A new design tool for Bayesian optimization: the

Gittins **index**

Ziv Scully
Cornell University

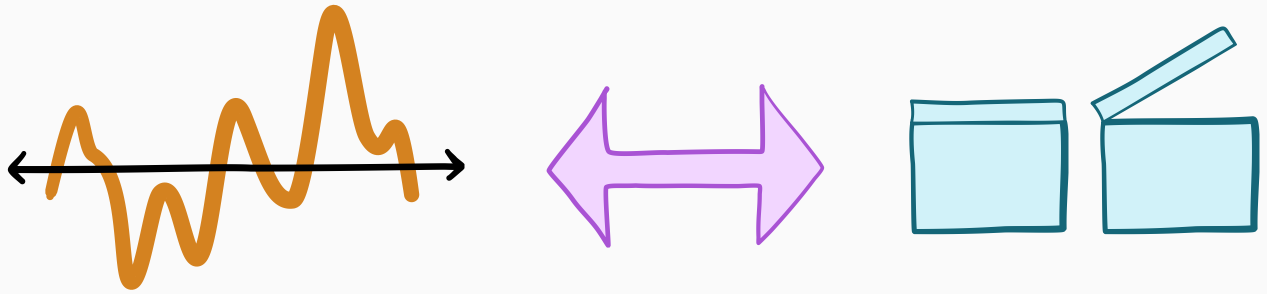

Collaborators

Collaborators

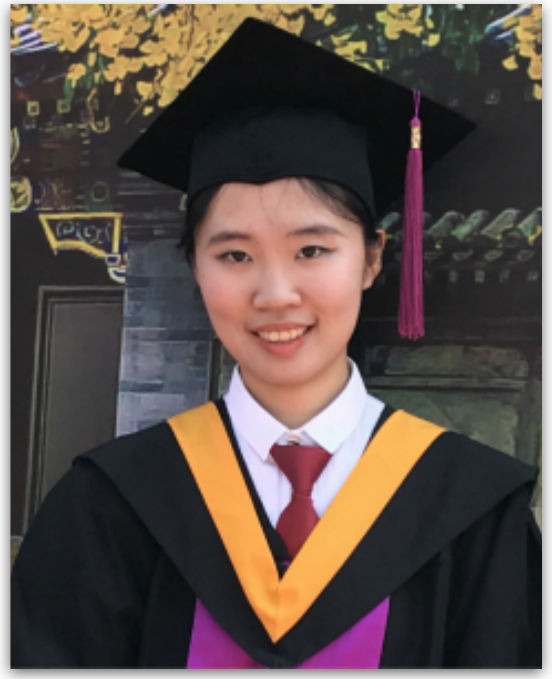

Qian Xie
Cornell

Alex Terenin
Cornell

Collaborators

NeurIPS 2024: *Cost-aware Bayesian Optimization via the Pandora's Box Gittins Index*

Raul Astudillo
Caltech

Peter Frazier
Cornell

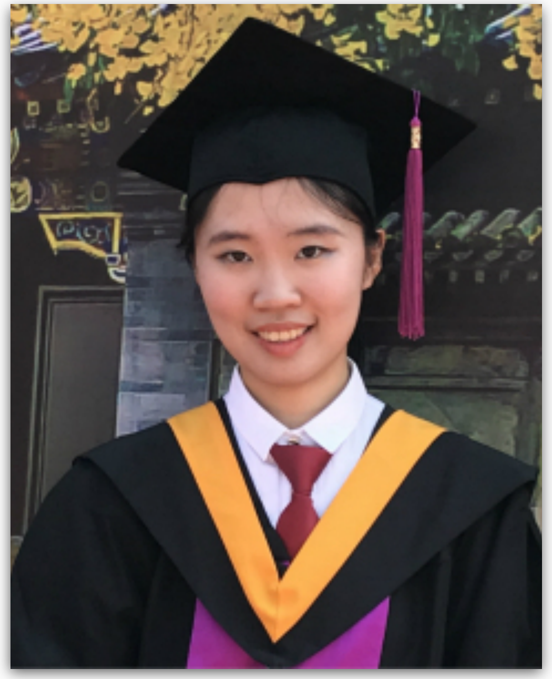

Qian Xie
Cornell

Alex Terenin
Cornell

Collaborators

NeurIPS 2024: *Cost-aware Bayesian Optimization via the Pandora's Box Gittins Index*

Raul Astudillo
Caltech

Peter Frazier
Cornell

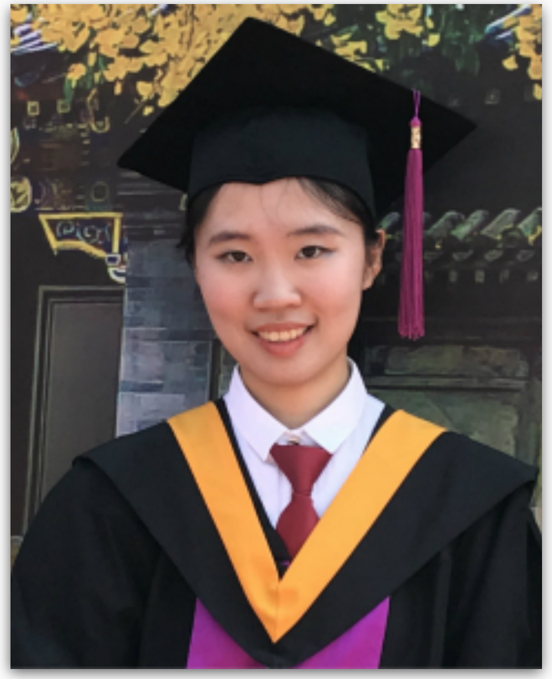

Qian Xie
Cornell

Alex Terenin
Cornell

Theo Brown
UCL/UKAE

Linda Cai
UC Berkeley

work in progress

What is **BayesOpt**?

Unknown function:

$f : [0, 1] \to \mathbb{R}$

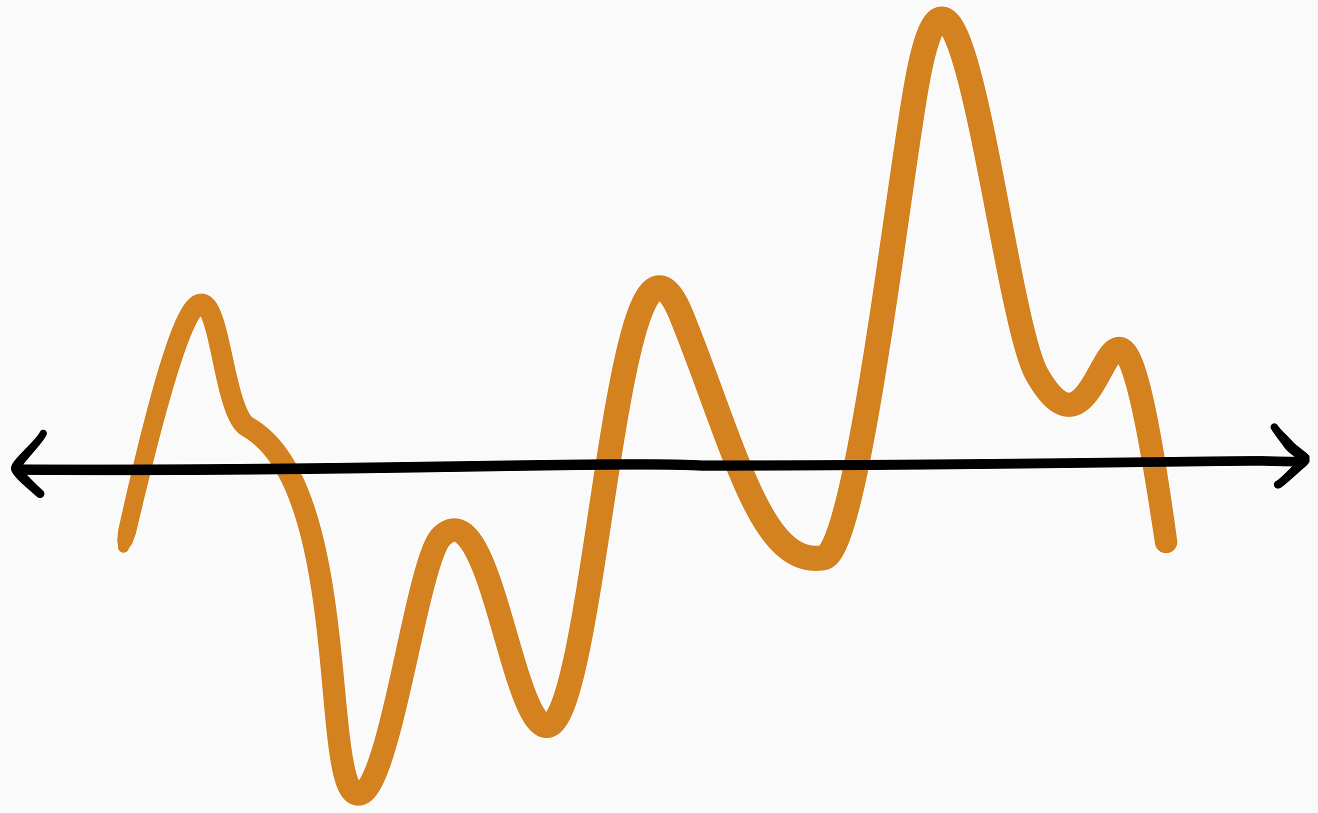

What is **BayesOpt**?

Assumption:
$f \sim \mathrm{GP}$

Unknown function:
$f : [0, 1] \to \mathbb{R}$

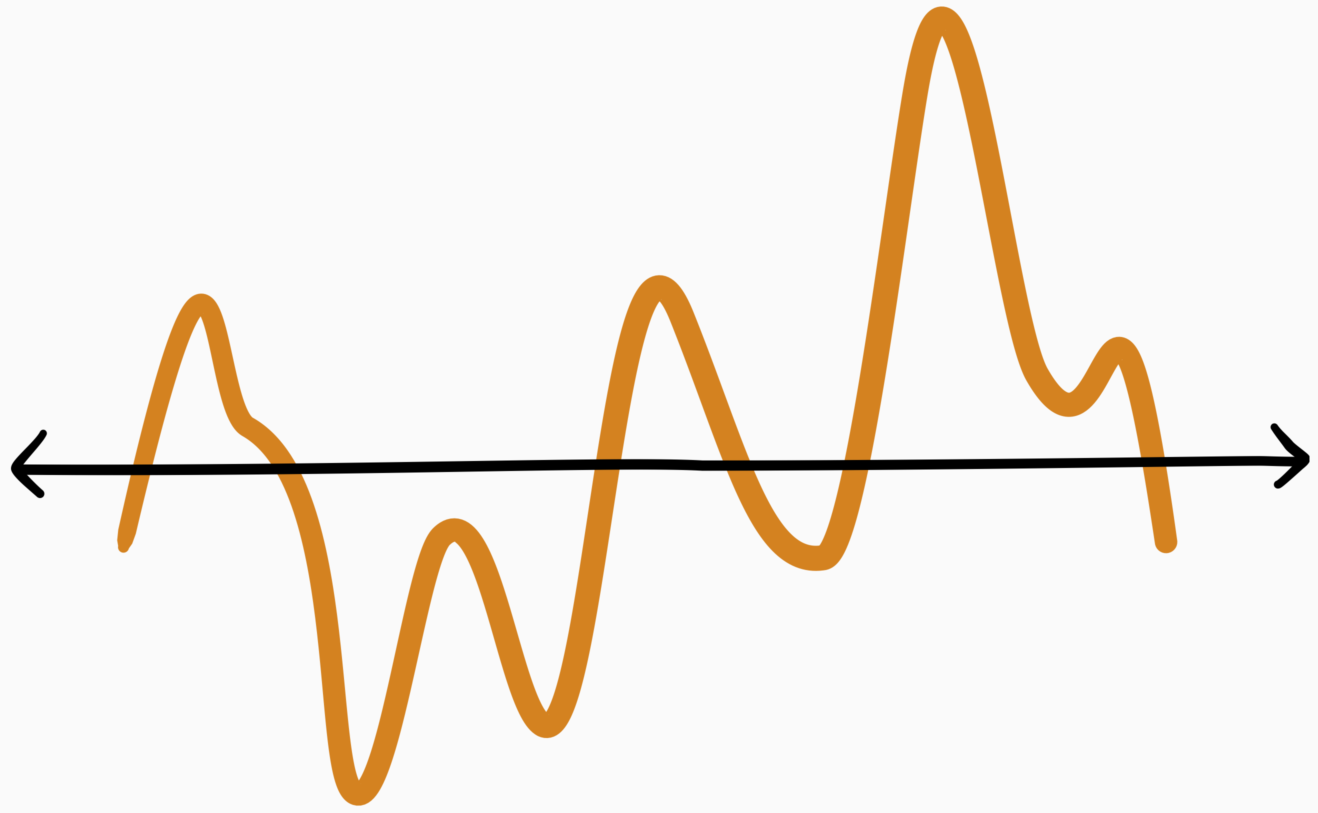

What is **BayesOpt**?

Assumption:
$f \sim \mathrm{GP}$

Unknown function:
$f : [0, 1] \to \mathbb{R}$

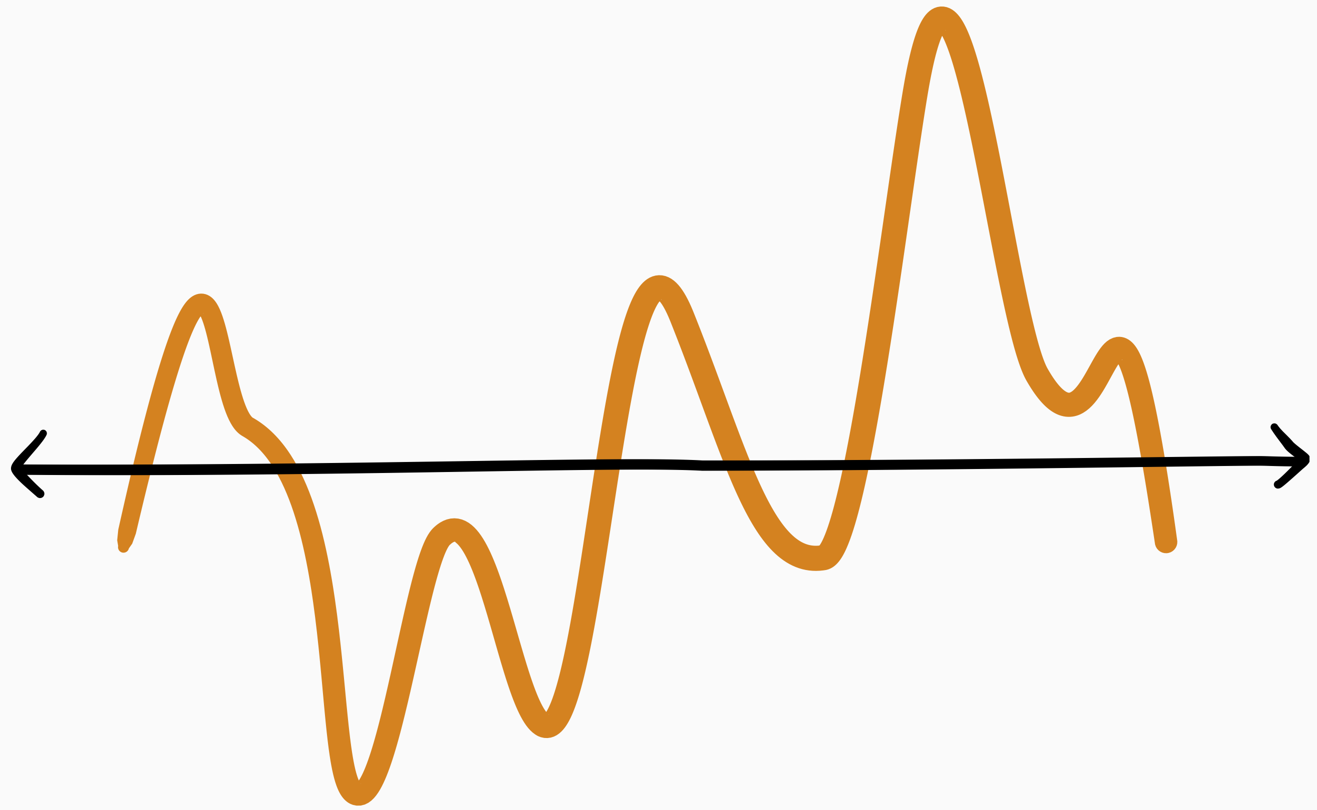

Goal: maximize

$$\mathbf{E}\Big[\max_{1 \le t \le T} f(x_t)\Big]$$

What is BayesOpt?

Assumption:
$f \sim \mathrm{GP}$

Unknown function:
$f : [0, 1] \to \mathbb{R}$

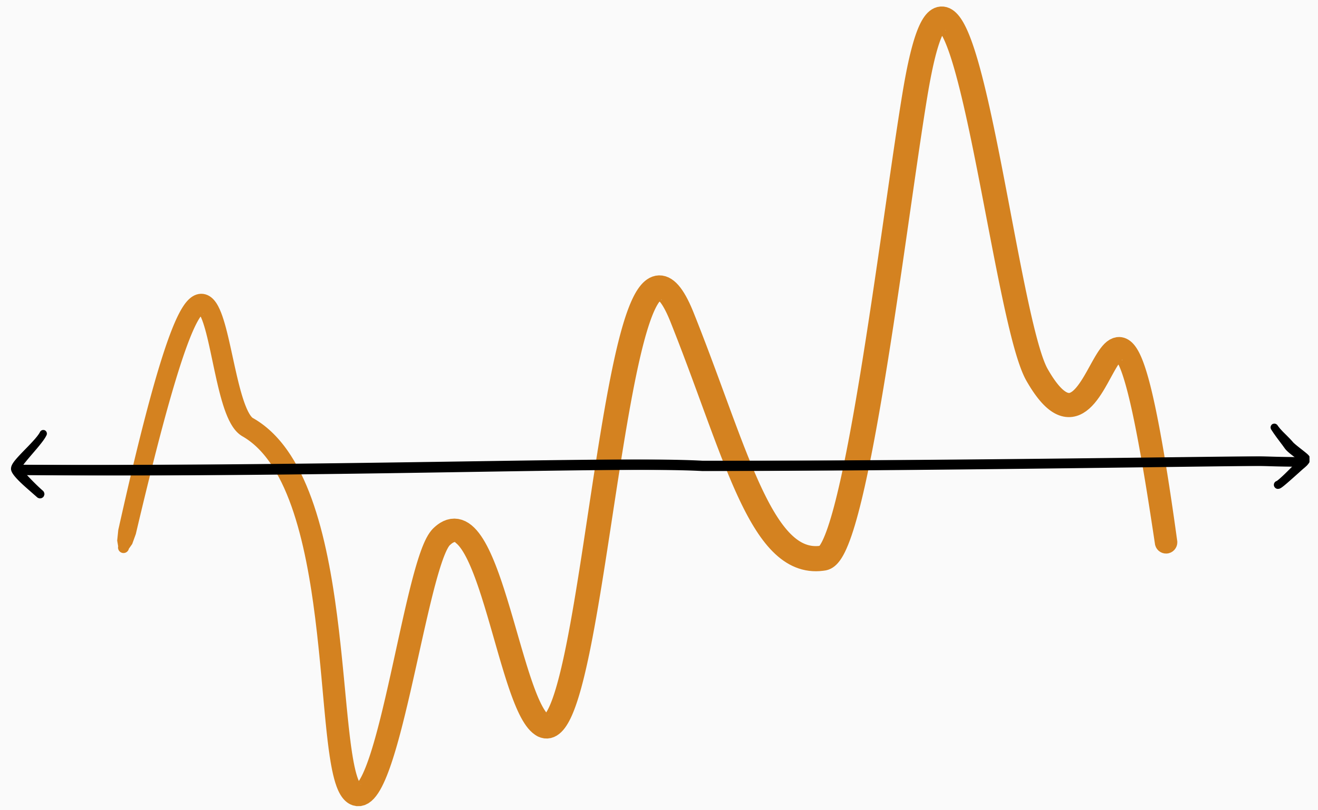

Goal: maximize

$$\mathbf{E}\Big[\max_{1 \leq t \leq T} f(x_t)\Big]$$

T fixed

x_t chosen

What happens during **BayesOpt**?

