The Gittins Index:

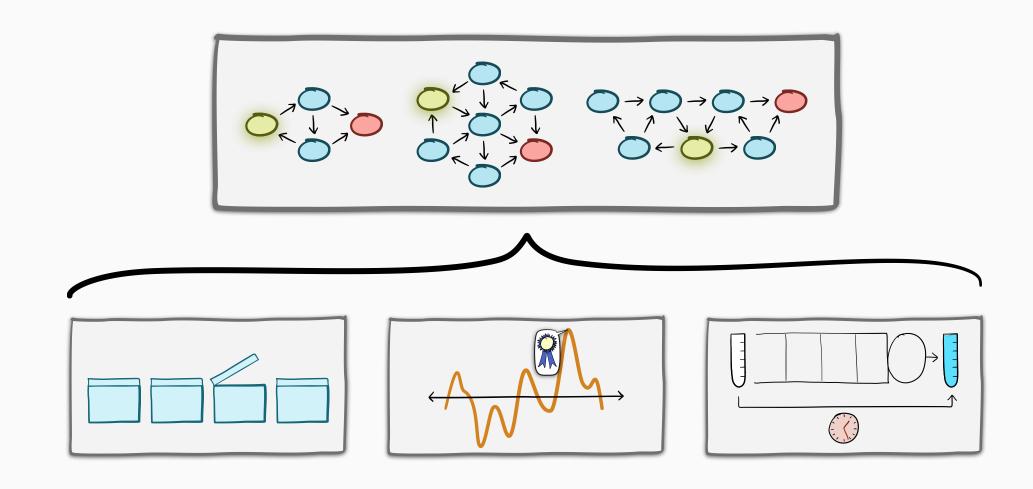
A Design Principle for Decision Making Under Uncertainty

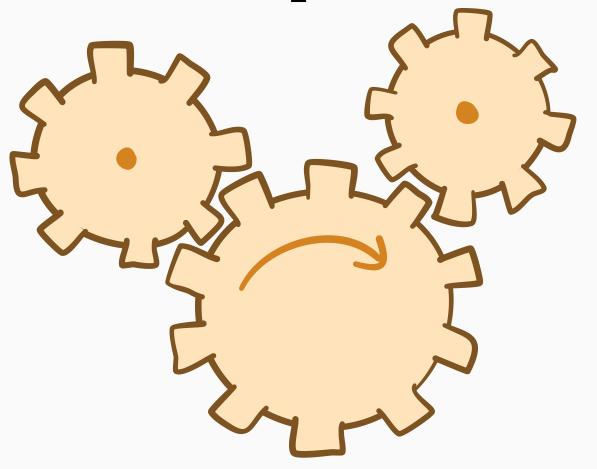
Ziv Scully

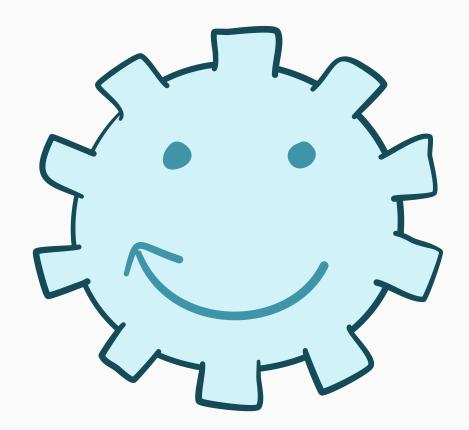
Cornell ORIE
https://ziv.codes

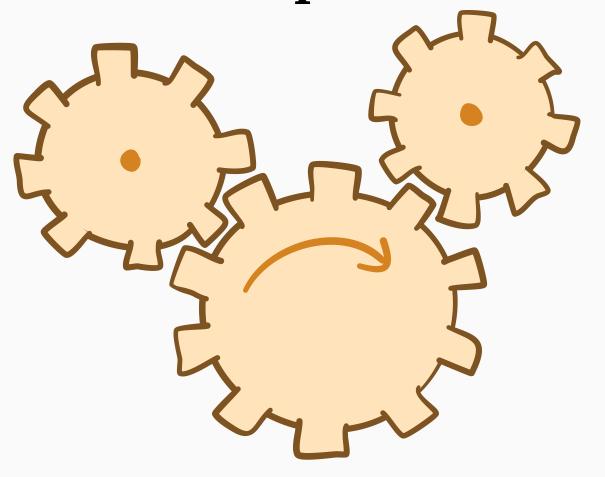
Alex Terenin

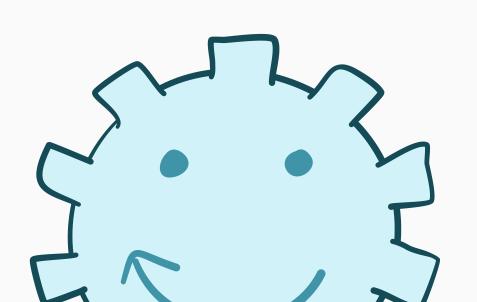
Cornell CDSES
https://avt.im

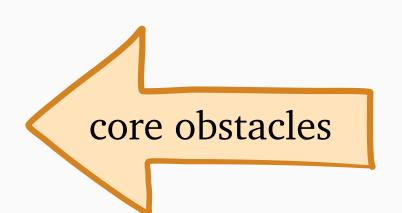


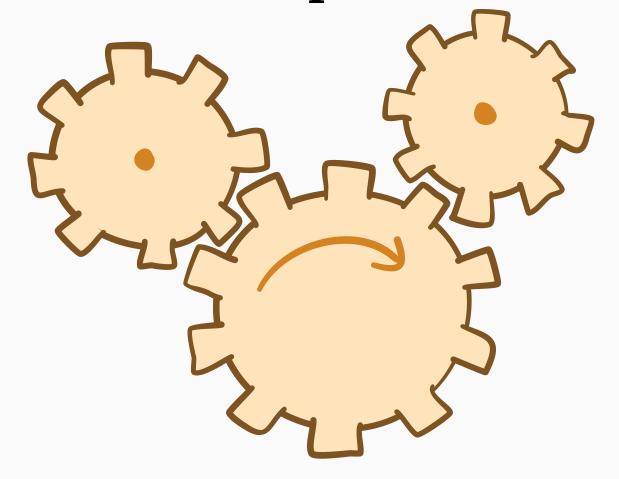


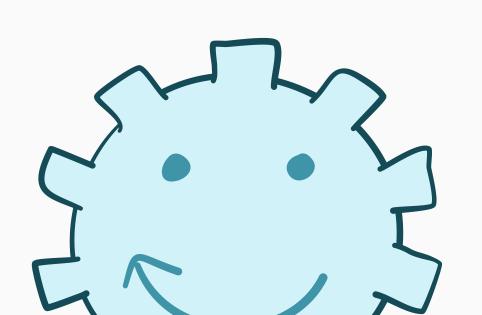


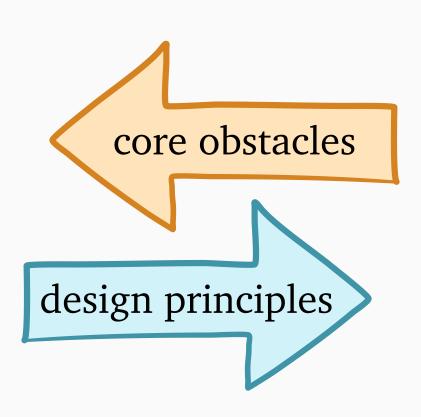


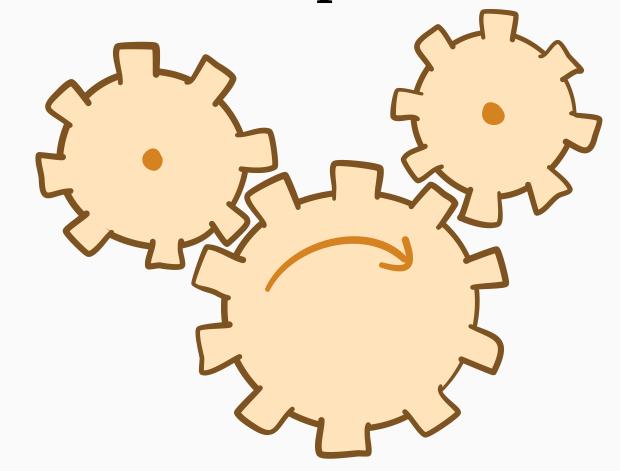


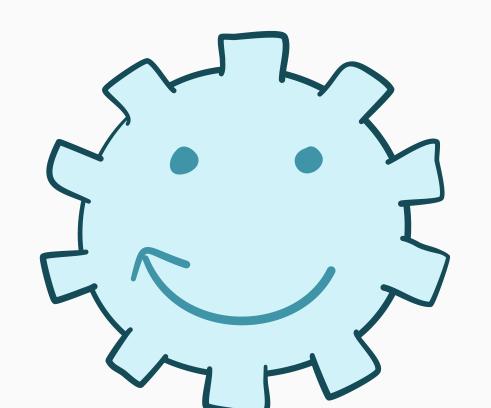


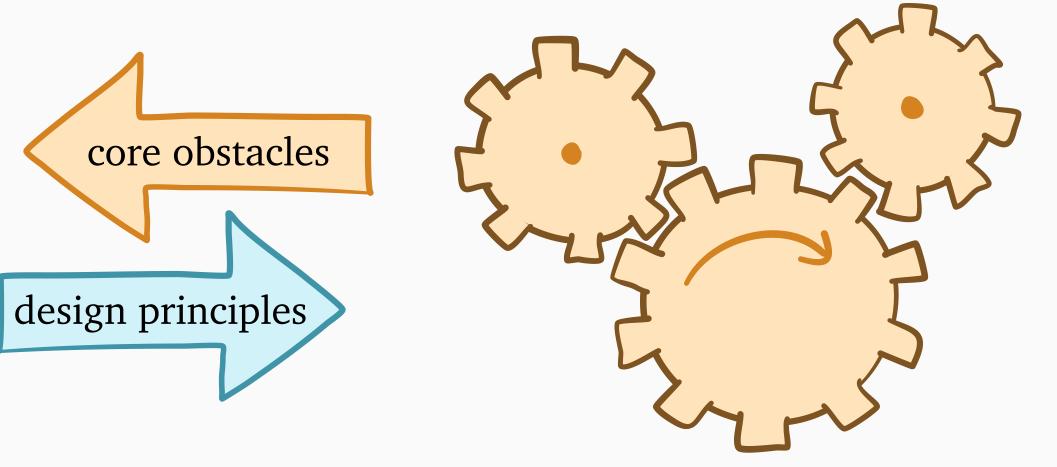




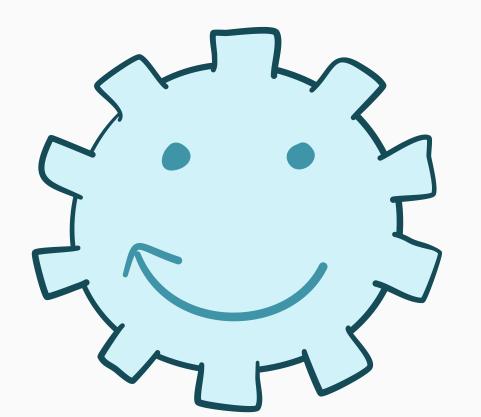




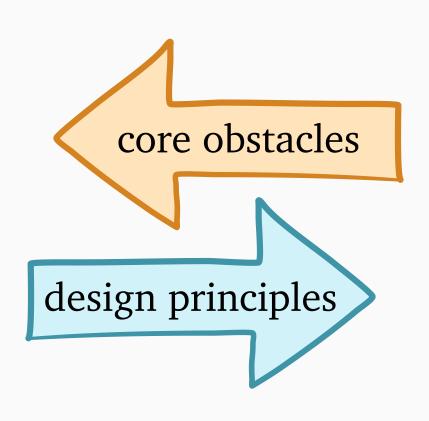


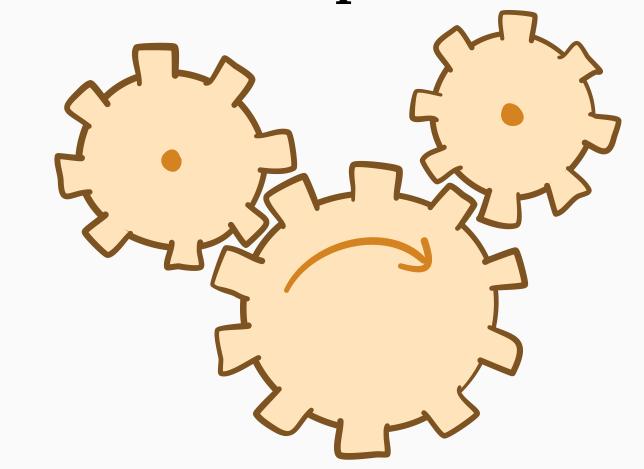


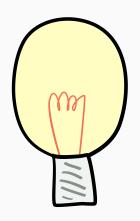




Practical problems

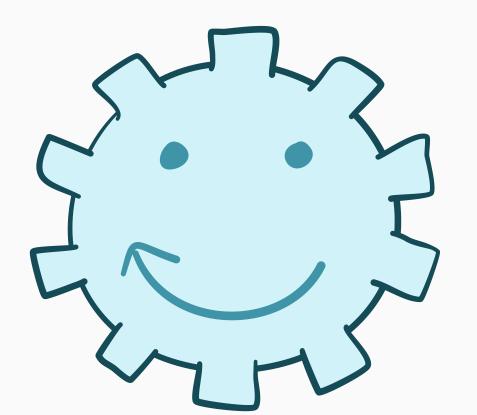




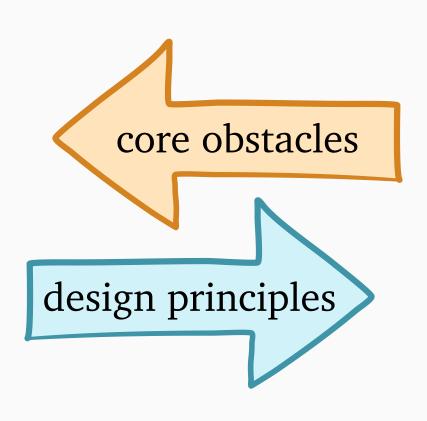


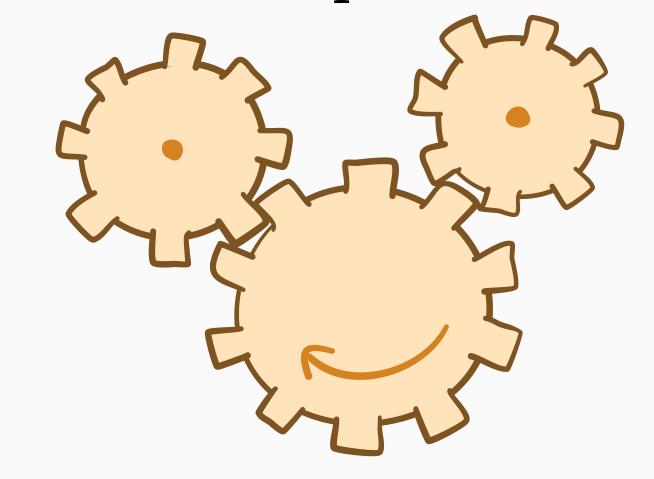
optimistic estimates
regularization
adaptive > nonadaptive
quantify information

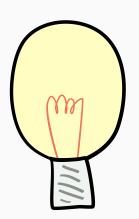




Practical problems

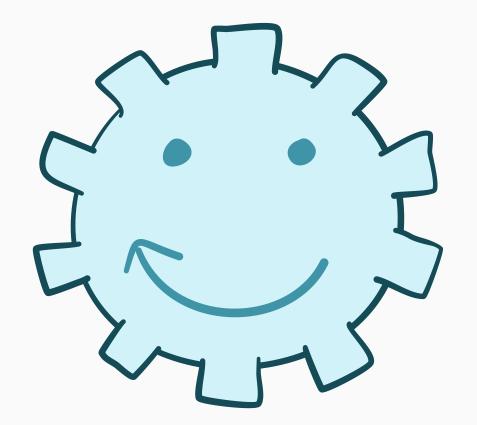




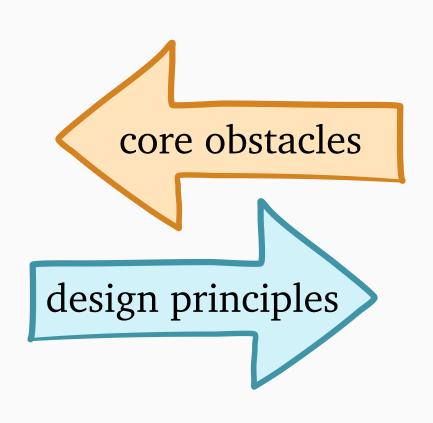


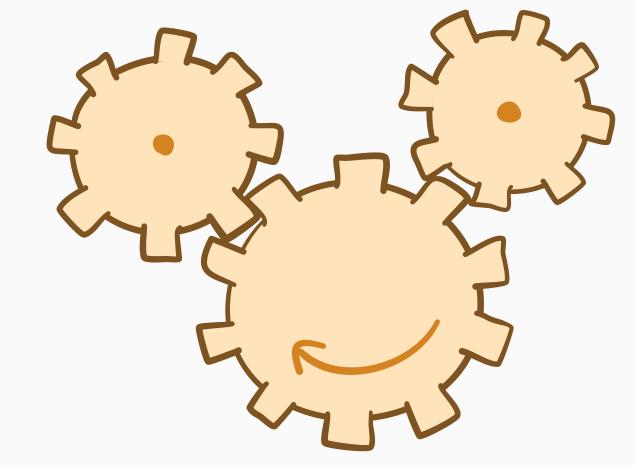
optimistic estimates
regularization
adaptive > nonadaptive
quantify information

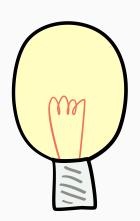




Practical problems

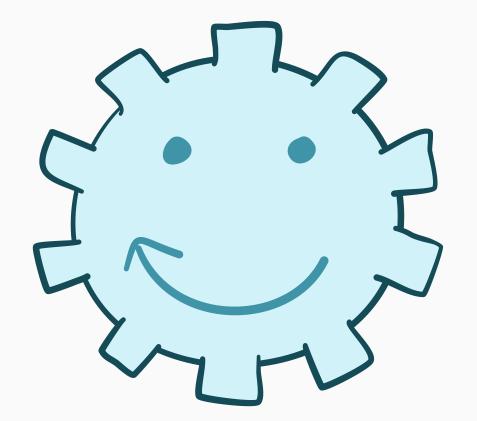




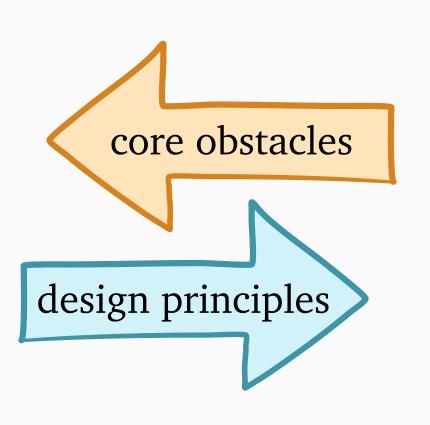


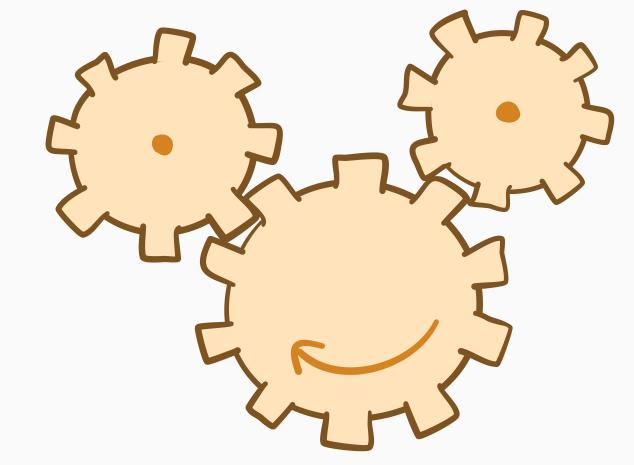
optimistic estimates
regularization
adaptive > nonadaptive
quantify information

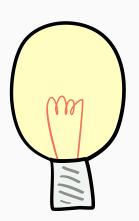
uncertainty with...



Practical problems



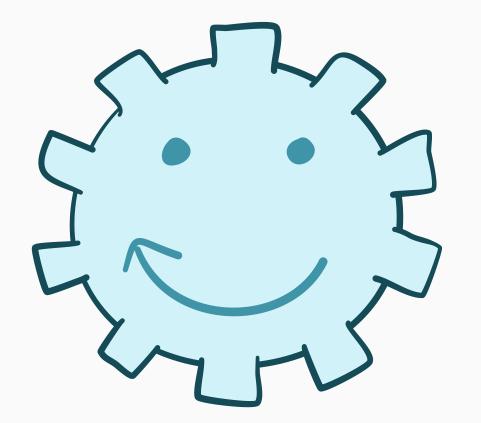




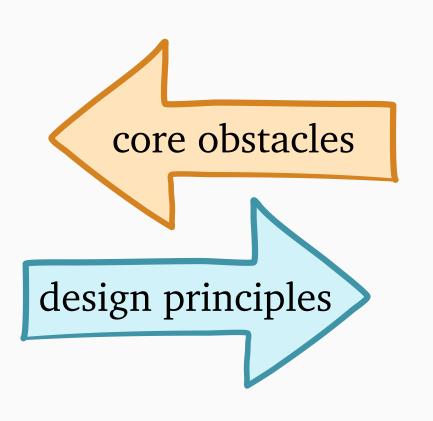
optimistic estimates
regularization
adaptive > nonadaptive
quantify information

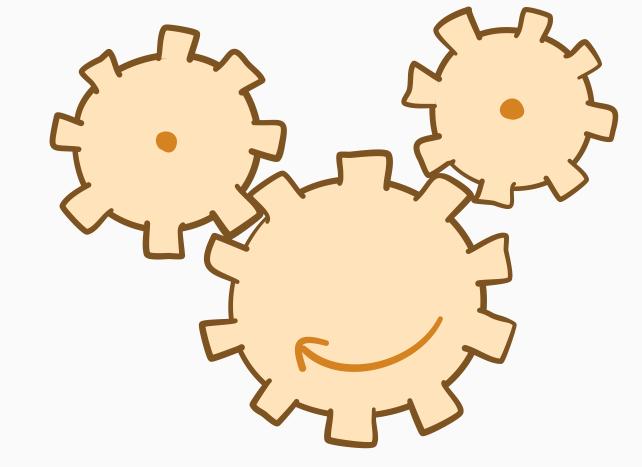
u

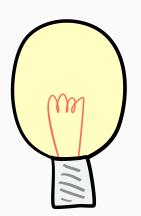
uncertainty with...
costly information



Practical problems



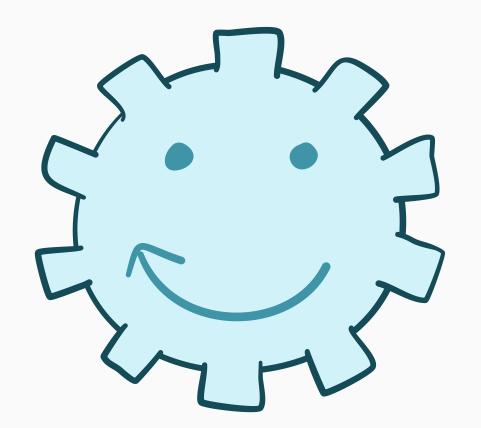




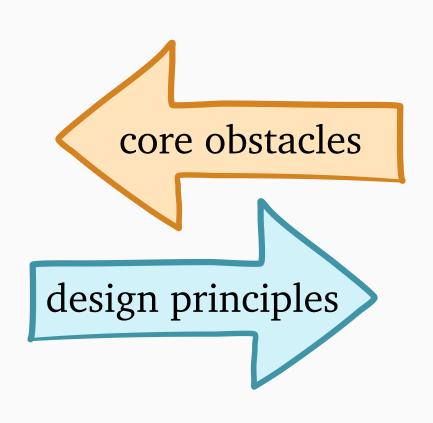
optimistic estimates
regularization
adaptive > nonadaptive
quantify information

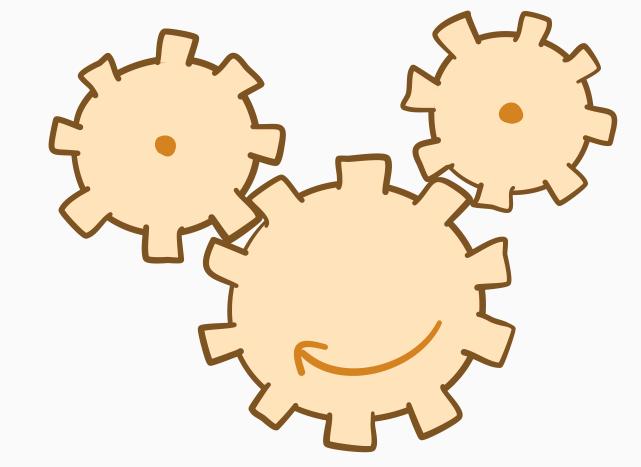


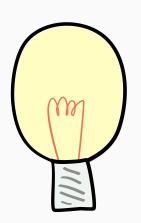
uncertainty with...
costly information
multi-stage uncertainty



Practical problems

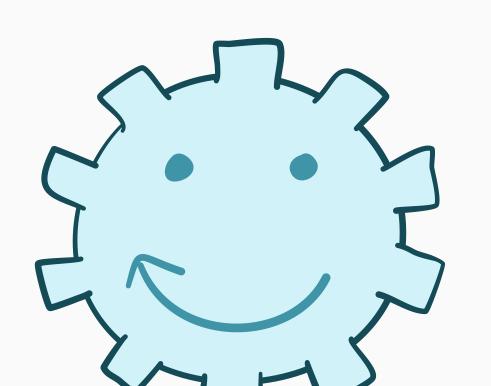




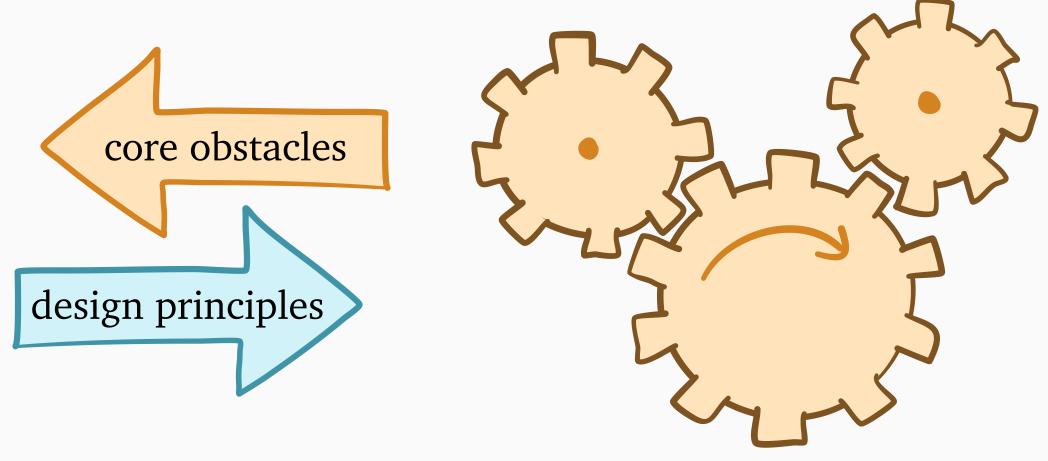


optimistic estimates
regularization
adaptive > nonadaptive
quantify information

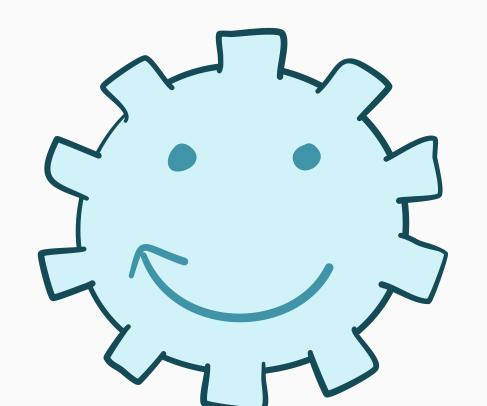
costly information multi-stage uncertainty heterogeneity

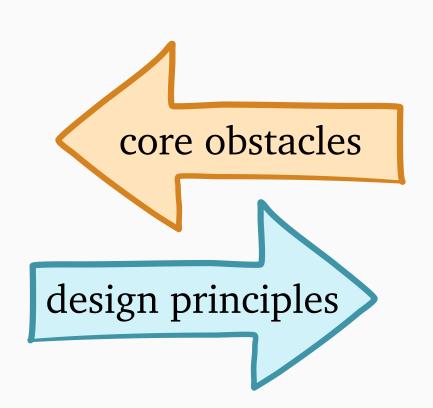


Practical problems

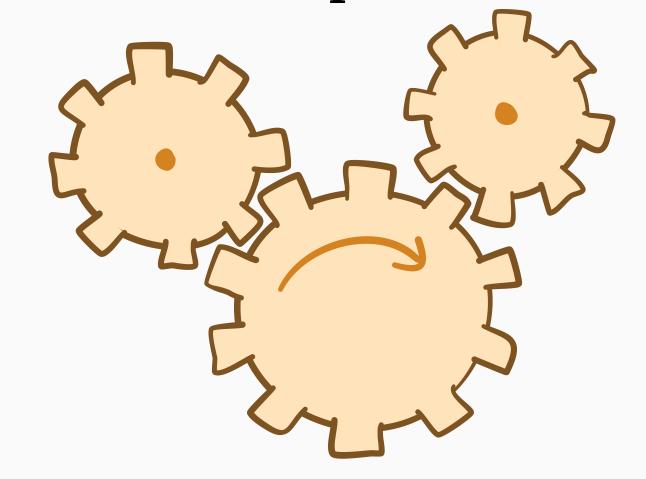


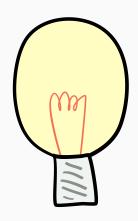
costly information multi-stage uncertainty heterogeneity





Practical problems



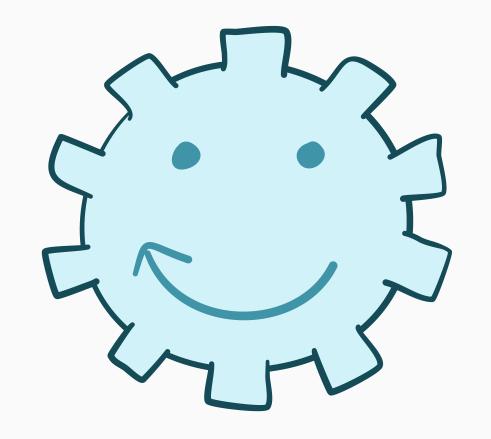


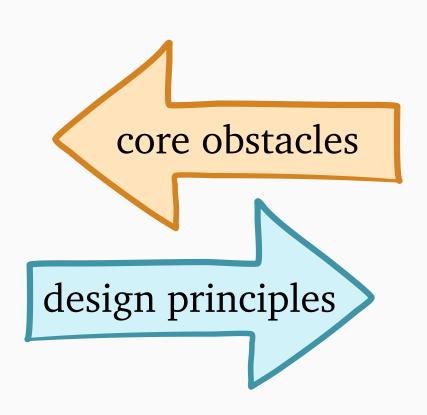
Gittins design principle:

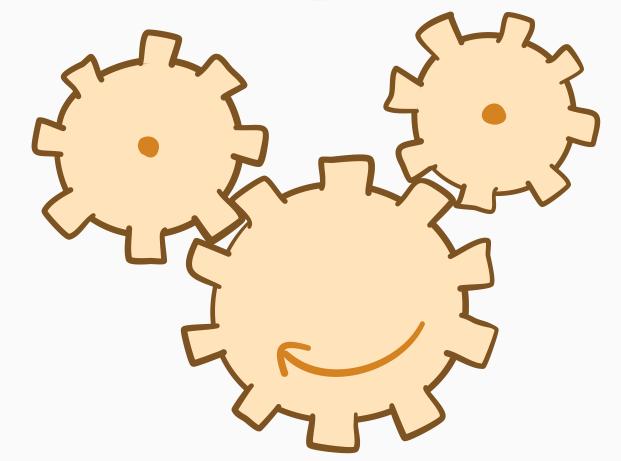
evaluate stochastic options by calibrating to deterministic options

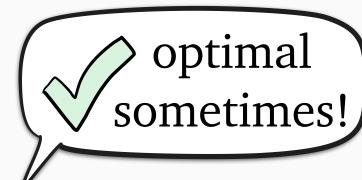
costly information multi-stage uncertainty heterogeneity

Practical problems



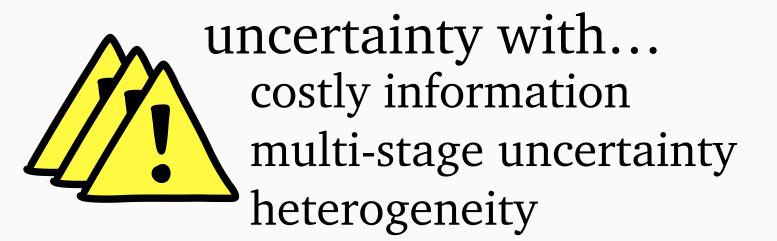




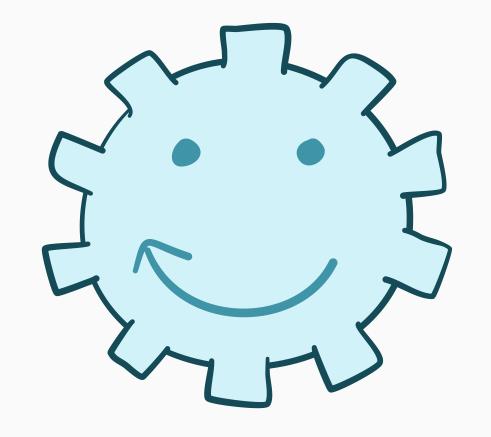


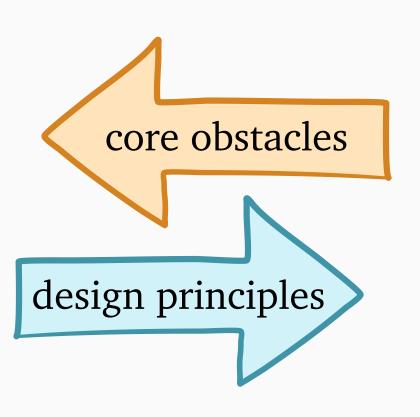


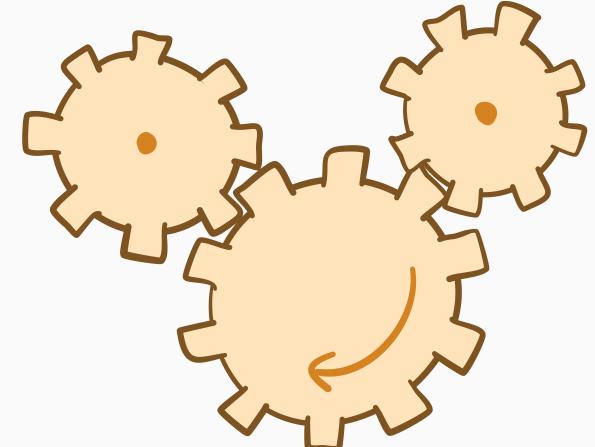
evaluate stochastic options by calibrating to deterministic options

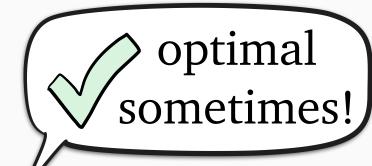


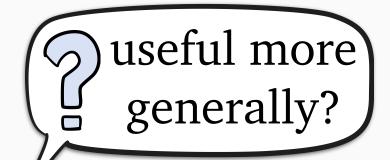
Practical problems





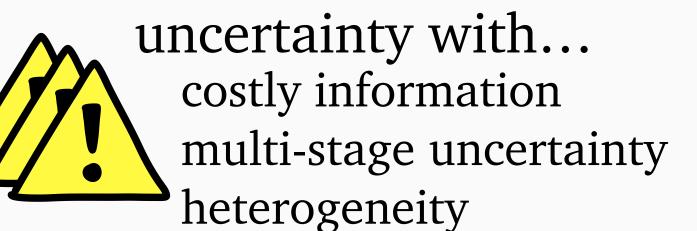




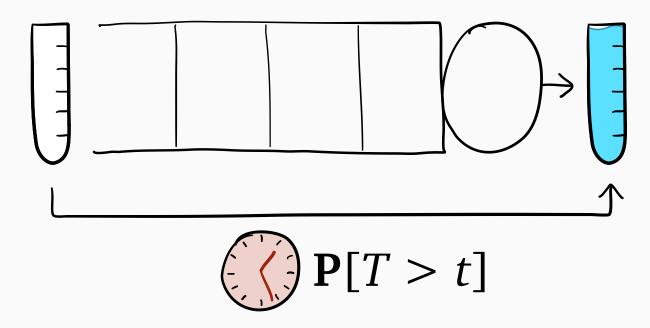




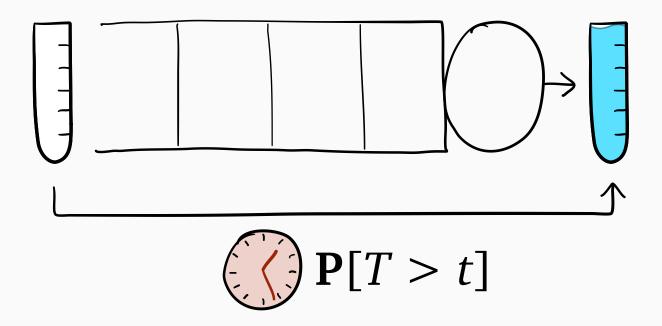
evaluate stochastic options by calibrating to deterministic options



Tail scheduling

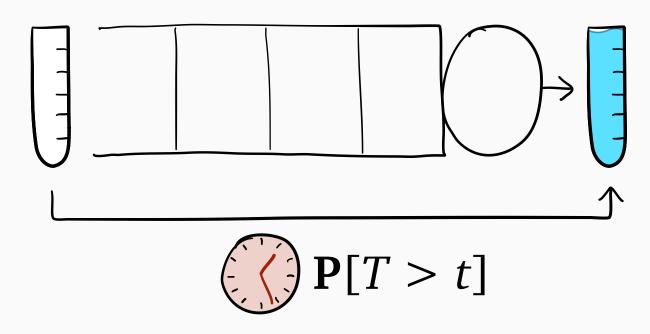


Tail scheduling

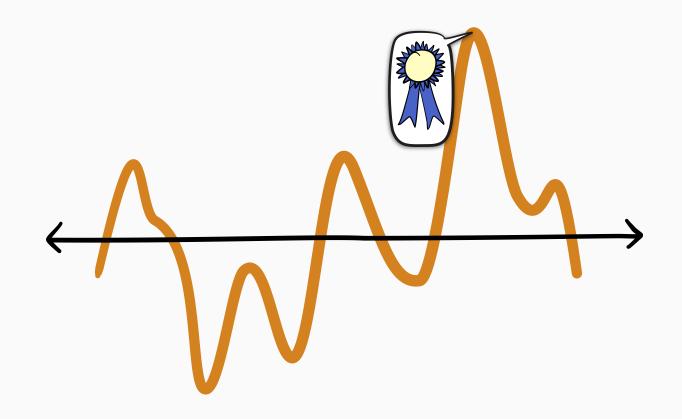


Goal: minimize probability of very long response time

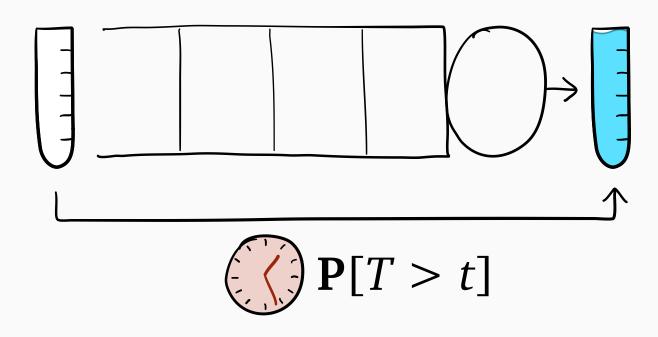
Tail scheduling



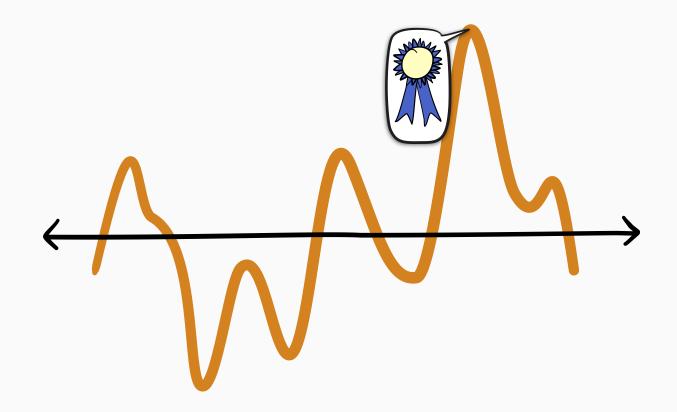
Goal: minimize probability of very long response time



Tail scheduling



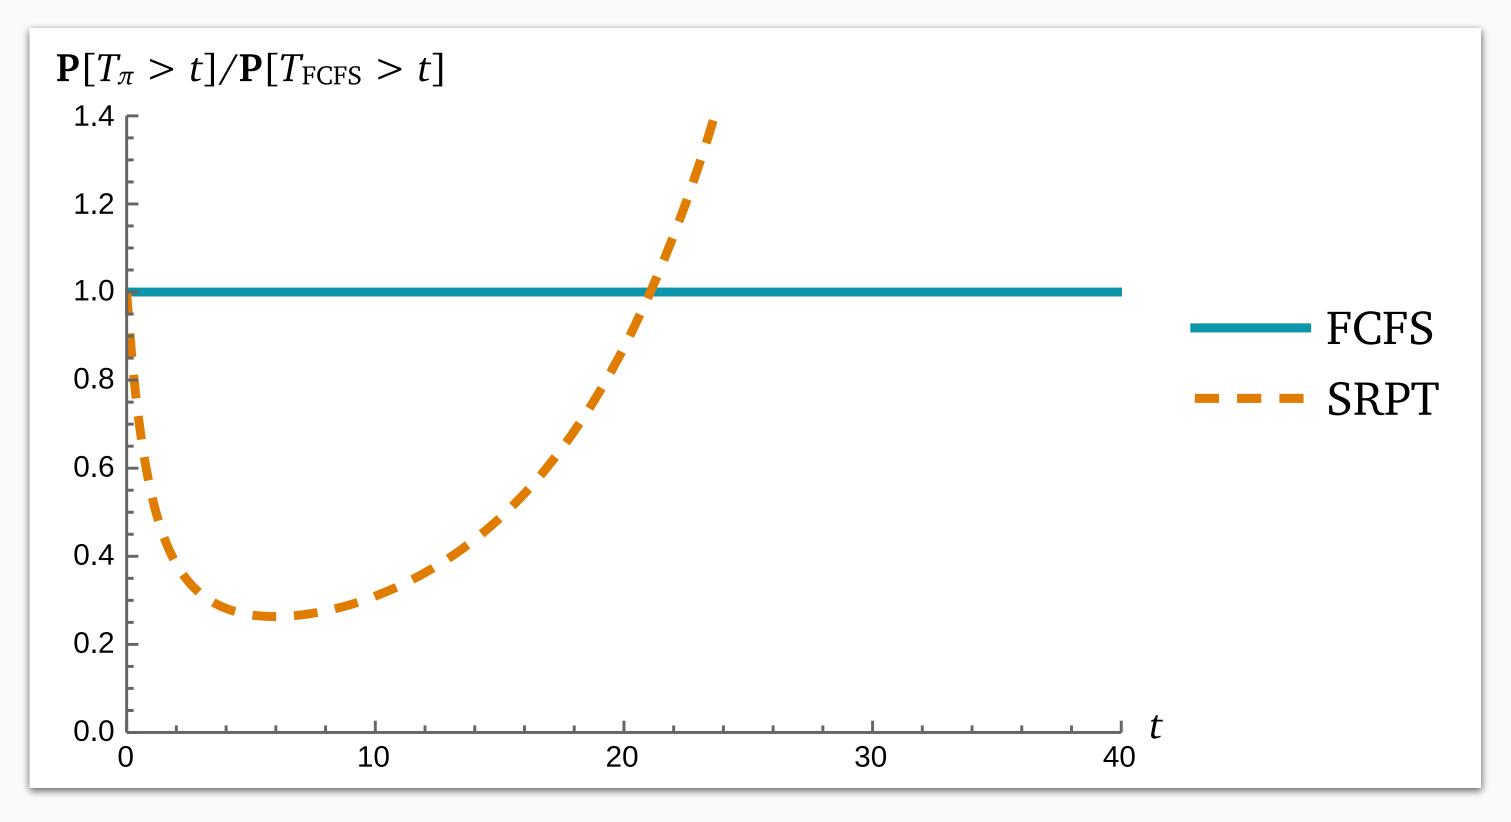
Goal: minimize probability of very long response time



Goal: find large function value with few function evaluations

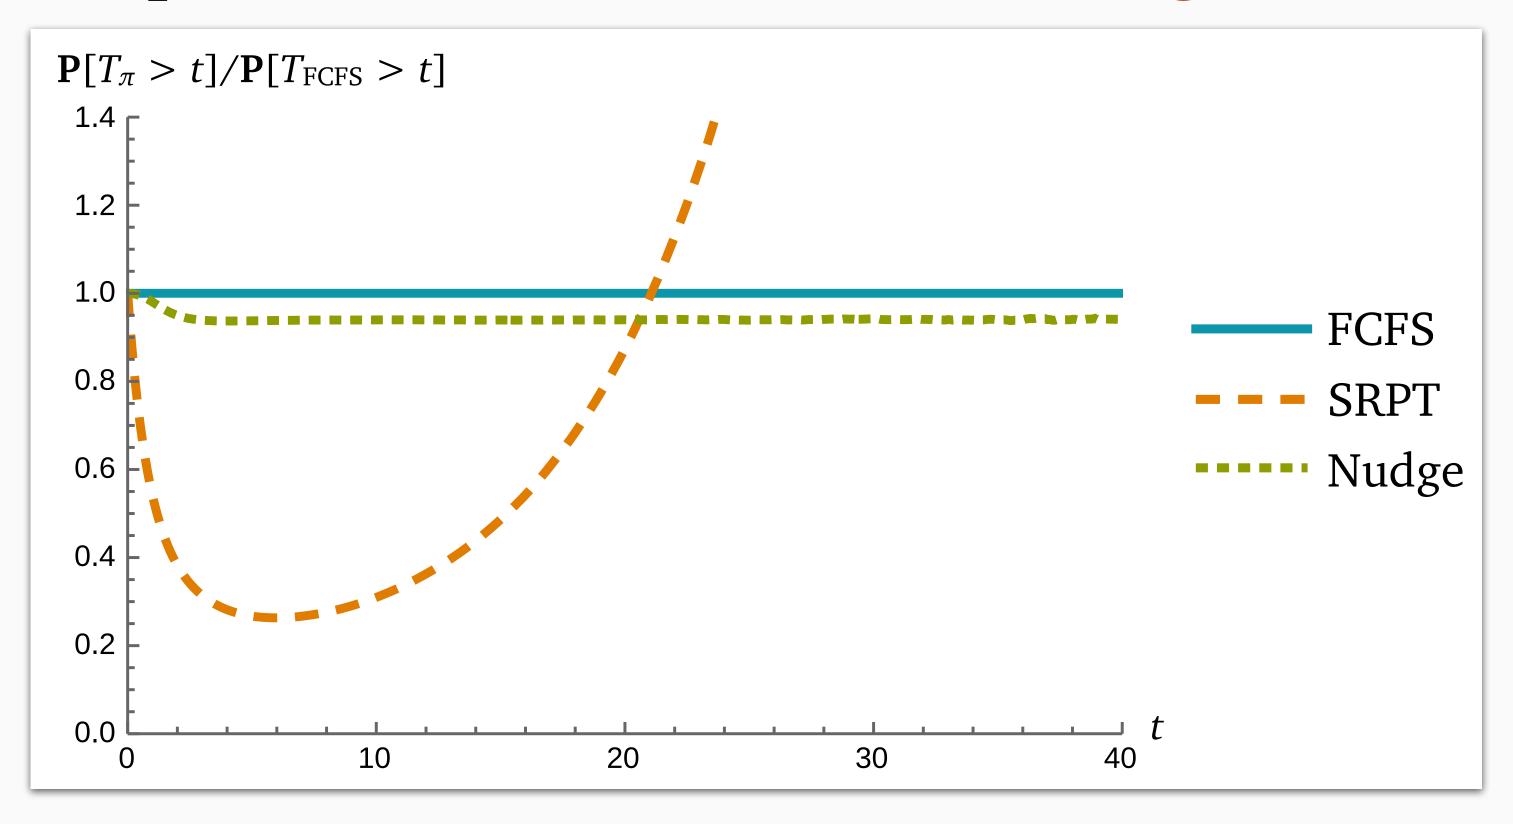


George Yu



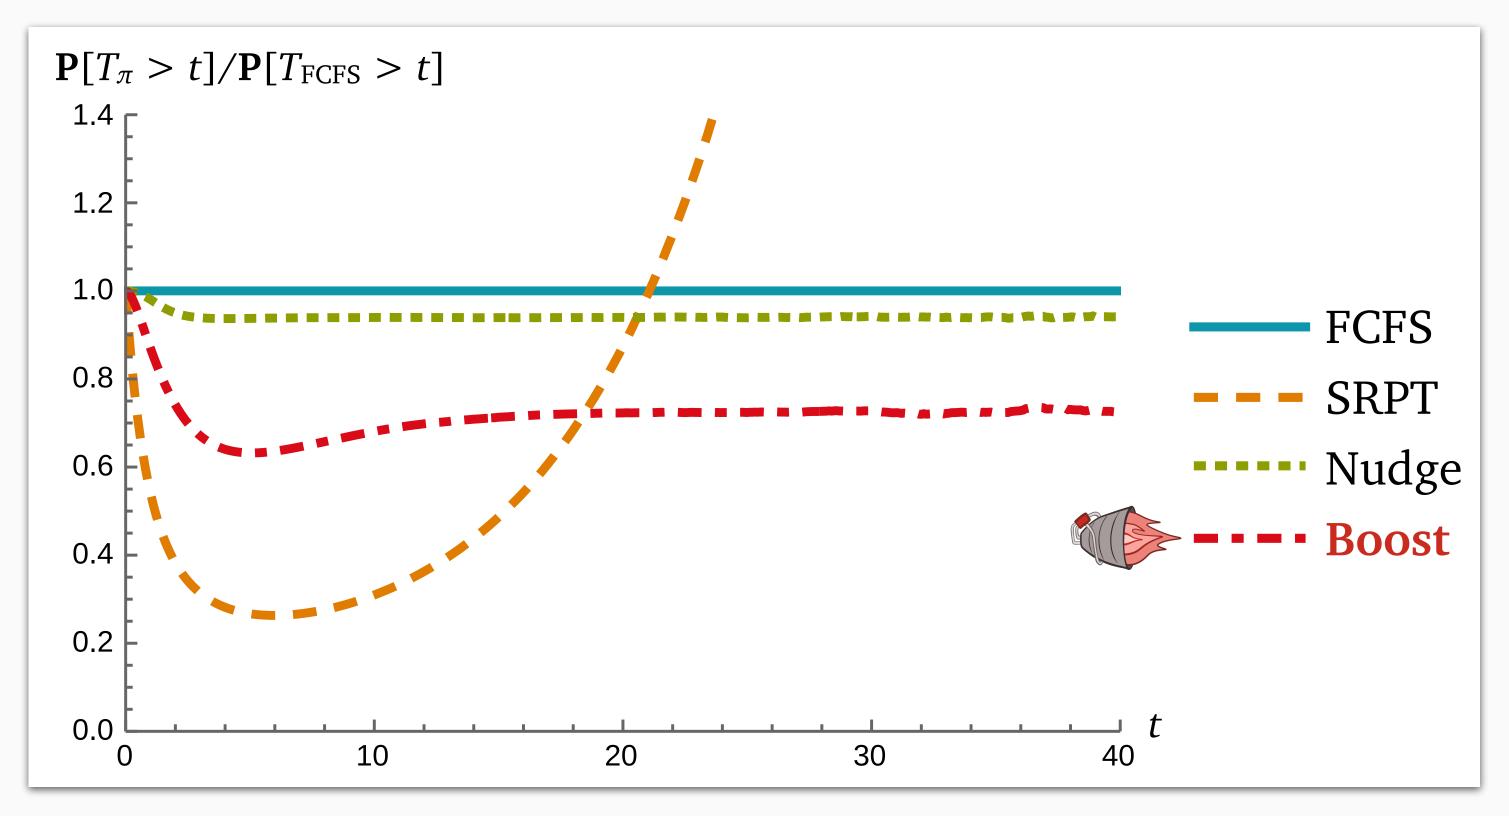


George Yu





George Yu

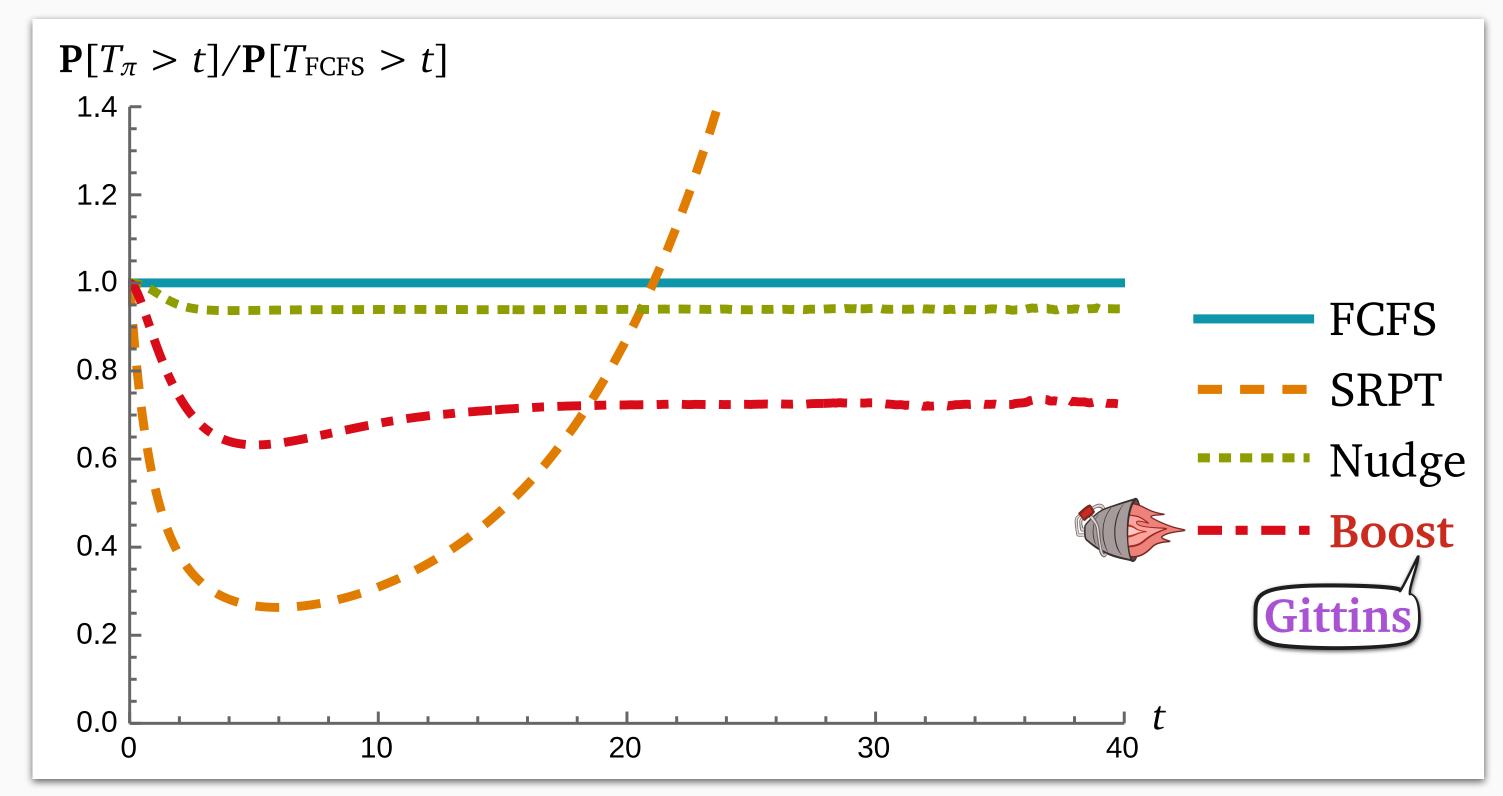




George Yu

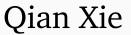


Amit Harley



See our SIGMETRICS papers [Yu & Scully, 2024; Harlev et al., 2025; Yu et al, 2026]



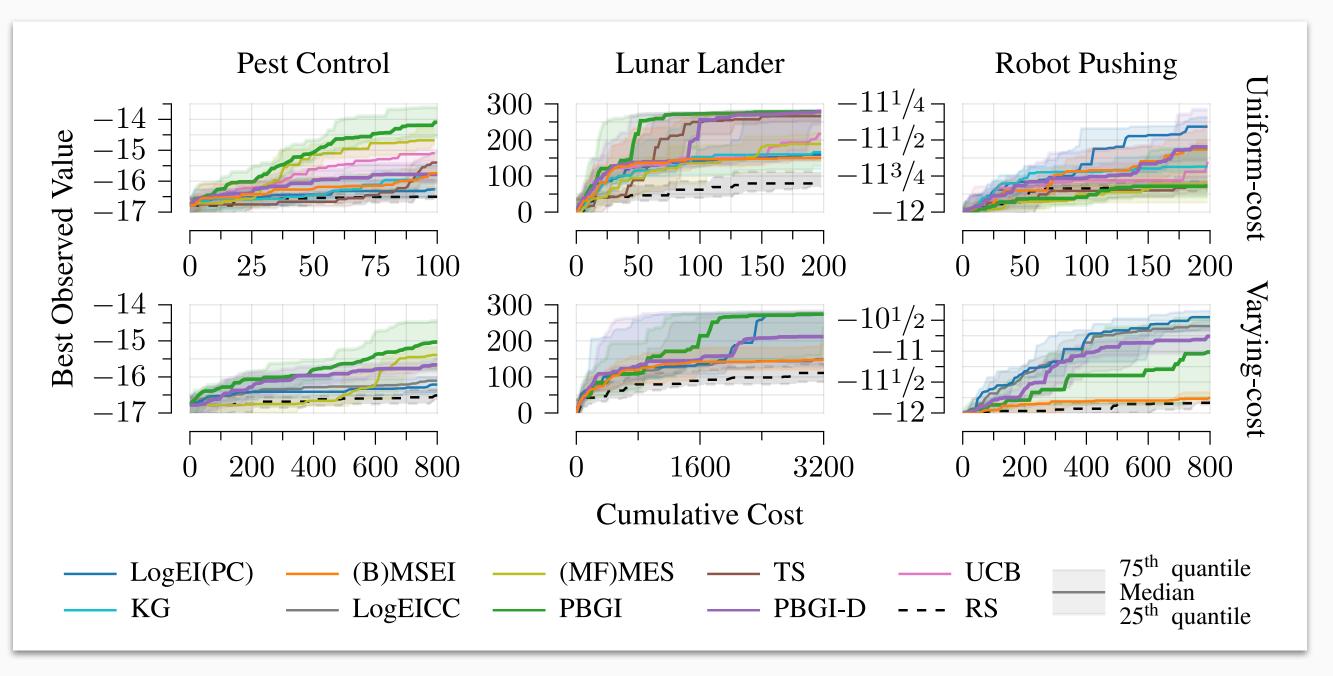




Raul Astudillo



Peter Frazier





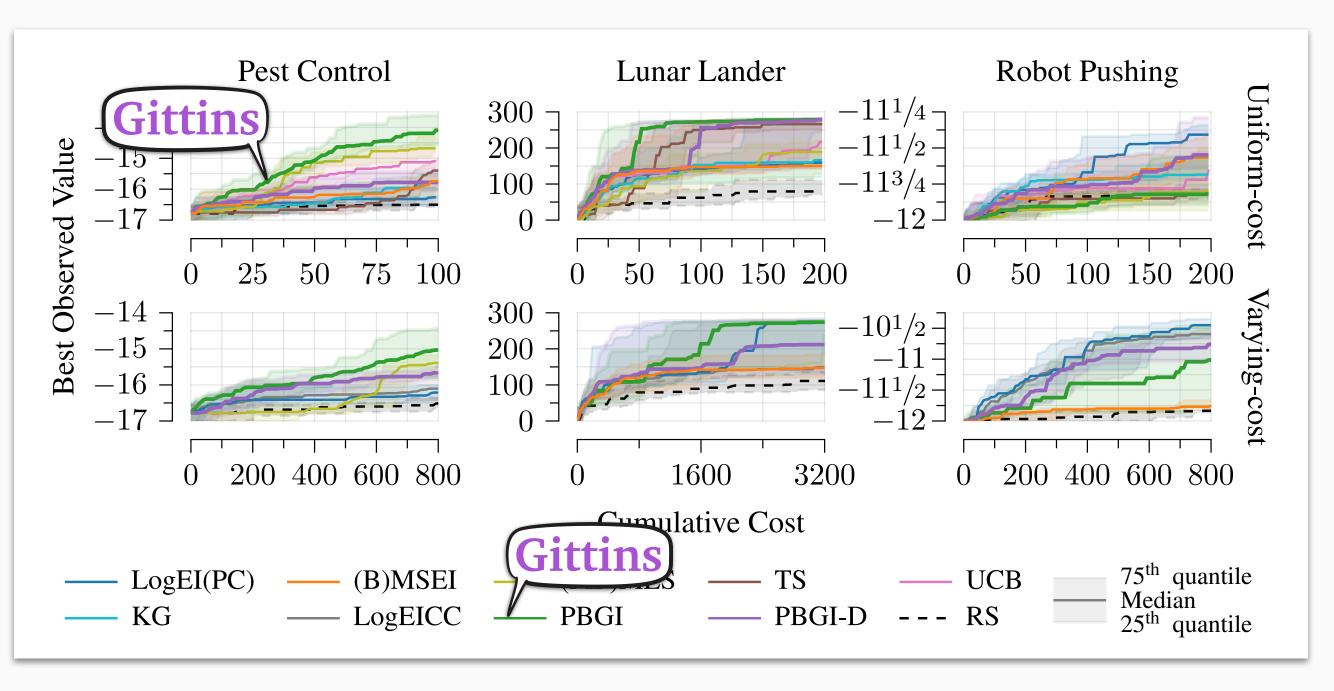




Raul Astudillo



Peter Frazier





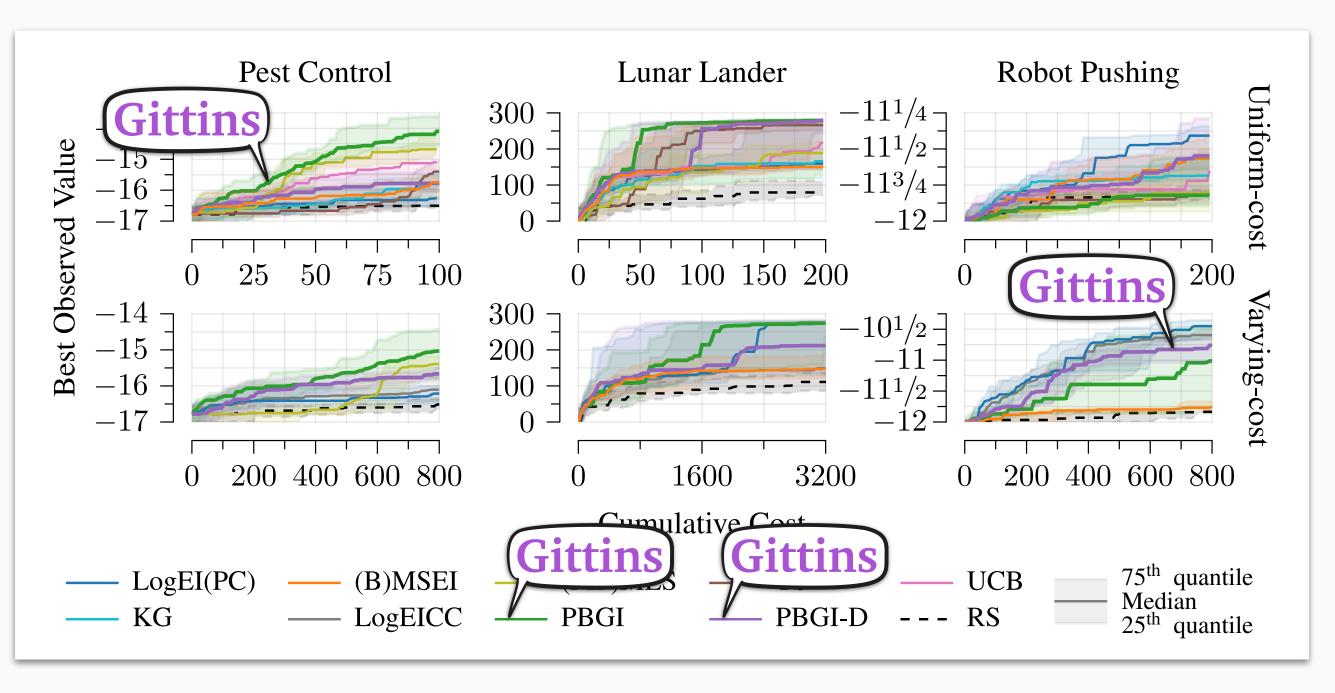


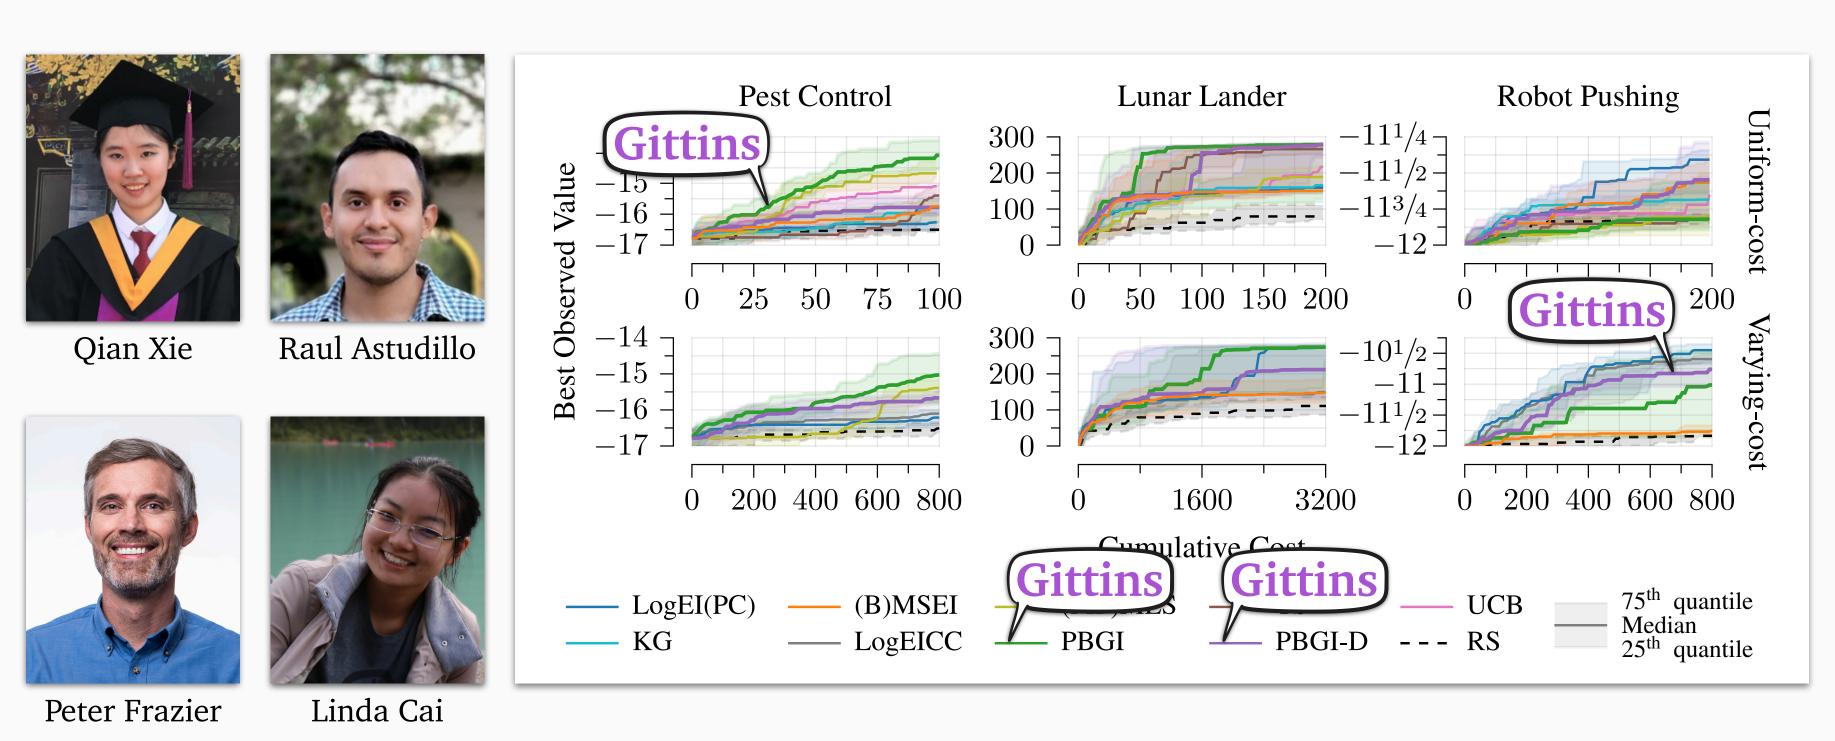


Raul Astudillo



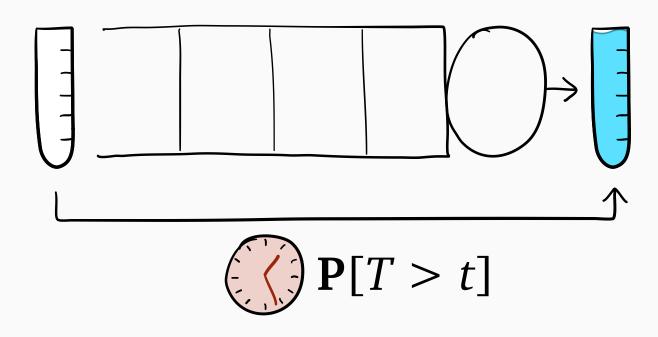
Peter Frazier



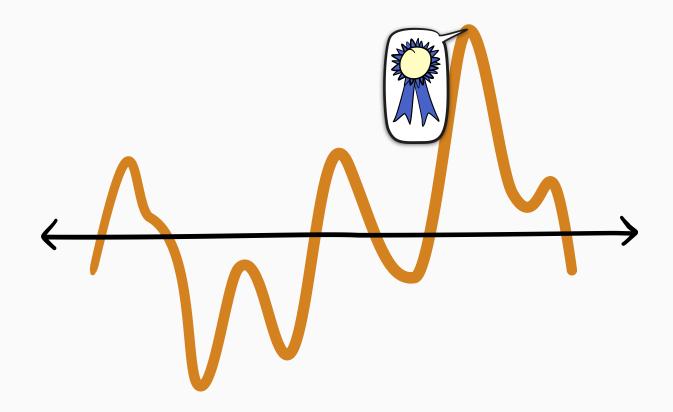


See our NeurIPS 2024 paper and arXiv preprint [Xie et al., 2024; Xie et al., 2025]

