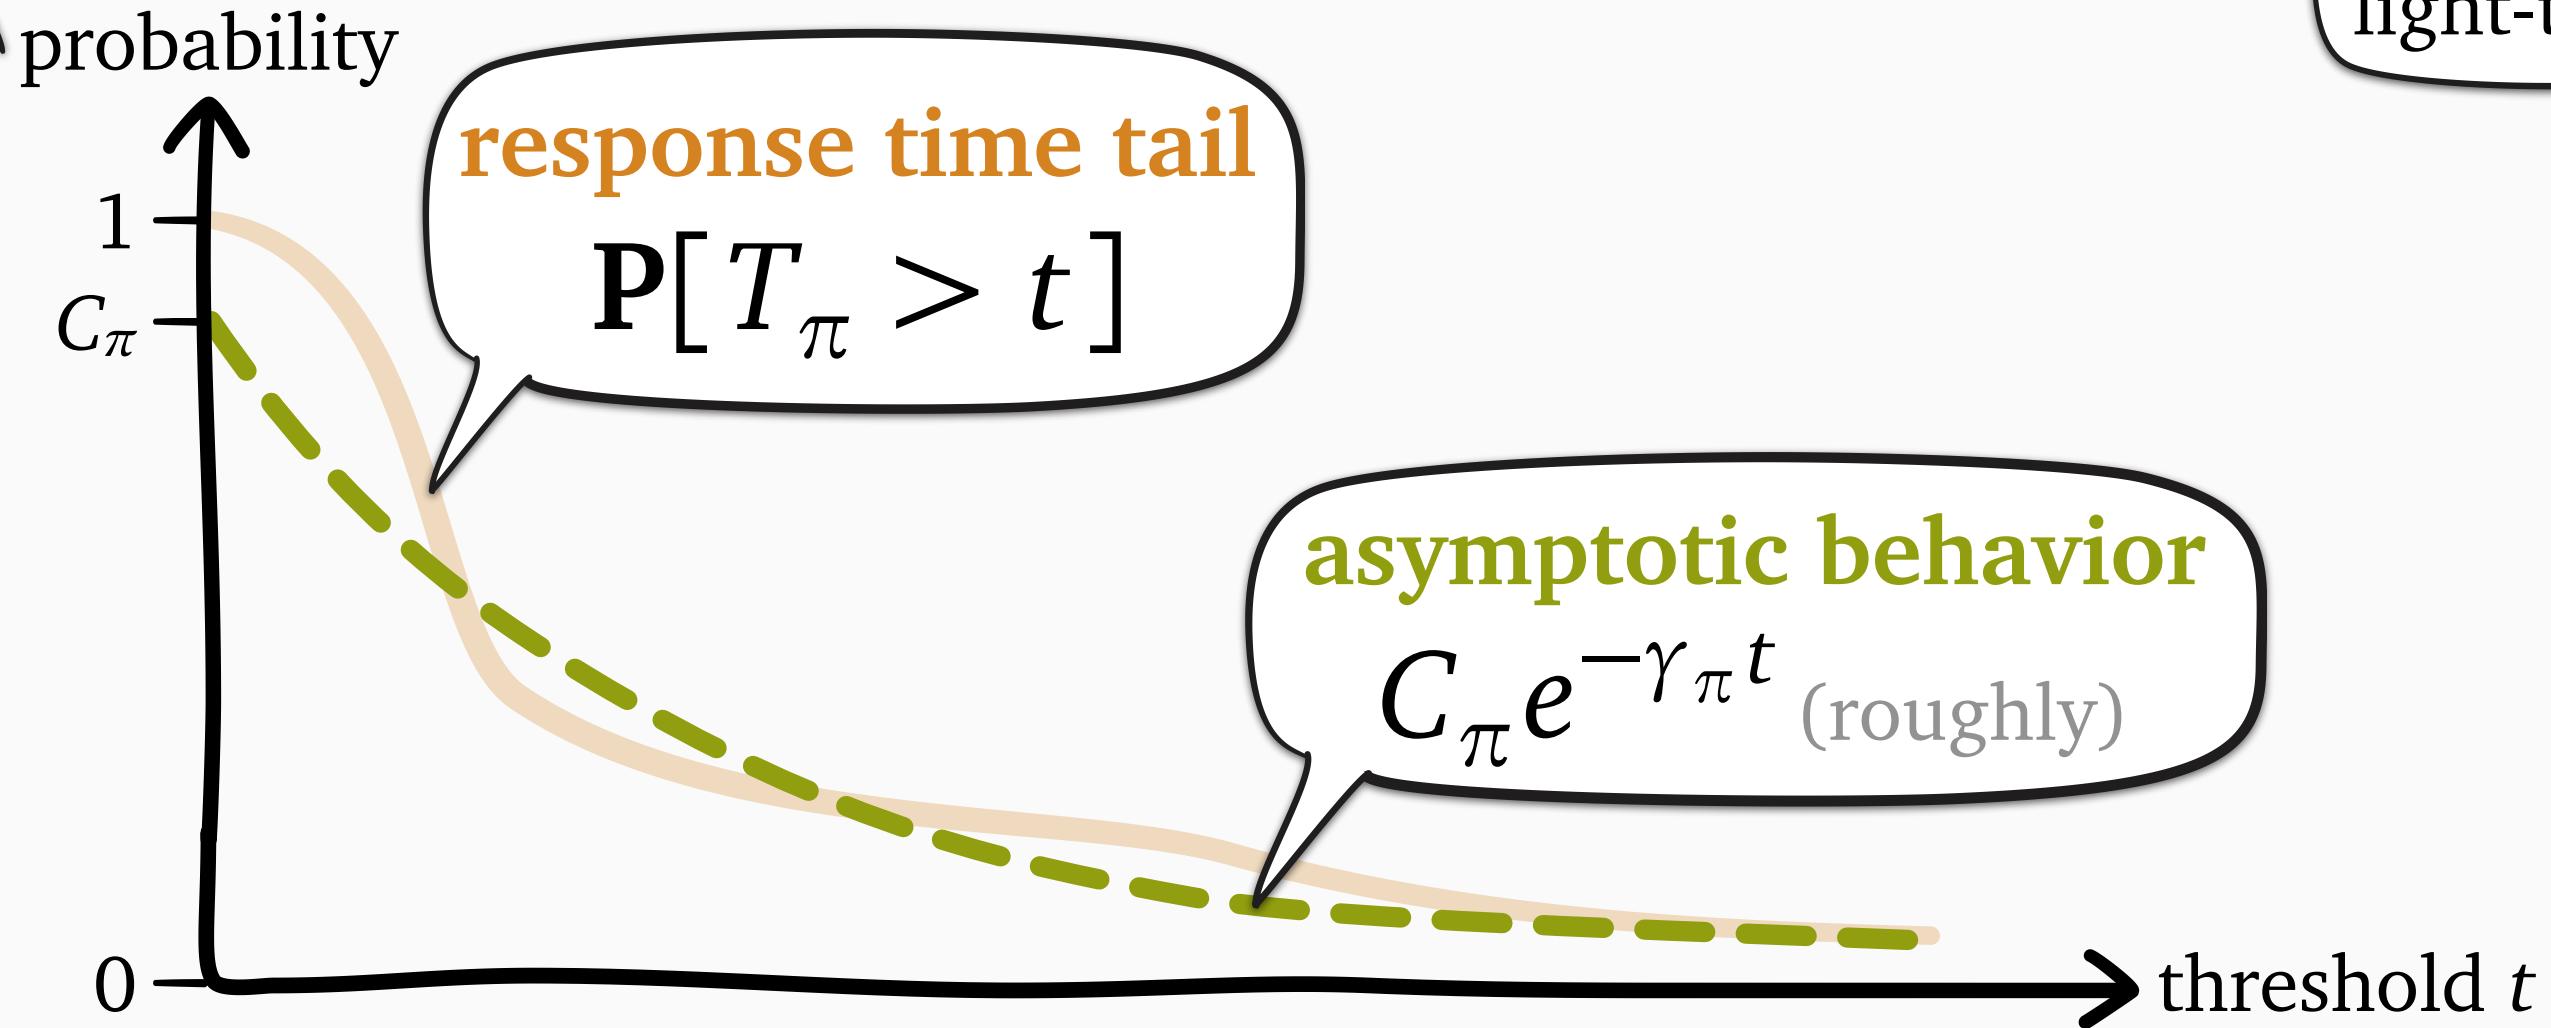
$\gamma_\pi = \text{decay rate of } \pi$