Time 0

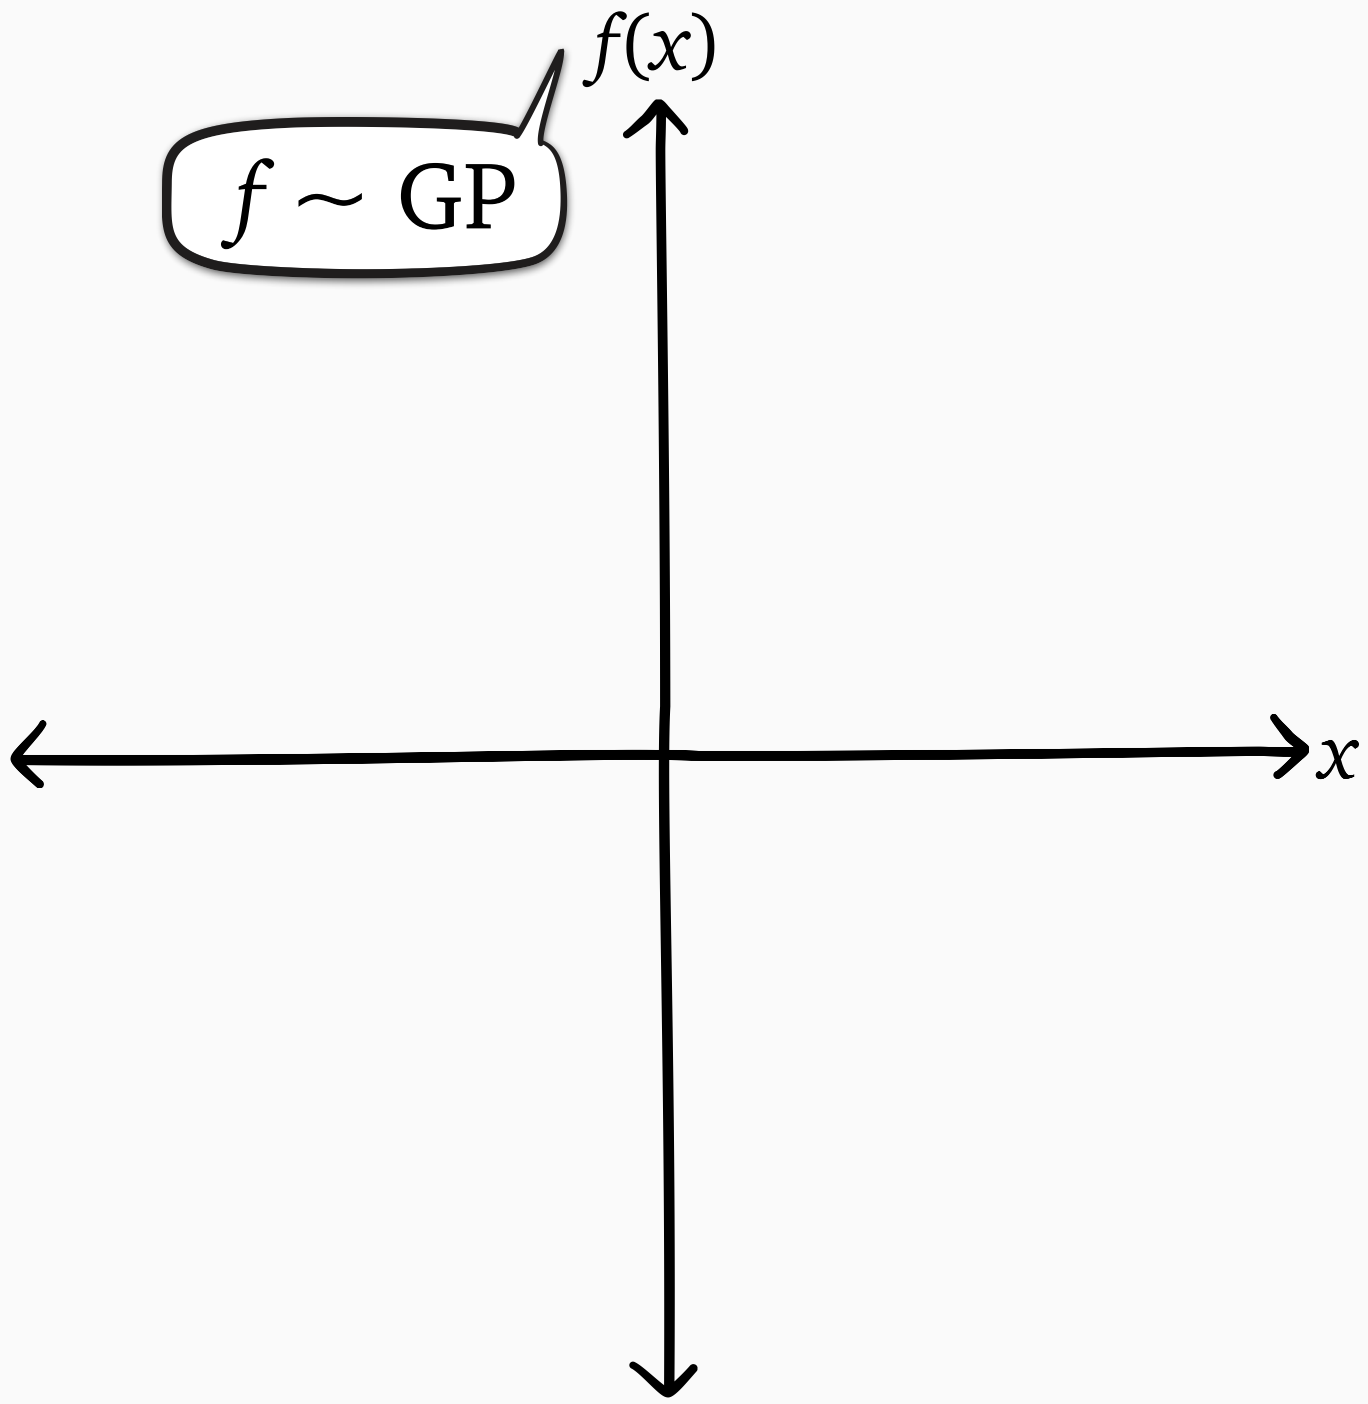

Time 1

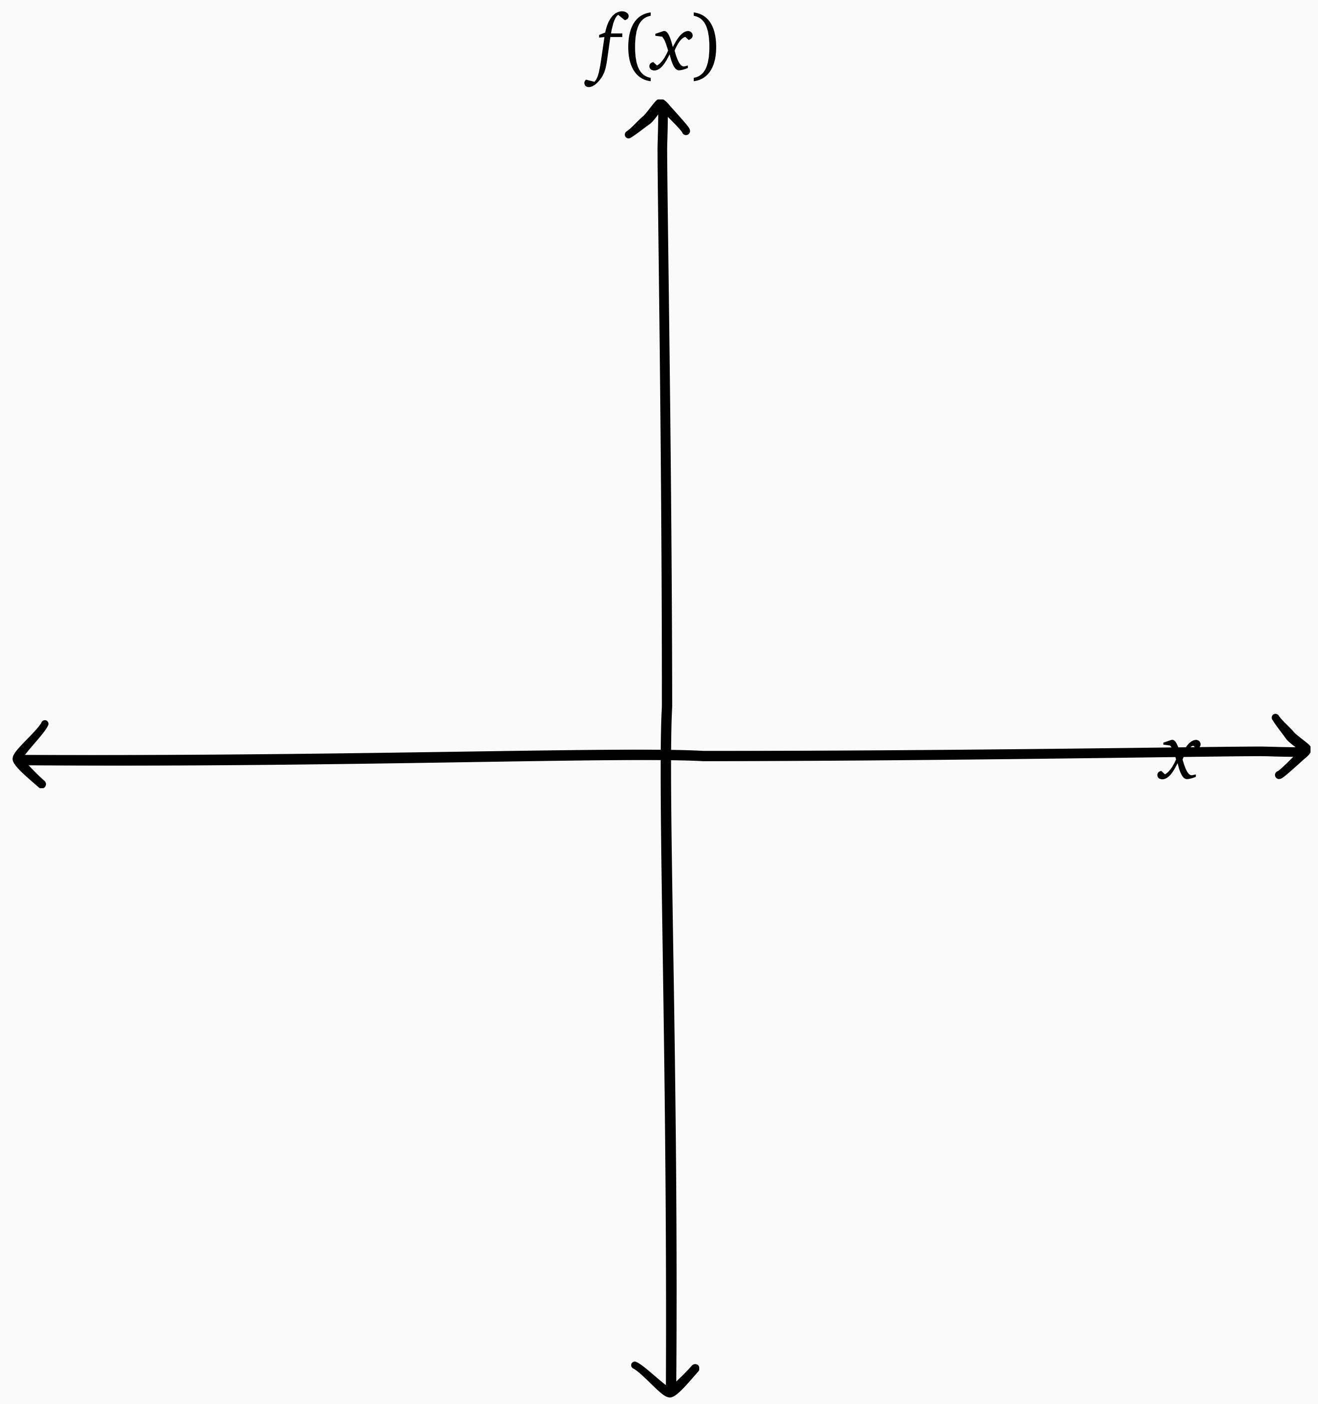

What happens during BayesOpt?

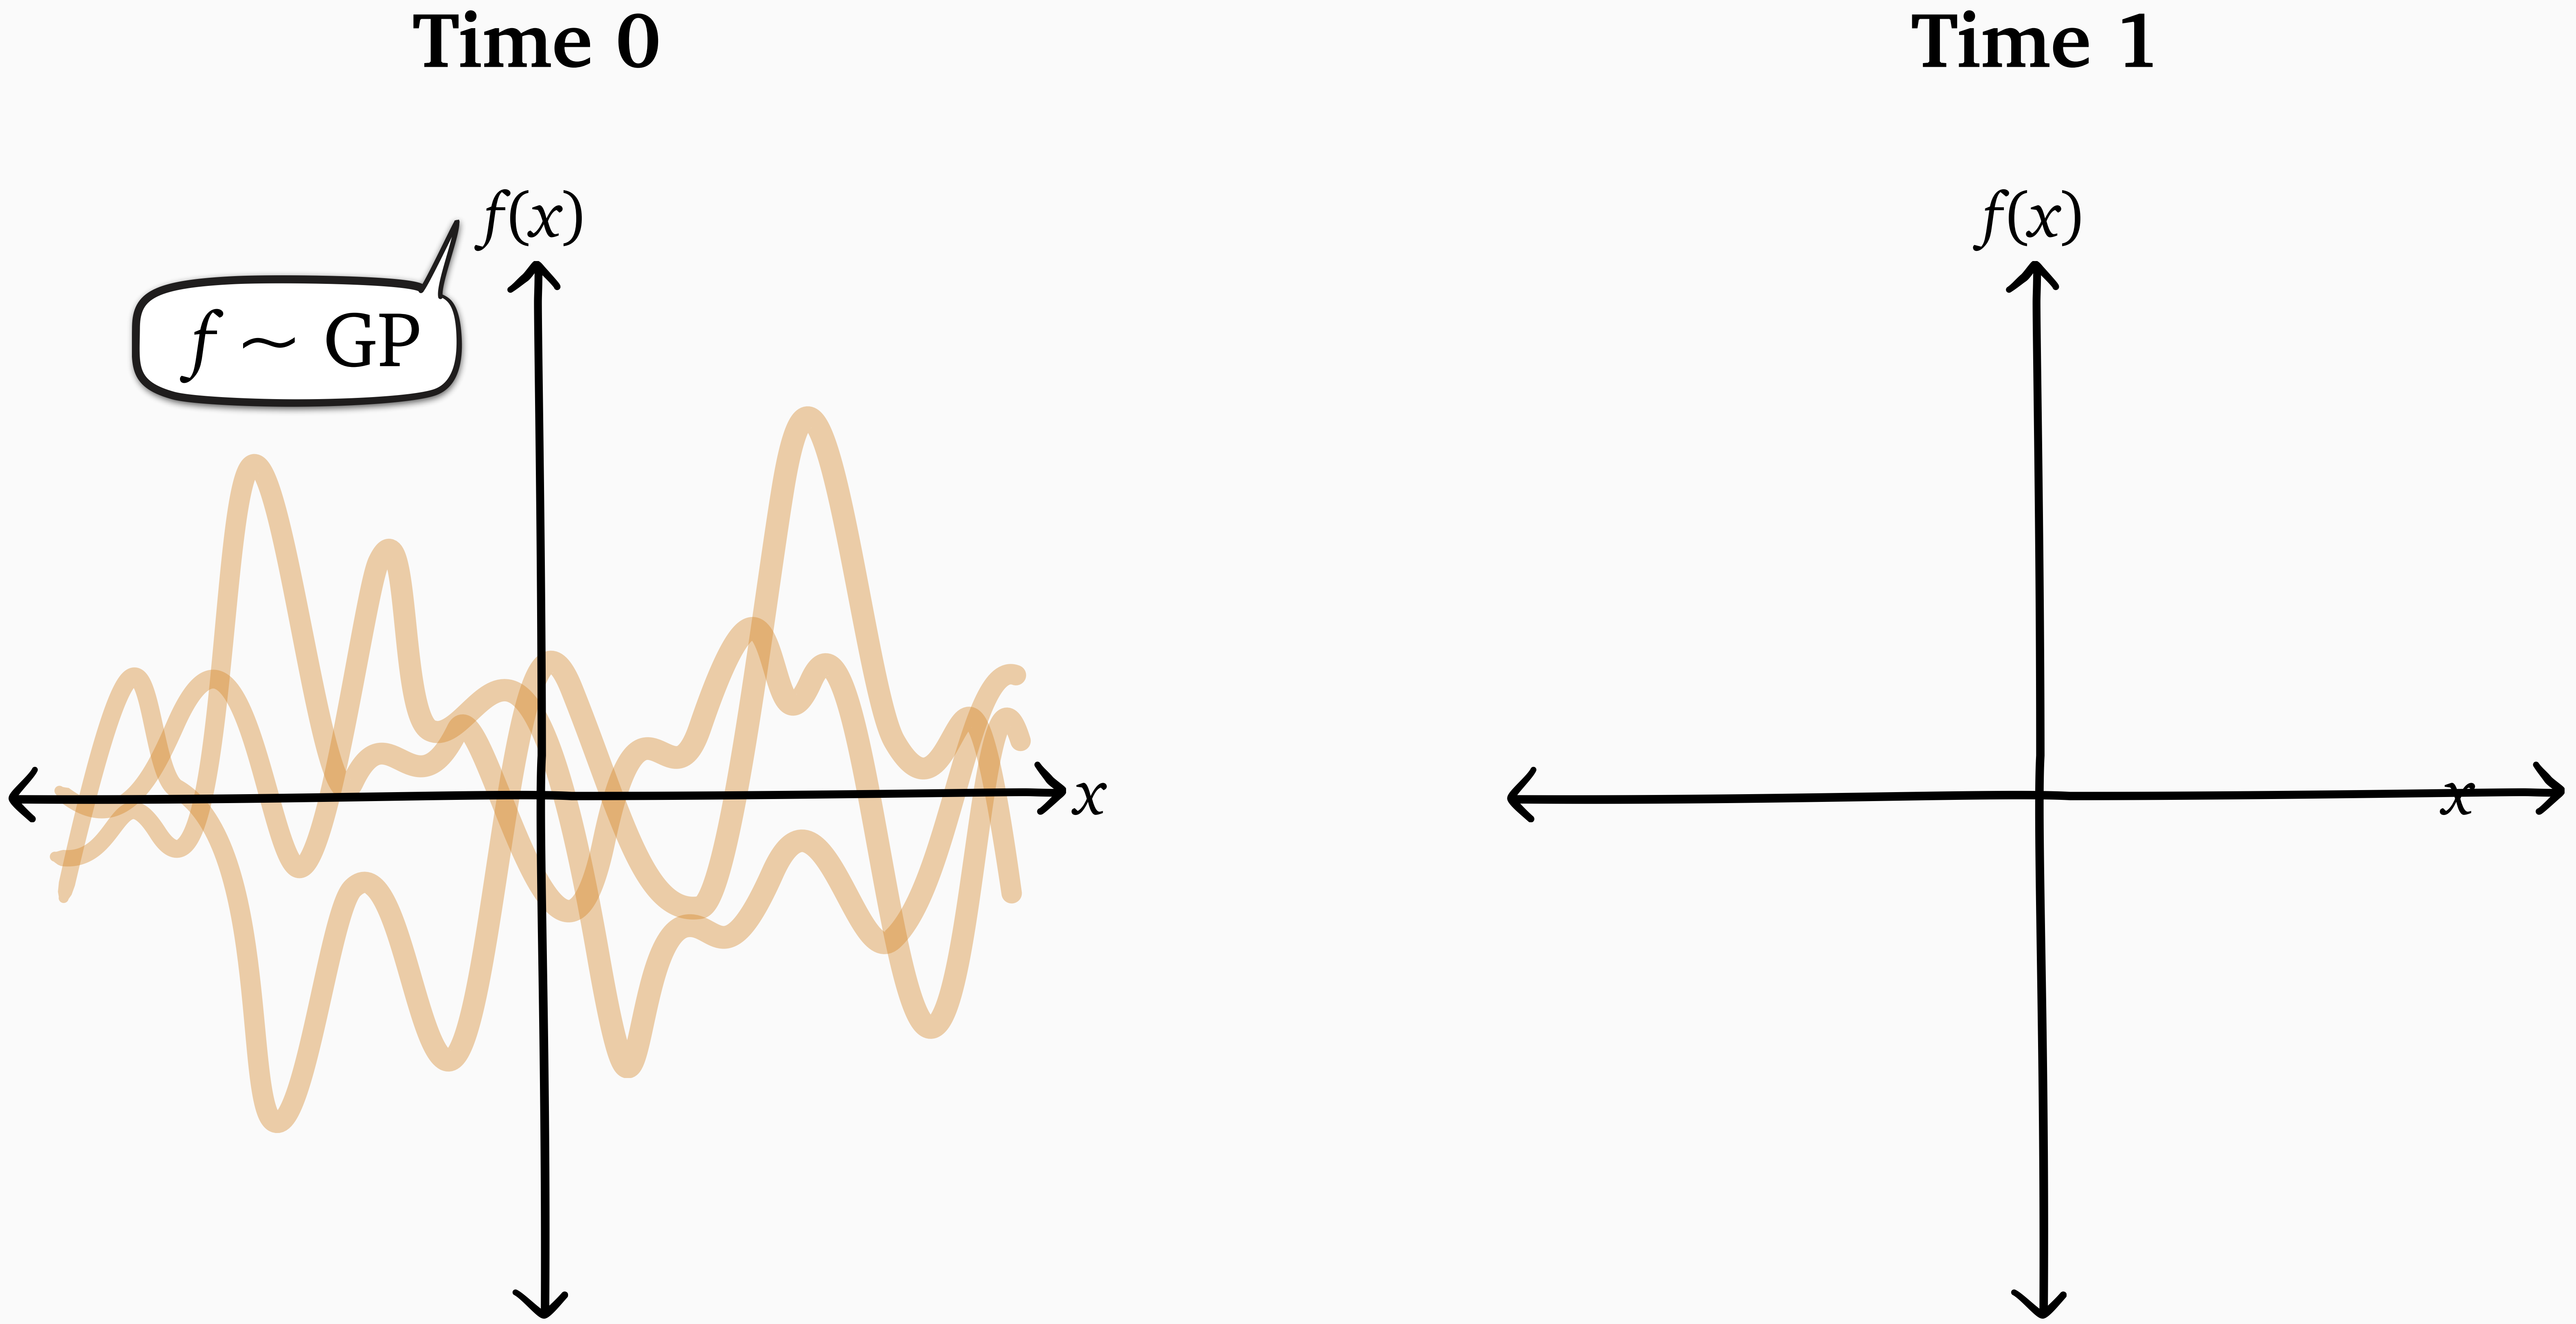

What happens during **BayesOpt**?