Tail scheduling

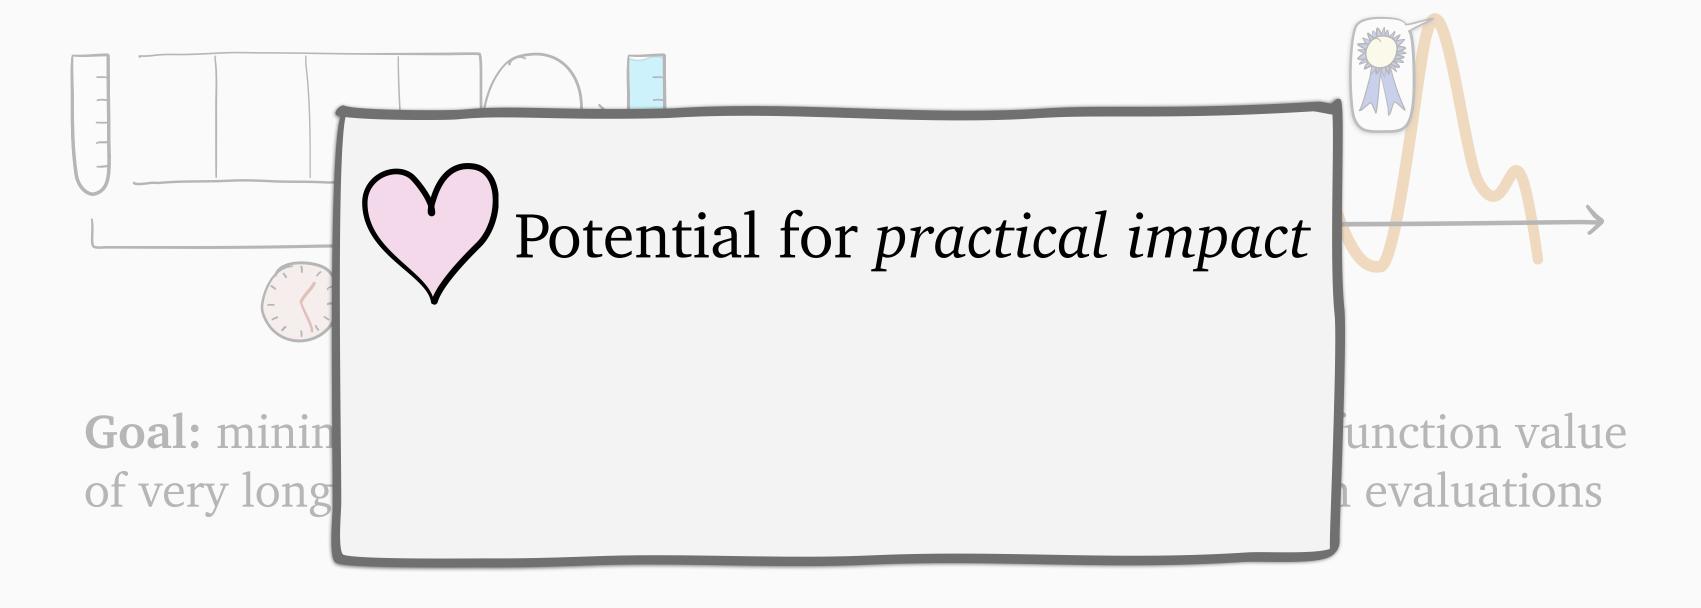


Goal: minimize probability of very long response time

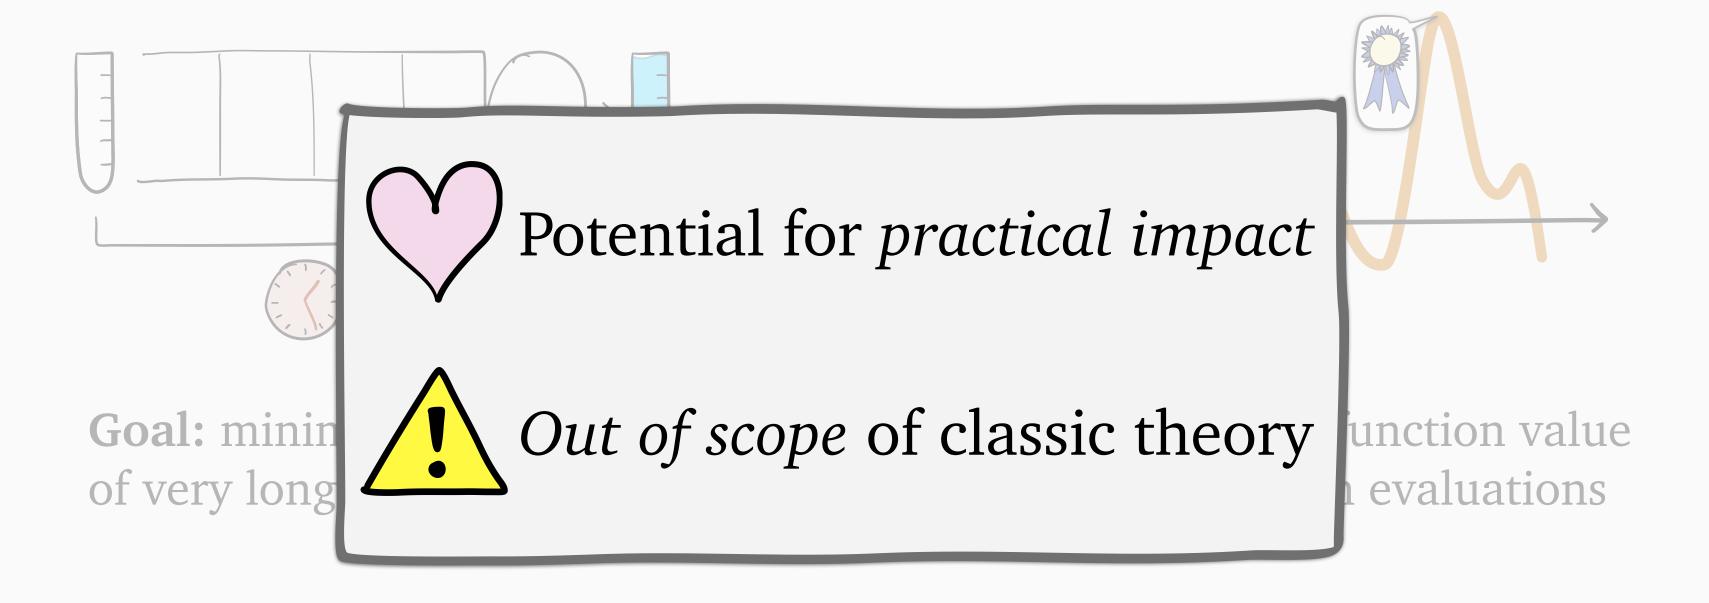


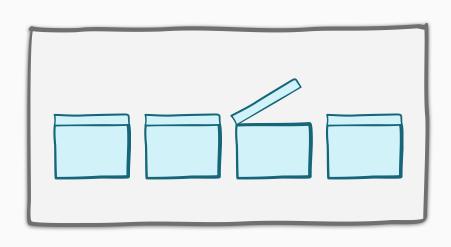
Goal: find large function value with few function evaluations

Tail scheduling

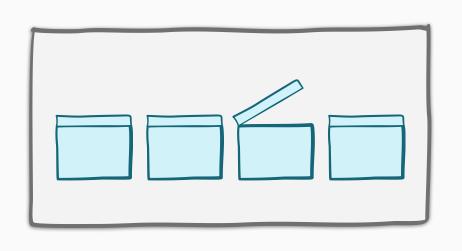


Tail scheduling

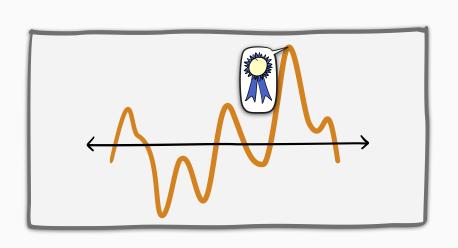




Introducing Gittins via Pandora's box

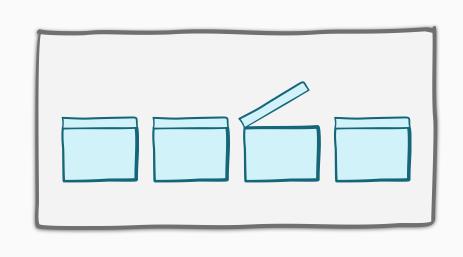


Introducing Gittins via Pandora's box

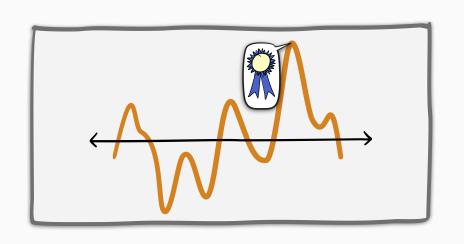


Part II:

Gittins for Bayesian optimization

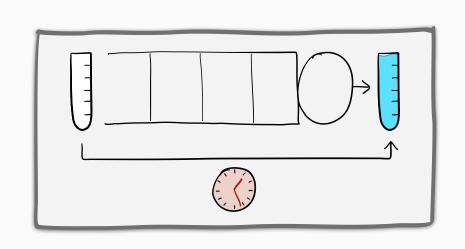


Introducing Gittins via Pandora's box



Part II:

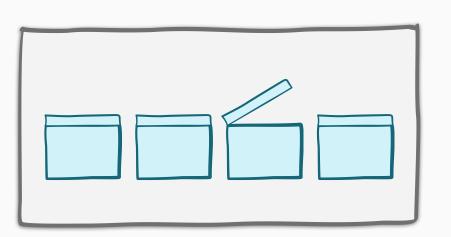
Gittins for Bayesian optimization

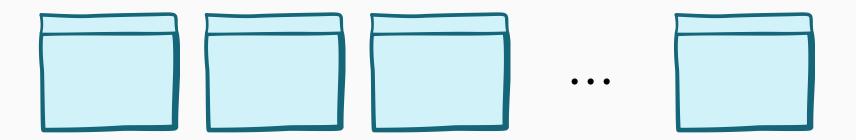


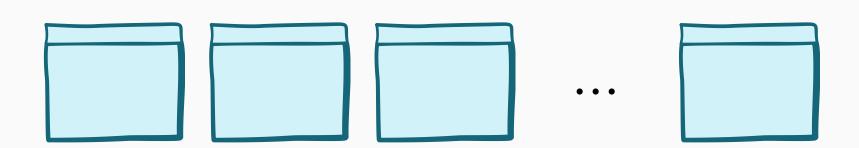
Part III:

Gittins for tail scheduling

Introducing Gittins via Pandora's box

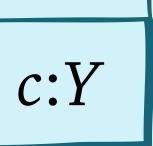


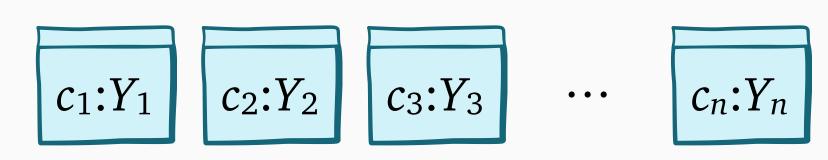




Each box:

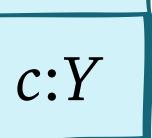
- Opening cost *c*
- Hidden reward *Y*

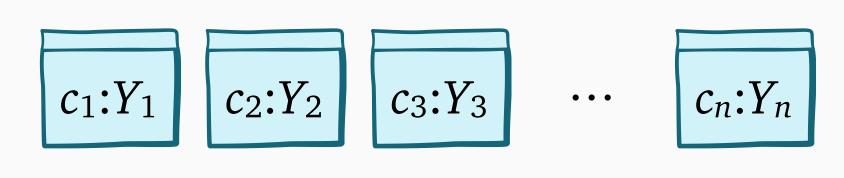




Each box:

- Opening cost *c*
- Hidden reward *Y*

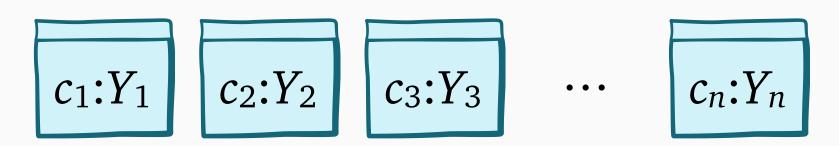




Each box:

- Opening cost *c*
- Hidden reward *Y* independent

c:Y



Decision process:

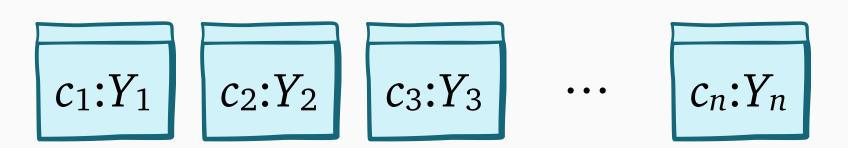
- Open boxes one at a time
- Stop by selecting open box

Each box:

- Opening cost *c*
- Hidden reward *Y*



c:Y



Decision process:

- Open boxes one at a time
- Stop by selecting open box

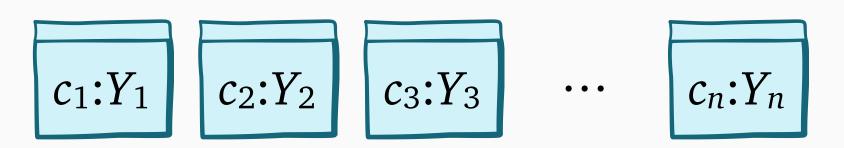
Goal: maximize
$$\mathbf{E} \left[Y_{\text{selected}} - \sum_{i \text{ opened}} c_i \right]$$

Each box:

- Opening cost *c*
- Hidden reward *Y*



c:Y



Decision process:

- Open boxes one at a time
- Stop by selecting open box

Goal: maximize
$$\mathbf{E} \left[Y_{\text{selected}} - \sum_{i \text{ opened}} c_i \right]$$

Each box:

- Opening cost *c*
- c:Y Hidden reward Y independent



$$c_1:Y_1$$
 $c_2:Y_2$ $c_3:Y_3$ \cdots $c_n:Y_n$

Decision process:

- · Open boxes one at a time
- Stop by selecting open box

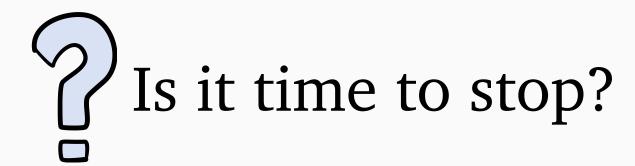
Goal: maximize
$$\mathbf{E} \left[Y_{\text{selected}} - \sum_{i \text{ opened}} c_i \right]$$

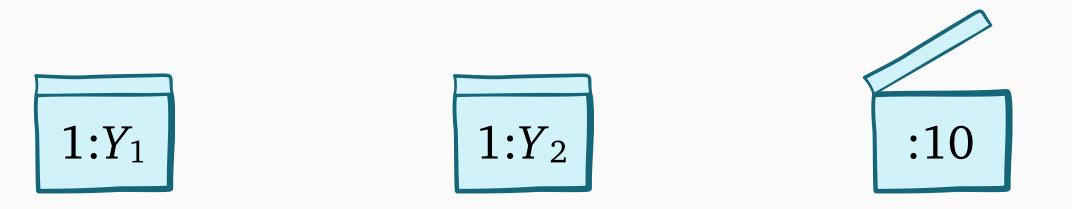
Each box:

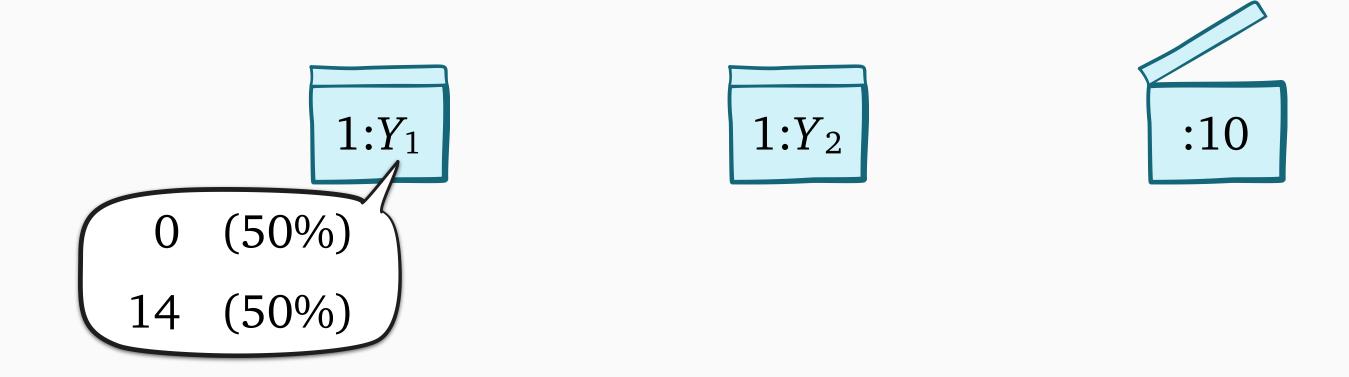
- Opening cost *c*
- Hidden reward *Y* independent

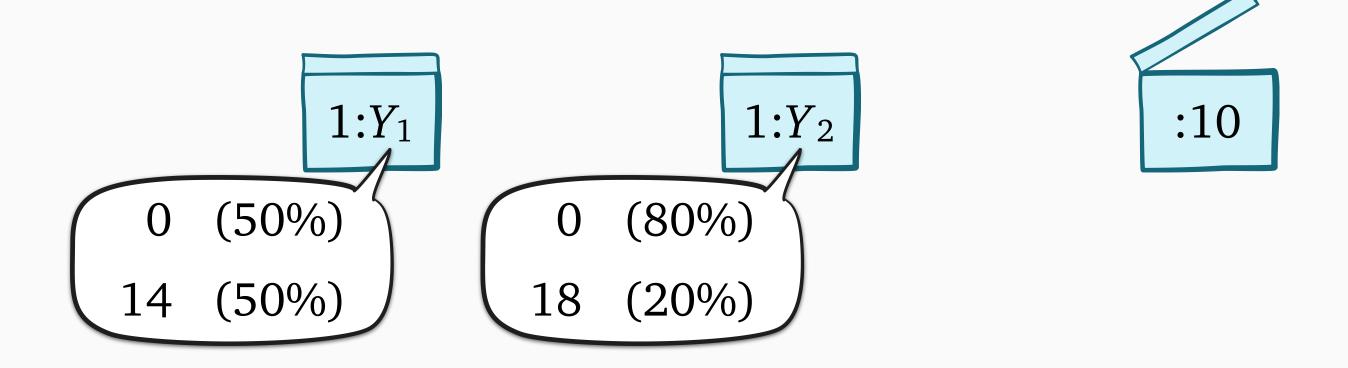






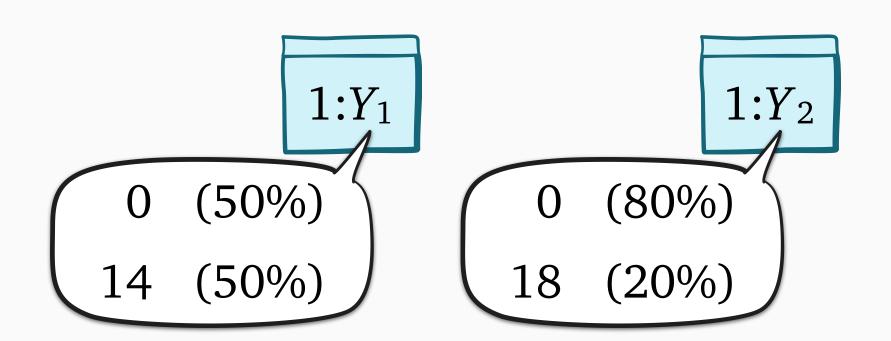


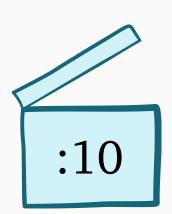






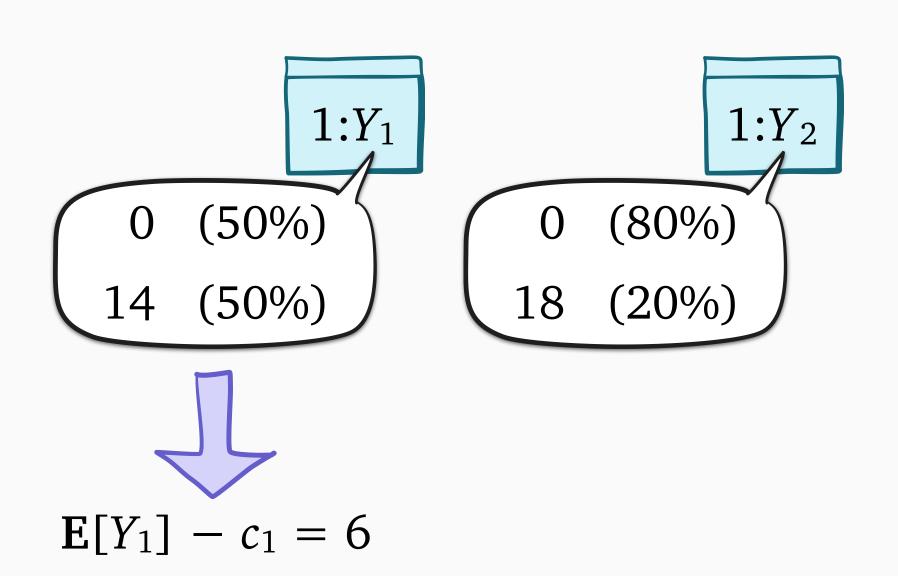
Use expected *net value*? $\mathbf{E}[Y] - c$

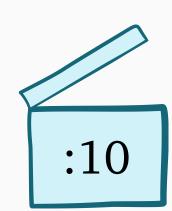






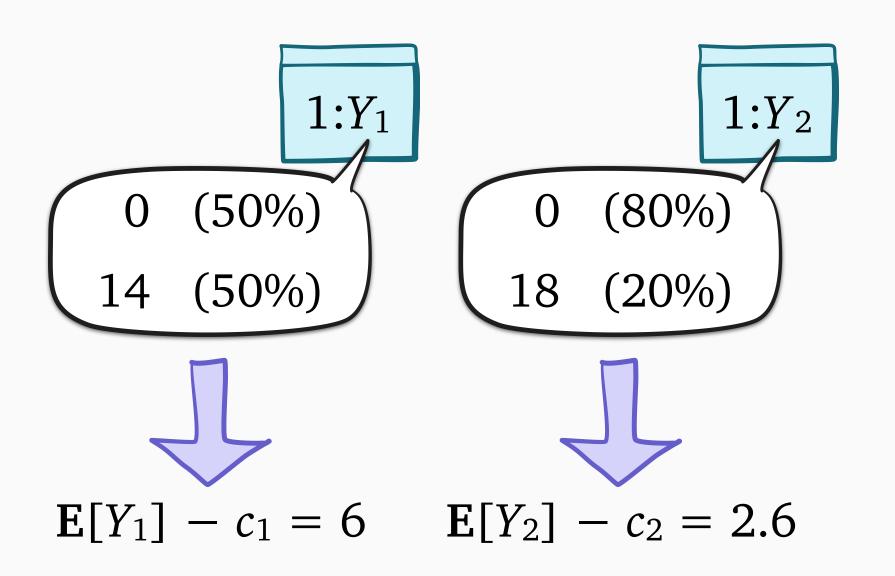
Use expected *net value*? $\mathbf{E}[Y] - c$

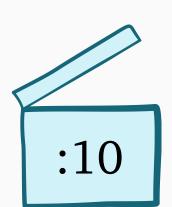






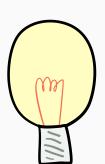
Use expected *net value*? $\mathbf{E}[Y] - c$



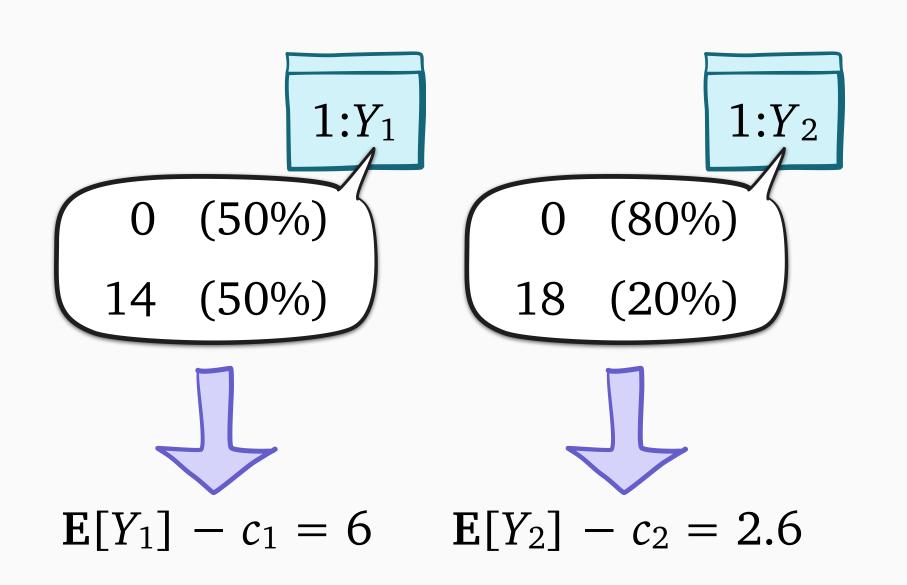


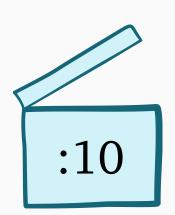


Use expected *net value*? E[Y] - c



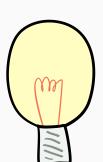
$$EI(Y, r) = E[max(Y - r, 0)]$$



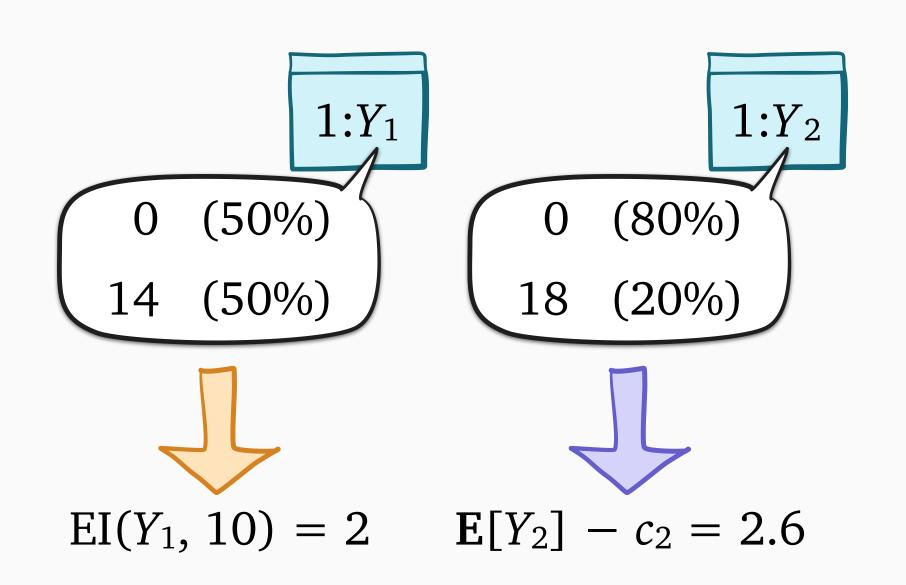




Use expected *net value*? E[Y] - c



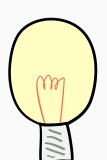
$$EI(Y, r) = E[max(Y - r, 0)]$$



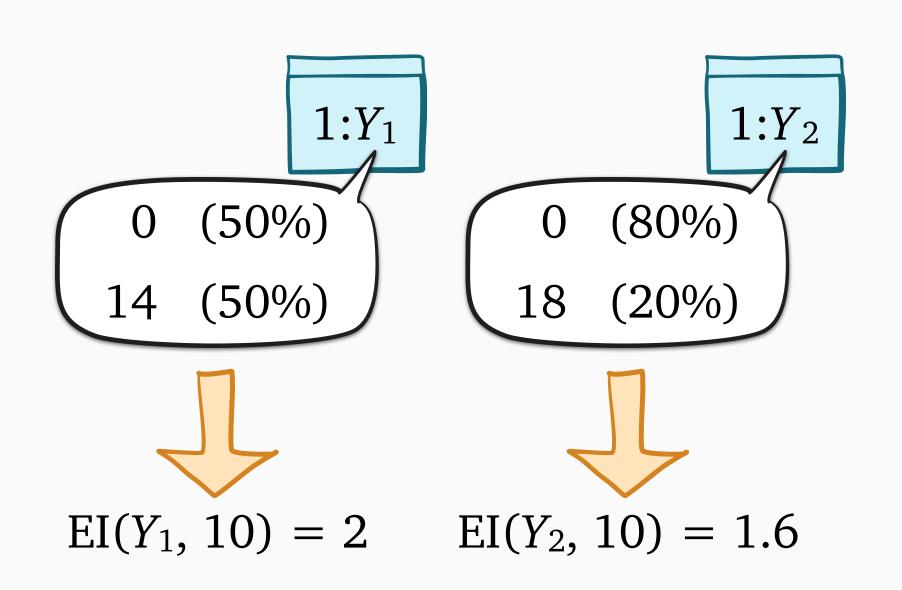


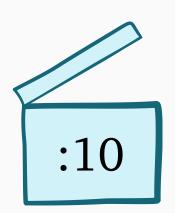


Use expected *net value*? E[Y] - c



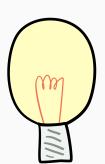
$$EI(Y, r) = E[max(Y - r, 0)]$$





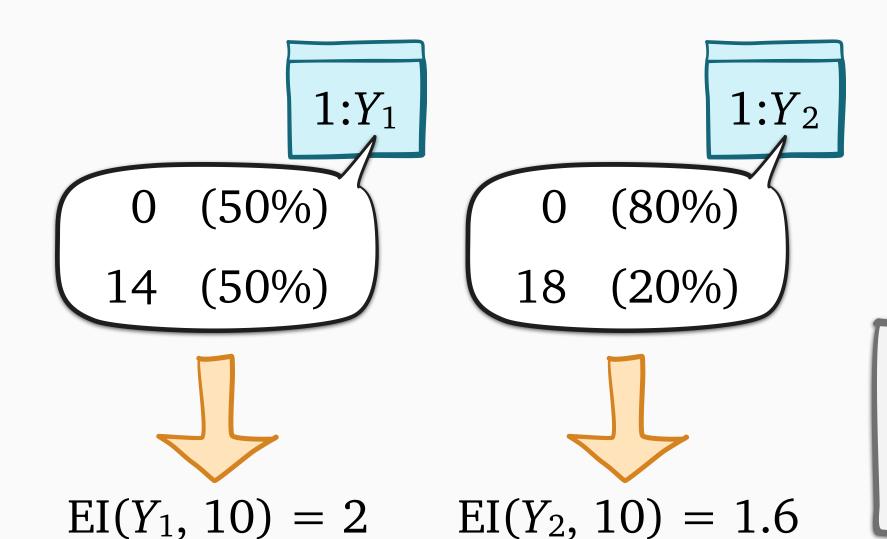


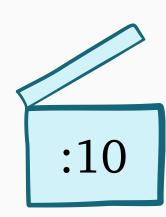
Use expected *net value*? $\mathbf{E}[Y] - c$



Expected improvement of Y over r:

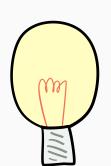
$$EI(Y, r) = E[max(Y - r, 0)]$$





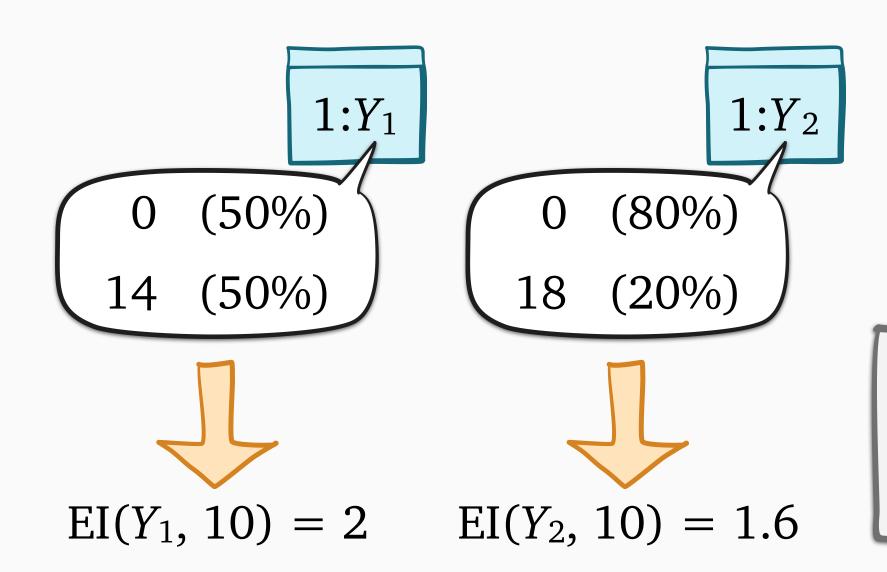
• Both boxes have $EI(Y_i, 10) > c_i$

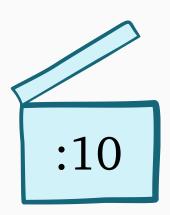




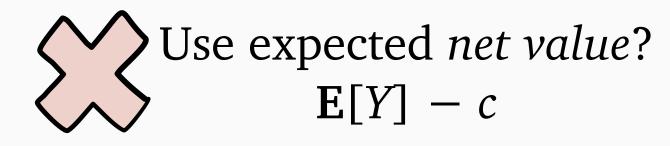
Expected improvement of Y over r:

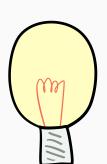
$$EI(Y, r) = E[max(Y - r, 0)]$$



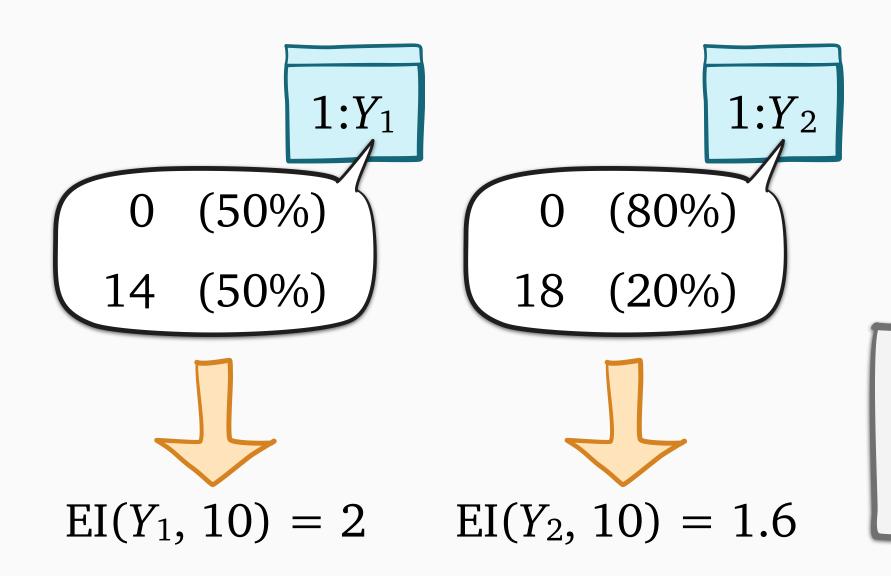


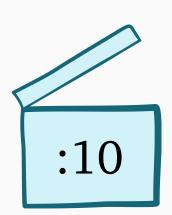
• Both boxes have $EI(Y_i, 10) > c_i$



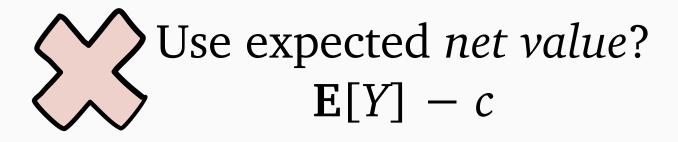


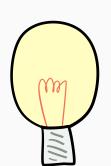
$$EI(Y, r) = E[max(Y - r, 0)]$$



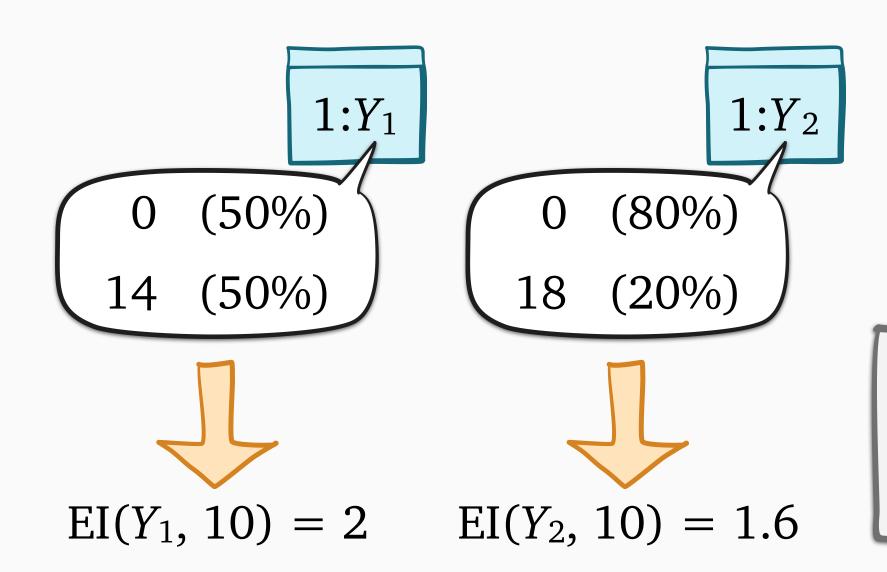


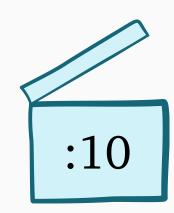
- Both boxes have $EI(Y_i, 10) > c_i$
- Box 1 has better EI





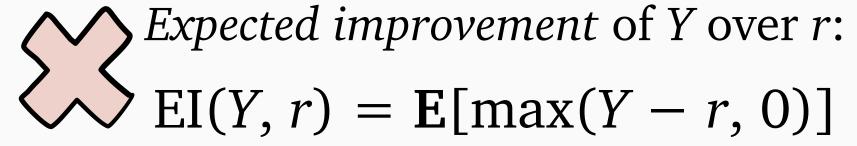
$$EI(Y, r) = E[max(Y - r, 0)]$$

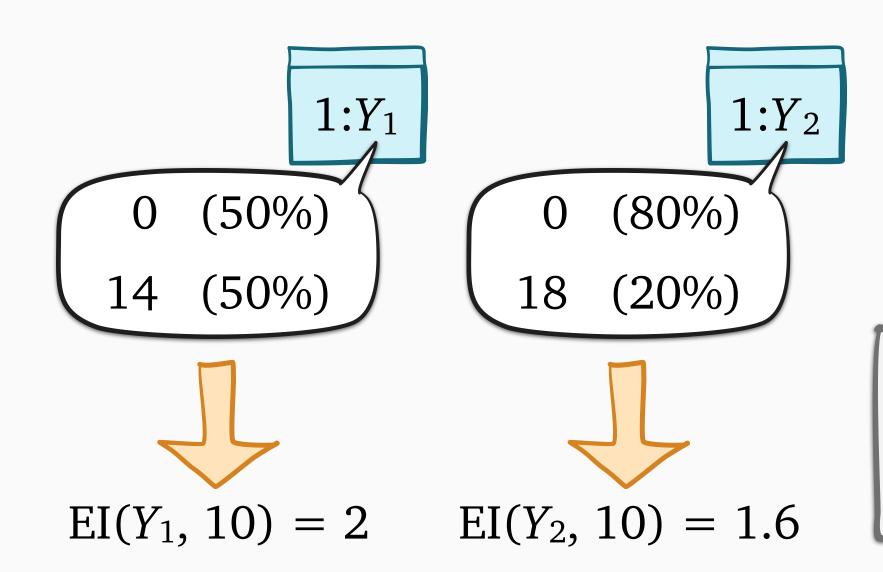


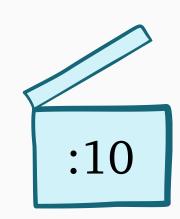


- Both boxes have $EI(Y_i, 10) > c_i$
- Box 1 has better EI
- Optimal action: open box 2!





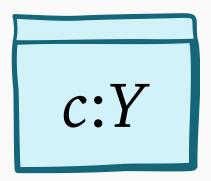




- Both boxes have $EI(Y_i, 10) > c_i$
- Box 1 has better EI
- Optimal action: open box 2!

Step 1: *rate* each box separately

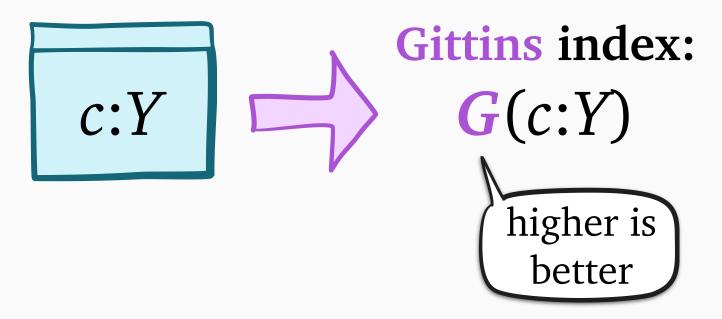
Step 1: *rate* each box separately



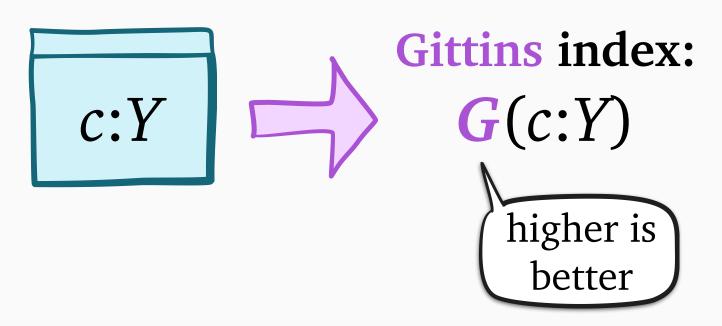
Step 1: *rate* each box separately

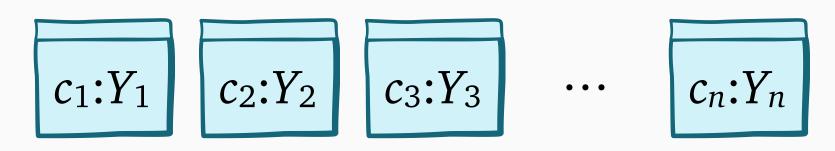


Step 1: *rate* each box separately

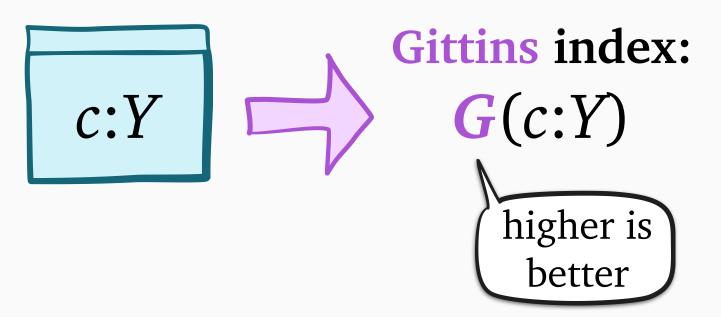


Step 1: *rate* each box separately

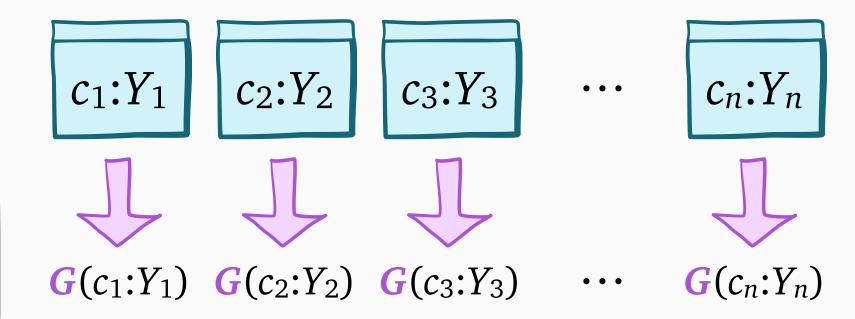




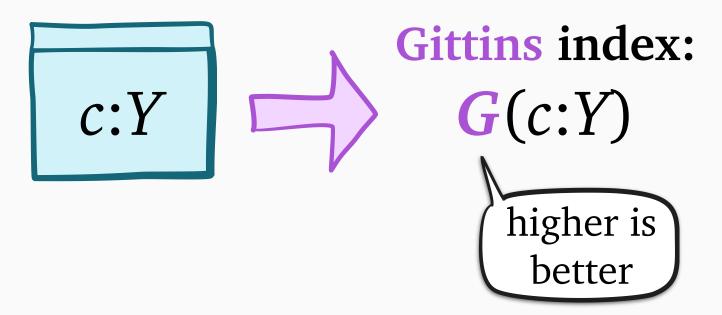
Step 1: *rate* each box separately



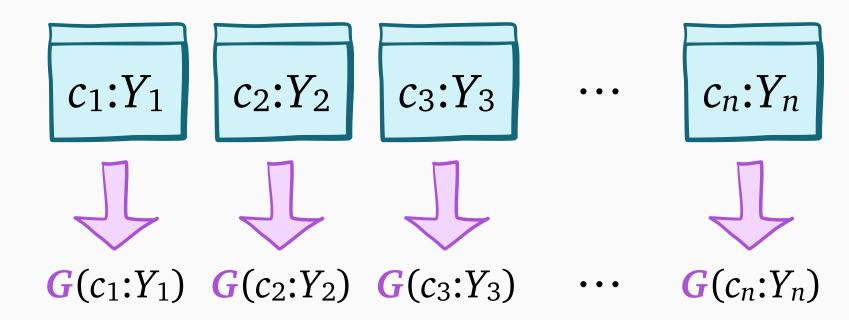
Step 2: *act* on box of best rating



Step 1: *rate* each box separately



Step 2: act on box of best rating



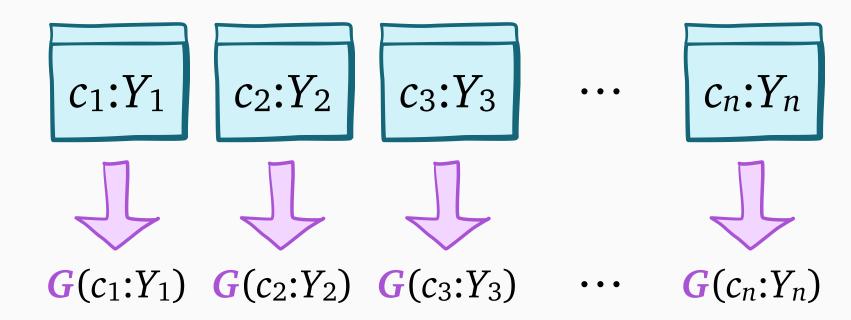
Gittins policy: if box of max Gittins index is...

- closed: open it
- open: select it

Step 1: *rate* each box separately

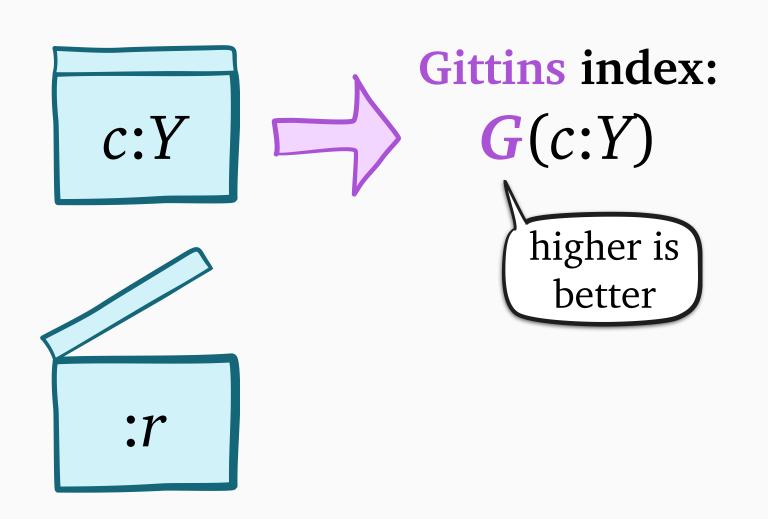


Step 2: act on box of best rating

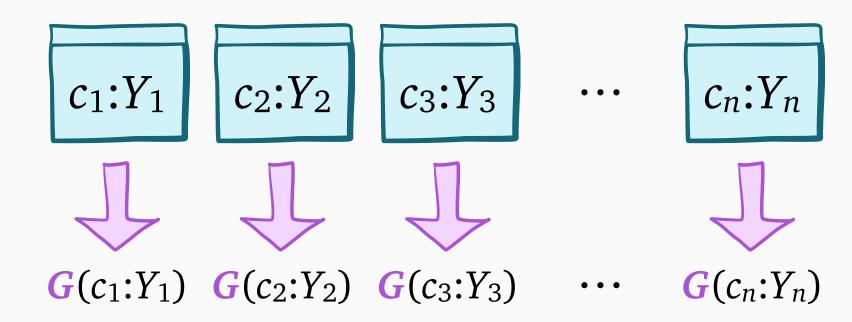


Gittins policy: if box of max Gittins index is...

Step 1: *rate* each box separately

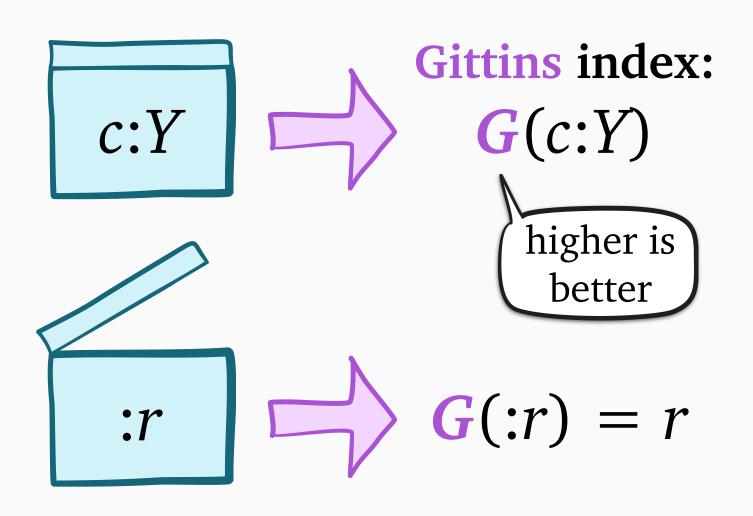


Step 2: act on box of best rating

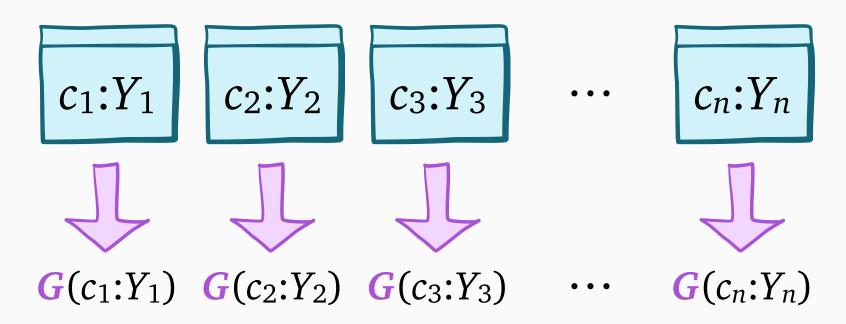


Gittins policy: if box of max Gittins index is...

Step 1: *rate* each box separately

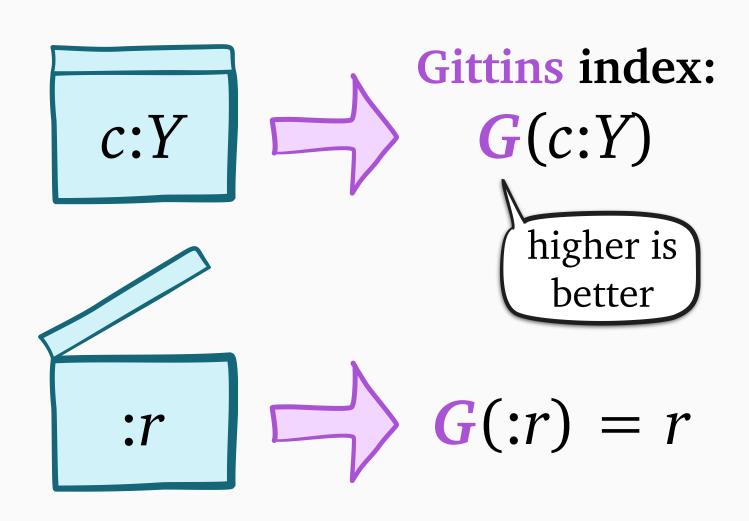


Step 2: act on box of best rating



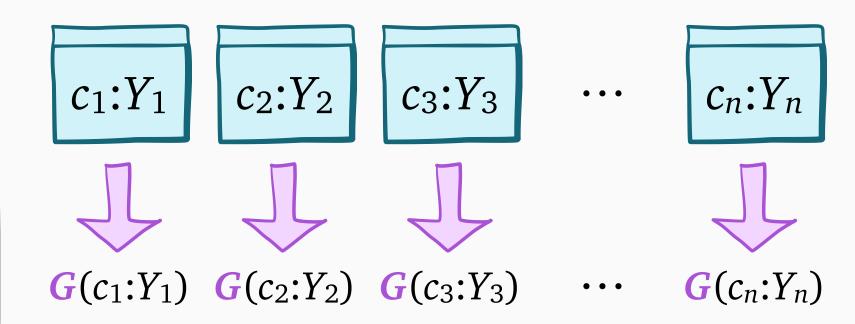
Gittins policy: if box of max Gittins index is...

Step 1: *rate* each box separately



Theorem: [Weitzman, 1979] the Gittins policy is optimal

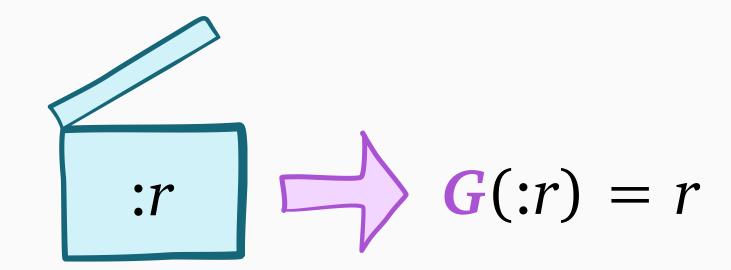
Step 2: act on box of best rating



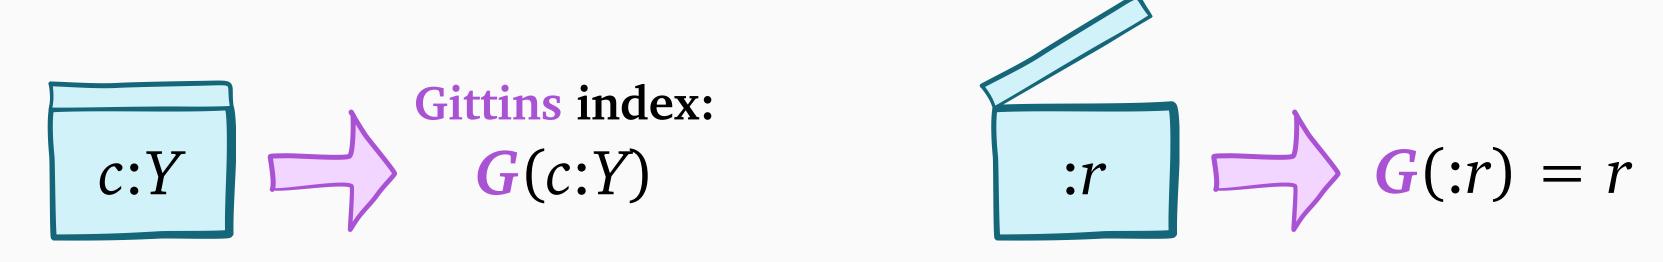
Gittins policy: if box of max Gittins index is...

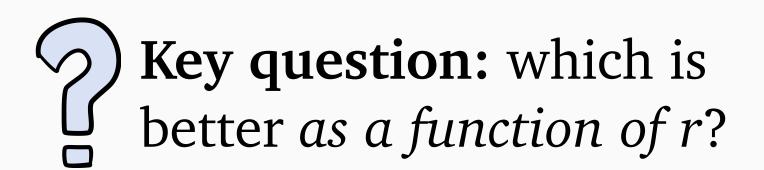
Gittins index of a box



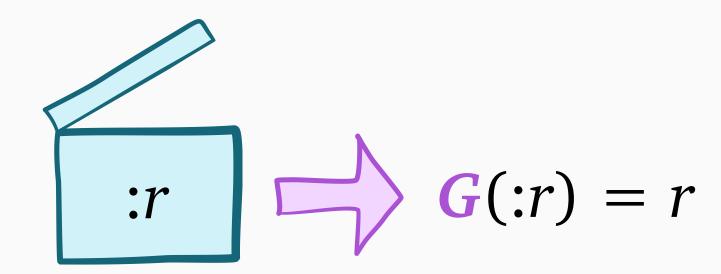


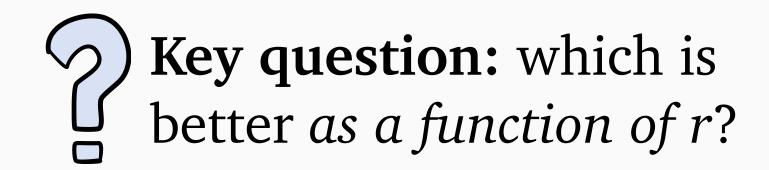
Gittins index of a box





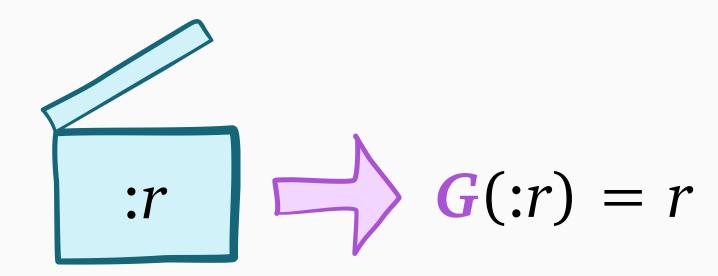


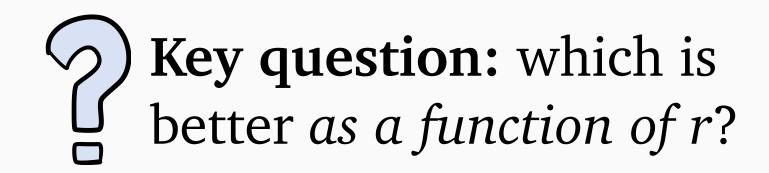


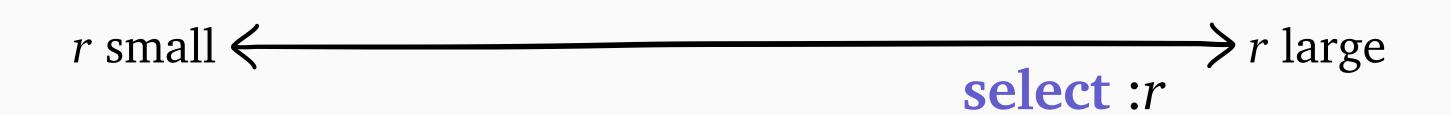


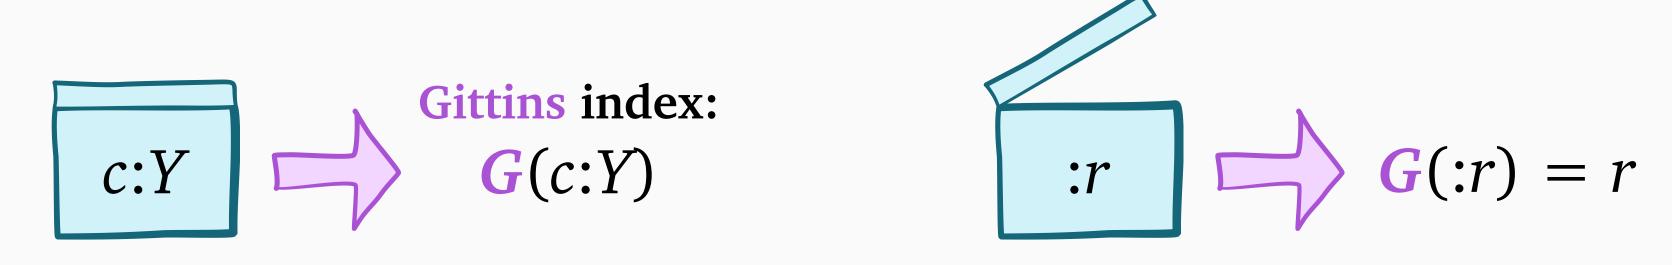
 $r \text{ small} \longleftrightarrow r \text{ large}$

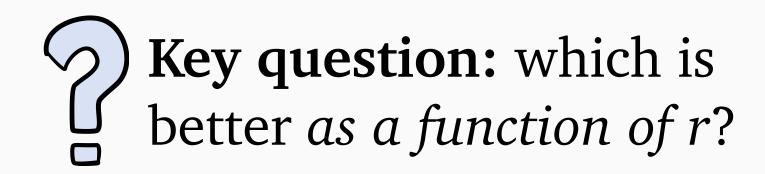


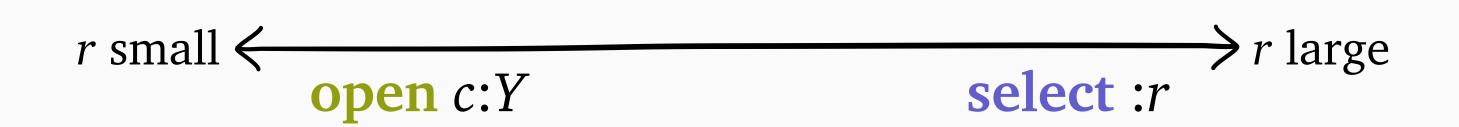




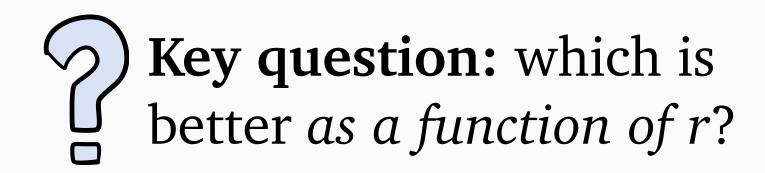


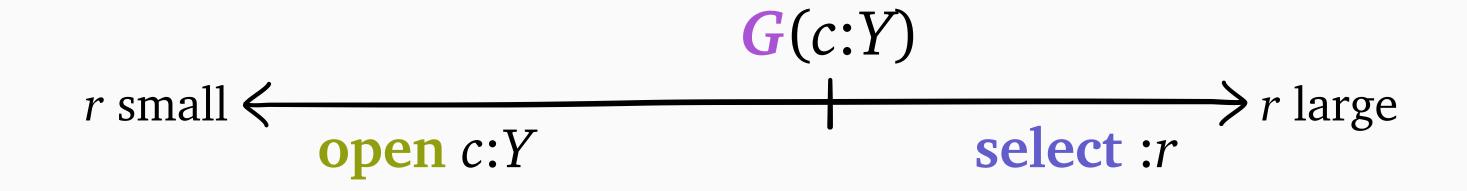


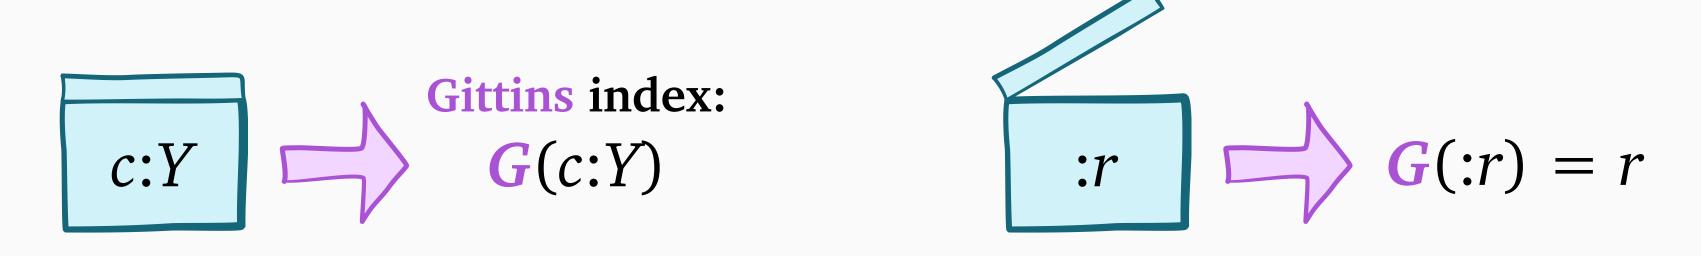


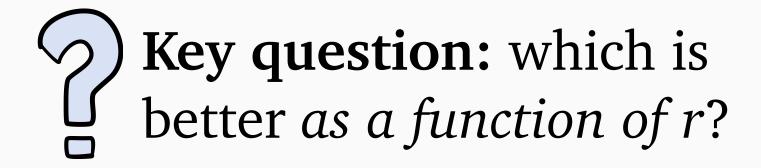


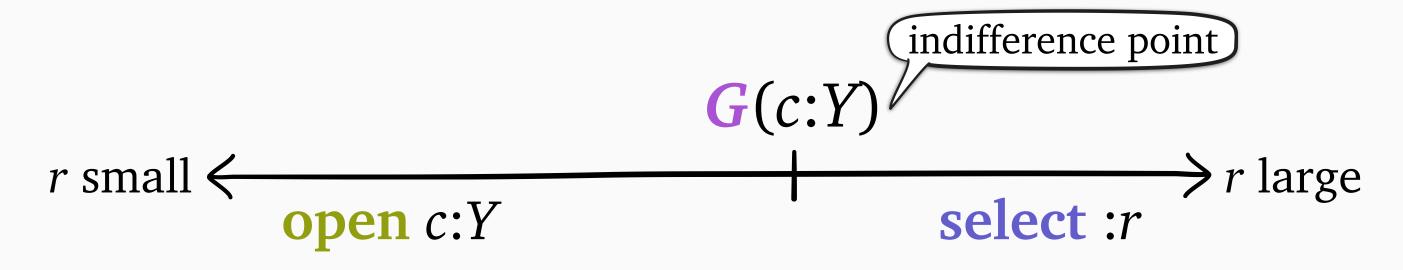






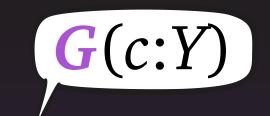






The *Gittins index* of *c*:*Y* is the value of *r* such that we're indifferent between "open *c*:*Y*" and "select :*r*"



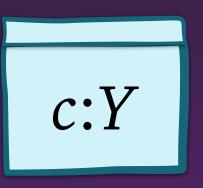


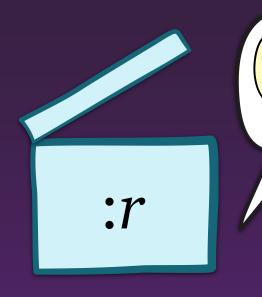
The *Gittins index* of *c*:*Y* is the value of *r* such that we're indifferent between "open *c*:*Y*" and "select :*r*"





The *Gittins index* of *c*:*Y* is the value of *r* such that we're indifferent between "open *c*:*Y*" and "select :*r*"



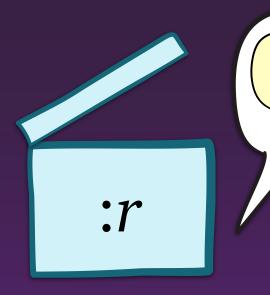


evaluate stochastic options by calibrating to deterministic options

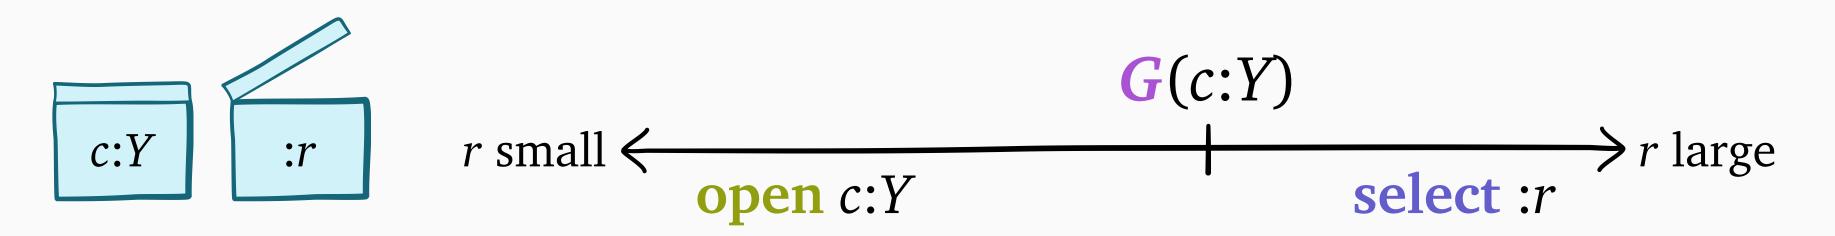
G(???)