$C_\pi = \text{tail constant of } \pi$

depends on
policy π

Asymptotic response time tail

when S is
light-tailed



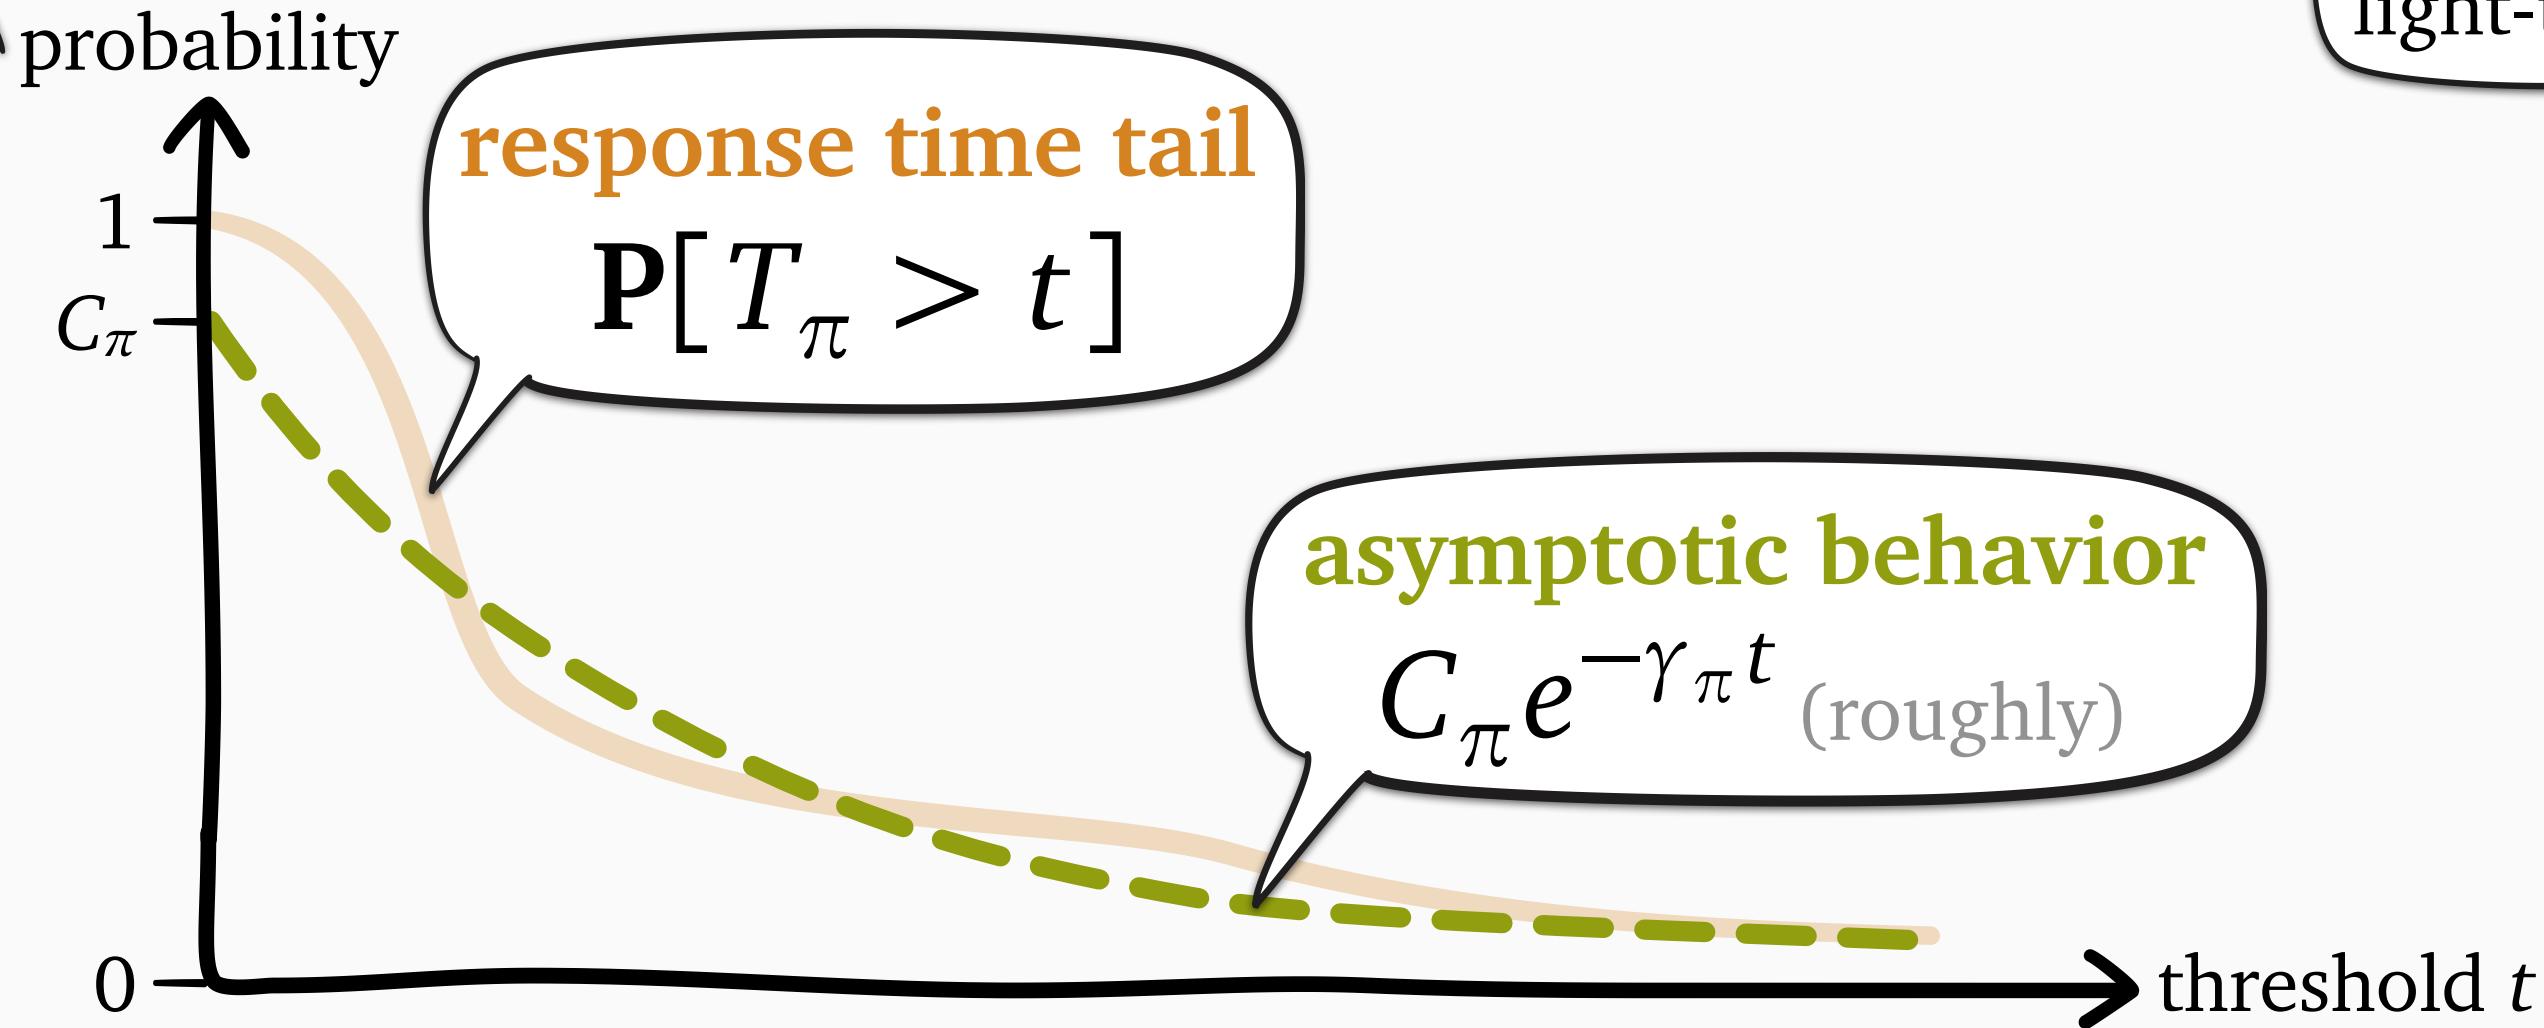
Weak optimality:
optimal γ_π

$\gamma_\pi = \text{decay rate of } \pi$
 $C_\pi = \text{tail constant of } \pi$

depends on
policy π

Asymptotic response time tail

when S is
light-tailed



Weak optimality:
optimal γ_π

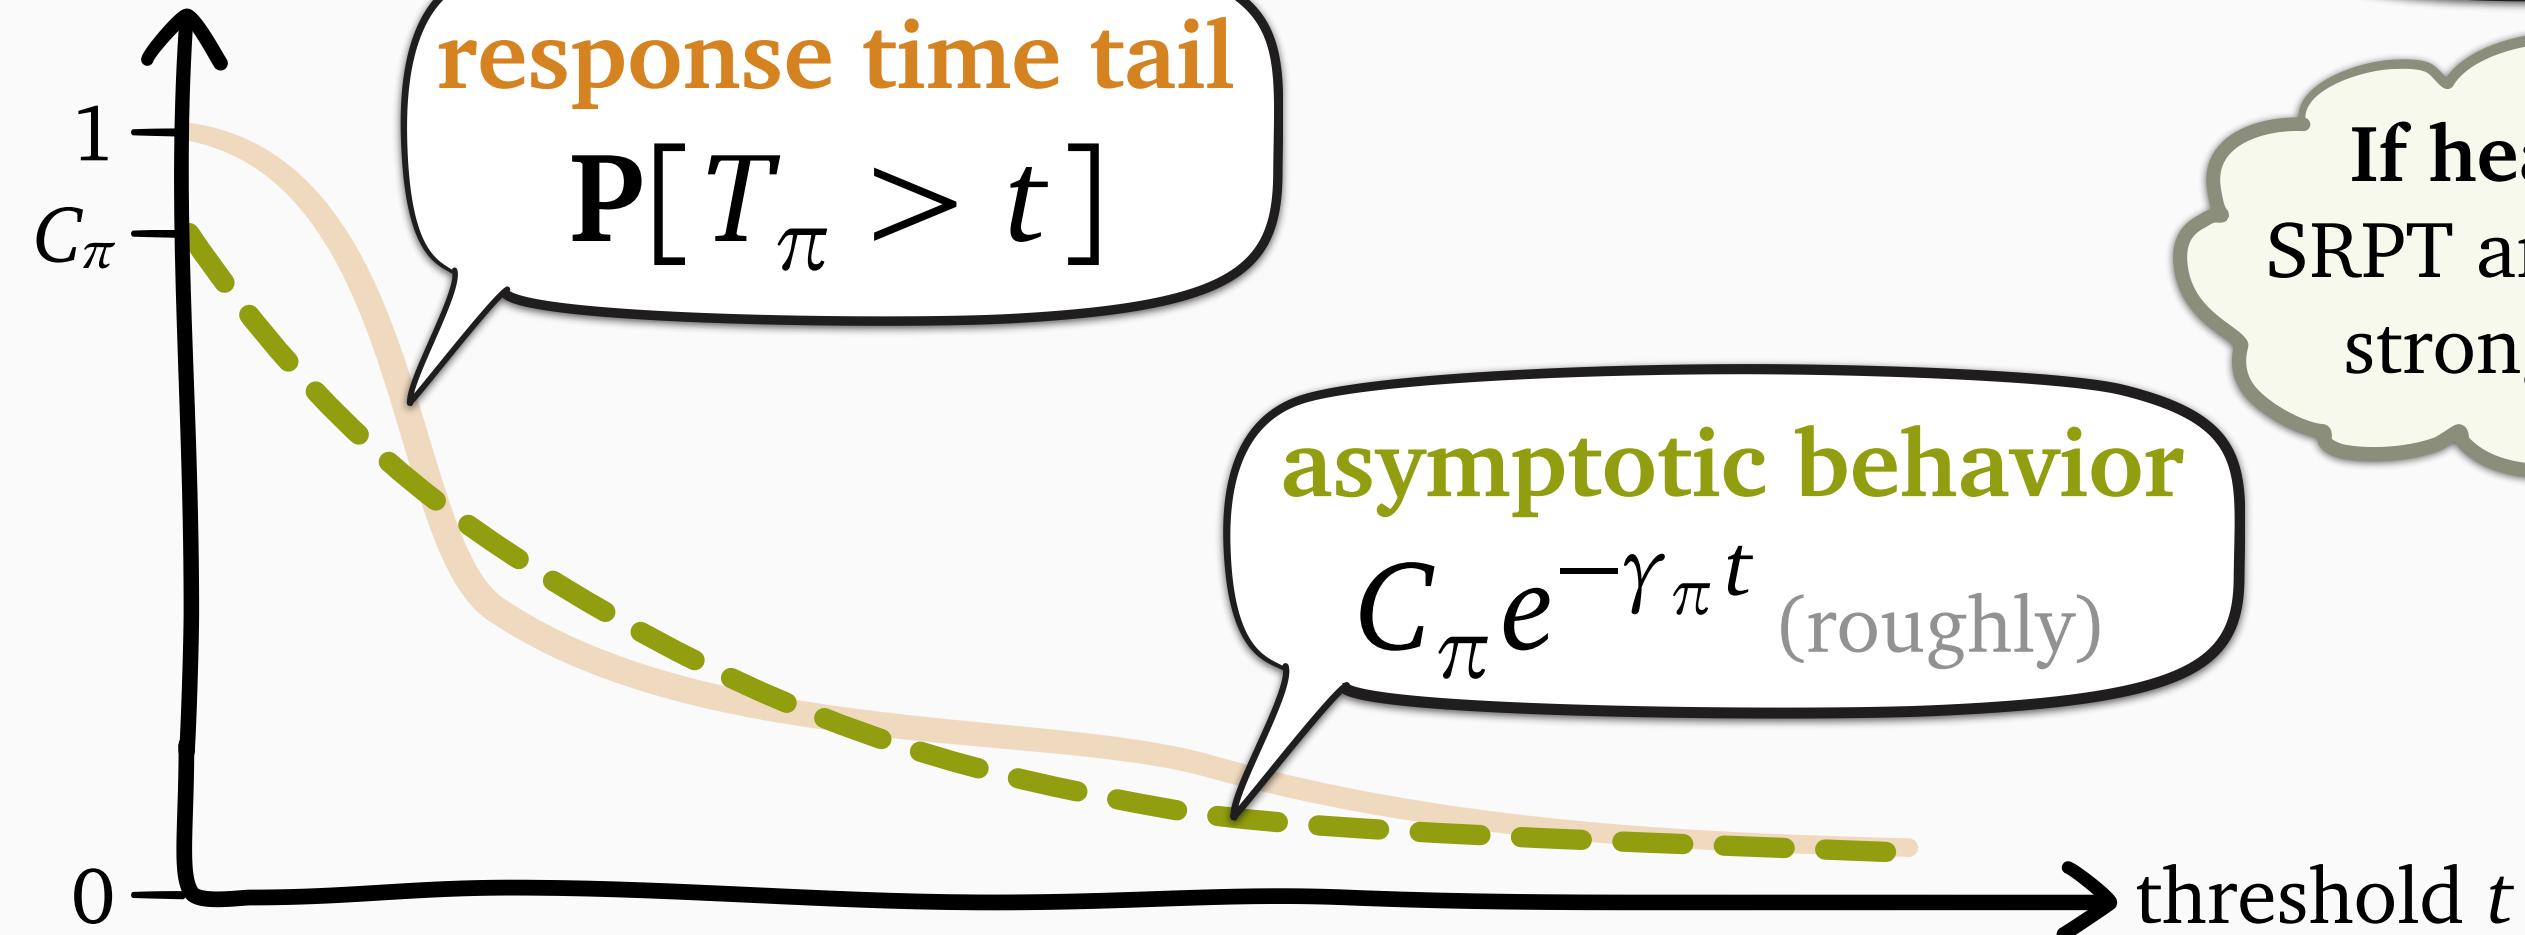
$\gamma_\pi = \text{decay rate of } \pi$
 $C_\pi = \text{tail constant of } \pi$

