

Uniform Bounds for **Scheduling**
with **Job Size Estimates**

Ziv Scully


Cornell

Isaac Grosf


Carnegie Mellon → Northwestern

Michael Mitzenmacher


Harvard



algorithms
with
predictions




scheduling algorithms
with
predictions




scheduling algorithms

with

job size predictions






scheduling algorithms

with

job size predictions



Twist: *stochastic* setting

How do we schedule
to minimize delay
with noisy size estimates?

How do we schedule
to minimize delay
with noisy size estimates?



What are scheduling and delay?

How do we schedule
to minimize delay
with noisy size estimates?



What are scheduling and delay?

What job size noise model?

How do we schedule
to minimize delay
with noisy size estimates?




What are scheduling and delay?


What job size noise model?

What can we hope to achieve?


What are scheduling and delay?



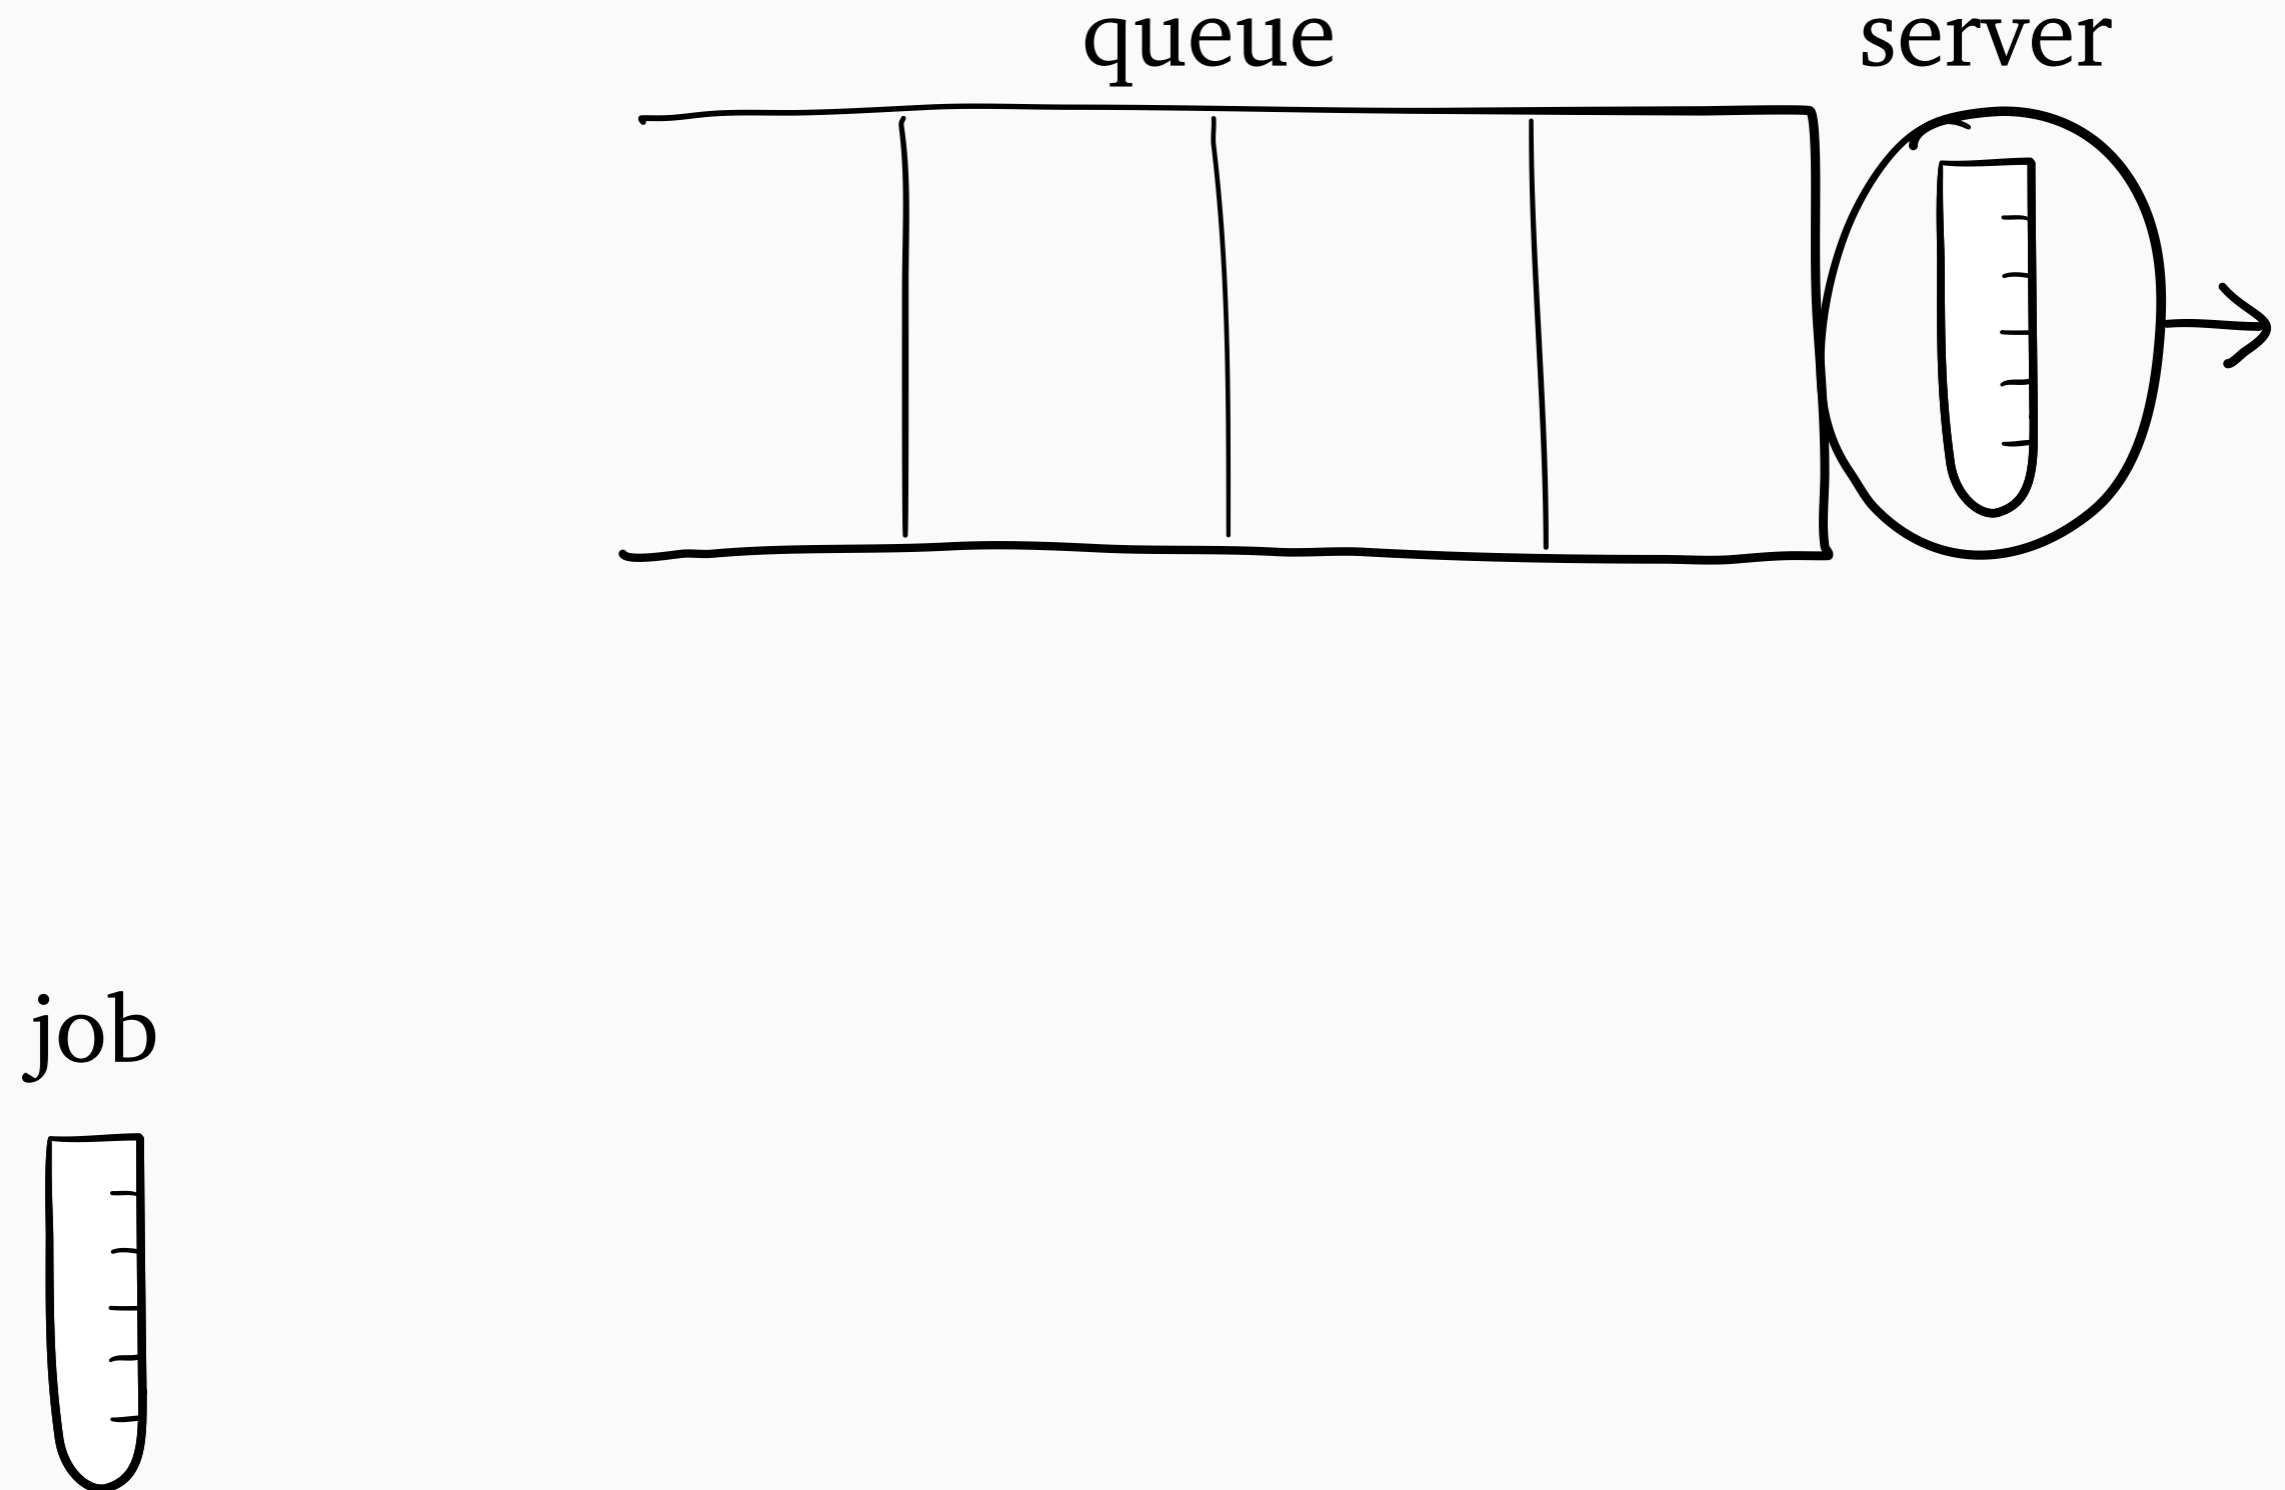
What are scheduling and delay?





What are scheduling and delay?





What are scheduling and delay?





What are scheduling and delay?




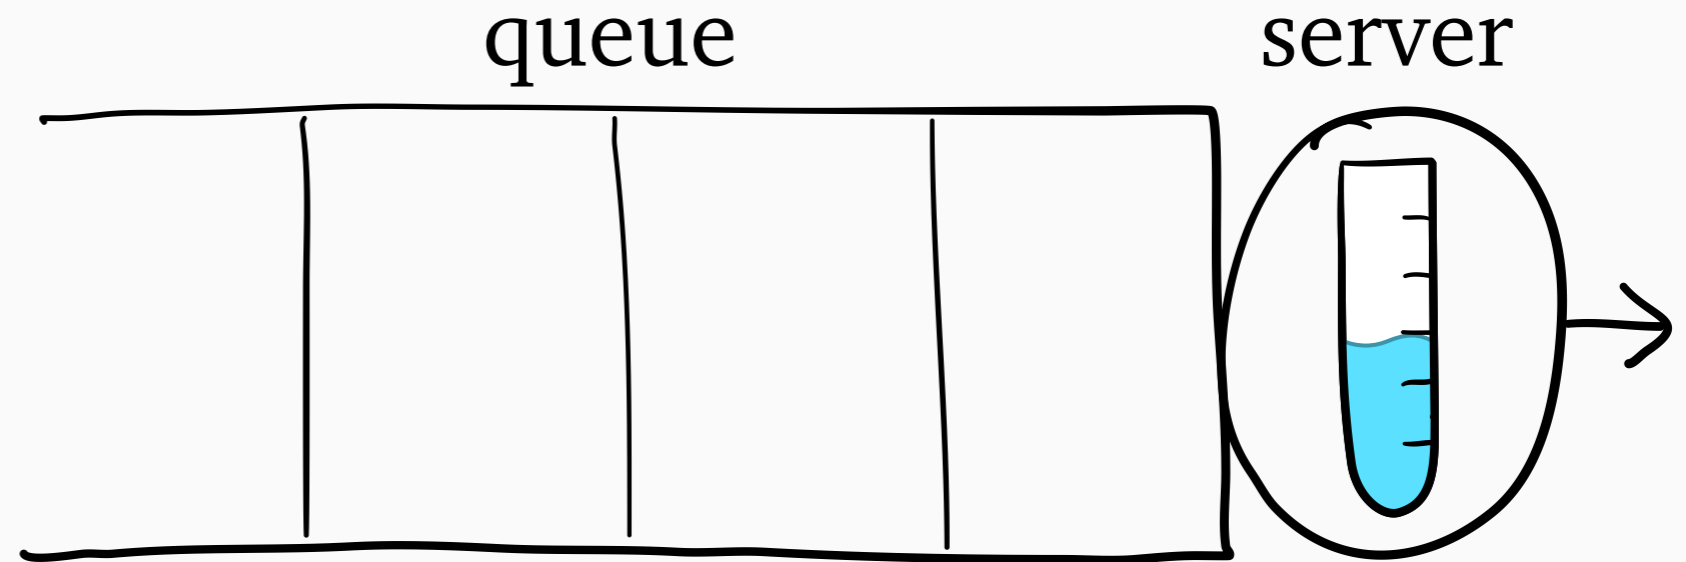
What are scheduling and delay?



