Optimal Multiserver Scheduling

with Unknown Job Sizes in Heavy Traffic

Ziv Scully
Isaac Grosof
Mor Harchol-Balter

Carnegie Mellon University

M/G/k

Optimal Multiserver Scheduling

with Unknown Job Sizes in Heavy Traffic

Ziv Scully
Isaac Grosof
Mor Harchol-Balter

Carnegie Mellon University

M/G/1 Queue

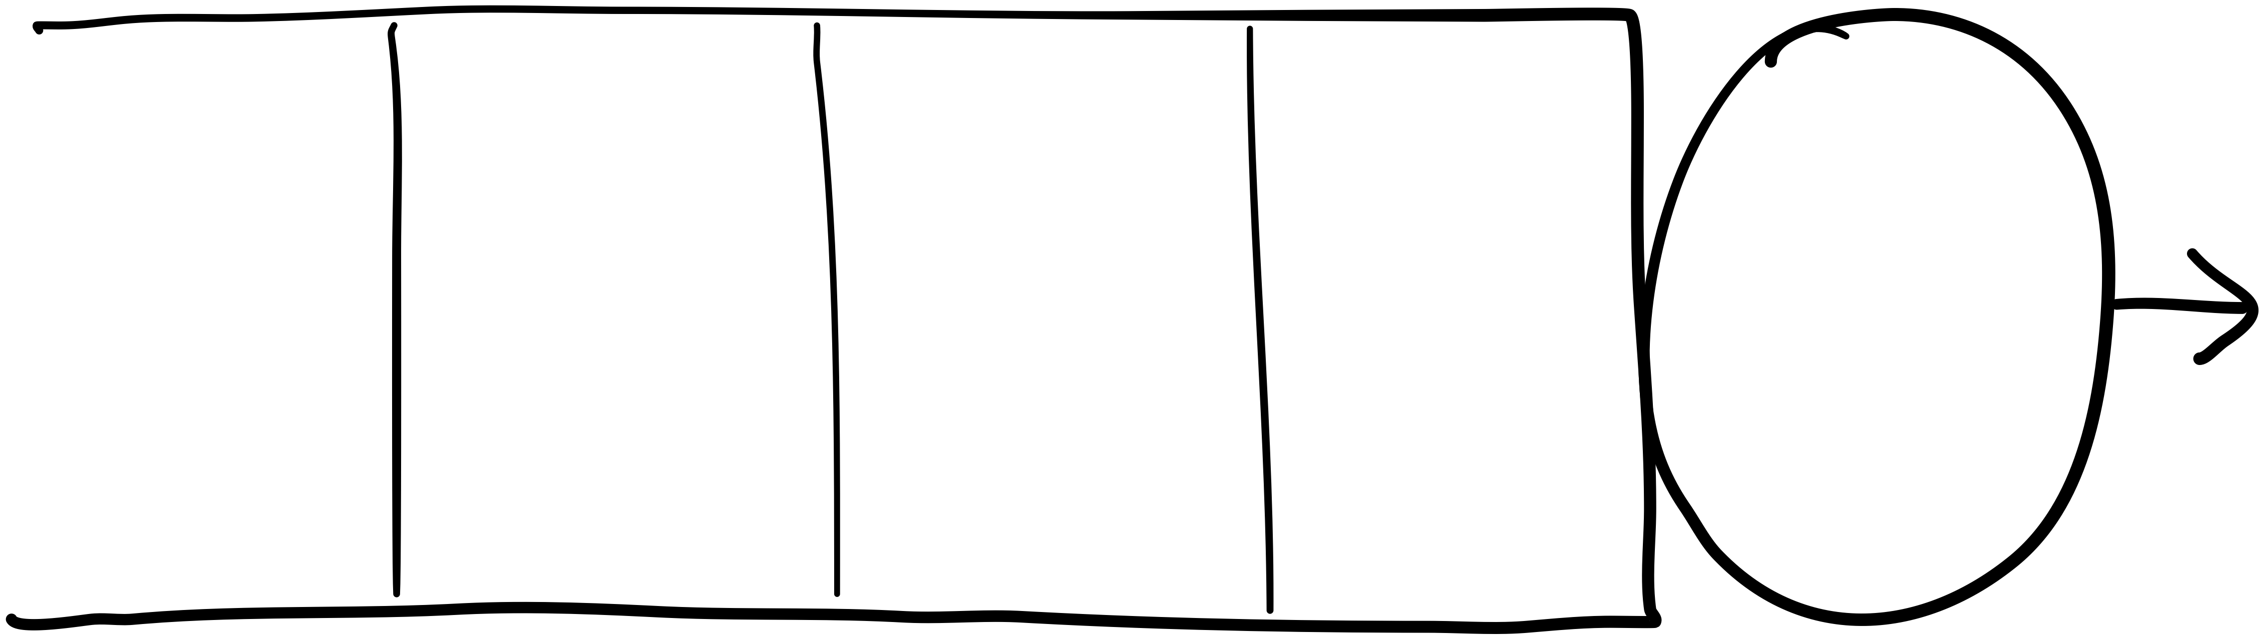

M/G/1 Queue

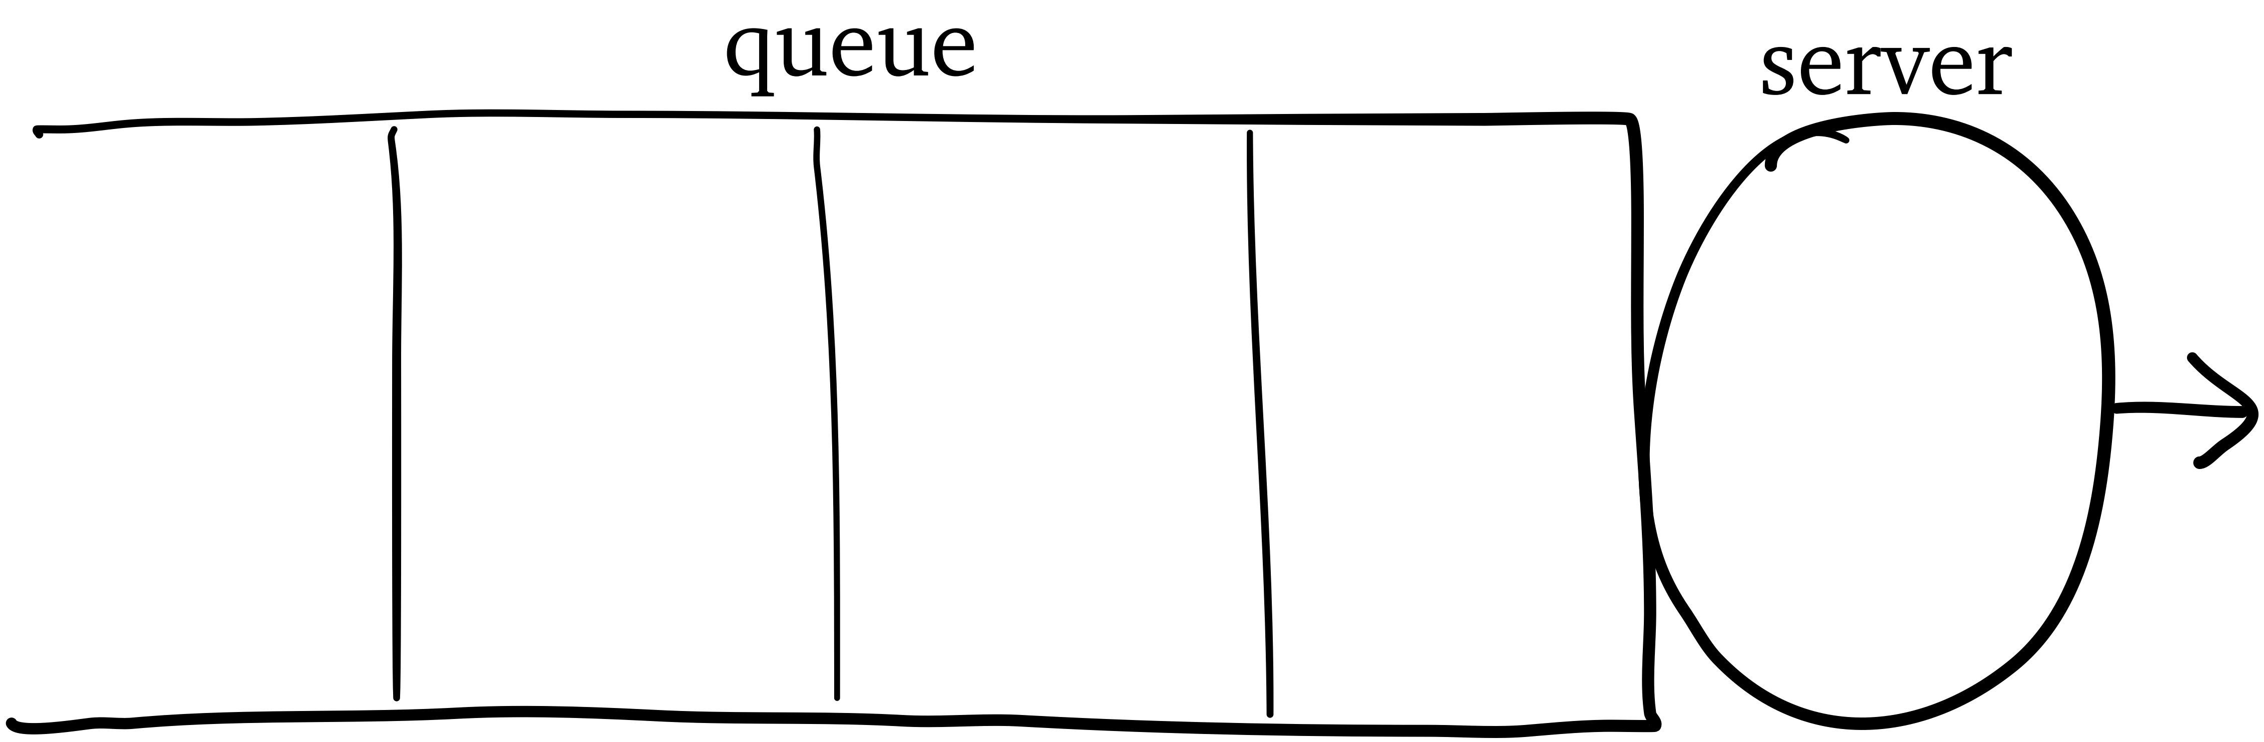

M/G/1 Queue

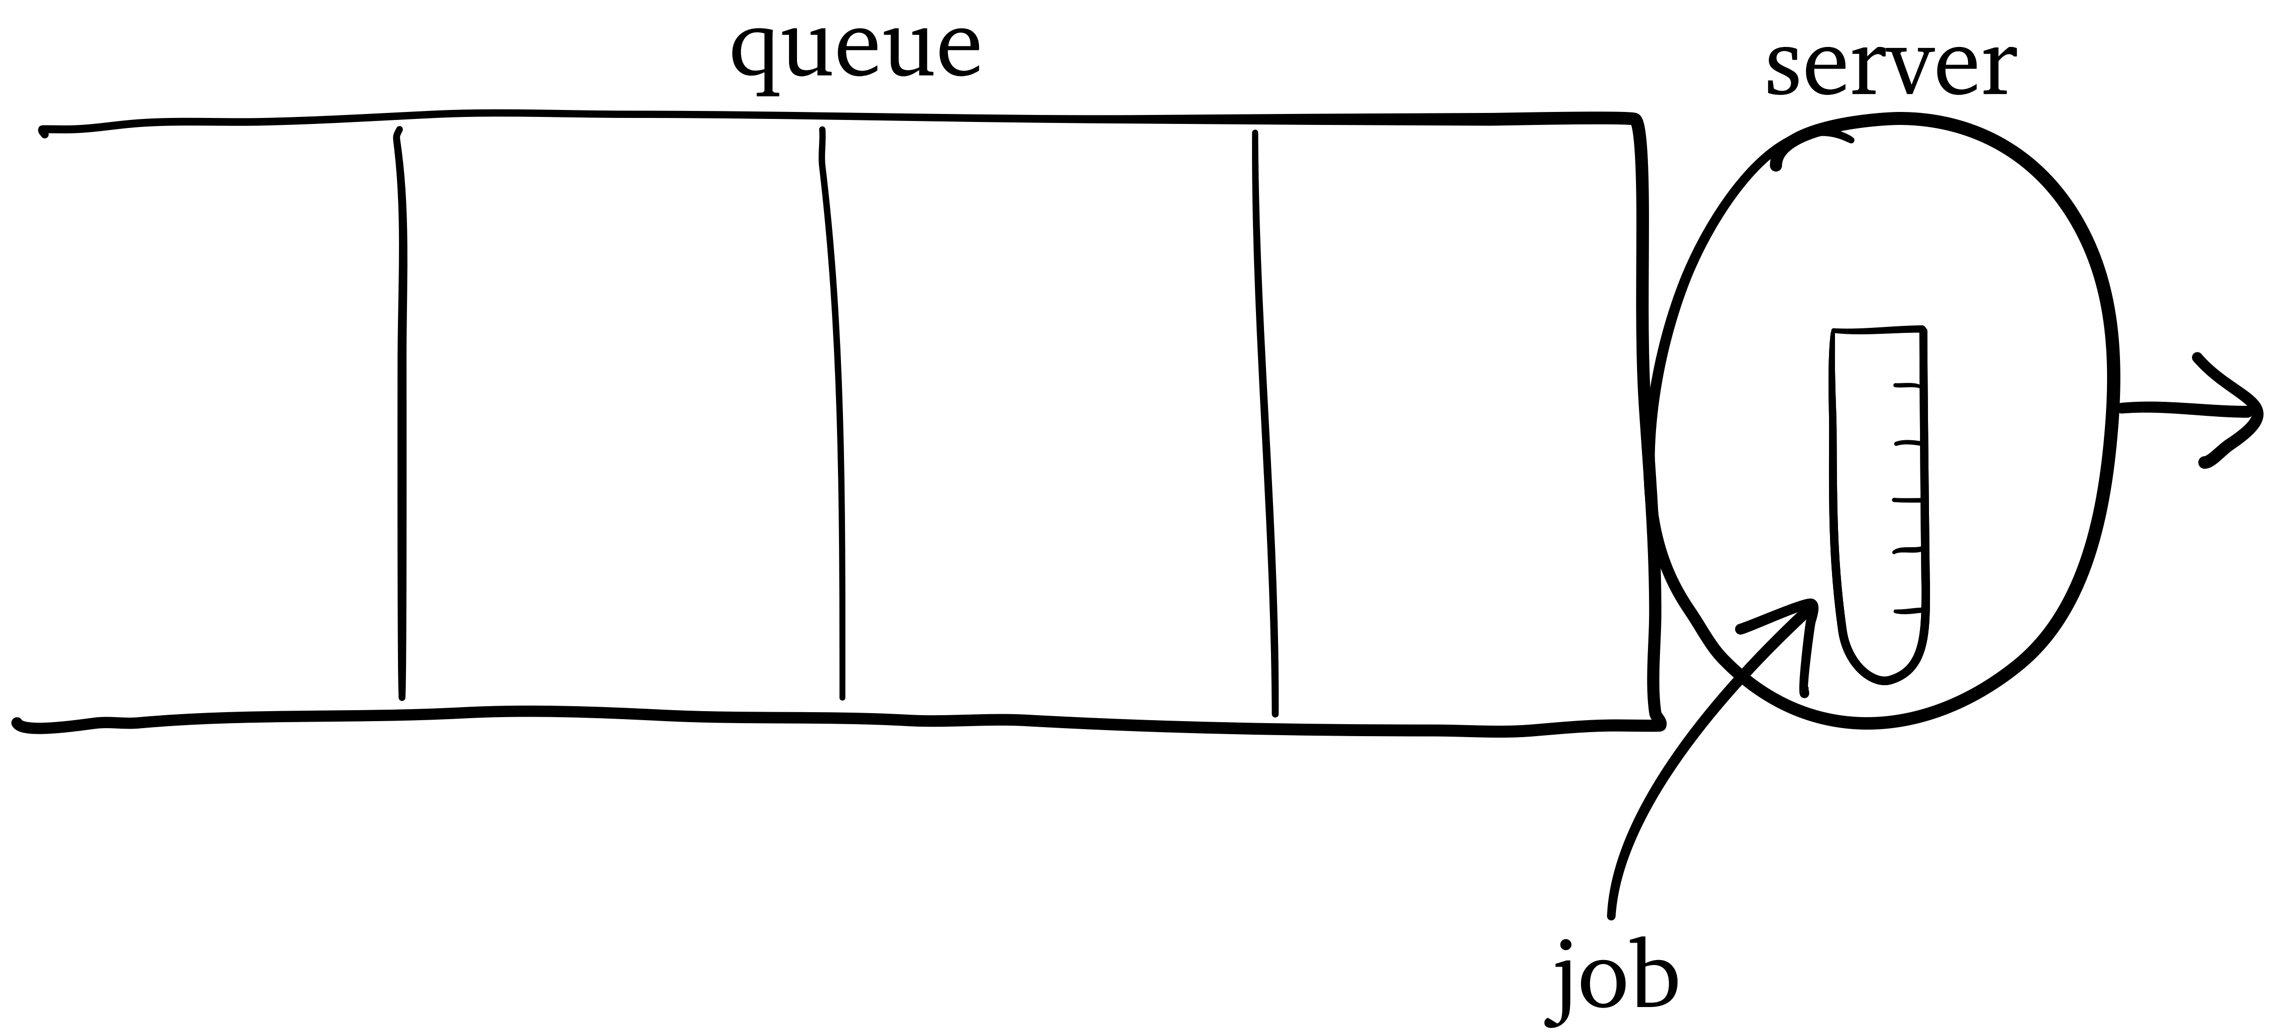

M/G/1 Queue

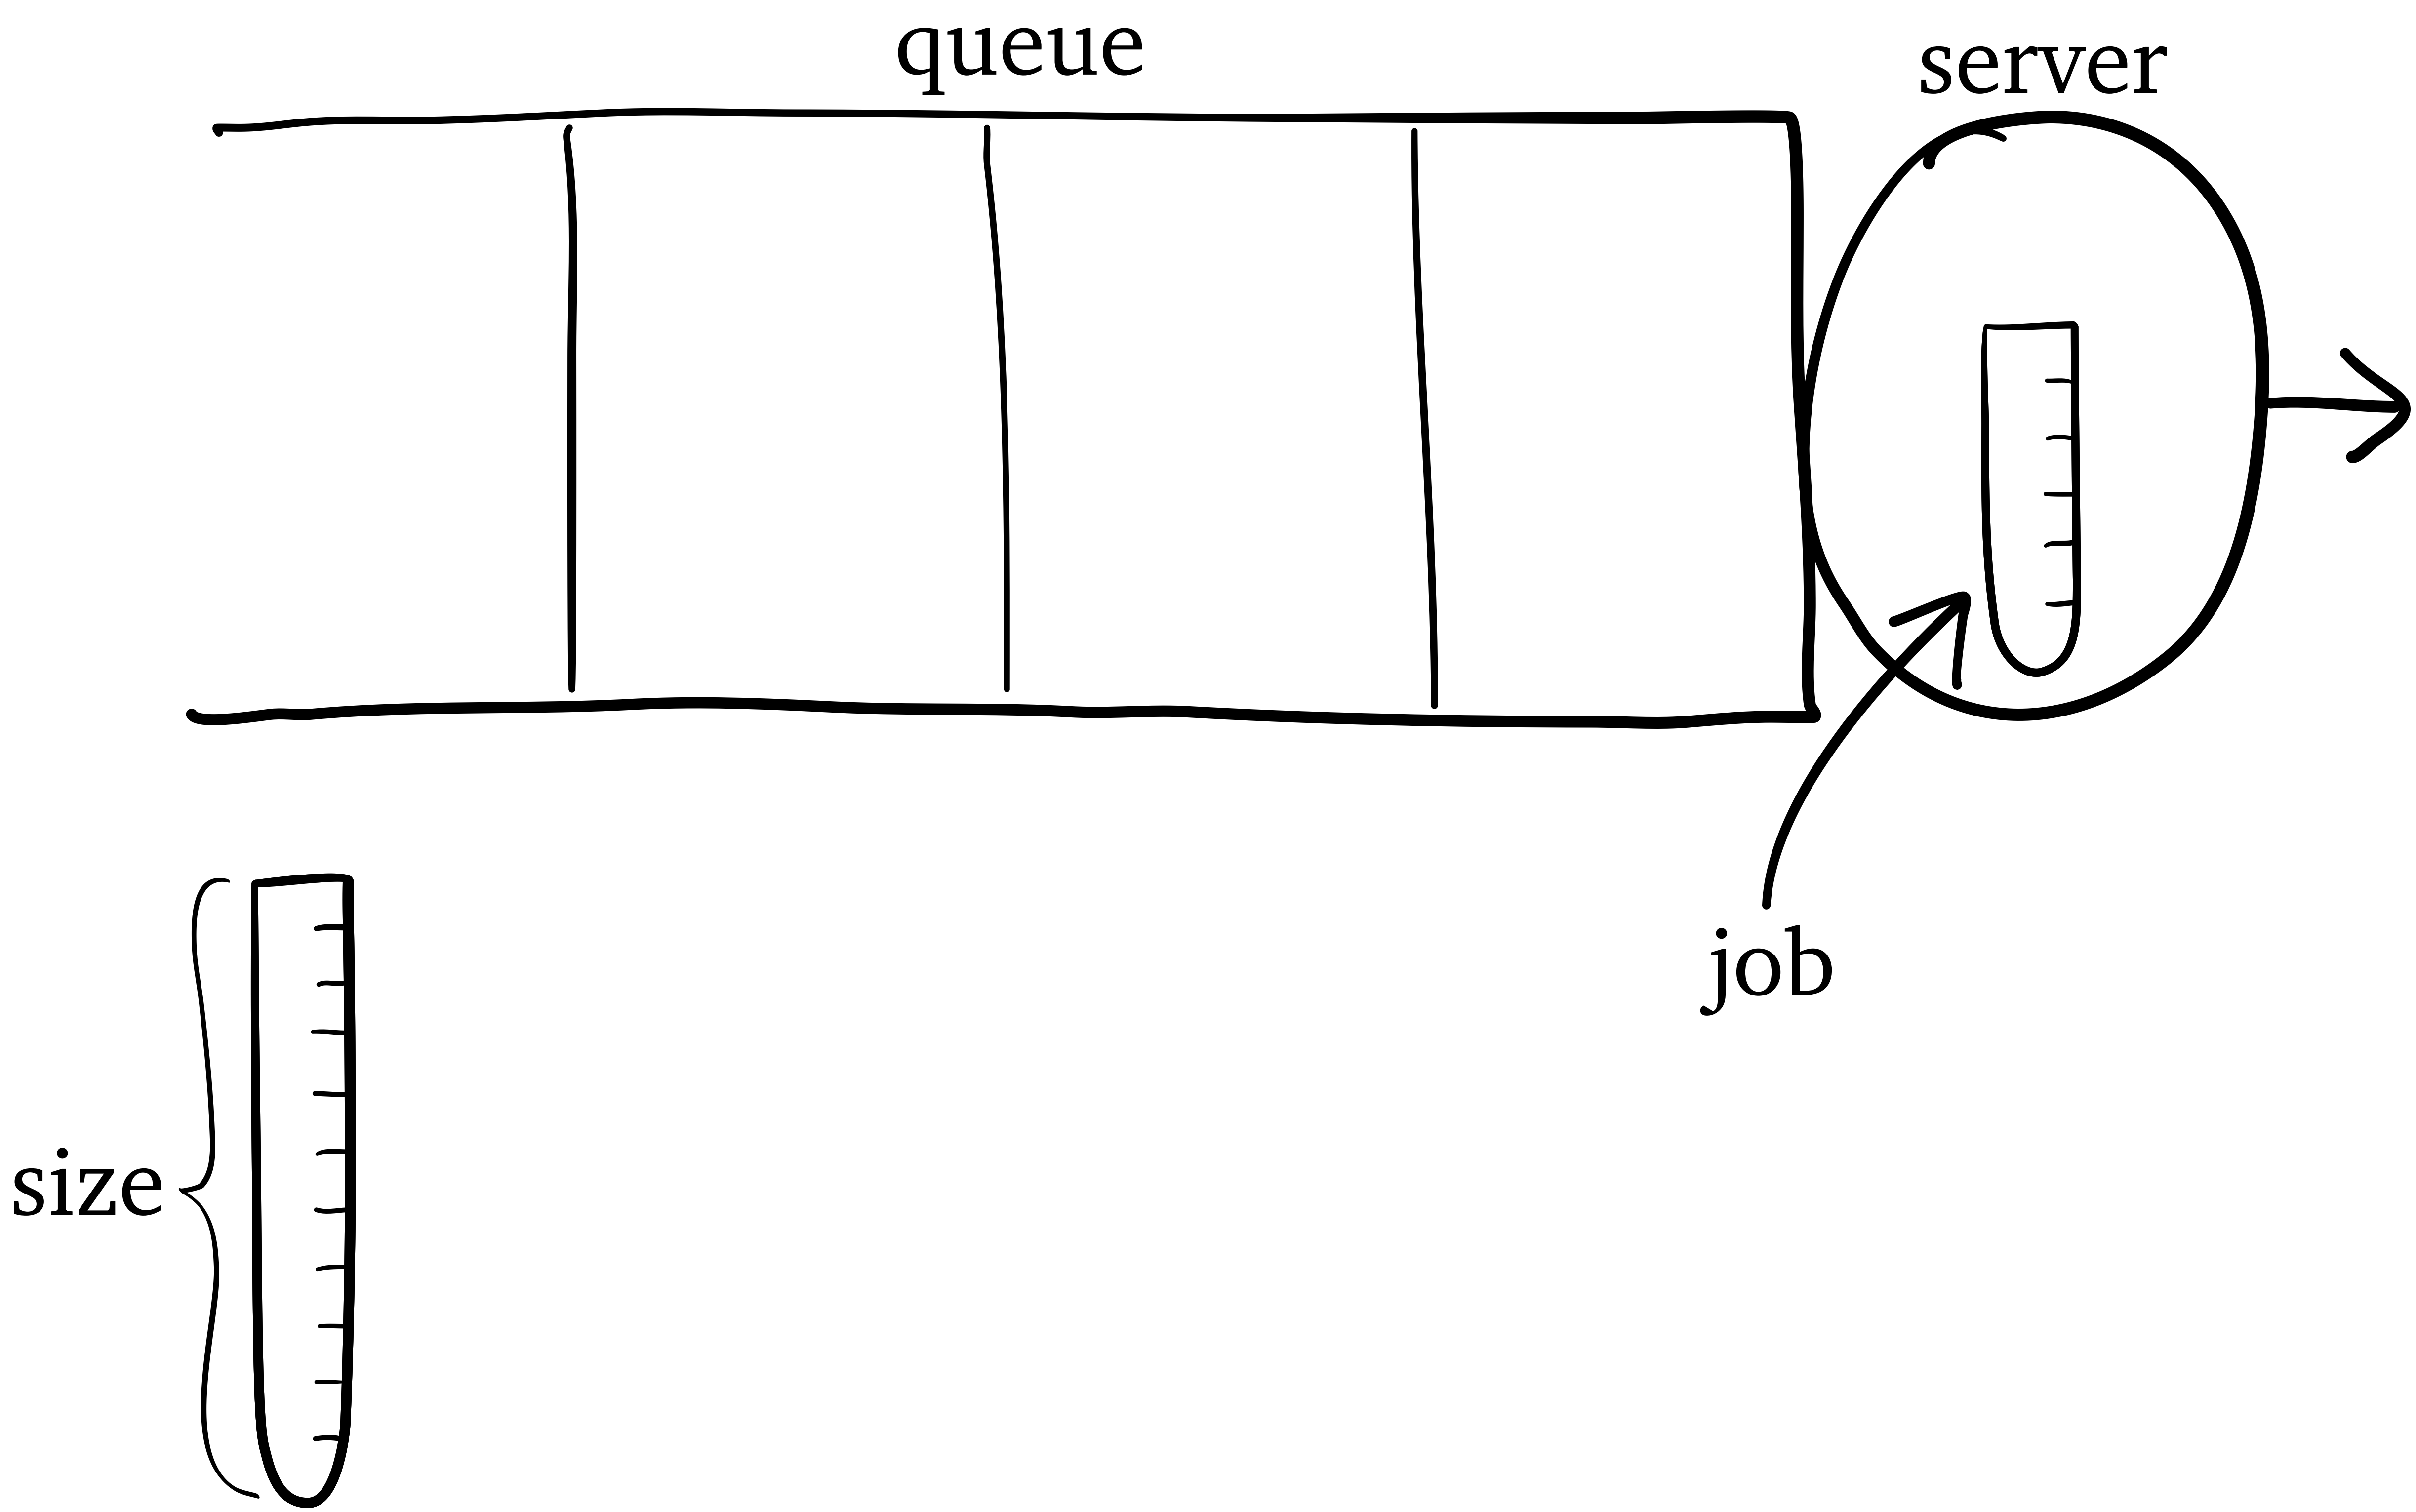

M/G/1 Queue

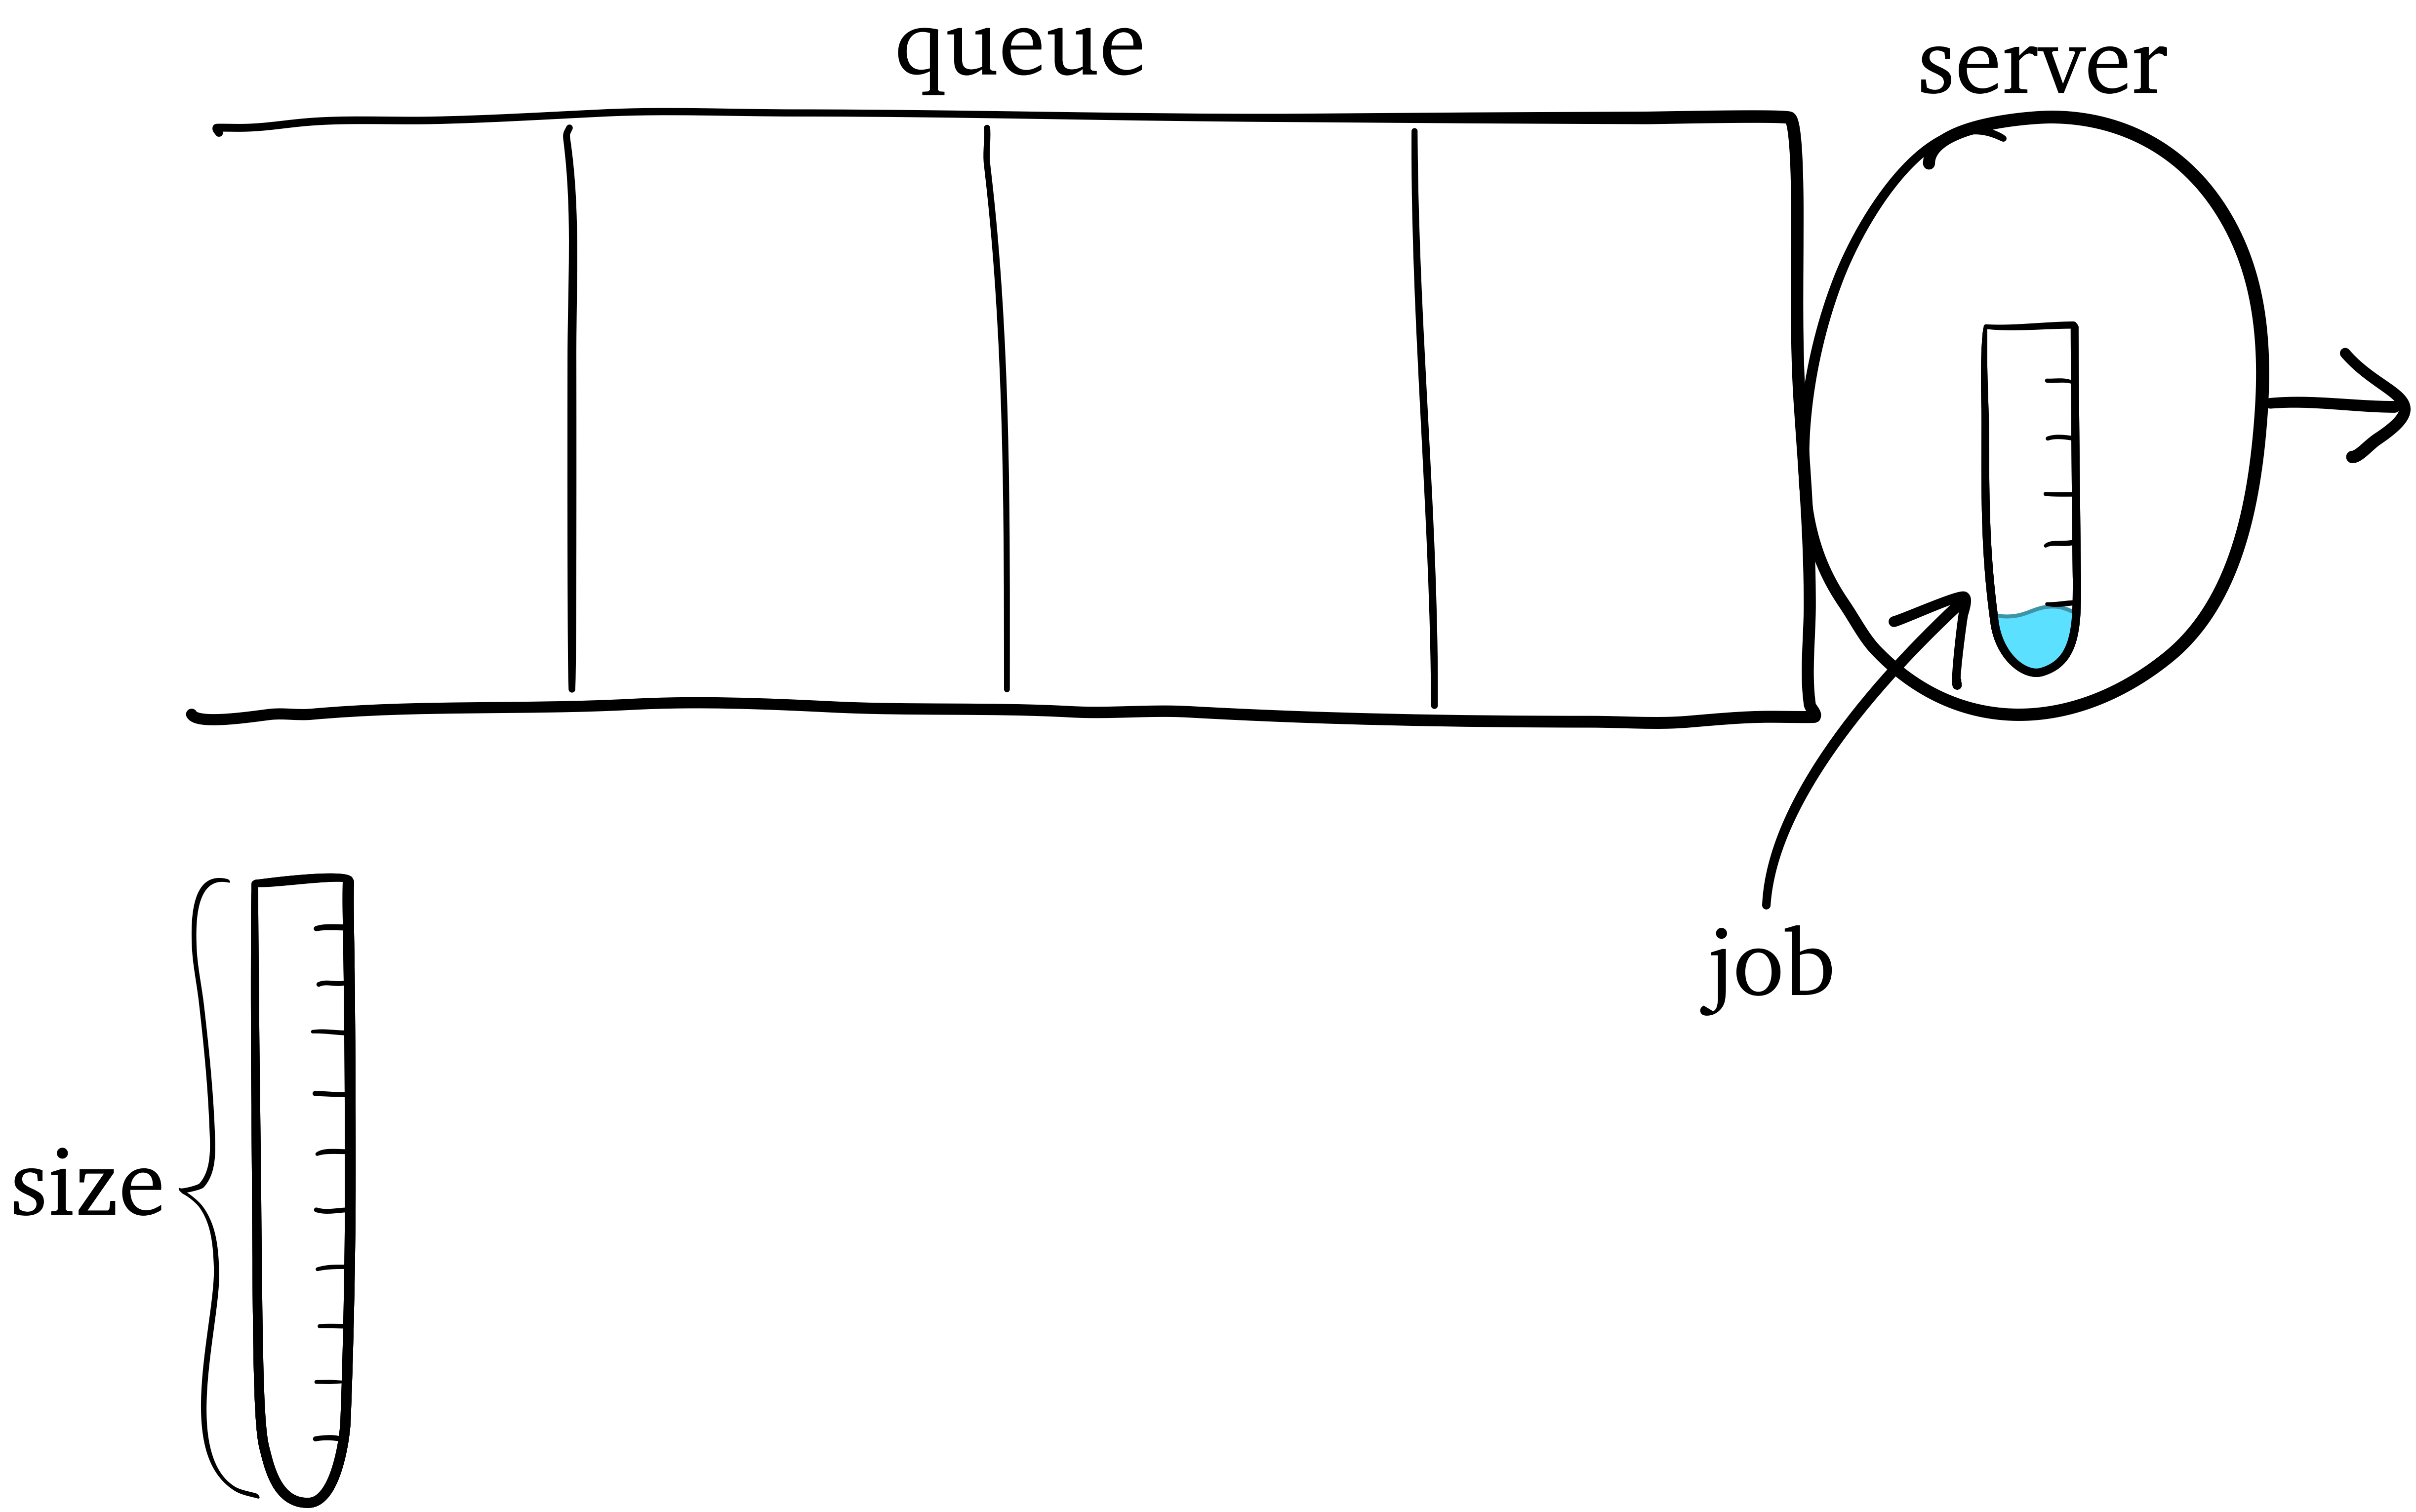

M/G/1 Queue

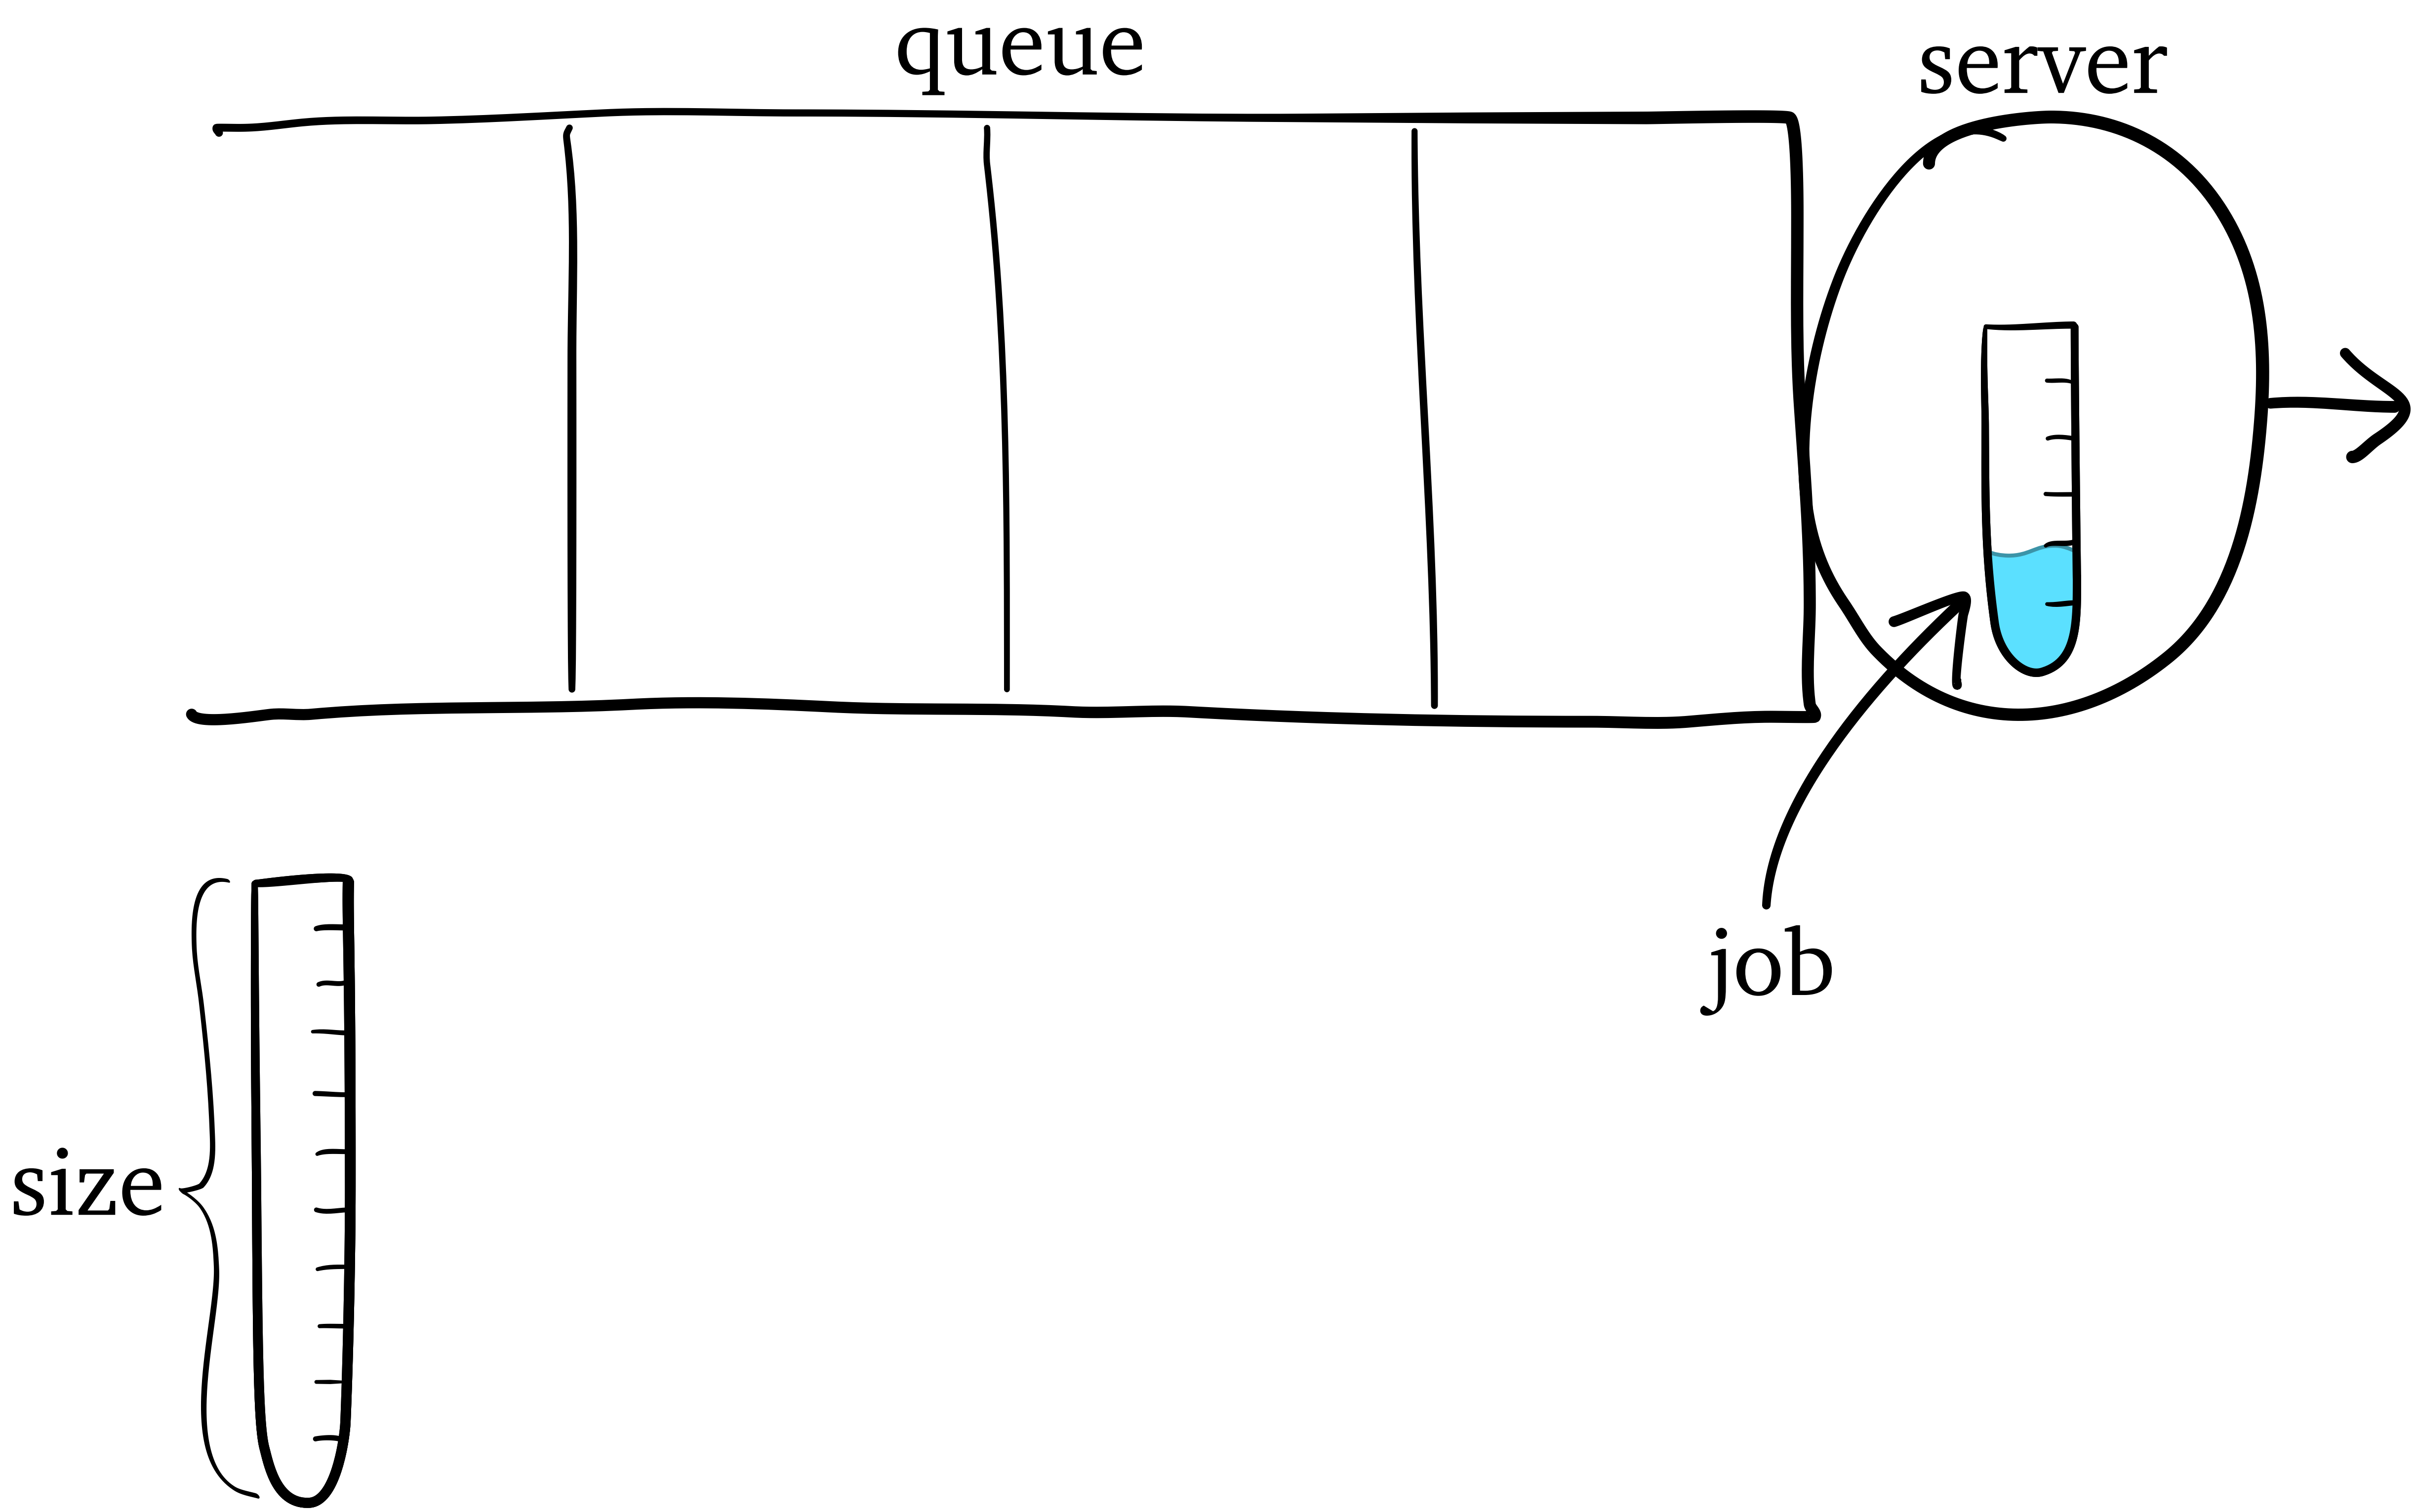

M/G/1 Queue

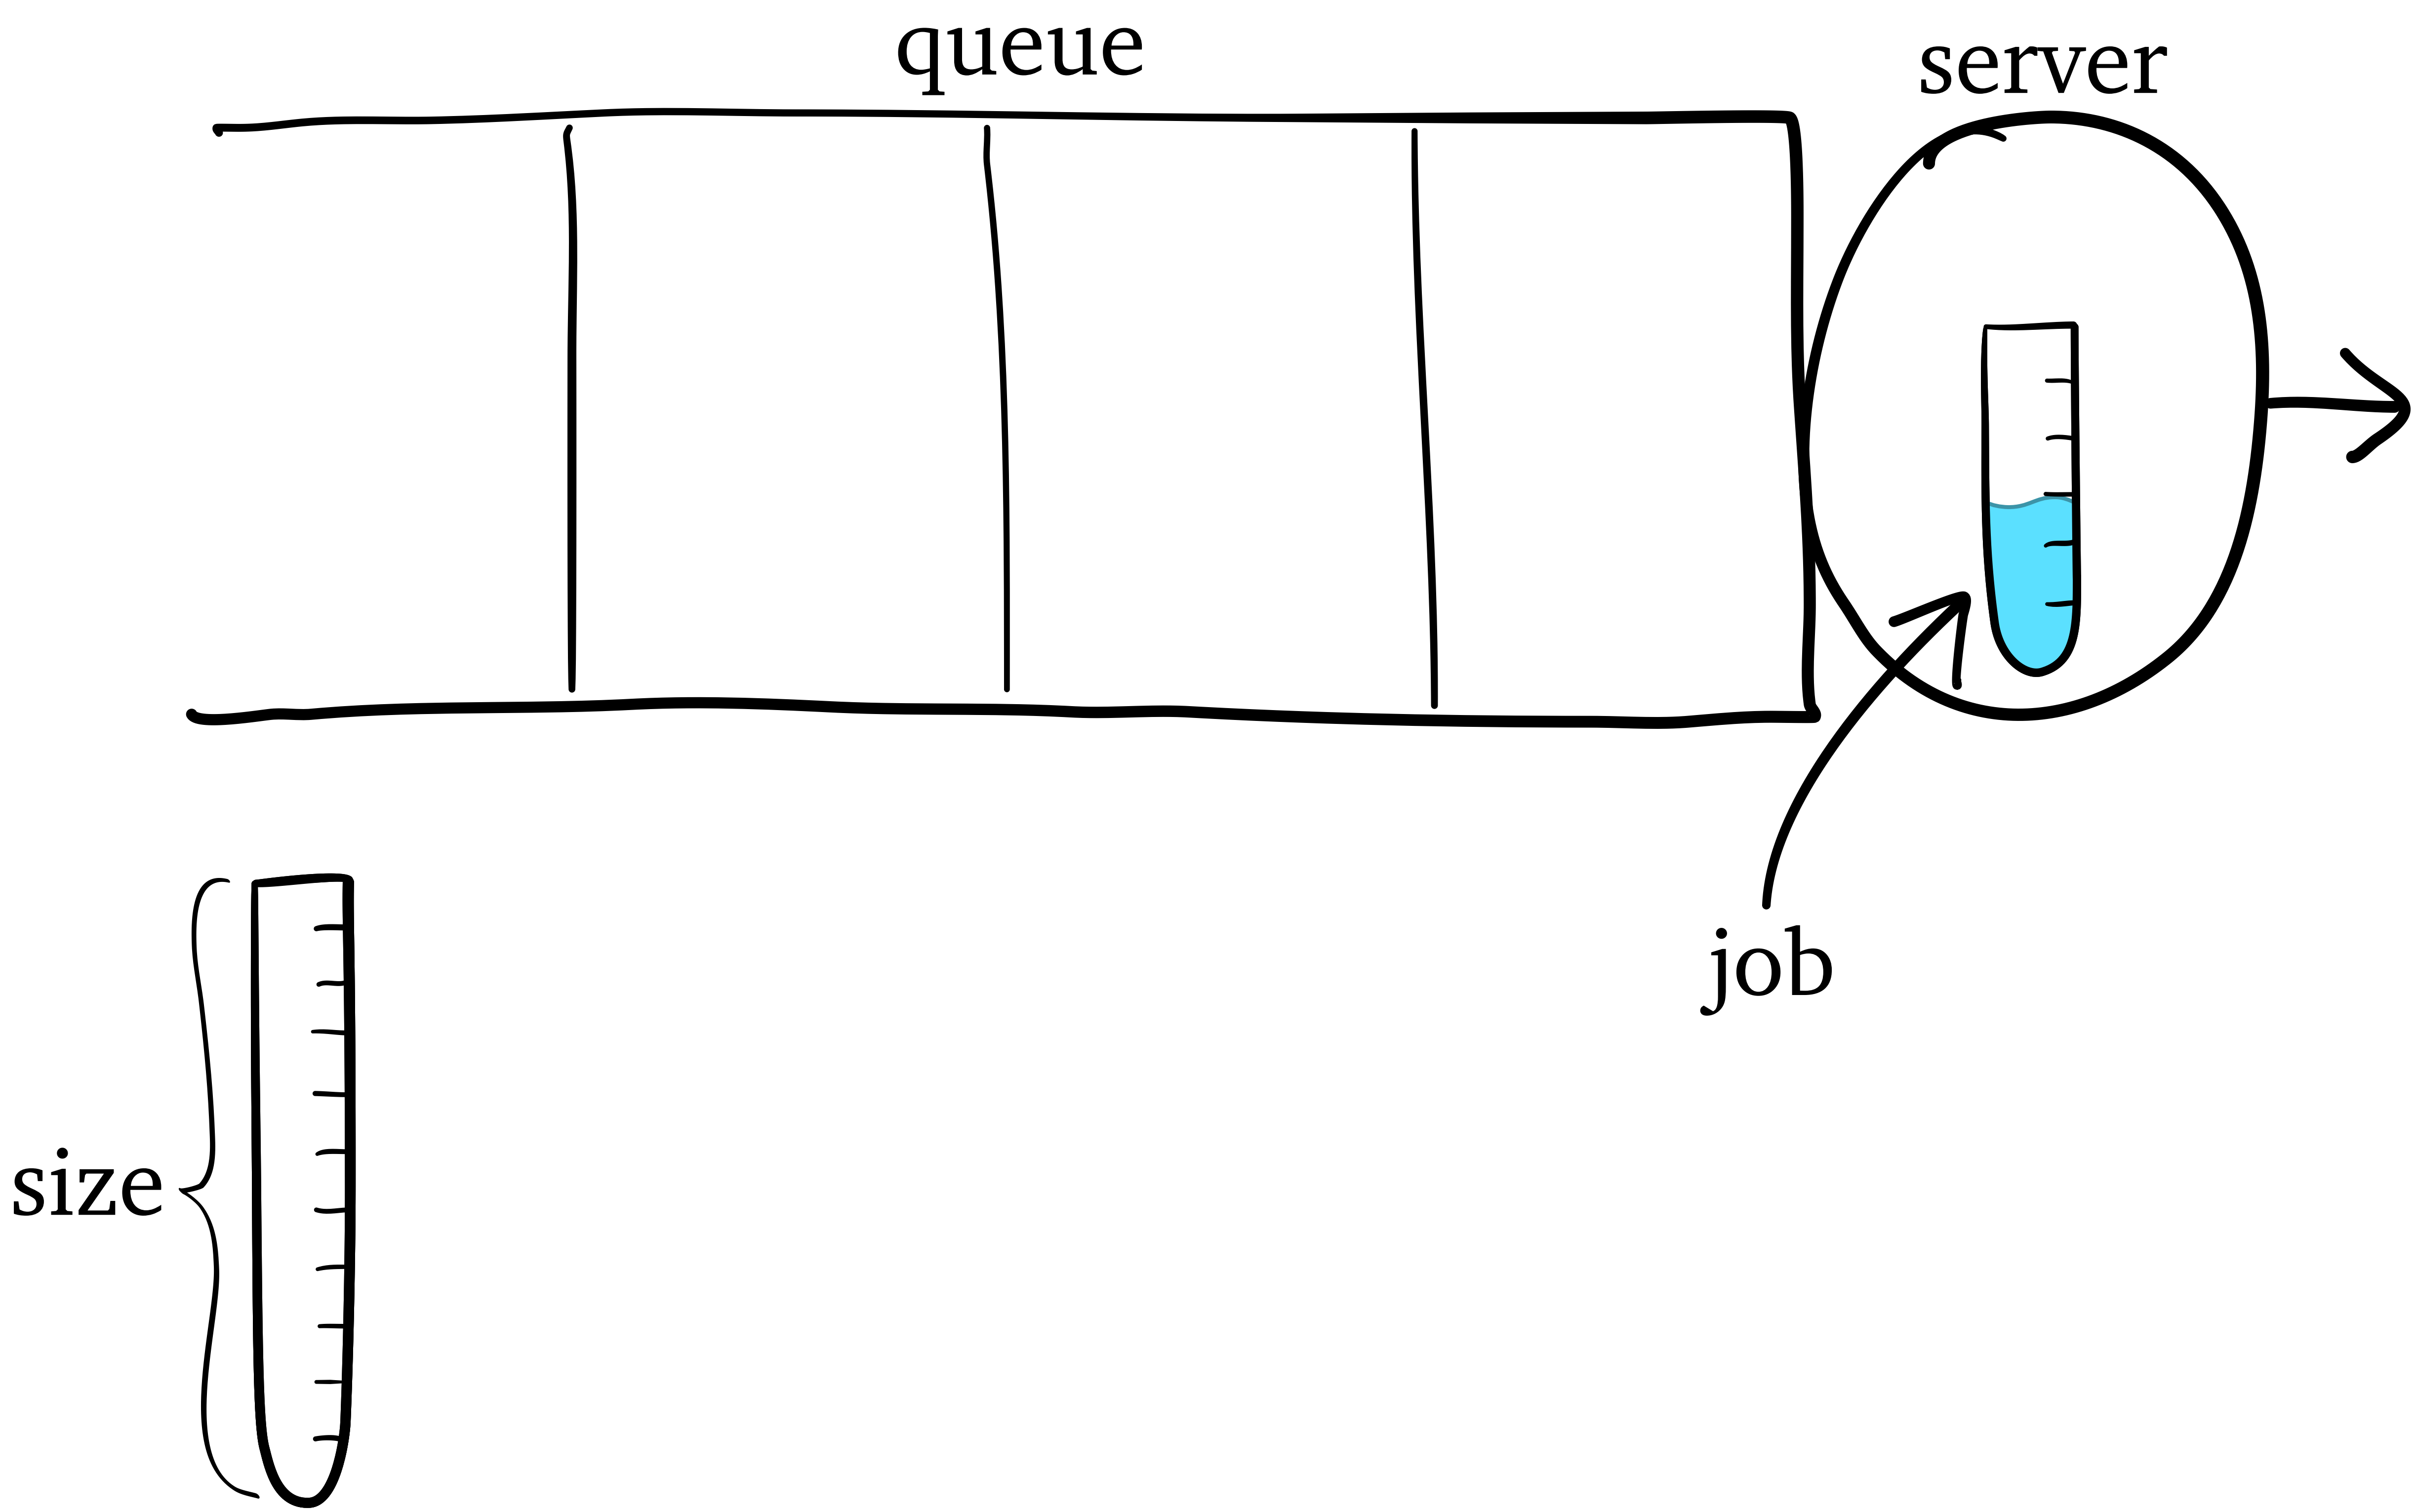

M/G/1 Queue

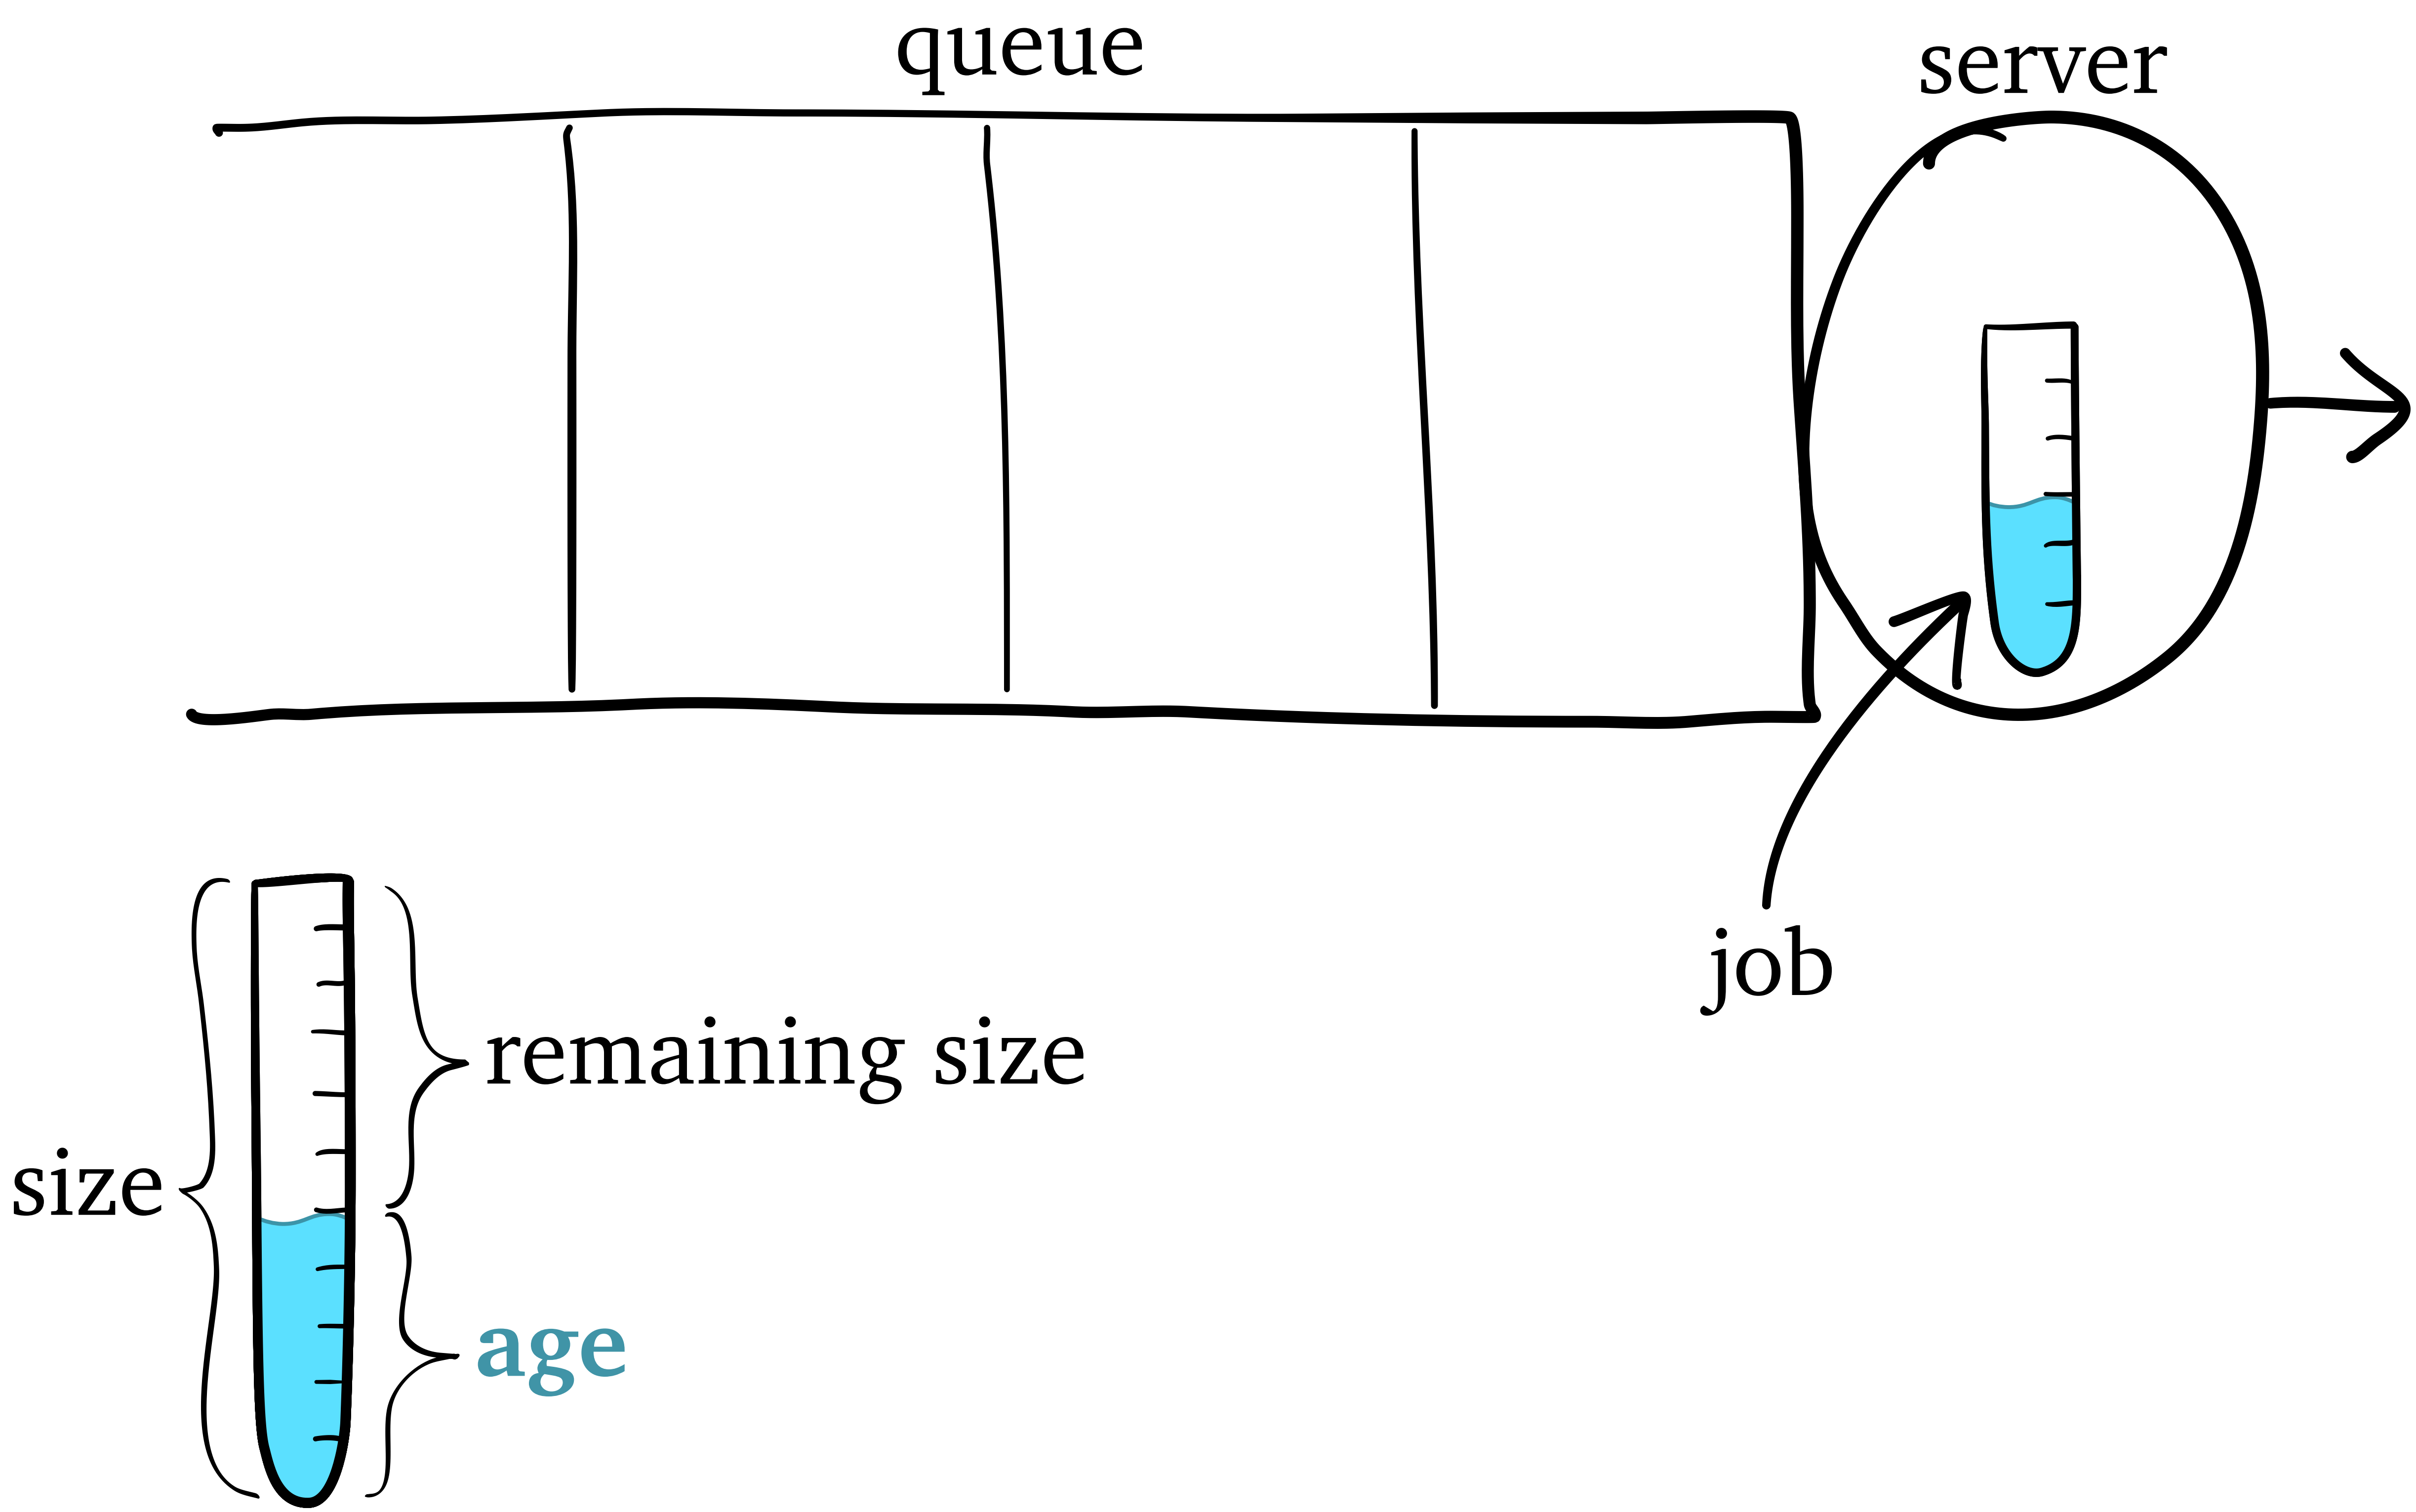

M/G/1 Queue

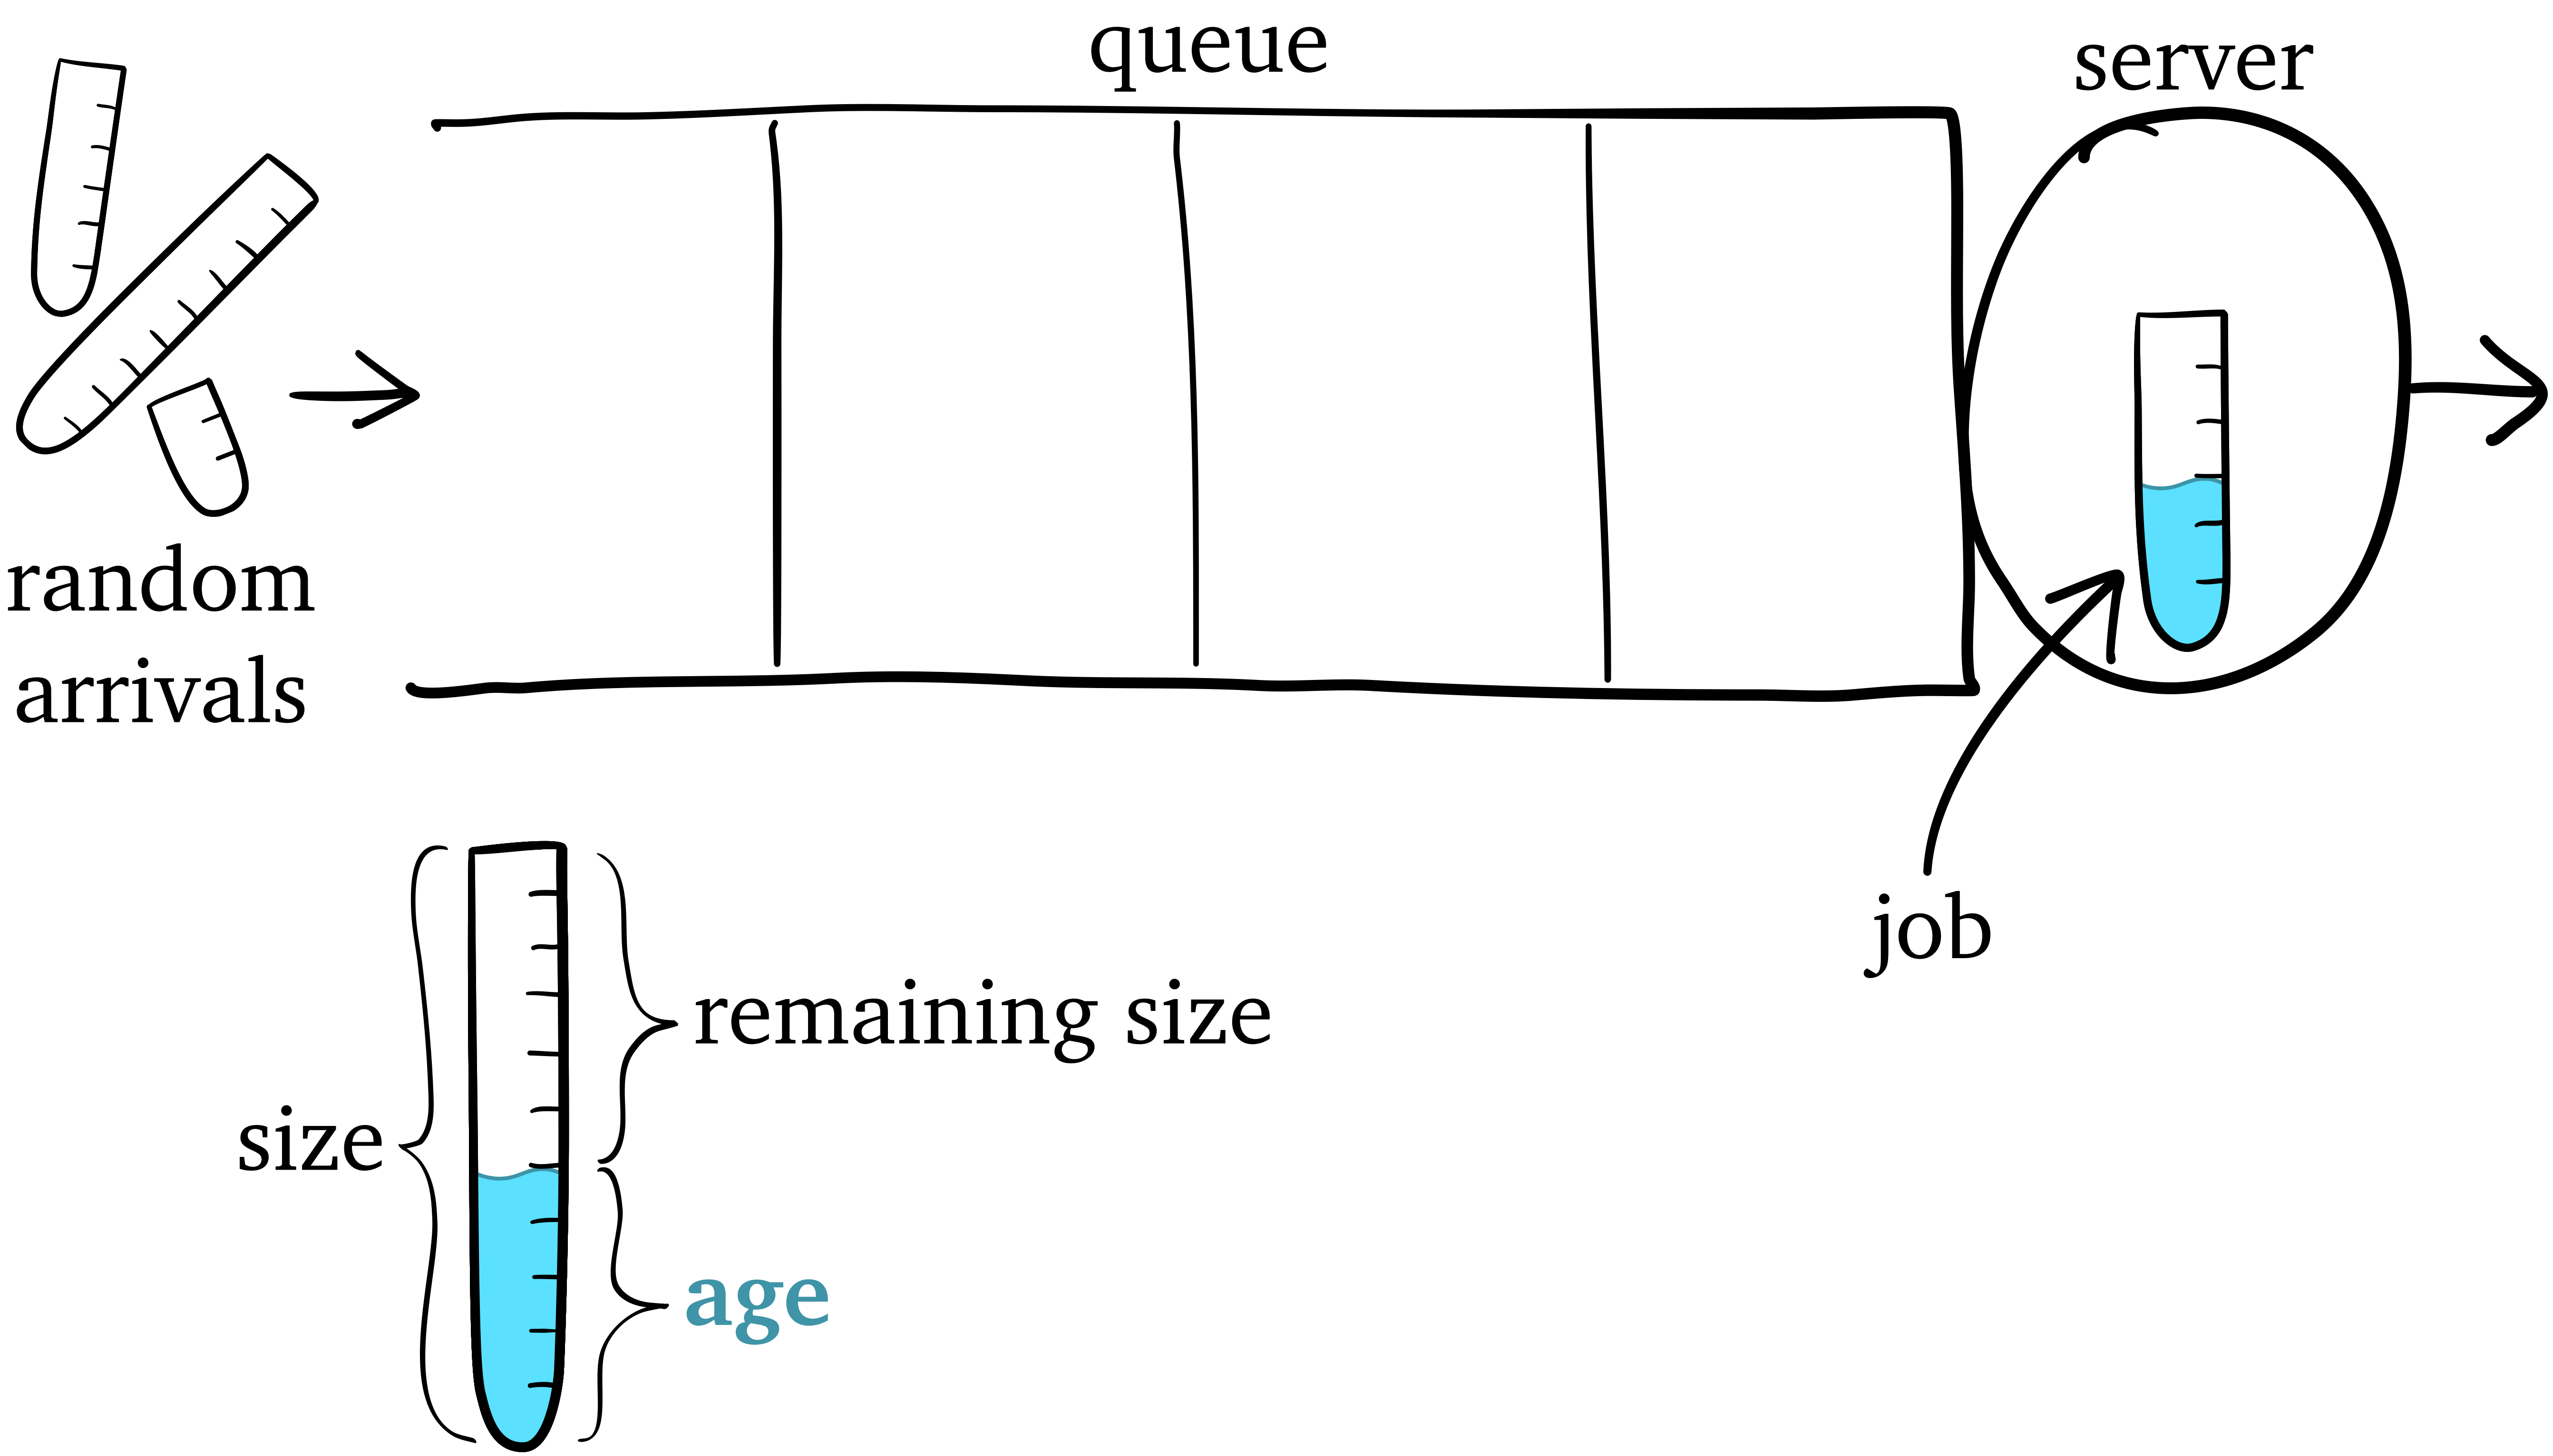

M/G/1 Queue

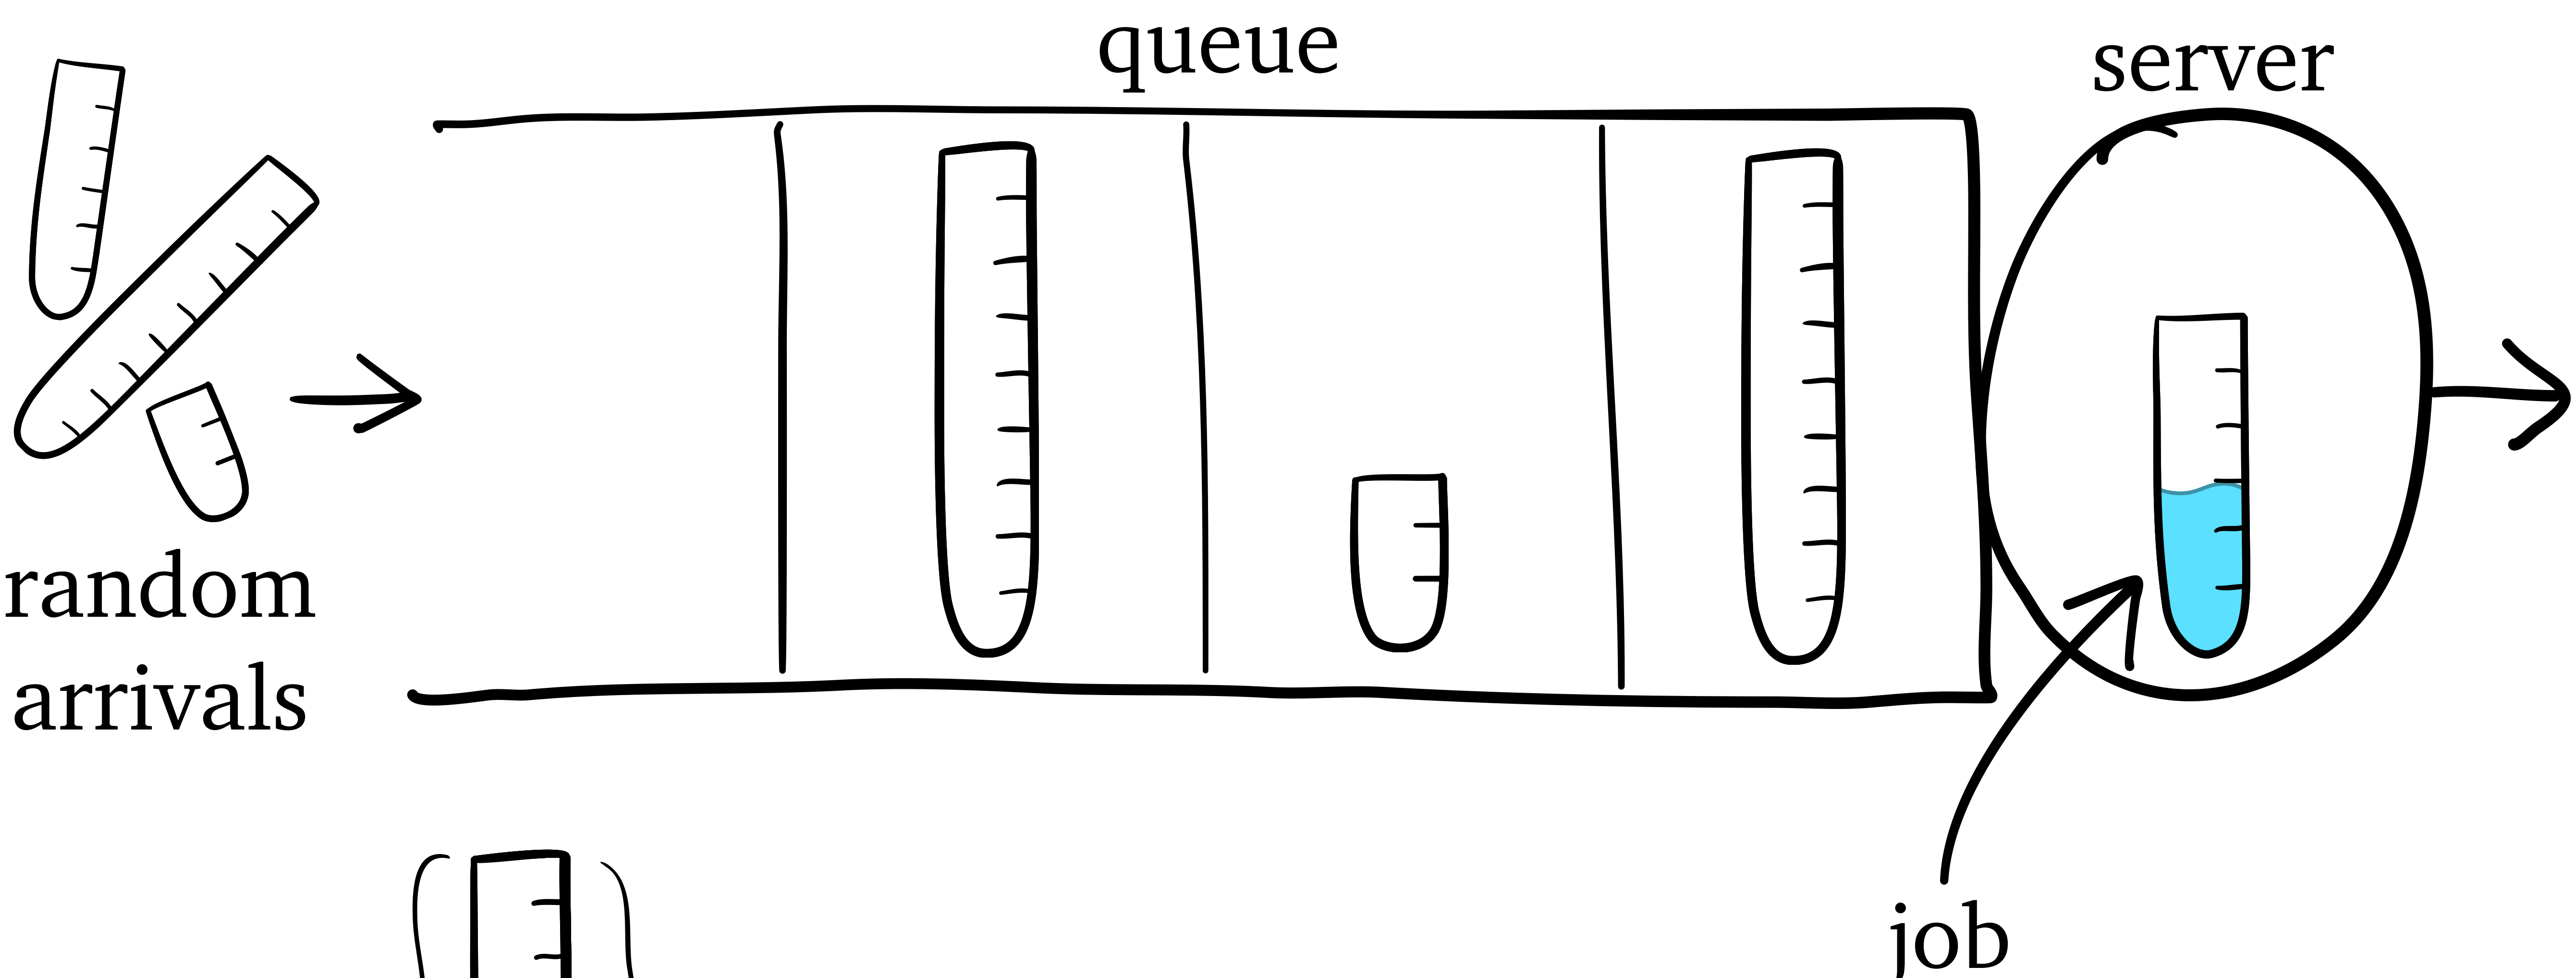

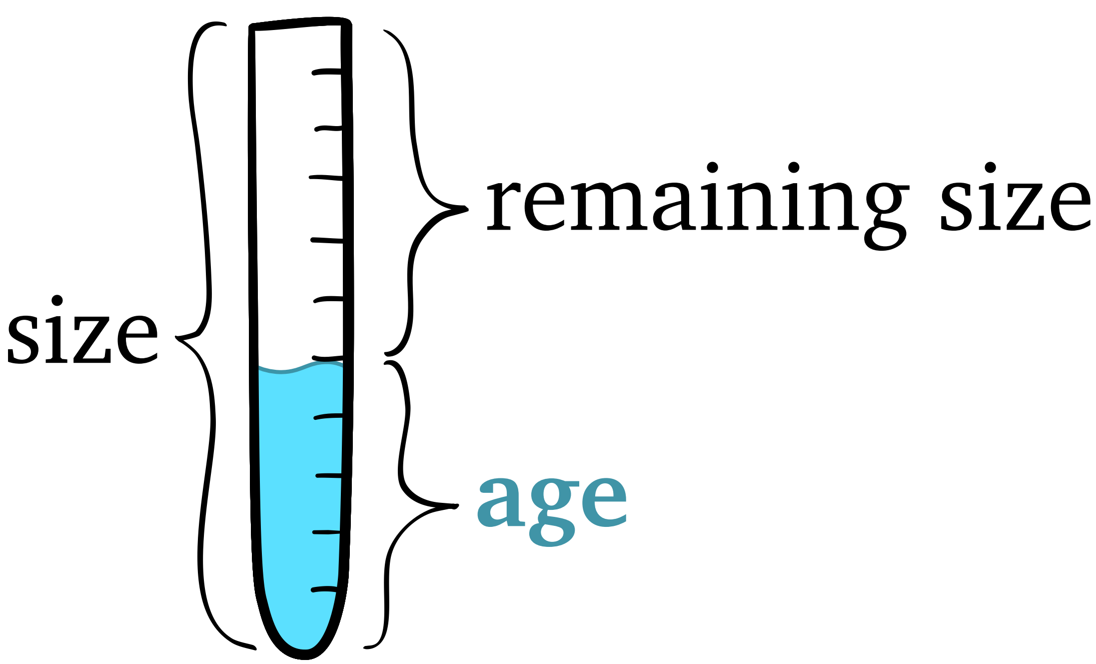

M/G/1 Queue

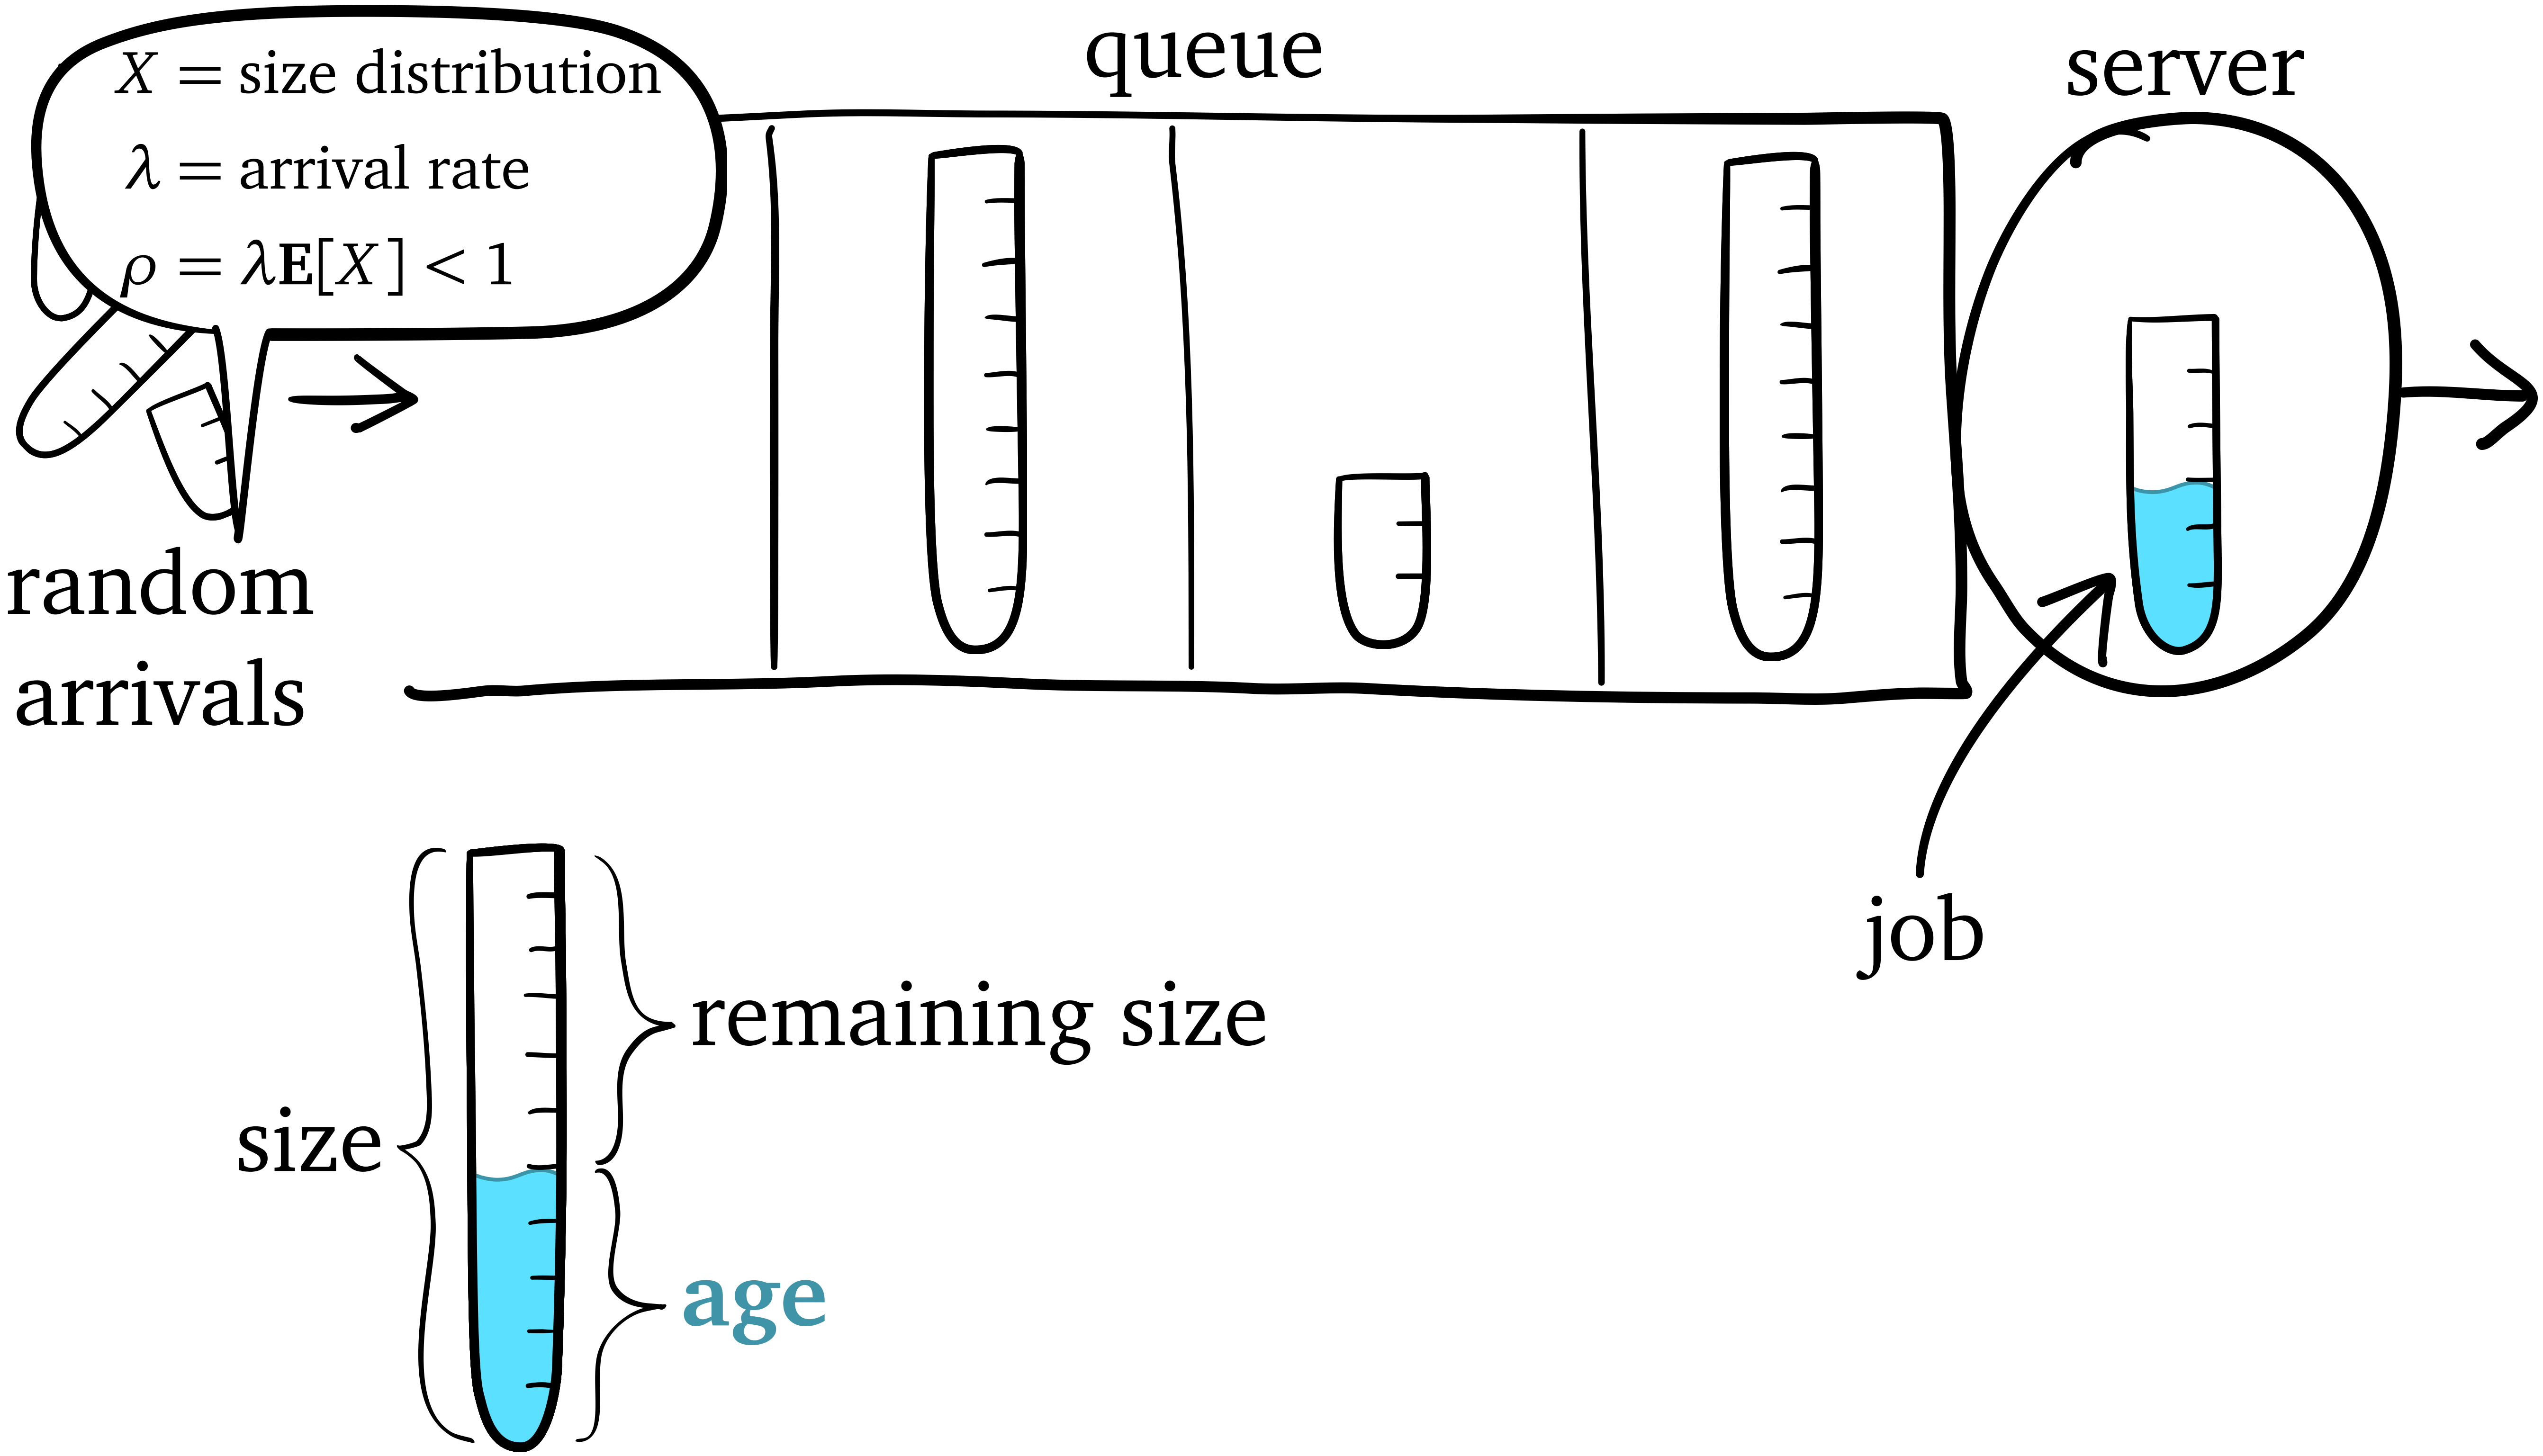

X = size distribution
λ = arrival rate
$\rho = \lambda \mathbf{E}[X] < 1$
queue
server
random arrivals
job
size
remaining size
age