Time 0

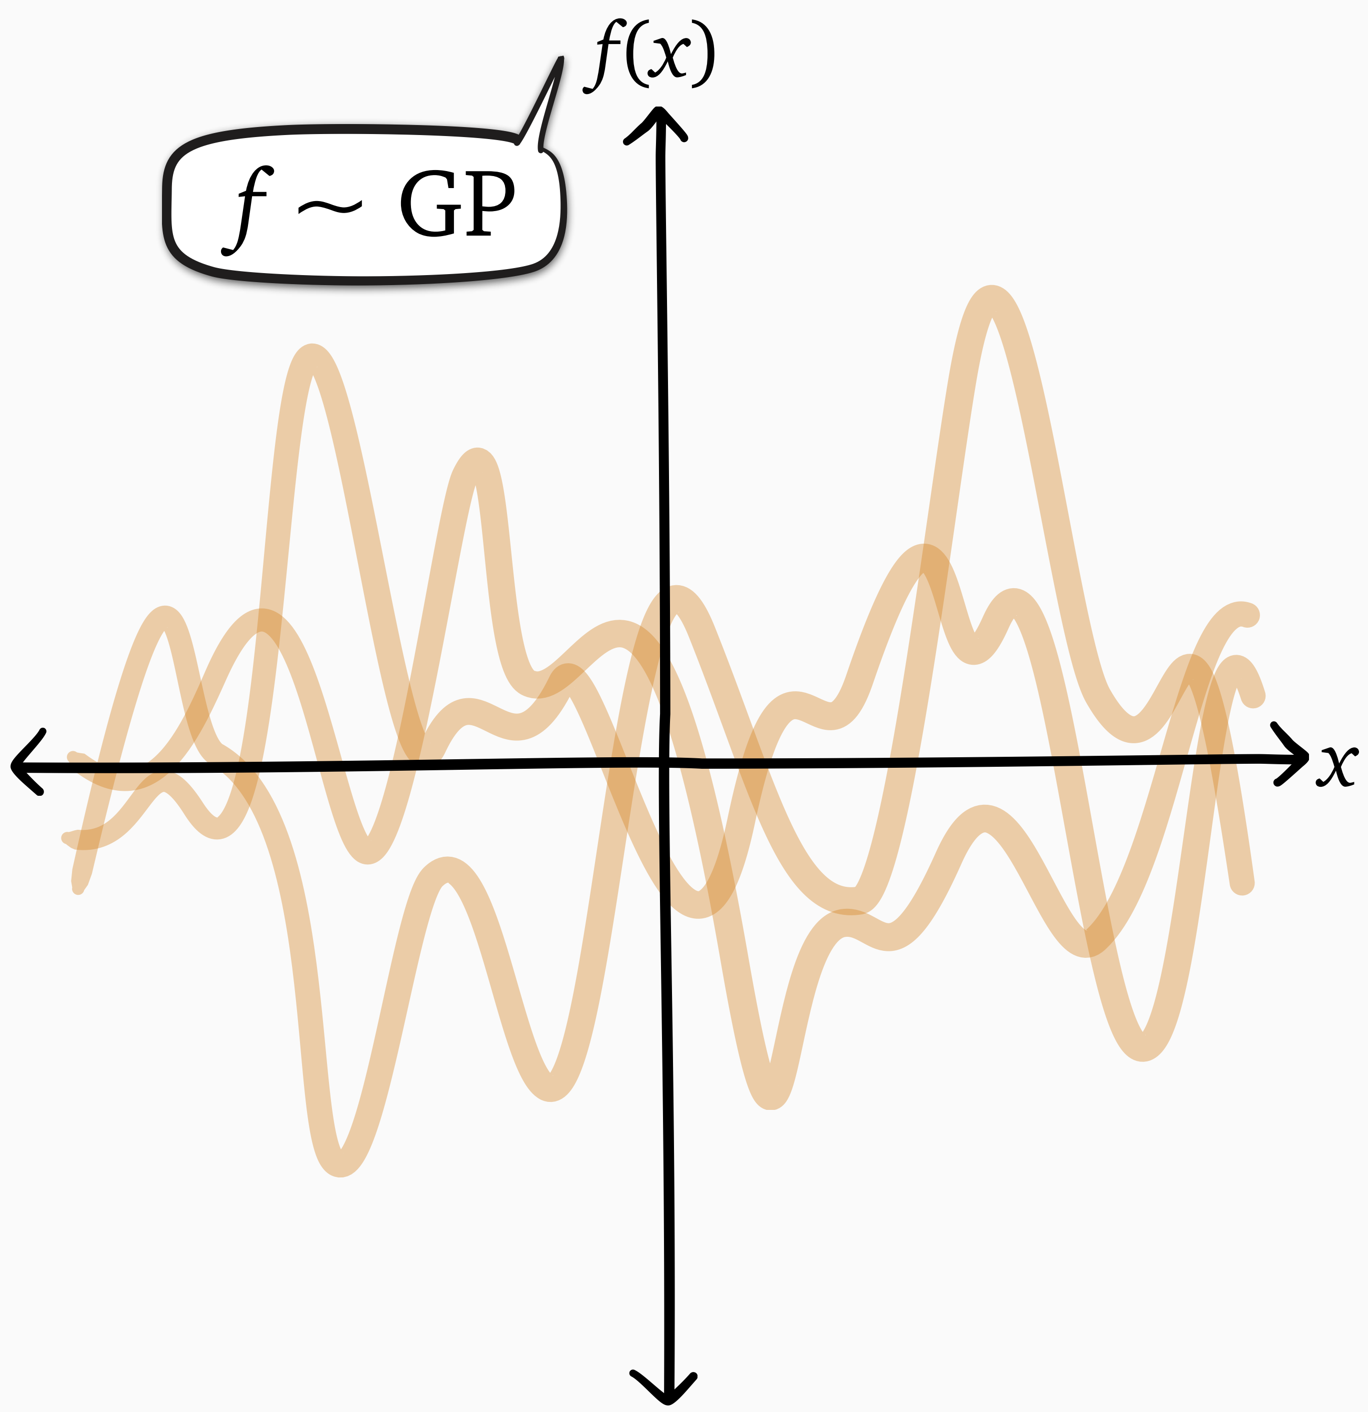

Time 1

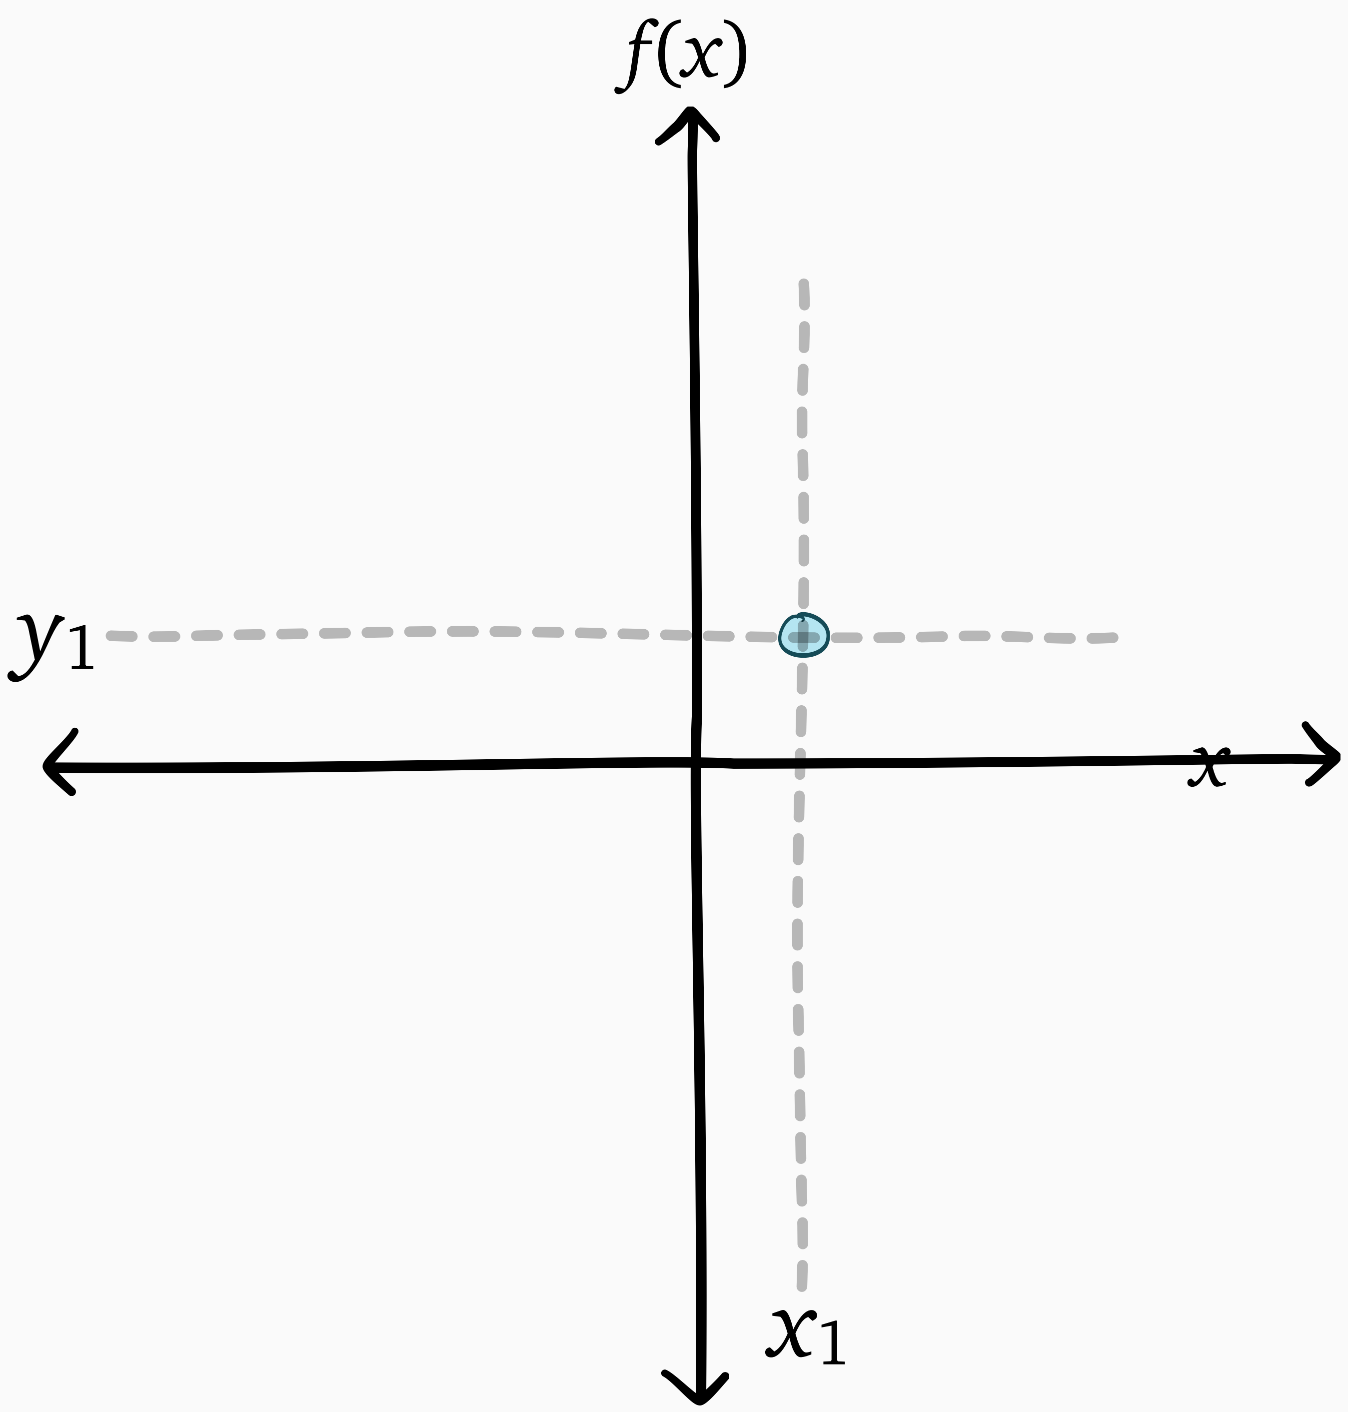

What happens during BayesOpt?