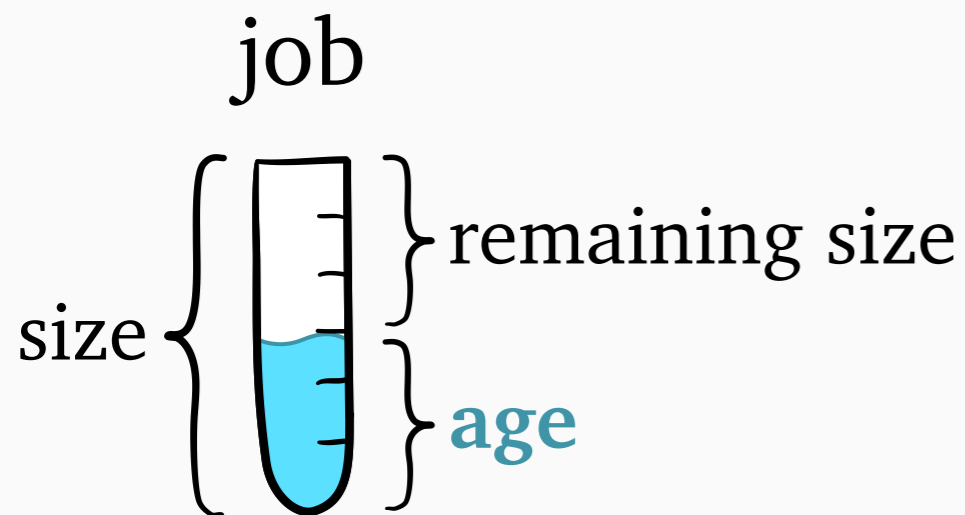

What are scheduling and delay?



What are scheduling and delay?




What are scheduling and delay?




What are scheduling and delay?


stochastic arrivals




queue



server




job




What are scheduling and delay?

stochastic arrivals




queue

server



job




What are scheduling and delay?


stochastic arrivals

queue

server



job

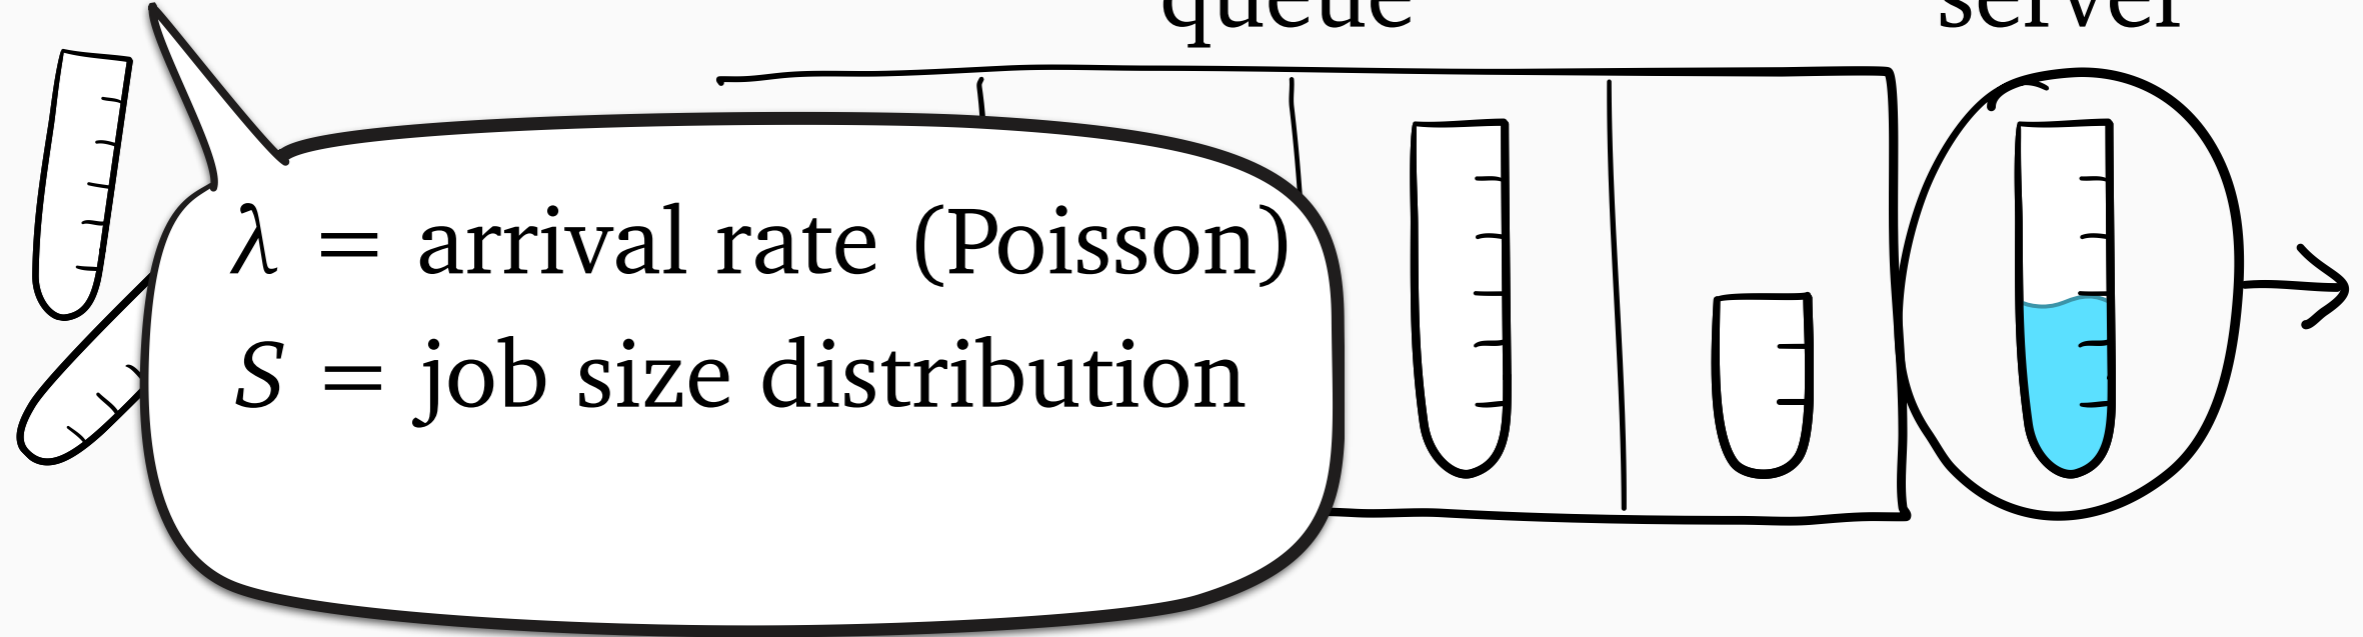


What are scheduling and delay?


stochastic arrivals

queue

server



job

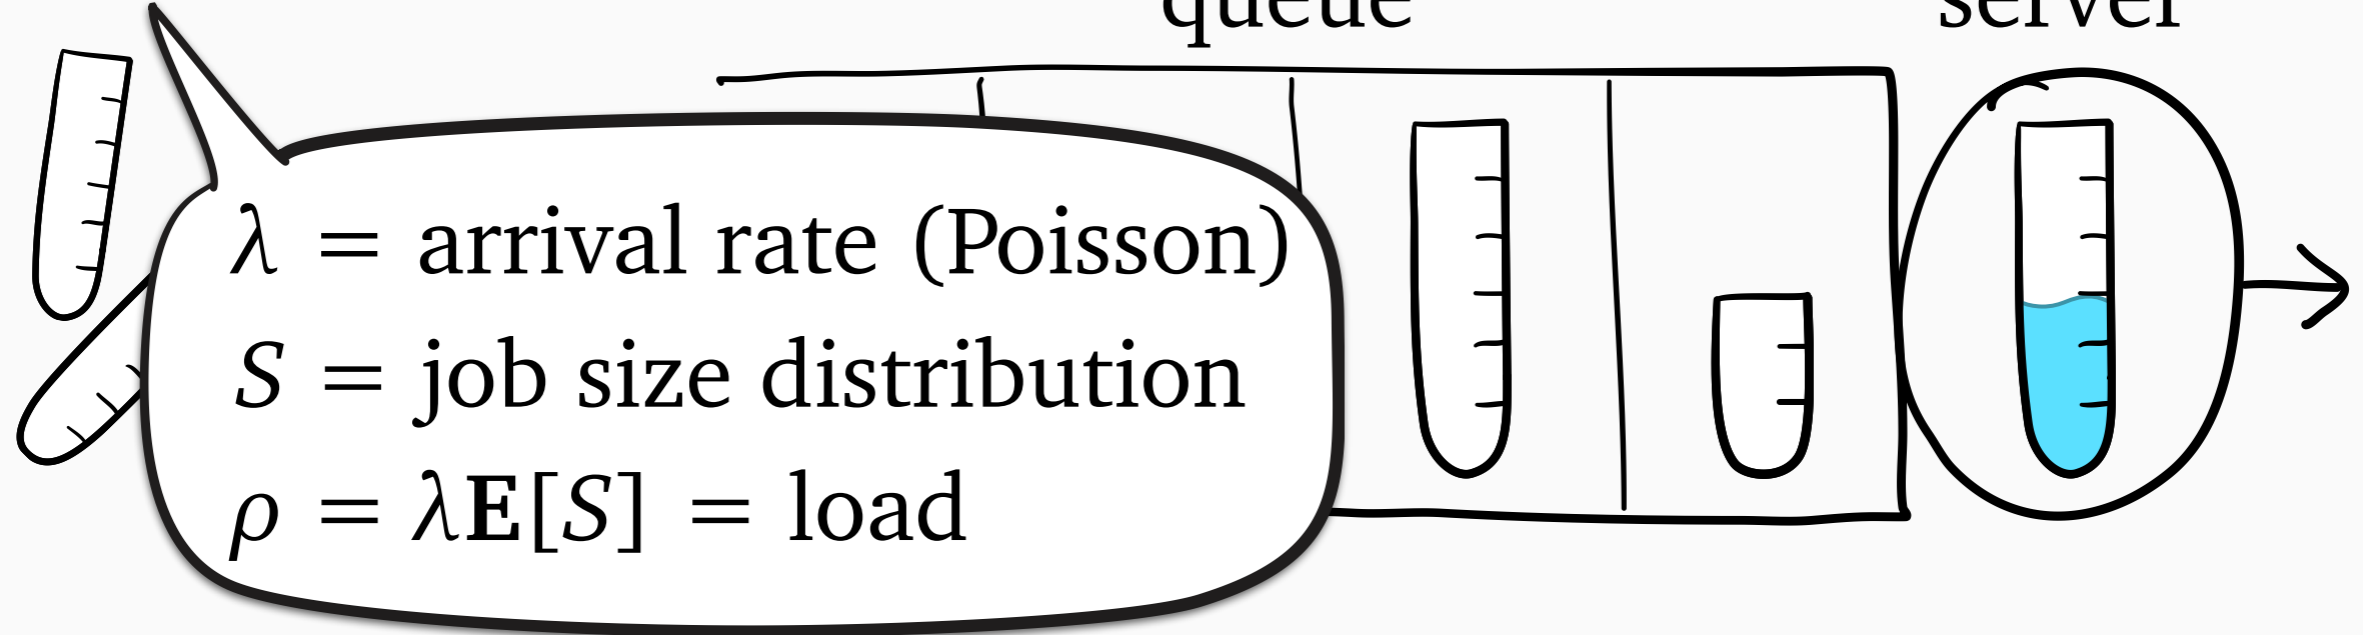


What are scheduling and delay?


stochastic arrivals

queue

server



job

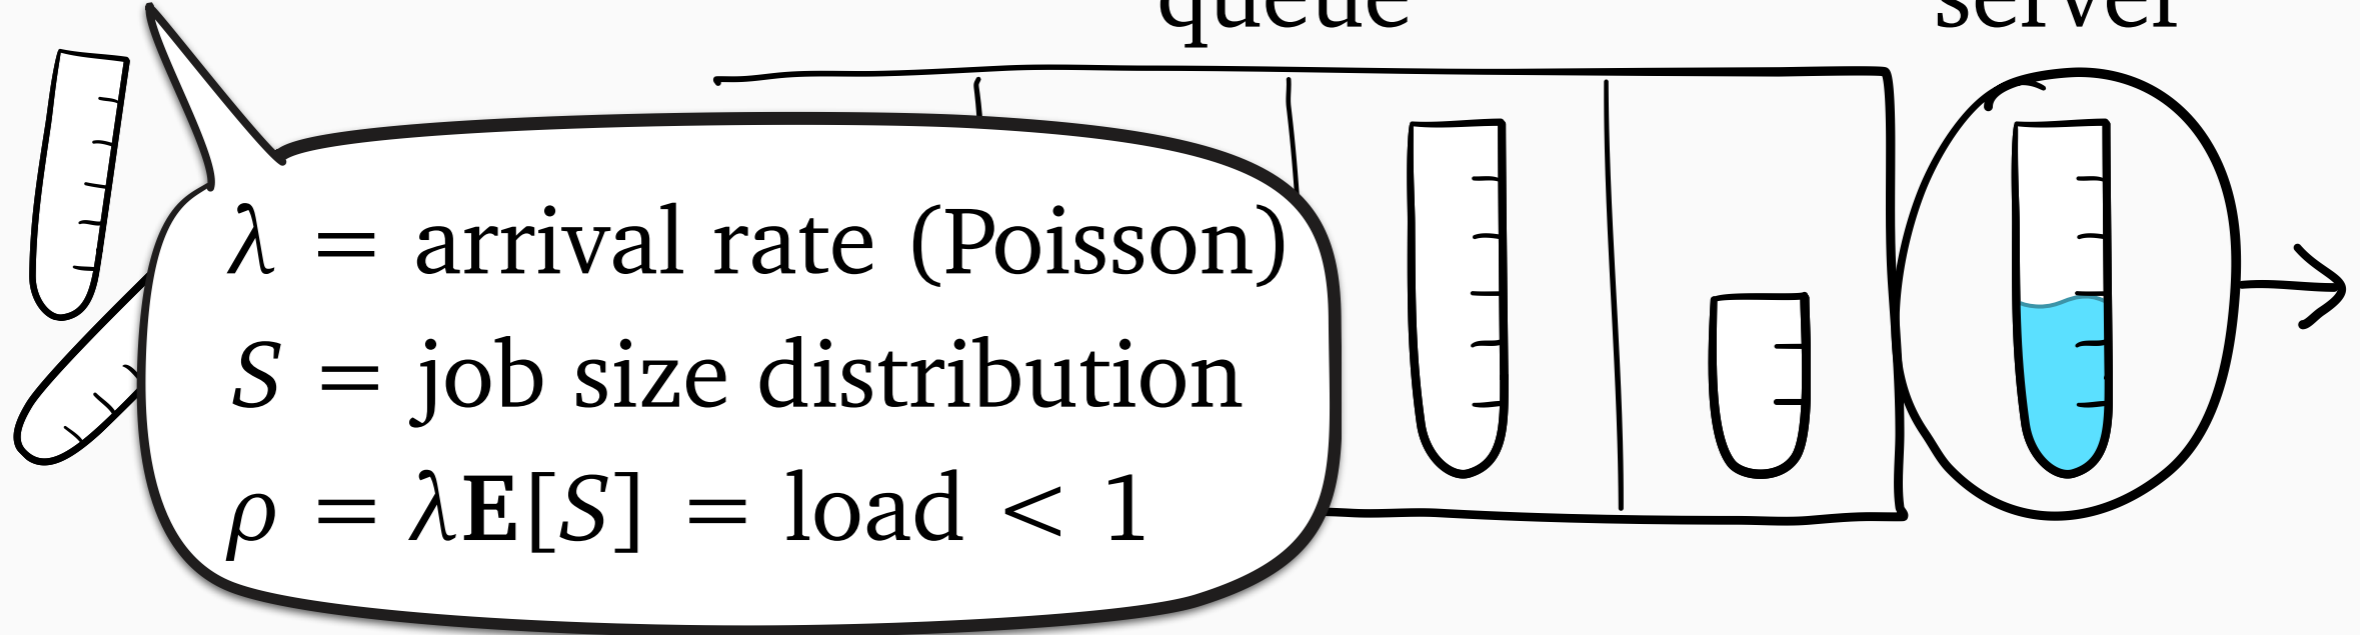


What are scheduling and delay?