M/G/1 Queue

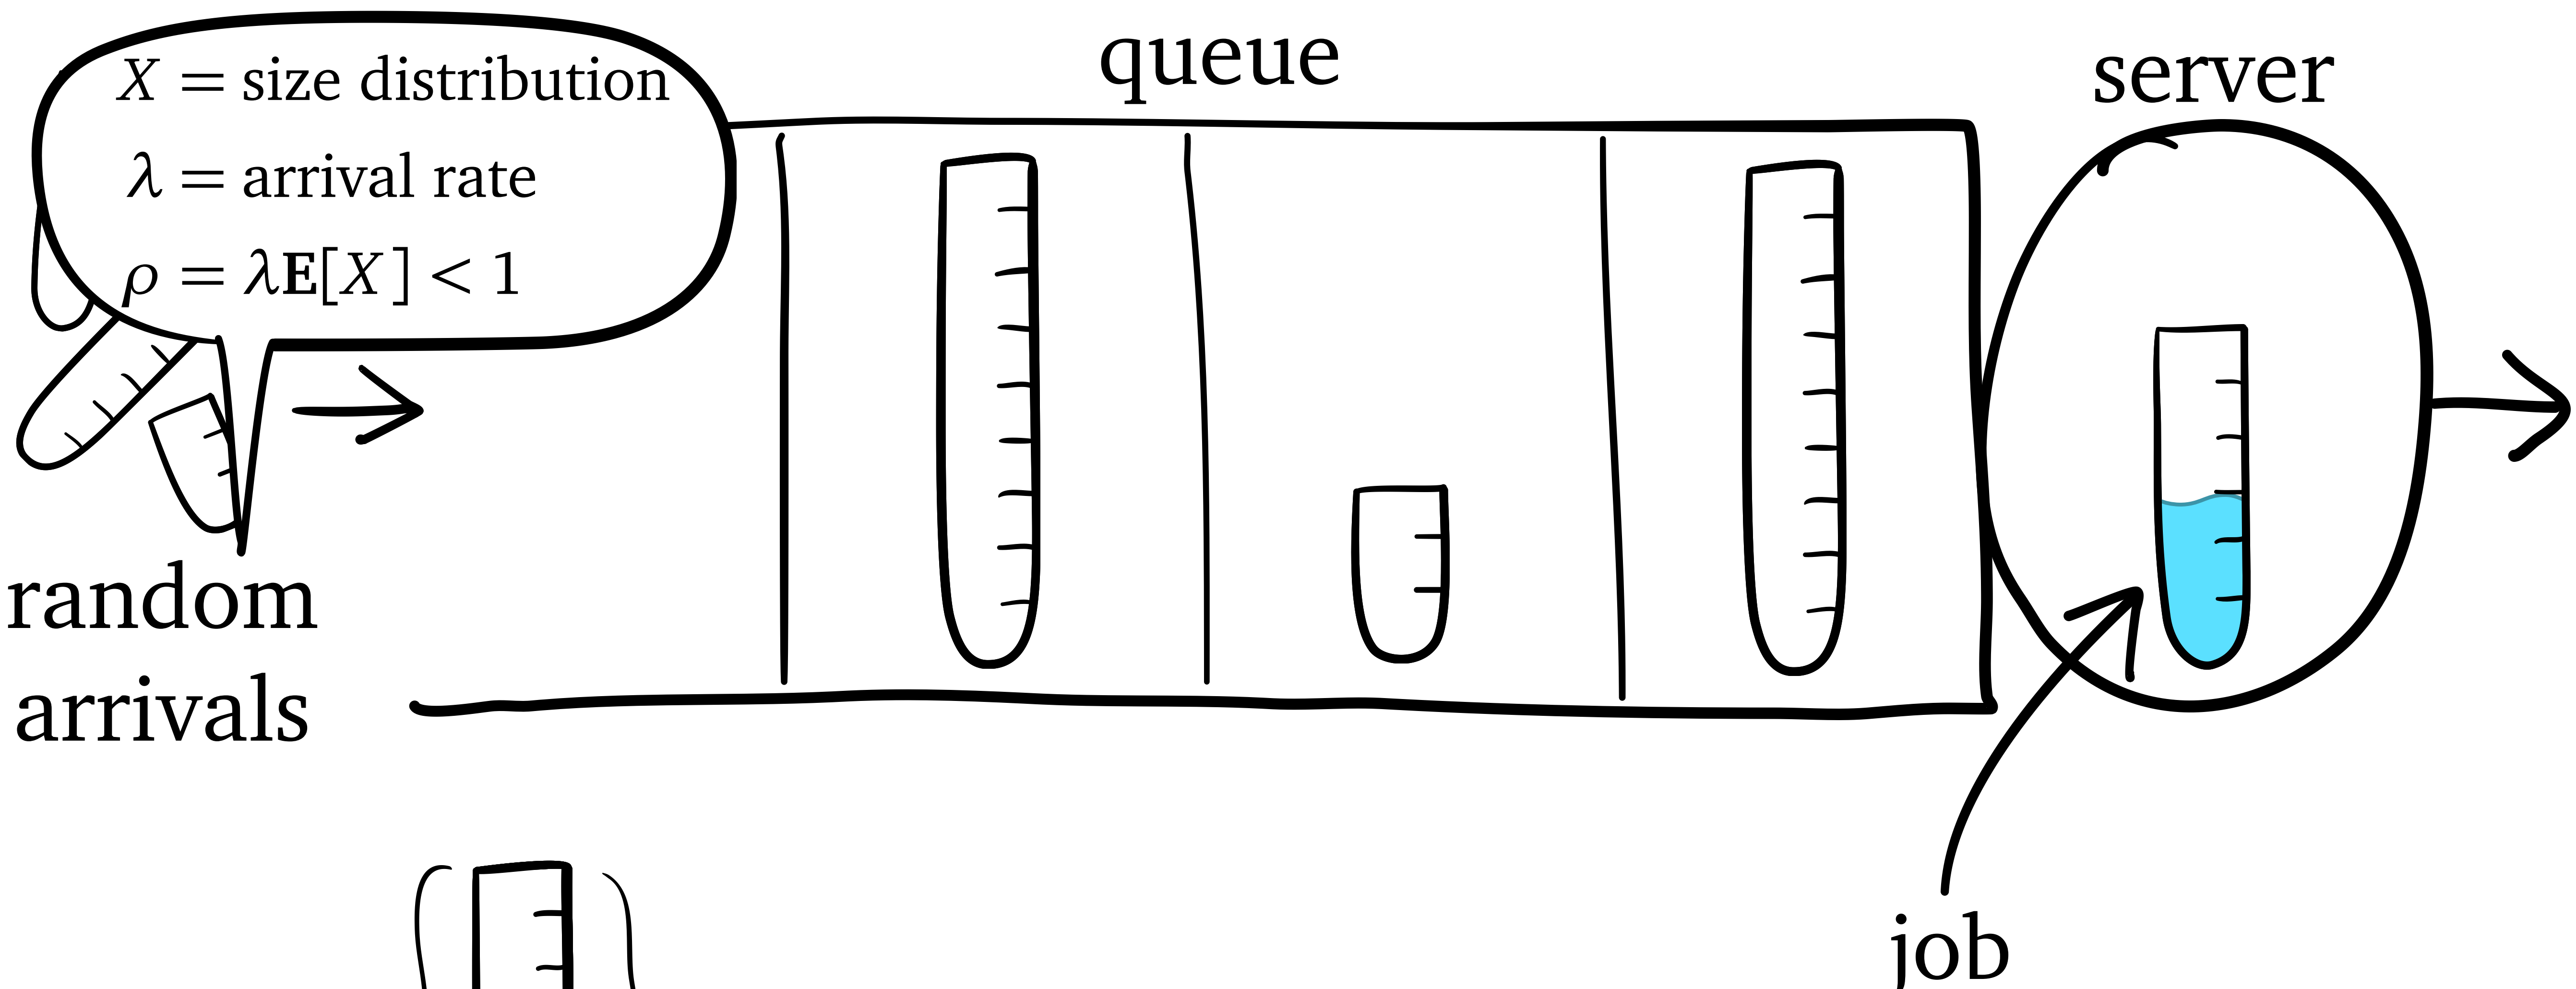

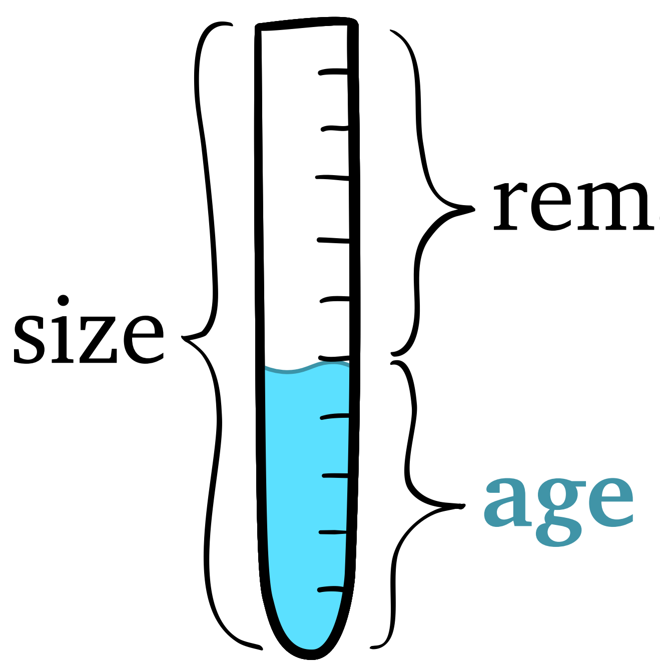

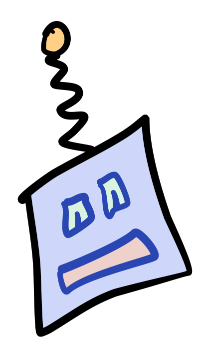

Scheduling policy:
picks which job to serve

M/G/1 Queue

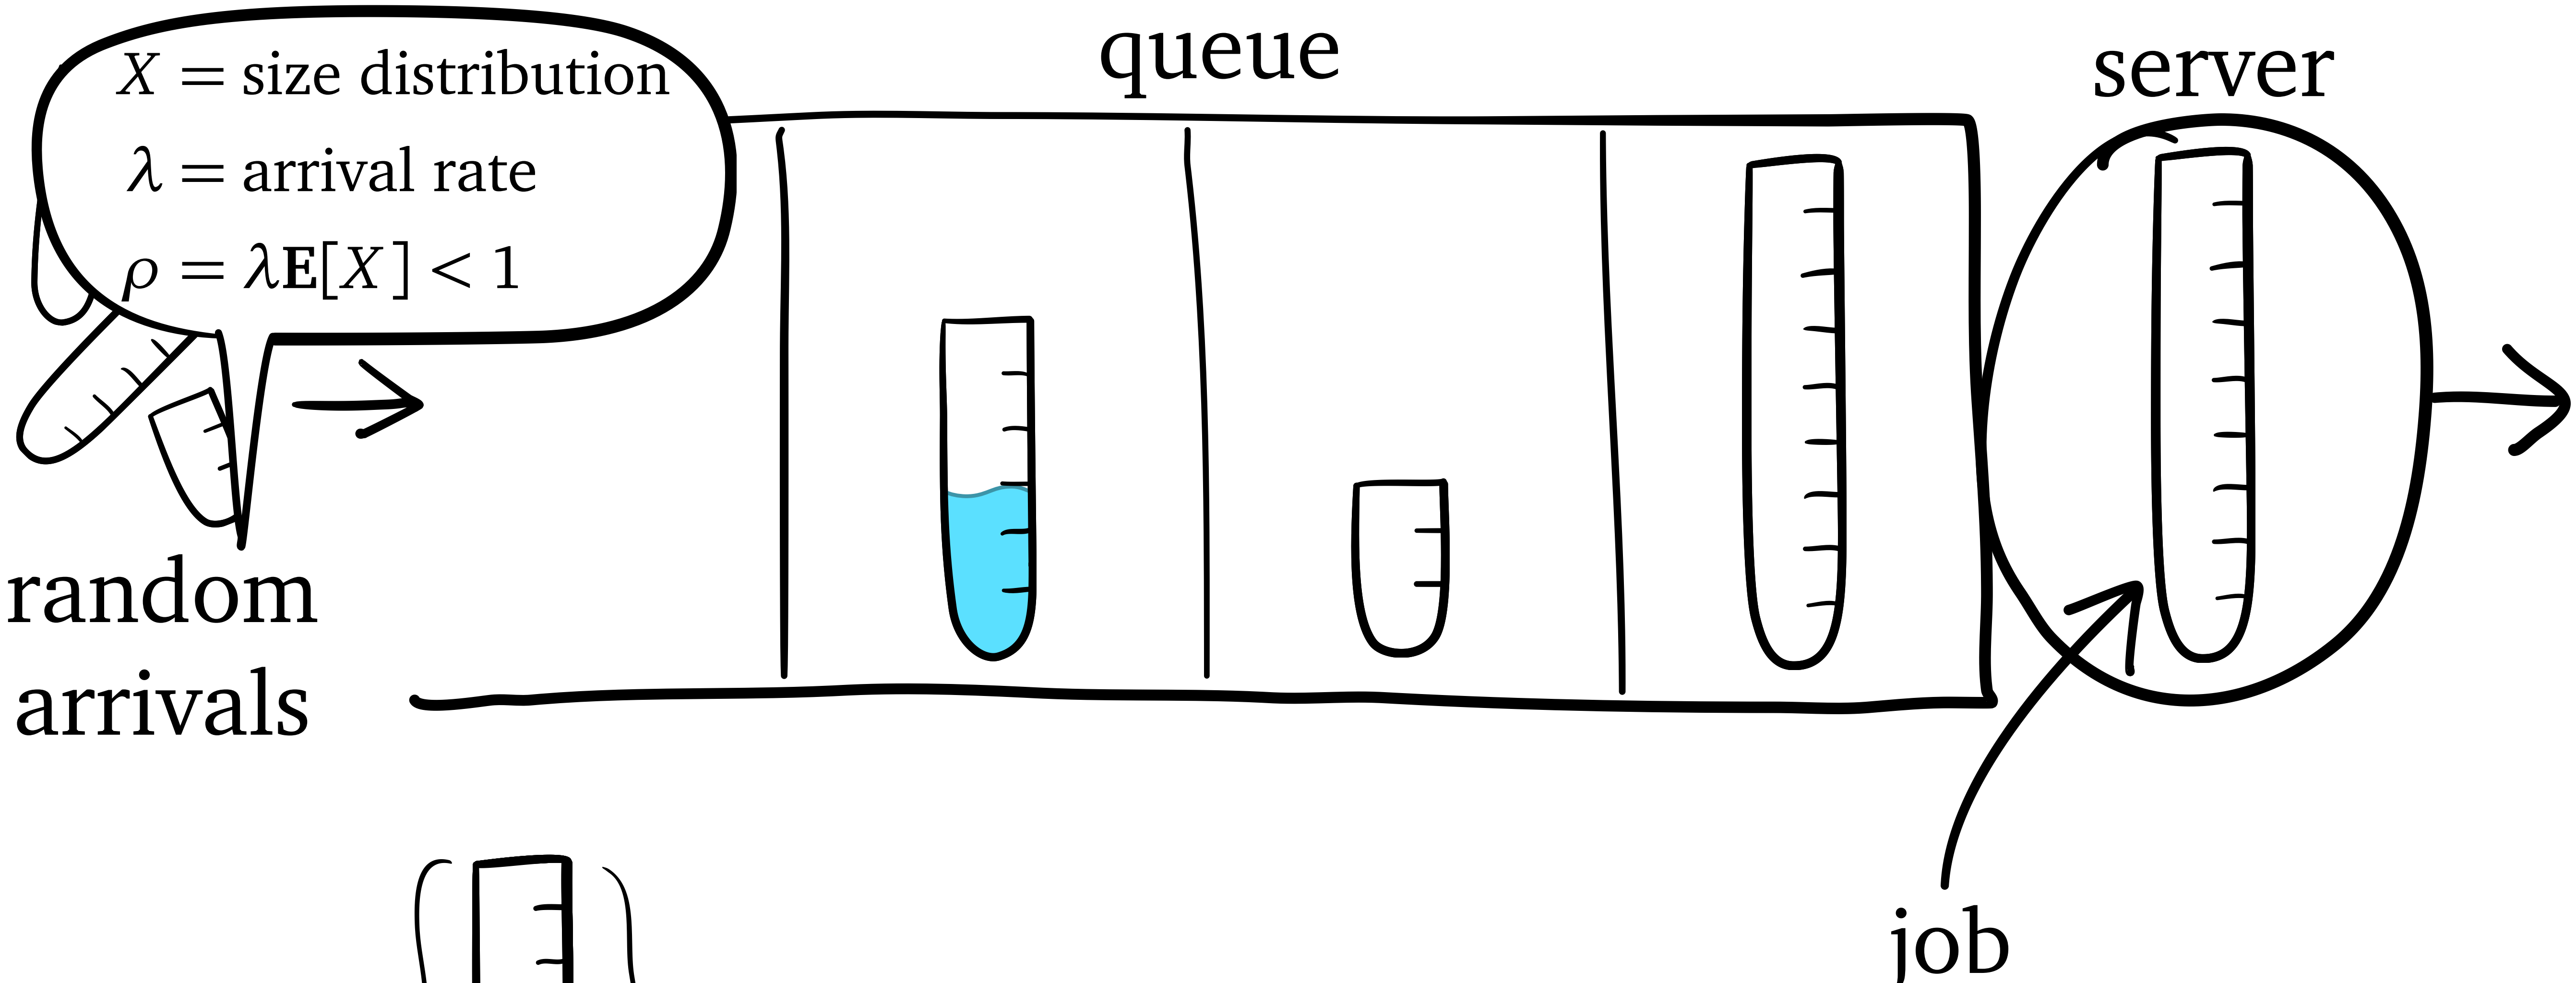

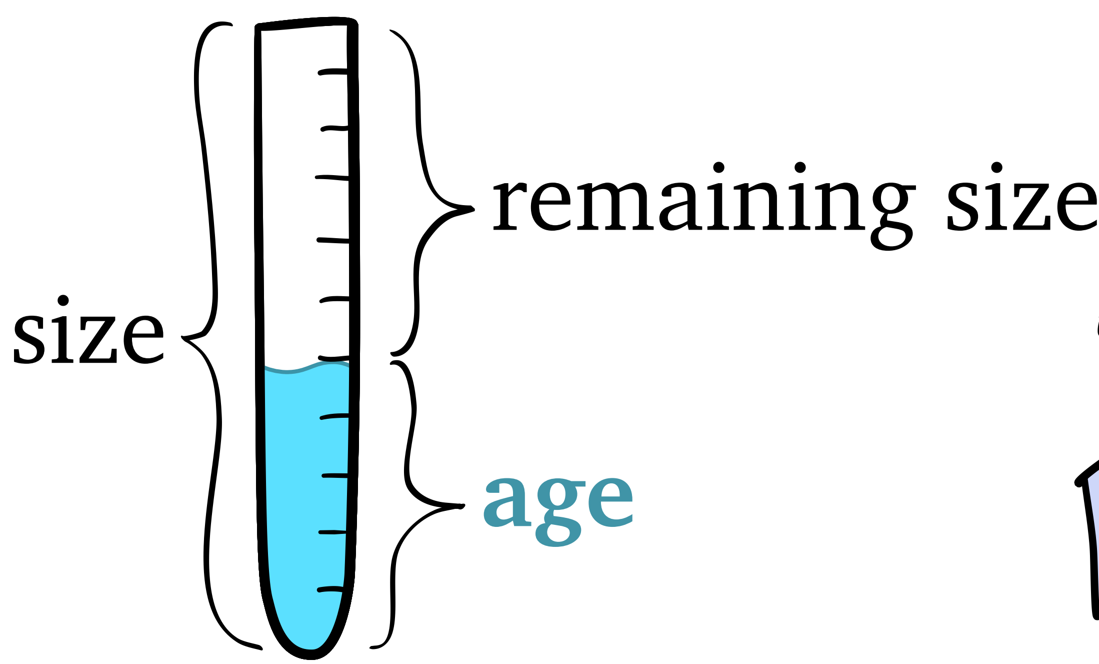

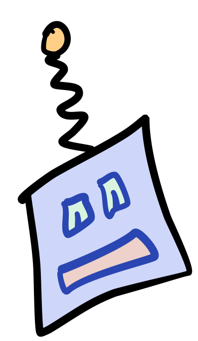

Scheduling policy:
picks which job to serve

M/G/1 Queue

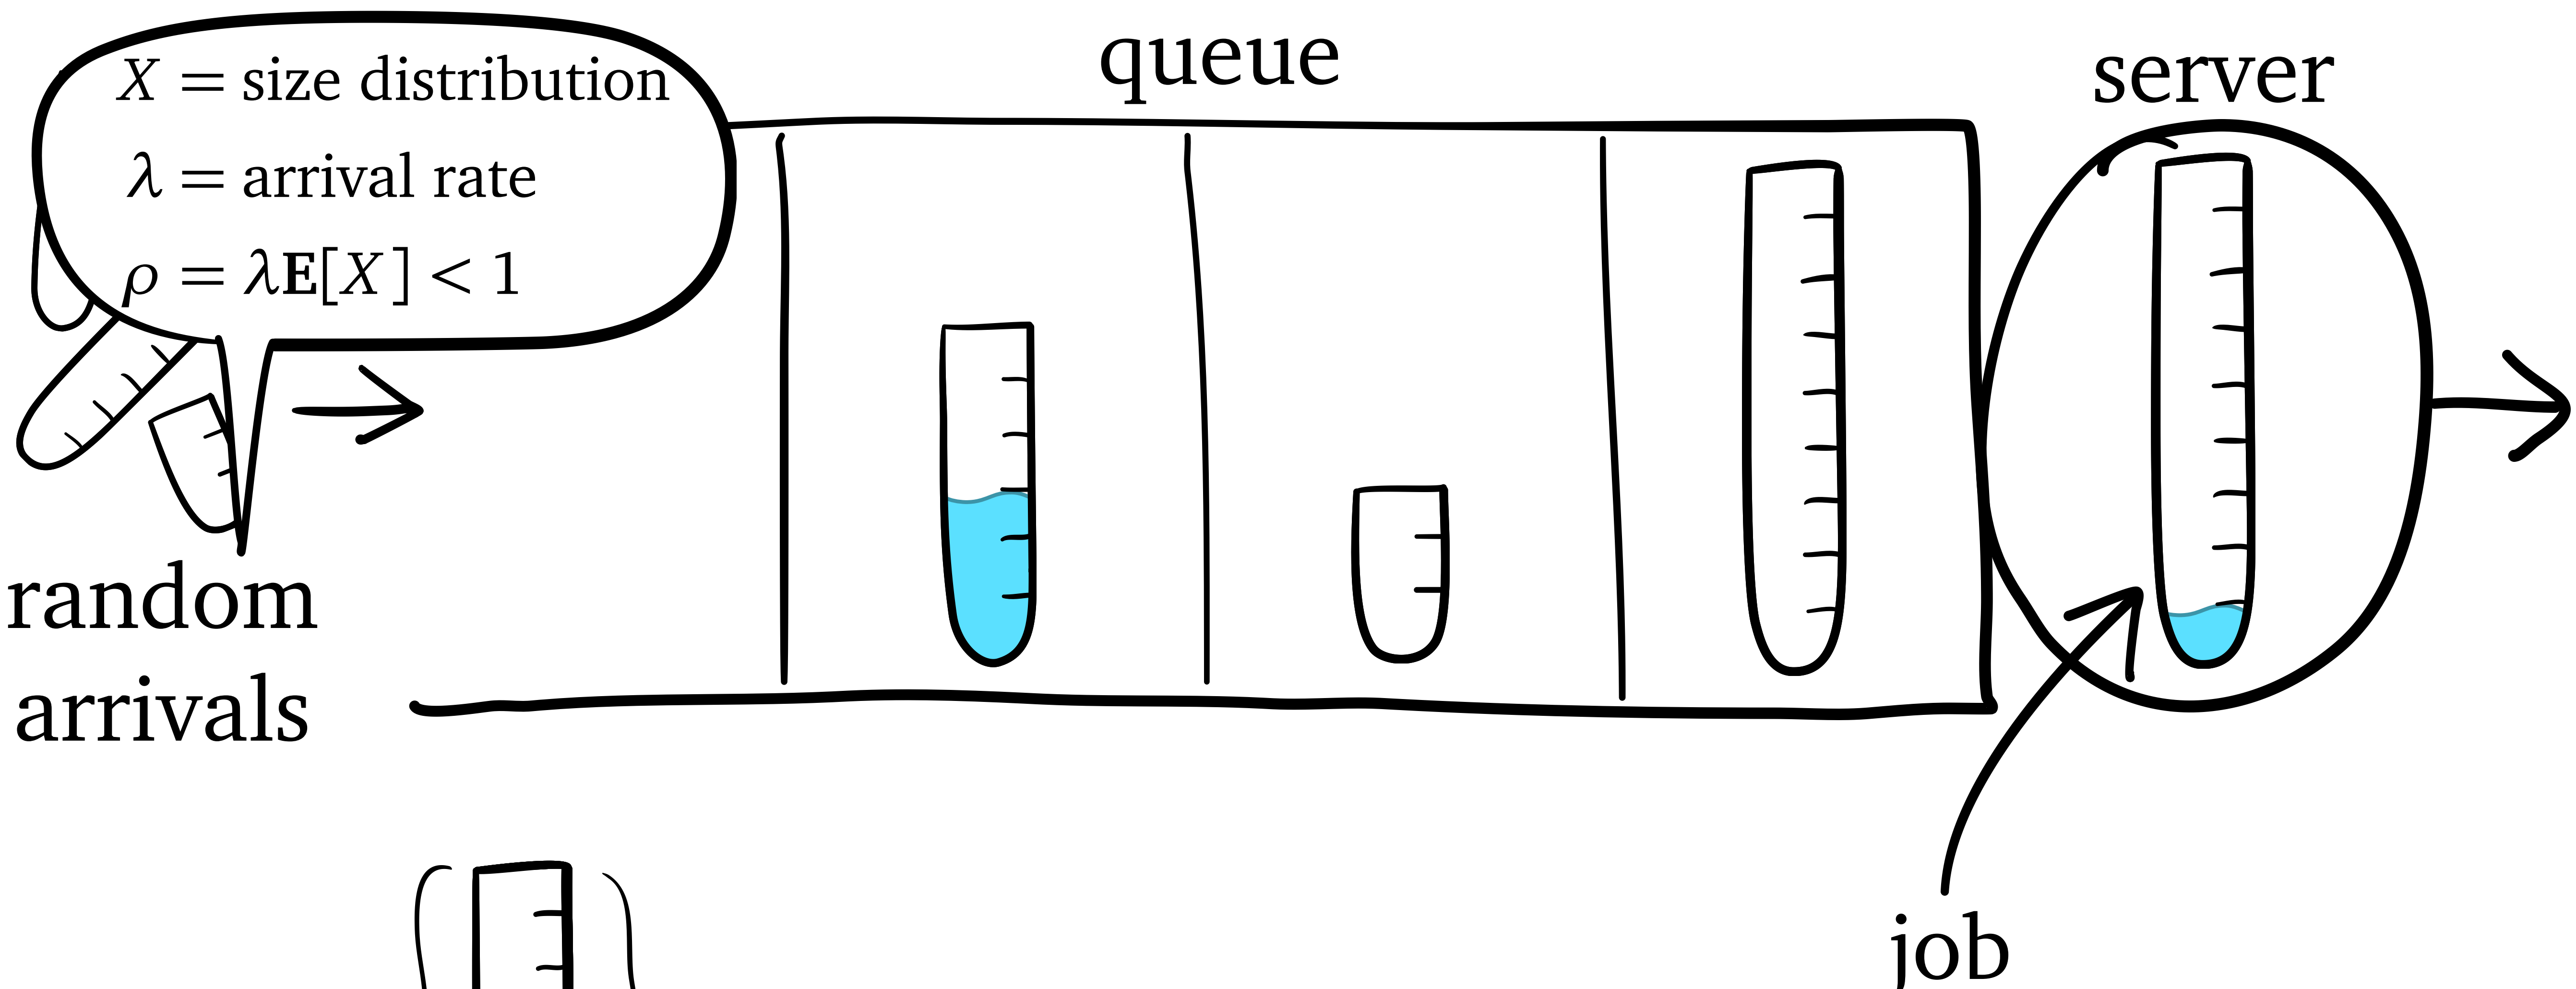

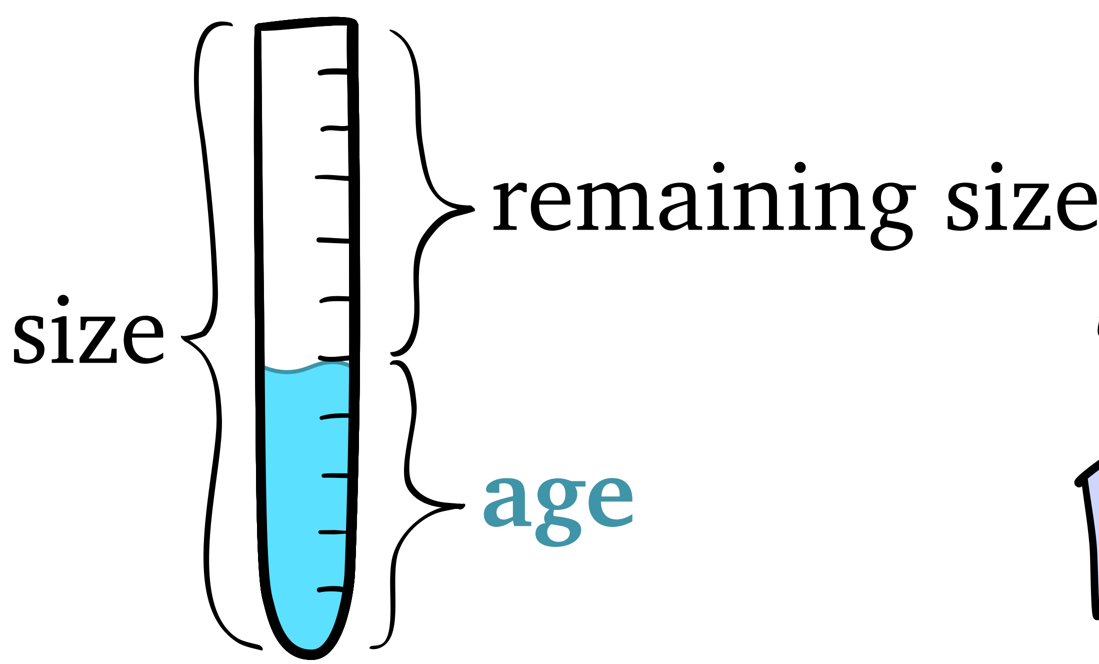

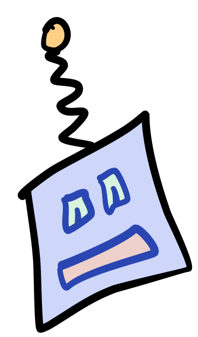

Scheduling policy:
picks which job to serve

M/G/1 Queue

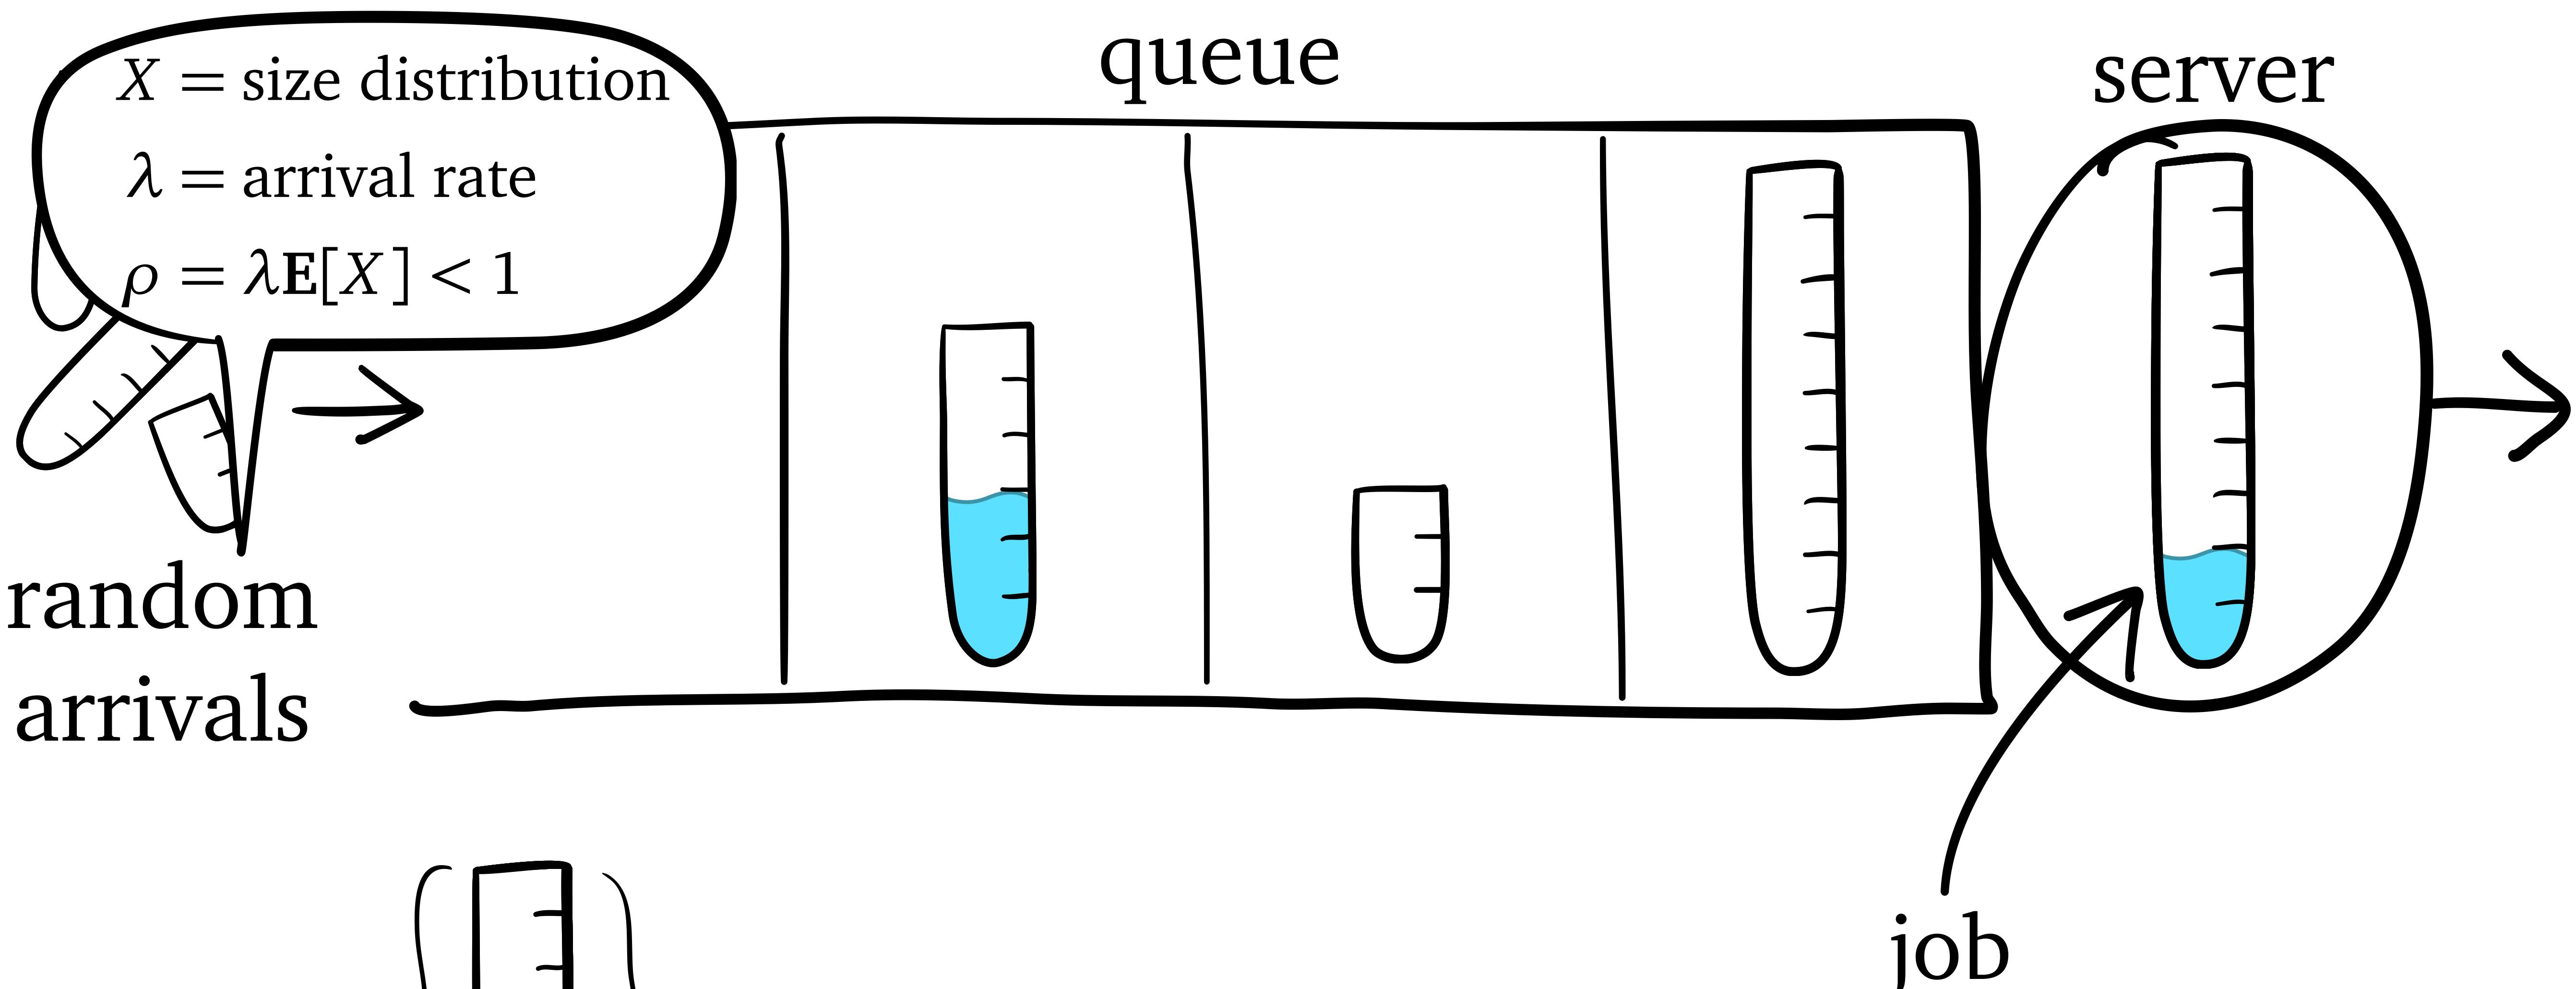

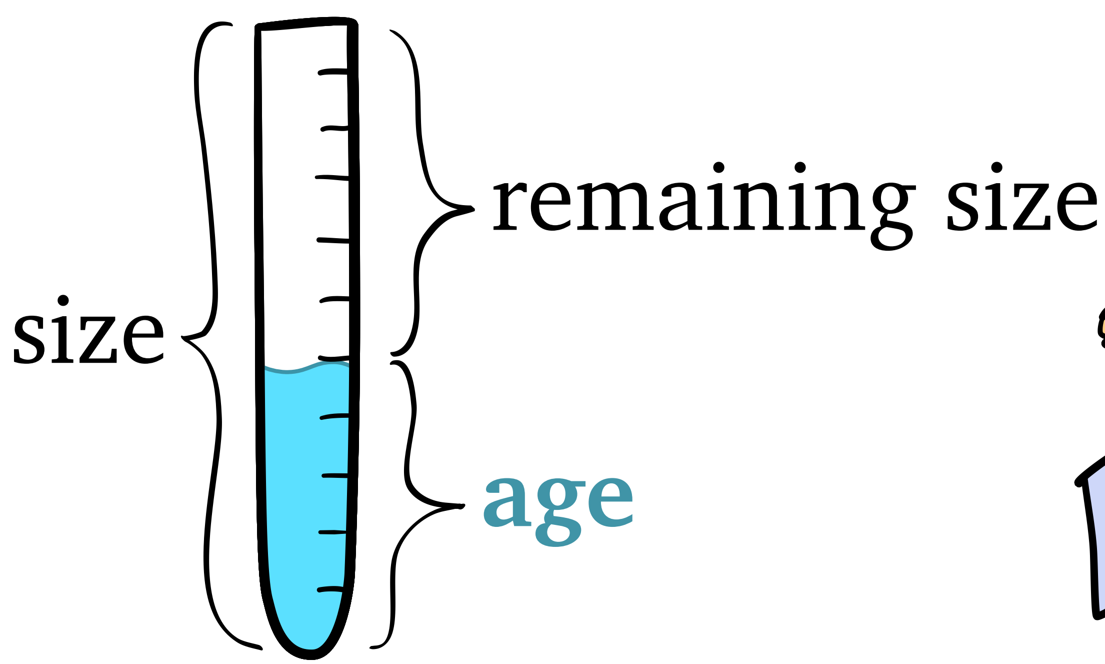

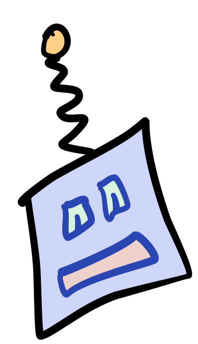

Scheduling policy:
picks which job to serve

M/G/1 Queue

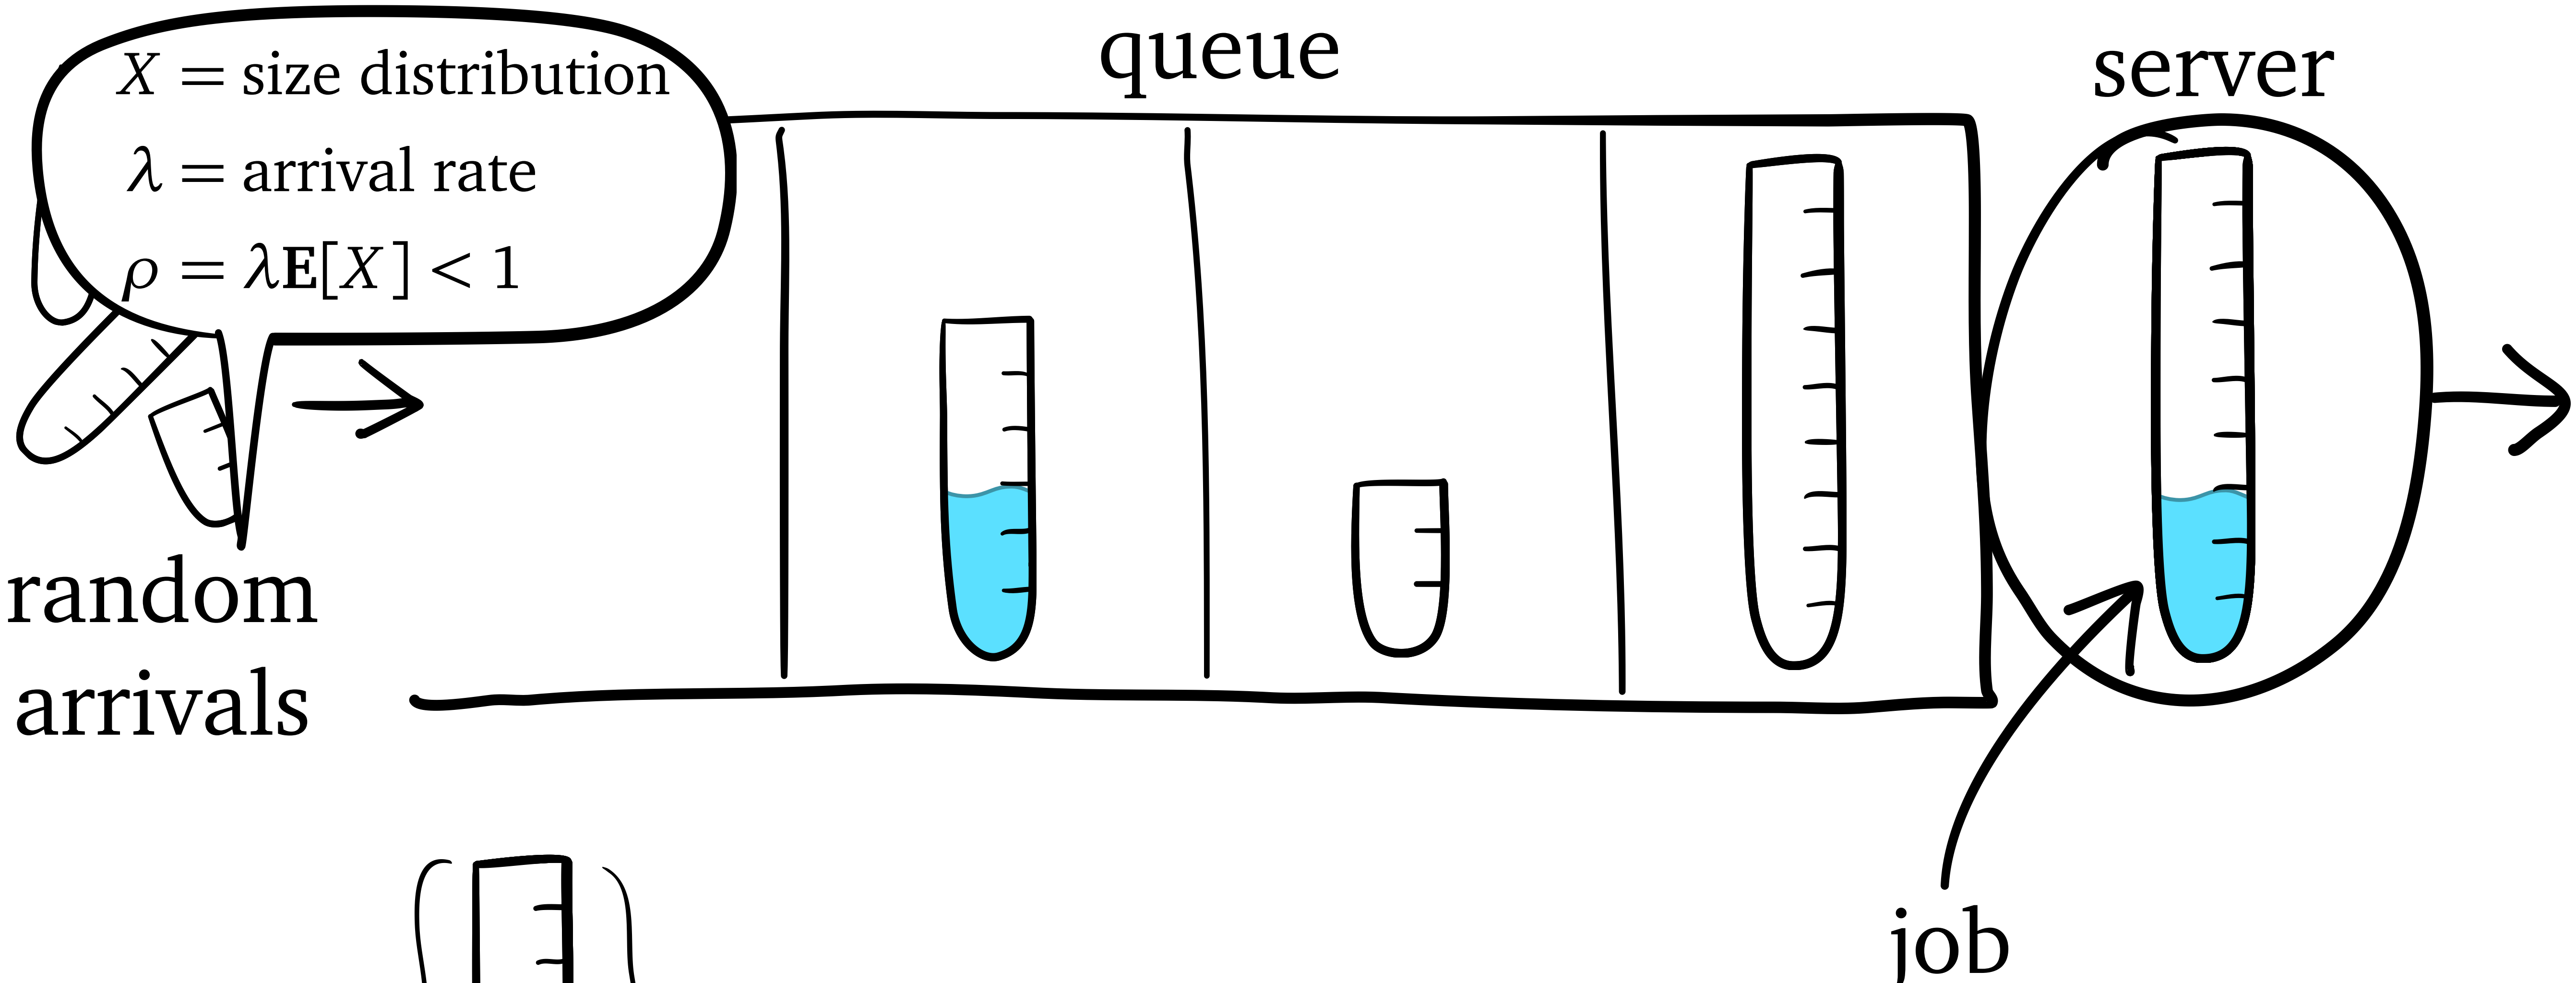

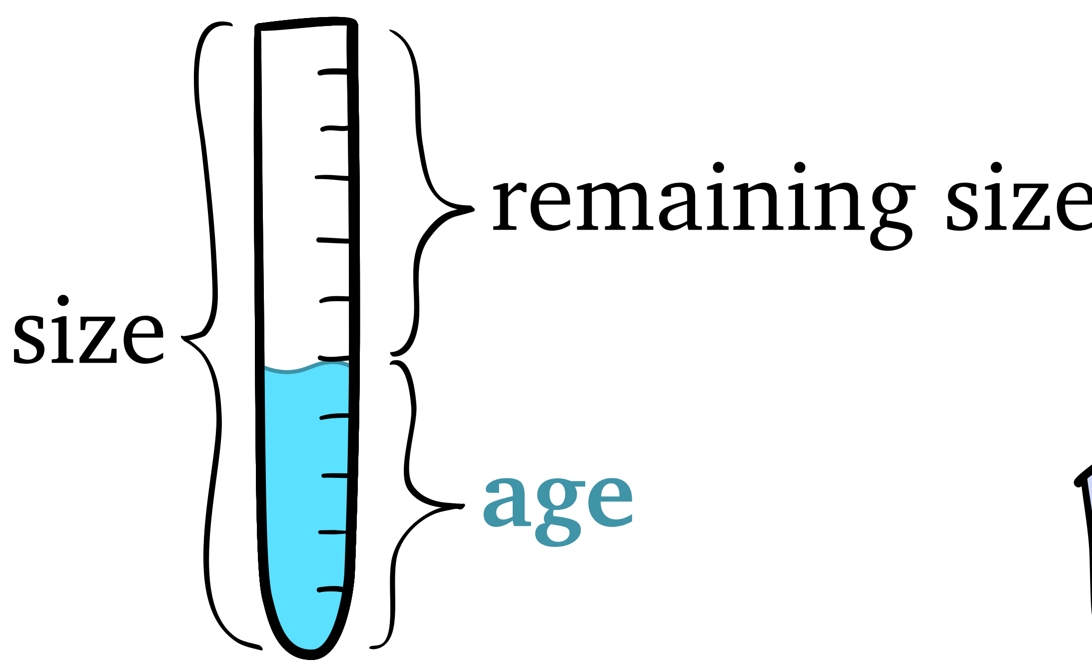

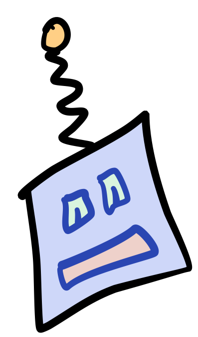

Scheduling policy:
picks which job to serve

M/G/k Queue

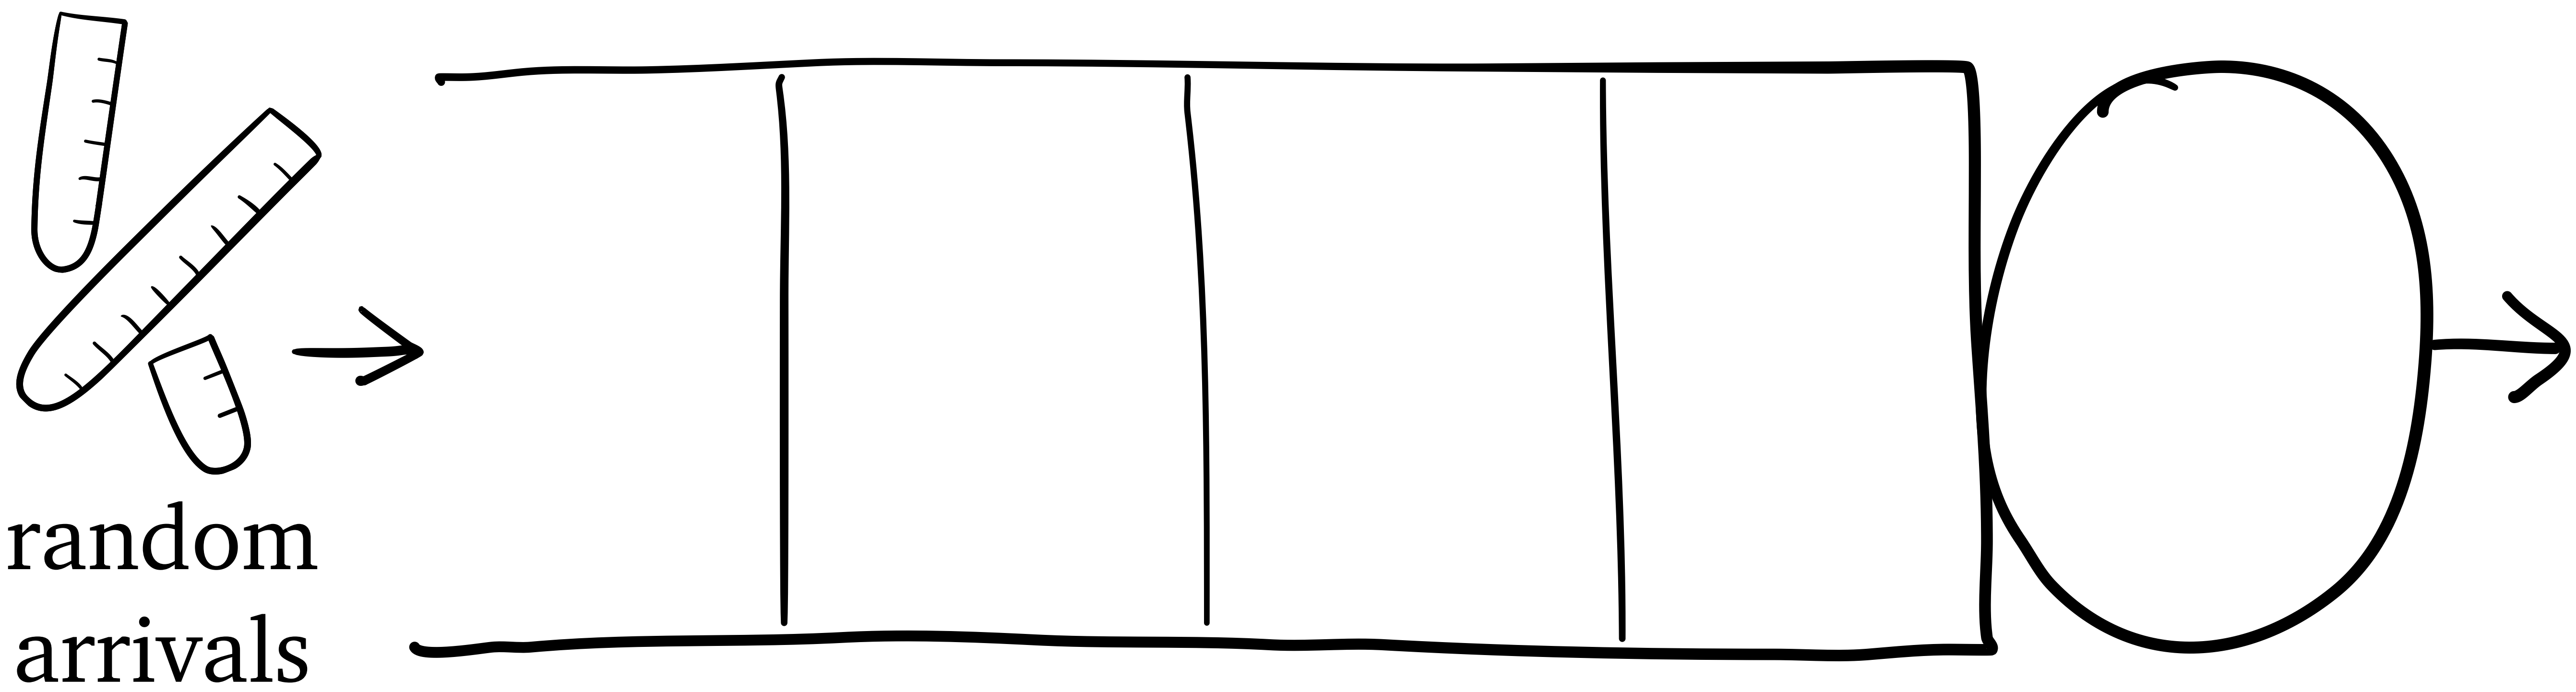

M/G/k Queue

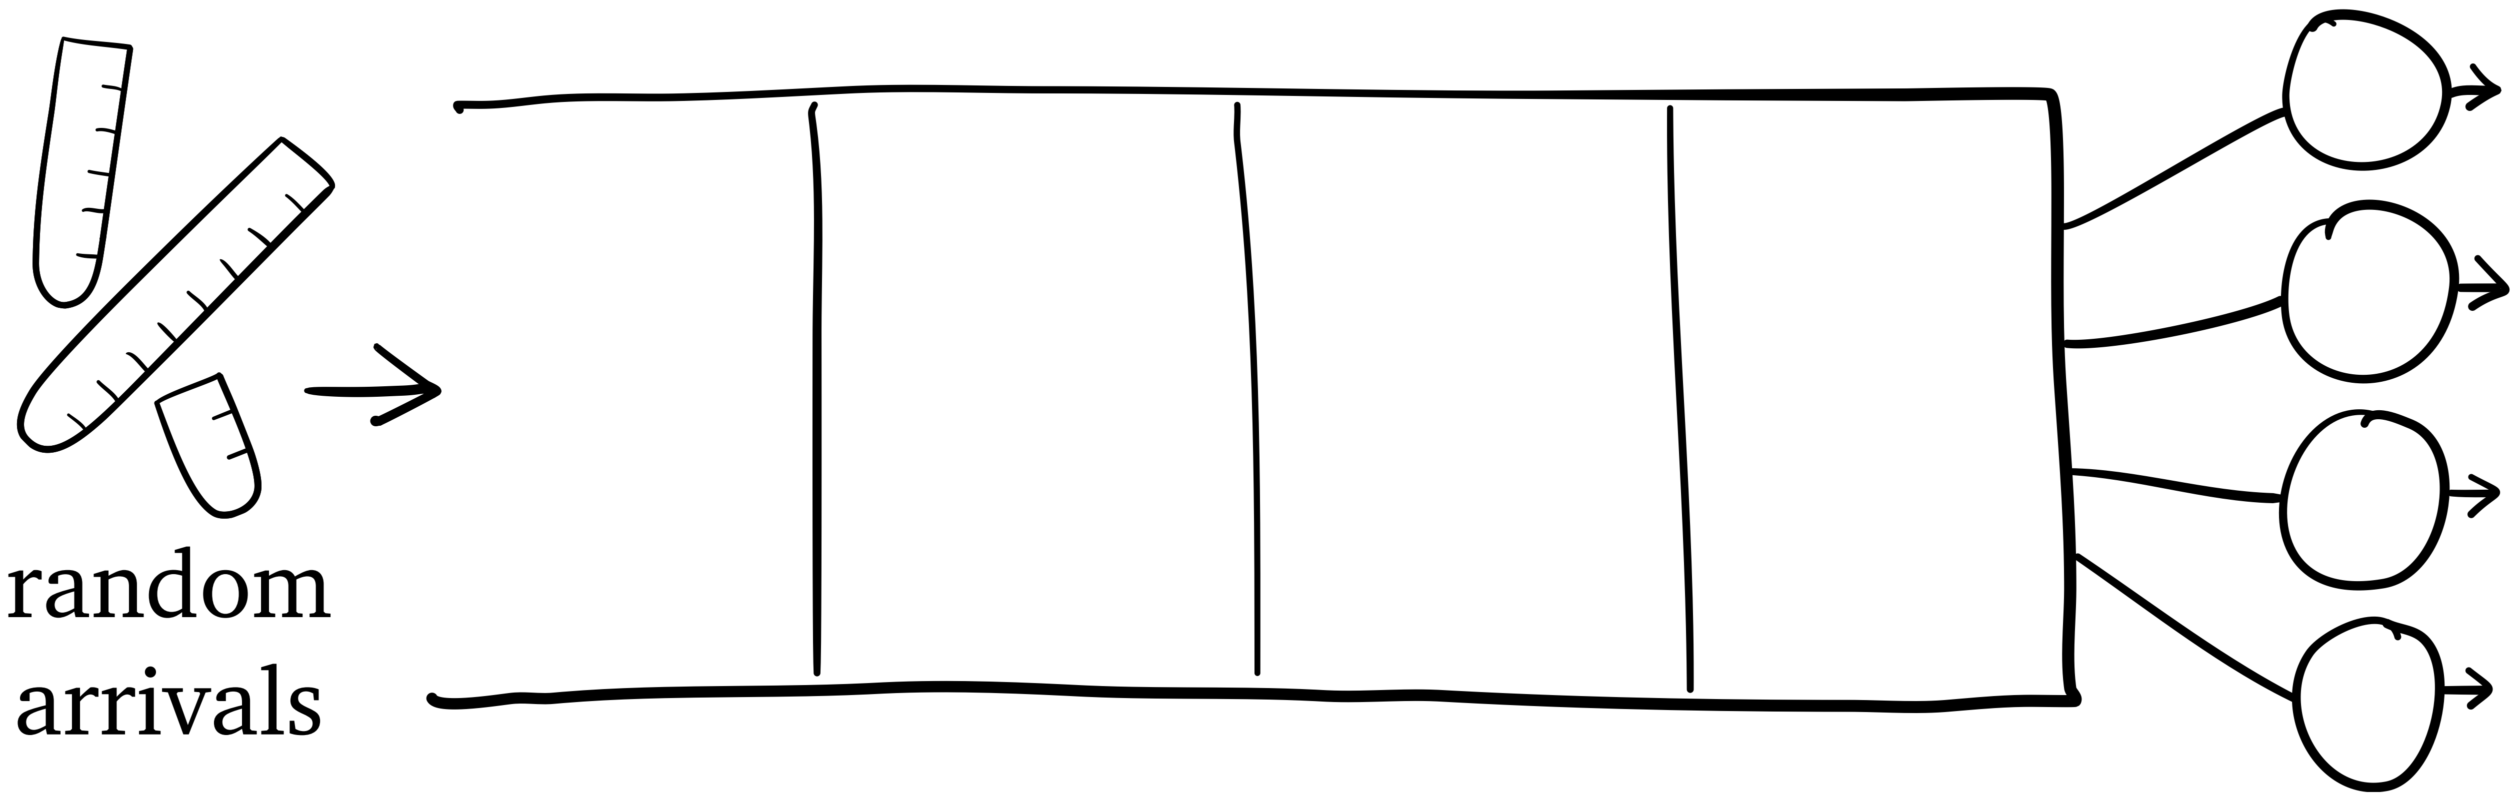

M/G/k Queue

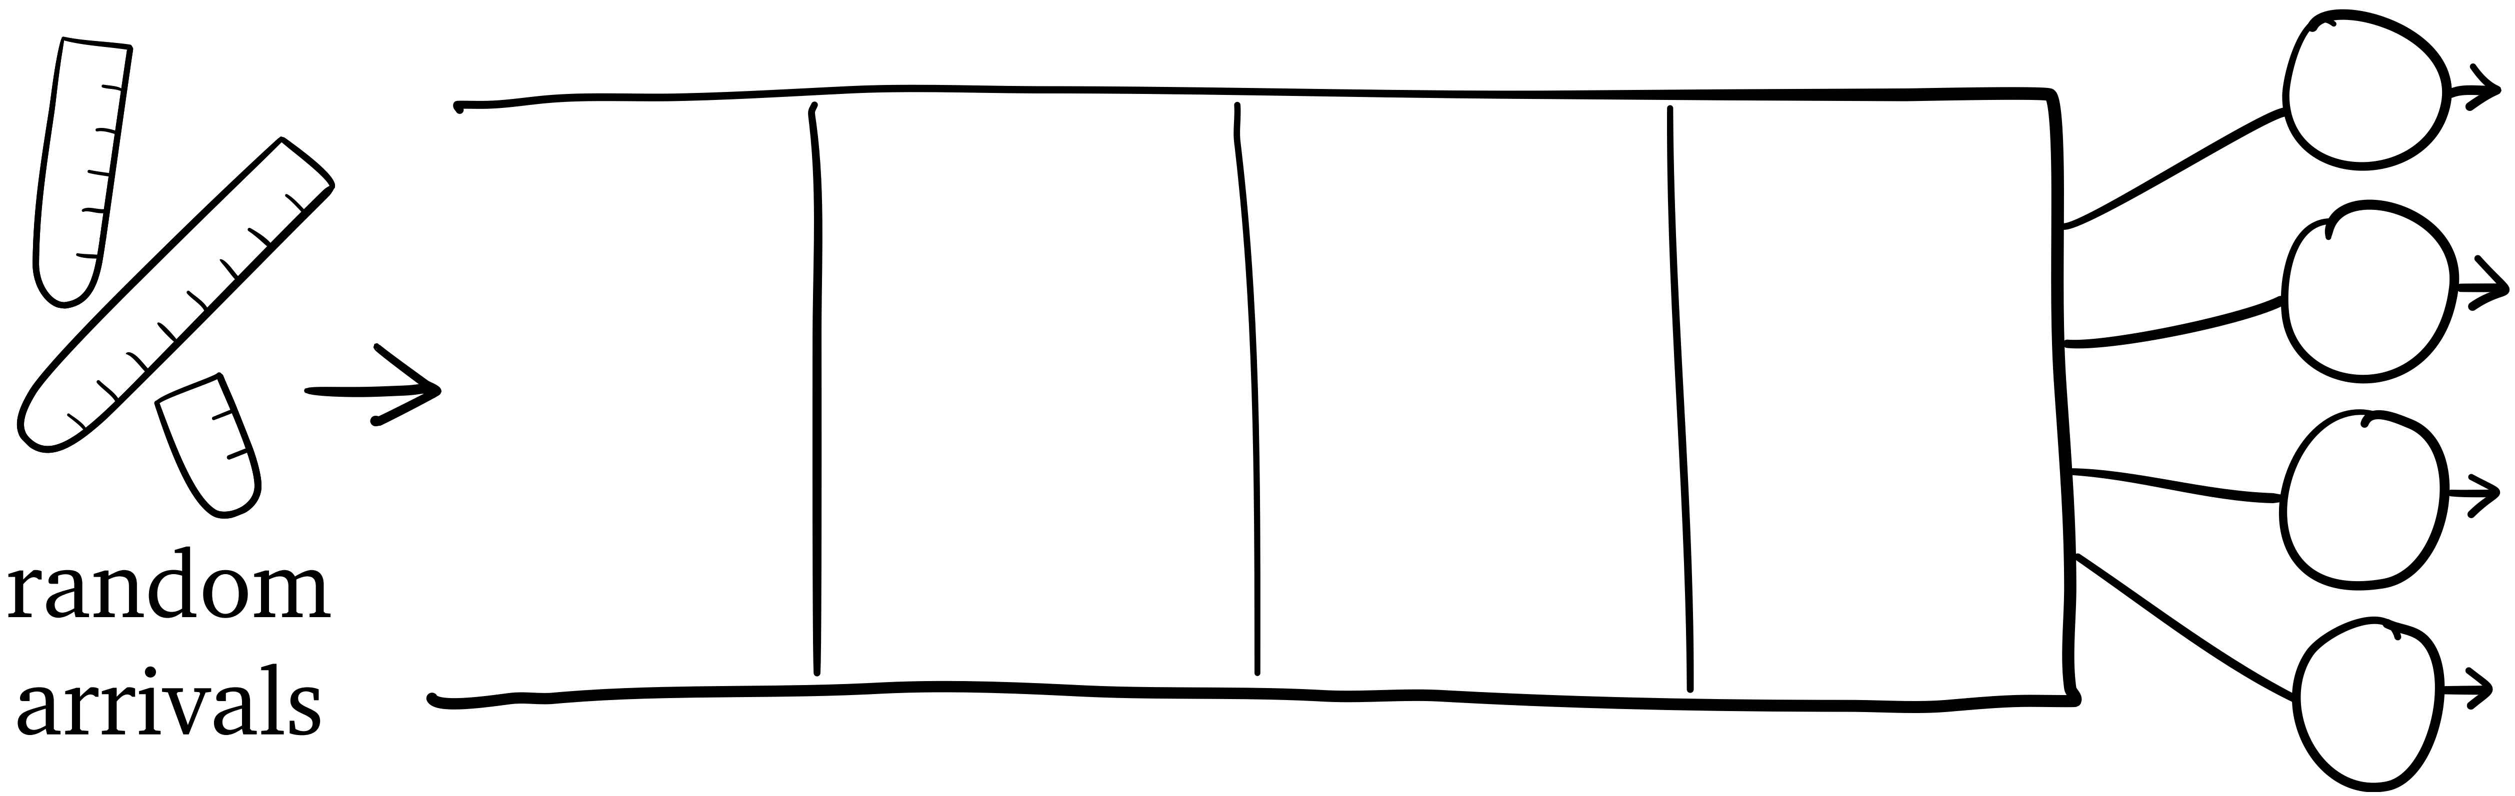

k servers, each speed $1/k$

M/G/k Queue

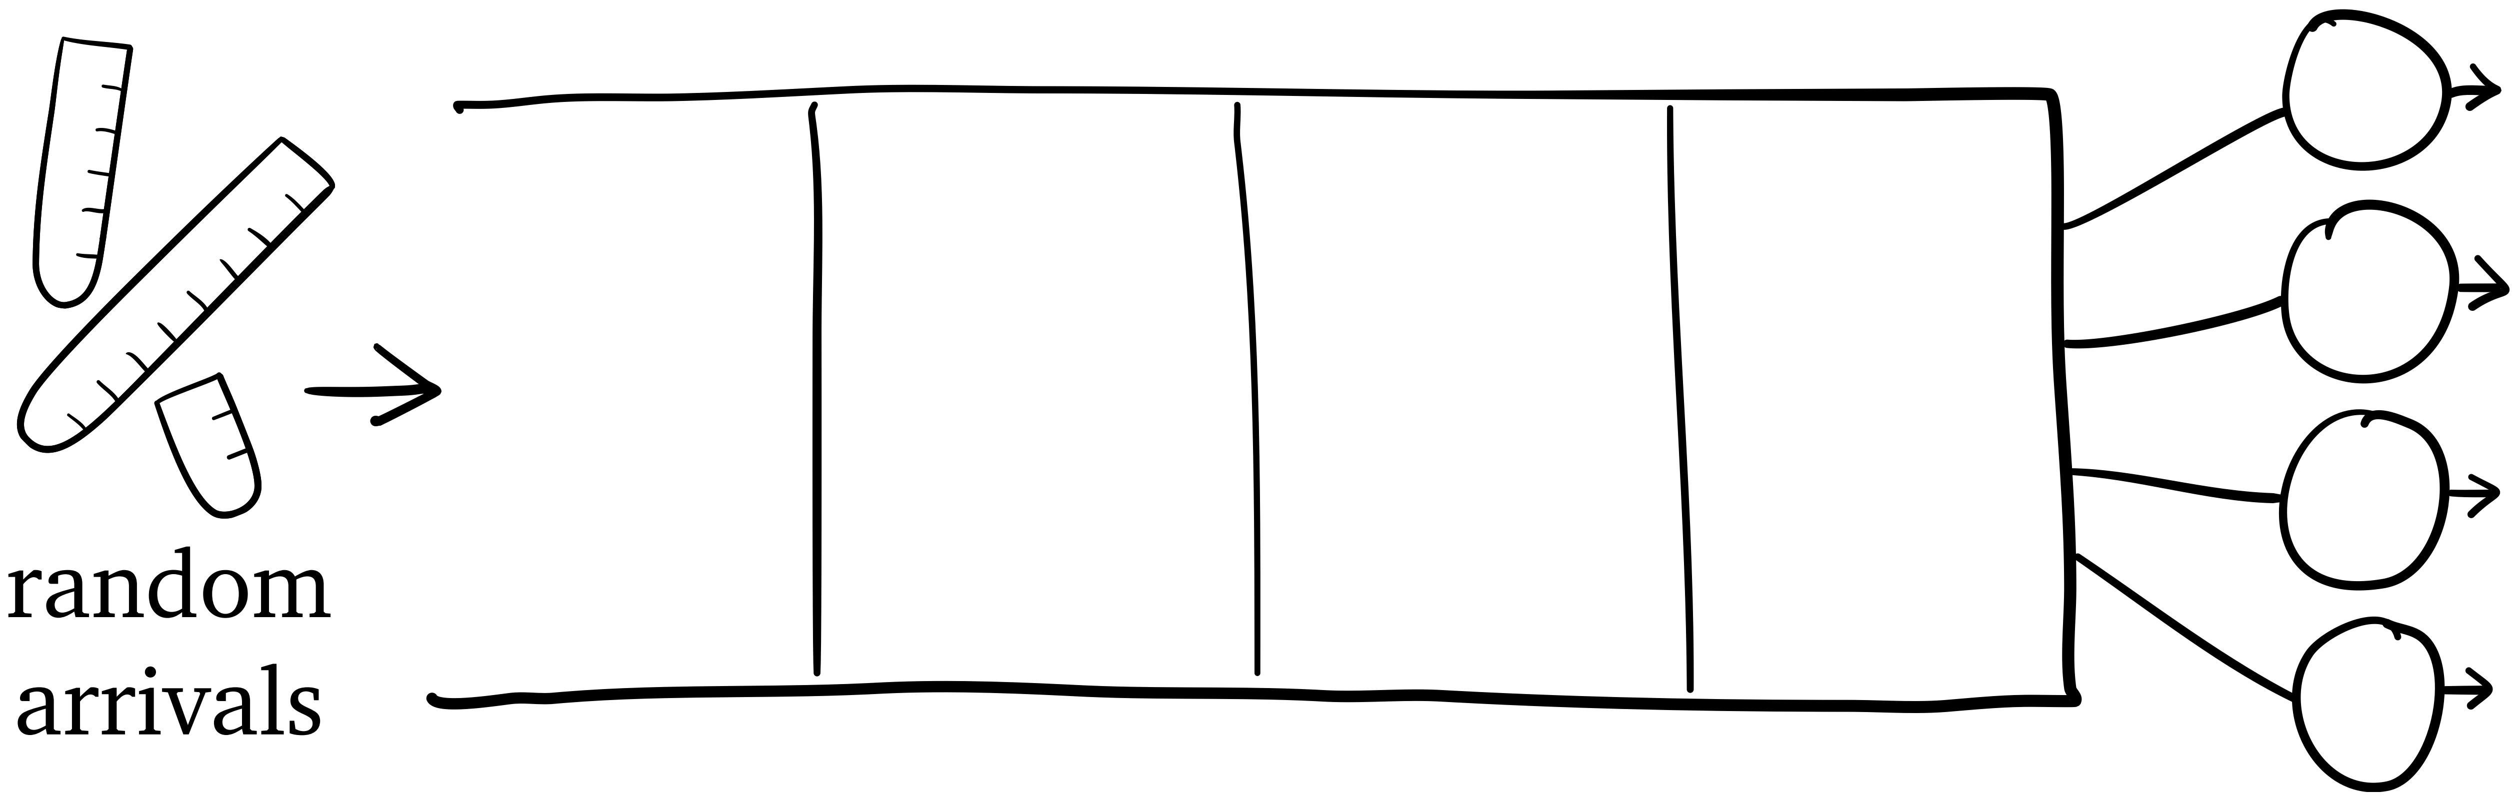

k servers, each speed $1/k$

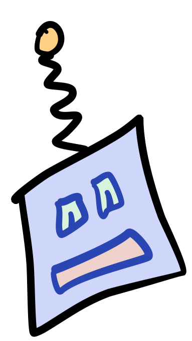

Scheduling policy:
picks which k jobs to serve

Response Time

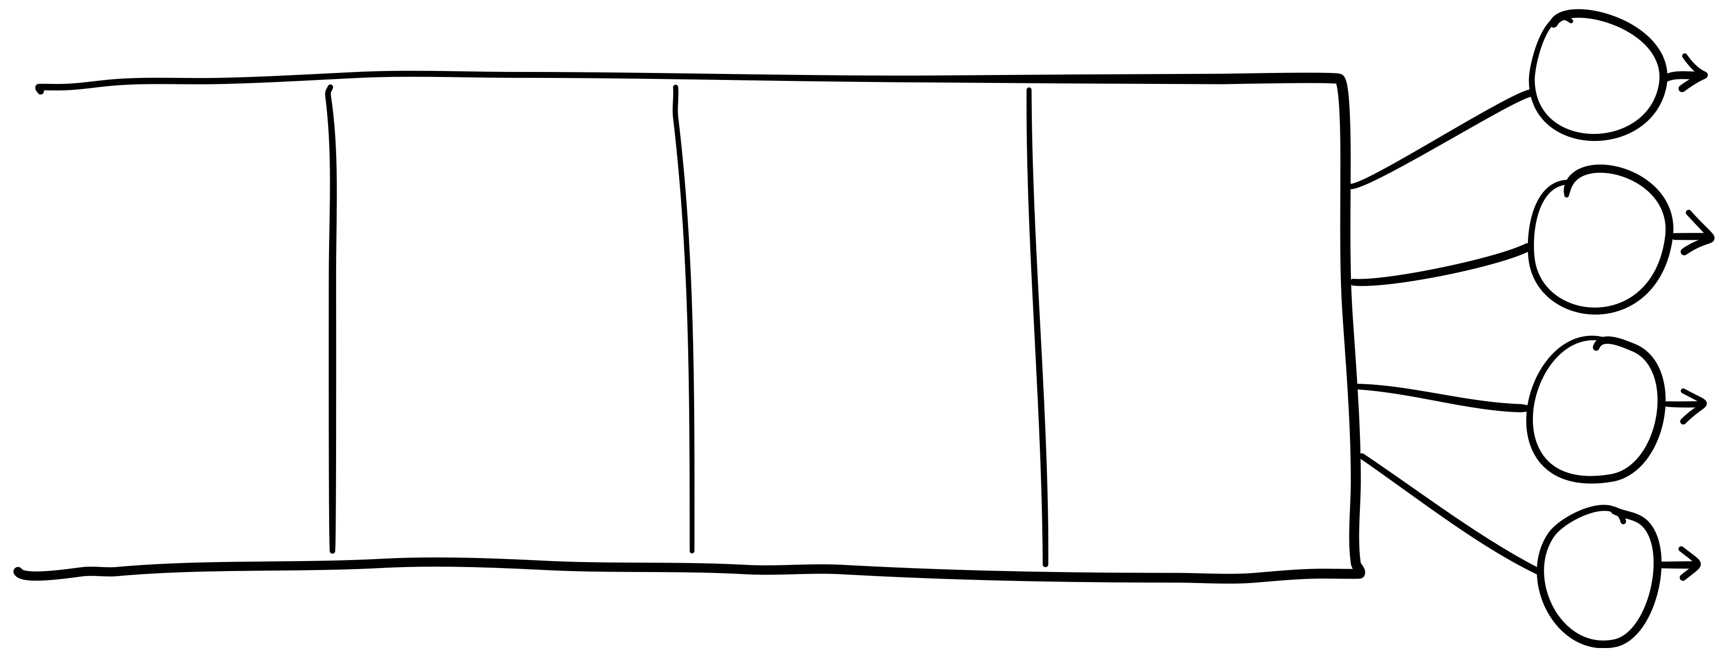

Response Time

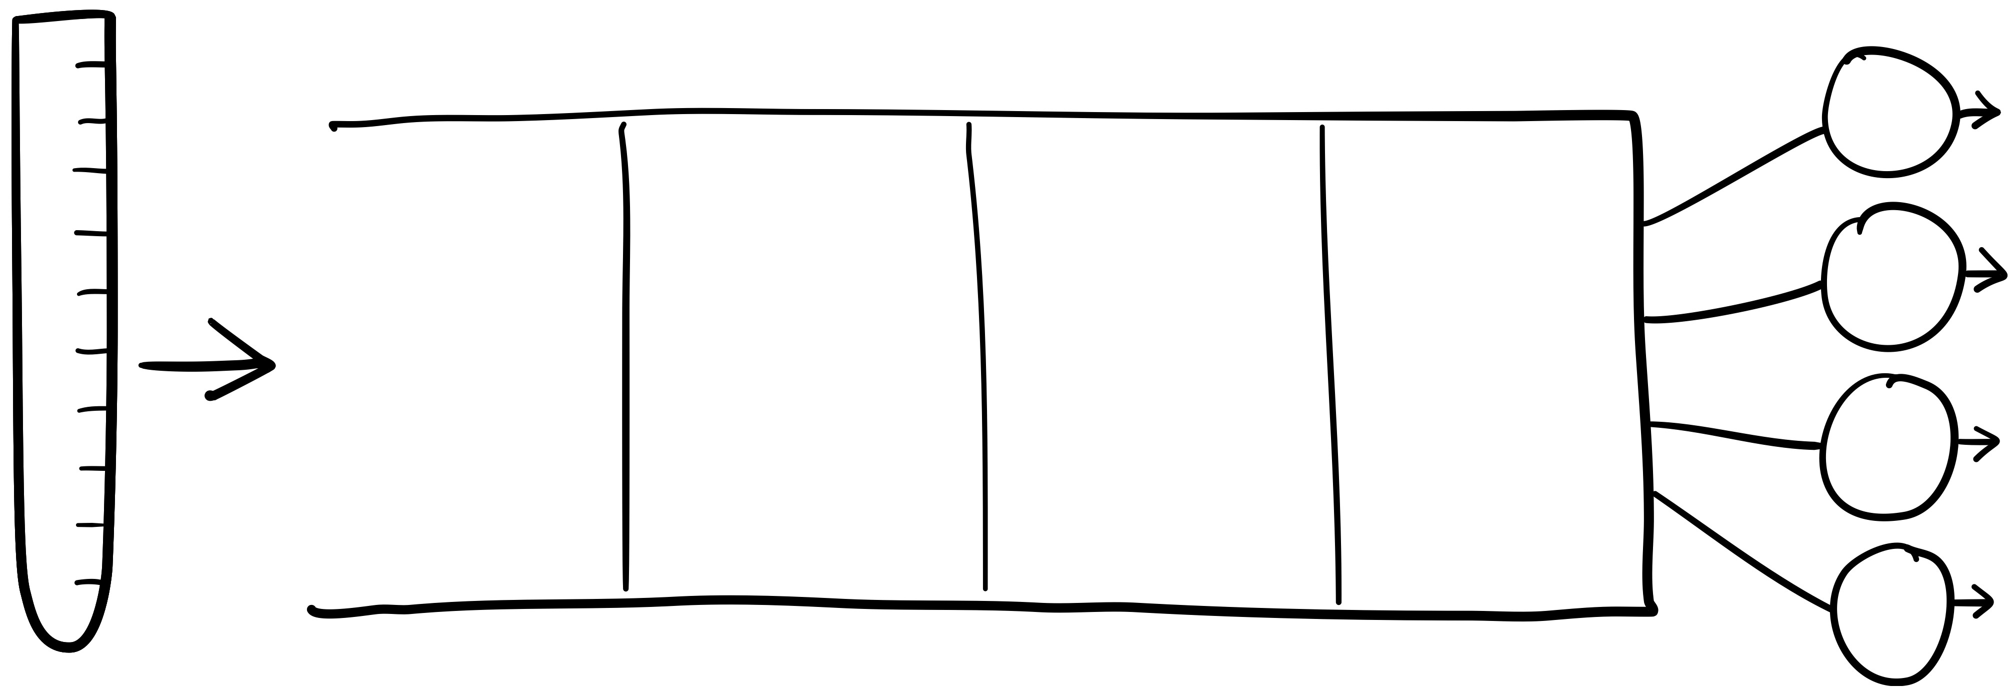

Response Time

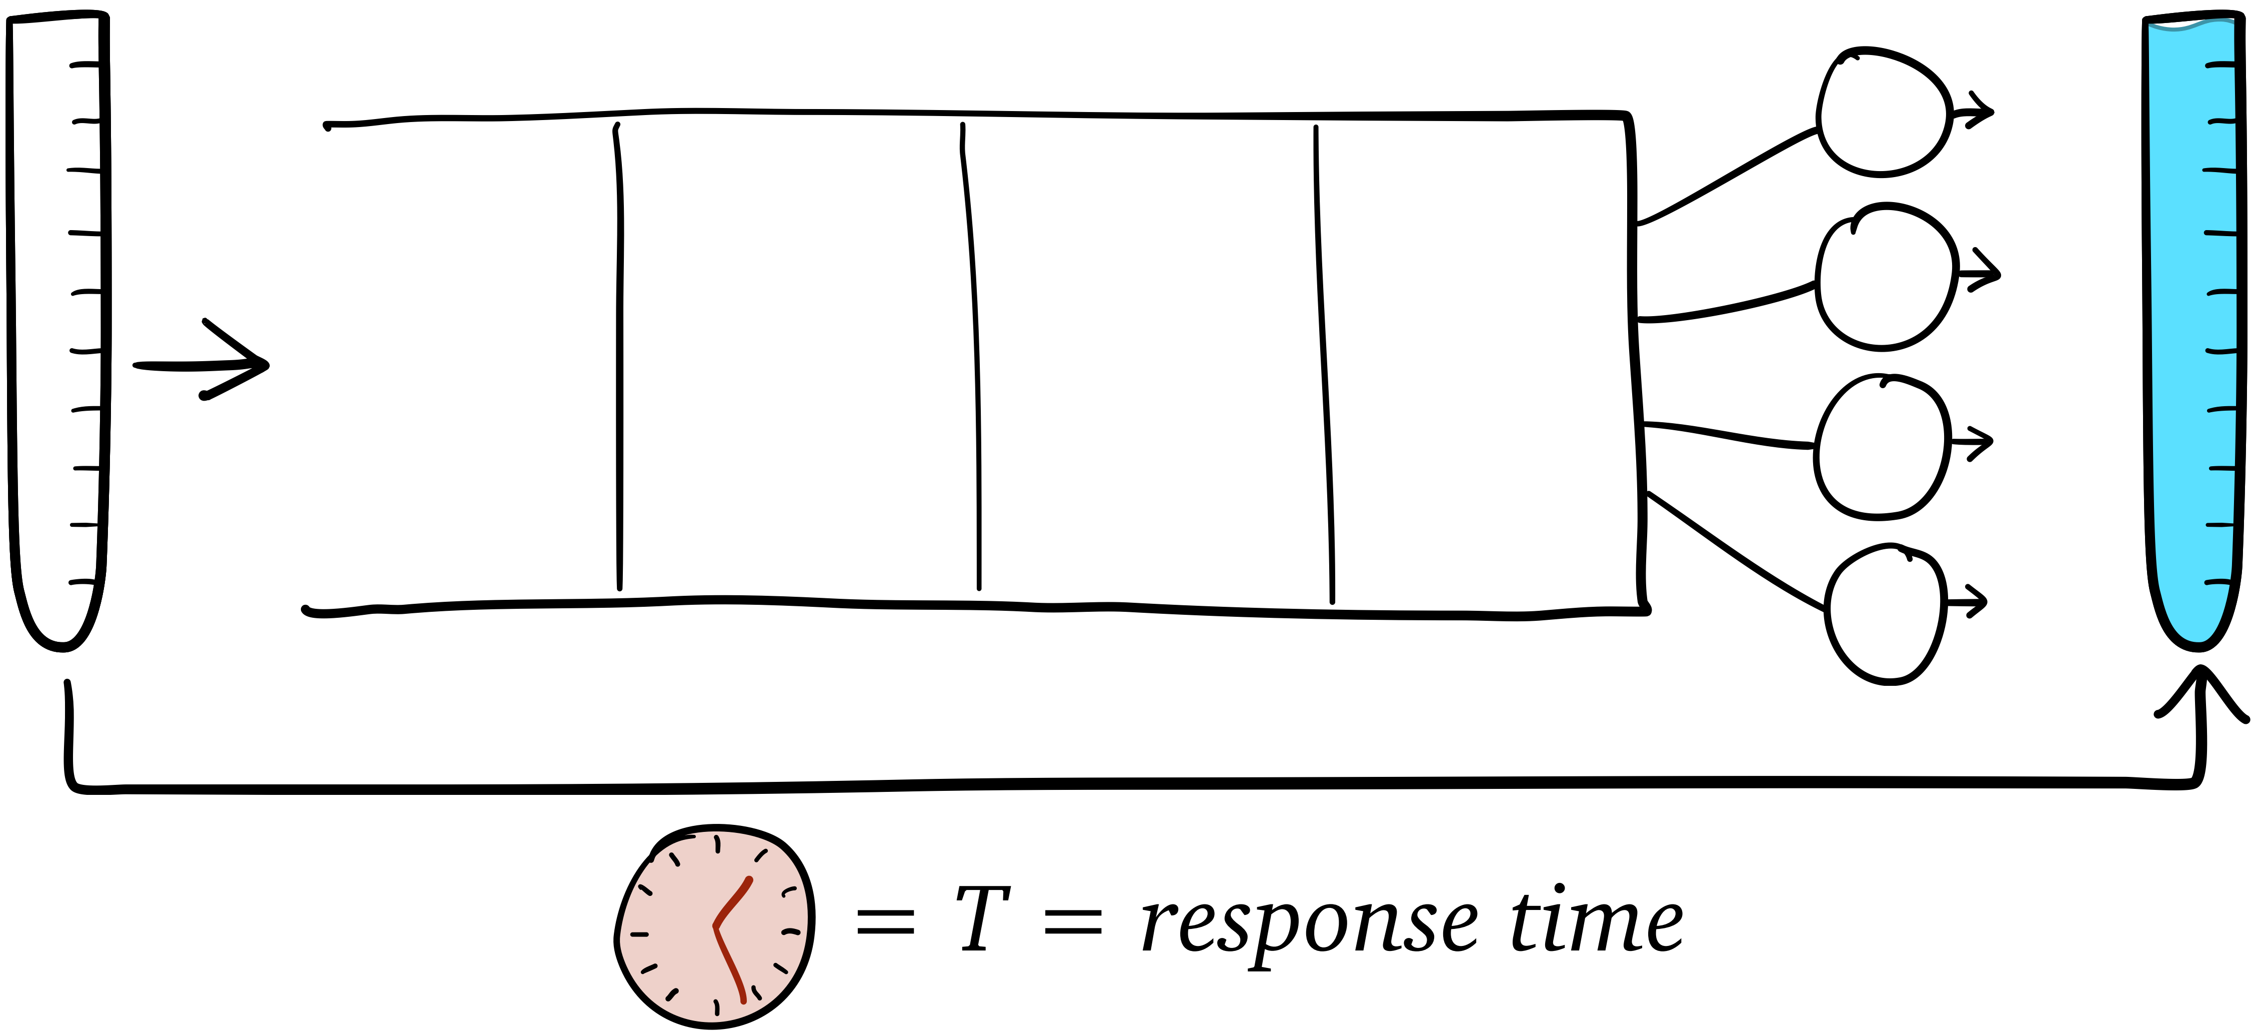

$= T =$ *response time*

Response Time

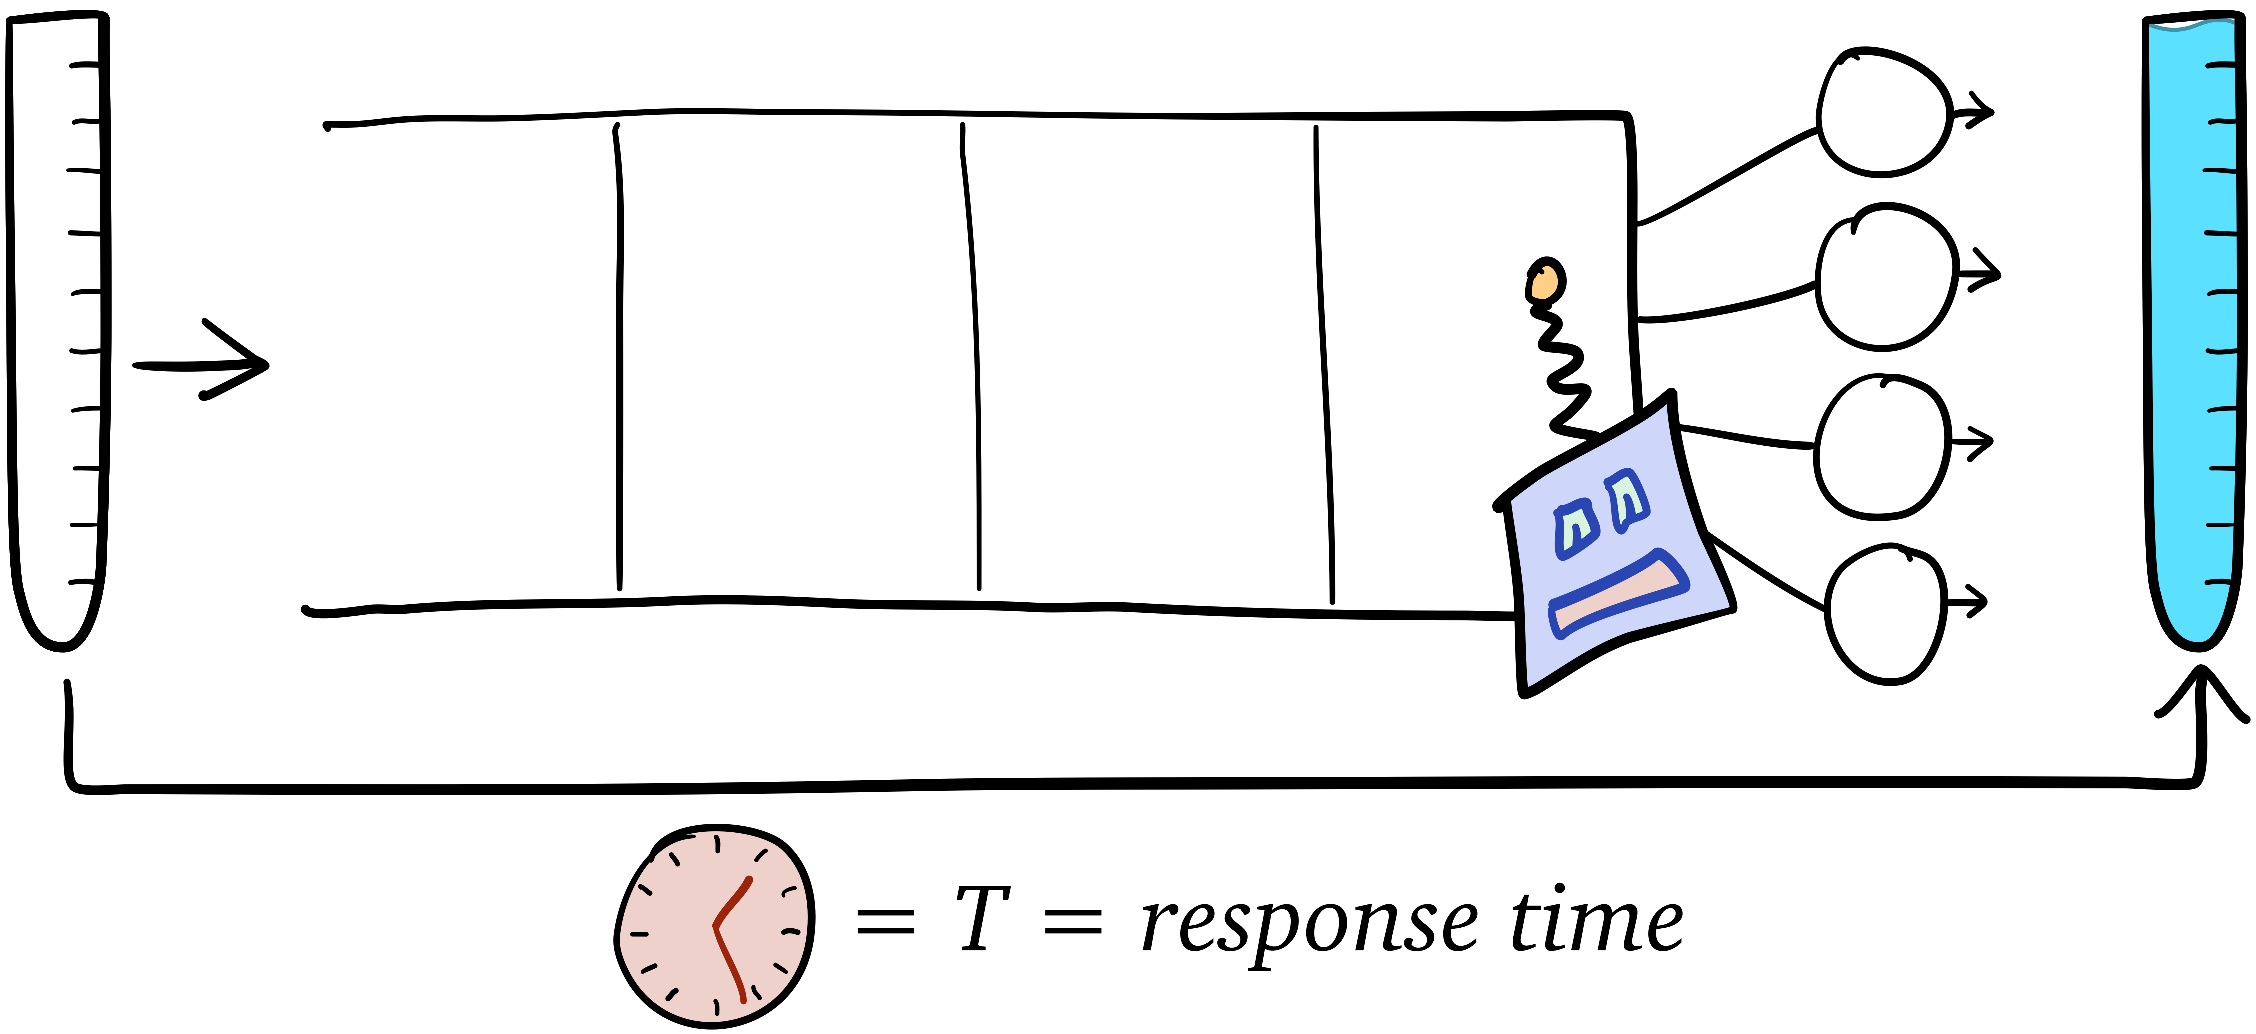

$= T =$ *response time*

Response Time

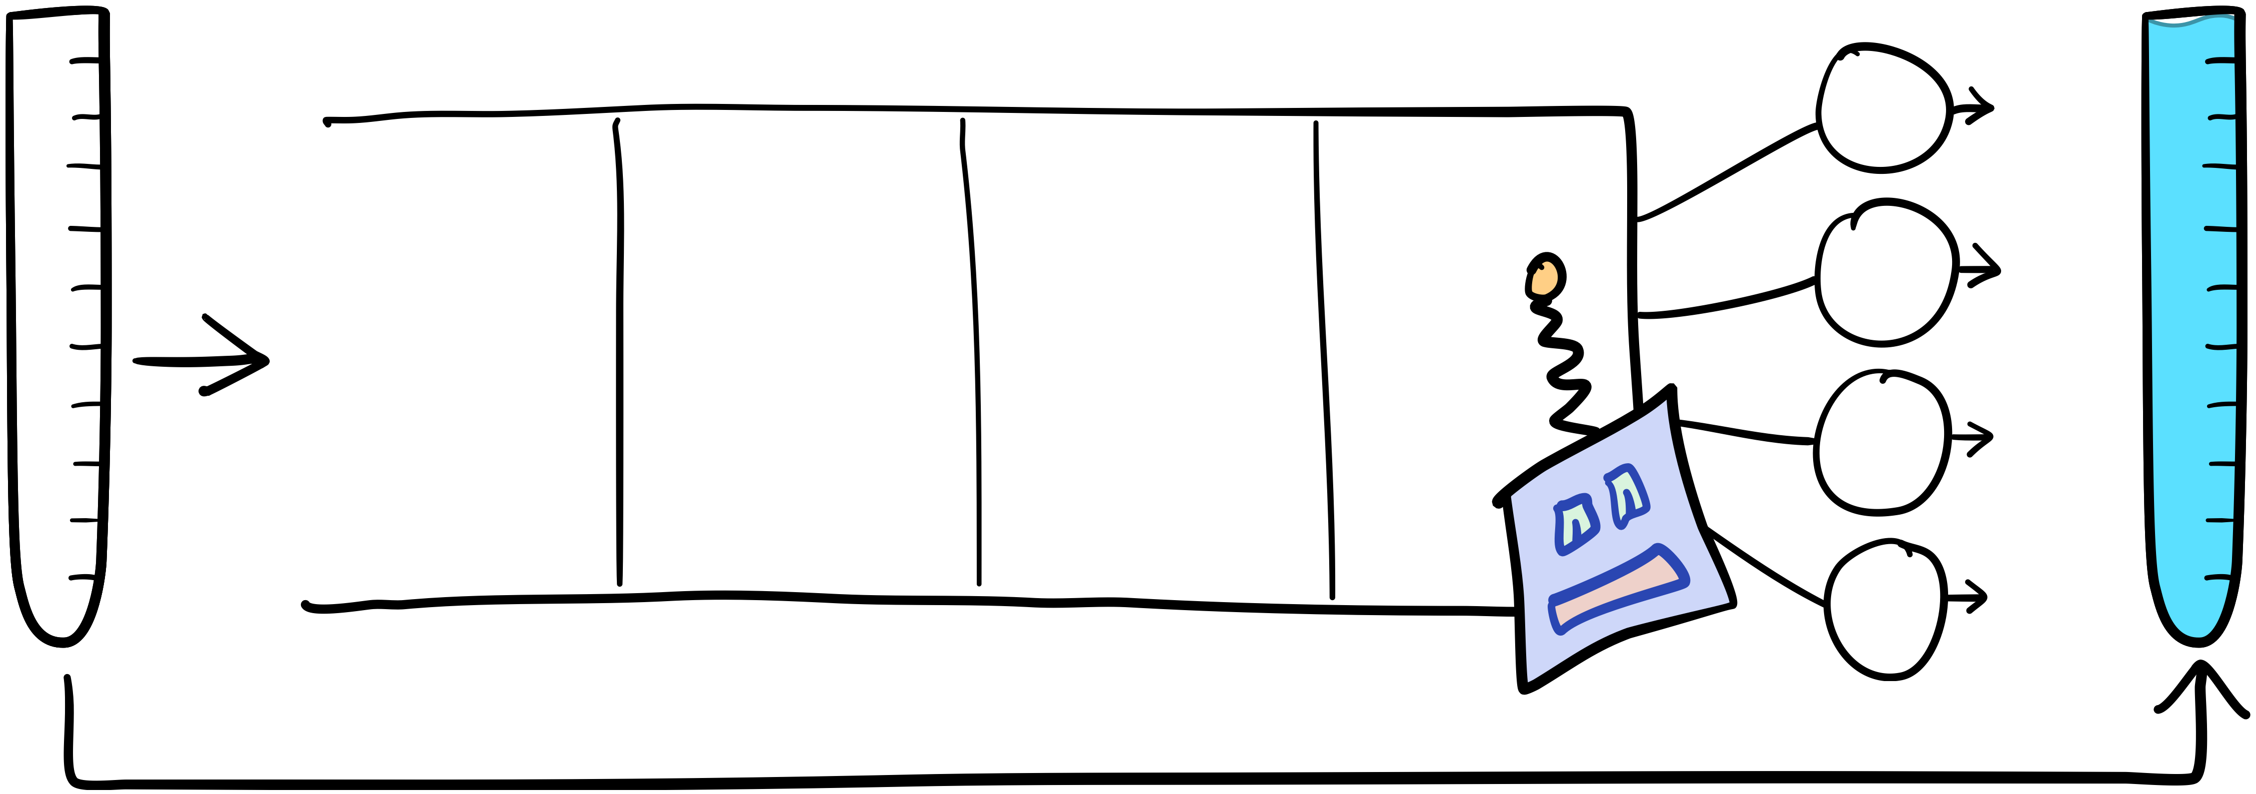

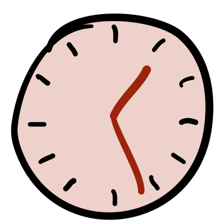 $= T =$ *response time*

Goal: schedule to minimize
mean response time $\mathbf{E}[T]$

Minimizing $\mathbf{E}[T]$

M/G/1	
M/G/k	

Minimizing $\mathbf{E}[T]$

	Known job sizes
M/G/1	
M/G/k	

Minimizing $\mathbf{E}[T]$

	Known job sizes
M/G/1	**SRPT**
M/G/k	

Minimizing $\mathbf{E}[T]$

	Known job sizes
M/G/1	**SRPT**
M/G/k	

serves job of least remaining size

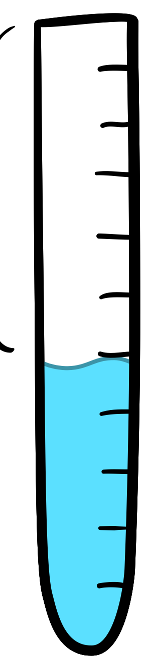

Minimizing $\mathbf{E}[T]$

	Known job sizes
M/G/1	**SRPT**
M/G/k	**SRPT-*k***

serves job of least remaining size

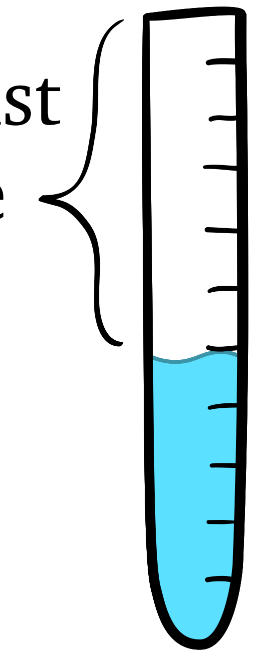

Minimizing $\mathbf{E}[T]$

	Known job sizes
M/G/1	**SRPT**
M/G/k	**SRPT-k**

in heavy traffic, X “finite variance”