Time 0

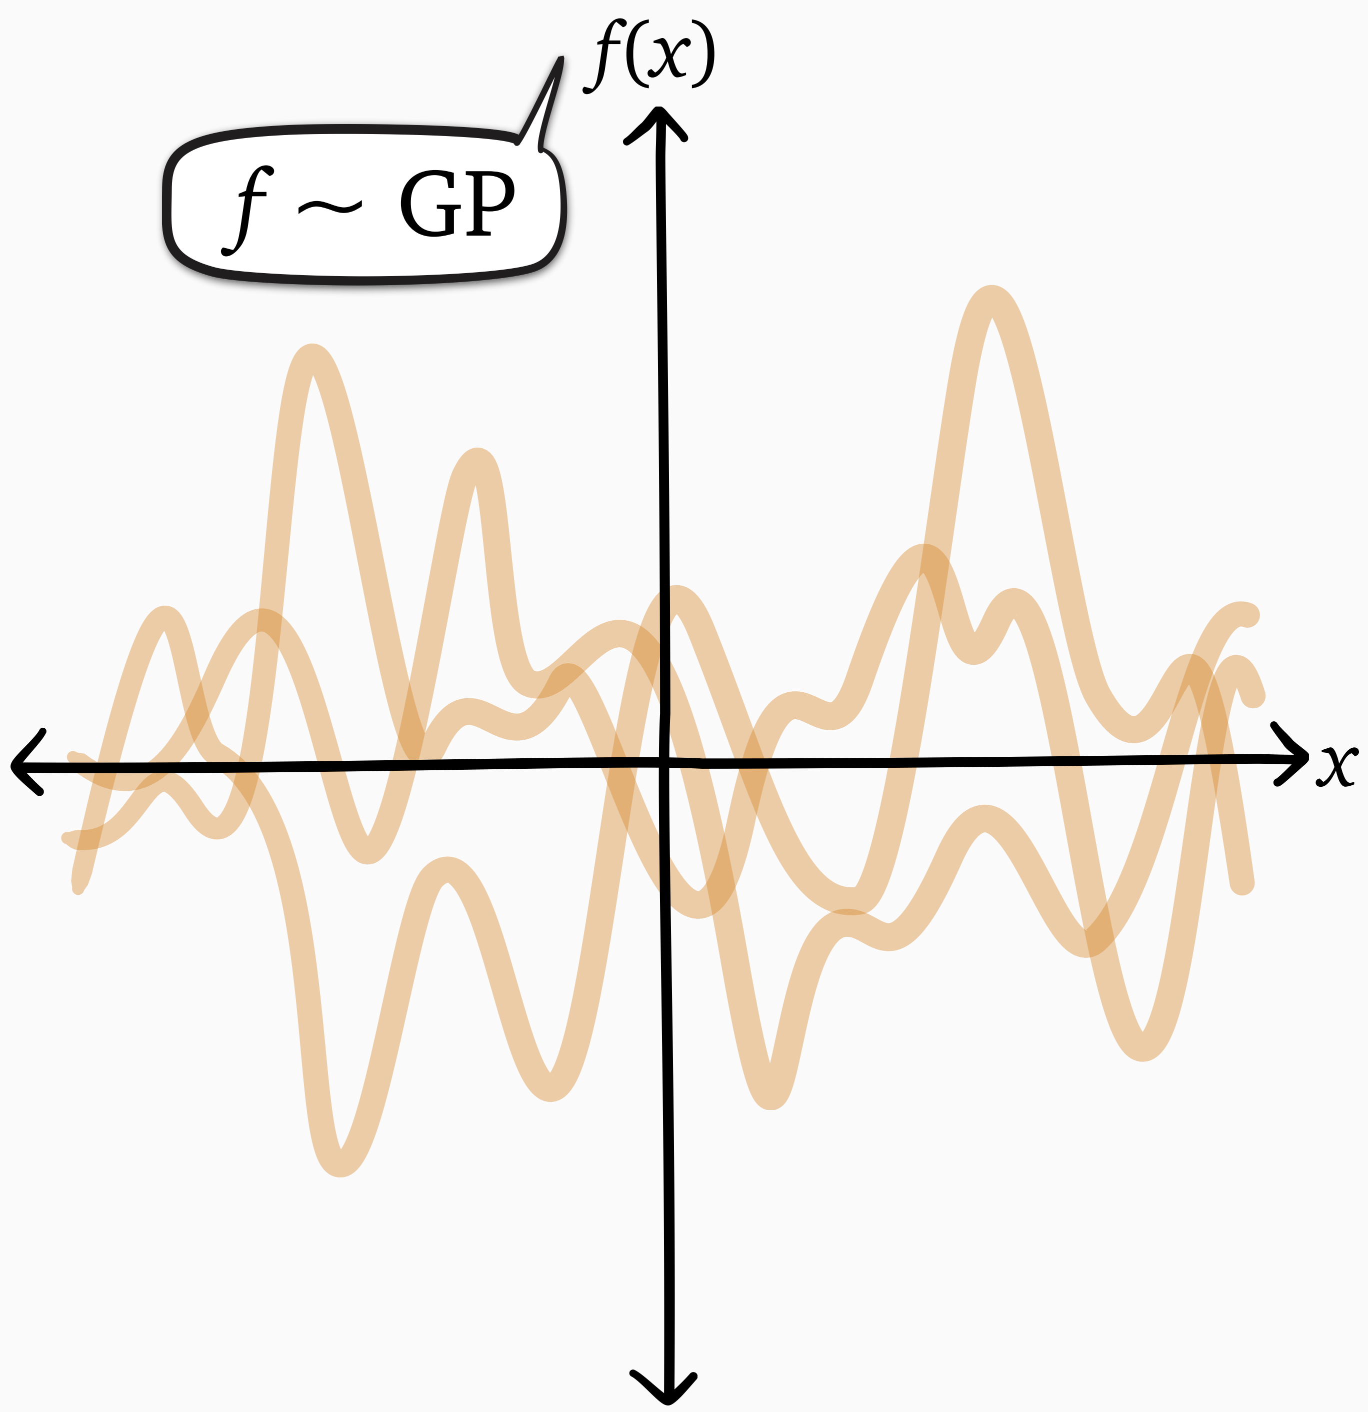

Time 1

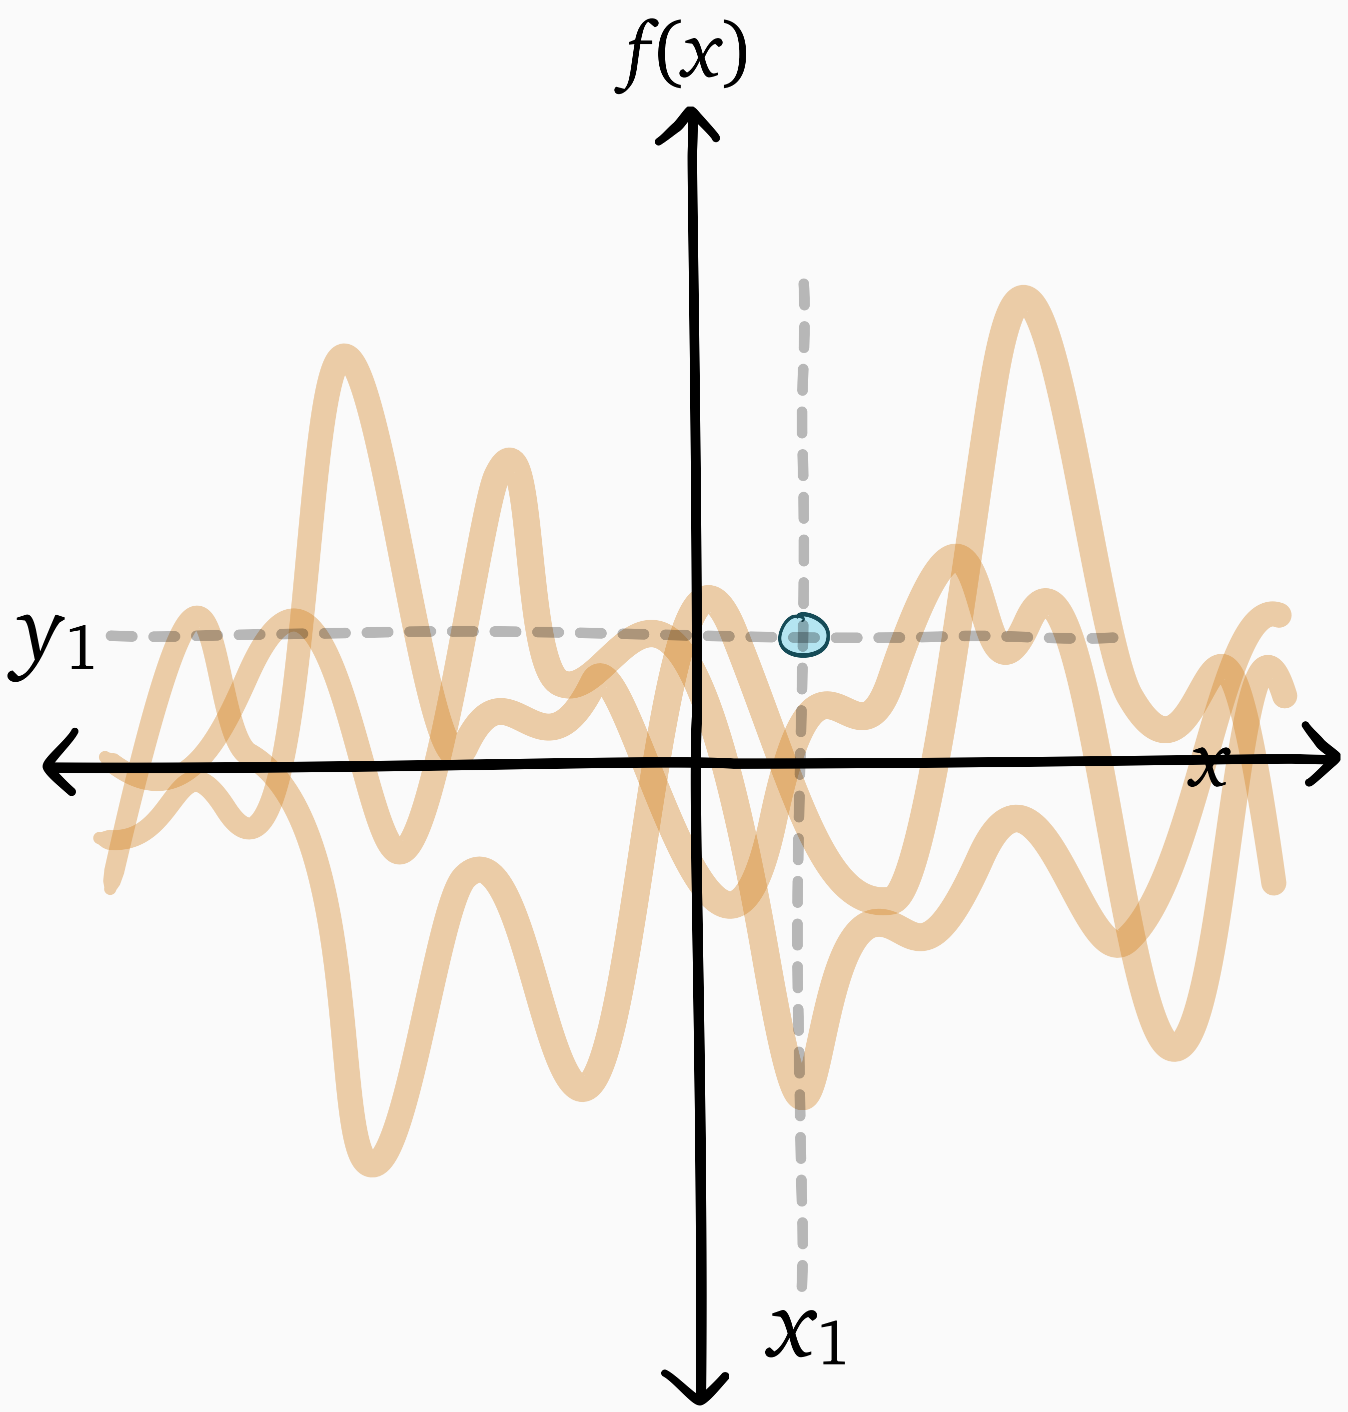

What happens during BayesOpt?

Time 0

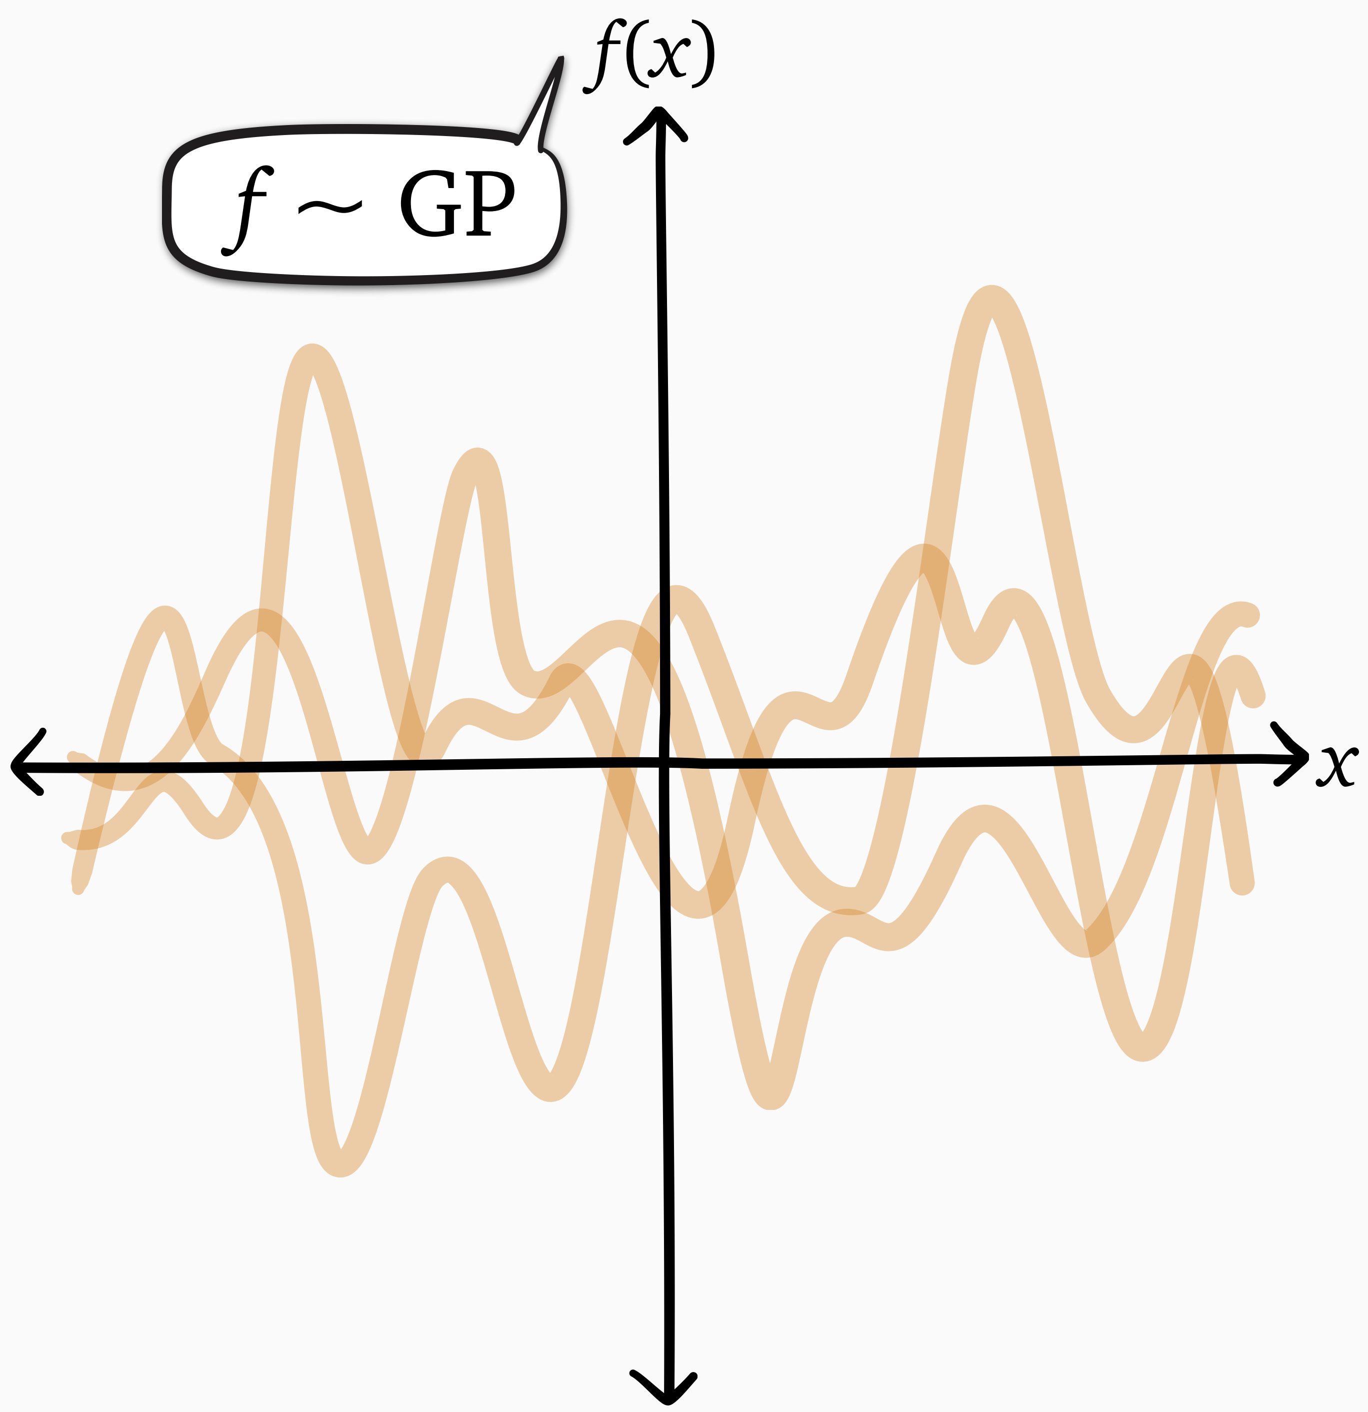

Time 1

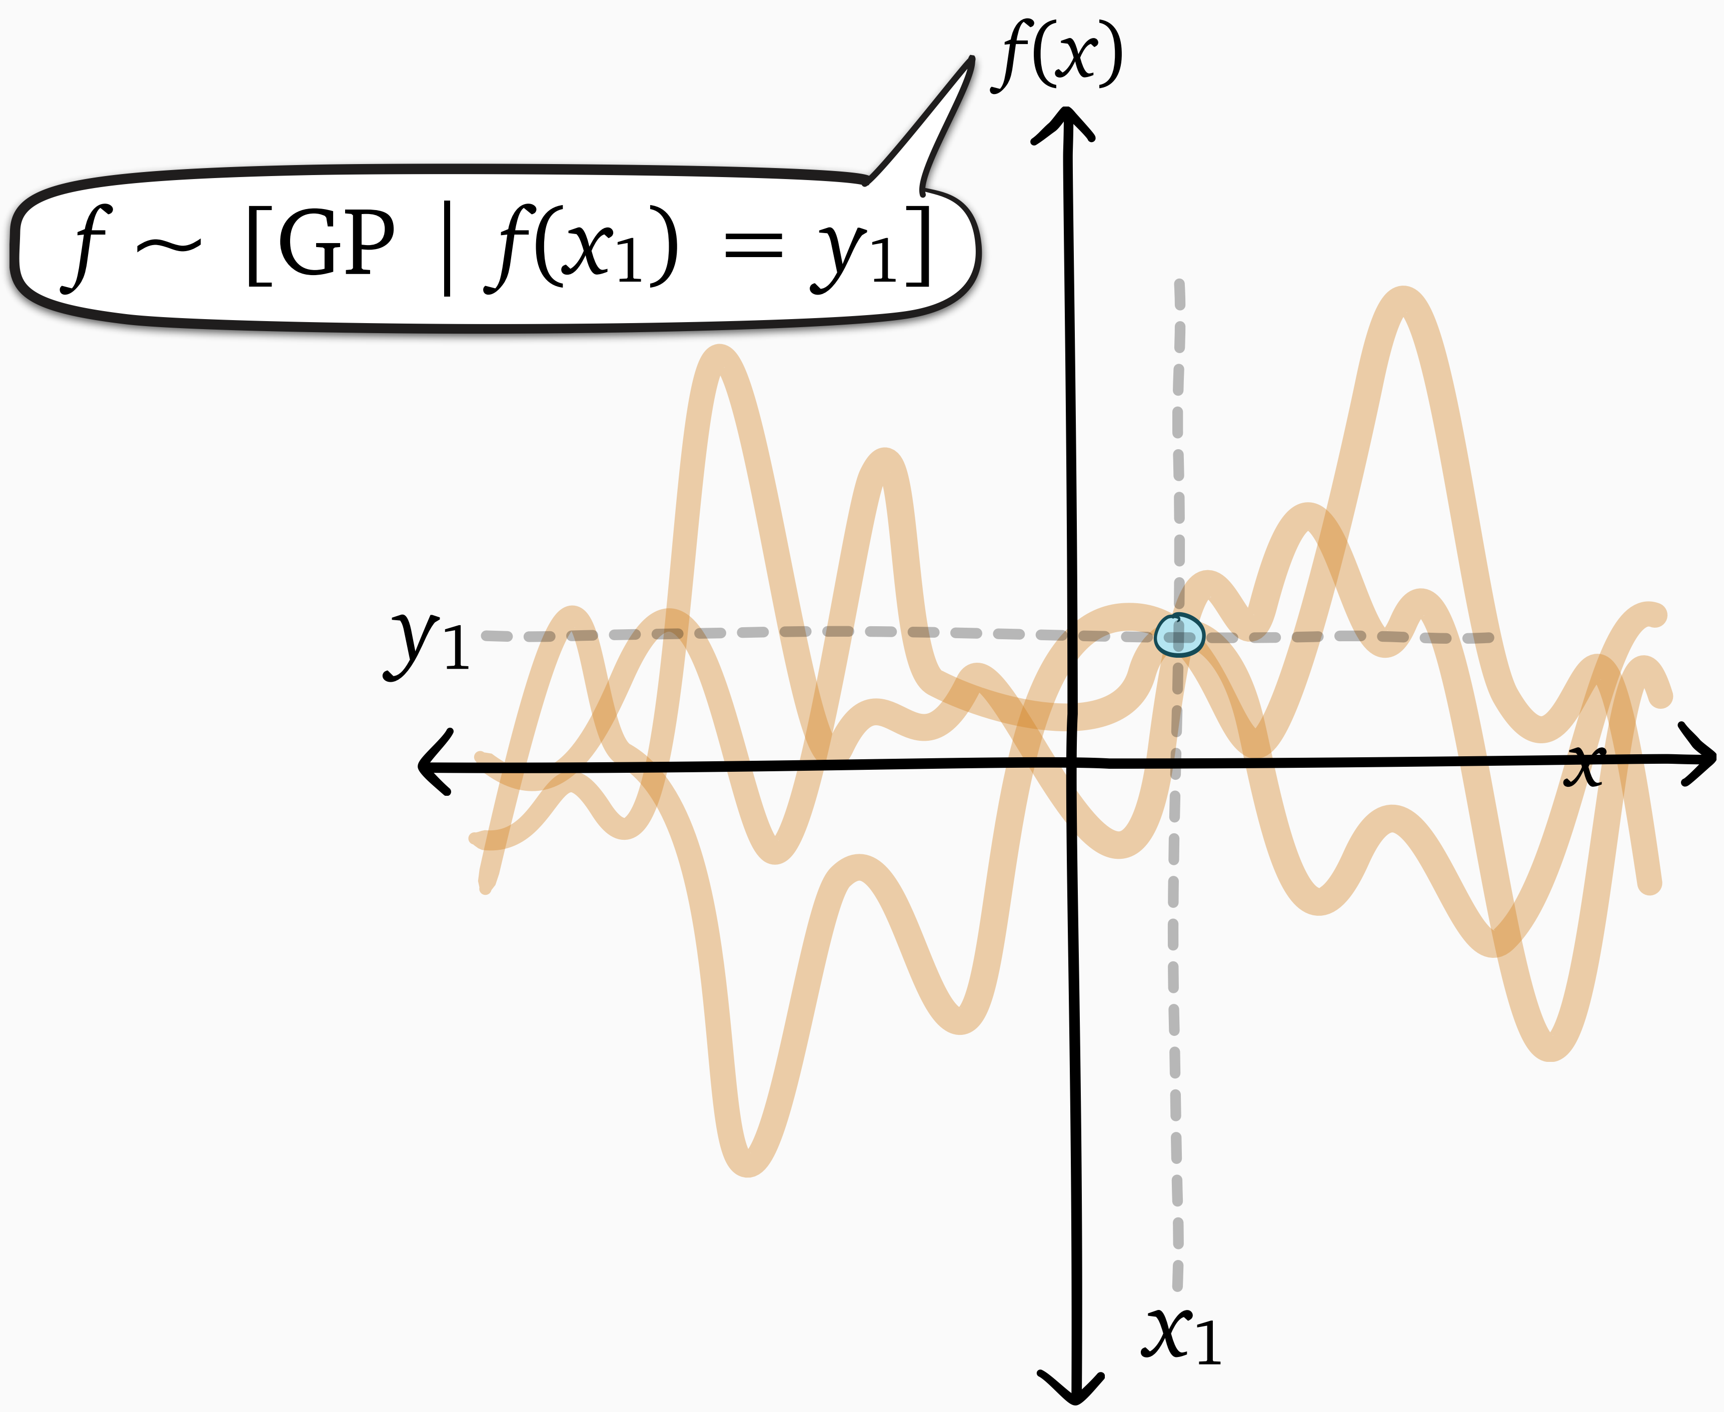

What happens during BayesOpt?