stochastic arrivals

queue

server




job




What are scheduling and delay?


stochastic arrivals




queue



server




job




What are scheduling and delay?

stochastic arrivals




queue

server




job




What are scheduling and delay?


stochastic arrivals




queue



server




job




What are scheduling and delay?

stochastic arrivals




queue

server




job




What are scheduling and delay?

stochastic arrivals




queue

server




job




What are scheduling and delay?

stochastic arrivals




“delay” = *response time*

job




What are scheduling and delay?

stochastic arrivals




queue

server




“delay” = *response time*

job




What are scheduling and delay?

stochastic arrivals




queue

server



“delay” = *response time*


job



Scheduling policy:
decides which job to serve


What are scheduling and delay?

stochastic arrivals




queue

server



“delay” = *response time*


job



Scheduling policy:
decides which job to serve


What are scheduling and delay?

stochastic arrivals




queue

server



“delay” = *response time*


job



Scheduling policy:
decides which job to serve


What are scheduling and delay?

stochastic arrivals




queue

server



“delay” = *response time*


job



Scheduling policy:
decides which job to serve


What are scheduling and delay?

stochastic arrivals




queue

server



“delay” = *response time*


job



Scheduling policy:
decides which job to serve


What are scheduling and delay?

stochastic arrivals




queue

server



“delay” = *response time*


job



Scheduling policy:
decides which job to serve


What are scheduling and delay?

stochastic arrivals




queue

server



“delay” = *response time*

job




Scheduling policy:
decides which job to serve

Optimizing delay


Optimizing delay

response time




Optimizing delay

stochastic arrival
process λ, S




response time




Optimizing delay

stochastic arrival
process λ, S




response time
distribution T



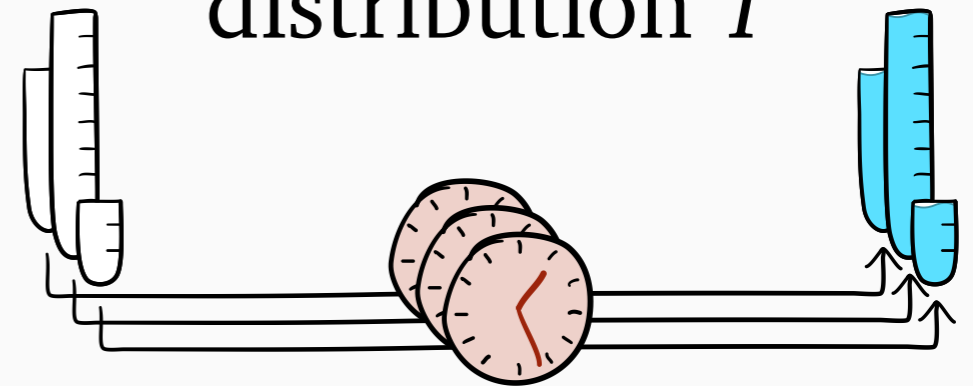
Optimizing delay

stochastic arrival
process λ, S




**scheduling
policy**

response time
distribution T




Optimizing delay

stochastic arrival
process λ, S



scheduling
policy


response time
distribution T



Goal: minimize *mean response time* $E[T]$


Optimizing delay

stochastic arrival
process λ, S



scheduling
policy

response time
distribution T




Goal: minimize *mean response time* $E[T]$

Warmup: if sizes known?


Optimizing delay

stochastic arrival
process λ, S



scheduling
policy

response time
distribution T




Goal: minimize *mean response time* $E[T]$

Warmup: if sizes known? **SRPT**


Optimizing delay

stochastic arrival
process λ, S



scheduling
policy

response time
distribution T




Goal: minimize *mean response time* $E[T]$


Warmup: if sizes known? **SRPT**

always serves job of
least remaining size

What job size noise model?



What job size noise model?



Model: (β, α) -bounded noise


true size s \Rightarrow **estimated** size $z \in [\beta s, \alpha s]$

What job size noise model?


below above

Model: (β, α) -bounded noise

true size s \Rightarrow **estimated** size $z \in [\beta s, \alpha s]$



What job size noise model?

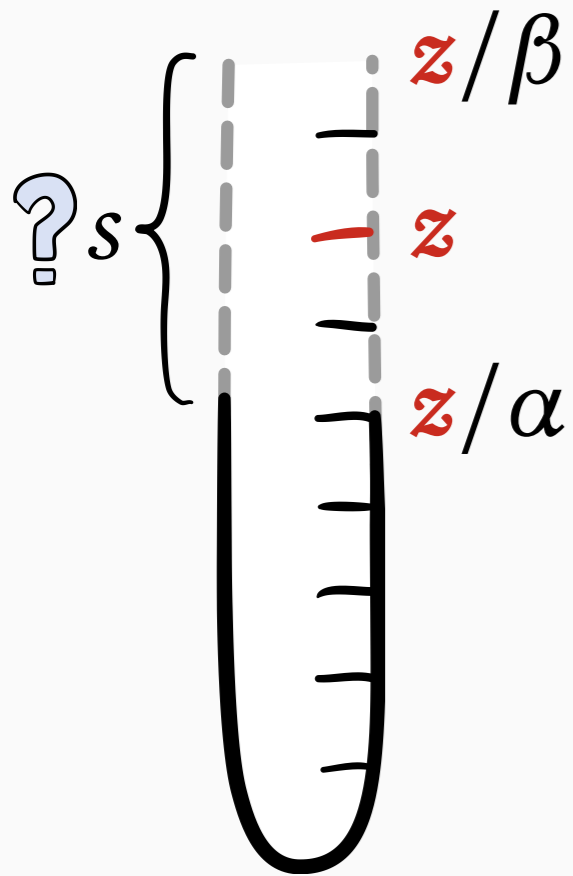


below above

Model: (β, α) -bounded noise

true size $s \Rightarrow$ **estimated** size $z \in [\beta s, \alpha s]$

What job size noise model?



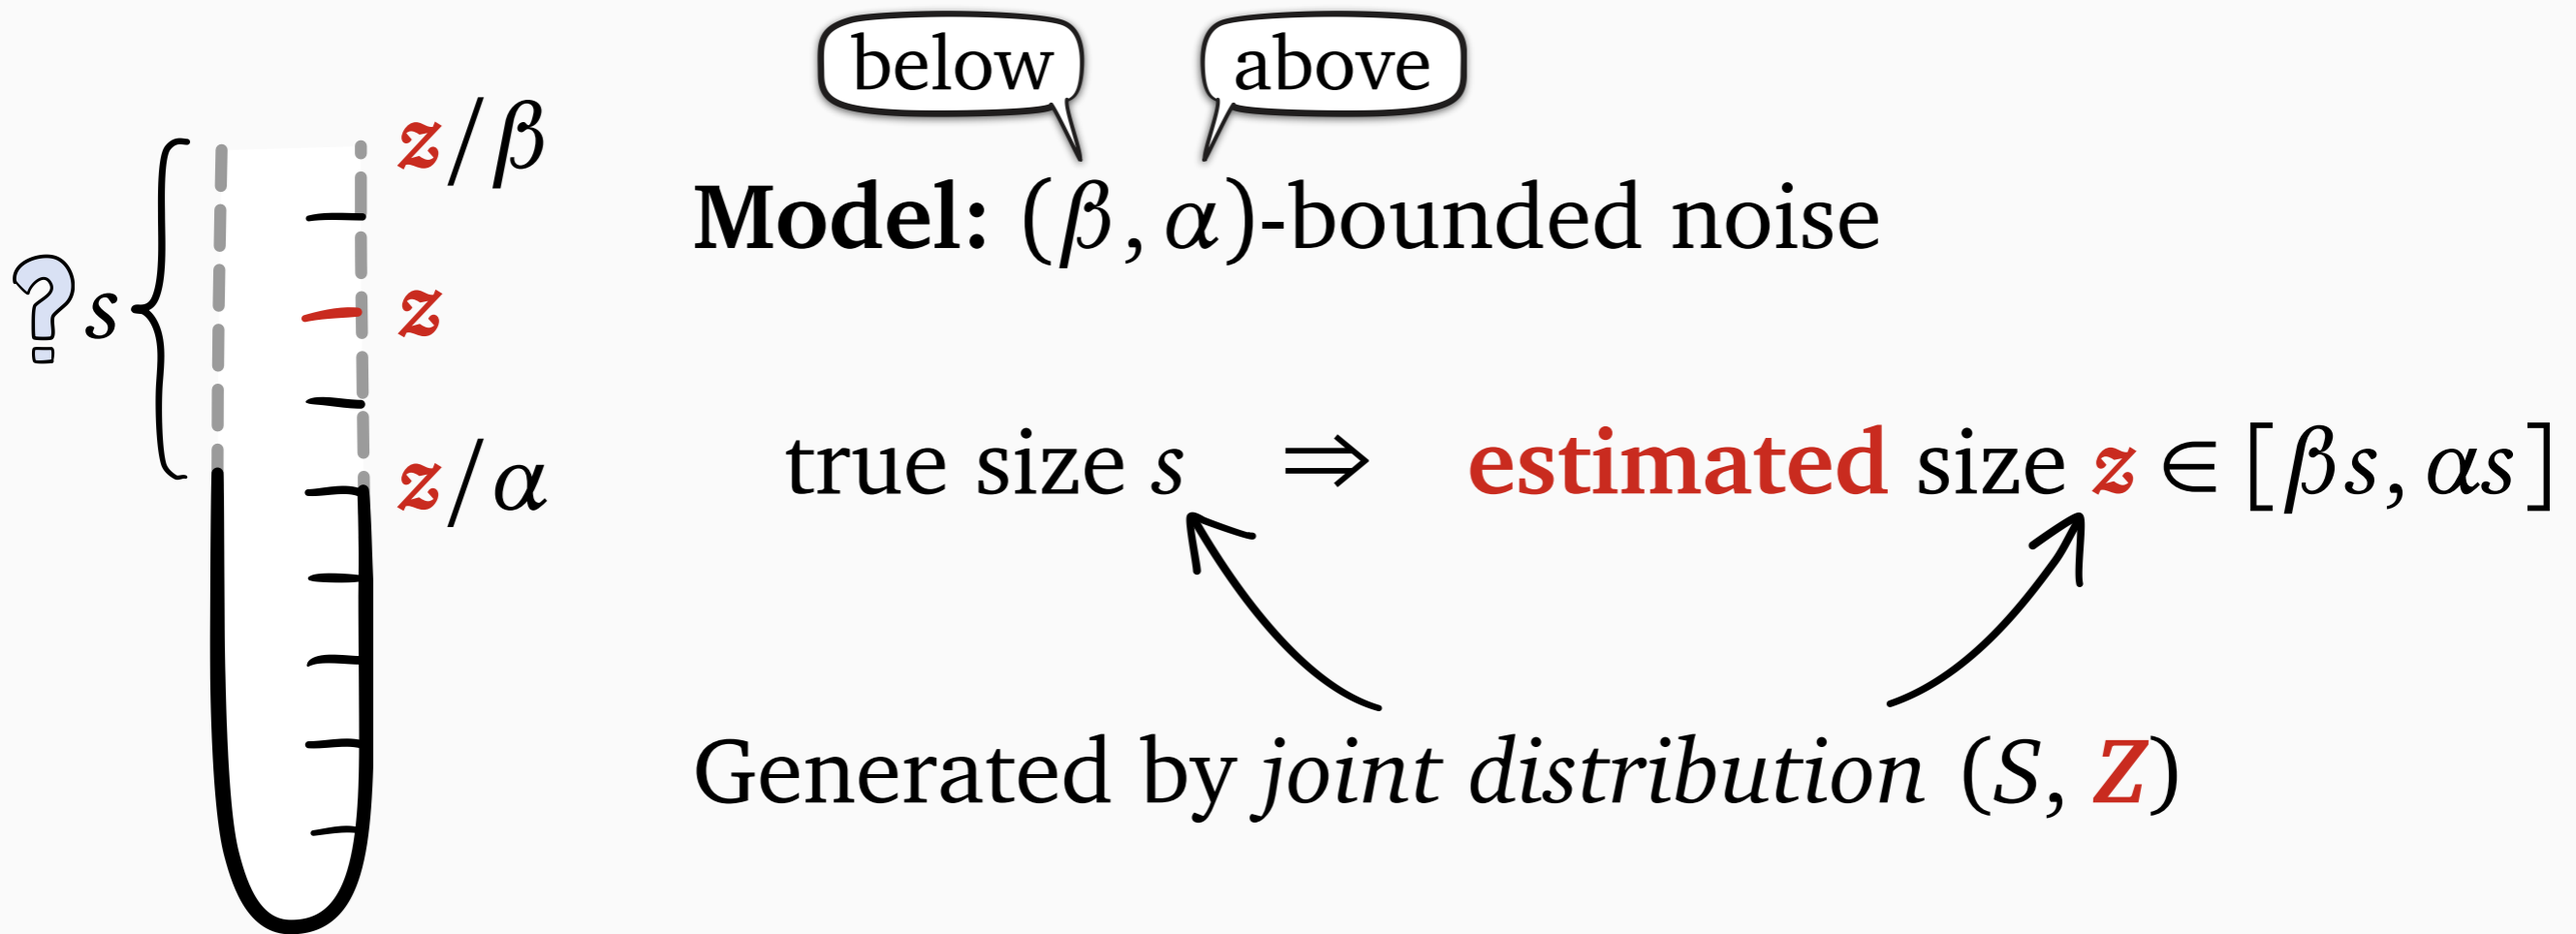
below above

Model: (β, α) -bounded noise

true size $s \Rightarrow$ **estimated** size $z \in [\beta s, \alpha s]$

Generated by *joint distribution* (S, Z)