serves job of least remaining size

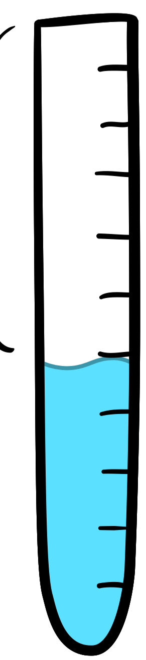

This talk:

unknown job sizes

Unknown Job Sizes

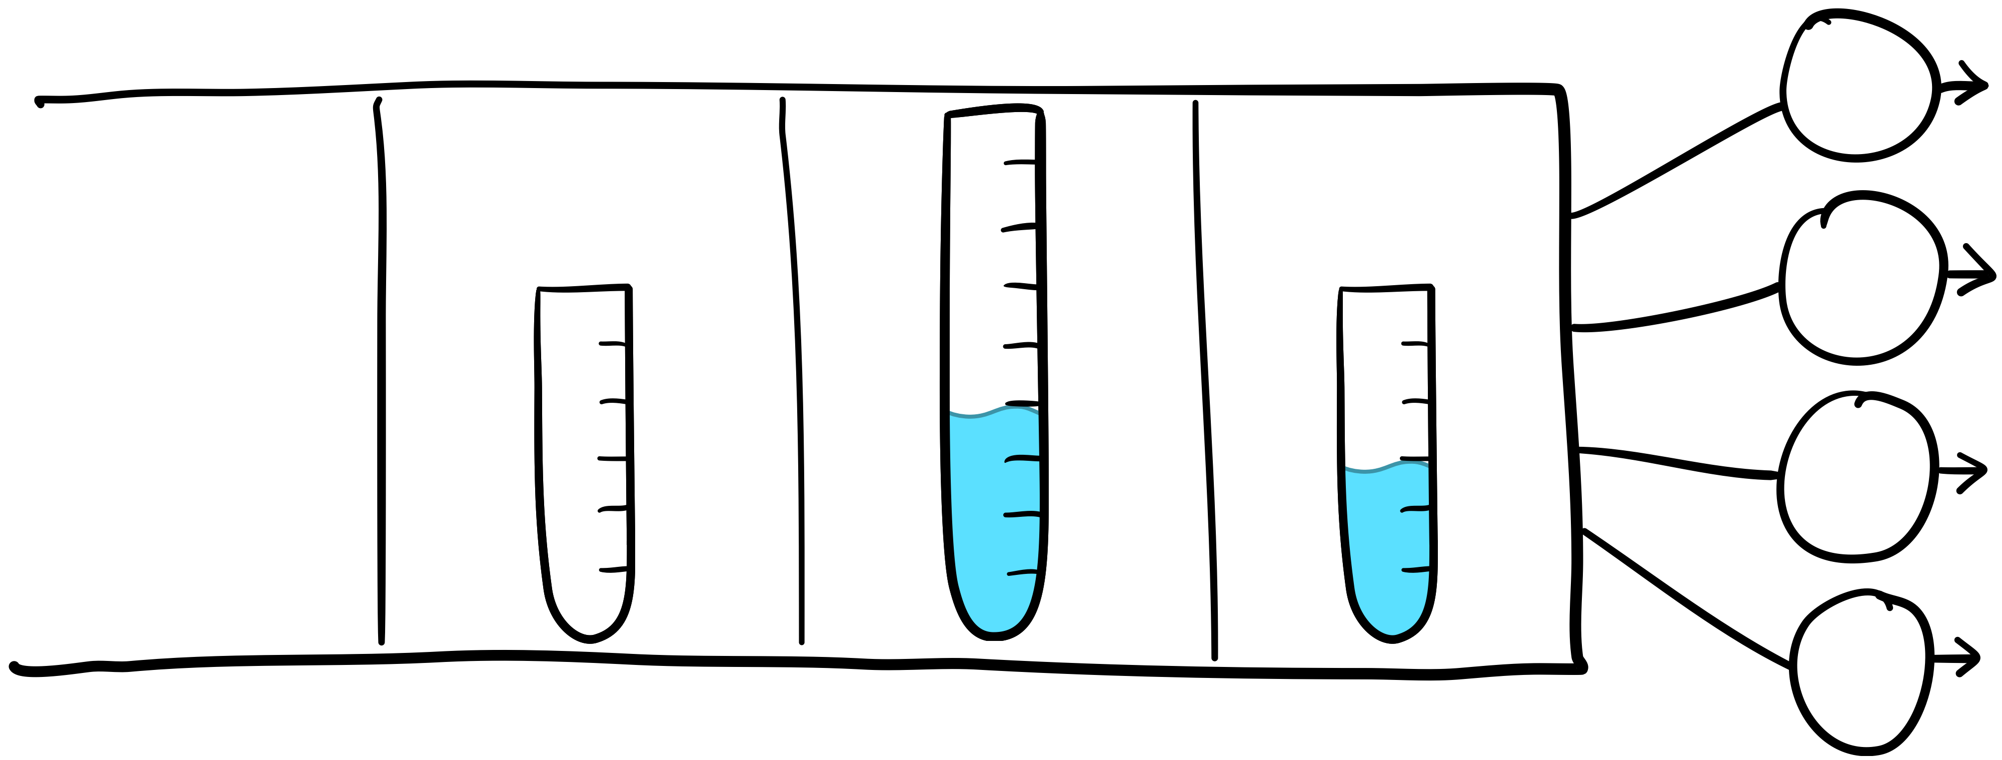

Unknown Job Sizes

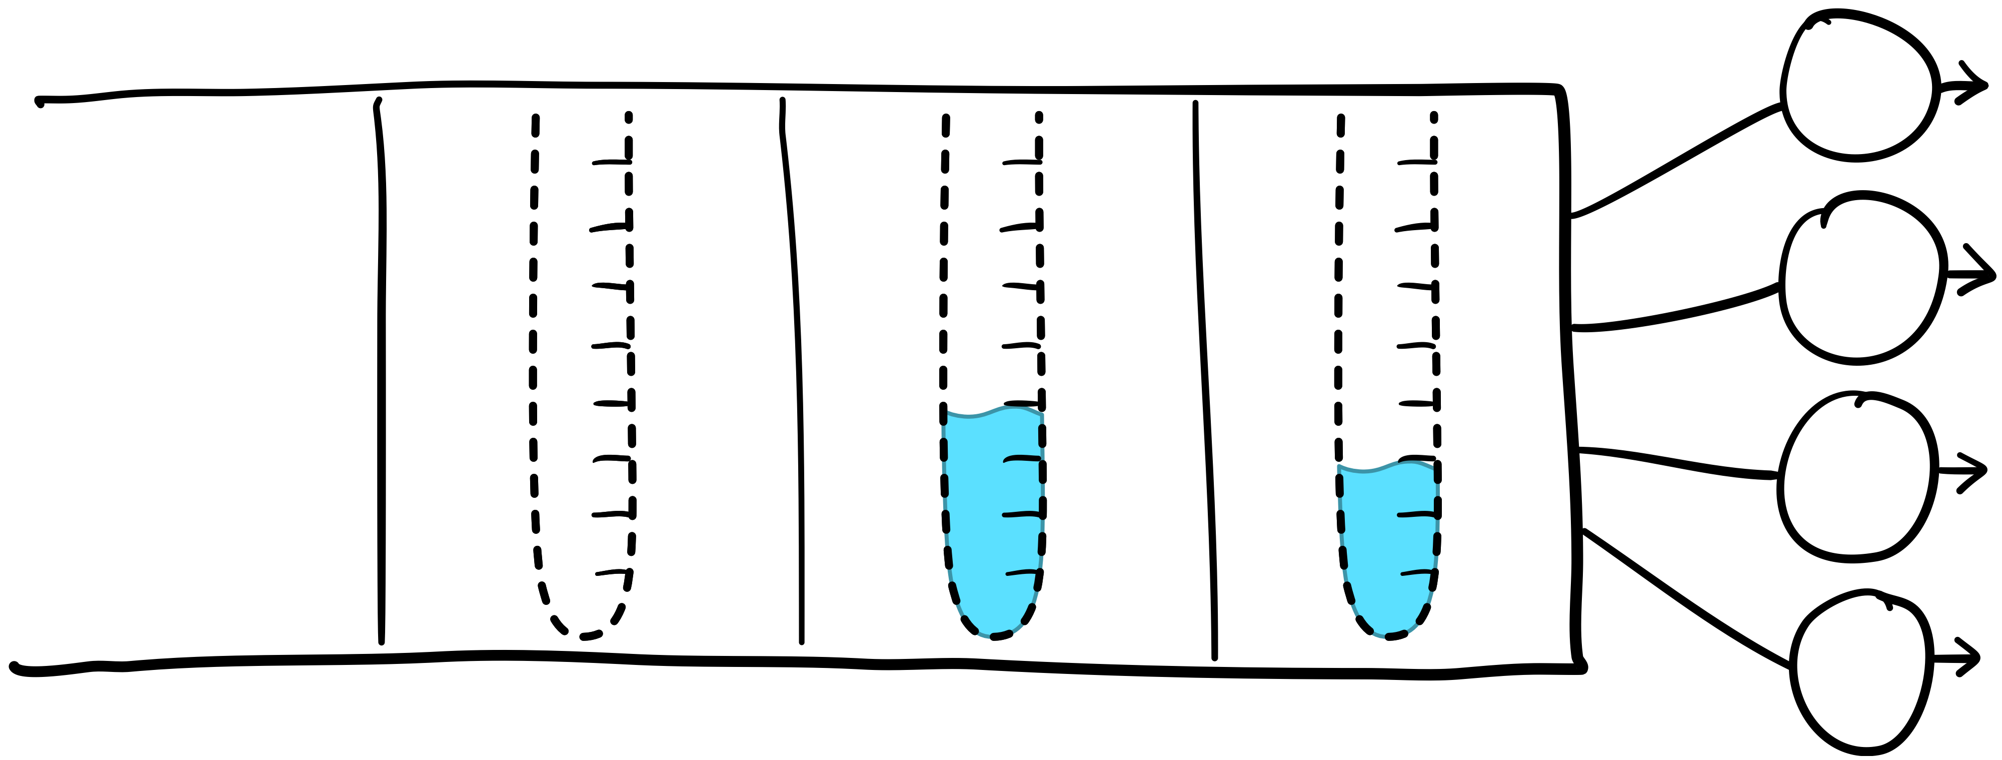

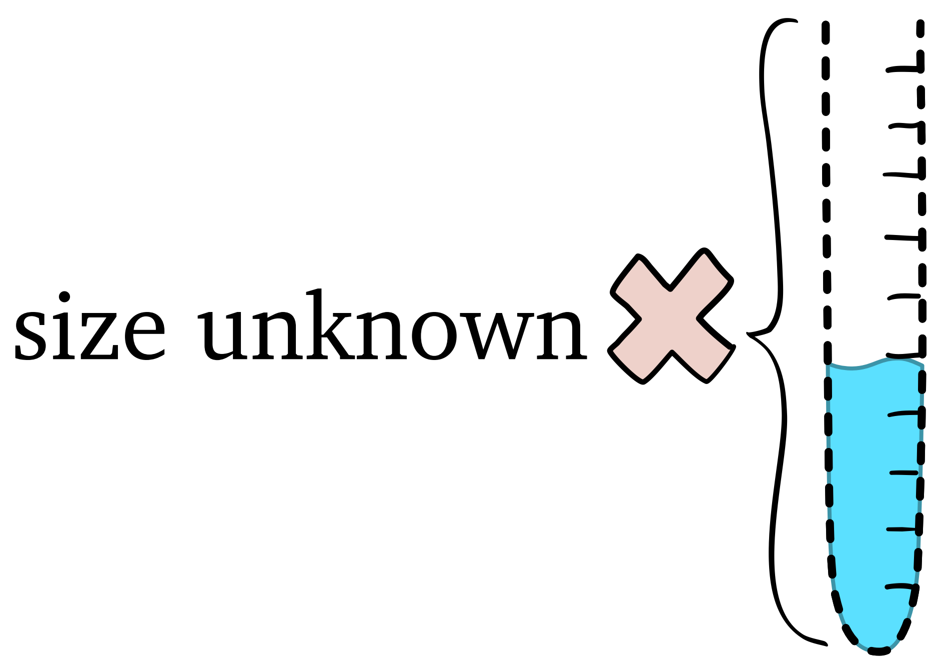

Unknown Job Sizes

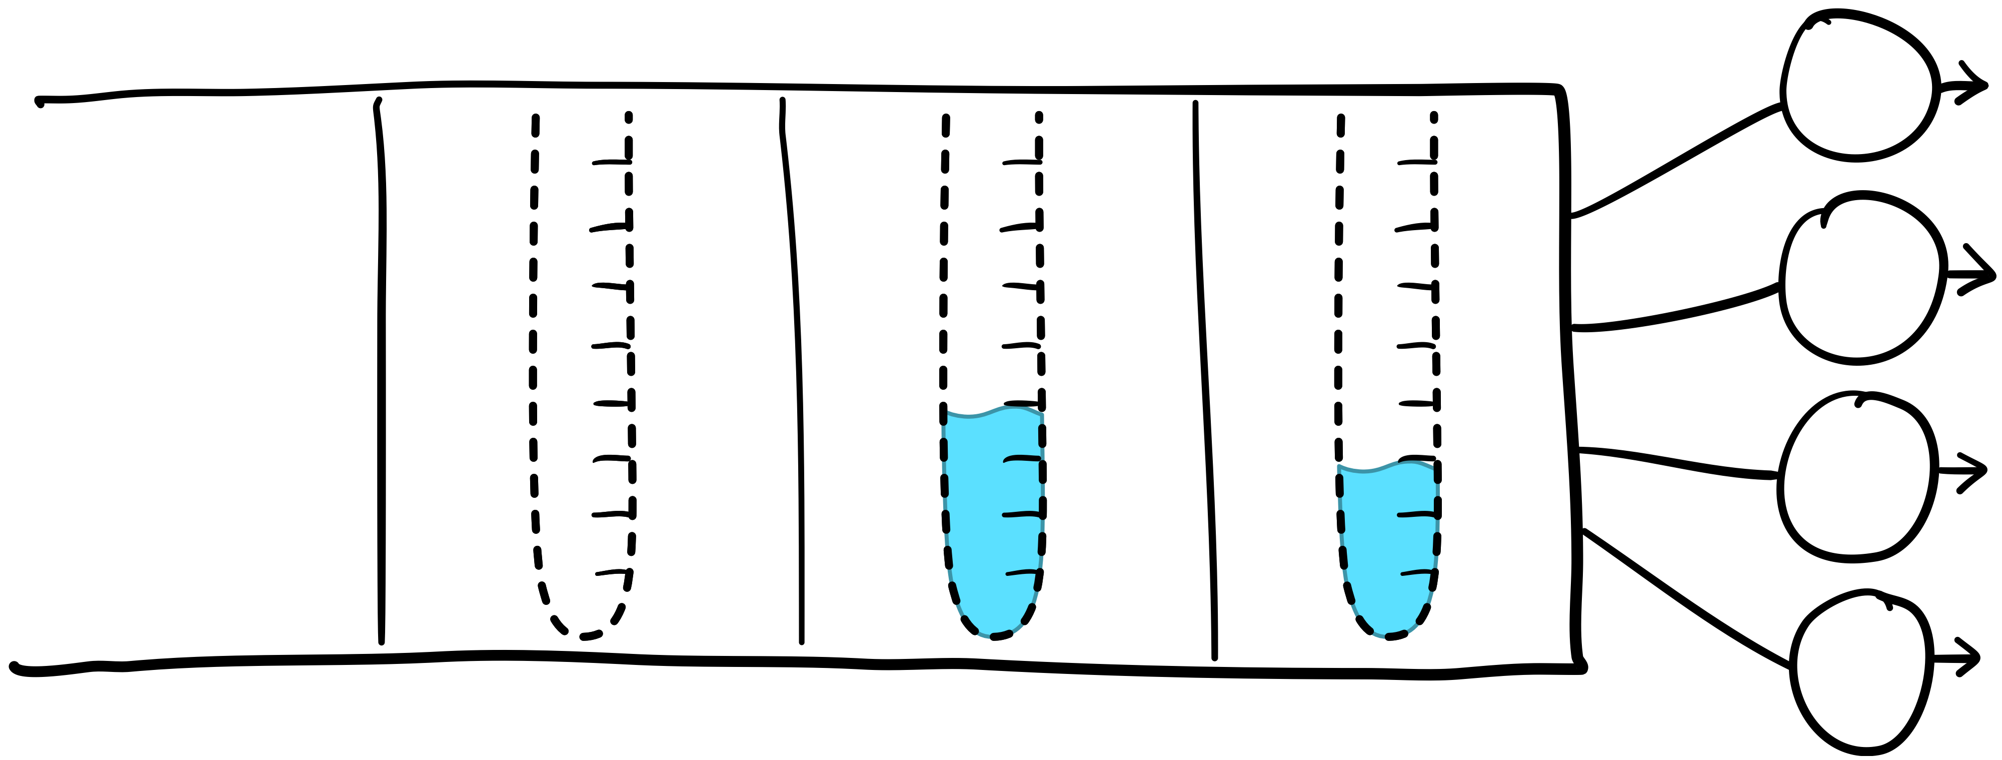

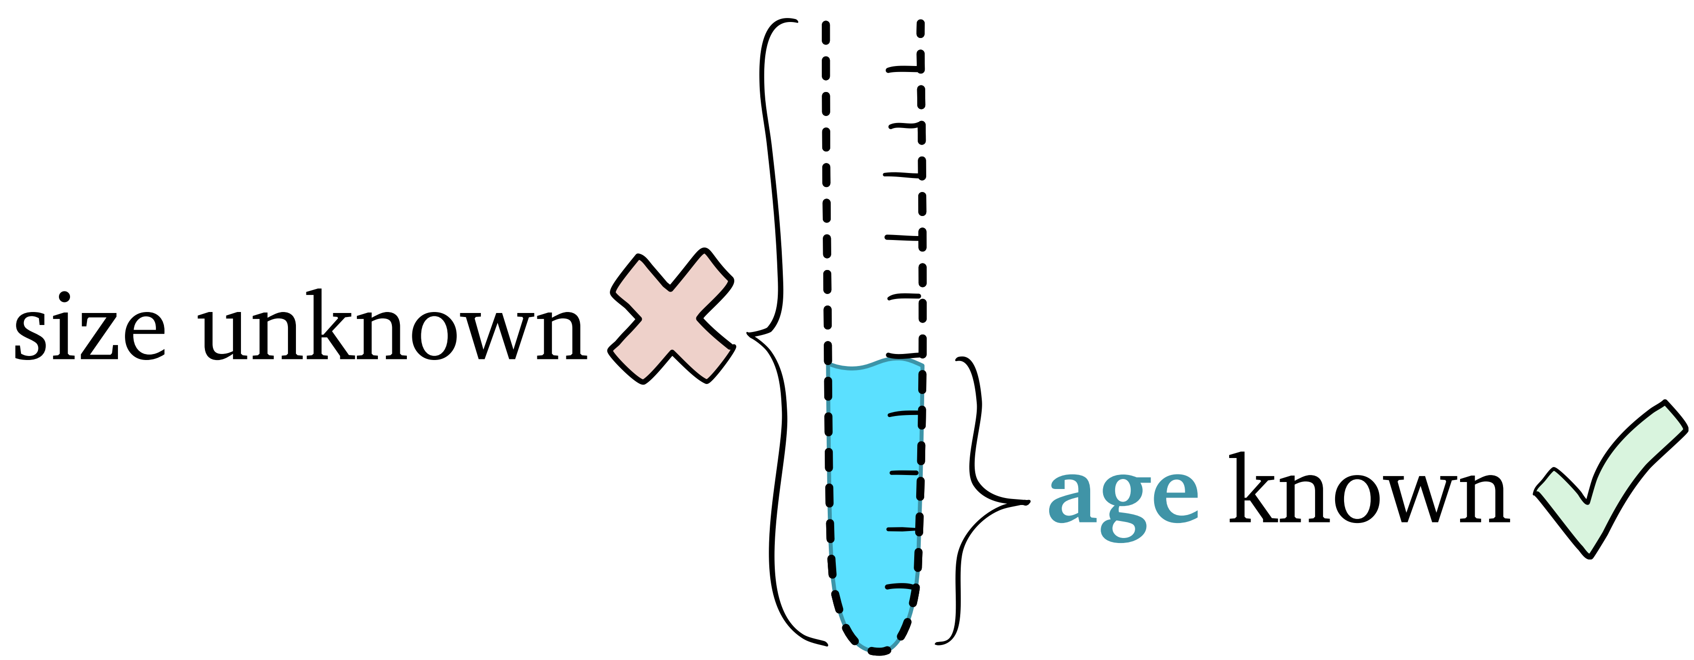

Unknown Job Sizes

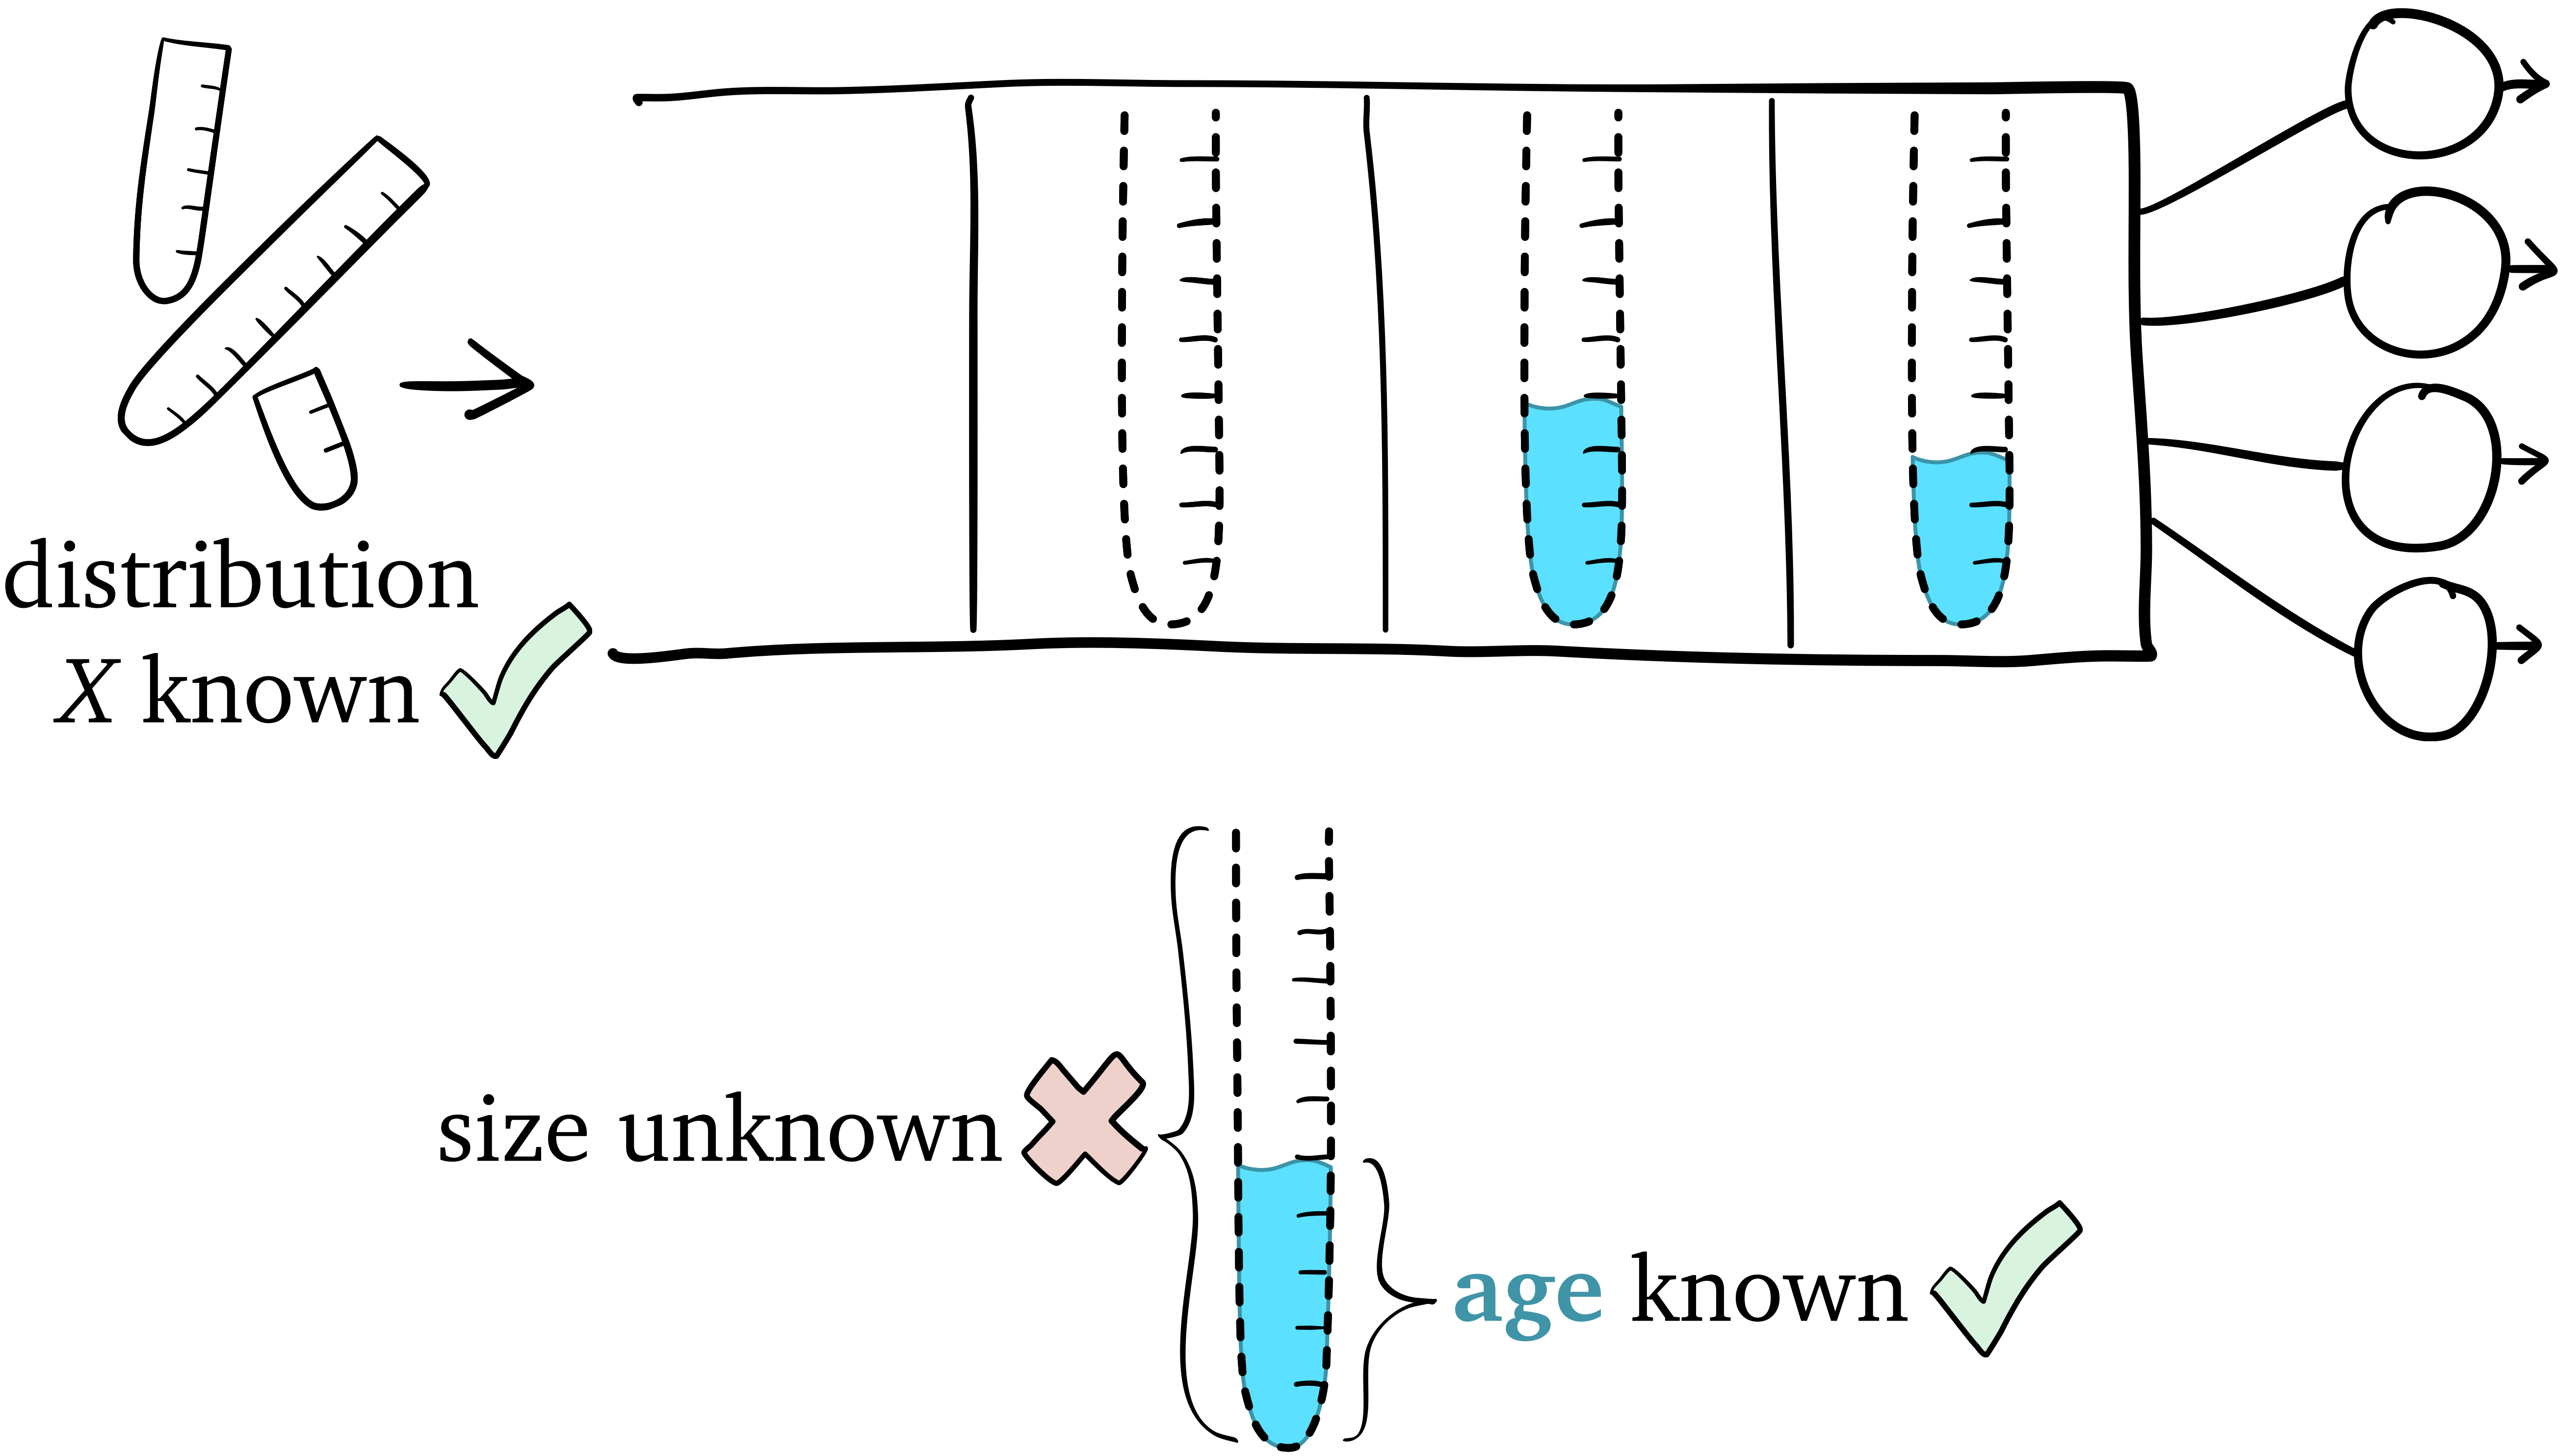

Optimal for M/G/1: **Gittins**

Optimal for M/G/1: **Gittins**

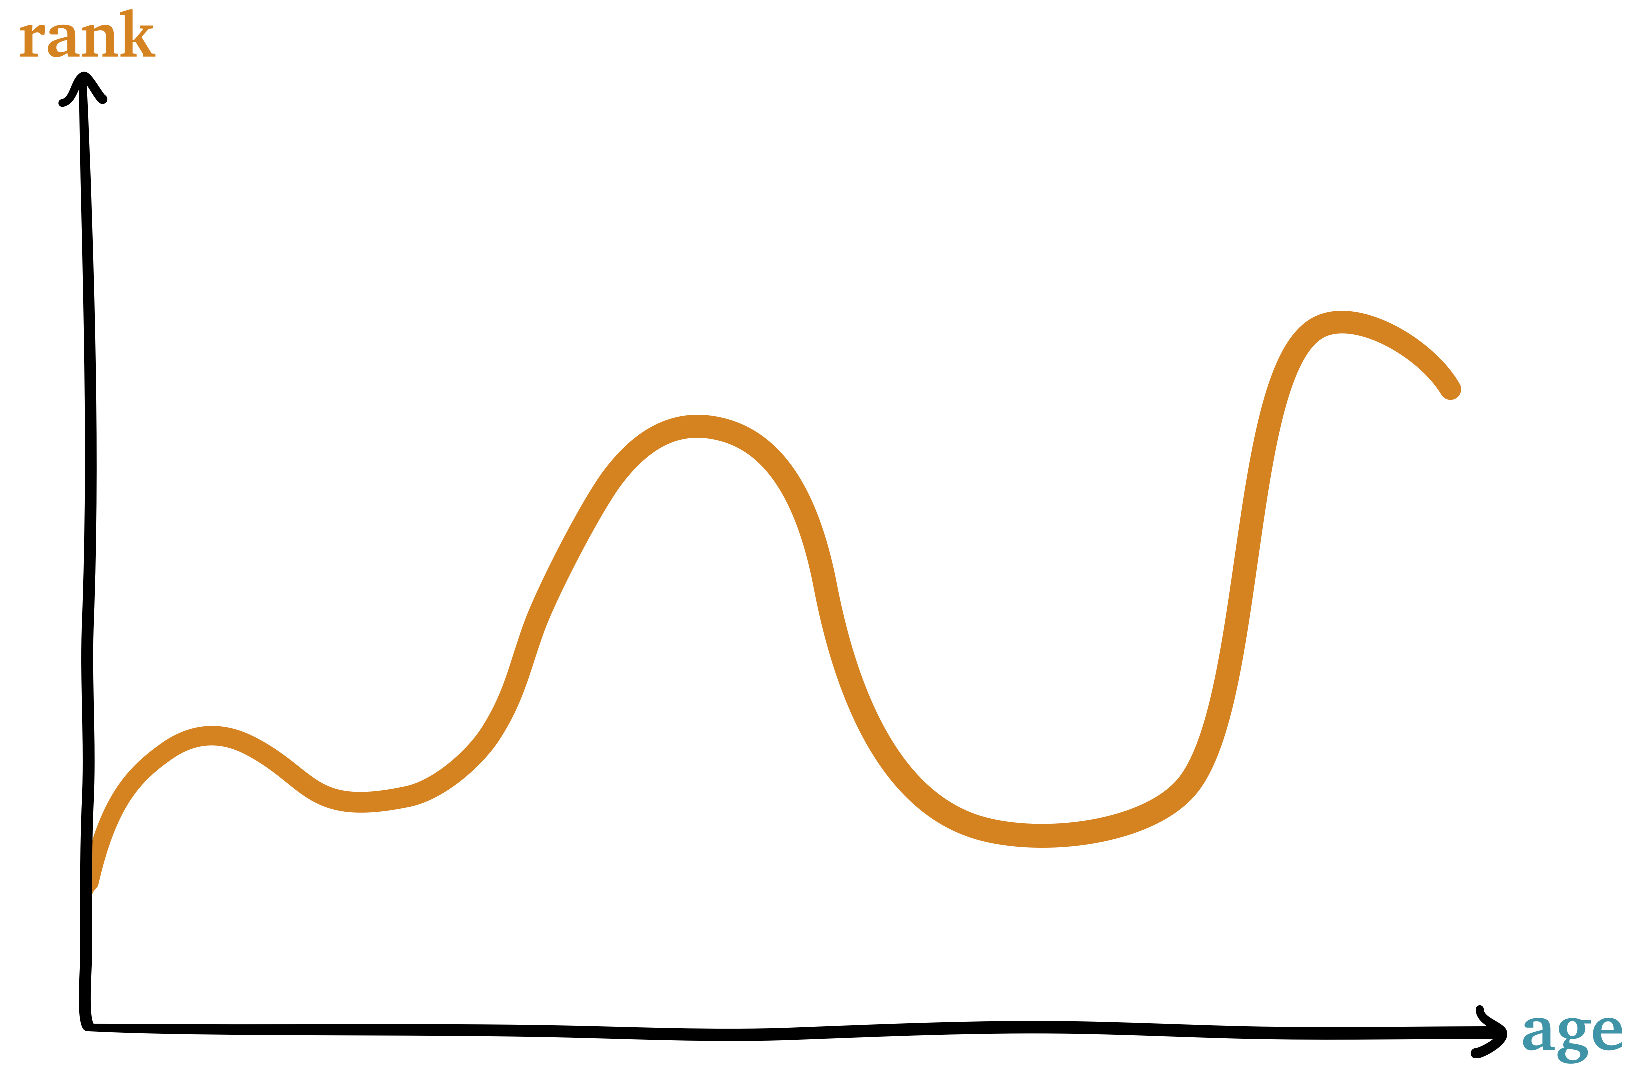

Optimal for M/G/1: **Gittins**

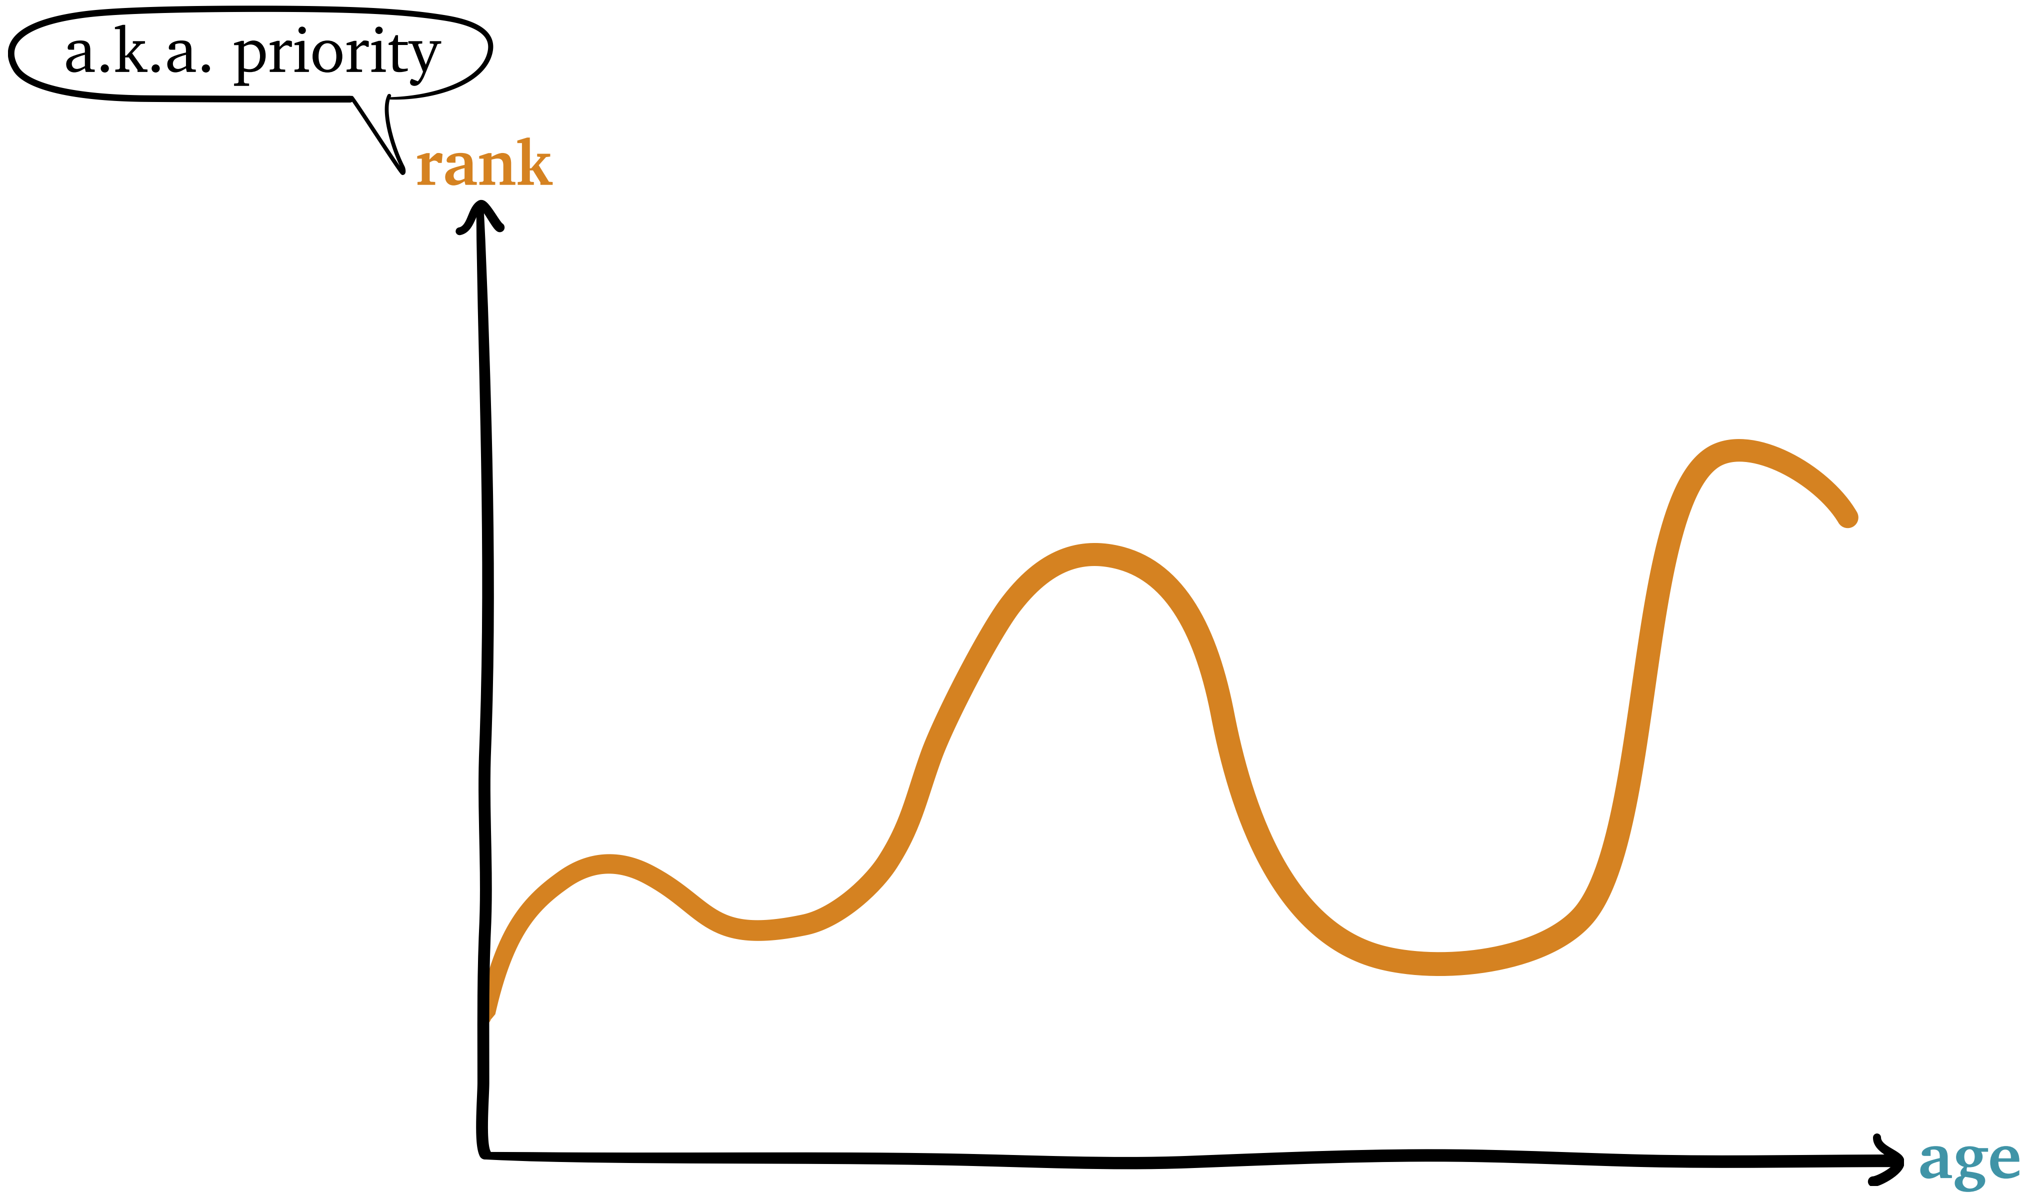

Optimal for M/G/1: **Gittins**

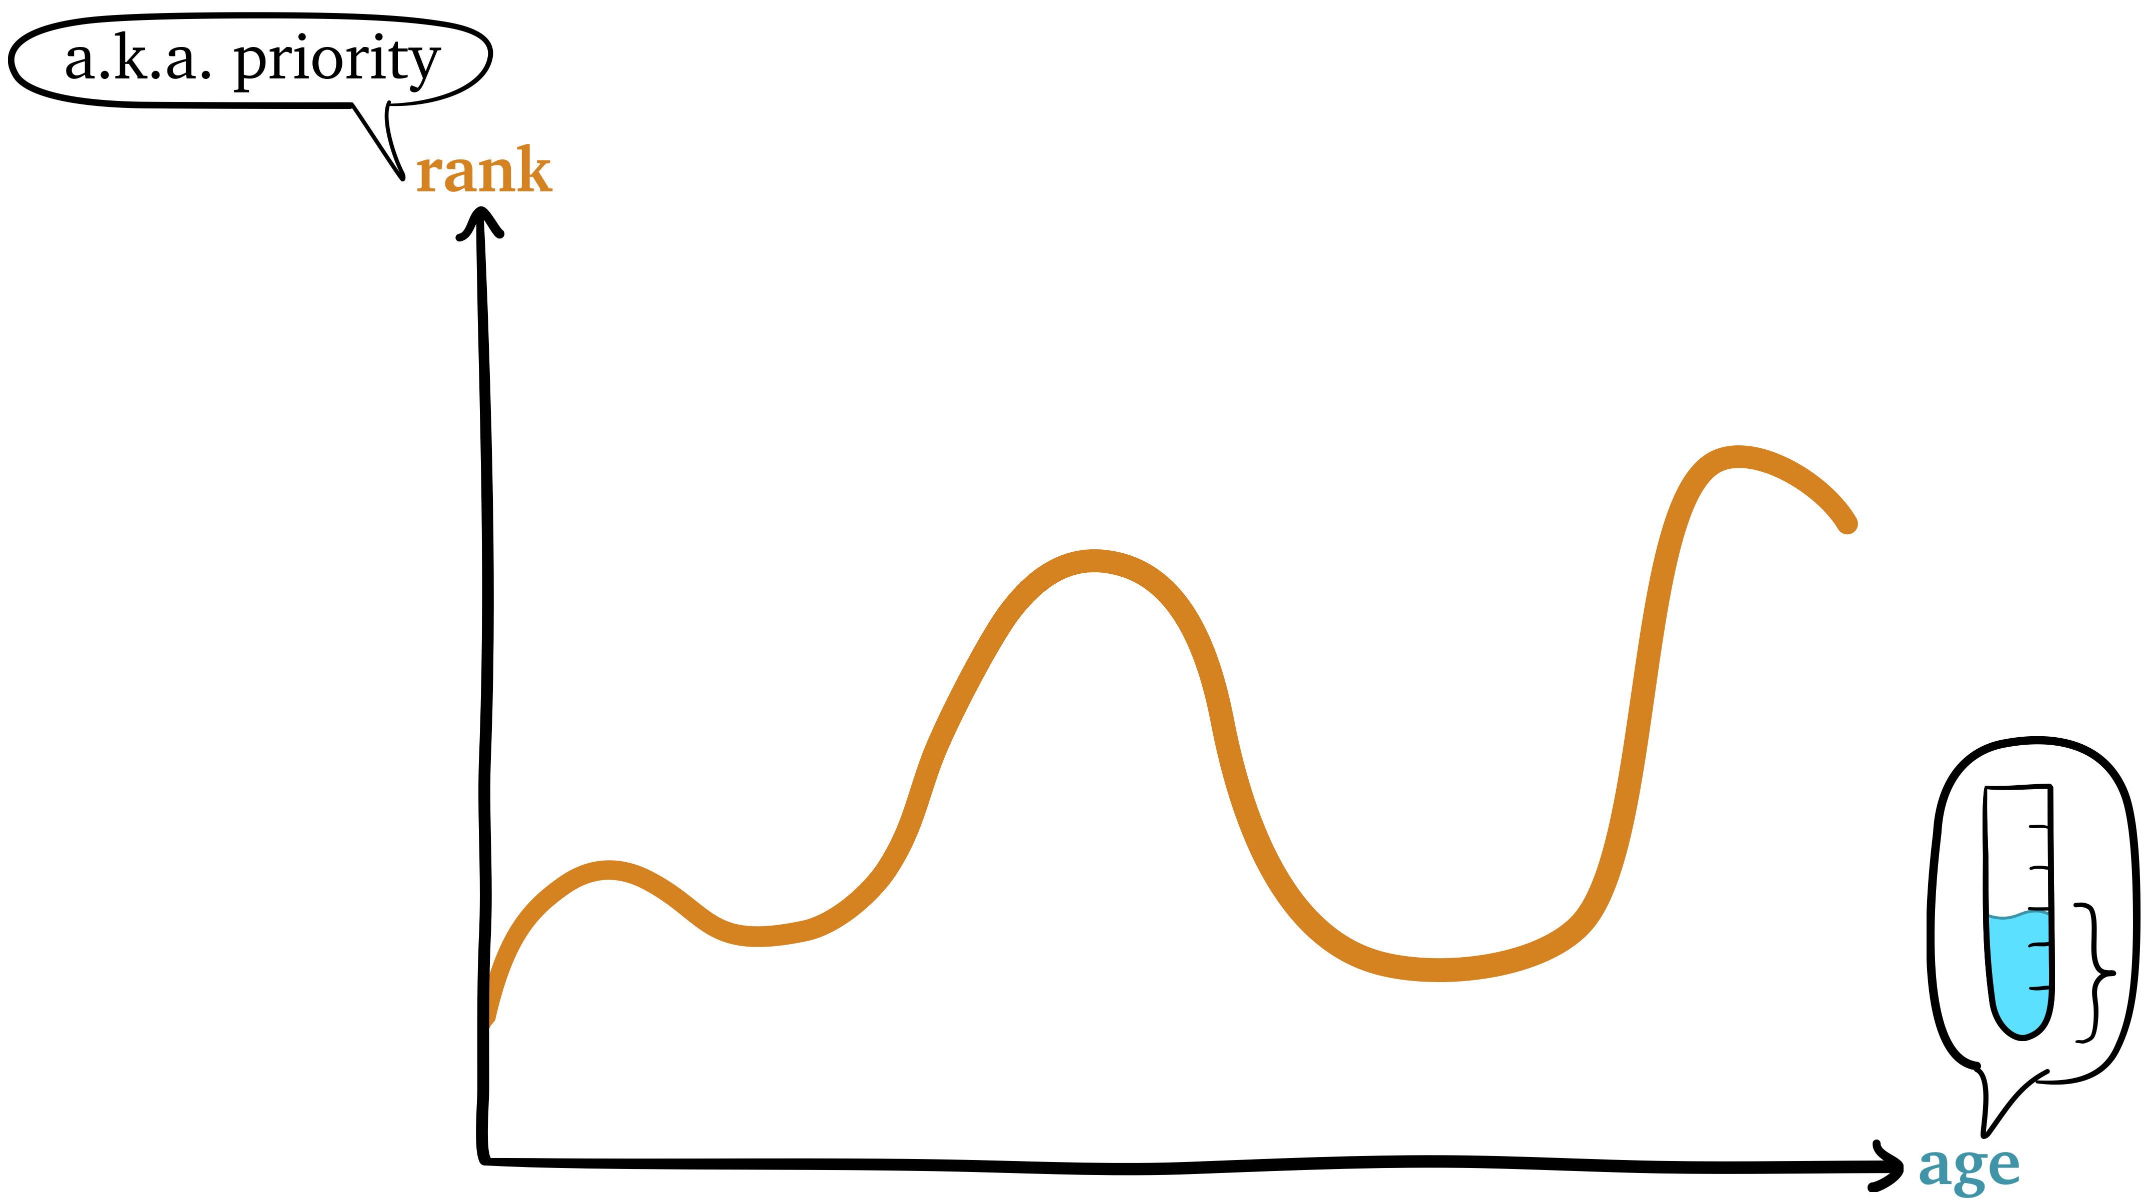

Optimal for M/G/1: **Gittins**

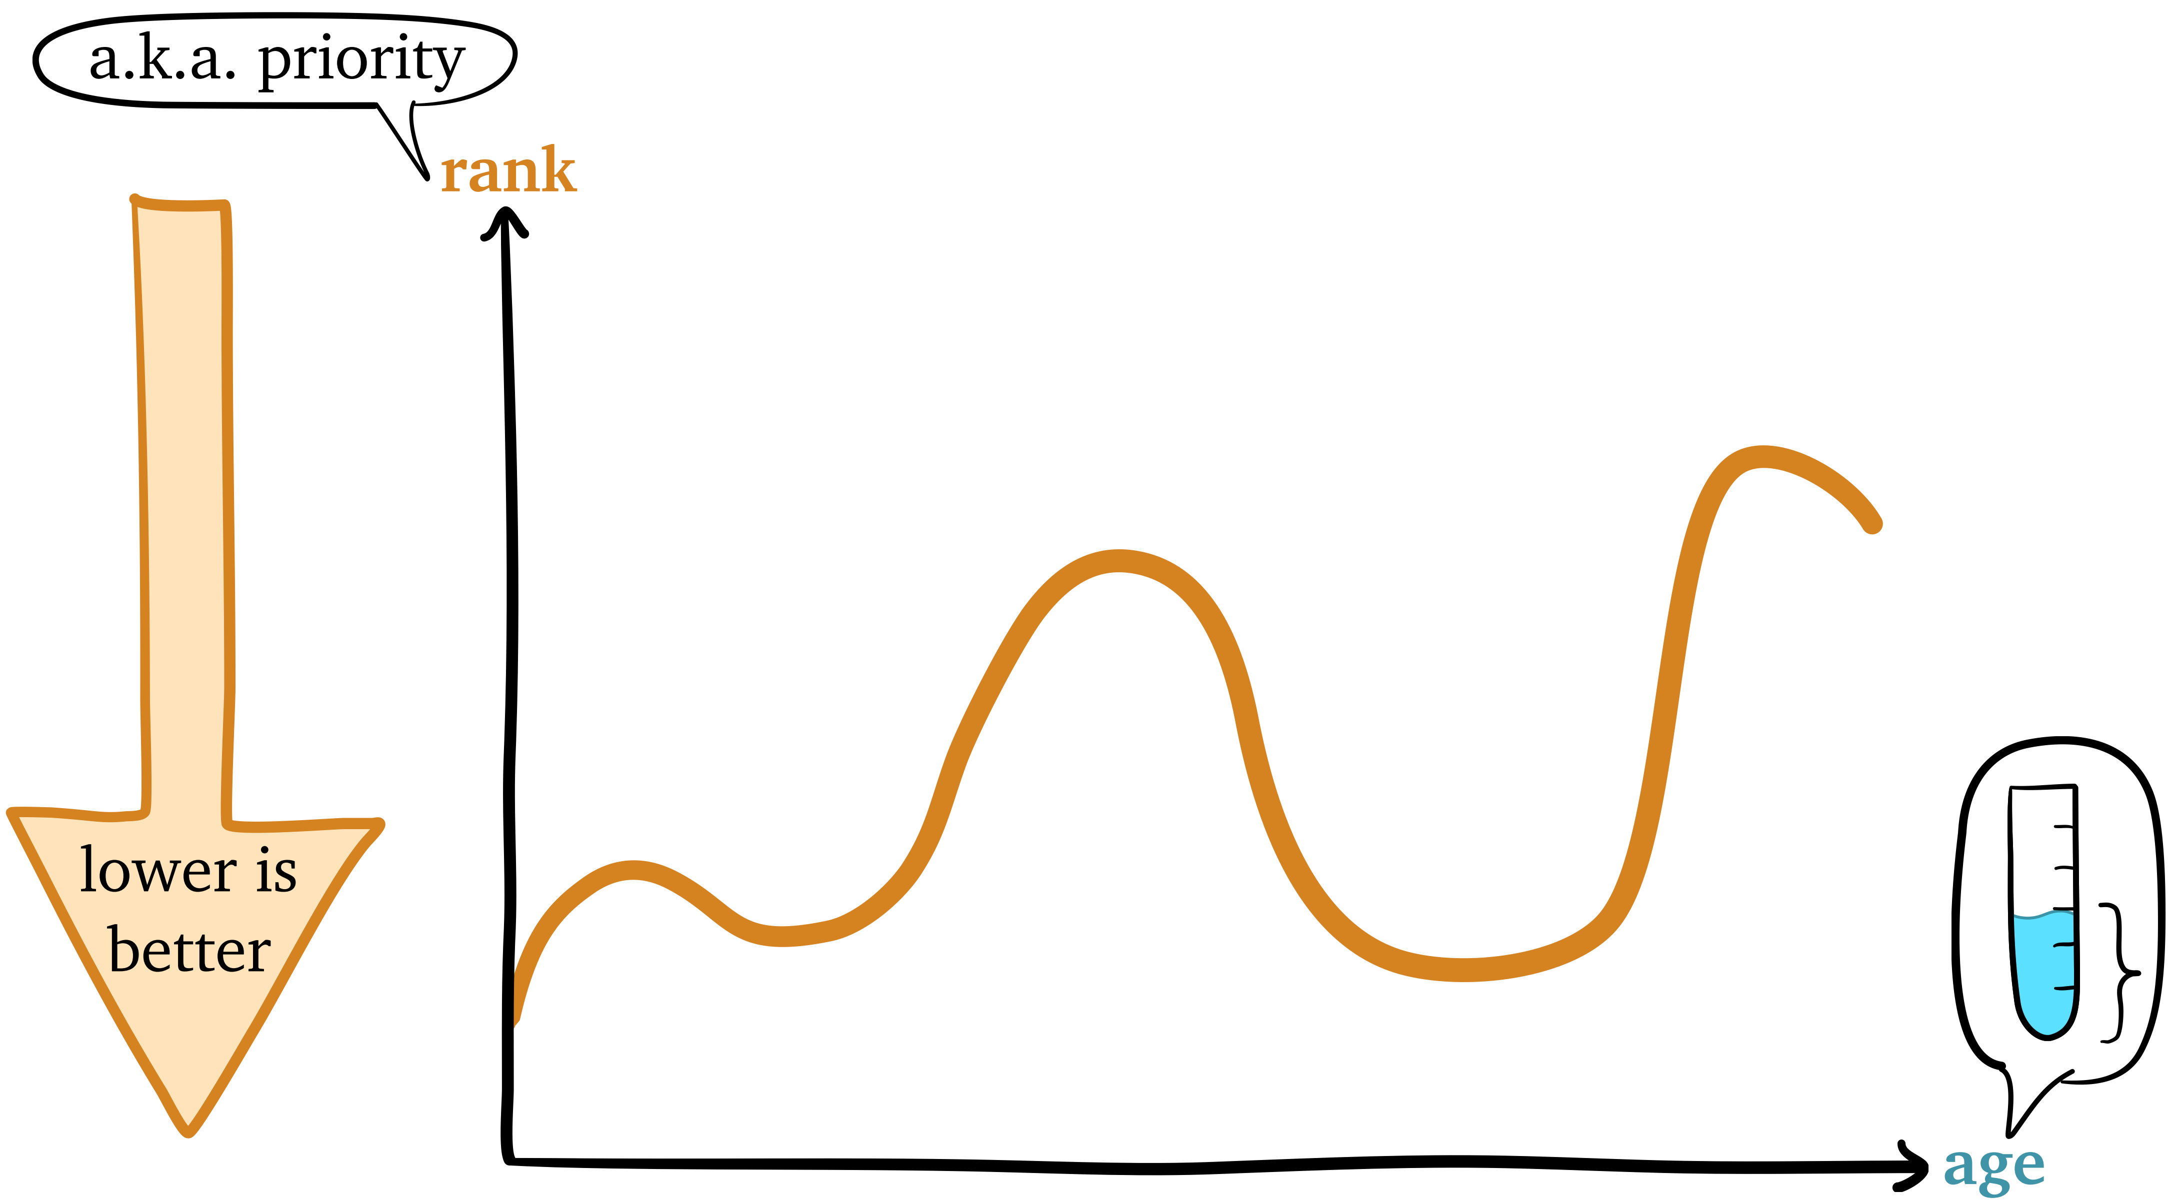

Optimal for M/G/1: Gittins

$$r_{\text{Gittins}}(a) = \inf_{b>a} \frac{\mathbf{E}[\min\{X, b\} \mid X > a]}{\mathbf{P}[X \le b \mid X > a]}$$

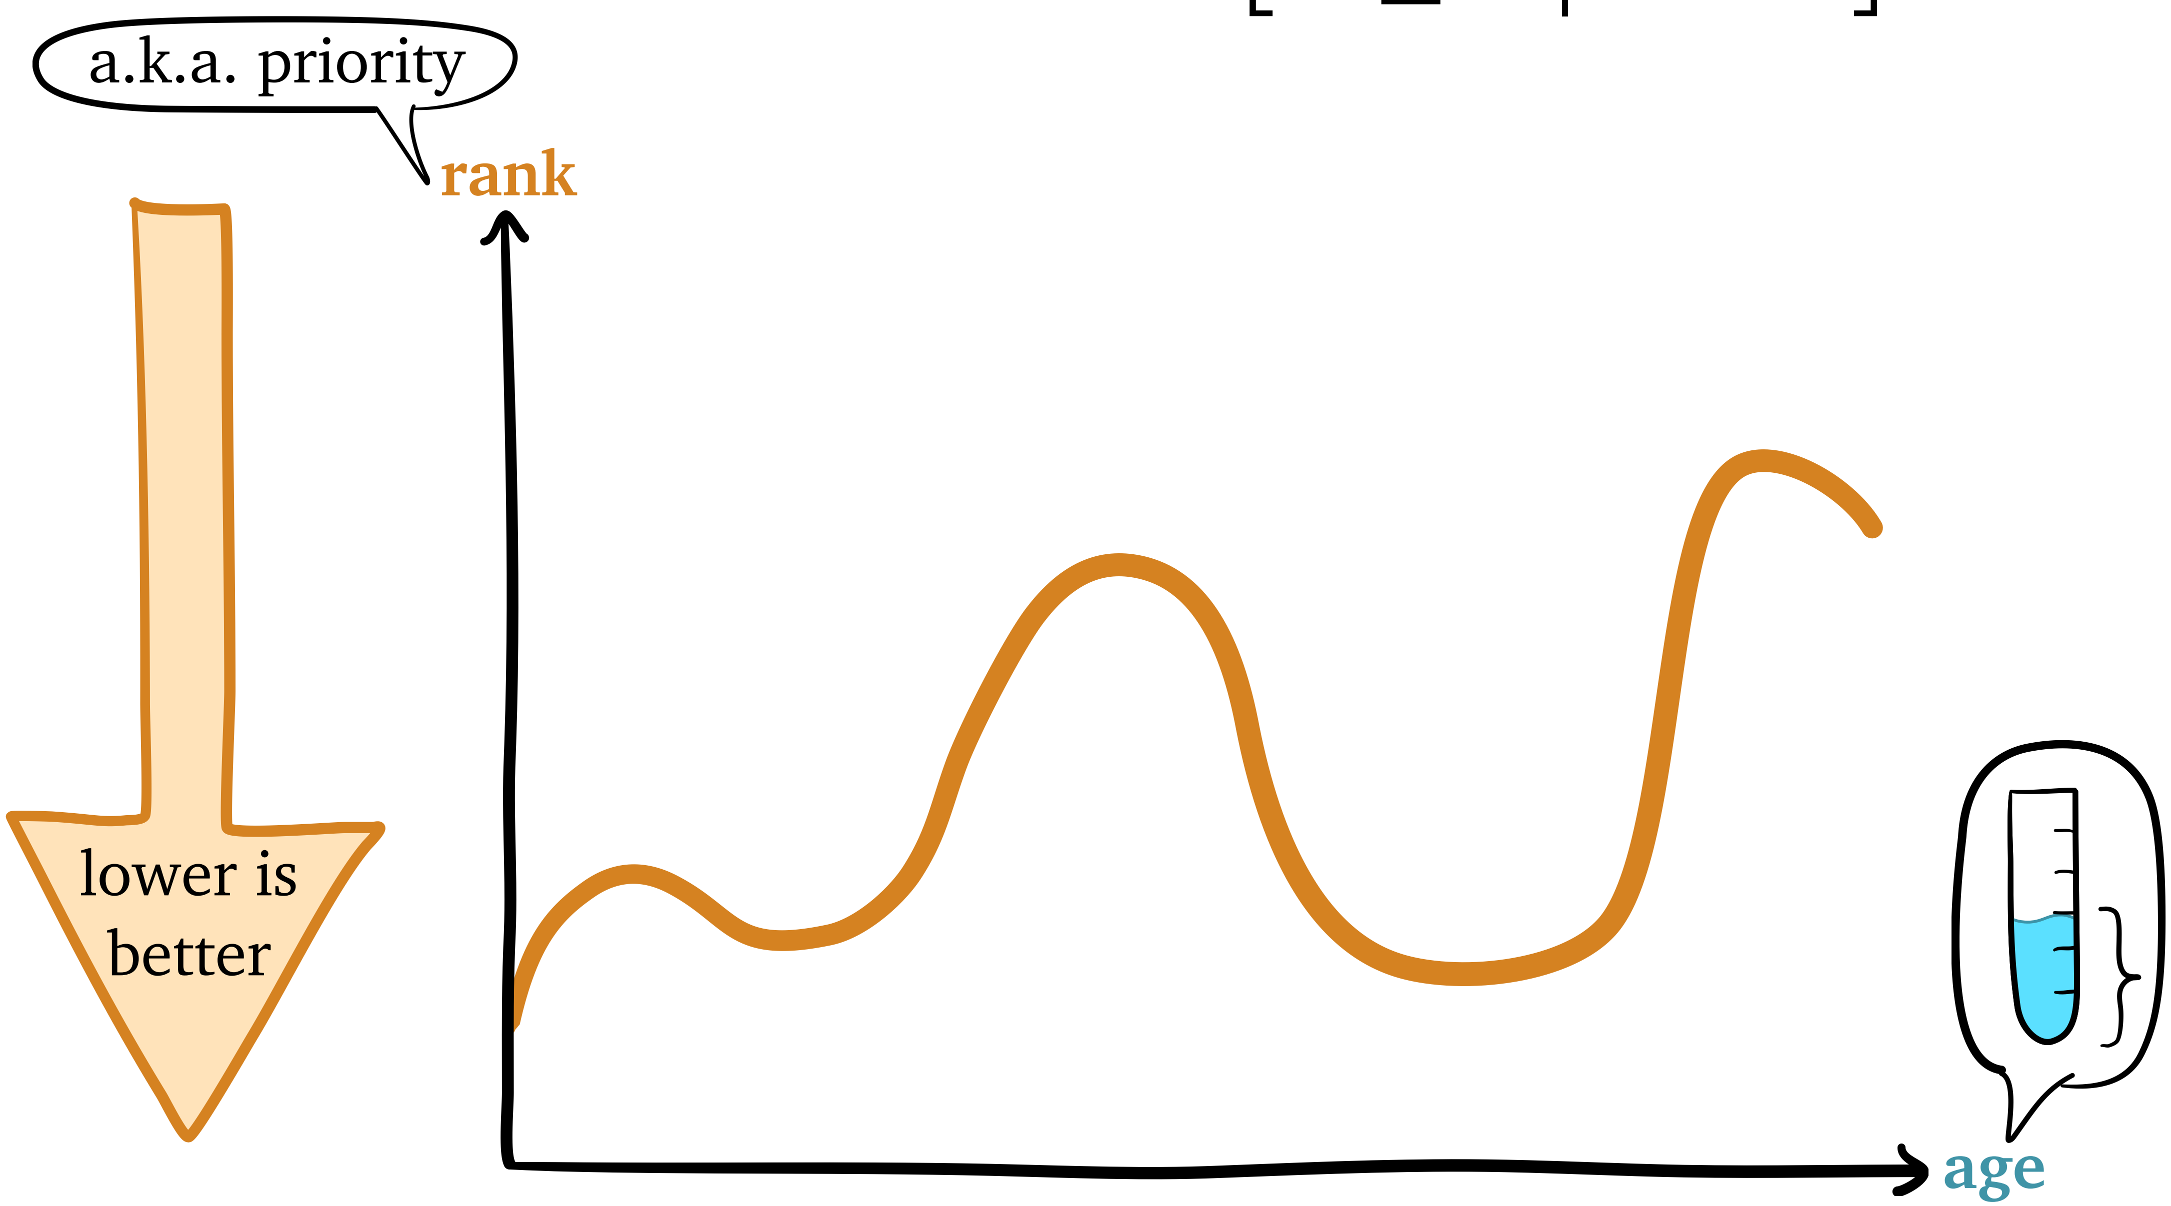

Optimal for M/G/1: Gittins

$$r_{\text{Gittins}}(a) = \inf_{b>a} \frac{\mathbf{E}[\min\{X, b\} \mid X > a]}{\mathbf{P}[X \leq b \mid X > a]}$$

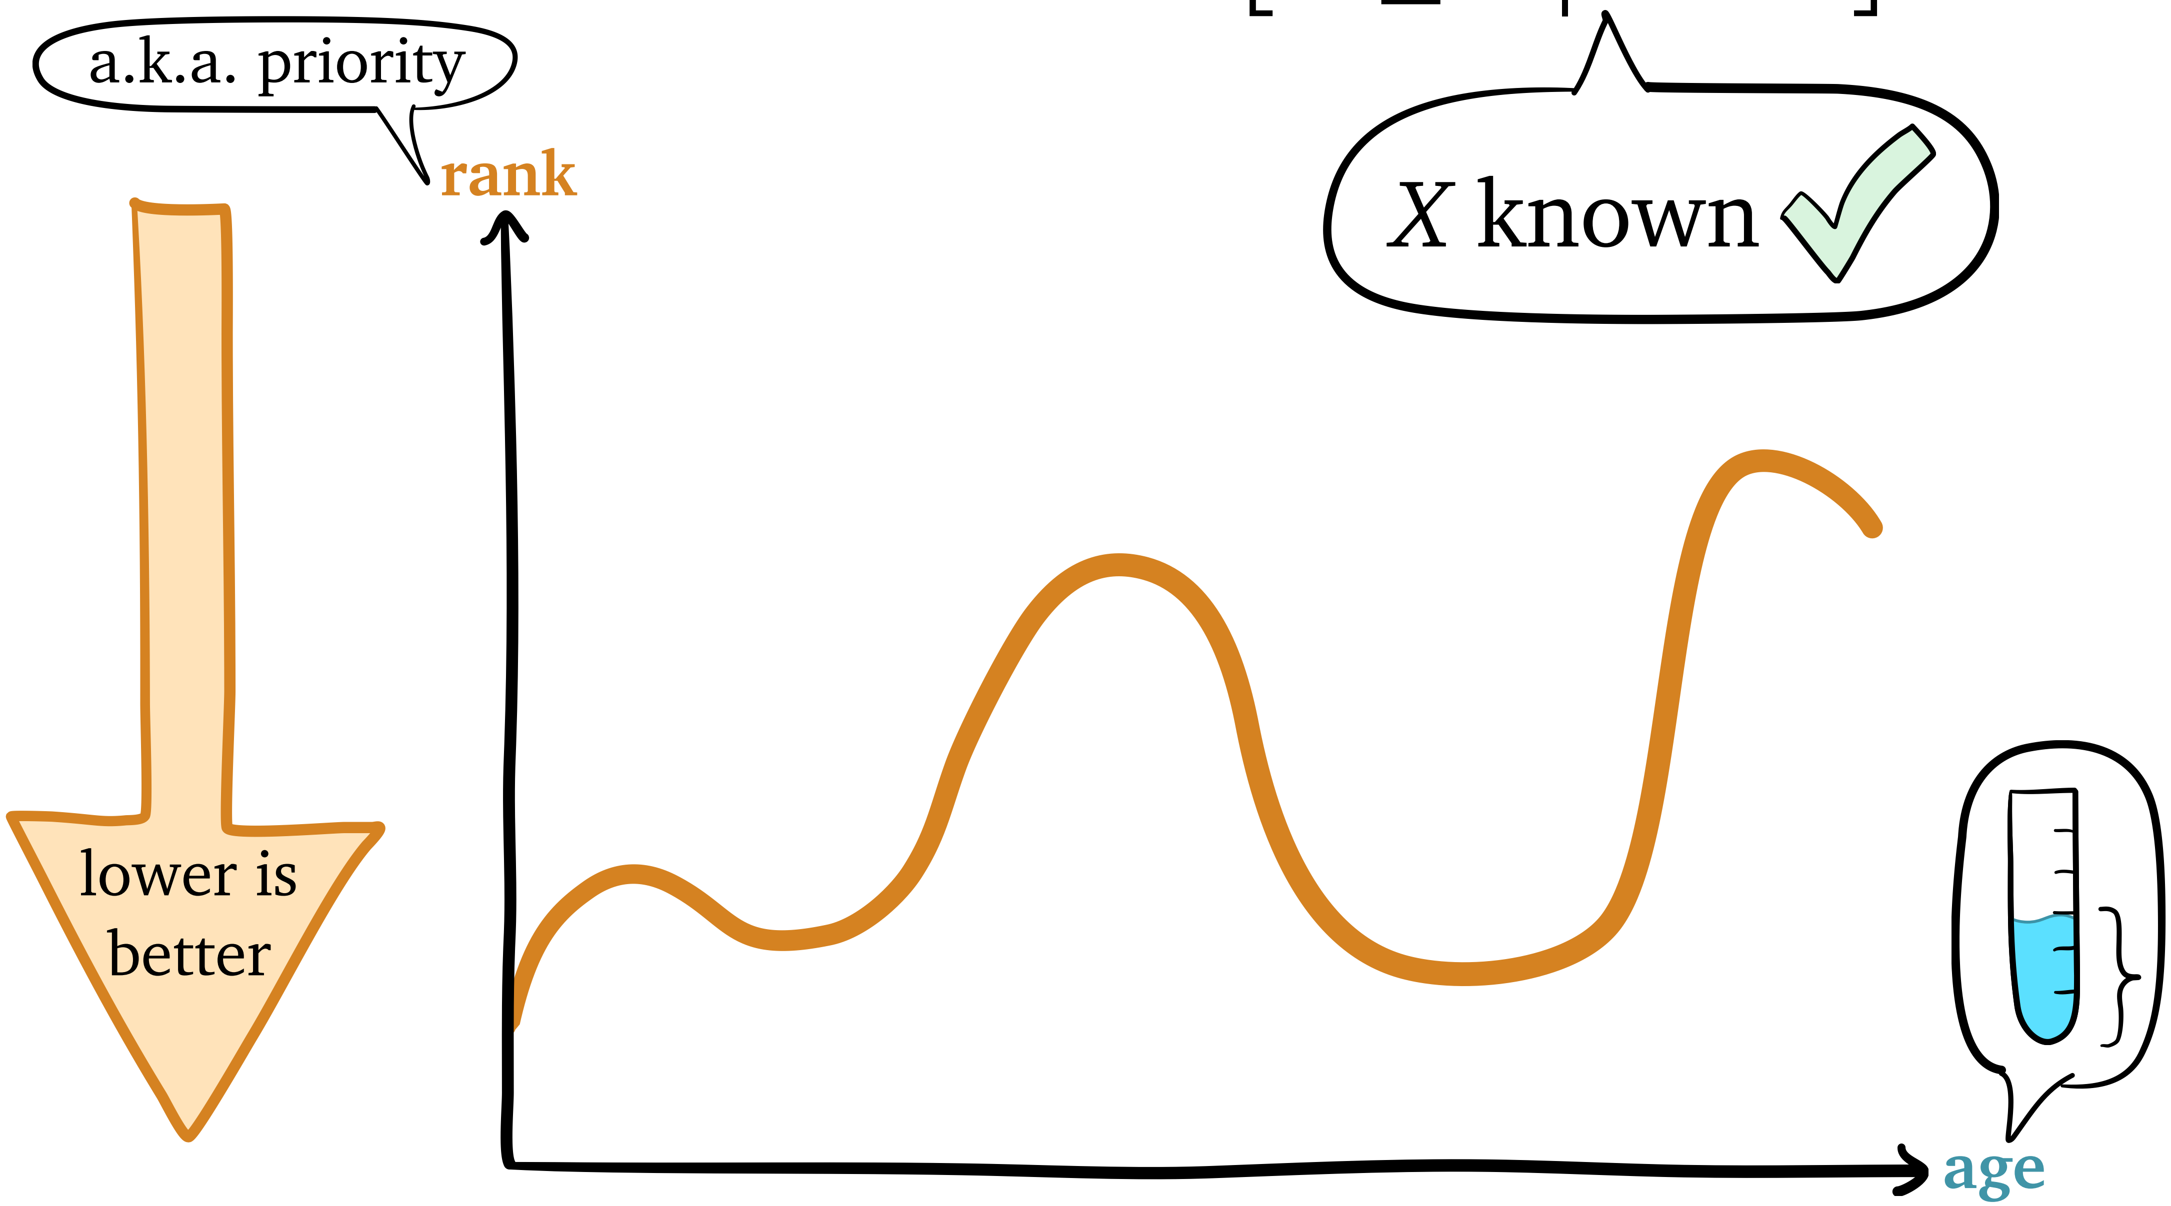

Optimal for M/G/1: Gittins

$$r_{\text{Gittins}}(a) = \inf_{b>a} \frac{\mathbf{E}[\min\{X, b\} \mid X > a]}{\mathbf{P}[X \le b \mid X > a]}$$

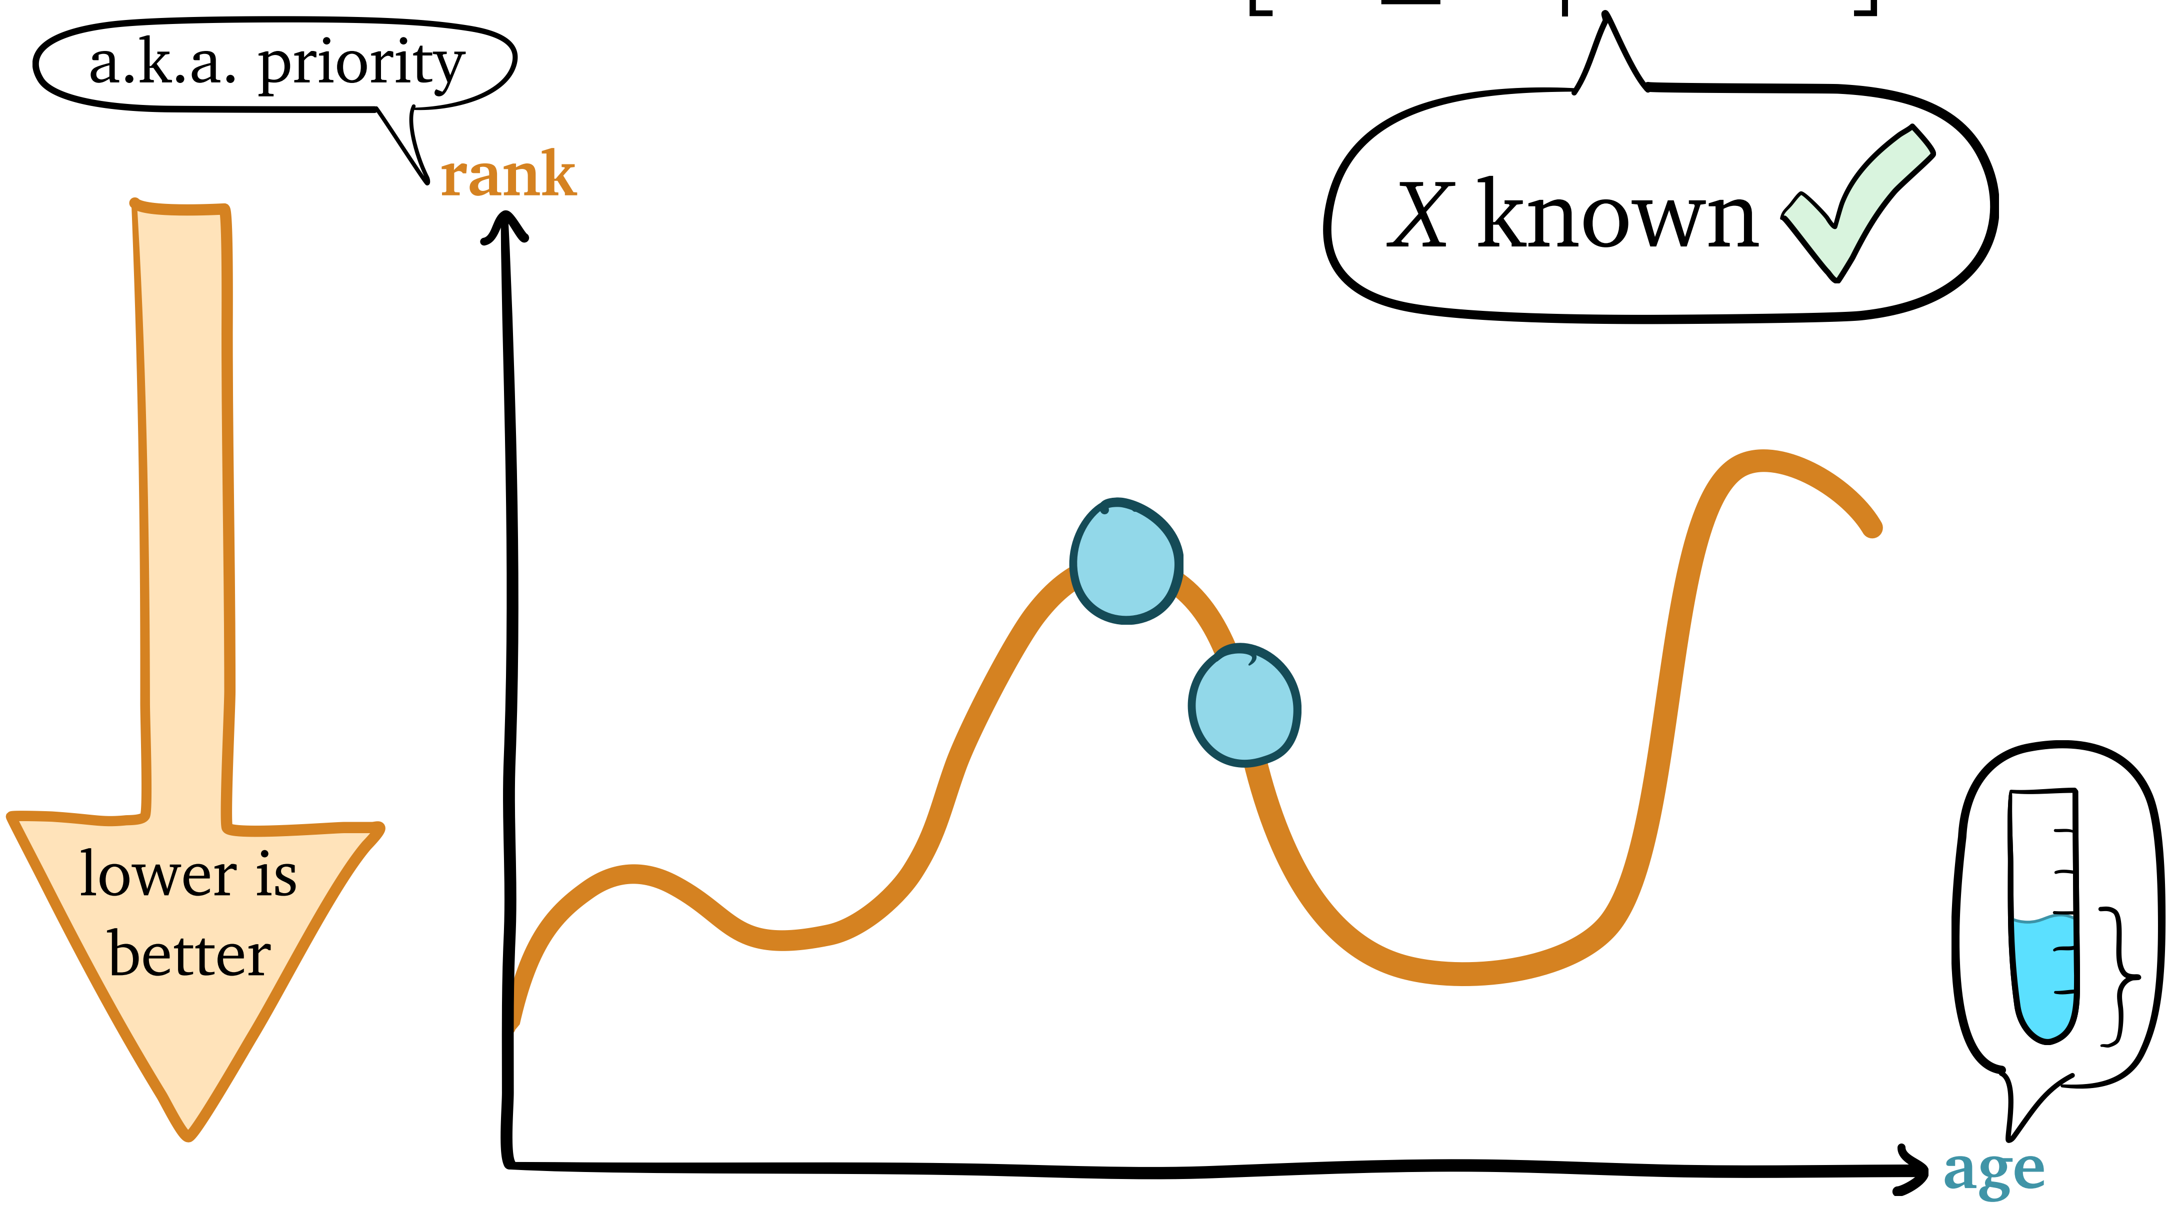

Optimal for M/G/1: Gittins

$$r_{\text{Gittins}}(a) = \inf_{b>a} \frac{\mathbf{E}[\min\{X, b\} \mid X > a]}{\mathbf{P}[X \le b \mid X > a]}$$

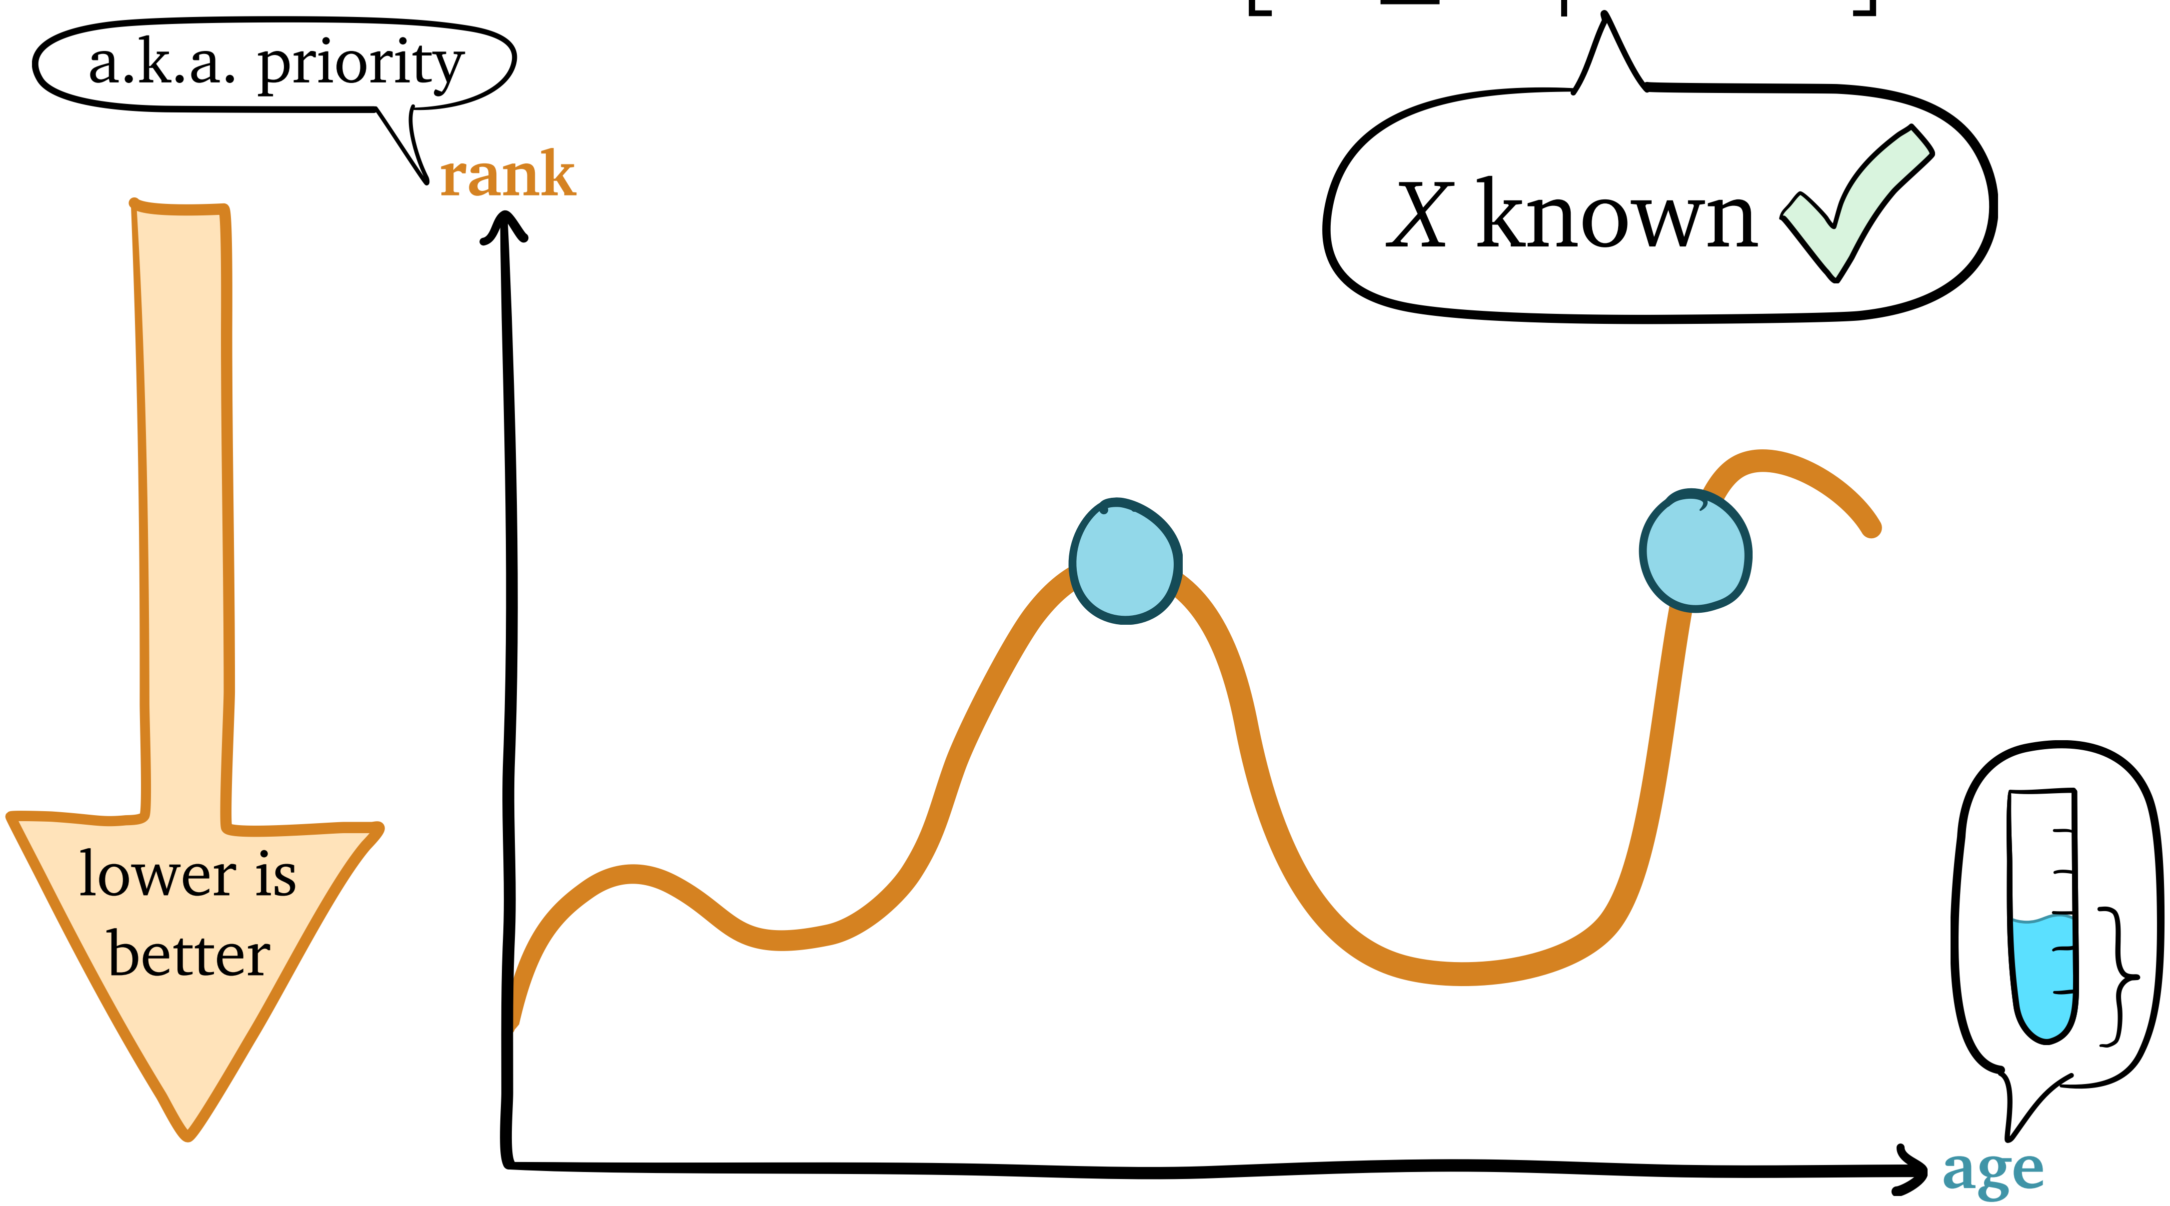

Optimal for M/G/1: Gittins

$$r_{\text{Gittins}}(a) = \inf_{b>a} \frac{\mathbf{E}[\min\{X, b\} \mid X > a]}{\mathbf{P}[X \leq b \mid X > a]}$$

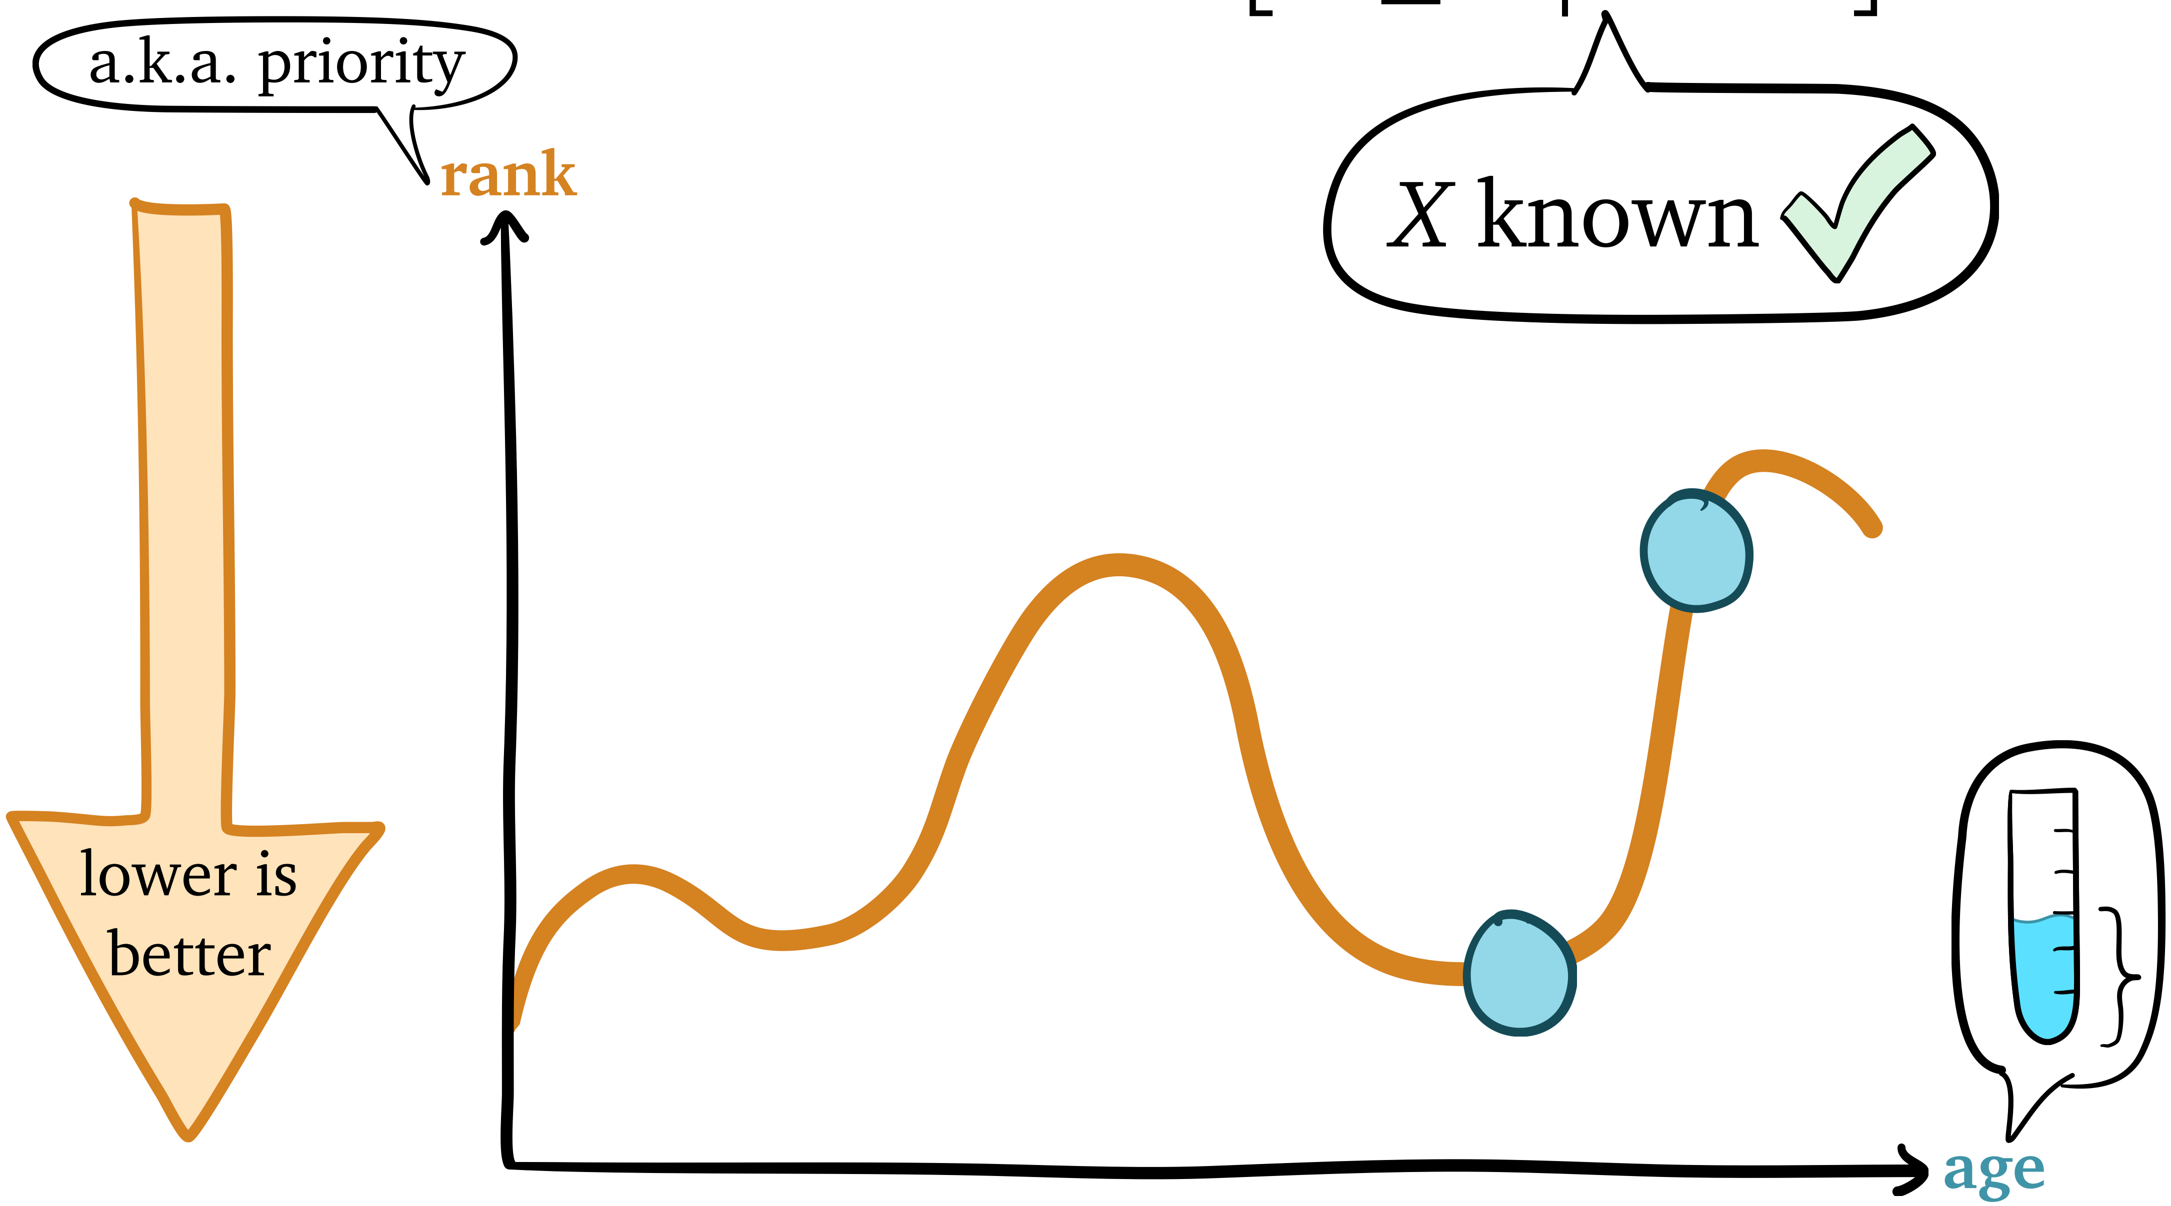

Minimizing $\mathbf{E}[T]$

	Known job sizes	Unknown job sizes
M/G/1	**SRPT**	
M/G/k	**SRPT-*k***	

in heavy traffic,
X “finite variance”

Minimizing **E**[T]

	Known job sizes	Unknown job sizes
M/G/1	**SRPT**	**Gittins**
M/G/k	**SRPT-k**	

in heavy traffic, X "finite variance"

Minimizing $\mathbf{E}[T]$

	Known job sizes	Unknown job sizes
M/G/1	**SRPT**	**Gittins**
M/G/k	**SRPT-k**	**???**

in heavy traffic, X "finite variance"

Minimizing $\mathbf{E}[T]$

	Known job sizes	Unknown job sizes
M/G/1	**SRPT**	**Gittins**
M/G/k	**SRPT-*k***	**???**

in heavy traffic, X "finite variance"

Does **Gittins-*k*** work?