Goal: maximize

$$\mathbf{E}\left[\max_{1 \le t \le T} f(x_t)\right]$$

Question: how to choose x_t?

Time 1

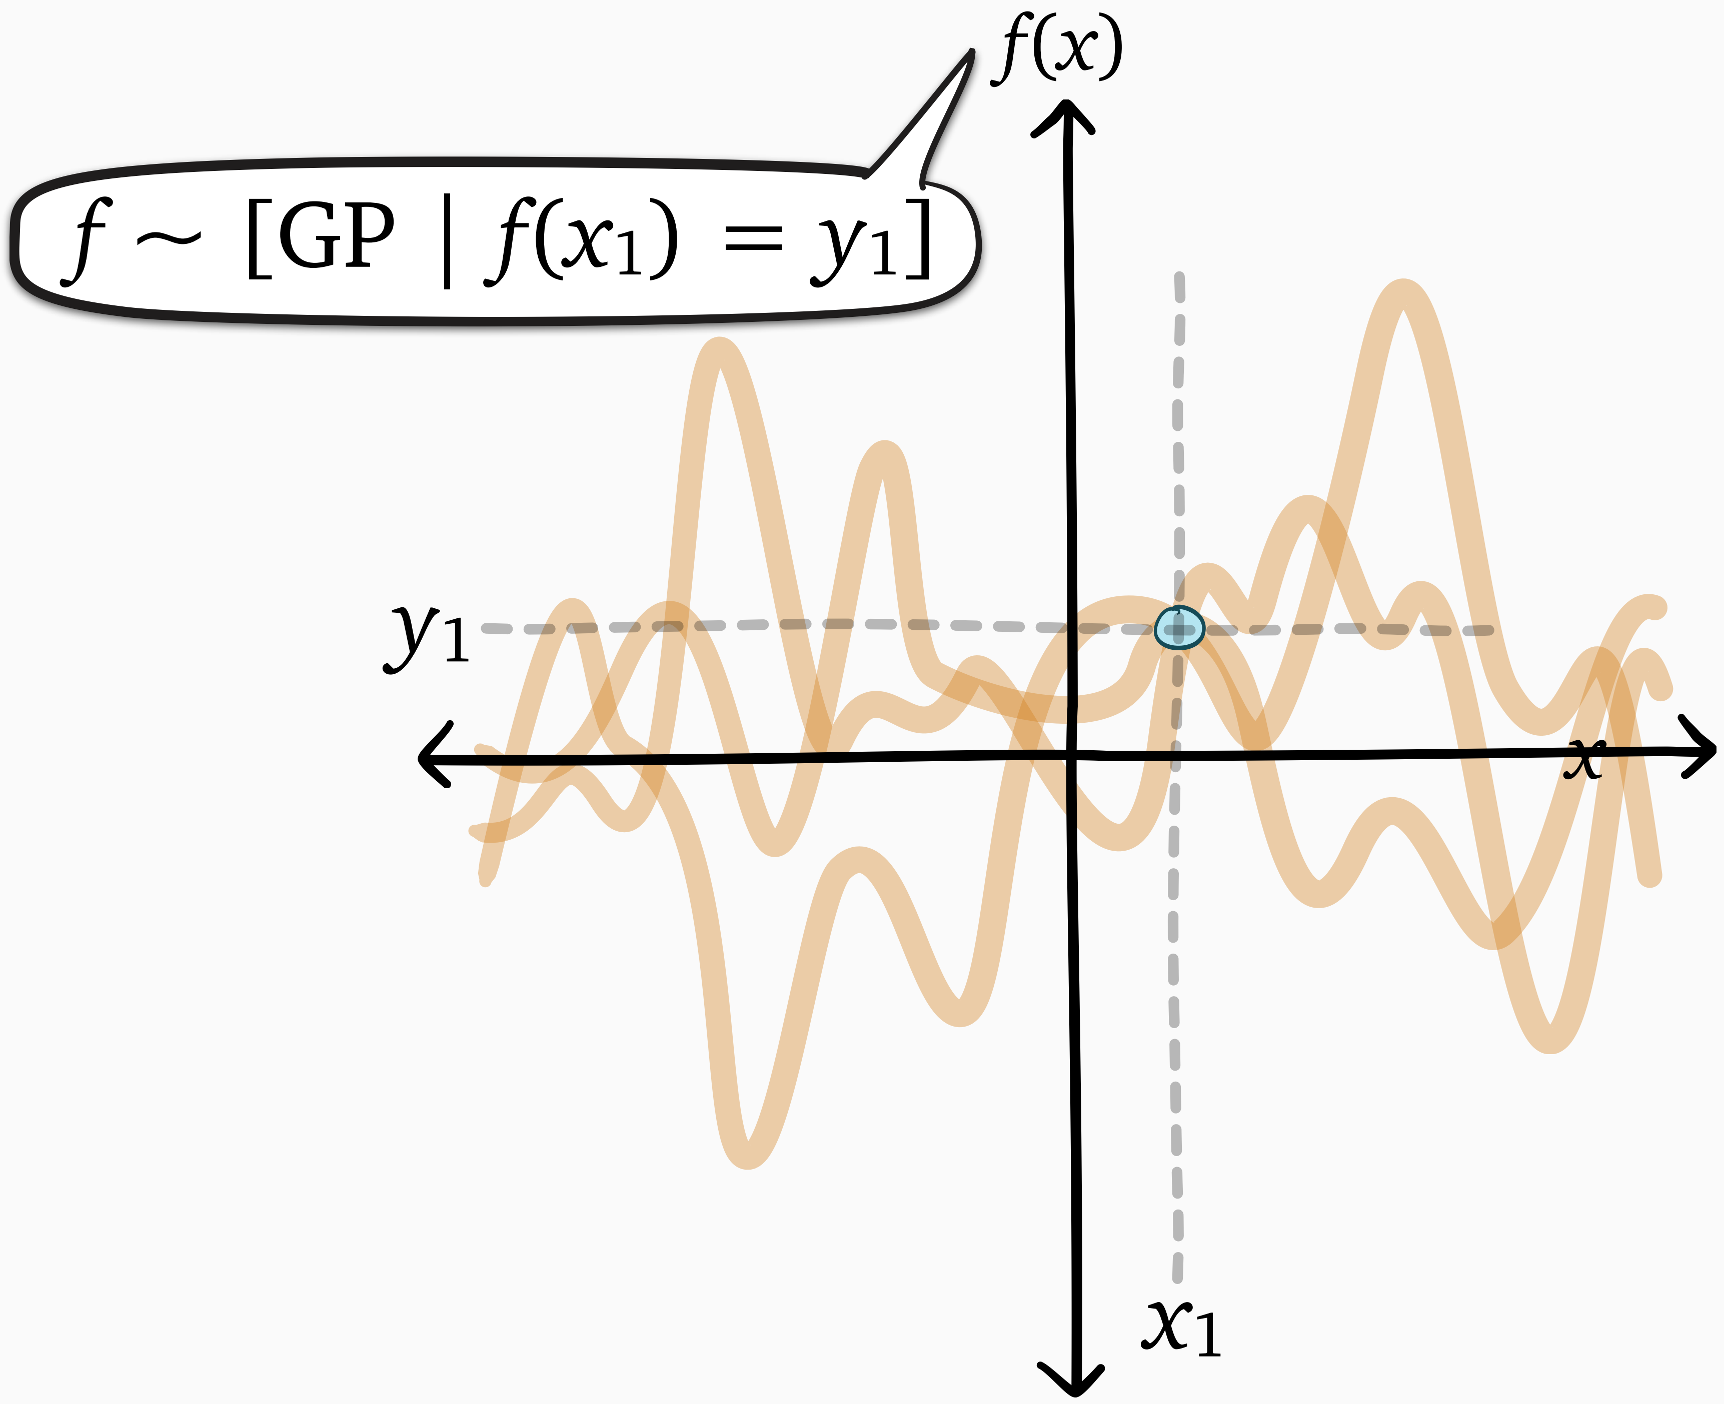

What happens during BayesOpt?

Goal: maximize

$$\mathbf{E}\left[\max_{1\le t\le T} f(x_t)\right]$$

Question: how to choose x_t?

- **expected improvement (EI)**
- upper confidence bound (UCB)
- Thompson sampling (TS)
- knowledge gradient (KG)
- entropy methods, lookahead, …