Strong optimality:
optimal γ_π and C_π

Asymptotic response time tail

depends on policy π

probability

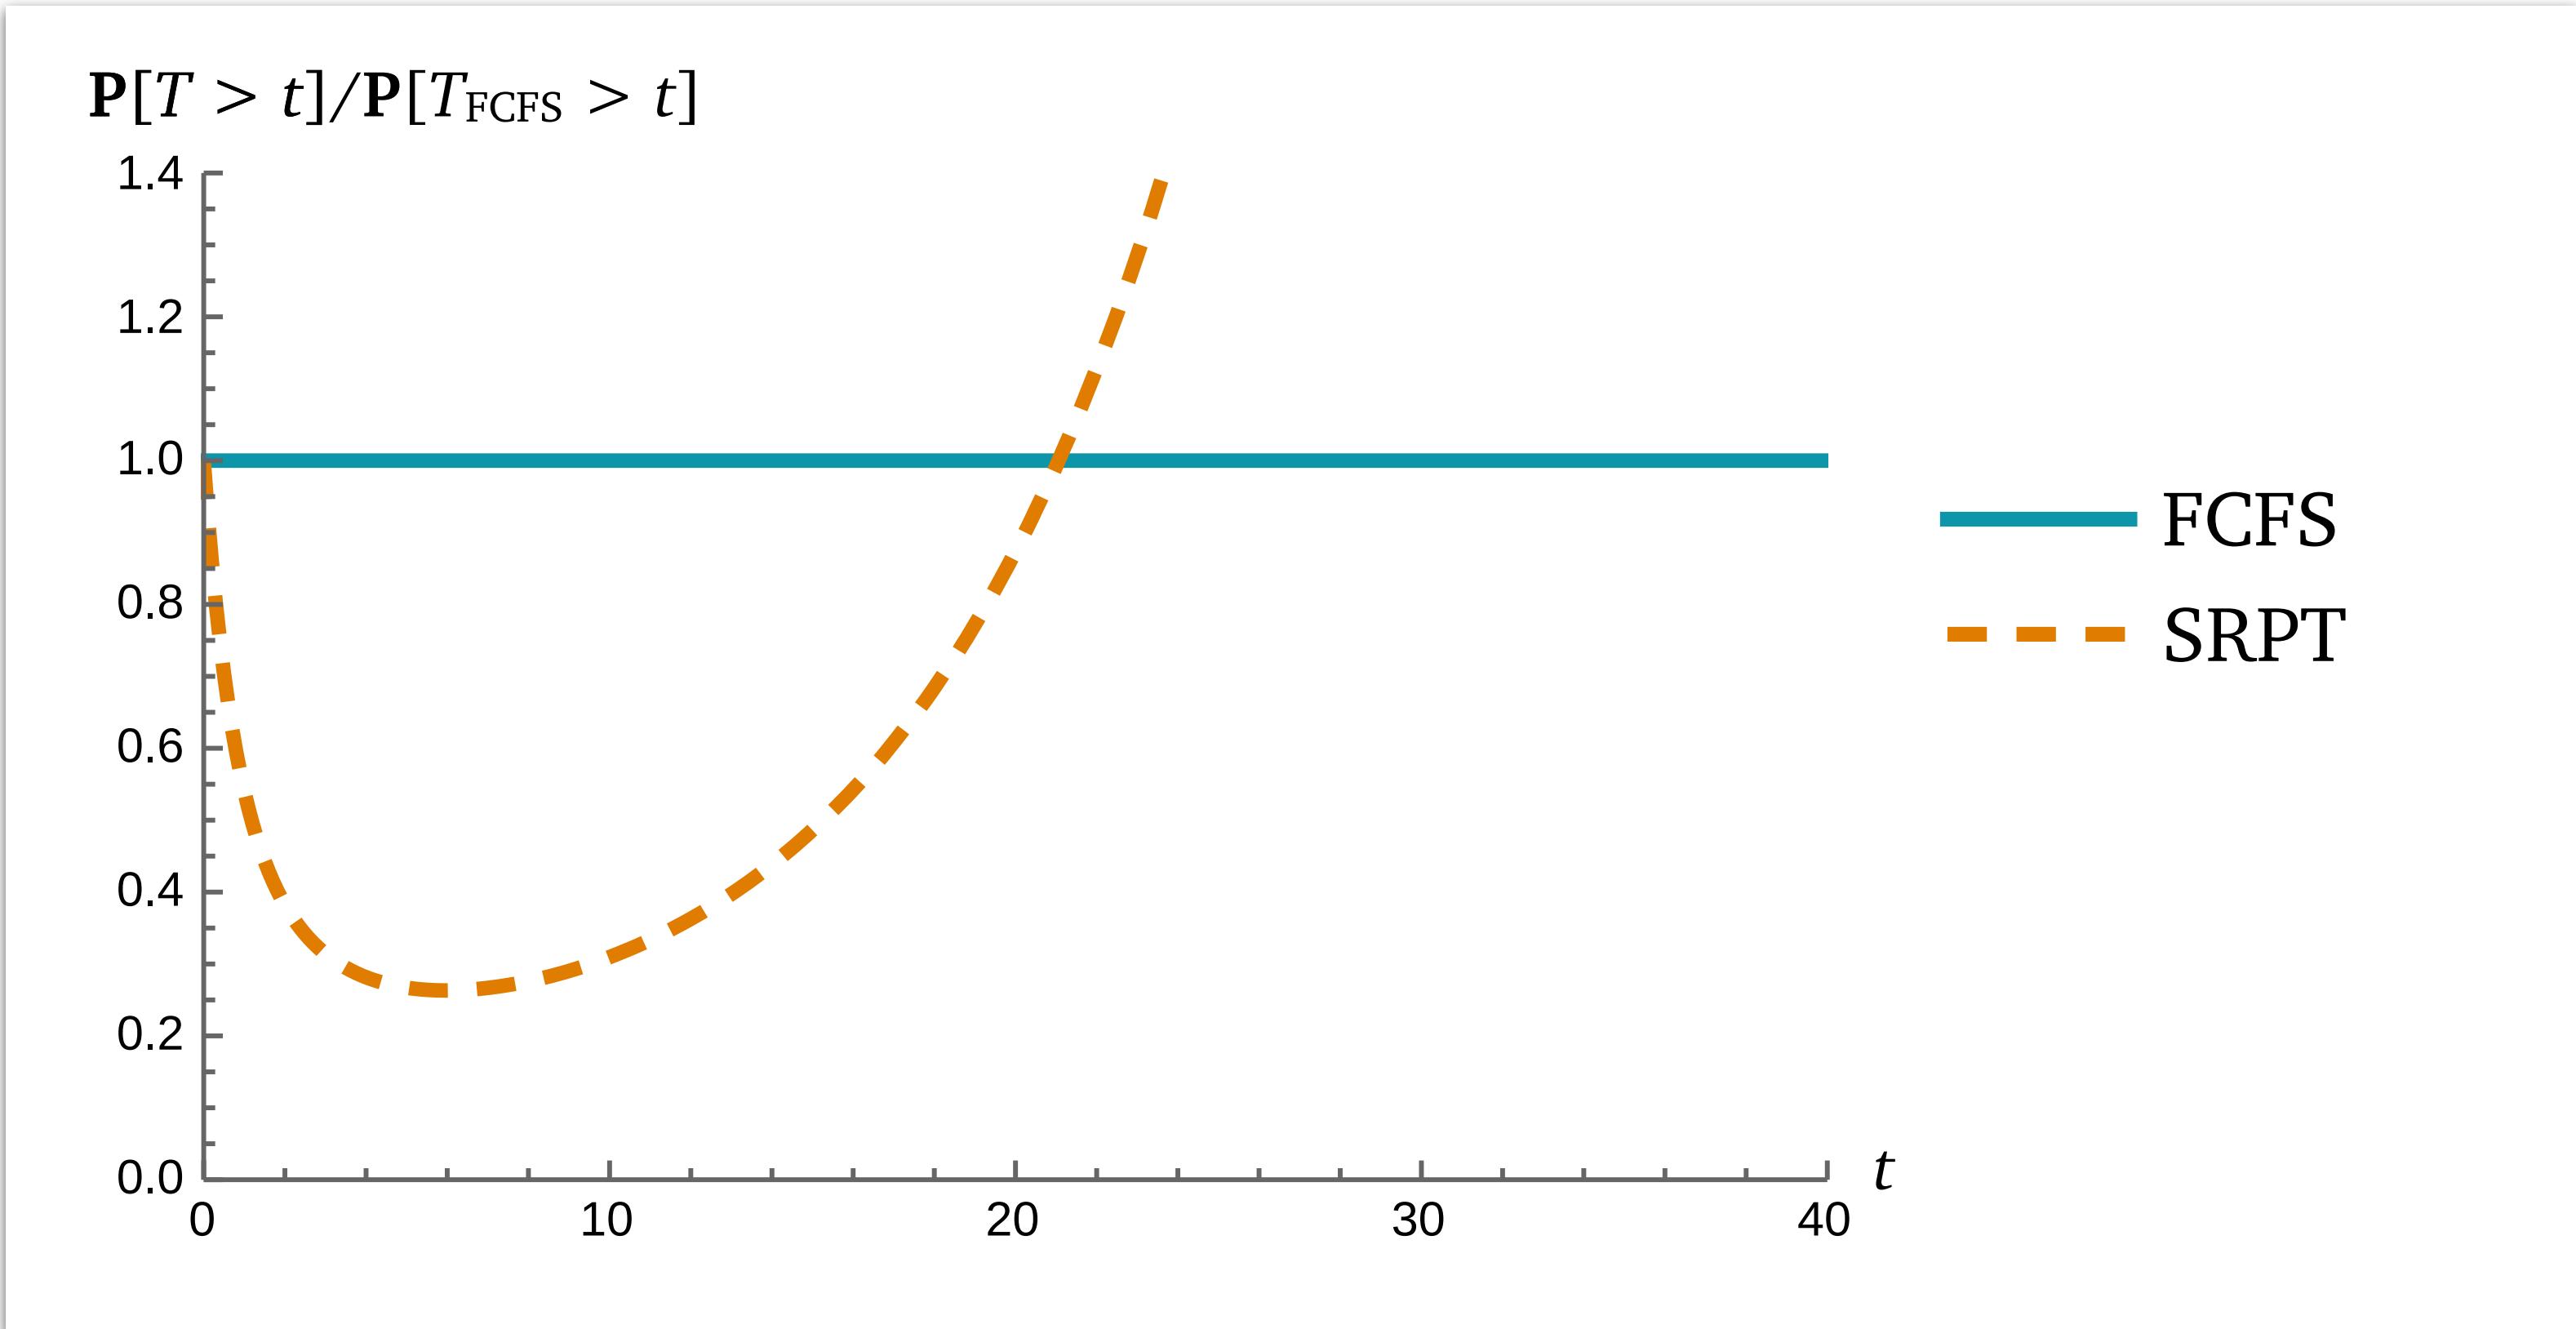


Weak optimality:
optimal γ_π

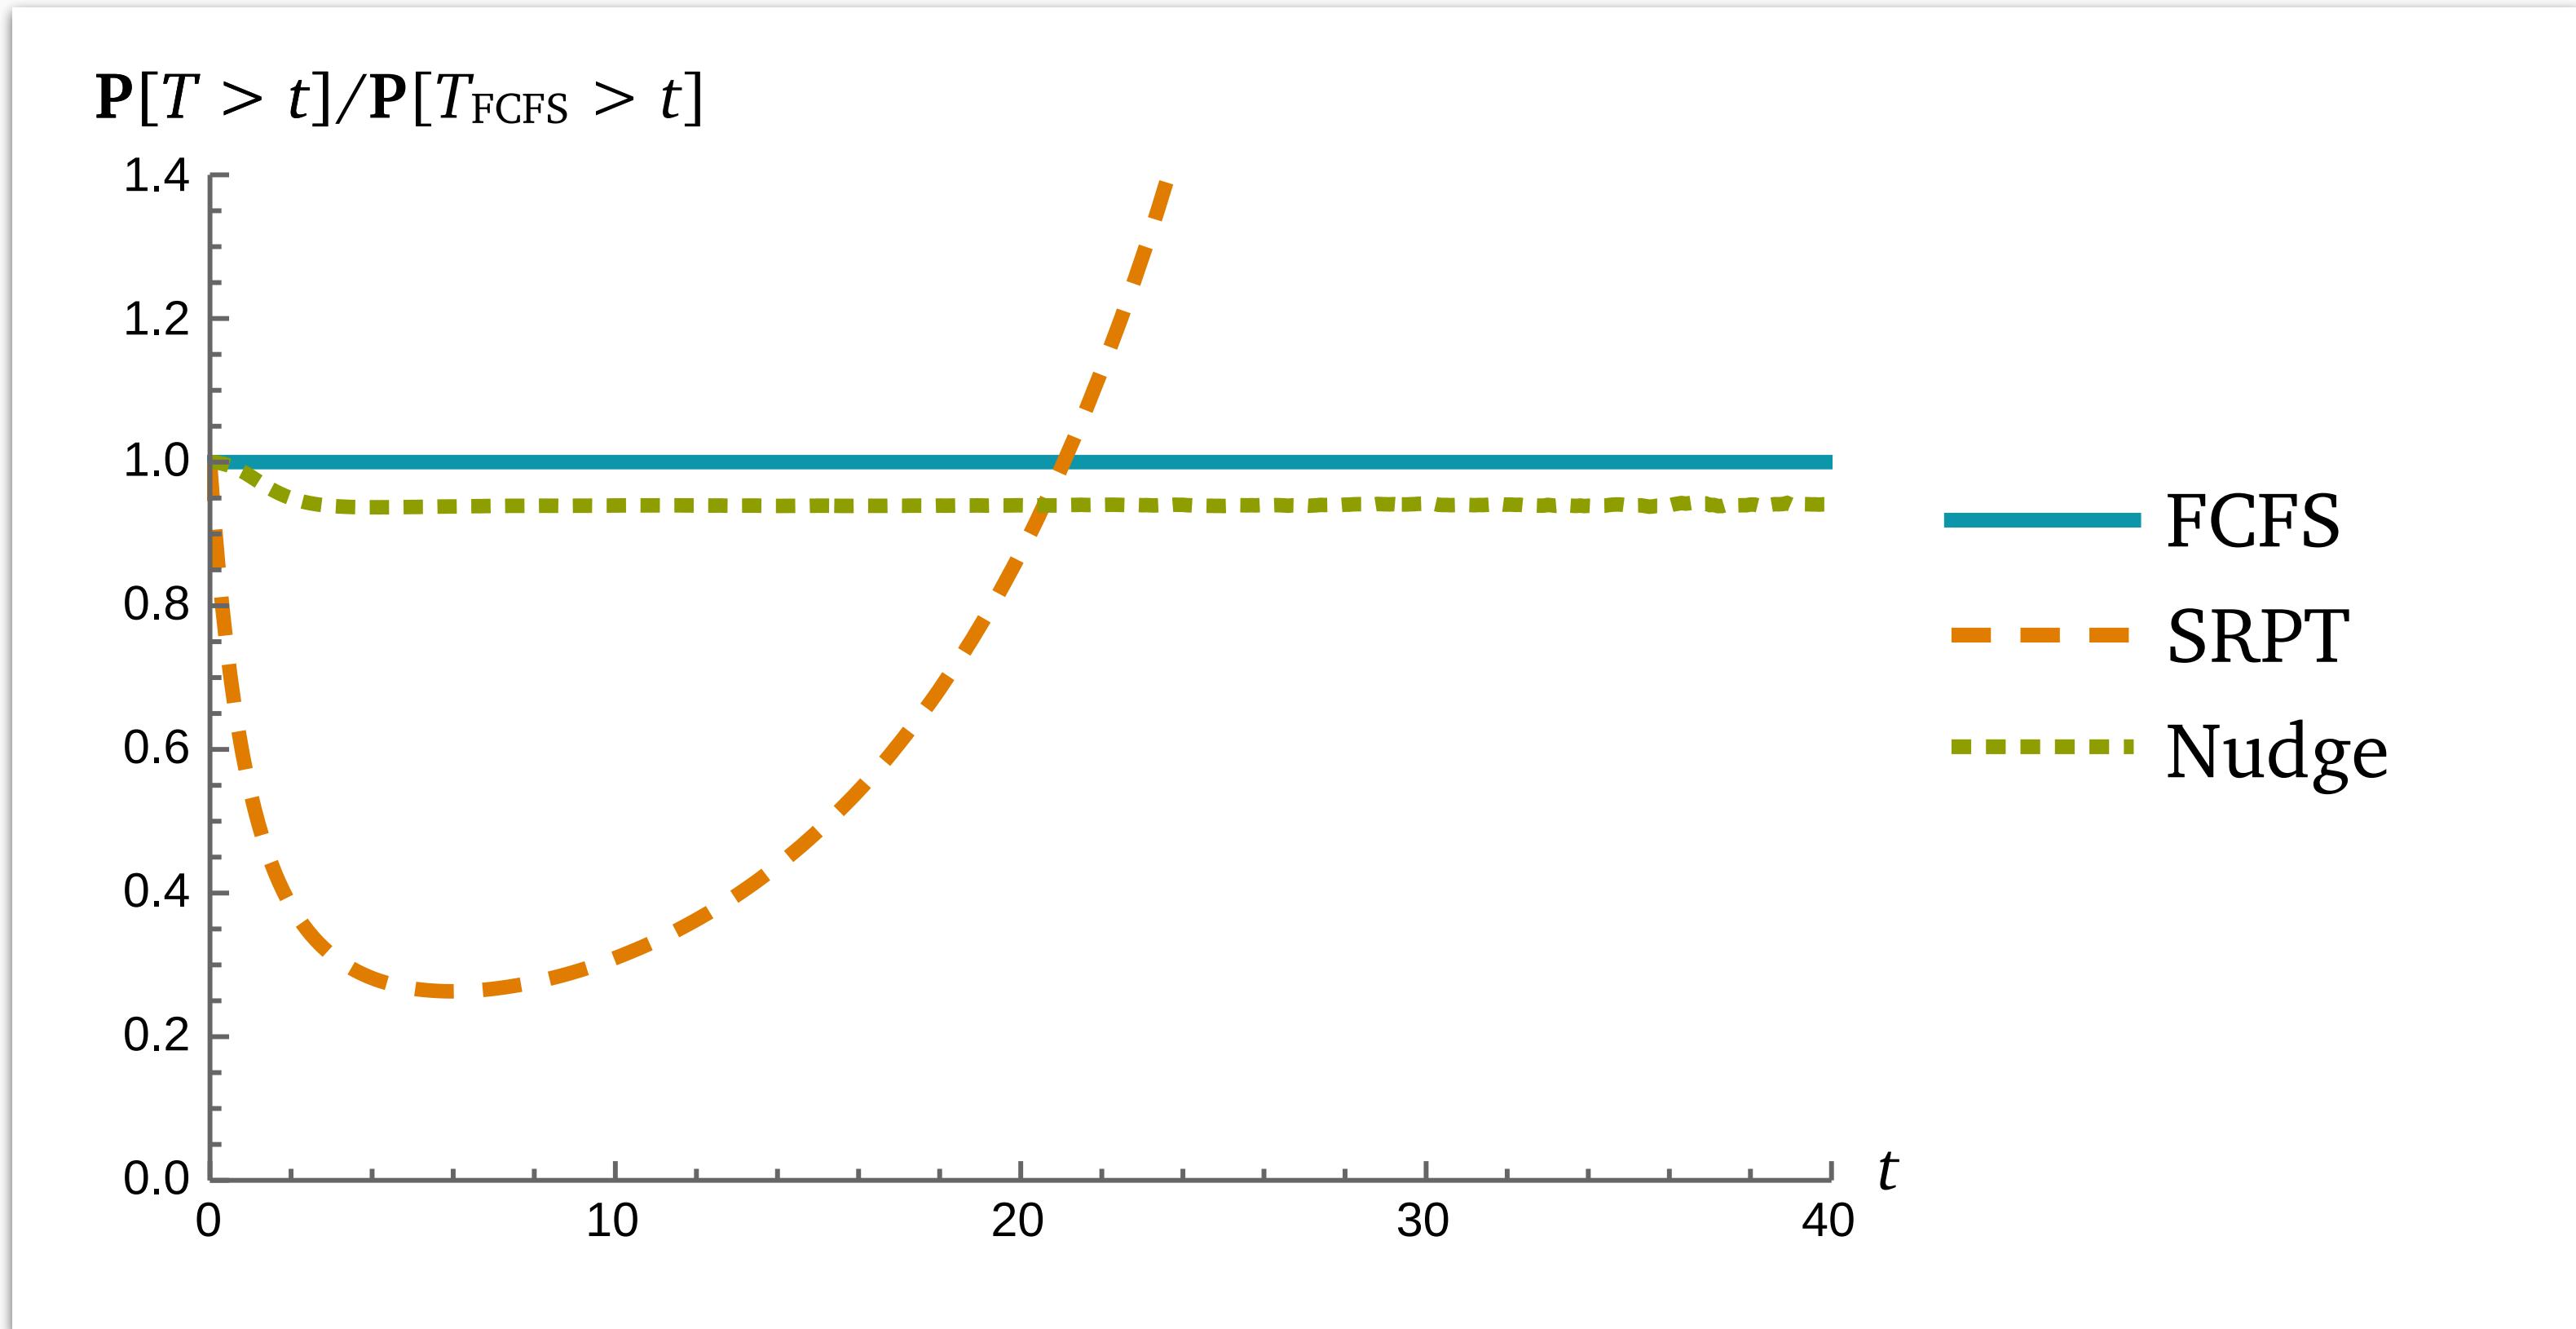
$\gamma_\pi = \text{decay rate of } \pi$
 $C_\pi = \text{tail constant of } \pi$