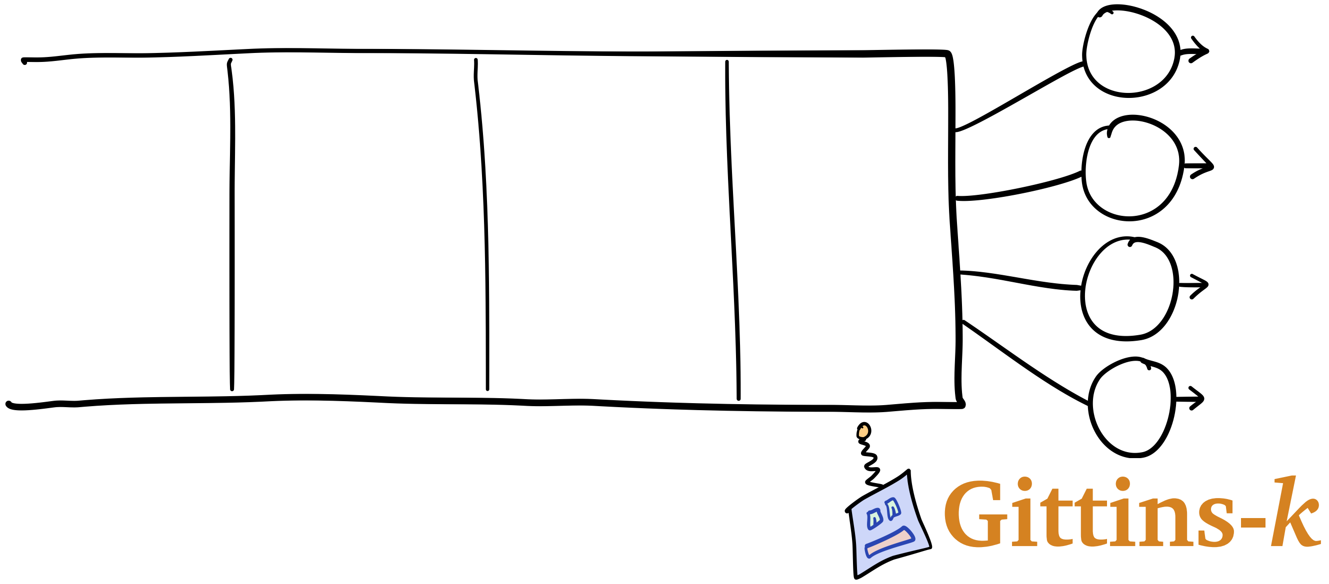
Gittins-*k*

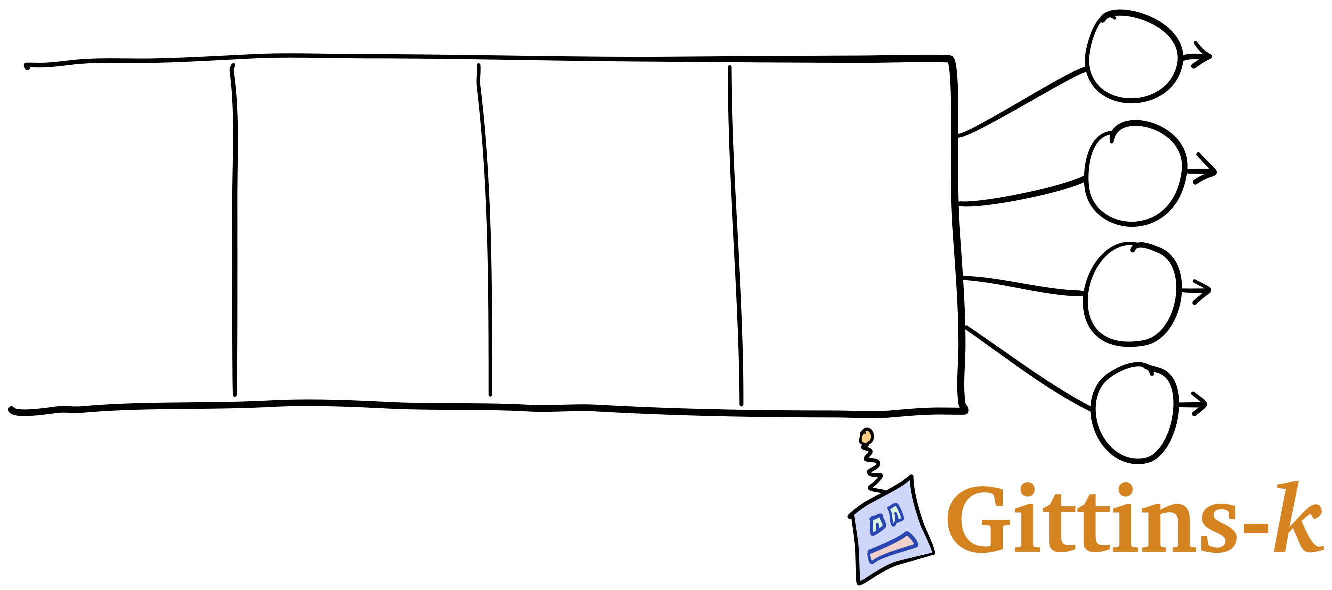
Gittins-k

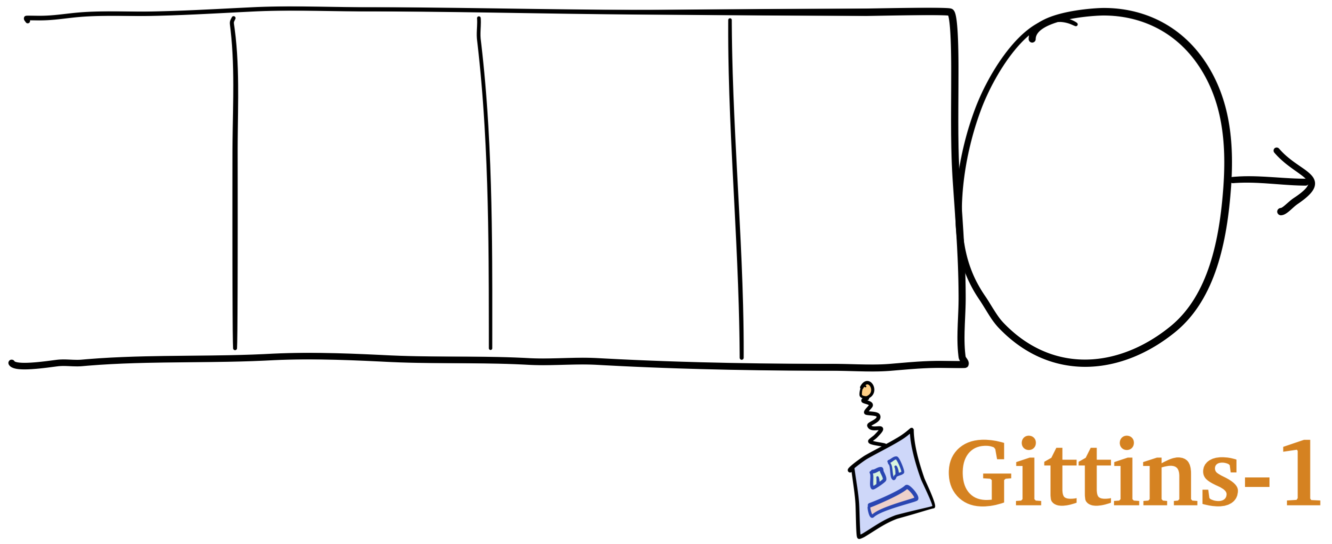
Gittins-1

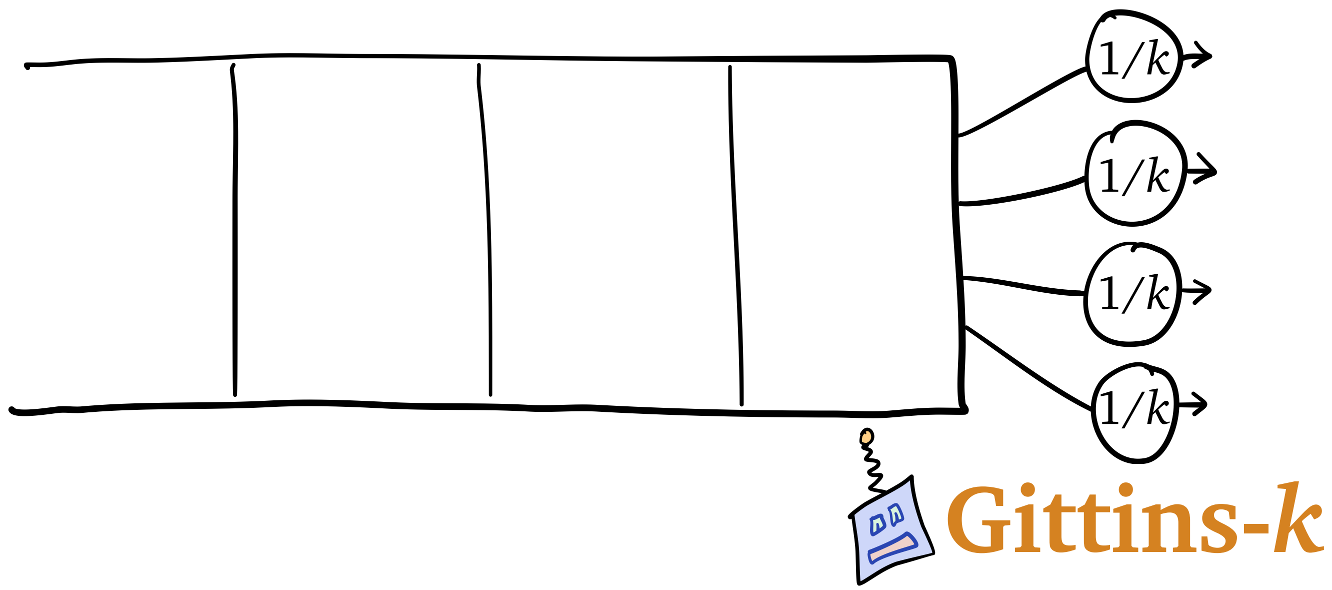
1/k
1/k
1/k
1/k
Gittins-k

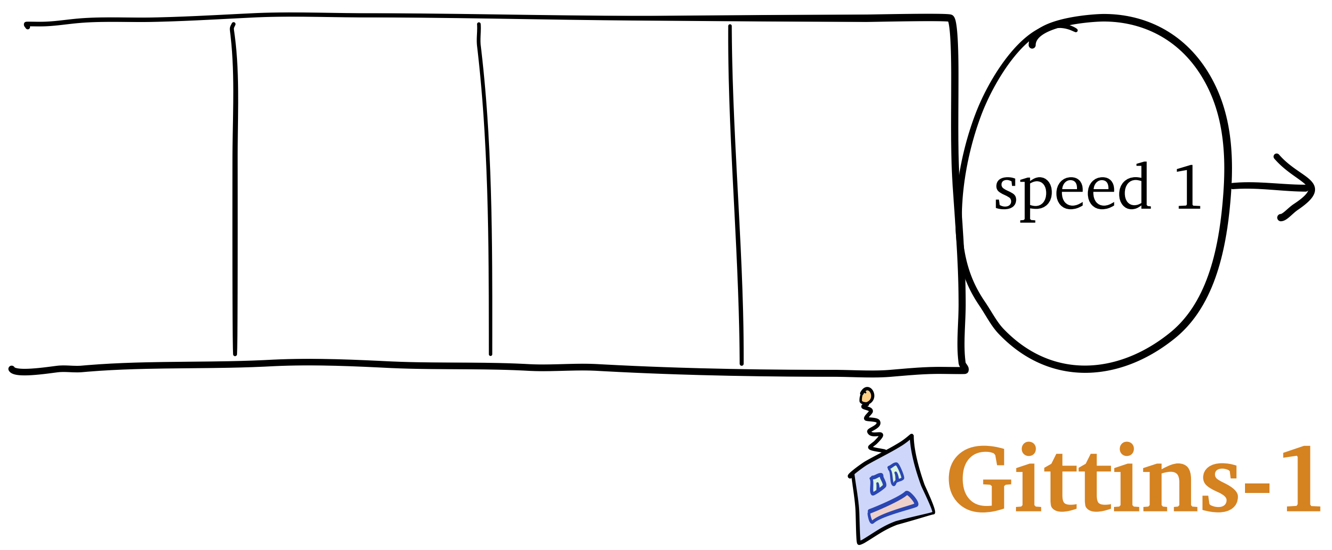
speed 1
Gittins-1

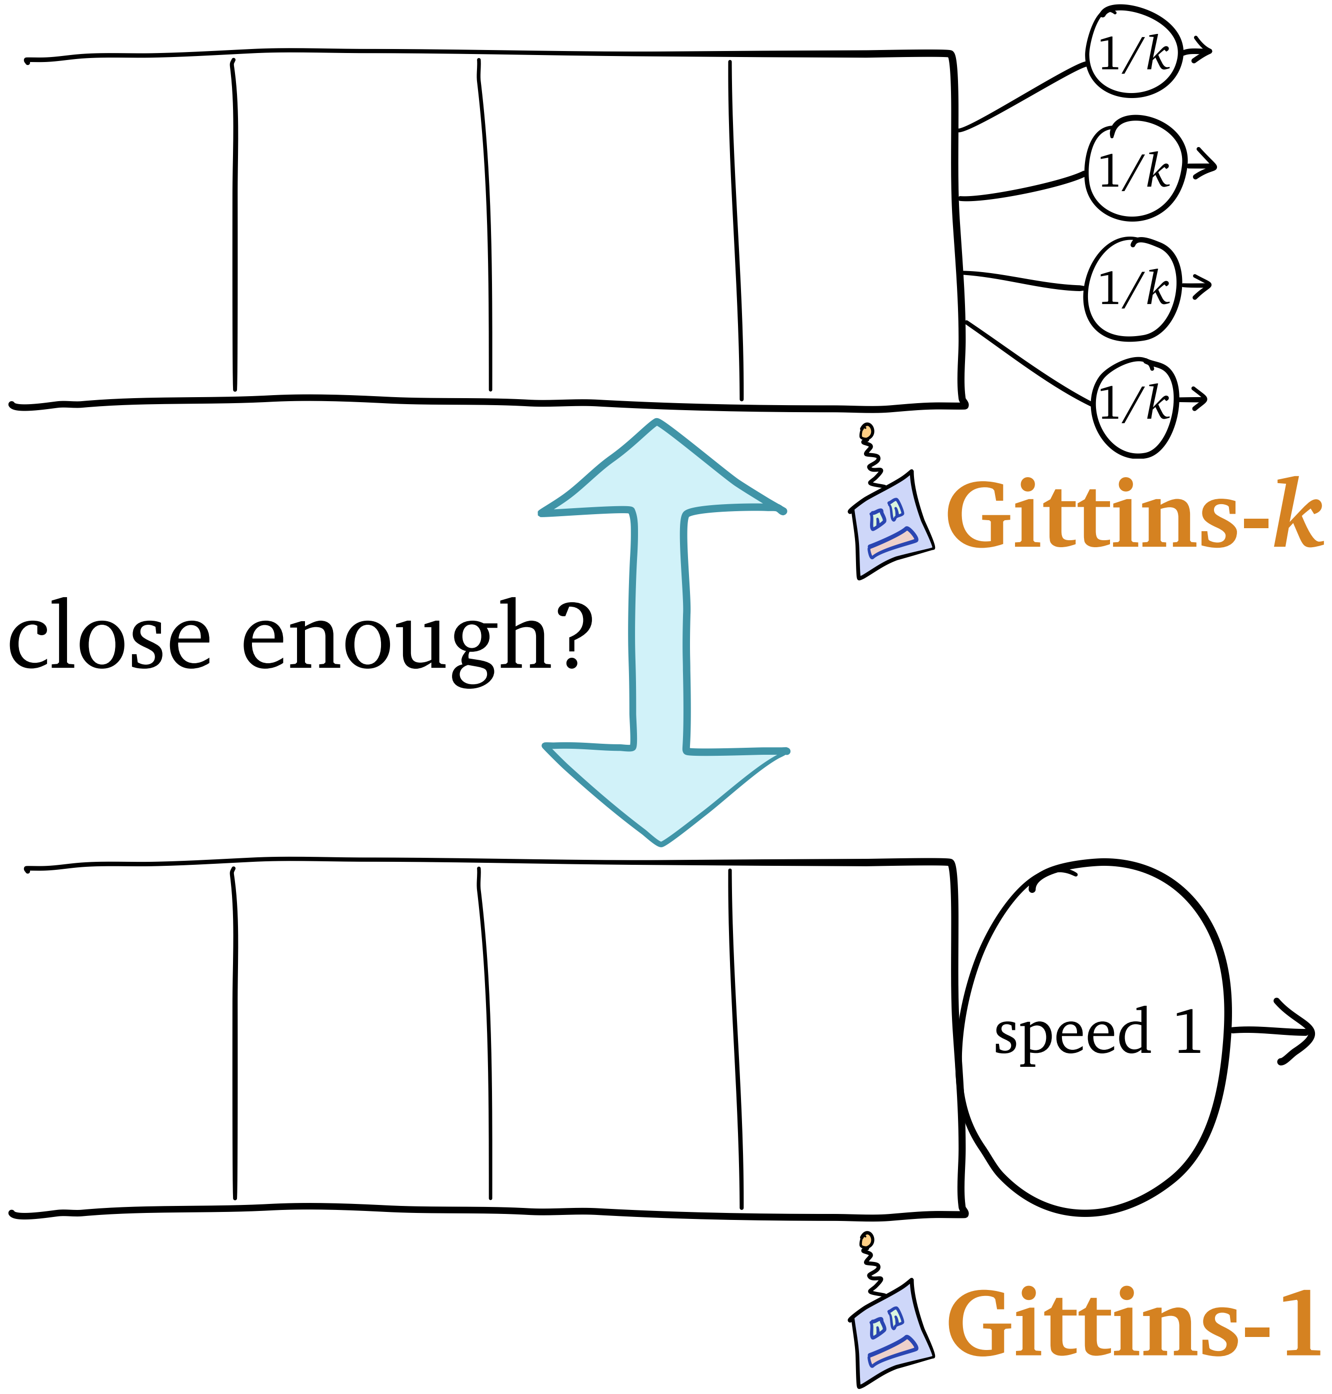
1/k
1/k
1/k
1/k
Gittins-k
close enough?
speed 1
Gittins-1

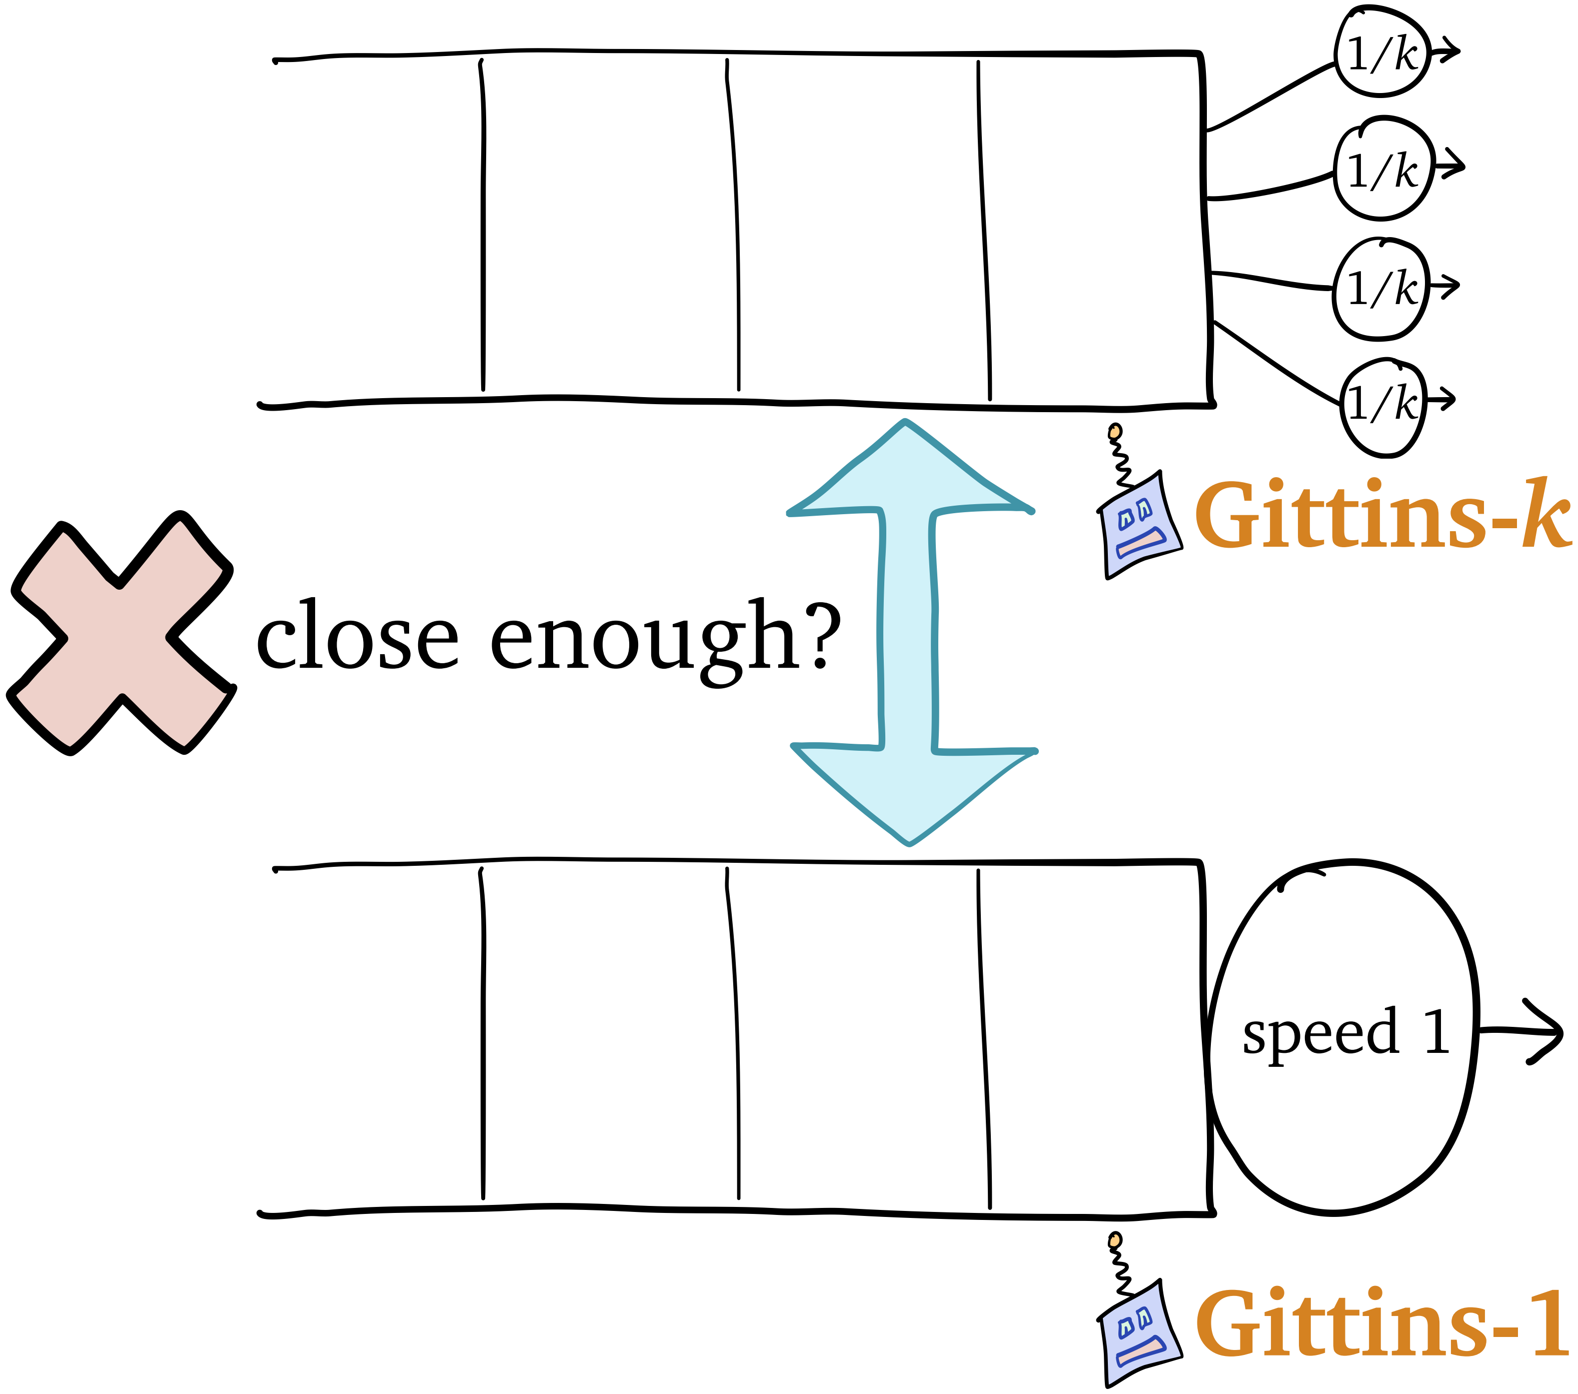
1/k
1/k
1/k
1/k
Gittins-k
close enough?
speed 1
Gittins-1

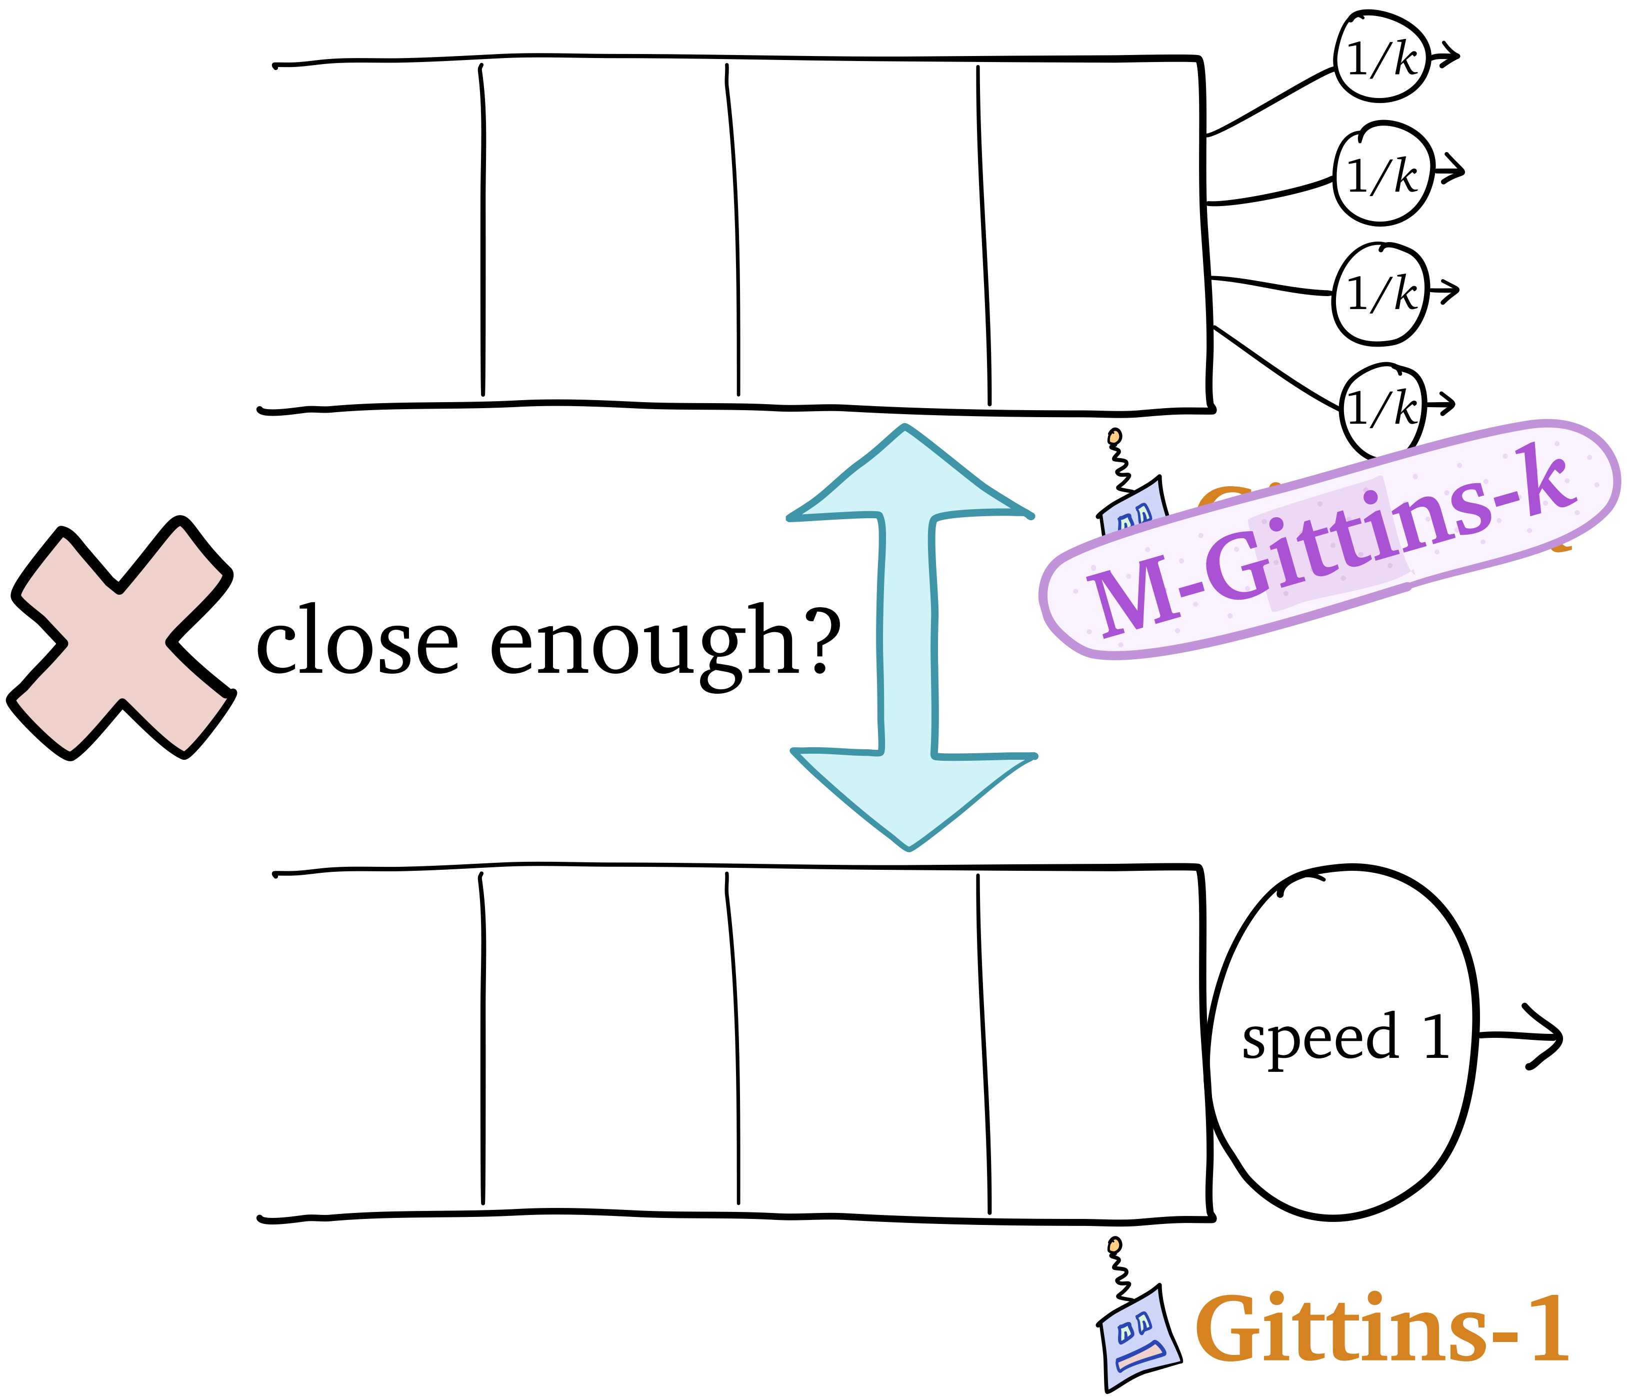

1/k
1/k
1/k
1/k
M-Gittins-k
close enough?
speed 1
Gittins-1

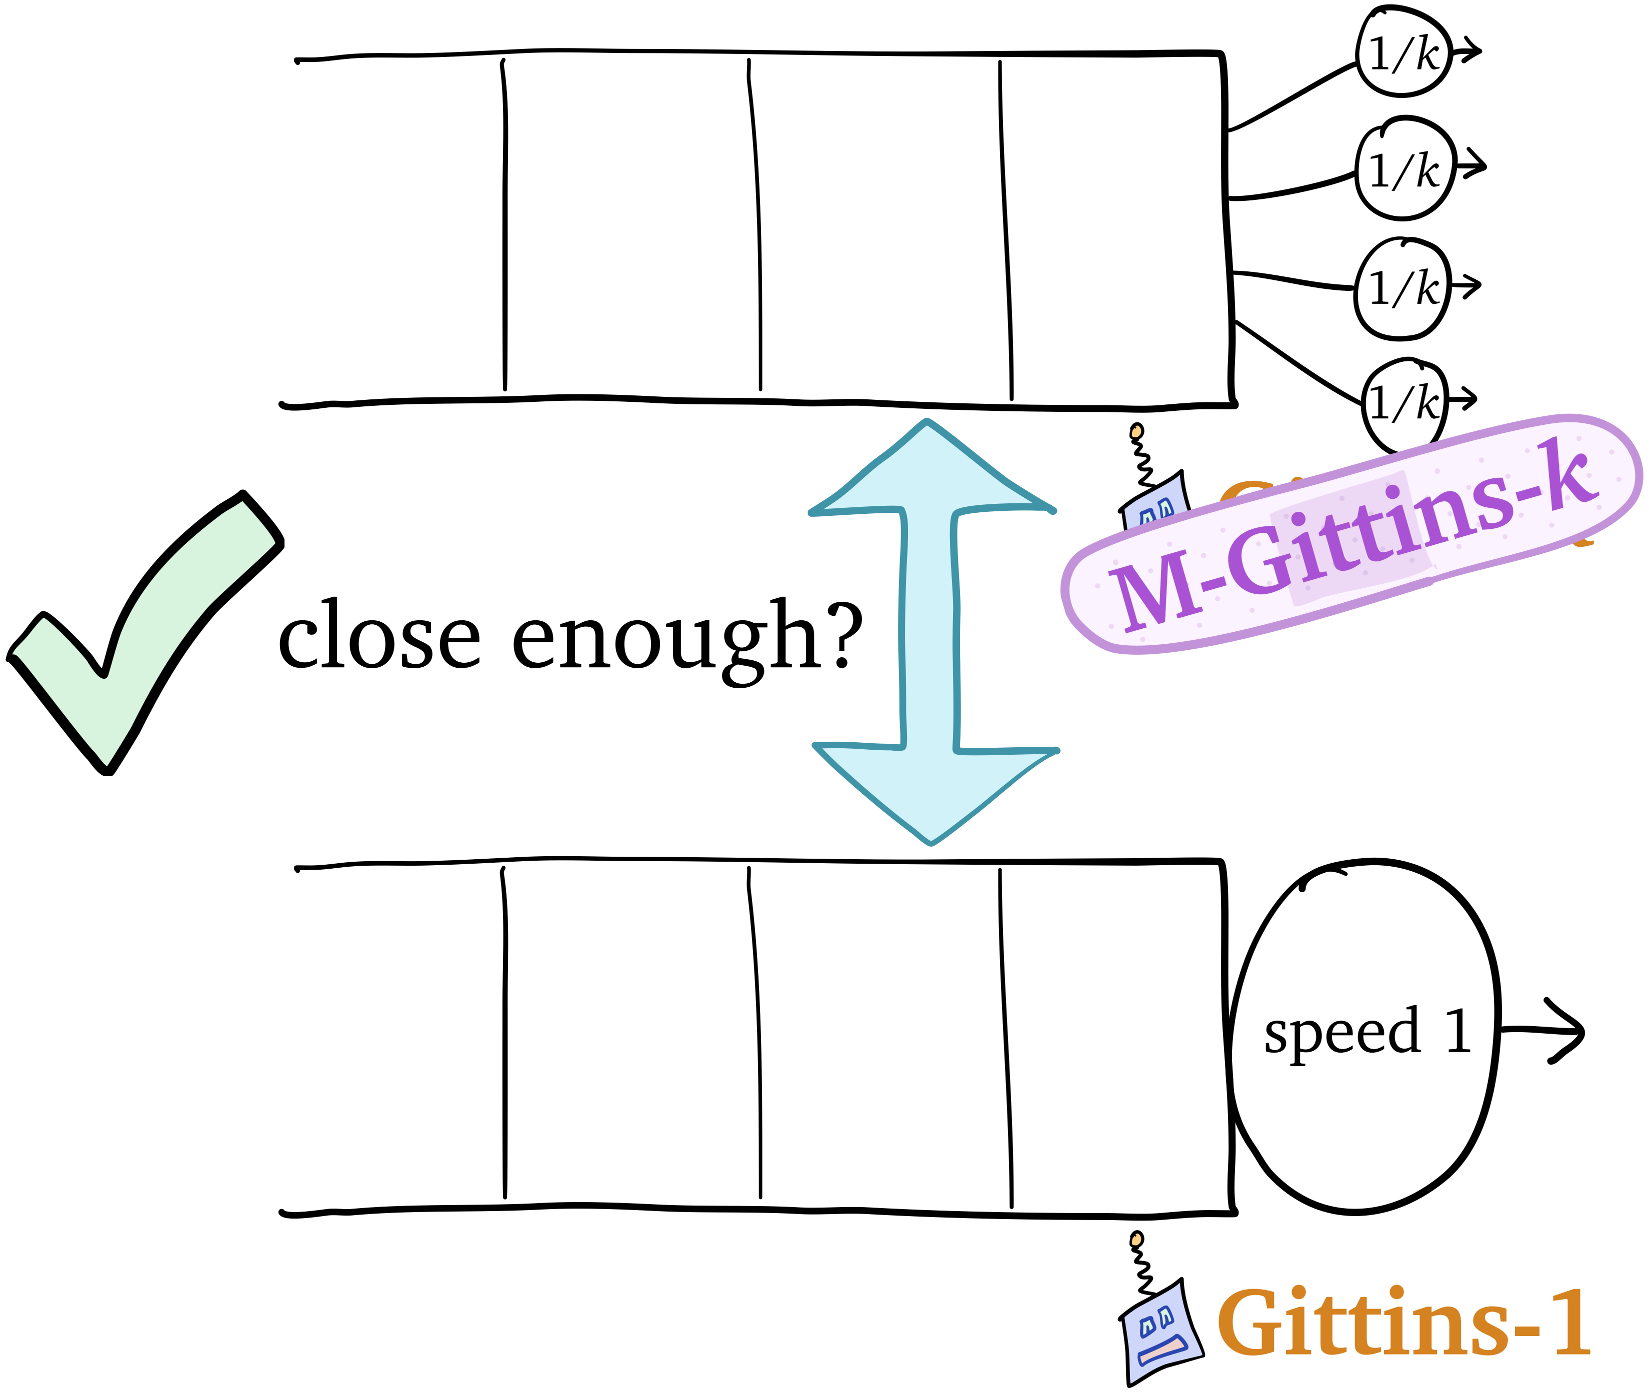
1/k
1/k
1/k
1/k
M-Gittins-k
close enough?
speed 1
Gittins-1

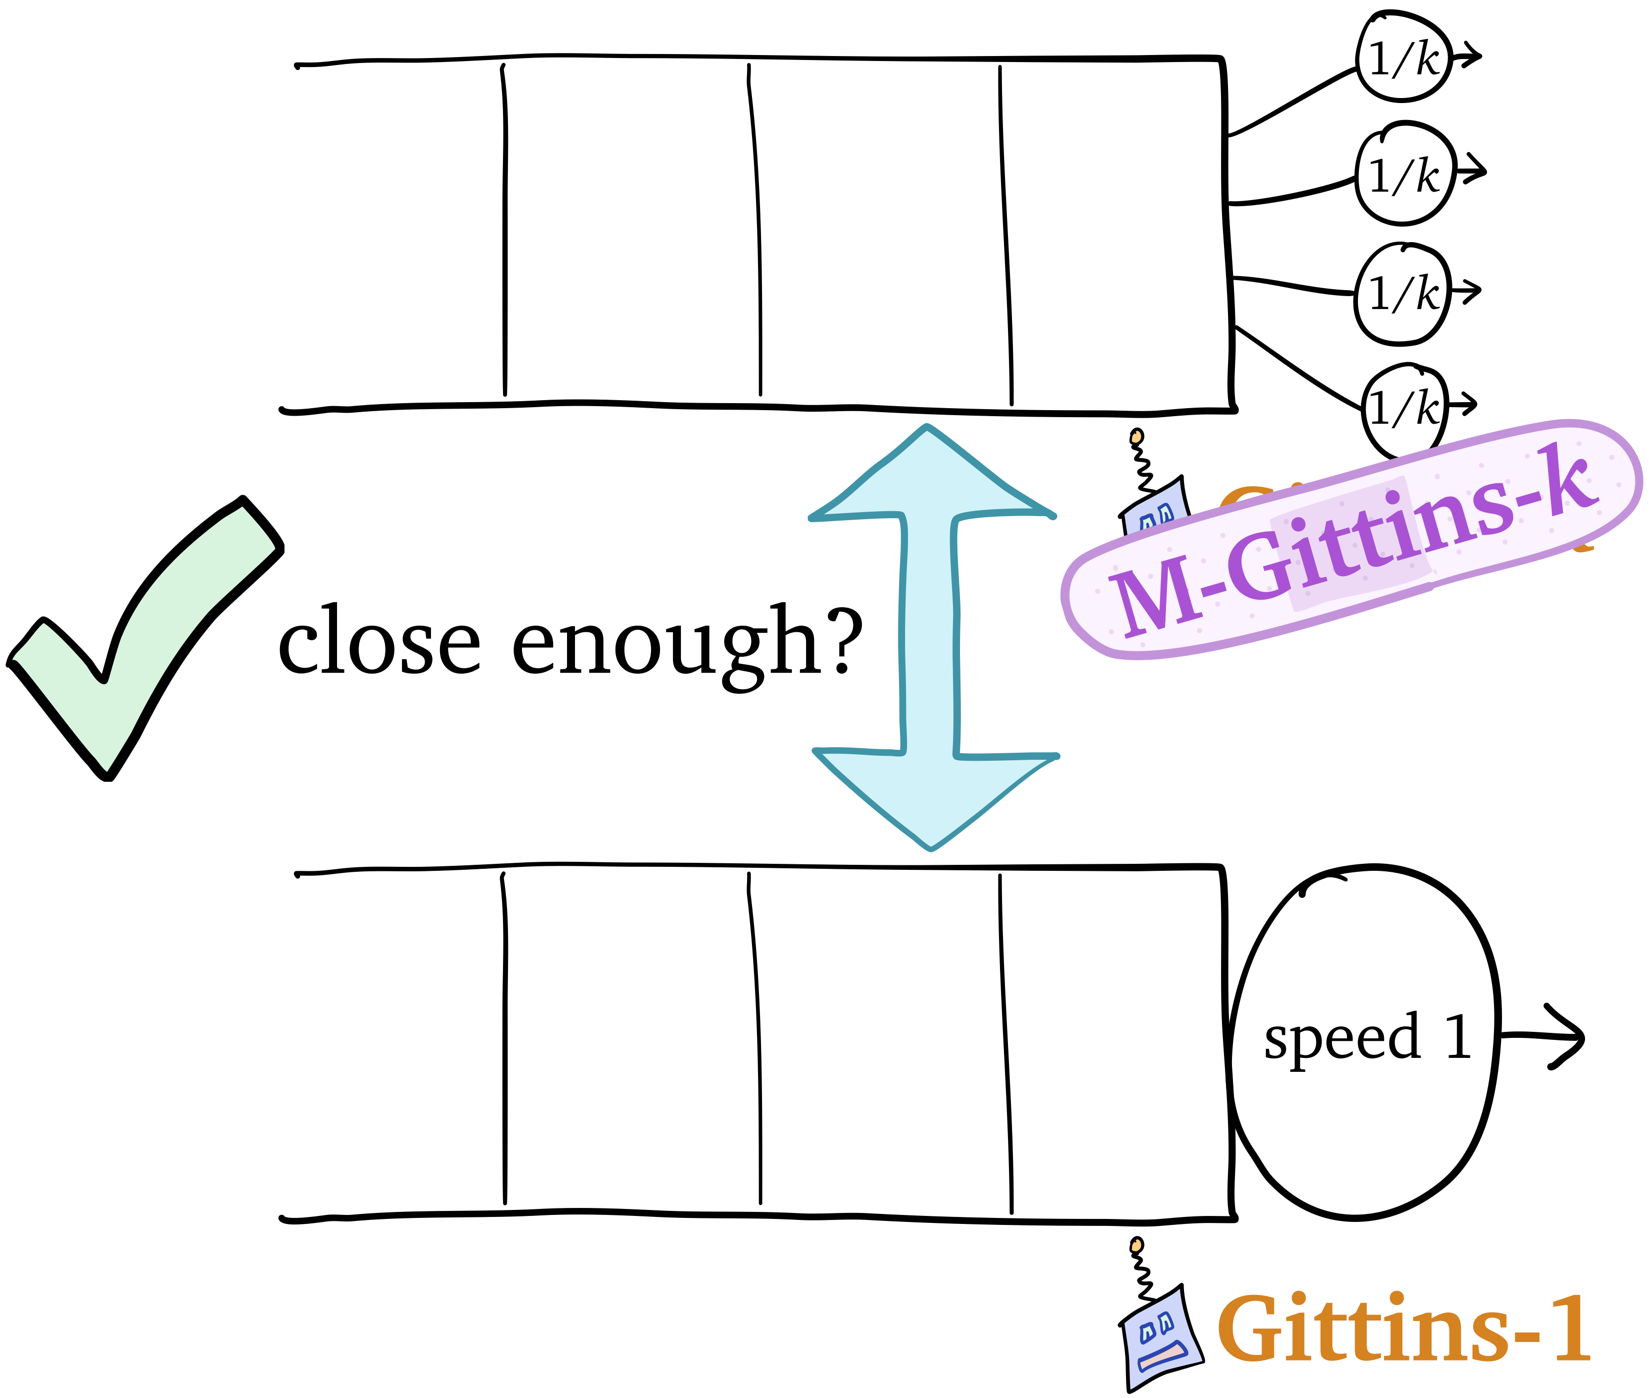

Theorem:

$$\lim_{\rho \to 1} \frac{\mathbf{E}[T_{\text{M-Gittins-}k}]}{\mathbf{E}[T_{\text{Gittins-1}}]} = 1$$

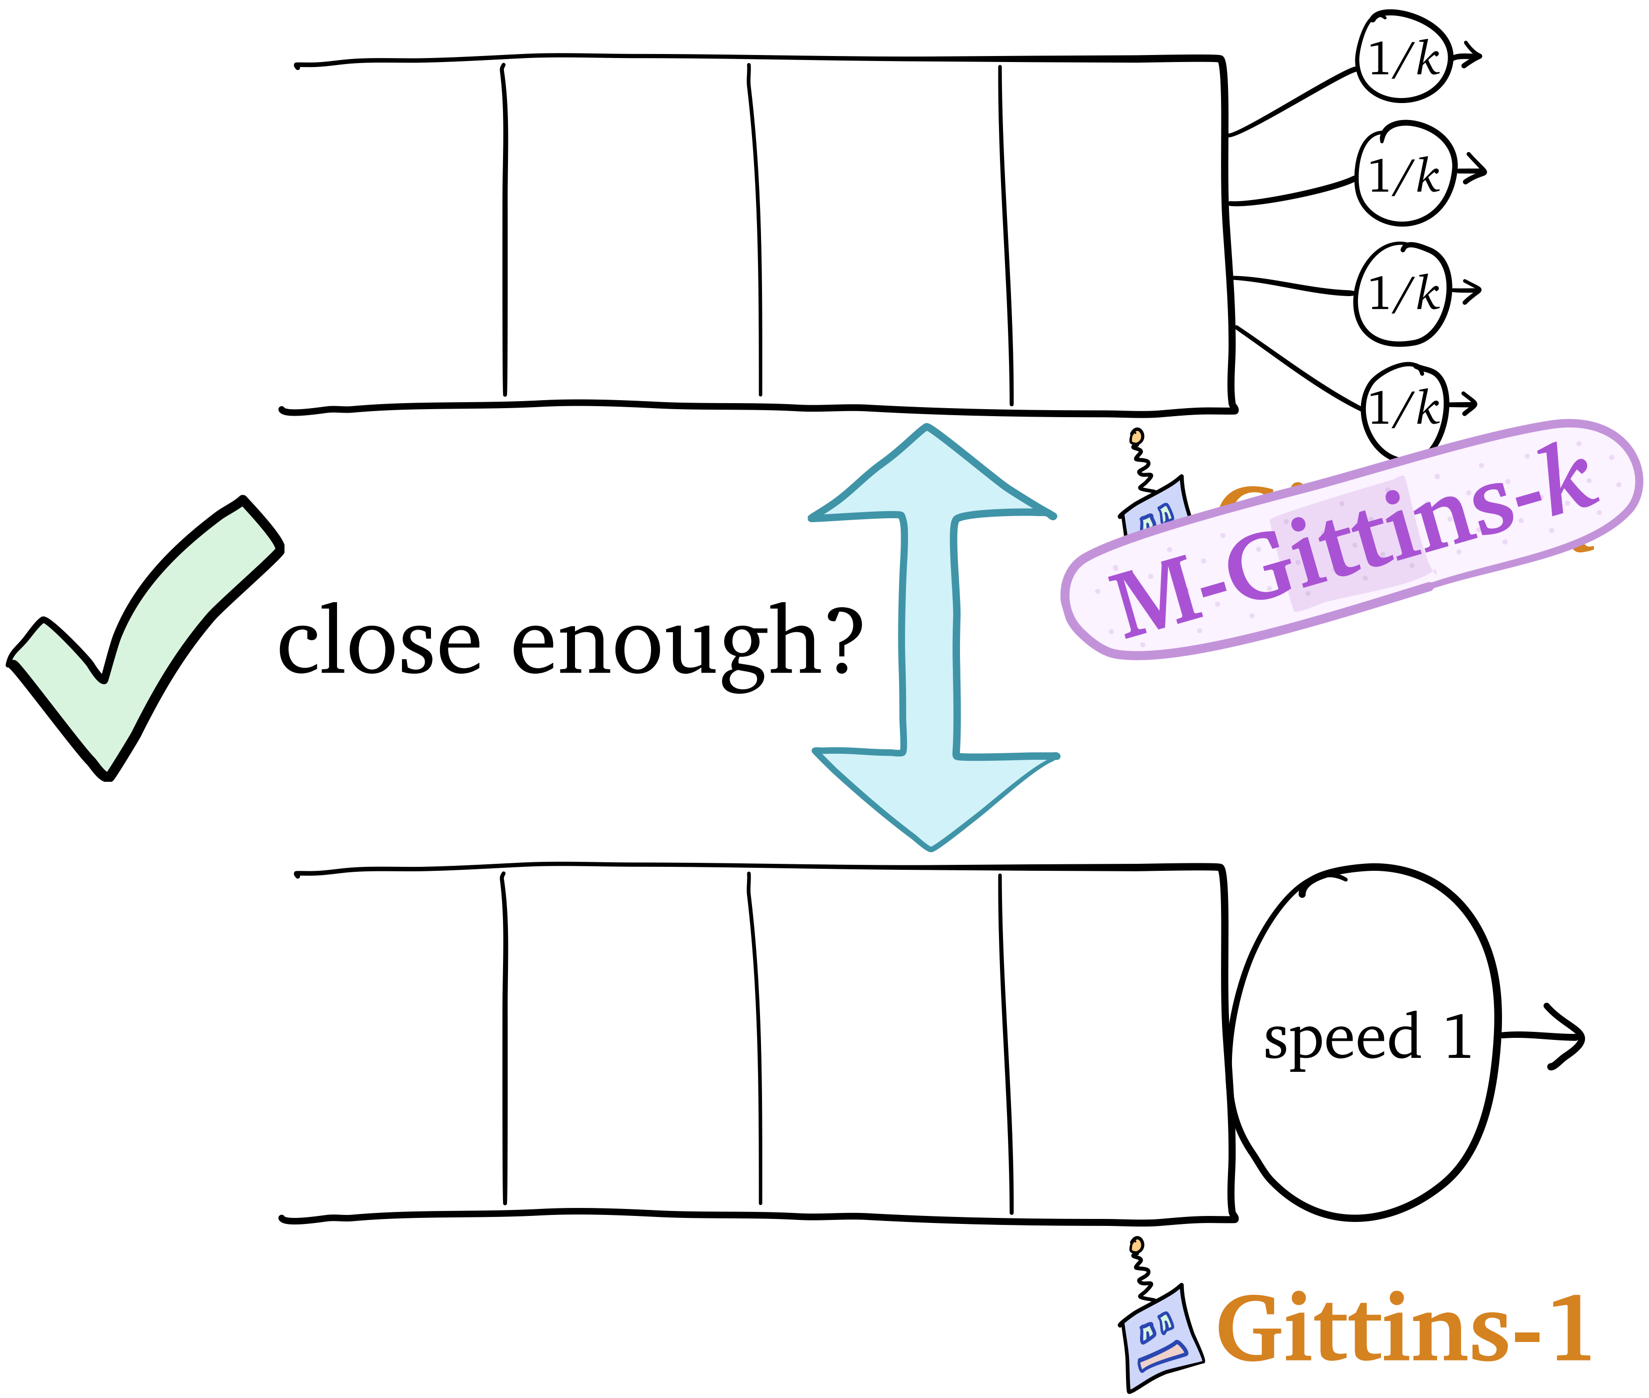

Theorem:

$$\lim_{\rho \to 1} \frac{\mathbf{E}[T_{\text{M-Gittins-}k}]}{\mathbf{E}[T_{\text{Gittins-1}}]} = 1$$