Time 1

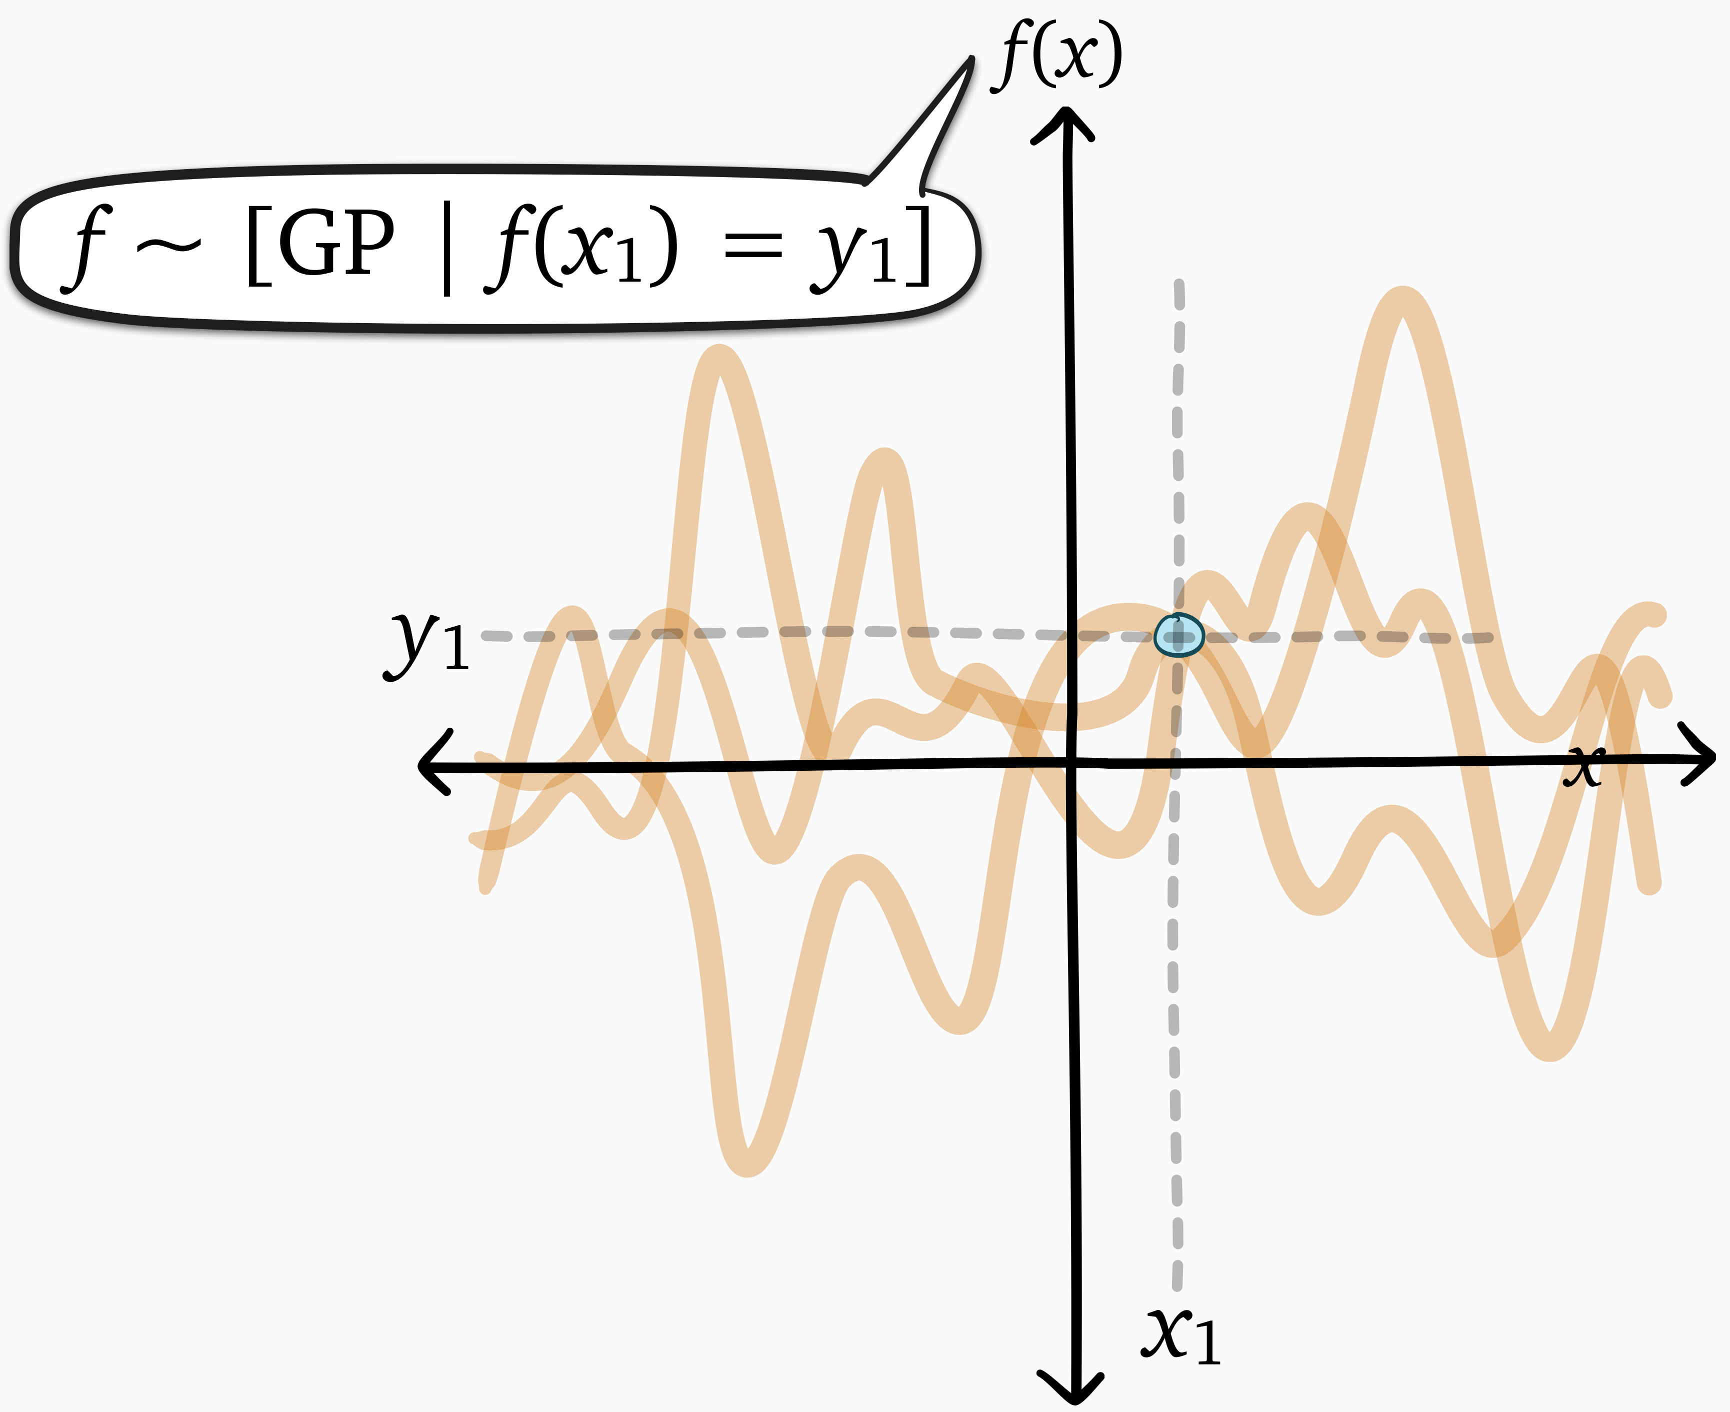

Expected improvement for BayesOpt

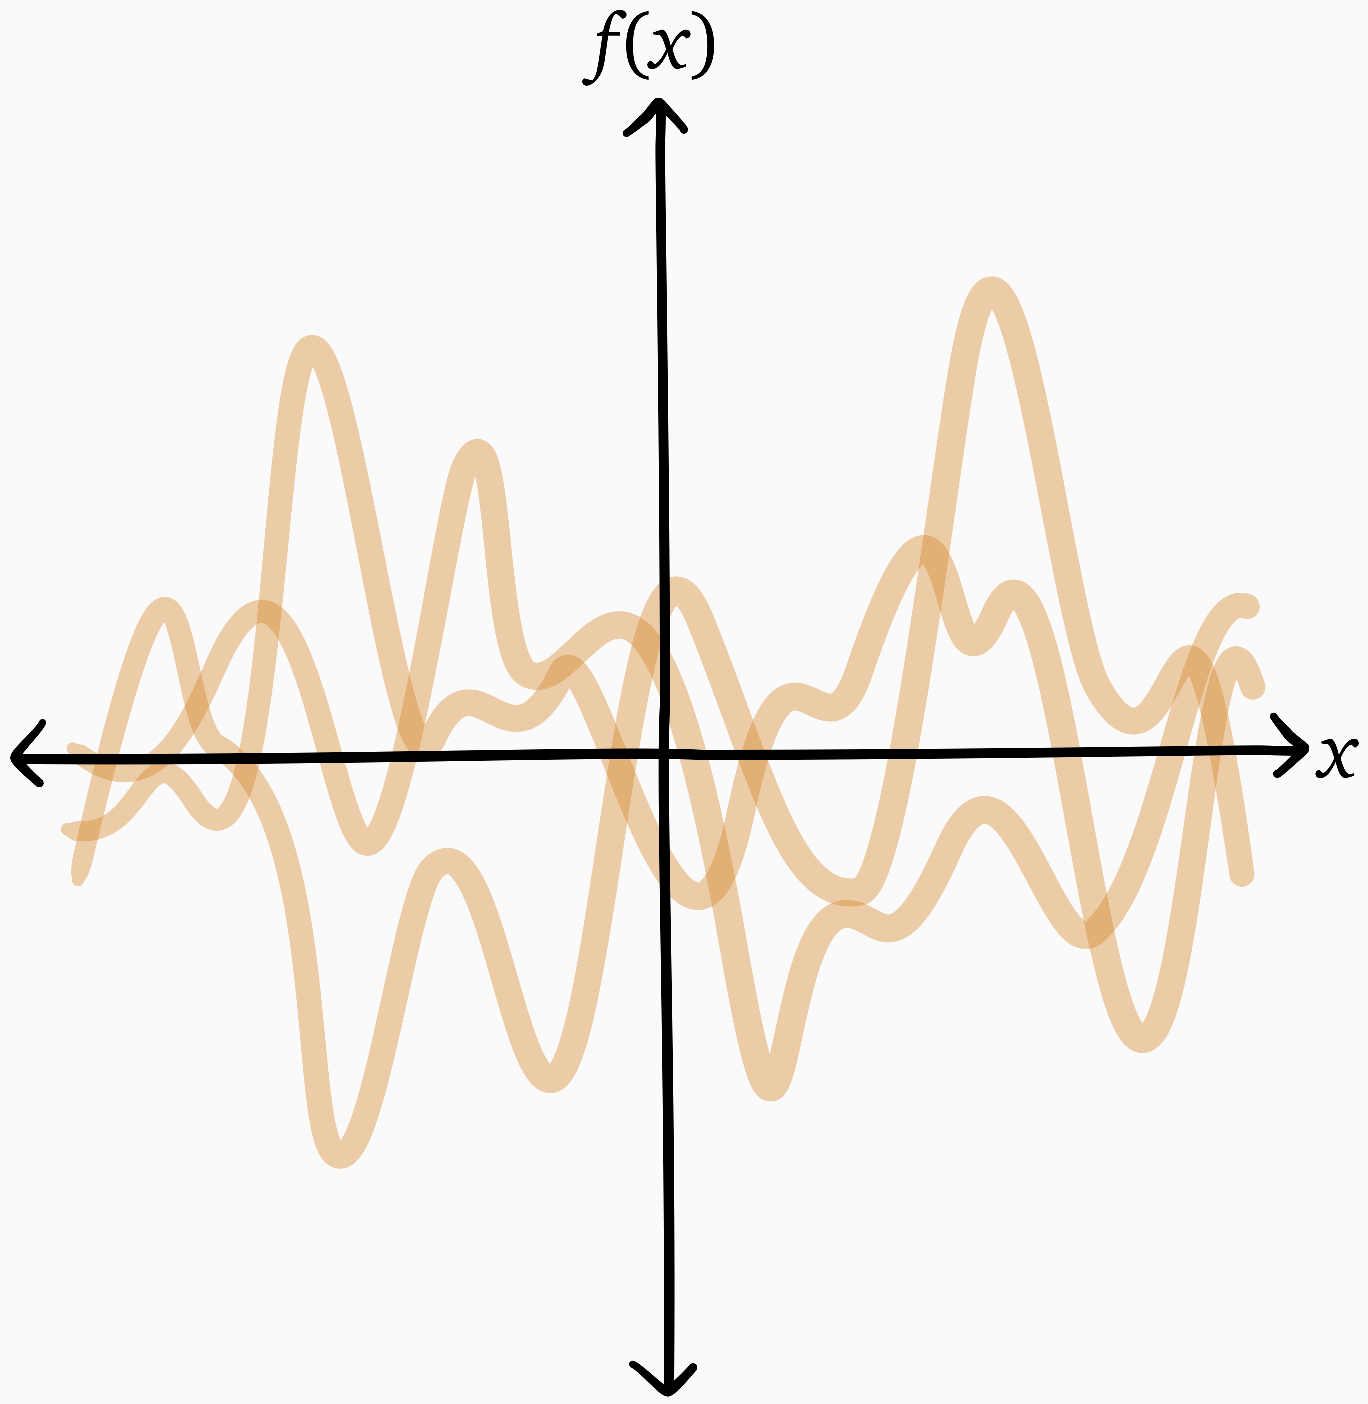

Expected improvement for BayesOpt

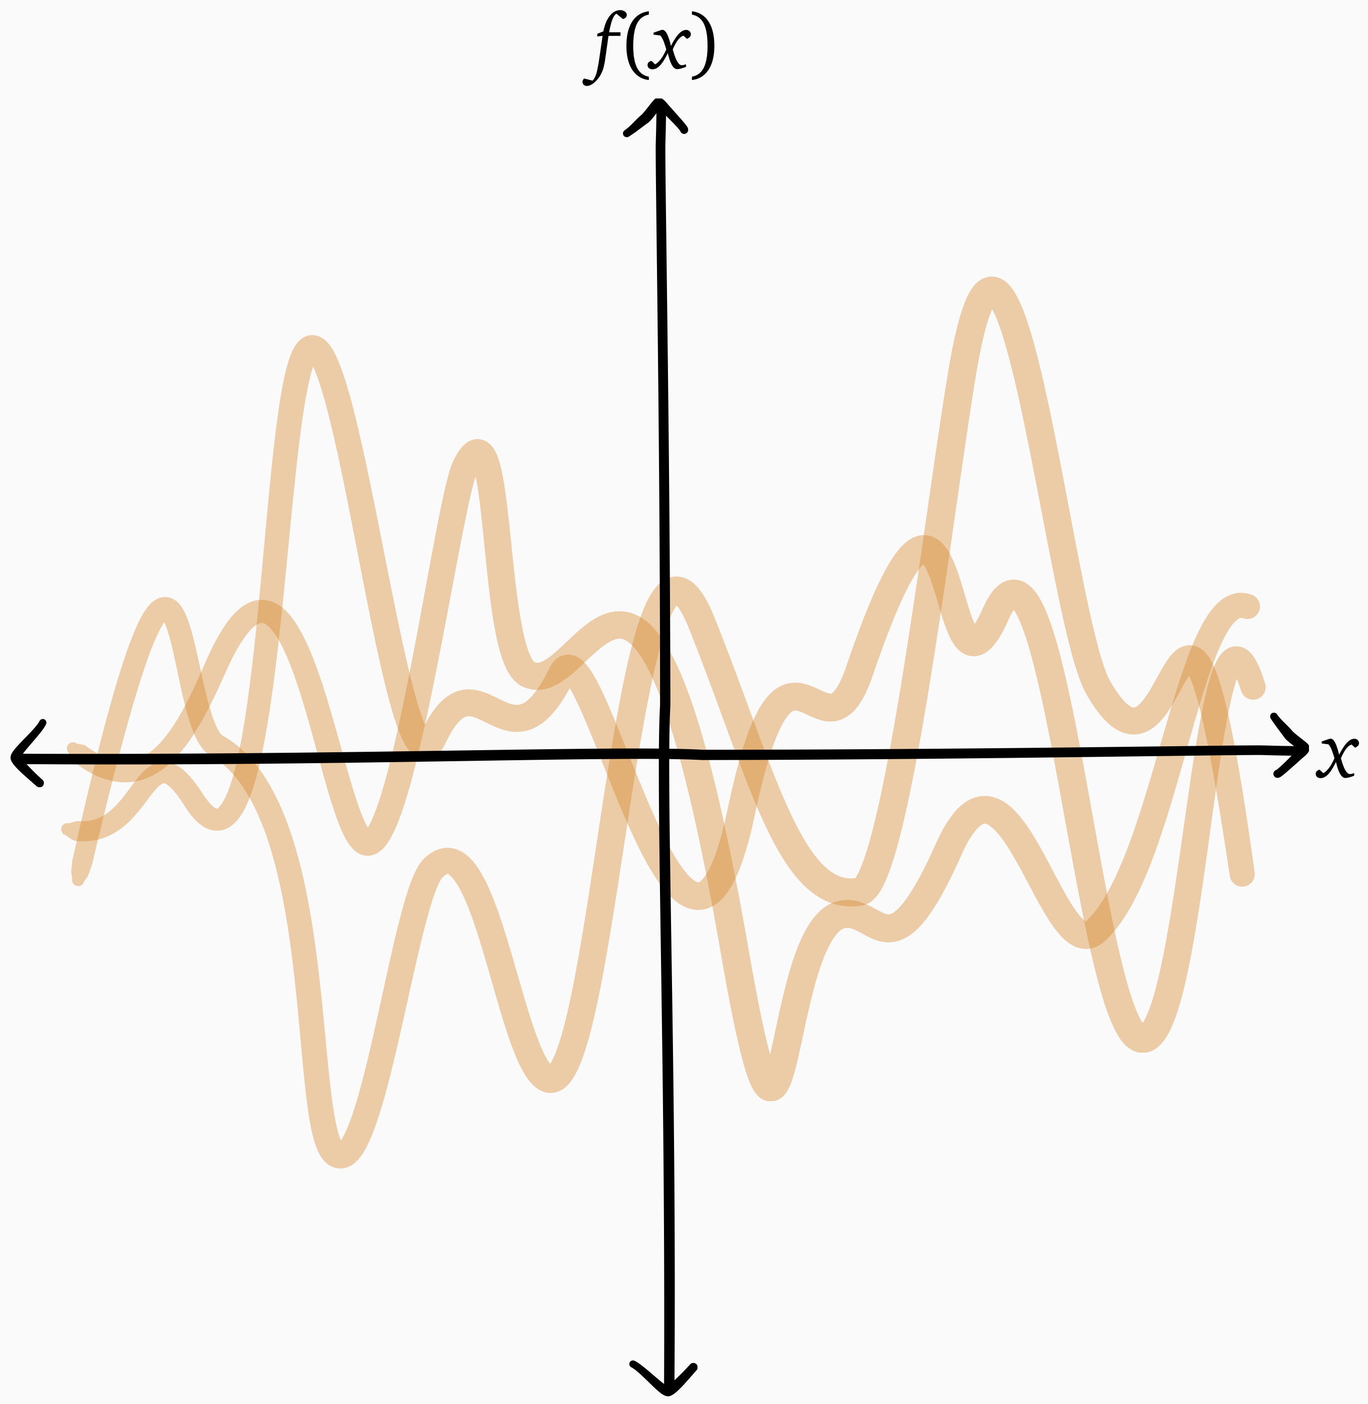

Expected improvement of Y over r:

$$\mathrm{EI}(Y, r) = \mathbf{E}[(Y - r)^+]$$

Expected improvement for BayesOpt

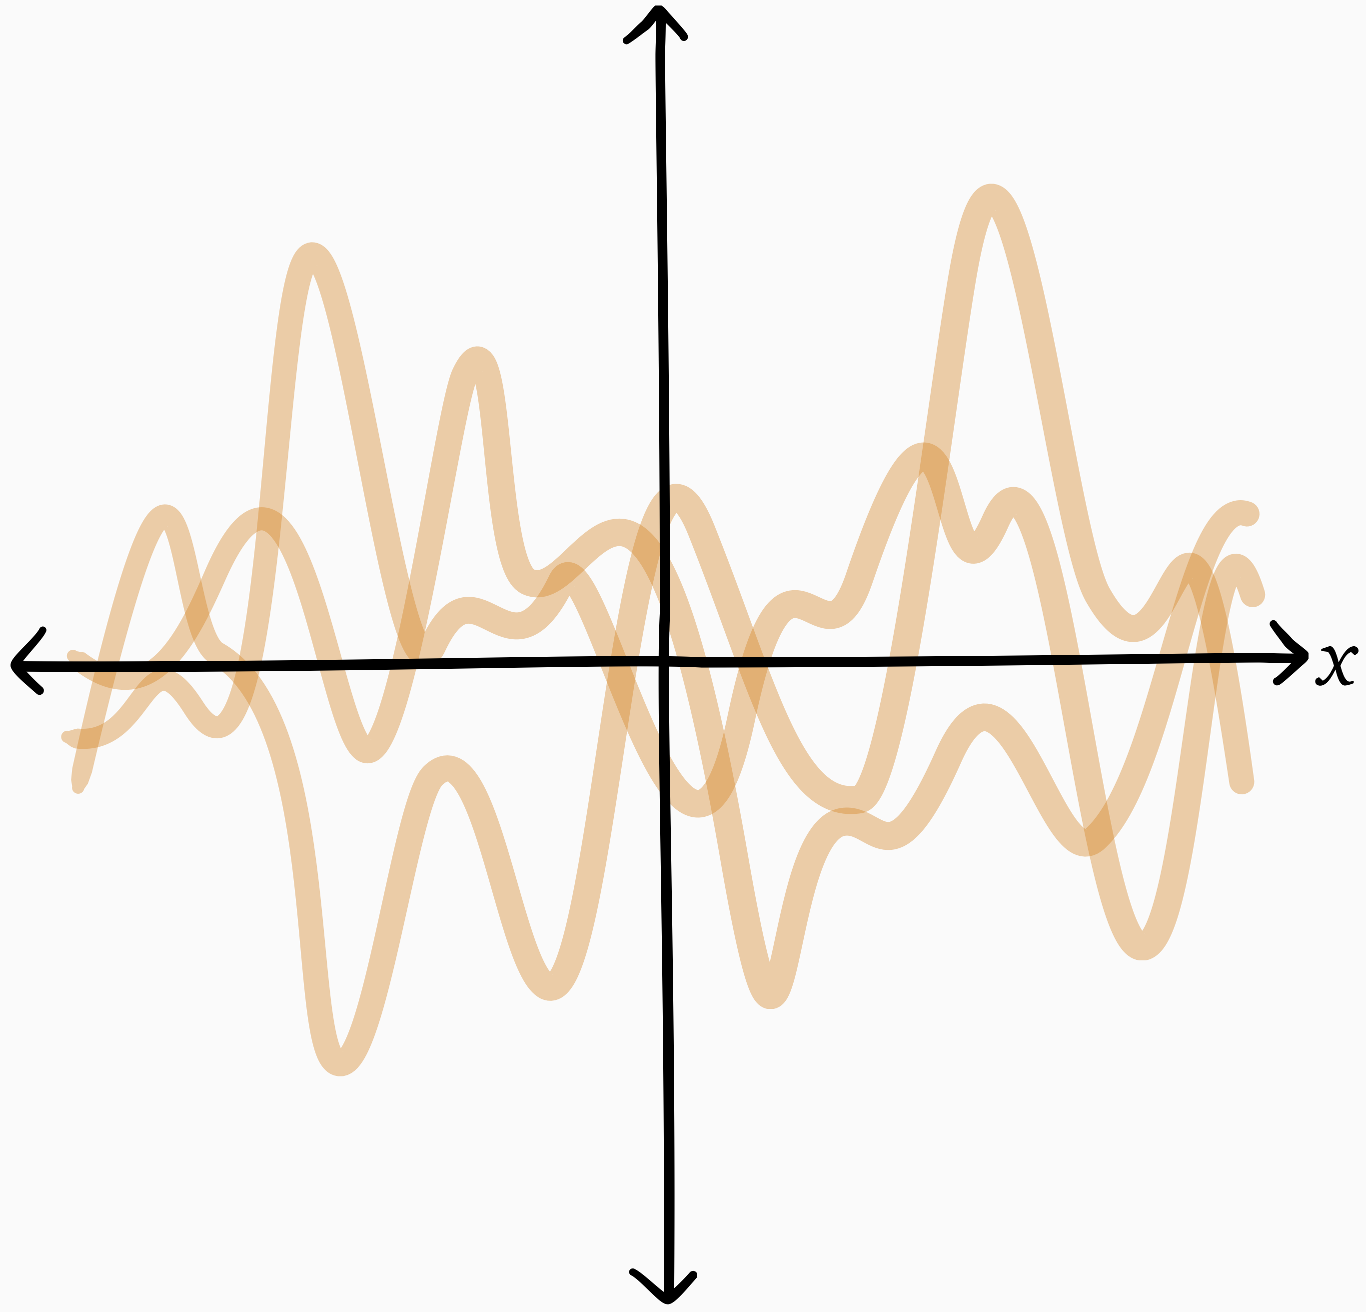

Expected improvement of Y over r:

$$\mathrm{EI}(Y, r) = \mathbf{E}[(Y - r)^+]$$

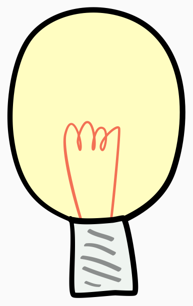

Maximize EI with r = best value so far

$$x_t = \arg\max_x \mathrm{EI}\Big(f(x),\ \max_{1 \le u < t} y_u\Big)$$

(with $f \sim$ [GP | data])

Expected improvement for BayesOpt

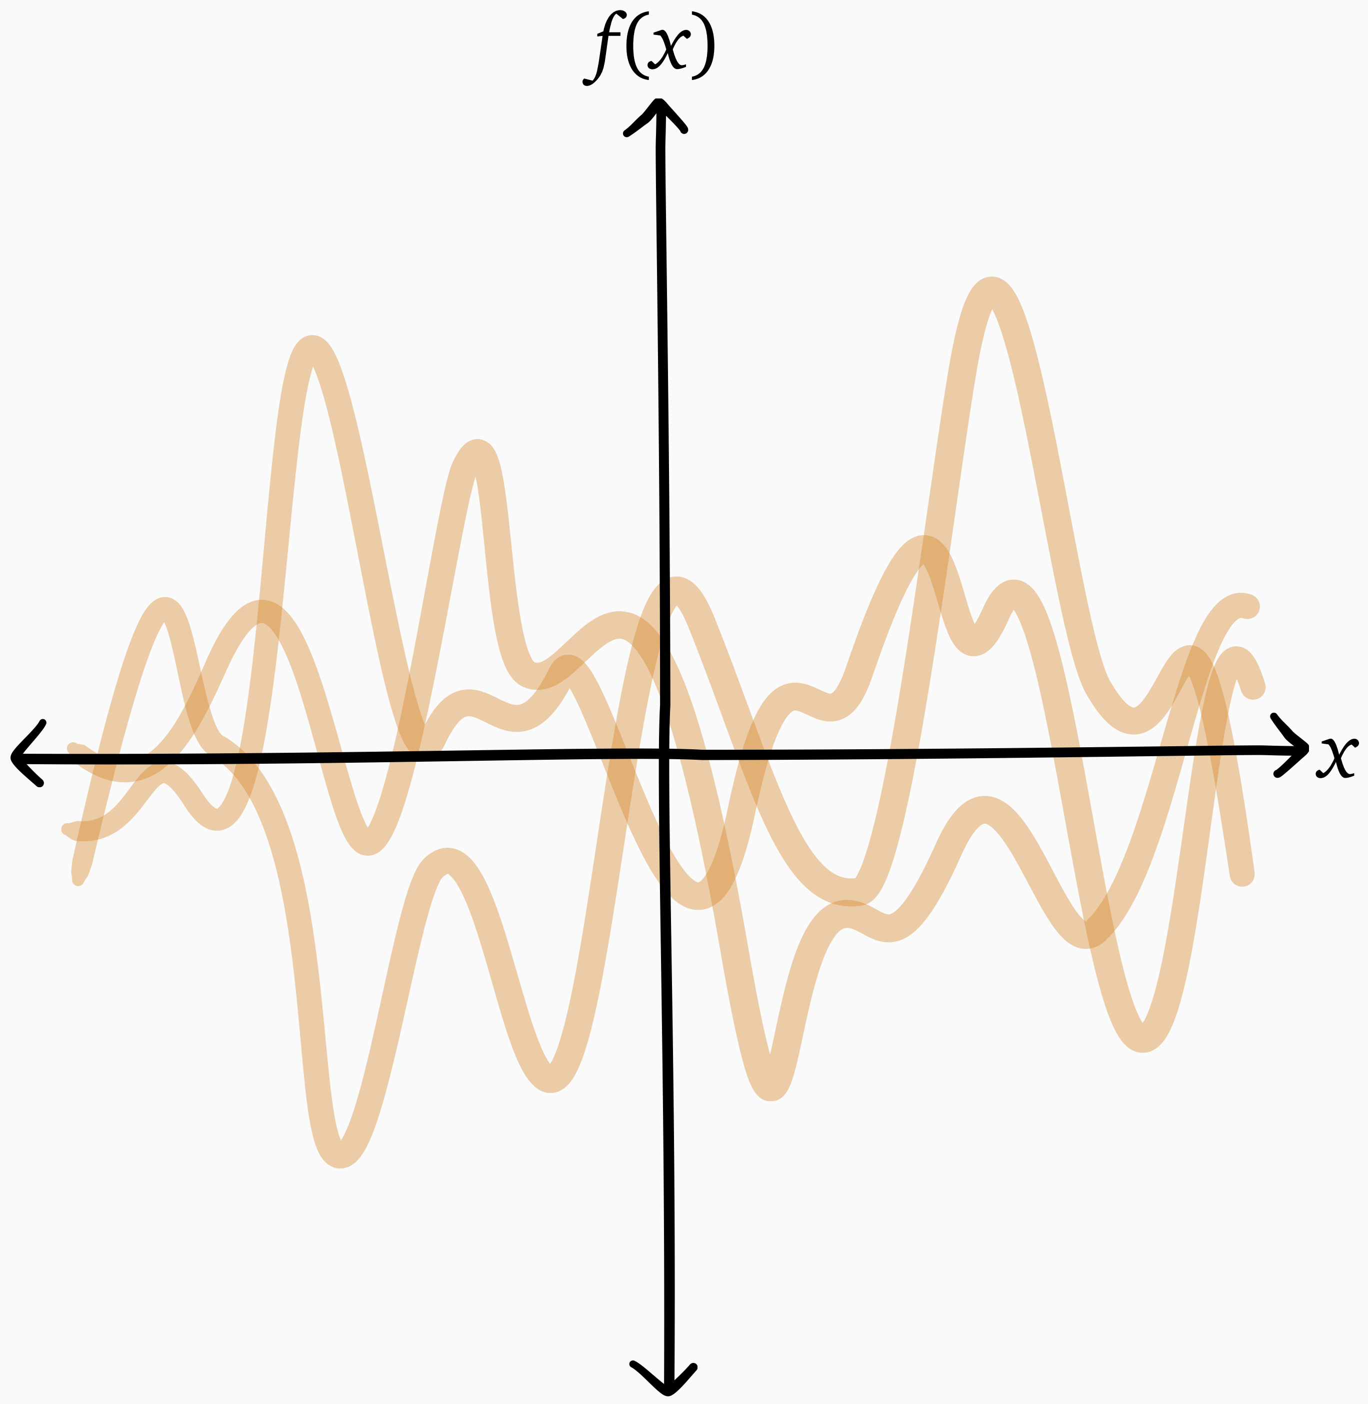

Expected improvement of Y over r:

$$\mathrm{EI}(Y, r) = \mathbf{E}[(Y - r)^+]$$

EI *function*

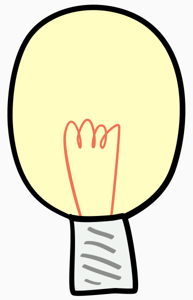

Maximize **EI** with r = best value so far

$$x_t = \arg\max_x \mathrm{EI}\big(f(x), \max_{1 \le u < t} y_u\big)$$

(with $f \sim$ [GP | data])