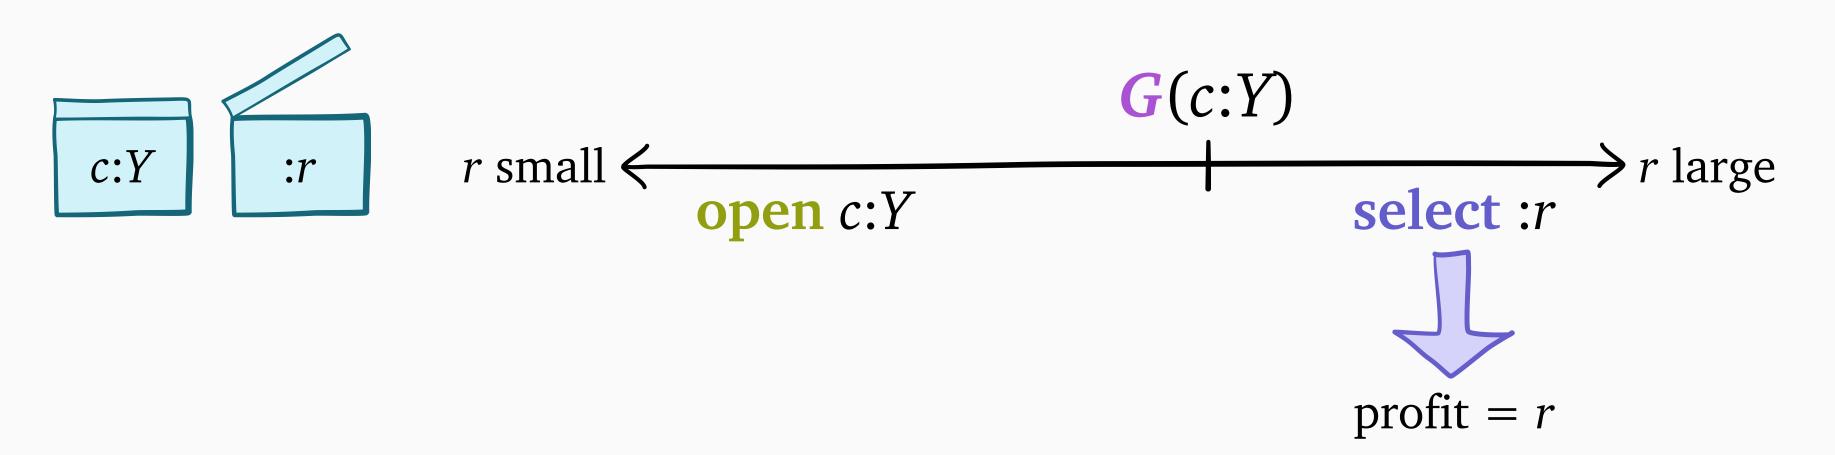
The *Gittins index* of ??? is the value of r such that we're indifferent between "act on ???" and "select :r"

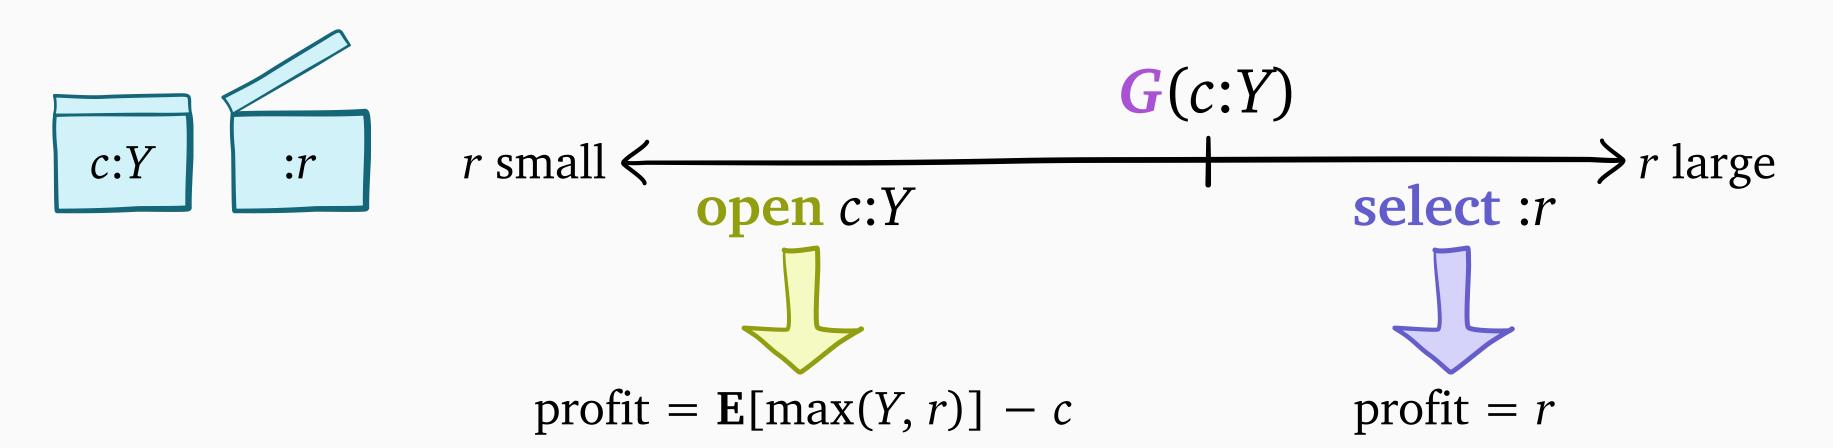


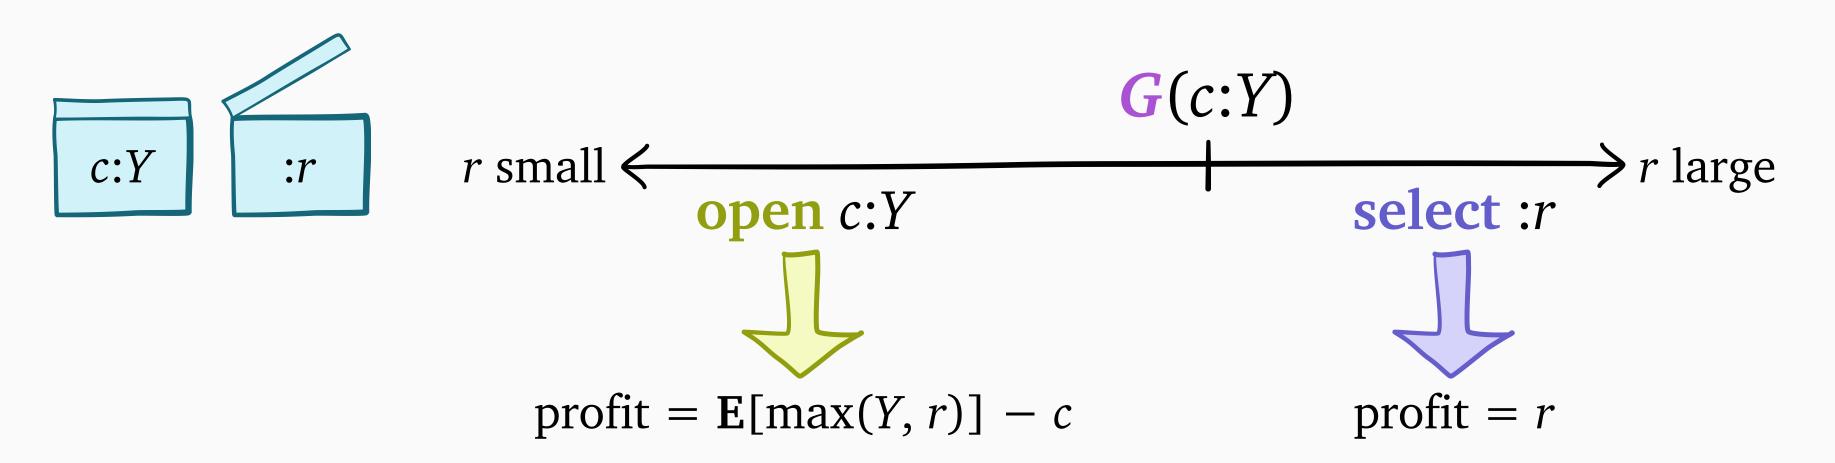


evaluate stochastic options by calibrating to deterministic options

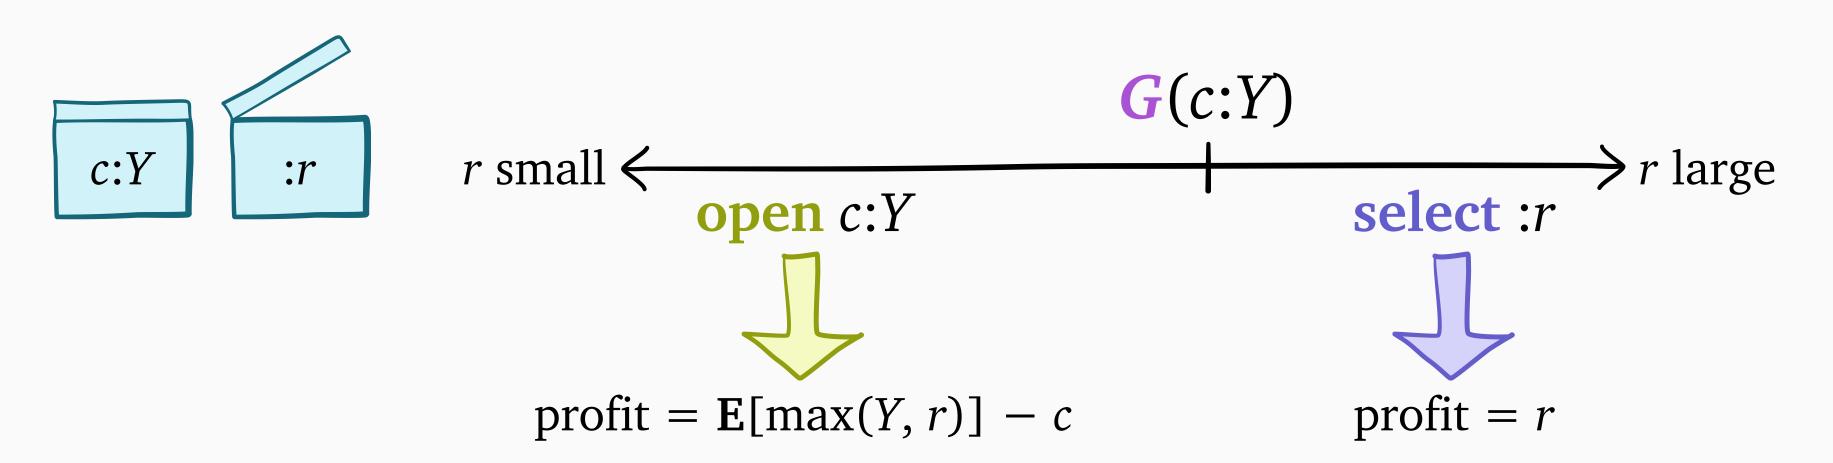




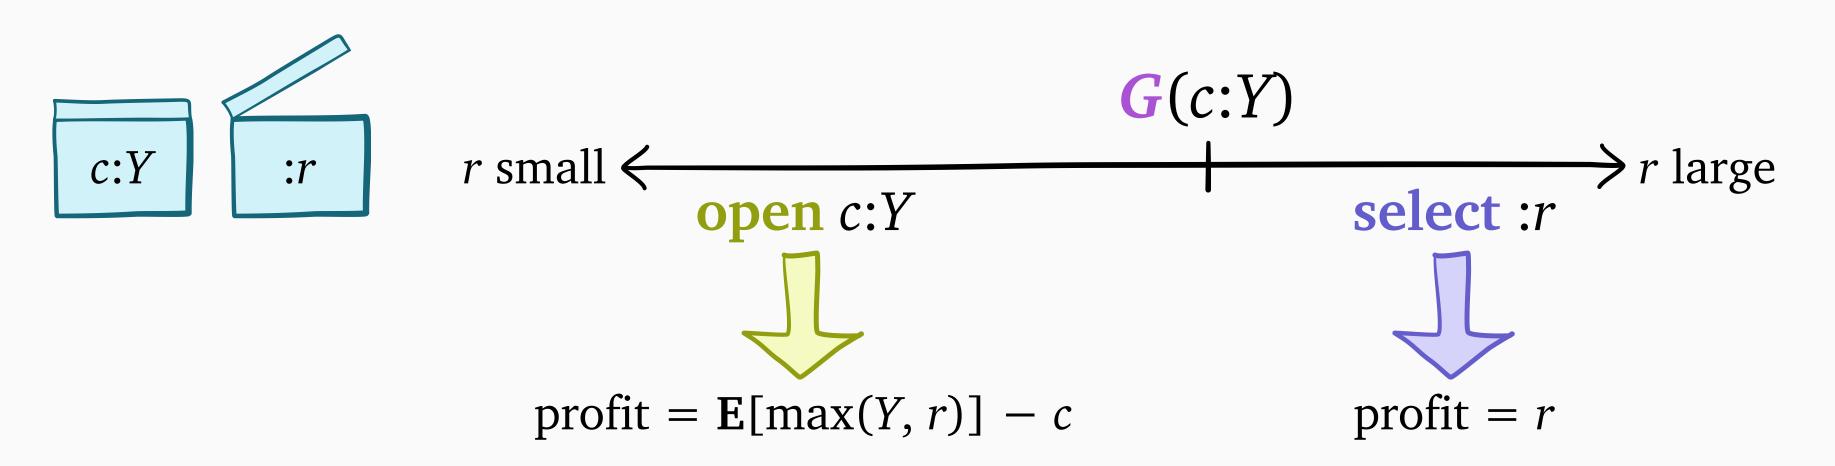




Indifference point r = G(c:Y)profit of open c:Y = profit of select :r

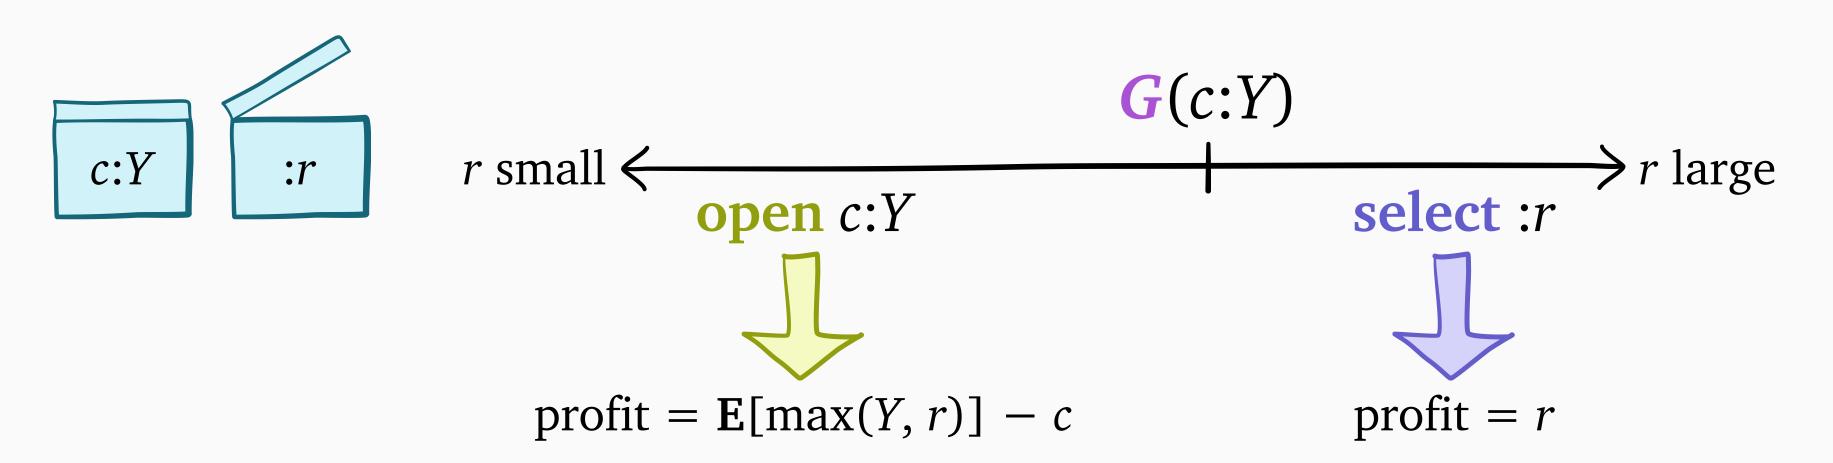


Indifference point r = G(c:Y) $E[\max(Y, r)] - c = \text{profit of select } : r$



Indifference point
$$r = G(c:Y)$$

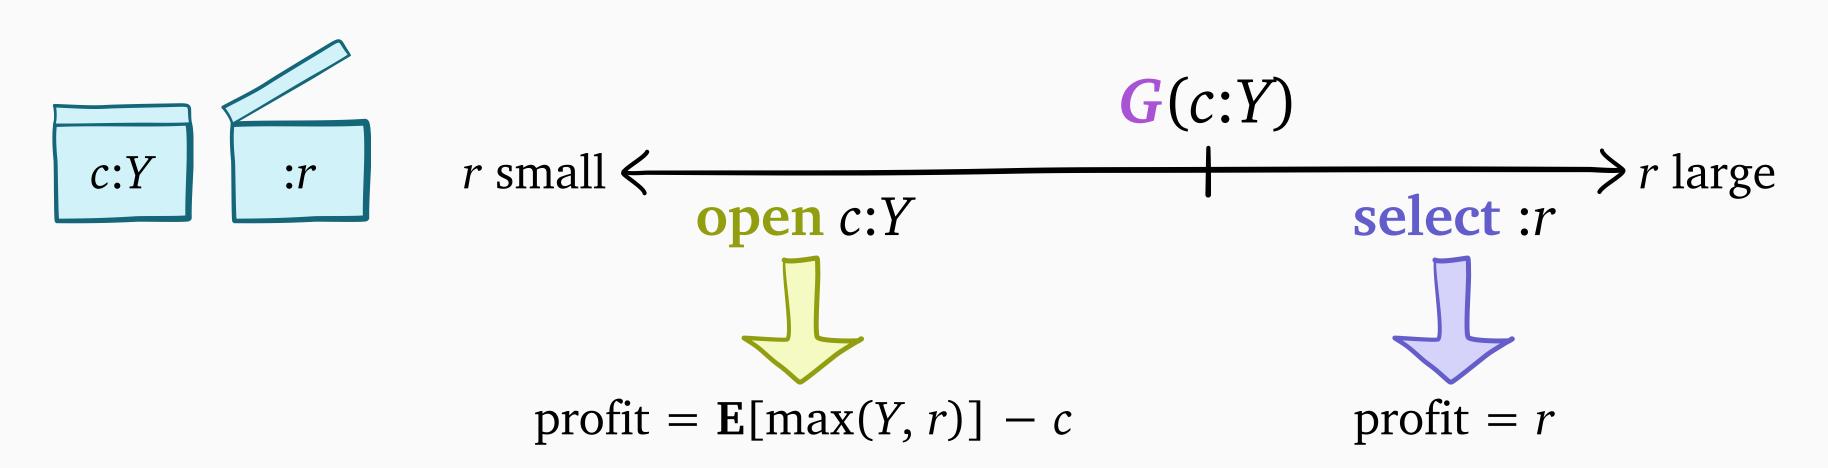
$$E[\max(Y, r)] - c = r$$



Indifference point
$$r = G(c:Y)$$

$$\mathbf{E}[\max(Y, r)] - c = r$$

$$\Rightarrow$$
 E[max($Y - r, 0$)] = c

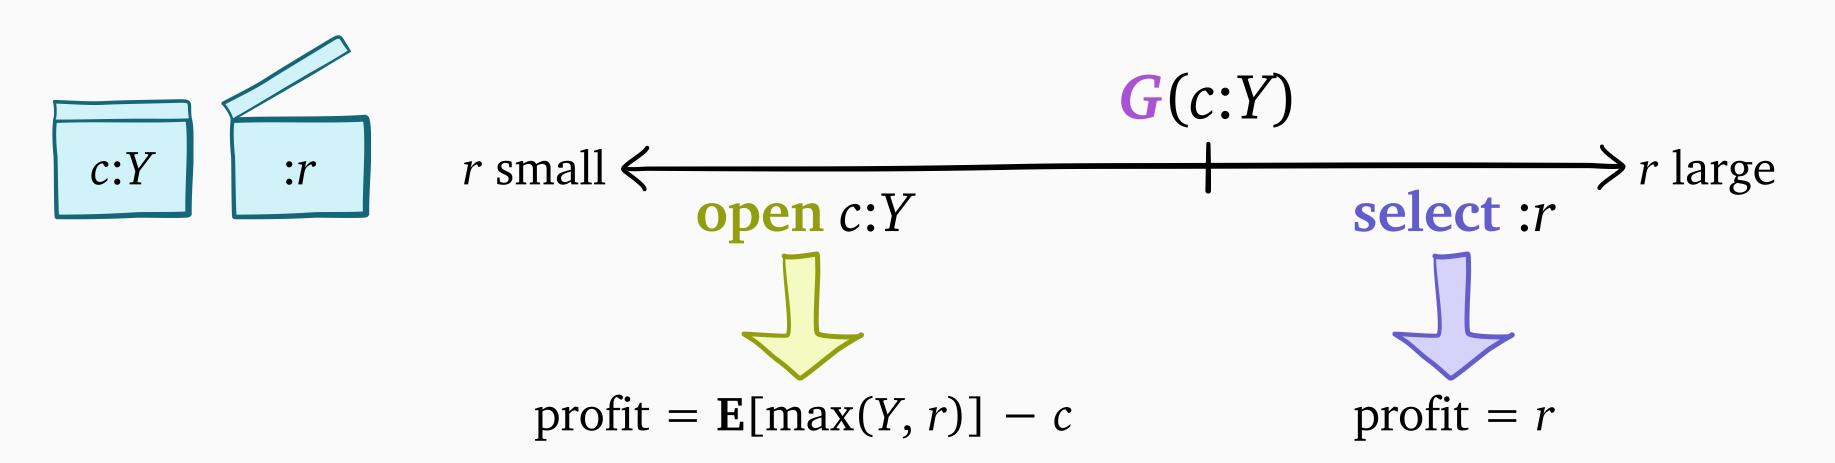


Indifference point
$$r = G(c:Y)$$

$$\mathbf{E}[\max(Y, r)] - c = r$$

$$\Rightarrow \mathbf{E}[\max(Y - r, 0)] = c$$

$$\underbrace{\mathbf{E}[(Y, r)]}$$



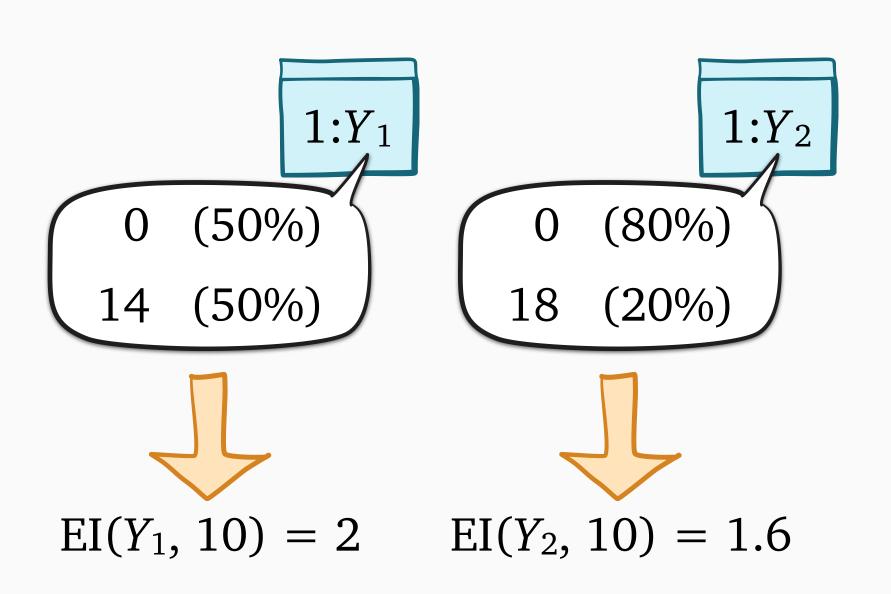
Indifference point
$$r = G(c:Y)$$

$$\mathbf{E}[\max(Y, r)] - c = r$$

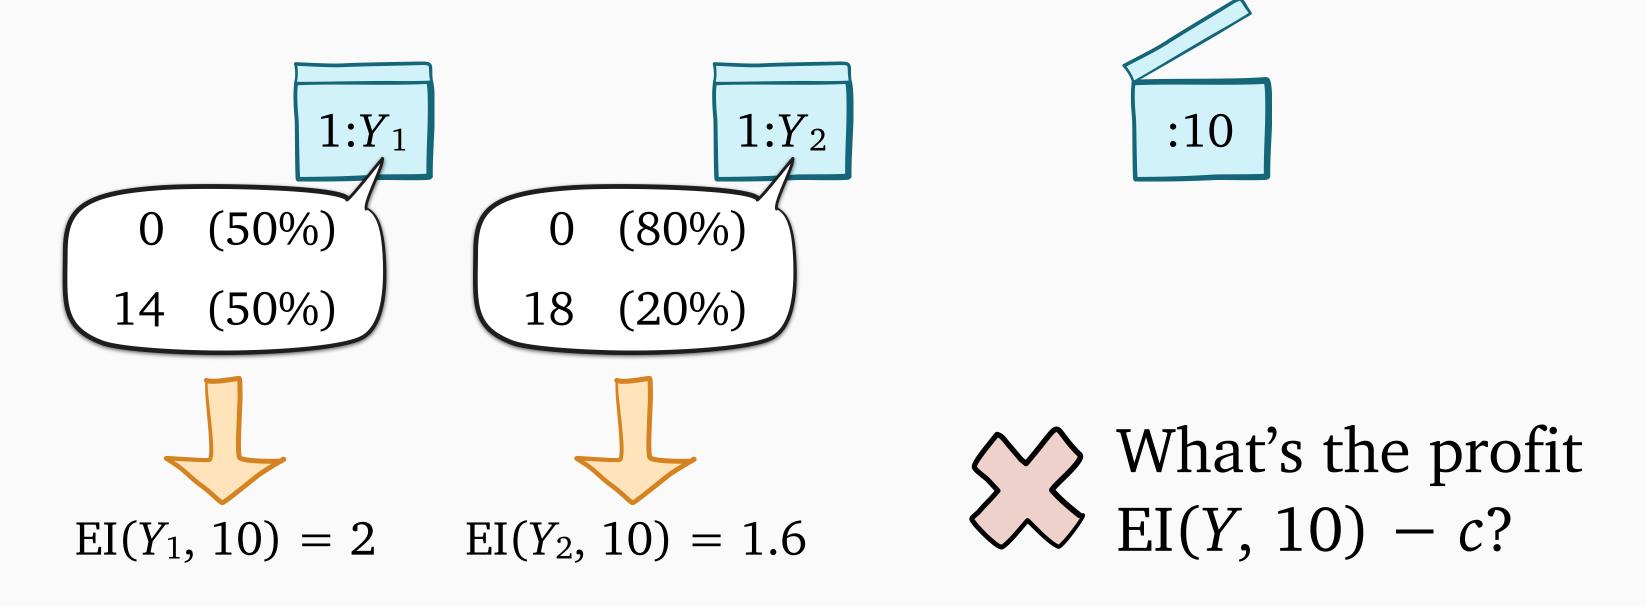
$$\Rightarrow \mathbf{E}[\max(Y - r, 0)] = c$$

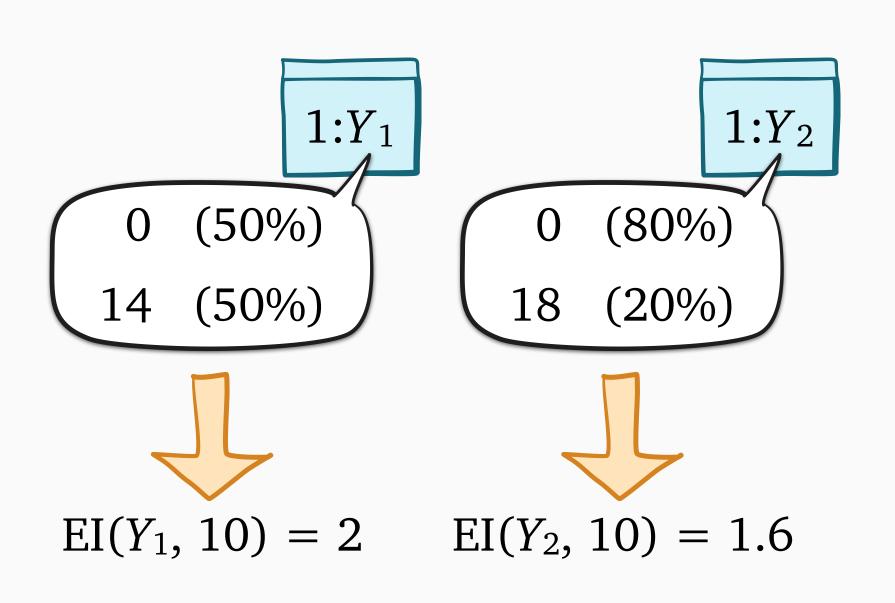
$$\underbrace{\mathsf{EI}(Y, r)}$$

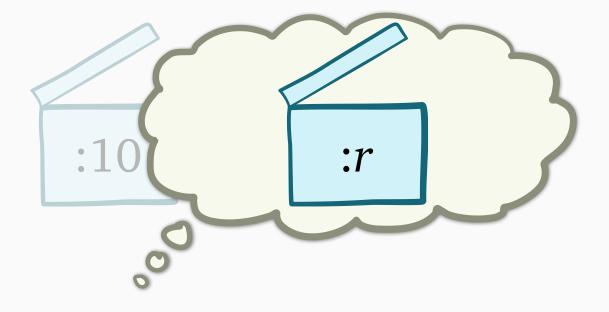
Definition:
$$G(c:Y)$$
 is solution r to $EI(Y, r) = c$







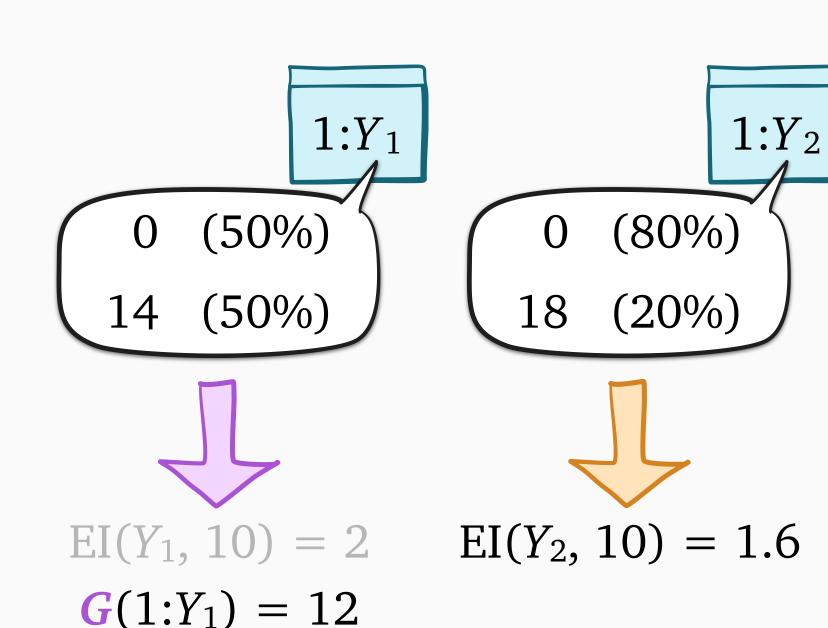


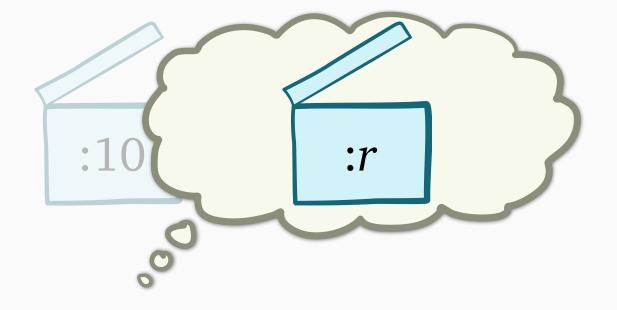


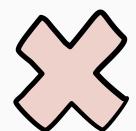


What's the profit EI(Y, 10) - c?



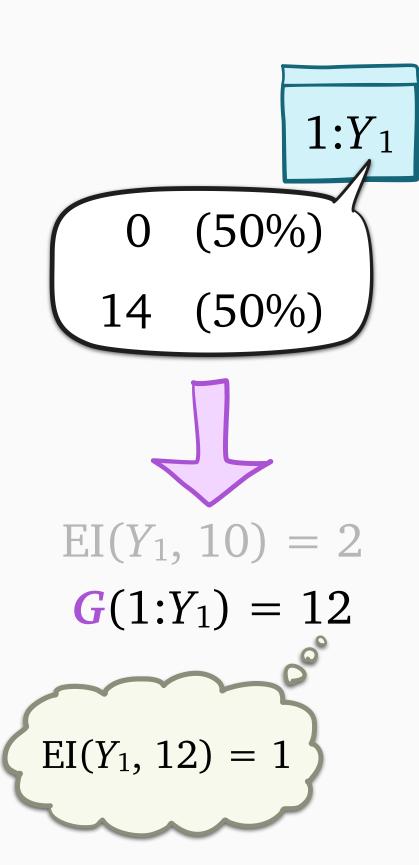


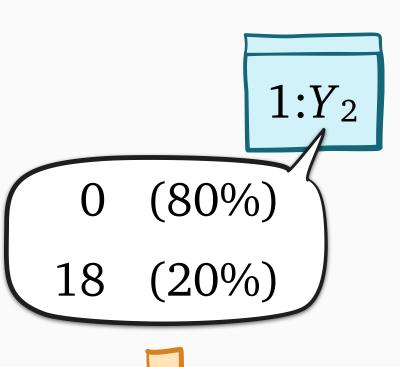


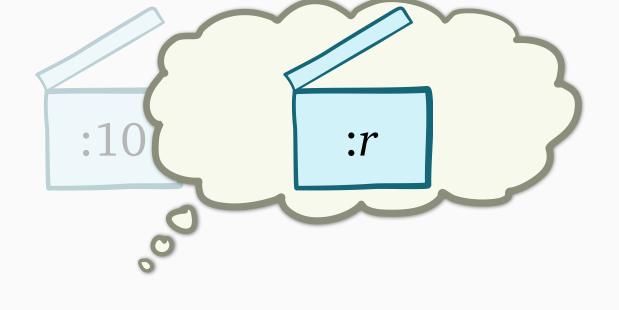


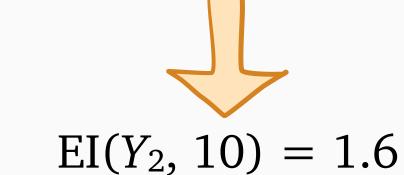
What's the profit EI(Y, 10) - c?

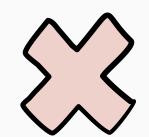






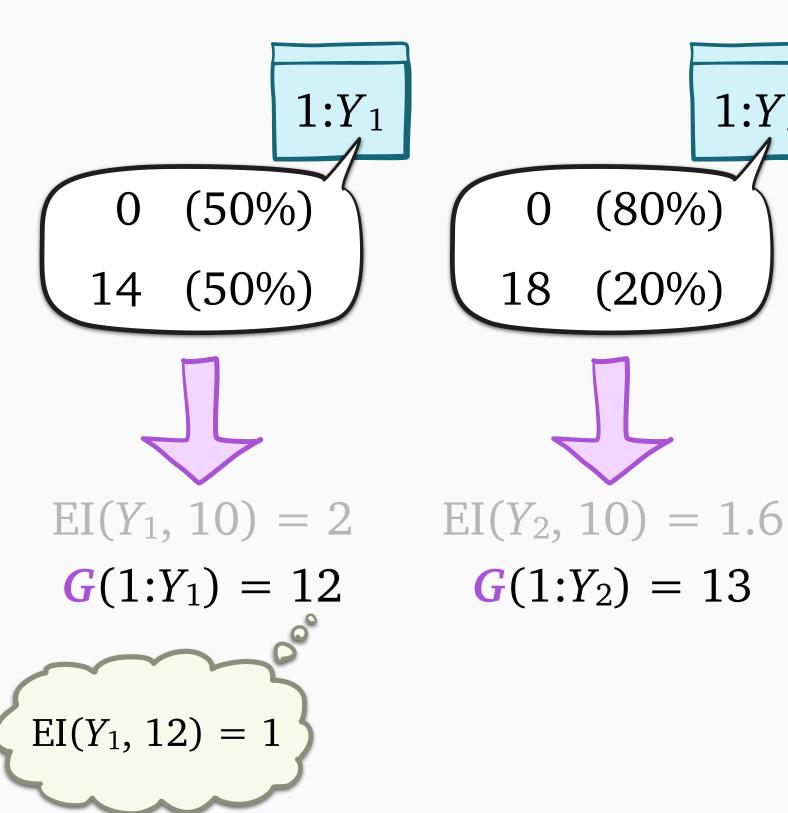


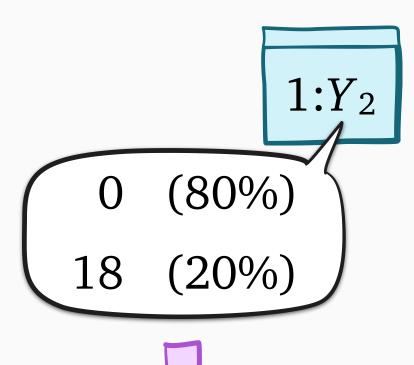


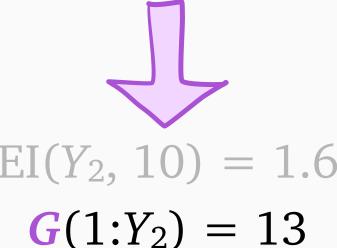


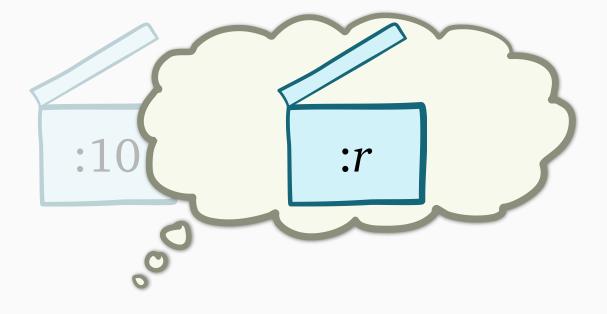
What's the profit EI(Y, 10) - c?

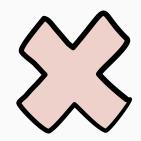






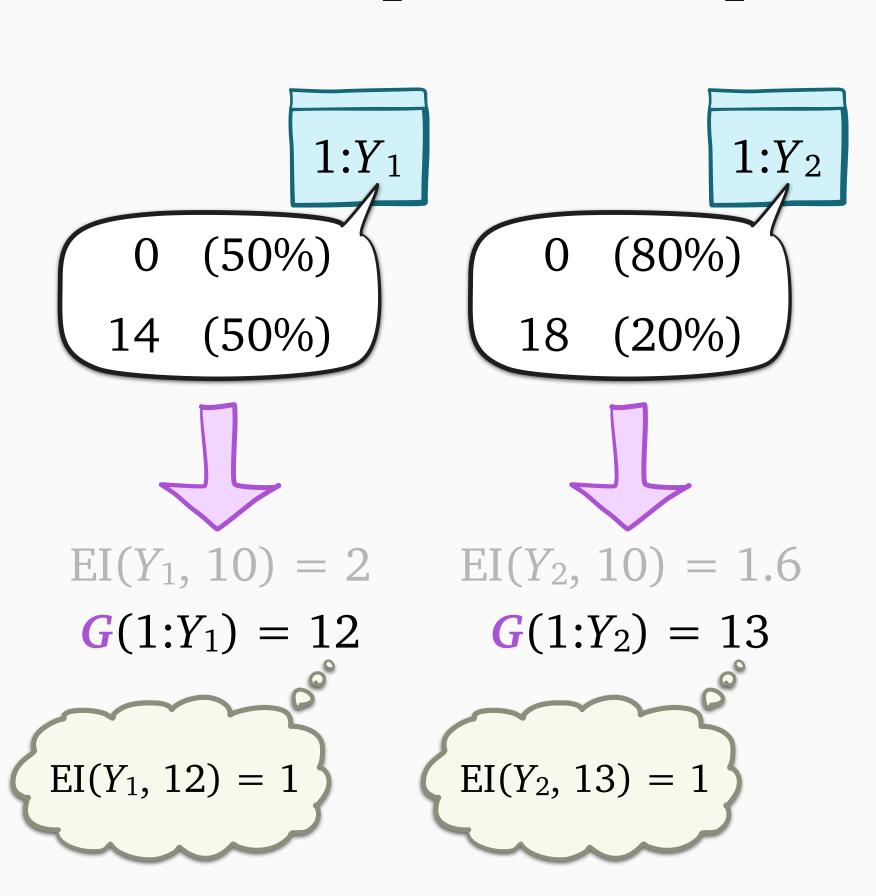


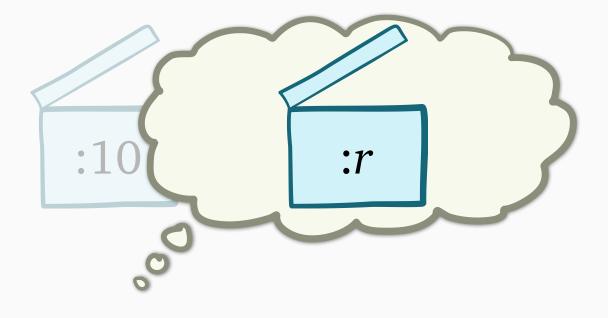


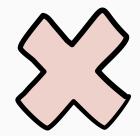


What's the profit $\rightarrow EI(Y, 10) - c?$









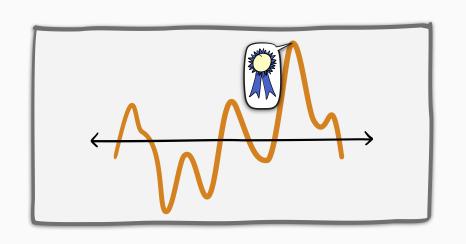
What's the profit EI(Y, 10) - c?





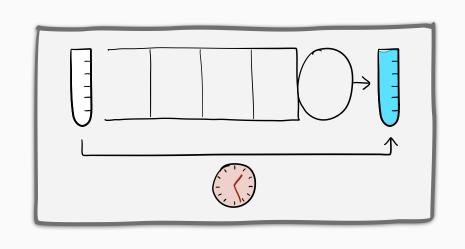
Part I:

Introducing Gittins via Pandora's box



Part II:

Gittins for Bayesian optimization

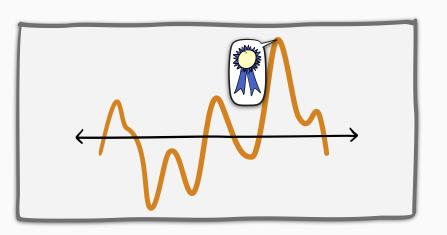


Part III:

Gittins for tail scheduling

Part II:

Gittins for Bayesian optimization



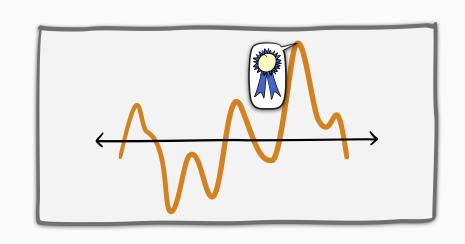
Click here for slides

https://presentations.avt.im/2025-10-27-Gittins-Index-Tutorial/



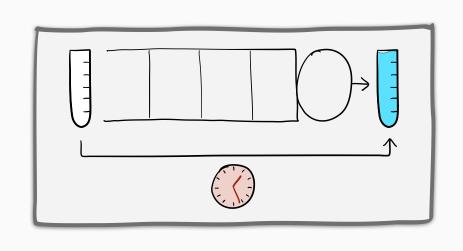
Part I:

Introducing Gittins via Pandora's box



Part II:

Gittins for Bayesian optimization



Part III:

Gittins for tail scheduling and more



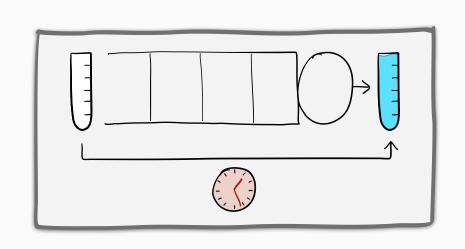
Part I:

Introducing Gittins via Pandora's box



Part II: NEW!

Gittins for Bayesian optimization

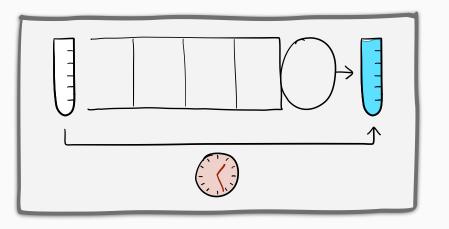


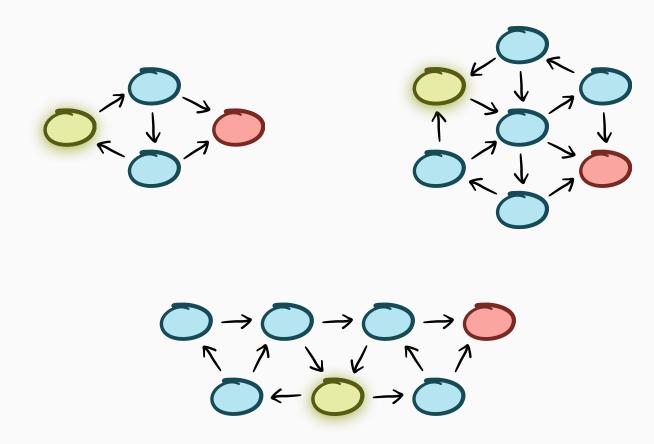
Part III:

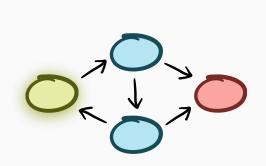
Gittins for tail scheduling and more

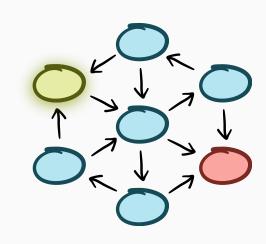
Part III:

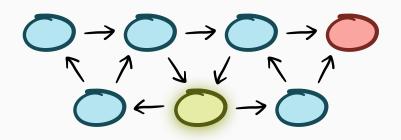
Gittins for tail scheduling



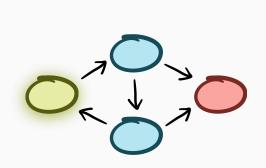


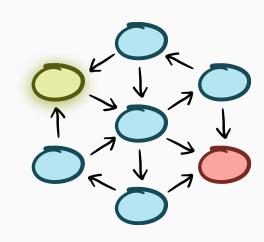


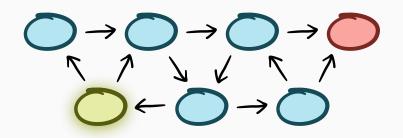




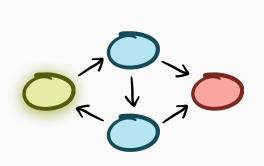
- Advance one Markov chain every step
- Get reward based on its current state
- Stop when one Markov chain terminates

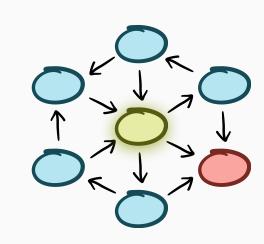


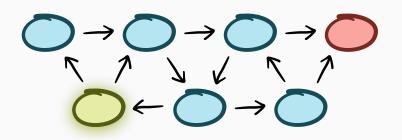




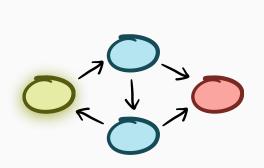
- Advance one Markov chain every step
- Get reward based on its current state
- Stop when one Markov chain terminates

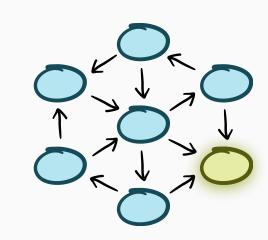


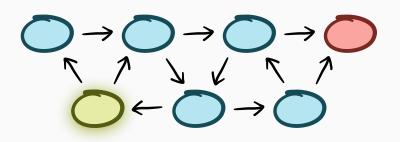




- Advance one Markov chain every step
- Get reward based on its current state
- Stop when one Markov chain terminates

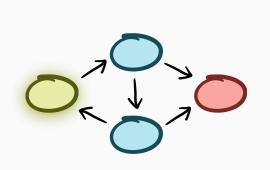


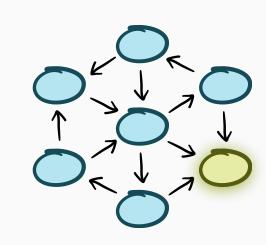


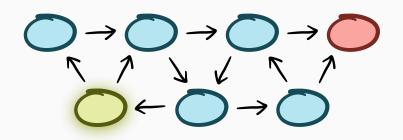


- Advance one Markov chain every step
- Get reward based on its current state
- Stop when one Markov chain terminates

Markov chain selection (MCS)





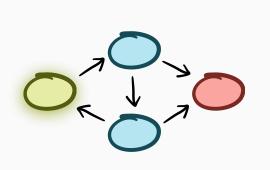


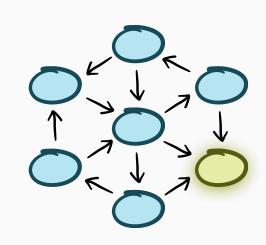
Decision process:

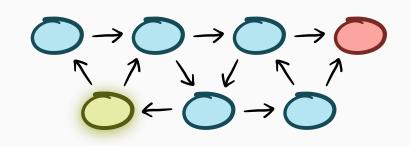
- Advance one Markov chain every step
- Get reward based on its current state
- Stop when one Markov chain terminates

Goal: maximize **E**[total reward]

Markov chain selection (MCS)







Decision process:

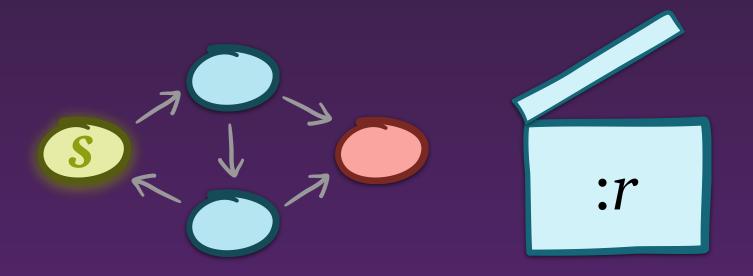
- Advance one Markov chain every step
- Get reward based on its current state
- Stop when one Markov chain terminates

Goal: maximize **E**[total reward]



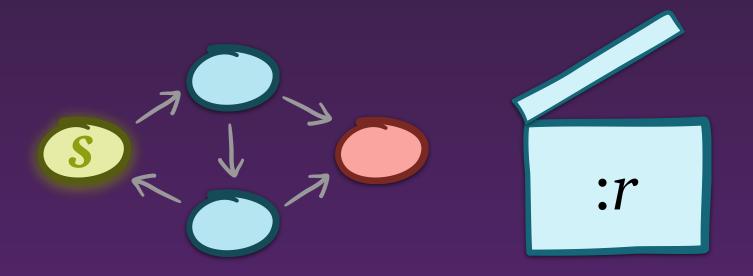
Optimal policy: advance Markov chain whose state has greatest Gittins index

The *Gittins index* of state *s* is the value of *r* such that we're indifferent between "act on *s*" and "get *r* now"



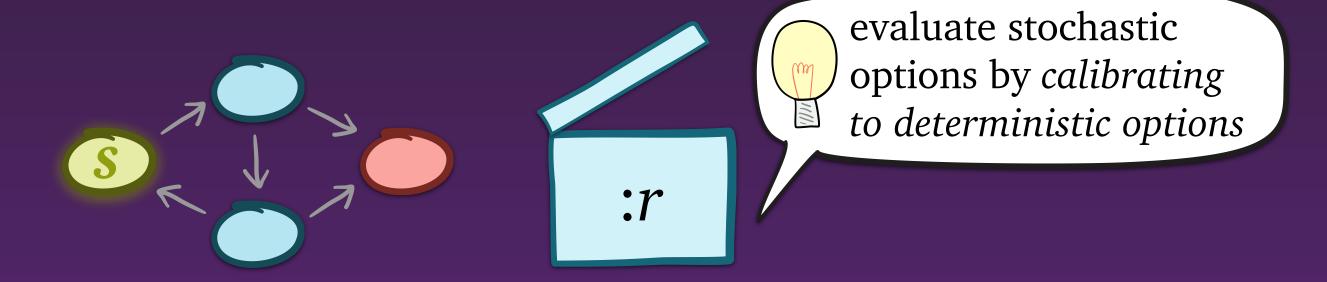


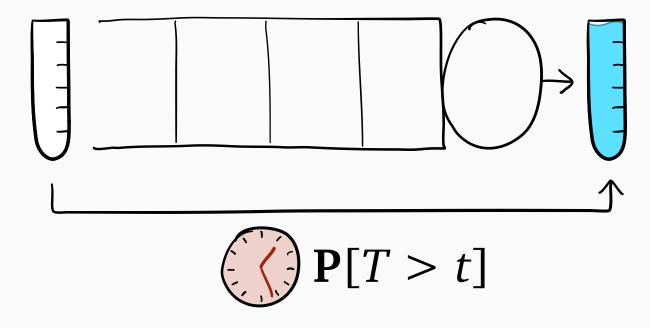
The *Gittins index* of state *s* is the value of *r* such that we're indifferent between "act on *s*" and "get *r* now"



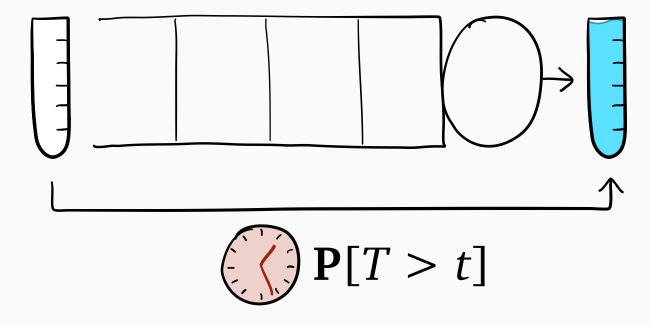


The *Gittins index* of state *s* is the value of *r* such that we're indifferent between "act on *s*" and "get *r* now"



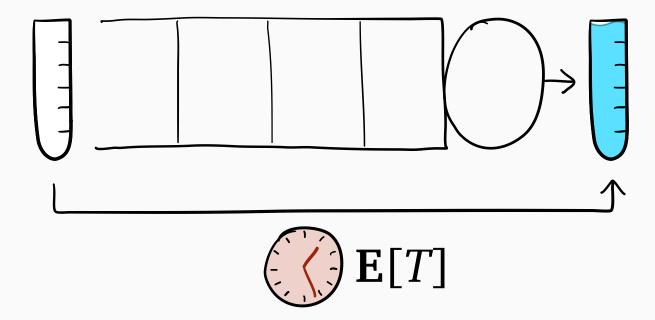


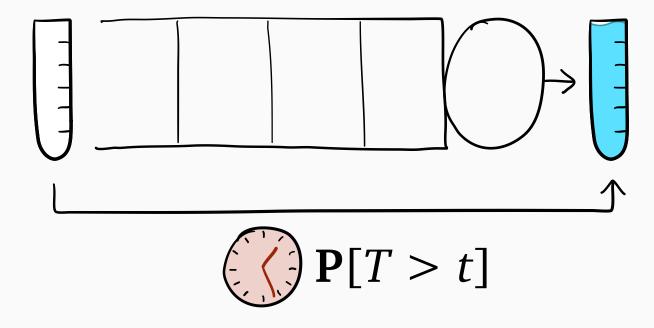
Goal: minimize probability of very long response time



Goal: minimize probability of very long response time

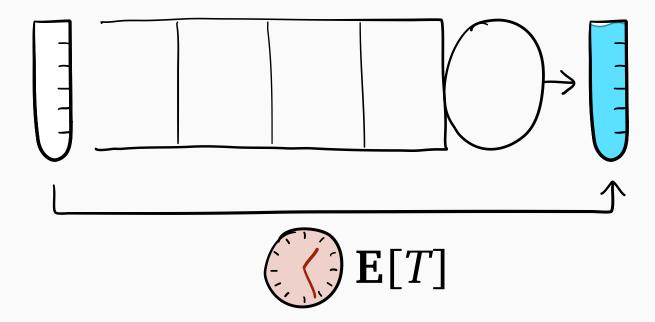
Mean scheduling



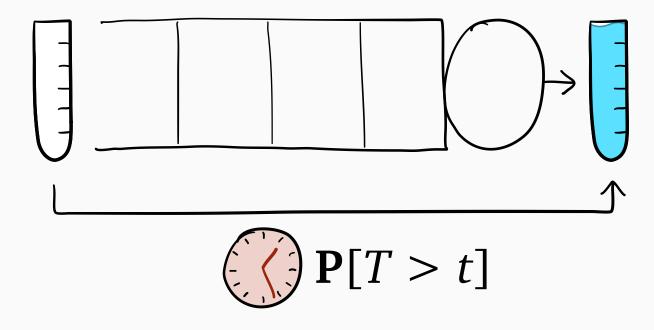


Goal: minimize probability of very long response time

Mean scheduling



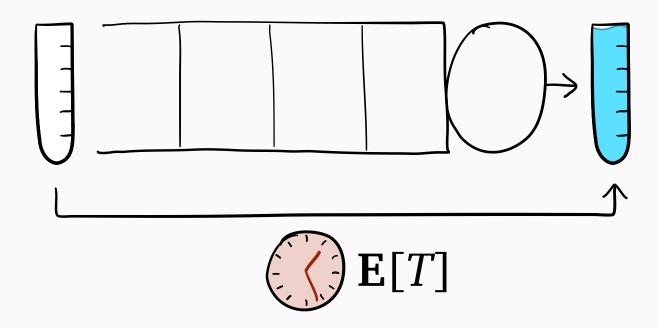
Goal: minimize average response time (or # jobs)



Goal: minimize probability of very long response time

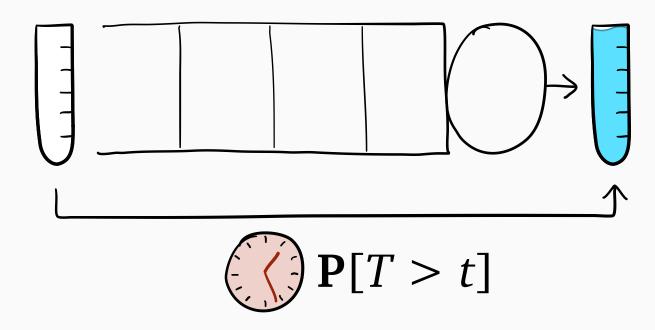


Mean scheduling



Goal: minimize average response time (or # jobs)

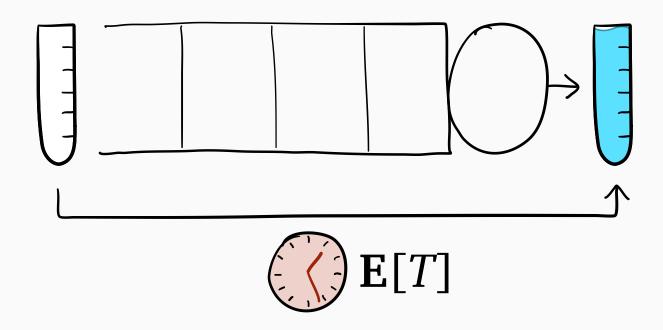




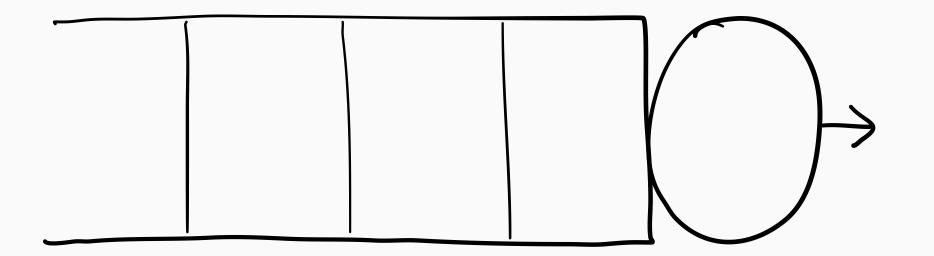
Goal: minimize probability of very long response time

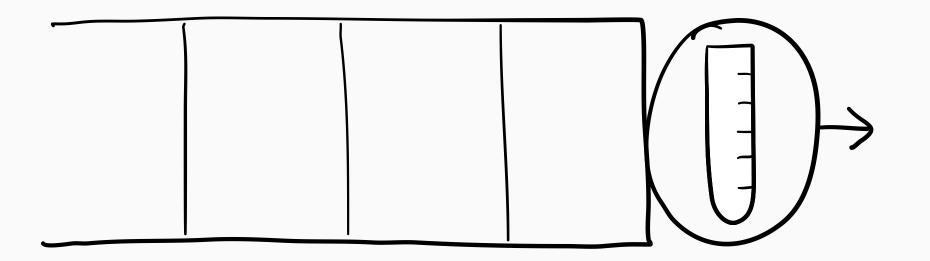


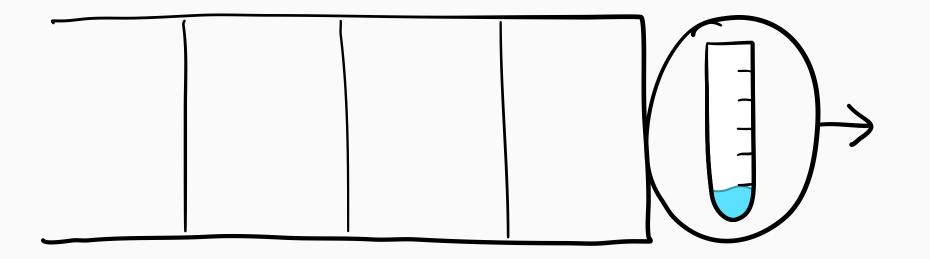
Mean scheduling

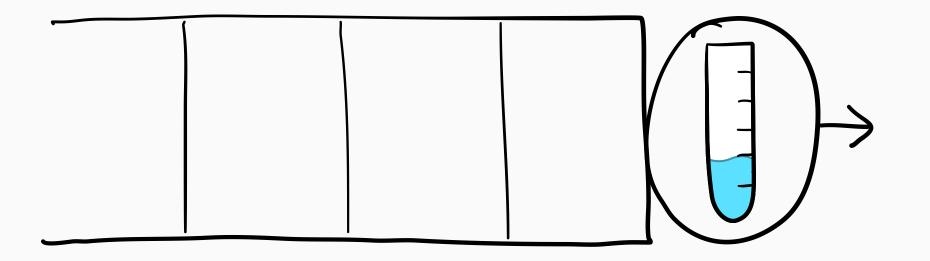


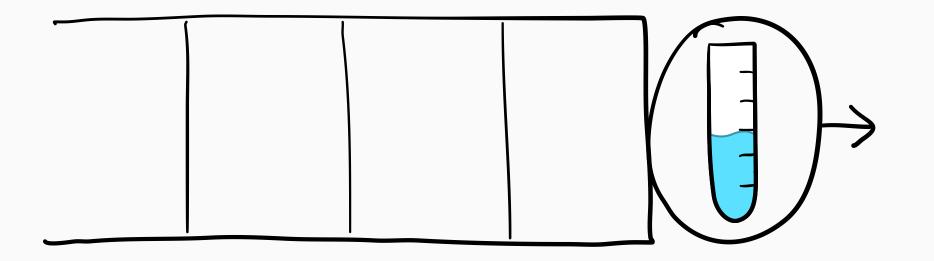
Goal: minimize average response time (or # jobs)