Strong optimality:
optimal γ_π and C_π

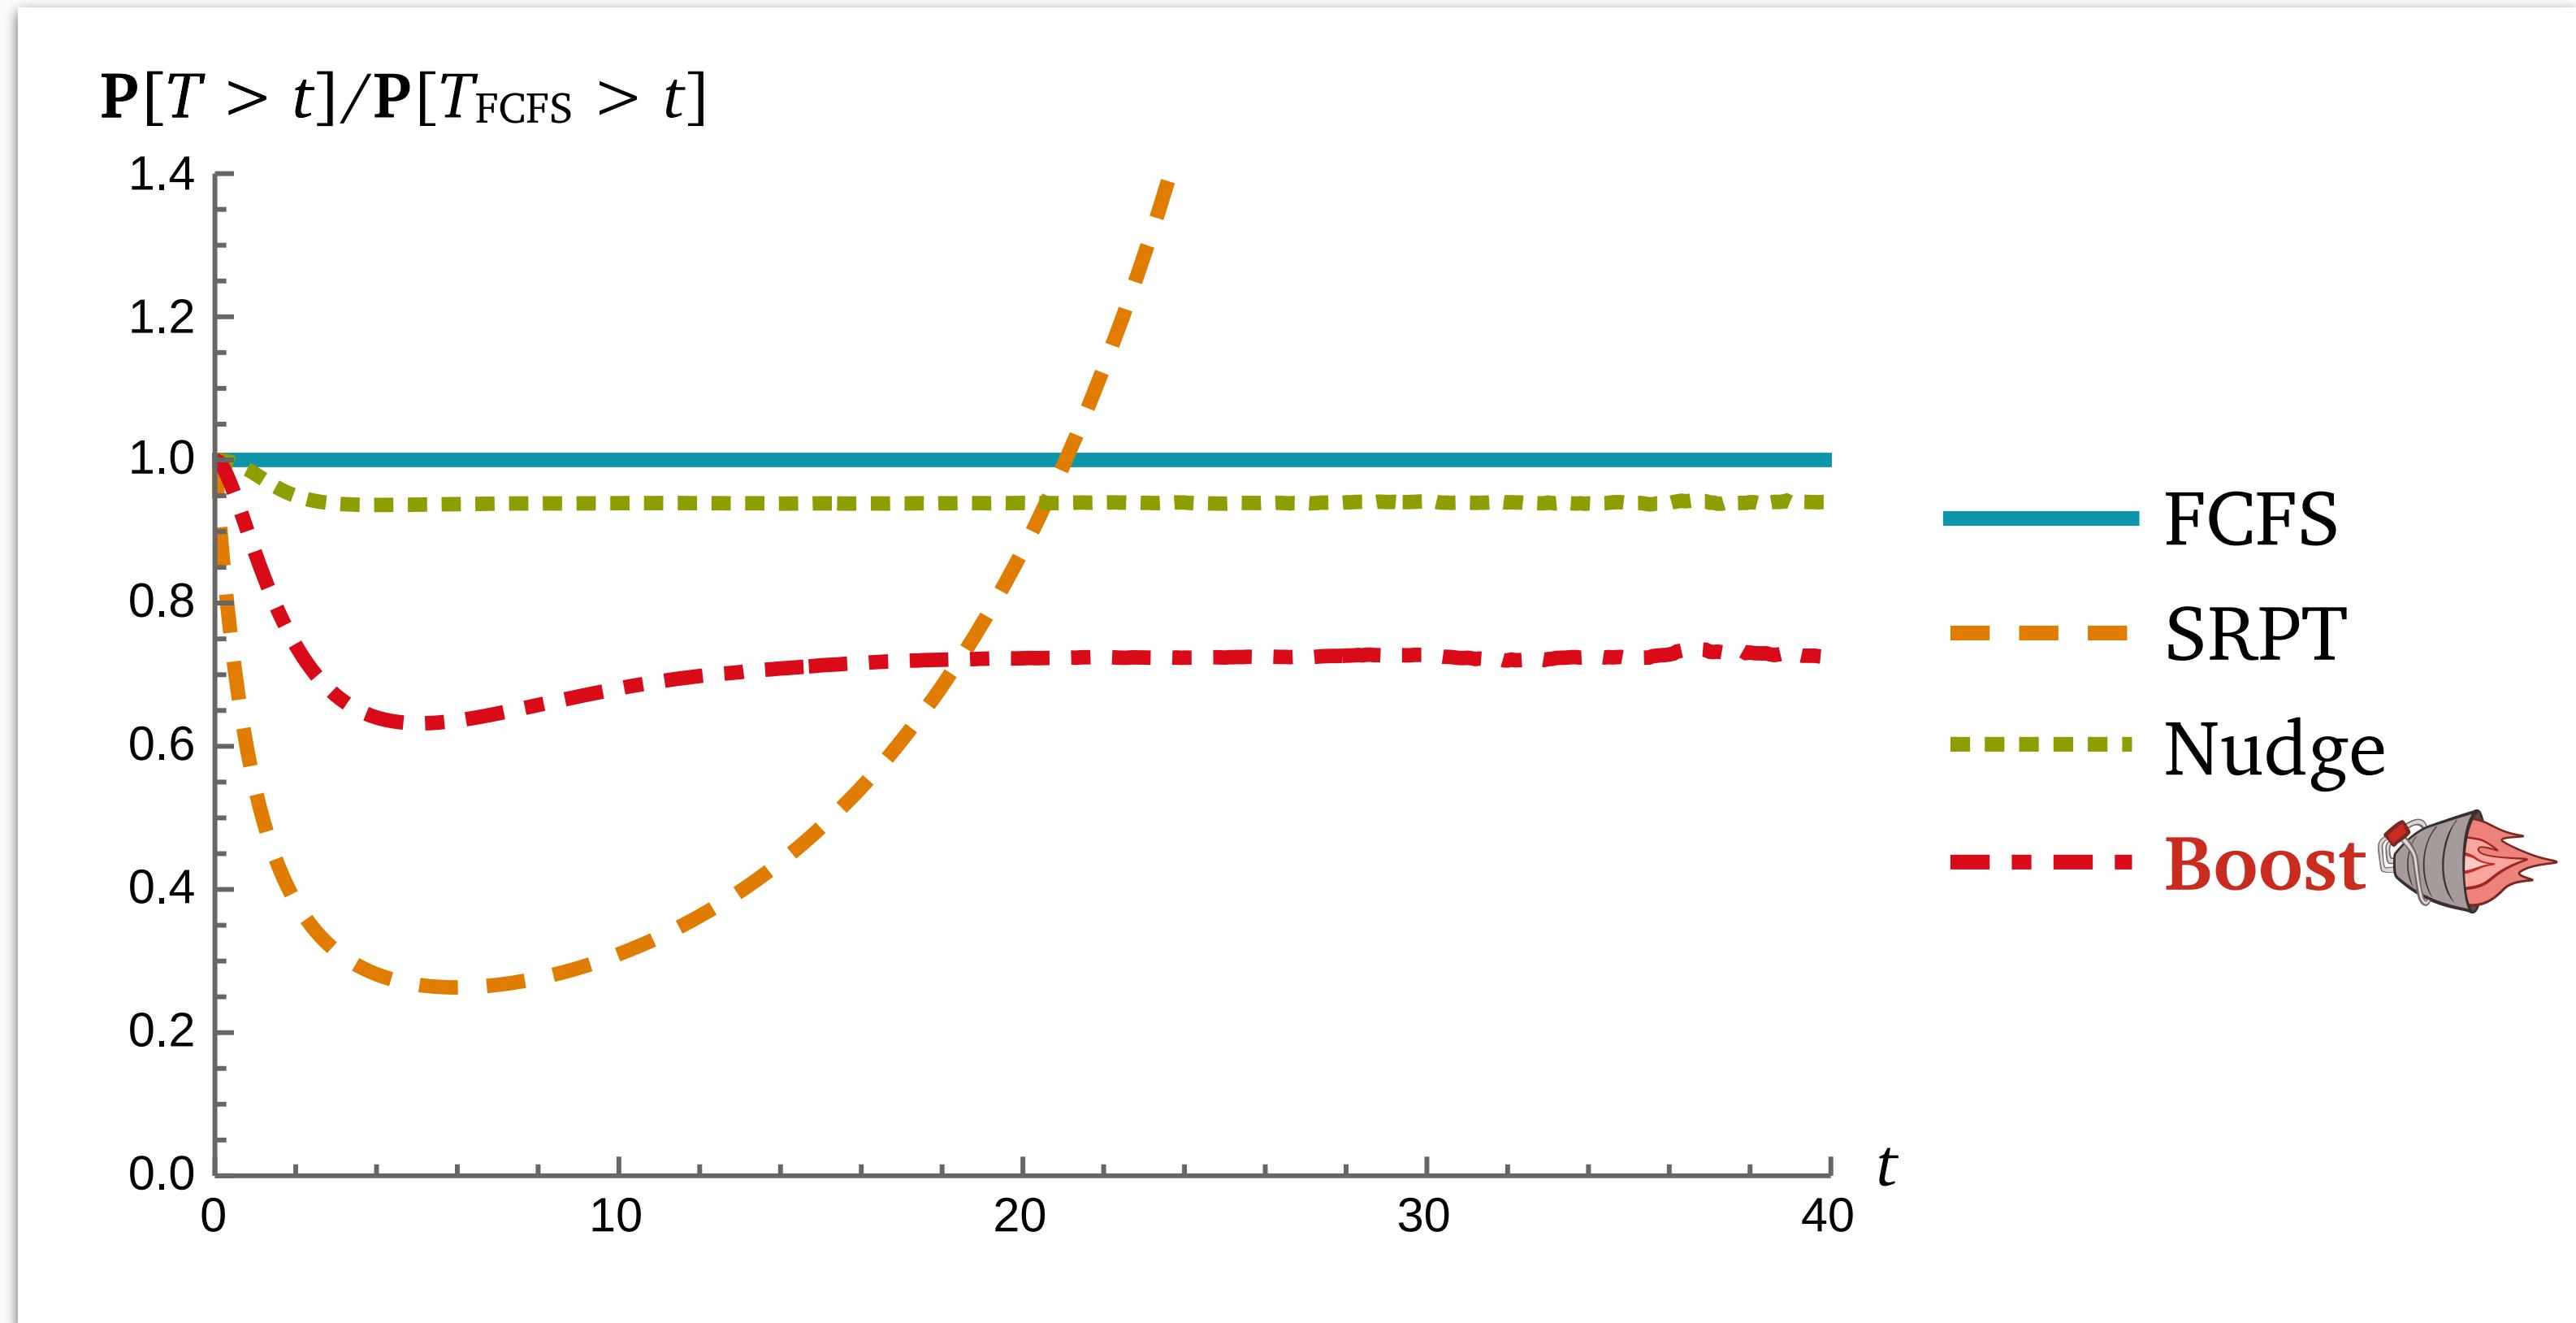
Optimizing the tail constant C



Optimizing the tail constant C



Optimizing the tail constant C



Boost policy



Scheduling rule: always serve job of
*minimum **boosted** arrival time*

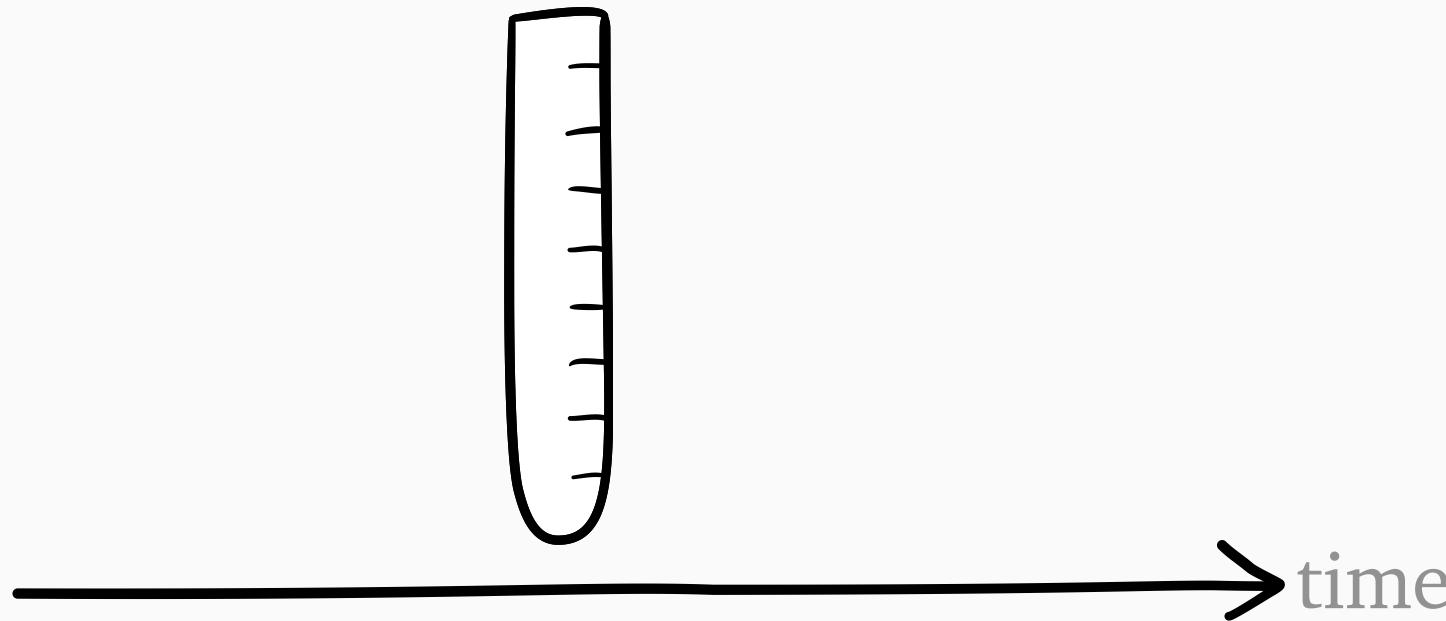


boosted arrival time
= arrival time – **boost(size s)**

Boost policy



Scheduling rule: always serve job of
minimum boosted arrival time

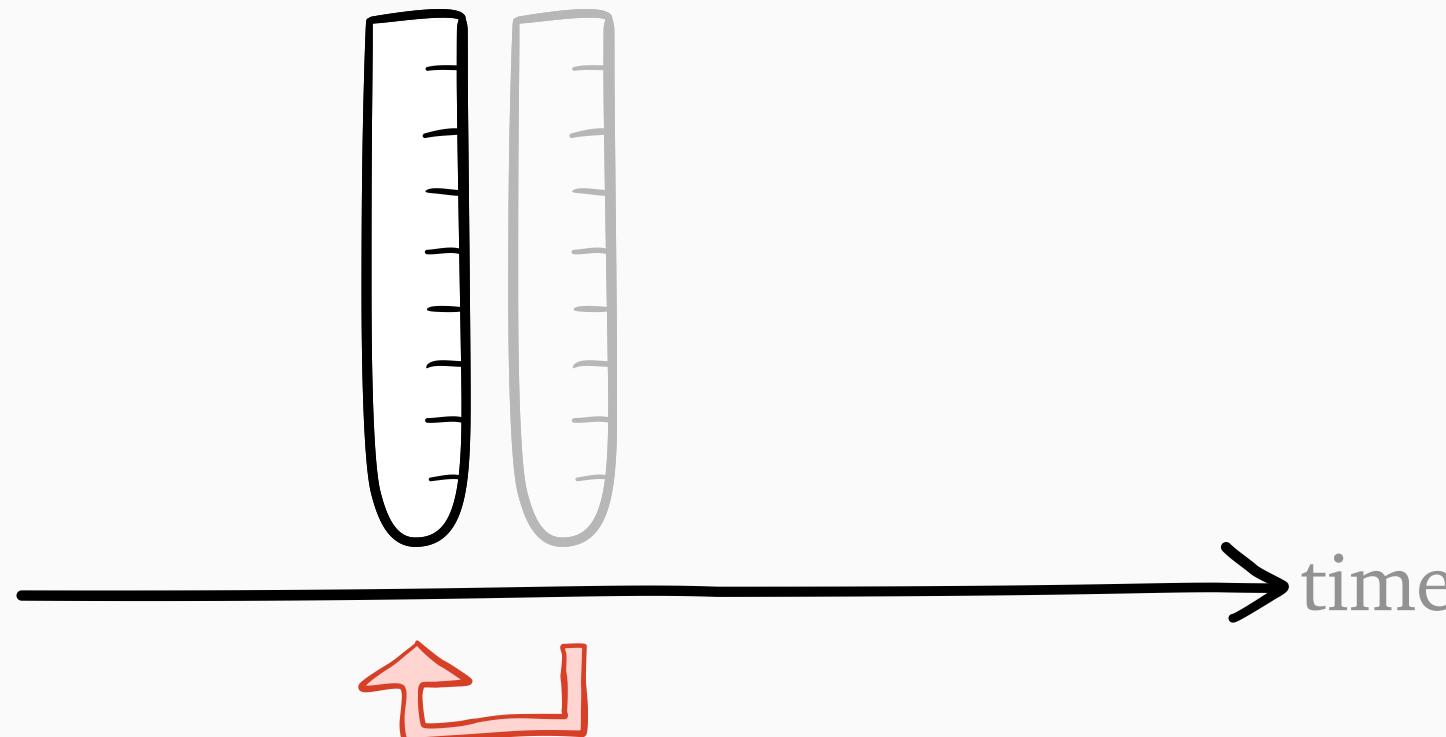


boosted arrival time
= arrival time – **boost**(size s)