1. What's wrong with **Gittins-*k***?
2. What is **M-Gittins-*k***?
3. How does **M-Gittins-*k*** help?

Analyzing **Gittins-1**

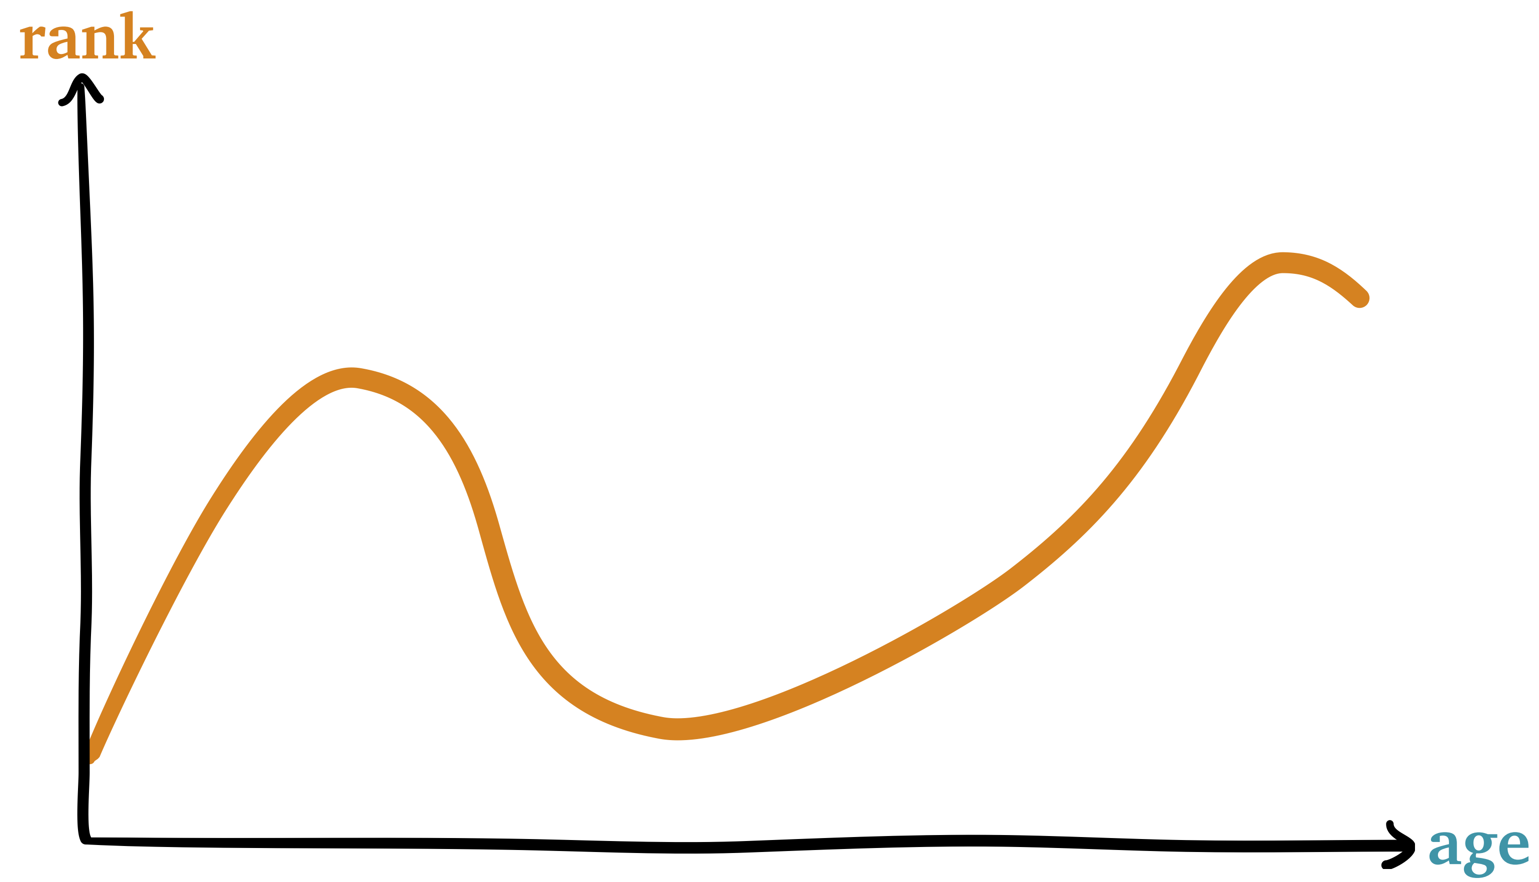

Analyzing **Gittins-1**

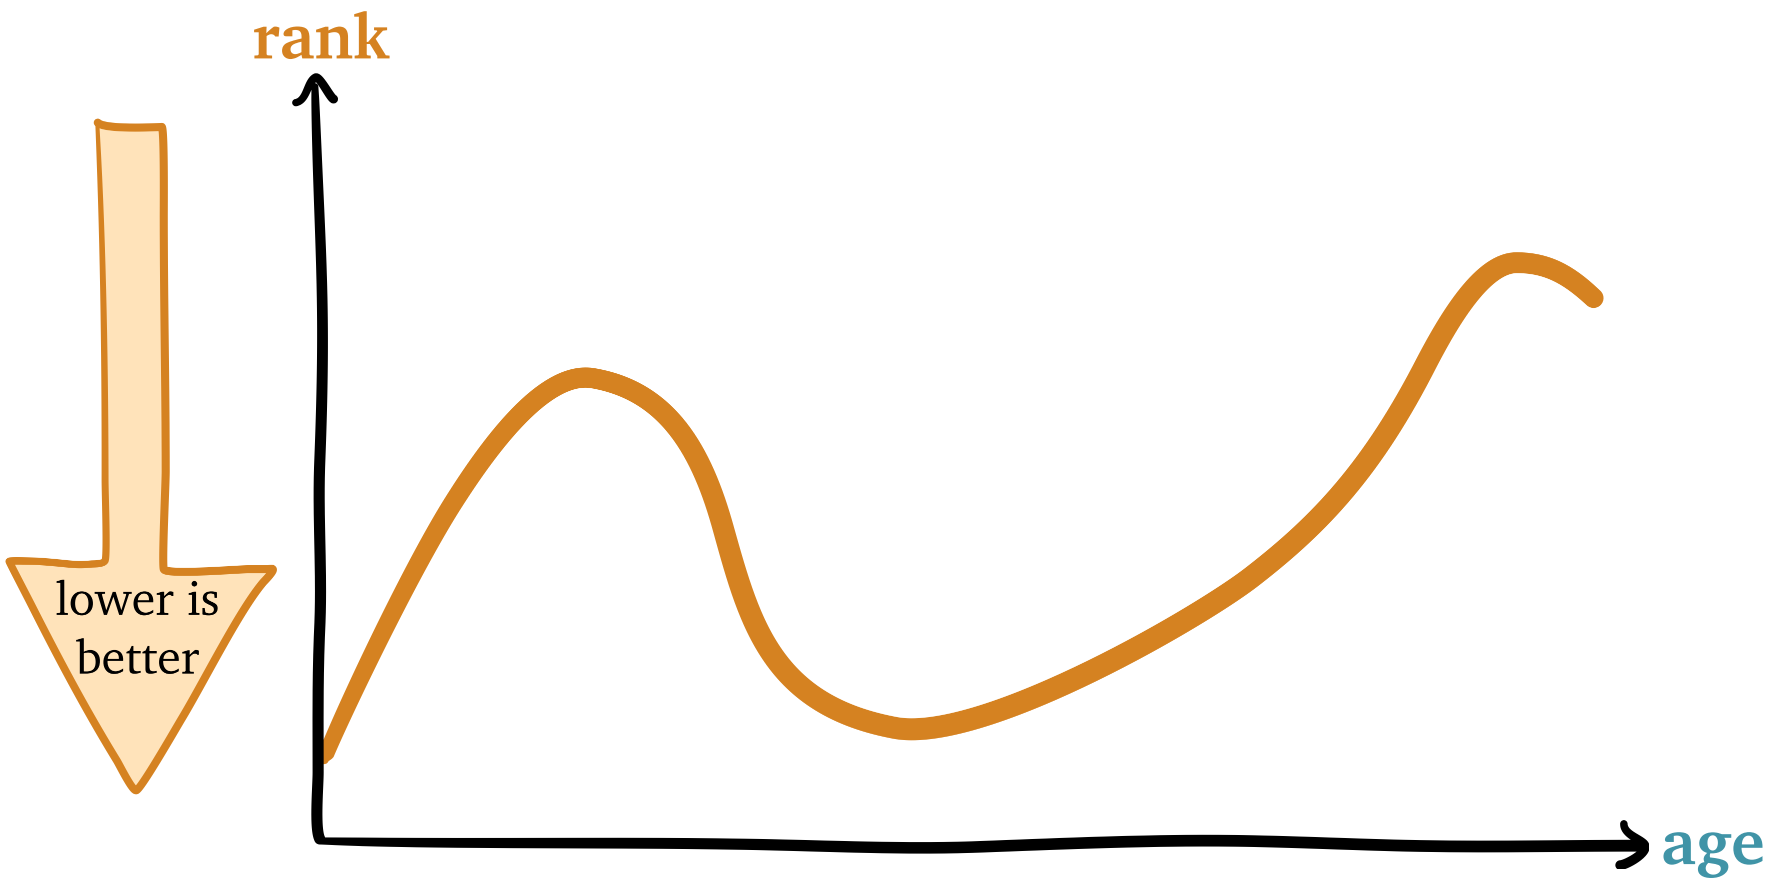

Analyzing **Gittins-1**

Suppose I'm a job of size x

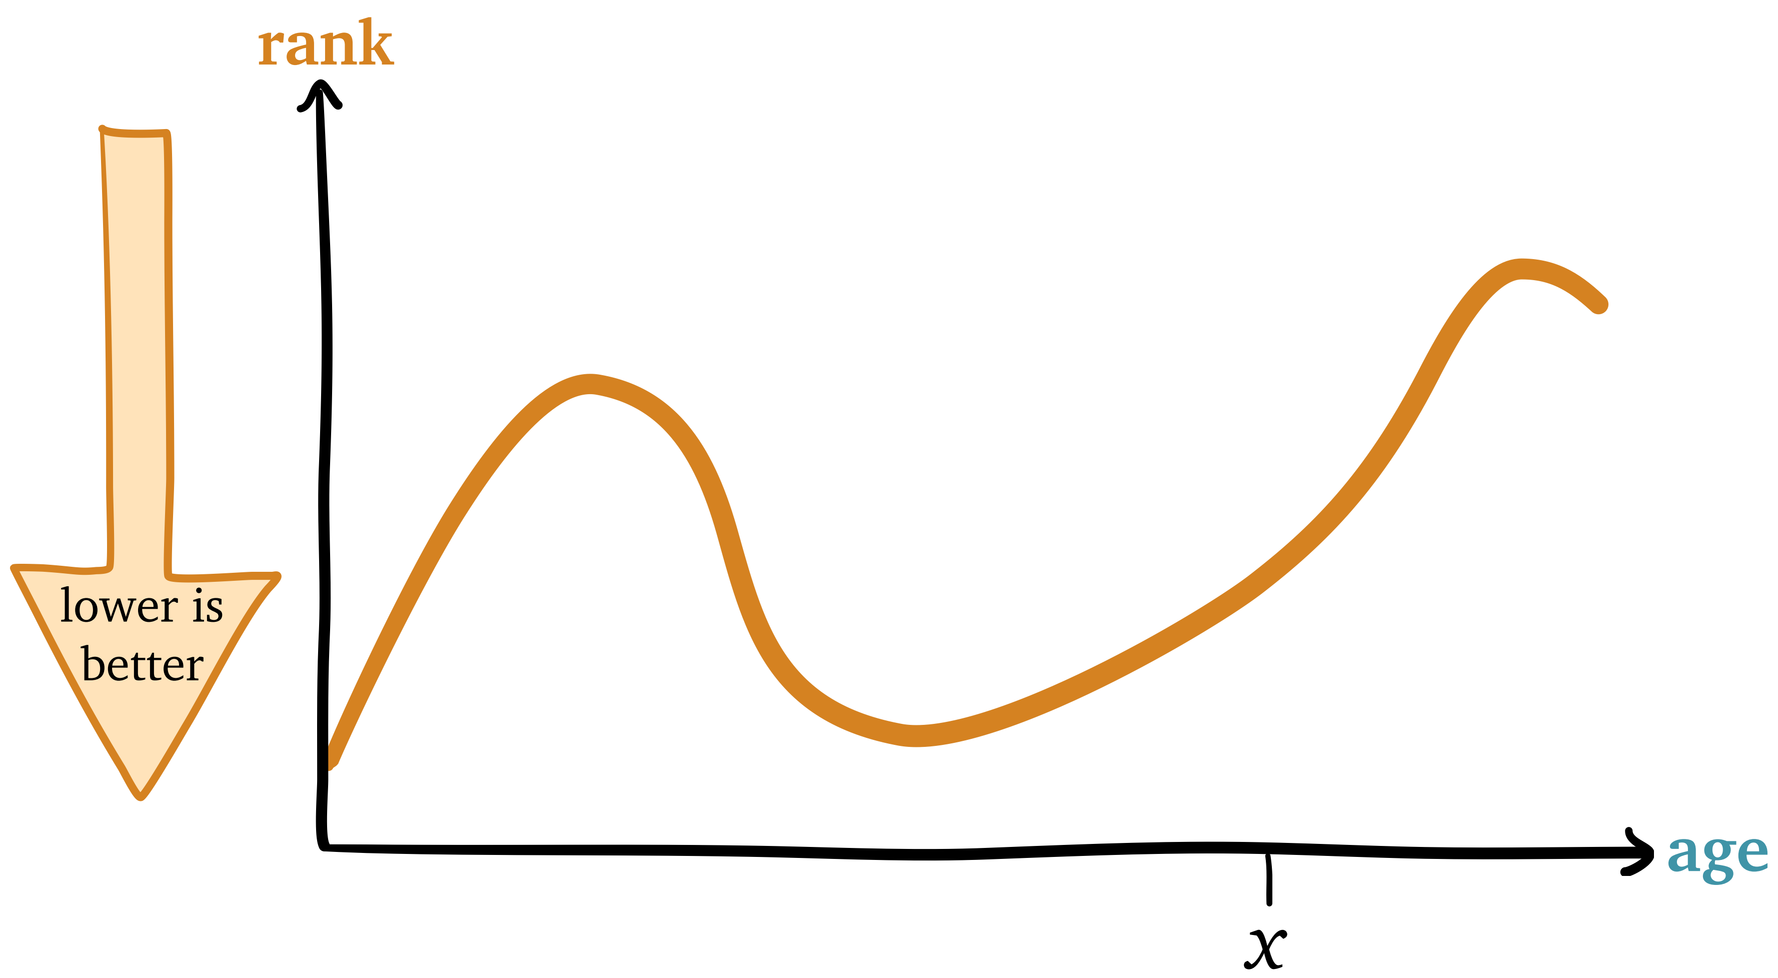

Analyzing Gittins-1

Suppose I'm a job of size x

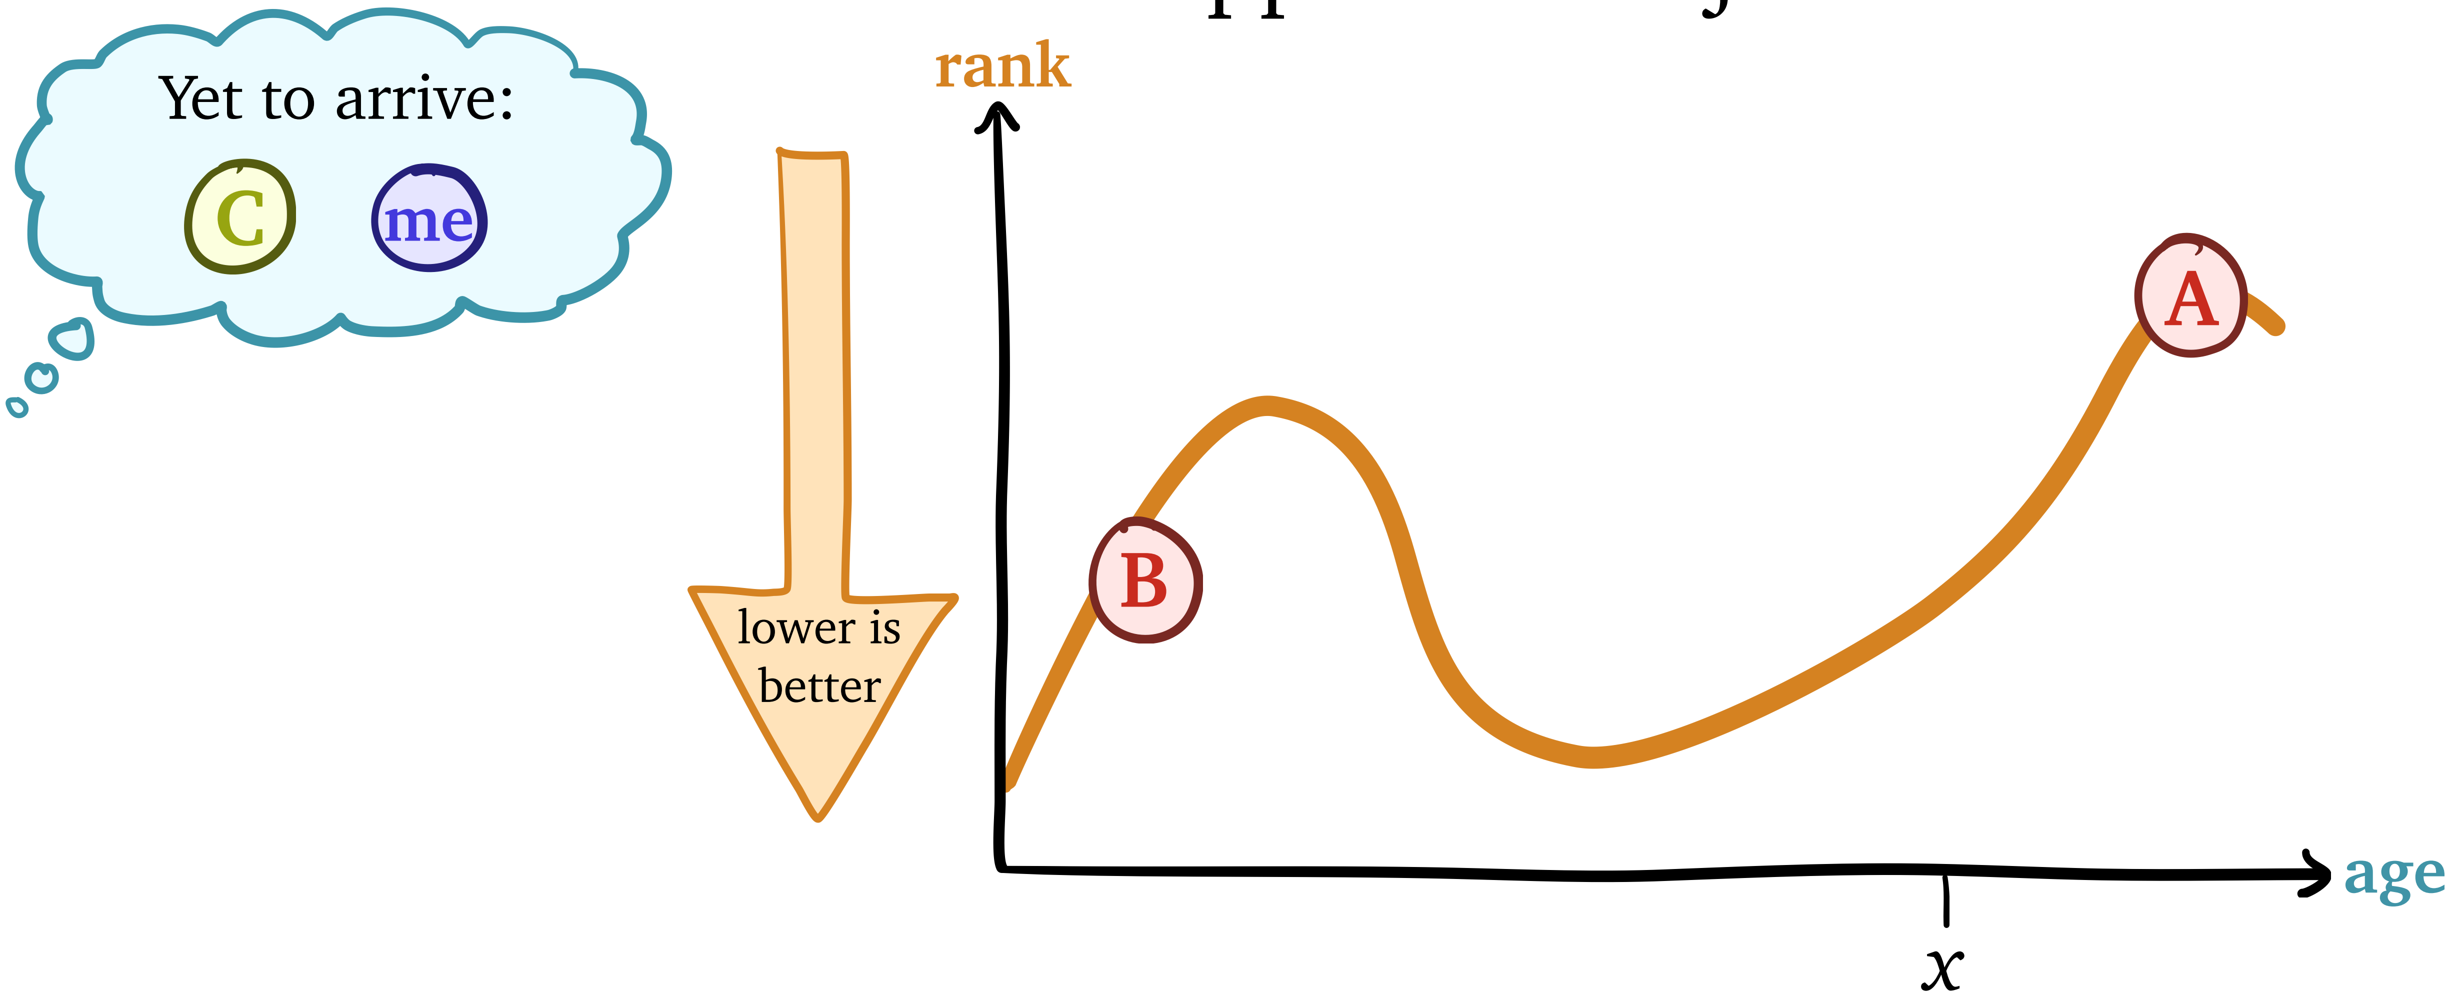

Analyzing **Gittins-1**

Suppose I'm a job of size x

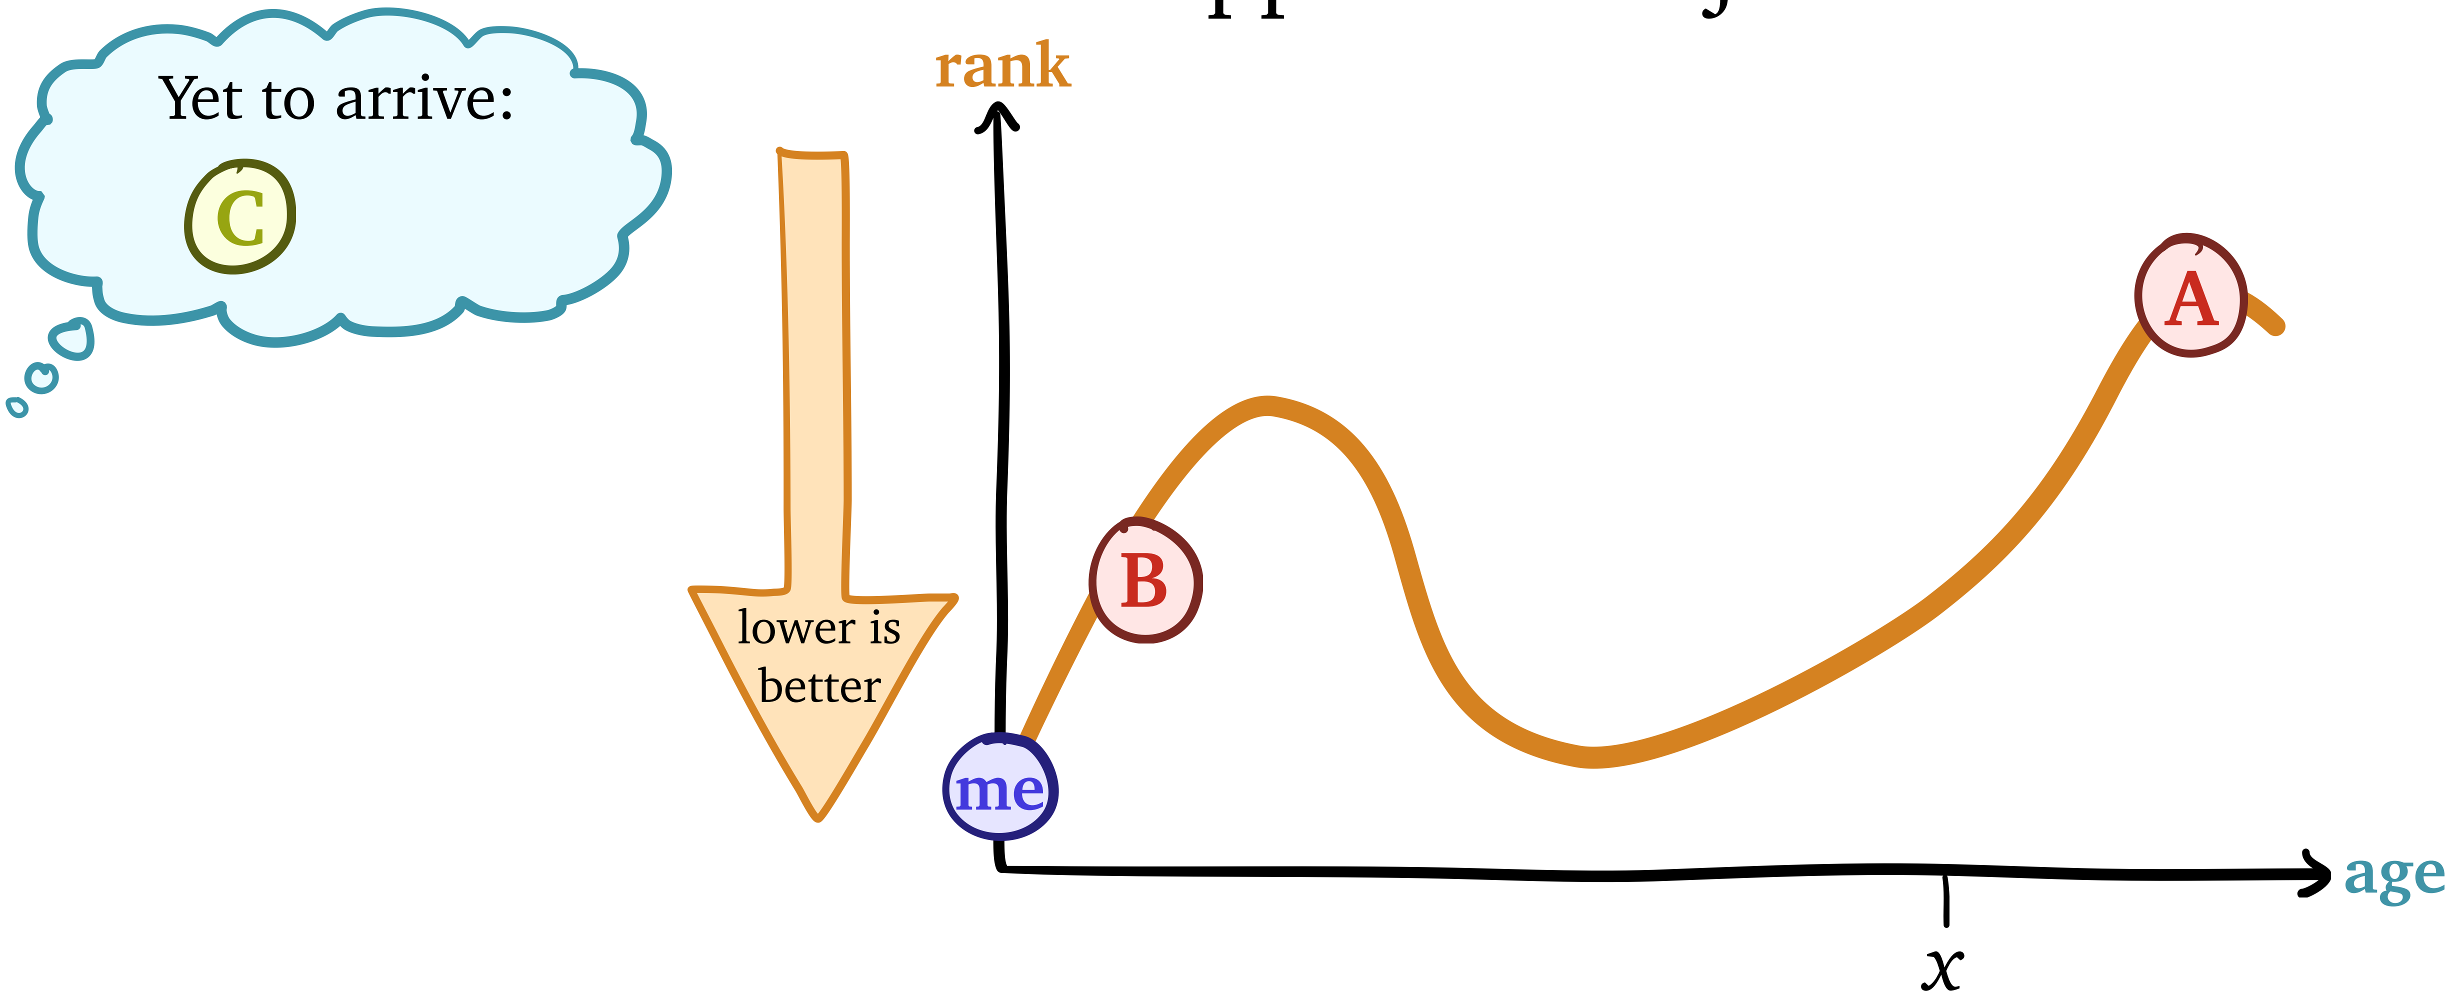

Analyzing **Gittins-1**

Suppose I'm a job of size x

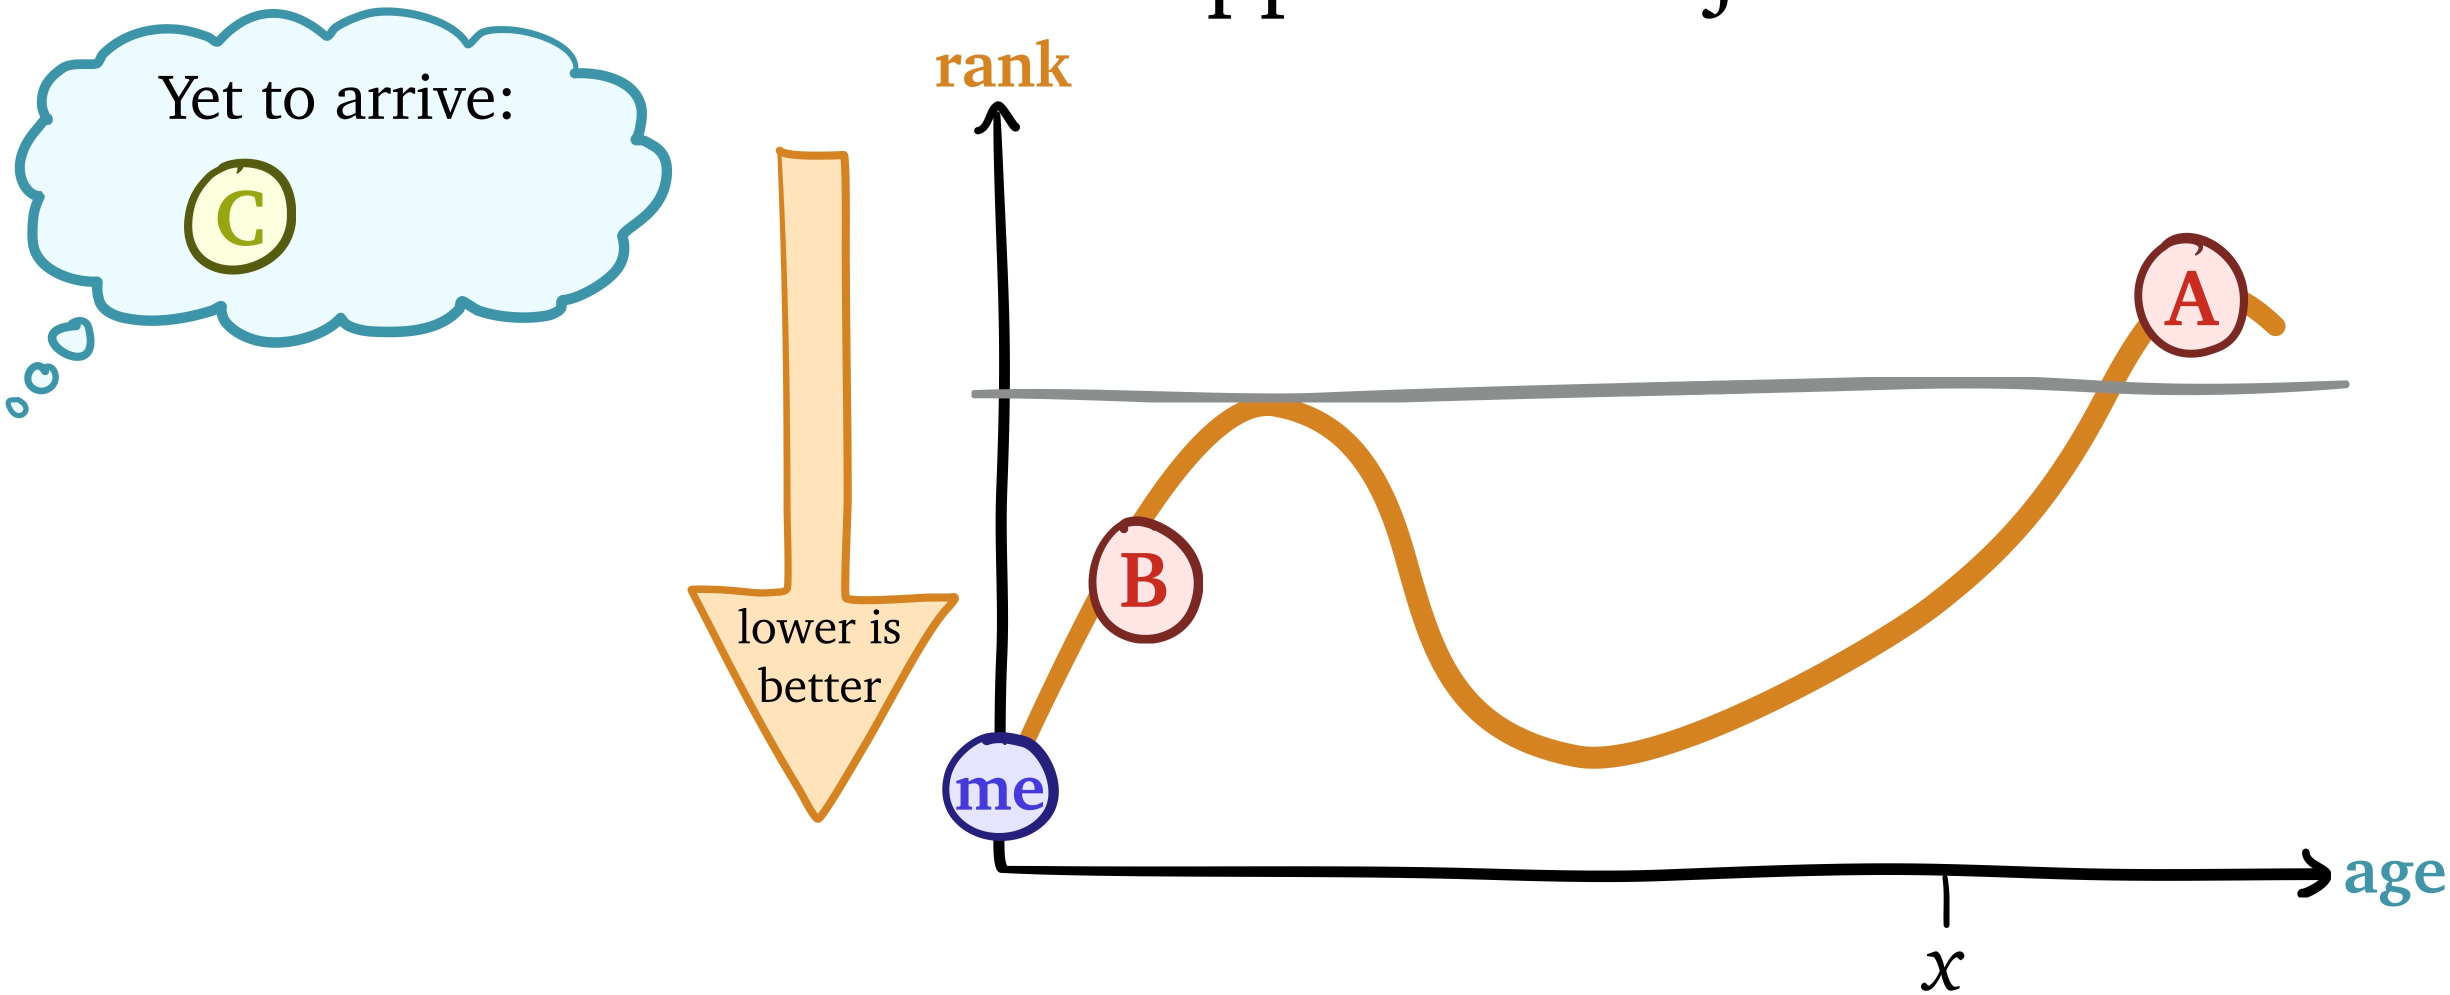

Analyzing **Gittins-1**

Suppose I'm a job of size x

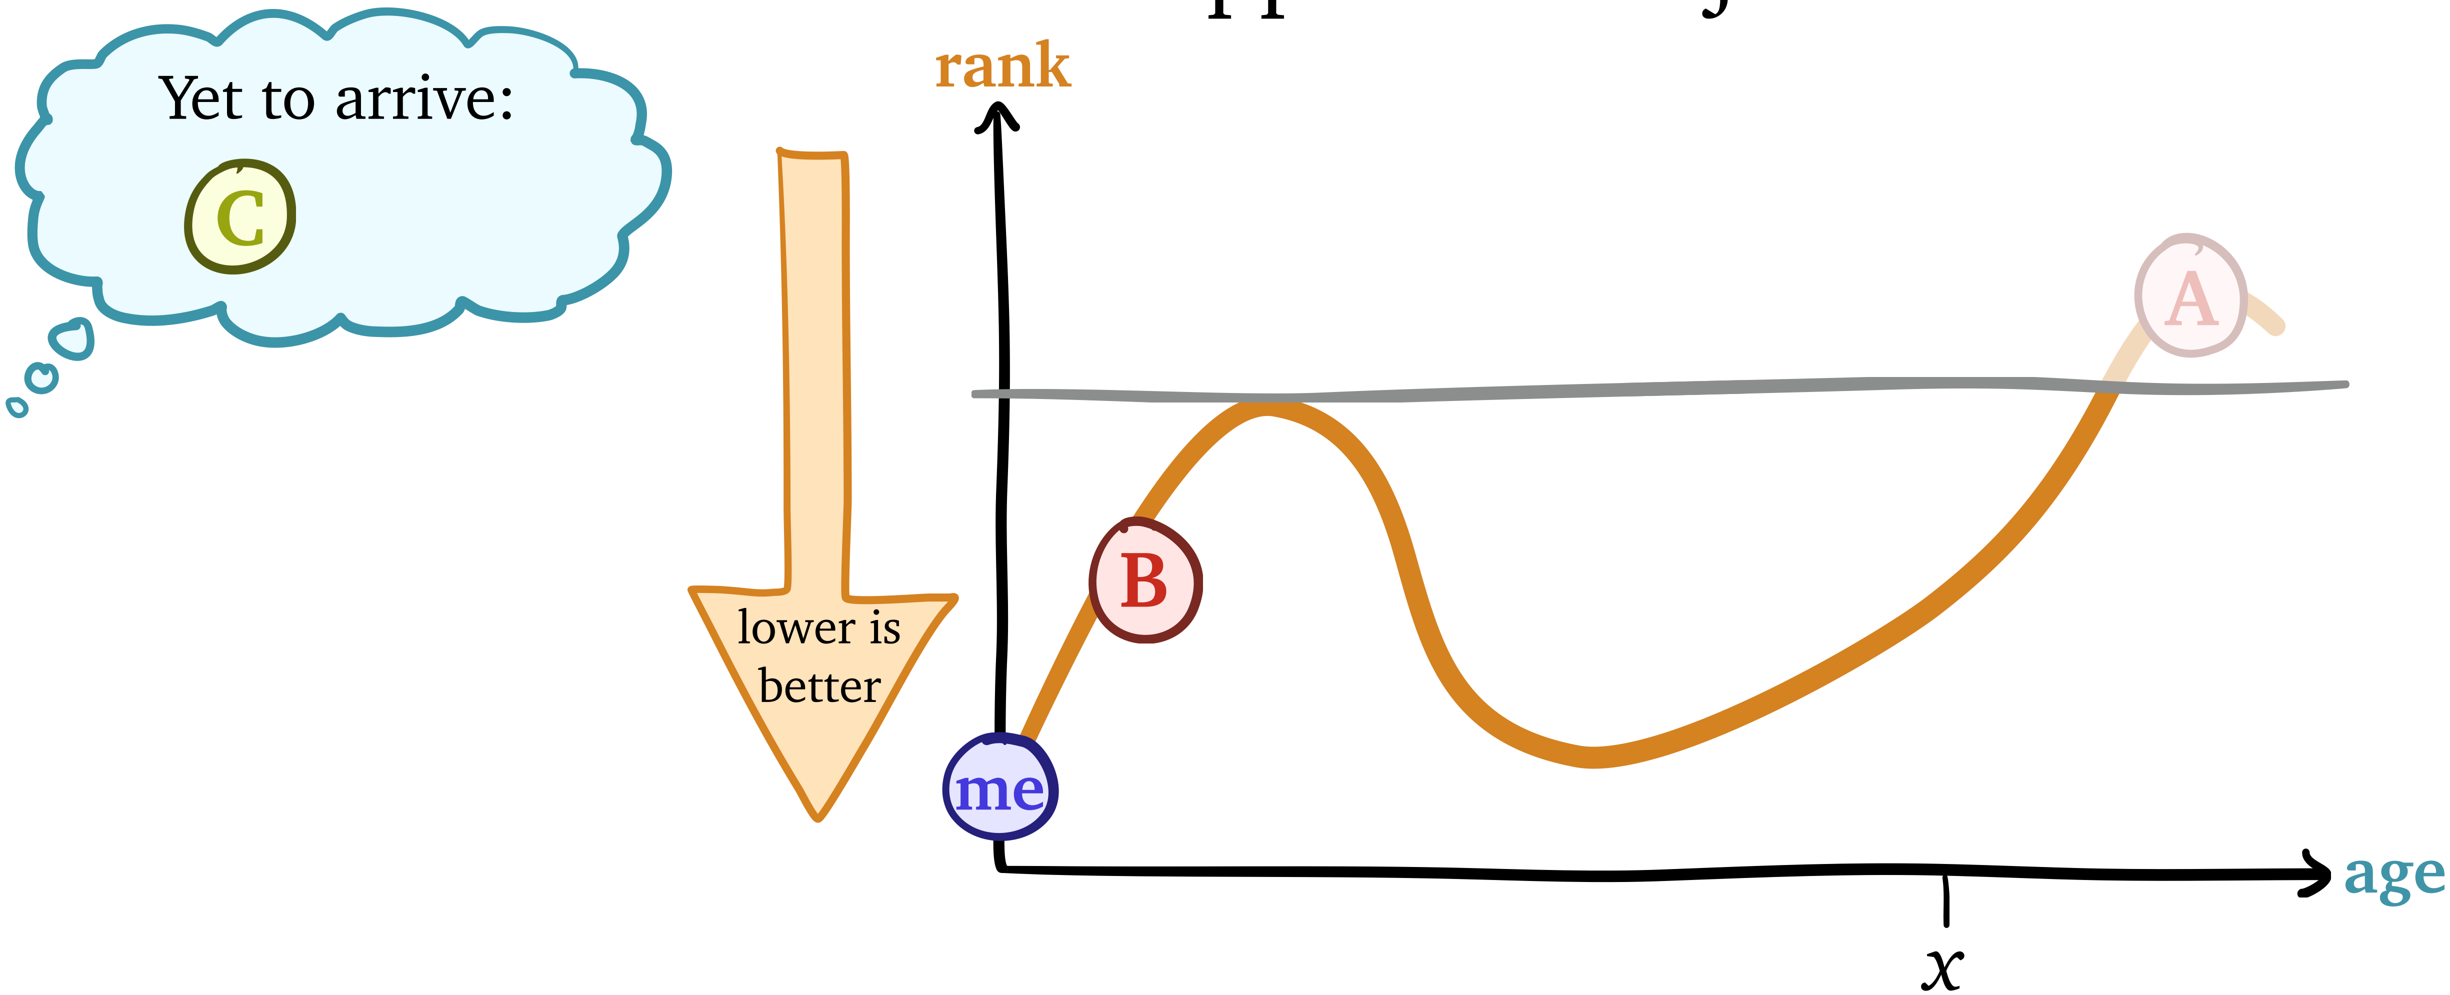

Analyzing **Gittins-1**

Suppose I'm a job of size x

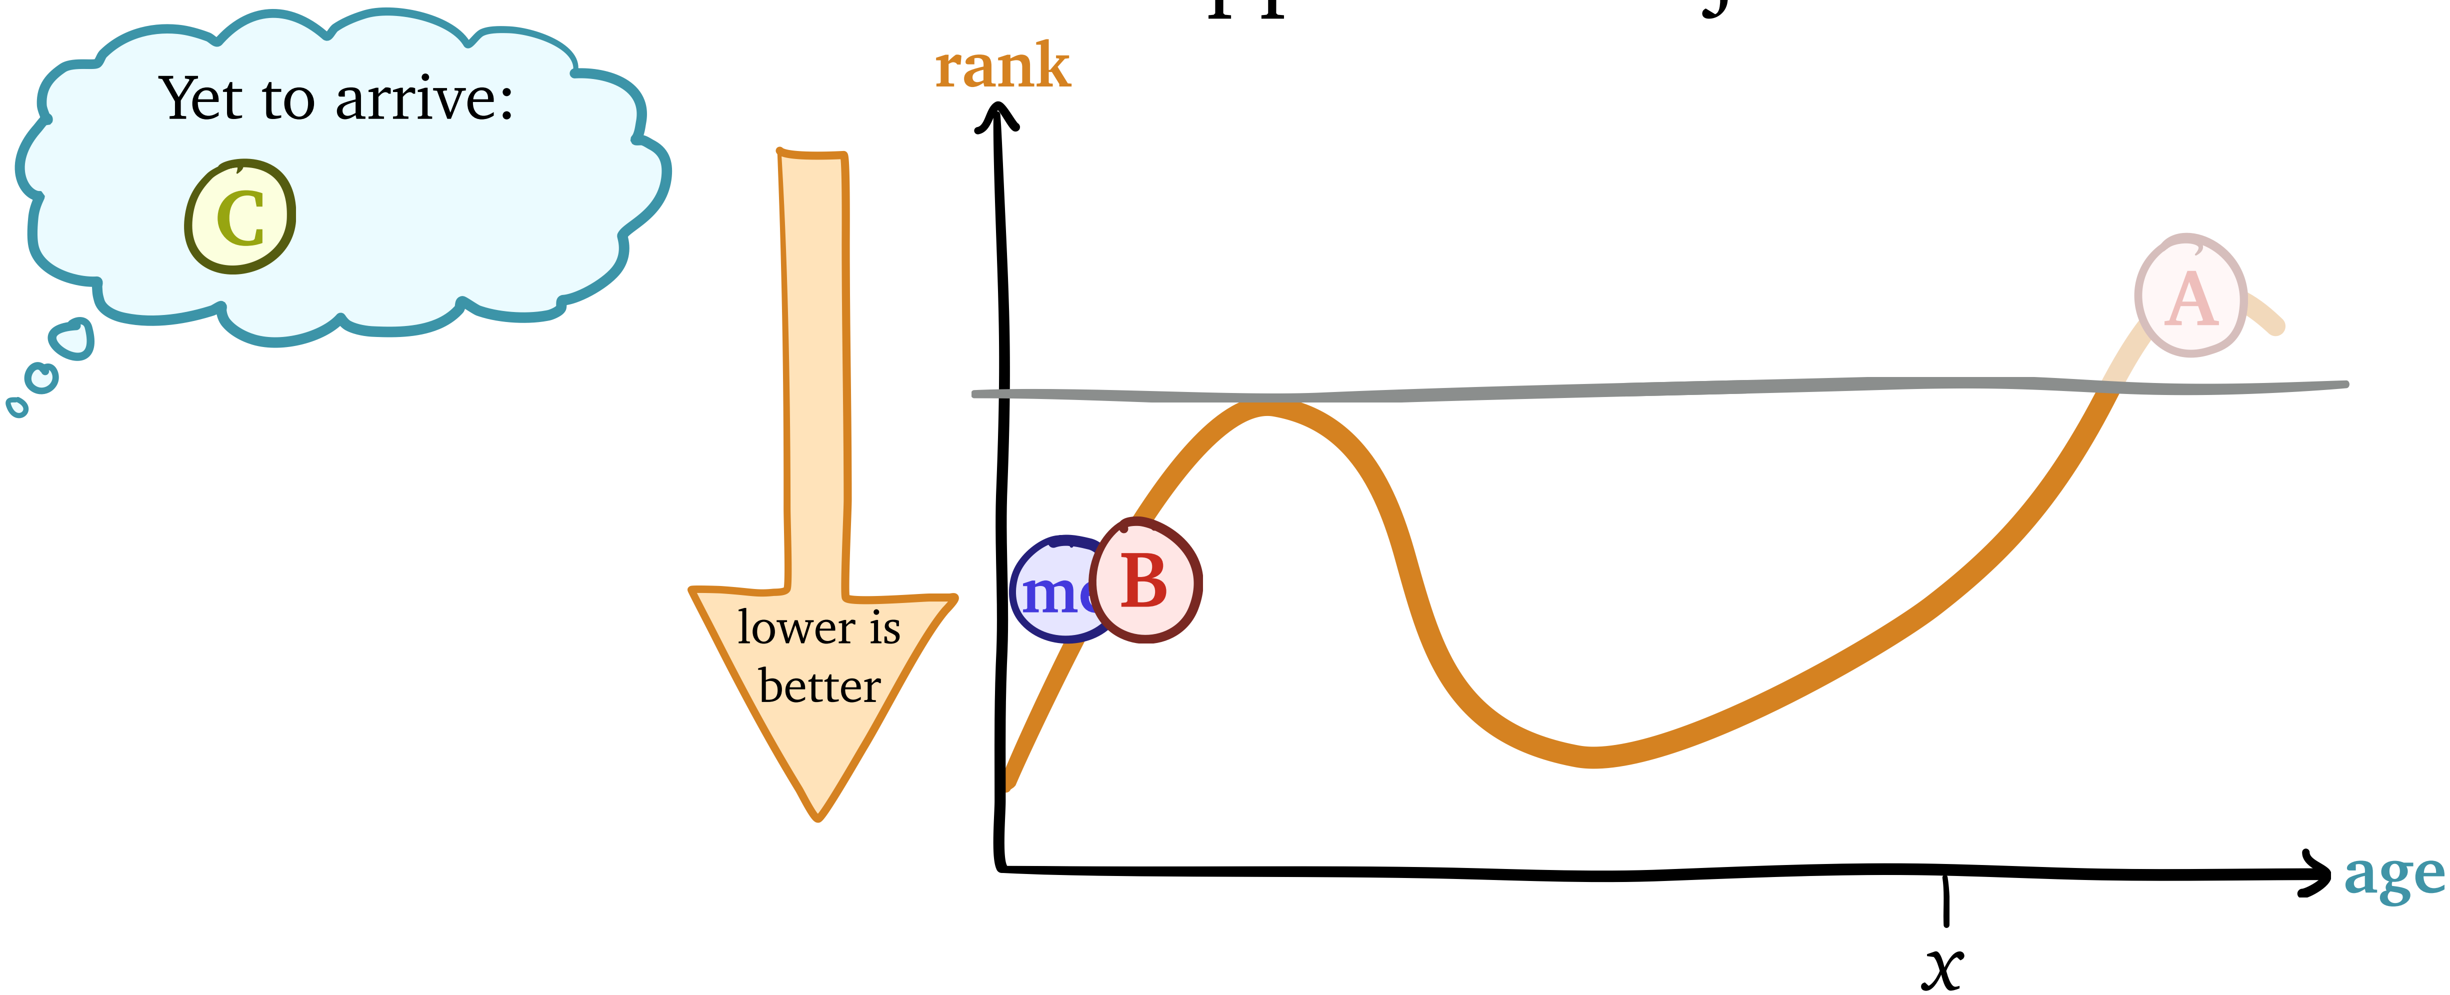

Analyzing **Gittins-1**

Suppose I'm a job of size x

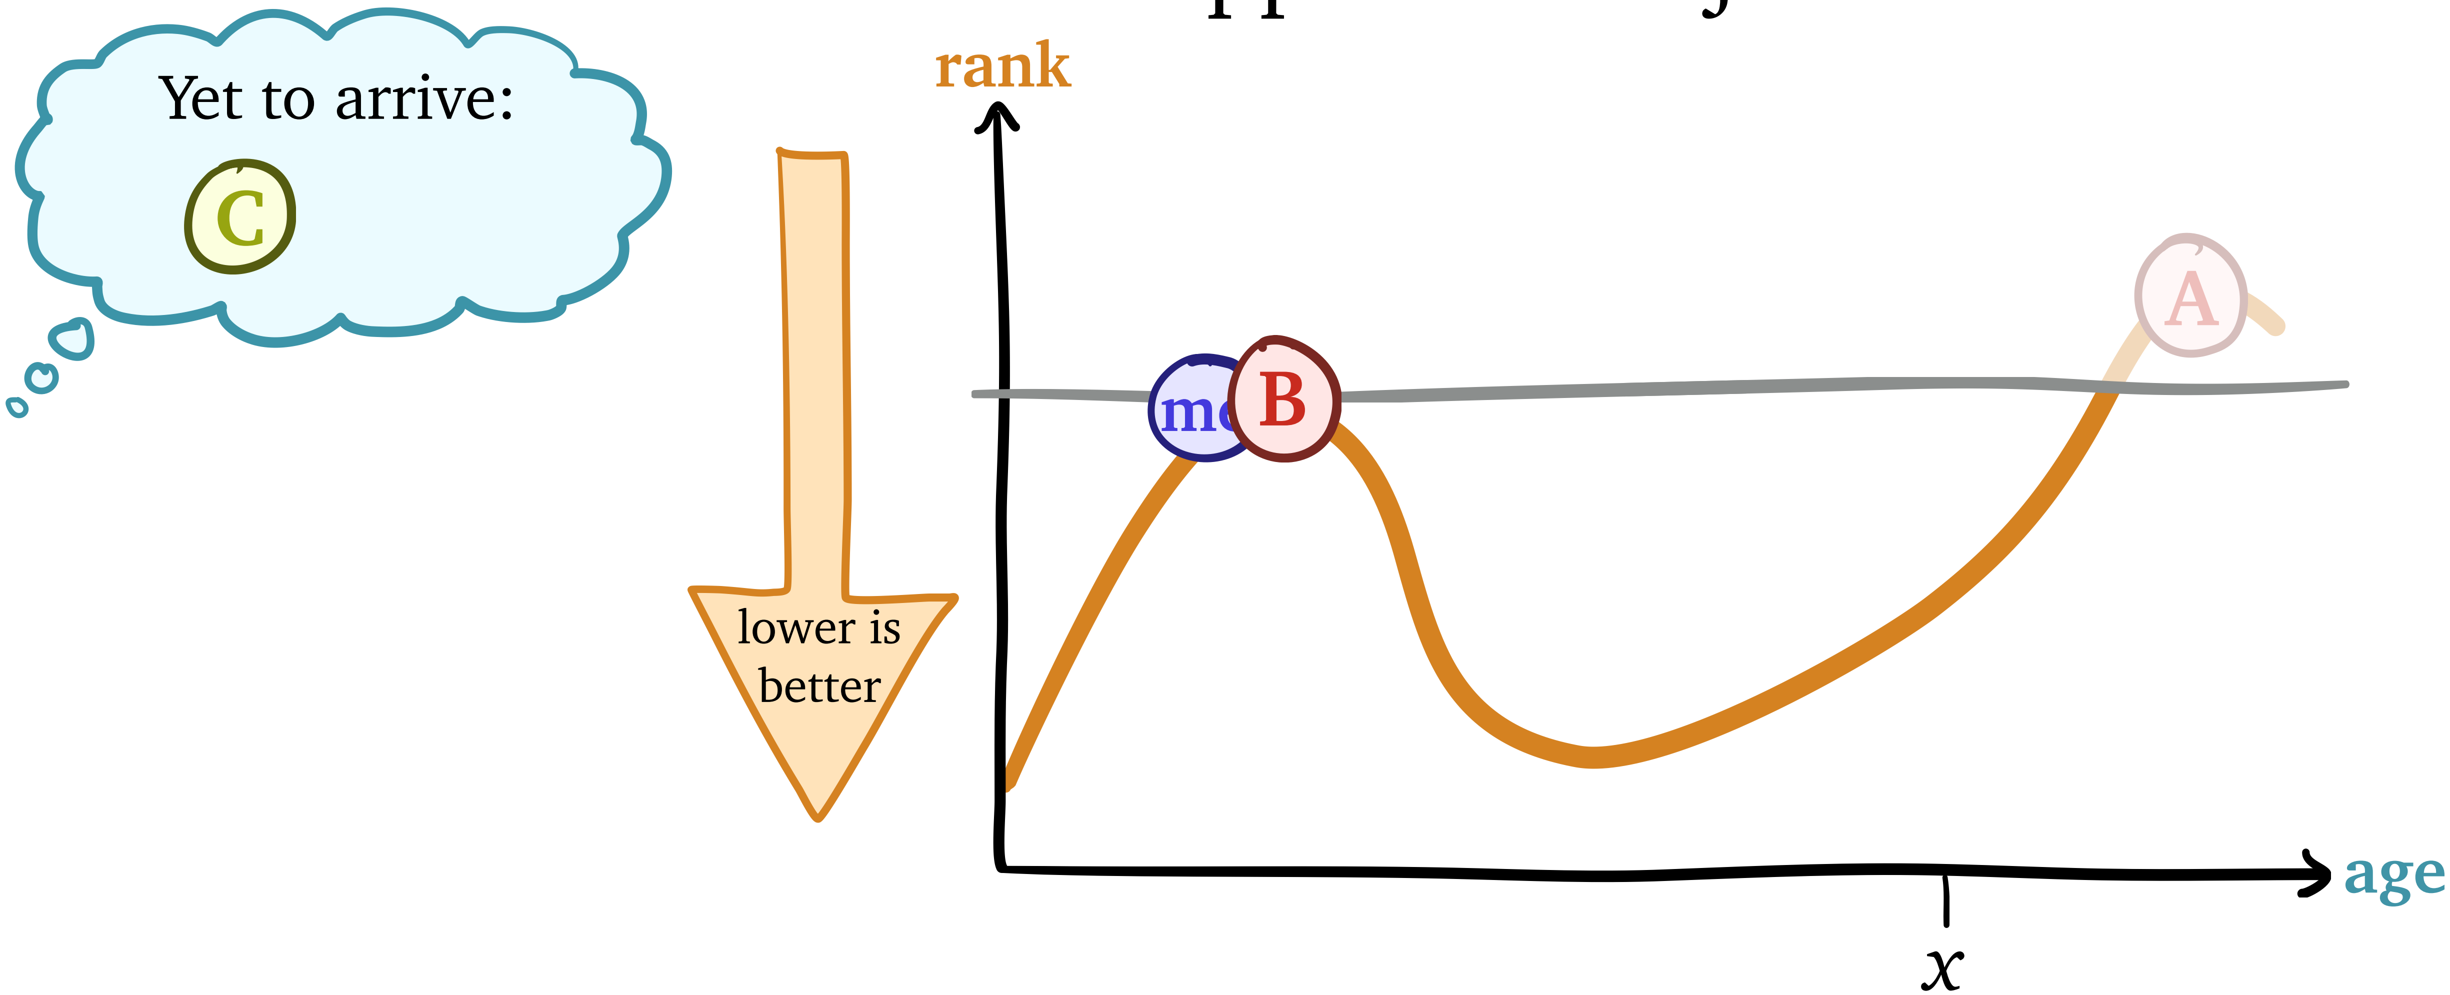

Analyzing **Gittins-1**

Suppose I'm a job of size x

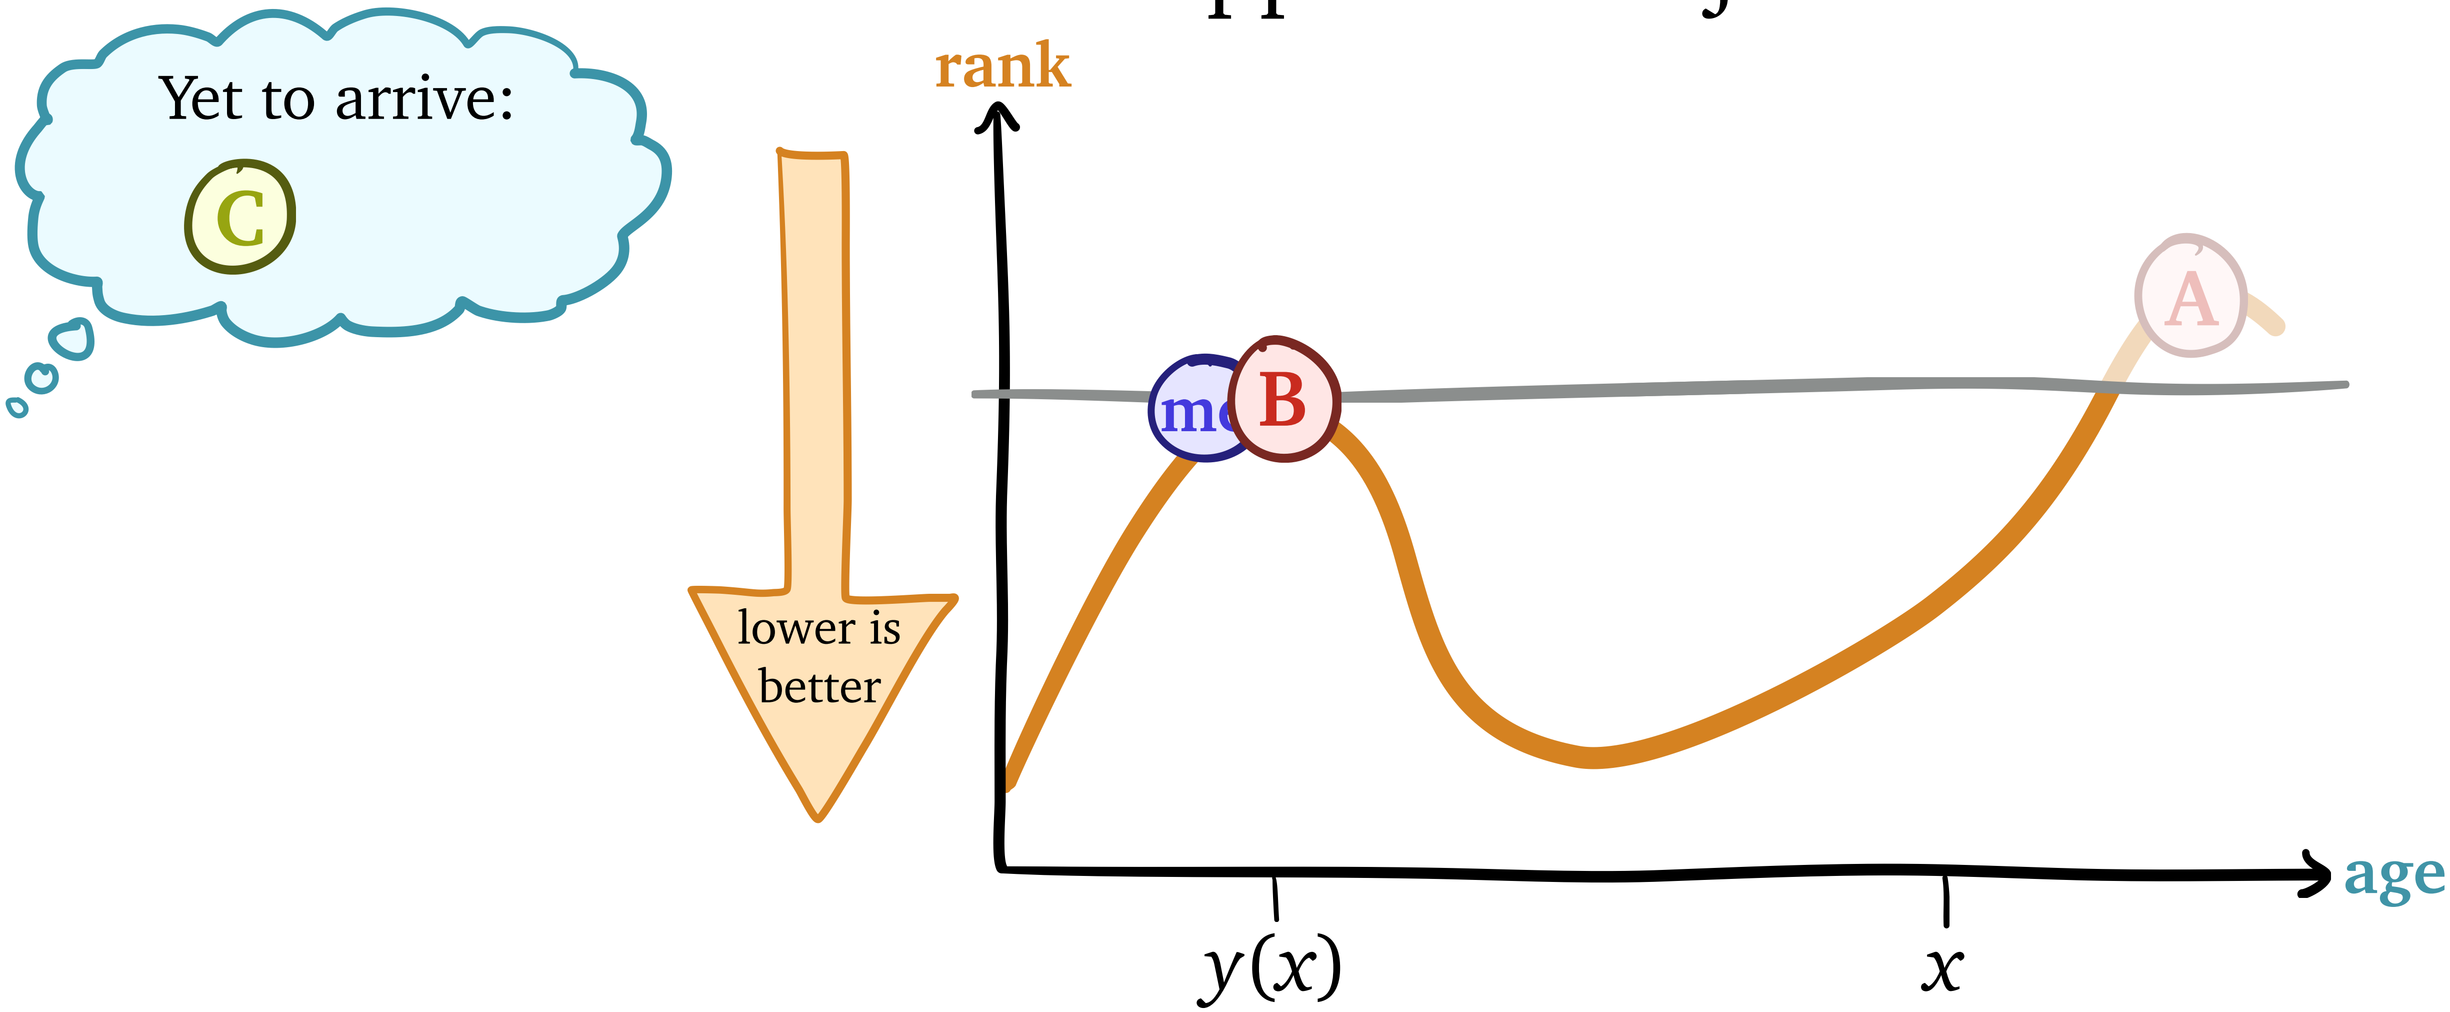

Analyzing **Gittins-1**

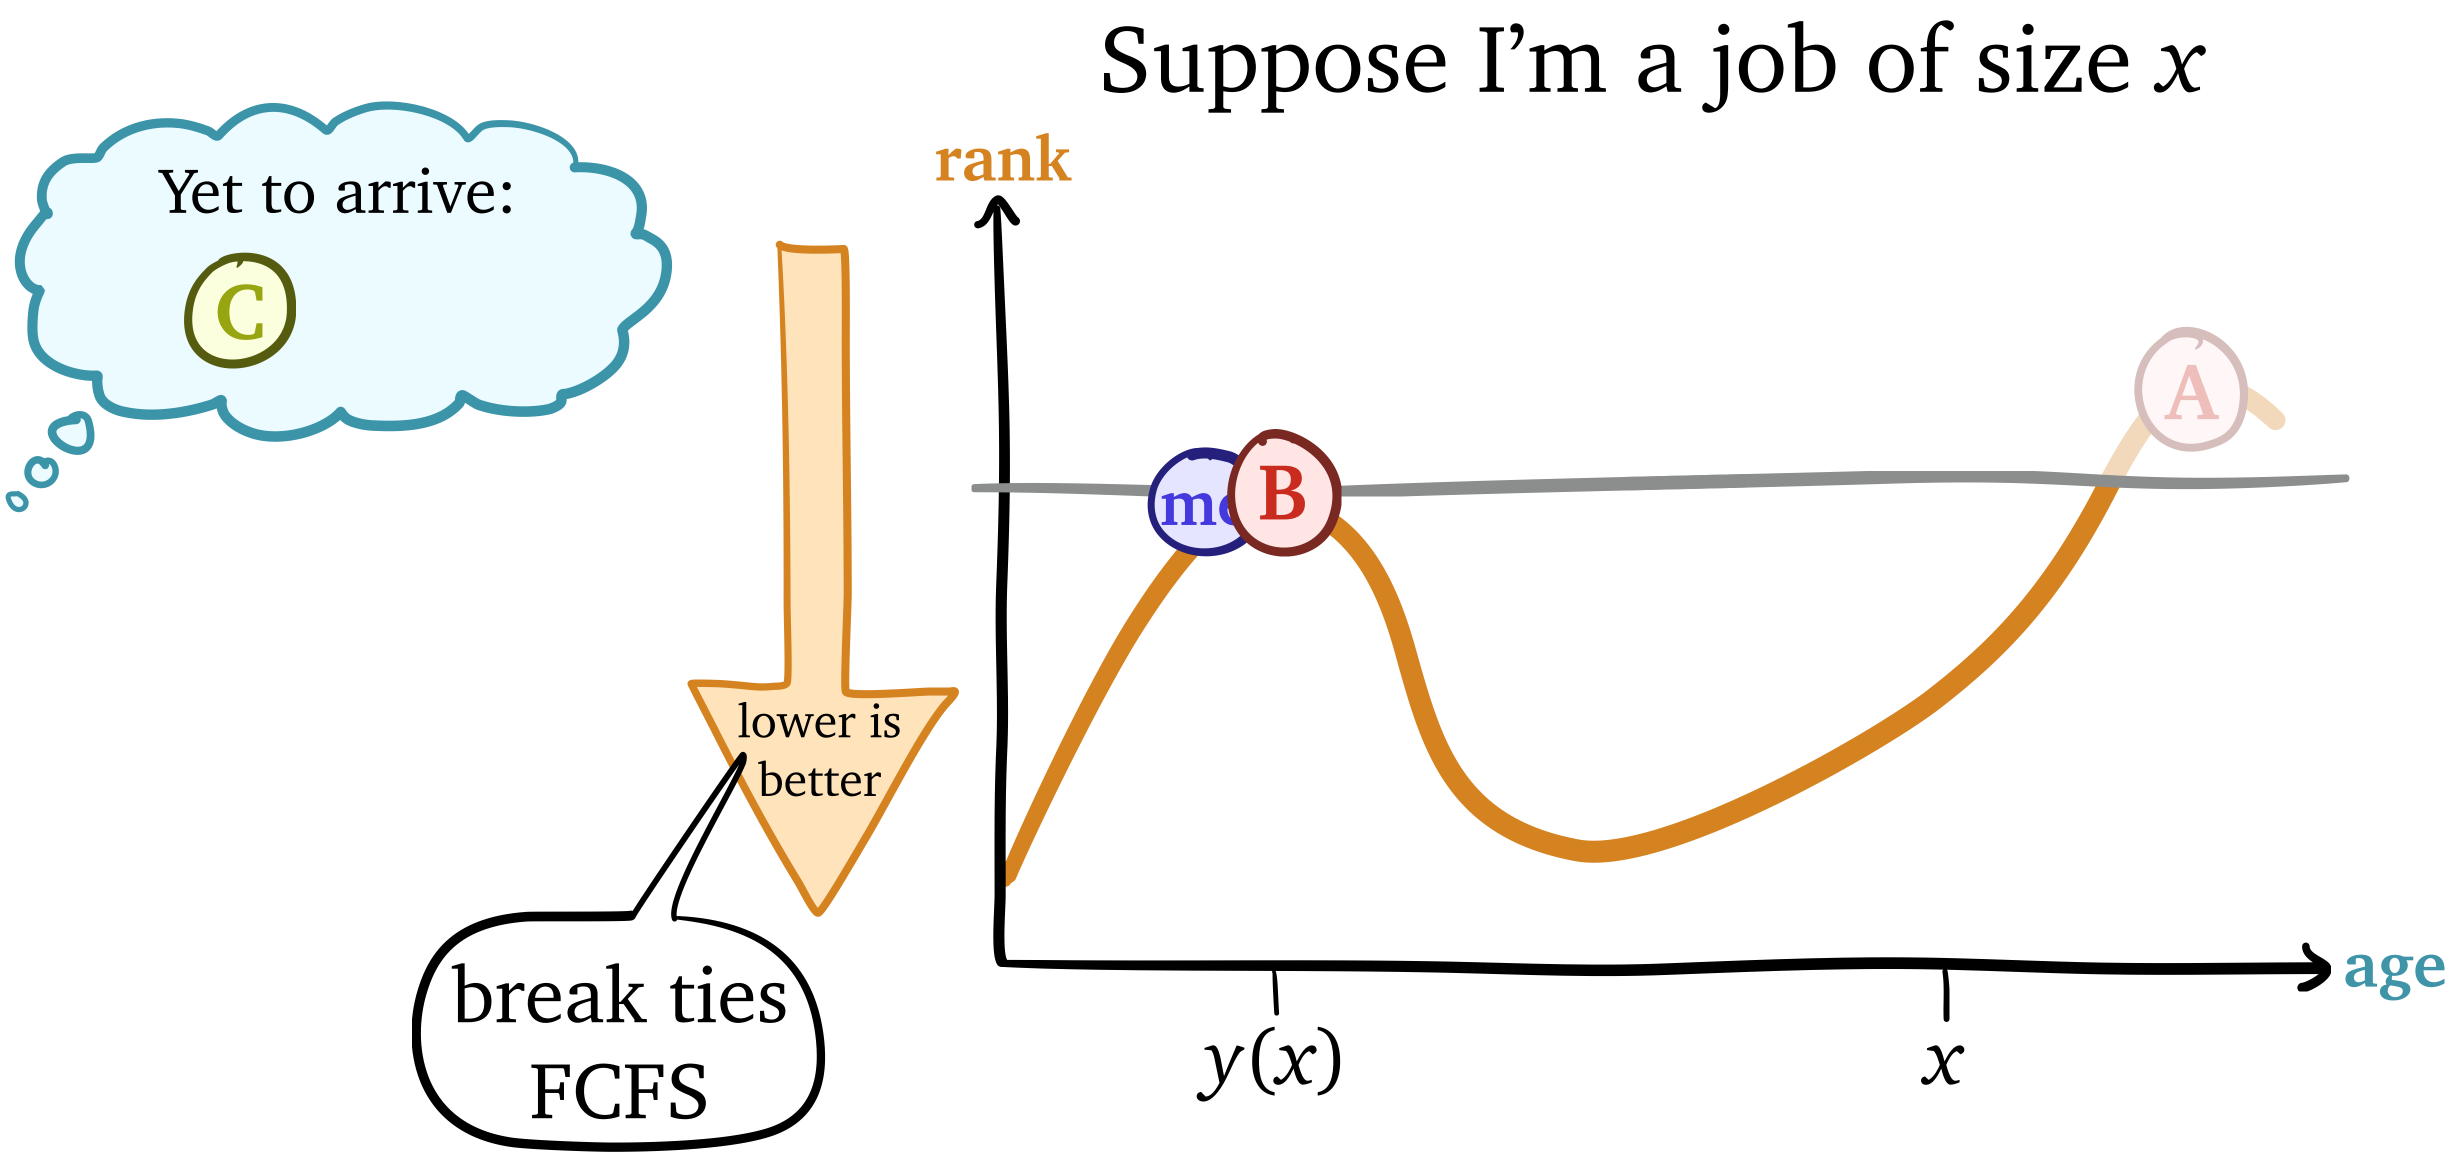

Analyzing **Gittins-1**

Suppose I'm a job of size x

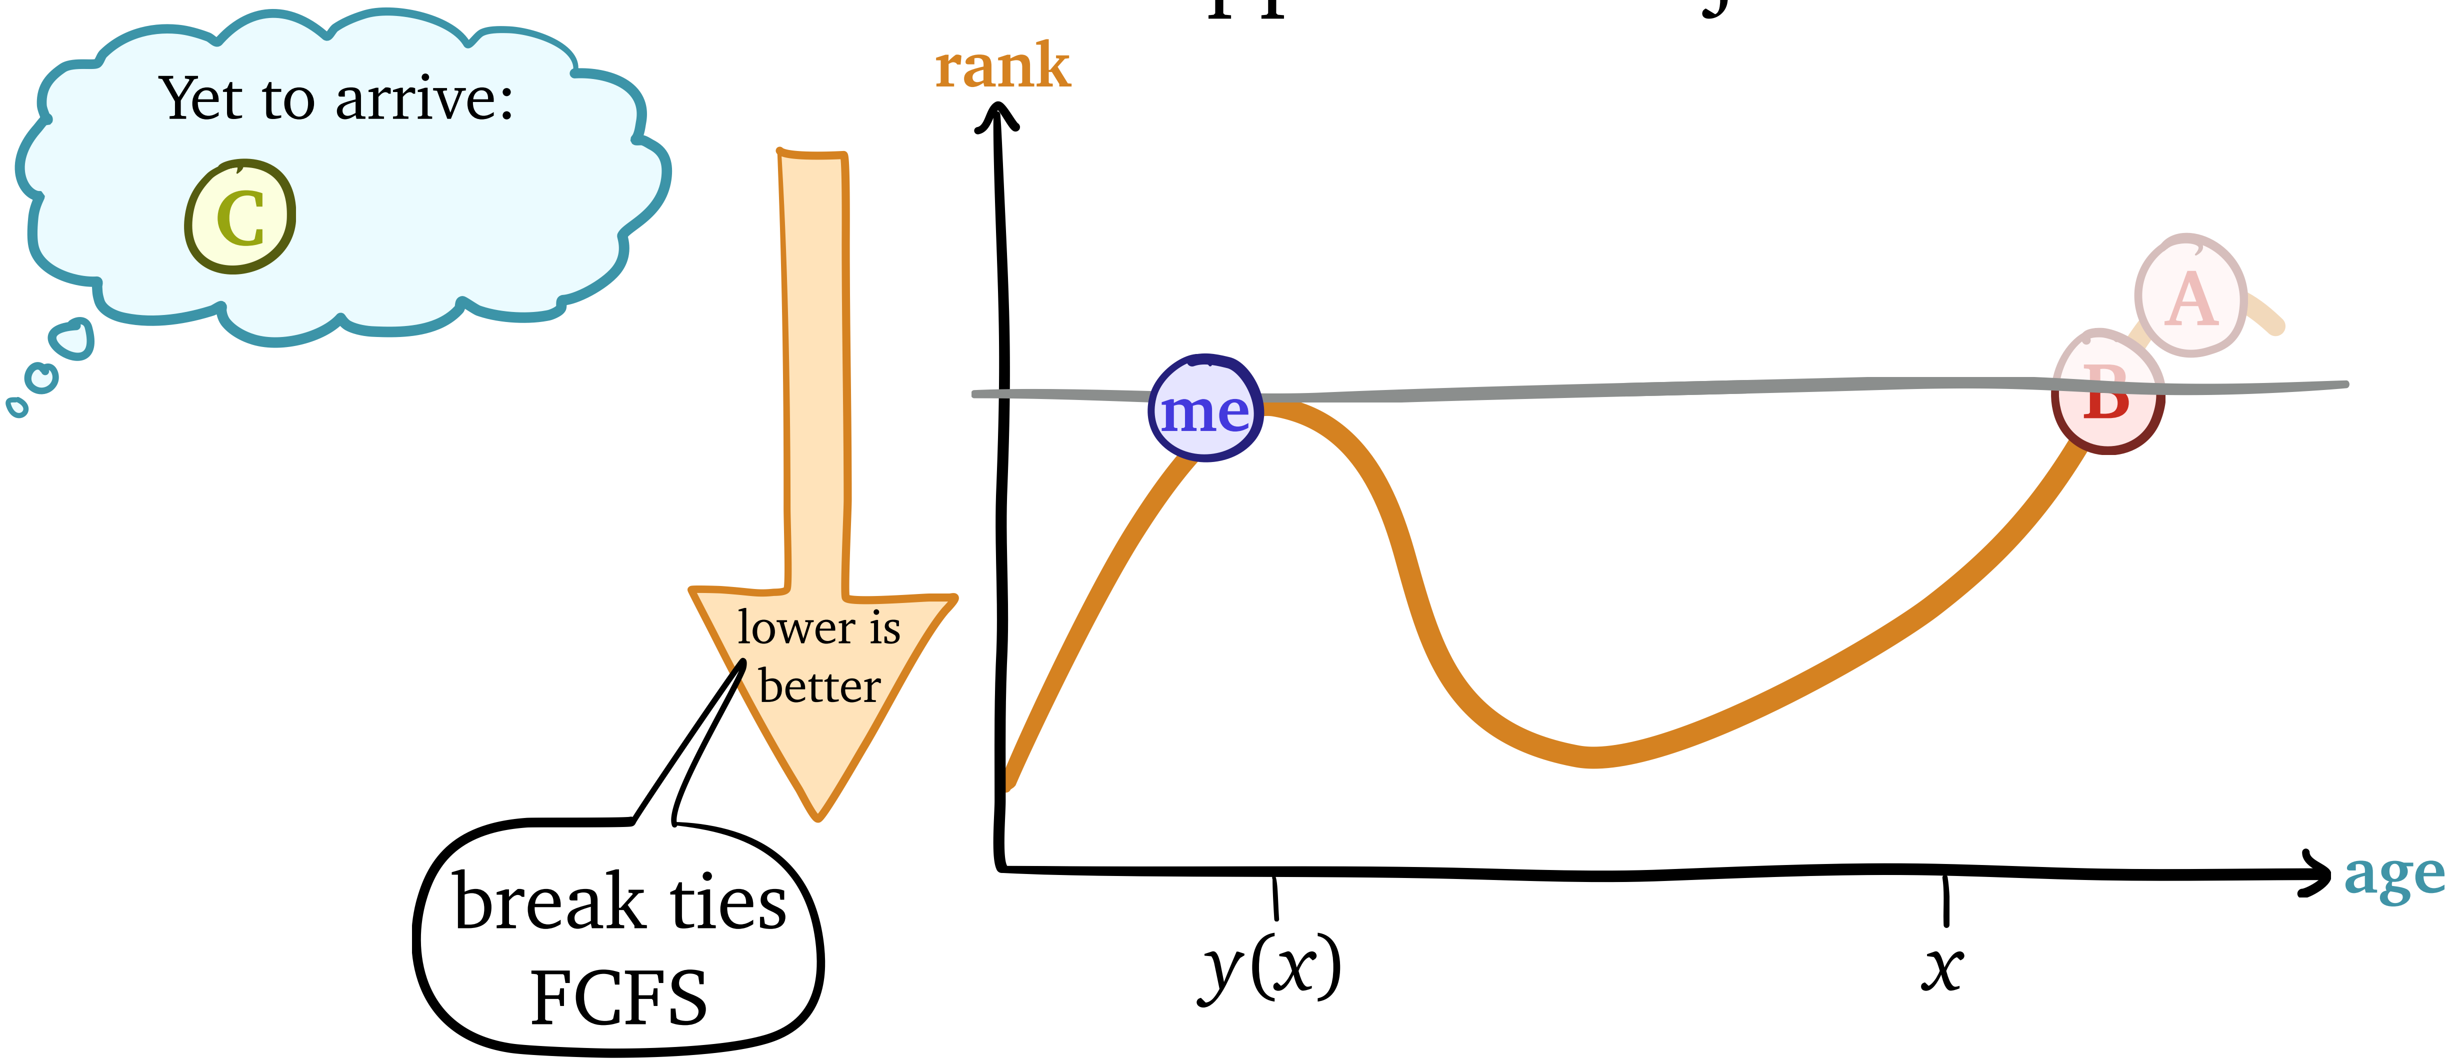

Analyzing **Gittins-1**

Suppose I'm a job of size x

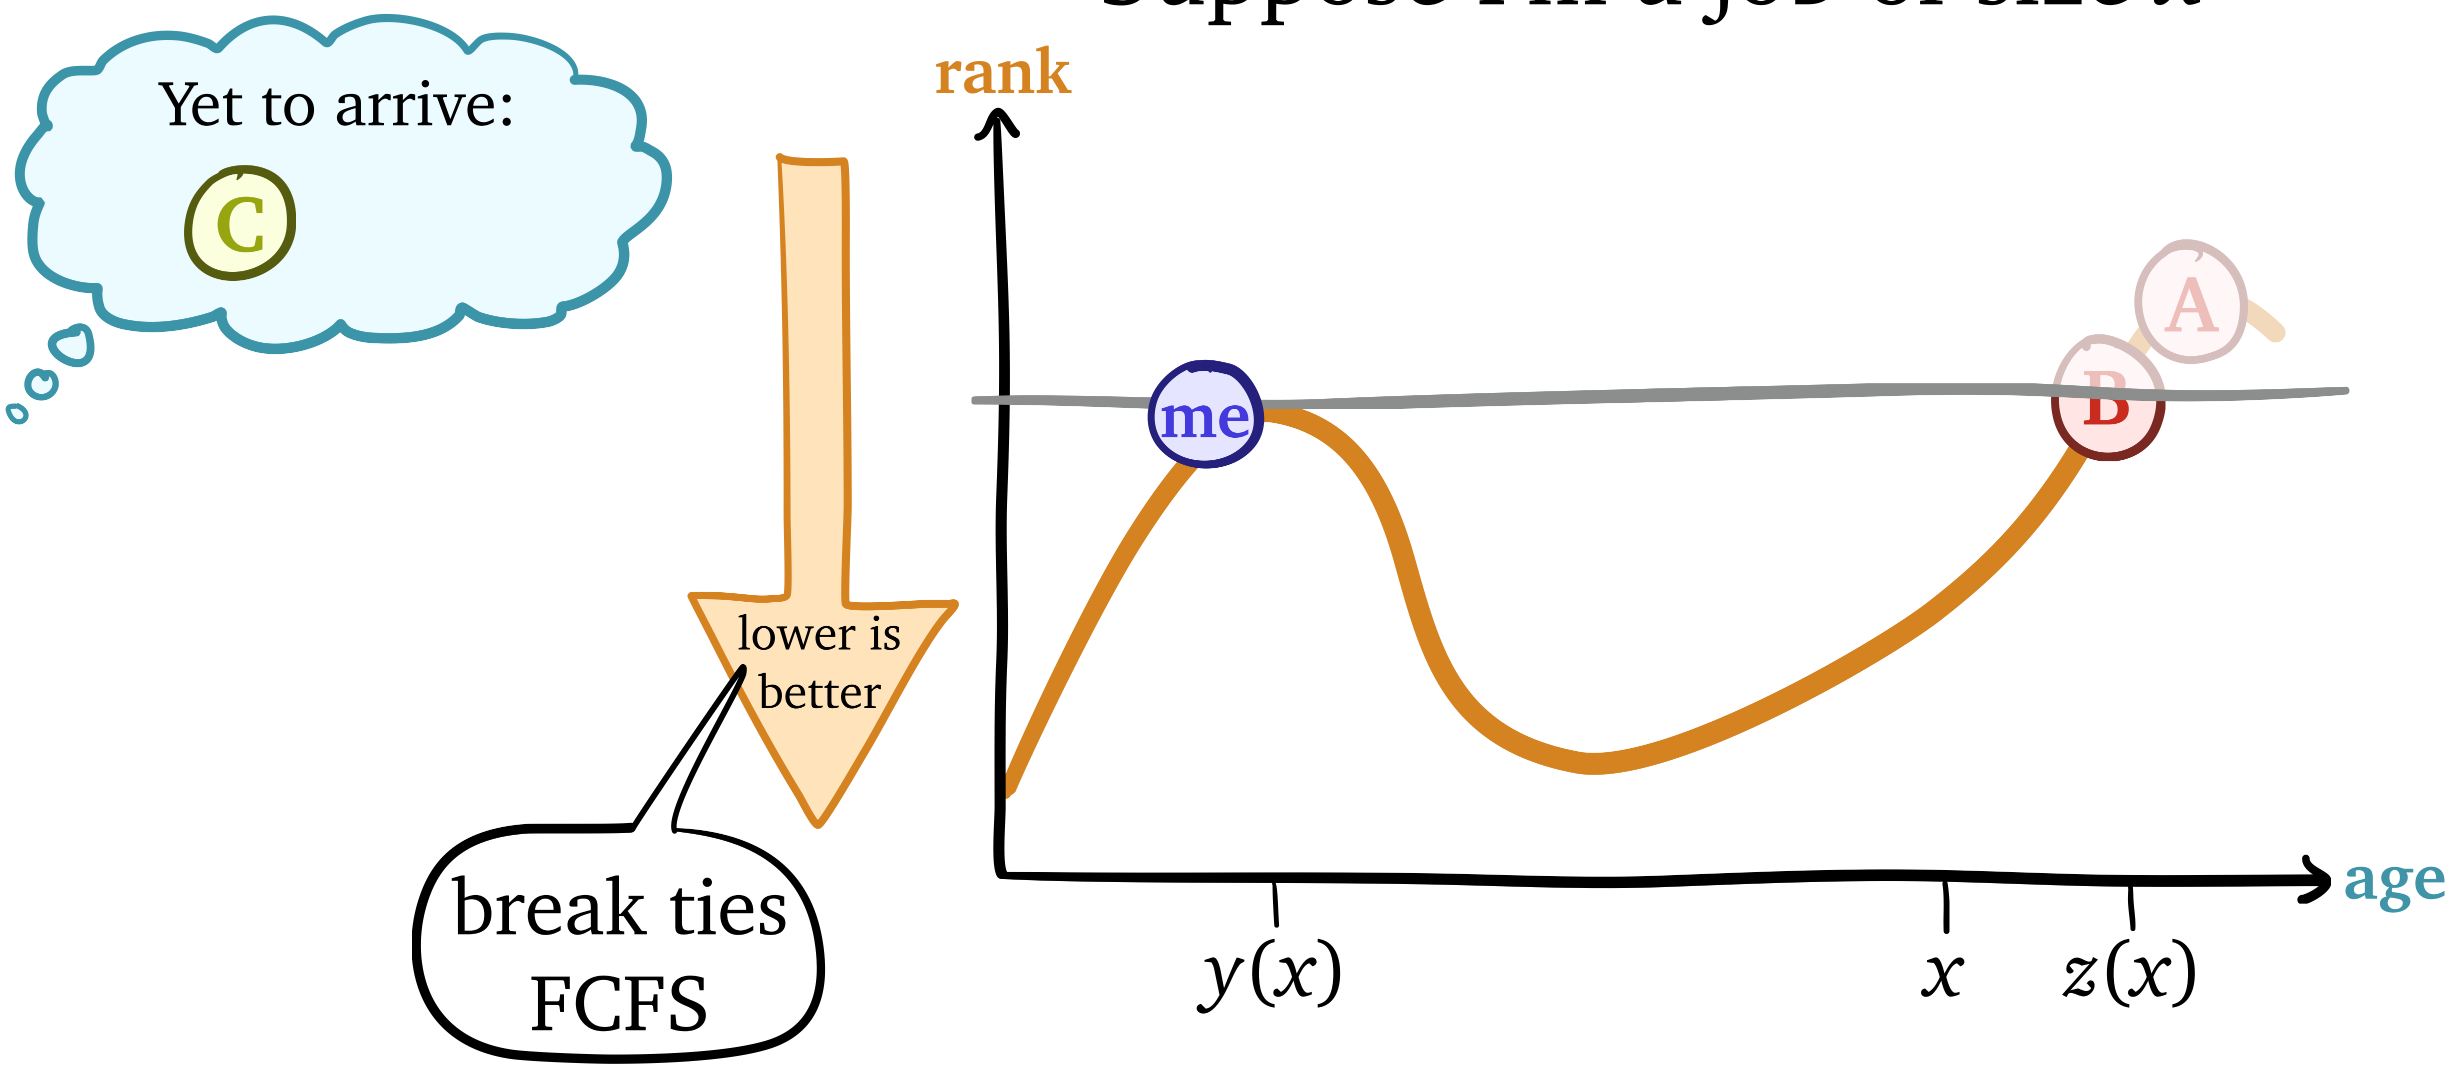

Analyzing **Gittins-1**

Suppose I'm a job of size x

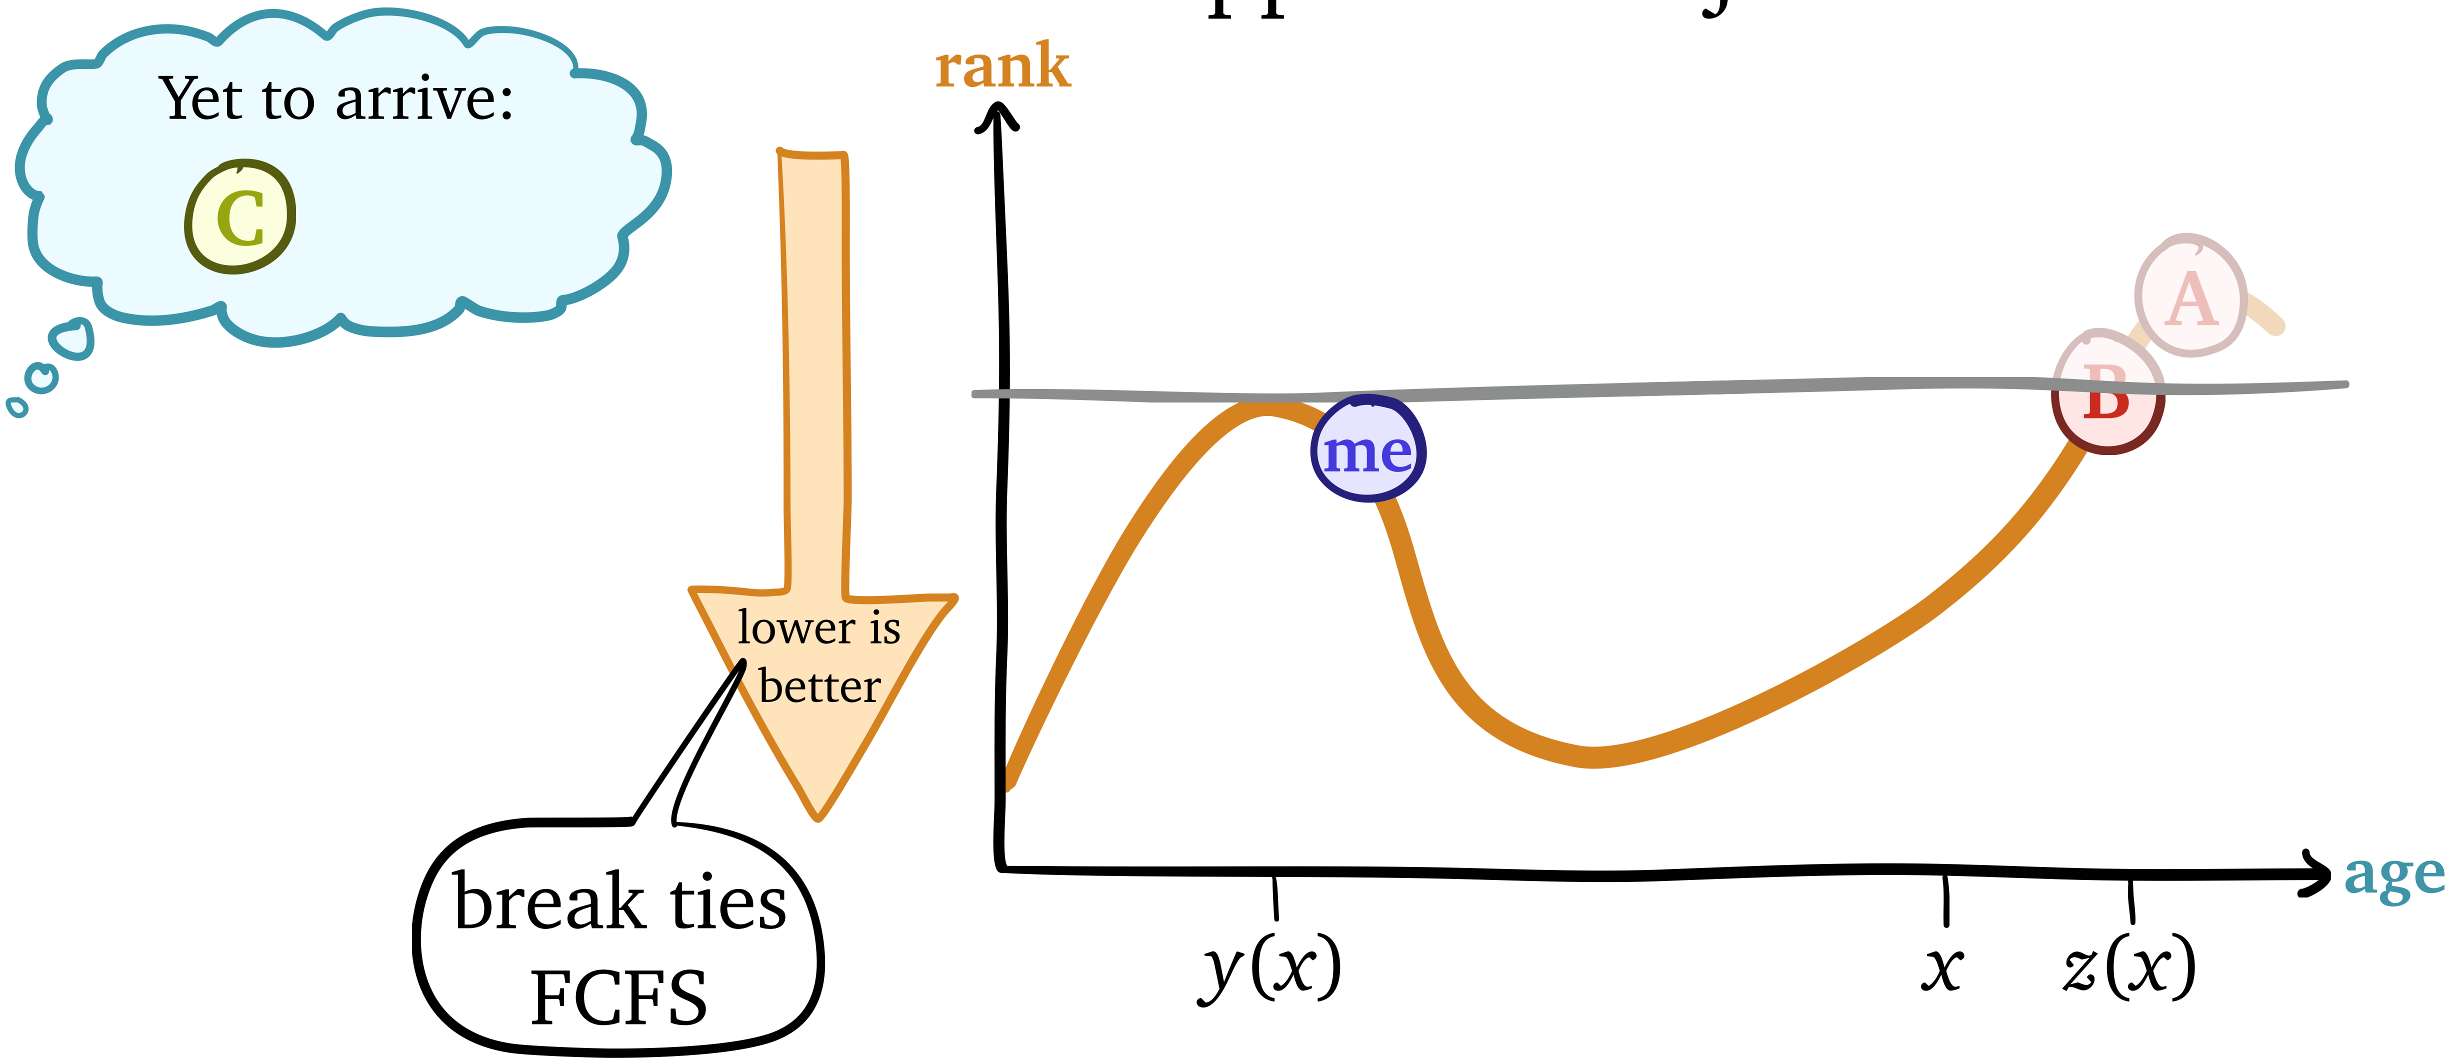

Analyzing **Gittins-1**

Suppose I'm a job of size x

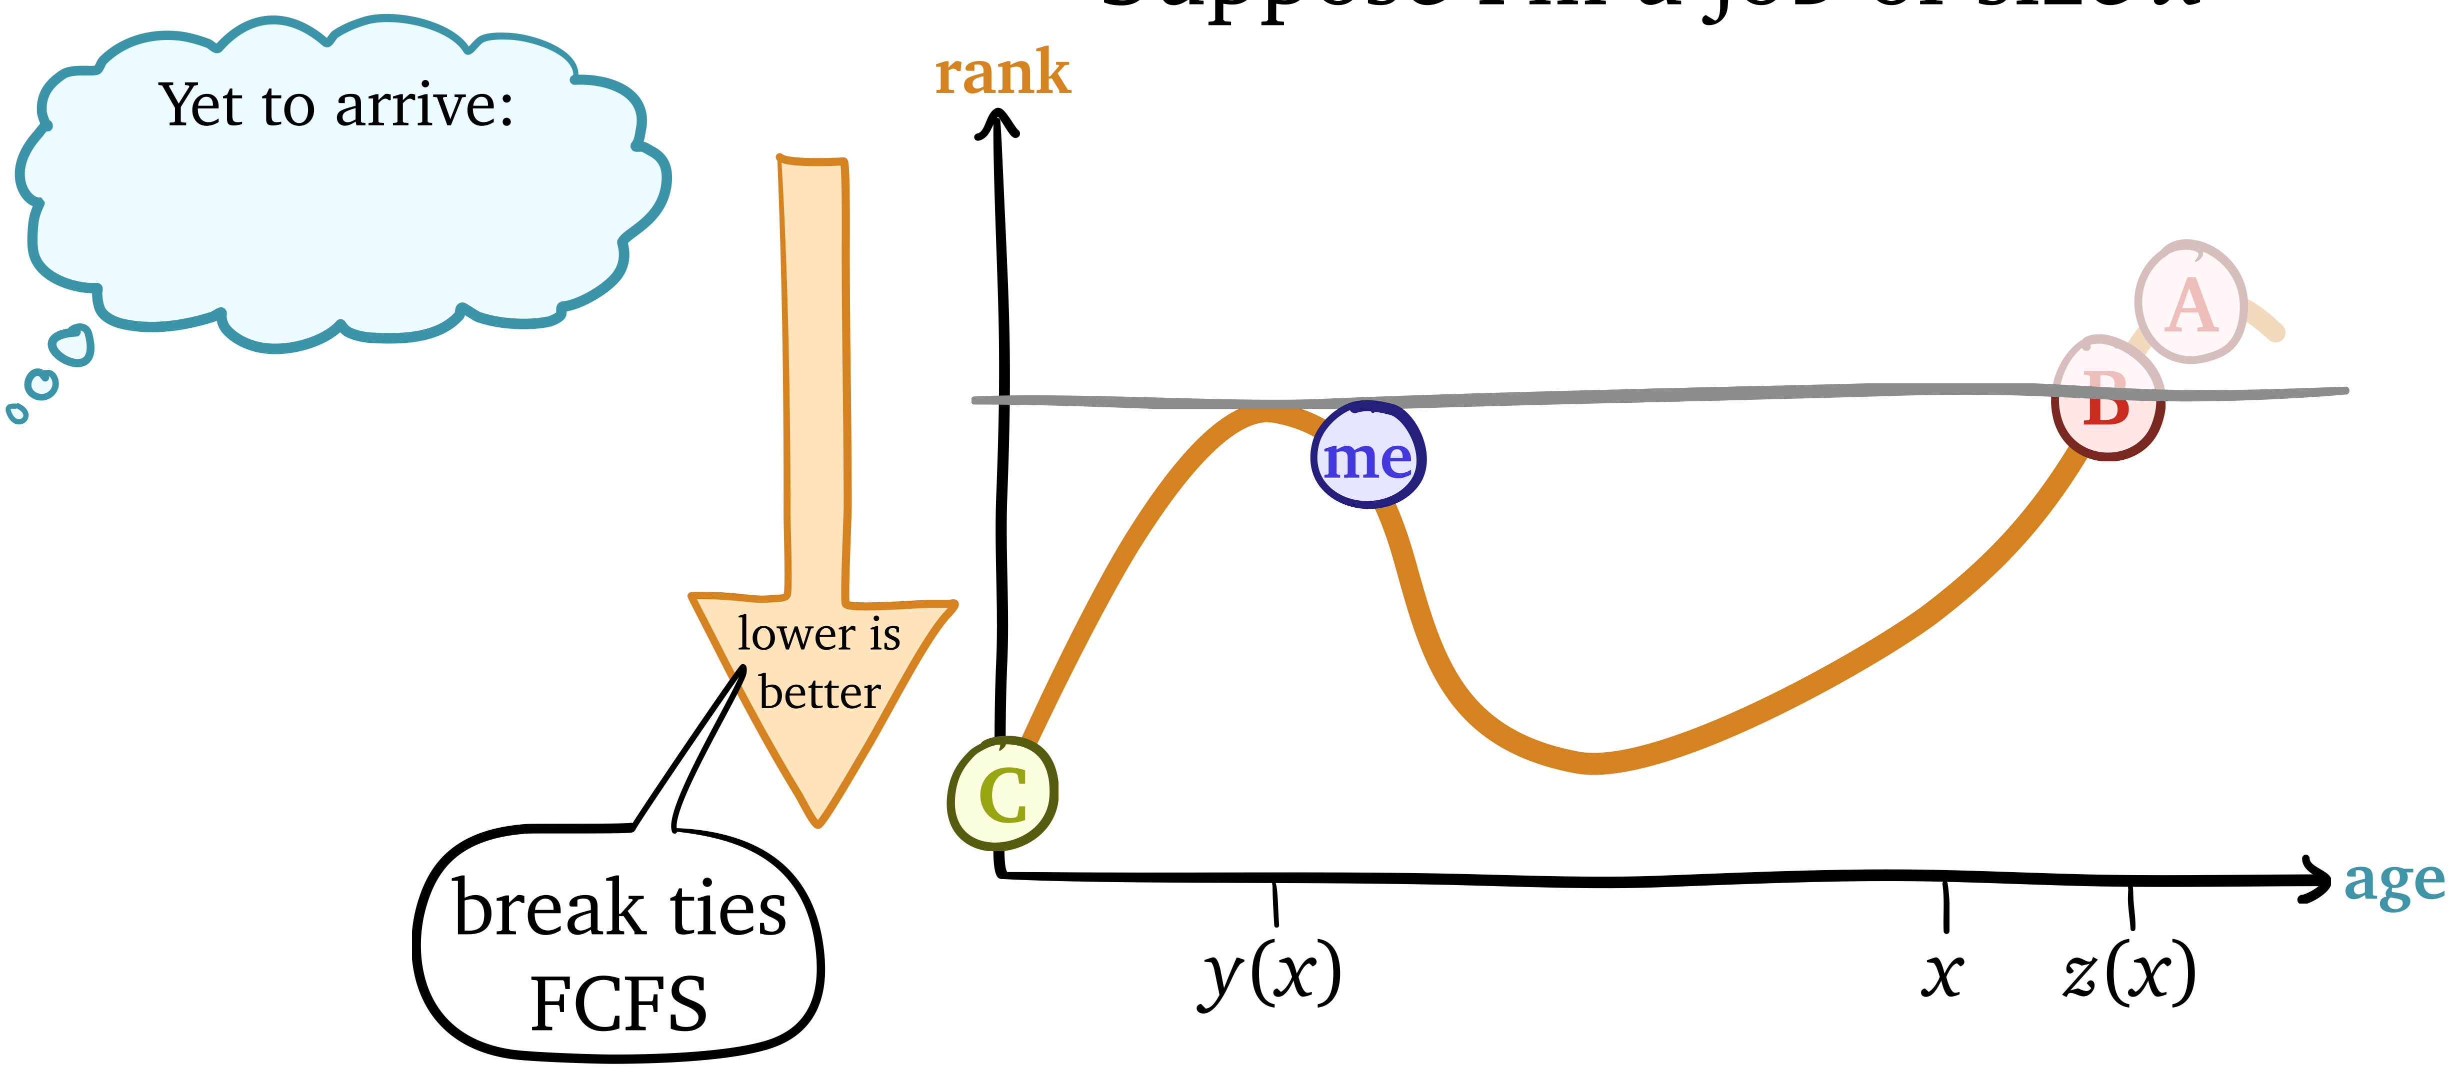

Analyzing **Gittins-1**

Suppose I'm a job of size x

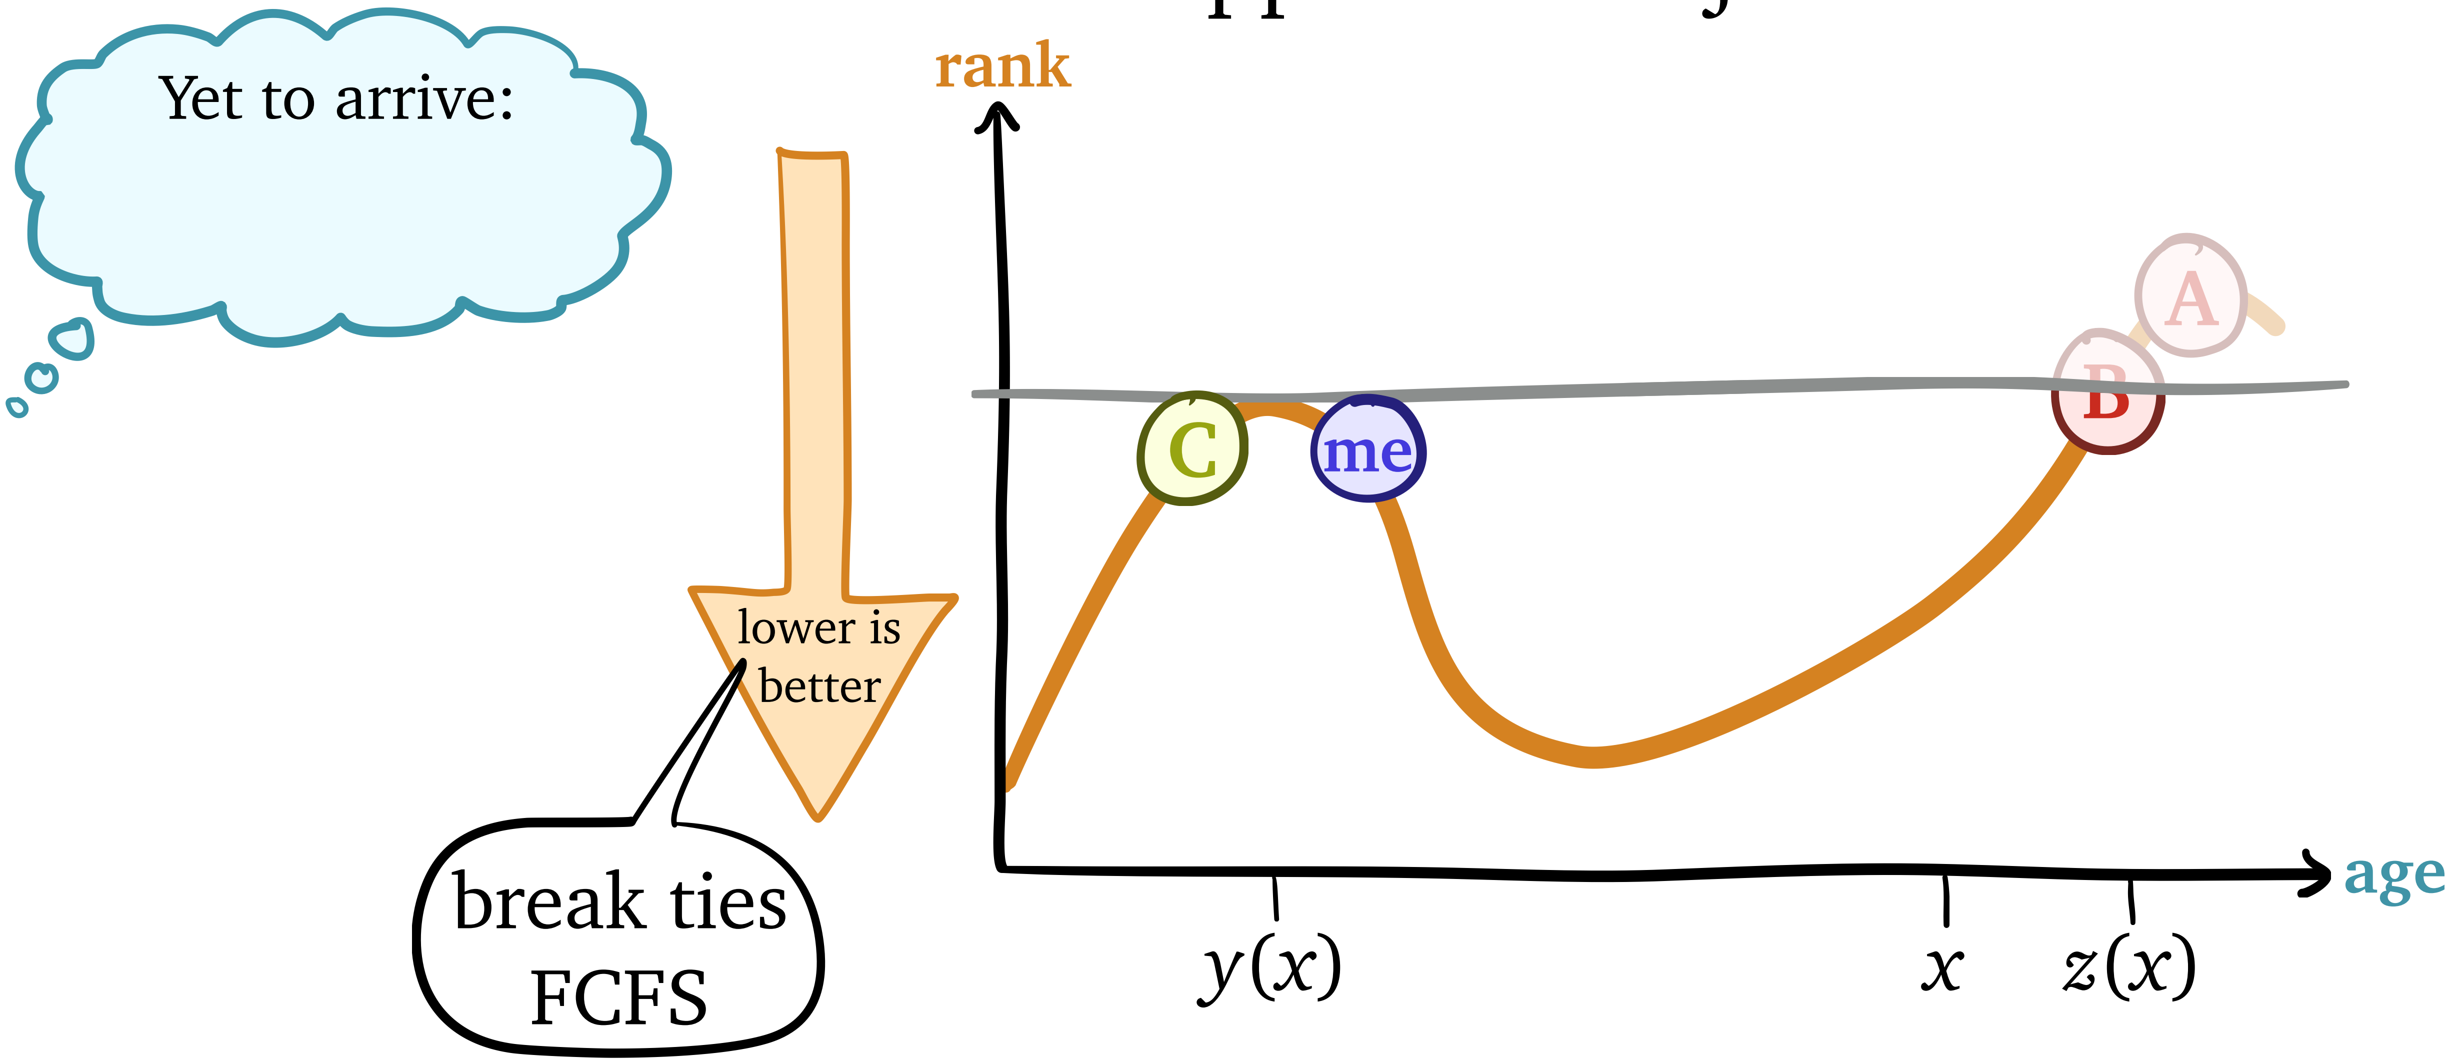

Analyzing **Gittins-1**

Suppose I'm a job of size x

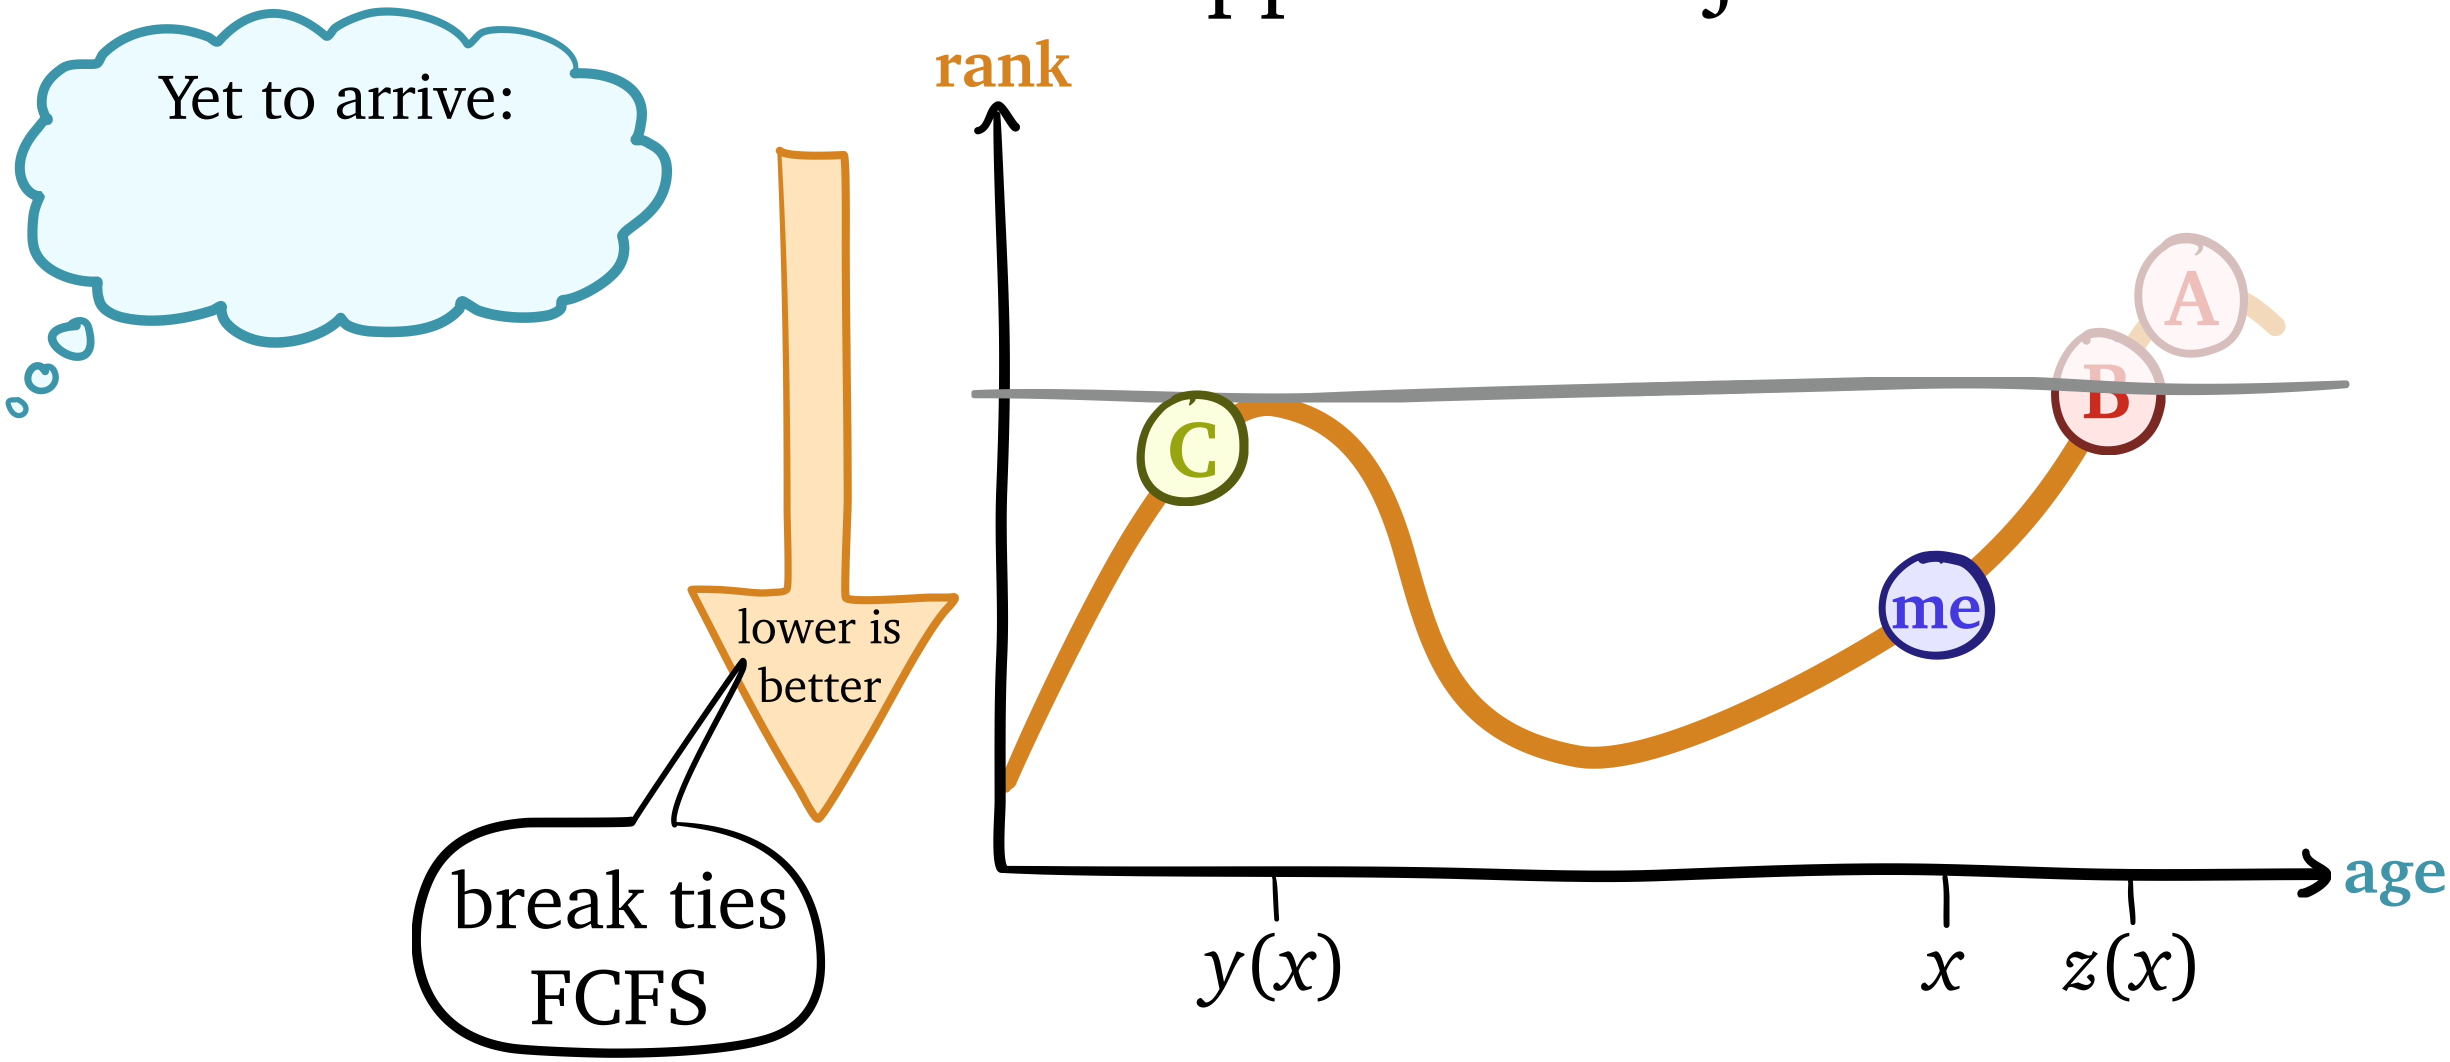

Analyzing **Gittins-1**

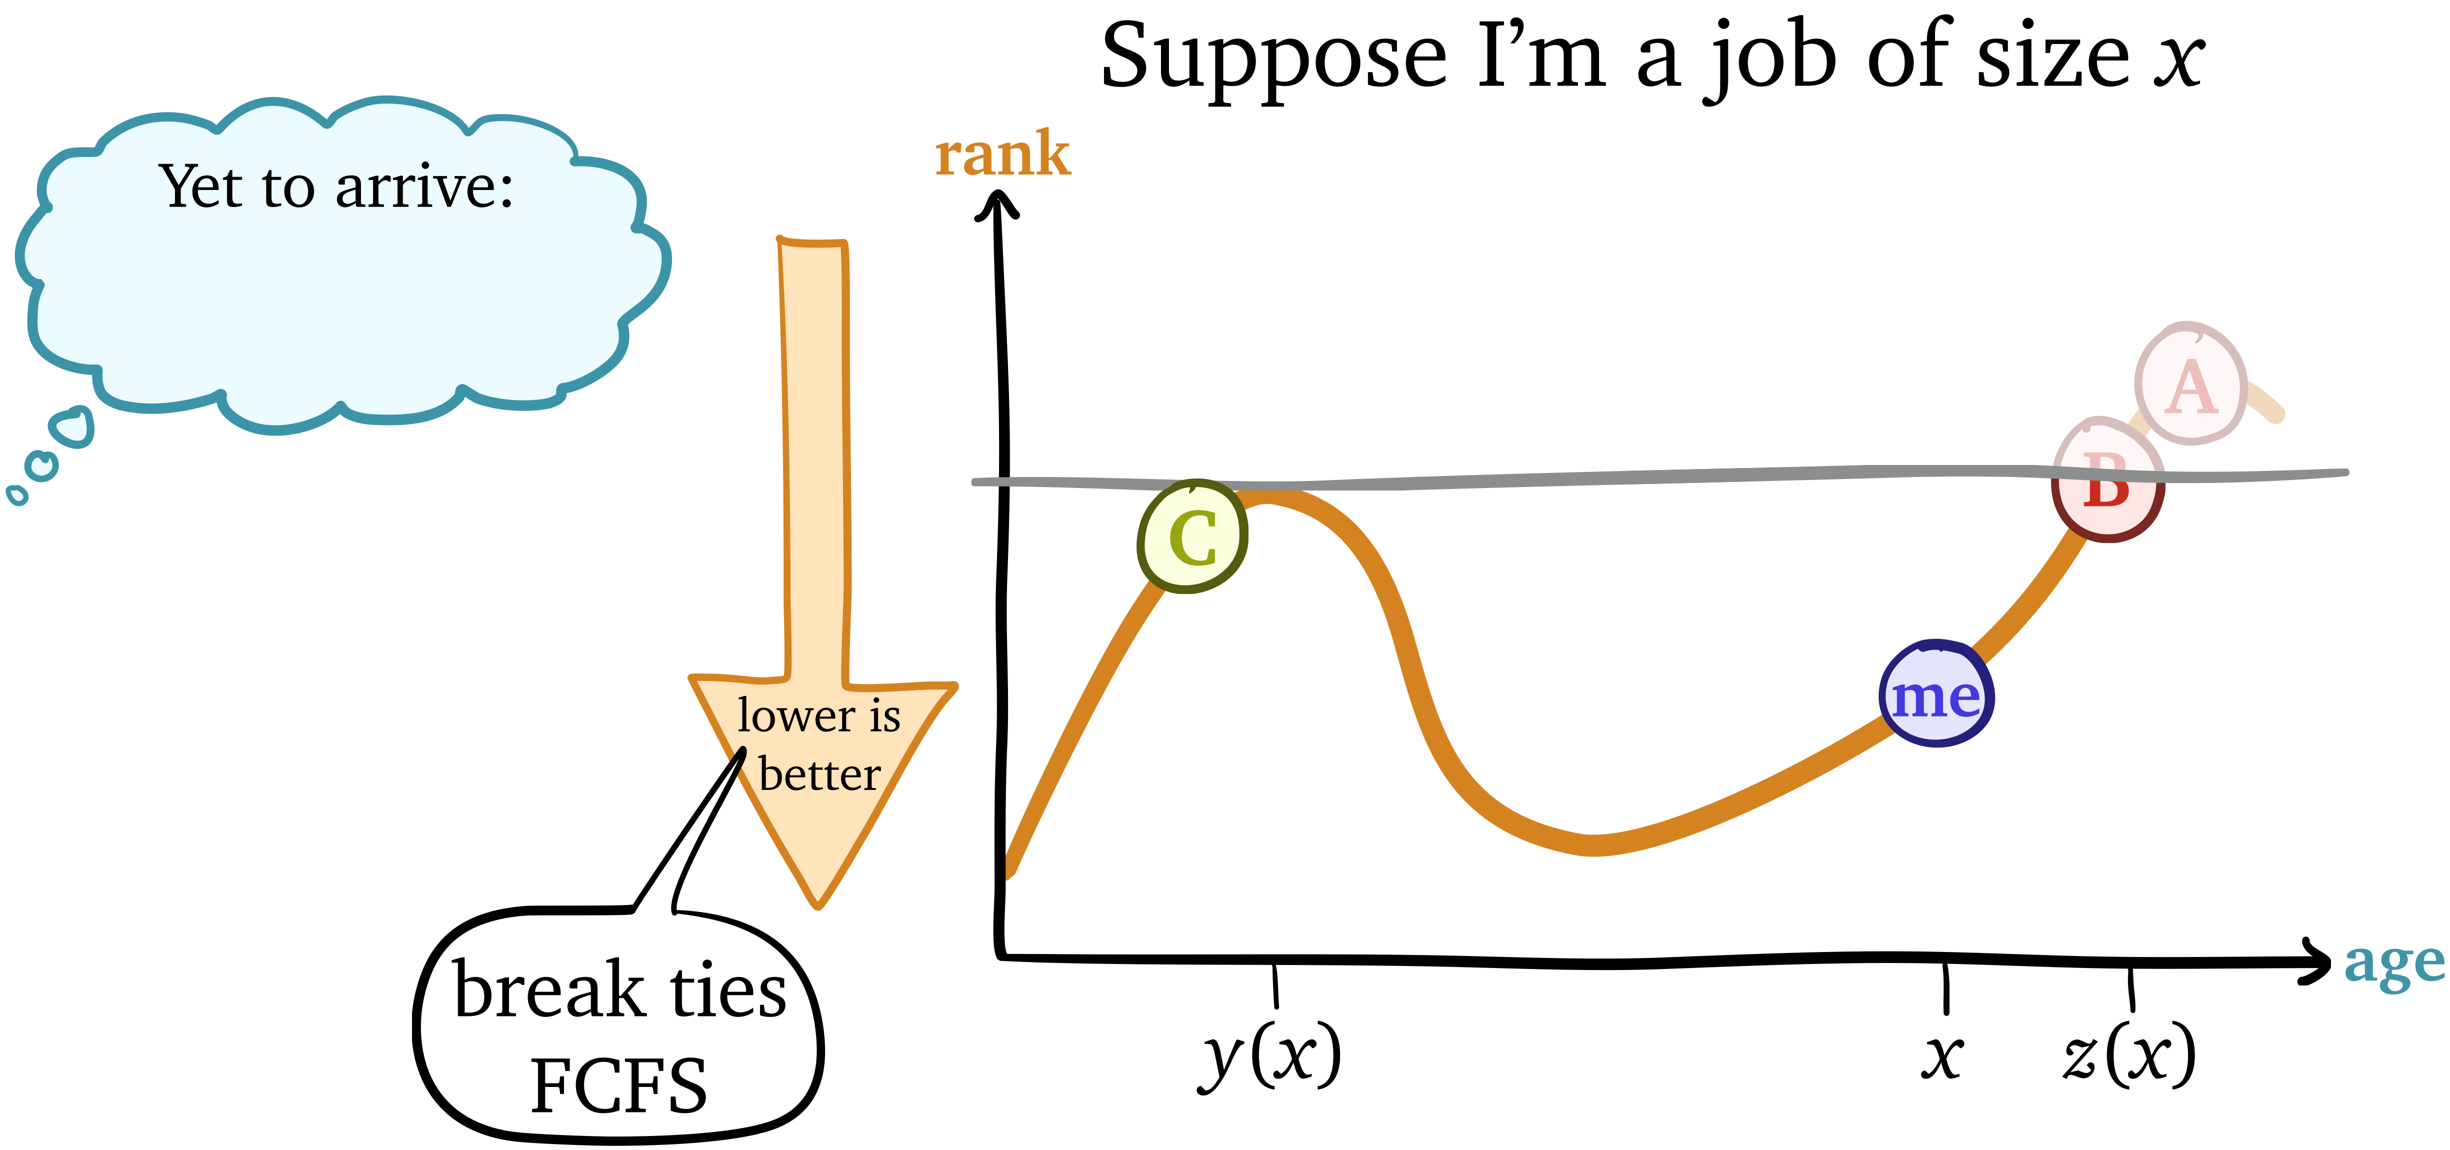

I can ignore { *old* jobs (**A** & **B**), *new* jobs (**C**) } after age { ???, ??? }