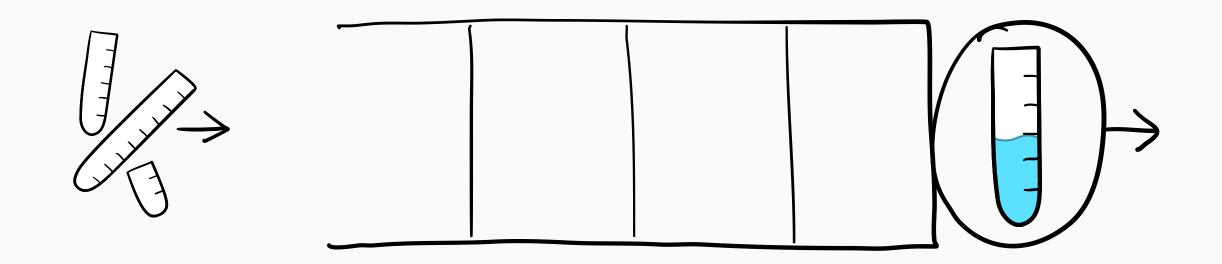


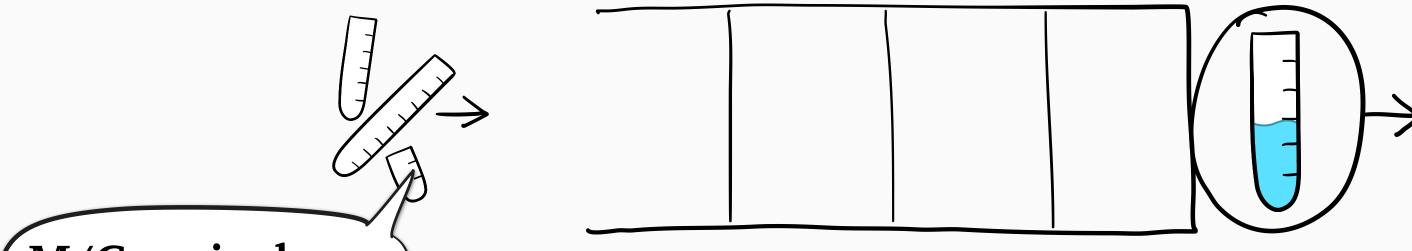




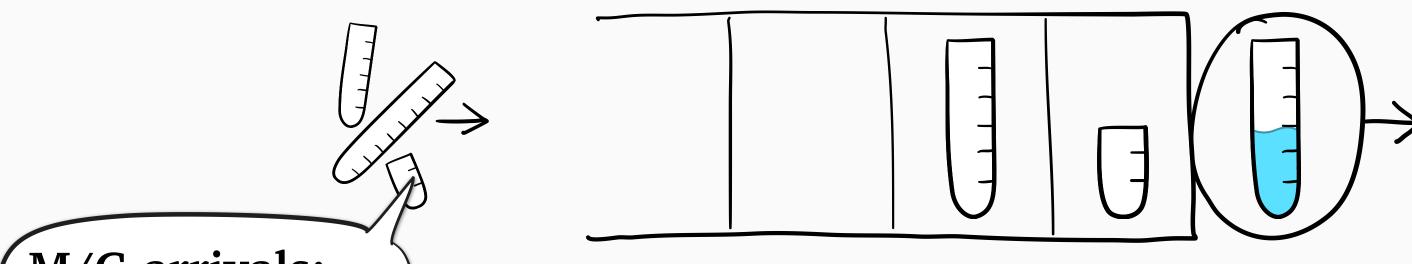




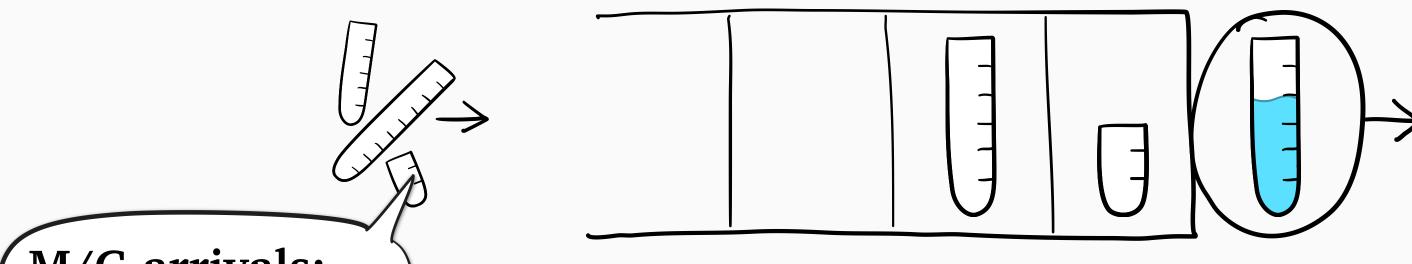




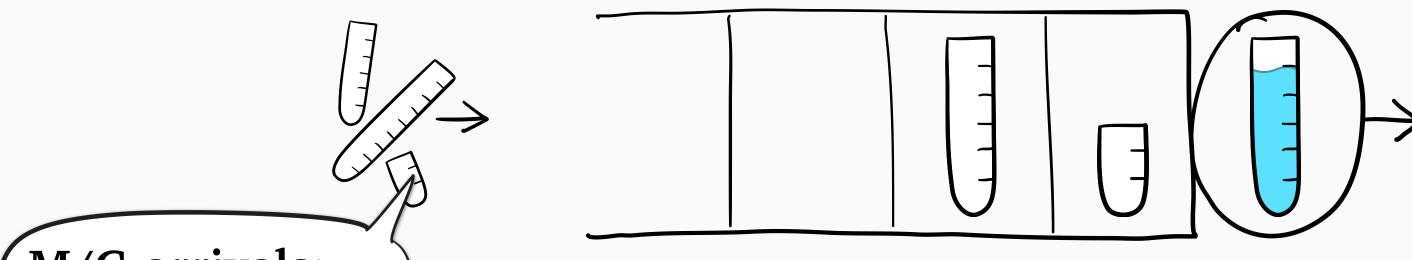
- arrival rate λ
- size distribution Z



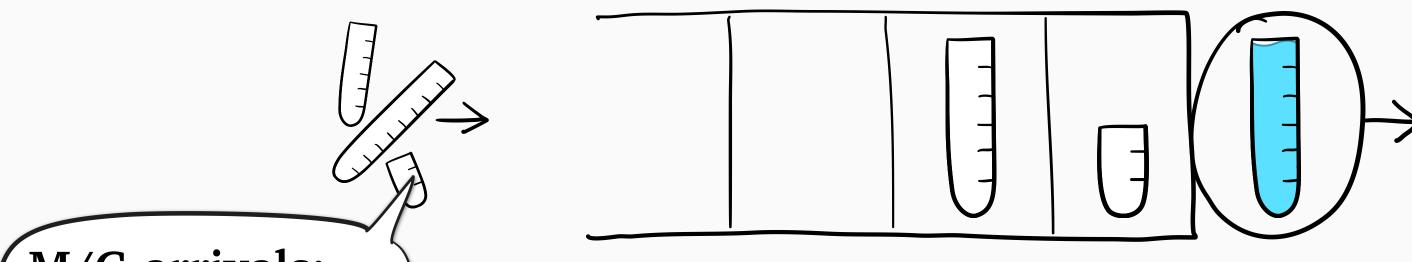
- arrival rate λ
- size distribution Z



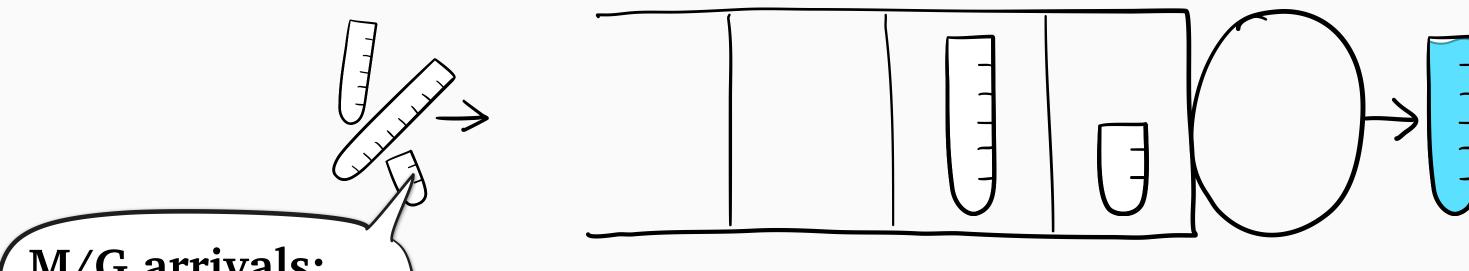
- arrival rate λ
- size distribution Z



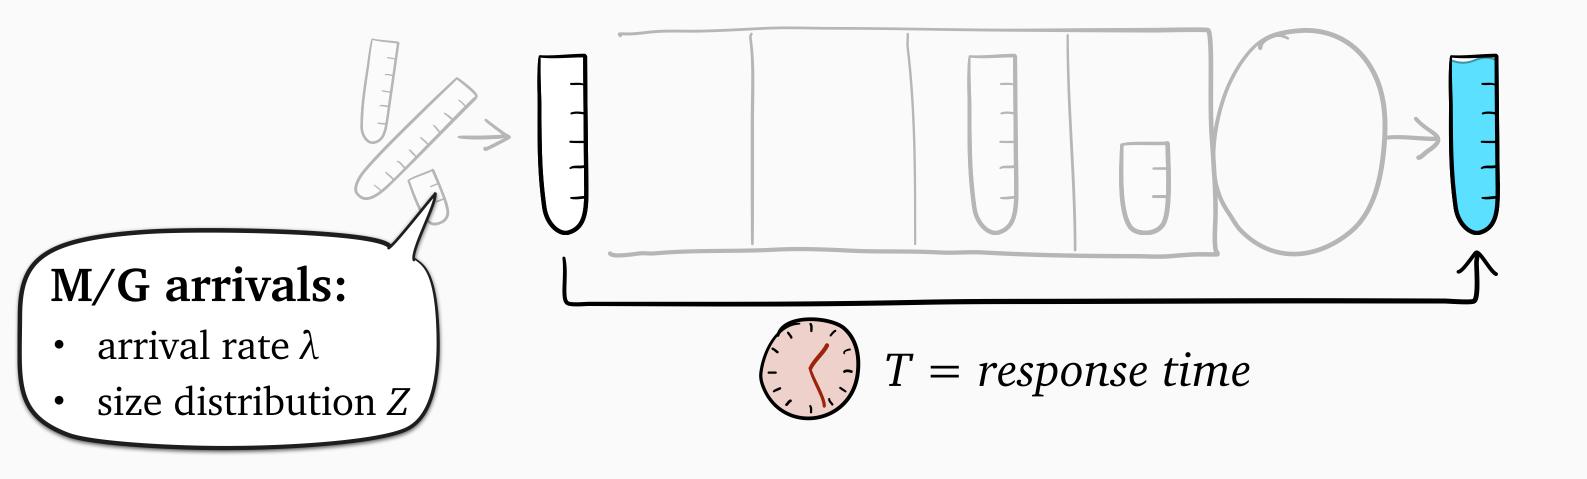
- arrival rate λ
- size distribution Z_{λ}

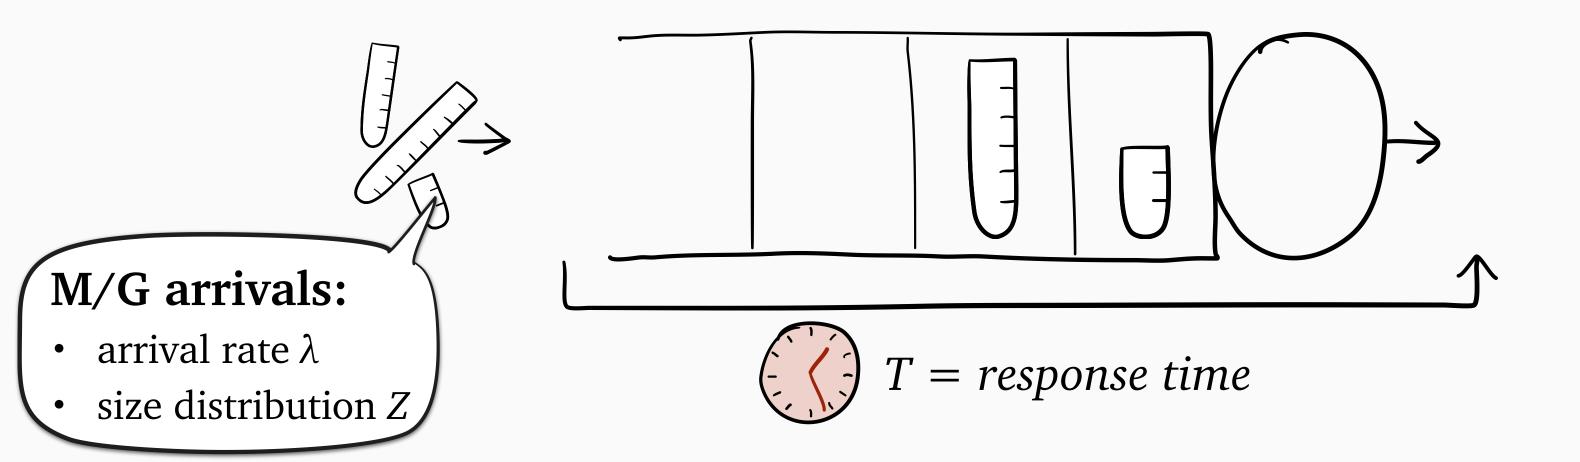


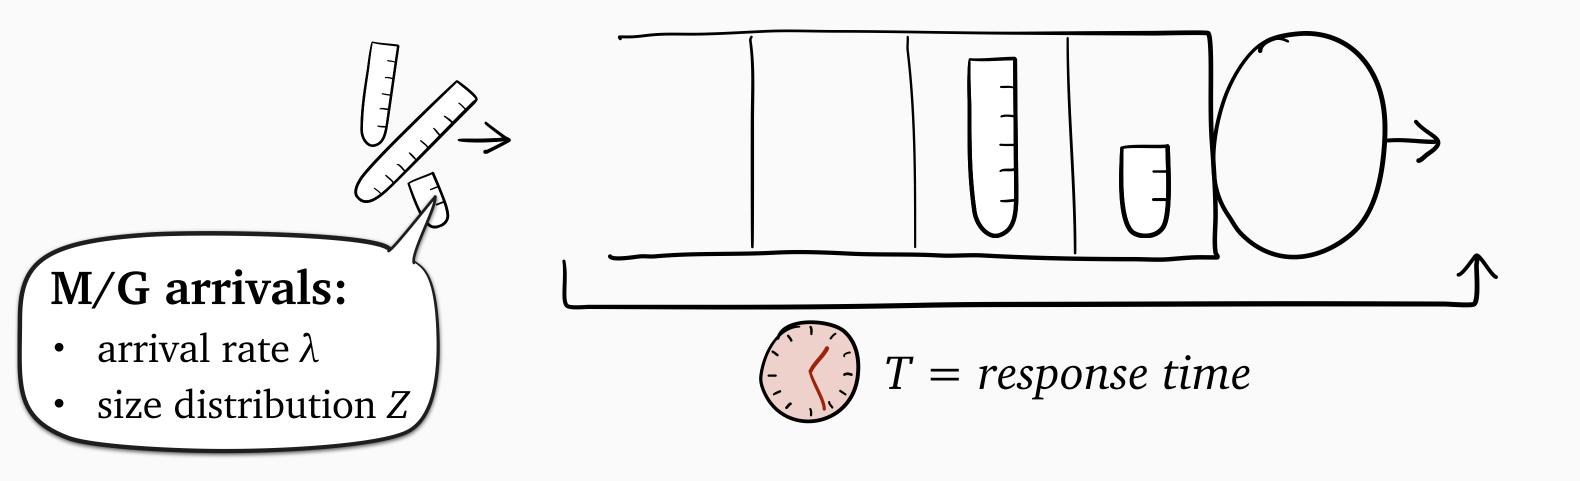
- arrival rate λ
- size distribution Z_{λ}



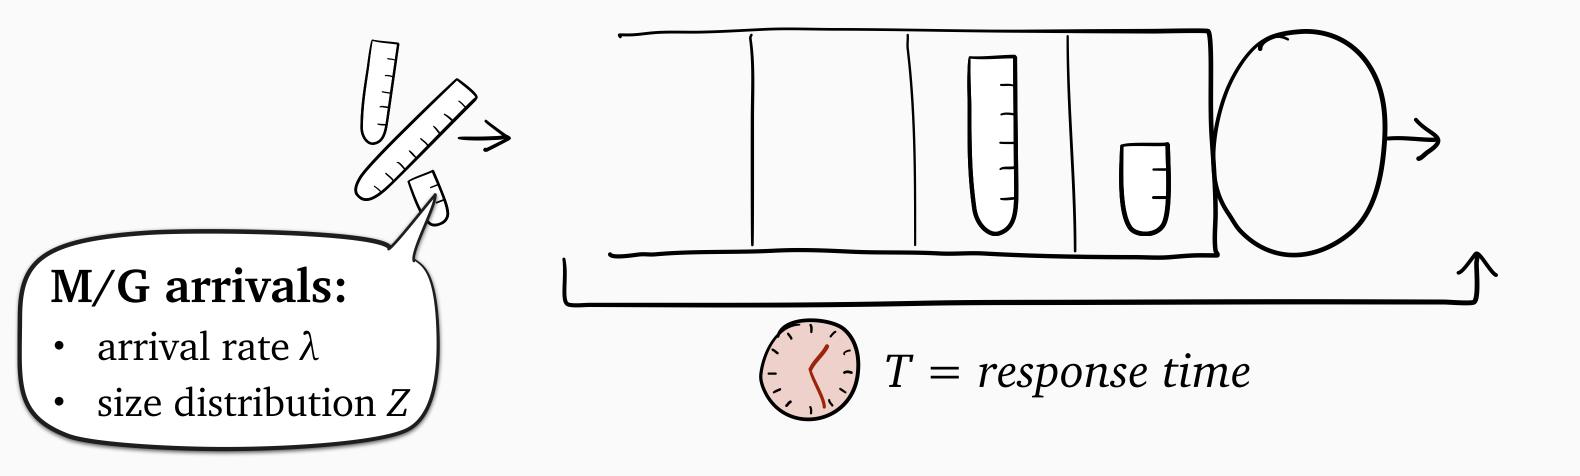
- arrival rate λ
- size distribution Z_{λ}





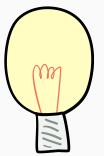




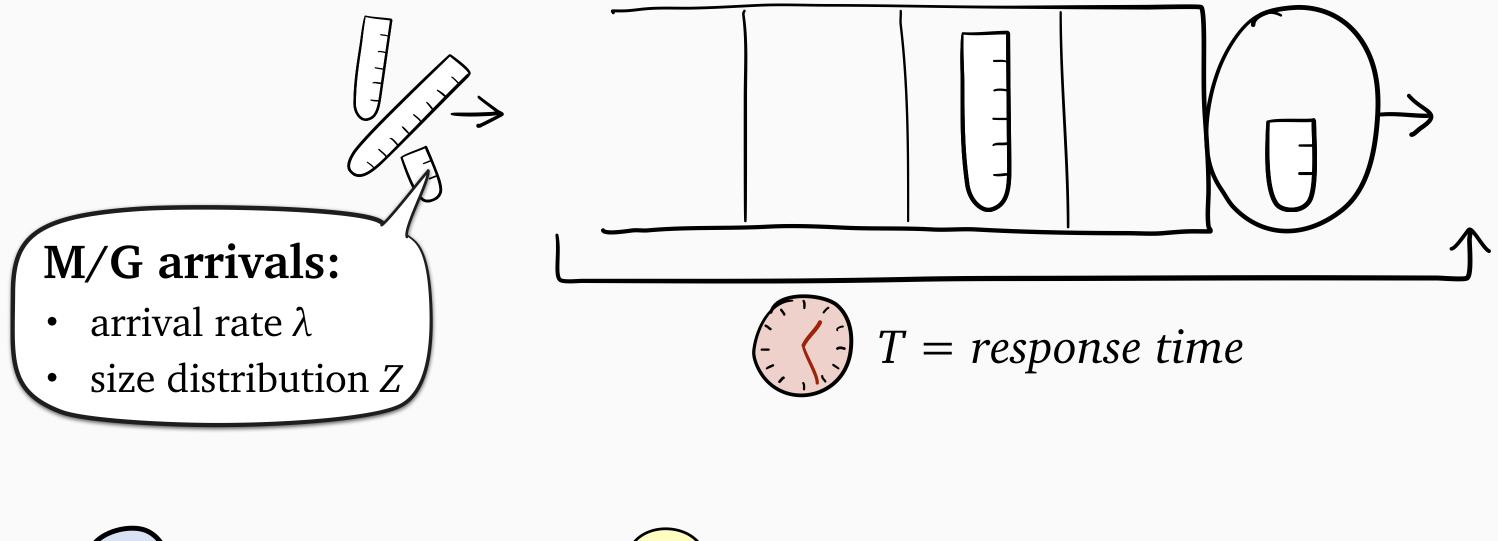




Minimize $\mathbf{E}[T]$?

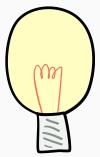


Serve short jobs before long jobs



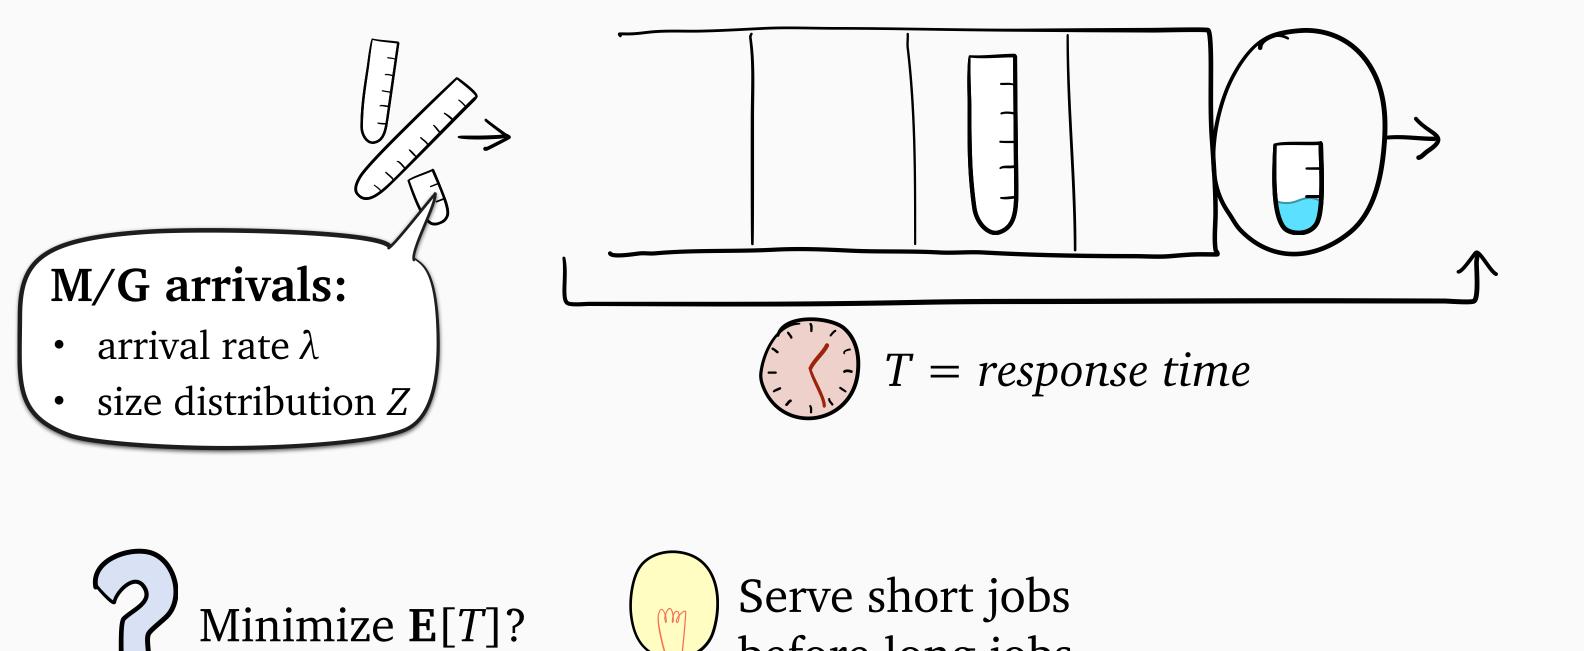


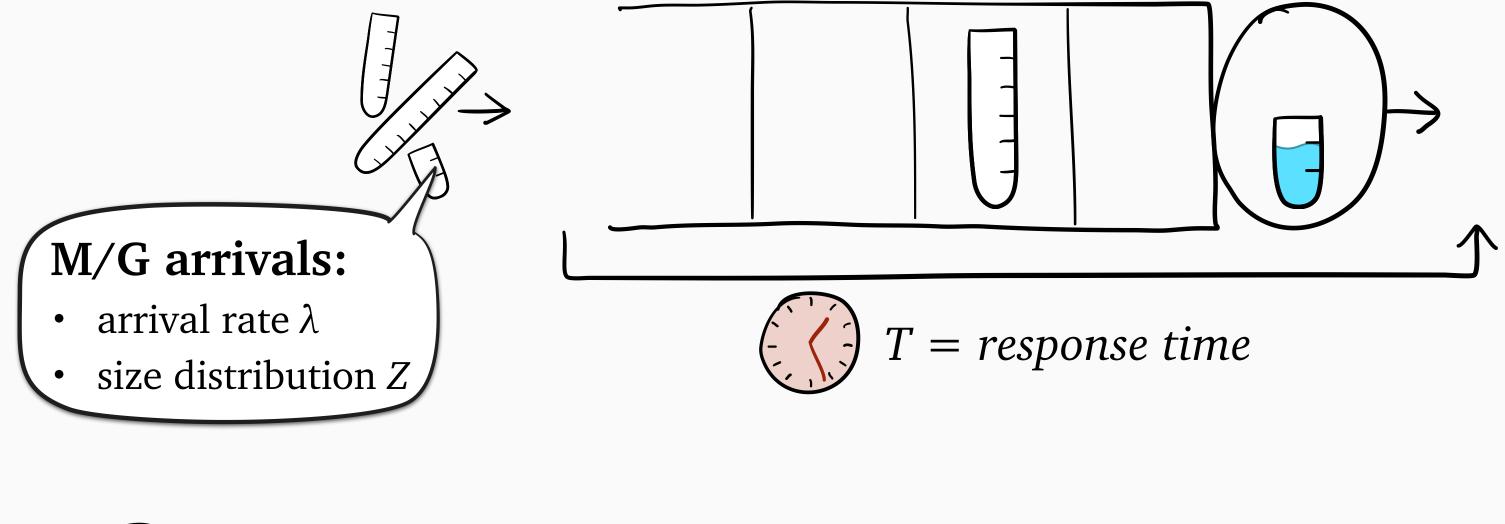
Minimize E[T]?



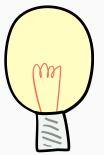
Serve short jobs before long jobs

before long jobs



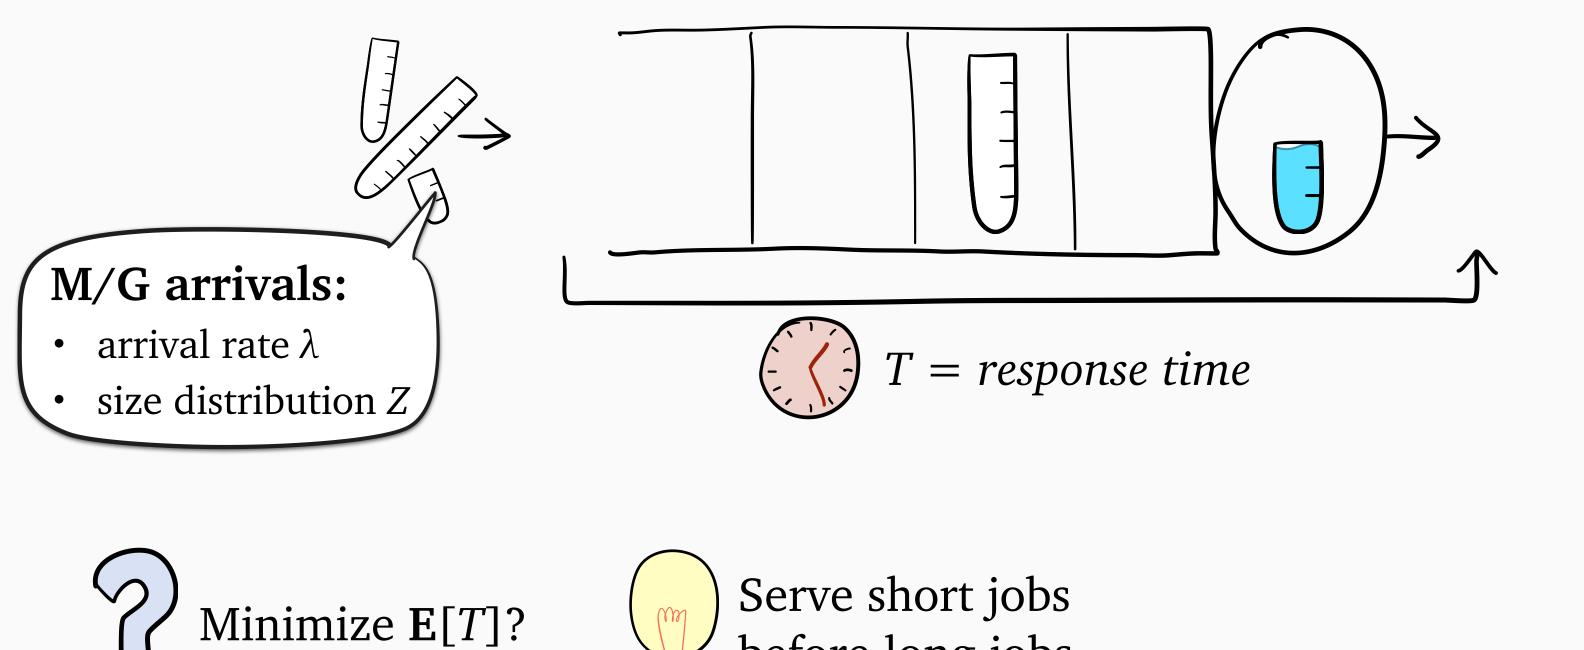


Minimize E[T]?

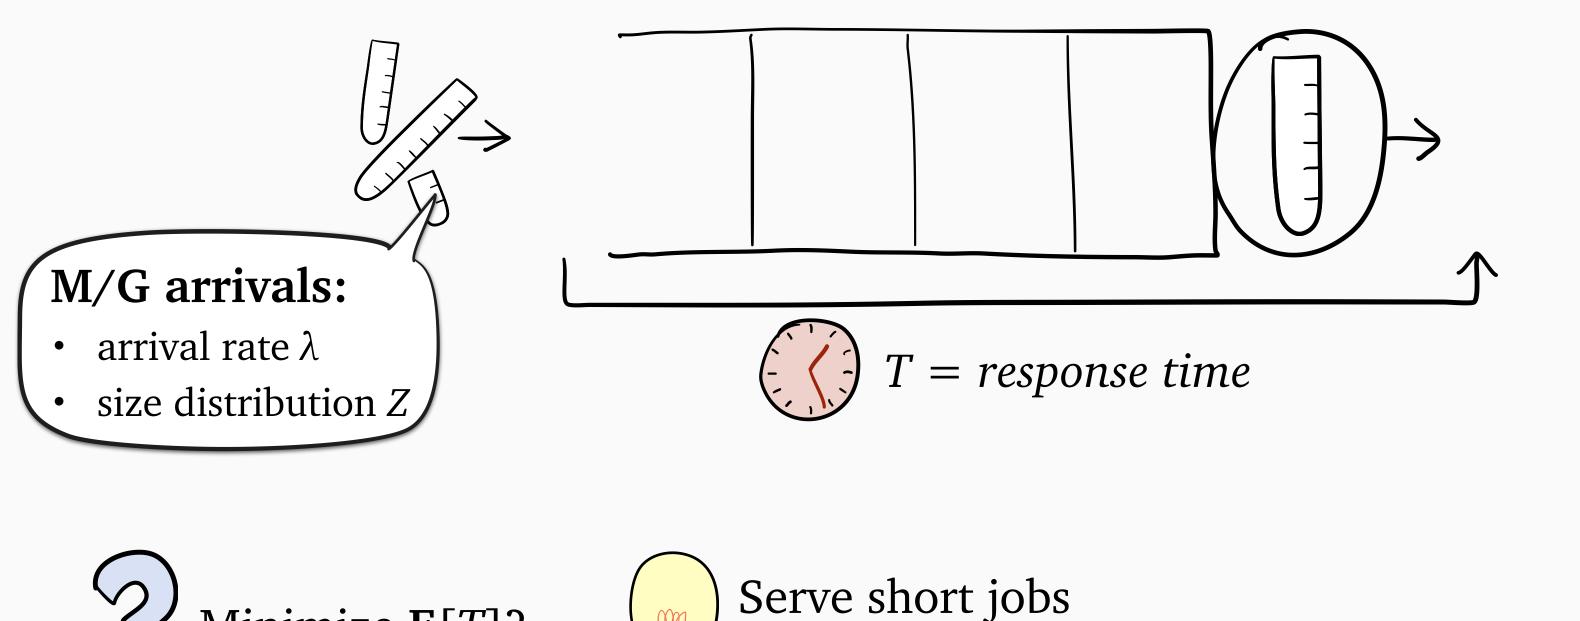


Serve short jobs before long jobs

before long jobs

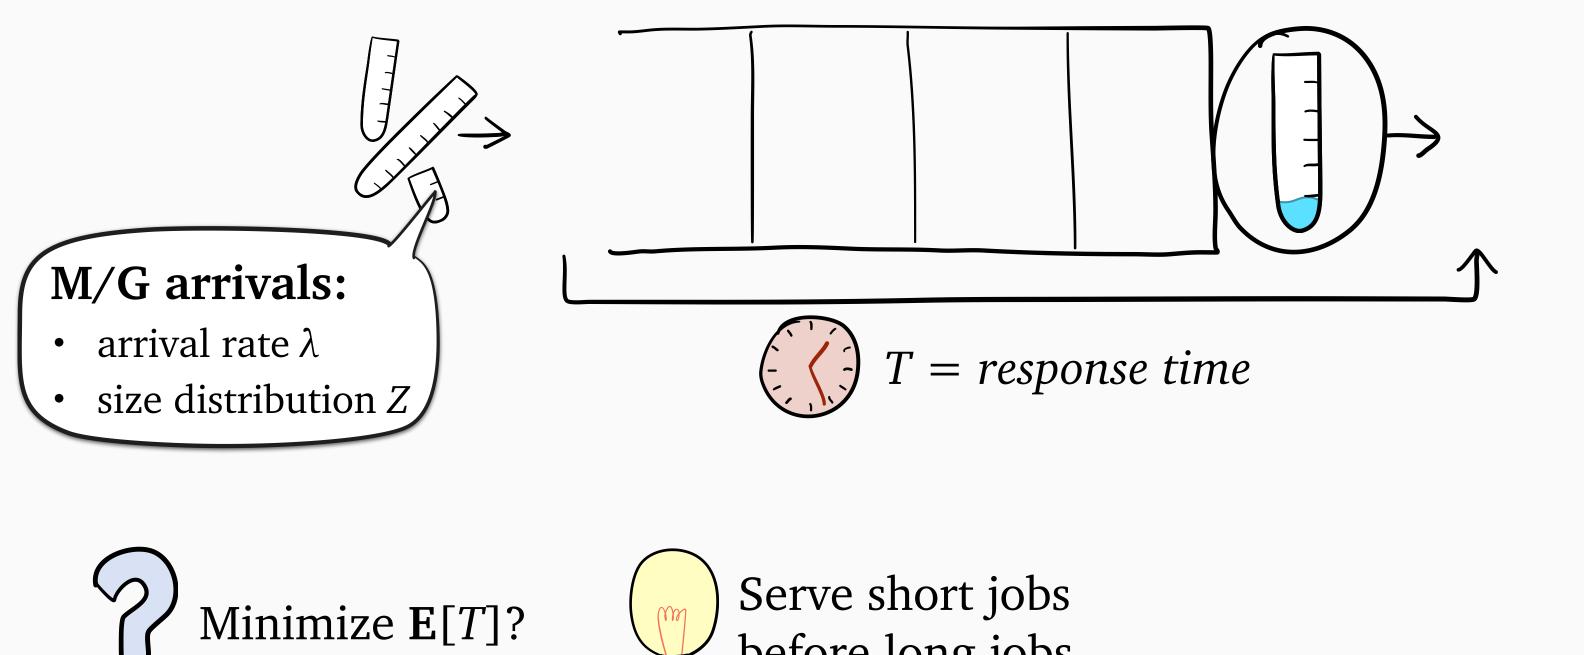


before long jobs

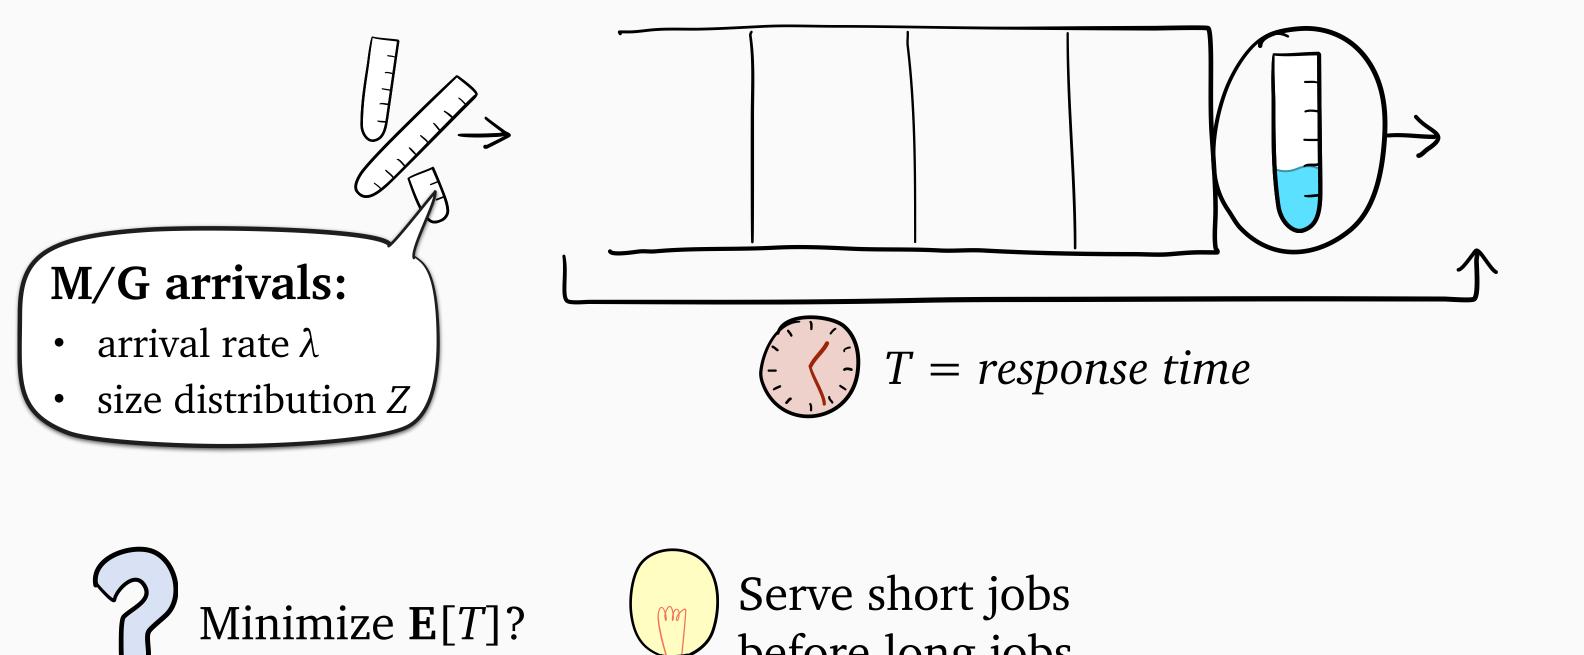


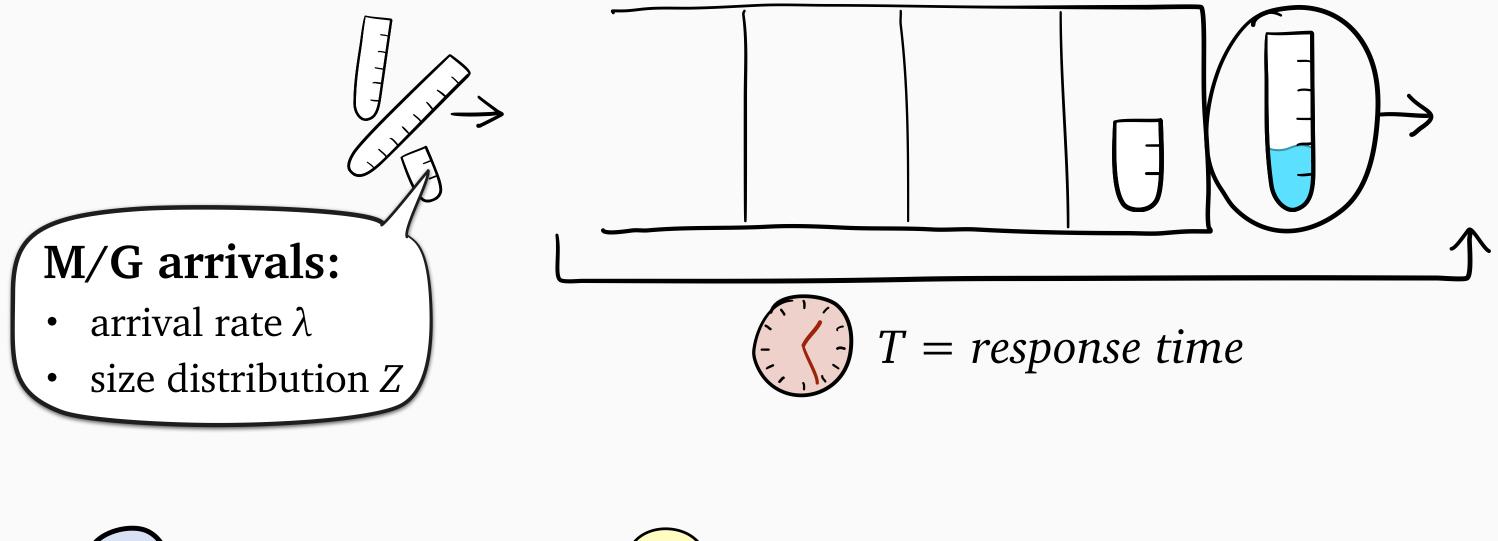
Minimize E[T]?

before long jobs



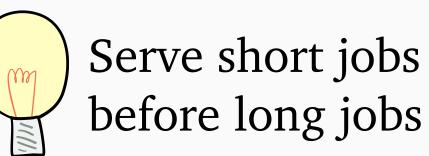
before long jobs



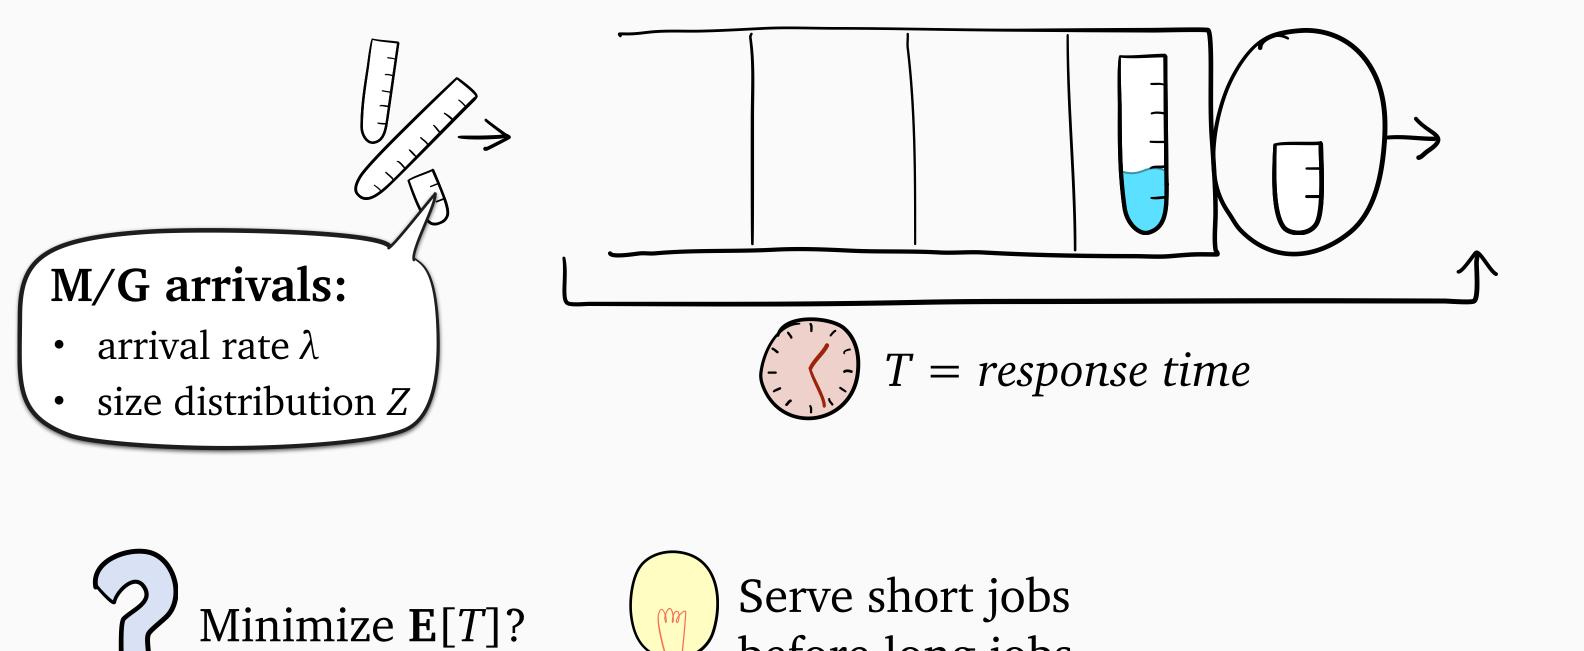




Minimize E[T]?

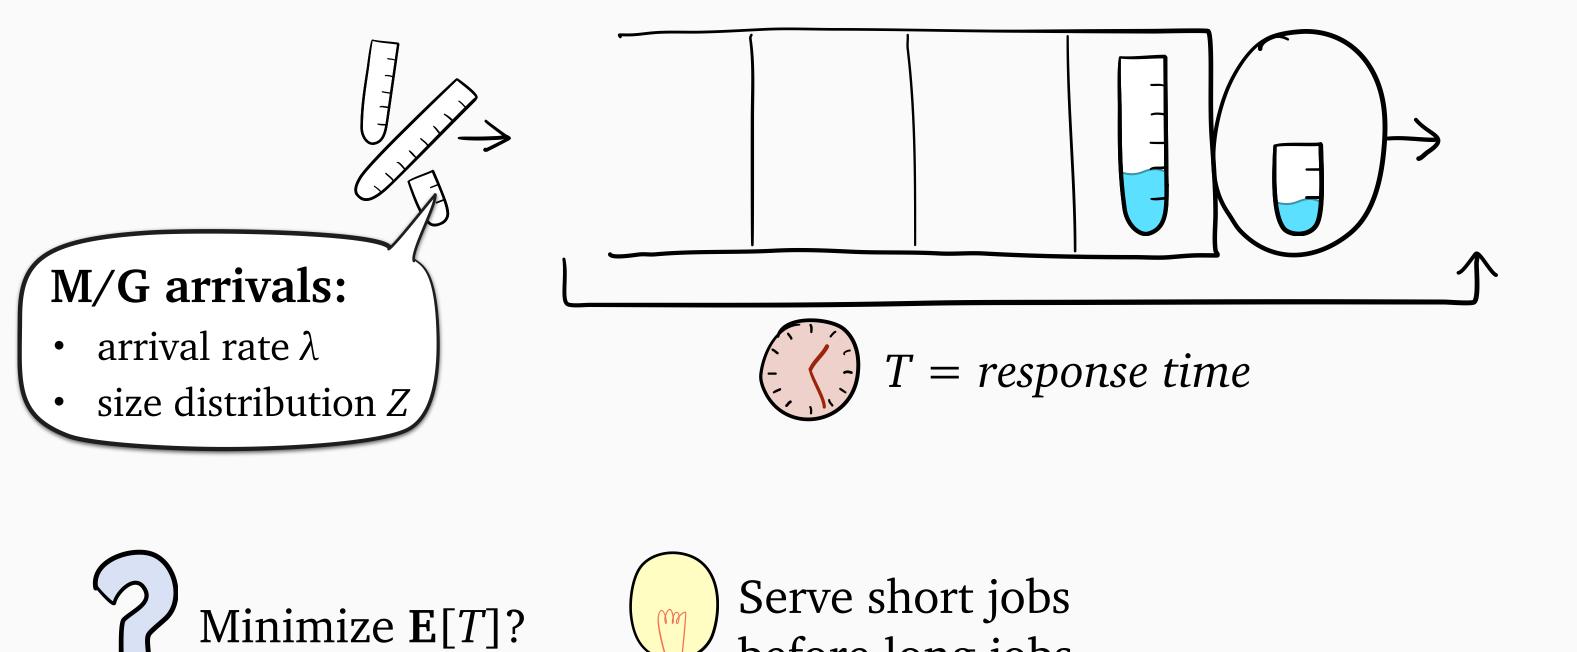


before long jobs

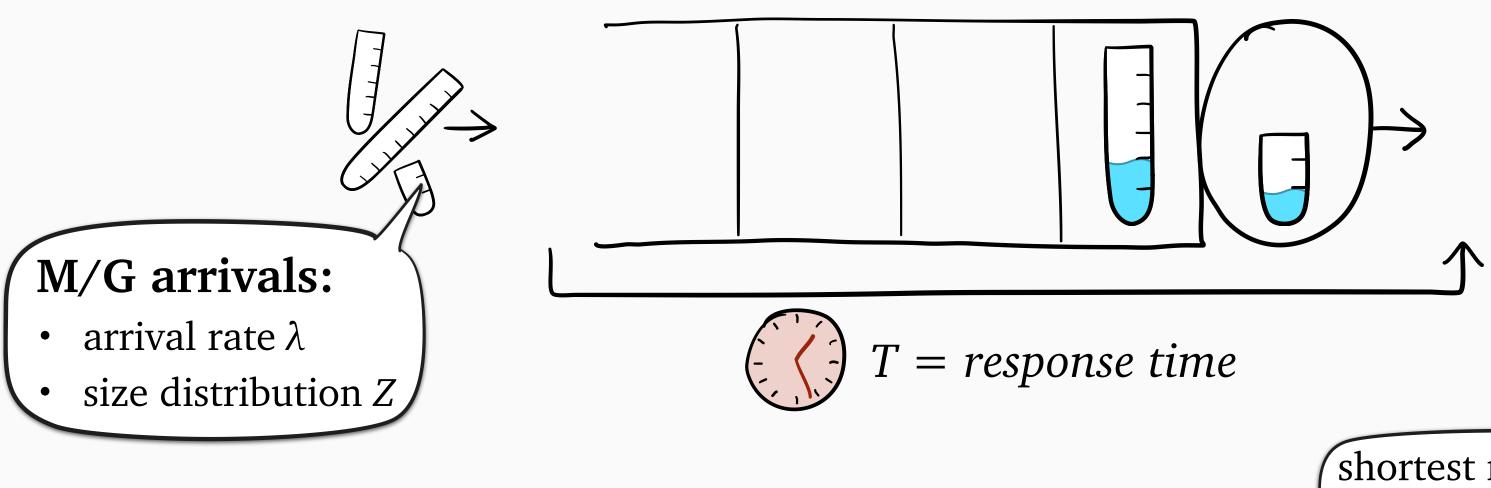


Mean scheduling in the M/G/1

before long jobs



Mean scheduling in the M/G/1





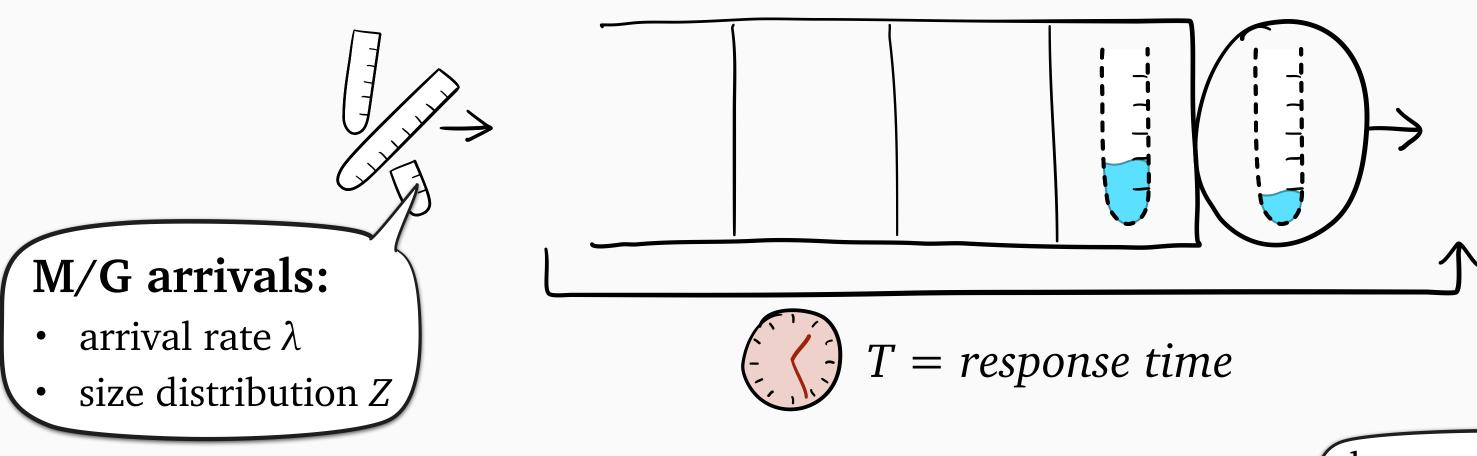
Minimize $\mathbf{E}[T]$?



Serve short jobs before long jobs



shortest remaining processing time





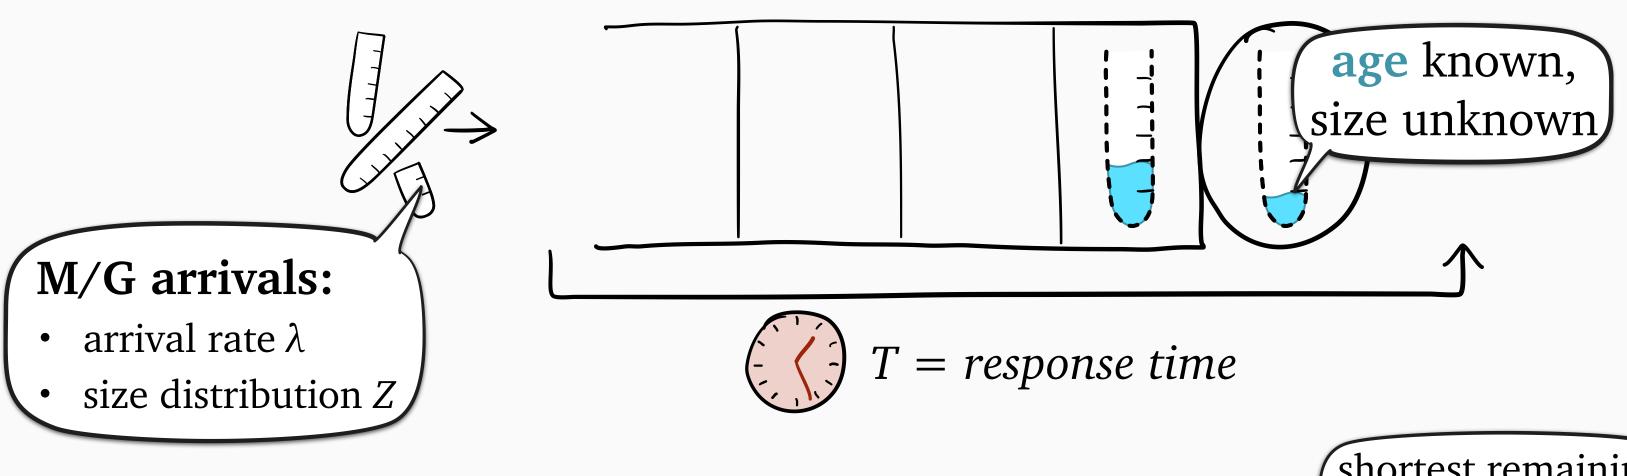
Minimize E[T]?



Serve short jobs before long jobs



shortest remaining processing time





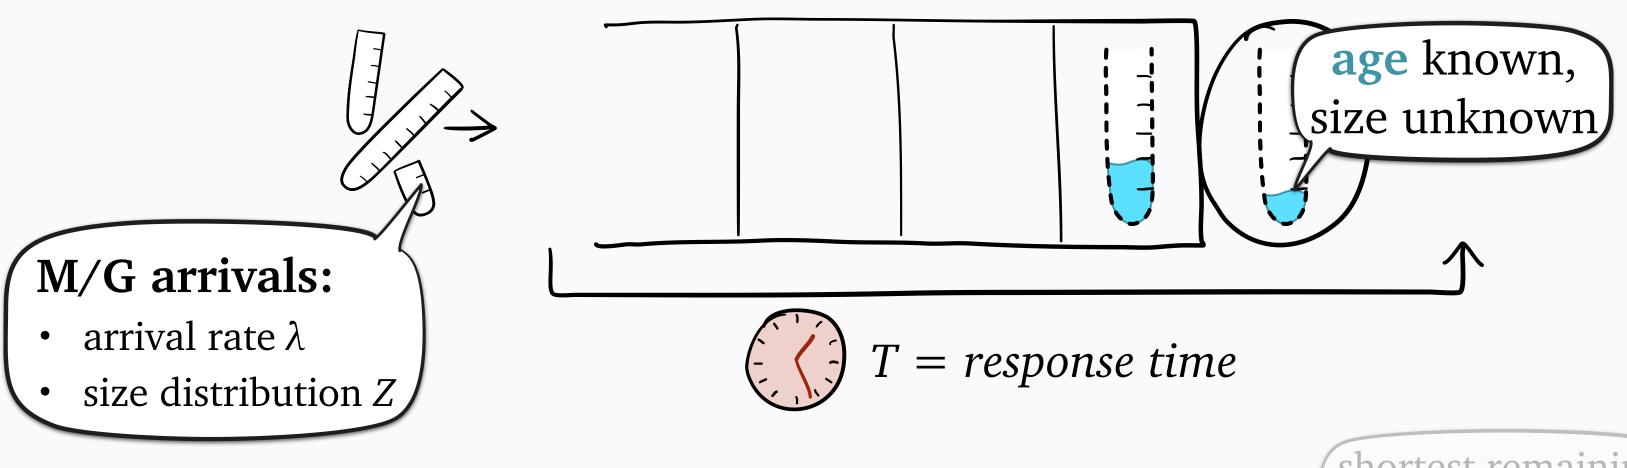
Minimize E[T]?



Serve short jobs before long jobs

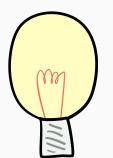


shortest remaining processing time

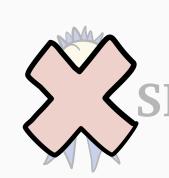




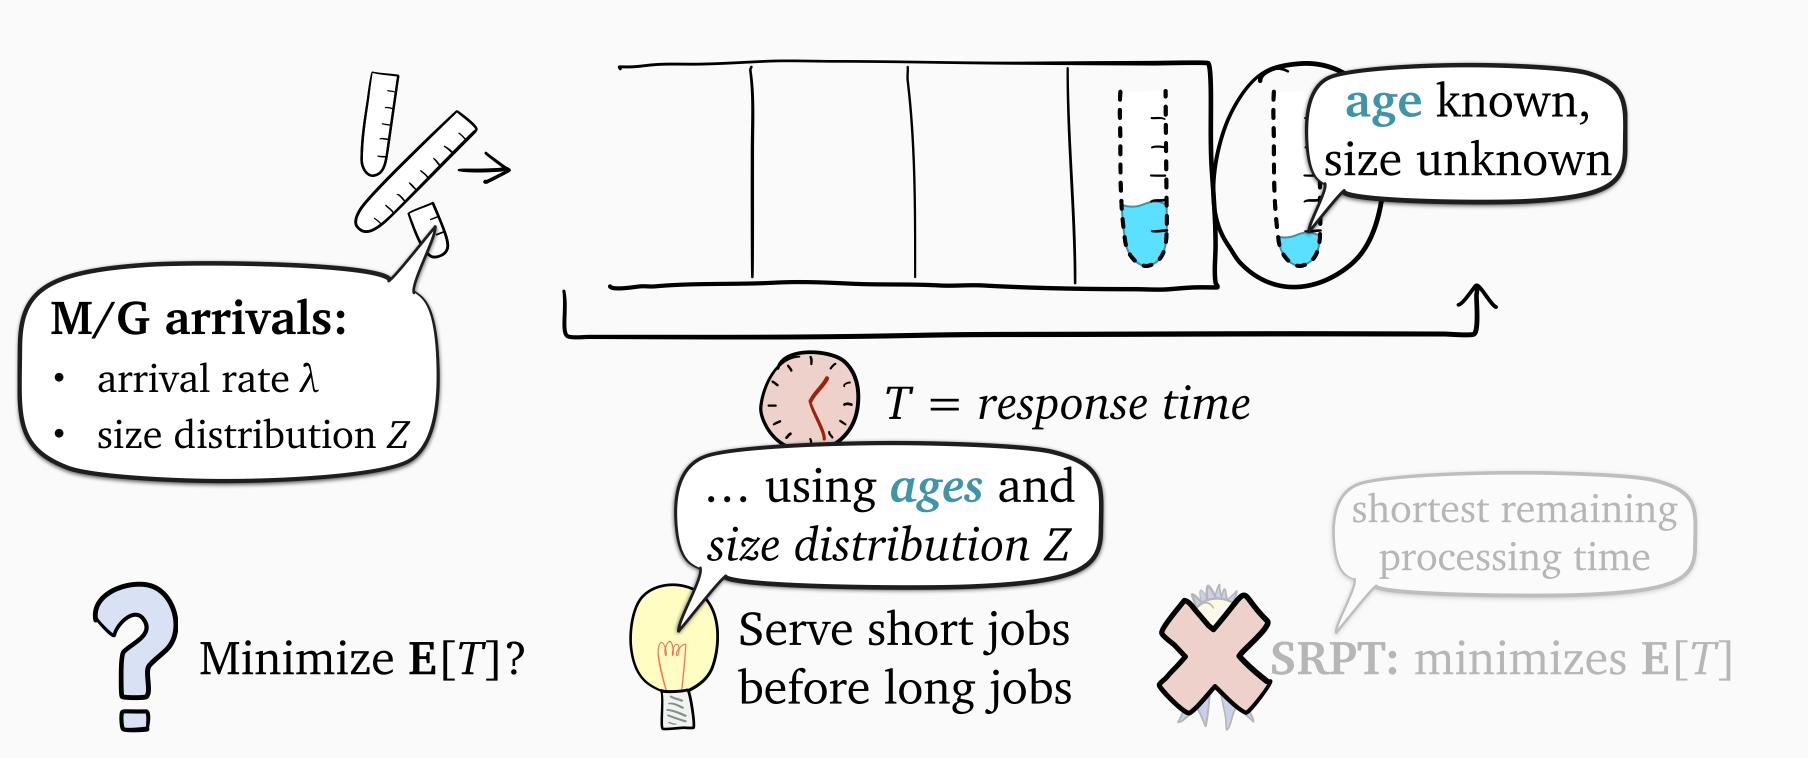
Minimize E[T]?