Boost policy



Scheduling rule: always serve job of
minimum boosted arrival time

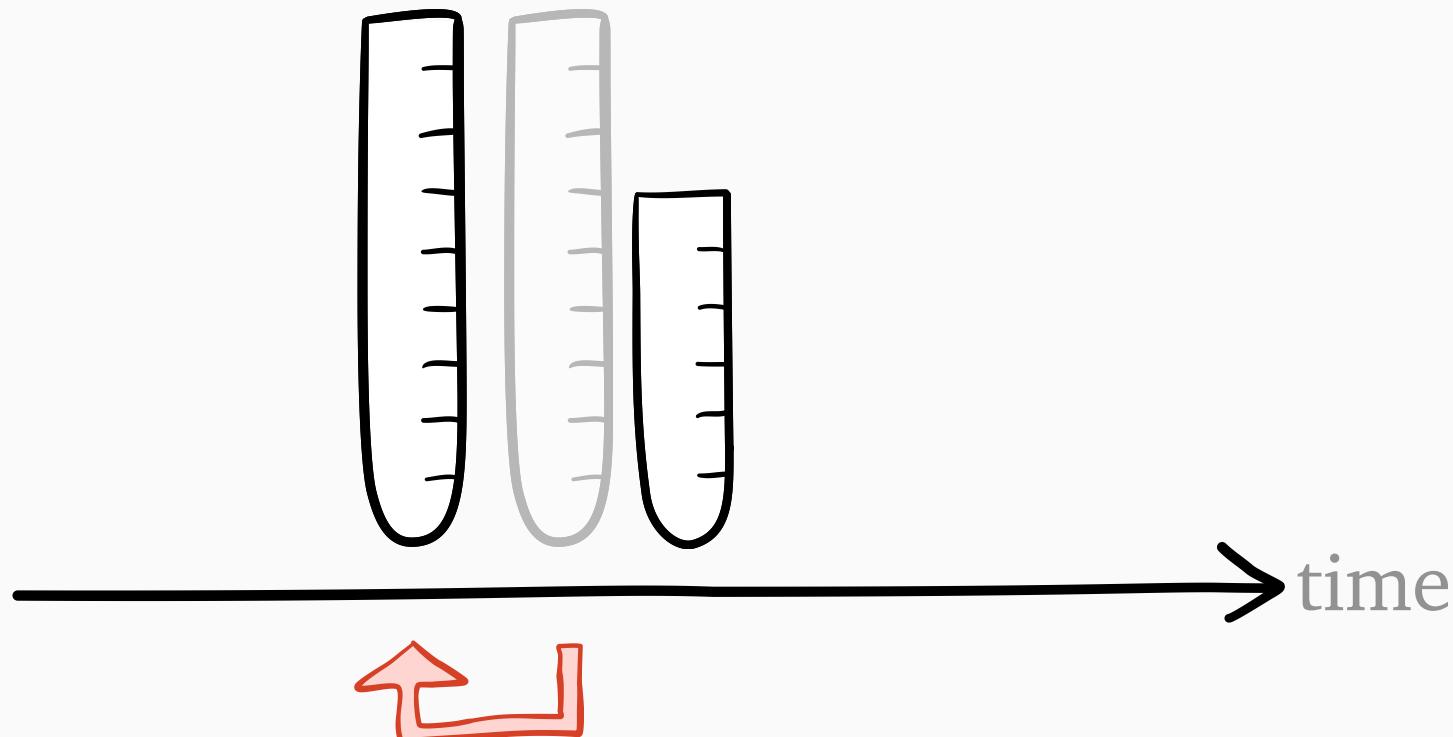


boosted arrival time
= arrival time – **boost**(size s)

Boost policy



Scheduling rule: always serve job of
minimum boosted arrival time

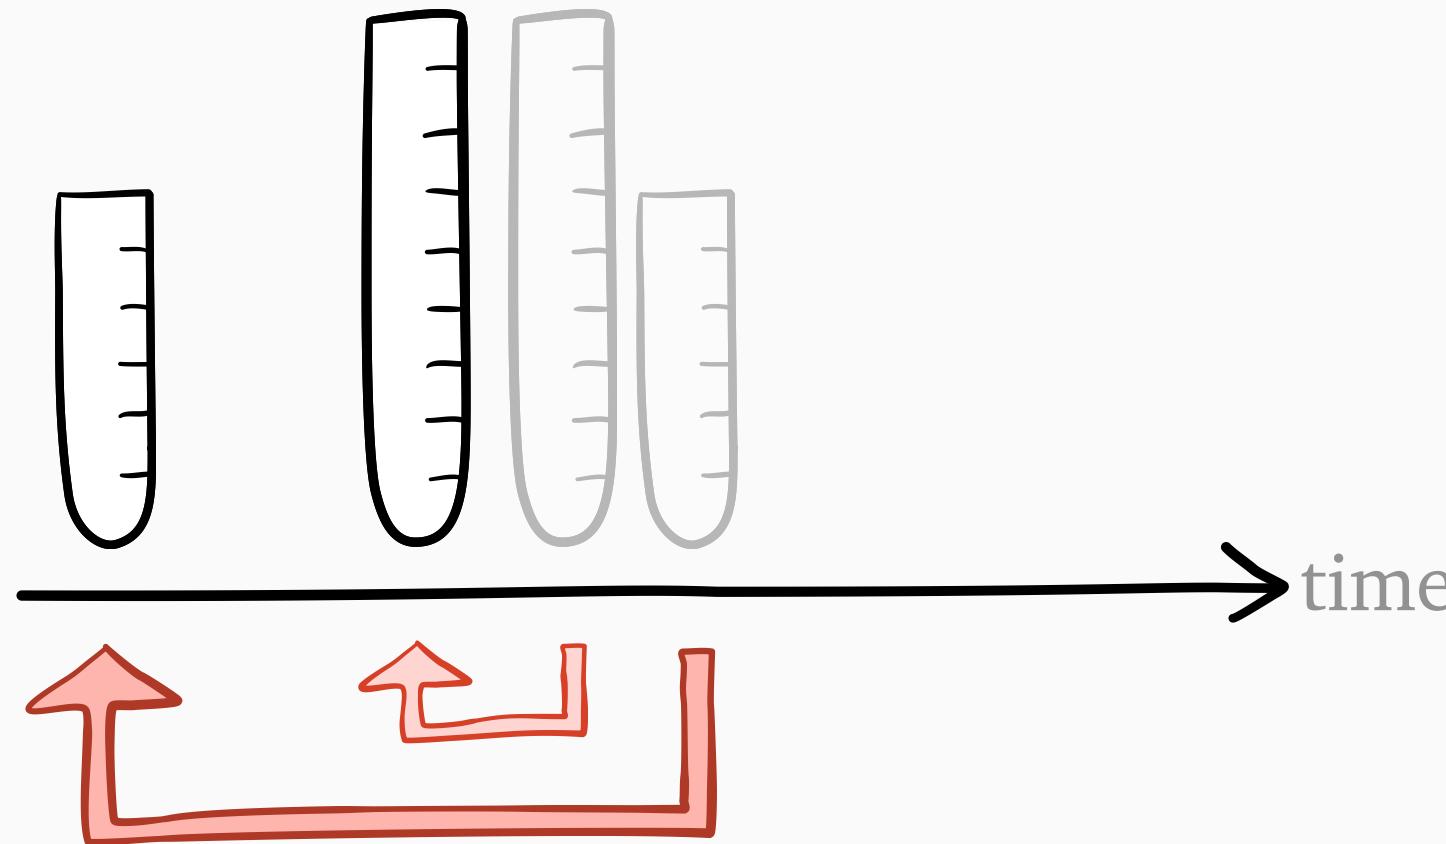


boosted arrival time
= arrival time – **boost**(size s)

Boost policy



Scheduling rule: always serve job of
minimum boosted arrival time

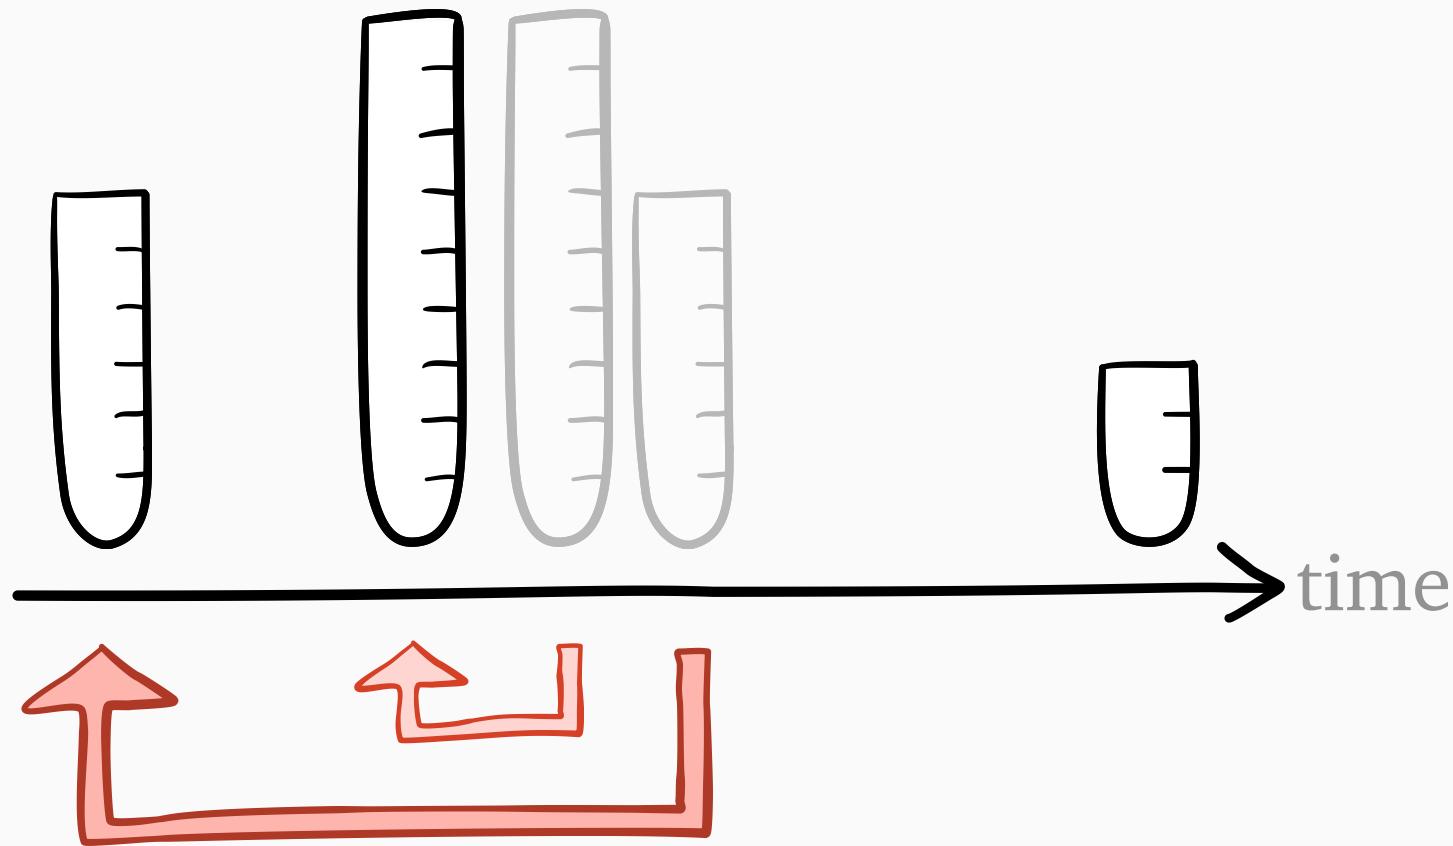


boosted arrival time
= arrival time – **boost**(size s)

Boost policy



Scheduling rule: always serve job of
minimum boosted arrival time

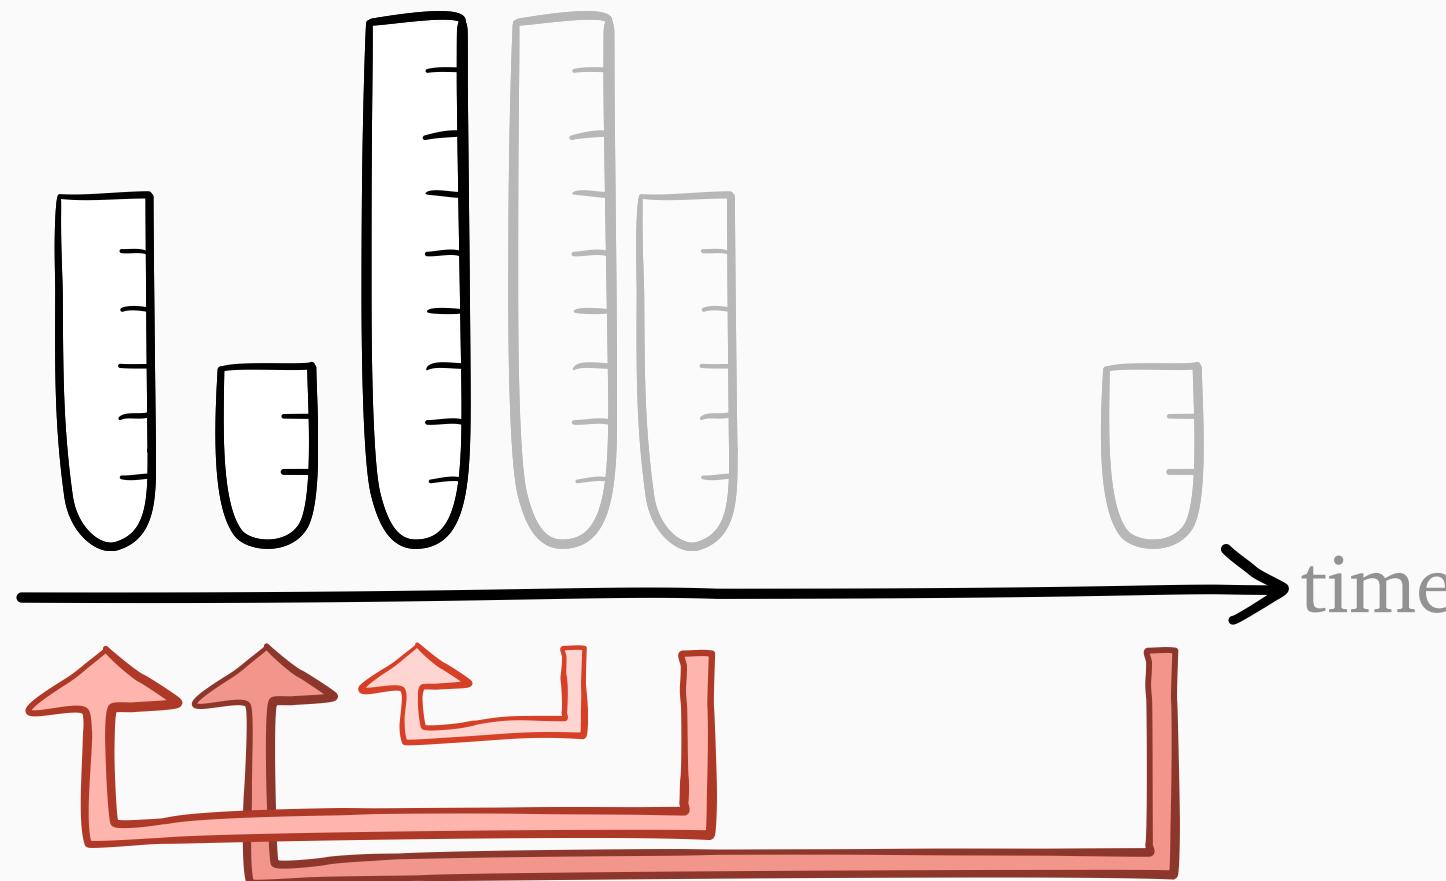


boosted arrival time
= arrival time – **boost**(size s)