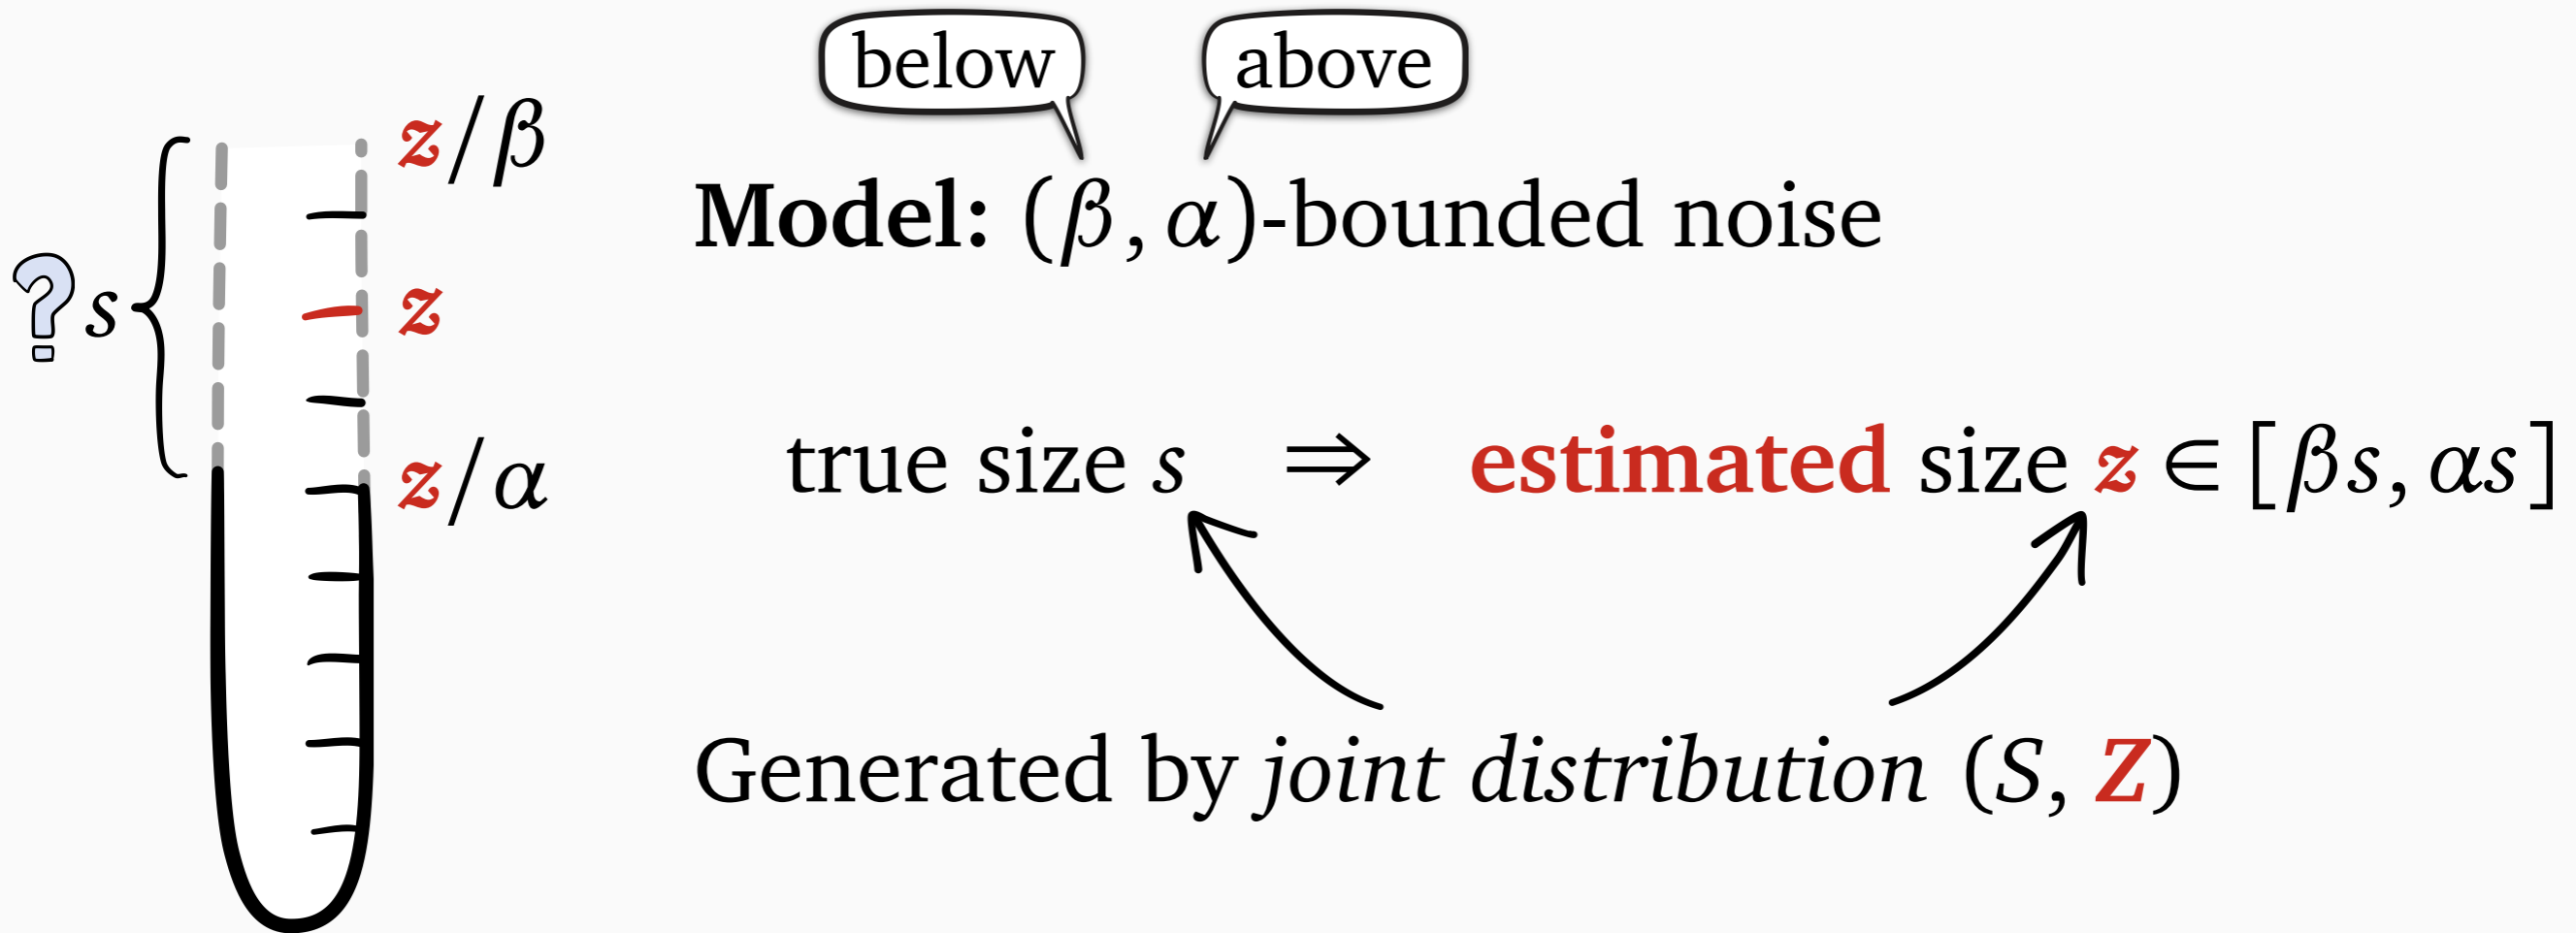
What job size noise model?



Goal: design a policy with “good” $E[T]$ for

- *any* joint distribution (S, \mathbf{Z})
- *any* values of α, β

What job size noise model?

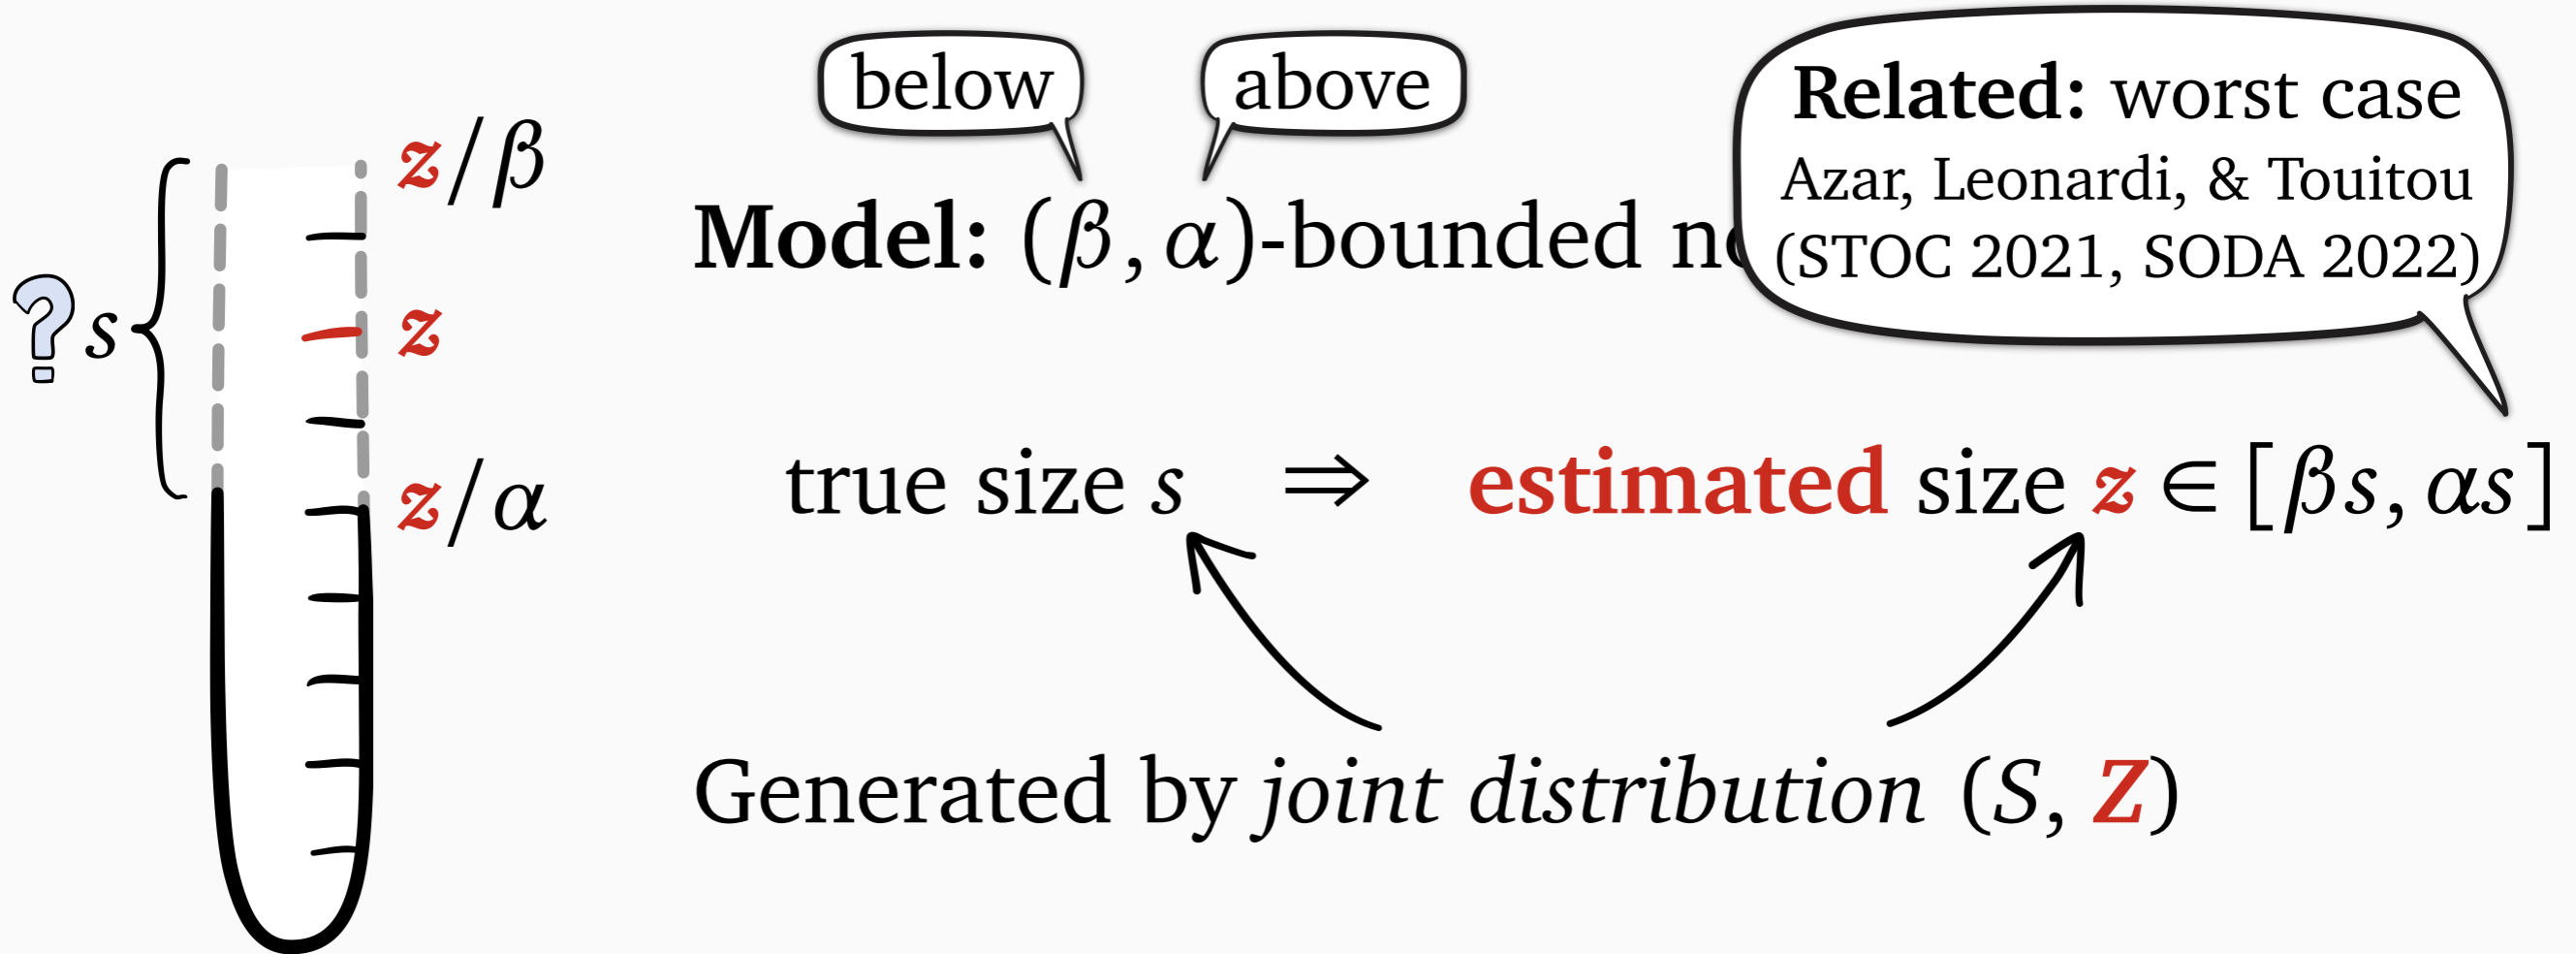


Goal: design a policy with “good” $E[T]$ for

- *any* joint distribution (S, Z)
- *any* values of α, β

semi-adversarial

What job size noise model?