Expected improvement for BayesOpt

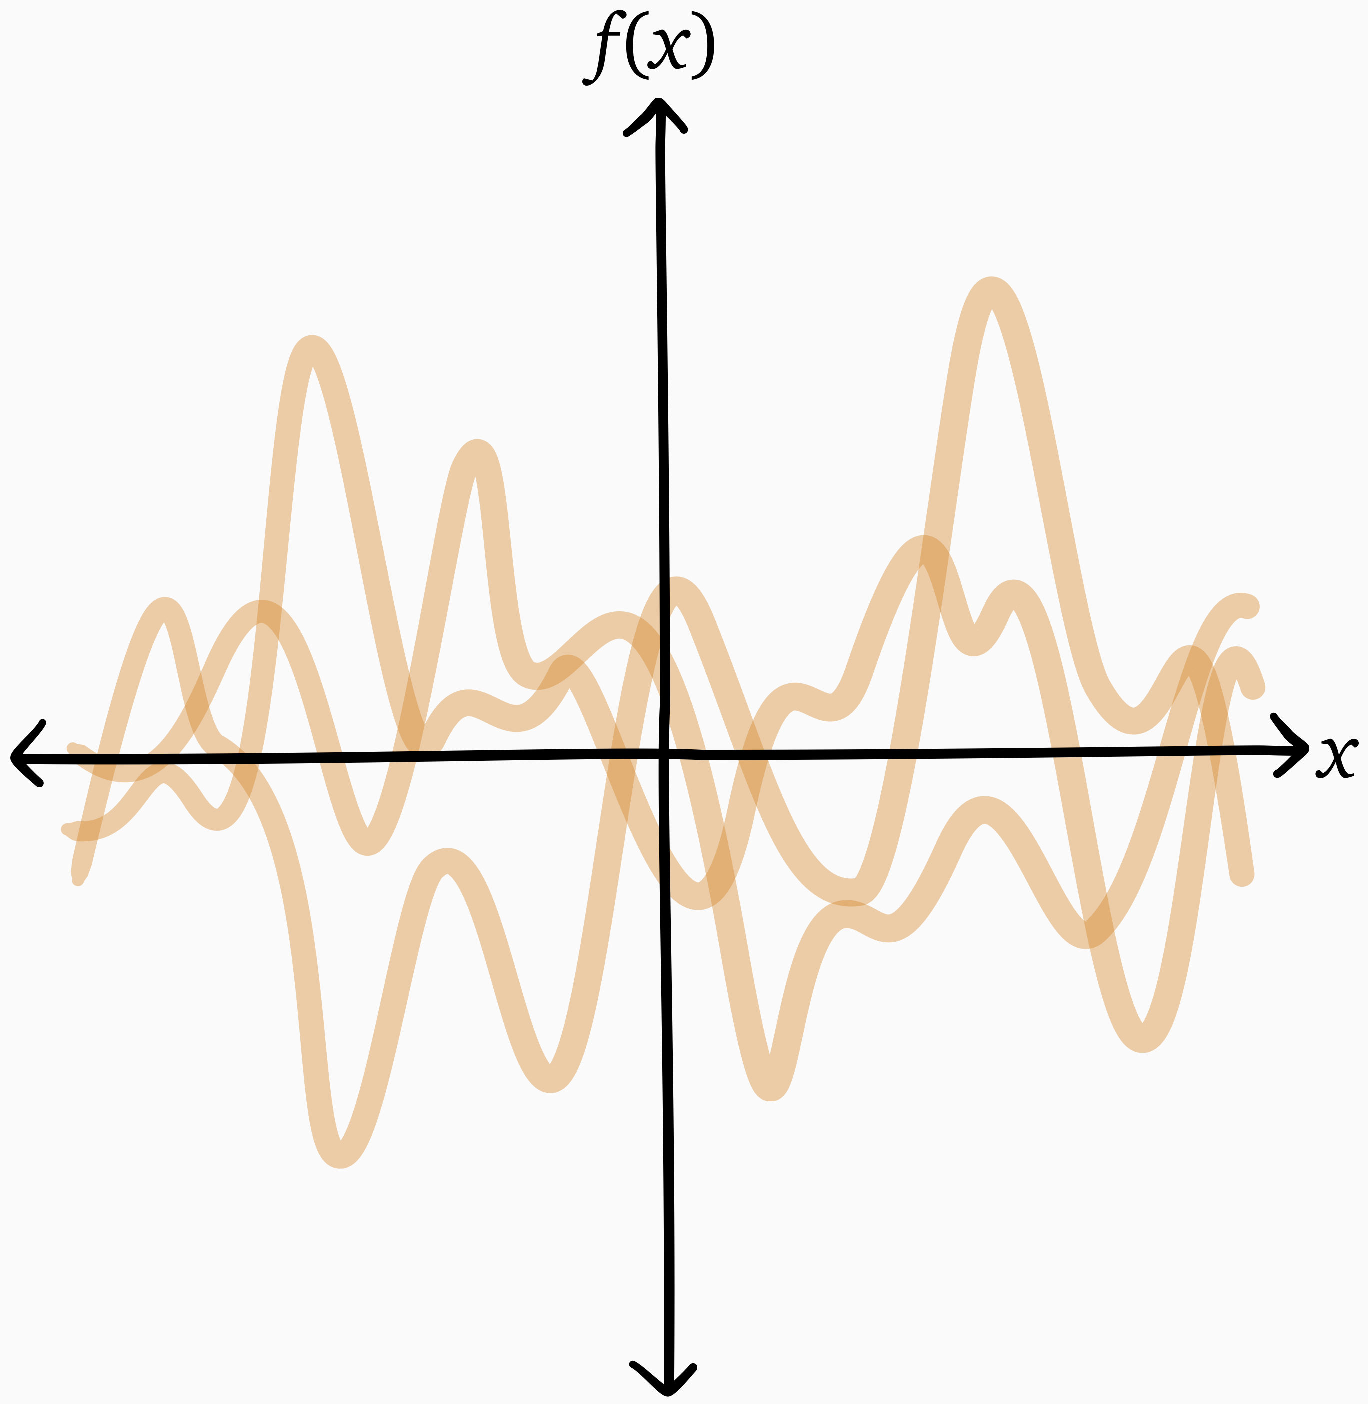

Expected improvement of Y over r:

$$\mathrm{EI}(Y, r) = \mathbf{E}[(Y - r)^+]$$

EI *function*

Maximize EI with r = best value so far

$$x_t = \arg\max_x \mathrm{EI}\Big(f(x),\ \max_{1 \le u < t} y_u\Big)$$

(with $f \sim$ [GP | data])

EI *policy*

Expected improvement for BayesOpt

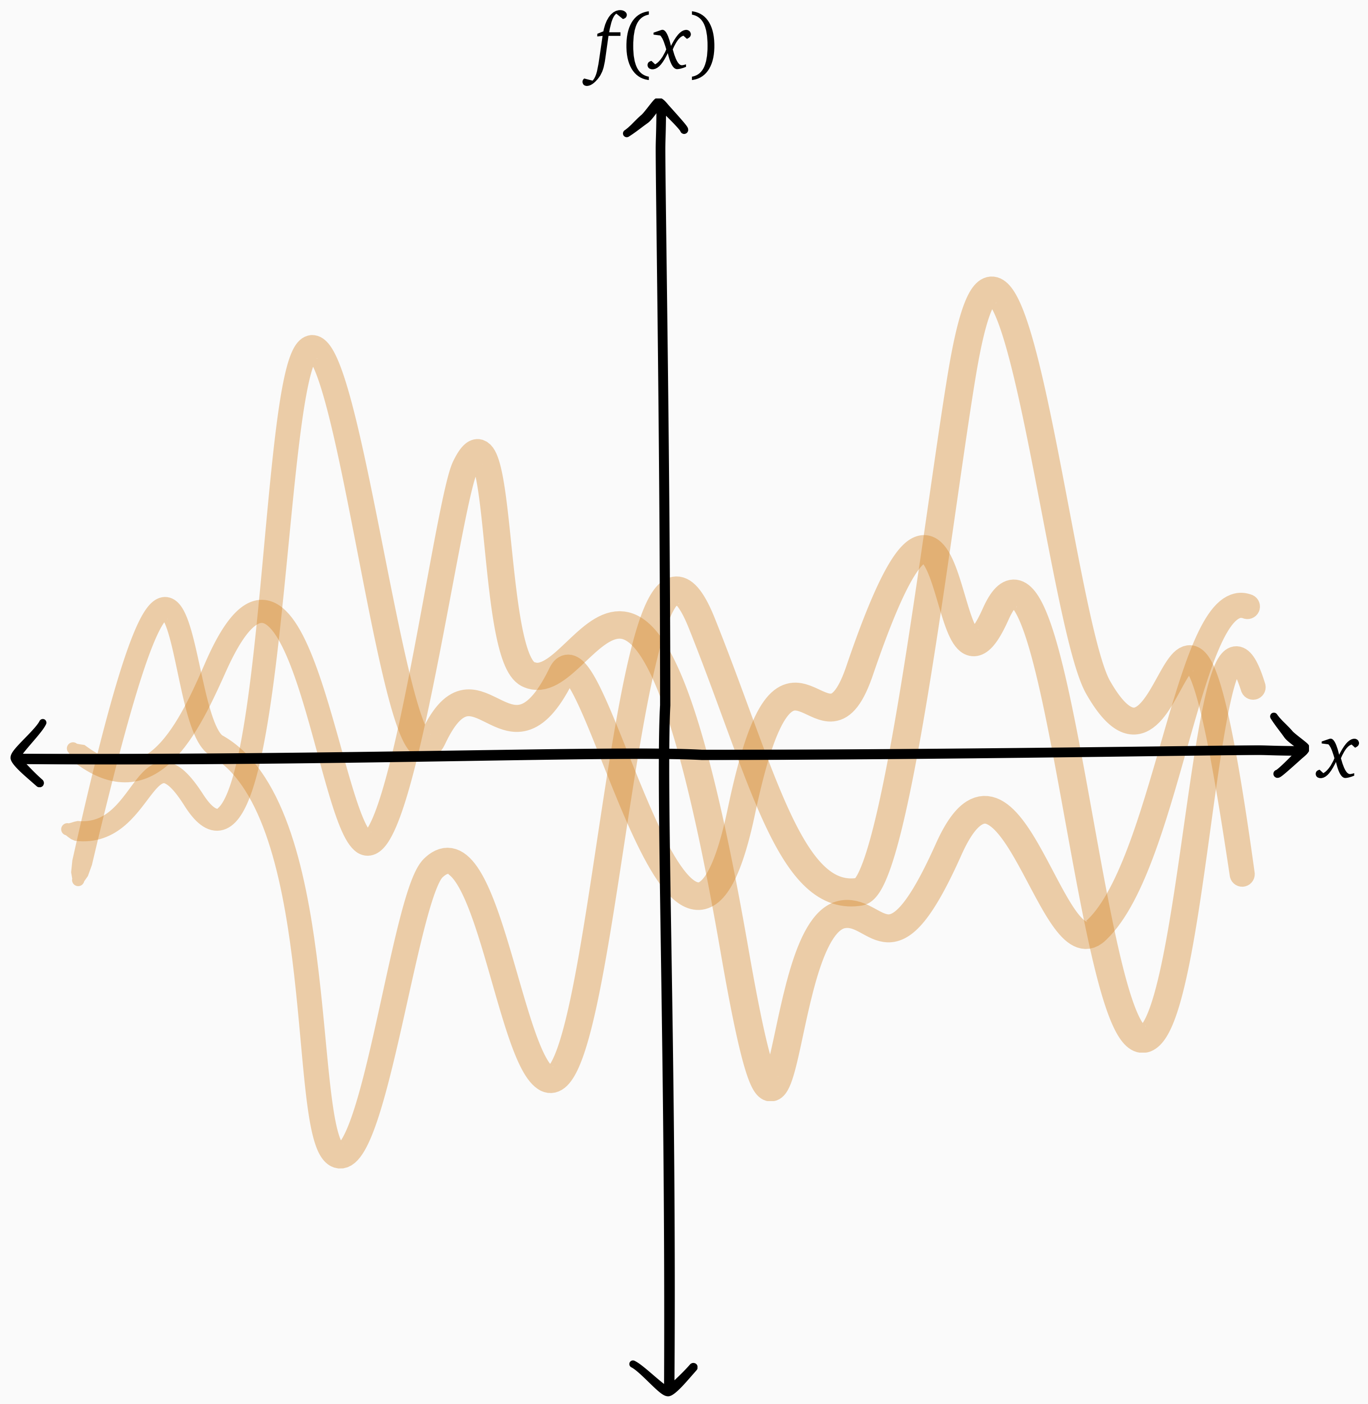

Expected improvement of Y over r:

$$\mathrm{EI}(Y, r) = \mathbf{E}[(Y - r)^{+}]$$

EI *function*

Maximize **EI** with r = best value so far

$$x_t = \arg\max_x \mathrm{EI}\big(f(x),\ \max_{1 \le u < t} y_u\big)$$

EI *policy*

(with $f \sim [\mathrm{GP} \mid \mathrm{data}]$)

BayesOpt

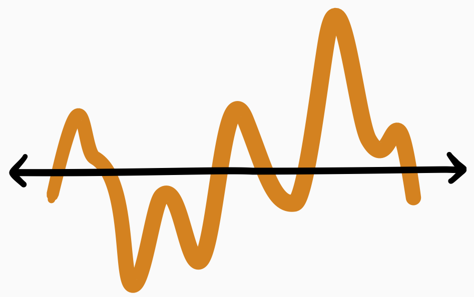 BayesOpt 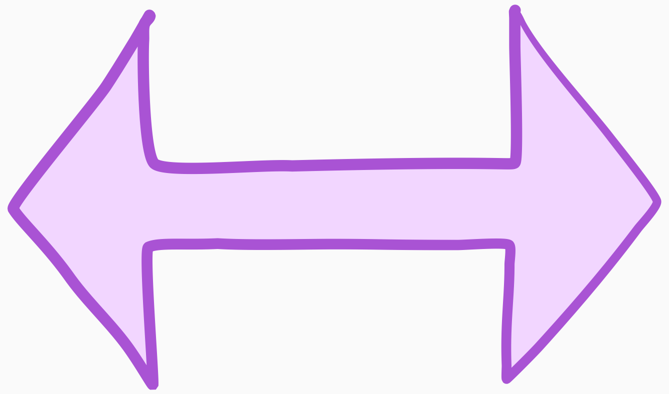 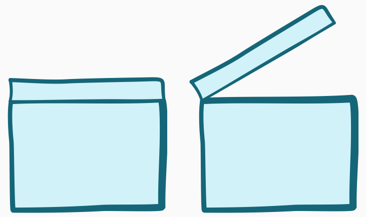 Pandora's box

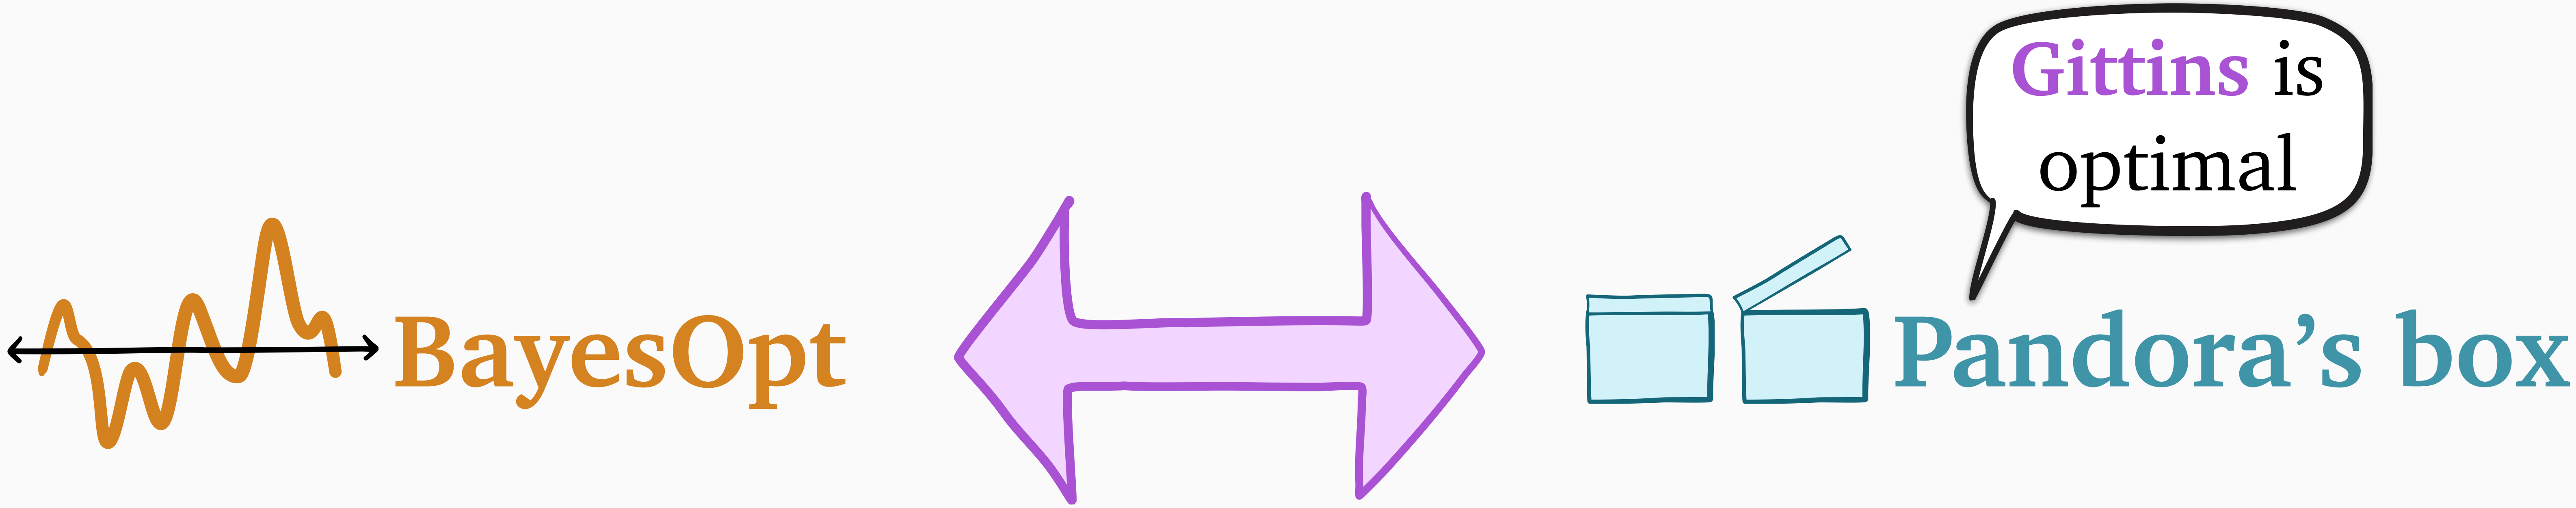
BayesOpt
Pandora's box
Gittins is optimal