Boost policy



Scheduling rule: always serve job of
minimum boosted arrival time

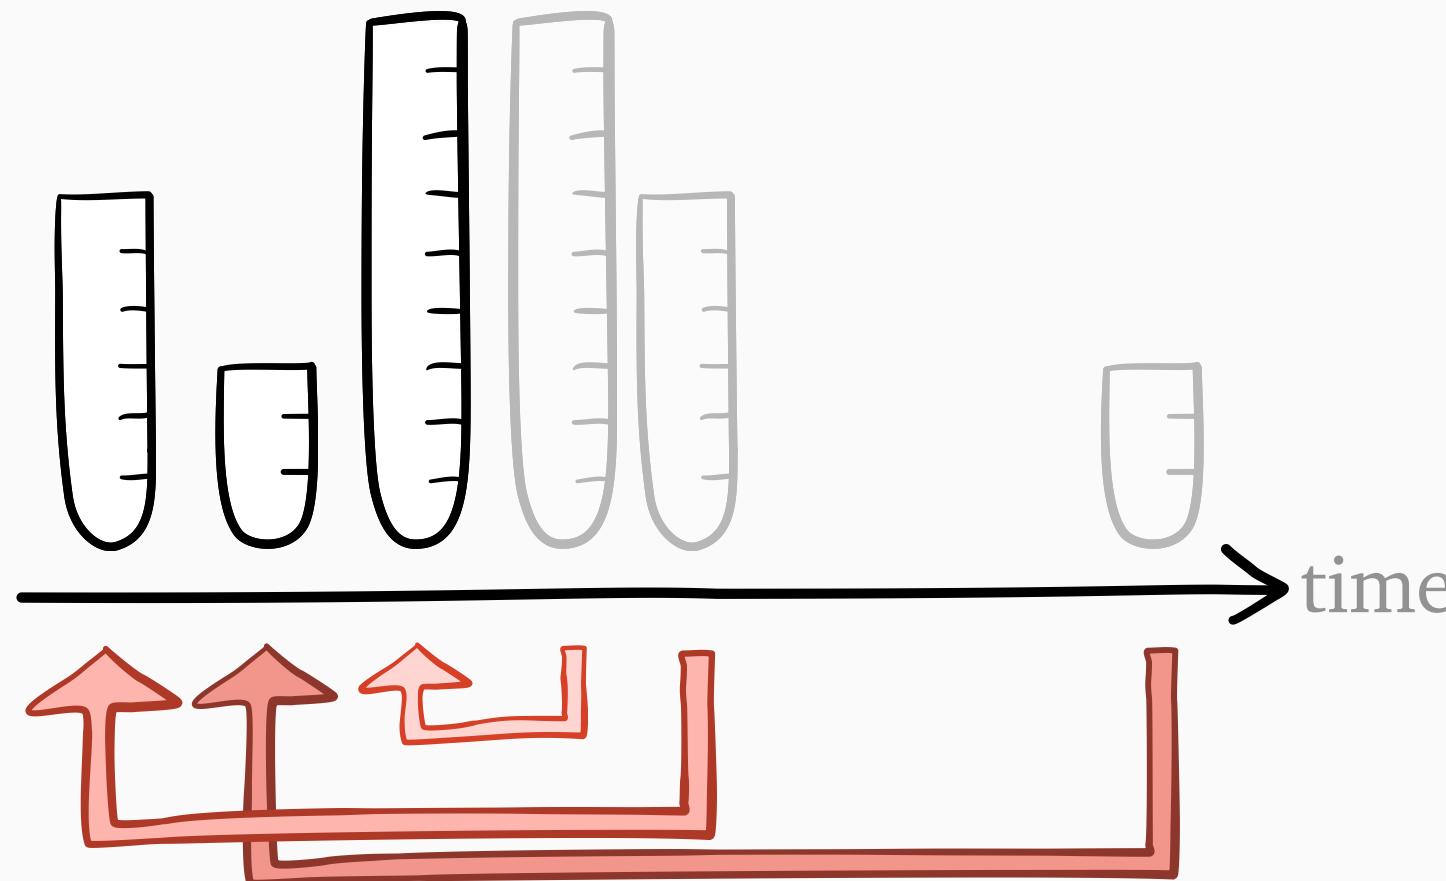


boosted arrival time
= arrival time – **boost**(size s)

Boost policy



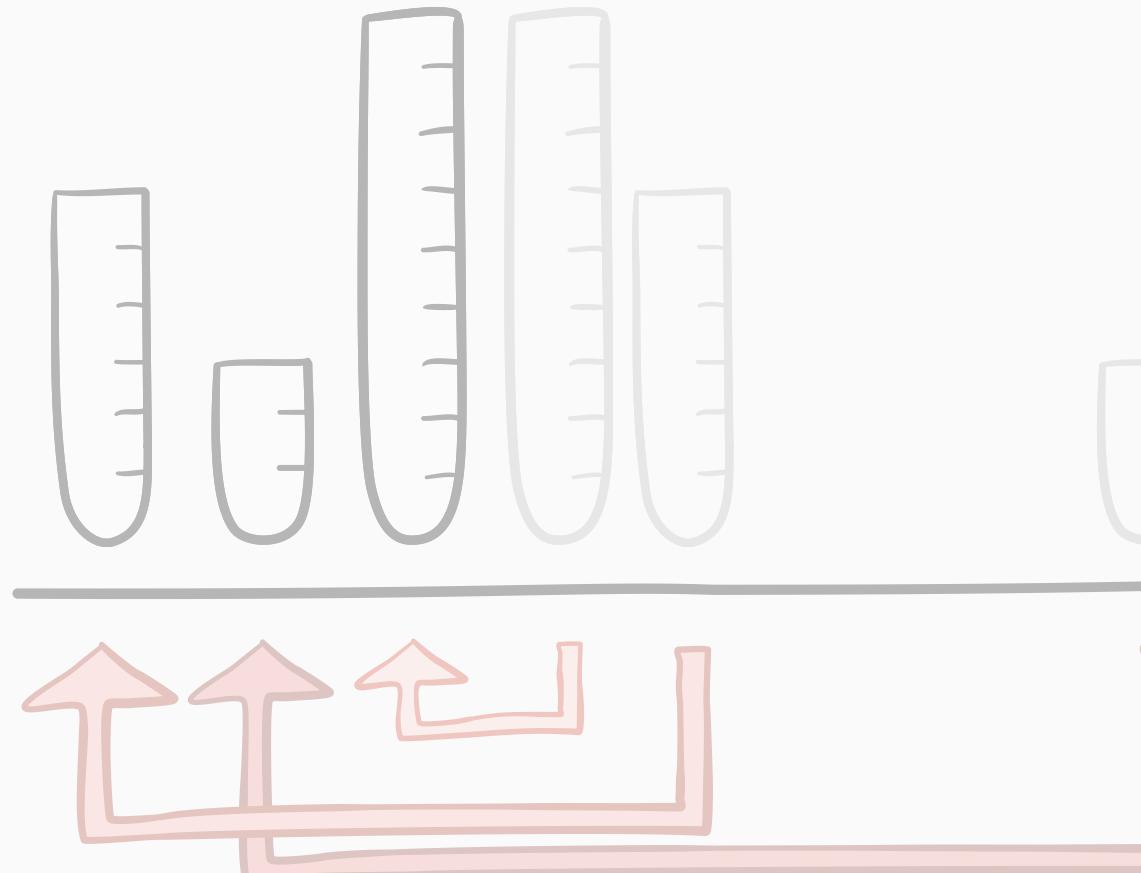
Scheduling rule: always serve job of minimum **boosted** arrival time



boosted arrival time
= arrival time – **boost**(size s)

$$\frac{1}{\gamma} \log \frac{1}{1 - e^{-\gamma s}}$$

Boost policy



Scheduling rule: always serve job of
minimum boosted arrival time



arrival time
arrival time – **boost**(size s)

$$\frac{1}{\gamma} \log \frac{1}{1 - e^{-\gamma s}}$$

Tail optimization questions

Tail optimization questions



From M/G/1 to M/G/k

Tail optimization questions



From M/G/1 to M/G/k



Minimizing tail of
slowdown $\mathbf{P}[T/S > x]$

Tail optimization questions

- ? From M/G/1 to M/G/k
- ? Minimizing tail of
slowdown $\mathbf{P}[T/S > x]$
- ? Revisiting “slightly”
heavy-tailed case ($\alpha = 4$)

Tail optimization questions



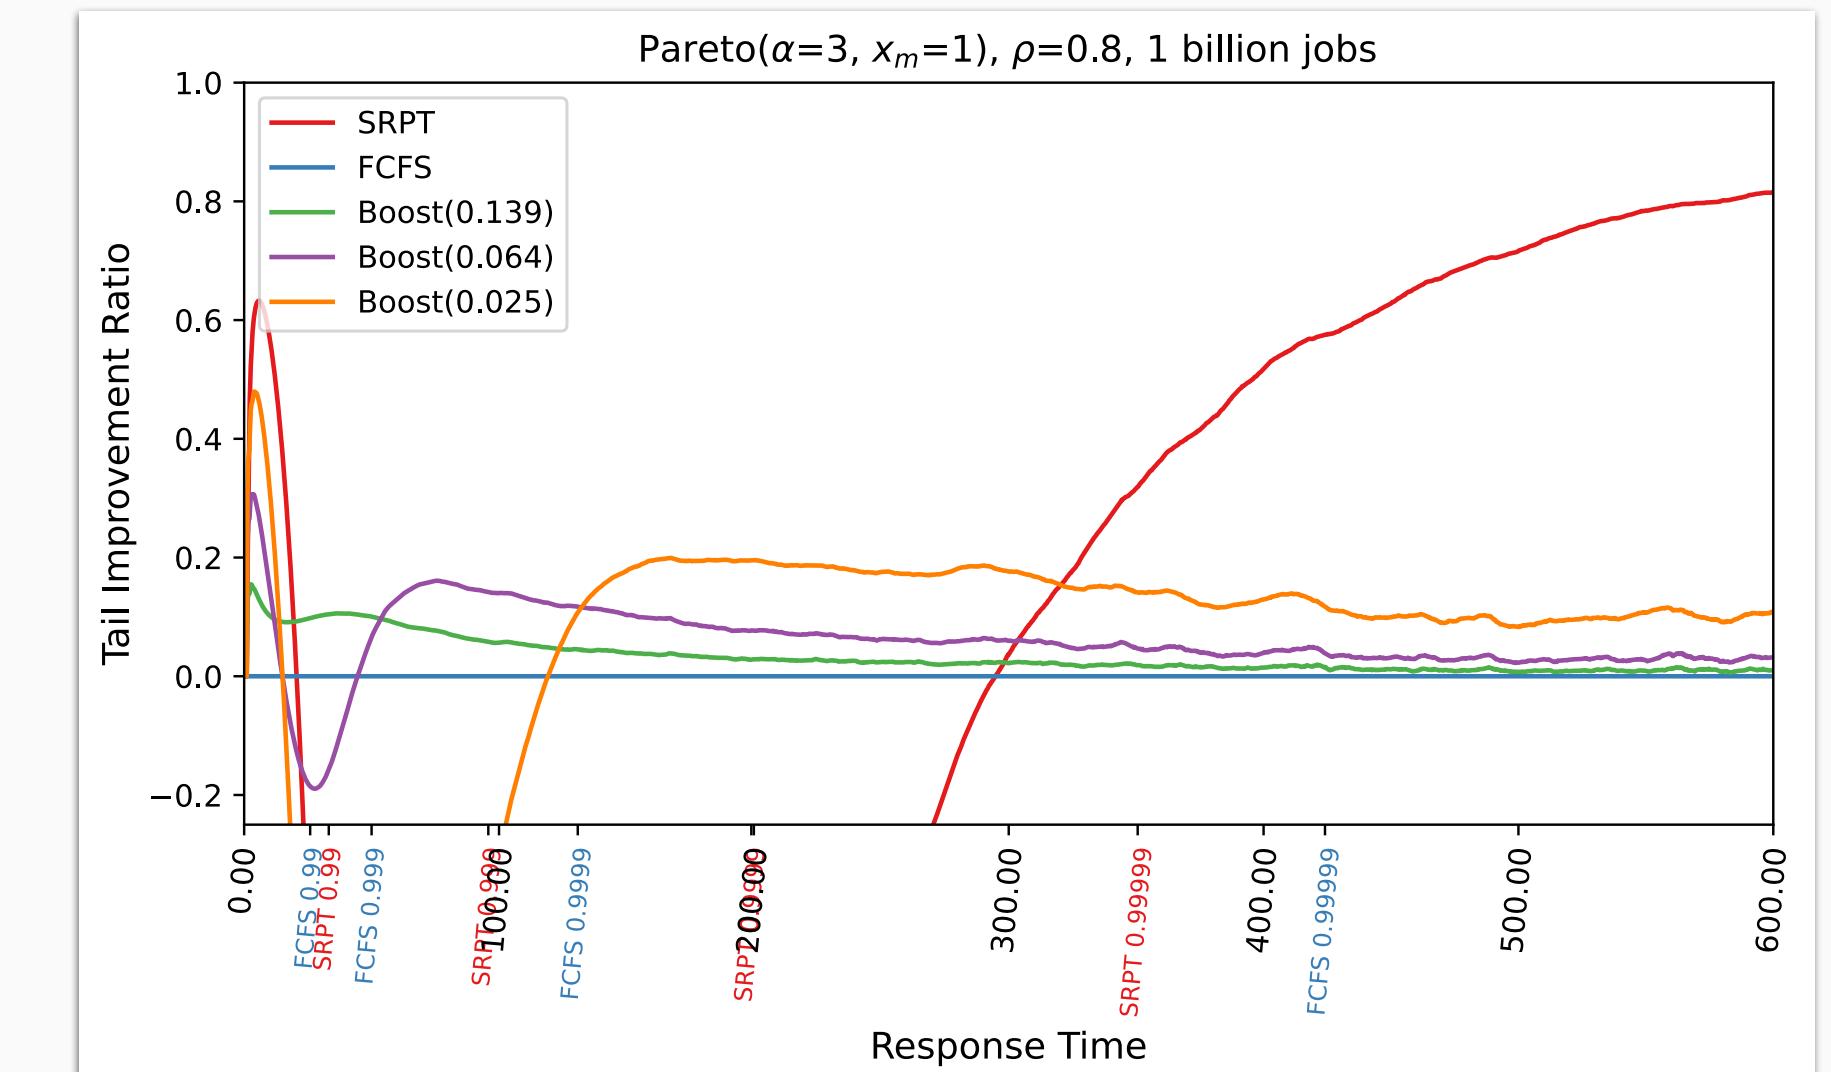
From M/G/1 to M/G/k



Minimizing tail of
slowdown $P[T/S > x]$



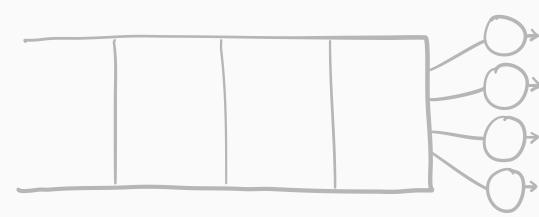
Revisiting “slightly”
heavy-tailed case ($\alpha = 4$)