Analyzing **Gittins-1**

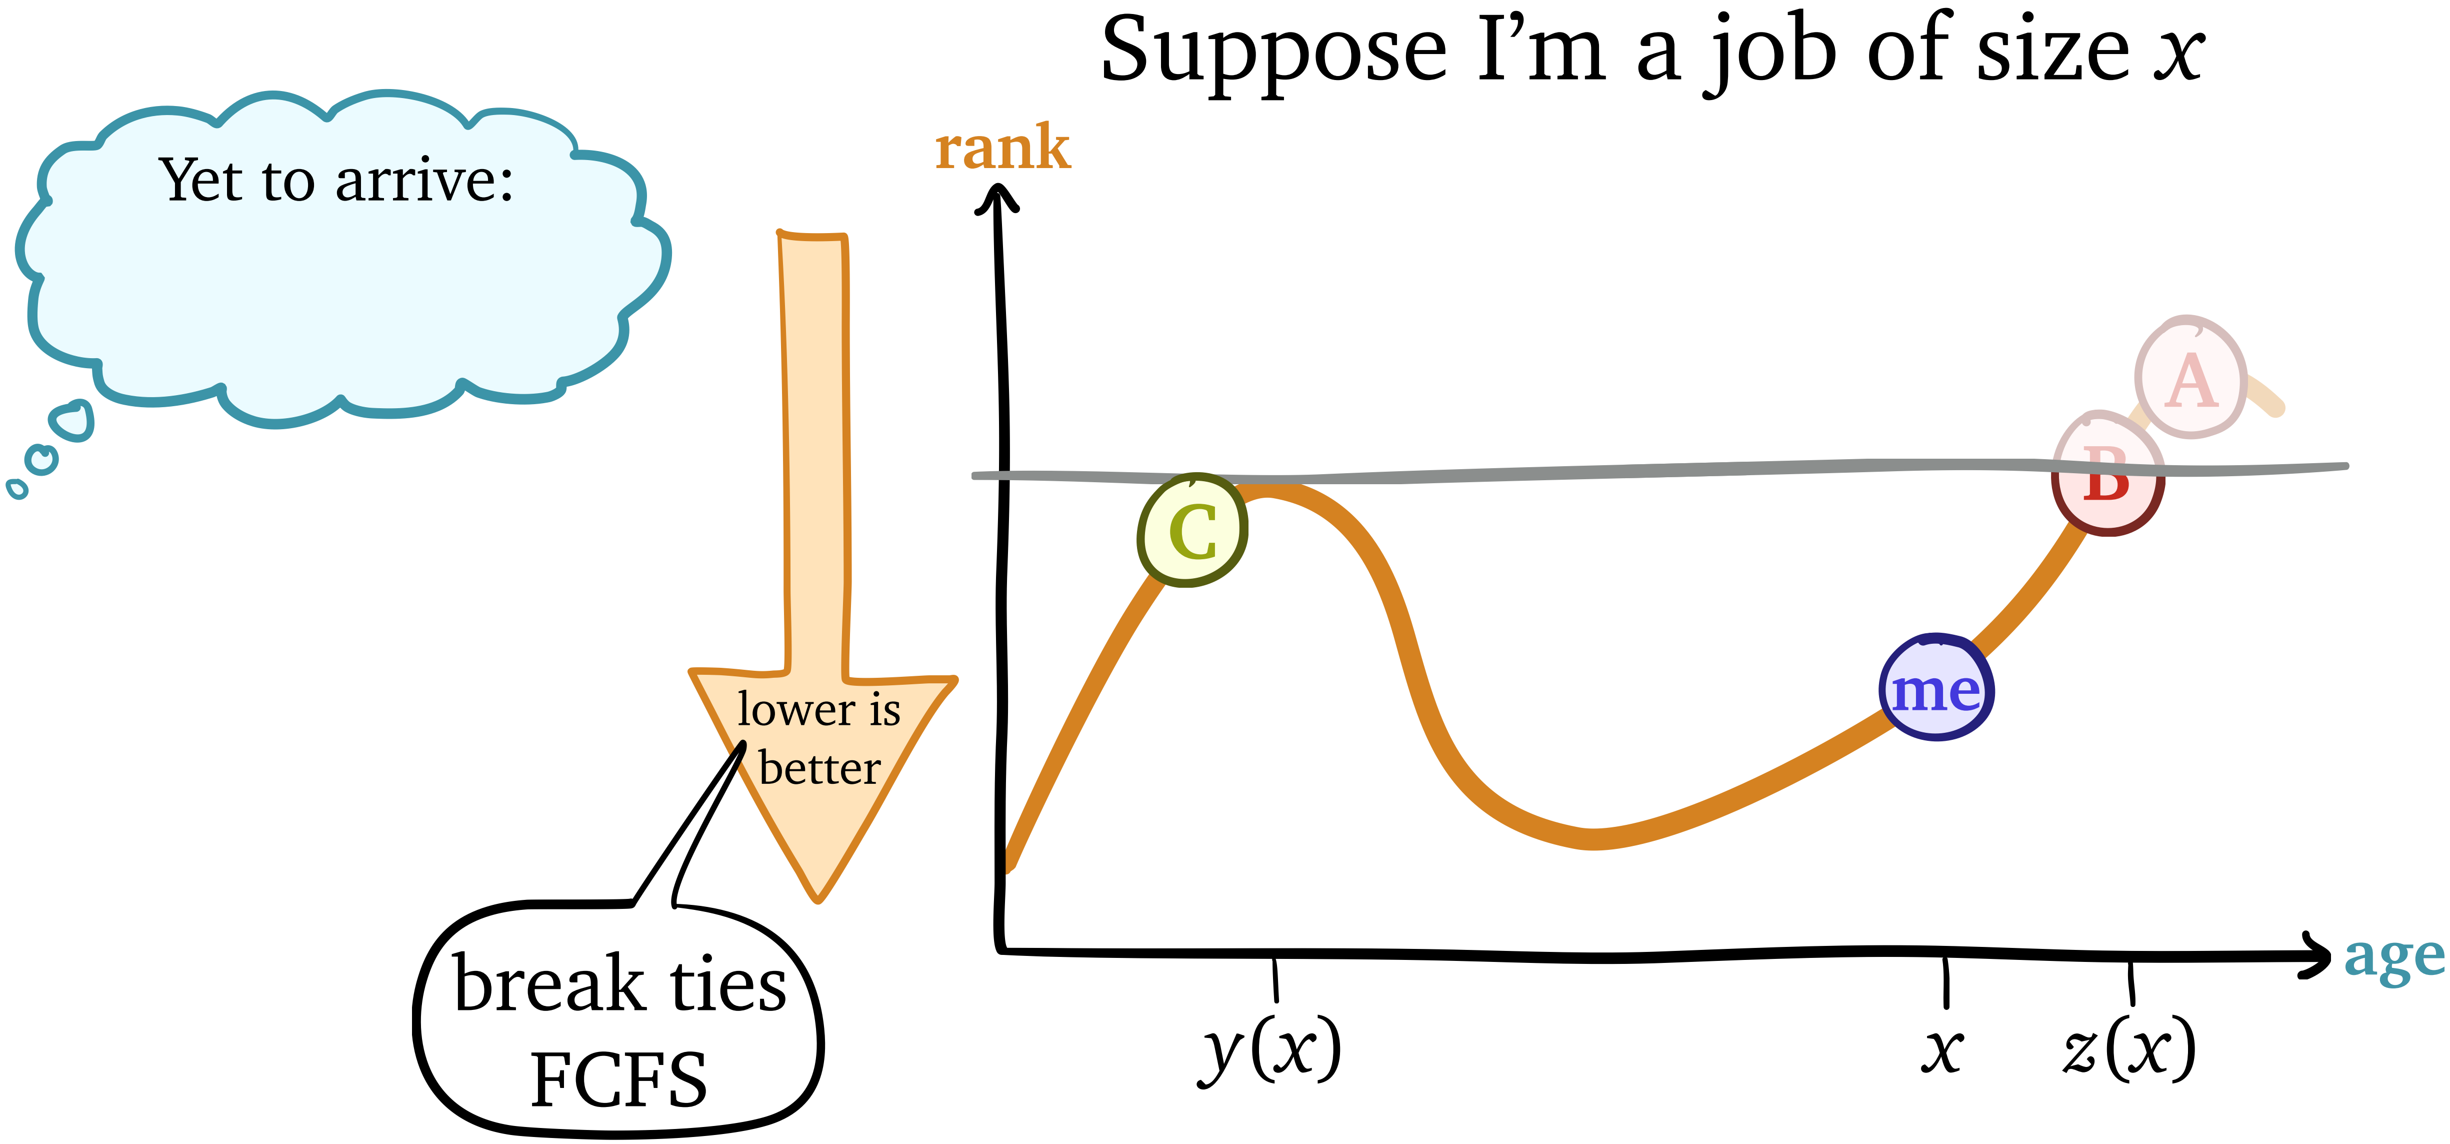

I can ignore { *old* jobs (**A** & **B**) / *new* jobs (**C**) } after age { $z(x)$ / **???** }

Analyzing **Gittins-1**

Suppose I'm a job of size x

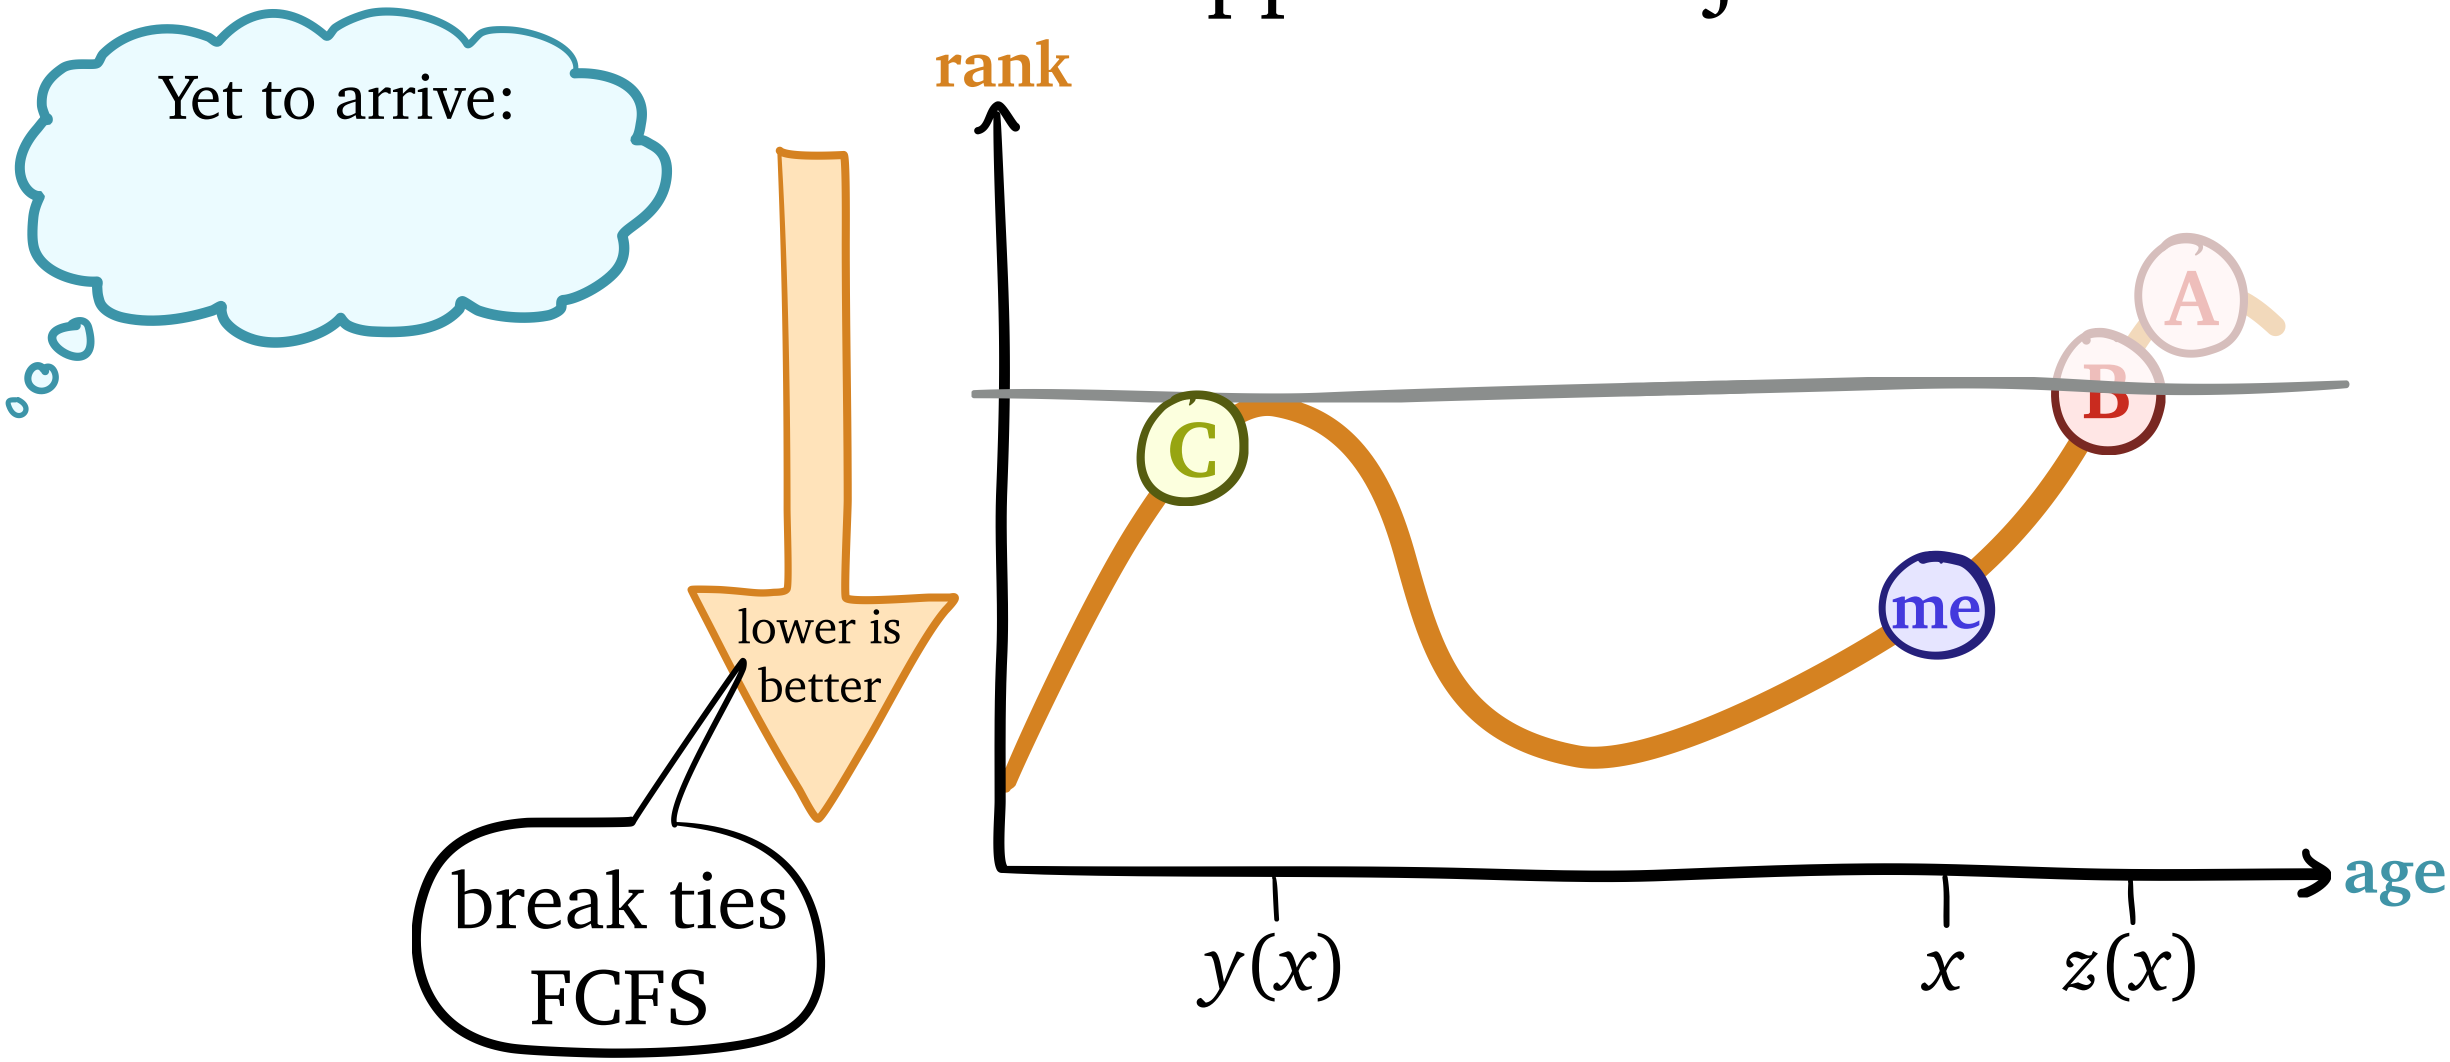

I can ignore $\left\{ \begin{array}{c} \textit{old} \text{ jobs (A \& B)} \\ \\ \textit{new} \text{ jobs (C)} \end{array} \right\}$ after age $\left\{ \begin{array}{c} z(x) \\ \\ y(x) \end{array} \right\}$

What goes wrong for **Gittins-*k***?

Analyzing **Gittins-*k***

Suppose I'm a job of size x

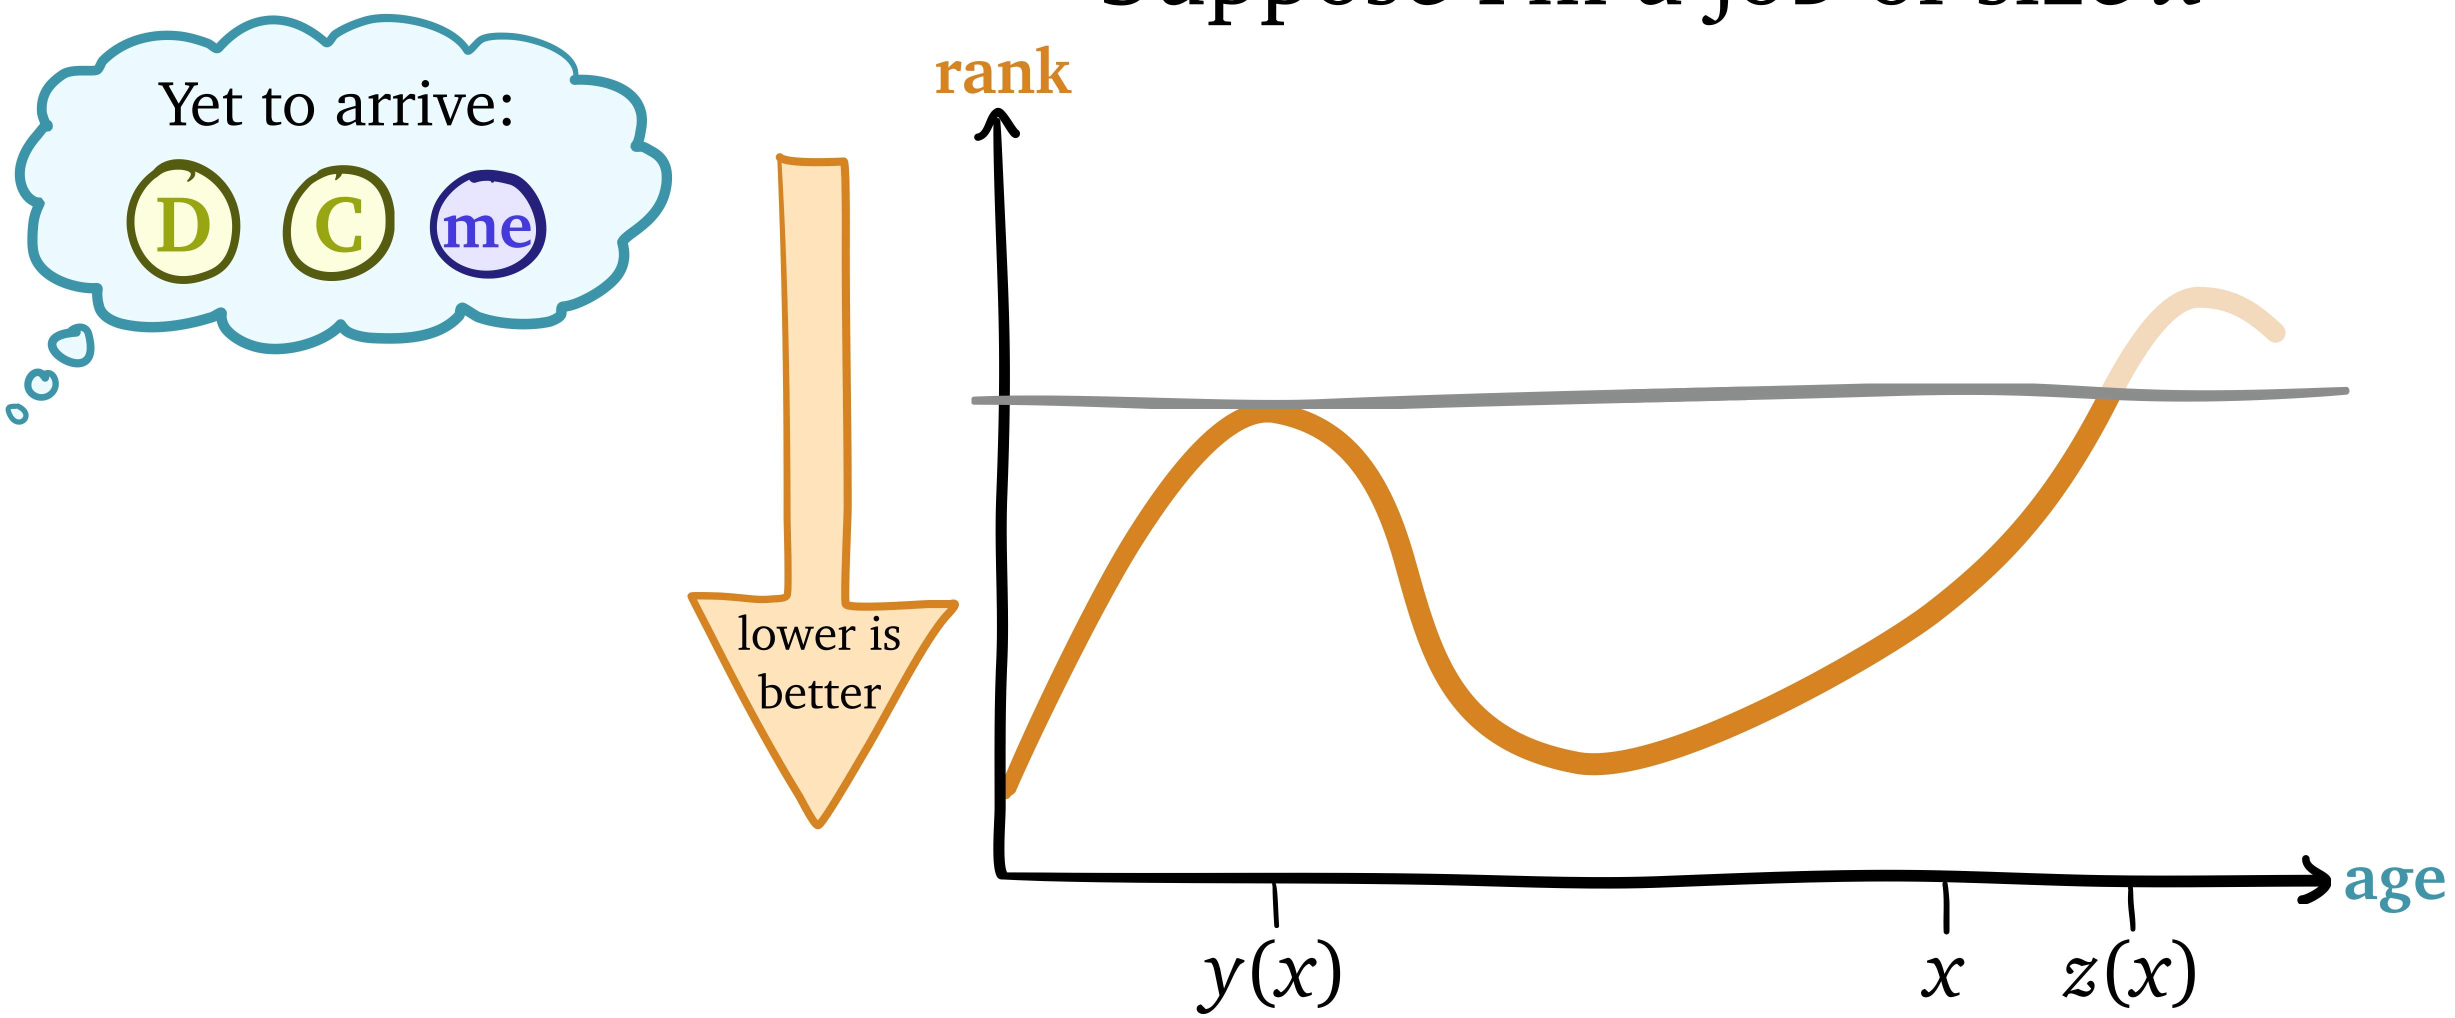

Analyzing **Gittins-*k***

Suppose I'm a job of size x

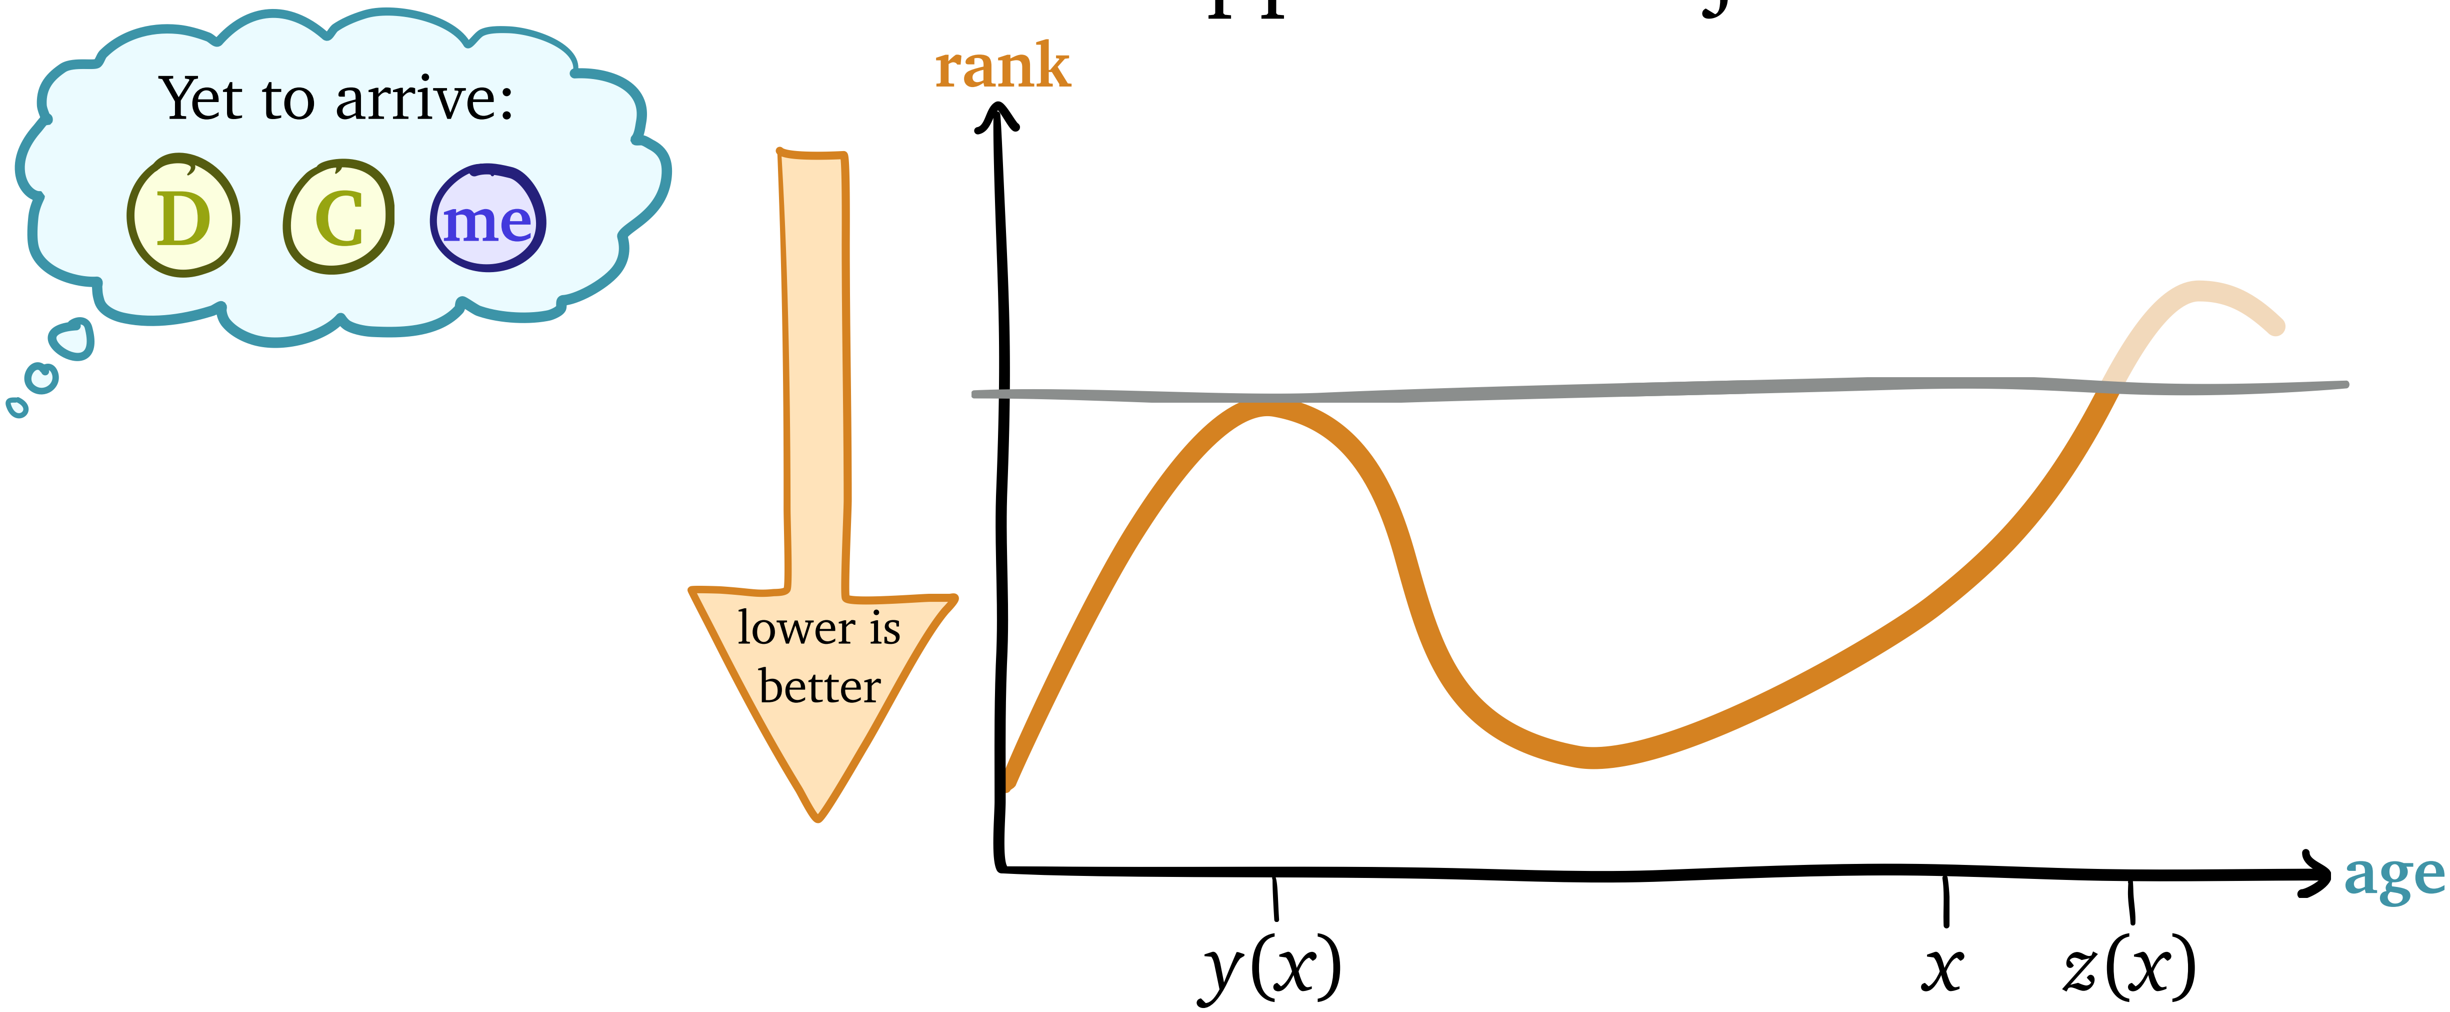

I can ignore { *old* jobs / *new* jobs (C & D) } after age { ??? / ??? }

Analyzing **Gittins-*k***

Suppose I'm a job $k = 2$

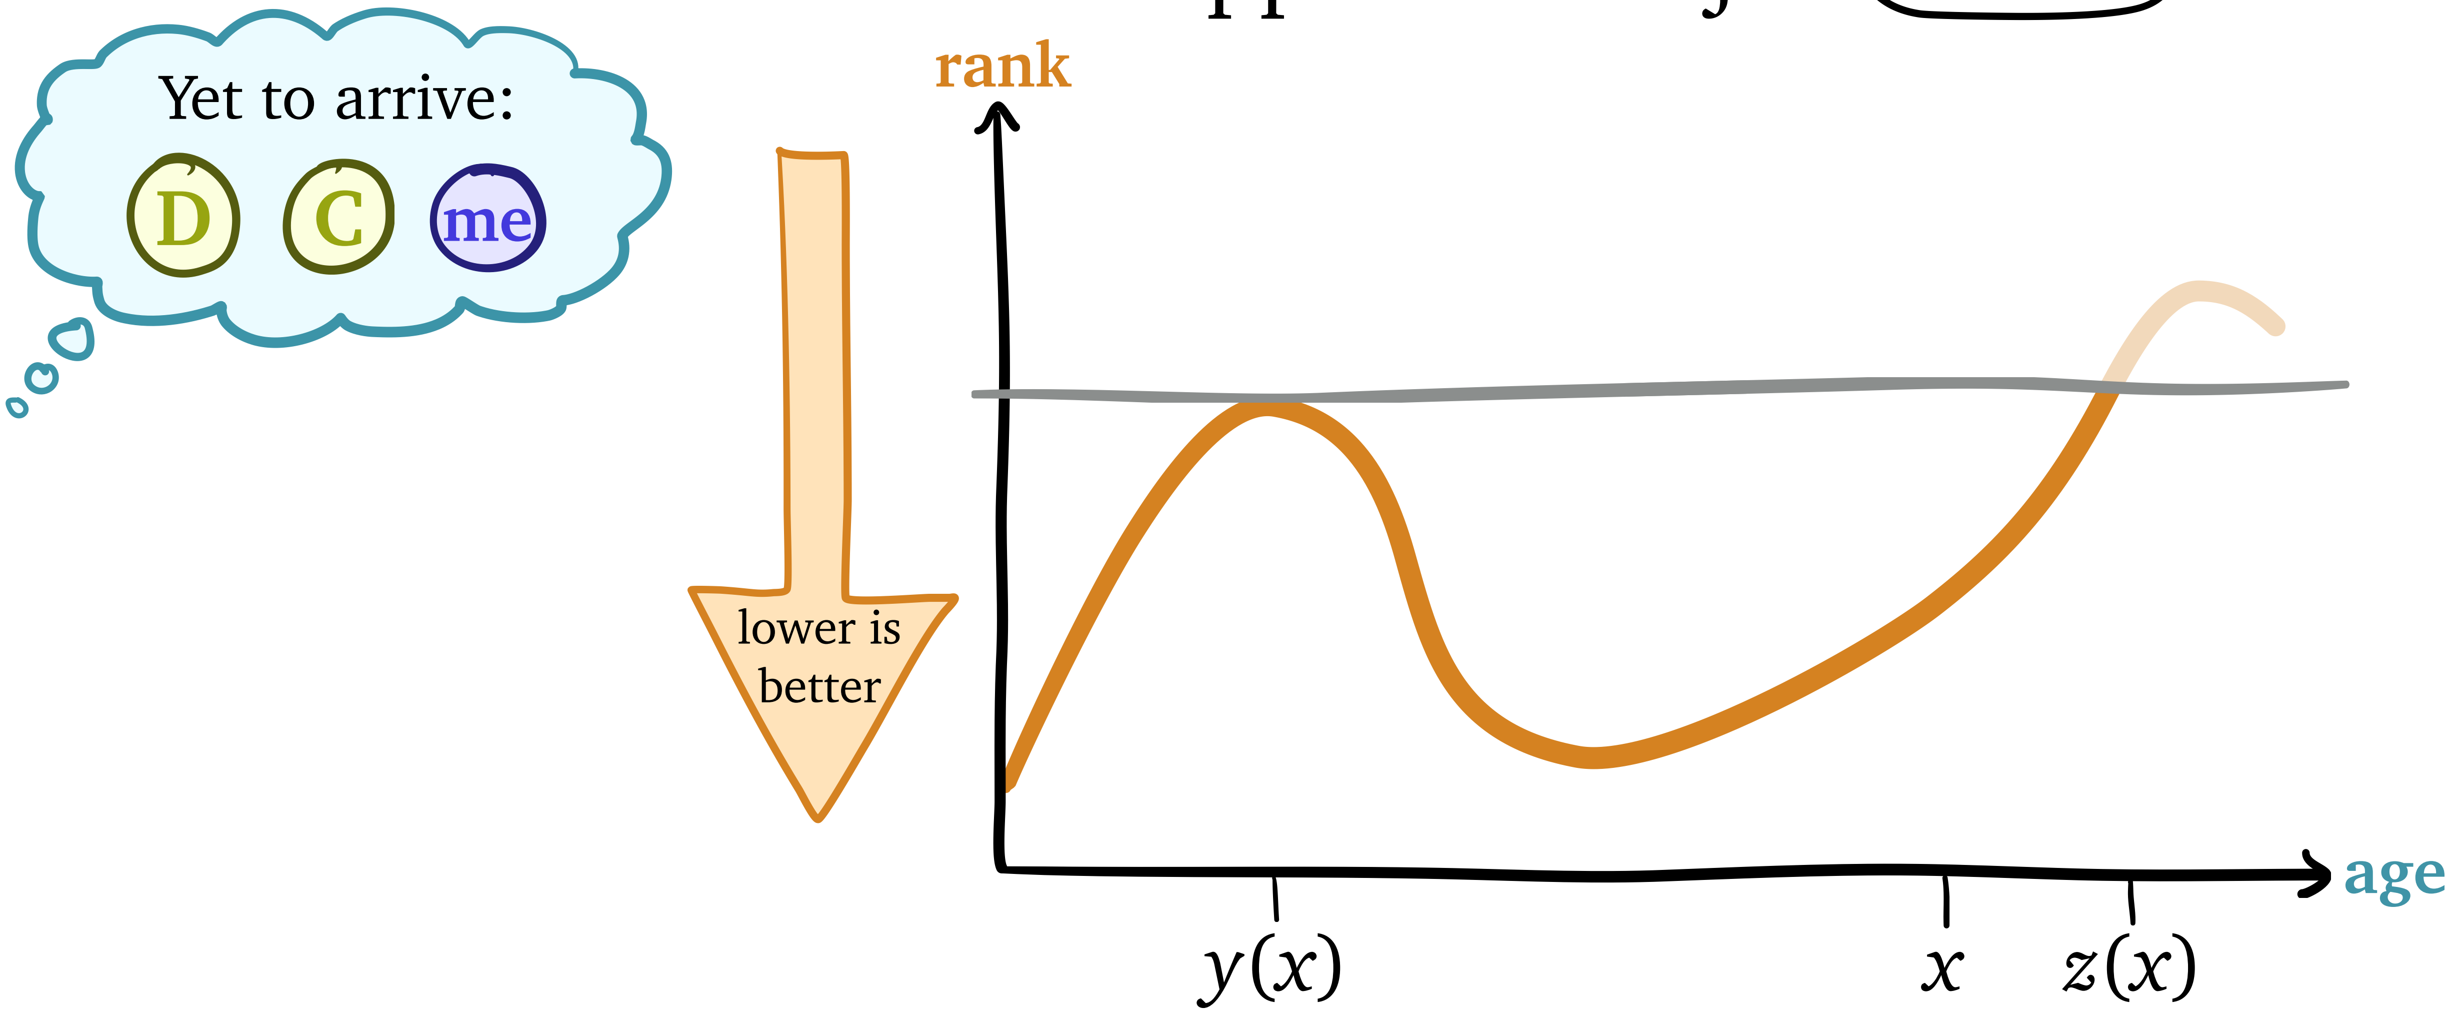

I can ignore { *old* jobs / *new* jobs (C & D) } after age { ??? / ??? }

Analyzing **Gittins-*k***

Suppose I'm a job x ($k = 2$)

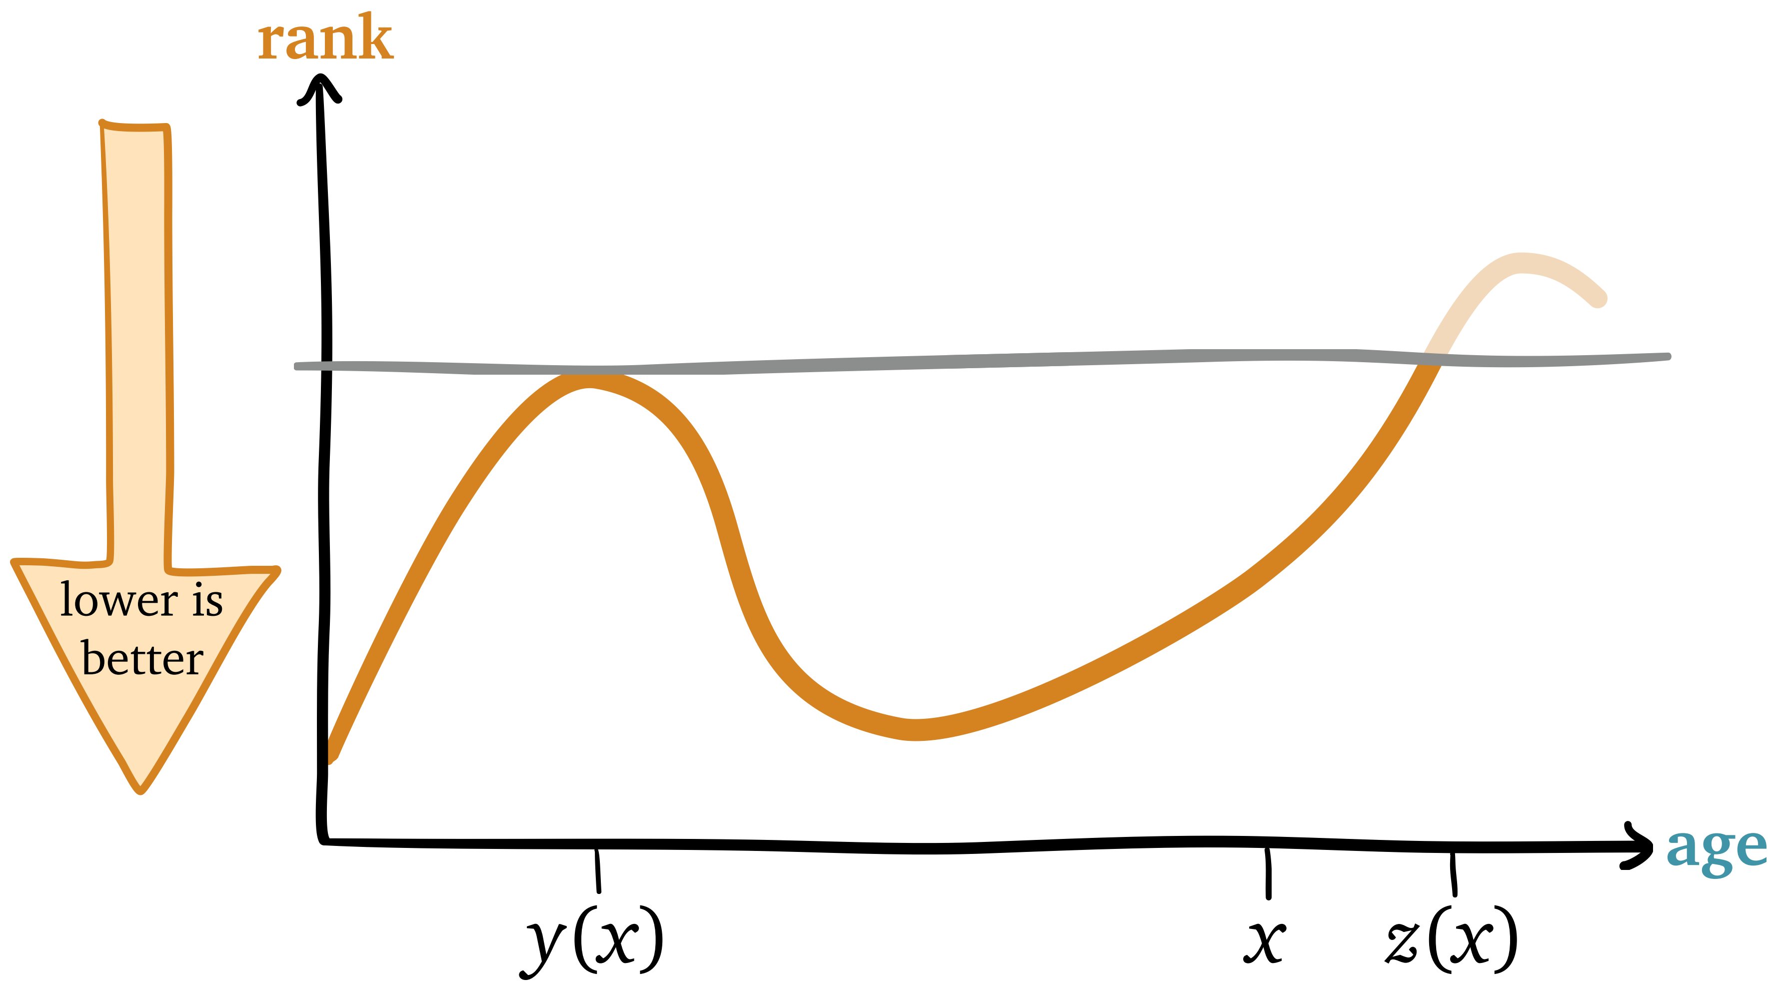

I can ignore { *old* jobs after age $z(x)$; *new* jobs (C & D) after age ??? }

Analyzing **Gittins-*k***

Suppose I'm a job x ($k = 2$)

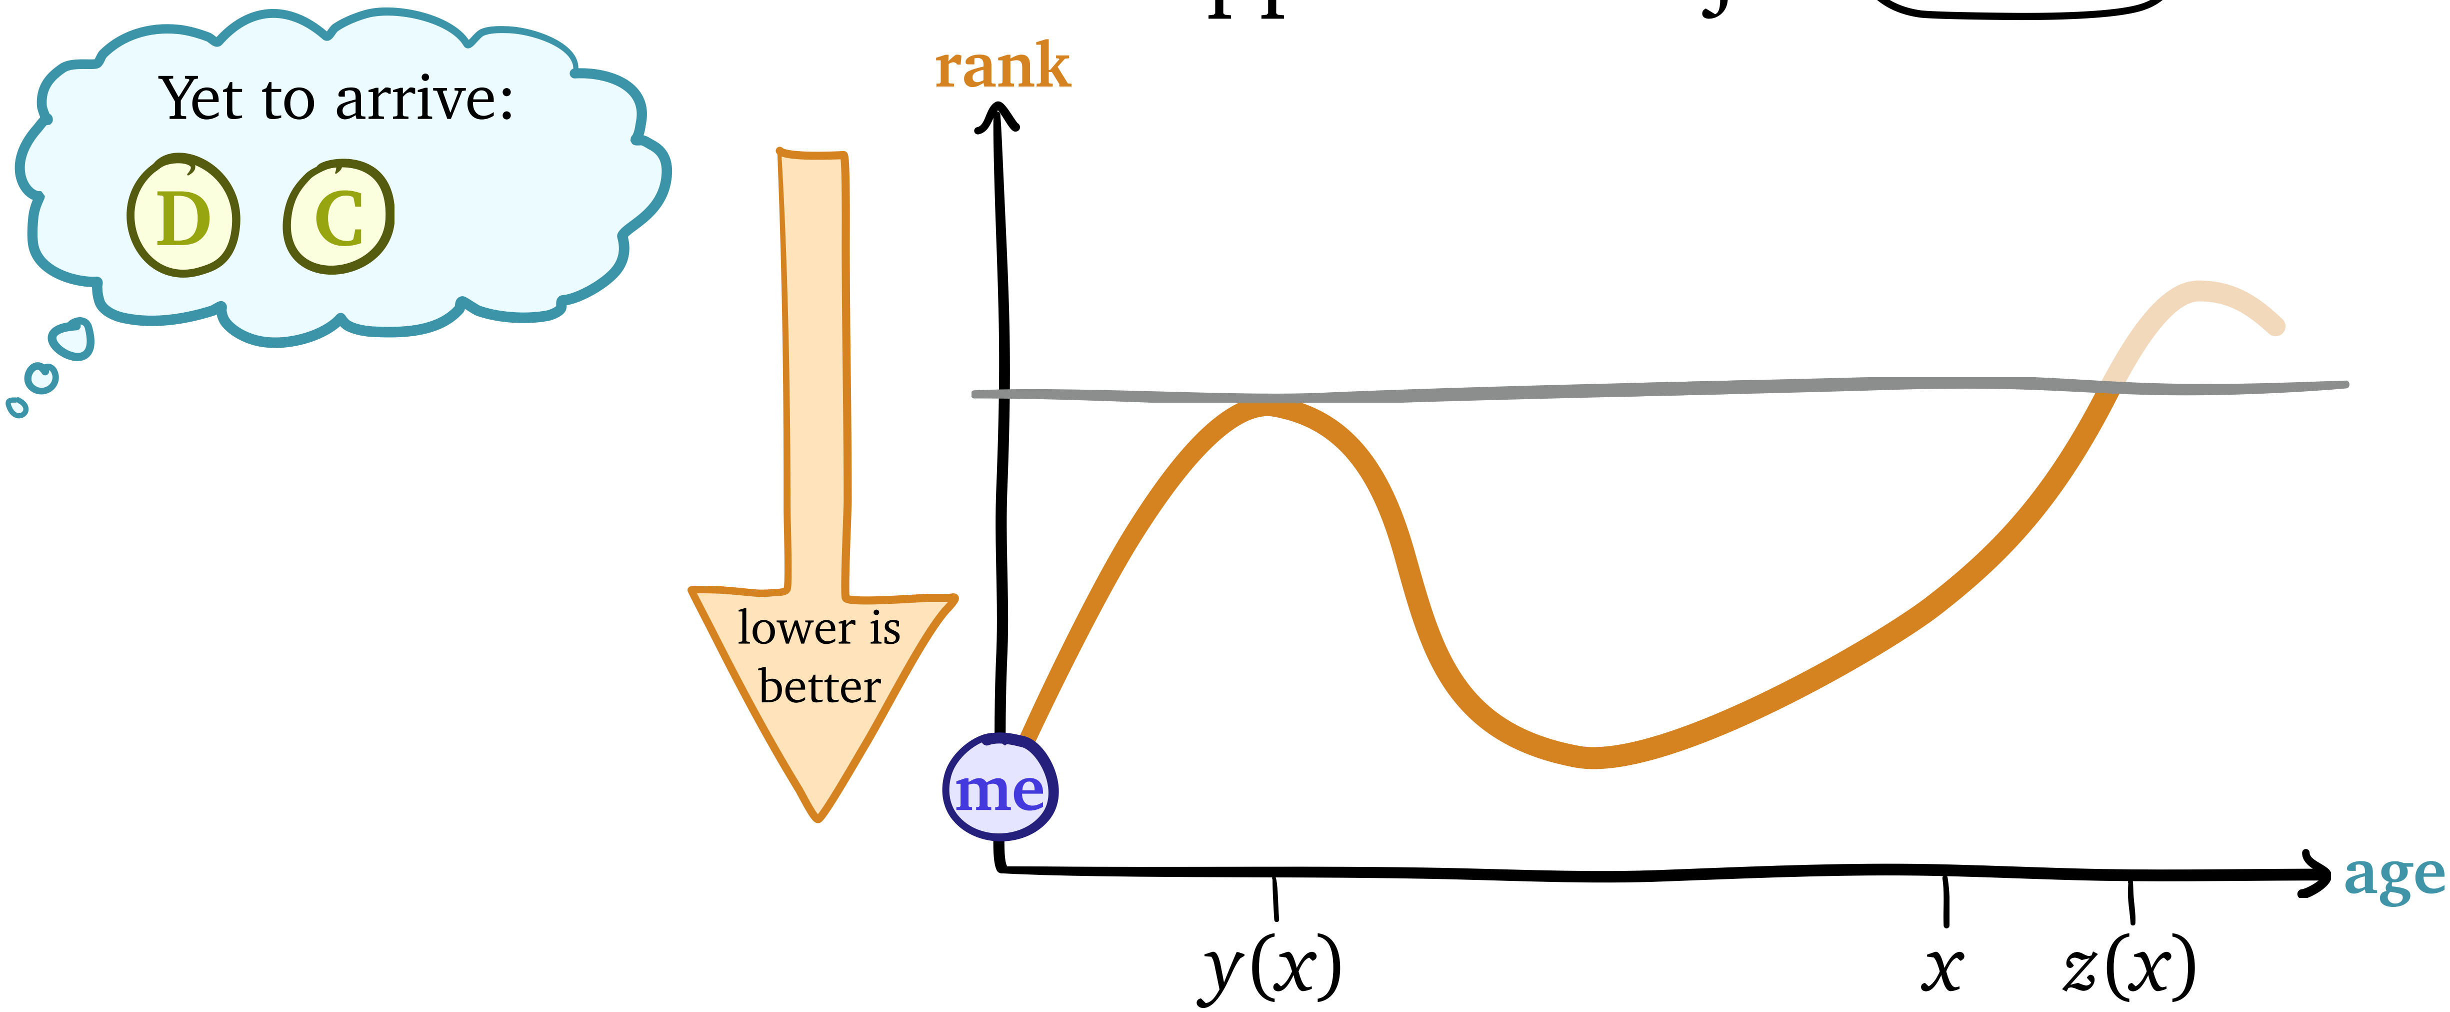

I can ignore *old* jobs after age $z(x)$

I can ignore *new* jobs (C & D) after age ???

Analyzing **Gittins-k**

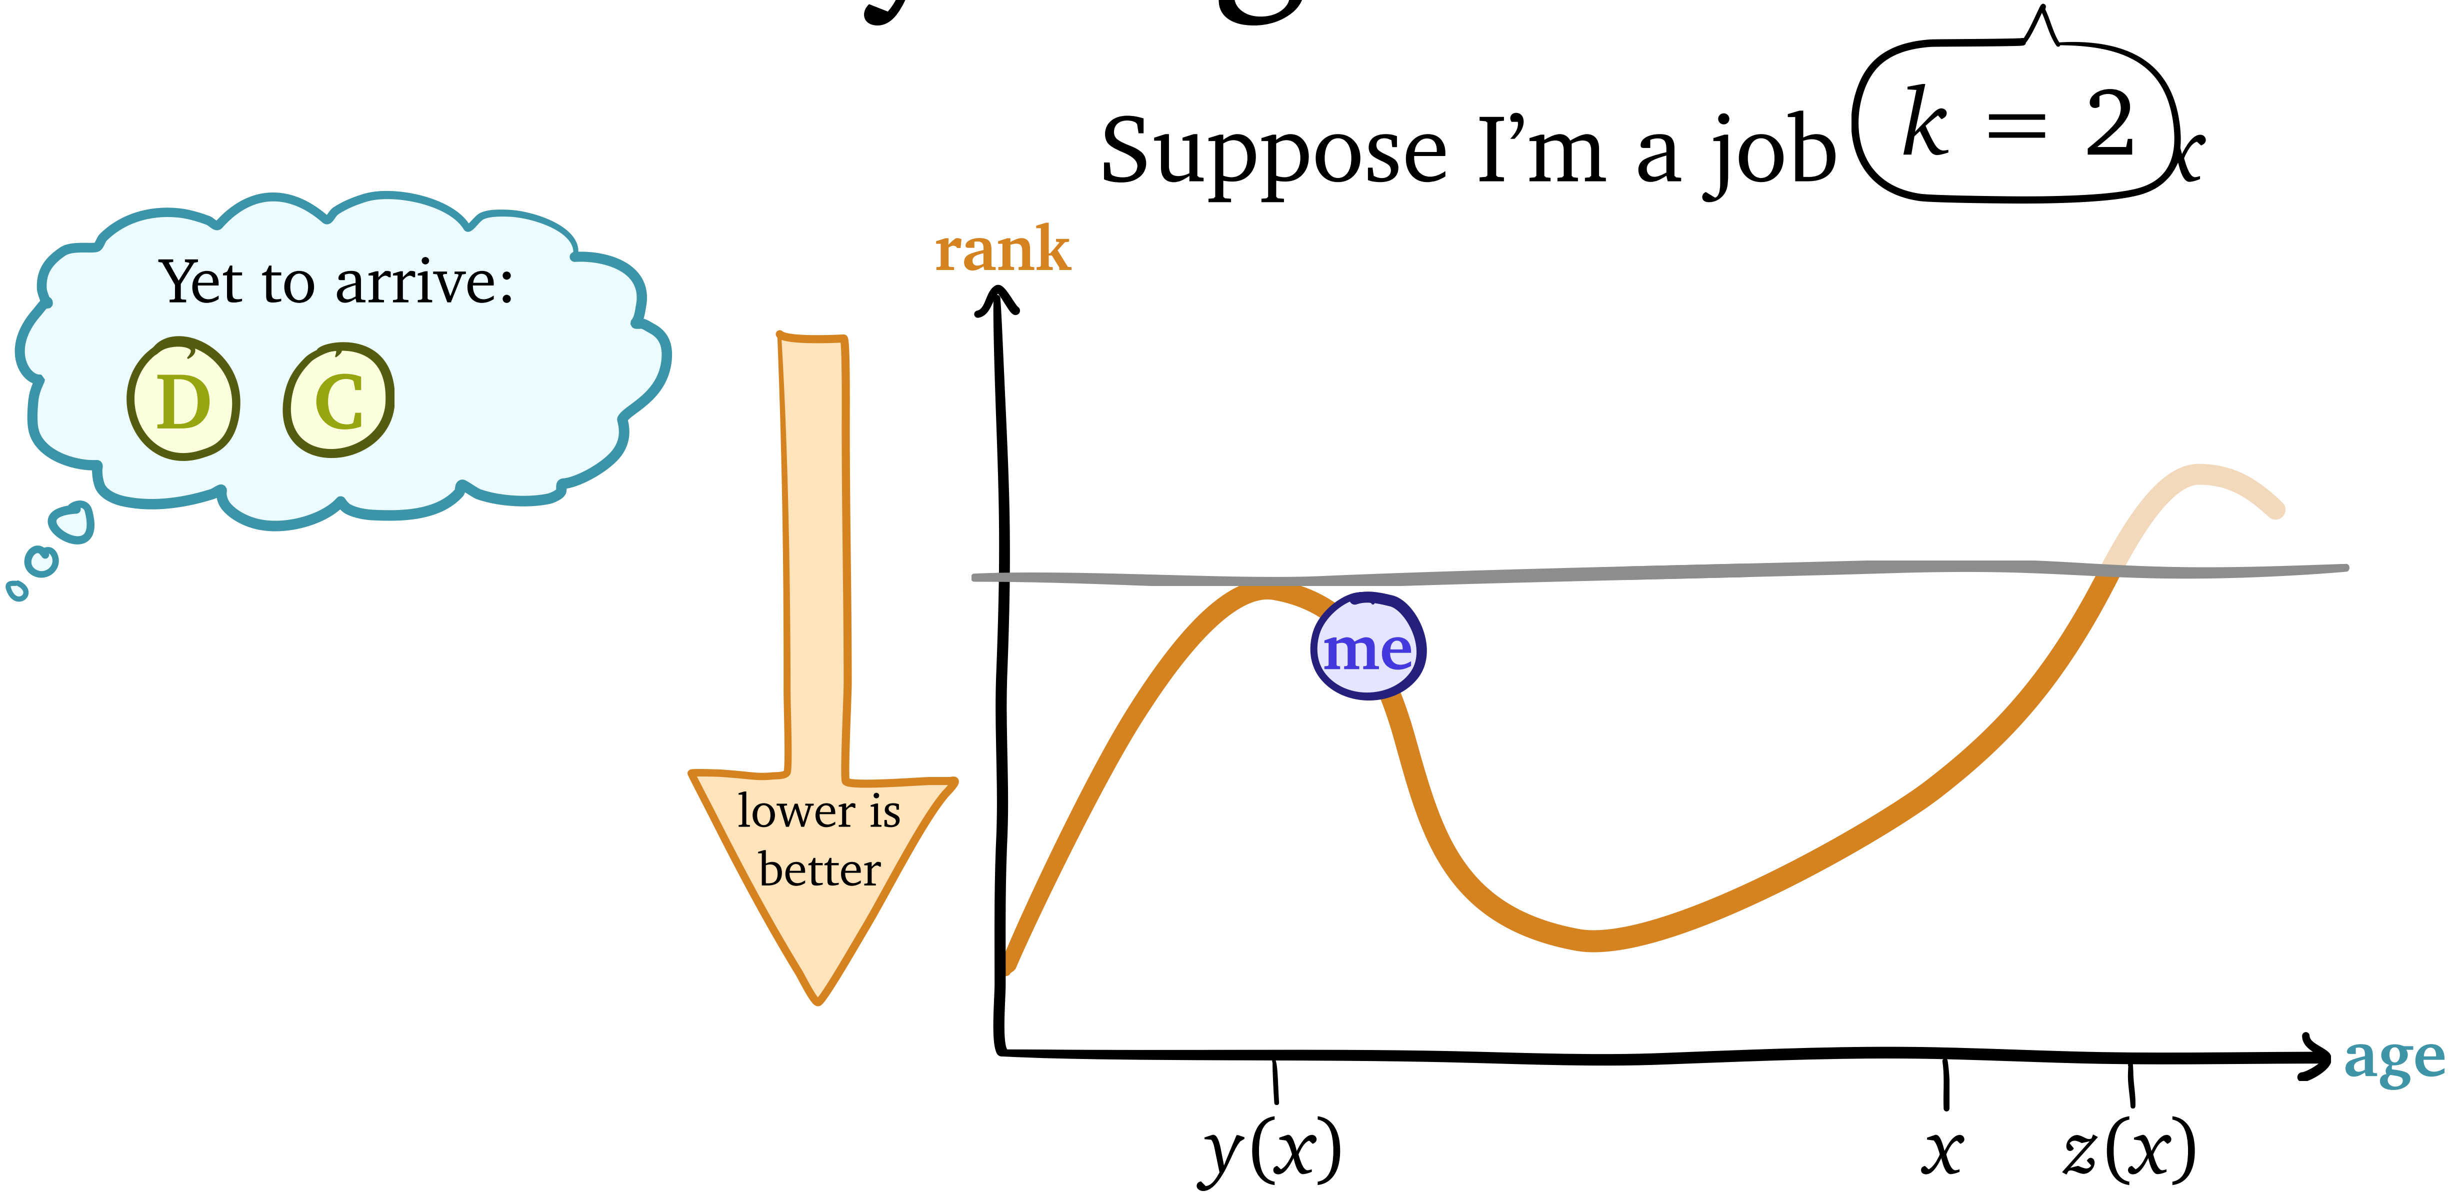

$$\text{I can ignore} \begin{Bmatrix} \textit{old}\text{ jobs} \\ \textit{new}\text{ jobs (C \& D)} \end{Bmatrix} \text{after age} \begin{Bmatrix} z(x) \\ ??? \end{Bmatrix}$$

Analyzing Gittins-k

Suppose I'm a job x

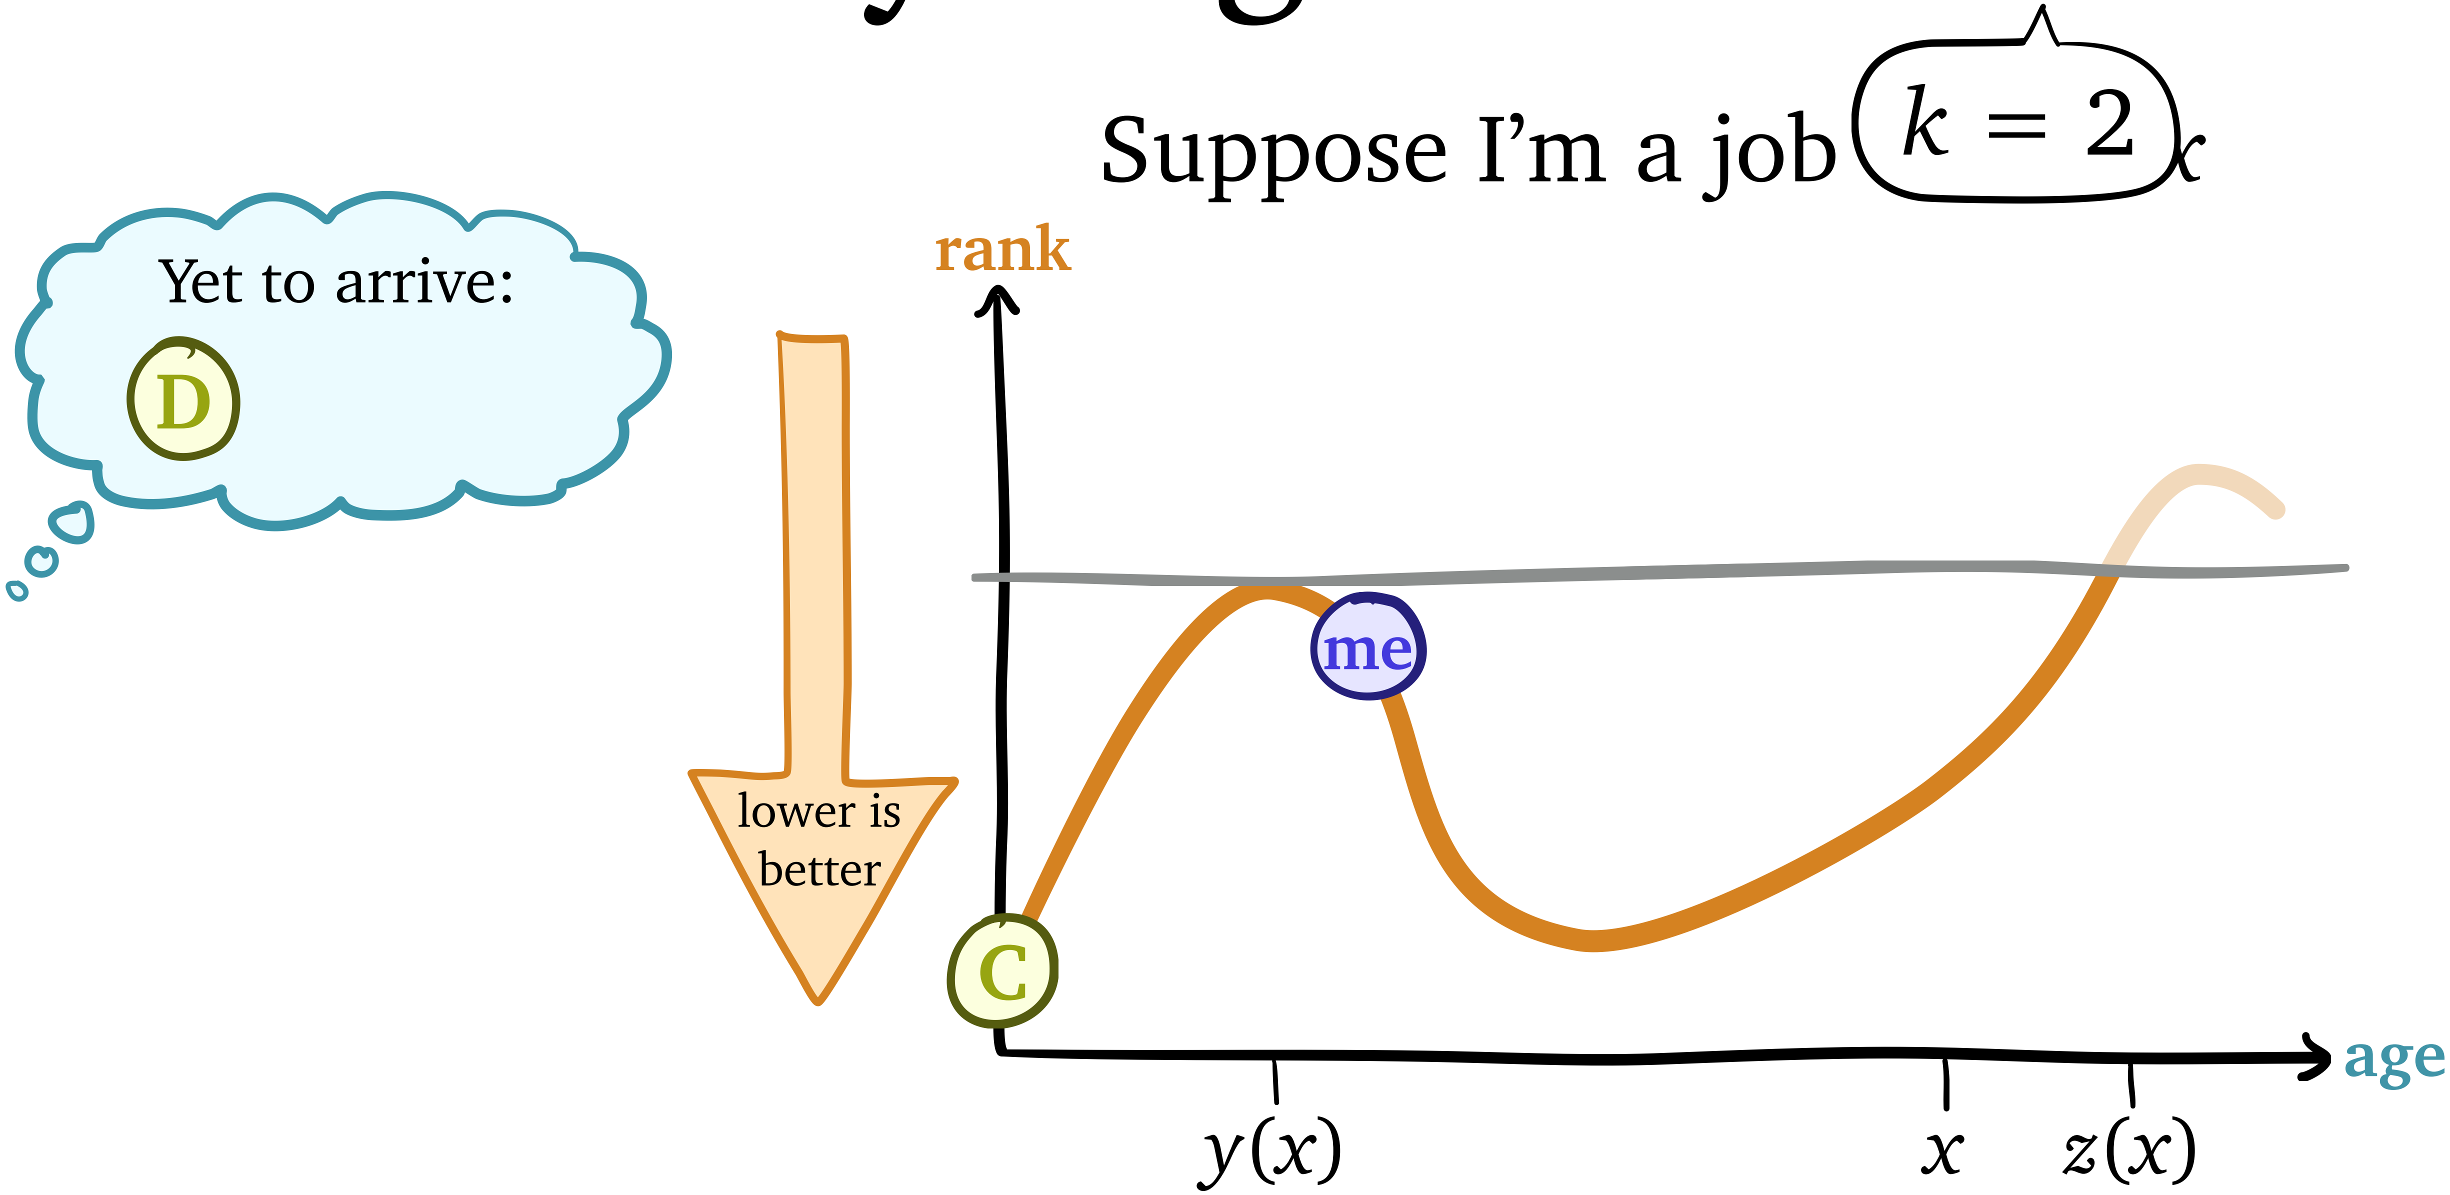

I can ignore $\left\{ \begin{array}{c} \textit{old} \text{ jobs} \\ \textit{new} \text{ jobs (C \& D)} \end{array} \right\}$ after age $\left\{ \begin{array}{c} z(x) \\ ??? \end{array} \right\}$

Analyzing **Gittins-*k***

Suppose I'm a job x

$k = 2$

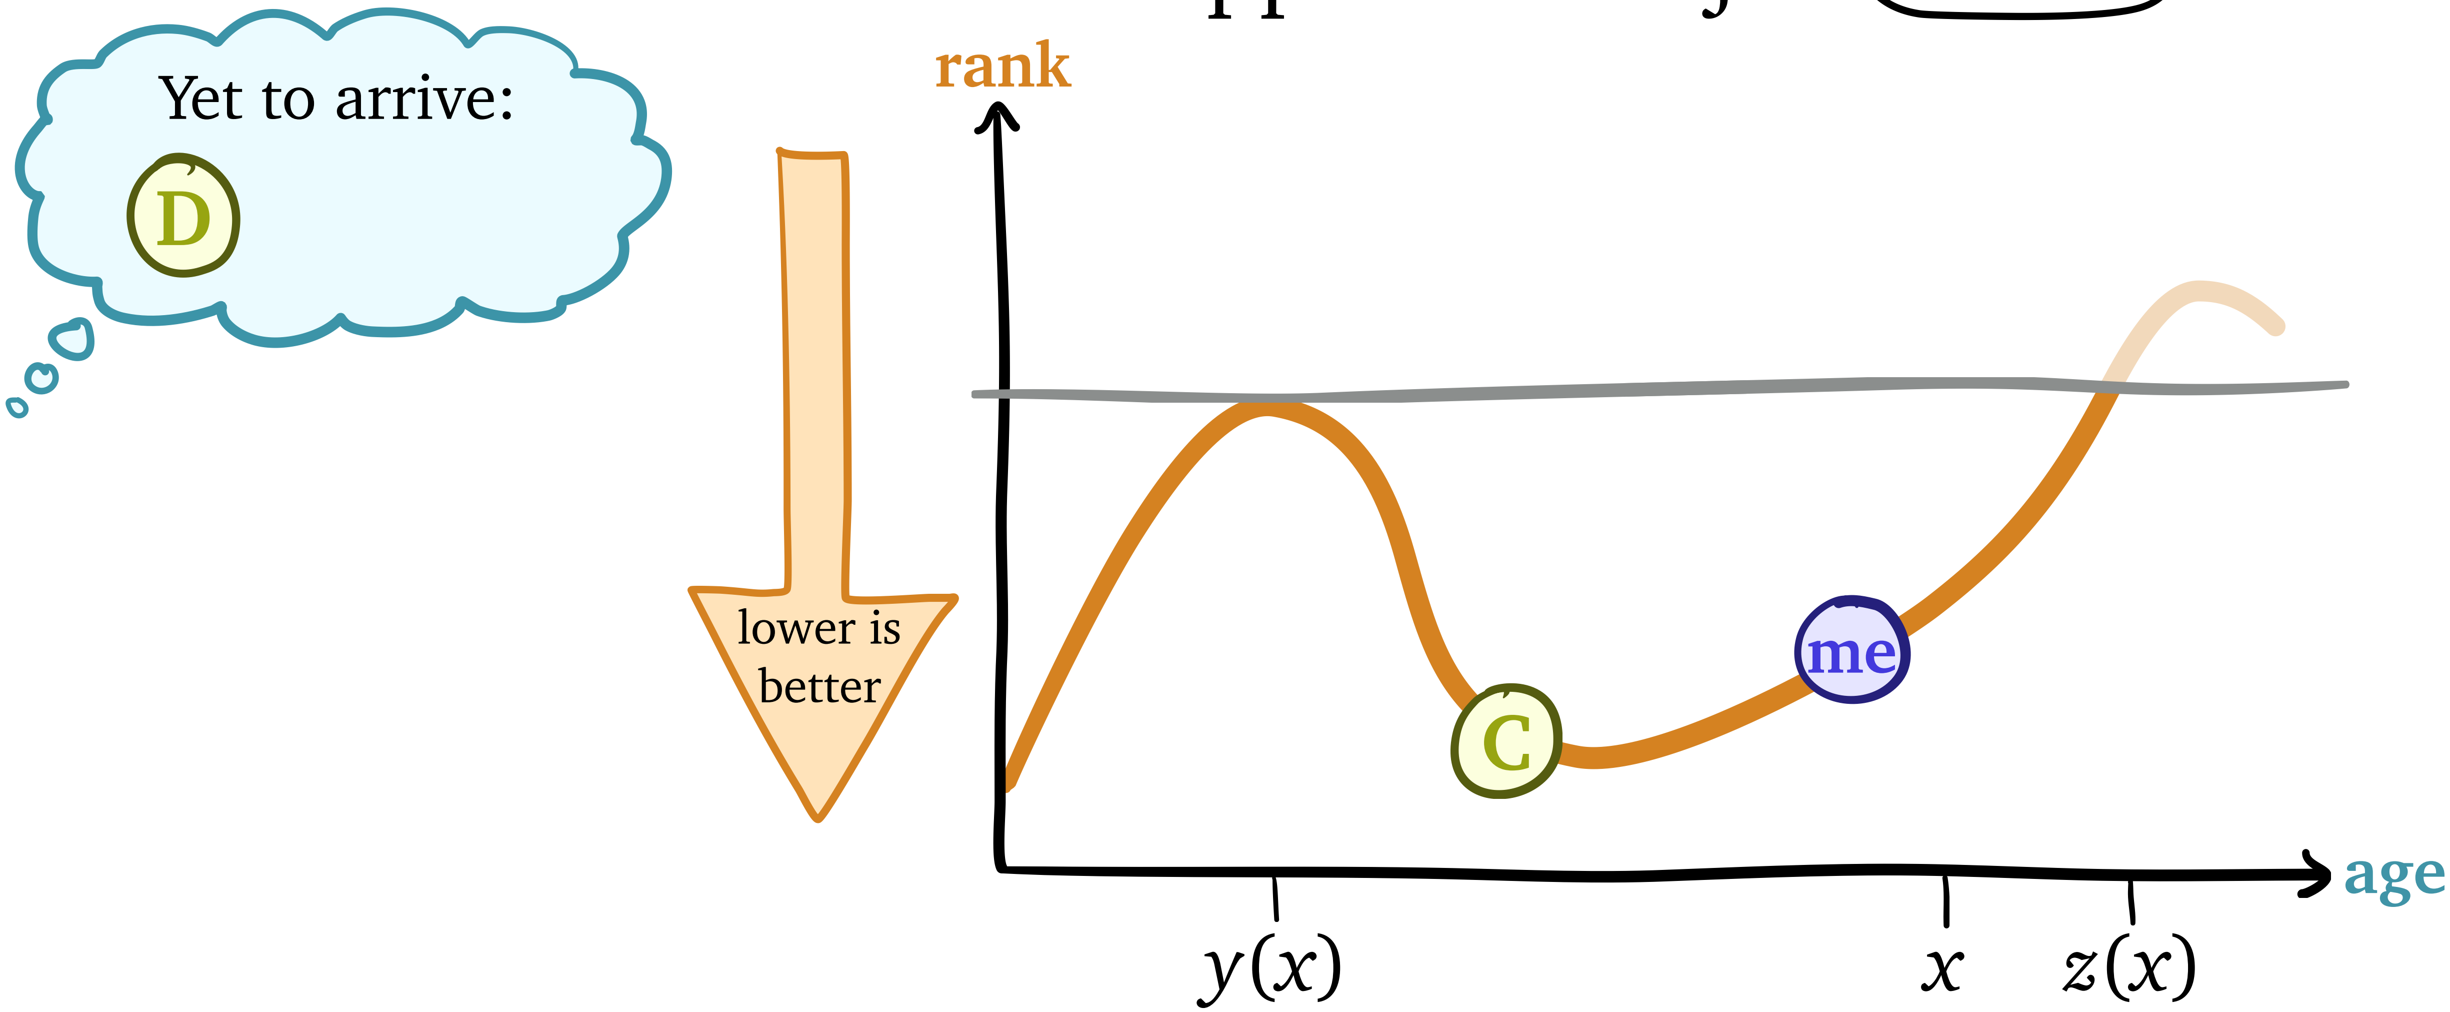

I can ignore *old* jobs after age $z(x)$

I can ignore *new* jobs (**C** & **D**) after age **???**

Analyzing Gittins-*k*

$k = 2$

Suppose I’m a job x

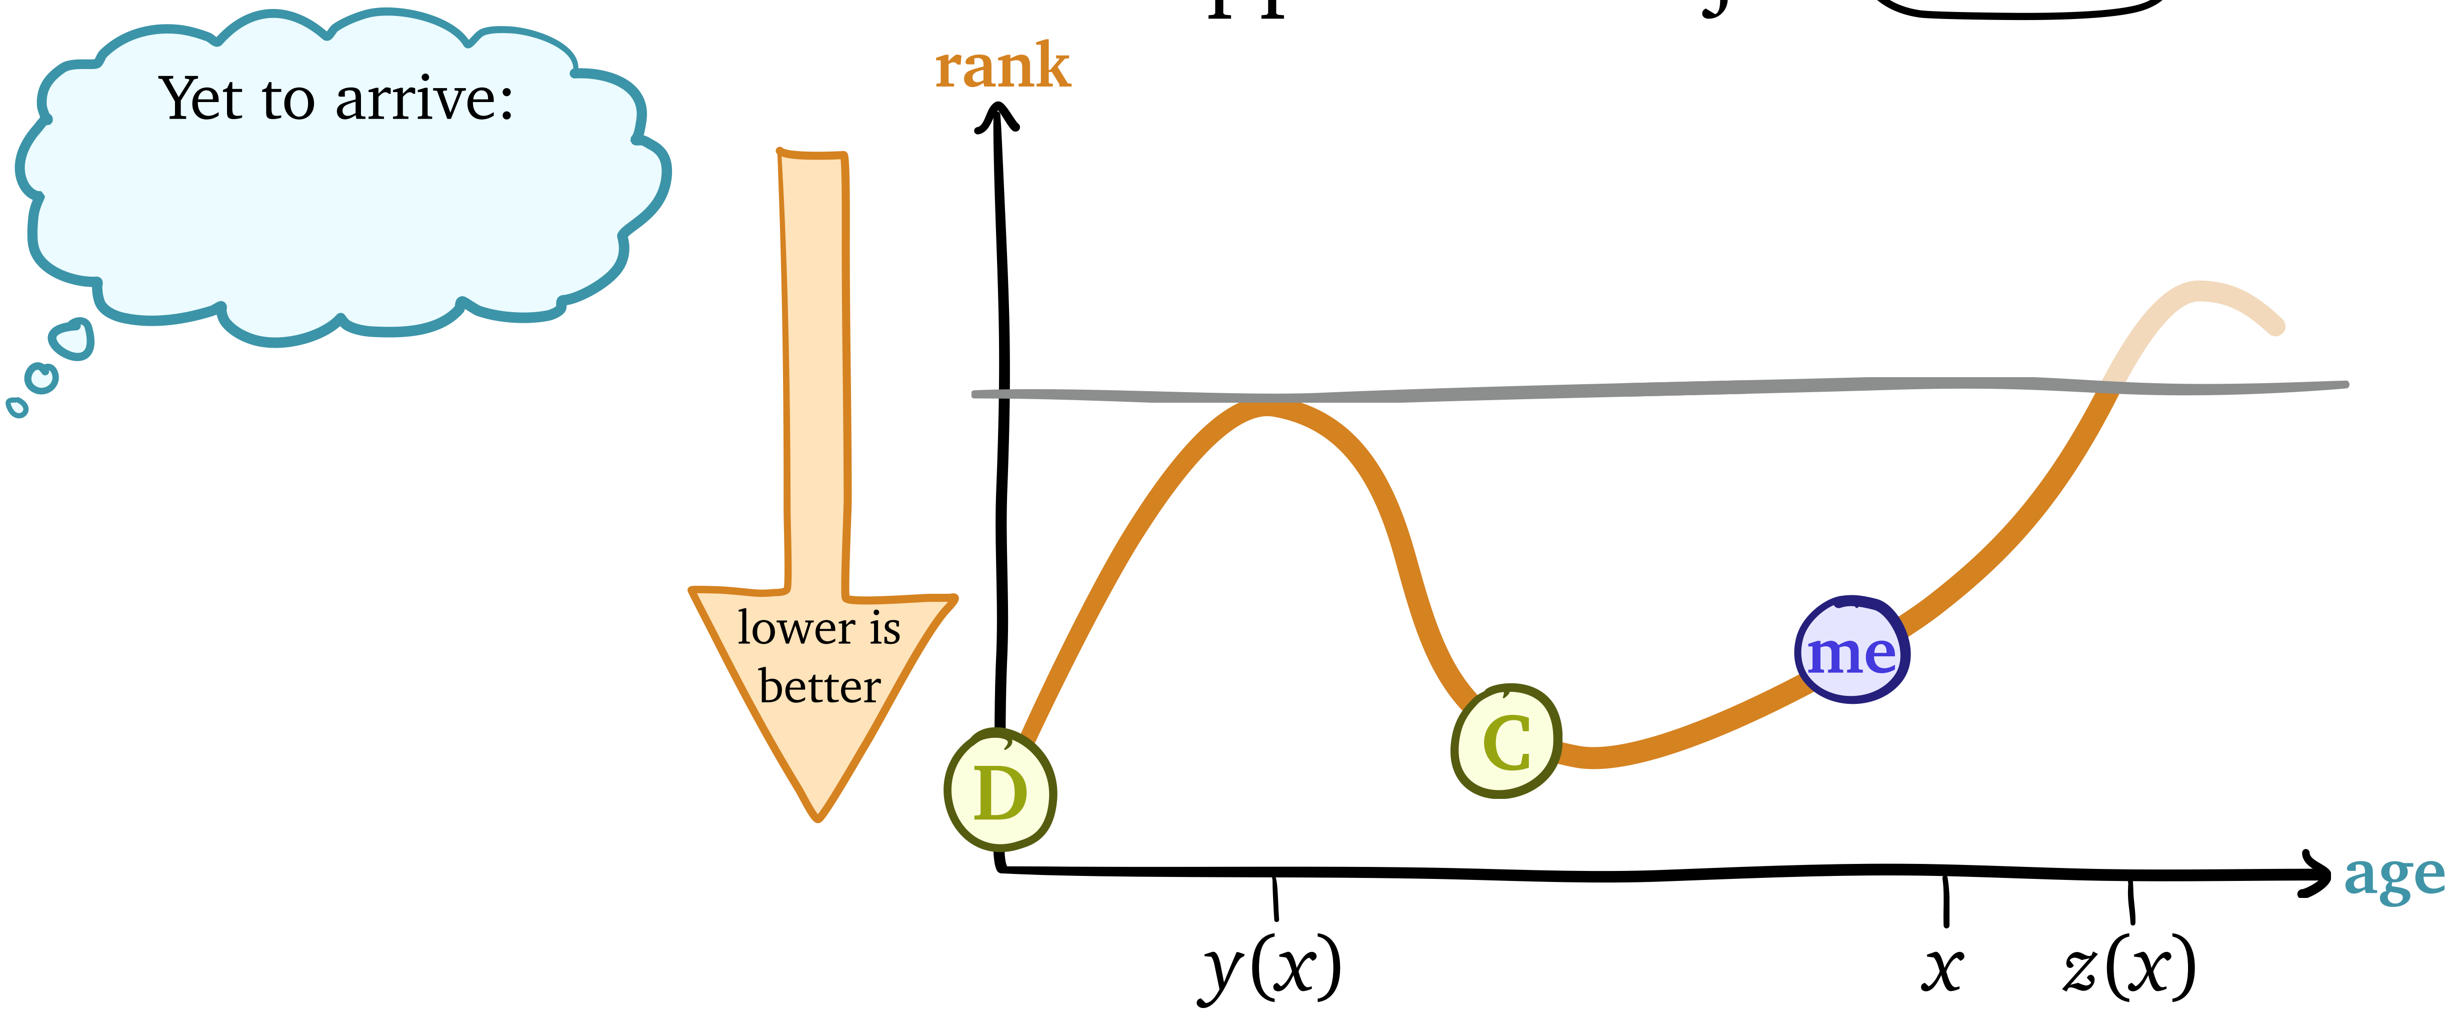

I can ignore { *old* jobs / *new* jobs (**C** & **D**) } after age { $z(x)$ / **???** }

Analyzing **Gittins-*k***

Suppose I'm a job $k = 2$ x

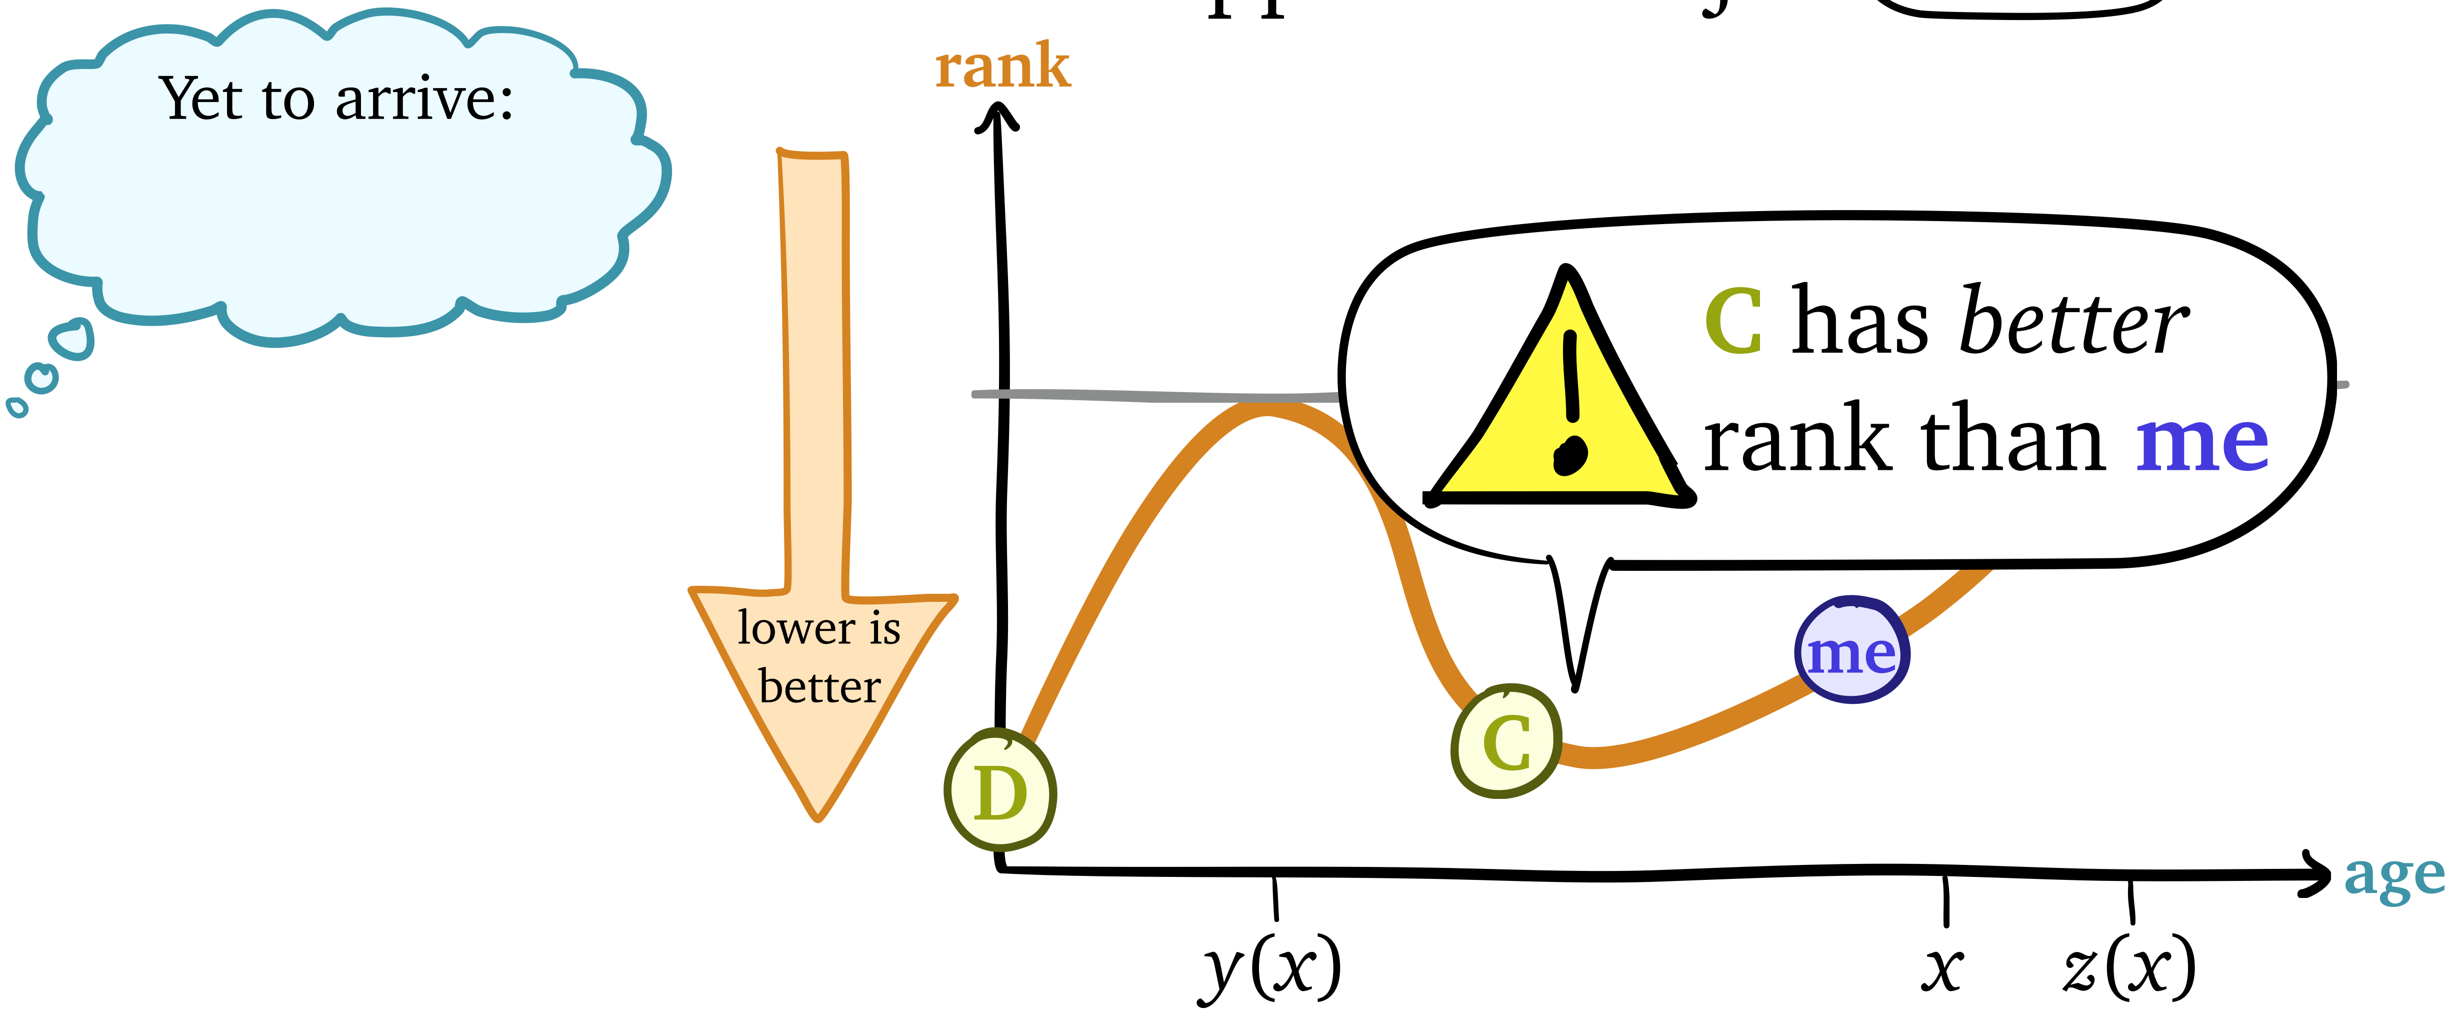

I can ignore { *old* jobs / *new* jobs (**C** & **D**) } after age { $z(x)$ / **???** }