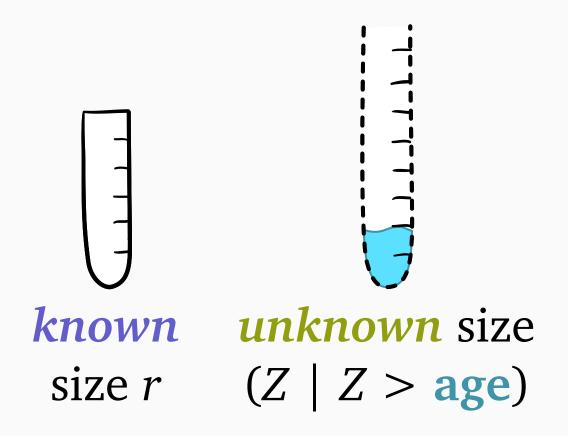


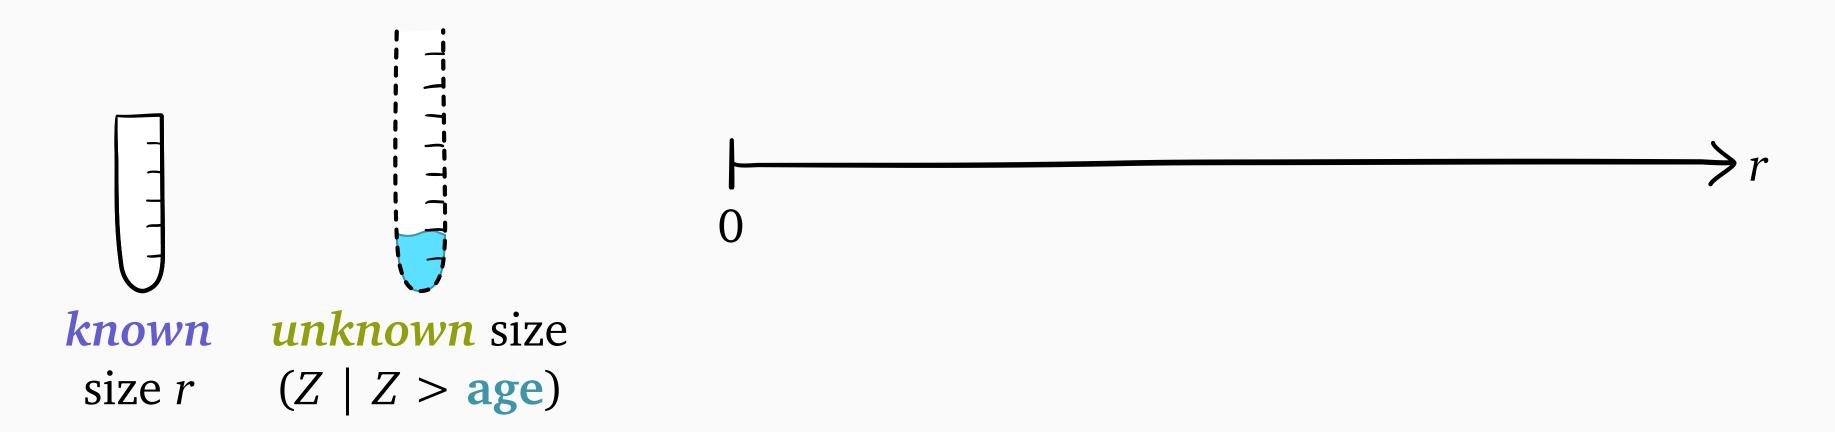
Serve short jobs before long jobs

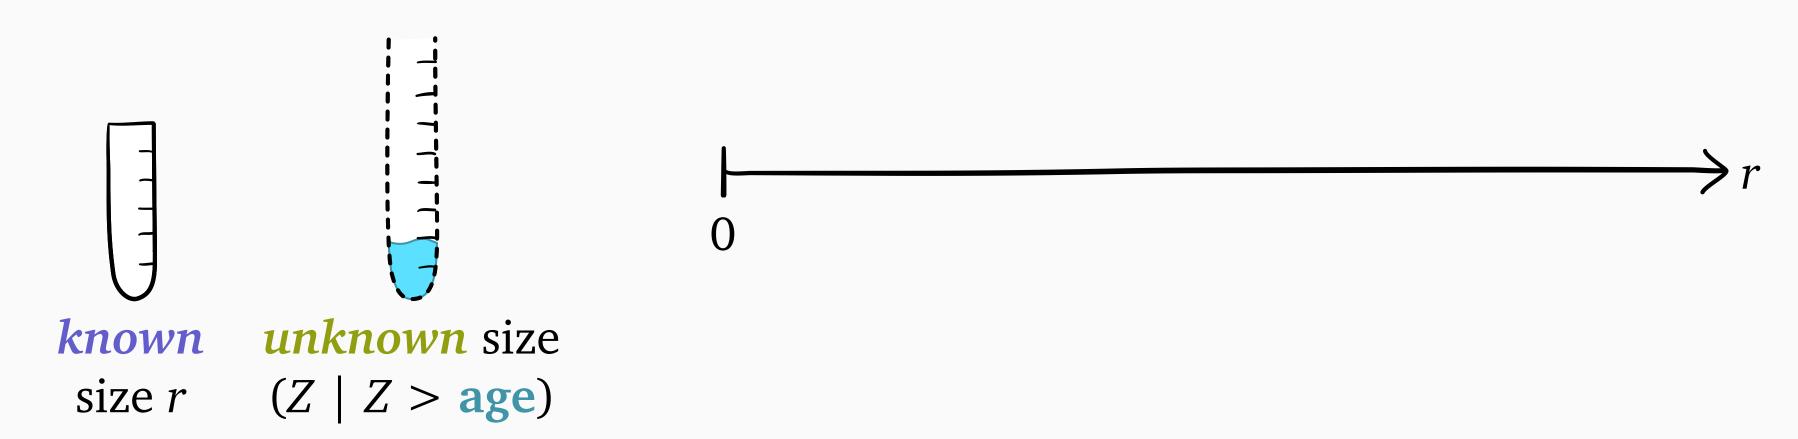


shortest remaining processing time

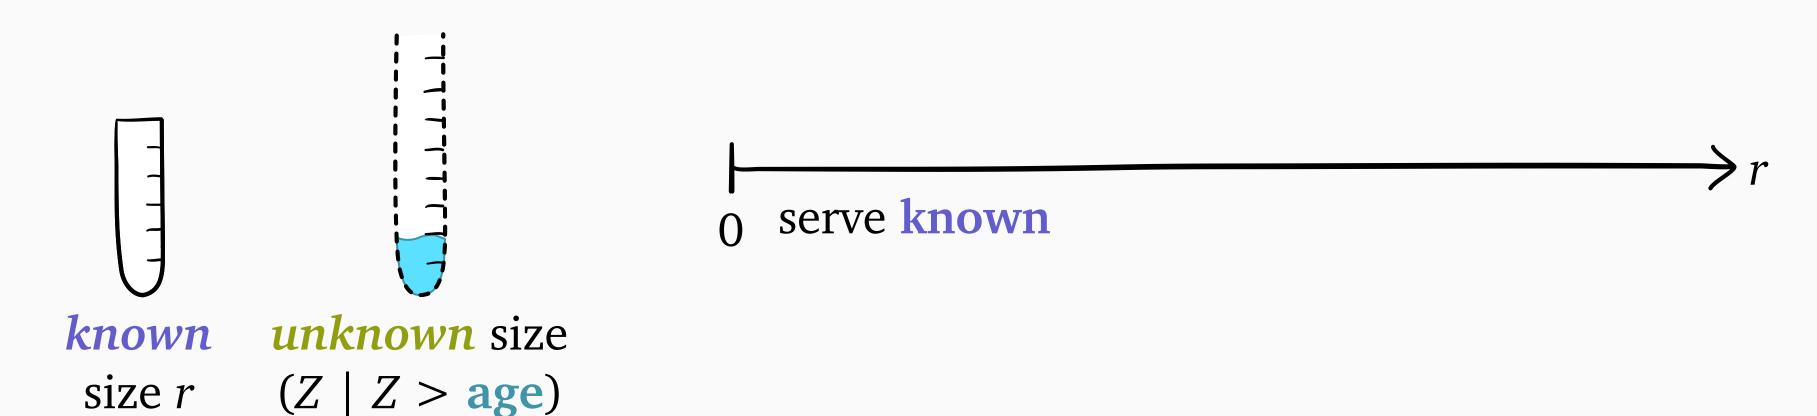


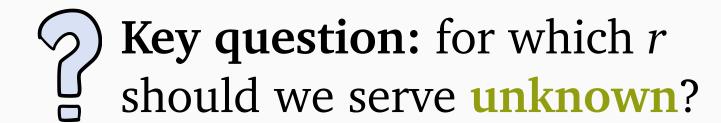


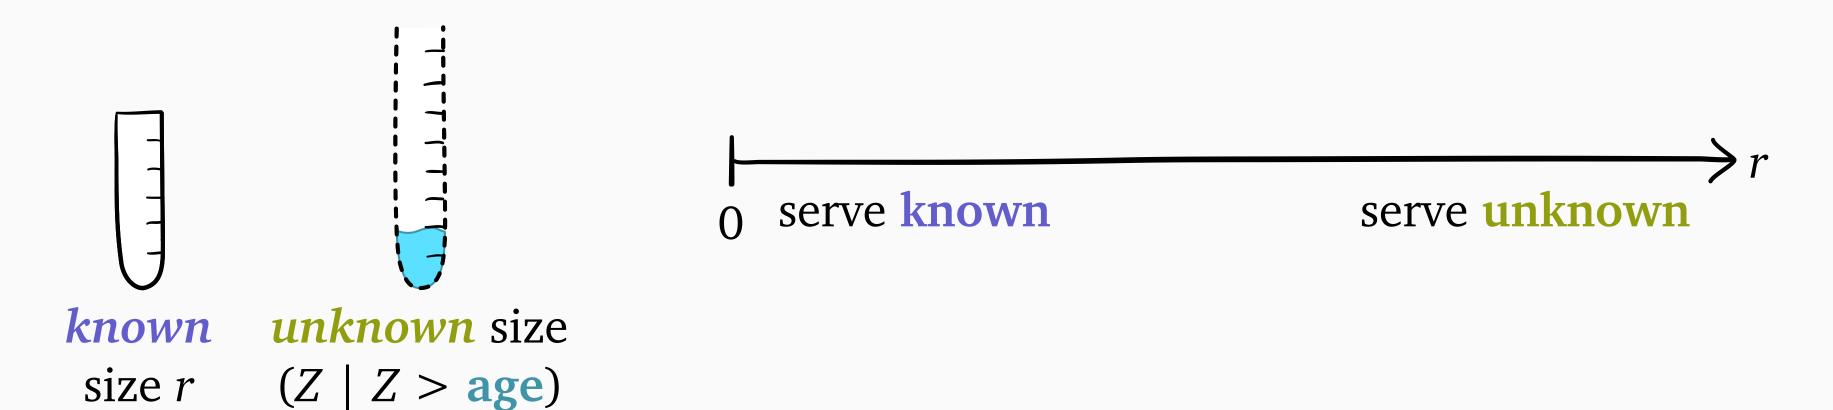




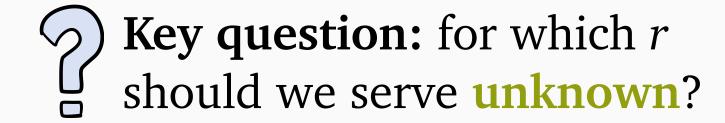


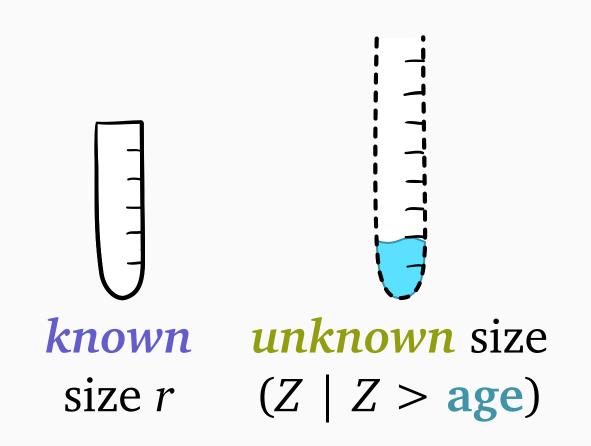


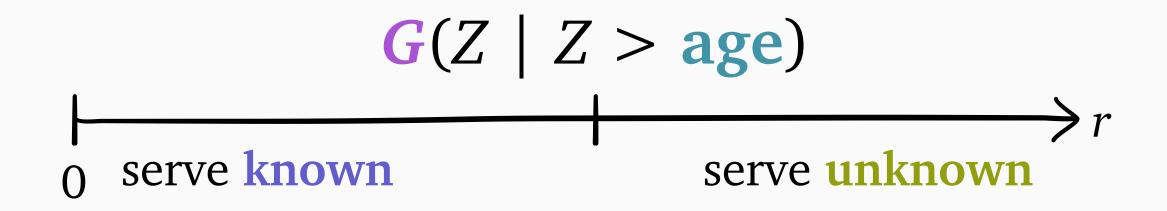


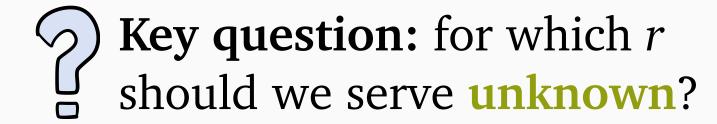


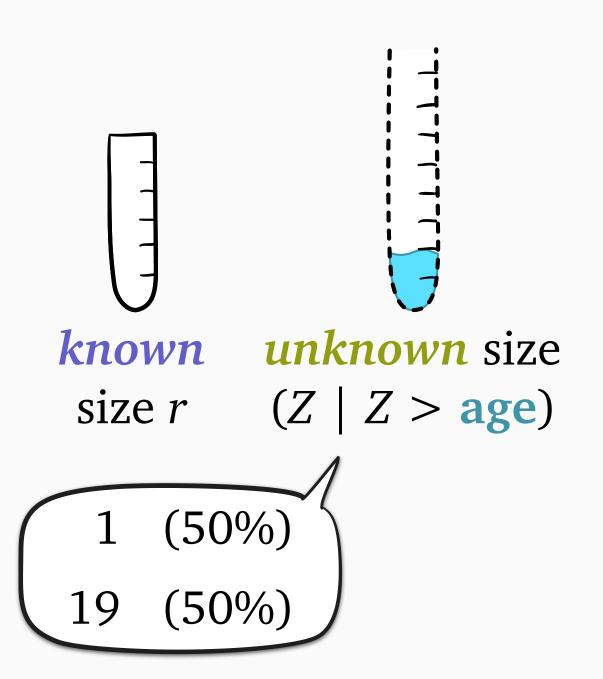
size r

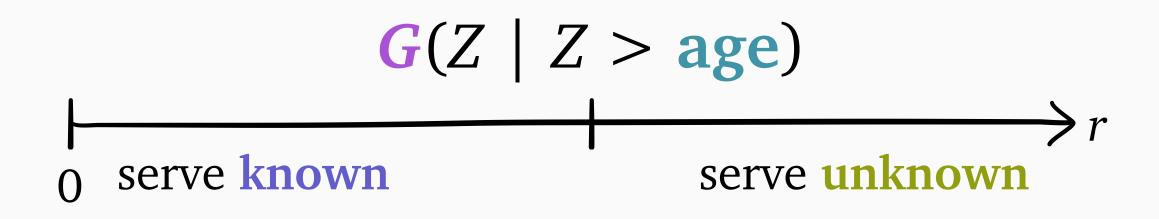


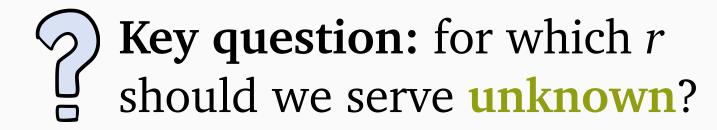


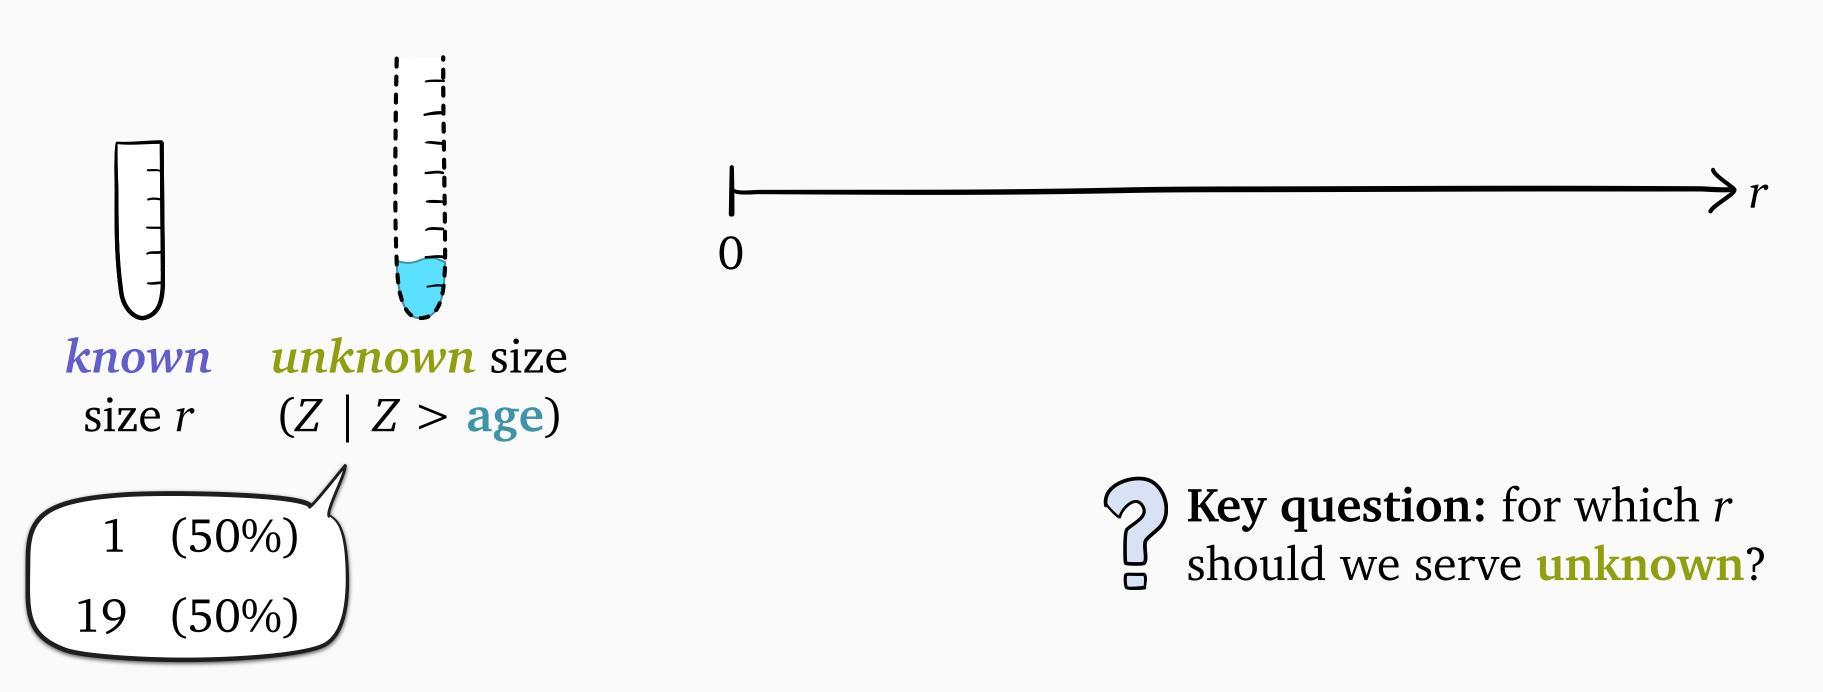


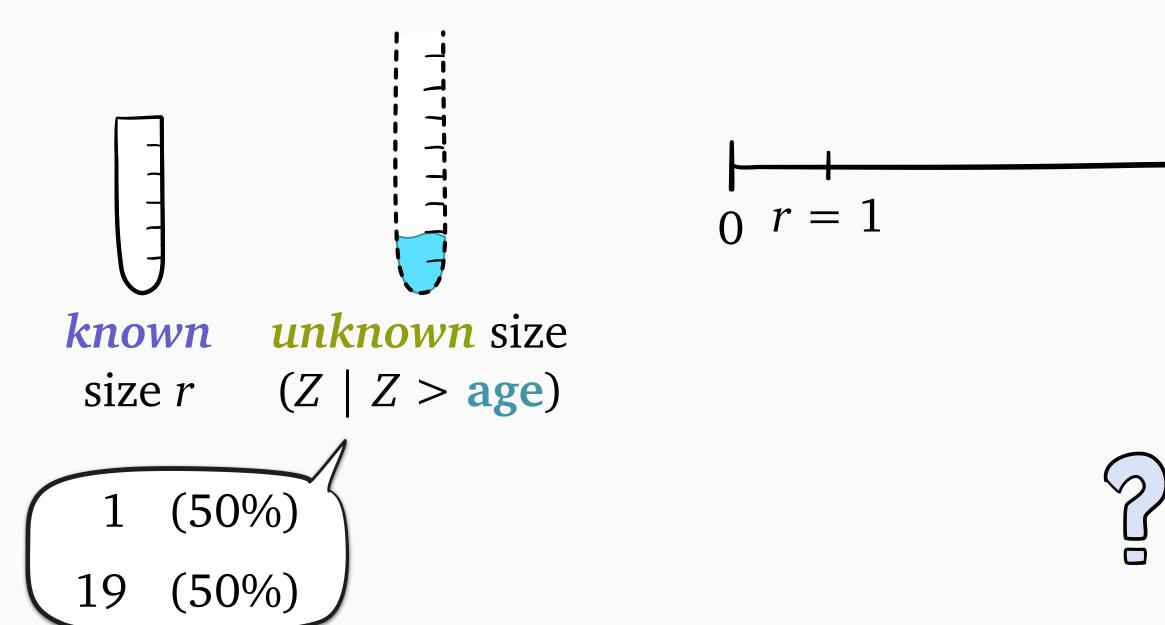




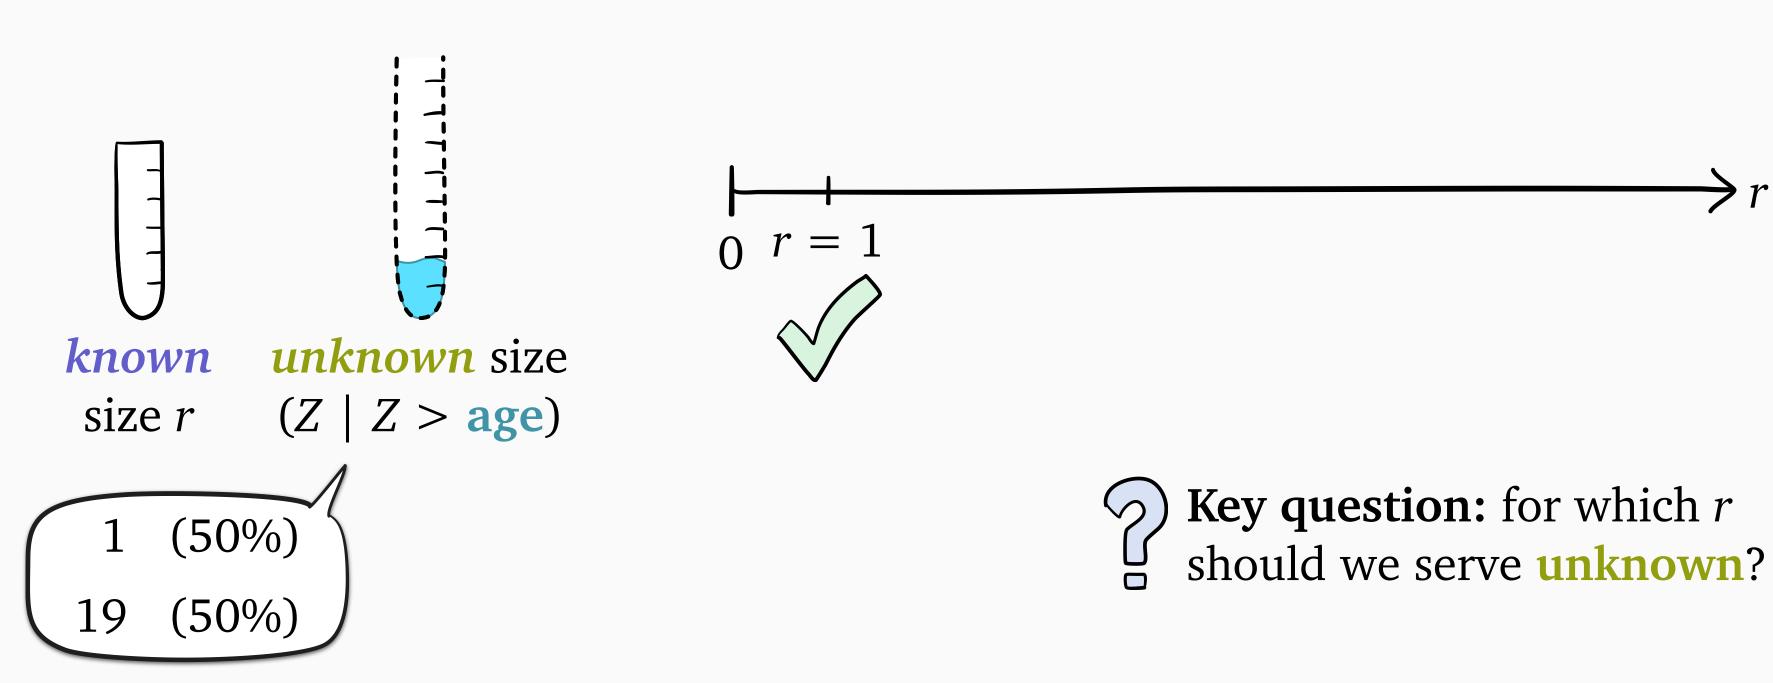


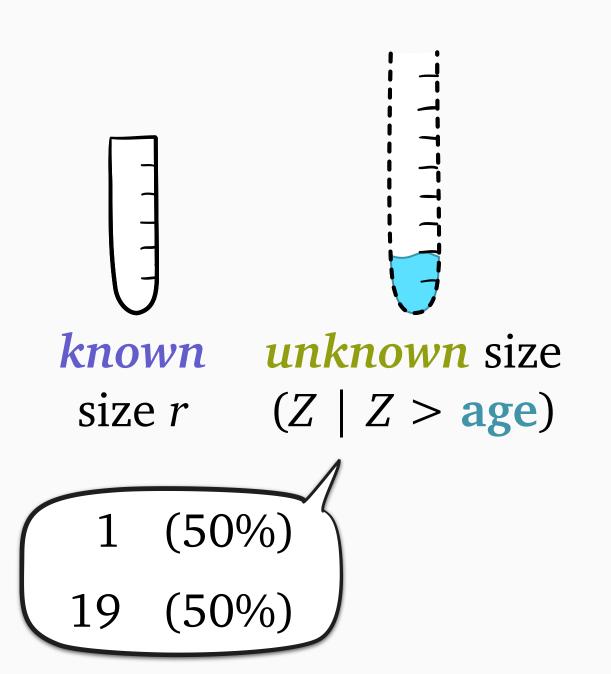


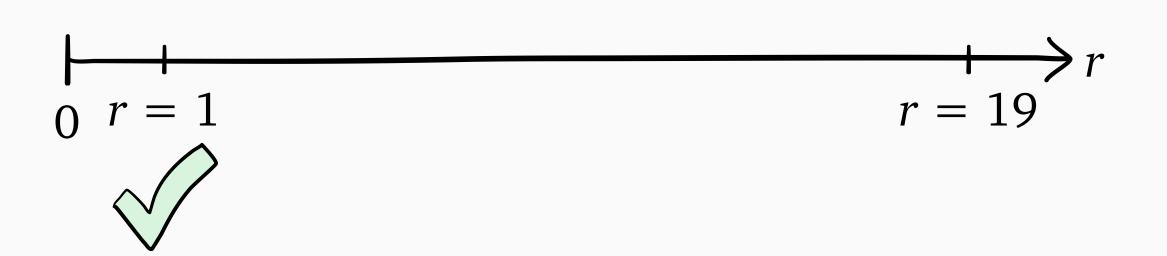


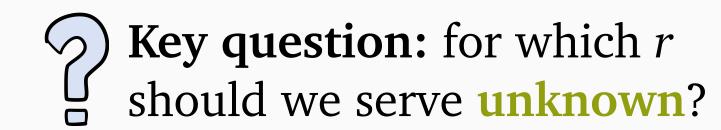


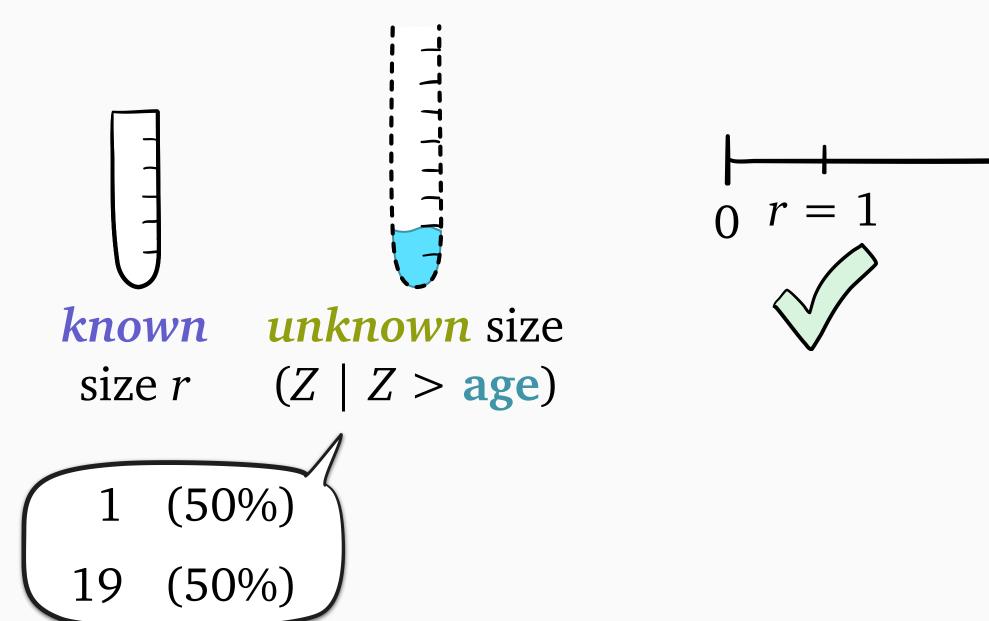


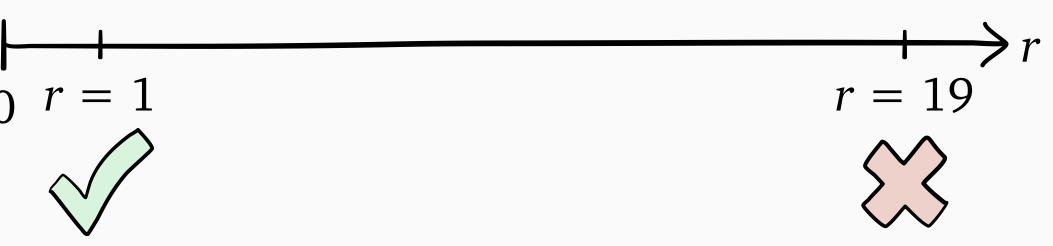


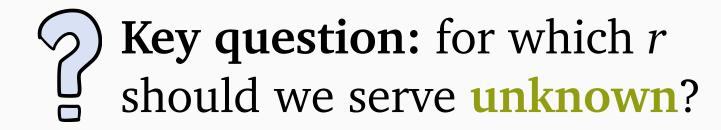


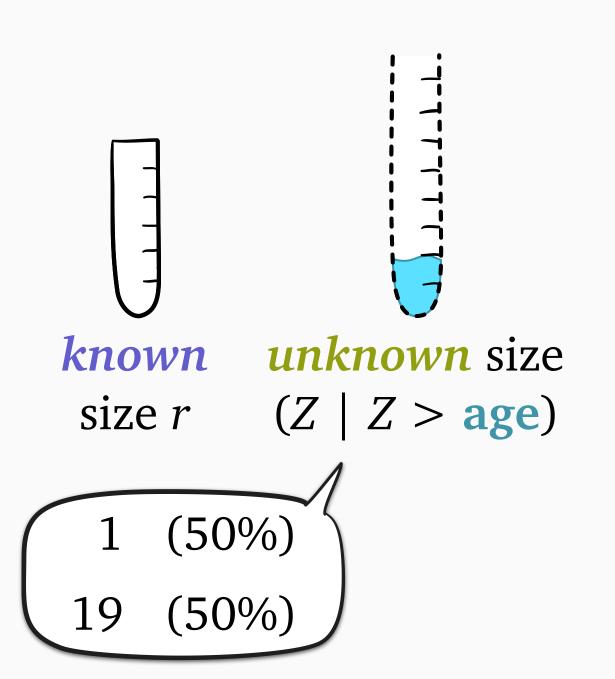


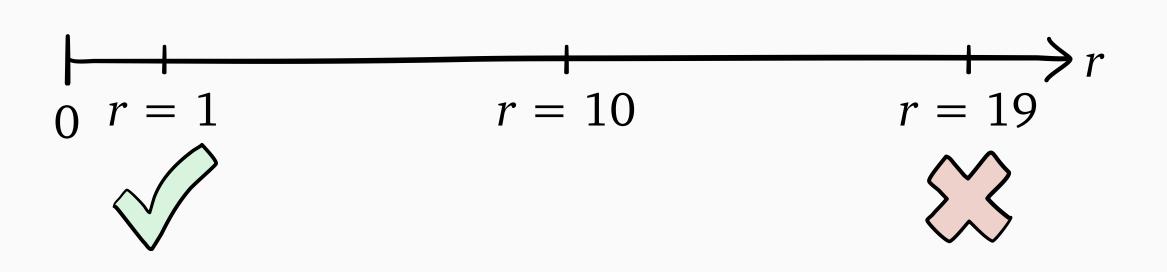


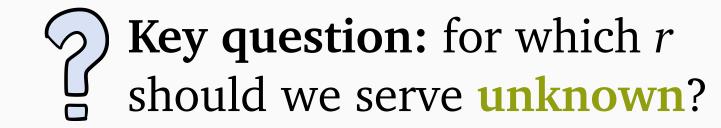


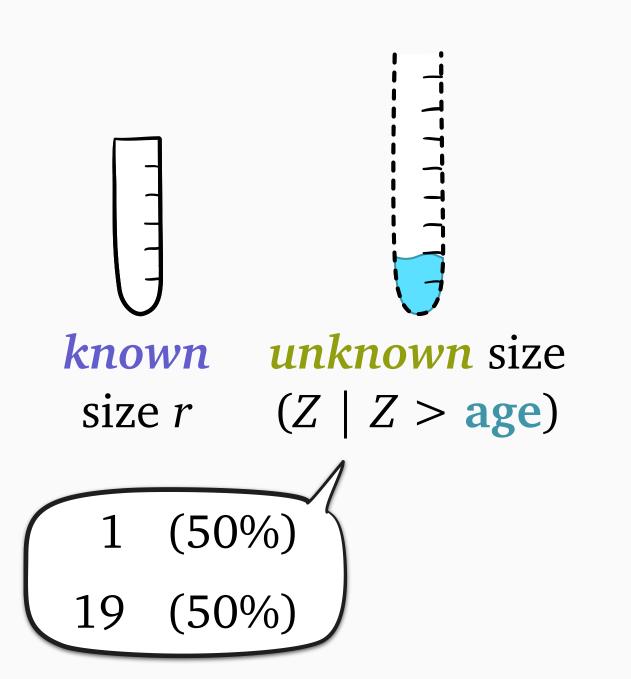


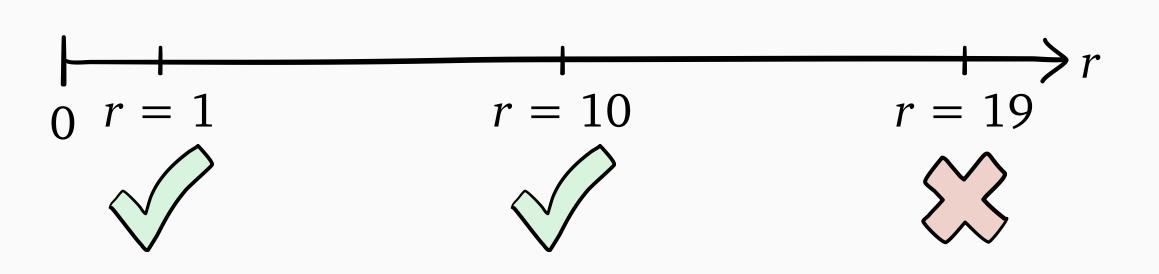


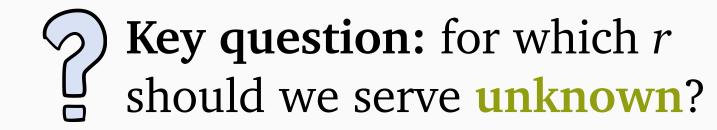


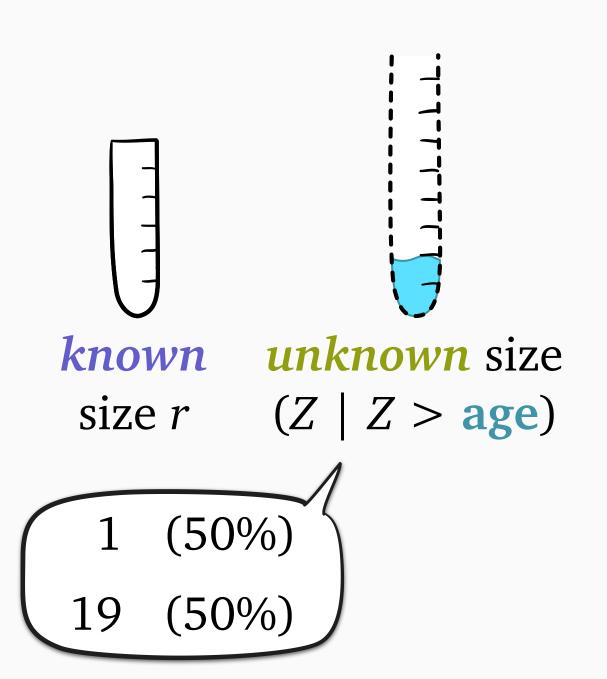


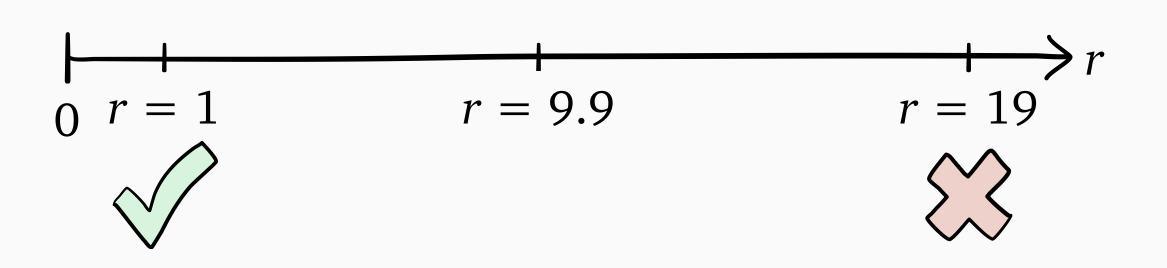


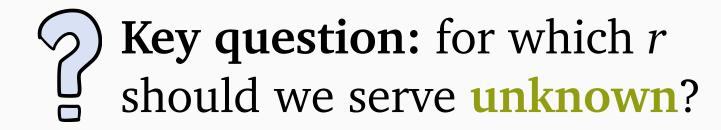


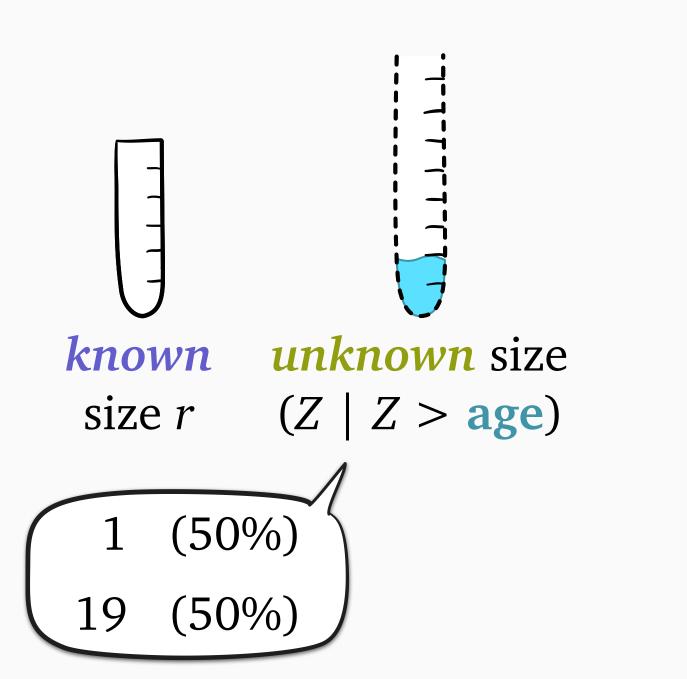


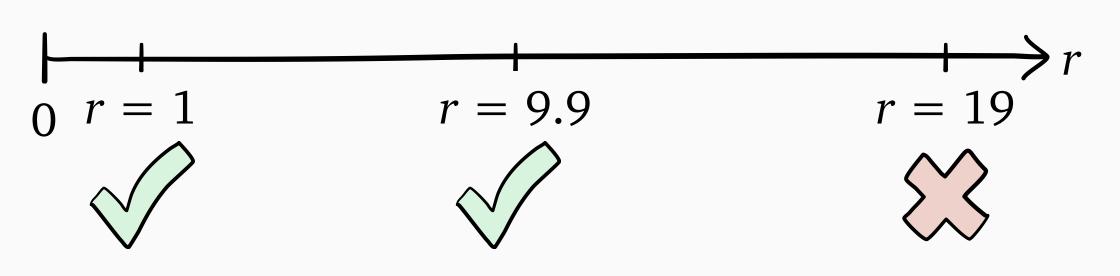


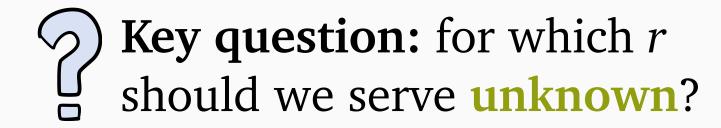


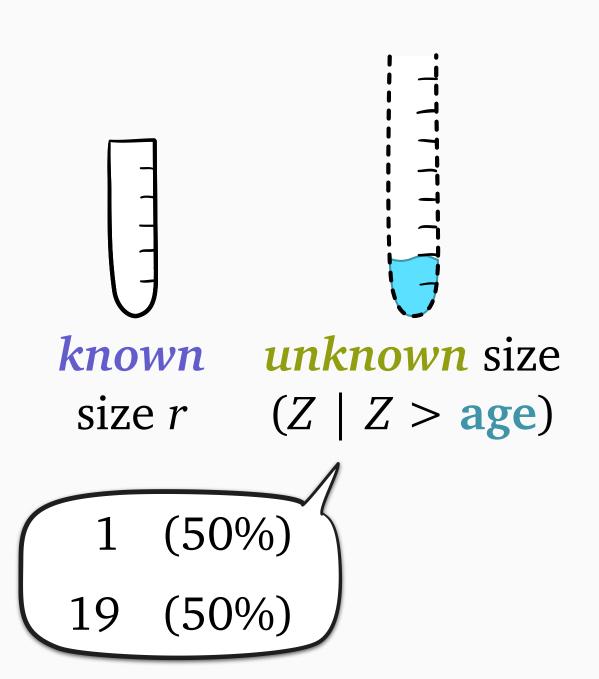


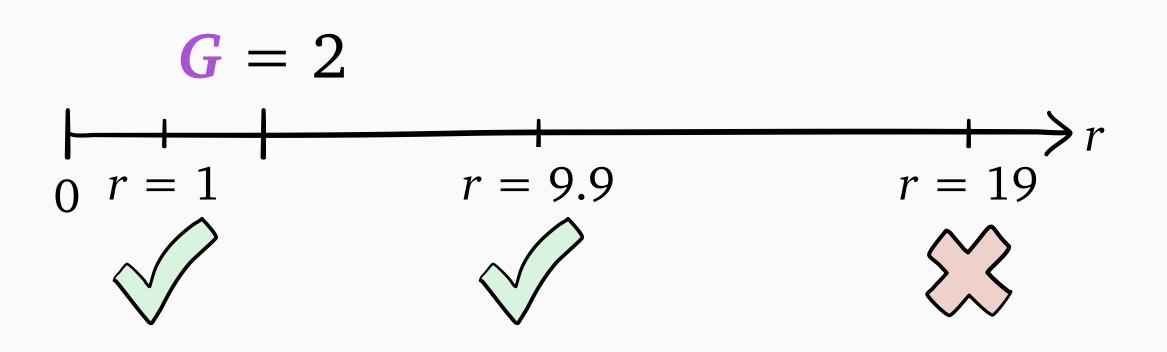


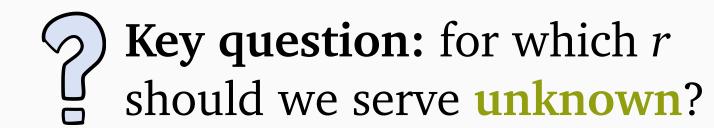


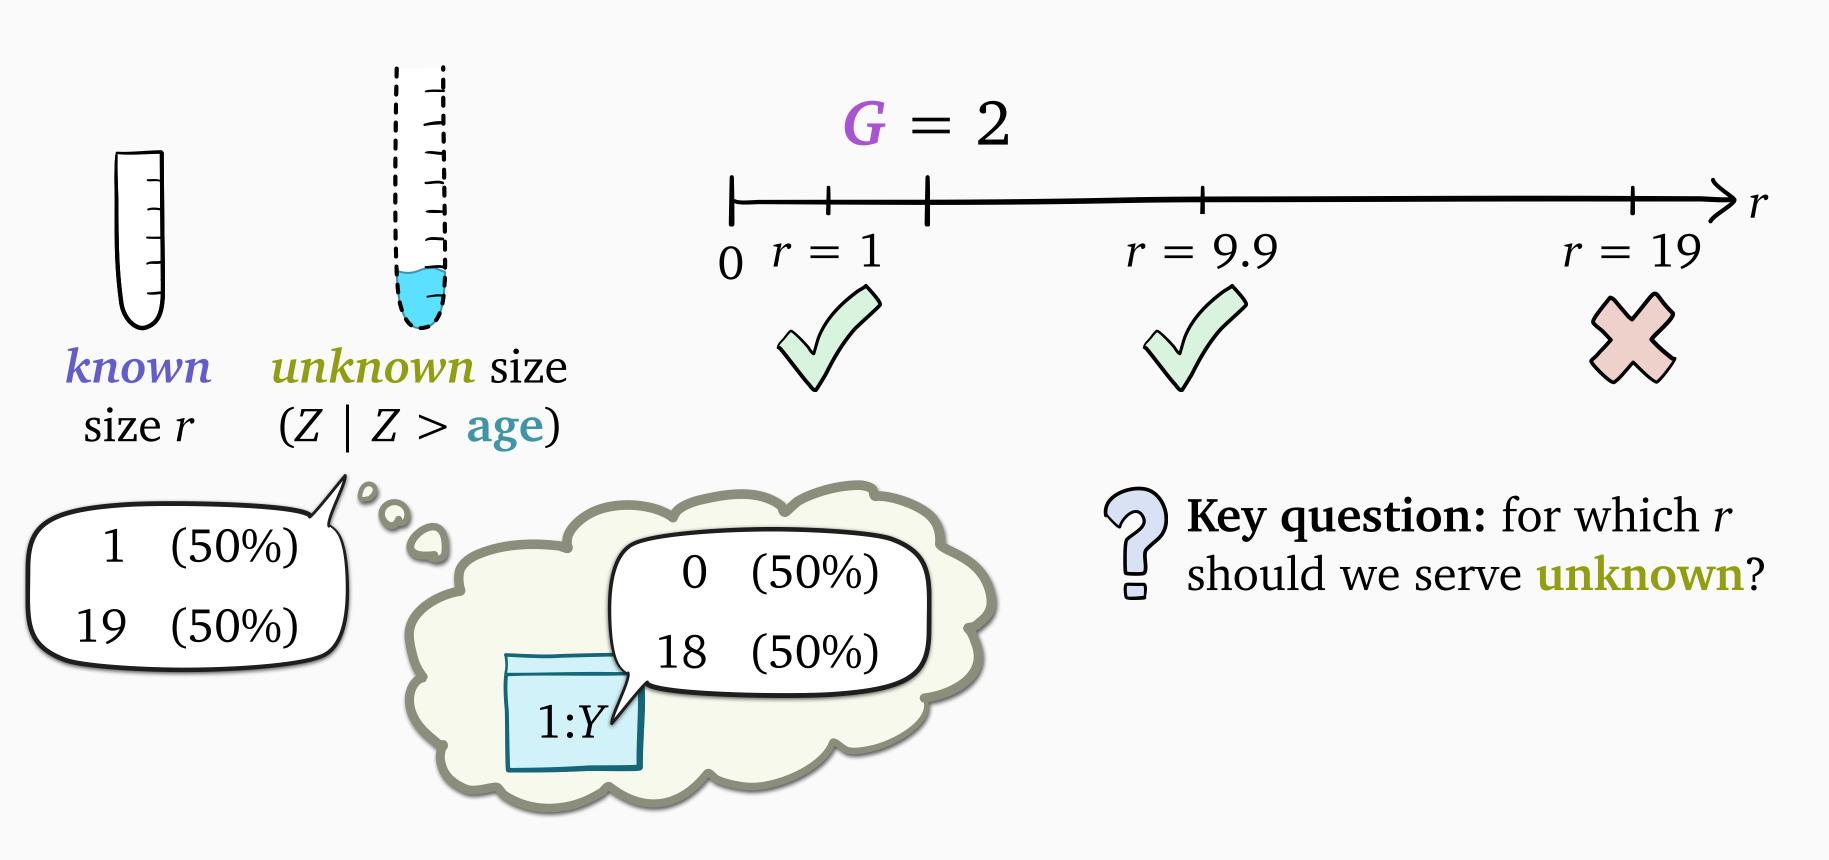


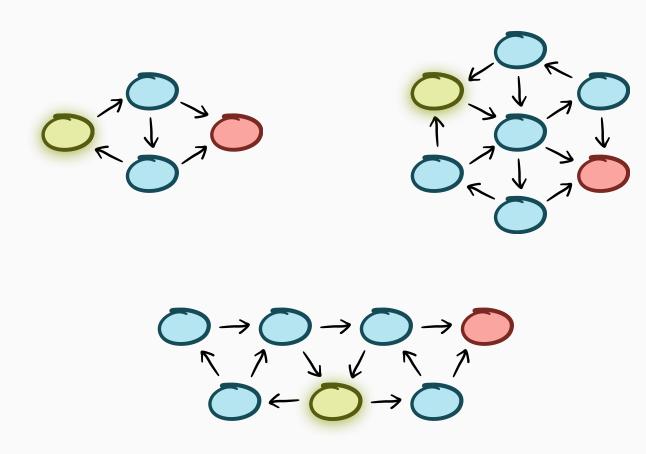


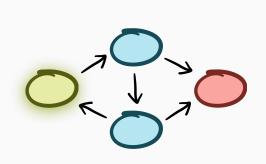


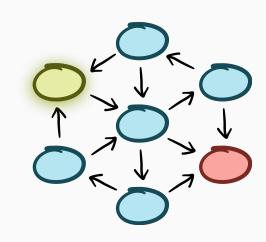


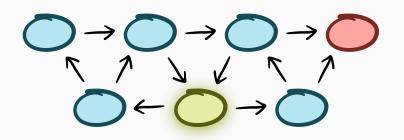






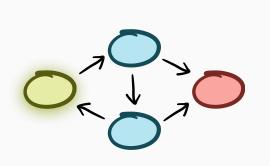


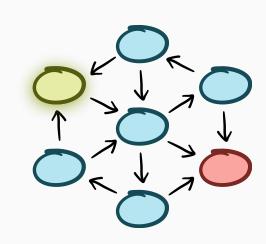


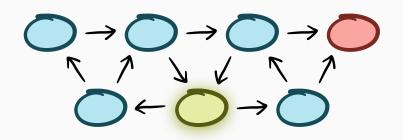


Decision process:

- Advance one Markov chain every step
- Get reward based on its current state
- Stop when one Markov chain terminates



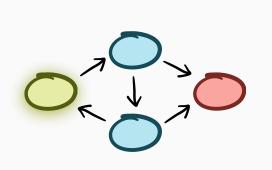


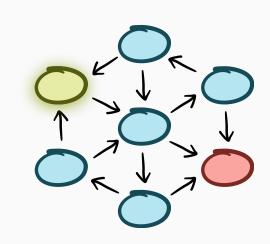


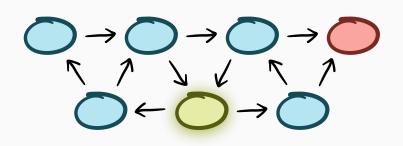
Decision process:

- Advance one Markov chain every step
- Get reward based on its current state
- Stop when one Markov chain terminates

Goal: maximize **E**[total reward]





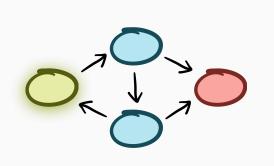


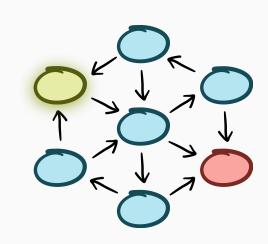
Decision process:

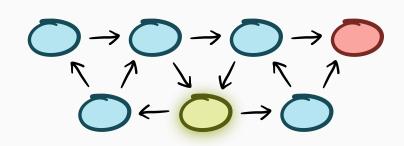
- Advance one Markov chain every step
- Get reward based on its current state
- Stop when one Markov chain terminates

Goal: maximize **E**[total reward]

Job scheduling:







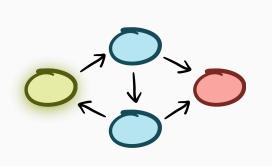
Decision process:

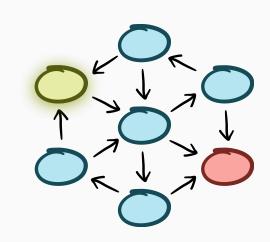
- Advance one Markov chain every step
- Get reward based on its current state
- Stop when one Markov chain terminates

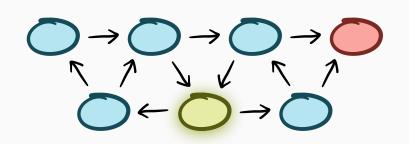
Goal: maximize **E**[total reward]

Job scheduling:

• All transitions incur reward -1







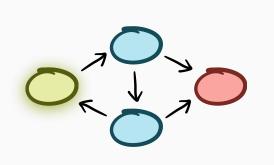
Decision process:

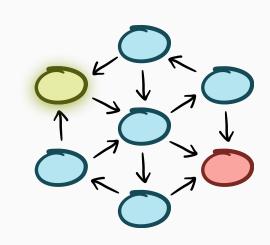
- Advance one Markov chain every step
- Get reward based on its current state
- Stop when one Markov chain terminates

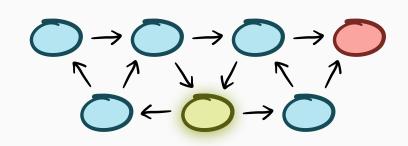
Goal: maximize **E**[total reward]

Job scheduling:

- All transitions incur reward -1
- **Terminal state** = "job complete"







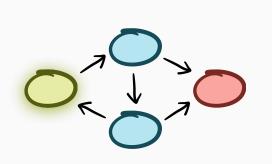
Decision process:

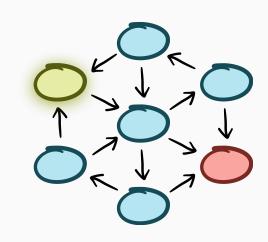
- Advance one Markov chain every step
- Get reward based on its current state
- Stop when one Markov chain terminates

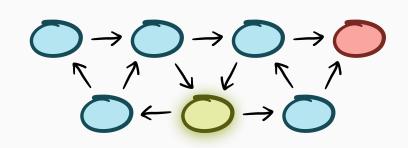
Goal: maximize **E**[total reward]

Job scheduling:

- All transitions incur reward -1
- **Terminal state** = "job complete"
- Minimizing **E**[time to first completion]







Decision process:

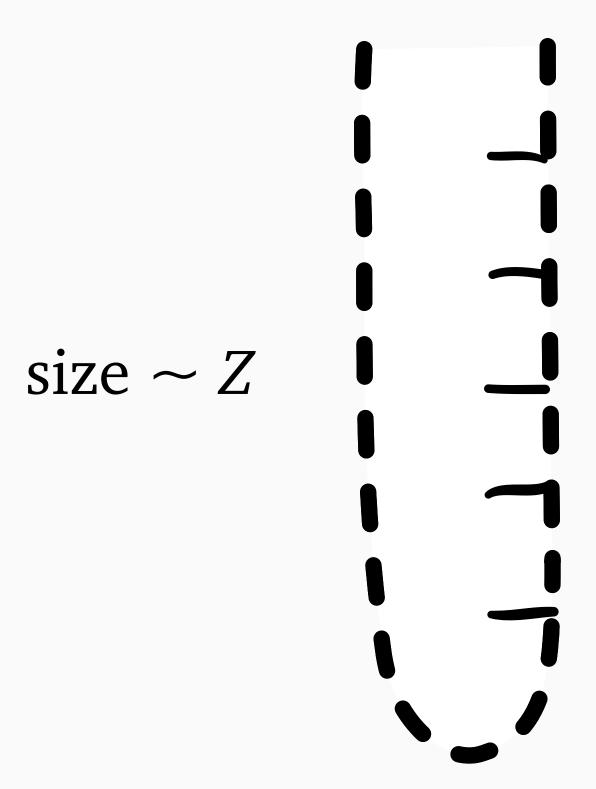
- Advance one Markov chain every step
- Get reward based on its current state
- Stop when one Markov chain terminates

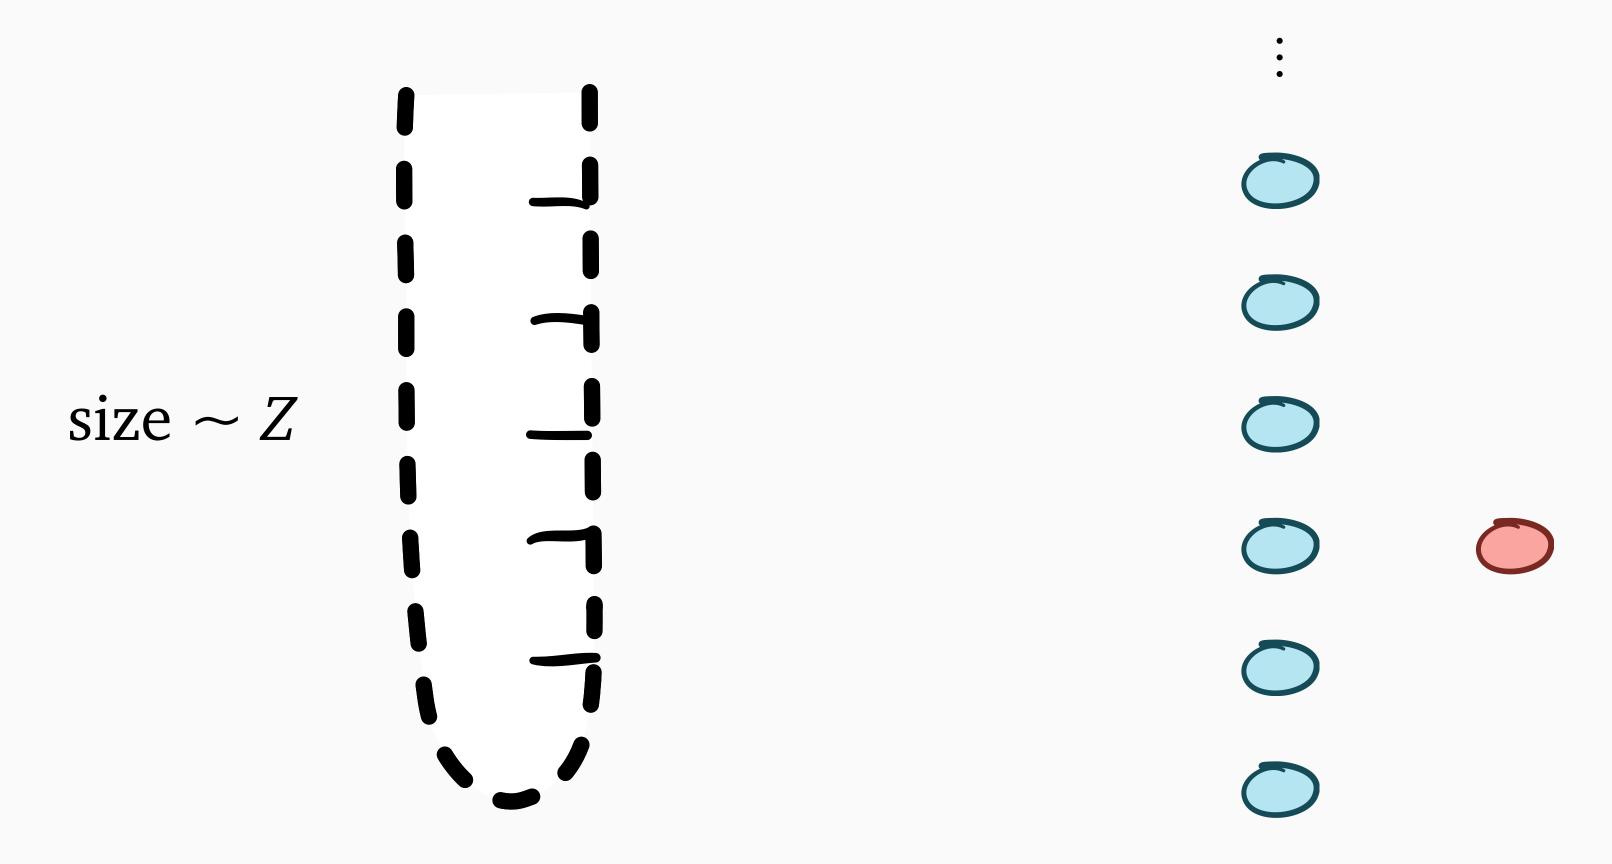
Goal: maximize E[total reward]

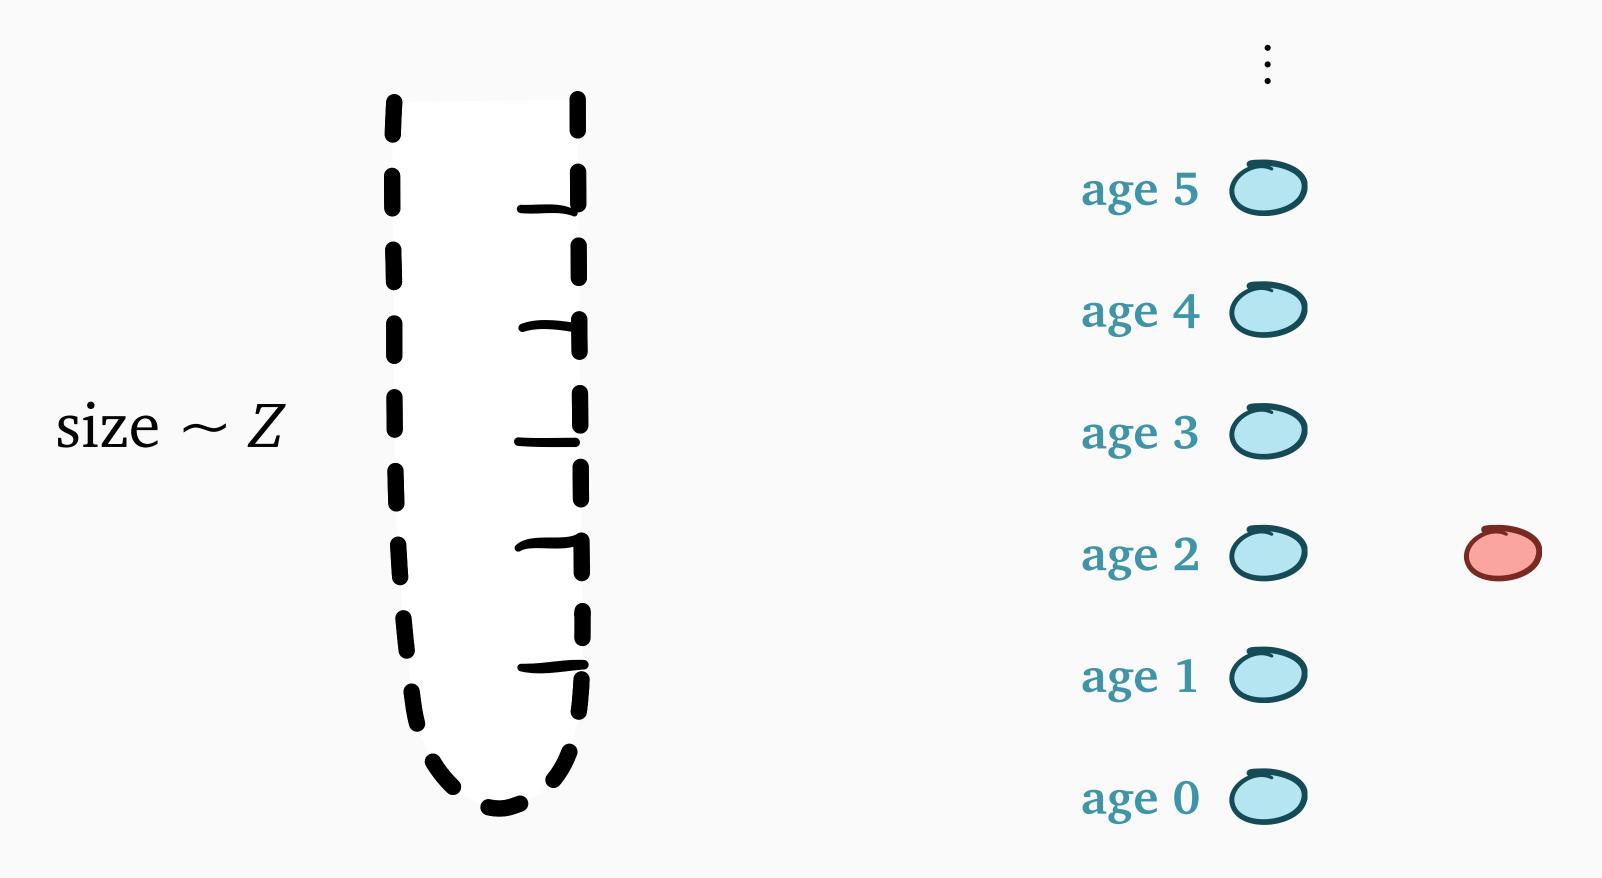
Job scheduling:

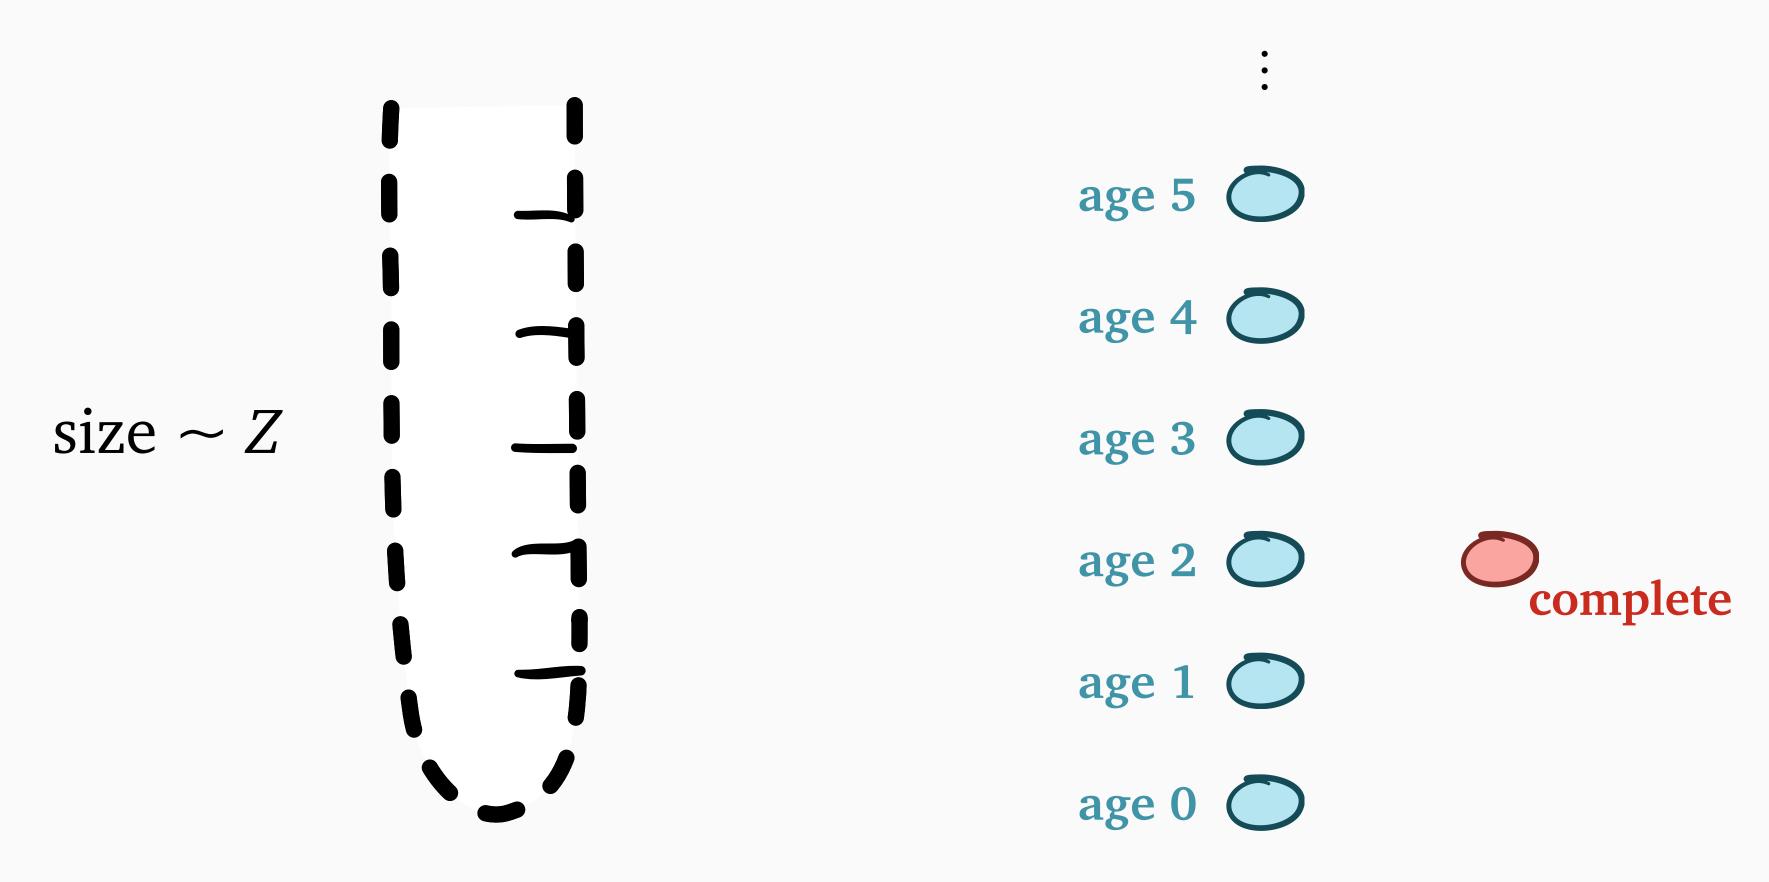
- All transitions incur reward -1
- **Terminal state** = "job complete"
- Minimizing E[time to first completion]