Goal: design a policy with “good” $E[T]$ for

- *any* joint distribution (S, Z)
- *any* values of α, β

semi-adversarial

What can we hope to achieve?

What can we hope to achieve?

C -consistent: $\frac{\mathbb{E}[T_P]}{\mathbb{E}[T_{\text{SRPT}}]} \rightarrow C$ as $\alpha, \beta \rightarrow 1$

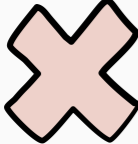
What can we hope to achieve?

C -consistent: $\frac{\mathbf{E}[T_P]}{\mathbf{E}[T_{\text{SRPT}}]} \rightarrow C$ as $\alpha, \beta \rightarrow 1$

R -robust: $\frac{\mathbf{E}[T_P]}{\mathbf{E}[T_{\text{SRPT}}]} \leq R$ for all α, β

What can we hope to achieve?

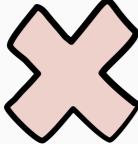
C -consistent: $\frac{\mathbb{E}[T_P]}{\mathbb{E}[T_{\text{SRPT}}]} \rightarrow C$ as $\alpha, \beta \rightarrow 1$

 R -robust: $\frac{\mathbb{E}[T_P]}{\mathbb{E}[T_{\text{SRPT}}]} \leq R$ for all α, β

What can we hope to achieve?

C -consistent: $\frac{\mathbb{E}[T_P]}{\mathbb{E}[T_{\text{SRPT}}]} \rightarrow C$ as $\alpha, \beta \rightarrow 1$

G -graceful: $\frac{\mathbb{E}[T_P]}{\mathbb{E}[T_{\text{SRPT}}]} \leq G \cdot \frac{\alpha}{\beta}$ for all α, β

 R -robust: $\frac{\mathbb{E}[T_P]}{\mathbb{E}[T_{\text{SRPT}}]} \leq R$ for all α, β

What can we hope to achieve?

✓ C -consistent: $\frac{\mathbb{E}[T_P]}{\mathbb{E}[T_{\text{SRPT}}]} \rightarrow C$ as $\alpha, \beta \rightarrow 1$

✓ G -graceful: $\frac{\mathbb{E}[T_P]}{\mathbb{E}[T_{\text{SRPT}}]} \leq G \cdot \frac{\alpha}{\beta}$ for all α, β

✗ R -robust: $\frac{\mathbb{E}[T_P]}{\mathbb{E}[T_{\text{SRPT}}]} \leq R$ for all α, β

Our contribution: *first policy P*
that's consistent and graceful

- $G = 3.5$
- $C = 1$

What can we hope to achieve?

✓ C -consistent: $\frac{\mathbb{E}[T_P]}{\mathbb{E}[T_{\text{SRPT}}]} \rightarrow C$ as $\alpha, \beta \rightarrow 1$

✓ G -graceful: $\frac{\mathbb{E}[T_P]}{\mathbb{E}[T_{\text{SRPT}}]} \leq G \cdot \frac{\alpha}{\beta}$ for all α, β

✗ R -robust: $\frac{\mathbb{E}[T_P]}{\mathbb{E}[T_{\text{SRPT}}]} \leq R$ for all α, β

Our contribution: *first policy* P 
that's consistent and graceful

- $G = 3.5$
- $C = 1$



What can we hope to achieve?

✓ C -consistent: $\frac{\mathbb{E}[T_P]}{\mathbb{E}[T_{\text{SRPT}}]} \rightarrow C$ as $\alpha, \beta \rightarrow 1$

✓ G -graceful: $\frac{\mathbb{E}[T_P]}{\mathbb{E}[T_{\text{SRPT}}]} \leq G \cdot \frac{\alpha}{\beta}$ for all α, β

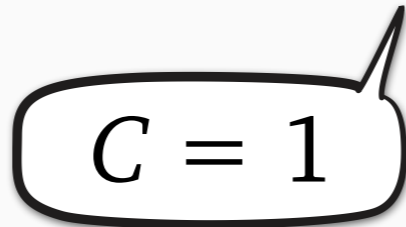
✗ R -robust: $\frac{\mathbb{E}[T_P]}{\mathbb{E}[T_{\text{SRPT}}]} \leq R$ for all α, β

Our contribution: first policy P 
that is impossible in worst case
Azar, Leonardi, & Touitou
(STOC 2021) graceful

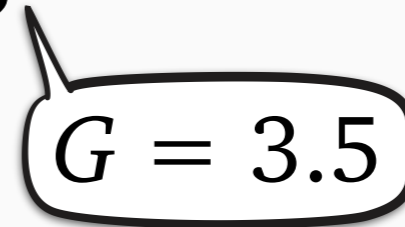
- $G = 5.5$
- $C = 1$



Our contribution: *first policy* P
that's consistent and graceful

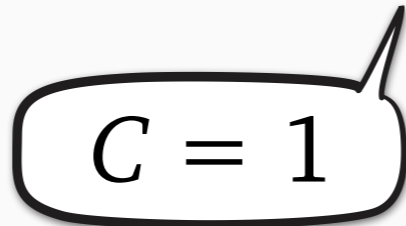


$C = 1$

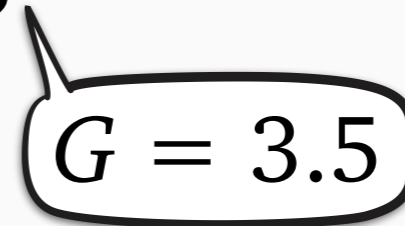


$G = 3.5$

Our contribution: *first policy* **P**
that's consistent and graceful



$C = 1$



$G = 3.5$



What is the new policy?

Our contribution: *first policy* P
that's consistent and graceful

$$C = 1$$

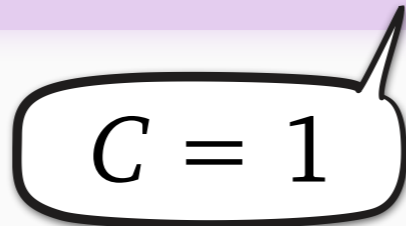
$$G = 3.5$$

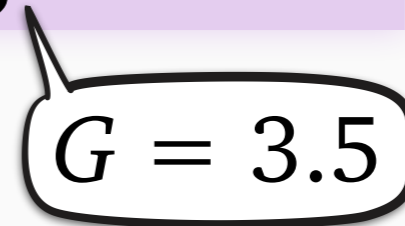


What is the new policy?

How do we bound its performance?

Our contribution: *first* policy P
that's consistent and graceful


$$C = 1$$


$$G = 3.5$$




What does “policy” mean?

What is the new policy?

How do we bound its performance?

Scheduling with **rank** functions




Scheduling with **rank** functions


rank




age a




Scheduling with **rank** functions




Scheduling with **rank** functions



Scheduling with **rank** functions




Scheduling with **rank** functions






Policy design space:
rank functions

What's the right **rank** function?



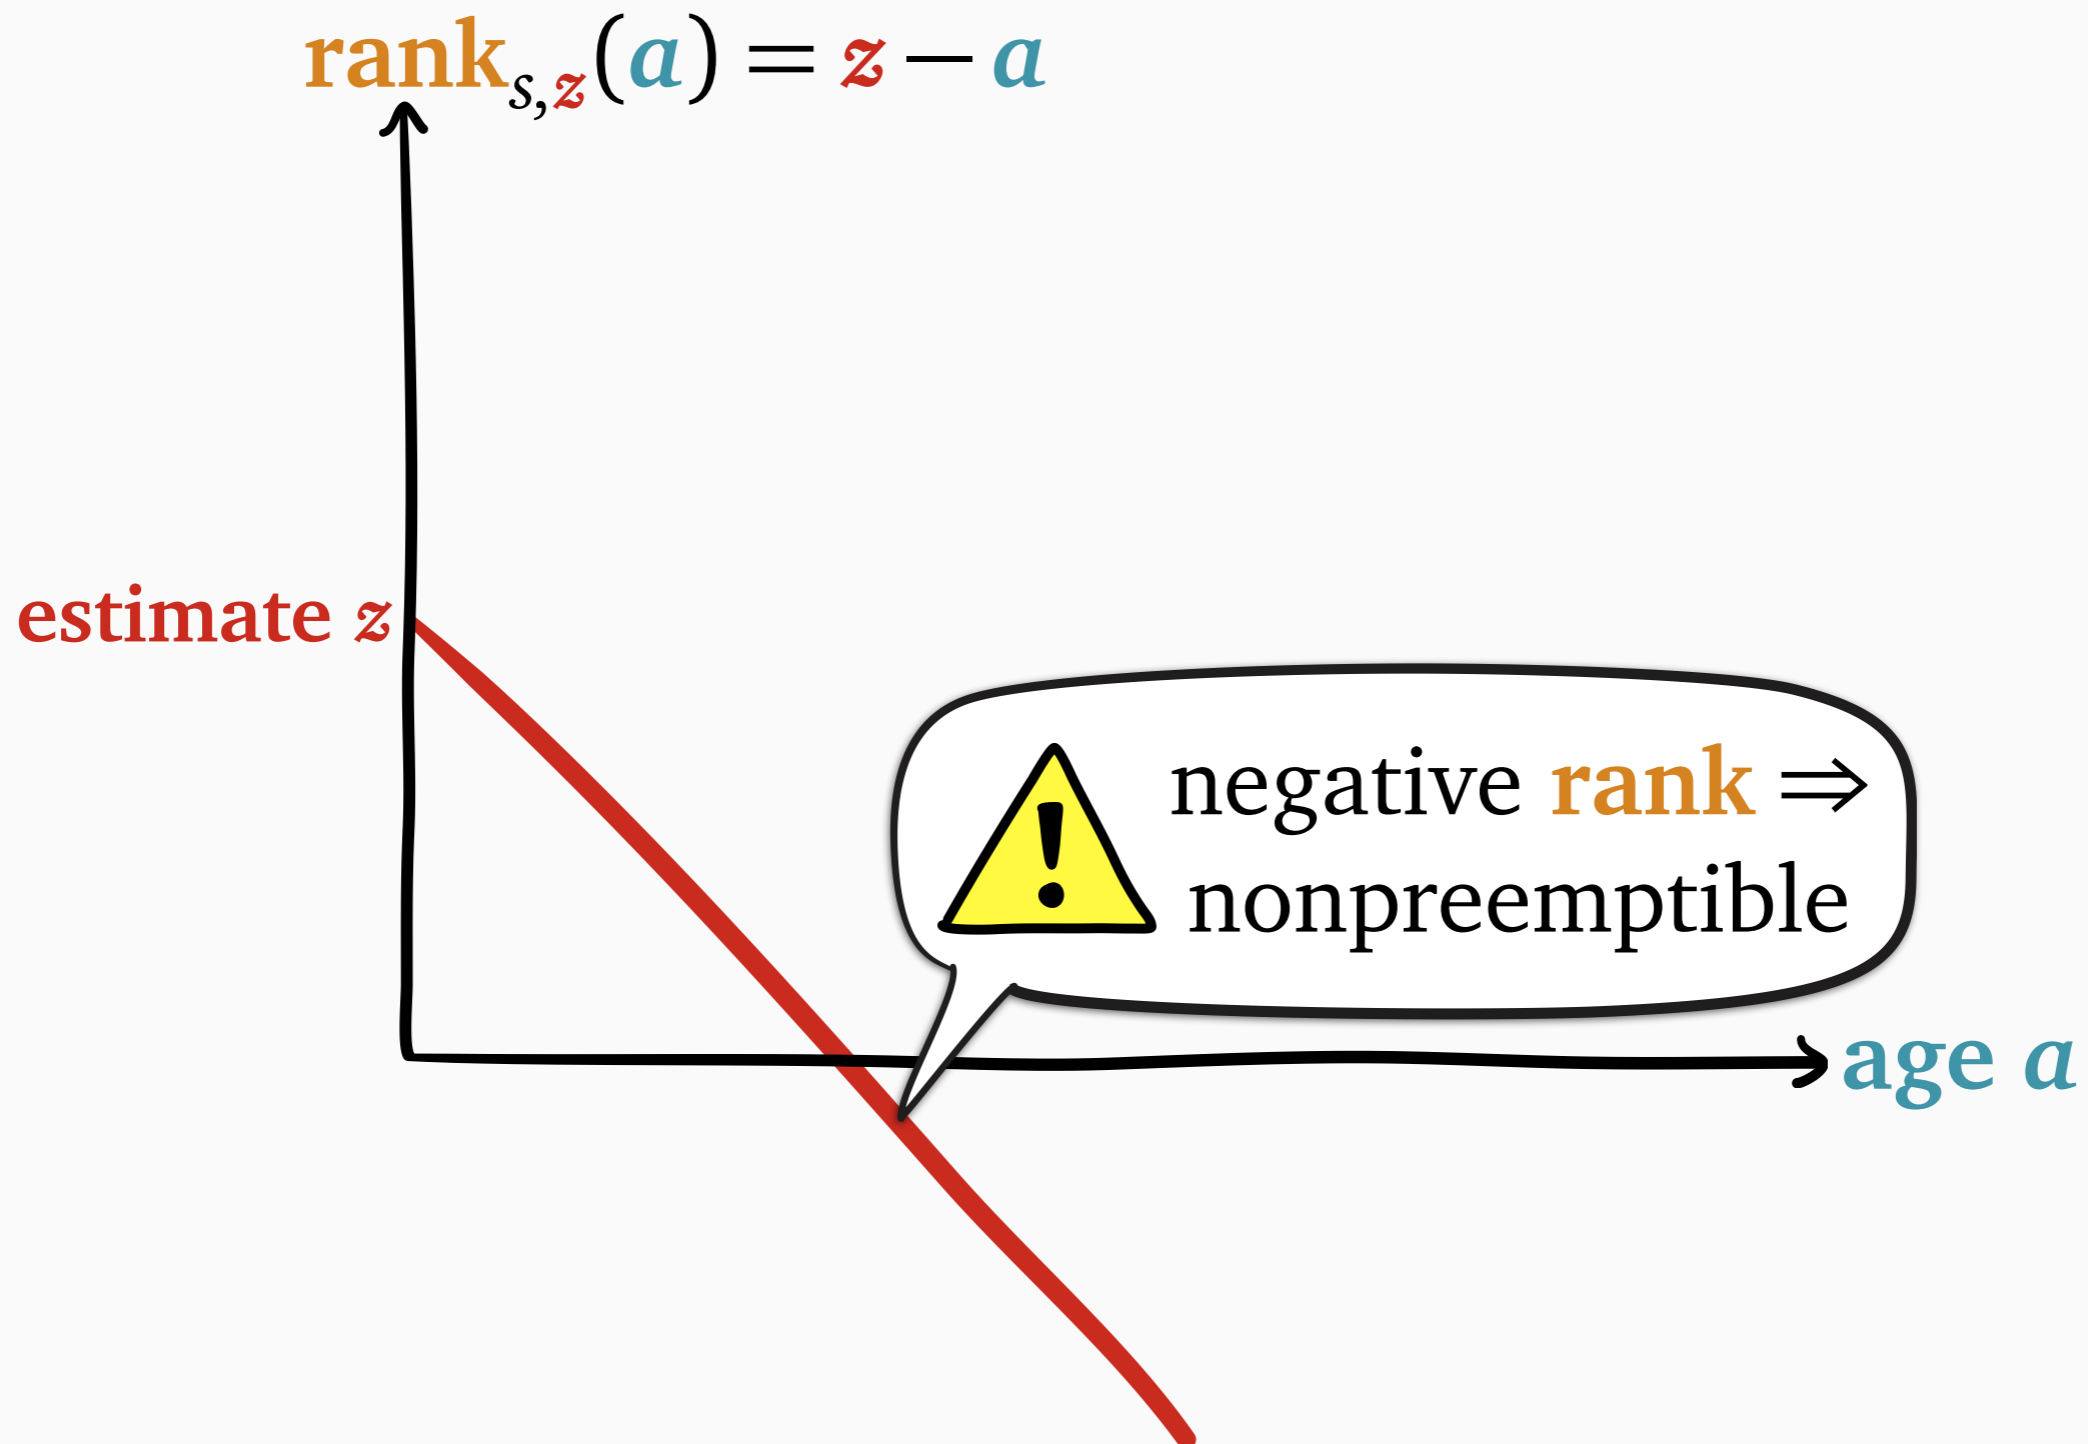
What's the right **rank** function?