Analyzing **Gittins-*k***

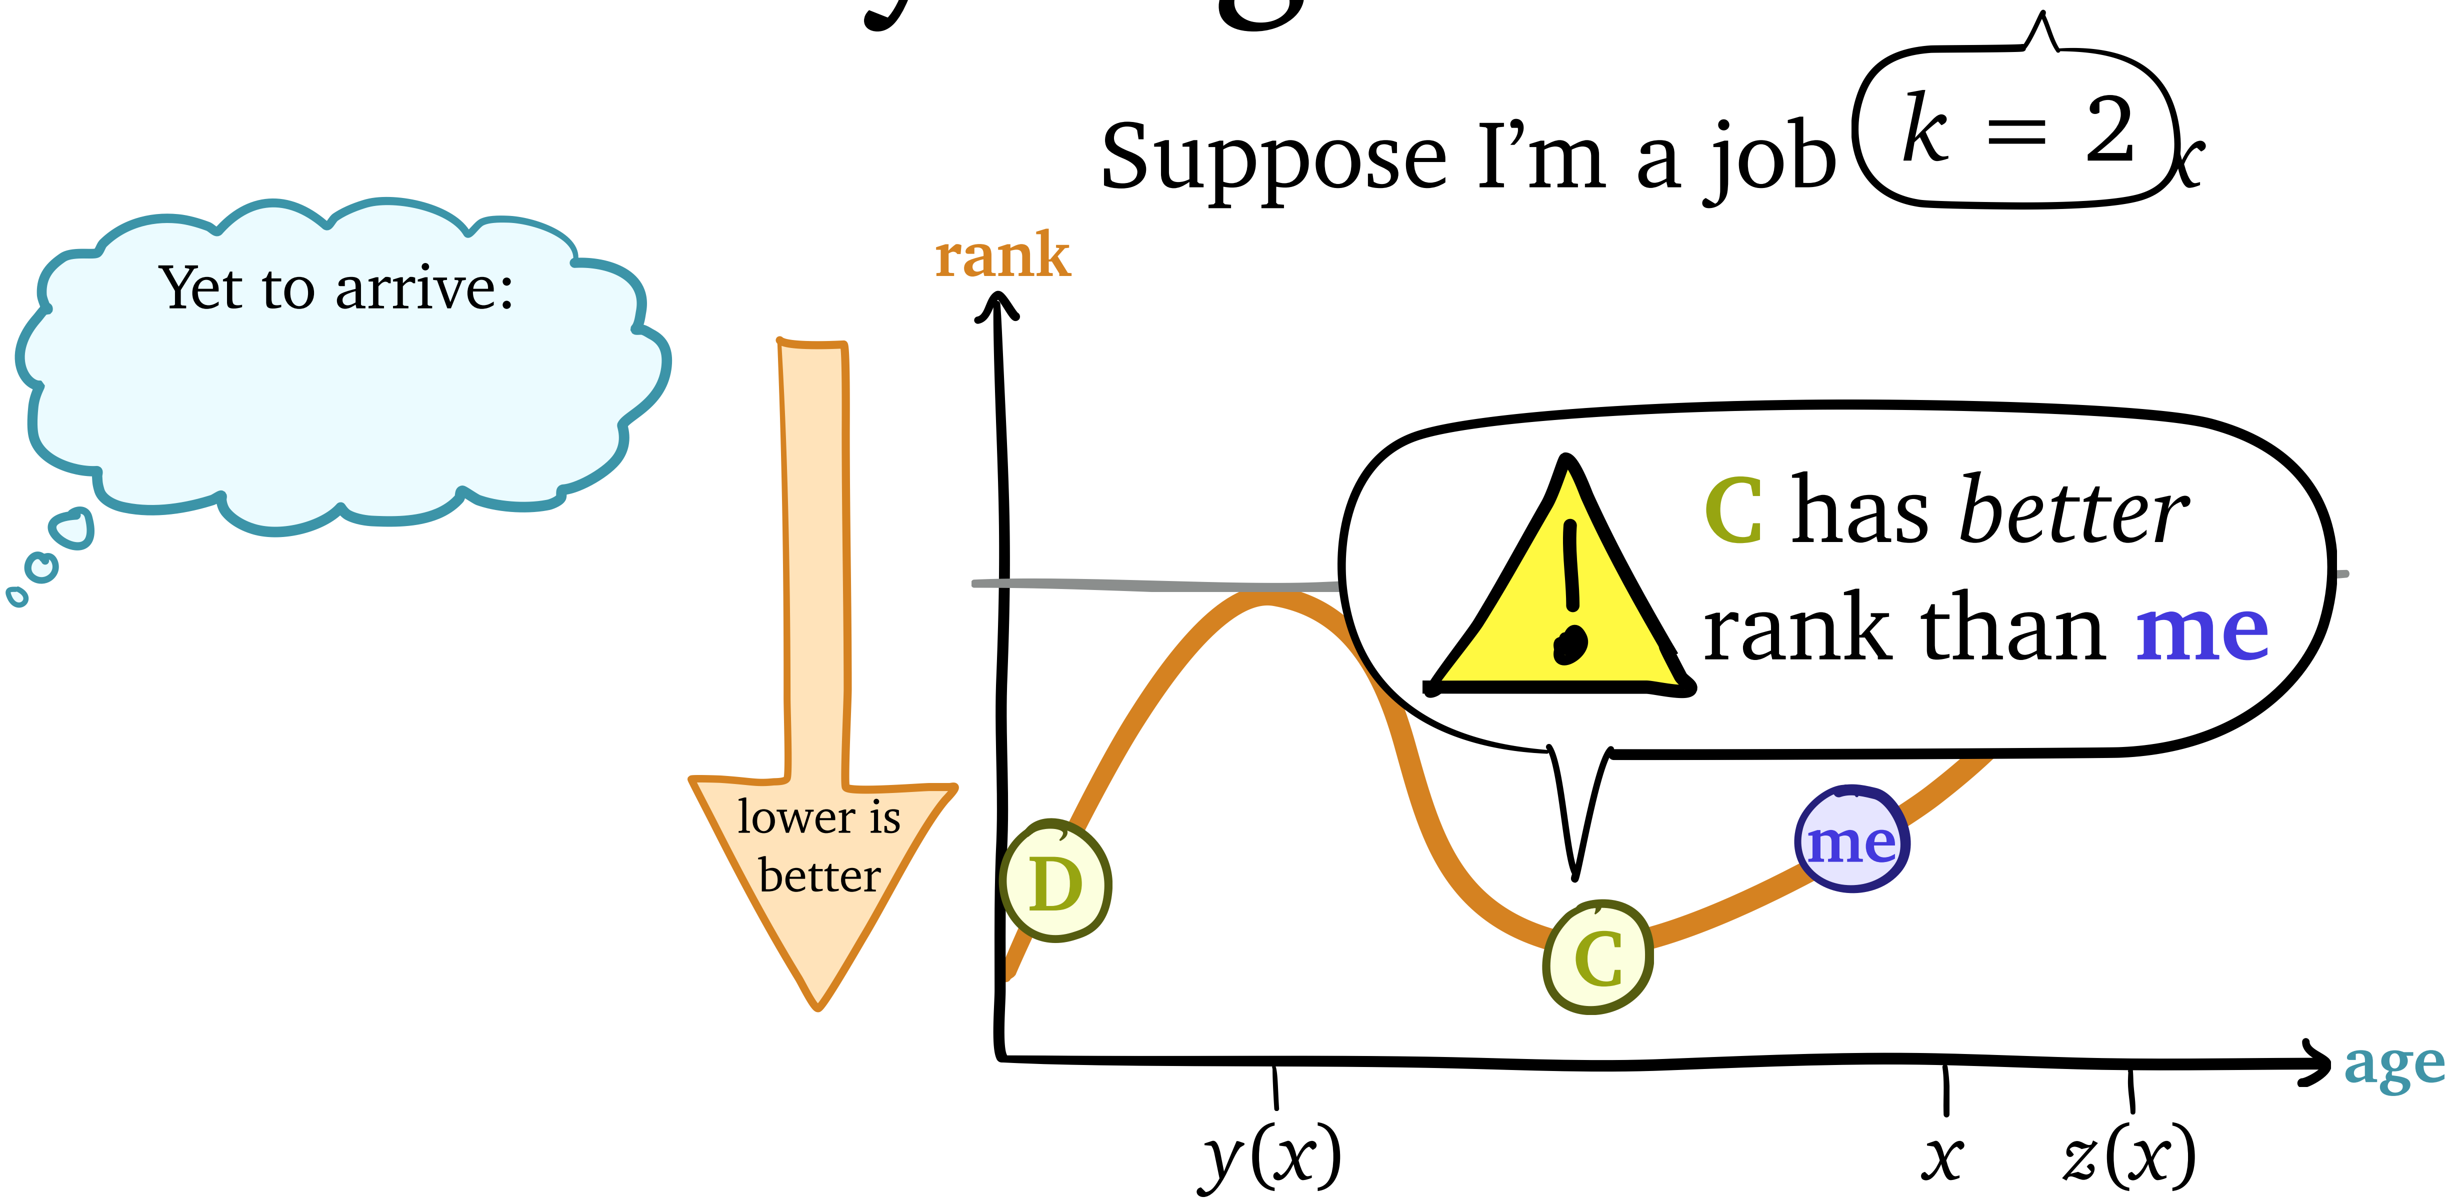

I can ignore { *old* jobs / *new* jobs (C & D) } after age { $z(x)$ / ??? }

Analyzing **Gittins-*k***

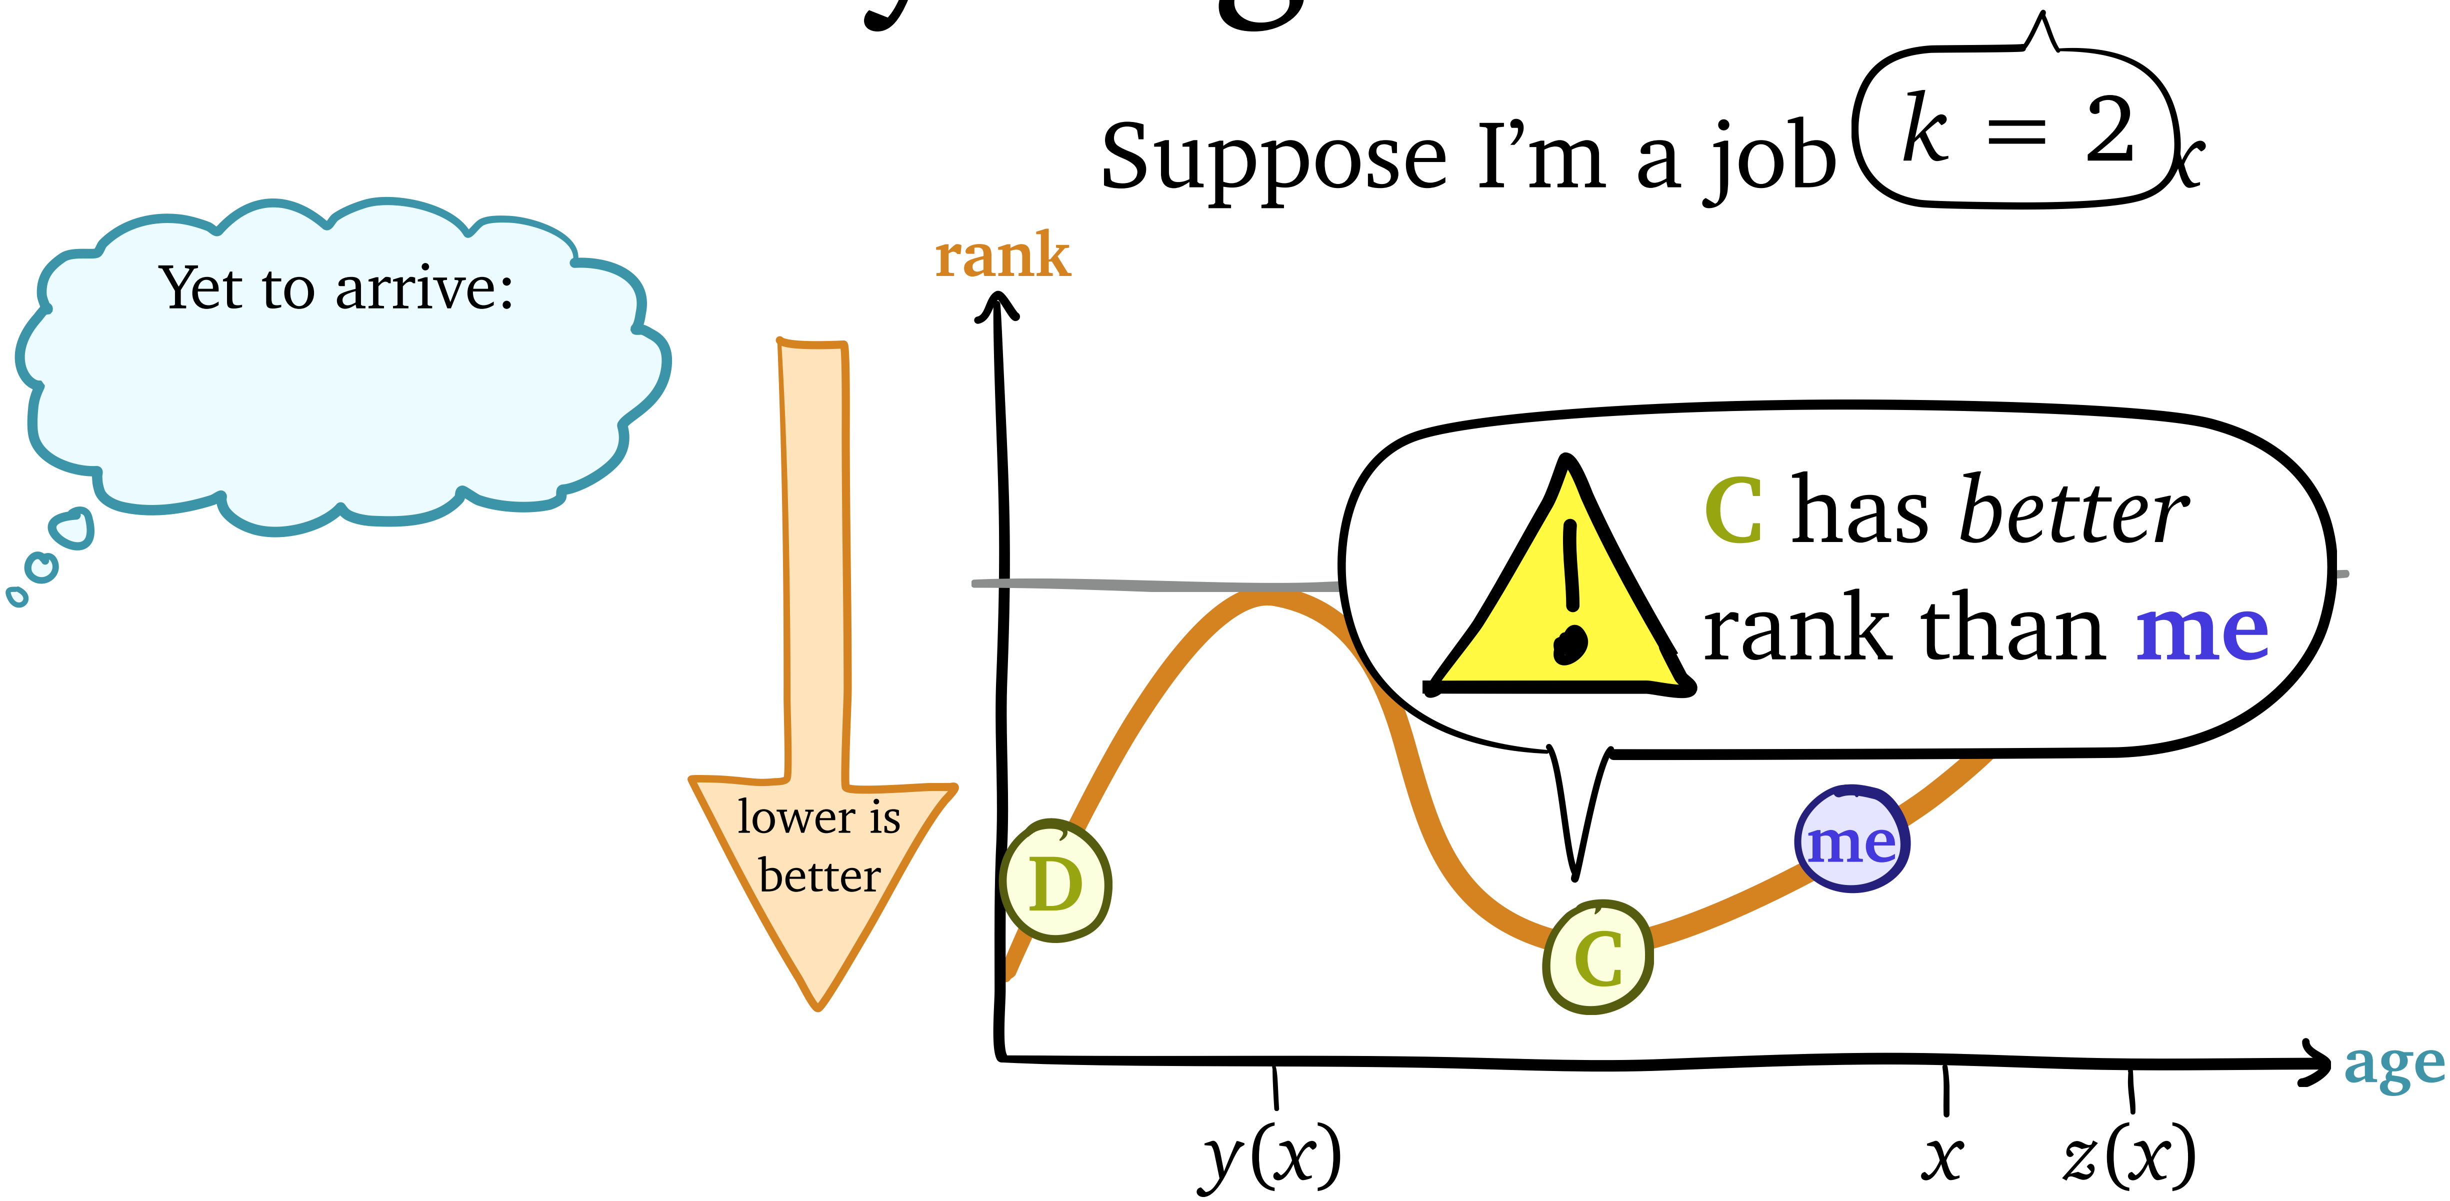

I can ignore { *old* jobs / *new* jobs (C & D) } after age { $z(x)$ / x }

Analyzing **Gittins-*k***

Suppose I'm a job $k = 2$ x

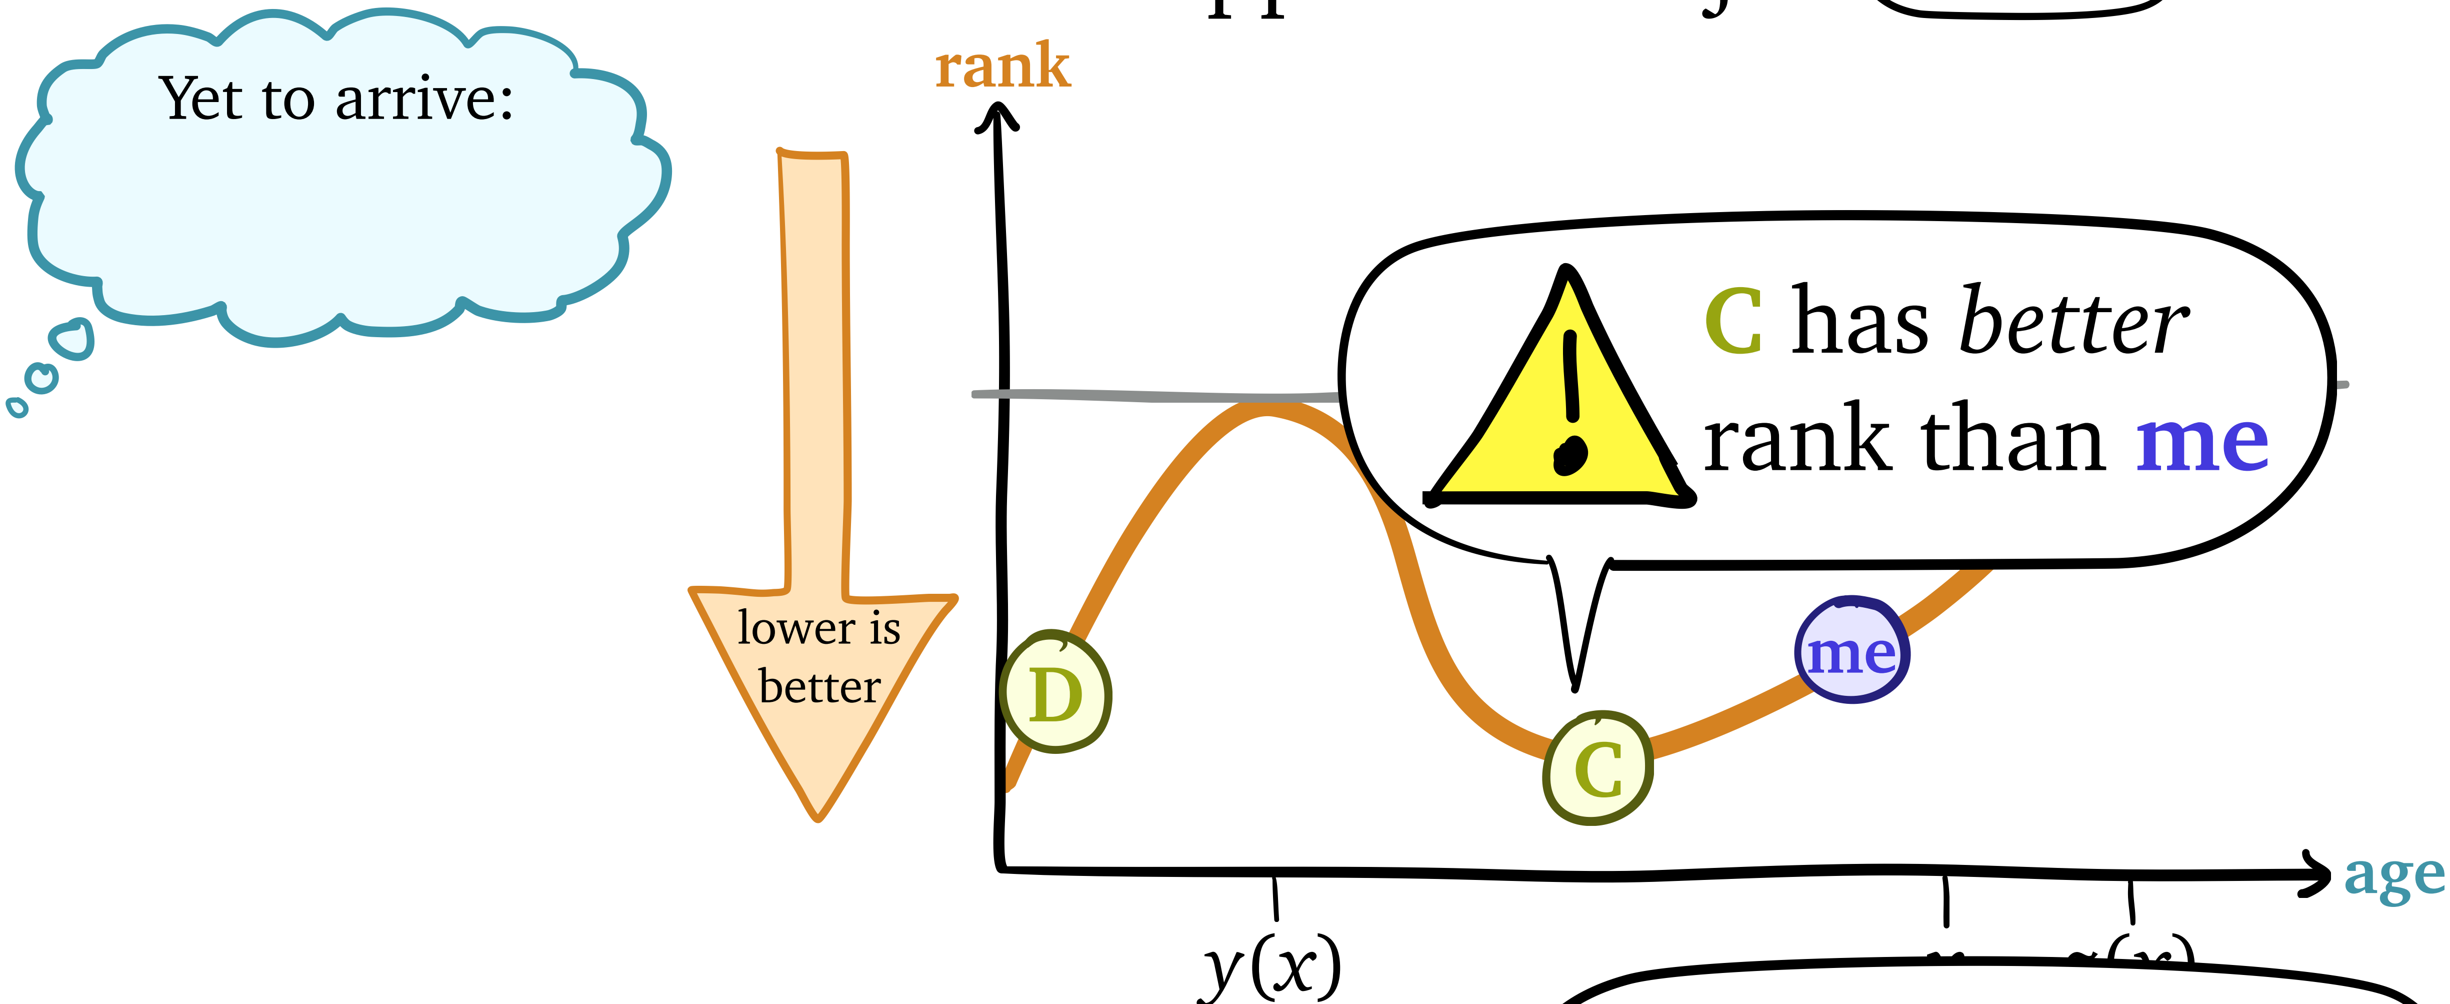

Was $y(x)$ for **Gittins-1**

I can ignore { *old* jobs / *new* jobs (C & D) } after age { x }

Need a version of **Gittins**
without waves

New Policy: **M-Gittins**

New Policy: M-Gittins

monotonic

New Policy: M-Gittins

monotonic

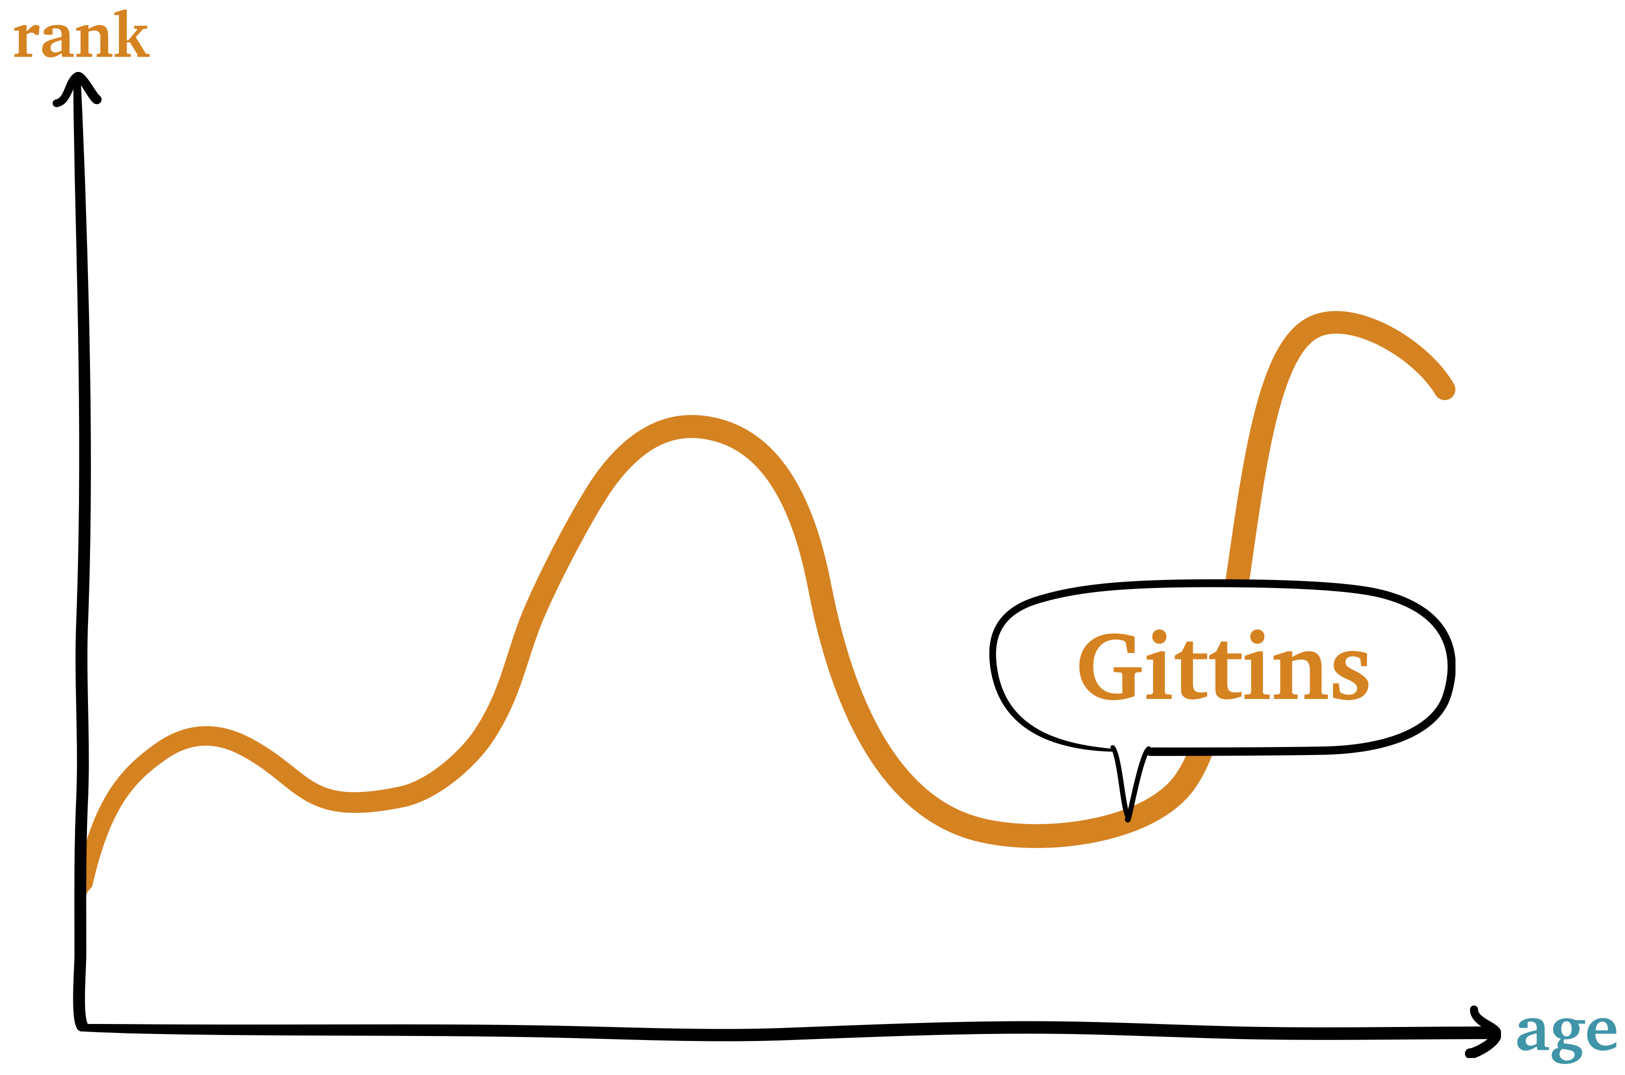

New Policy: M-Gittins

monotonic

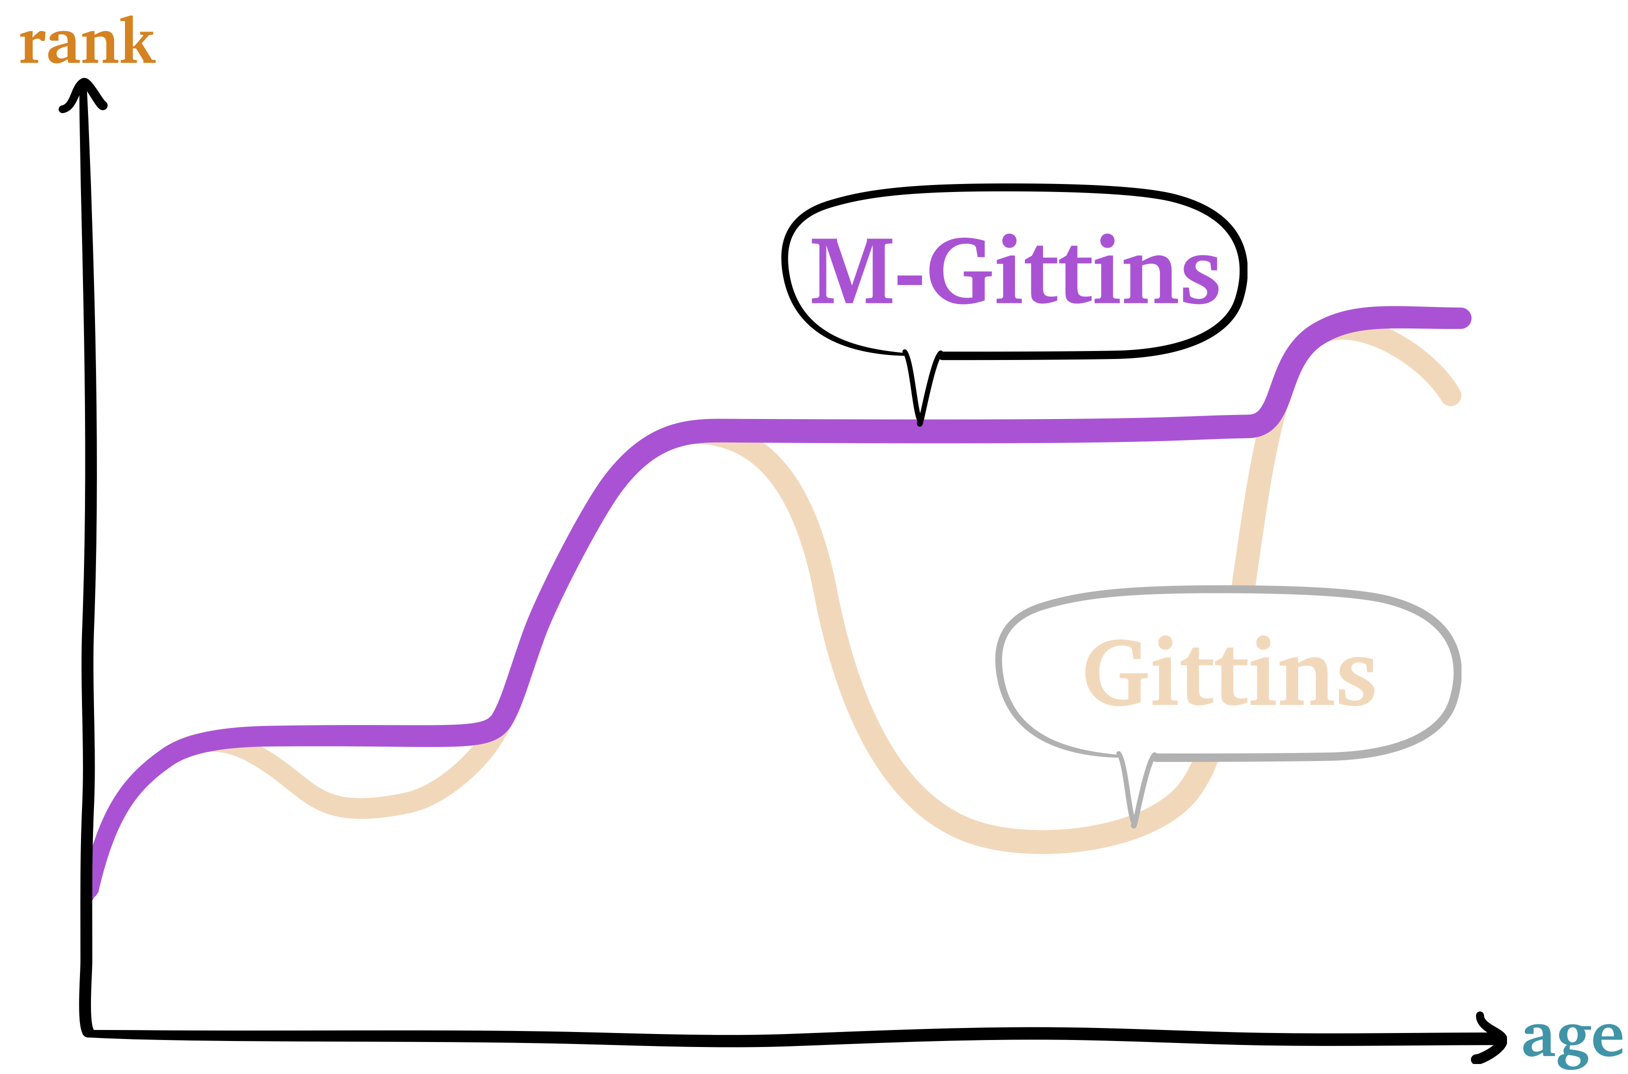

New Policy: M-Gittins

monotonic

$$r_{\text{M-Gittins}}(a) = \max_{0 \le b \le a} r_{\text{Gittins}}(b)$$

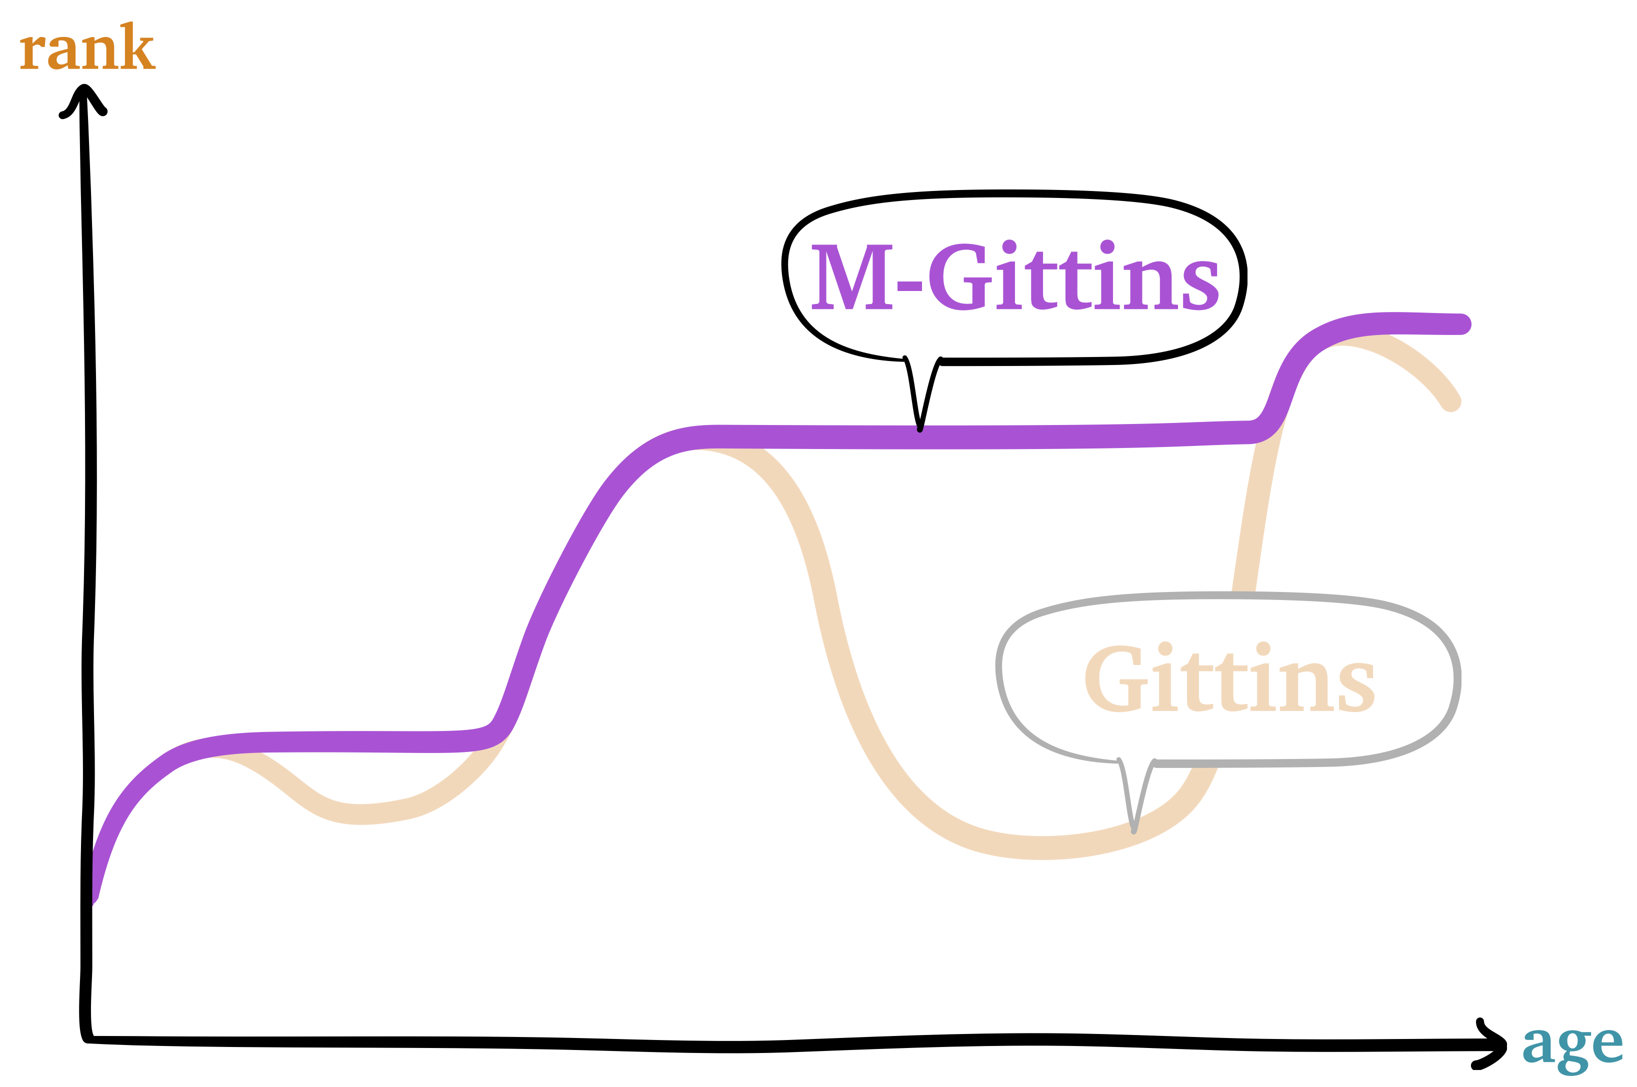

M-Gittins-*k* Saves the Day

Suppose I'm a job of size x

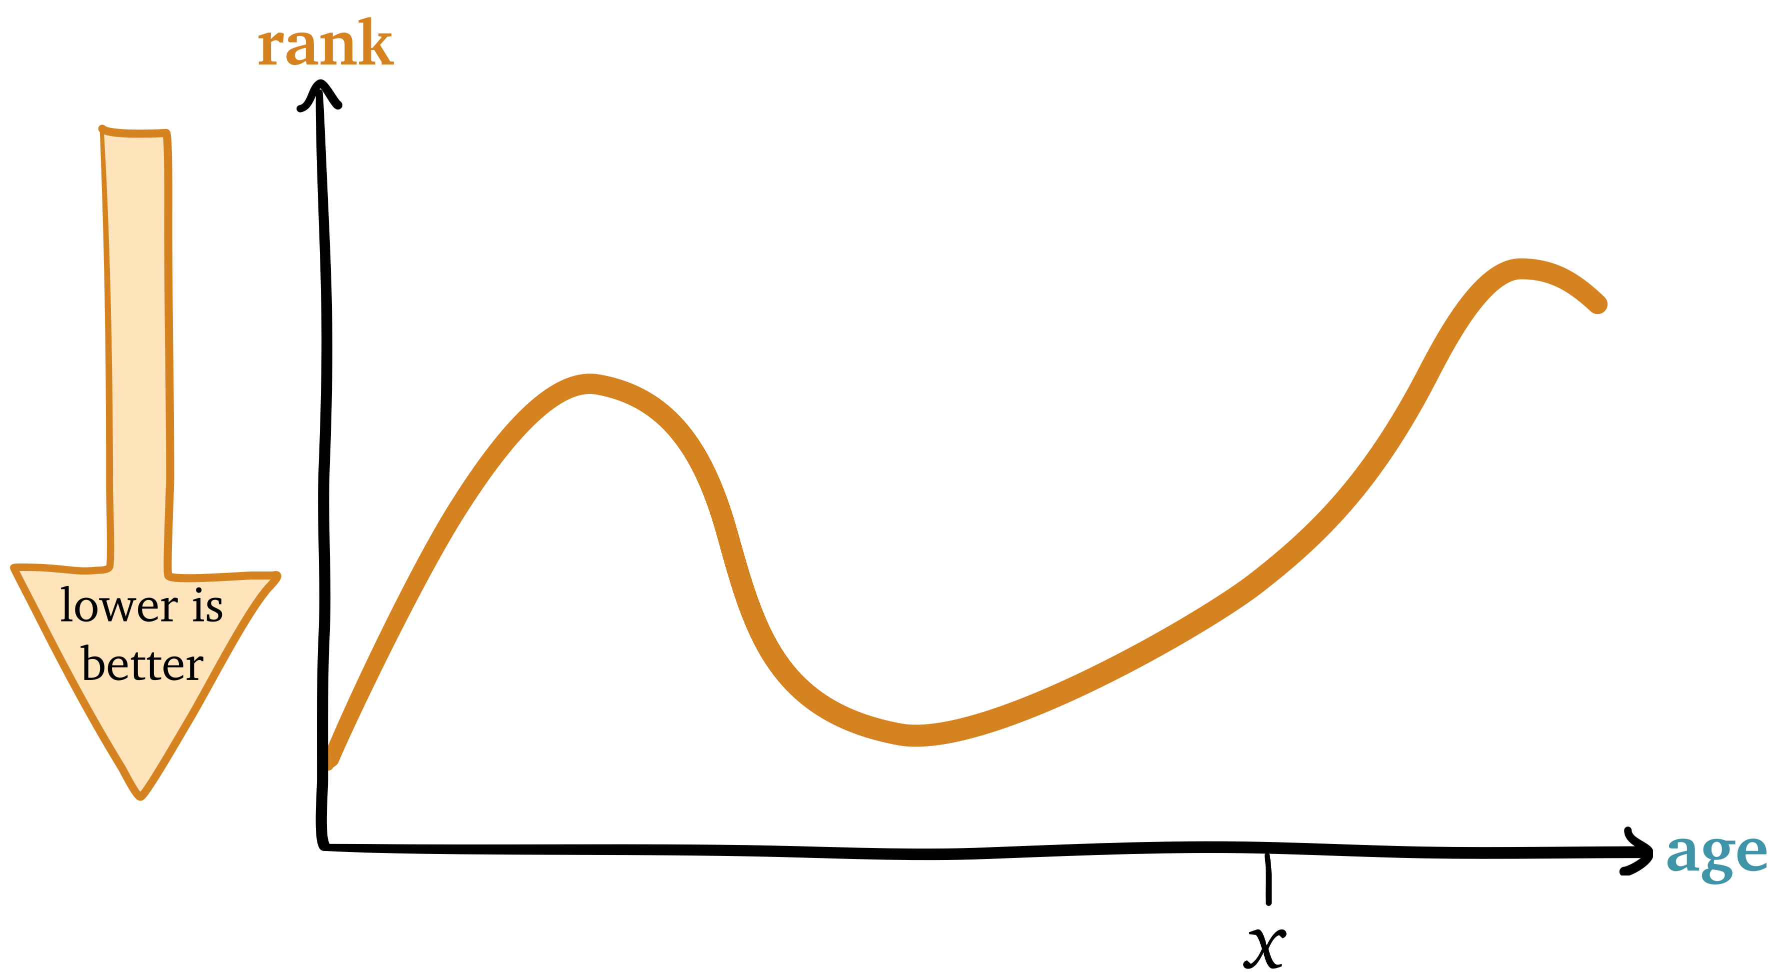

M-Gittins-*k* Saves the Day

Suppose I'm a job of size x

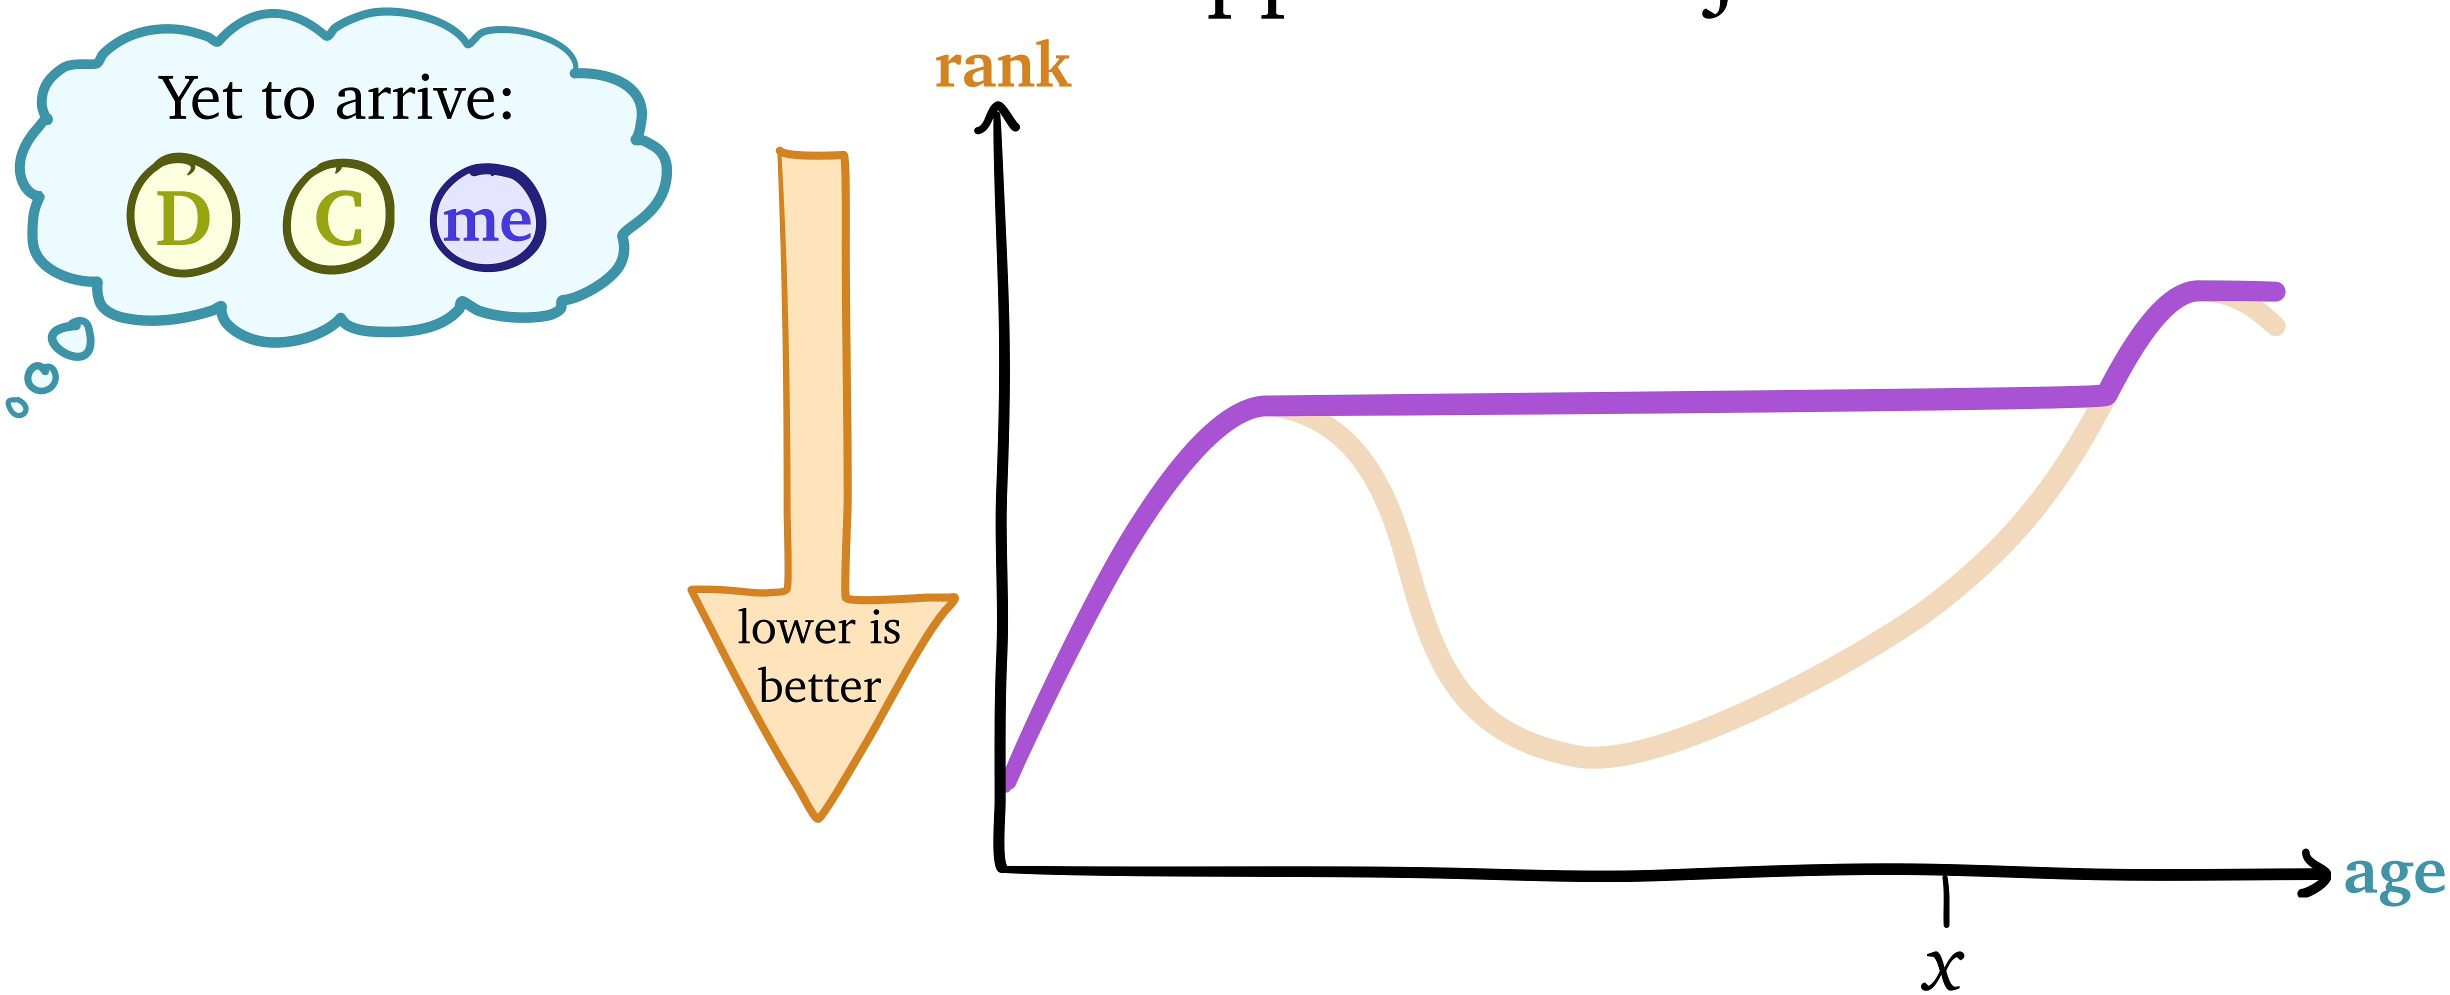

M-Gittins-*k* Saves the Day

Suppose I'm a job of size x

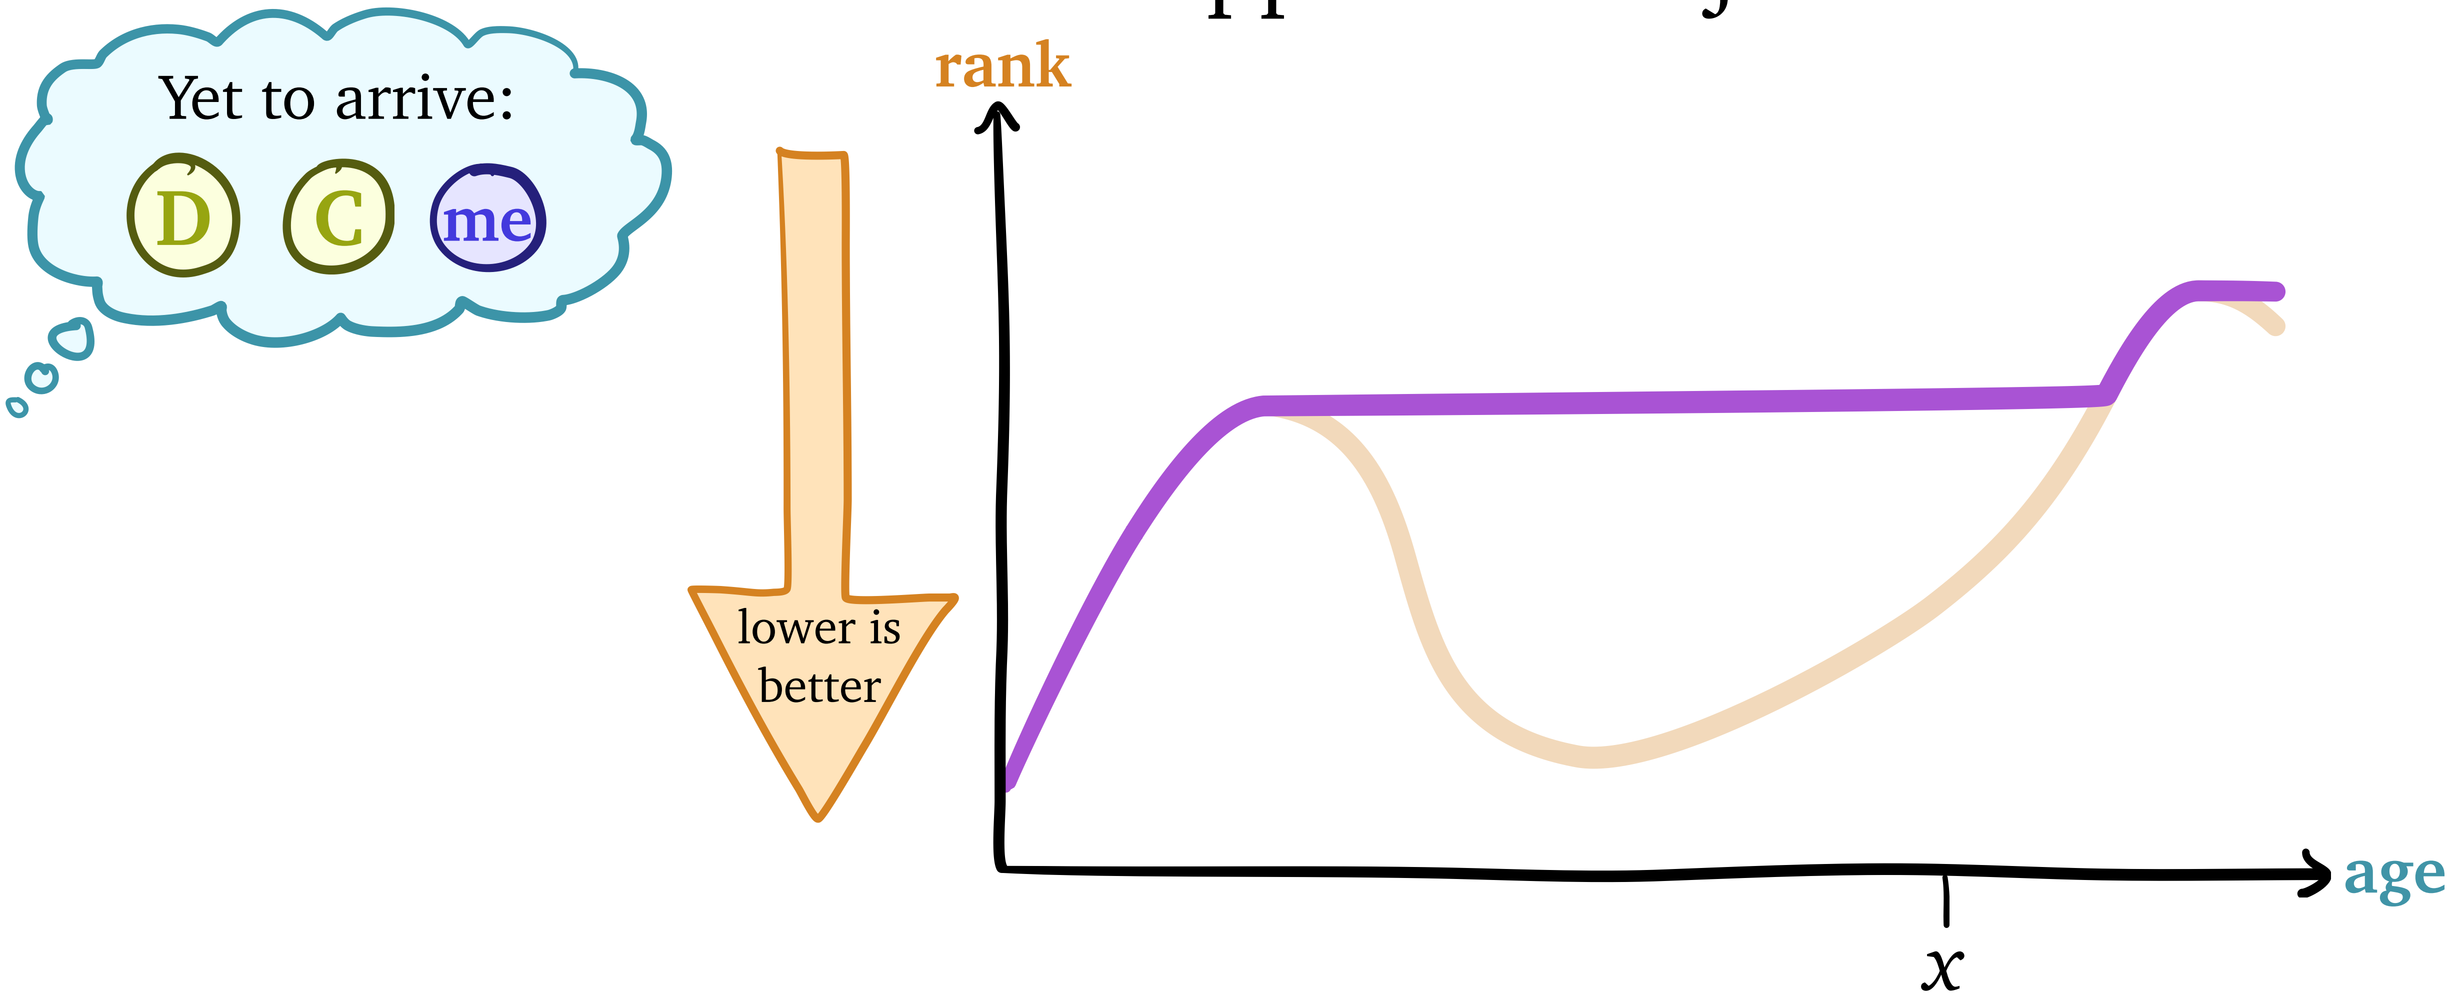

I can ignore { *old* jobs / *new* jobs (C & D) } after age { ??? / ??? }

M-Gittins-*k* Saves the Day

$k = 2$

ppose I’m a job of size x

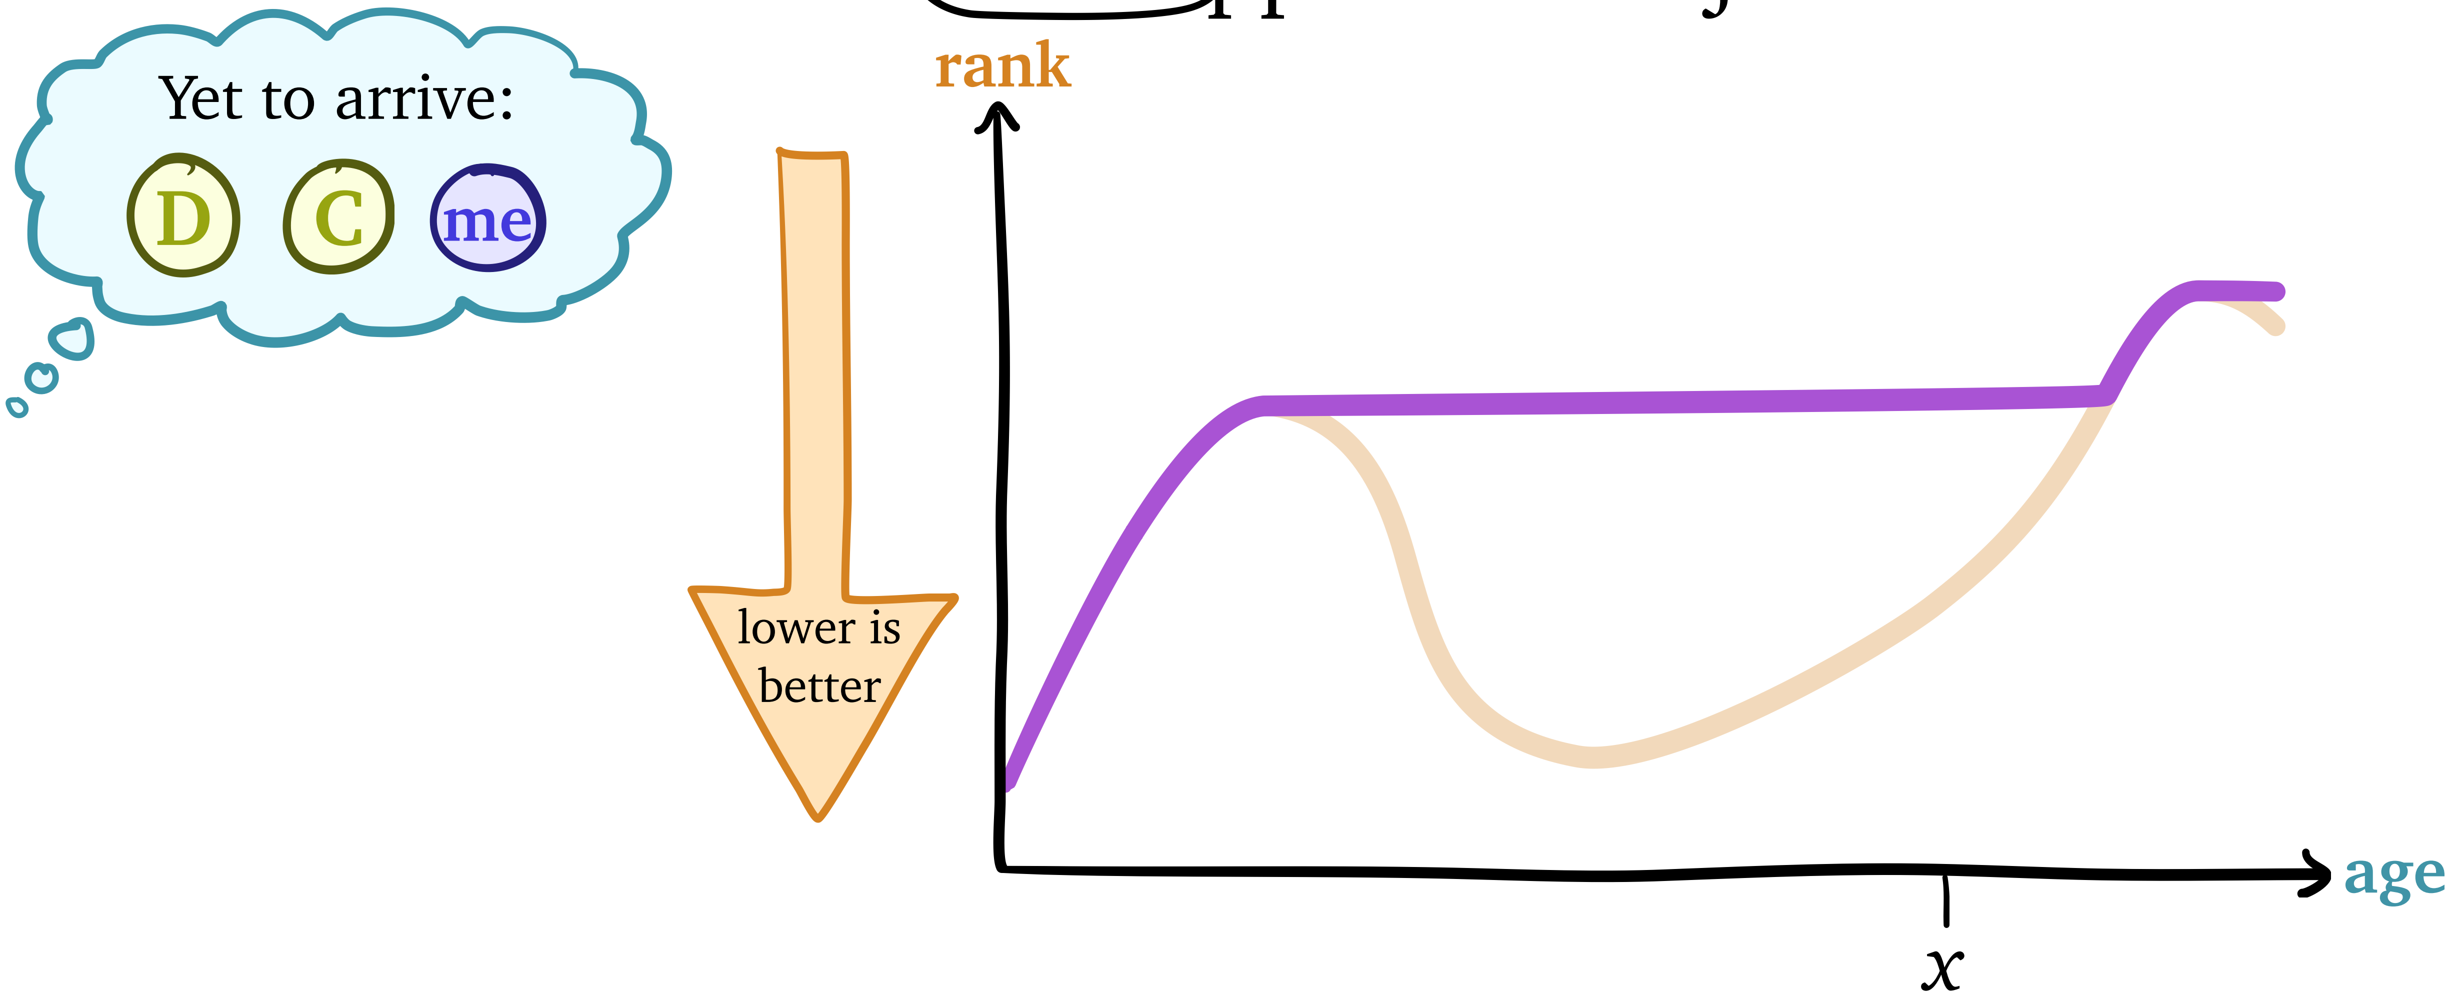

I can ignore { *old* jobs / *new* jobs (C & D) } after age { ??? / ??? }

M-Gittins-k Saves the Day

ppose I'm a job of size x

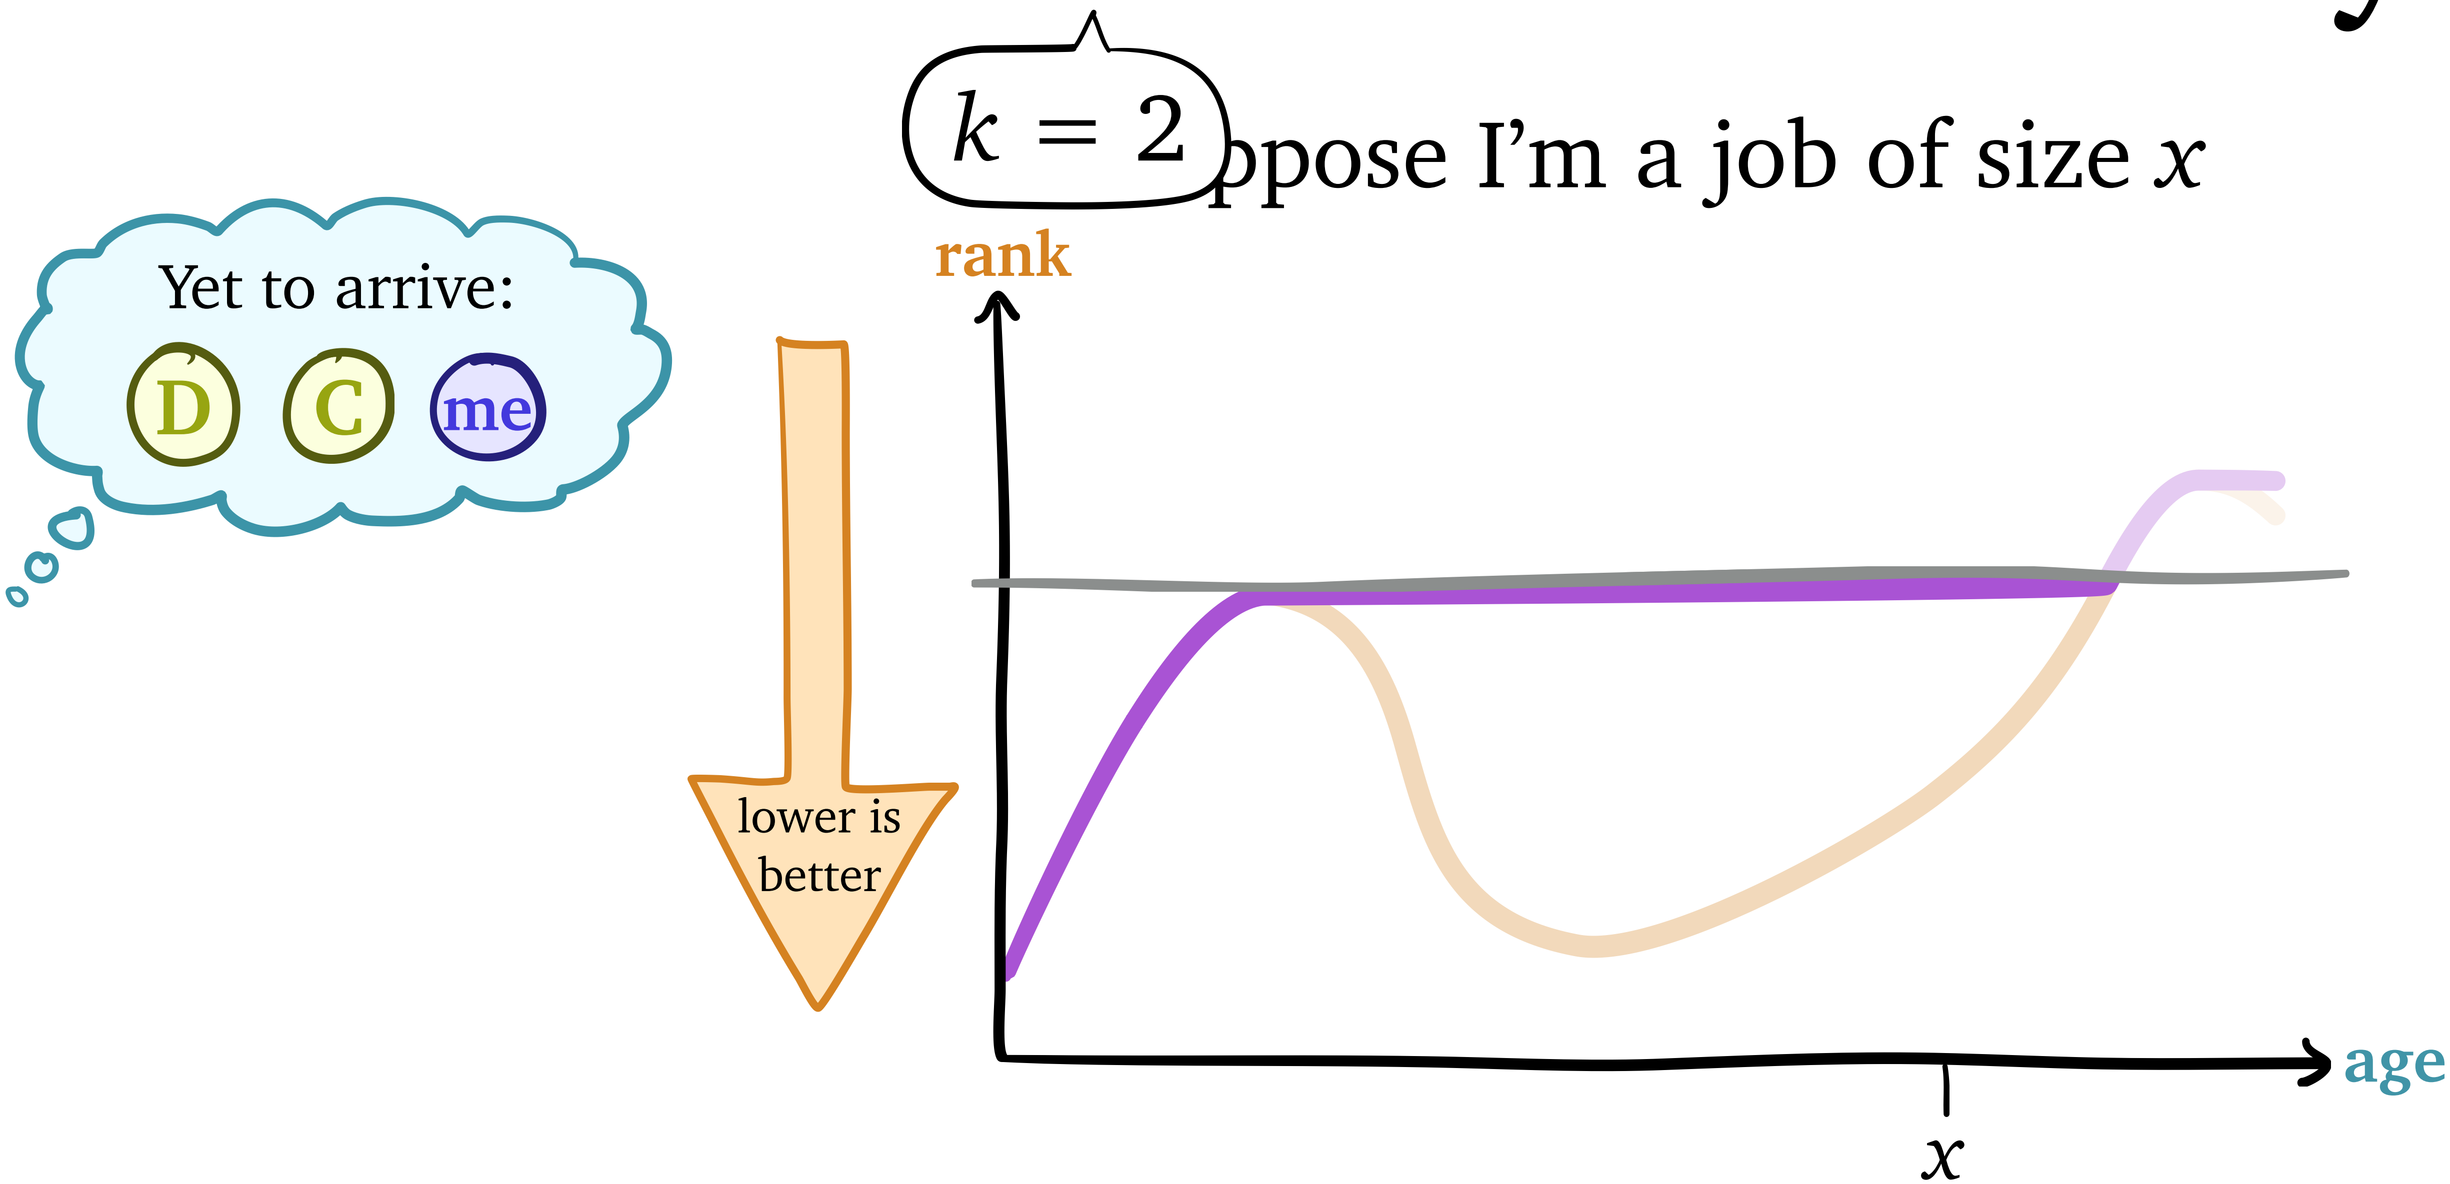

I can ignore { *old* jobs / *new* jobs (C & D) } after age { ??? / ??? }

M-Gittins-k Saves the Day

$k = 2$

ppose I’m a job of size x

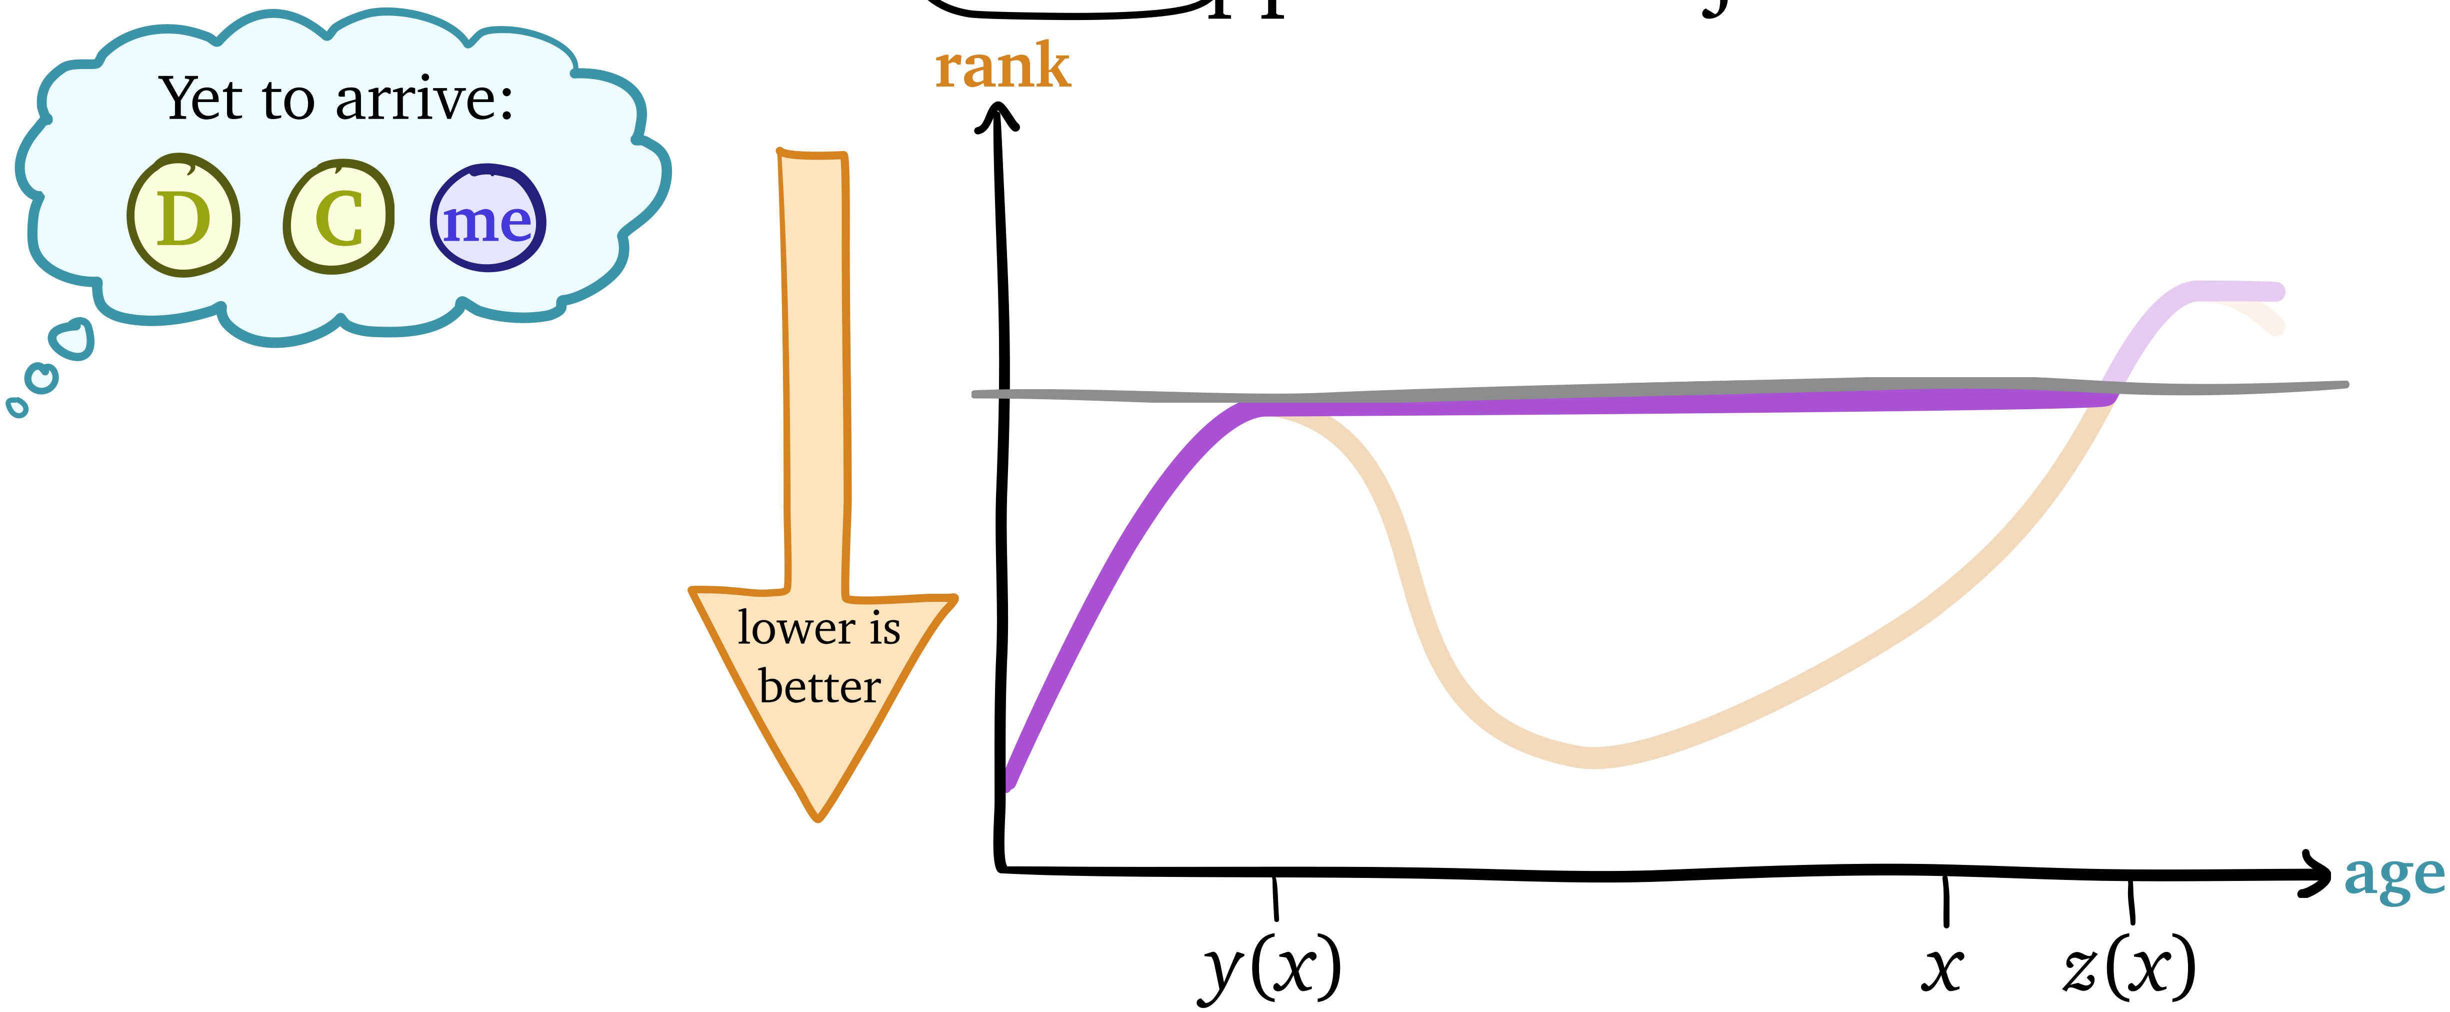

I can ignore { *old* jobs / *new* jobs (C & D) } after age { ??? / ??? }

M-Gittins-*k* Saves the Day

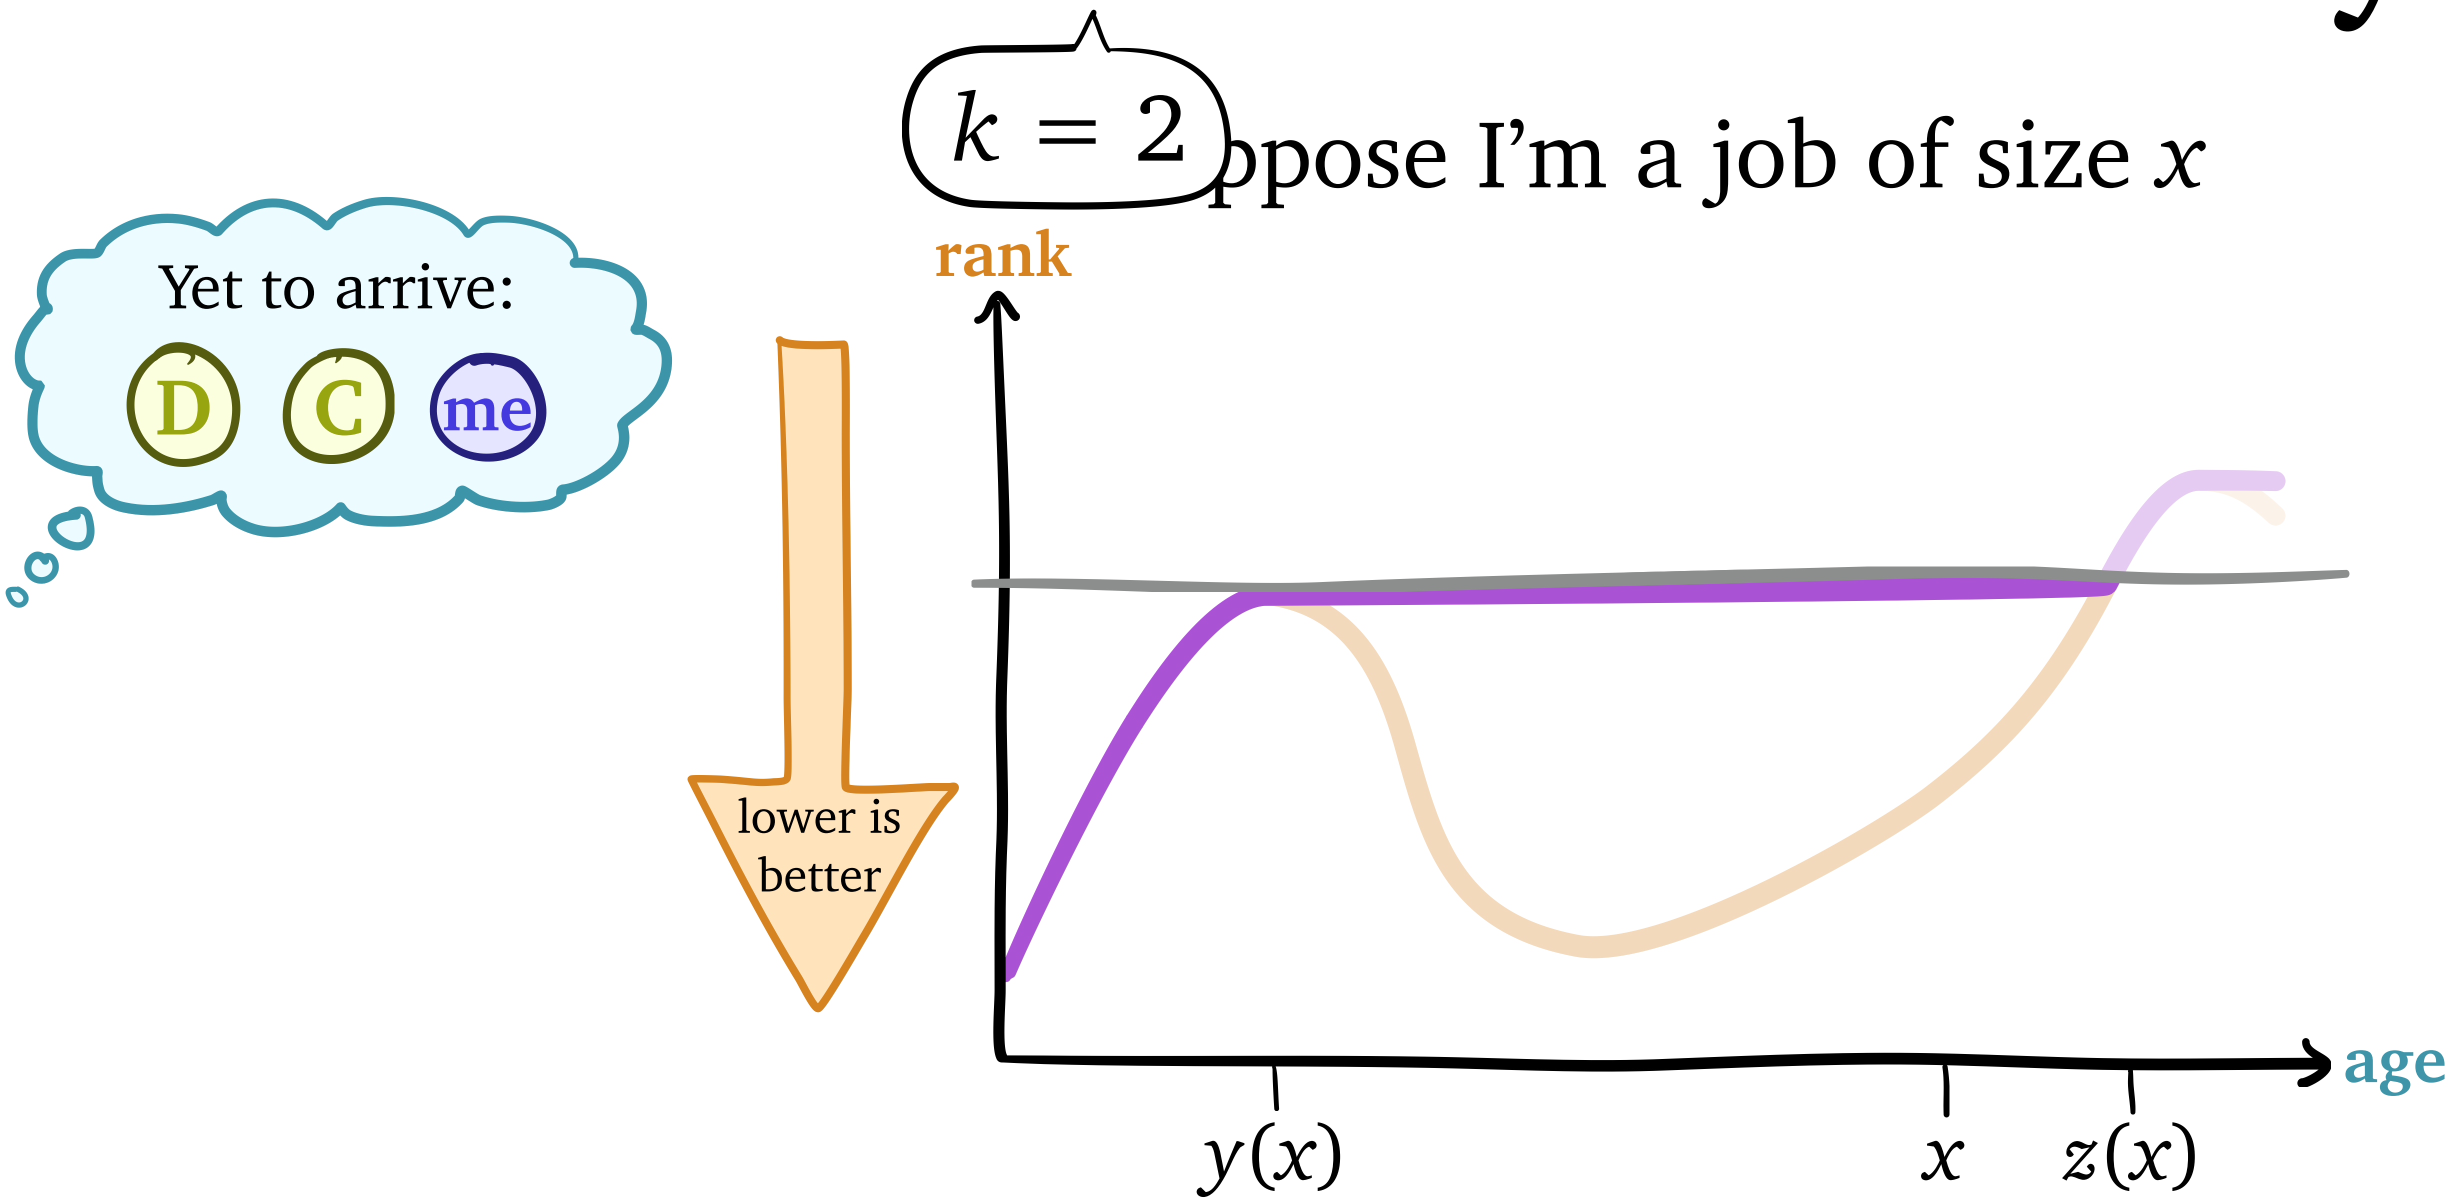

I can ignore { *old* jobs / *new* jobs (C & D) } after age { $z(x)$ / ??? }

M-Gittins-k Saves the Day

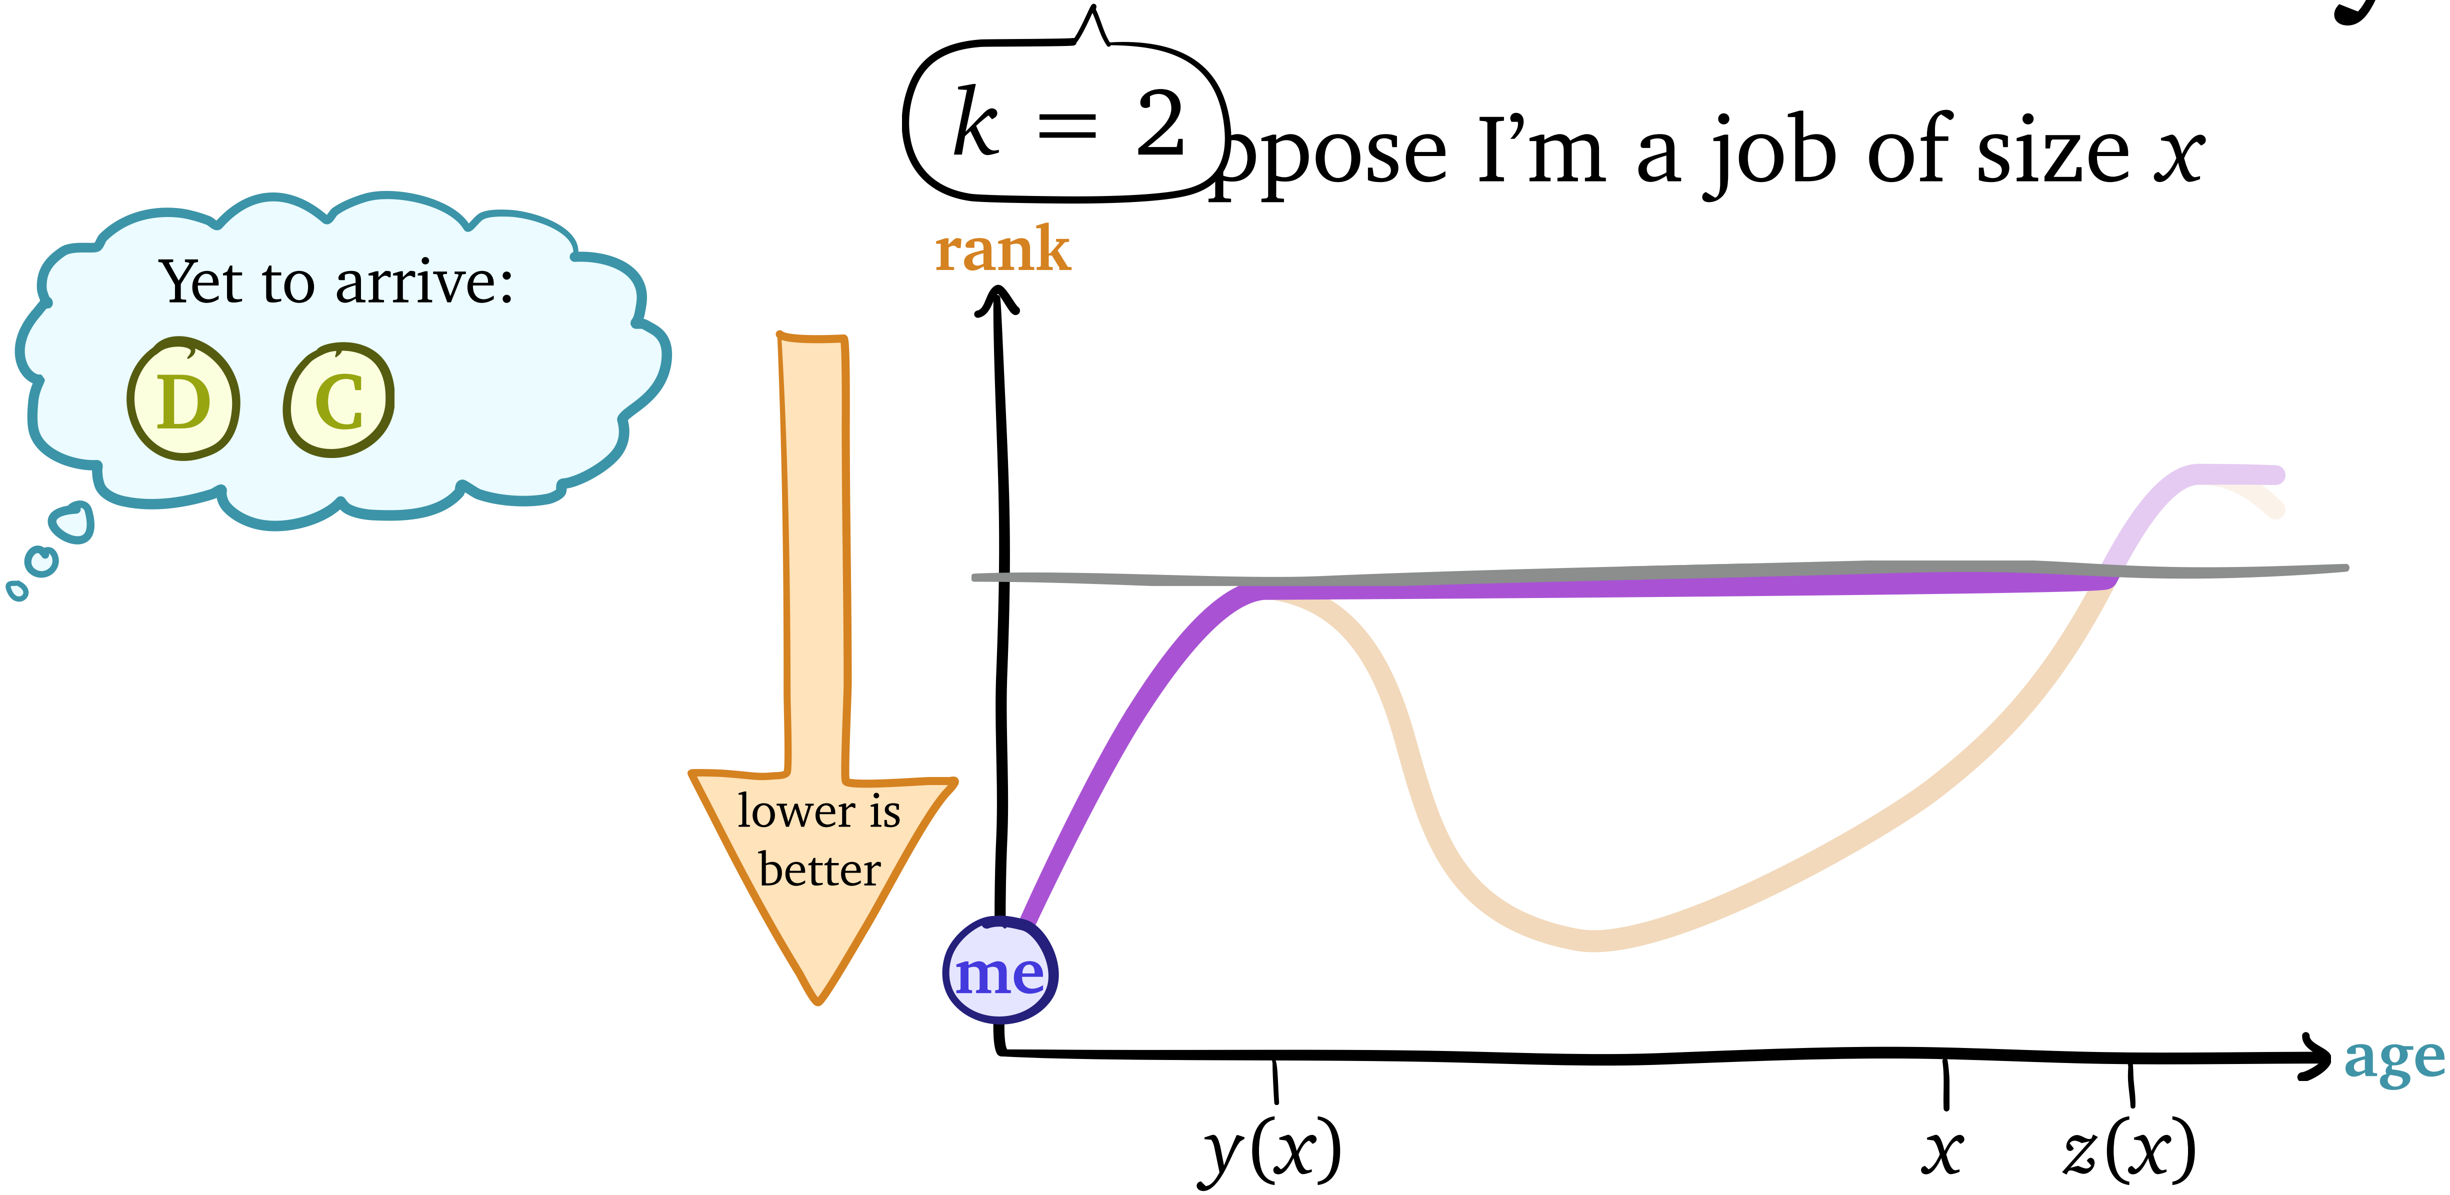

I can ignore $\left\{ \begin{matrix} \textit{old} \text{ jobs} \\ \textit{new} \text{ jobs (C \& D)} \end{matrix} \right\}$ after age $\left\{ \begin{matrix} z(x) \\ ??? \end{matrix} \right\}$