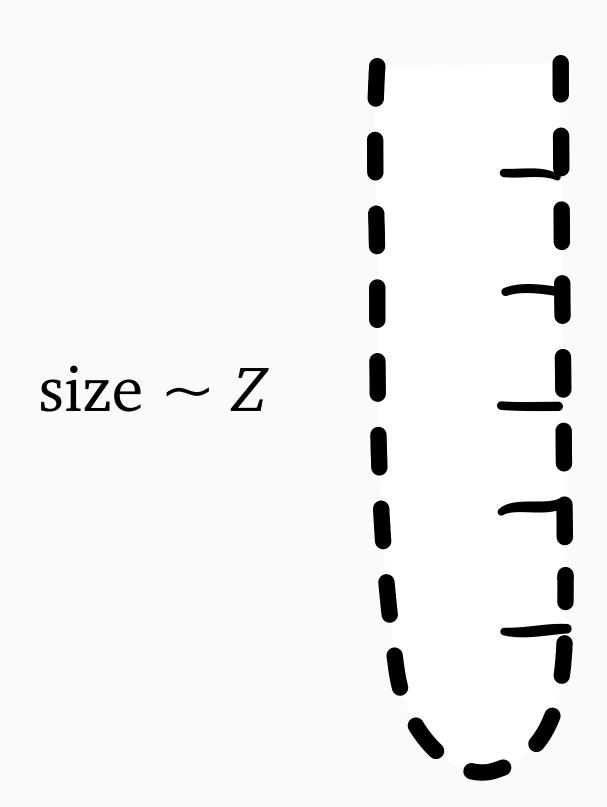
Want to minimize, for all k, $\mathbf{E}[\text{time to }k\text{th completion}]$

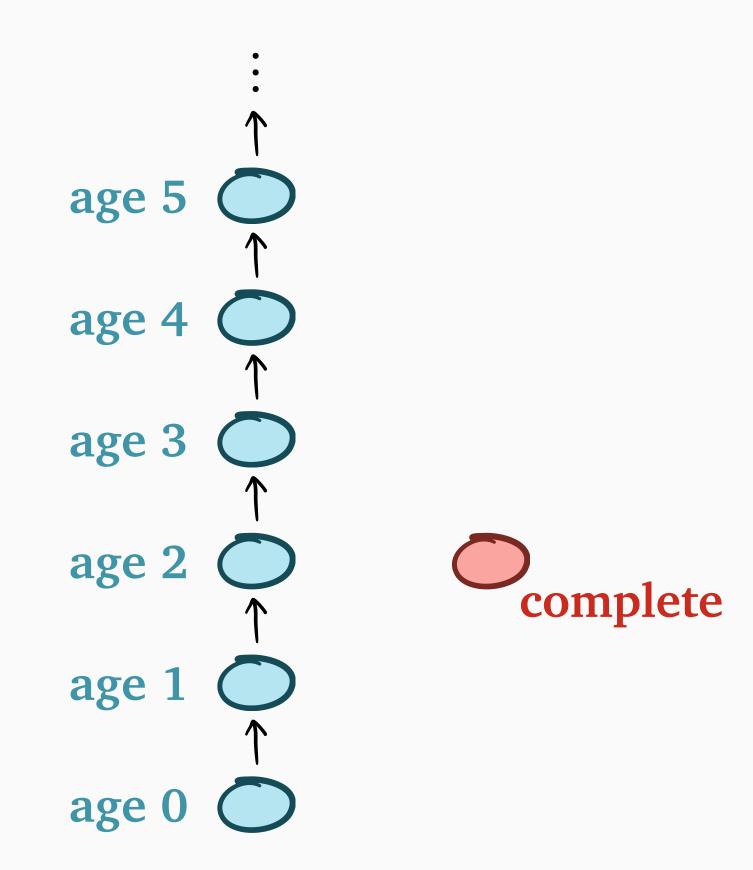


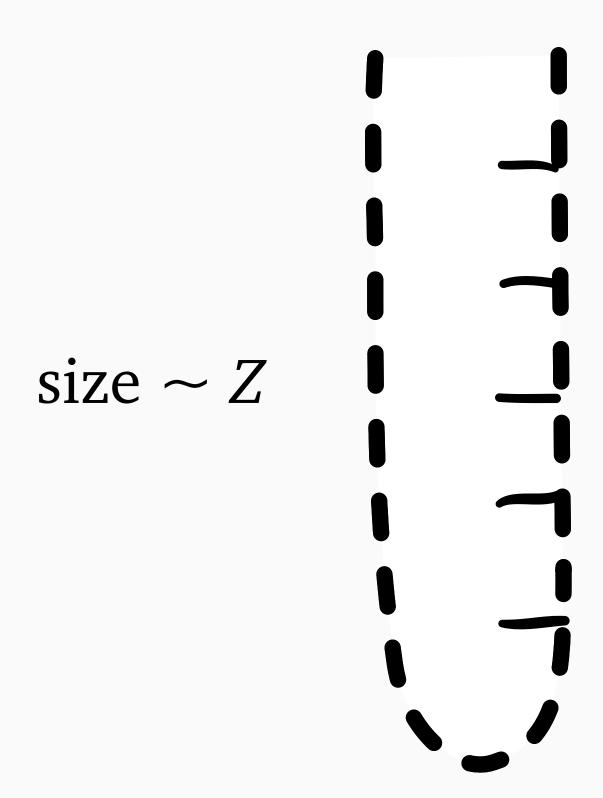


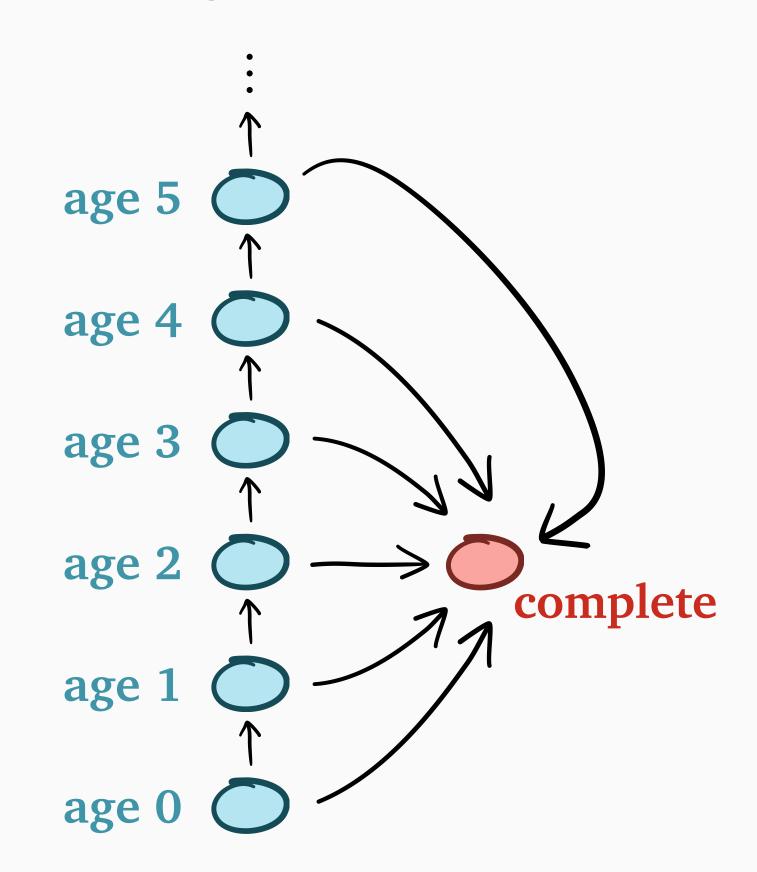


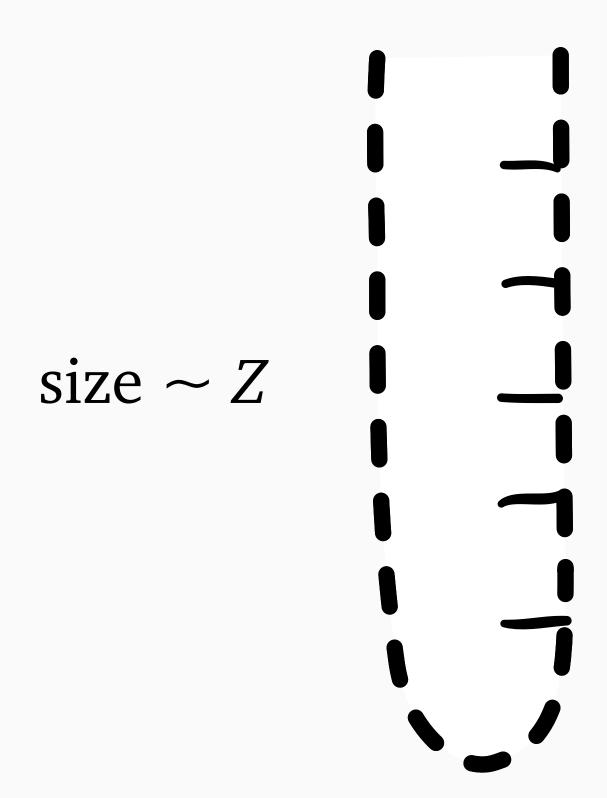


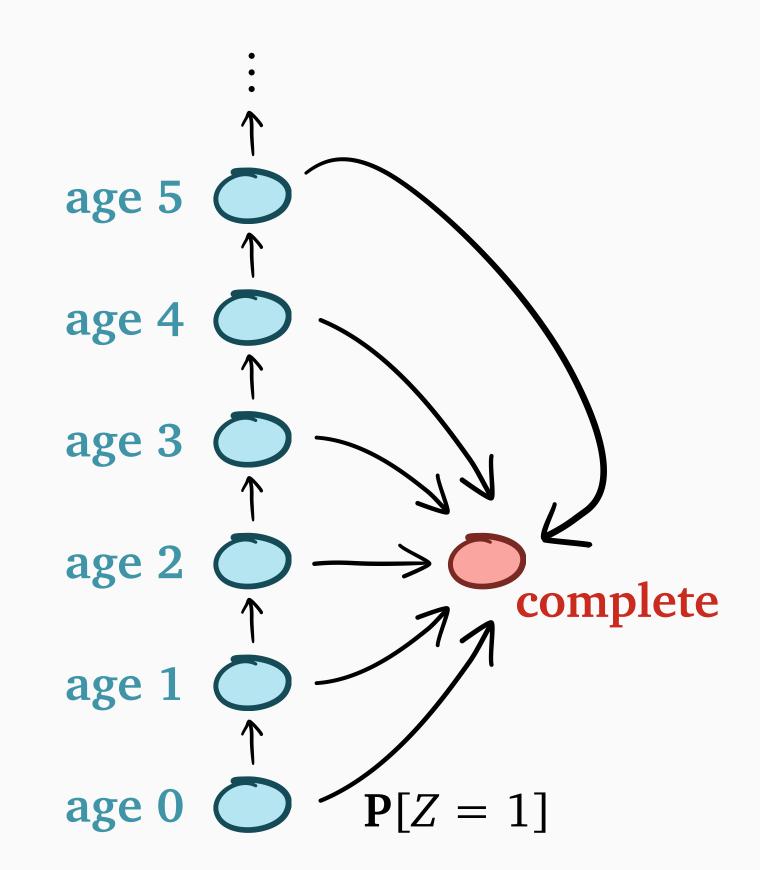


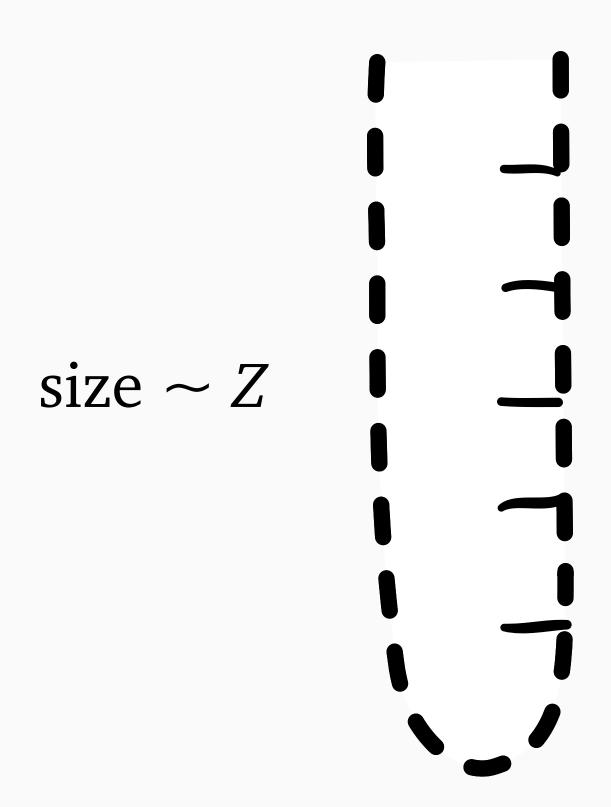


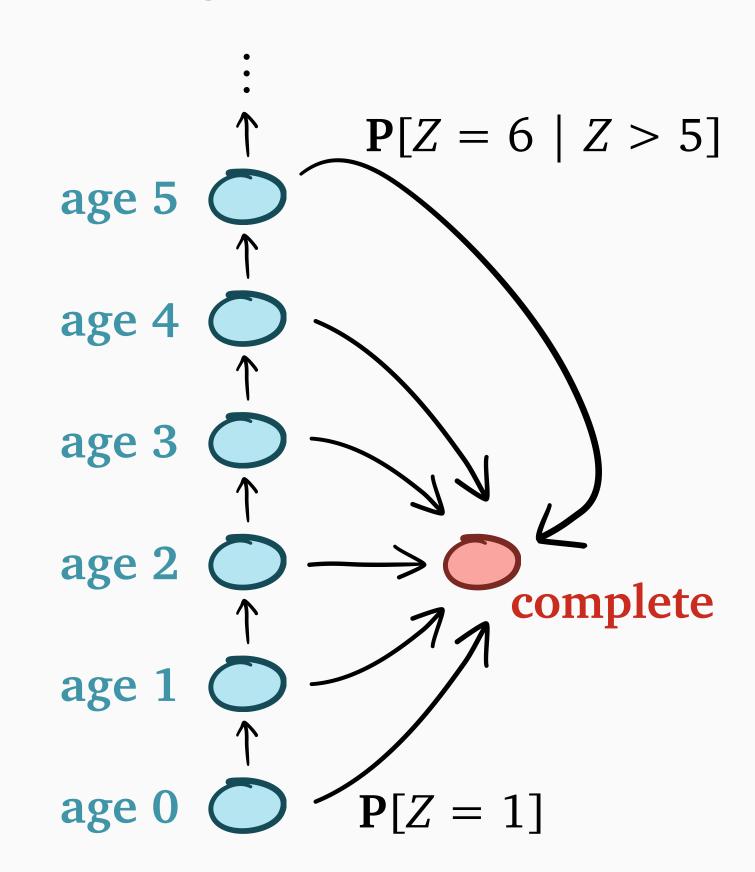


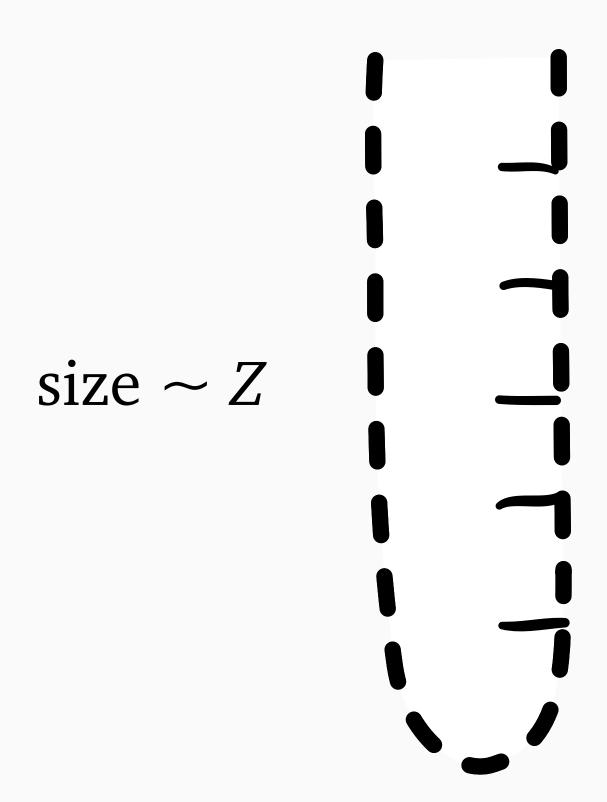


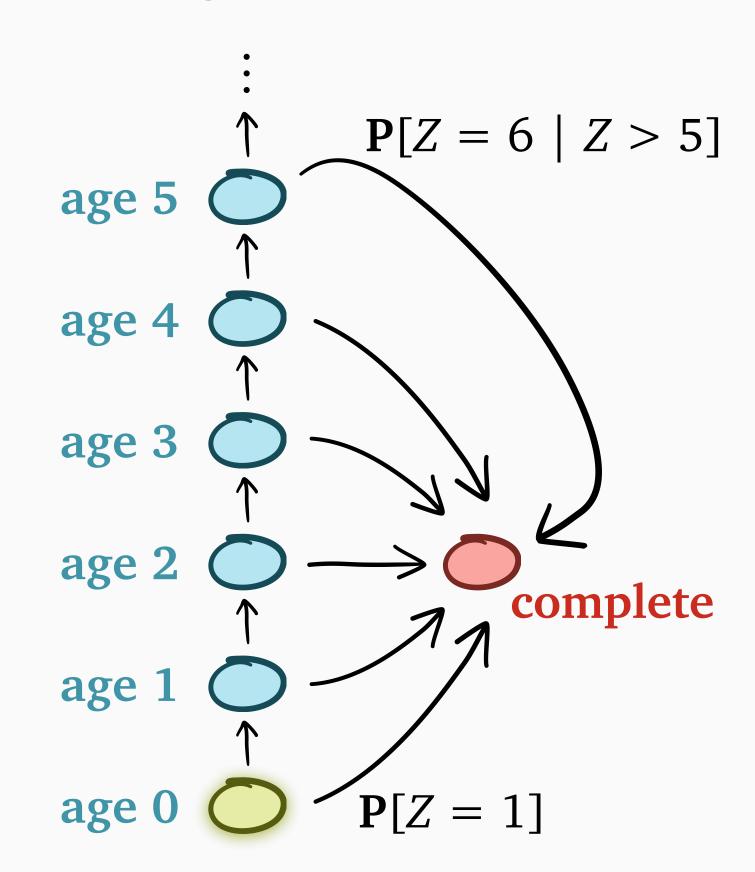


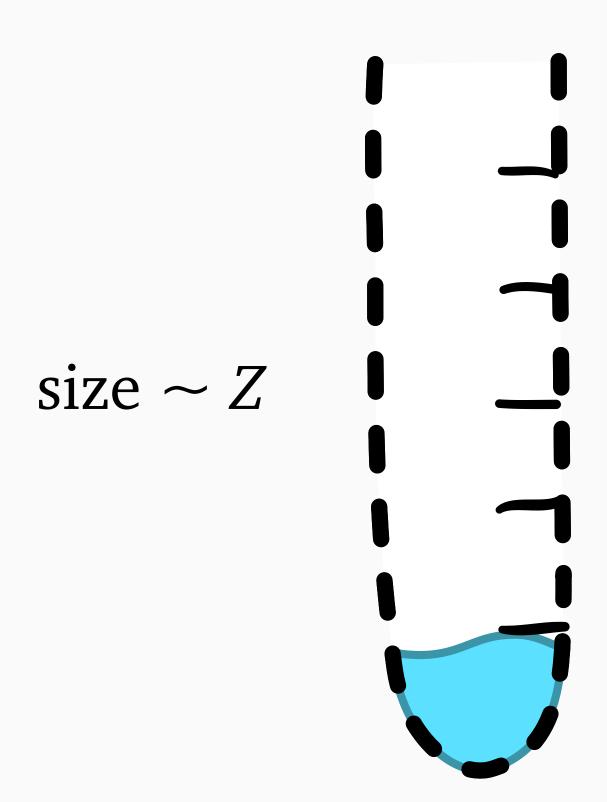


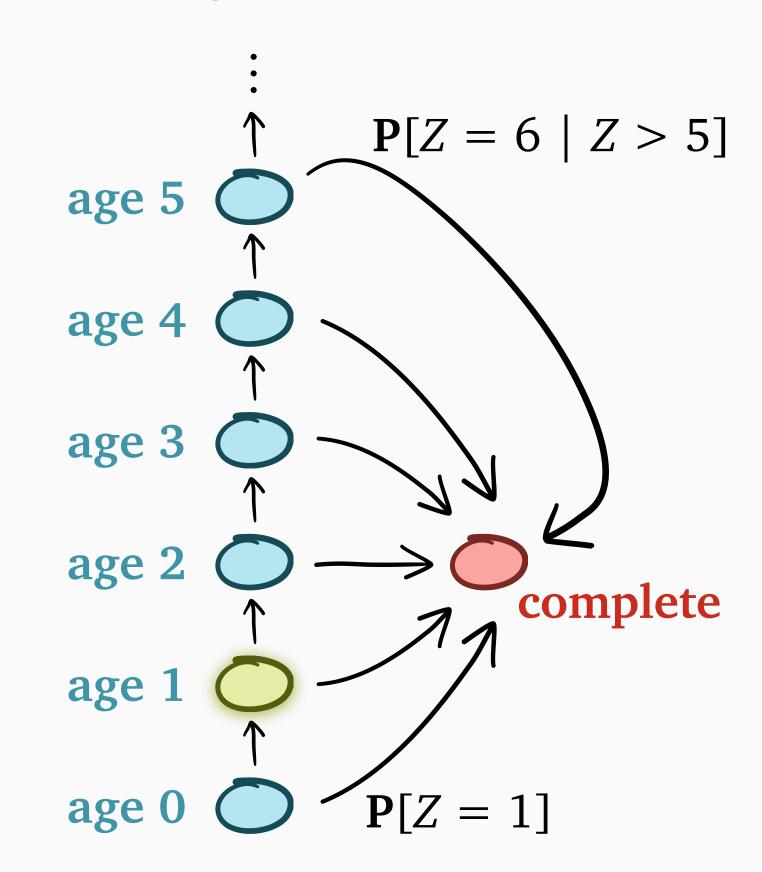


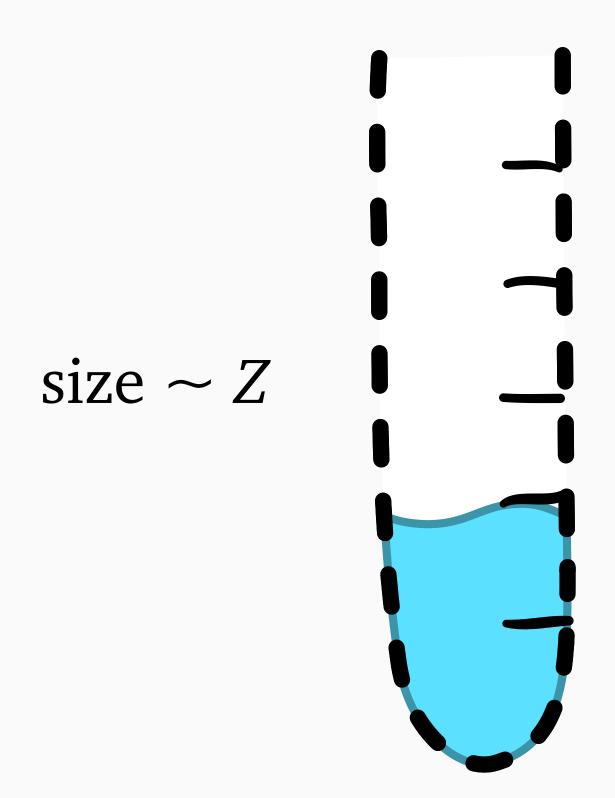


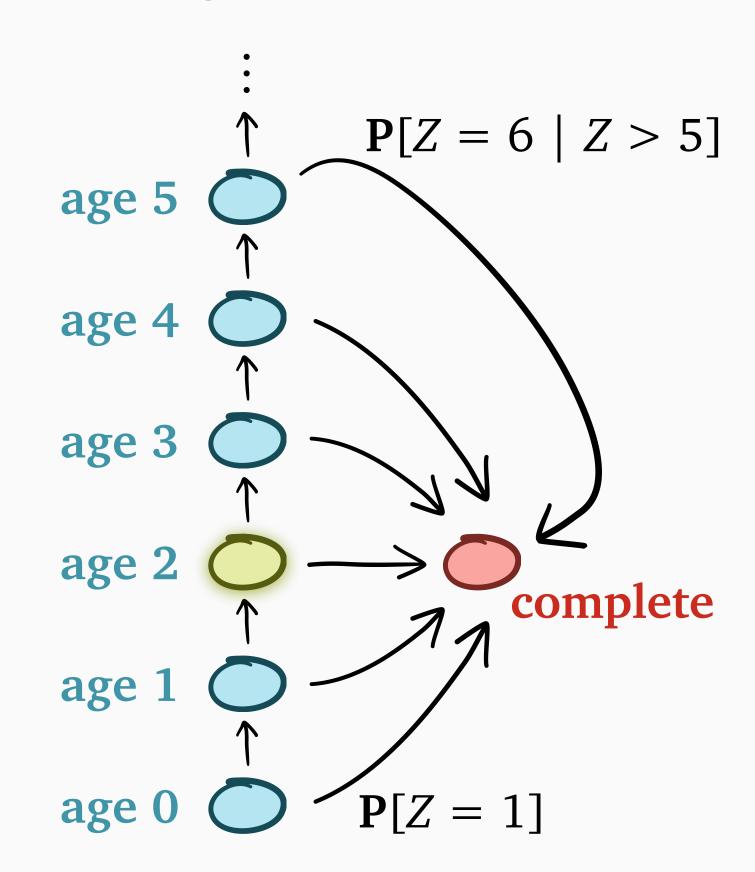


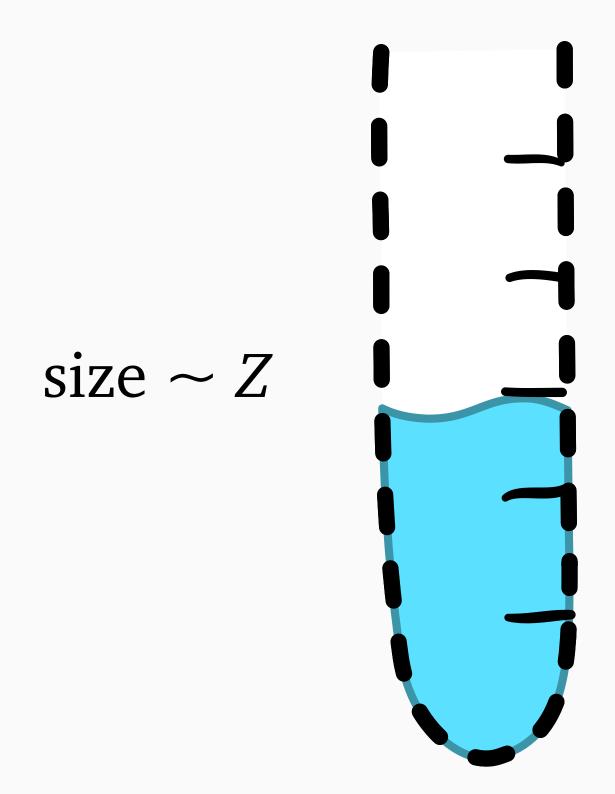


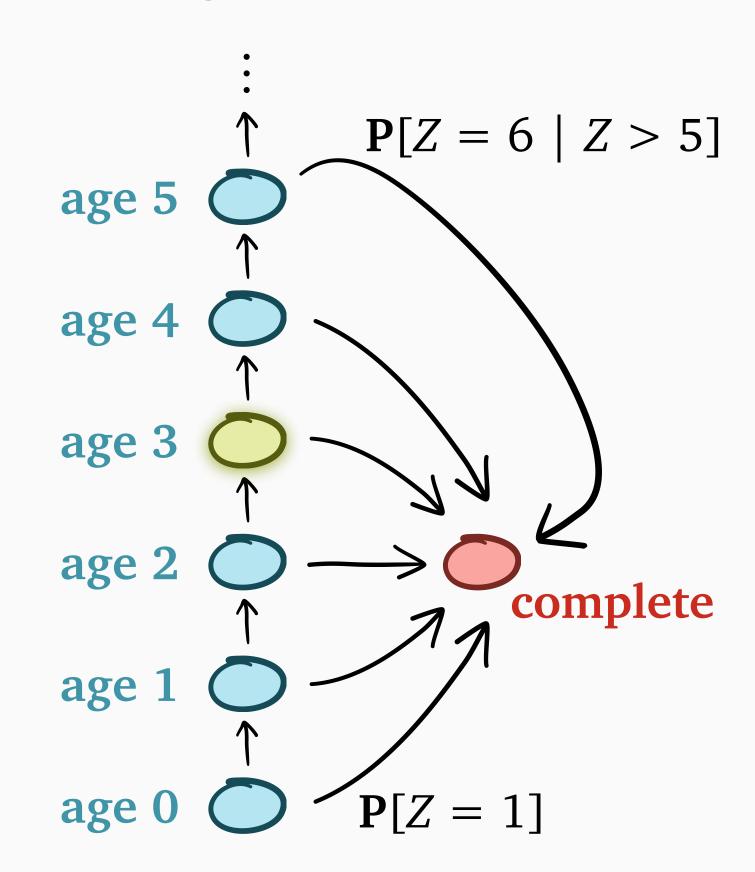


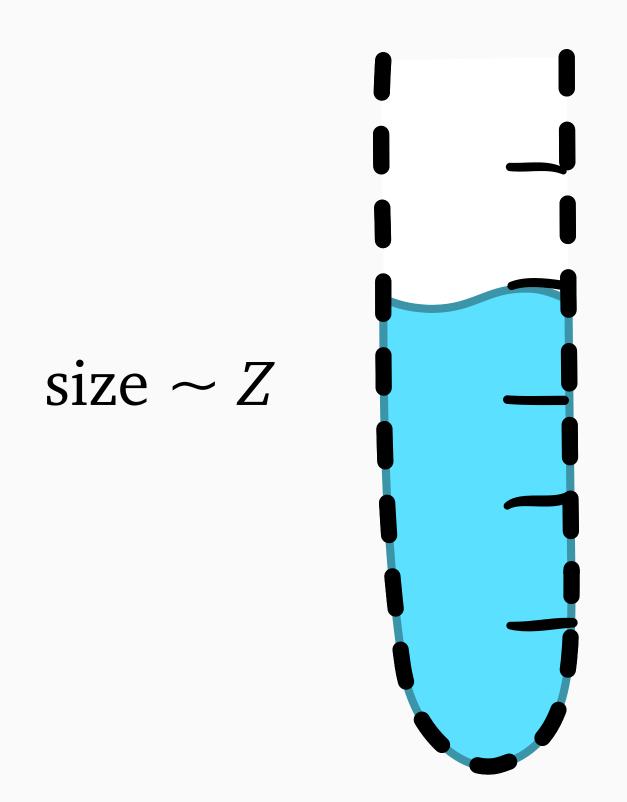


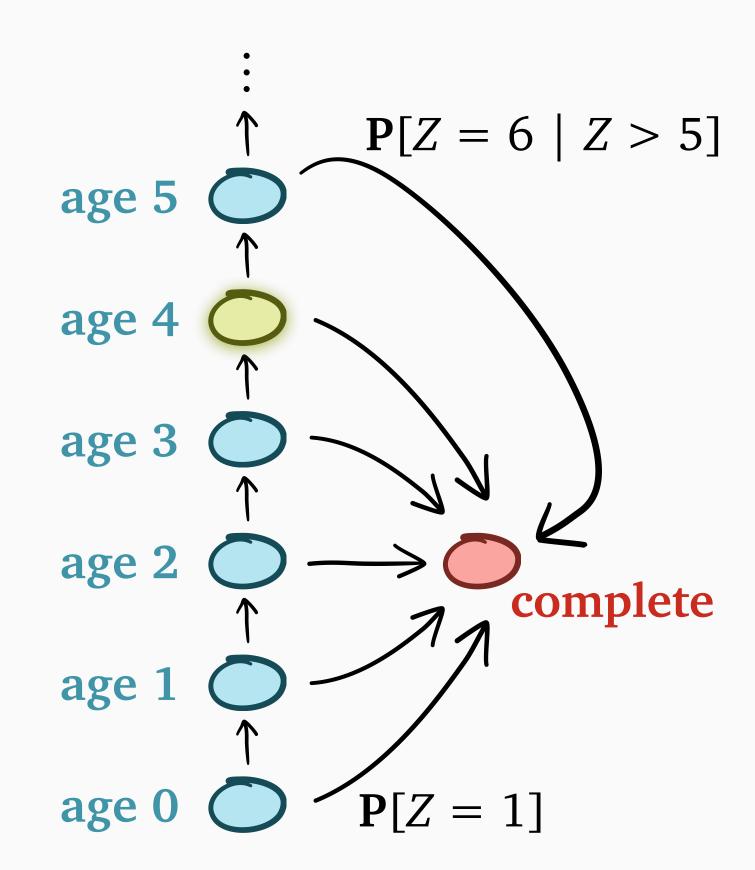


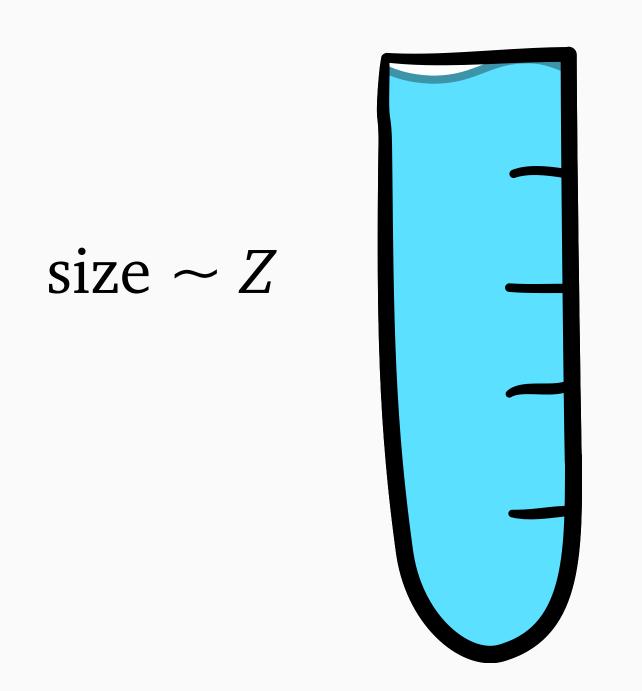


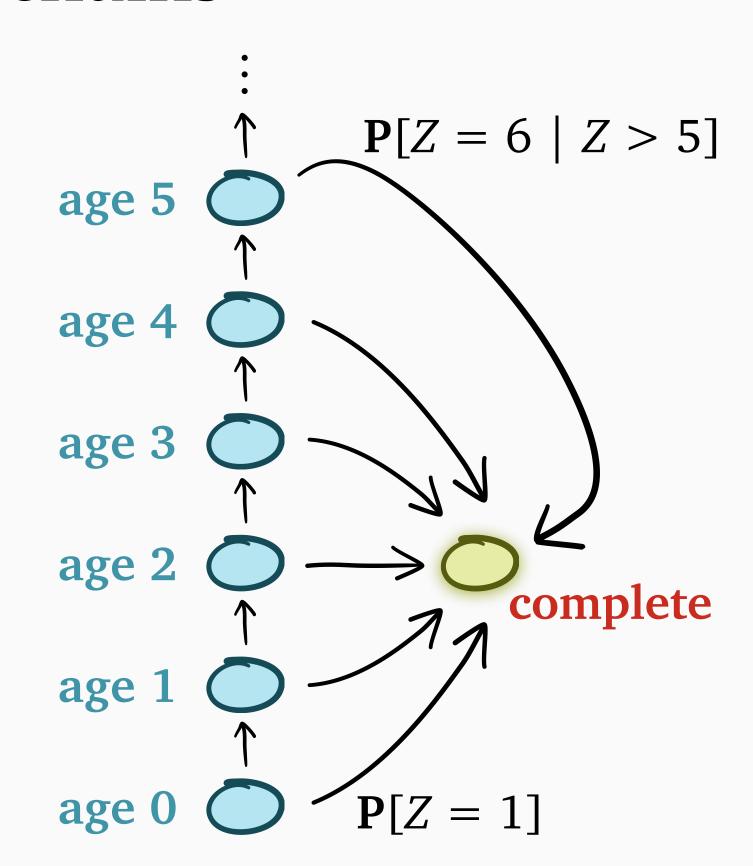


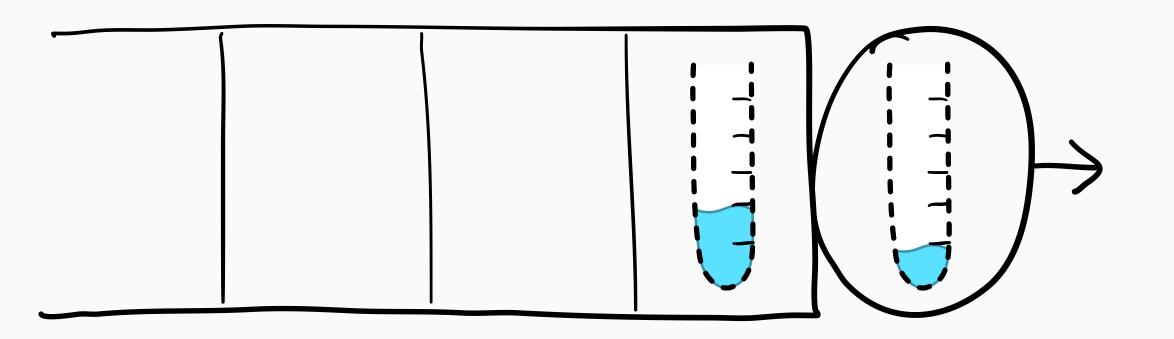


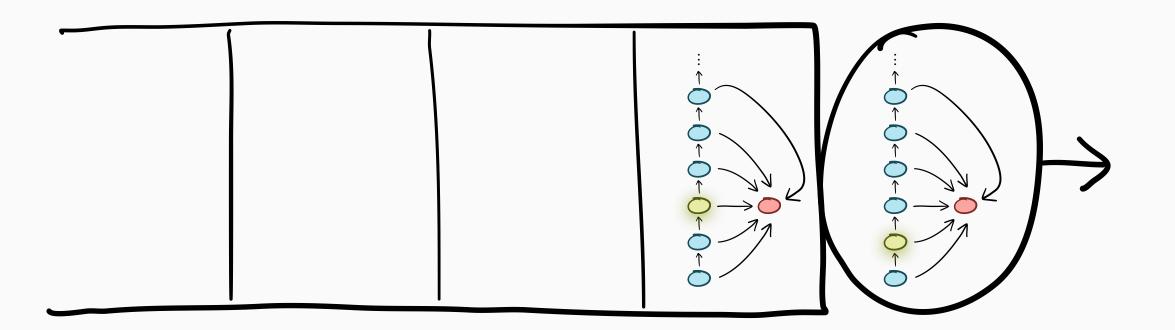


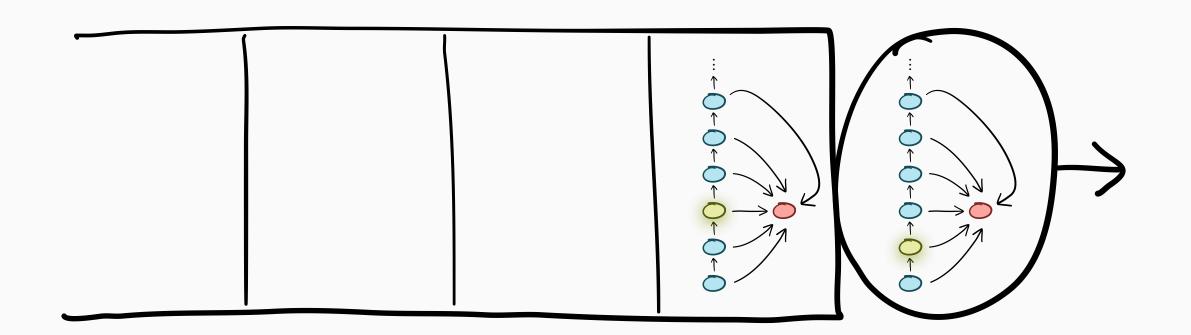


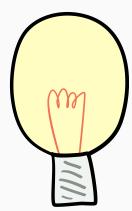


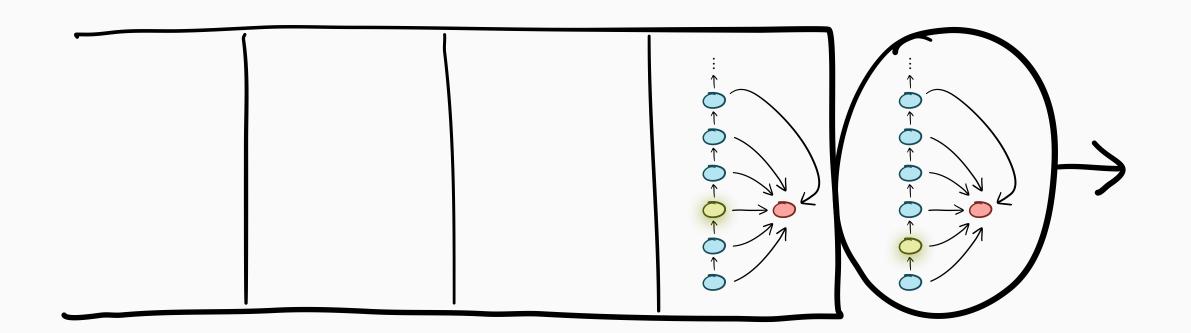






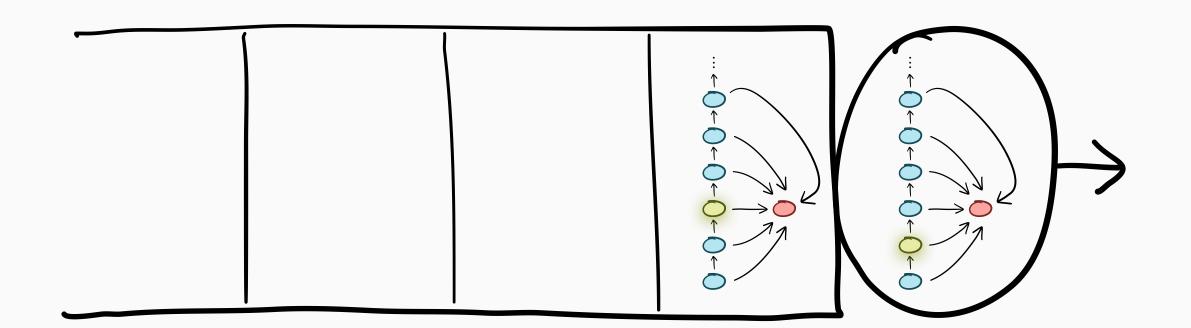


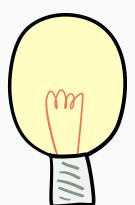




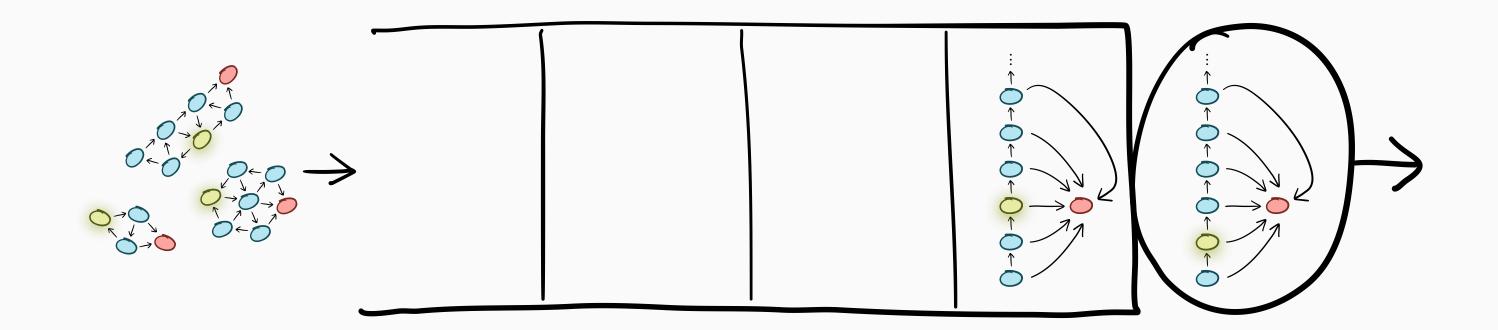


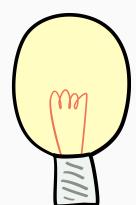
transition probabilities = info about *jobs in general*



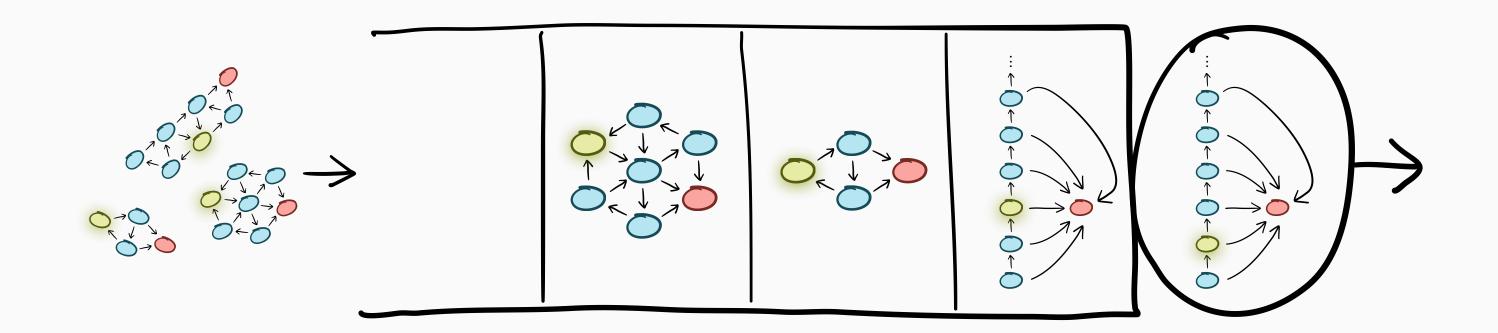


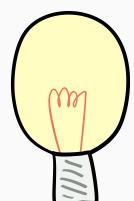
transition probabilities = info about jobs in general current state = info about this specific job



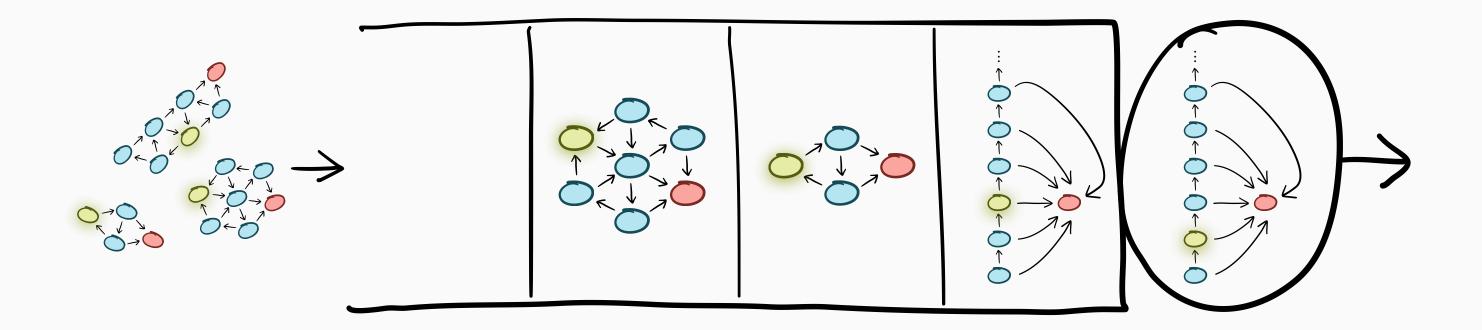


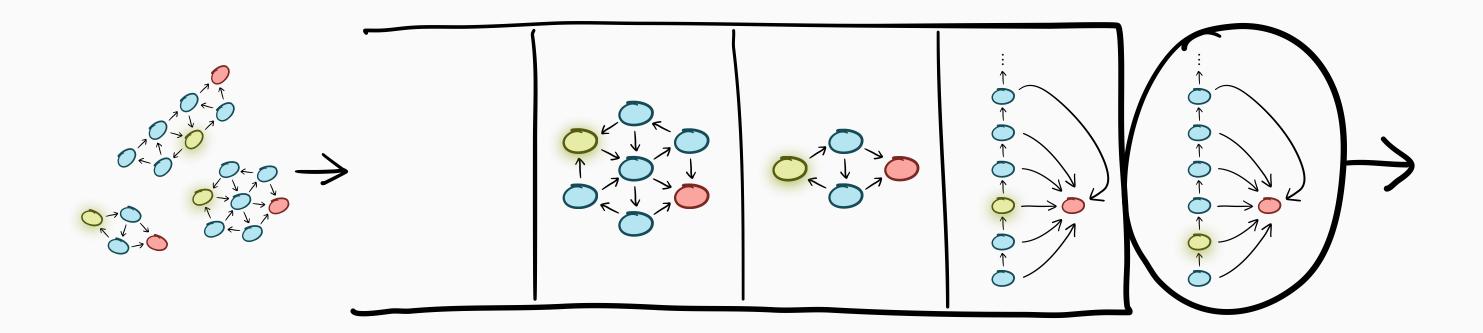
transition probabilities = info about jobs in general current state = info about this specific job



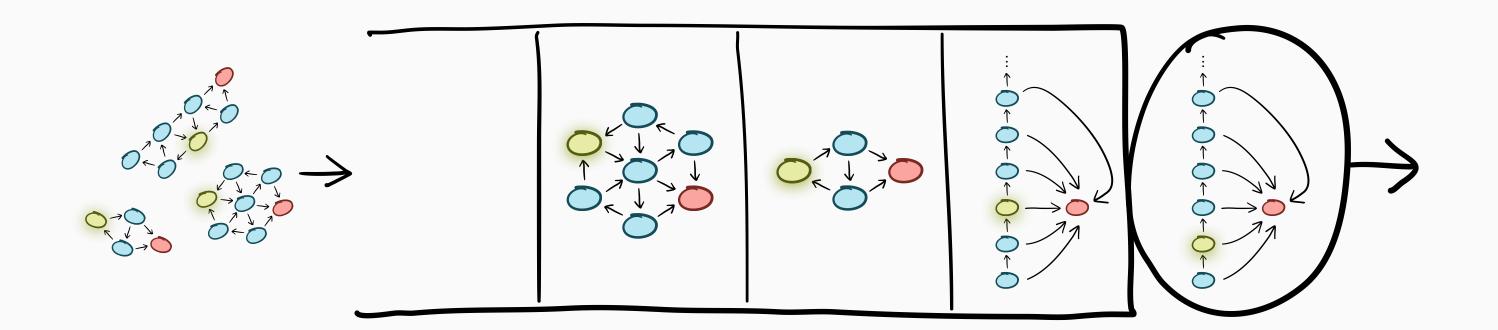


transition probabilities = info about jobs in general current state = info about this specific job



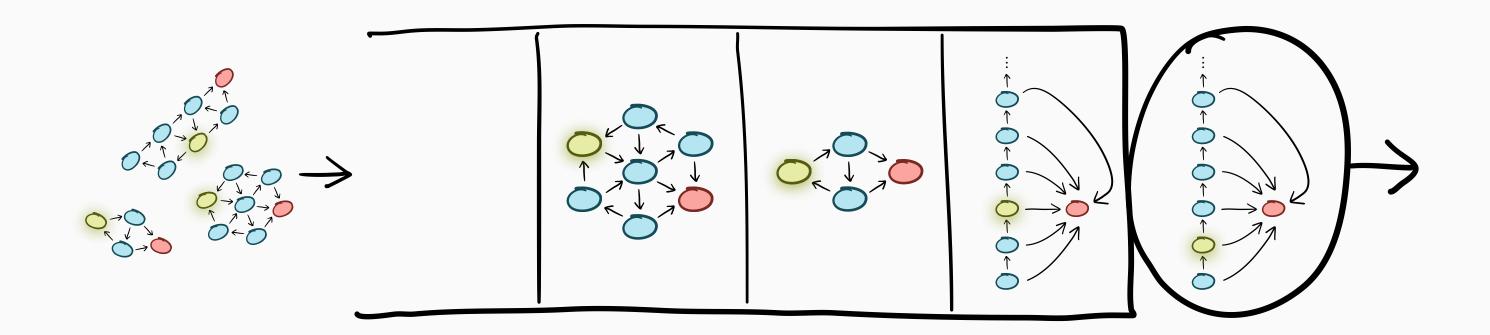








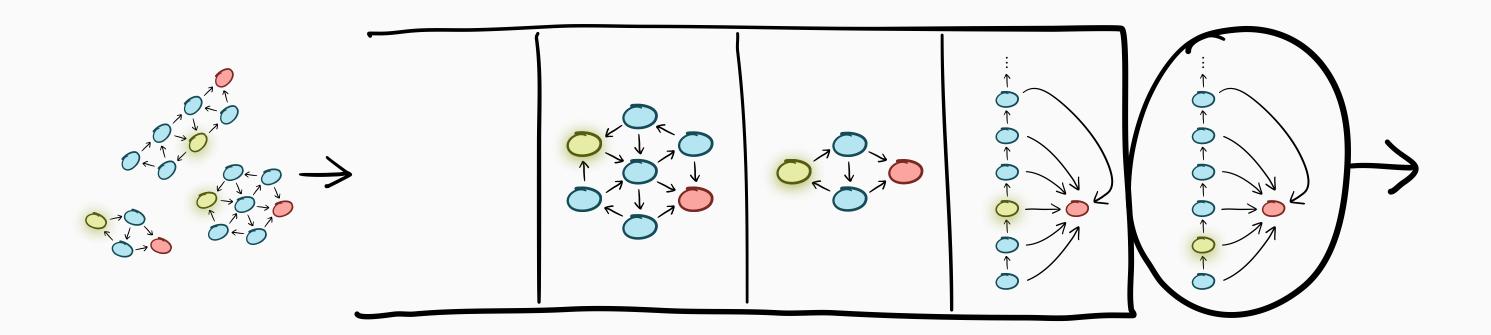




Gittins optimality needs
Poisson arrivals

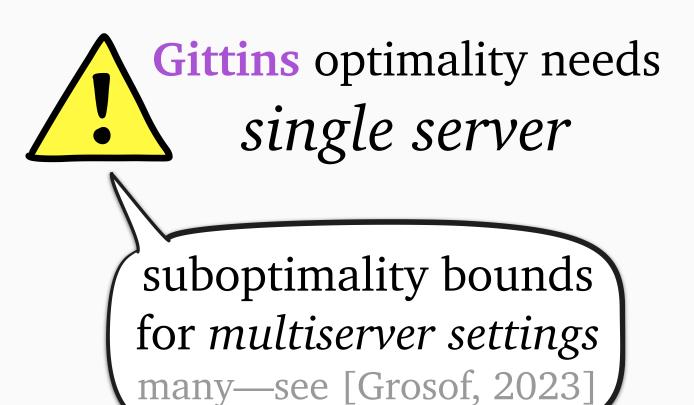
suboptimality bound for *G/G arrivals* [Hong & Scully, 2023]



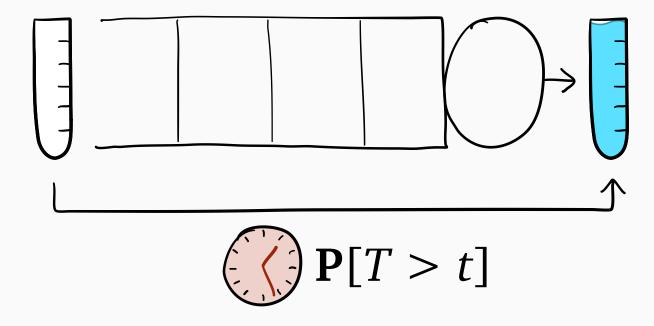


Gittins optimality needs
Poisson arrivals

suboptimality bound for *G/G arrivals* [Hong & Scully, 2023]

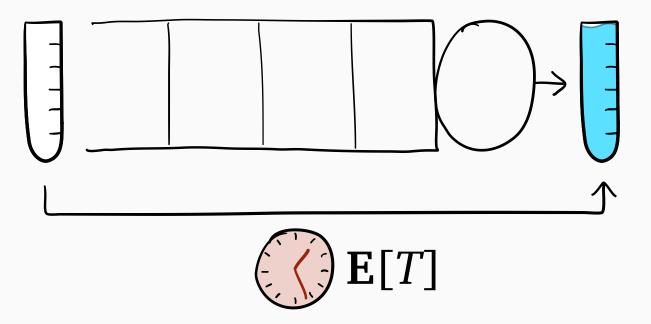


Tail scheduling



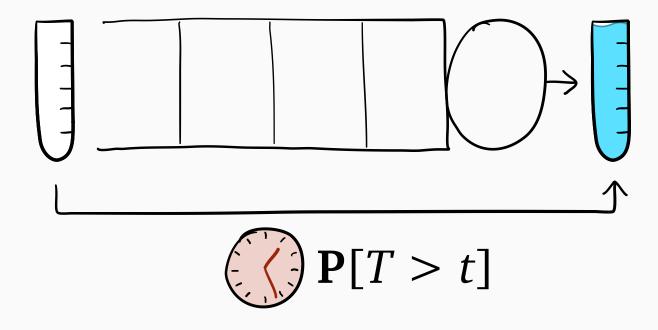
Goal: minimize probability of very long response time

Mean scheduling



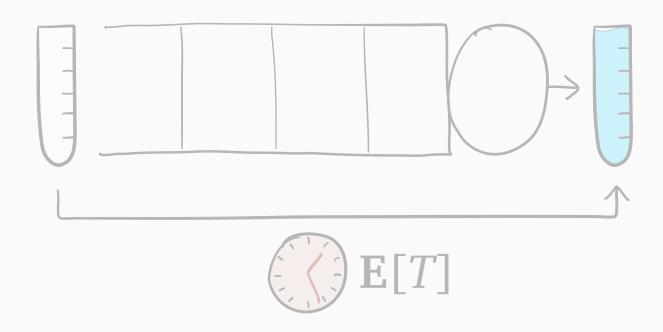
Goal: minimize average response time (or # jobs)

Tail scheduling

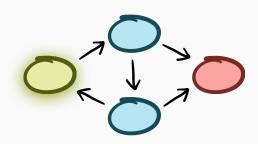


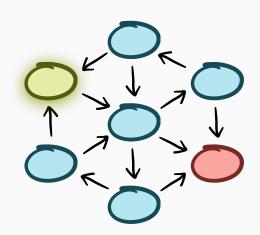
Goal: minimize probability of very long response time

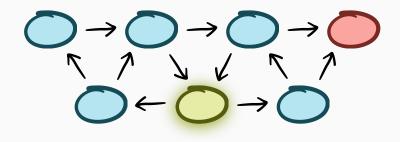




Goal: minimize average response time (or # jobs)

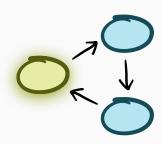


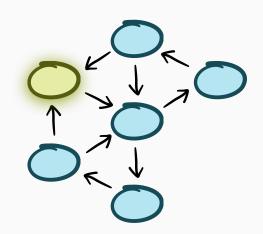


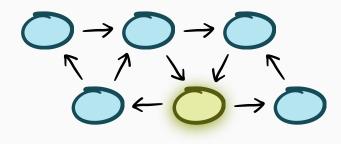


Decision process:

- Advance one Markov chain every step
- Get reward based on its current state
- Stop when one Markov chain terminates



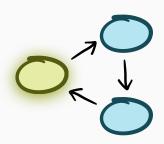


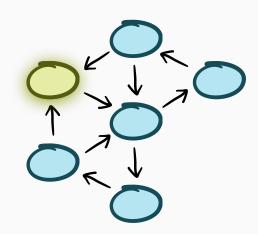


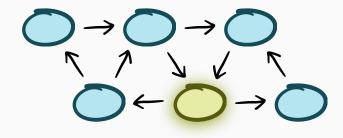
Decision process:

- Advance one Markov chain every step
- Get reward based on its current state





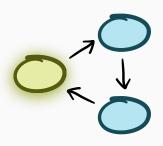


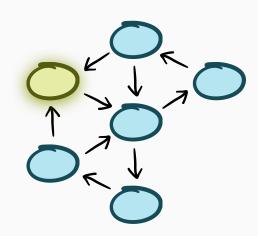


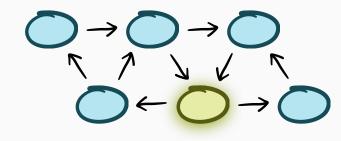
Decision process:

- Advance one Markov chain every step
- Get reward based on its current state
- Stop when one Markov chain terminates

Goal: maximize **E**[**discounted** reward]?







Decision process:

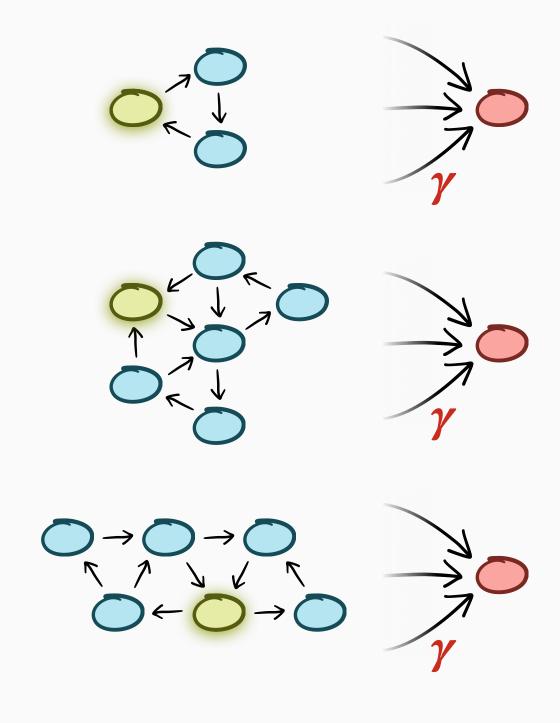
- Advance one Markov chain every step
- Get reward based on its current state
- Stop when one Markov chain terminates

Goal: maximize **E**[discounted reward]?



Use constant termination probability

 γ = **discount** factor



Decision process:

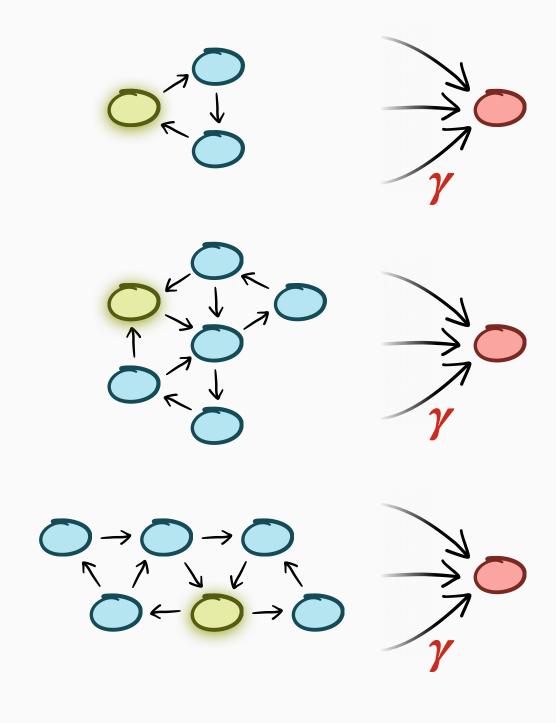
- Advance one Markov chain every step
- Get reward based on its current state
- Stop when one Markov chain terminates

Goal: maximize **E**[discounted reward]?



Use constant termination probability

 γ = discount factor



Decision process:

- Advance one Markov chain every step
- Get reward based on its current state
- Stop when one Markov chain terminates

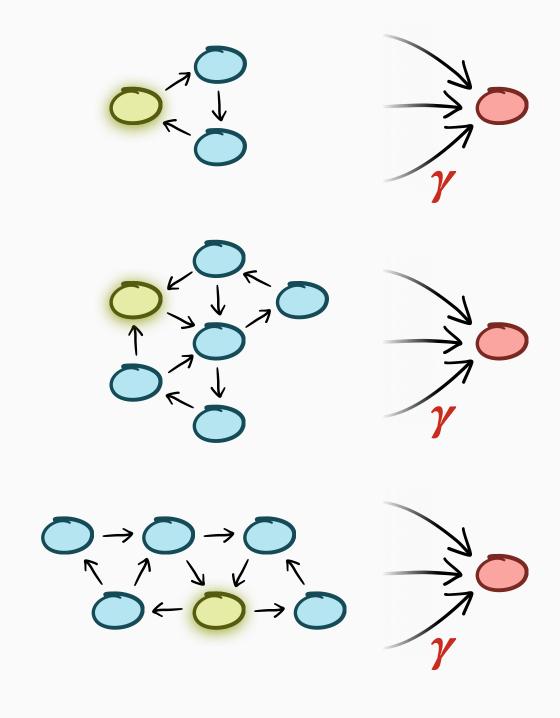
Goal: maximize **E**[discounted reward]?





Use constant termination probability

 γ = discount factor



Decision process:

- Advance one Markov chain every step
- Get reward based on its current state

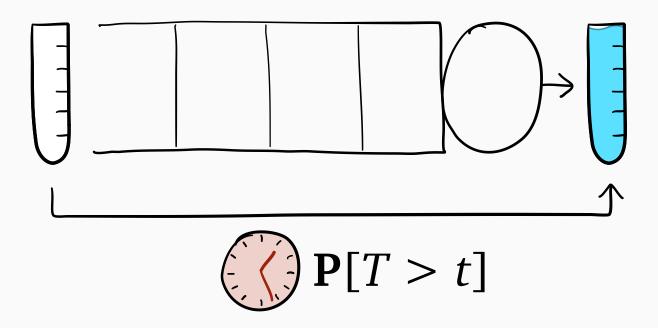


Goal: maximize **E**[discounted reward]?