Naive





What's the right **rank** function?

Naive



What's the right **rank** function?




What's the right **rank** function?

Checkmark


$$\text{rank}_{s,z}(a) = |z - a|$$

estimate z




Not these...

Naive




What's the right **rank** function?

Checkmark



Not these...

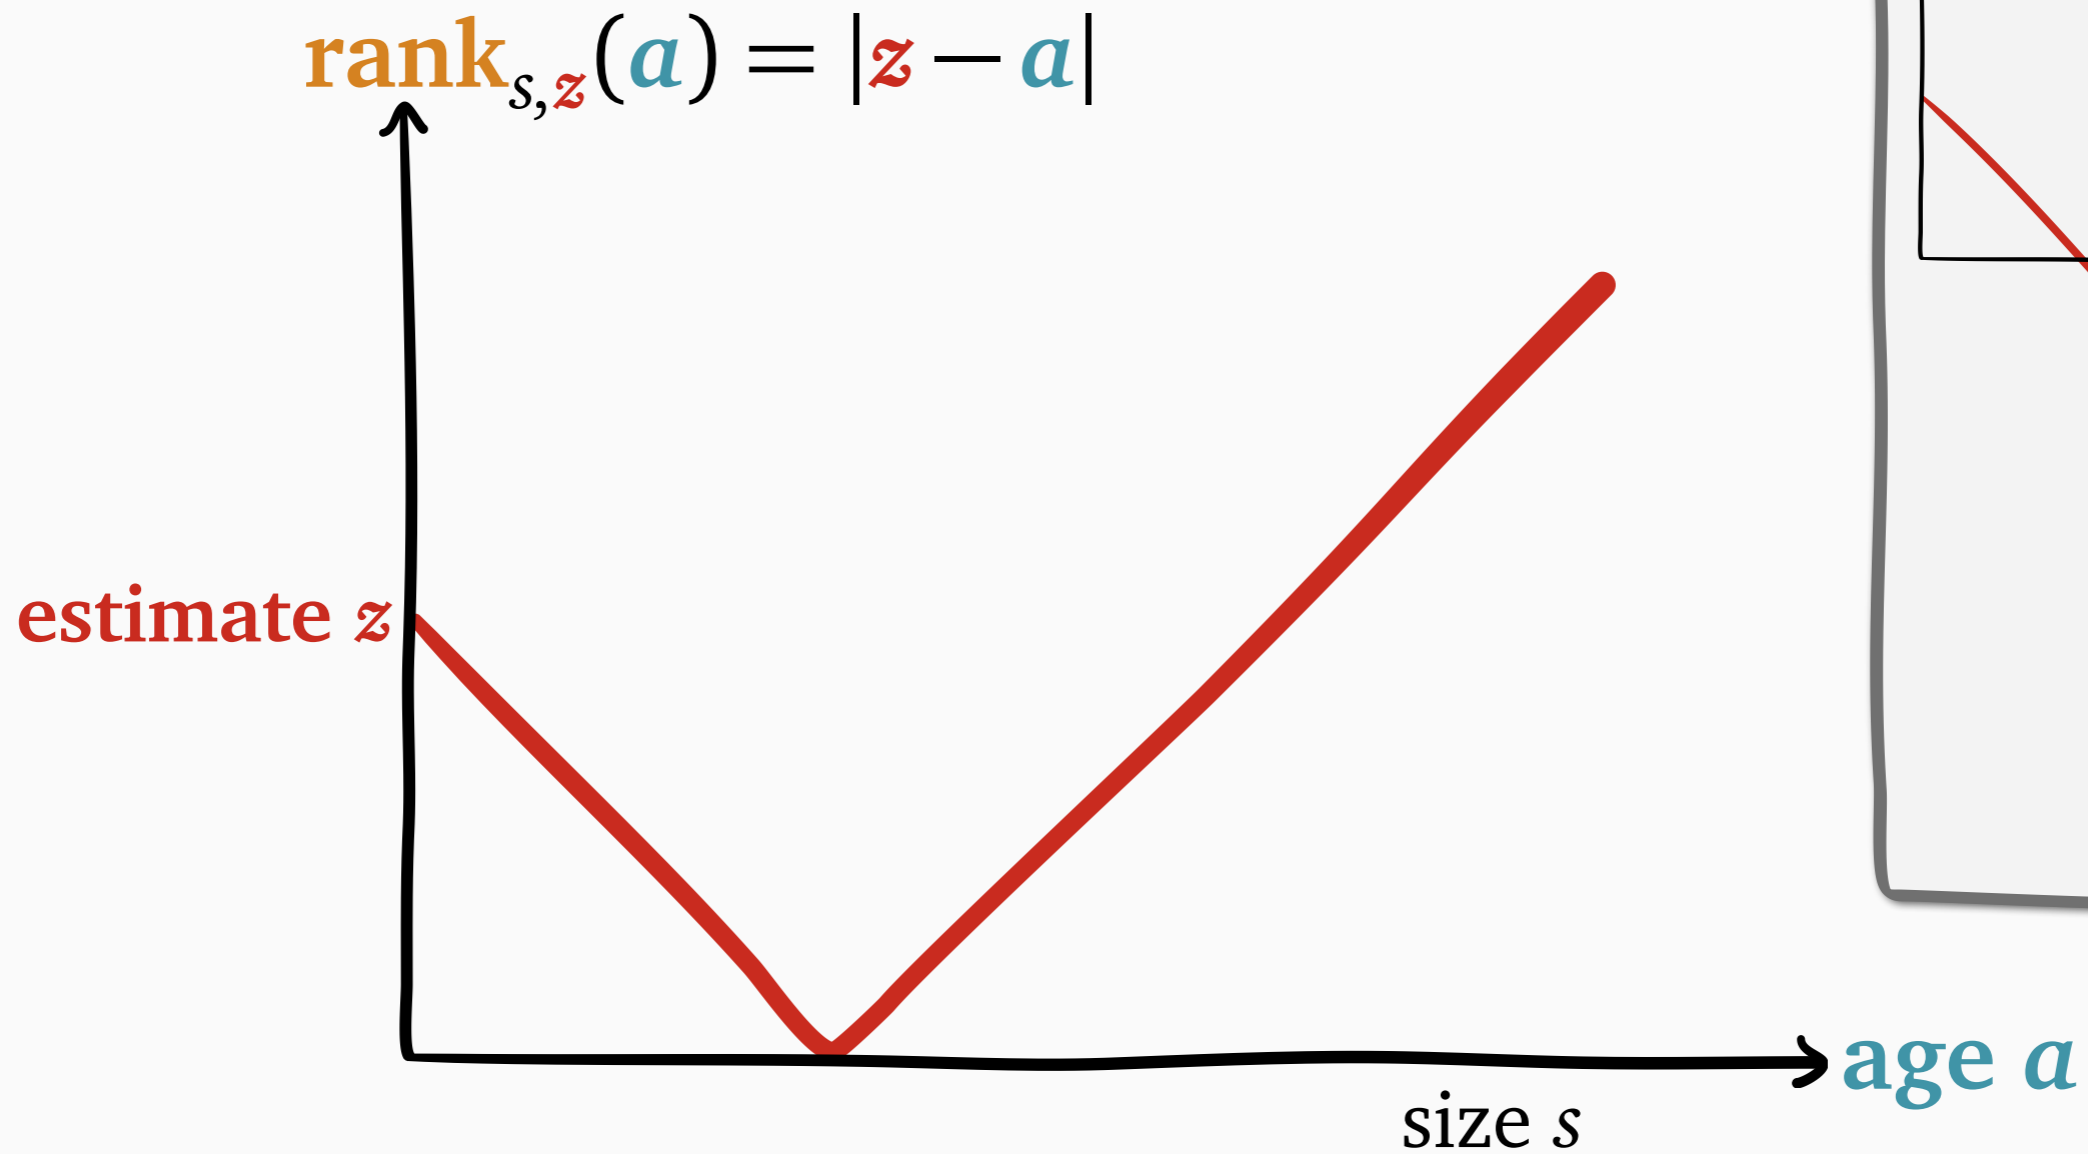
Naive



What if $\beta < \frac{1}{2}$?


What's the right **rank** function?

Checkmark



Not these...


Naive



What if $\beta < \frac{1}{2}$?


What's the right **rank** function?

Checkmark



Not these...

Naive



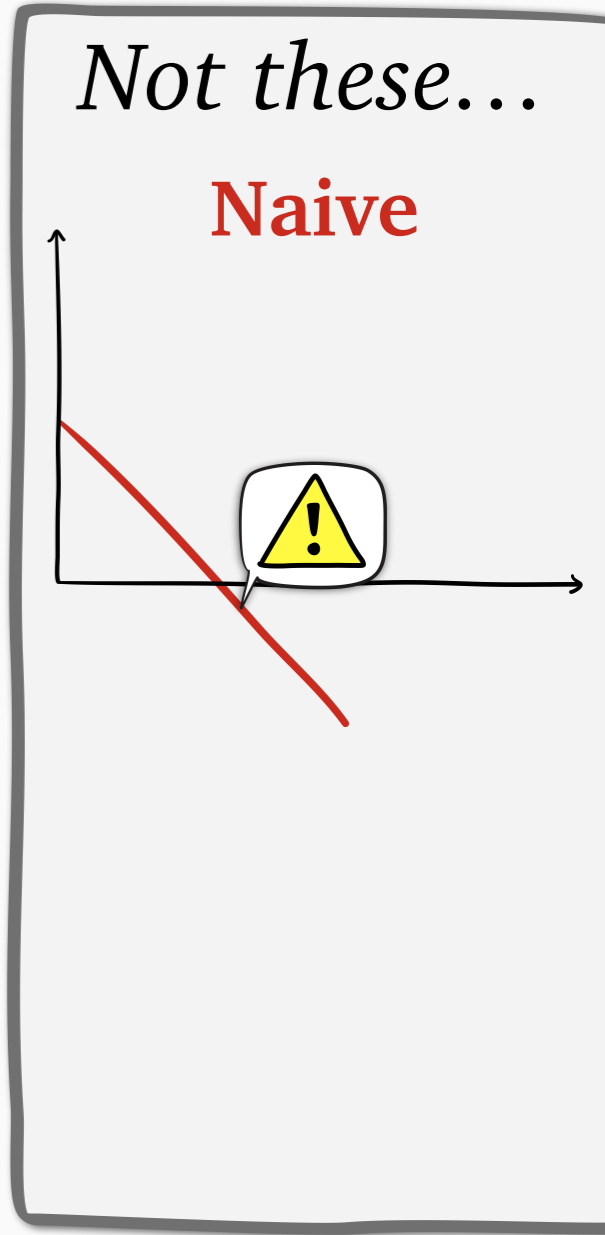

What if $\beta < \frac{1}{2}$?

What's the right **rank** function?

Checkmark



$$\text{rank}_{s,z}(a) = |z - a|$$

new worst **rank** \Rightarrow
preemption likely 



What if $\beta < \frac{1}{2}$?


What's the right **rank** function?



What's the right **rank** function?


Radical

$$\text{rank}_{s,z}(a) = \min\{|z - a|, z\}$$




Not these...

Naive




Checkmark



What's the right **rank** function?


Radical

$$\text{rank}_{s,z}(a) = \min\{|z - a|, z\}$$




Not these...