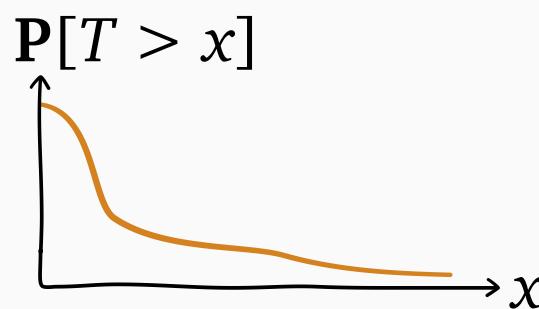
Hard scheduling questions



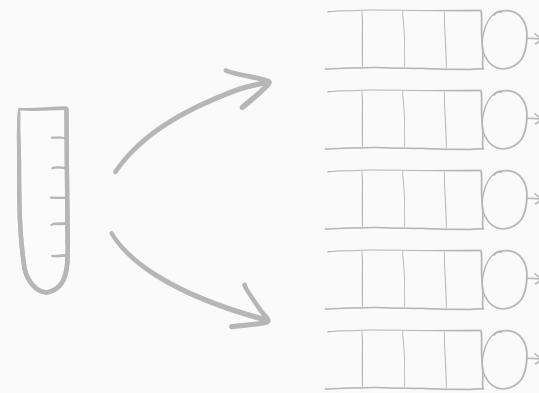
*What if we have
unknown job sizes?*



*What if there are
multiple servers?*



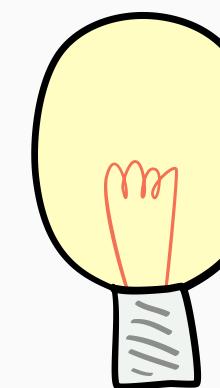
***What if we want to optimize
tails instead of means?***



*What if we can only use
dispatching without
fancy scheduling?*

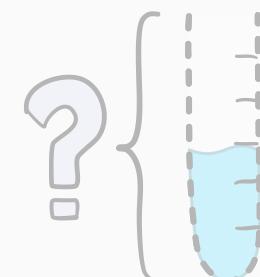
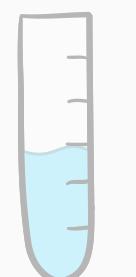


***What are the underlying
theoretical tools?***

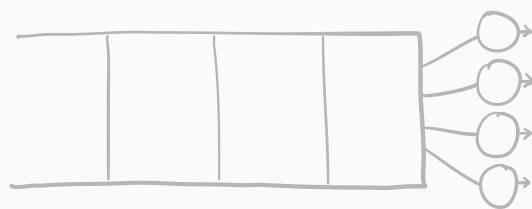


***What are the resulting
practical lessons?***

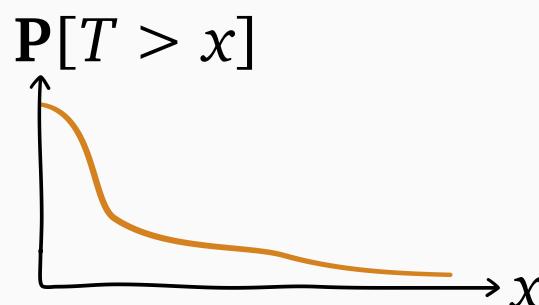
Hard scheduling questions



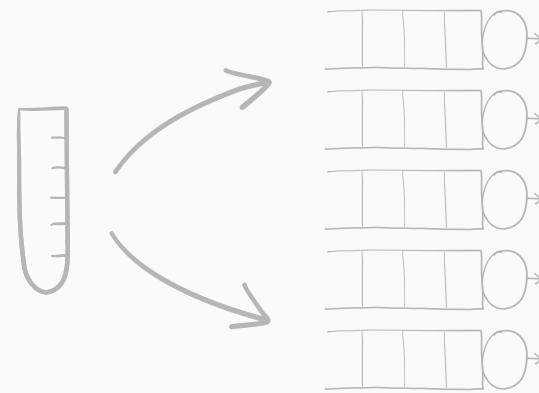
*What if we have
unknown job sizes?*



*What if there are
multiple servers?*



***What if we want to optimize
tails instead of means?***

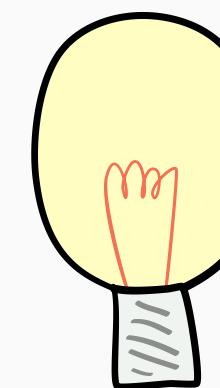


*What if we can only use
dispatching without
fancy scheduling?*



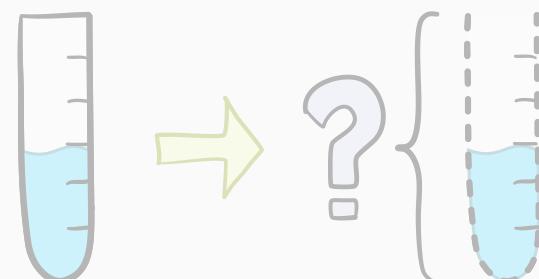
minimize C
 \Updownarrow
minimize “ $E[e^{\gamma T}]$ ”

***What are the underlying
theoretical tools?***

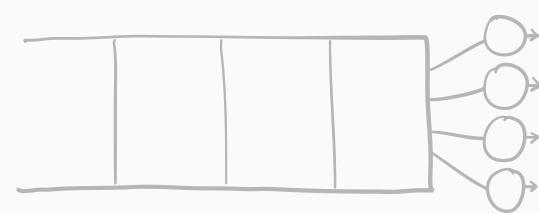


***What are the resulting
practical lessons?***

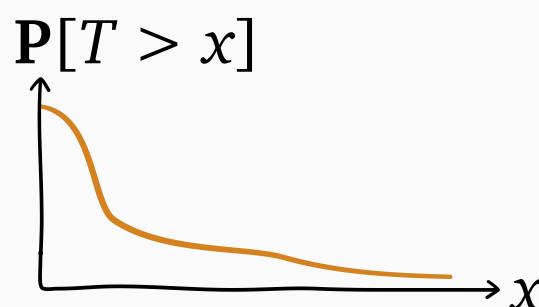
Hard scheduling questions



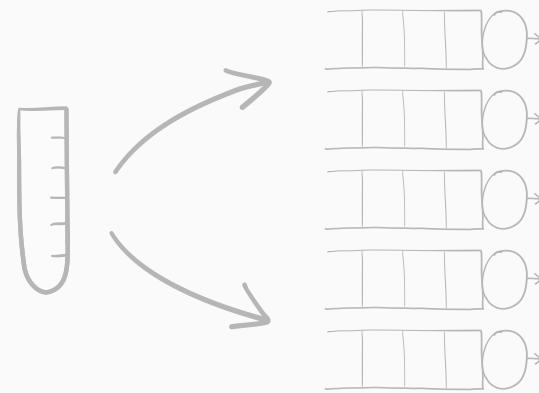
*What if we have
unknown job sizes?*



*What if there are
multiple servers?*



***What if we want to optimize
tails instead of means?***

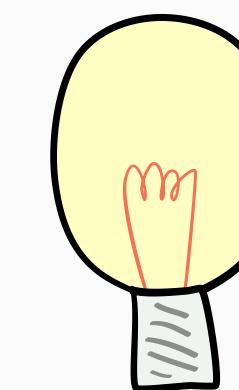


*What if we can only use
dispatching without
fancy scheduling?*



minimize C
 \Updownarrow
minimize “ $E[e^{\gamma T}]$ ”

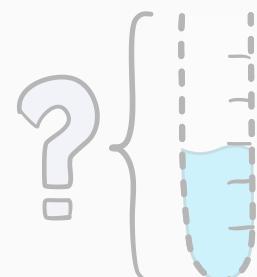
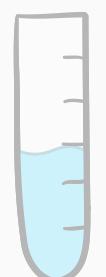
***What are the underlying
theoretical tools?***



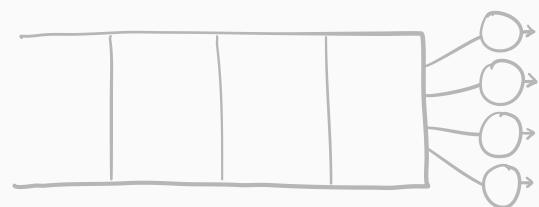
***What are the resulting
practical lessons?***

- Use arrival times and sizes
- Try **Boost** instead of FCFS!
- Try measuring $E[e^{\gamma T}]$?

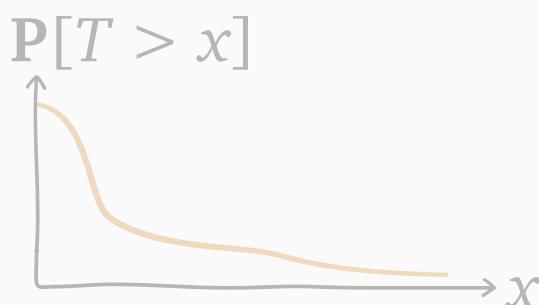
Hard scheduling questions



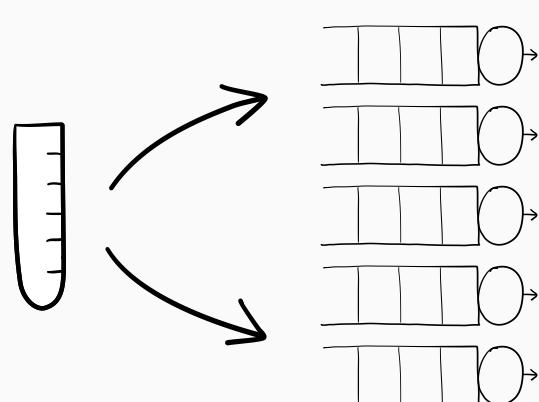
*What if we have
unknown job sizes?*



*What if there are
multiple servers?*



*What if we want to optimize
tails instead of means?*

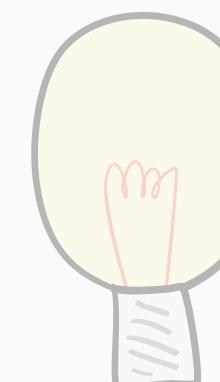


***What if we can only use
dispatching without
fancy scheduling?***



minimize C
 \Leftrightarrow
minimize " $E[e^{\gamma T}]$ "

*What are the underlying
theoretical tools?*

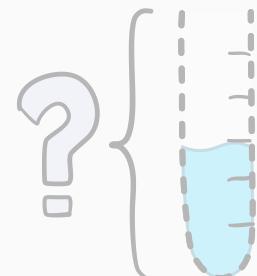
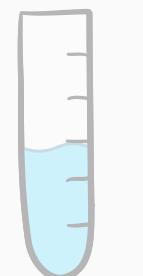


*What are the resulting
practical lessons?*

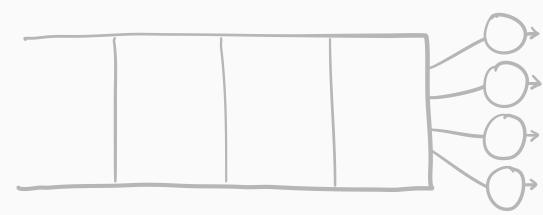
- Use arrival times and sizes
- Try **Boost** instead of FCFS!
- Try measuring $E[e^{\gamma T}]$?

... stay tuned for Natalie's talk!

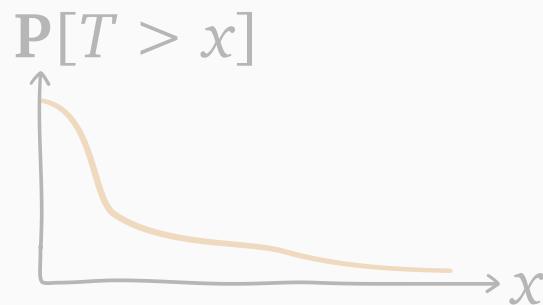
Hard scheduling questions



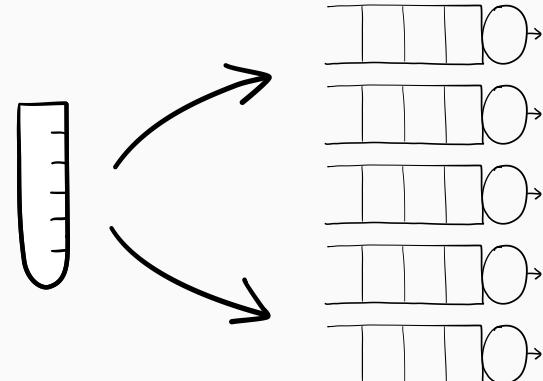
*What if we have
unknown job sizes?*



*What if there are
multiple servers?*



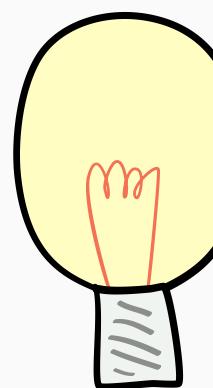
*What if we want to optimize
tails instead of means?*



***What if we can only use
dispatching without
fancy scheduling?***

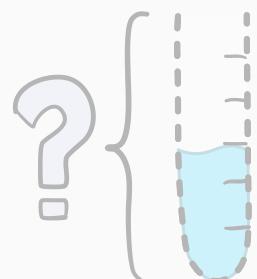
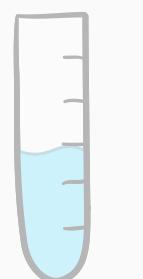


***What are the underlying
theoretical tools?***

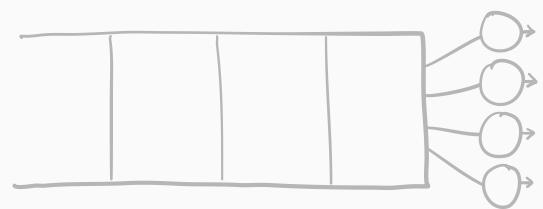


***What are the resulting
practical lessons?***

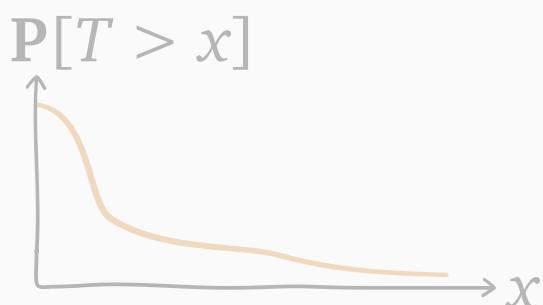
Hard scheduling questions



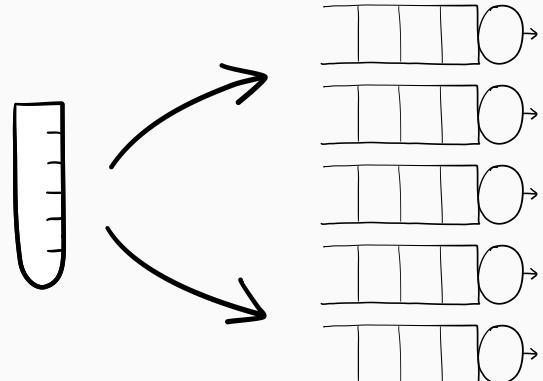
*What if we have
unknown job sizes?*



*What if there are
multiple servers?*



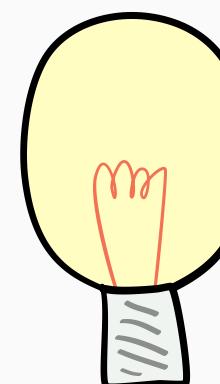
*What if we want to optimize
tails instead of means?*



***What if we can only use
dispatching without
fancy scheduling?***



***What are the underlying
theoretical tools?***



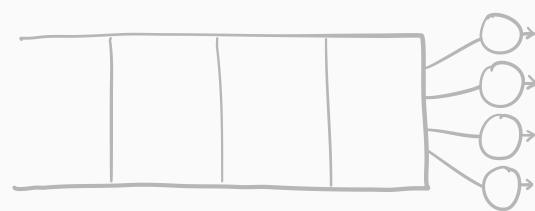
***What are the resulting
practical lessons?***

 ***newish Continuous state
space collapse***

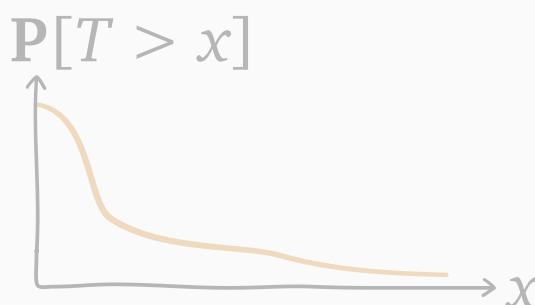
Hard scheduling questions



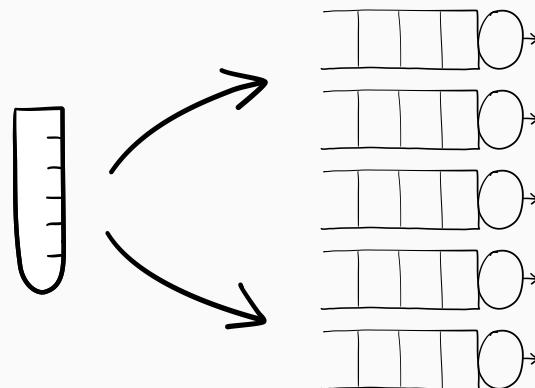
*What if we have
unknown job sizes?*



*What if there are
multiple servers?*



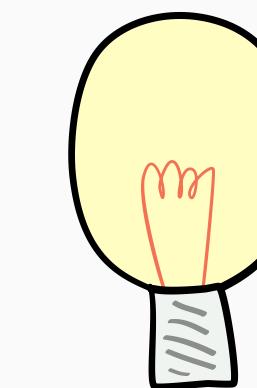
*What if we want to optimize
tails instead of means?*



***What if we can only use
dispatching without
fancy scheduling?***



***What are the underlying
theoretical tools?***



***What are the resulting
practical lessons?***

- Simulate priority classes
- Maintain work imbalance using closed-loop control

But wait, there's more!