M-Gittins-*k* Saves the Day

$k = 2$

…ppose I'm a job of size x

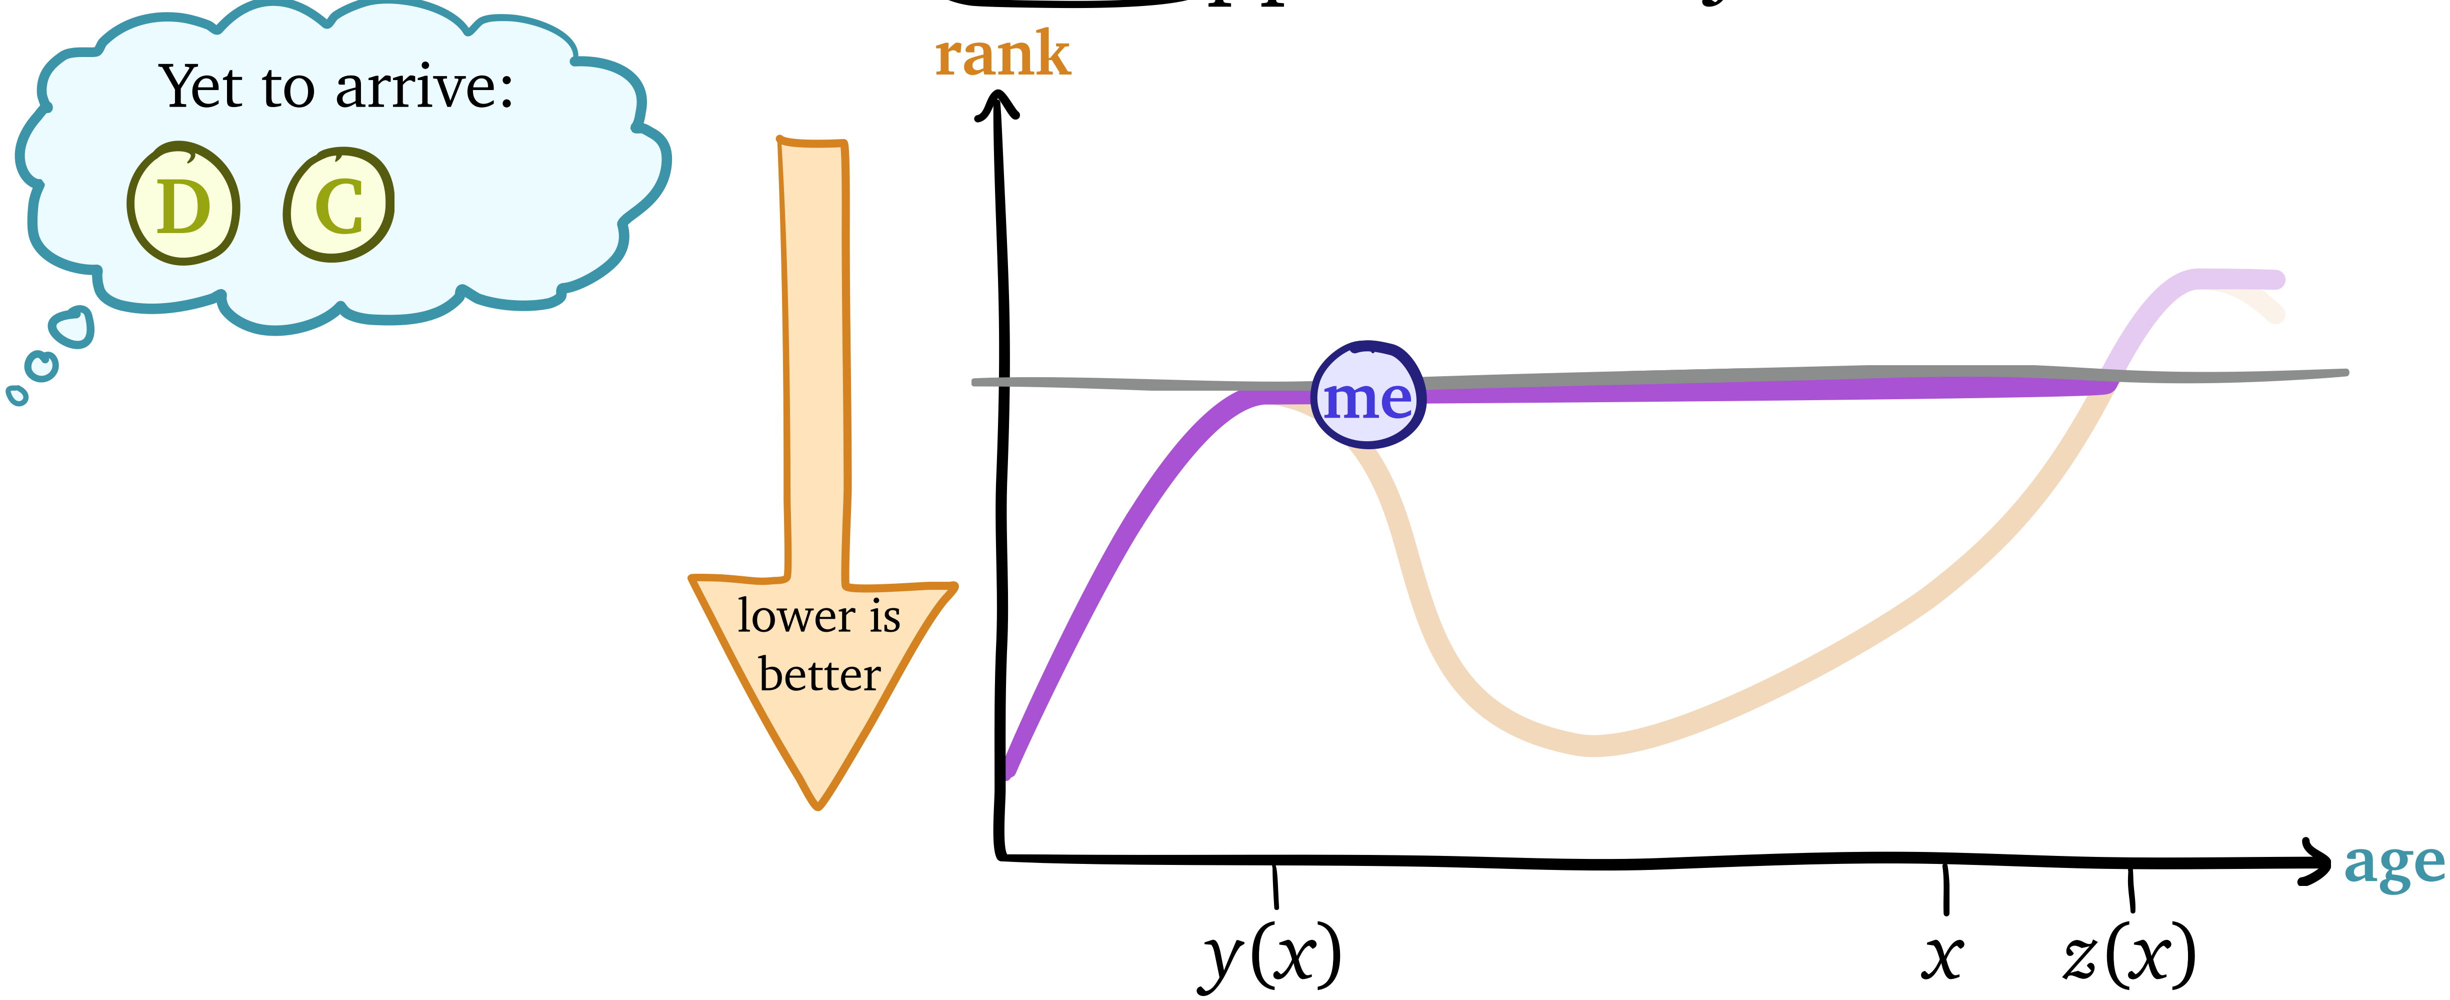

I can ignore { *old* jobs / *new* jobs (C & D) } after age { $z(x)$ / ??? }

M-Gittins-k Saves the Day

$k = 2$

…ppose I'm a job of size x

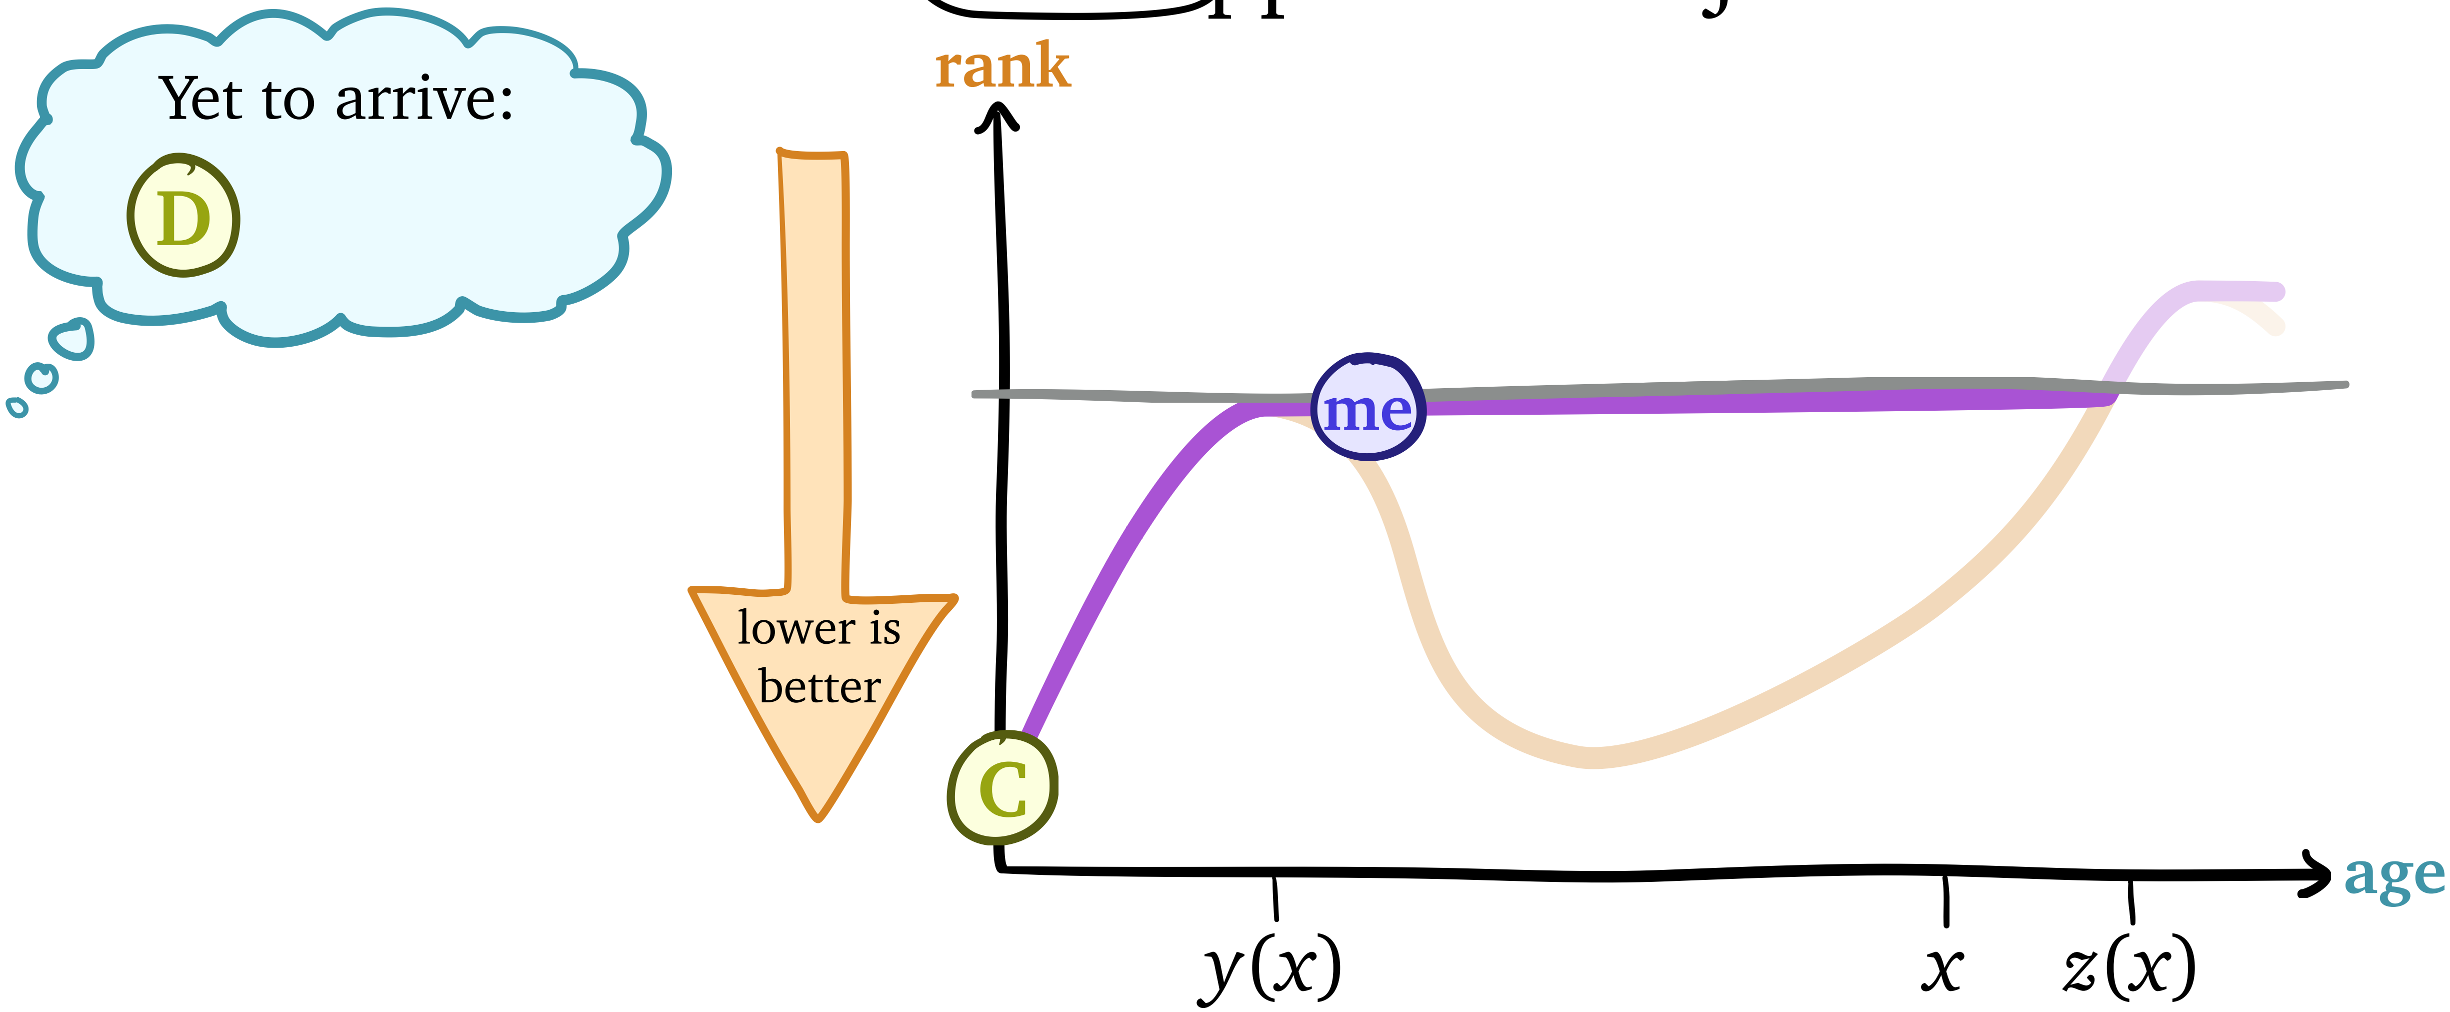

I can ignore { *old* jobs / *new* jobs (C & D) } after age { $z(x)$ / ??? }

M-Gittins-k Saves the Day

$k = 2$

ppose I’m a job of size x

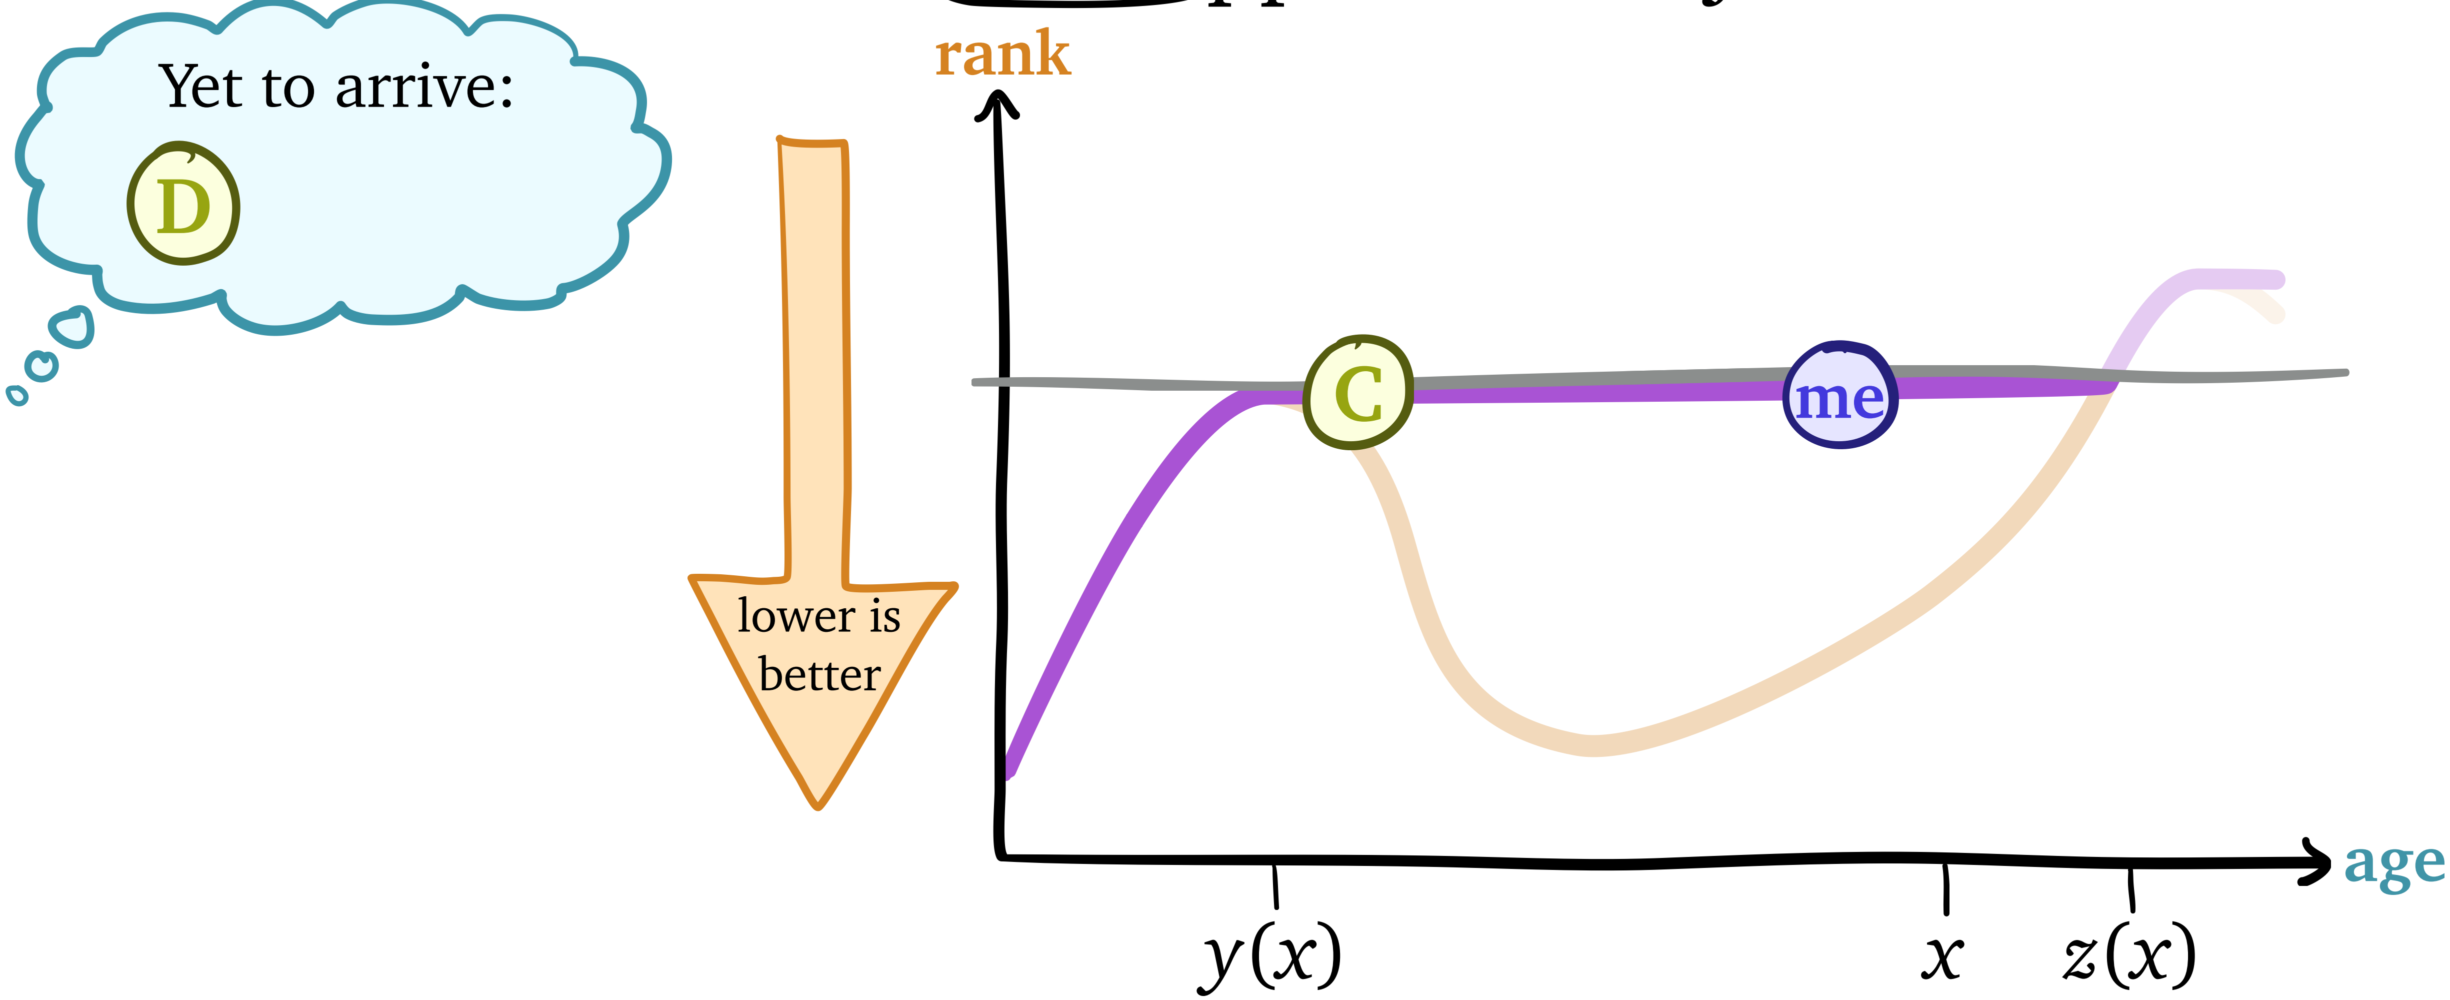

I can ignore *old* jobs after age $z(x)$

I can ignore *new* jobs (C & D) after age ???

M-Gittins-k Saves the Day

$k = 2$

ppose I’m a job of size x

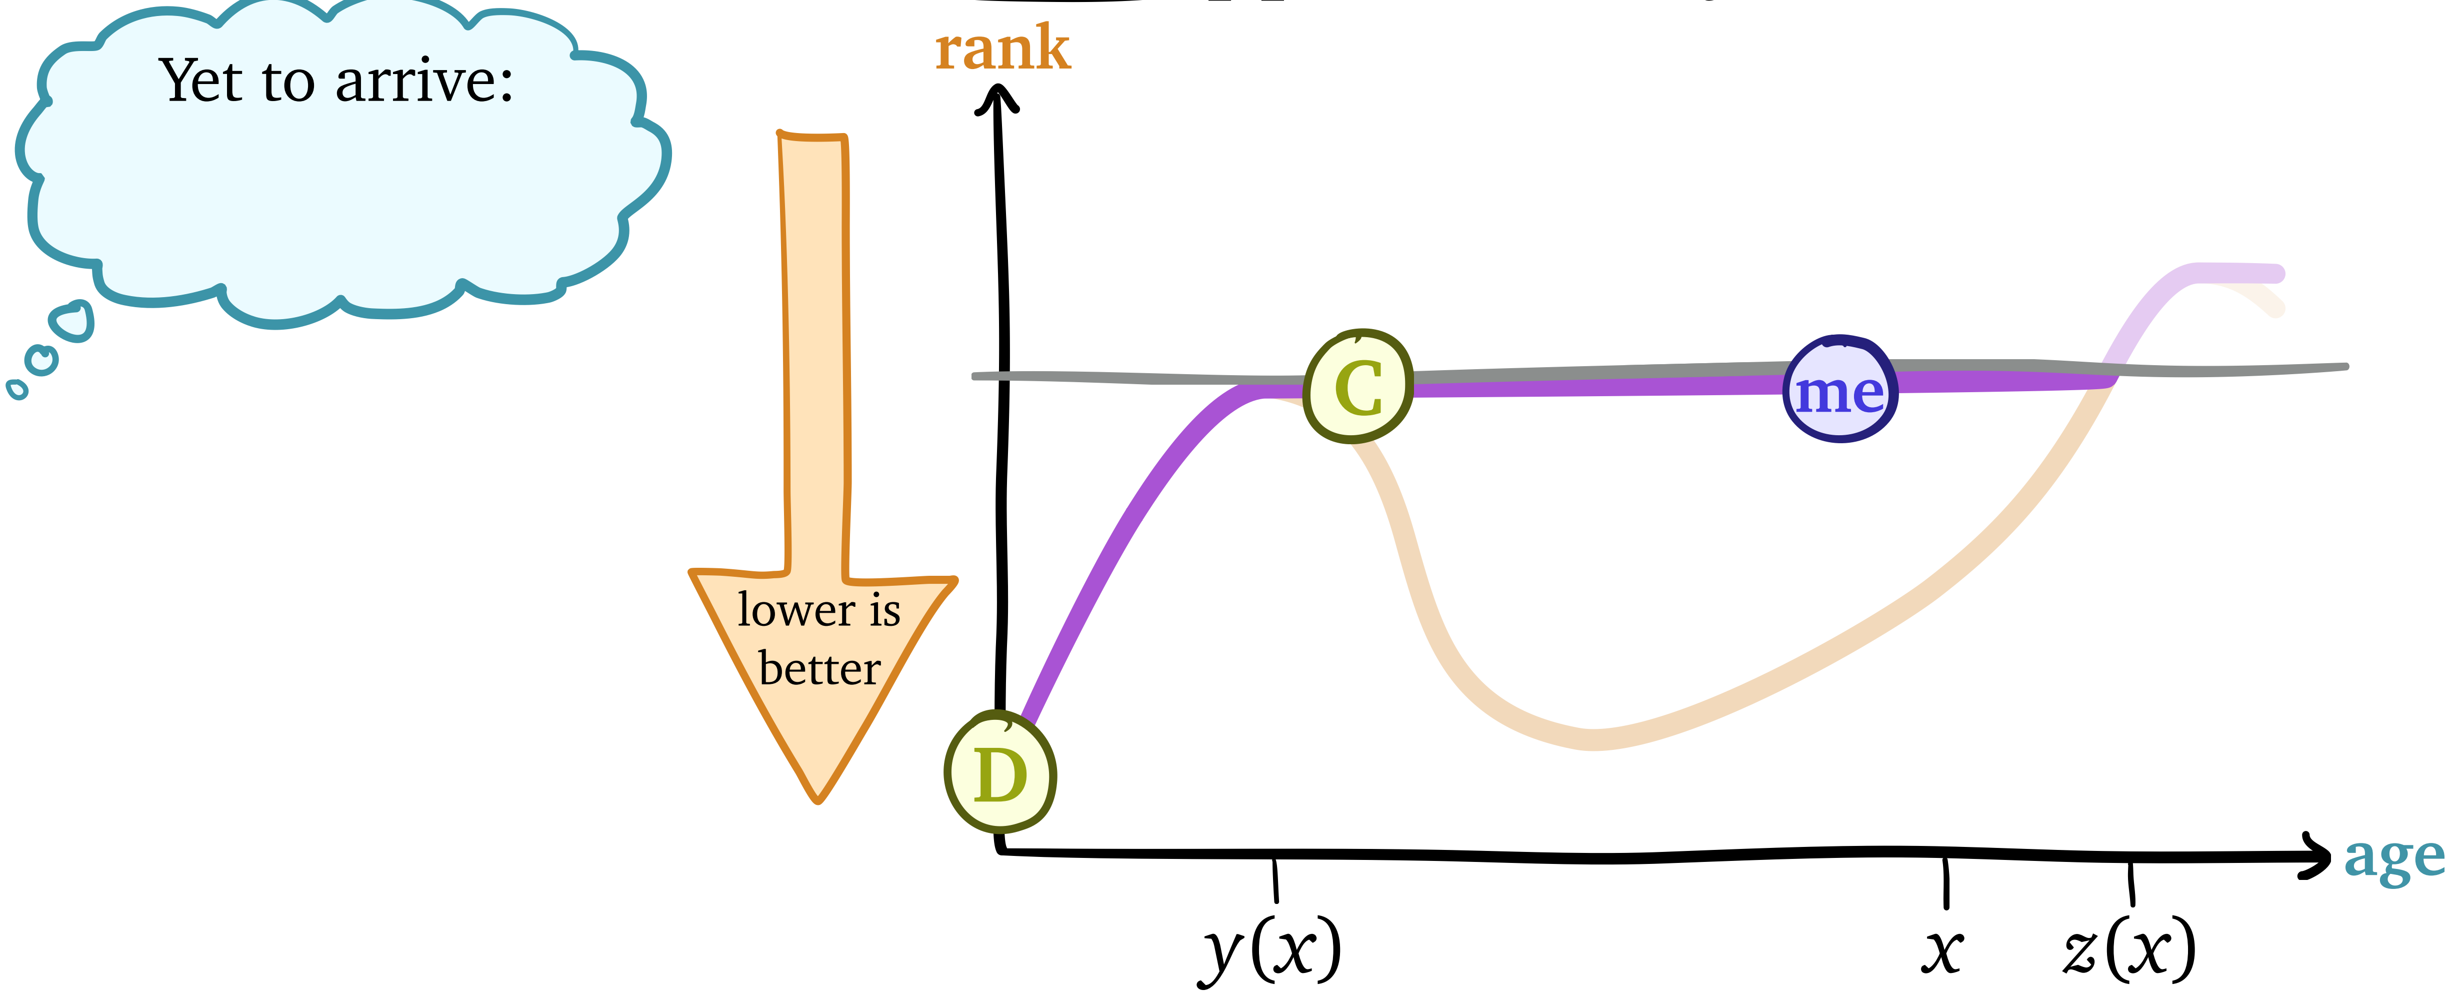

I can ignore $\left\{ \begin{array}{c} \textit{old} \text{ jobs} \\ \textit{new} \text{ jobs (C \& D)} \end{array} \right\}$ after age $\left\{ \begin{array}{c} z(x) \\ ??? \end{array} \right\}$

M-Gittins-k Saves the Day

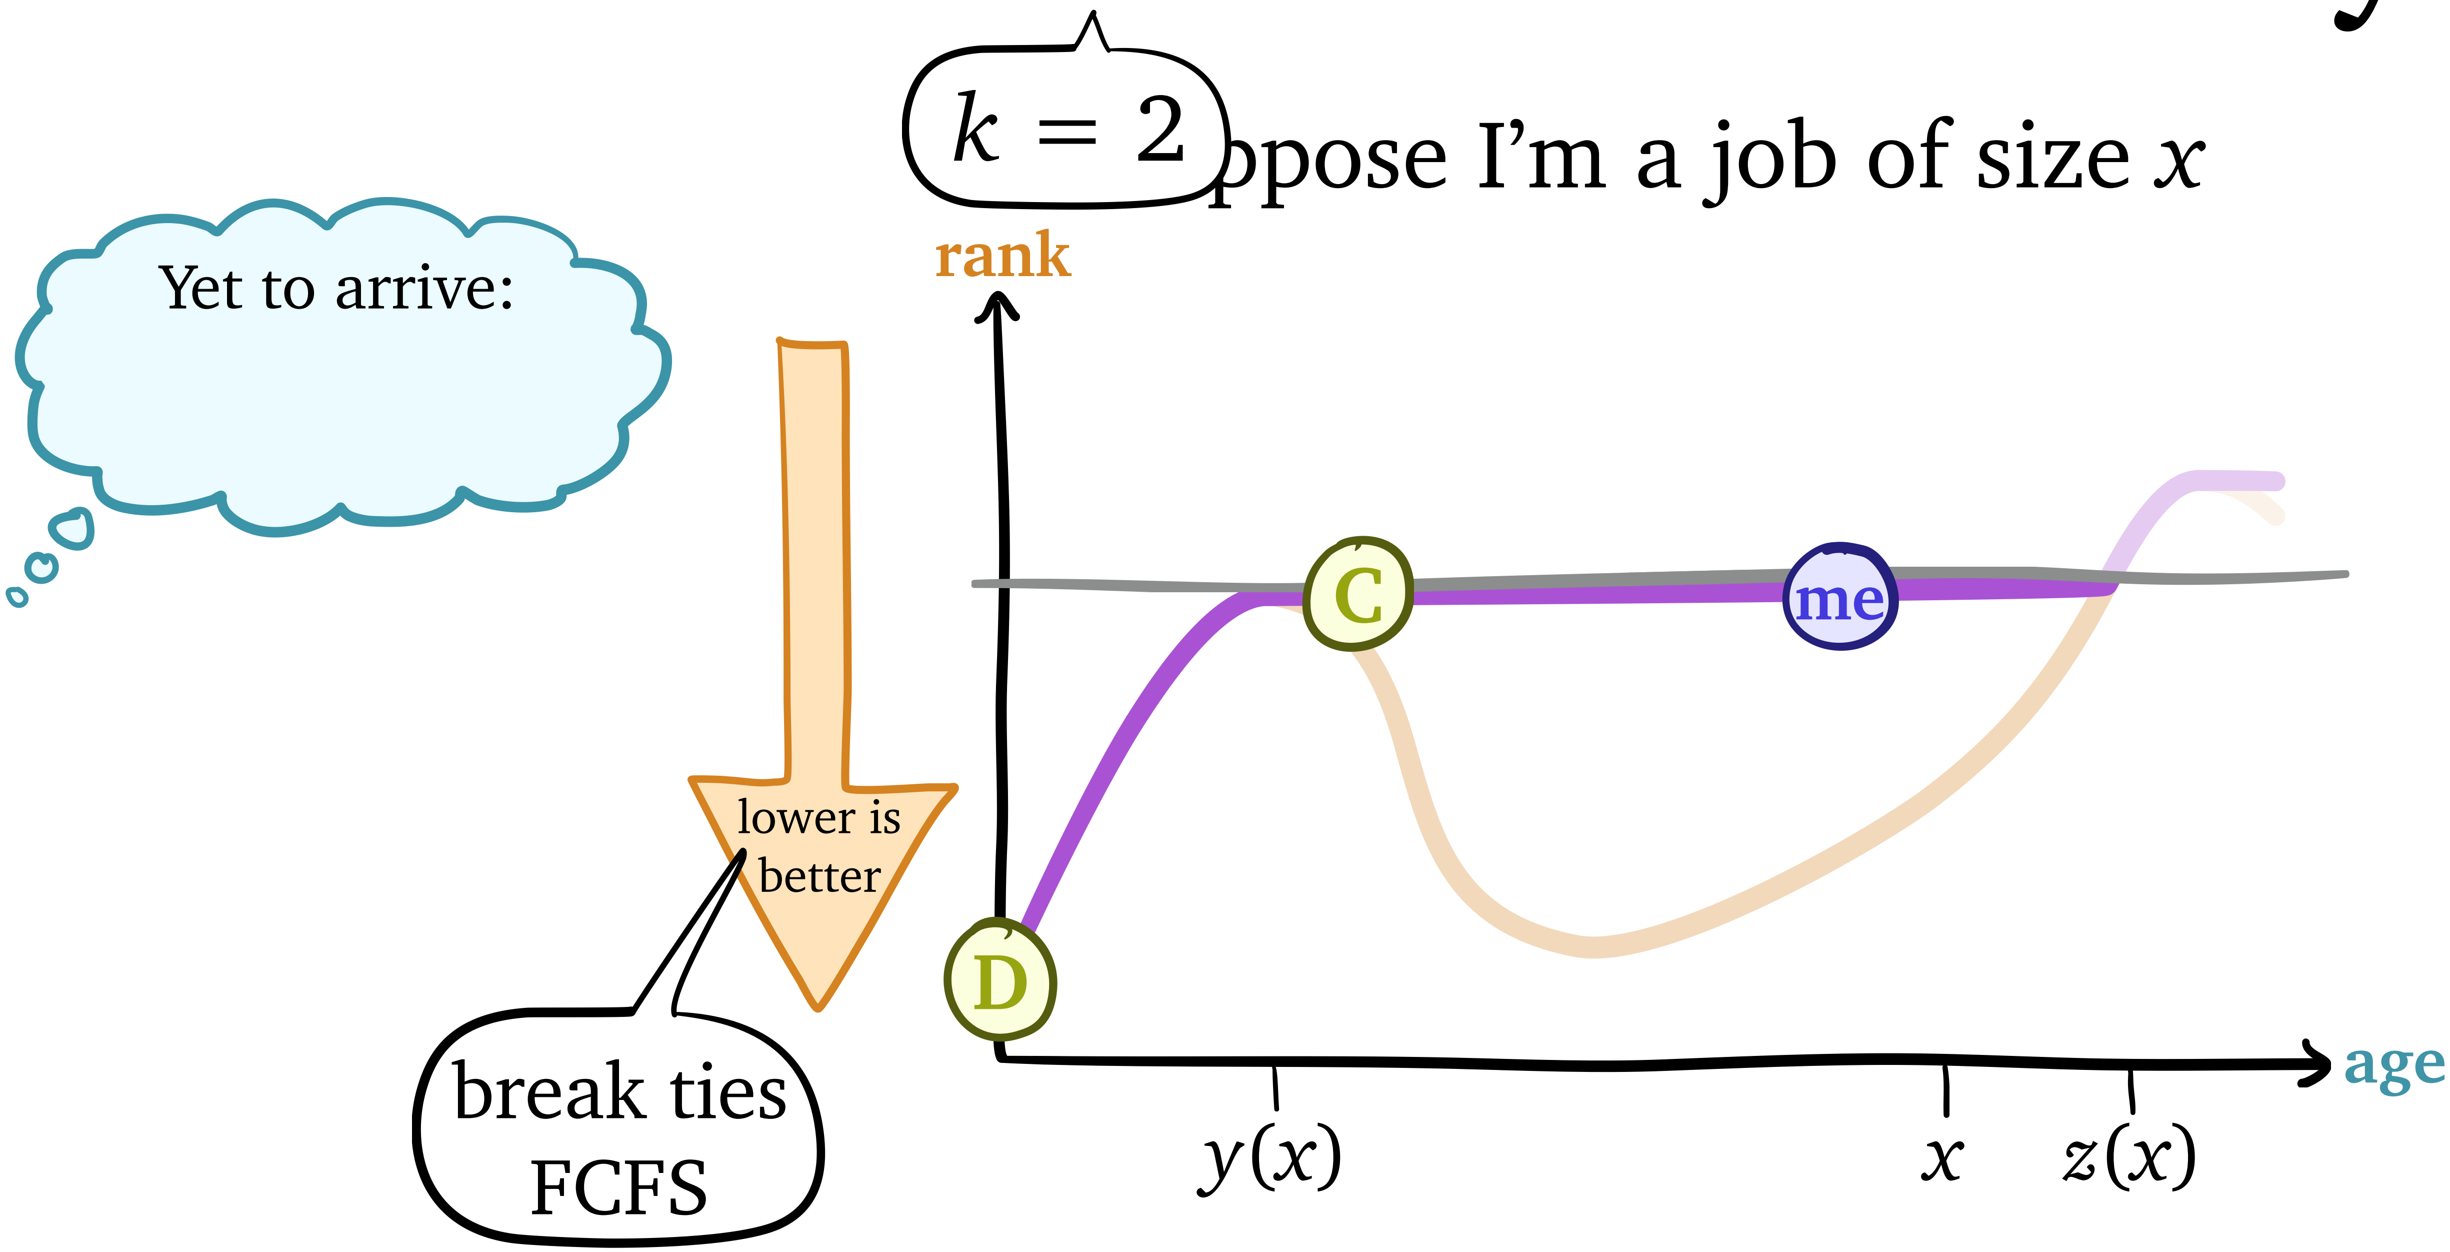

I can ignore { *old* jobs / *new* jobs (C & D) } after age { $z(x)$ / ??? }

M-Gittins-*k* Saves the Day

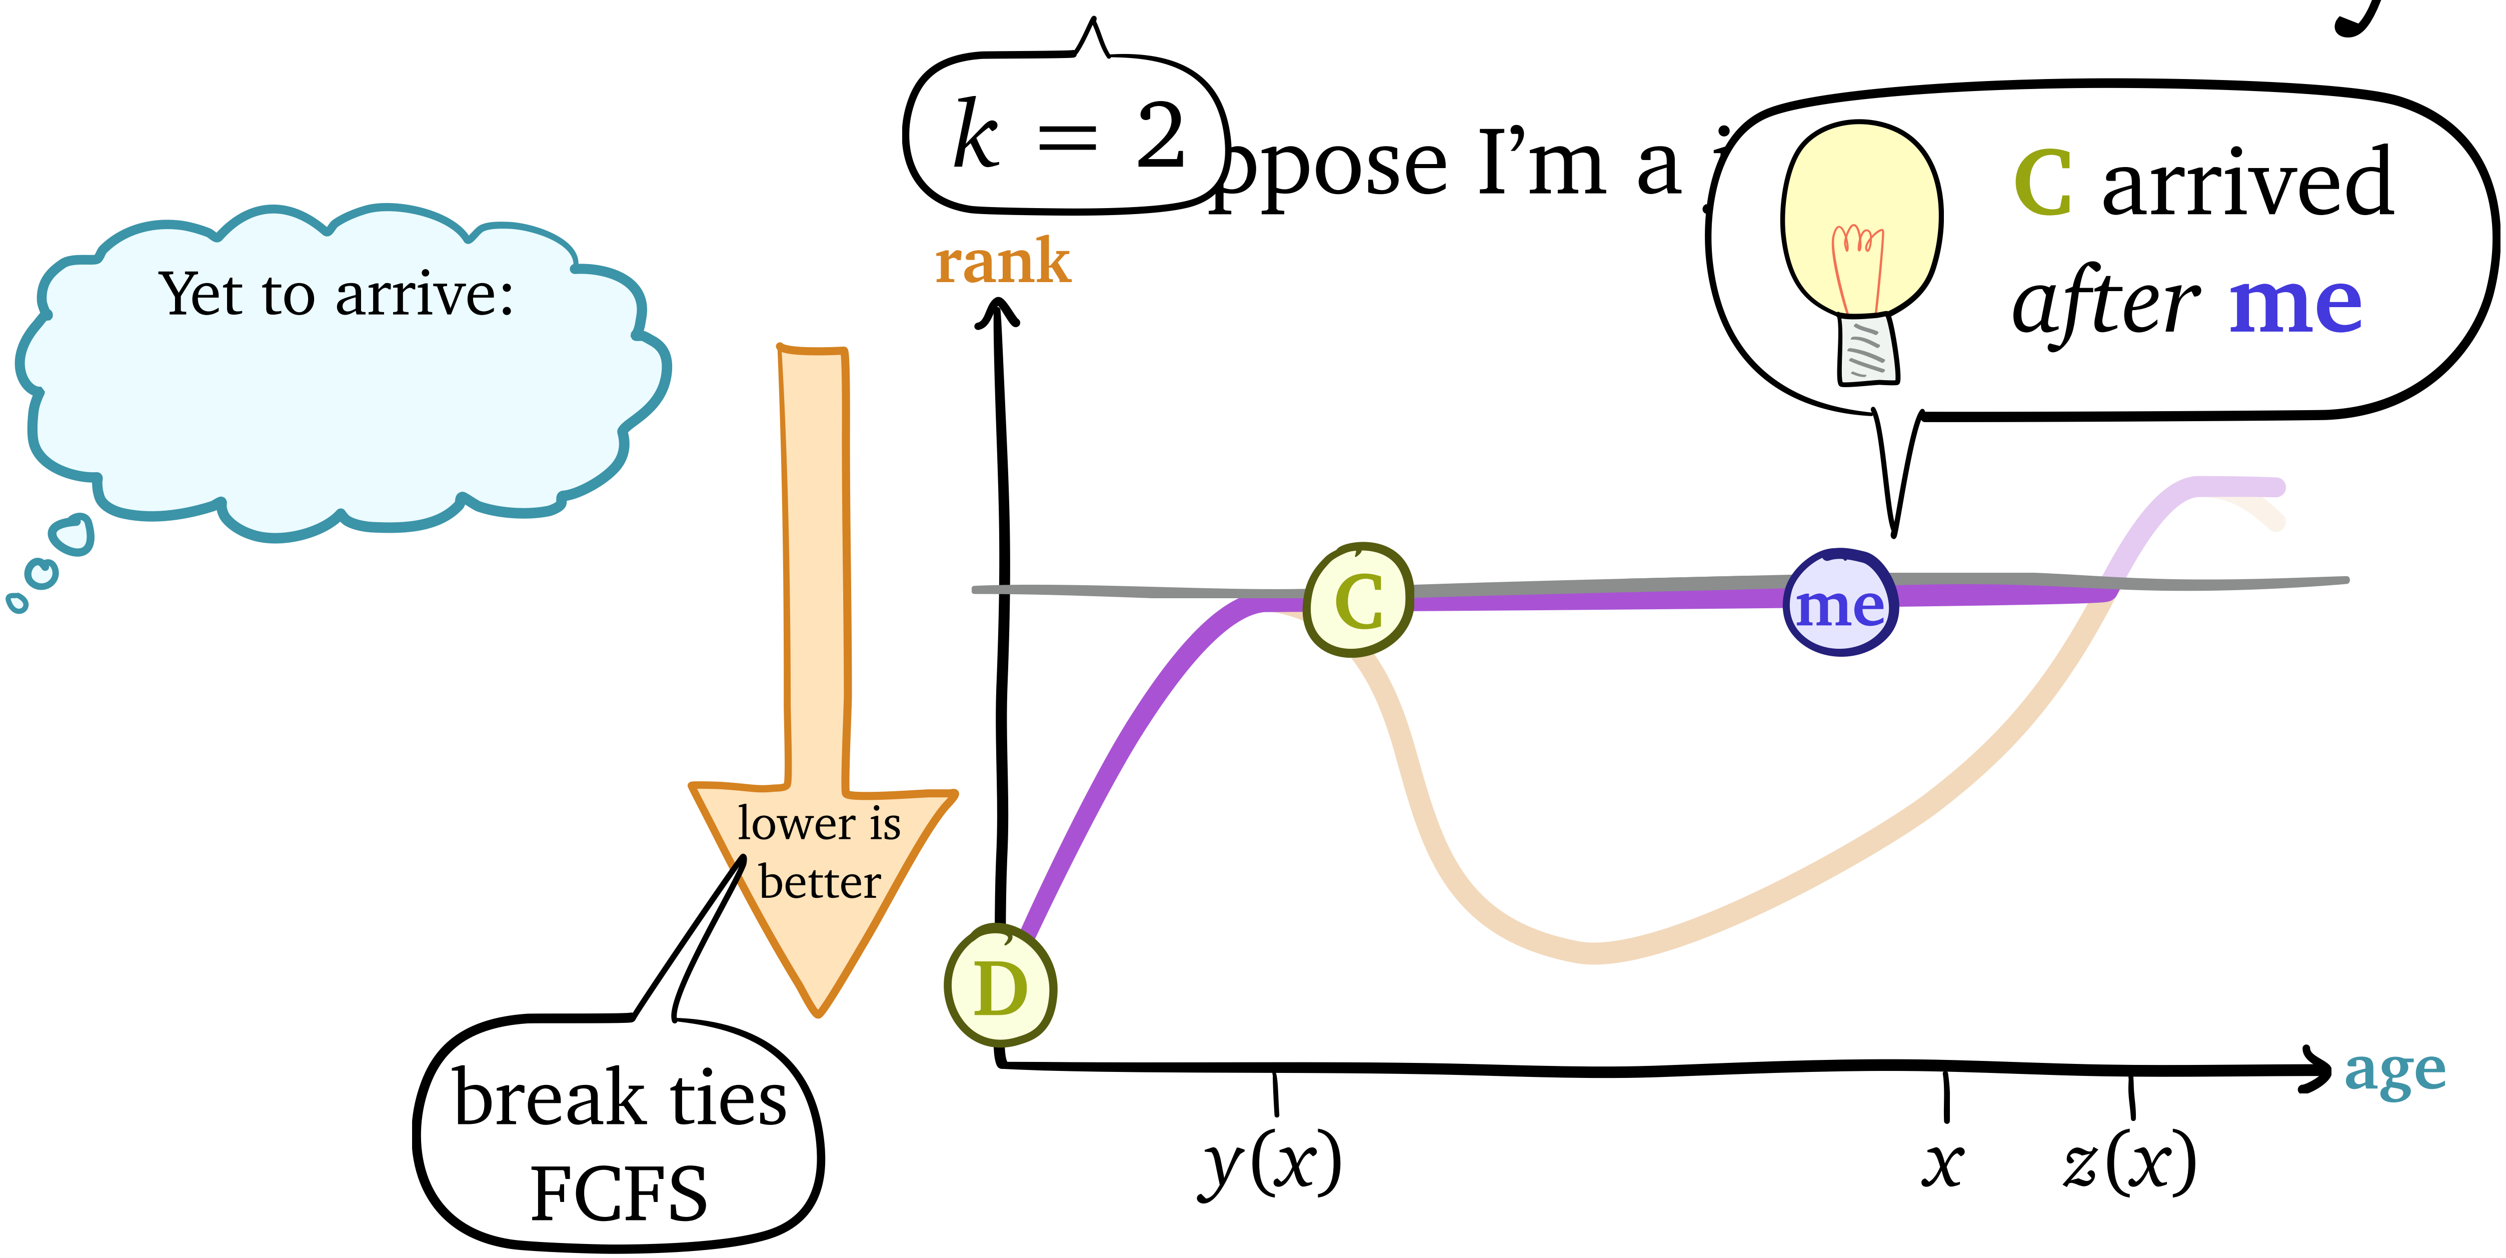

I can ignore { *old* jobs / *new* jobs (C & D) } after age { $z(x)$ / ??? }

M-Gittins-*k* Saves the Day

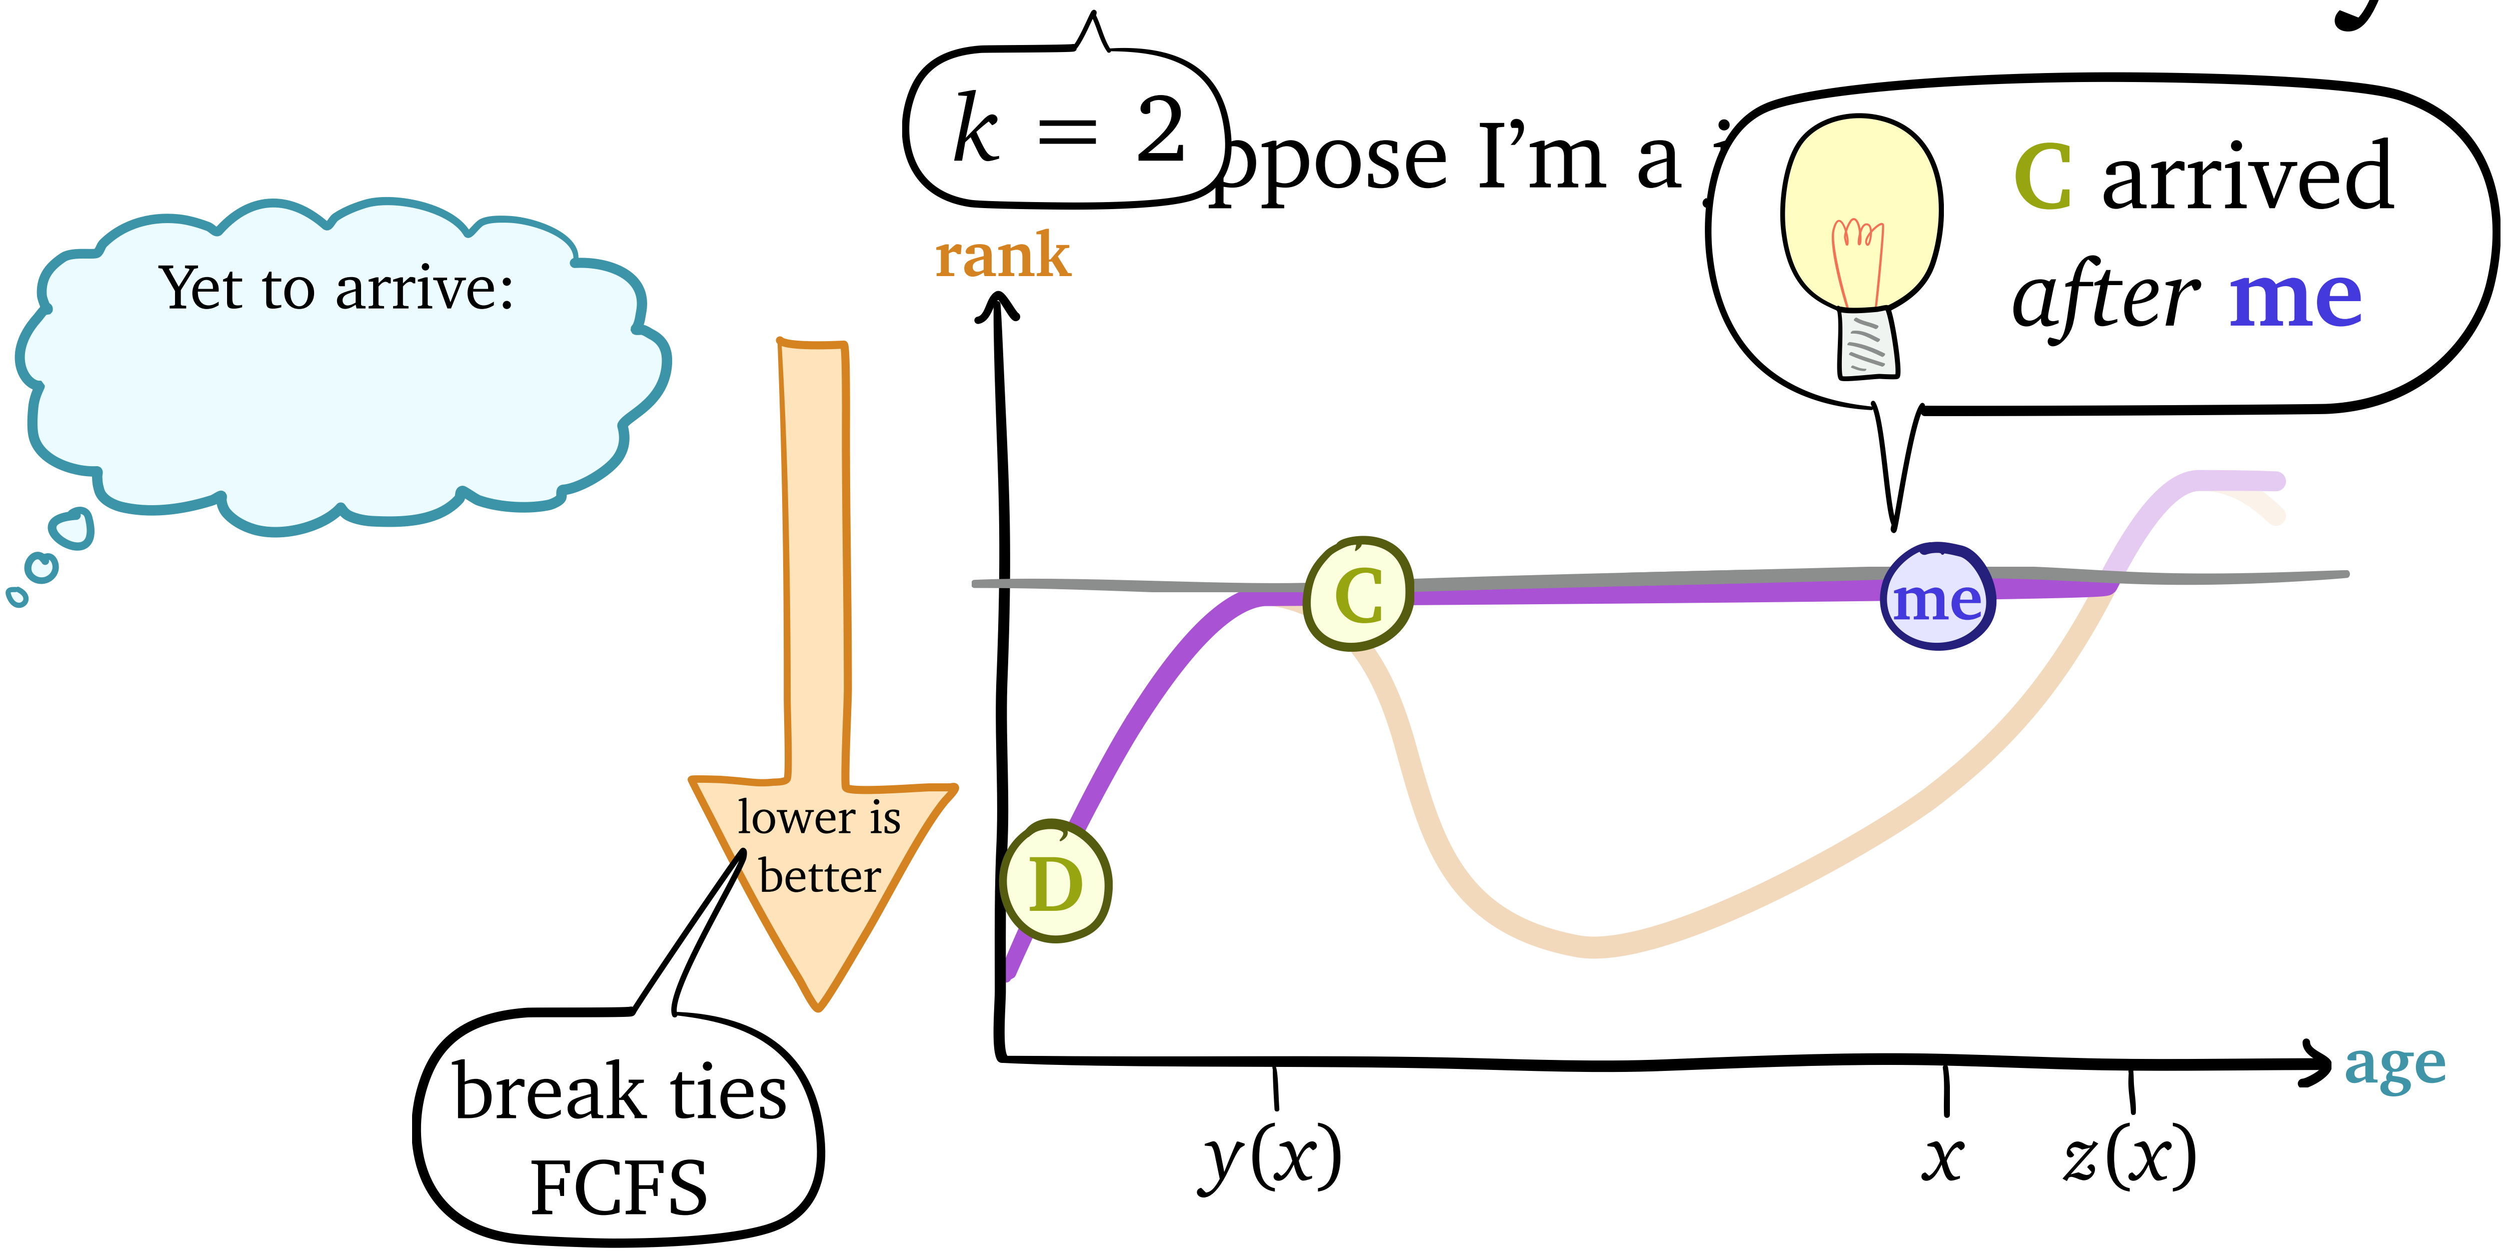

I can ignore { *old* jobs after age $z(x)$; *new* jobs (C & D) after age ??? }

M-Gittins-k Saves the Day

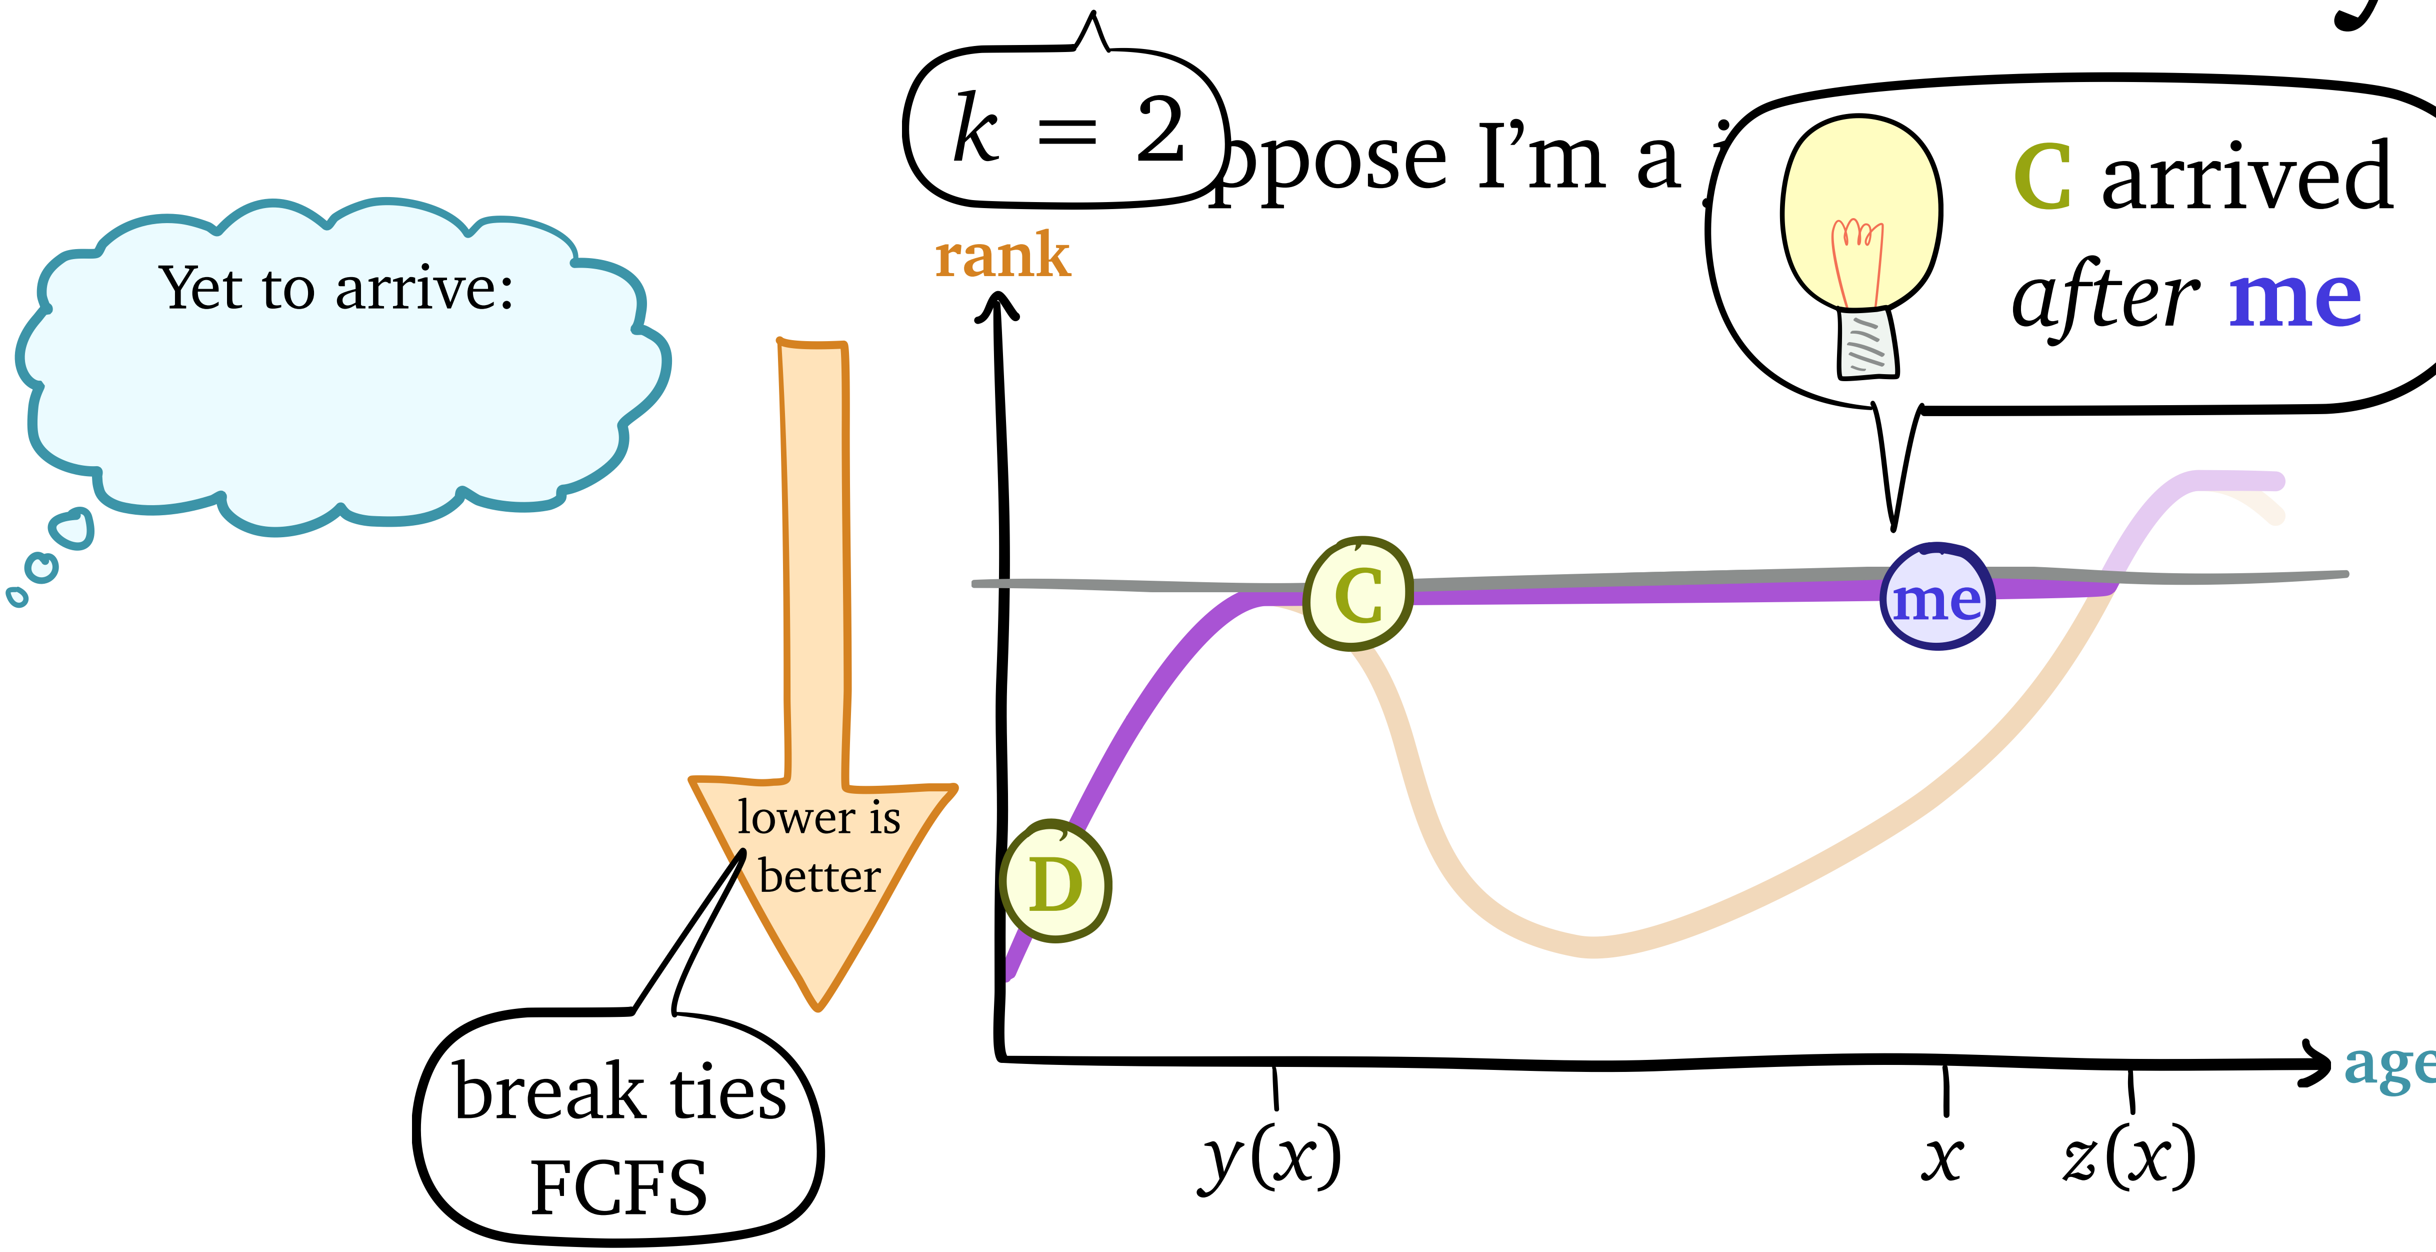

I can ignore $\left\{\begin{array}{c}\textit{old}\text{ jobs}\\ \textit{new}\text{ jobs (C \& D)}\end{array}\right\}$ after age $\left\{\begin{array}{c}z(x)\\ y(x)\end{array}\right\}$

M-Gittins-*k* Saves the Day

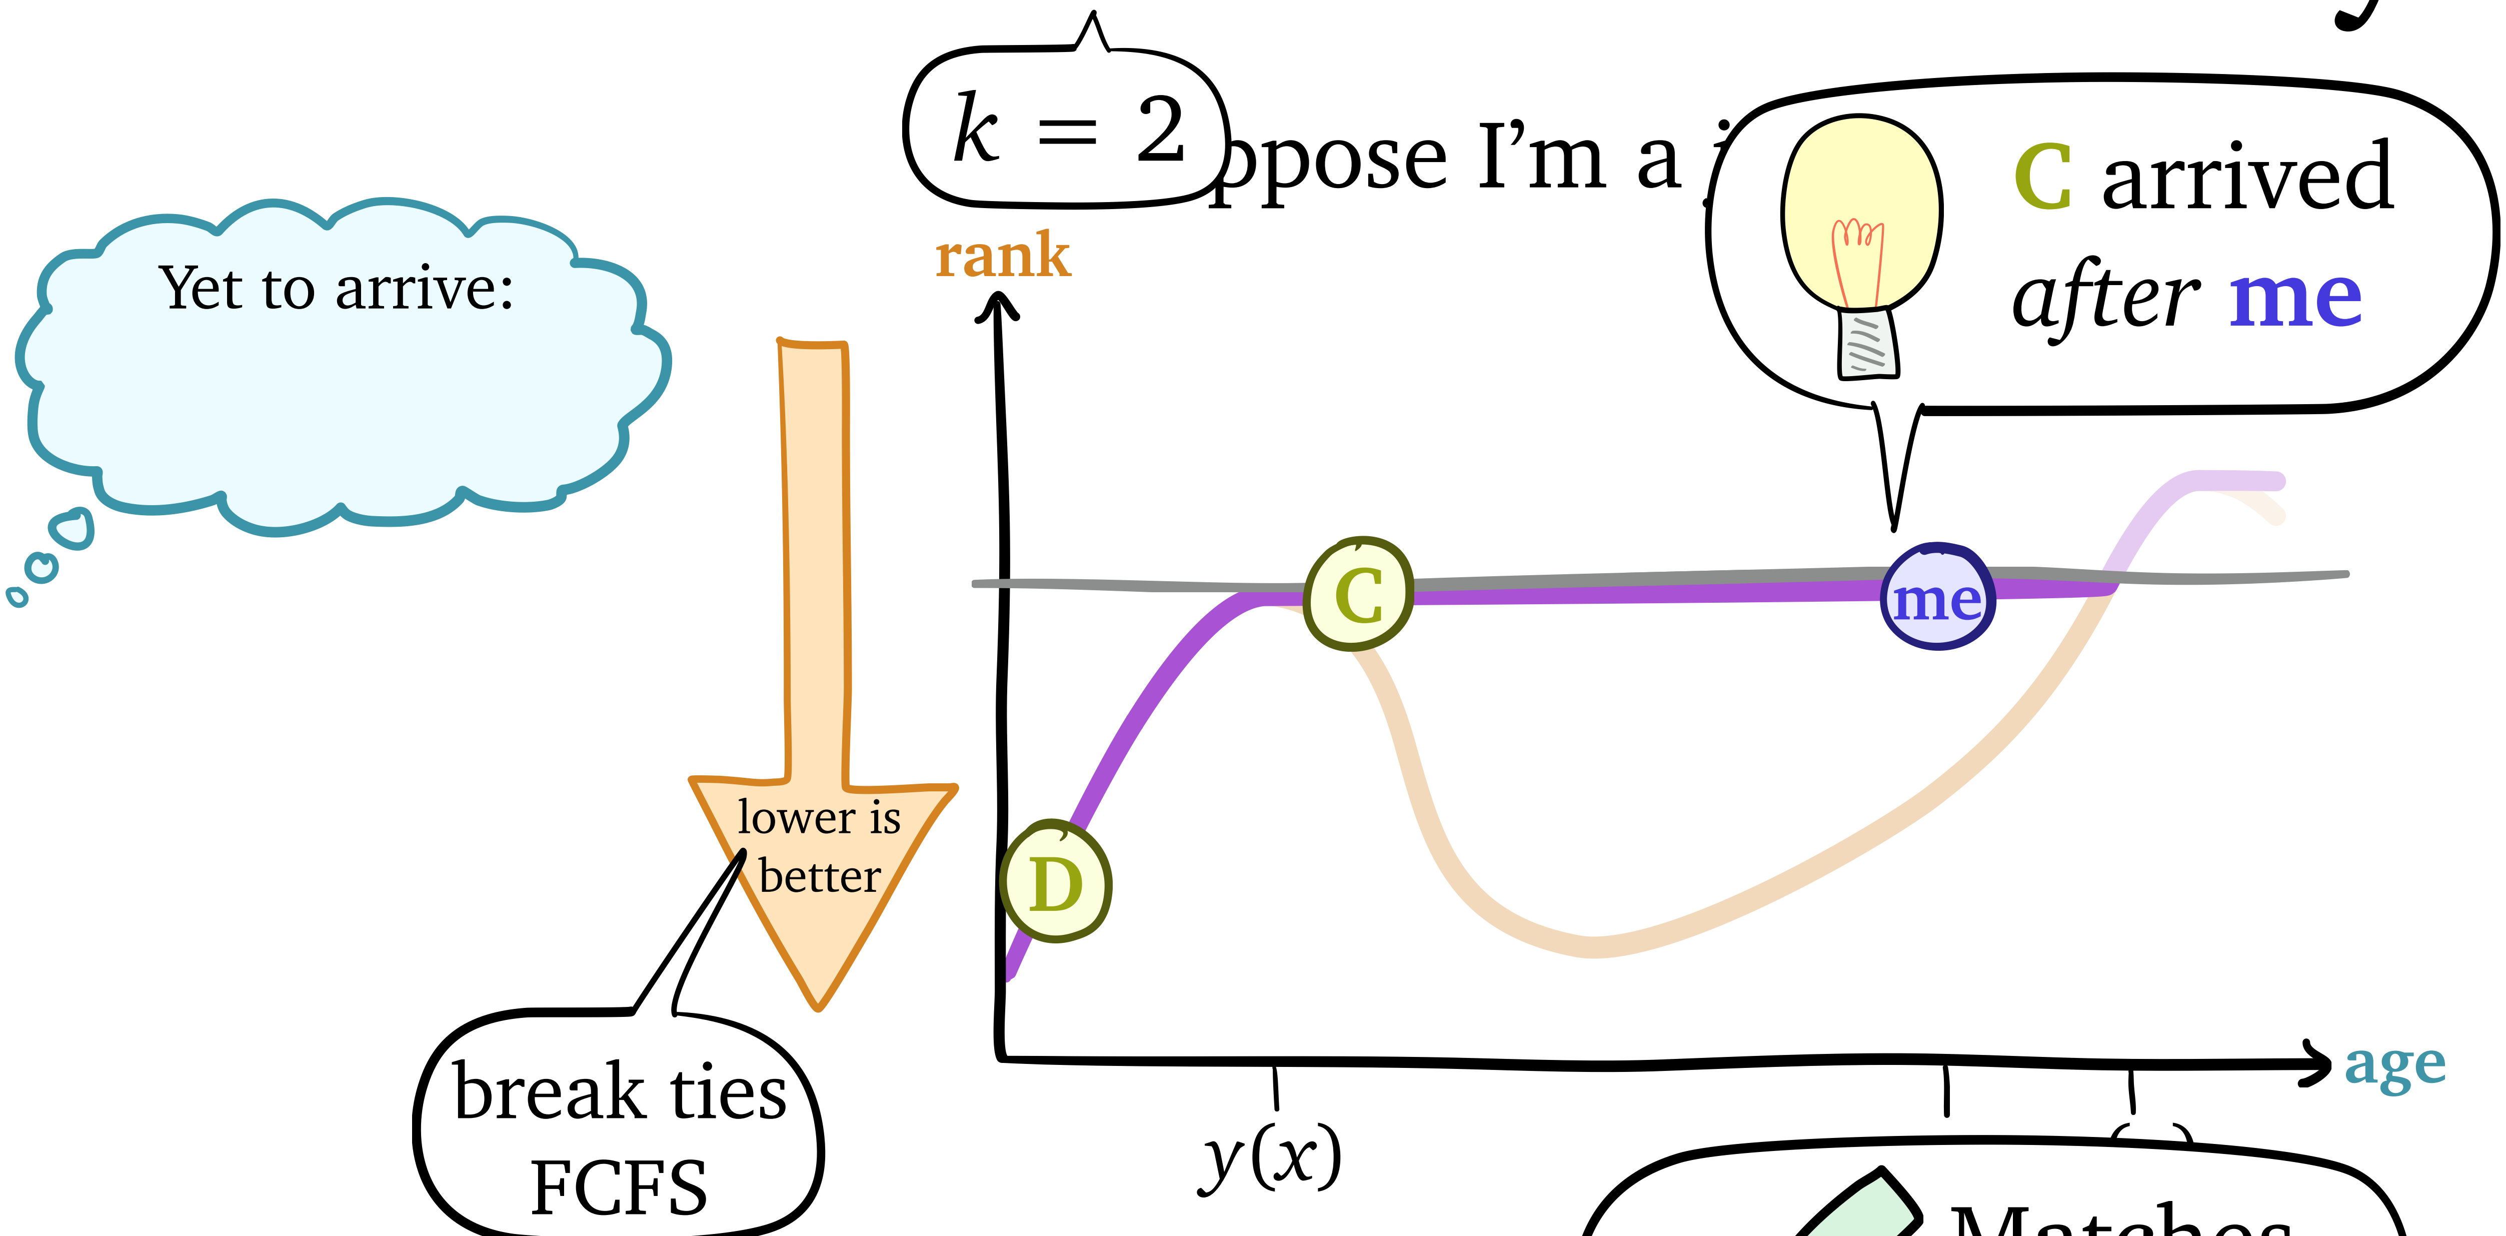

I can ignore { *old* jobs / *new* jobs (C & D) } after age { $y(x)$ }

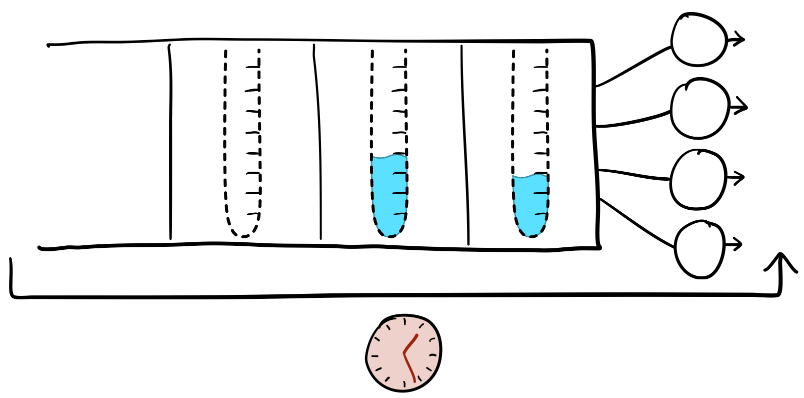

Goal: minimize heavy-traffic $\mathbf{E}[T]$ in M/G/k with unknown job sizes

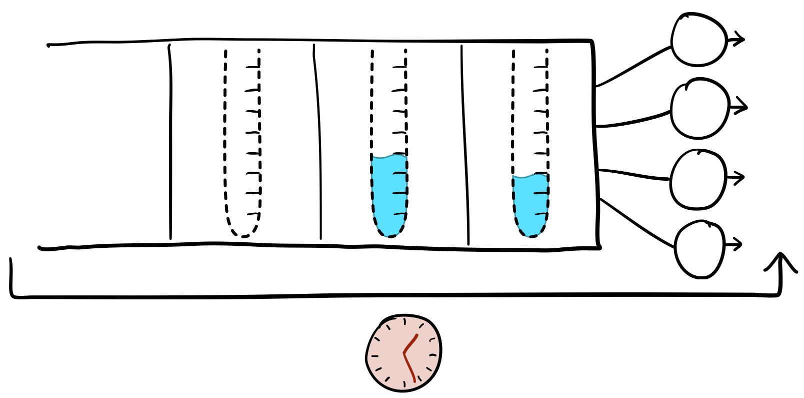

Goal: minimize heavy-traffic $\mathbf{E}[T]$ in M/G/k with unknown job sizes

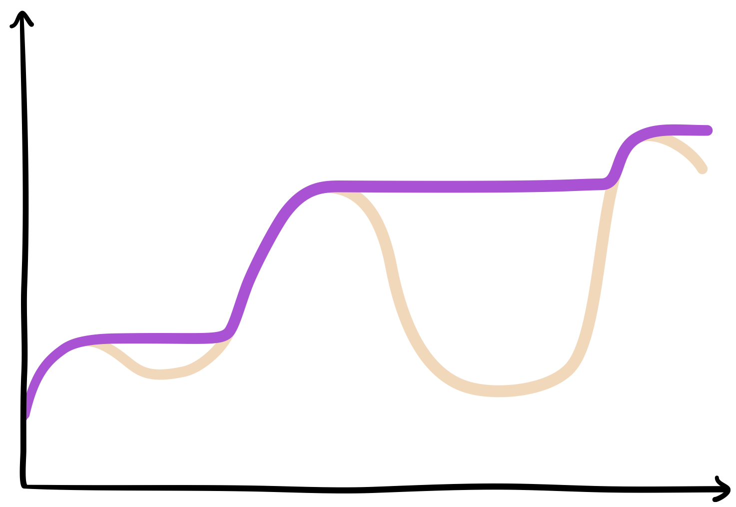

Key idea: new *monotonic* variant of **Gittins**, namely **M-Gittins**

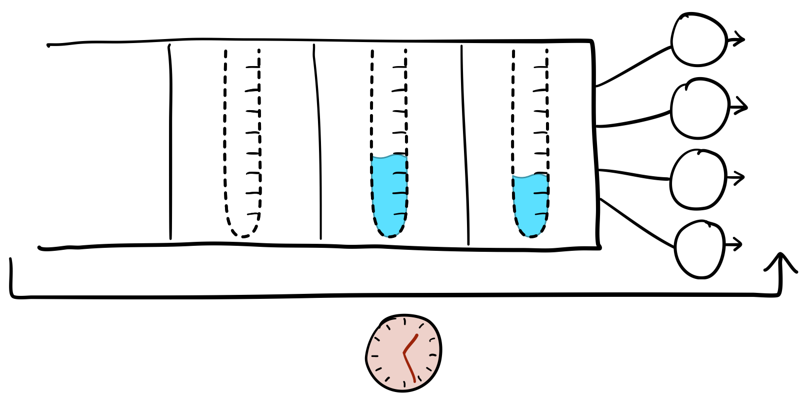

Goal: minimize heavy-traffic $\mathbf{E}[T]$ in M/G/k with unknown job sizes

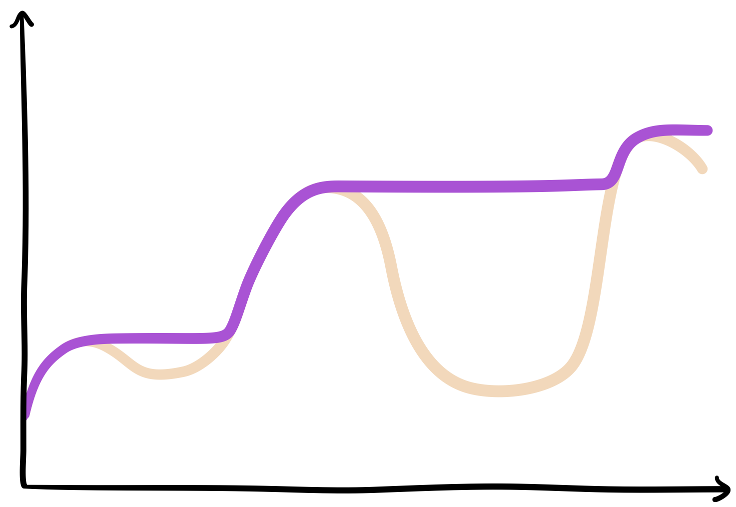

Key idea: new *monotonic* variant of **Gittins**, namely **M-Gittins**

Theorem: $\lim_{\rho \to 1} \frac{\mathbf{E}[T_{\text{M-Gittins-}k}]}{\mathbf{E}[T_{\text{Gittins-1}}]} = 1$

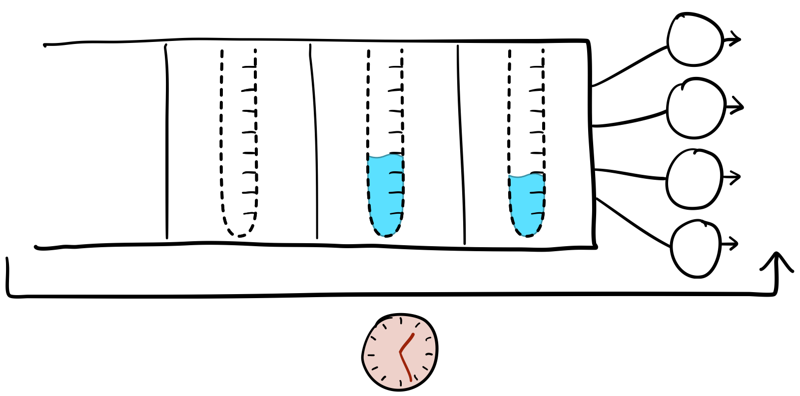

Goal: minimize heavy-traffic $\mathbf{E}[T]$ in M/G/k with unknown job sizes

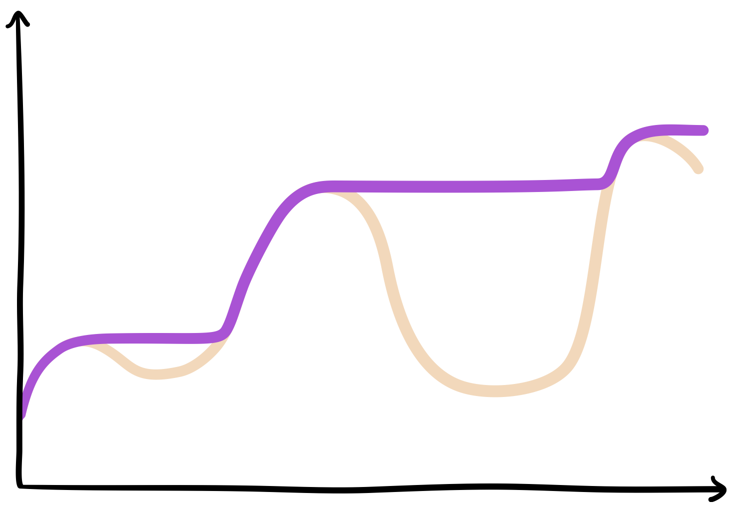

Key idea: new *monotonic* variant of **Gittins**, namely **M-Gittins**

Theorem: $\lim_{\rho \to 1} \frac{\mathbf{E}[T_{\textbf{M-Gittins-}k}]}{\mathbf{E}[T_{\textbf{Gittins-1}}]} = 1$

Get in touch: `zscully@cs.cmu.edu`

Bonus Slides

Main Results

Suppose X is heavy-tailed with finite variance

Main Results

Suppose X is heavy-tailed with finite variance

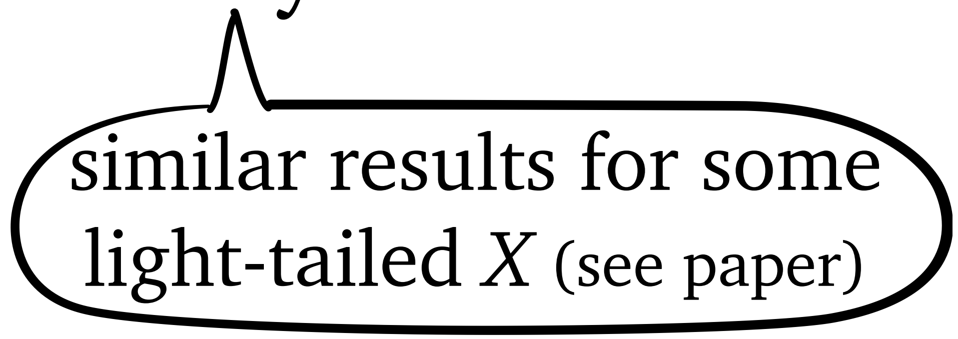

Main Results

Suppose X is heavy-tailed with finite variance

similar results for some light-tailed X (see paper)

Step 1: link **M-Gittins-k** to **Gittins-1**

Step 2: analyze heavy-traffic **Gittins-1**

Main Results

Suppose X is heavy-tailed with finite variance

similar results for some light-tailed X (see paper)

Step 1: link **M-Gittins-*k*** to **Gittins-1**

$$\mathbf{E}[T_{\text{M-Gittins-}k}] \le \mathbf{E}[T_{\text{Gittins-1}}] + k \cdot O\left(\log \frac{1}{1-\rho}\right)$$

Step 2: analyze heavy-traffic **Gittins-1**

Main Results

Suppose X is heavy-tailed with finite variance

similar results for some light-tailed X (see paper)

Step 1: link **M-Gittins-*k*** to **Gittins-1**

$$\mathbf{E}[T_{\text{M-Gittins-}k}] \leq \mathbf{E}[T_{\text{Gittins-1}}] + k \cdot O\left(\log \frac{1}{1-\rho}\right)$$

Step 2: analyze heavy-traffic **Gittins-1**

$$\mathbf{E}[T_{\text{Gittins-1}}] = \omega\left(\log \frac{1}{1-\rho}\right)$$

Main Results

Suppose X is heavy-tailed with finite variance

similar results for some light-tailed X (see paper)

Step 1: link **M-Gittins-k** to **Gittins-1**

$$\mathbf{E}[T_{\text{M-Gittins-}k}] \leq \mathbf{E}[T_{\text{Gittins-1}}] + k \cdot O\left(\log \frac{1}{1-\rho}\right)$$

Step 2: analyze heavy-traffic **Gittins-1**

$$\mathbf{E}[T_{\text{Gittins-1}}] = \omega\left(\log \frac{1}{1-\rho}\right)$$

$$\Theta\left(\frac{1}{1-\rho} \Bigg/ \max_{0 \leq b \leq \bar{F}_e^{-1}(1-\rho)} \mathbf{E}[X - b \mid X > b]\right)$$

Heavy-Traffic Optimality

Heavy-Traffic Optimality

Theorem:
M-Gittins-*k* is *heavy-traffic optimal* in the M/G/*k*, specifically

$$\lim_{\rho \to 1} \frac{\mathbf{E}[T_{\textbf{M-Gittins-}k}]}{\mathbf{E}[T_{\textbf{Gittins-1}}]} = 1,$$

Heavy-Traffic Optimality

Theorem:
M-Gittins-k is *heavy-traffic optimal* in the M/G/k, specifically

$$\lim_{\rho \to 1} \frac{\mathbf{E}[T_{\textbf{M-Gittins-}k}]}{\mathbf{E}[T_{\textbf{Gittins-1}}]} = 1,$$

if X is in *any* of the following classes:

Heavy-Traffic Optimality

Theorem:
M-Gittins-k is *heavy-traffic optimal* in the M/G/k, specifically

$$\lim_{\rho \to 1} \frac{\mathbf{E}[T_{\textbf{M-Gittins-}k}]}{\mathbf{E}[T_{\textbf{Gittins-1}}]} = 1,$$

if X is in *any* of the following classes:

- bounded

Heavy-Traffic Optimality

Theorem:
M-Gittins-*k* is *heavy-traffic optimal* in the M/G/*k*, specifically

$$\lim_{\rho \to 1} \frac{\mathbf{E}[T_{\text{M-Gittins-}k}]}{\mathbf{E}[T_{\text{Gittins-1}}]} = 1,$$

if X is in *any* of the following classes:

- bounded
- "finite-variance heavy-tailed"
  (*O*-regularly varying with Matuszewska indices less than –2)

Heavy-Traffic Optimality

Theorem:
M-Gittins-k is *heavy-traffic optimal* in the M/G/k, specifically

$$\lim_{\rho \to 1} \frac{\mathbf{E}[T_{\text{M-Gittins-}k}]}{\mathbf{E}[T_{\text{Gittins-1}}]} = 1,$$

if X is in *any* of the following classes:

- bounded
- "finite-variance heavy-tailed"
  (*O*-regularly varying with Matuszewska indices less than –2)
- MDA(Λ) with "quasi-decreasing hazard rate", e.g. $h(x) = \Theta(x^{-\gamma})$

Heavy-Traffic Optimality

Theorem:
M-Gittins-k is *heavy-traffic optimal* in the M/G/k, specifically

$$\lim_{\rho \to 1} \frac{\mathbf{E}[T_{\text{M-Gittins-}k}]}{\mathbf{E}[T_{\text{Gittins-1}}]} = 1,$$

if X is in *any* of the following classes:

- bounded
- "finite-variance heavy-tailed"
  (*O*-regularly varying with Matuszewska indices less than –2)
- MDA(Λ) with "quasi-decreasing hazard rate", e.g. $h(x) = \Theta(x^{-\gamma})$

exponential,
log-normal,
Weibull…

M/G/1 Heavy-Traffic Scaling

M/G/1 Heavy-Traffic Scaling

Theorem:

"infinite variance"

"finite variance"

M/G/1 Heavy-Traffic Scaling

Theorem:

"infinite variance" — If $X \in \mathrm{OR}(-2,-1)$, then

$$\mathbf{E}[T_{\textbf{Gittins-1}}] = \Theta\left(\log \frac{1}{1-\rho}\right),$$

"finite variance"

M/G/1 Heavy-Traffic Scaling

Theorem:

"infinite variance" — If $X \in \mathrm{OR}(-2,-1)$, then

$$\mathbf{E}[T_{\text{Gittins-1}}] = \Theta\left(\log \frac{1}{1-\rho}\right),$$

"finite variance" — and if $X \in \mathrm{OR}(-\infty,-2) \cup \mathrm{MDA}(\Lambda) \cup \mathrm{ENBUE}$, then

$$\mathbf{E}[T_{\text{Gittins-1}}] = \Theta\left(\frac{1}{1-\rho} \Big/ \max_{0 \le b \le \overline{F}_e^{-1}(1-\rho)} \mathbf{E}[X - b \mid X > b]\right).$$