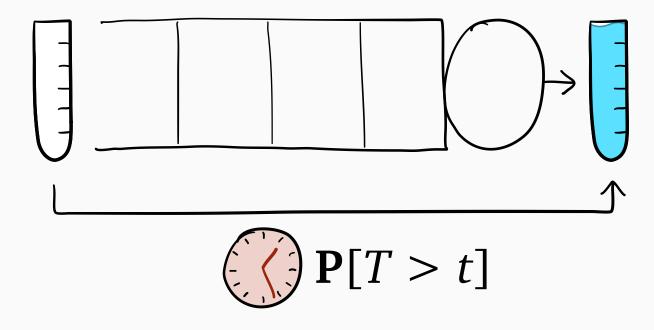


Tail scheduling



Goal: minimize probability of very long response time

Tail scheduling



Goal: minimize probability of very long response time

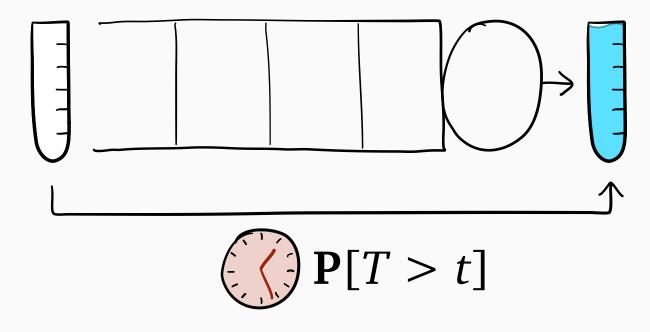
Moment generating function fact:

$$\mathbf{P}[T > t] \sim Ce^{-\theta t} \quad (t \to \infty)$$

$$\updownarrow$$

$$\mathbf{E}[e^{(1-\varepsilon)\theta T}] \sim \frac{C}{\varepsilon} \quad (\varepsilon \to 0)$$

Tail scheduling



Goal: minimize probability of very long response time

Moment generating function fact:

$$\mathbf{P}[T > t] \sim Ce^{-\theta t} \quad (t \to \infty)$$

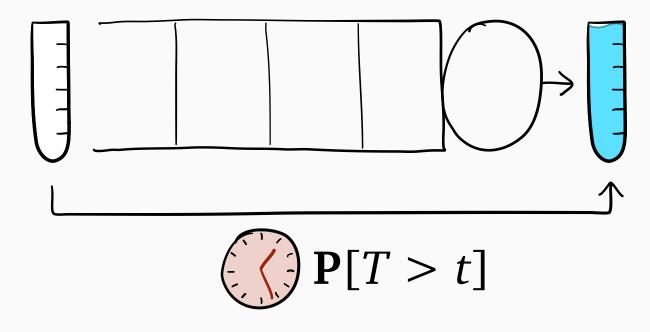
$$\updownarrow$$

$$\mathbf{E}[e^{(1-\varepsilon)\theta T}] \sim \frac{C}{\varepsilon} \quad (\varepsilon \to 0)$$



On completion, job pays cost 1, which *inflates* during time *T*

Tail scheduling



Goal: minimize probability of very long response time

Moment generating function fact:

$$\mathbf{P}[T > t] \sim Ce^{-\theta t} \quad (t \to \infty)$$

$$\updownarrow$$

$$\mathbf{E}[e^{(1-\varepsilon)\theta T}] \sim \frac{C}{\varepsilon} \quad (\varepsilon \to 0)$$



On completion, job pays cost 1, which *inflates* during time *T*



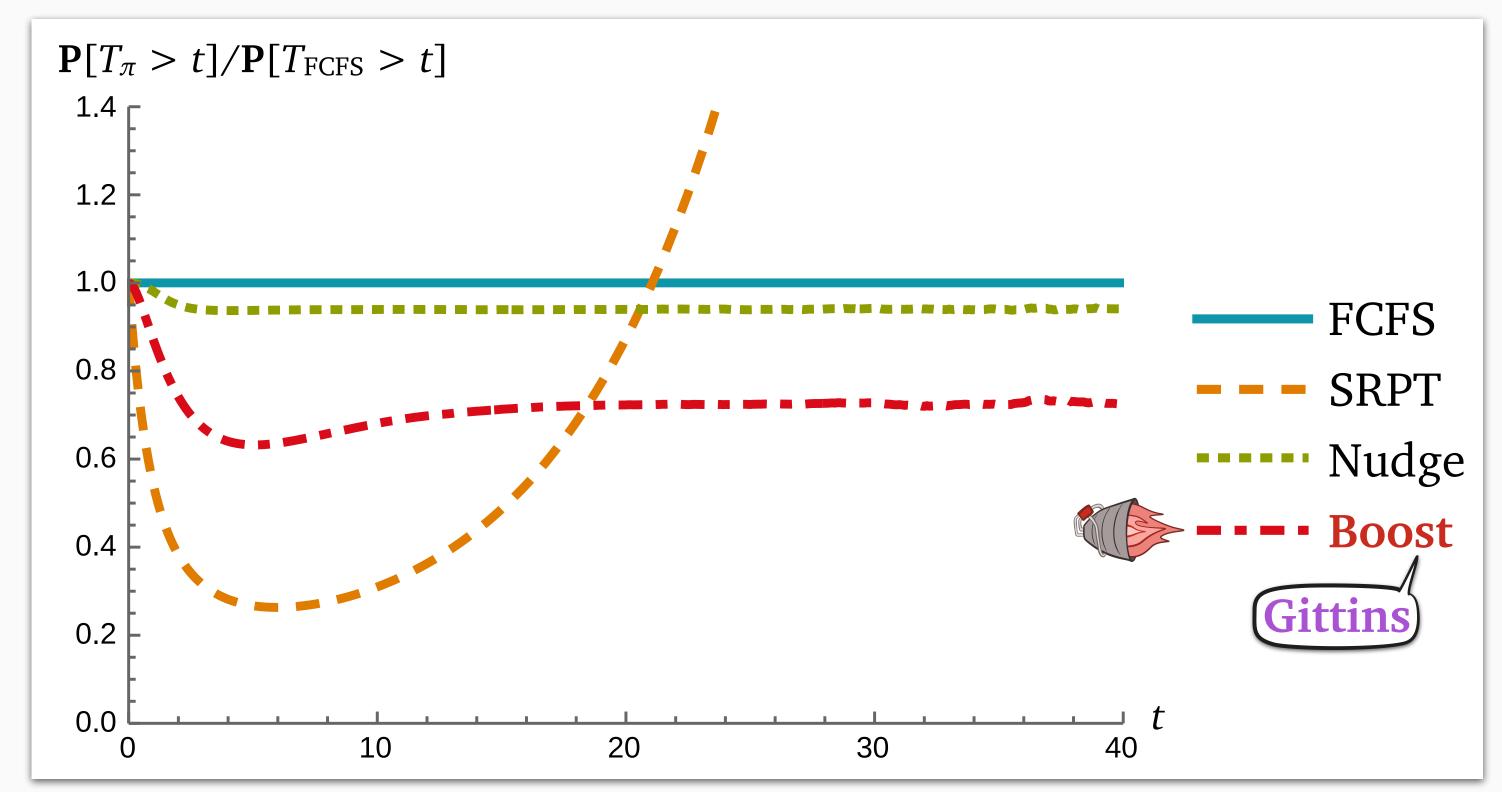
Empirical results: tail scheduling



George Yu



Amit Harley



See our SIGMETRICS papers [Yu & Scully, 2024; Harlev et al., 2025; Yu et al, 2026]



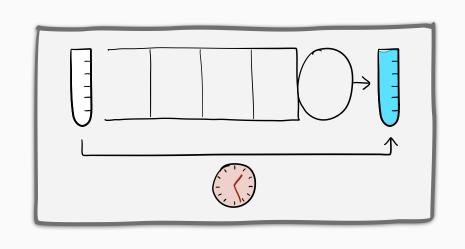
Part I:

Introducing Gittins via Pandora's box



Part II:

Gittins for Bayesian optimization



Part III:

Gittins for tail scheduling and more



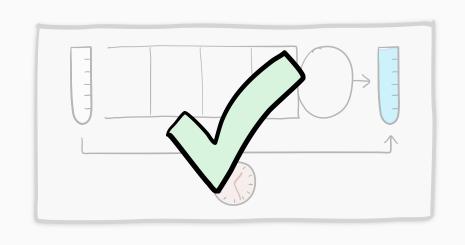
Part I:

Introducing Gittins via Pandora's box



Part II:

Gittins for Bayesian optimization



Part III:

Gittins for tail scheduling and more





• *Computation*, especially for continuous state spaces



- *Computation*, especially for continuous state spaces
- Multiple actions within each process, i.e. MDP selection (a.k.a. bandit superprocesses)



- *Computation*, especially for continuous state spaces
- Multiple actions within each process, i.e. MDP selection (a.k.a. bandit superprocesses)

e.g. **Pandora's box** with optional inspection



- *Computation*, especially for continuous state spaces
- Multiple actions within each process,
 i.e. MDP selection (a.k.a. bandit superprocesses)
- "Restless" process state changes occurring even when not actively advanced

e.g. **Pandora's box** with optional inspection



- *Computation*, especially for continuous state spaces
- Multiple actions within each process,
 i.e. MDP selection (a.k.a. bandit superprocesses)
- "Restless" process state changes occurring even when not actively advanced

more generally,
weakly coupled MDPs

e.g. Pandora's box

with optional inspection



- *Computation*, especially for continuous state spaces
- Multiple actions within each process, i.e. MDP selection (a.k.a. bandit superprocesses)
- "Restless" process state changes occurring even when not actively advanced
- *Model inference*, or another path to a data-driven version of **Gittins**

e.g. **Pandora's box** with optional inspection

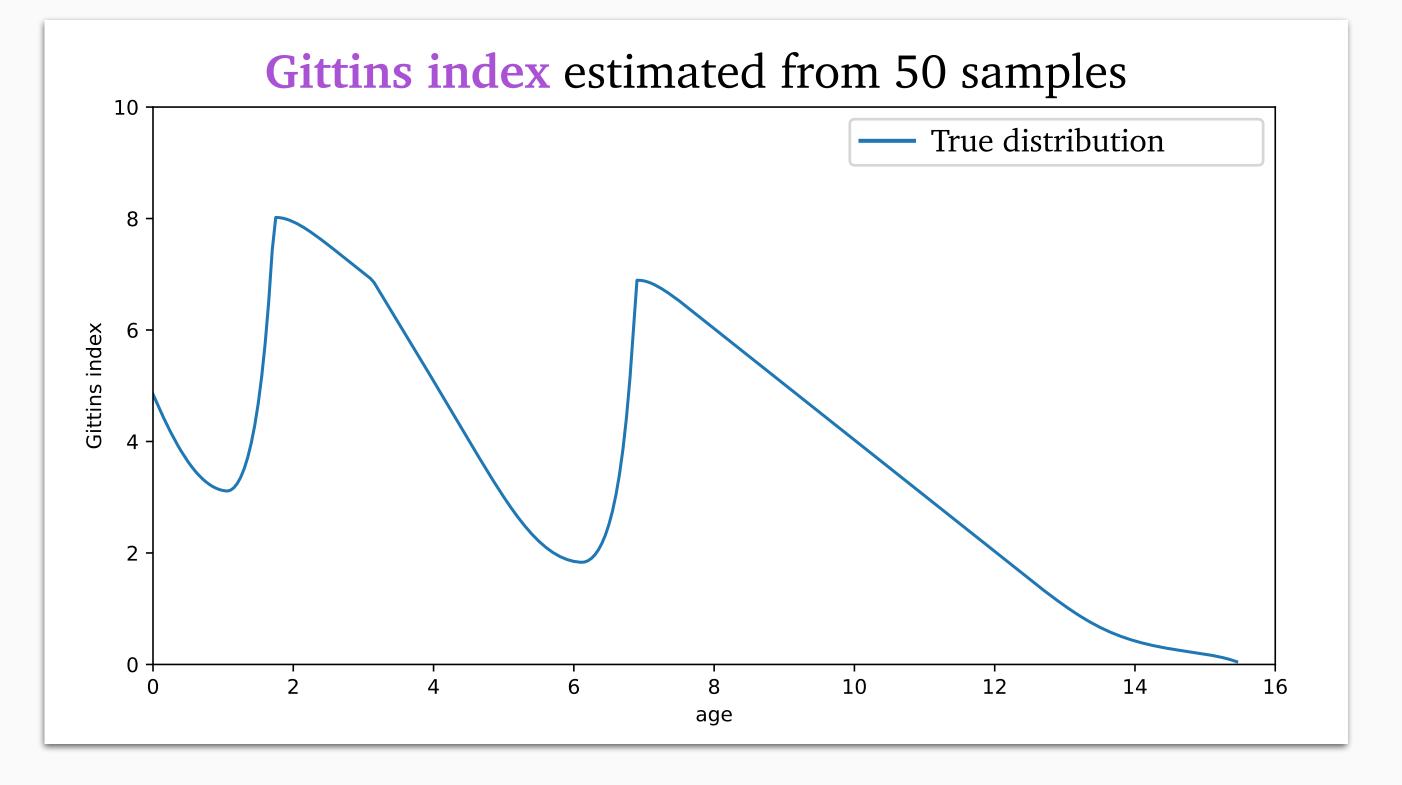
more generally, weakly coupled MDPs



Shefali Ramakrishna



Amit Harlev

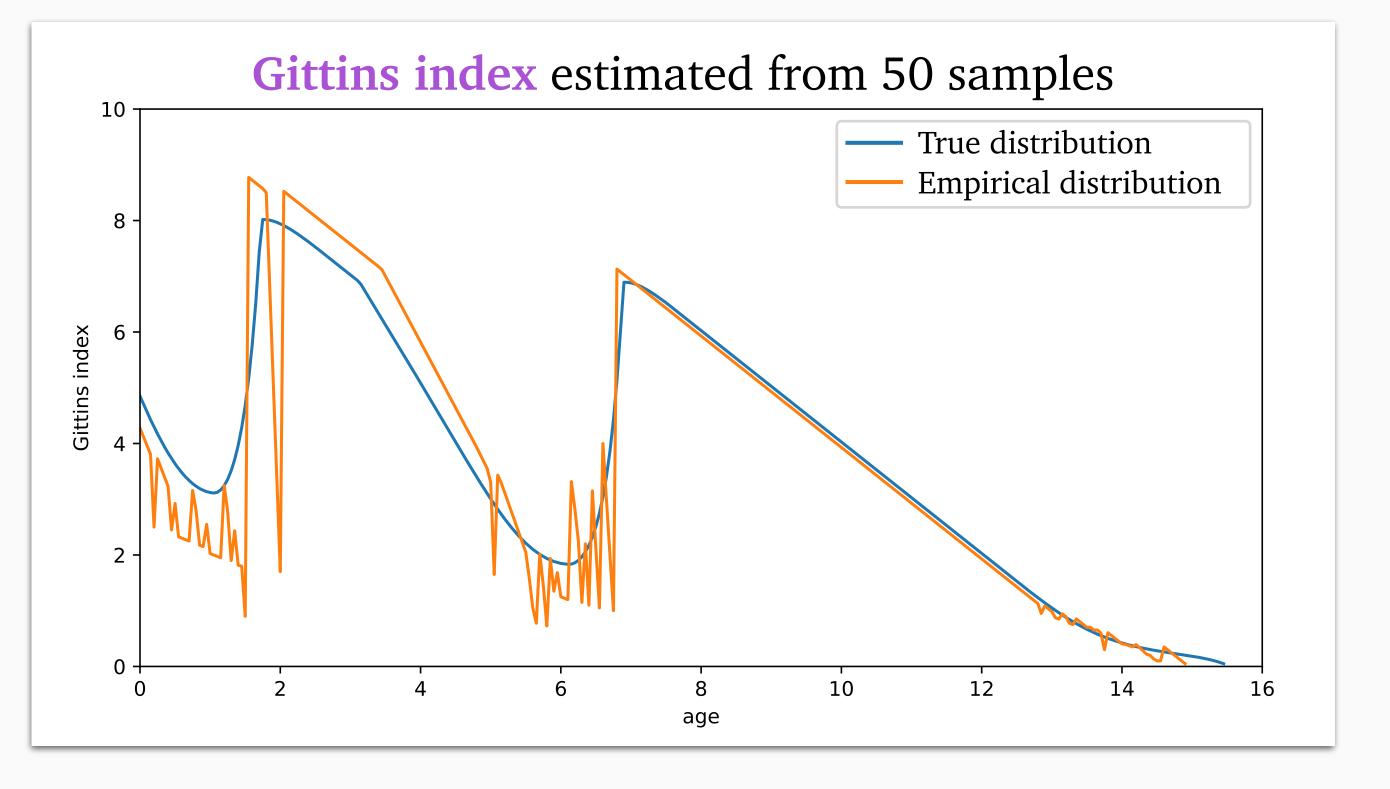




Shefali Ramakrishna



Amit Harlev

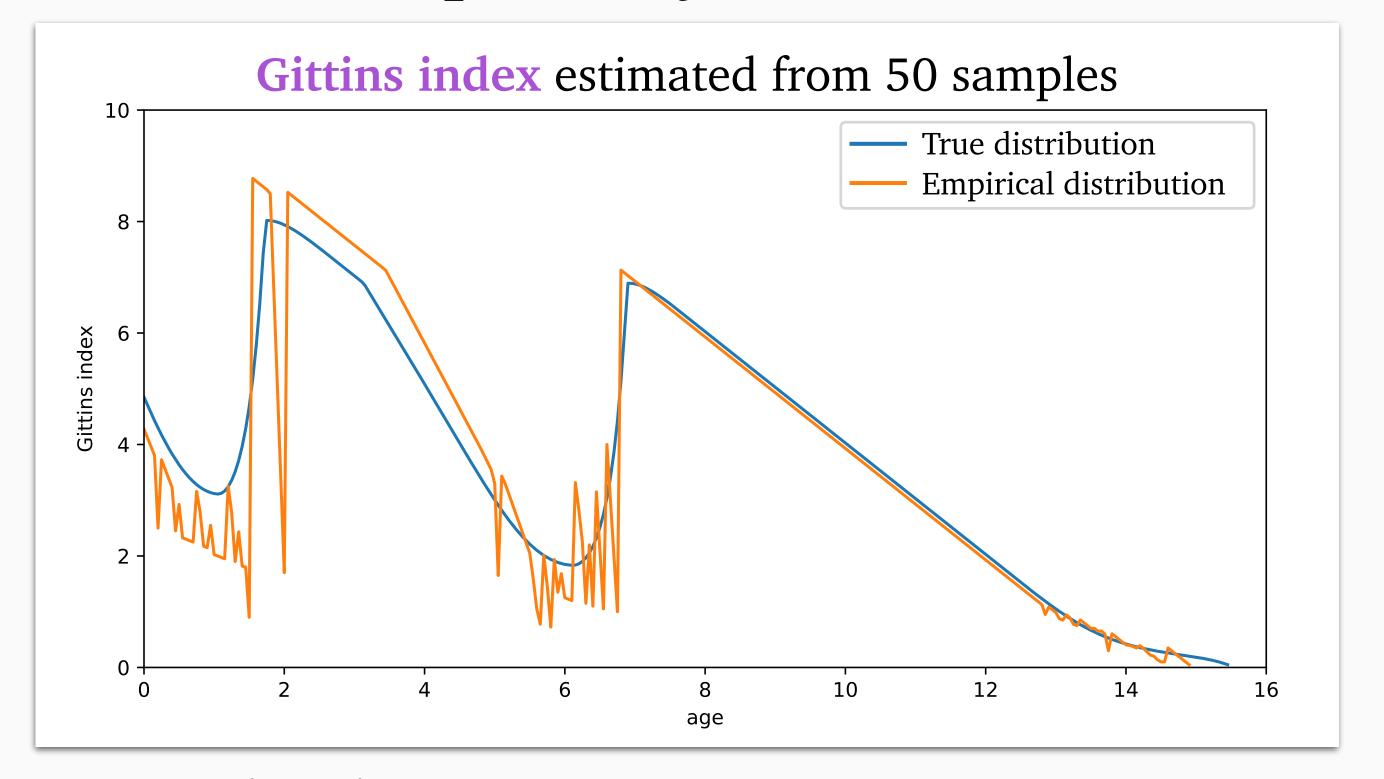




Shefali Ramakrishna



Amit Harlev



Discrete sizes: empirical works

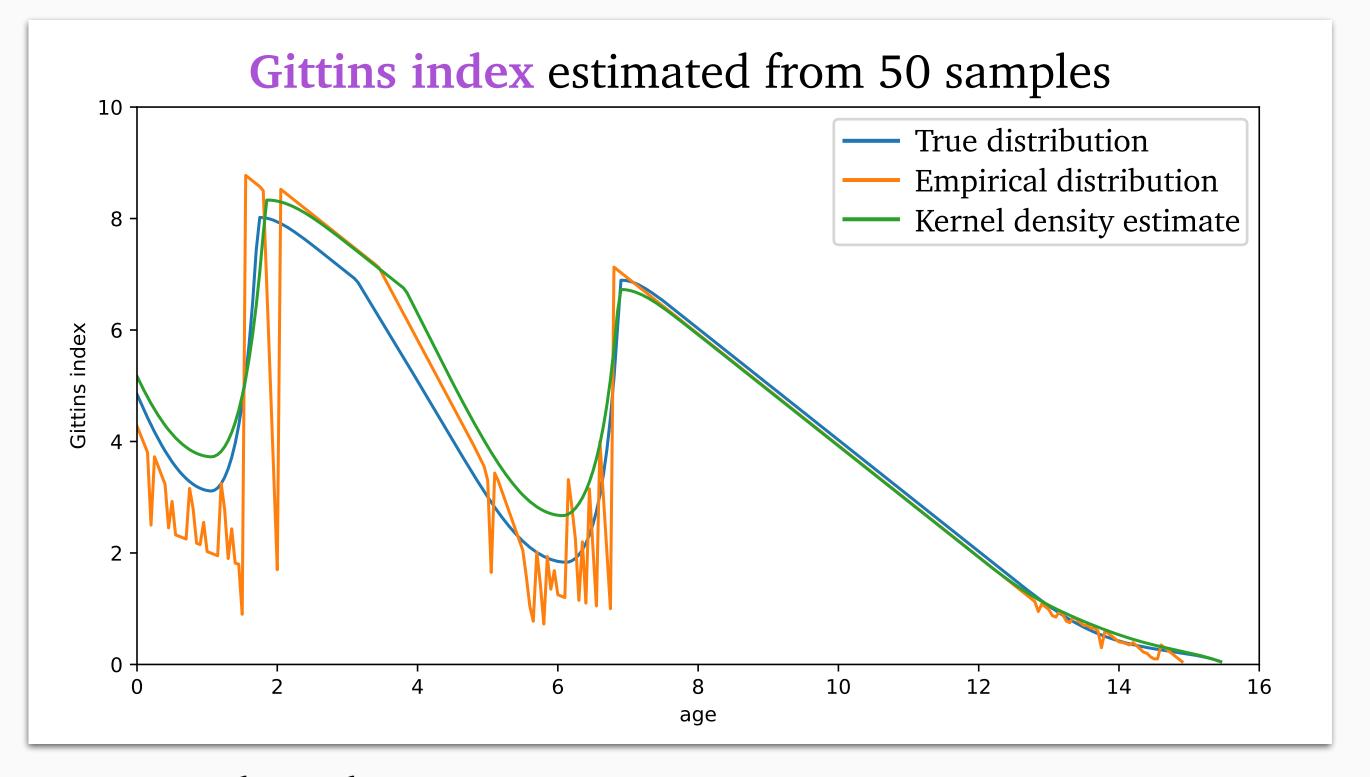
[Ramakrishna et al., 2024]



Shefali Ramakrishna



Amit Harlev



Discrete sizes: empirical works

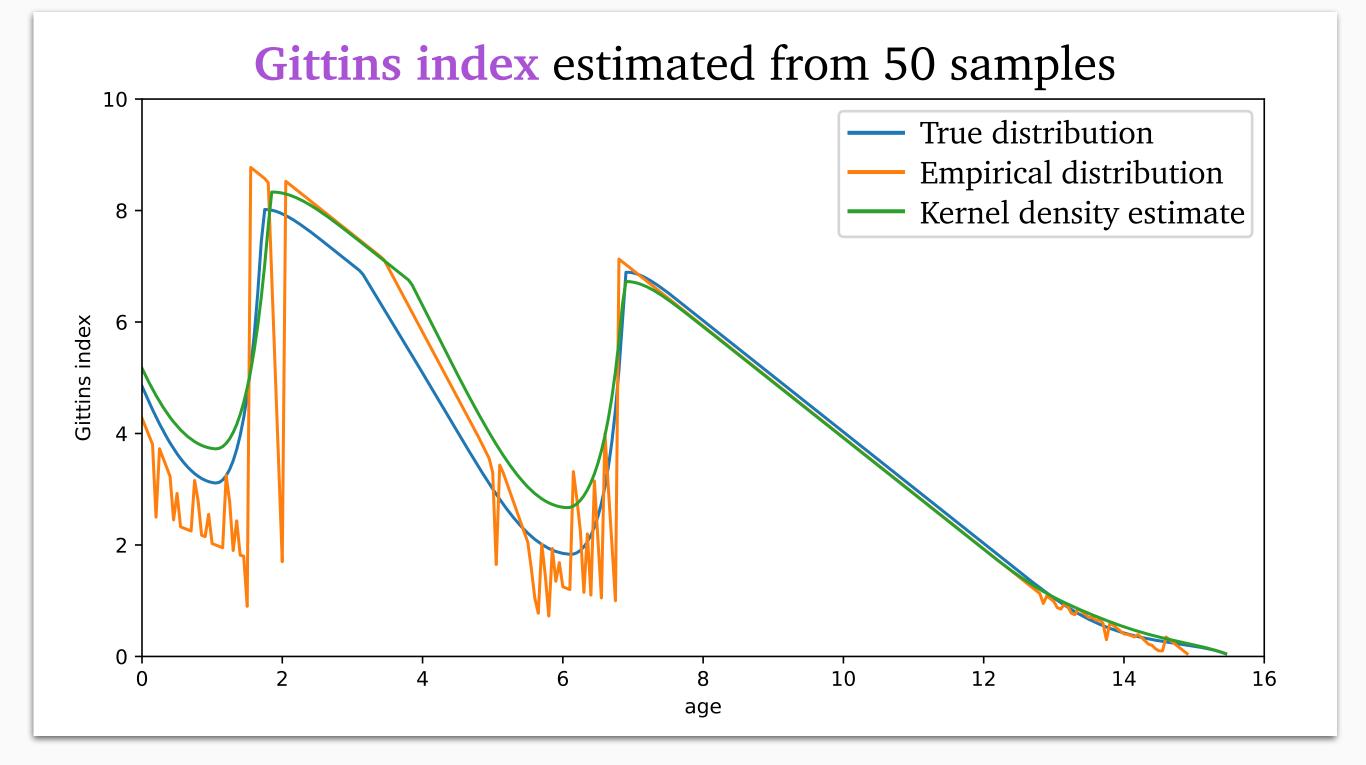
[Ramakrishna et al., 2024]



Shefali Ramakrishna



Amit Harlev



Discrete sizes: empirical works

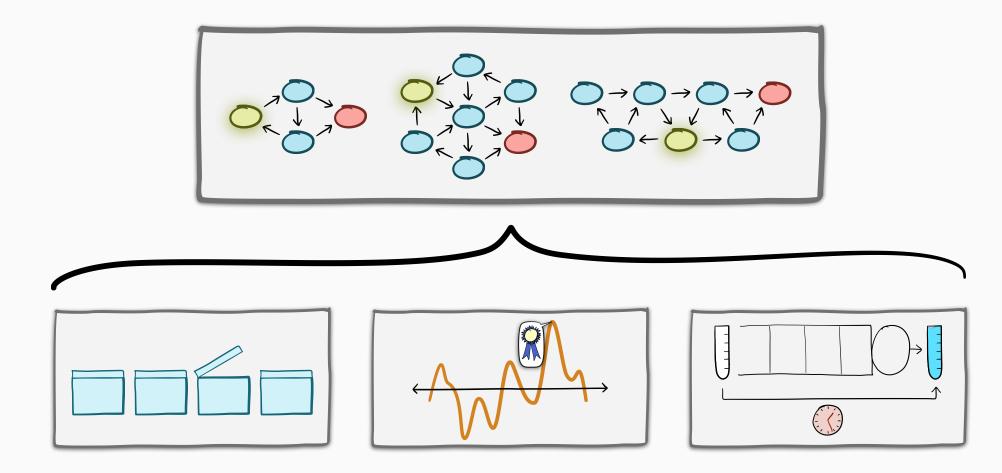
[Ramakrishna et al., 2024]

Continuous sizes: empirical still works

[Ramakrishna et al., in submission]

The Gittins Index:

A Design Principle for Decision Making Under Uncertainty



Ziv Scully
Cornell ORIE
https://ziv.codes

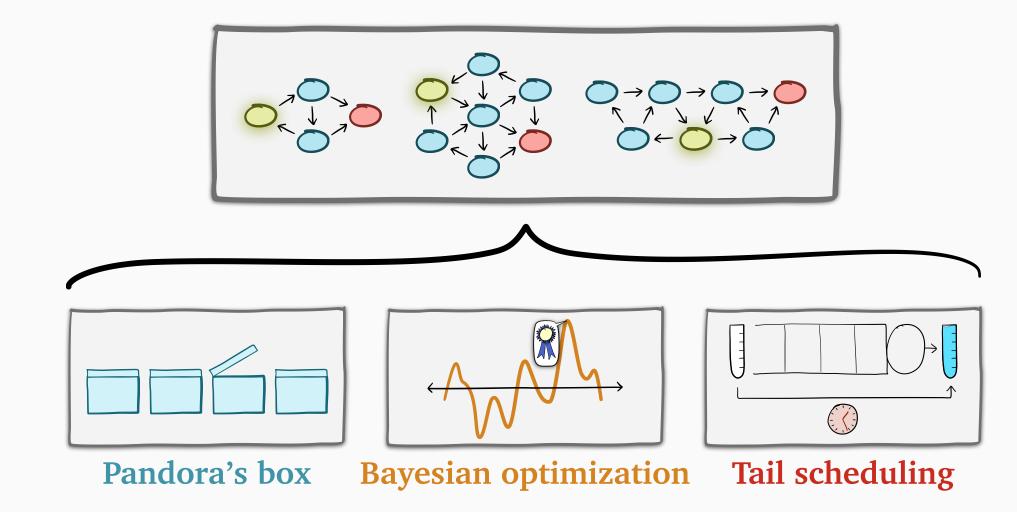
Alex Terenin

Cornell CDSES

https://avt.im

The Gittins Index:

A Design Principle for Decision Making Under Uncertainty



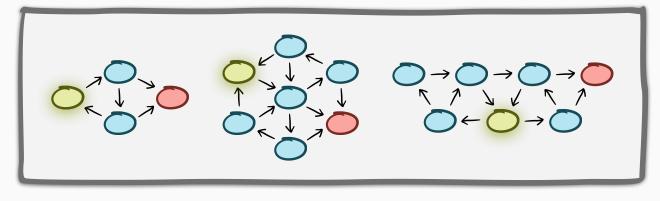
Ziv Scully
Cornell ORIE
https://ziv.codes

Alex Terenin
Cornell CDSES
https://avt.im

The Gittins Index:

A Design Principle for Decision Making Under Uncertainty

Markov chain selection

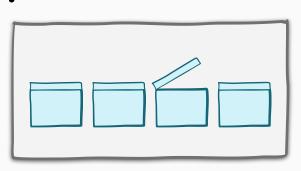


Ziv Scully

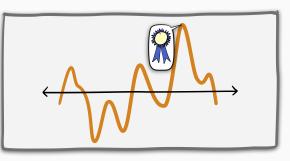
Cornell ORIE
https://ziv.codes

Alex Terenin

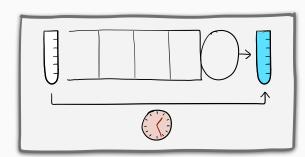
Cornell CDSES
https://avt.im



Pandora's box



Bayesian optimization

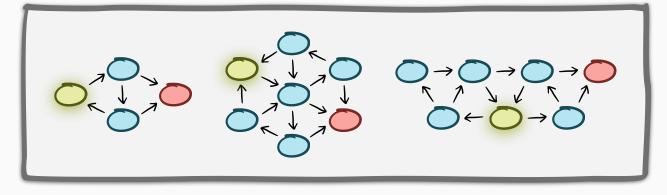


Tail scheduling

The Gittins Index

A Design Principle for Decision Making Under Uncertainty

Markov chain selection

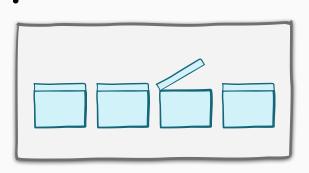


Ziv Scully

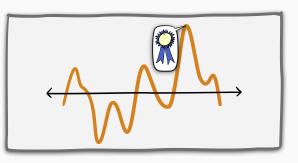
Cornell ORIE
https://ziv.codes

Alex Terenin

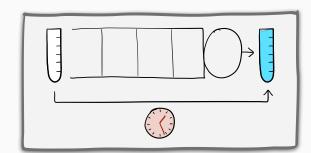
Cornell CDSES
https://avt.im



Pandora's box



Bayesian optimization



evaluate stochastic

options by calibrating

to deterministic options

Tail scheduling

The Gittins Index

A Design Principle for Decision Making Under Uncertainty

Pandora's box

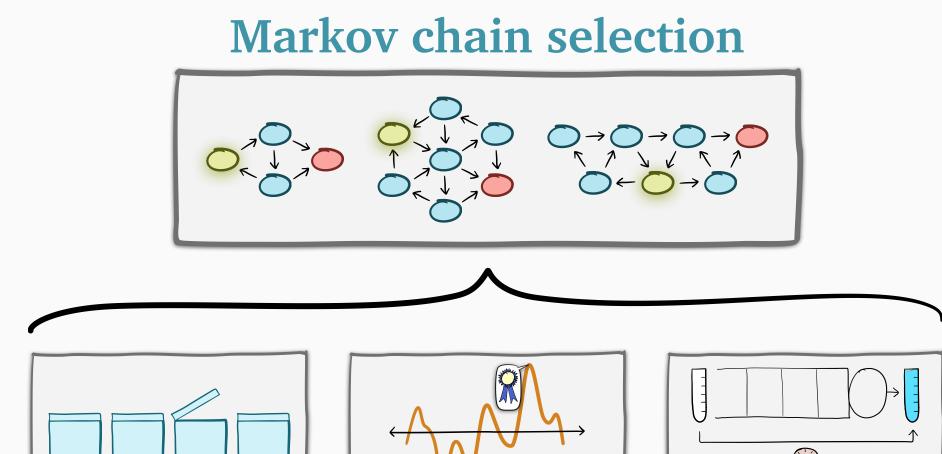


Ziv Scully

Cornell ORIE
https://ziv.codes

Alex Terenin

Cornell CDSES
https://avt.im



Bayesian optimization

evaluate stochastic

Tail scheduling

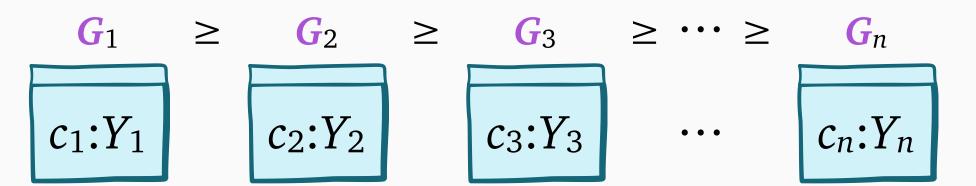
options by calibrating

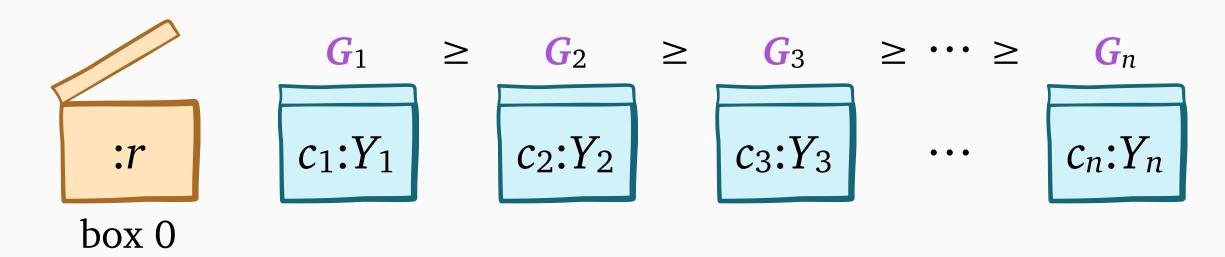
to deterministic options

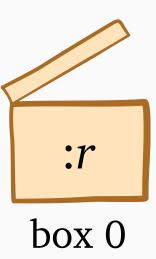
Bonus slides

- Optimality proof sketch
- Quizzes
- Pandora's box with optional inspection

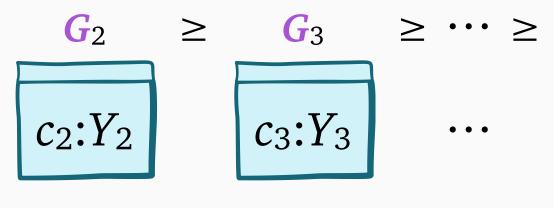
Optimality proof sketch

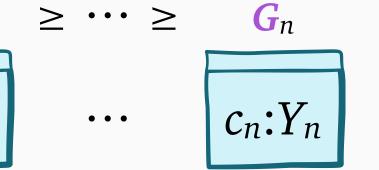




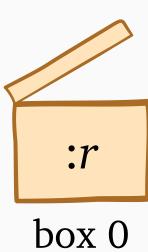


$$c_1:Y_1$$





•
$$r \geq G_1$$
:



 $c_1:Y_1$

 $G_1 \geq$

 $c_2:Y_2$

 G_2

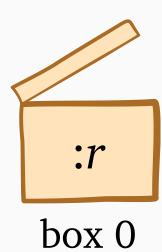
 $c_3:Y_3$

 \geq $G_3 \geq \cdots \geq$

 $c_n:Y_r$

 G_n

• $r \geq G_1$:



 $G_1 \geq c_1:Y_1$

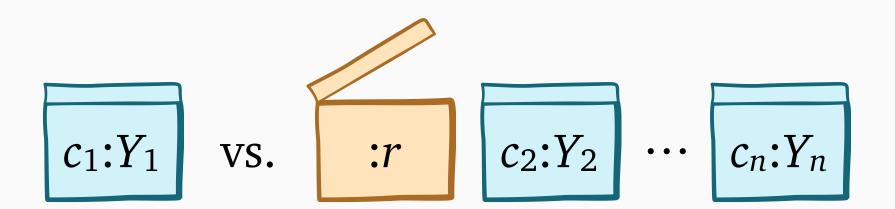
 $c_2:Y_2$

 G_2

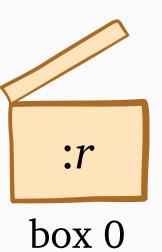
 $c_3:Y_3$

 \geq $G_3 \geq \cdots \geq$

 $c_n:Y_n$



• $r \ge G_1$:



 $G_1 \geq c_1:Y_1$

 $c_2:Y_2$

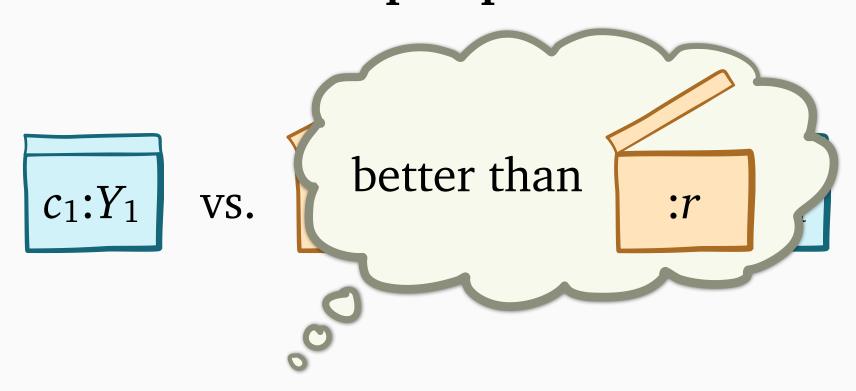
 G_2

 $c_3:Y_3$

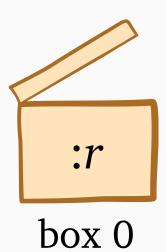
 \geq G_3 \geq \cdots \geq

 $c_n:Y_n$

 G_n



• $r \ge G_1$:



 $c_1:Y_1$

 $G_1 \geq$

 $c_2:Y_2$

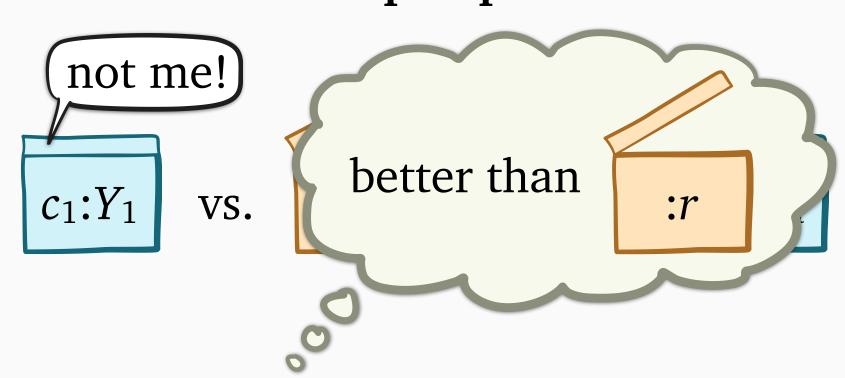
 G_2

 $c_3:Y_3$

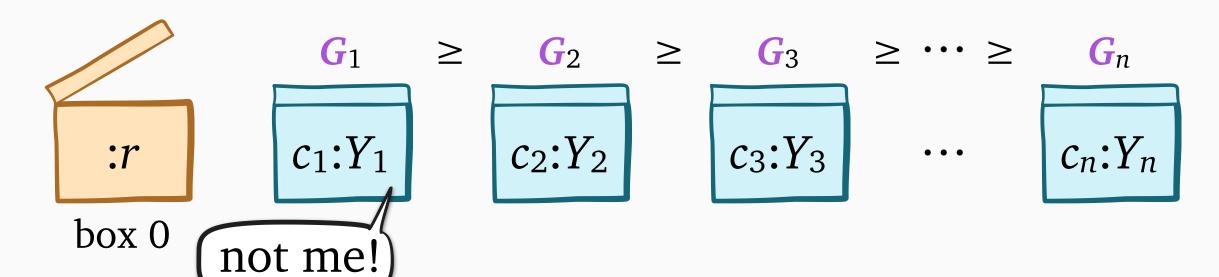
 \geq $G_3 \geq \cdots \geq$

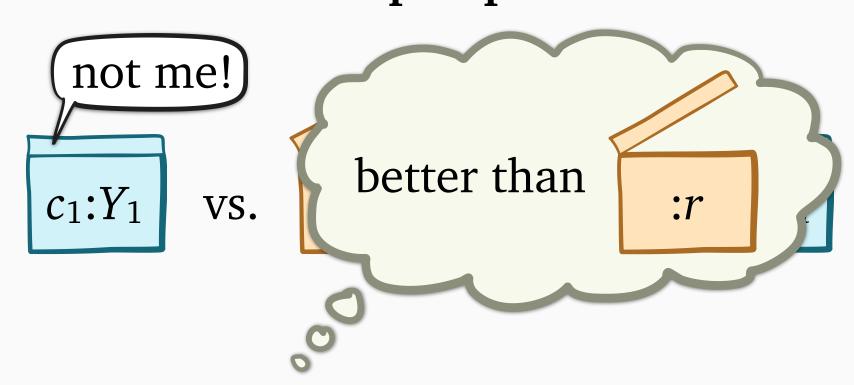
 $c_n:Y_n$

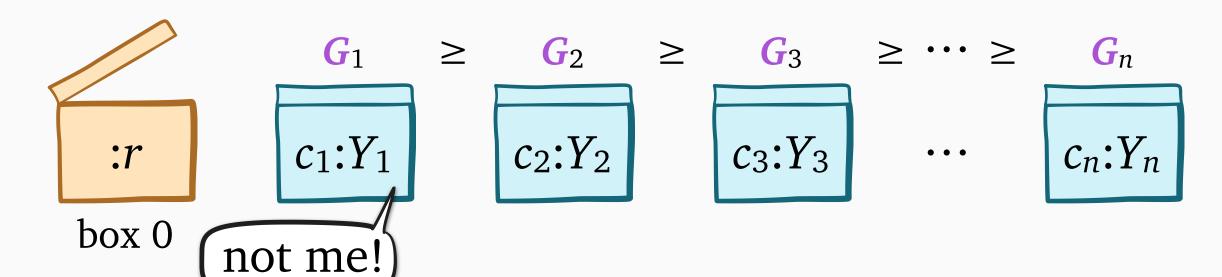
 G_n

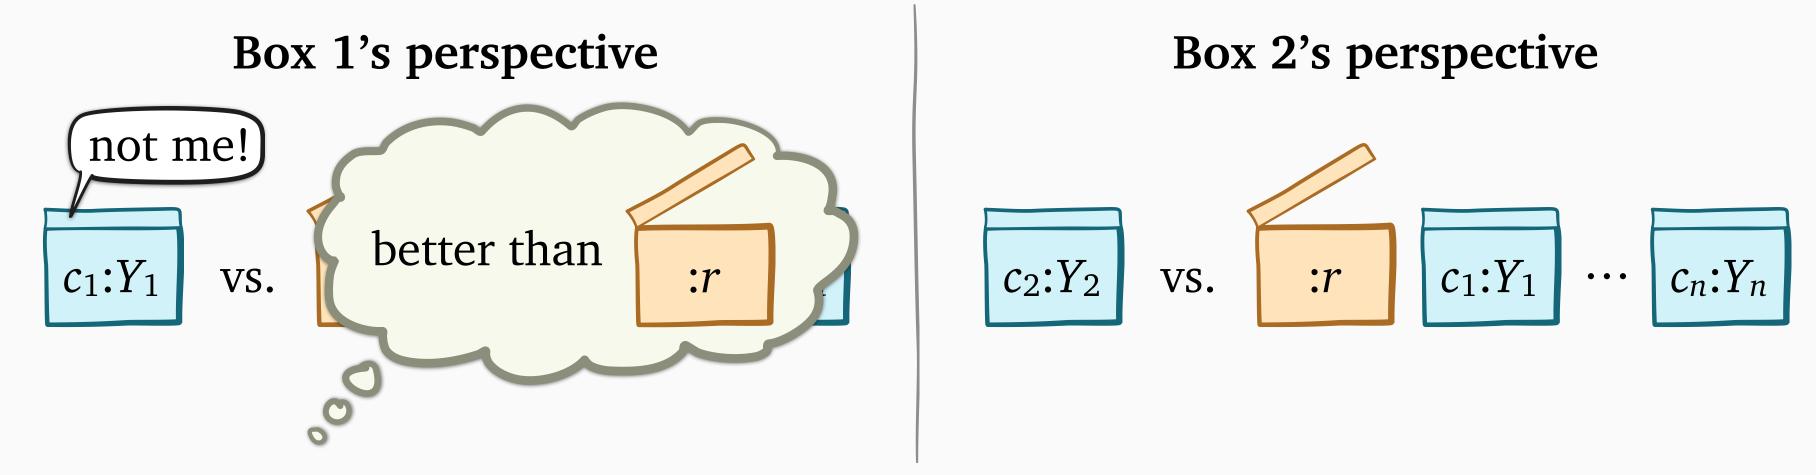


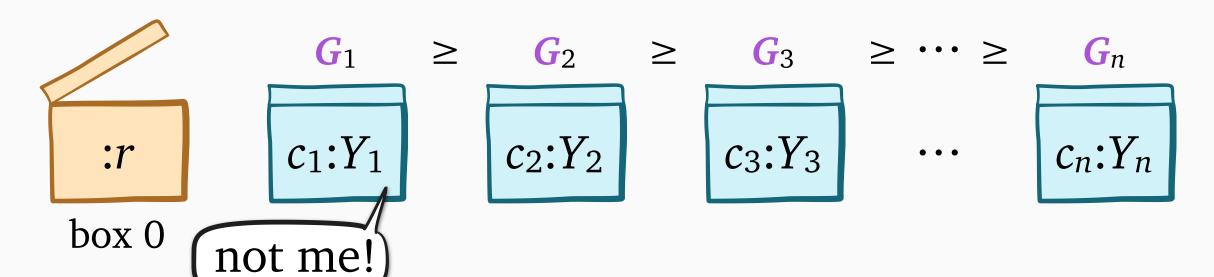
• $r \ge G_1$:

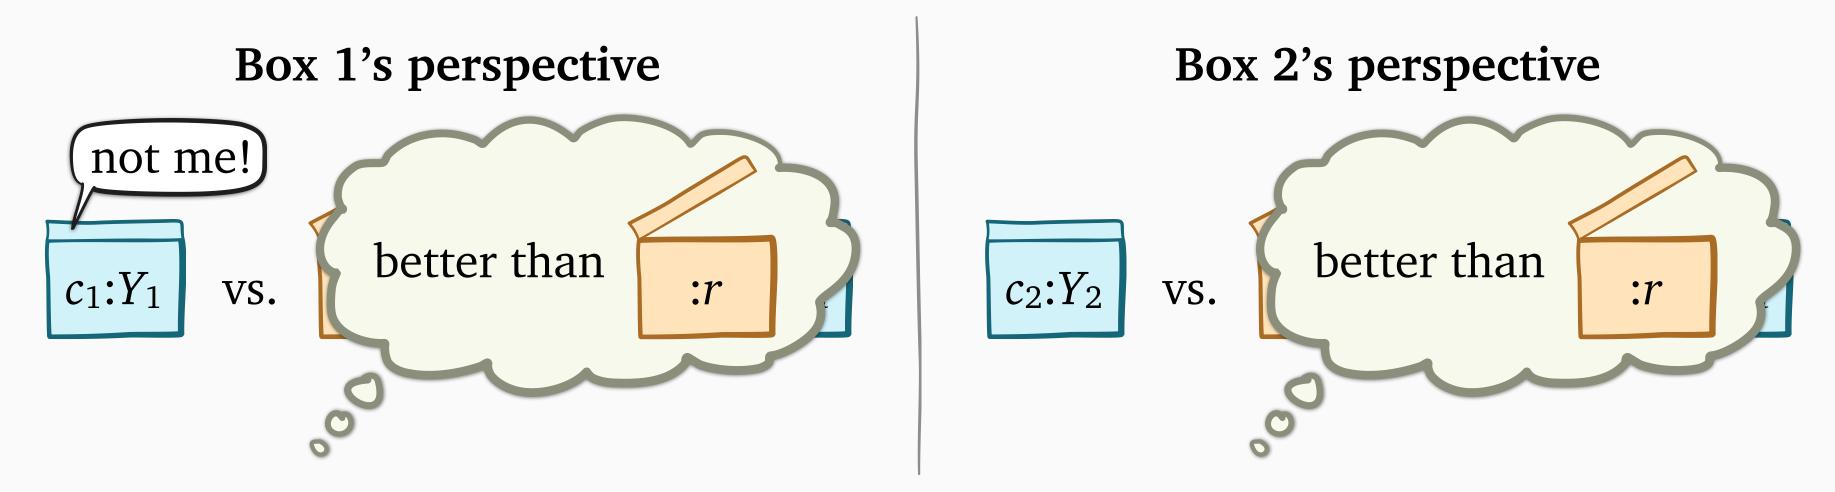


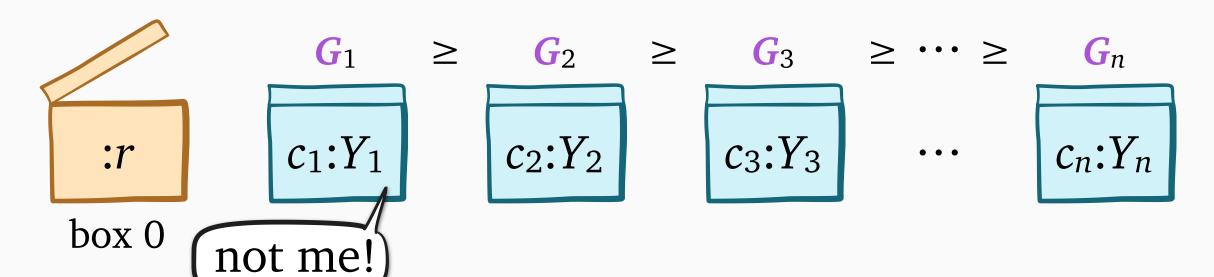


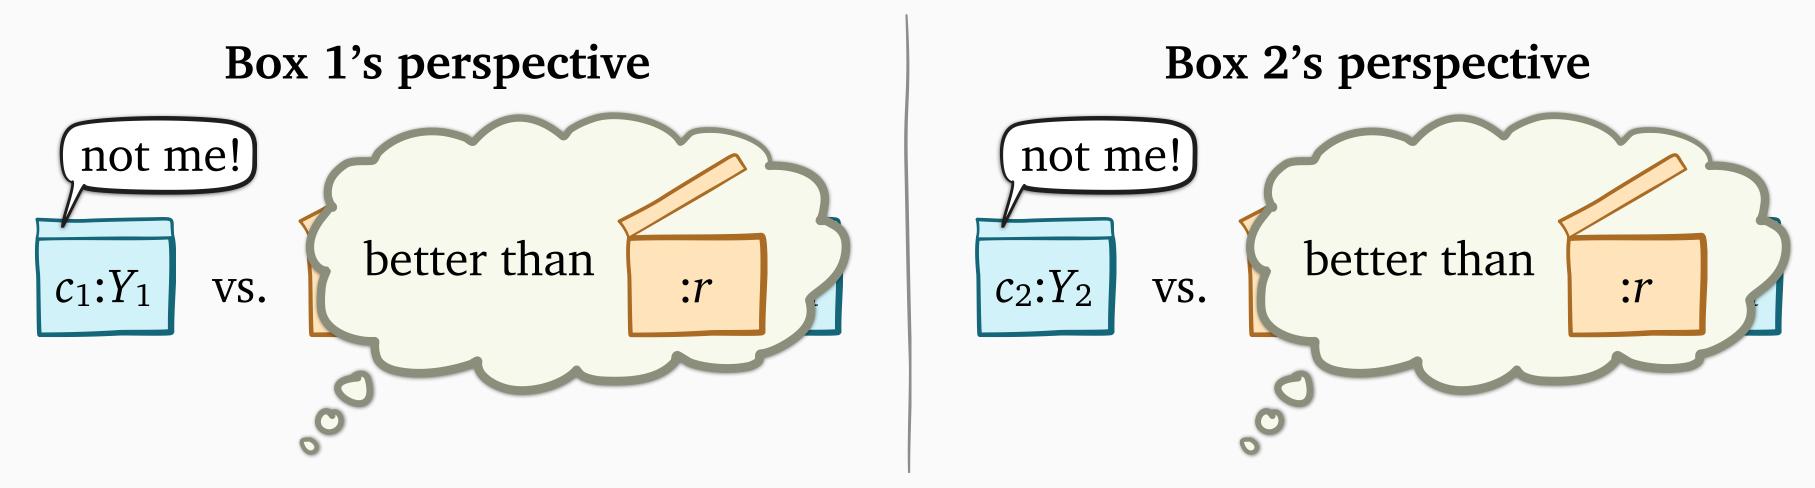


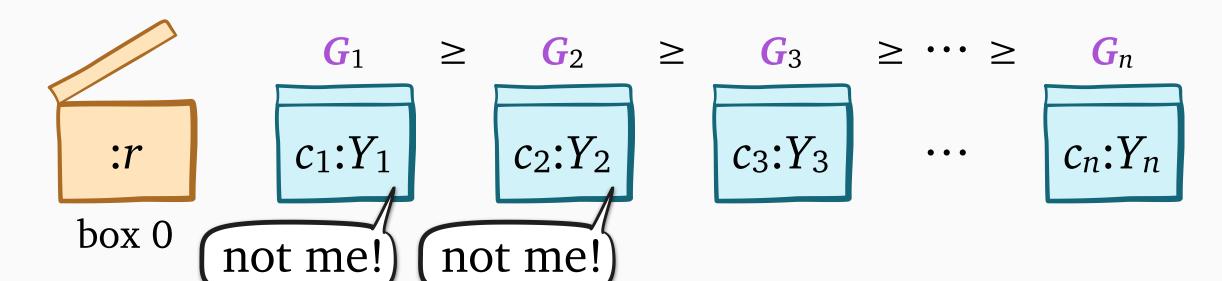


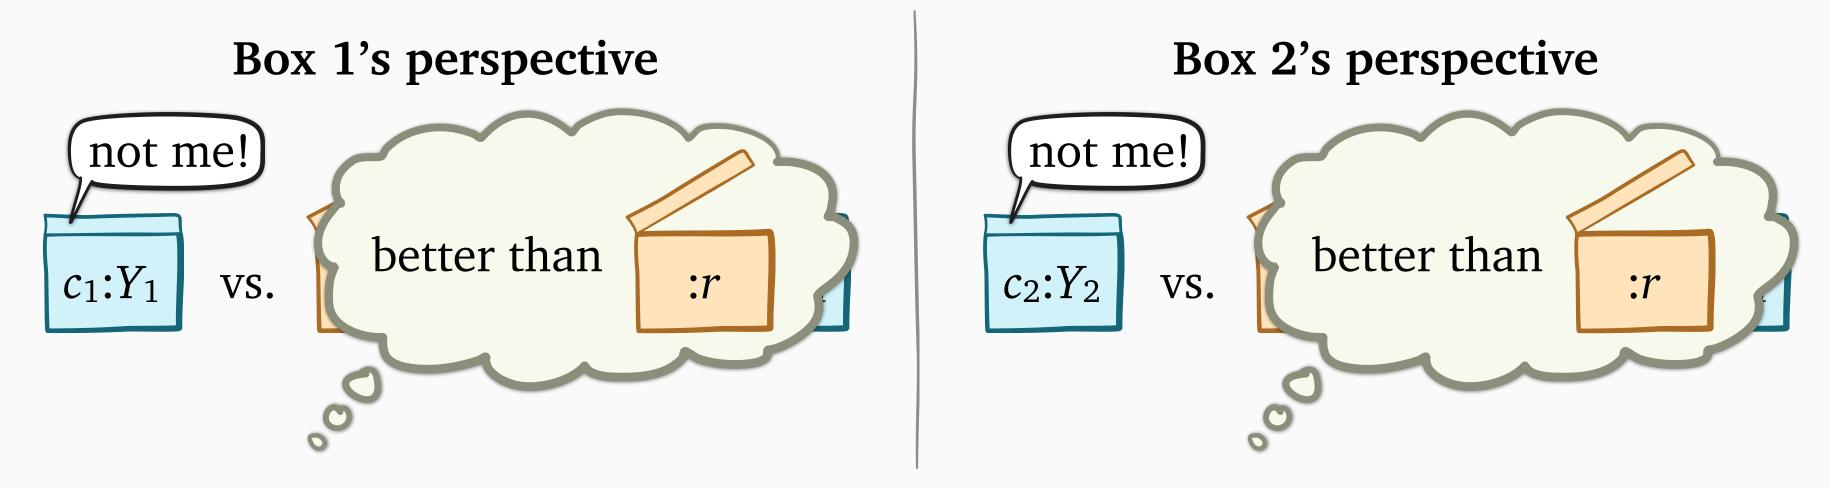


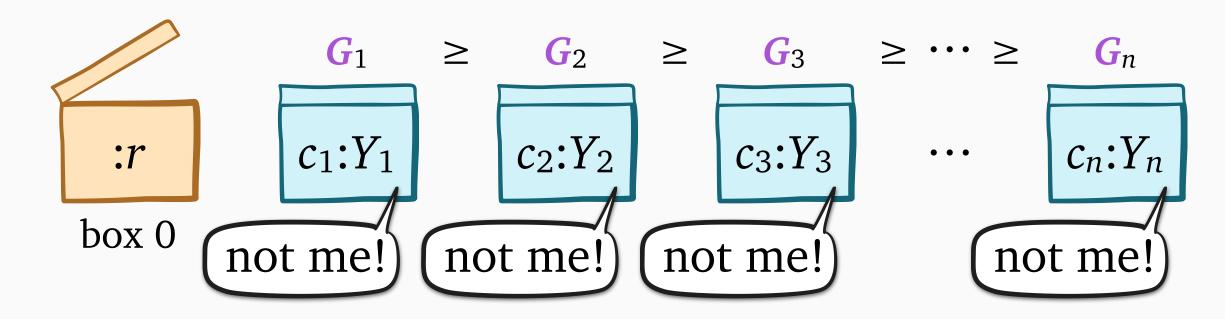


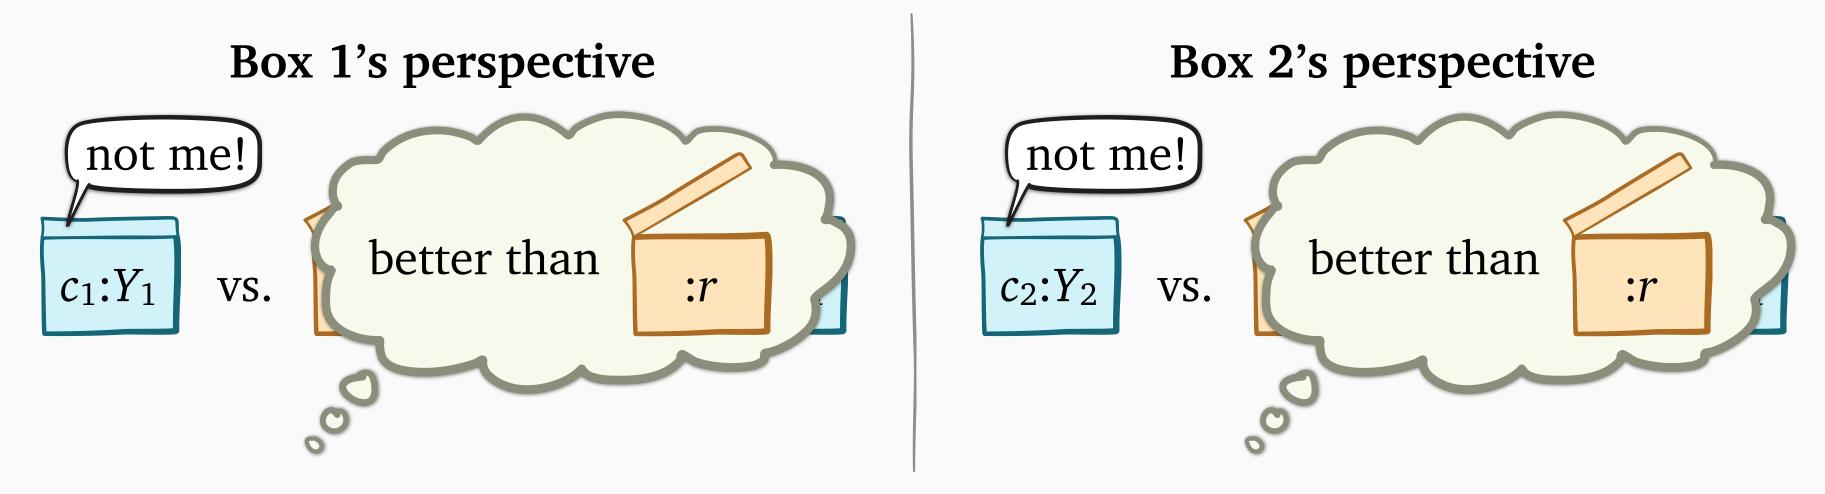




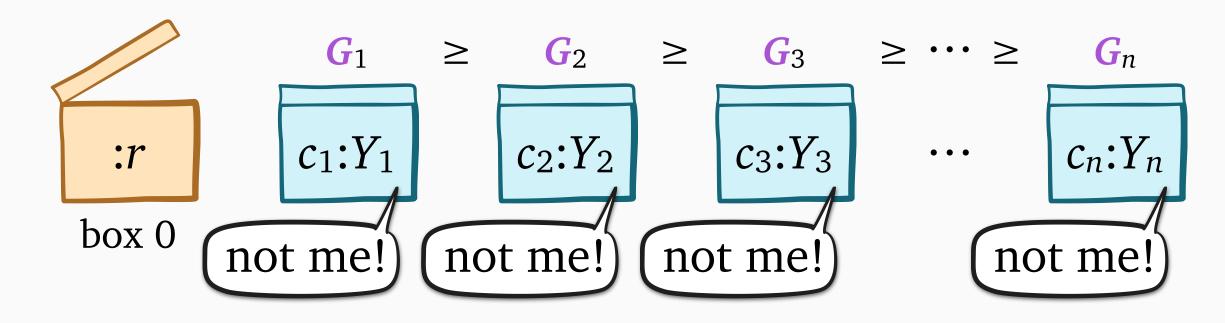


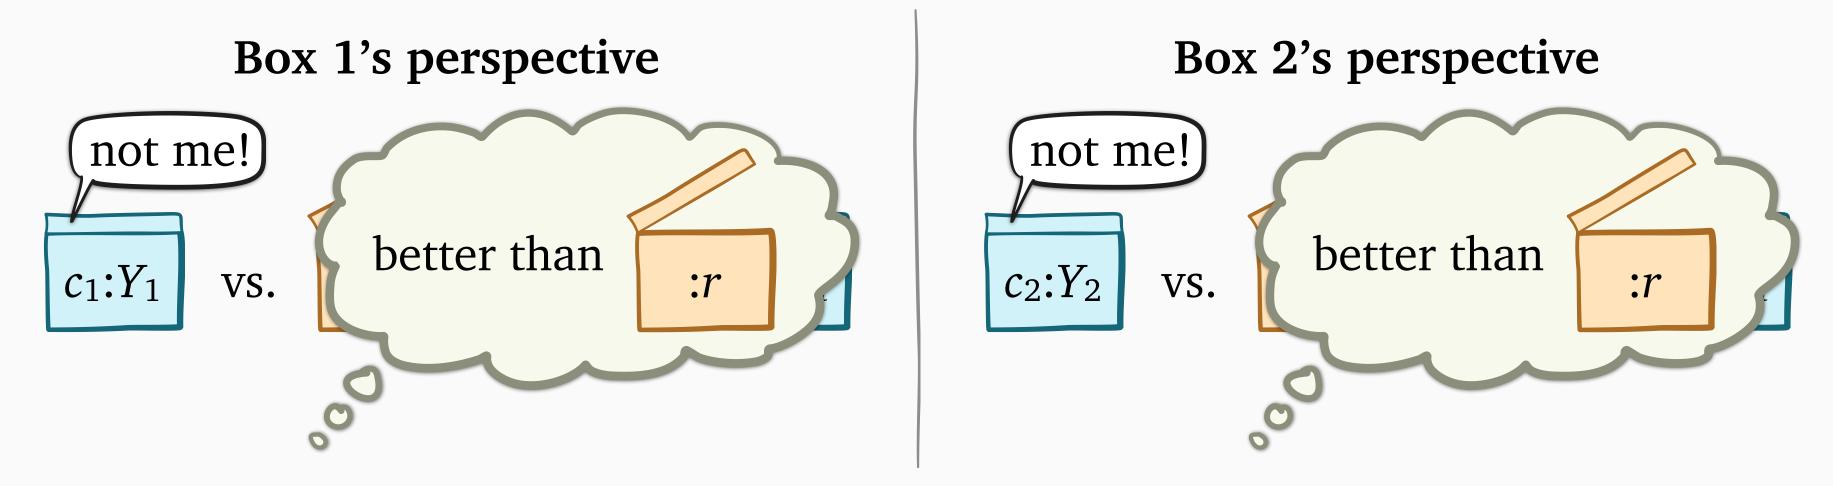




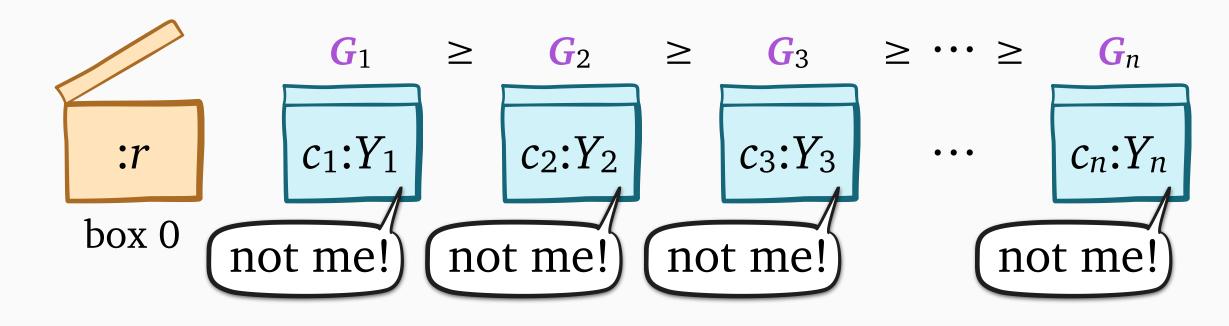


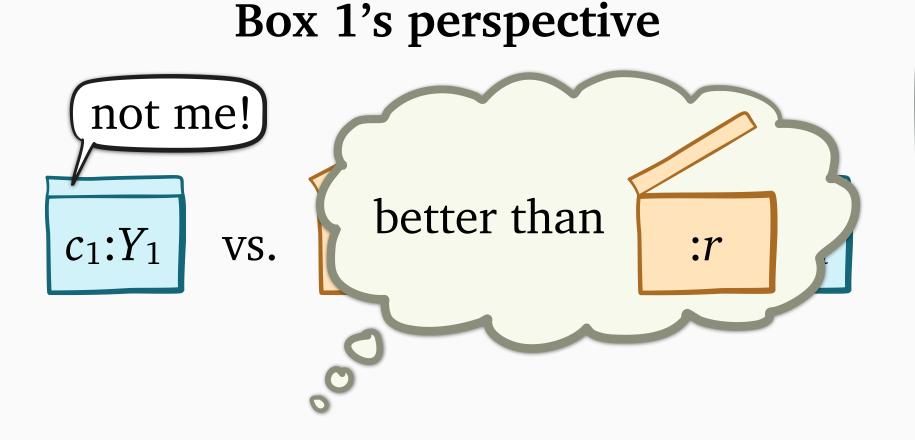
• $r \ge G_1$: select box 0

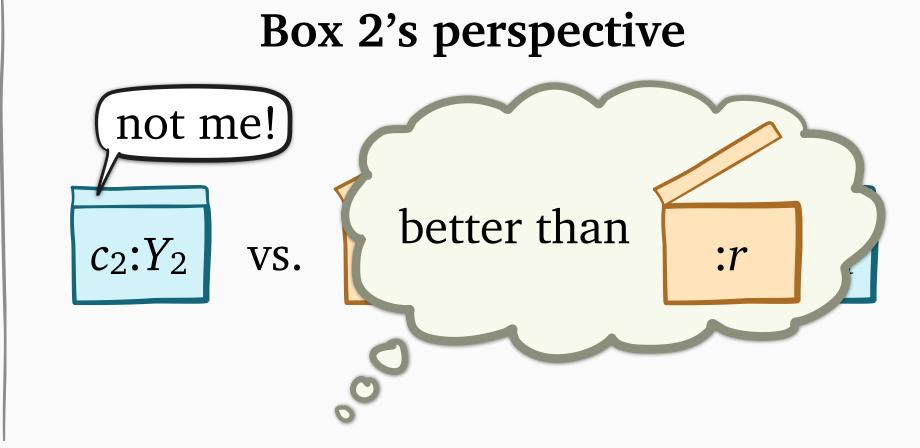




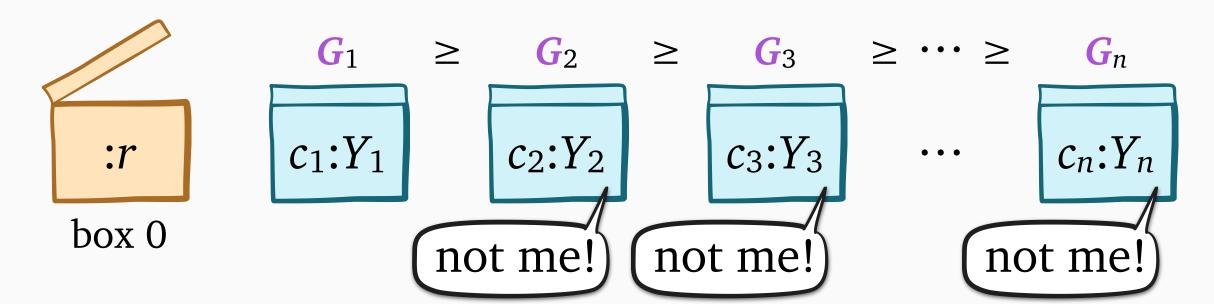
- $r \ge G_1$: select box 0
- $r = G_1$:



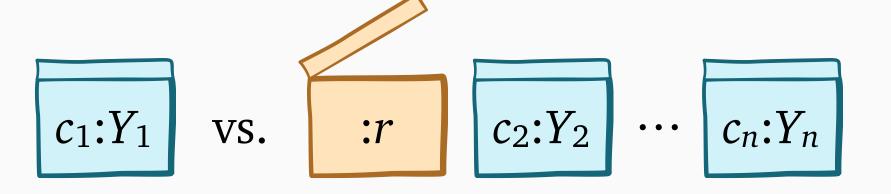


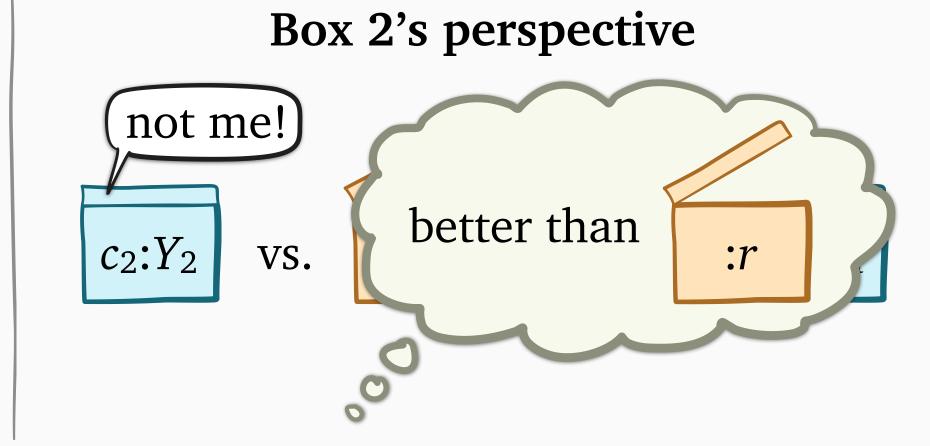


- $r \ge G_1$: select box 0
- $r = G_1$:

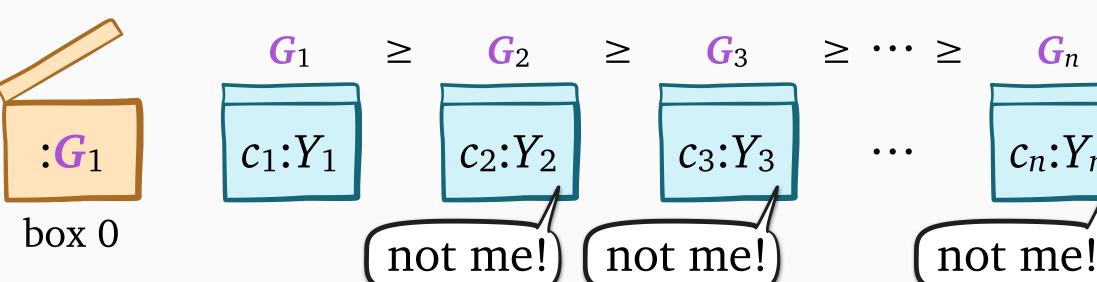


Box 1's perspective

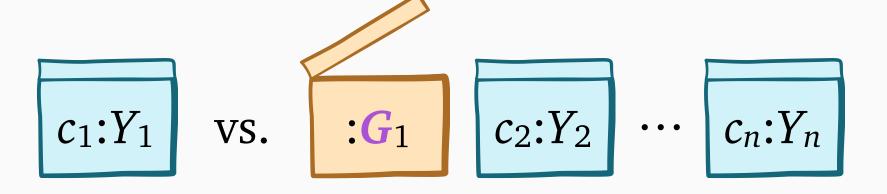




- $r \geq G_1$: select box 0
- $r = G_1$:



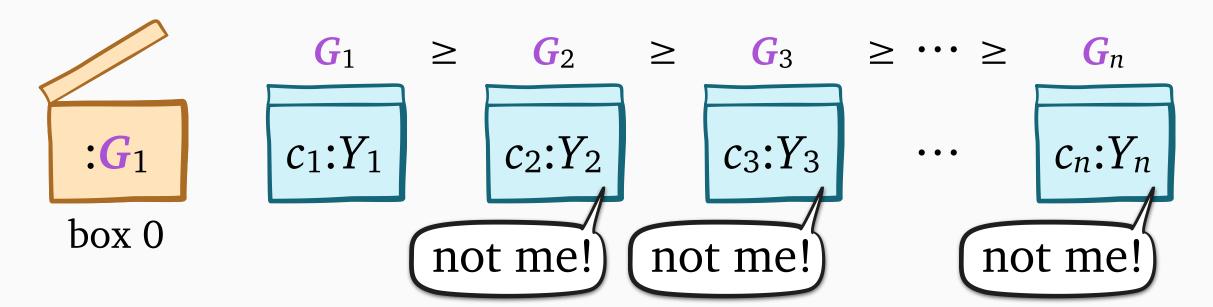
Box 1's perspective



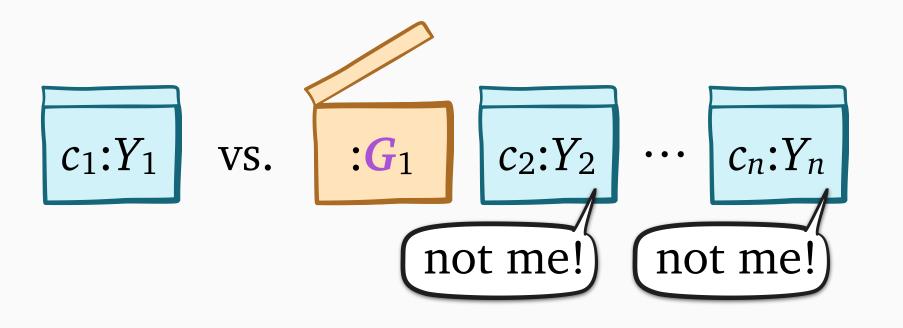
Box 2's perspective not me! better than $c_2:Y_2$ VS.

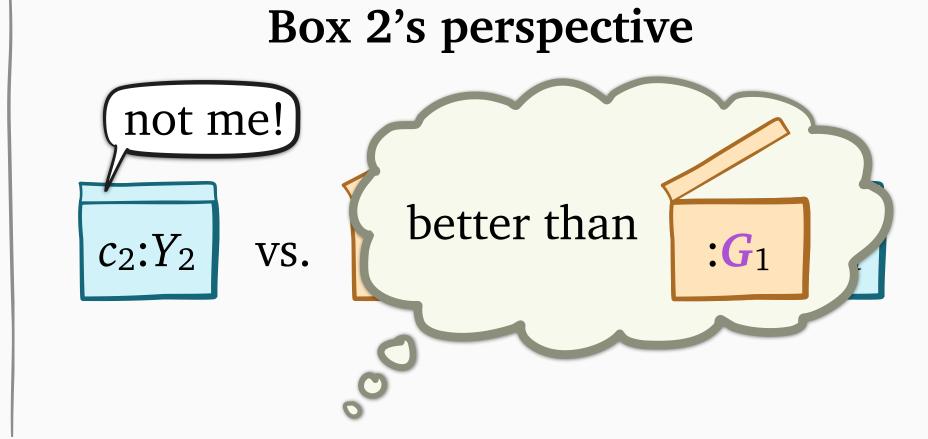
 $c_n:Y_n$

- $r \ge G_1$: select box 0
- $r = G_1$:

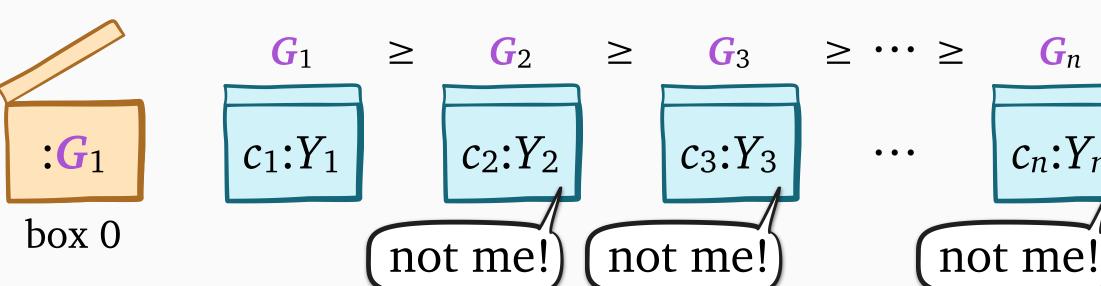


Box 1's perspective

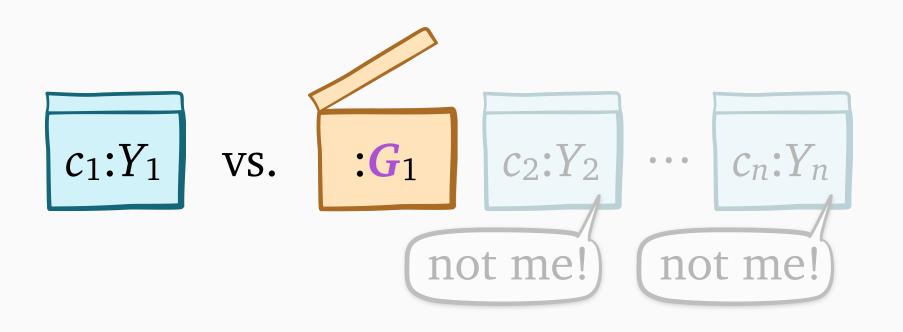


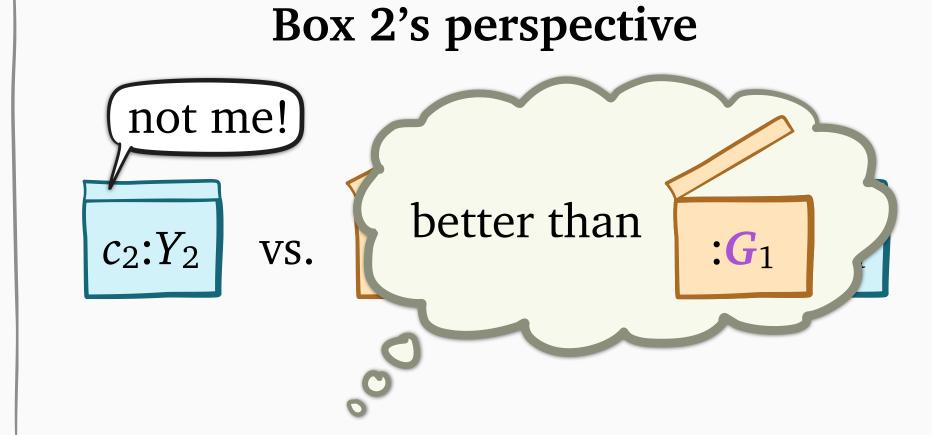


- $r \geq G_1$: select box 0
- $r = G_1$:



Box 1's perspective

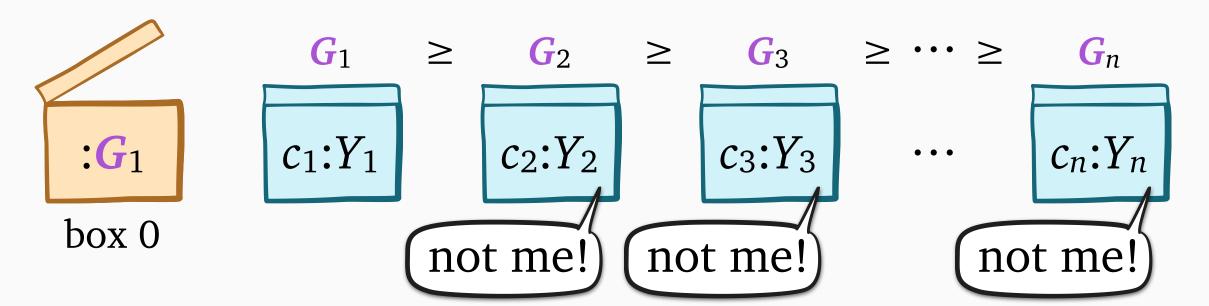




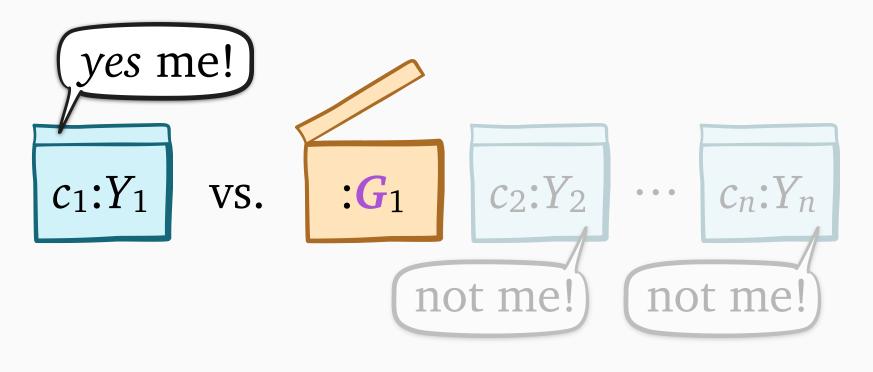
 G_n

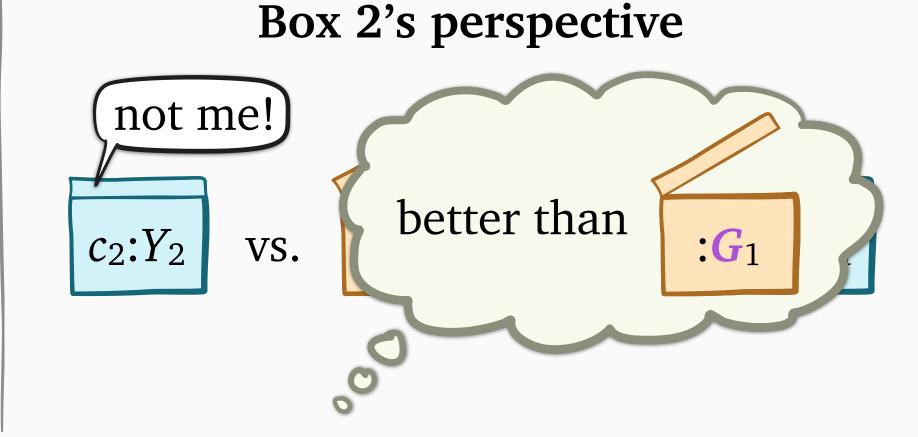
 $c_n:Y_n$

- $r \ge G_1$: select box 0
- $r = G_1$:

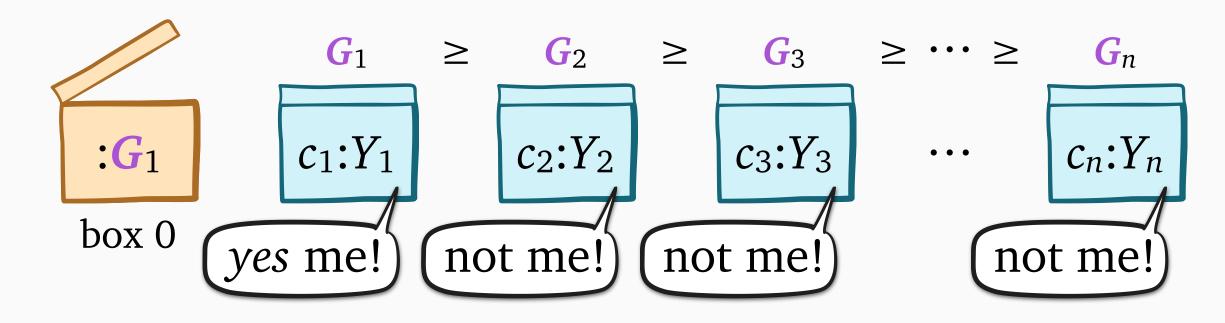


Box 1's perspective

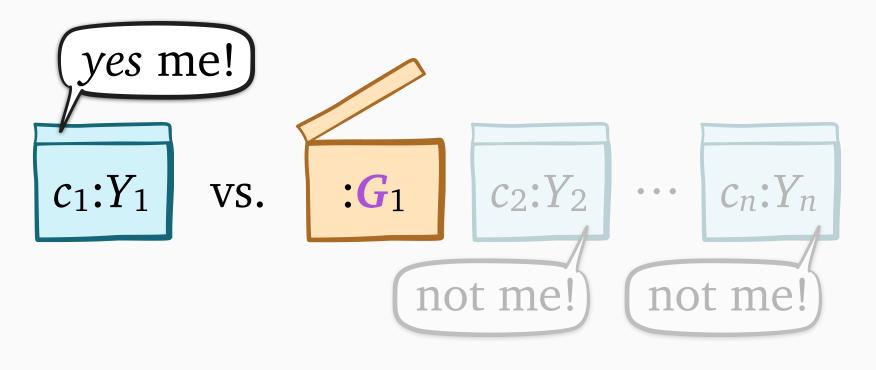




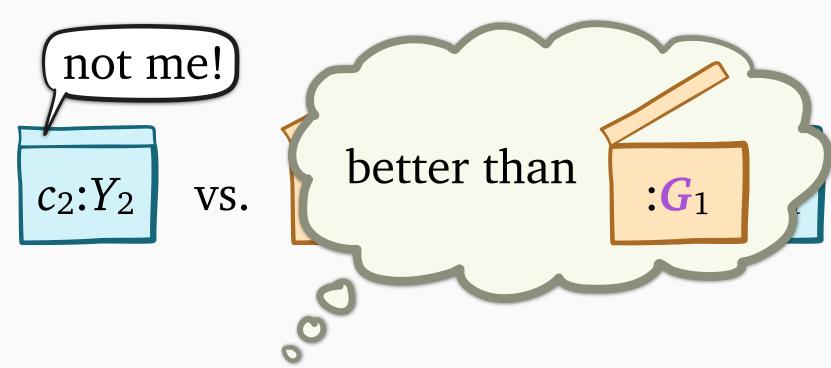
- $r \ge G_1$: select box 0
- $r = G_1$:



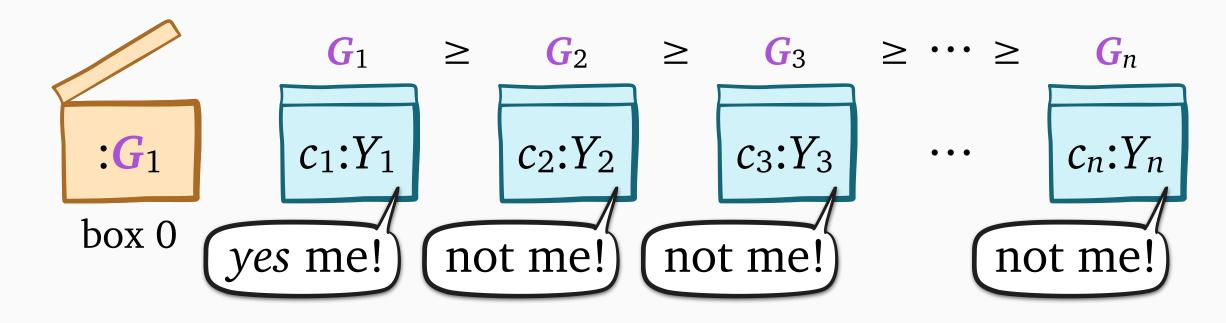
Box 1's perspective



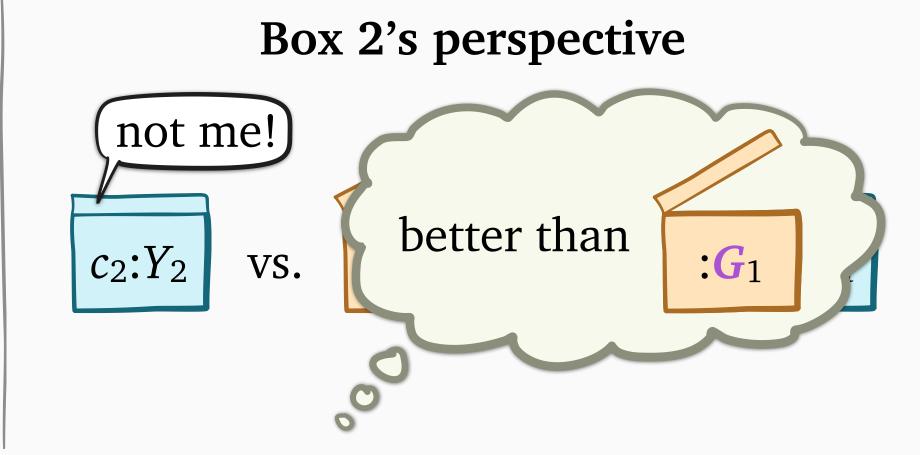
Box 2's perspective



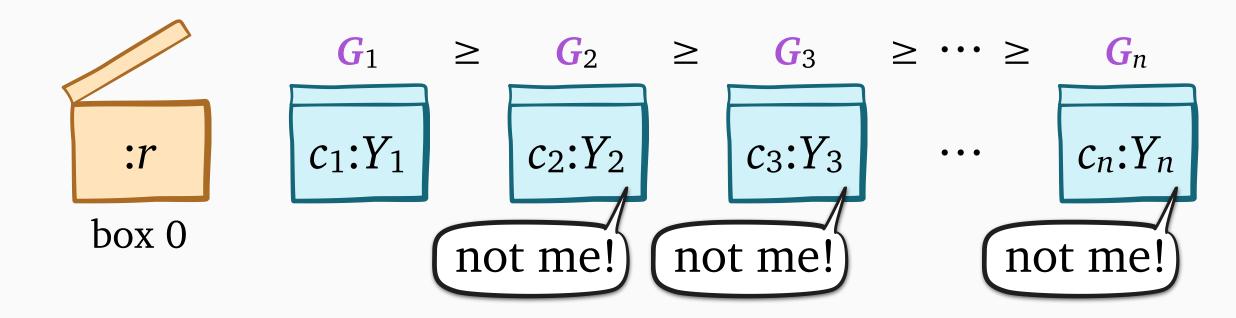
- $r \ge G_1$: select box 0
- $r = G_1$: open box 1



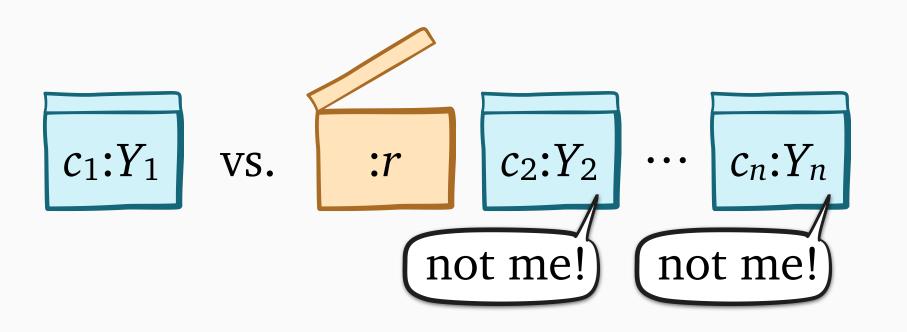
Box 1's perspective yes me! $c_1:Y_1$ vs. $c_2:Y_2$ ··· $c_n:Y_n$ not me! not me!



- $r \ge G_1$: select box 0
- $r = G_1$: open box 1
- $r < G_1$:

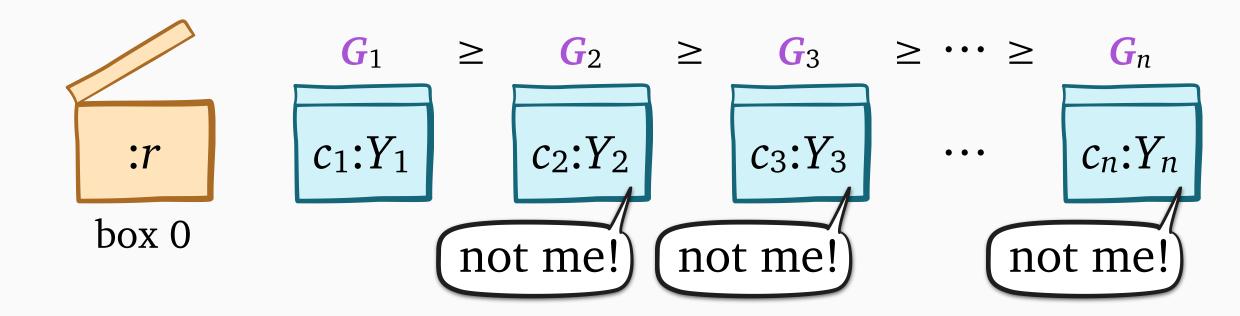


Box 1's perspective

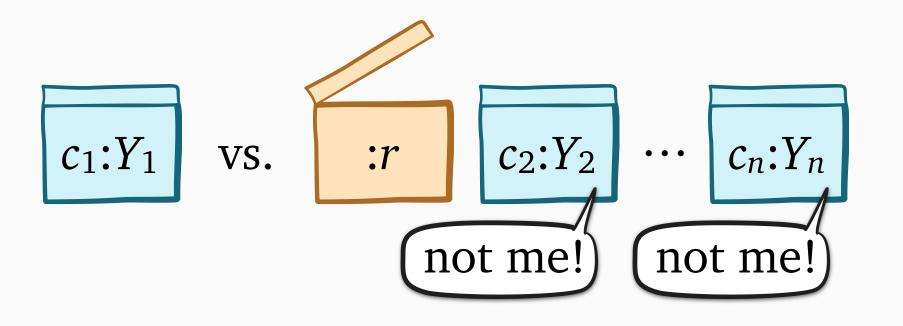


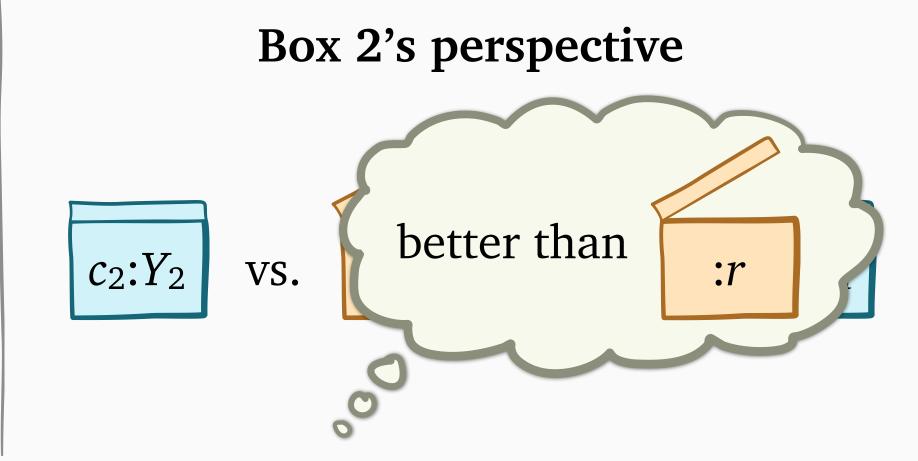
Box 2's perspective not me! c2:Y2 vs. better than :r

- $r \ge G_1$: select box 0
- $r = G_1$: open box 1
- $r < G_1$:

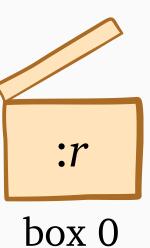


Box 1's perspective





- $r \geq G_1$: select box 0
- $r = G_1$: open box 1



 $c_1:Y_1$

 G_1

 $c_2:Y_2$

 G_2

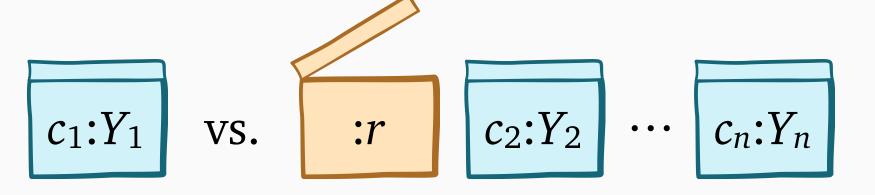
 \geq

 $c_3:Y_3$

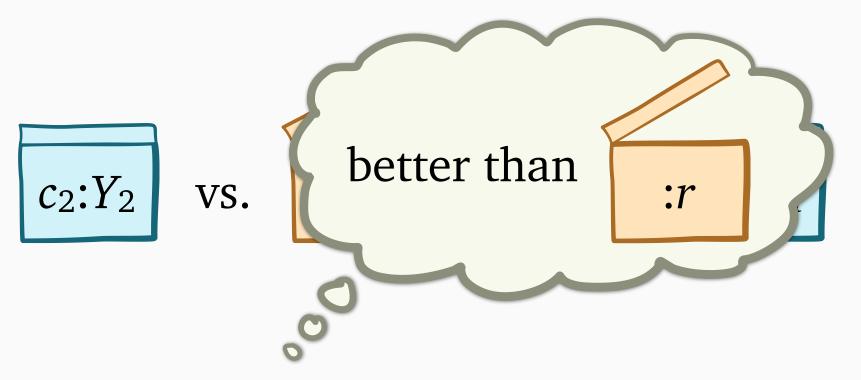
G₃

• $r < G_1$:

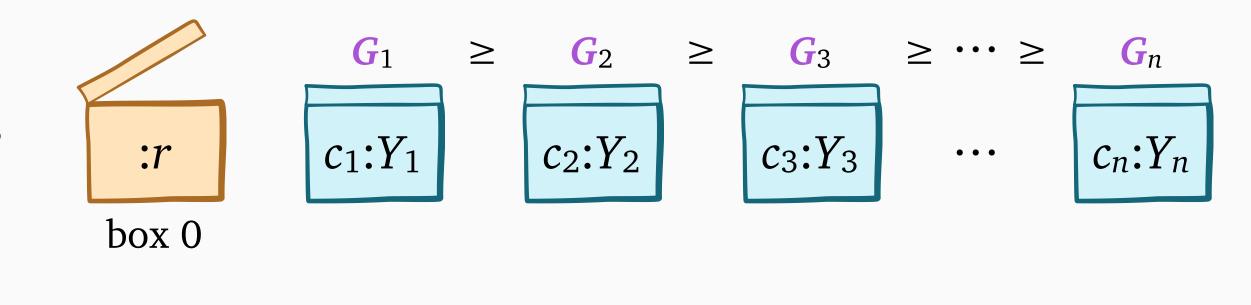
Box 1's perspective

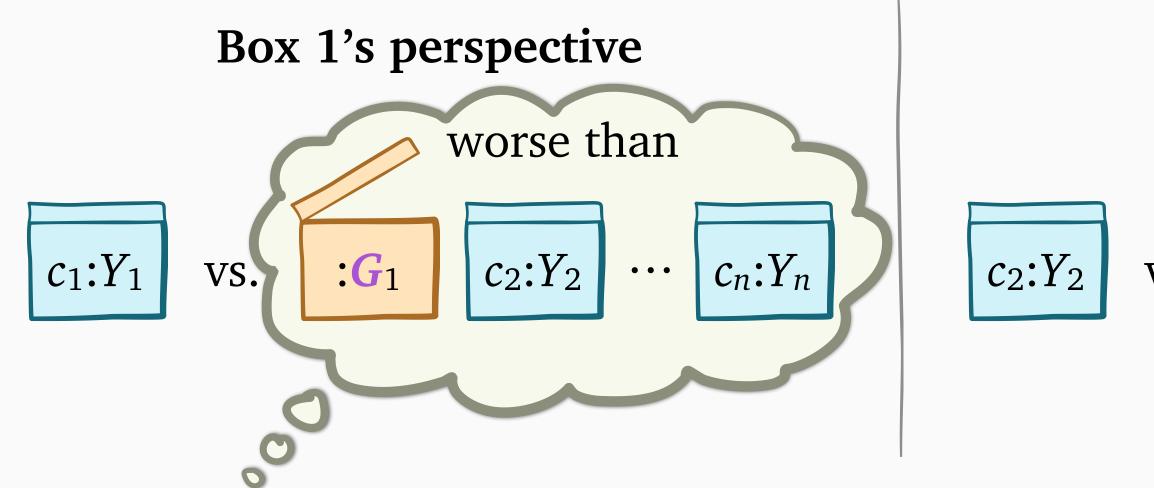


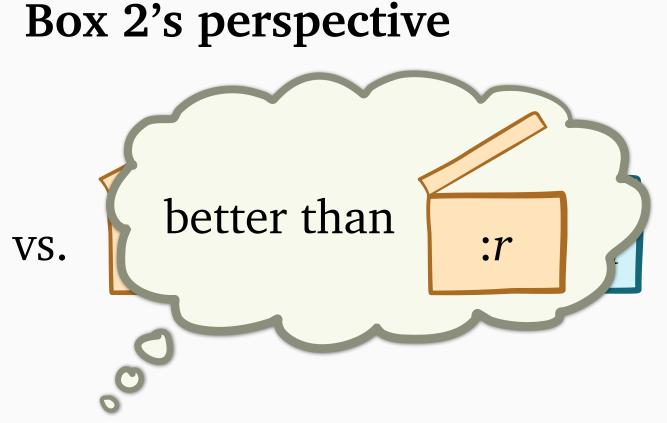
Box 2's perspective



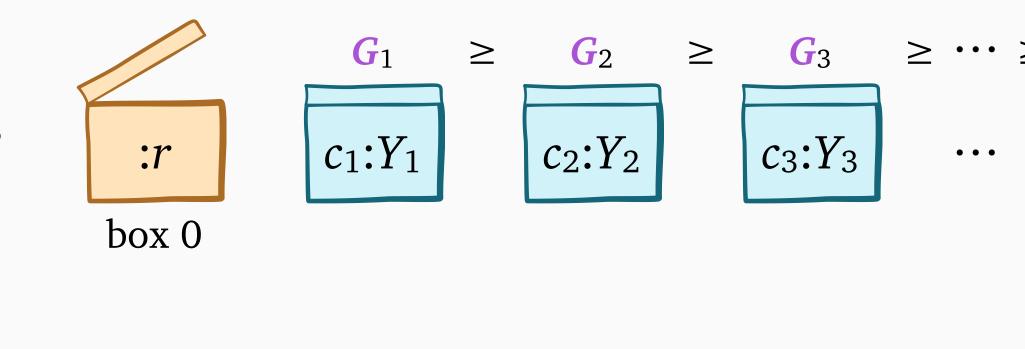
- $r \ge G_1$: select box 0
- $r = G_1$: open box 1
- $r < G_1$:

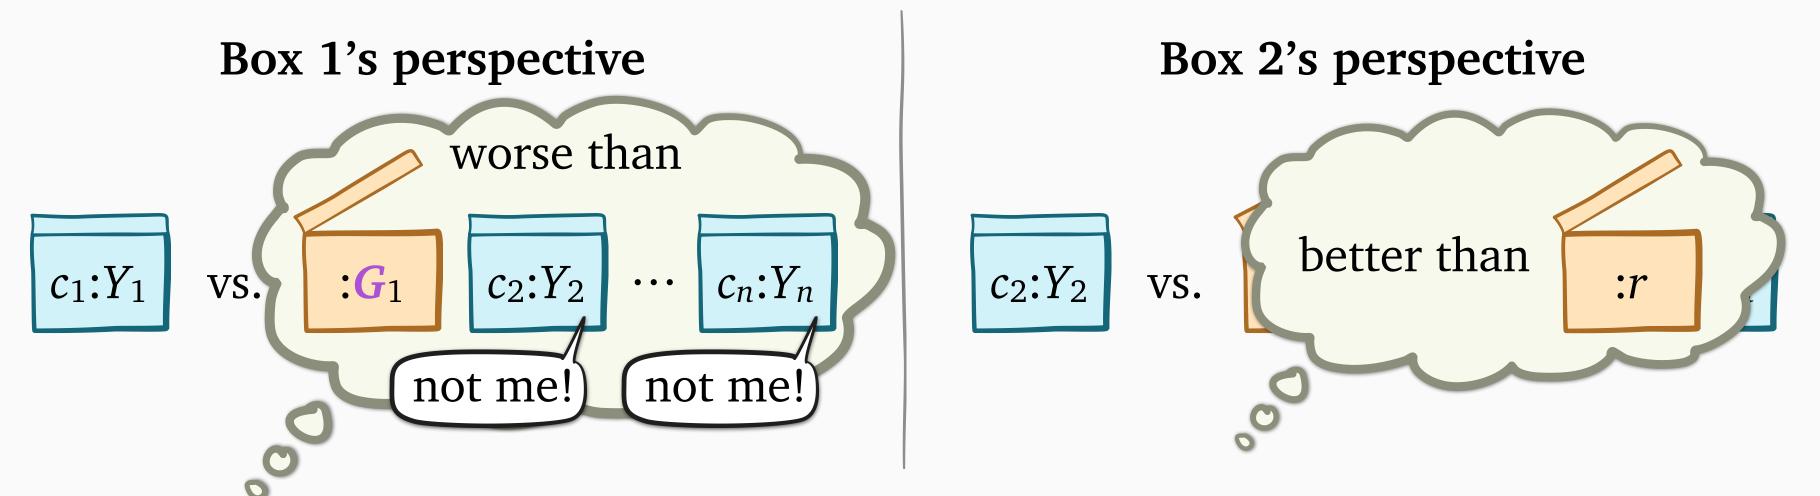






- $r \ge G_1$: select box 0
- $r = G_1$: open box 1
- $r < G_1$:

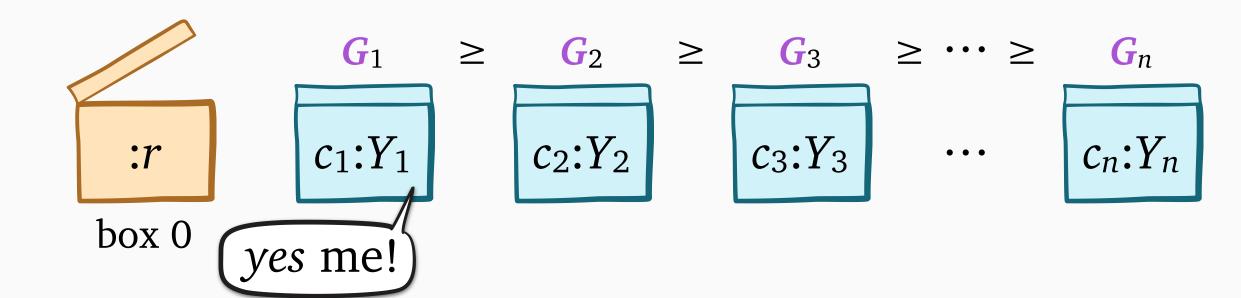


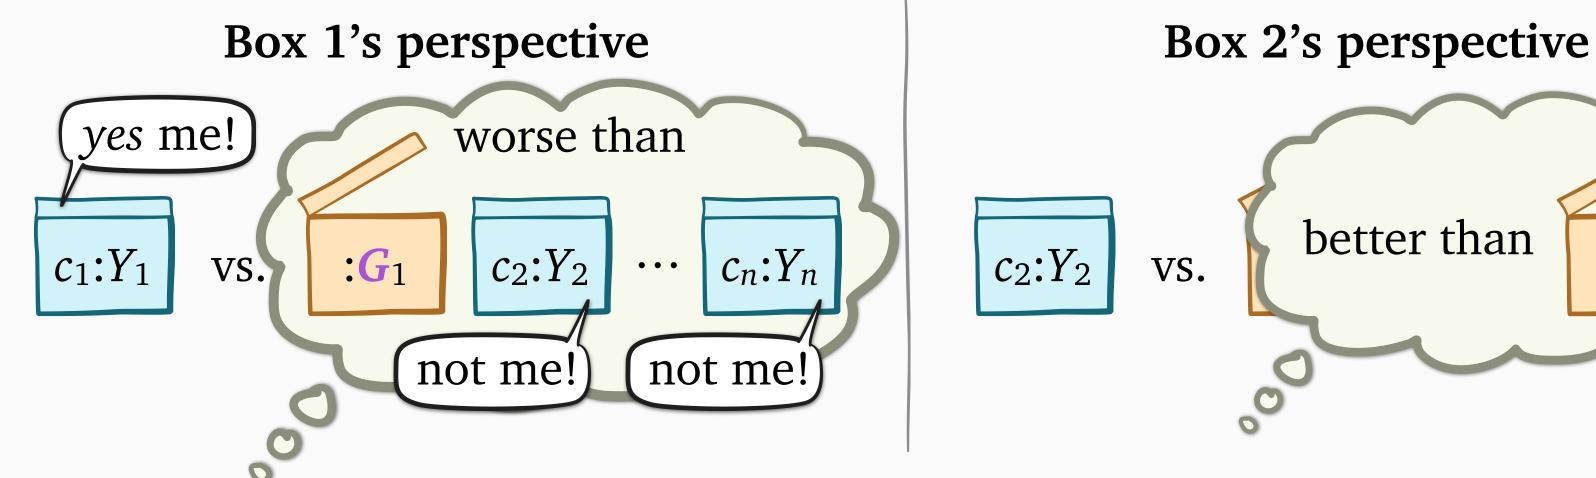


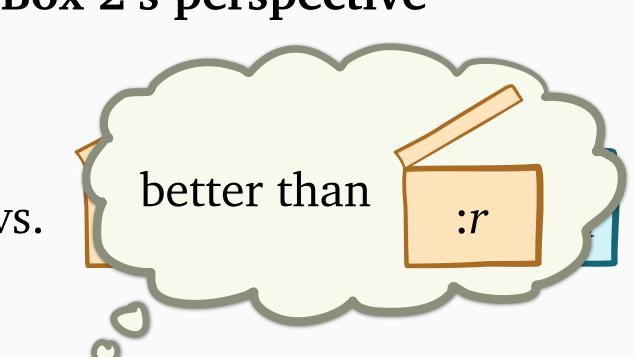
 G_n

 $c_n:Y_n$

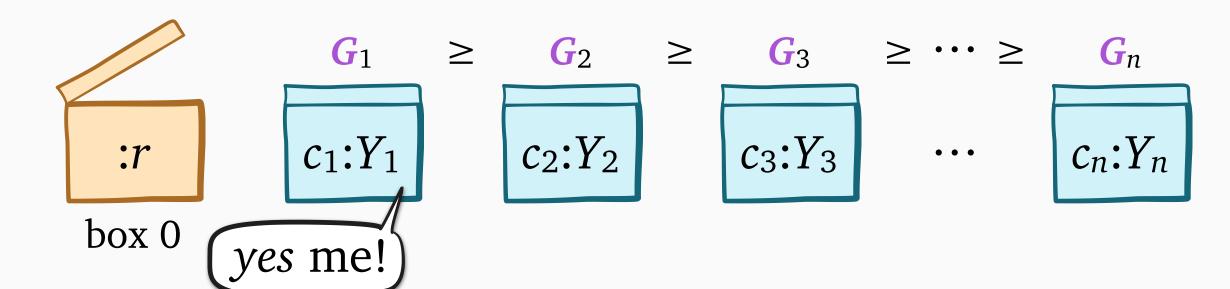
- $r \geq G_1$: select box 0
- $r = G_1$: open box 1
- $r < G_1$:

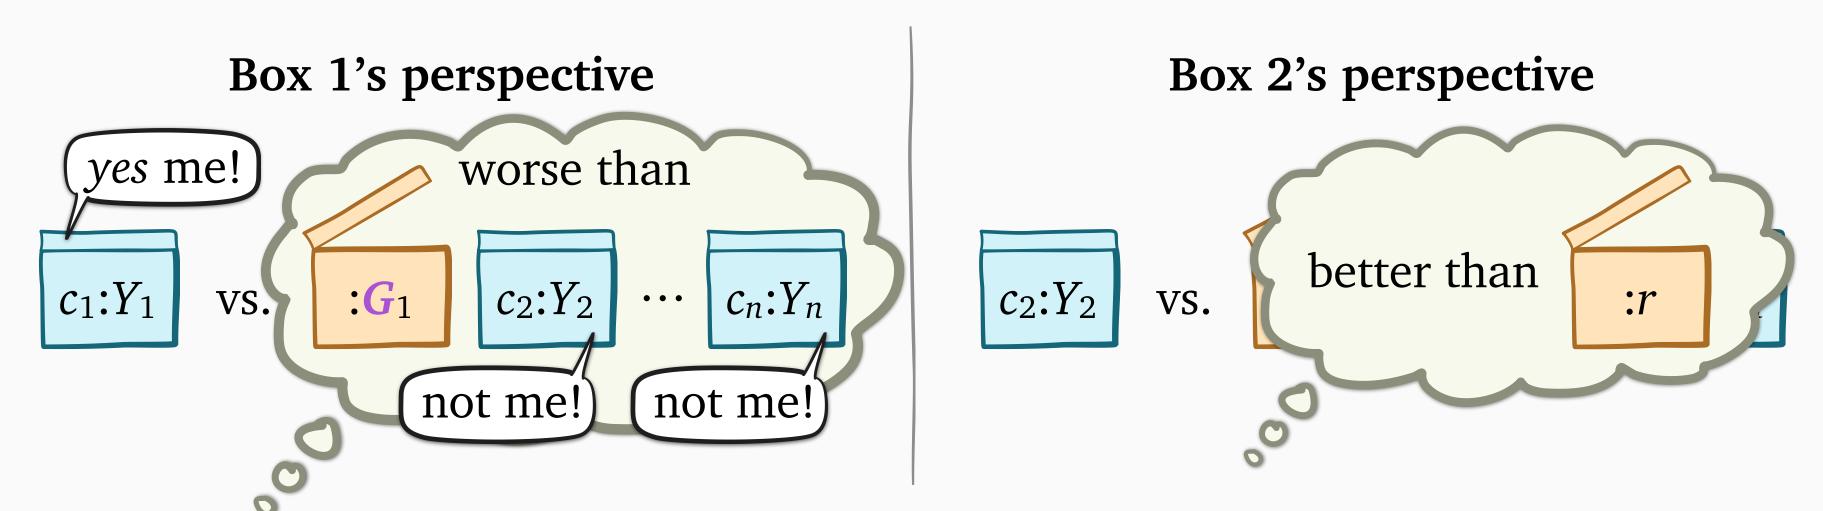




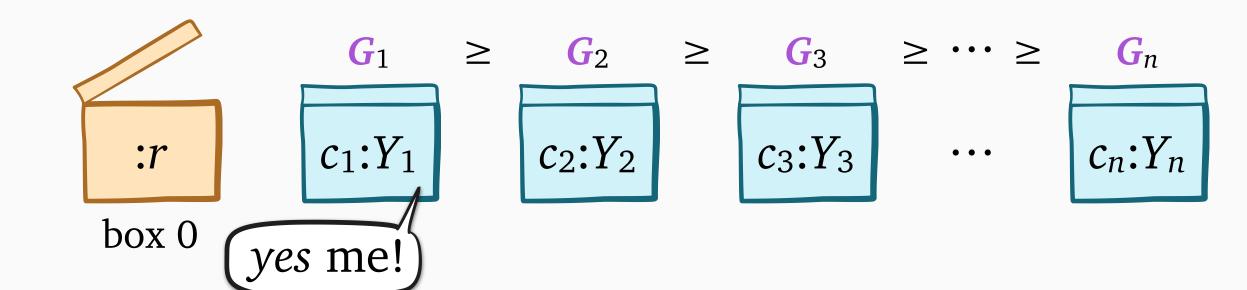


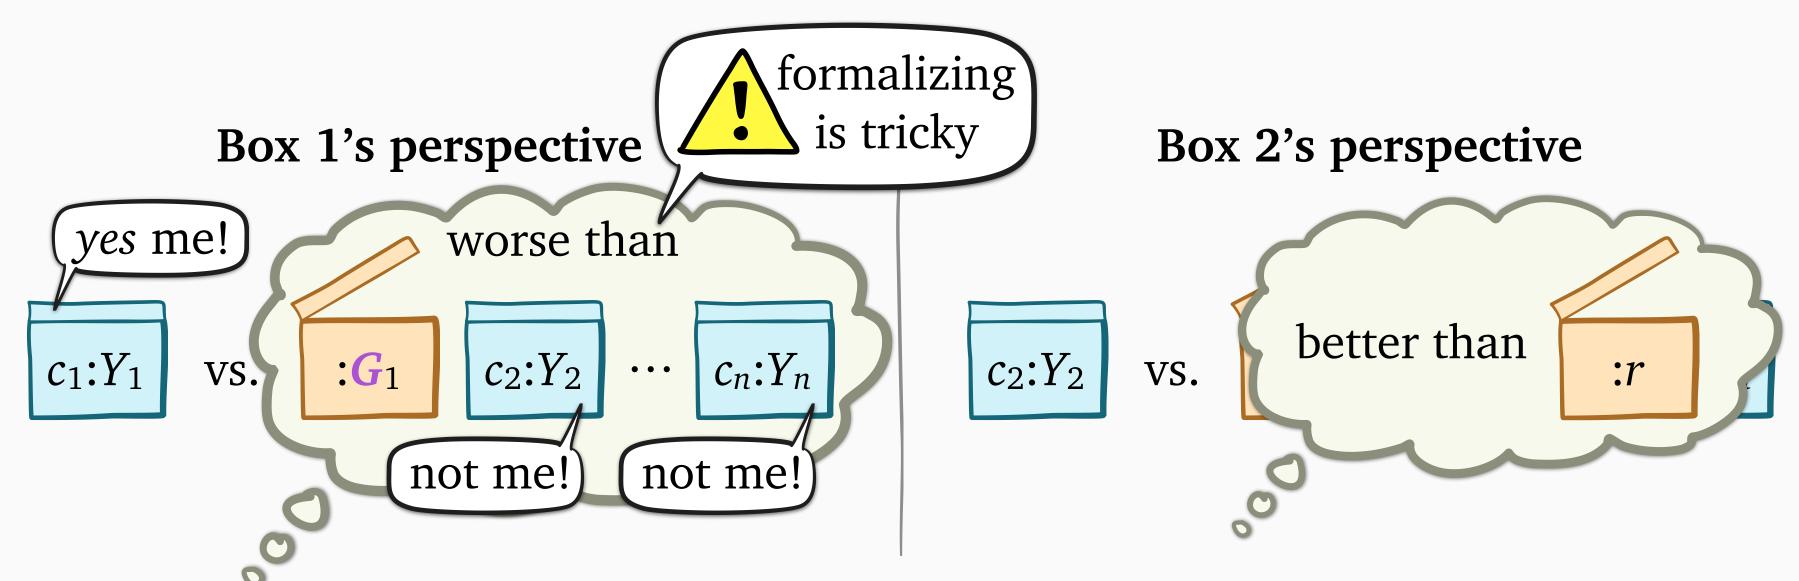
- $r \ge G_1$: select box 0
- $r = G_1$: open box 1
- $r < G_1$: open box 1





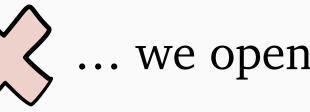
- $r \ge G_1$: select box 0
- $r = G_1$: open box 1
- $r < G_1$: open box 1





Quizzes

... we open *k* boxes at once?



 \dots we open k boxes at once?



 \dots we open k boxes at once?

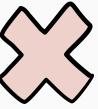
... we select *k* boxes at the end?



 \dots we open k boxes at once?



... we select *k* boxes at the end?

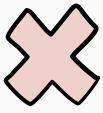


 \dots we open k boxes at once?



... we select *k* boxes at the end?

... we select a *spanning tree* of boxes at the end?



 \dots we open k boxes at once?



... we select *k* boxes at the end?



... we select a *spanning tree* of boxes at the end? [Singla, 2018; Gupta et al., 2019]



 \dots we open k boxes at once?

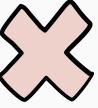


... we select *k* boxes at the end?



... we select a *spanning tree* of boxes at the end? [Singla, 2018; Gupta et al., 2019]

... we can open at most *n* boxes?



 \dots we open k boxes at once?



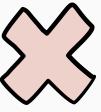
... we select *k* boxes at the end?



.. we select a *spanning tree* of boxes at the end? [Singla, 2018; Gupta et al., 2019]



.. we can open at most n boxes?



 \dots we open k boxes at once?



... we select *k* boxes at the end?



.. we select a *spanning tree* of boxes at the end? [Singla, 2018; Gupta et al., 2019]



... we can open at most n boxes?

... we can open at most *n* boxes in expectation?



 \dots we open k boxes at once?



... we select *k* boxes at the end?



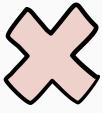
... we select a *spanning tree* of boxes at the end? [Singla, 2018; Gupta et al., 2019]



... we can open at most *n* boxes?



... we can open at most *n* boxes in expectation? [Aminian, Manshadi, & Niazadeh, 2025]



.. we open *k* boxes at once?



... we select *k* boxes at the end?



.. we select a *spanning tree* of boxes at the end? [Singla, 2018; Gupta et al., 2019]

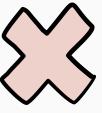


... we can open at most *n* boxes?



... we can open at most *n* boxes in expectation? [Aminian, Manshadi, & Niazadeh, 2025]

... we can select a closed box?



 \dots we open k boxes at once?



... we select *k* boxes at the end?



.. we select a *spanning tree* of boxes at the end? [Singla, 2018; Gupta et al., 2019]



... we can open at most *n* boxes?



... we can open at most *n* boxes in expectation? [Aminian, Manshadi, & Niazadeh, 2025]



... we can select a closed box?

[Fu, Li, & Liu, 2023; Beyhaghi & Cai, 2023]



 \dots we open k boxes at once?



... we select *k* boxes at the end?



... we select a *spanning tree* of boxes at the end? [Singla, 2018; Gupta et al., 2019]



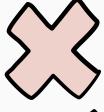
... we can open at most *n* boxes?



... we can open at most *n* boxes in expectation? [Aminian, Manshadi, & Niazadeh, 2025]

... we can select a closed box?

[Fu, Li, & Liu, 2023; Beyhaghi & Cai, 2023]



 \dots we open k boxes at once?



... we select *k* boxes at the end?



... we select a *spanning tree* of boxes at the end? [Singla, 2018; Gupta et al., 2019]



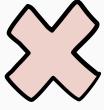
... we can open at most *n* boxes?



... we can open at most *n* boxes in expectation? [Aminian, Manshadi, & Niazadeh, 2025]

... we can select a closed box? [Fu, Li, & Liu, 2023; Beyhaghi & Cai, 2023]

... there are correlations between box values?



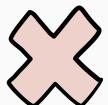
 \dots we open k boxes at once?



... we select *k* boxes at the end?



... we select a *spanning tree* of boxes at the end? [Singla, 2018; Gupta et al., 2019]



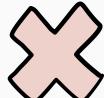
... we can open at most *n* boxes?



... we can open at most *n* boxes in expectation? [Aminian, Manshadi, & Niazadeh, 2025]



... we can select a closed box? [Fu, Li, & Liu, 2023; Beyhaghi & Cai, 2023]



... there are correlations between box values? [Gergatsouli & Tzamos, 2023]



 \dots we open k boxes at once?



... we select *k* boxes at the end?



... we select a *spanning tree* of boxes at the end? [Singla, 2018; Gupta et al., 2019]

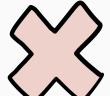


... we can open at most *n* boxes?



... we can open at most *n* boxes in expectation? [Aminian, Manshadi, & Niazadeh, 2025]





- .. there are correlations between box values? [Gergatsouli & Tzamos, 2023]
- ... there are multiple inspection steps?



 \dots we open k boxes at once?



... we select *k* boxes at the end?



... we select a *spanning tree* of boxes at the end? [Singla, 2018; Gupta et al., 2019]



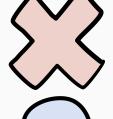
... we can open at most *n* boxes?



... we can open at most *n* boxes in expectation? [Aminian, Manshadi, & Niazadeh, 2025]



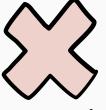
... we can select a closed box? [Fu, Li, & Liu, 2023; Beyhaghi & Cai, 2023]



... there are correlations between box values? [Gergatsouli & Tzamos, 2023]



... there are multiple inspection steps?



we open k boxes at once?



... we select *k* boxes at the end?



... we select a spanning tree of boxes at the end? [Singla, 2018; Gupta et al., 2019]



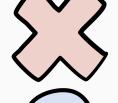
... we can open at most *n* boxes?



 \dots we can open at most n boxes in expectation? [Aminian, Manshadi, & Niazadeh, 2025]



... we can select a closed box? [Fu, Li, & Liu, 2023; Beyhaghi & Cai, 2023]

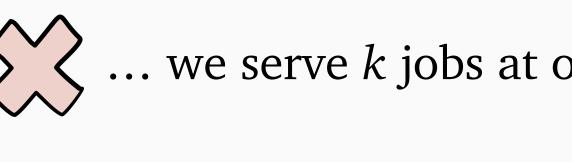


... there are correlations between box values? [Gergatsouli & Tzamos, 2023]



... there are multiple inspection steps?

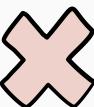
... we serve *k* jobs at once?



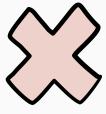


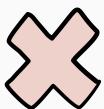
... jobs arrive over time (arbitrary)





... jobs arrive over time (arbitrary)



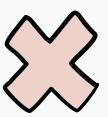


... jobs arrive over time (arbitrary)

... jobs arrive over time (Poisson)



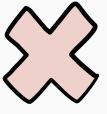
 \sum ... we serve k jobs at once?



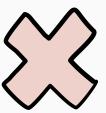
... jobs arrive over time (arbitrary)



... jobs arrive over time (Poisson)



 \sum ... we serve k jobs at once?

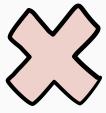


... jobs arrive over time (arbitrary)

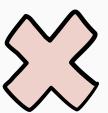


... jobs arrive over time (Poisson)

... Poisson arrivals affected by job in service?



 \dots we serve k jobs at once?



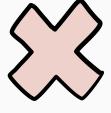
... jobs arrive over time (arbitrary)



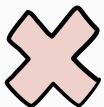
... jobs arrive over time (Poisson)



... Poisson arrivals affected by job in service?



 \dots we serve k jobs at once?

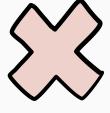


... jobs arrive over time (arbitrary)

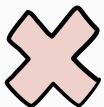


... jobs arrive over time (Poisson)

... Poisson arrivals affected by job in service?



 \dots we serve k jobs at once?



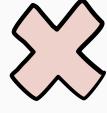
... jobs arrive over time (arbitrary)



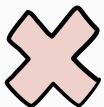
... jobs arrive over time (Poisson)

... Poisson arrivals affected by job in service?

... different jobs have different holding costs?



 \dots we serve k jobs at once?



... jobs arrive over time (arbitrary)



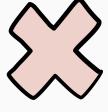
... jobs arrive over time (Poisson)



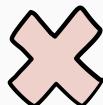
... Poisson arrivals affected by job in service?



... different jobs have different holding costs?



 \dots we serve k jobs at once?



... jobs arrive over time (arbitrary)



... jobs arrive over time (Poisson)



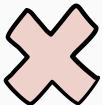
... Poisson arrivals affected by job in service?



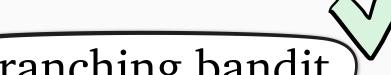
... holding costs change over time (arbitrary)?



 \dots we serve k jobs at once?



... jobs arrive over time (arbitrary)



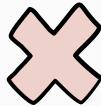
... jobs arrive over time (Poisson)



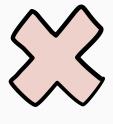
... Poisson arrivals affected by job in service?



... different jobs have different holding costs?



... holding costs change over time (arbitrary)?



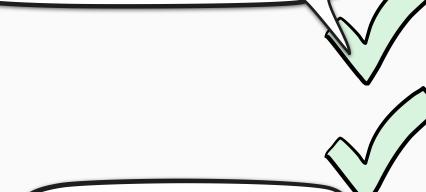
 \dots we serve k jobs at once?



... jobs arrive over time (arbitrary)



... jobs arrive over time (Poisson)



restless bandit

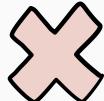
... Poisson arrivals affected by job in service?



... holding costs change over time (arbitrary)?



 \dots we serve k jobs at once?



... jobs arrive over time (arbitrary)



... jobs arrive over time (Poisson)

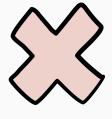


... Poisson arrivals affected by job in service?

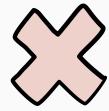
... different jobs have different holding costs?

... holding costs change over time (arbitrary)?

... holding costs change over time (convex)?



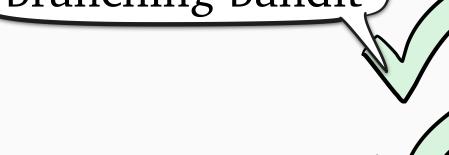
 \dots we serve k jobs at once?



... jobs arrive over time (arbitrary)



... jobs arrive over time (Poisson)



... Poisson arrivals affected by job in service?

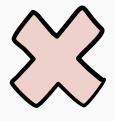


... different jobs have different holding costs?





... holding costs change over time (convex)?



 \dots we serve k jobs at once?



... jobs arrive over time (arbitrary)



... jobs arrive over time (Poisson)

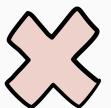


... Poisson arrivals affected by job in service?

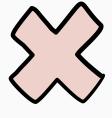


... different jobs have different holding costs?

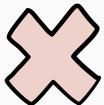




... holding costs change over time (convex)?



 \dots we serve k jobs at once?



... jobs arrive over time (arbitrary)



... jobs arrive over time (Poisson)

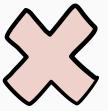


... Poisson arrivals affected by job in service?



... different jobs have different holding costs?





... holding costs change over time (convex)?





 \dots we serve k jobs at once?



... jobs arrive over time (arbitrary)



... jobs arrive over time (Poisson)



... Poisson arrivals affected by job in service?



... different jobs have different holding costs?



... holding costs change over time (arbitrary)?



... holding costs change over time (convex)?







 \dots we serve k jobs at once?



... jobs arrive over time (arbitrary)



... jobs arrive over time (Poisson)



... Poisson arrivals affected by job in service?



... different jobs have different holding costs?

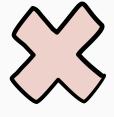


... holding costs change over time (arbitrary)?



... holding costs change over time (convex)?





 \dots we serve k jobs at once?



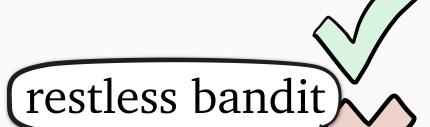
... jobs arrive over time (arbitrary)



... jobs arrive over time (Poisson)



... Poisson arrivals affected by job in service?



... different jobs have different holding costs?



... holding costs change ov



... and with discounted costs/rewards



... holding costs change over time (convex)?



Pandora's box with optional inspection

Pandora's box: is Gittins still optimal if...



 \dots we open k boxes at once?



... we select *k* boxes at the end?



... we select a *spanning tree* of boxes at the end? [Singla, 2018; Gupta et al., 2019]



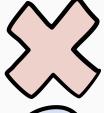
... we can open at most *n* boxes?



... we can open at most *n* boxes *in expectation*? [Aminian, Manshadi, & Niazadeh, 2025]



... we can select a closed box? [Fu, Li, & Liu, 2023; Beyhaghi & Cai, 2023]



... there are correlations between box values? [Gergatsouli & Tzamos, 2023]



... there are multiple inspection steps?

Pandora's box: is Gittins still optimal if...



 \dots we open k boxes at once?



... we select *k* boxes at the end?



... we select a *spanning tree* of boxes at the end? [Singla, 2018; Gupta et al., 2019]



... we can open at most *n* boxes?



... we can open at most *n* boxes in expectation? [Aminian, Manshadi, & Niazadeh, 2025]







... there are multiple inspection steps?