- G/G arrivals
- Generalized switch
- Parallelizable jobs and other complex service choices
- Multiserver jobs and other complex packing concerns
- Preemption overhead
- Large-system limits
- Stein's method

Hong & Scully (2024). doi:[10.1016/j.peva.2023.102377](https://doi.org/10.1016/j.peva.2023.102377).

Hurtado-Lange & Maguluri (2020). doi:[10.1287/stsy.2019.0056](https://doi.org/10.1287/stsy.2019.0056).

Berg (2022). https://bsb20.github.io/bsberg_phd_csd_2022.pdf.

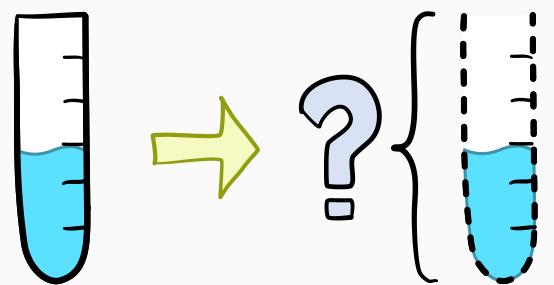
Grosof, Scully, Harchol-Balter, & Scheller-Wolf (2022). doi:[10.1145/3570612](https://doi.org/10.1145/3570612).

Ramakrishna & Scully (2024). doi:[10.1145/3695411.3695419](https://doi.org/10.1145/3695411.3695419).

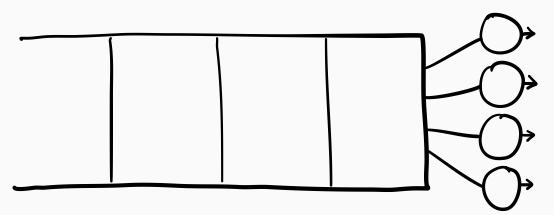
Hong & Wang (2022). doi:[10.1145/3492866.3549717](https://doi.org/10.1145/3492866.3549717).

Braverman (2023). doi:[10.1287/stsy.2022.0102](https://doi.org/10.1287/stsy.2022.0102).

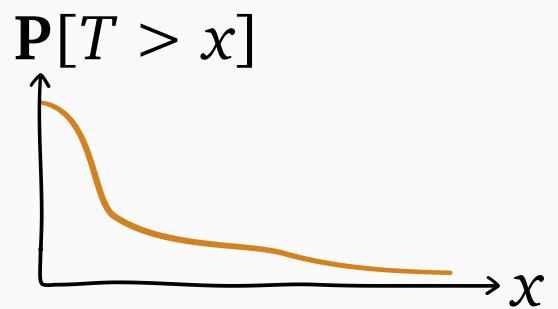
Hard scheduling questions



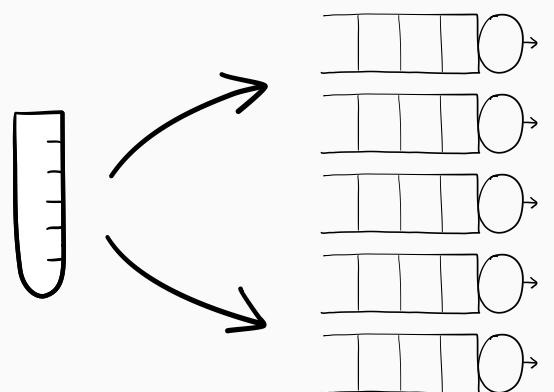
*What if we have
unknown job sizes?*



*What if there are
multiple servers?*

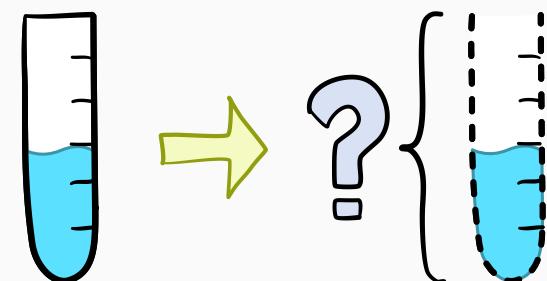


*What if we want to optimize
tails instead of means?*

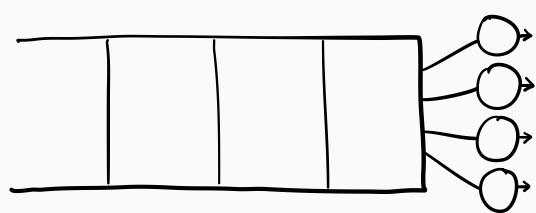


*What if we can only use
dispatching without
fancy scheduling?*

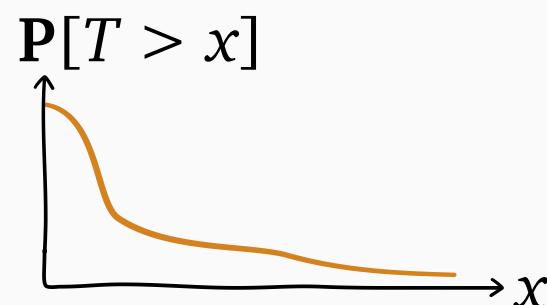
Hard scheduling questions



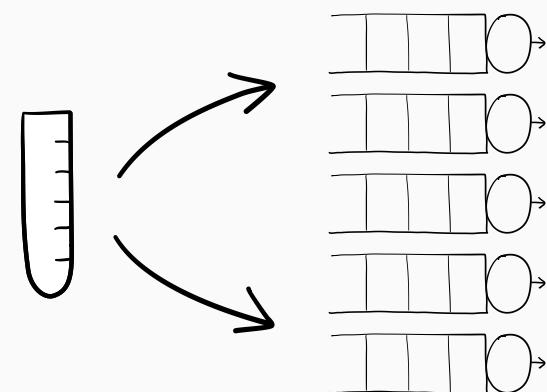
*What if we have
unknown job sizes?*



*What if there are
multiple servers?*



*What if we want to optimize
tails instead of means?*

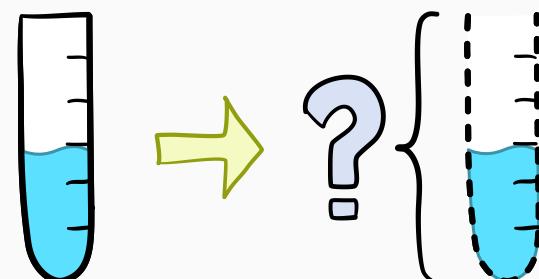


*What if we can only use
dispatching without
fancy scheduling?*

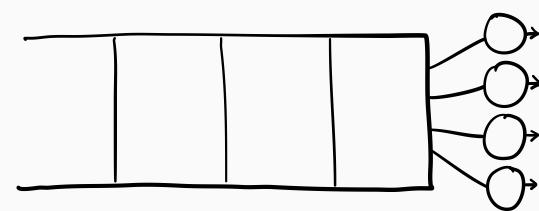


Theoretical tools

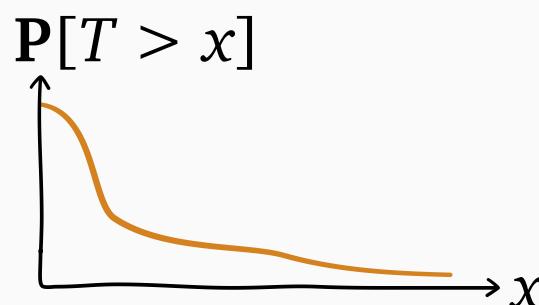
Hard scheduling questions



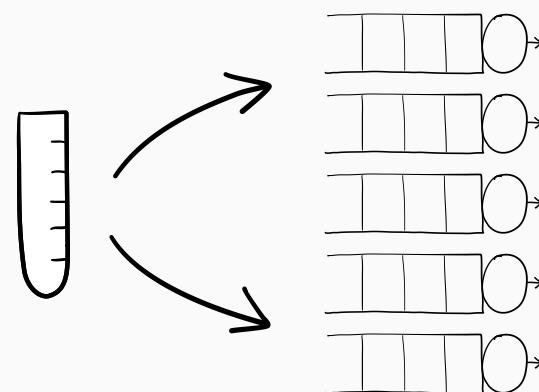
*What if we have
unknown job sizes?*



*What if there are
multiple servers?*



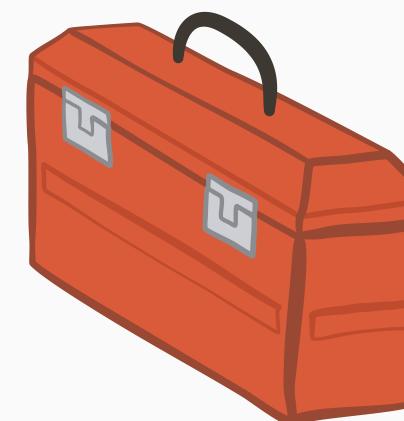
*What if we want to optimize
tails instead of means?*



*What if we can only use
dispatching without
fancy scheduling?*

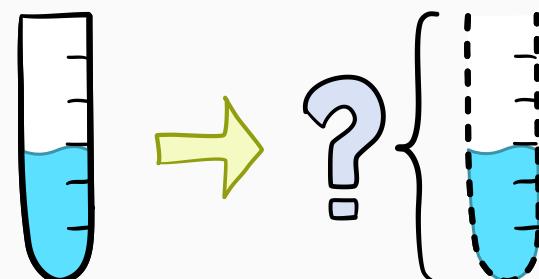


Work decomposition
State space collapse

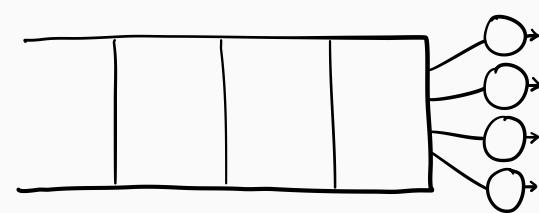


Theoretical tools

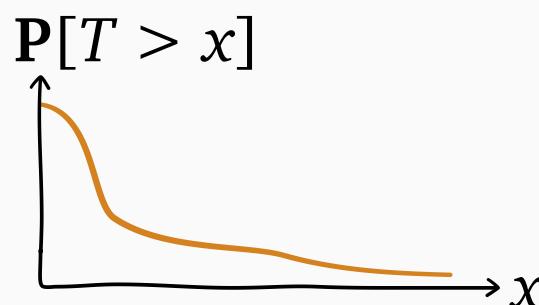
Hard scheduling questions



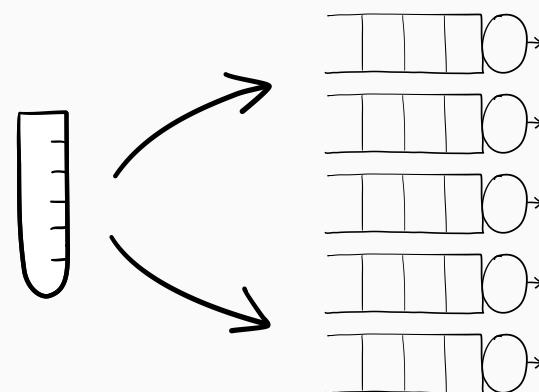
*What if we have
unknown job sizes?*



*What if there are
multiple servers?*



*What if we want to optimize
tails instead of means?*



*What if we can only use
dispatching without
fancy scheduling?*



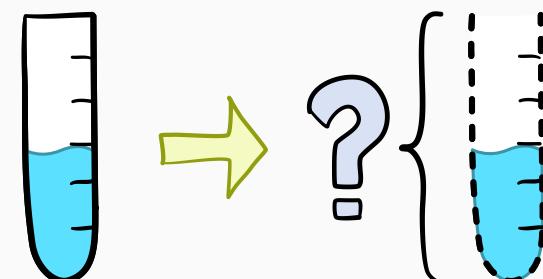
Work decomposition
State space collapse



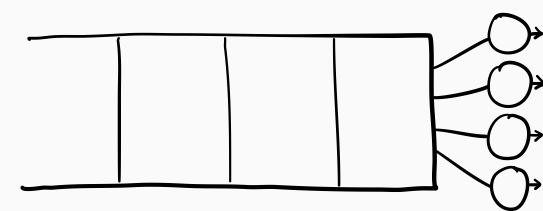
Theoretical tools
 Practical lessons



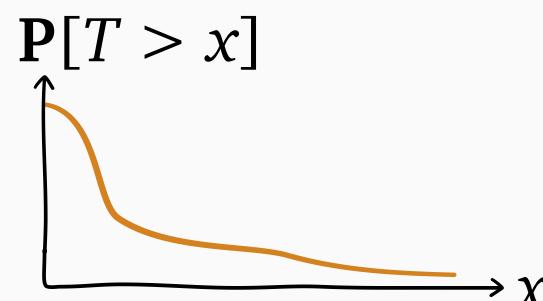
Hard scheduling questions



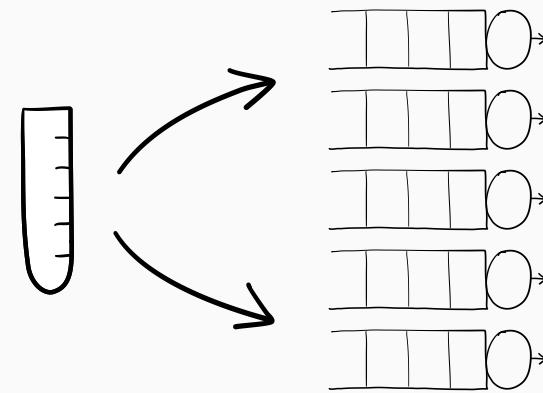
What if we have unknown job sizes?



What if there are multiple servers?



What if we want to optimize tails instead of means?



What if we can only use dispatching without fancy scheduling?



Work decomposition
State space collapse



PSJF is robust

Use size and arrival time

Closed-loop dispatching control



Theoretical tools
↓
Practical lessons
↓
↓
↓
↓
↓

Empirical Gittins works

k servers \approx 1 server