Naive



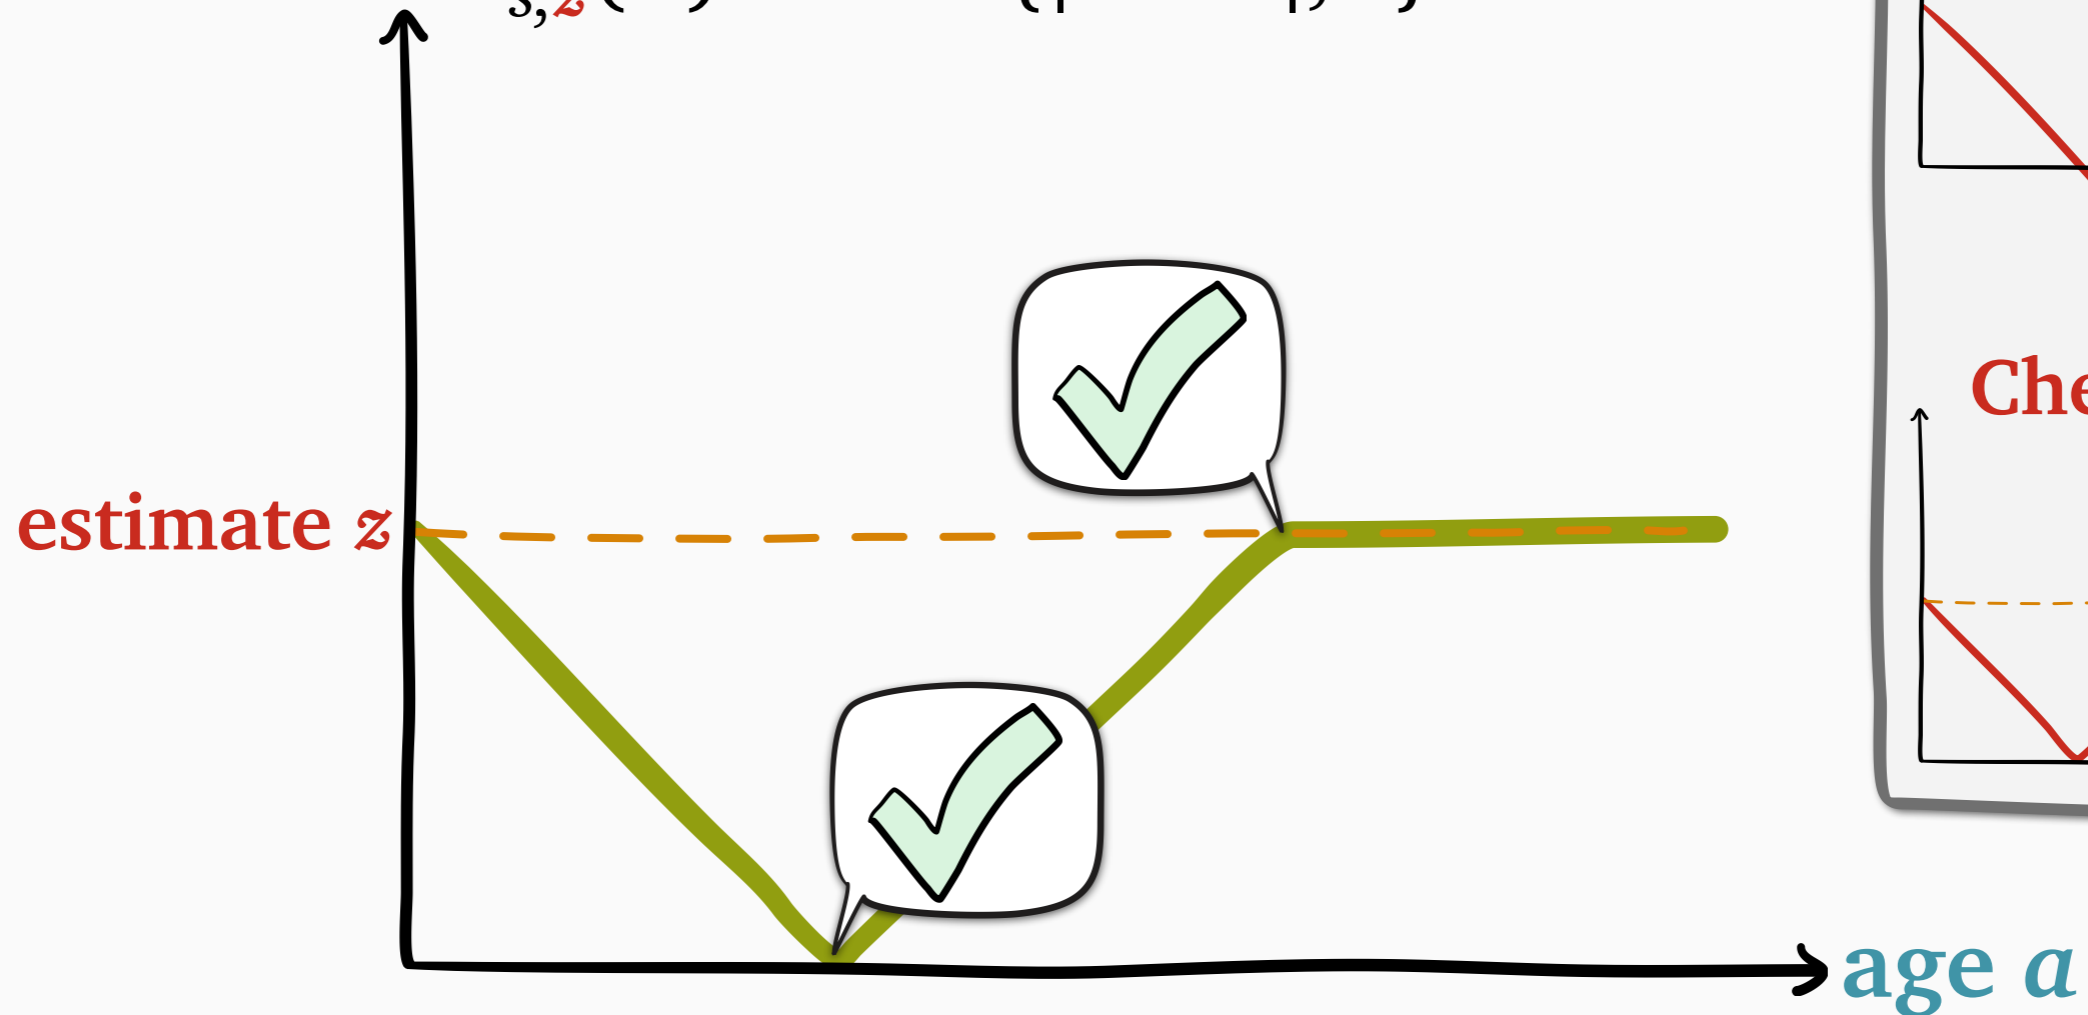
Checkmark



What's the right **rank** function?


Radical

$$\text{rank}_{s,z}(a) = \min\{|z - a|, z\}$$




Not these...

Naive




Checkmark



What's the right **rank** function?

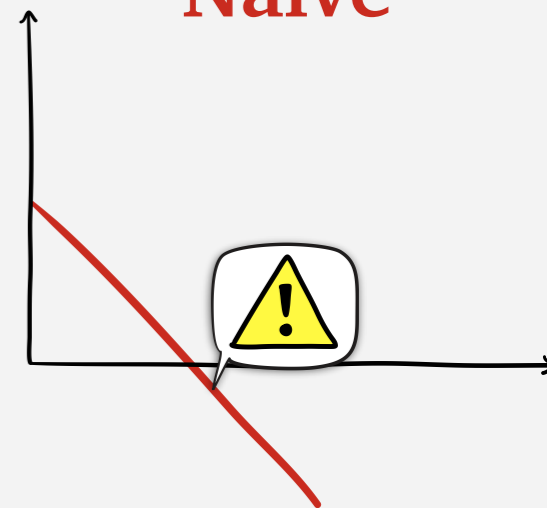
Radical

$$\text{rank}_{s,z}(a) = \min\{|z - a|, z\}$$




Not these...



Naive










Checkmark












Theorem: **Radical** is 1-consistent, 3.5-graceful


















Scully, Harchol-Balter,
& Scheller-Wolf
(SIGMETRICS 2018)





Scully, Harchol-Balter,
& Scheller-Wolf
(SIGMETRICS 2018)

rank **Scale**



WINE

Scully, Grosf,
& Harchol-Balter
(SIGMETRICS 2021)




Schedule **O**rdered by **A**ge-based **P**riority




SOAP

Schedule **O**rdered by **A**ge-based **P**riority

stochastic arrival
process λ , (S , Z)



any **rank** function








SOAP

Schedule **O**rdered by **A**ge-based **P**riority


stochastic arrival
process λ , (S , Z)











any **rank** function




response time
distribution T








rank functions
close enough



SOAP



WINE




WINE

Work Integral Number Equality



WINE

Work Integral Number Equality




number of jobs N



WINE

Work Integral Number Equality




number of jobs N

Little's Law:
 $E[T] = E[N]/\lambda$





WINE

Work Integral Number Equality



r -work $W(r)$




number of jobs N

Little's Law:
 $E[T] = E[N]/\lambda$





WINE

Work Integral Number Equality



r -work $W(r)$



number of jobs N

Little's Law:
 $E[T] = E[N]/\lambda$




What is r -work?


Get N from r -work?

Bound **Scale's** $E[T]$?


What is *r*-work?



What is *r*-work?



What is r -work?




Definition:


$W(r)$ = total remaining size of jobs
whose remaining size is $\leq r$

r -work

What is r -work?




Definition:




$W(r)$ = total remaining size of jobs
whose remaining size is $\leq r$

r -work

What is r -work?




Definition:



$W(r)$ = total remaining size of jobs
whose remaining size is $\leq r$

r -work

What is r -work?



Definition:

$W(r)$ = total remaining size of jobs
whose remaining size is $\leq r$

r -work



WINE: under *any* scheduling policy,

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



WINE: under *any* scheduling policy,

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



Lemma:

$$\mathbf{E}[W_{\text{Scale}}(r)] \leq \mathbf{E}[W_{\text{SRPT}}\left(\frac{\alpha}{\beta}r\right)]$$



WINE: under *any* scheduling policy,

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



Lemma:

$$\mathbf{E}[W_{\text{Scale}}(r)] \leq \mathbf{E}[W_{\text{SRPT}}(\frac{\alpha}{\beta}r)]$$



Theorem:

$$\mathbf{E}[T_{\text{Scale}}] \leq \frac{\alpha}{\beta} \mathbf{E}[T_{\text{SRPT}}]$$

also holds in
worst case



WINE: under *any* scheduling policy,

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



Lemma:

$$\mathbf{E}[W_{\text{Scale}}(r)] \leq \mathbf{E}[W_{\text{SRPT}}(\frac{\alpha}{\beta}r)]$$



Theorem:

$$\mathbf{E}[T_{\text{Scale}}] \leq \frac{\alpha}{\beta} \mathbf{E}[T_{\text{SRPT}}]$$

also holds in
worst case



WINE: under *any* scheduling policy,

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



Lemma:






$$\mathbf{E}[W_{\text{Scale}}(r)] \leq \mathbf{E}[W_{\text{SRPT}}\left(\frac{\alpha}{\beta}r\right)]$$







Theorem:

also holds with
“noisy scaling”

$$\mathbf{E}[T_{\text{Scale}}] \leq \frac{\alpha}{\beta} \mathbf{E}[T_{\text{SRPT}}]$$




WINE







Scale




WINE




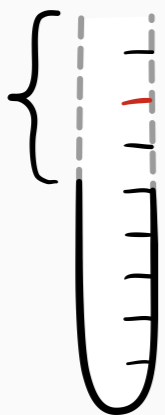


Scale

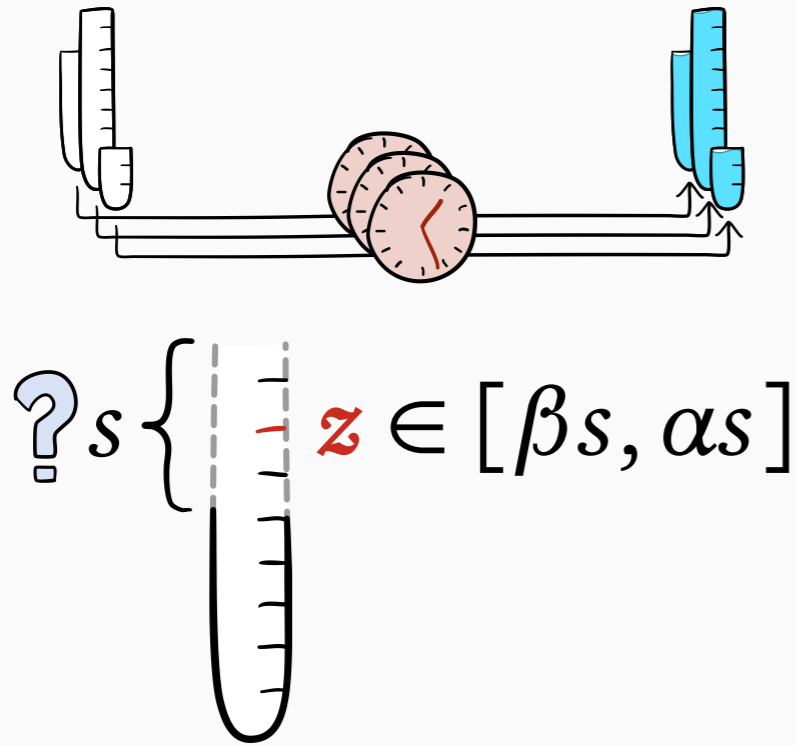


Problem: minimize $E[T]$
with noisy information

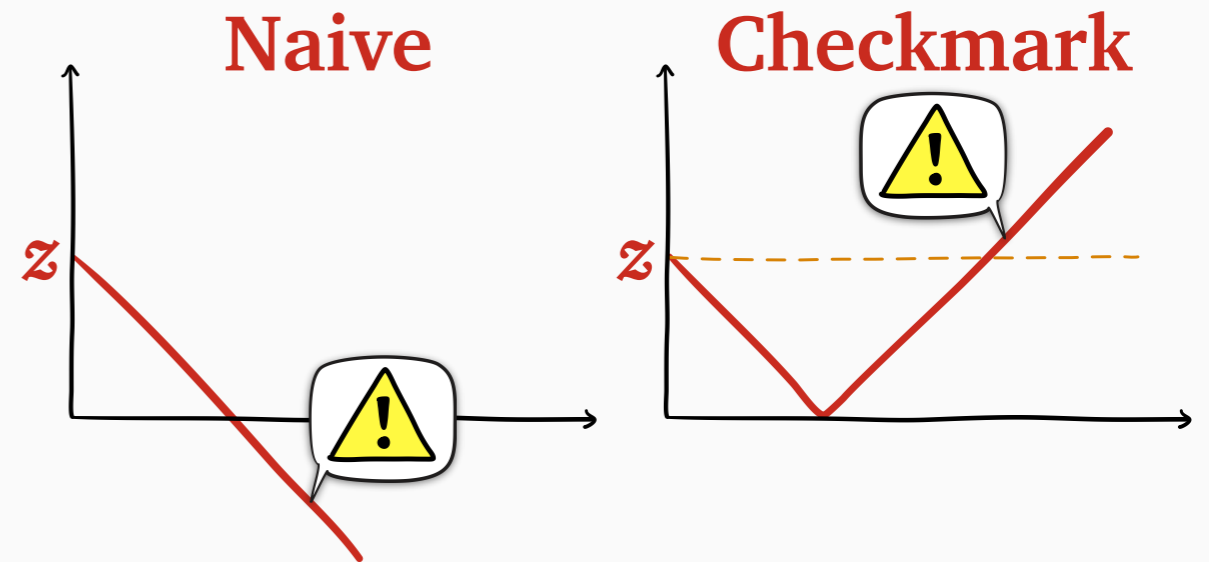


? s {  $z \in [\beta s, \alpha s]$


Problem: minimize $E[T]$
with noisy information




Obstacle: natural **rank**
functions perform badly




Problem: minimize $E[T]$
with noisy information




Solution: new policy,
Radical, with provably
bounded $E[T]$



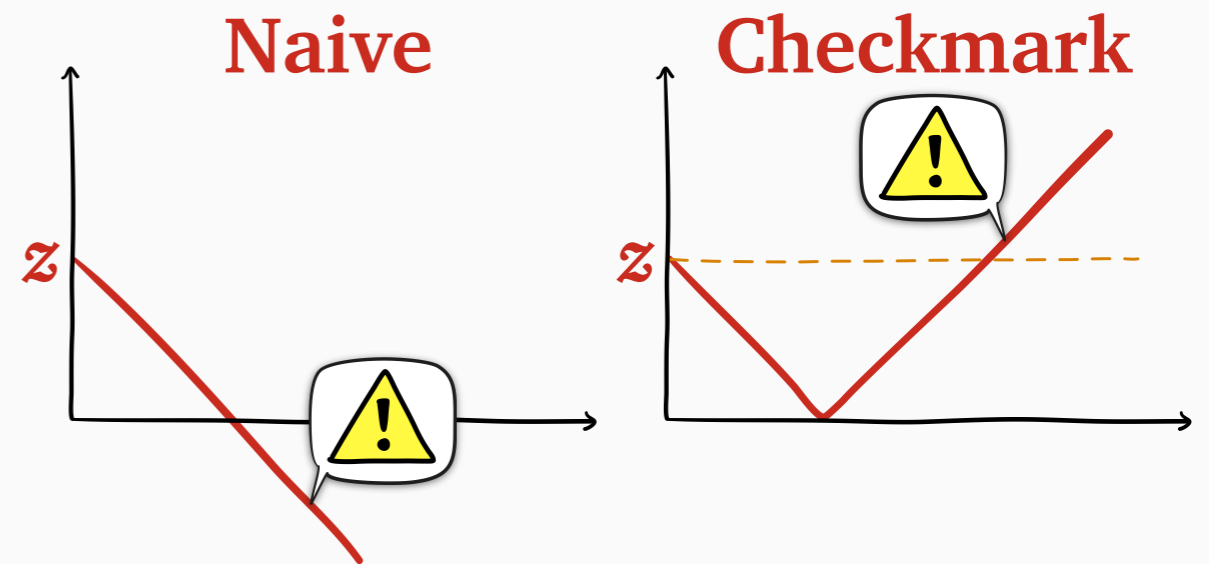
Obstacle: natural **rank**
functions perform badly



Problem: minimize $E[T]$
with noisy information




Obstacle: natural **rank**
functions perform badly




Solution: new policy,
Radical, with provably
bounded $E[T]$

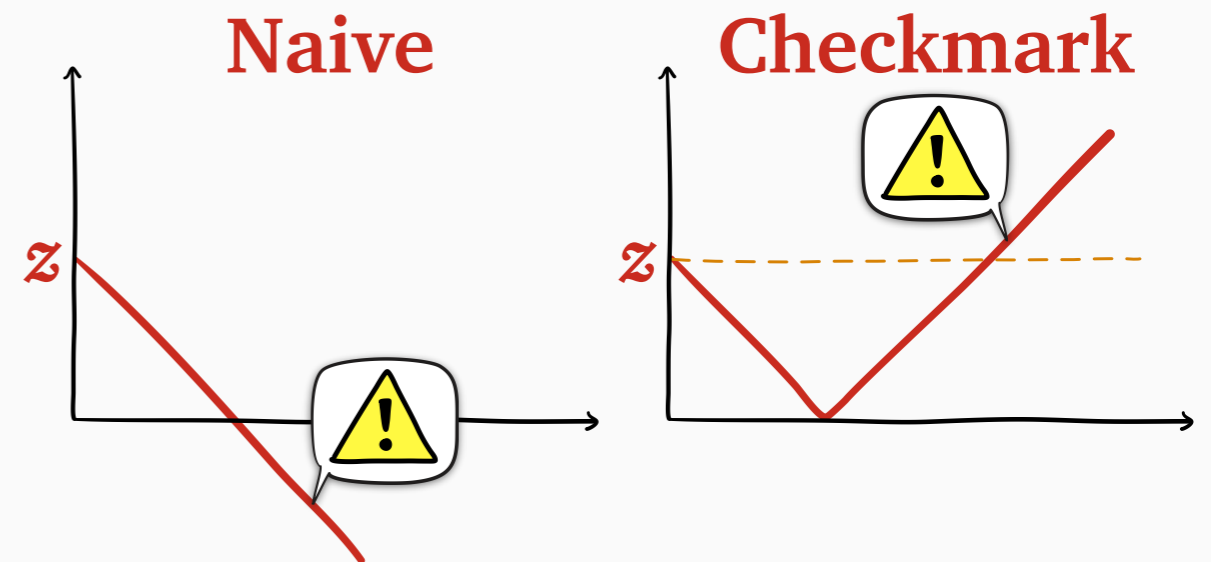
1-consistent,
3.5-graceful




Problem: minimize $E[T]$
with noisy information



Obstacle: natural **rank**
functions perform badly



Solution: new policy,
Radical, with provably
bounded $E[T]$



1-consistent,
3.5-graceful

Method: two new tools
from queueing theory



Consistency-robustness tradeoff?

✓ C -consistent: $\frac{\mathbf{E}[T_P]}{\mathbf{E}[T_{\text{SRPT}}]} \rightarrow C$ as $\alpha, \beta \rightarrow 1$

✓ G -graceful: $\frac{\mathbf{E}[T_P]}{\mathbf{E}[T_{\text{SRPT}}]} \leq G \cdot \frac{\alpha}{\beta}$ for all α, β

✗ R -robust: $\frac{\mathbf{E}[T_P]}{\mathbf{E}[T_{\text{SRPT}}]} \leq R$ for all α, β

Our contribution: *first policy P*
that's consistent and graceful

- $G = 3.5$
- $C = 1$

Consistency-robustness tradeoff?

✓ C -consistent: $\frac{\mathbf{E}[T_P]}{\mathbf{E}[T_{\text{SRPT}}]} \rightarrow C$ as $\alpha, \beta \rightarrow 1$

✓ G -graceful: $\frac{\mathbf{E}[T_P]}{\mathbf{E}[T_{\text{SRPT}}]} \leq G \cdot \frac{\alpha}{\beta}$ for all α, β

✗ R -robust: $\frac{\mathbf{E}[T_P]}{\mathbf{E}[T_{\text{SRPT}}]} \leq R$ for all α, β



Our contribution: *first policy P* that's consistent and graceful

- $G = 3.5$
- $C = 1$

Consistency-robustness tradeoff?

✓ C -consistent: $\frac{E[T_P]}{E[T_{SRPT}]} \rightarrow C$ as $\alpha, \beta \rightarrow 1$

✓ G -graceful: $\frac{E[T_P]}{E[T_{SRPT}]} \leq G \cdot \frac{\alpha}{\beta}$ for all α, β

✗ R -robust: $\frac{E[T_P]}{E[T_{SRPT}]} \leq R$ for all α, β




Our contribution
that's consistent




- $G = 3.5$
- $C = 1$


$k = 0.25, \sigma = 0.5$



$k = 0.25, \sigma = 1.0$




$k = 0.25, \sigma = 1.5$




— SRPT


$k = 0.25, \sigma = 0.5$



$k = 0.25, \sigma = 1.0$




$k = 0.25, \sigma = 1.5$




— SRPT
- - - Noisy SRPT


Noisy SRPT (Naive)




$k = 0.25, \sigma = 0.5$



$k = 0.25, \sigma = 1.0$





$k = 0.25, \sigma = 1.5$




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

Noisy SRPT (Naive)







Noisy SRPT (Naive)




Noisy PSJF (1.5-graceful)







Noisy SRPT (Naive)



Noisy PSJF (1.5-graceful)








Noisy PSJF
suffices

Noisy SRPT (Naive)

Noisy PSJF (1.5-graceful)








Noisy PSJF
suffices

ignore the
estimates

Noisy SRPT (Naive)

Noisy PSJF (1.5-graceful)






Noisy PSJF
suffices


either way

ignore the
estimates

Noisy SRPT (Naive)

Noisy PSJF (1.5-graceful)





Noisy PSJF suffices

either way



ignore the estimates

Noisy SR


.5-graceful)

?


Mix Noisy PSJF with Opt. w/out?



Problem: minimize $E[T]$
with noisy information




Solution: new policy,
Radical, with provably
bounded $E[T]$



1-consistent,
3.5-graceful

Obstacle: natural **rank**
functions perform badly



Method: two new tools
from queueing theory



Proof sketch

Lemma:

$$\mathbf{E}[W_{\text{Scale}}(r)] \leq \mathbf{E}[W_{\text{SRPT}}(\frac{\alpha}{\beta}r)]$$

Proof sketch

Lemma:

$$\mathbf{E}[W_{\text{SRPT}}(\mathbf{r})] \leq \mathbf{E}[W_{\text{Scale}}(\mathbf{r})] \leq \mathbf{E}\left[W_{\text{SRPT}}\left(\frac{\alpha}{\beta}\mathbf{r}\right)\right]$$

Proof sketch

Lemma:

$$\mathbf{E}[W_{\text{SRPT}}(\mathbf{r})] \leq \mathbf{E}[W_{\text{Scale}}(\mathbf{r})] \leq \mathbf{E}\left[W_{\text{SRPT}}\left(\frac{\alpha}{\beta}\mathbf{r}\right)\right]$$

Key steps:

Proof sketch

Lemma:


$$\mathbf{E}[W_{\text{SRPT}}(\mathbf{r})] \leq \mathbf{E}[W_{\text{Scale}}(\mathbf{r})] \leq \mathbf{E}\left[W_{\text{SRPT}}\left(\frac{\alpha}{\beta}\mathbf{r}\right)\right]$$

Key steps:

1. **SRPT** minimizes mean \mathbf{r} -work

Proof sketch

Lemma:



$$\mathbf{E}[W_{\text{SRPT}}(\mathbf{r})] \leq \mathbf{E}[W_{\text{Scale}}(\mathbf{r})] \leq \mathbf{E}\left[W_{\text{SRPT}}\left(\frac{\alpha}{\beta}\mathbf{r}\right)\right]$$

Key steps:

1. **SRPT** minimizes mean \mathbf{r} -work

Proof sketch

Lemma:



$$\mathbf{E}[W_{\text{SRPT}}(\mathbf{r})] \leq \mathbf{E}[W_{\text{Scale}}(\mathbf{r})] \leq \mathbf{E}\left[W_{\text{SRPT}}\left(\frac{\alpha}{\beta}\mathbf{r}\right)\right]$$

Key steps:

1. **SRPT** minimizes mean \mathbf{r} -work
2. **Scale** minimizes mean **noise-scaled- \mathbf{r}** -work

Proof sketch

Lemma:


$$\mathbf{E}[W_{\text{SRPT}}(\mathbf{r})] \leq \mathbf{E}[W_{\text{Scale}}(\mathbf{r})] \leq \mathbf{E}\left[W_{\text{SRPT}}\left(\frac{\alpha}{\beta}\mathbf{r}\right)\right]$$

Key steps:


1. **SRPT** minimizes mean \mathbf{r} -work
2. **Scale** minimizes mean **noise-scaled- \mathbf{r}** -work



filters using **Scale's rank**
instead of **SRPT's rank**

Proof sketch

Lemma:


$$\mathbf{E}[W_{\text{SRPT}}(\mathbf{r})] \leq \mathbf{E}[W_{\text{Scale}}(\mathbf{r})] \leq \mathbf{E}\left[W_{\text{SRPT}}\left(\frac{\alpha}{\beta}\mathbf{r}\right)\right]$$

Key steps:


1. **SRPT** minimizes mean \mathbf{r} -work
2. **Scale** minimizes mean **noise-scaled- \mathbf{r}** -work
3. Under any policy,



filters using **Scale's rank**
instead of **SRPT's rank**

Proof sketch


Lemma:


$$\mathbf{E}[W_{\text{SRPT}}(\mathbf{r})] \leq \mathbf{E}[W_{\text{Scale}}(\mathbf{r})] \leq \mathbf{E}\left[W_{\text{SRPT}}\left(\frac{\alpha}{\beta}\mathbf{r}\right)\right]$$

Key steps:

1. **SRPT** minimizes mean \mathbf{r} -work
2. **Scale** minimizes mean **noise-scaled- \mathbf{r}** -work
3. Under any policy,

$$\mathbf{r}\text{-work} \leq \mathbf{noise-scaled-}\alpha\mathbf{r}\text{-work} \leq \frac{\alpha}{\beta}\mathbf{r}\text{-work}$$



filters using **Scale's rank** instead of **SRPT's rank**

Proof sketch


Lemma:

$$\mathbb{E}[W_{\text{SRPT}}(r)] \leq \mathbb{E}[W_{\text{Scale}}(r)] \leq \mathbb{E}\left[W_{\text{SRPT}}\left(\frac{\alpha}{\beta}r\right)\right]$$

Key steps:

1. **SRPT** minimizes mean r -work
2. **Scale** minimizes mean **noise-scaled- r -work**
3. Under any policy,

$$r\text{-work} \leq \text{noise-scaled-}\alpha r\text{-work} \leq \frac{\alpha}{\beta}r\text{-work}$$



filters using **Scale's rank** instead of **SRPT's rank**