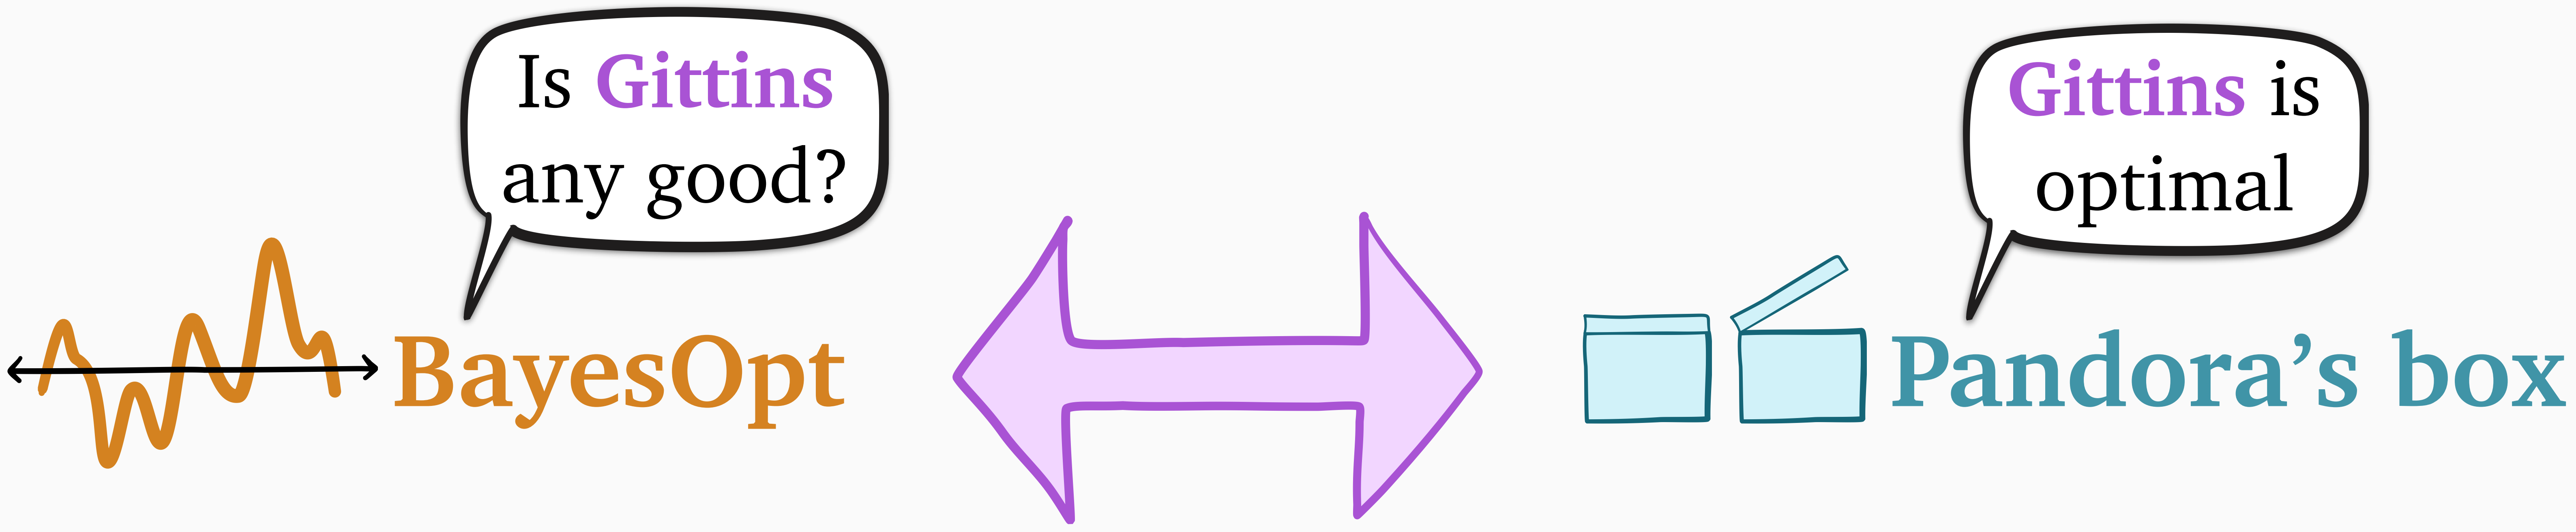
Is Gittins any good?
BayesOpt
Gittins is optimal
Pandora's box

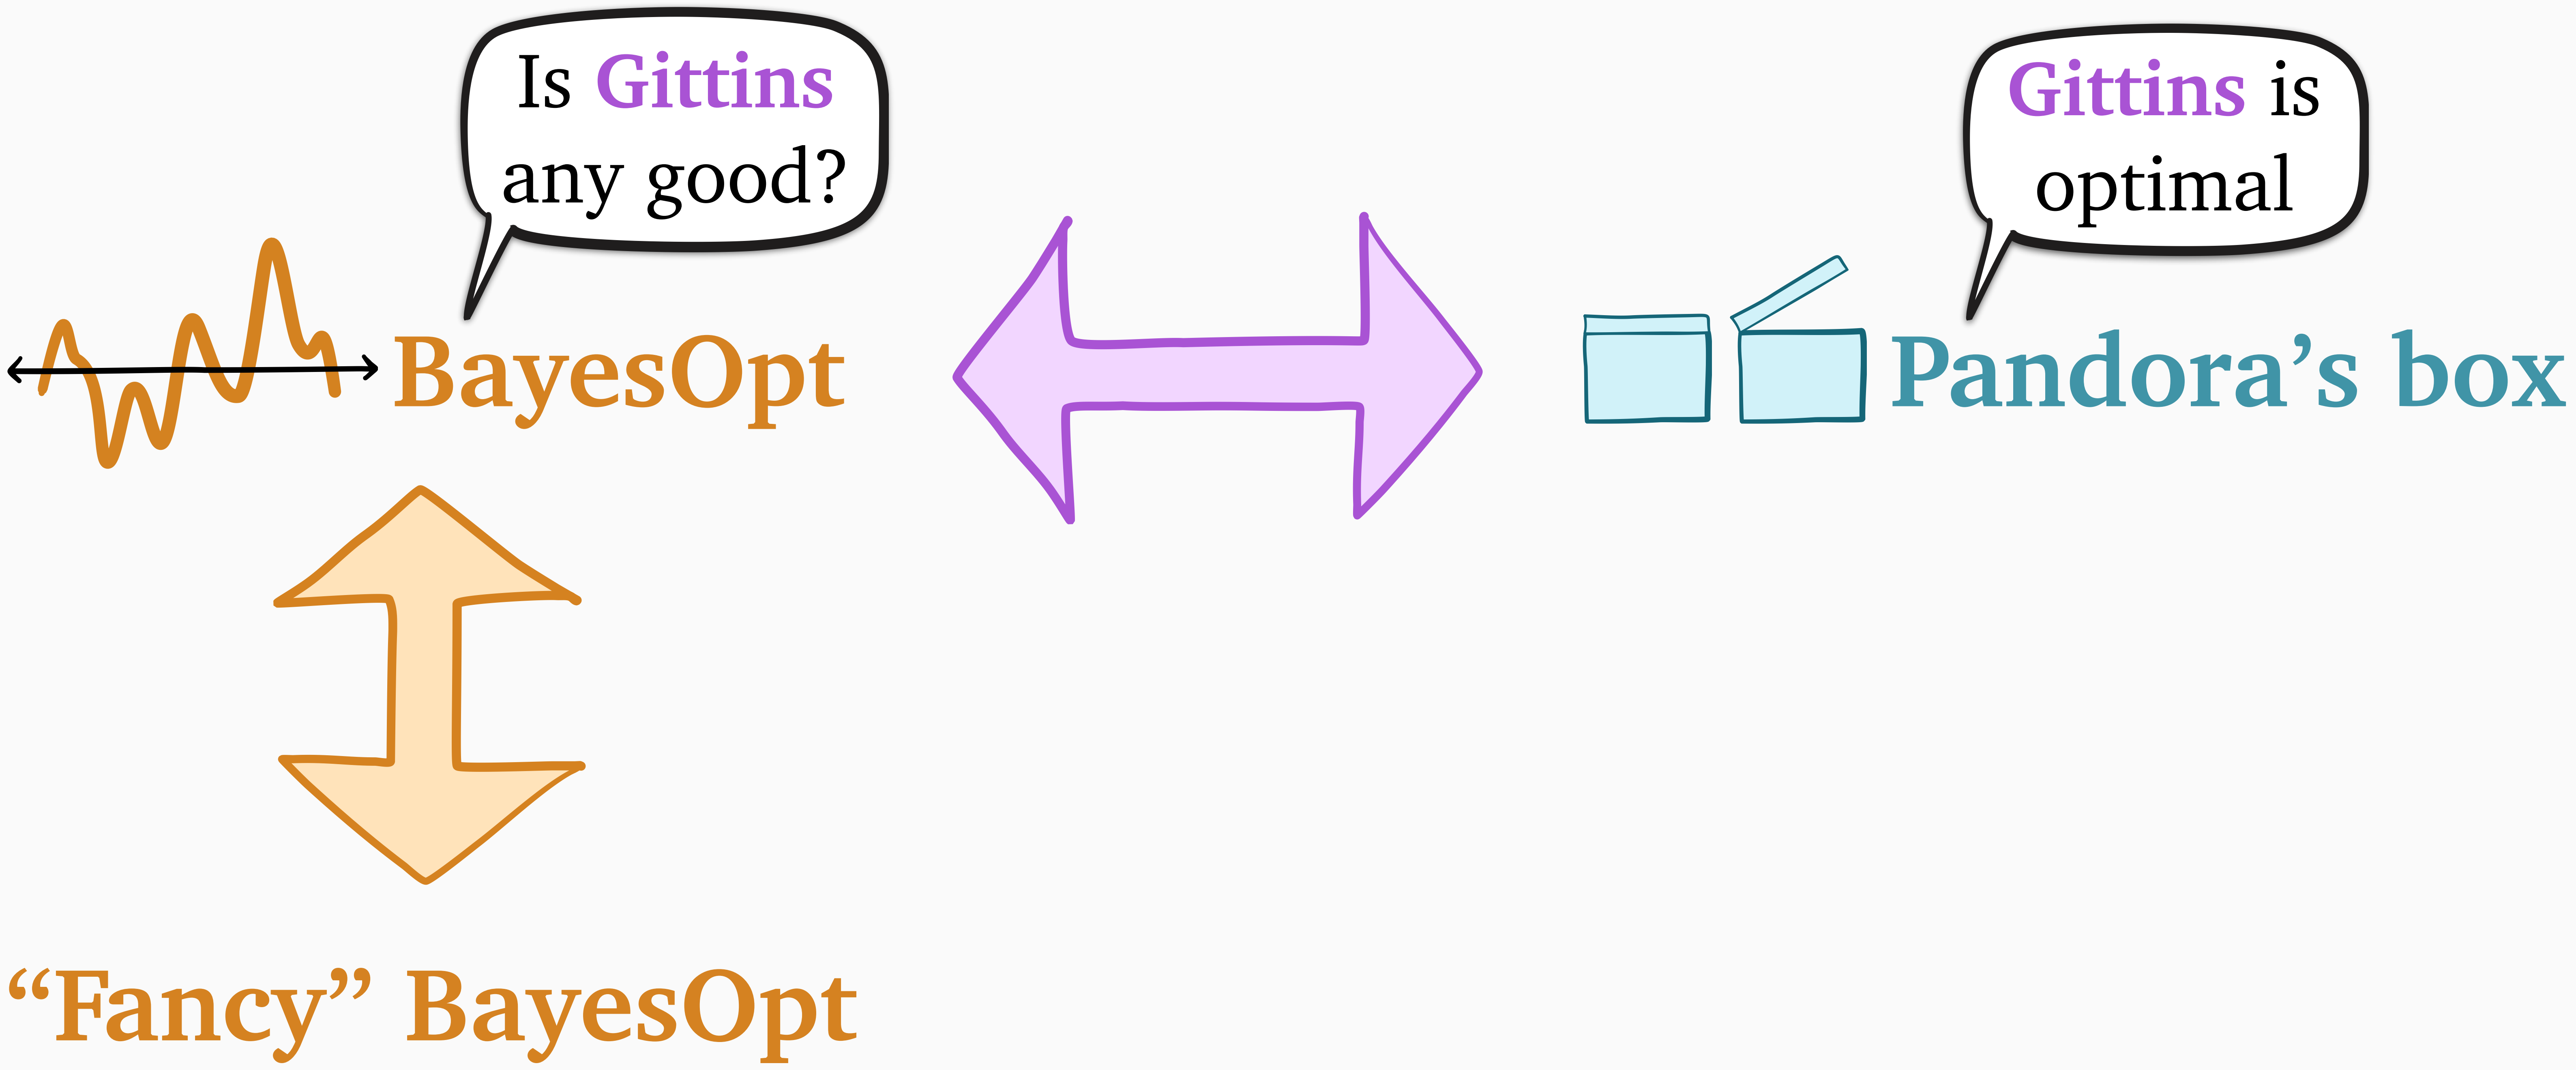
Is Gittins any good?
BayesOpt
Gittins is optimal
Pandora's box
"Fancy" BayesOpt

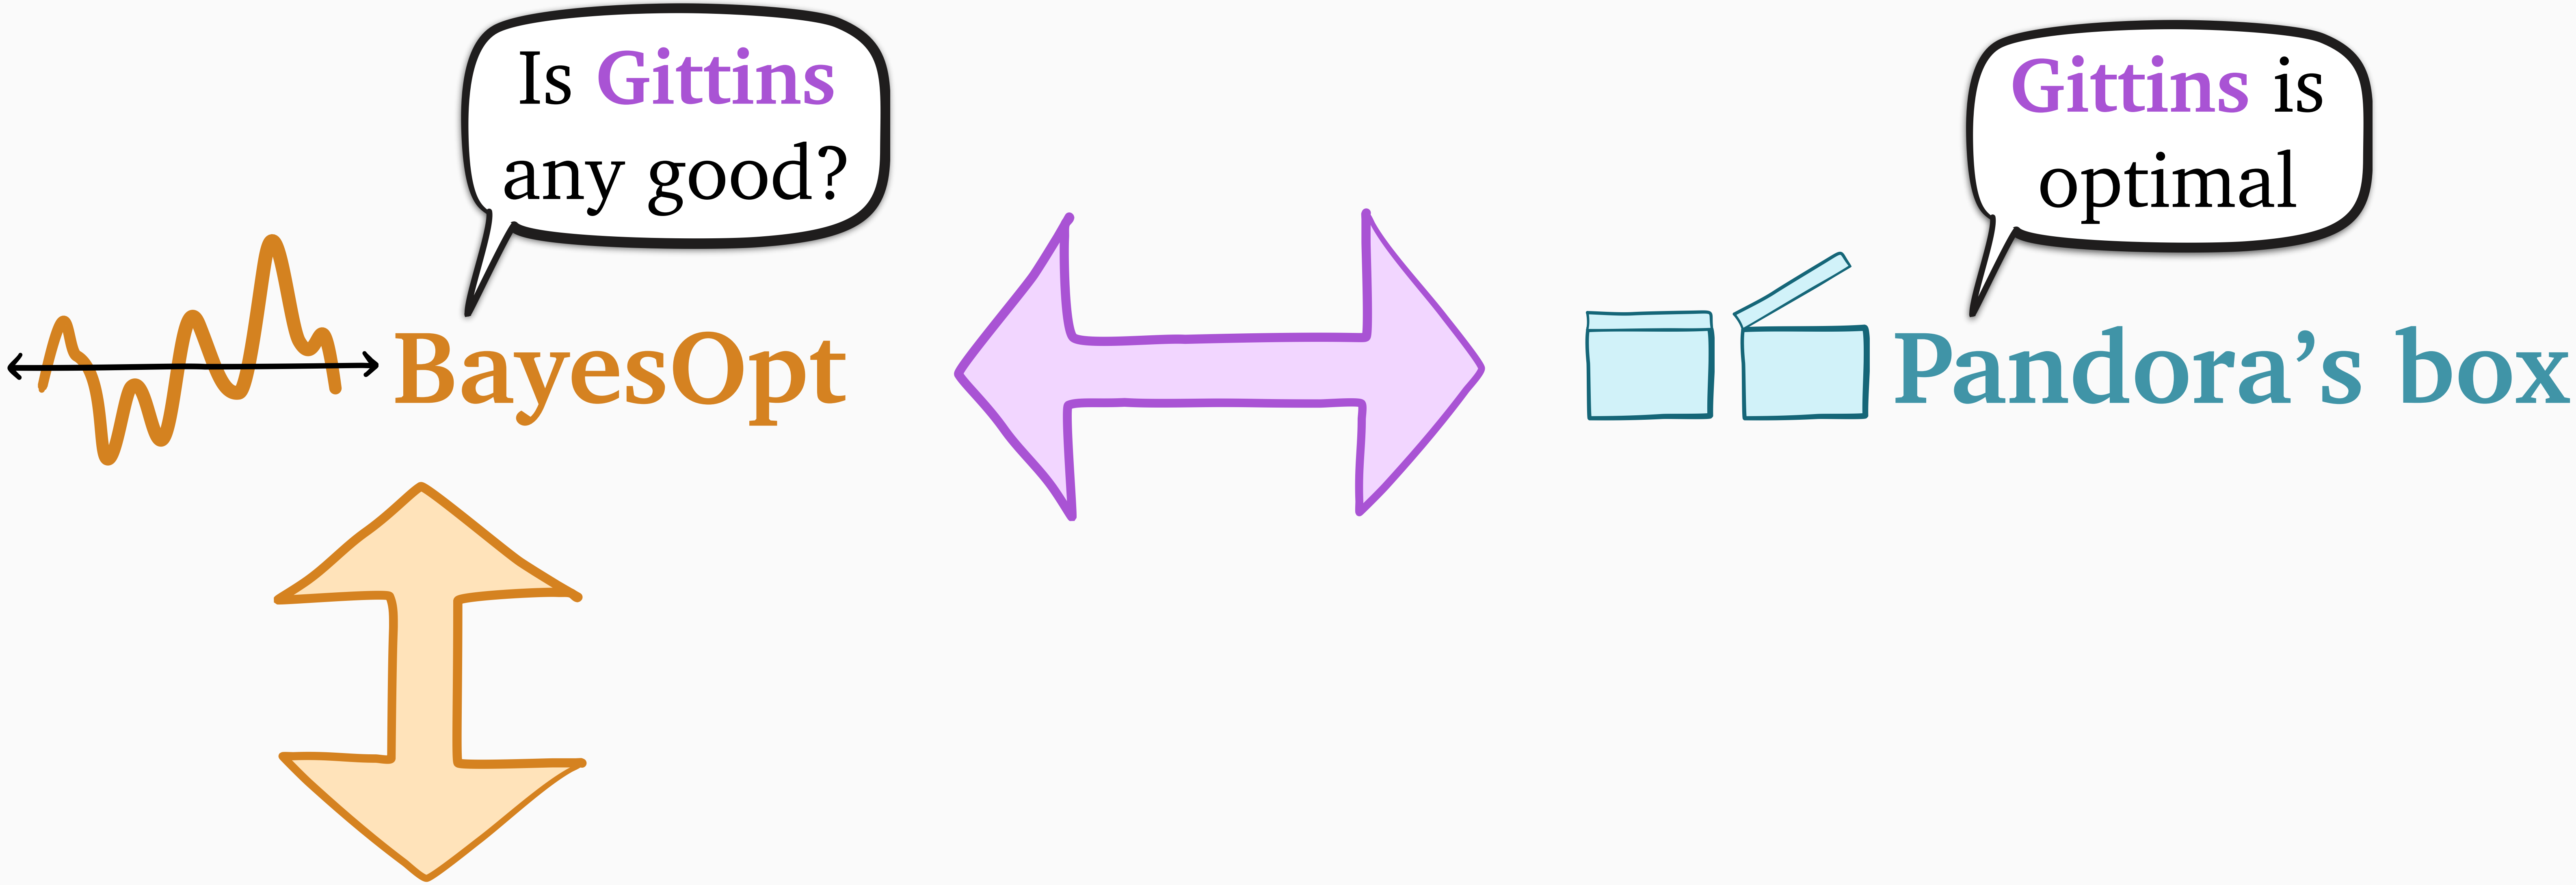

"Fancy" BayesOpt

- Multistage inspection
- Heterogeneous costs
- Partial feedback
- Automatic stopping

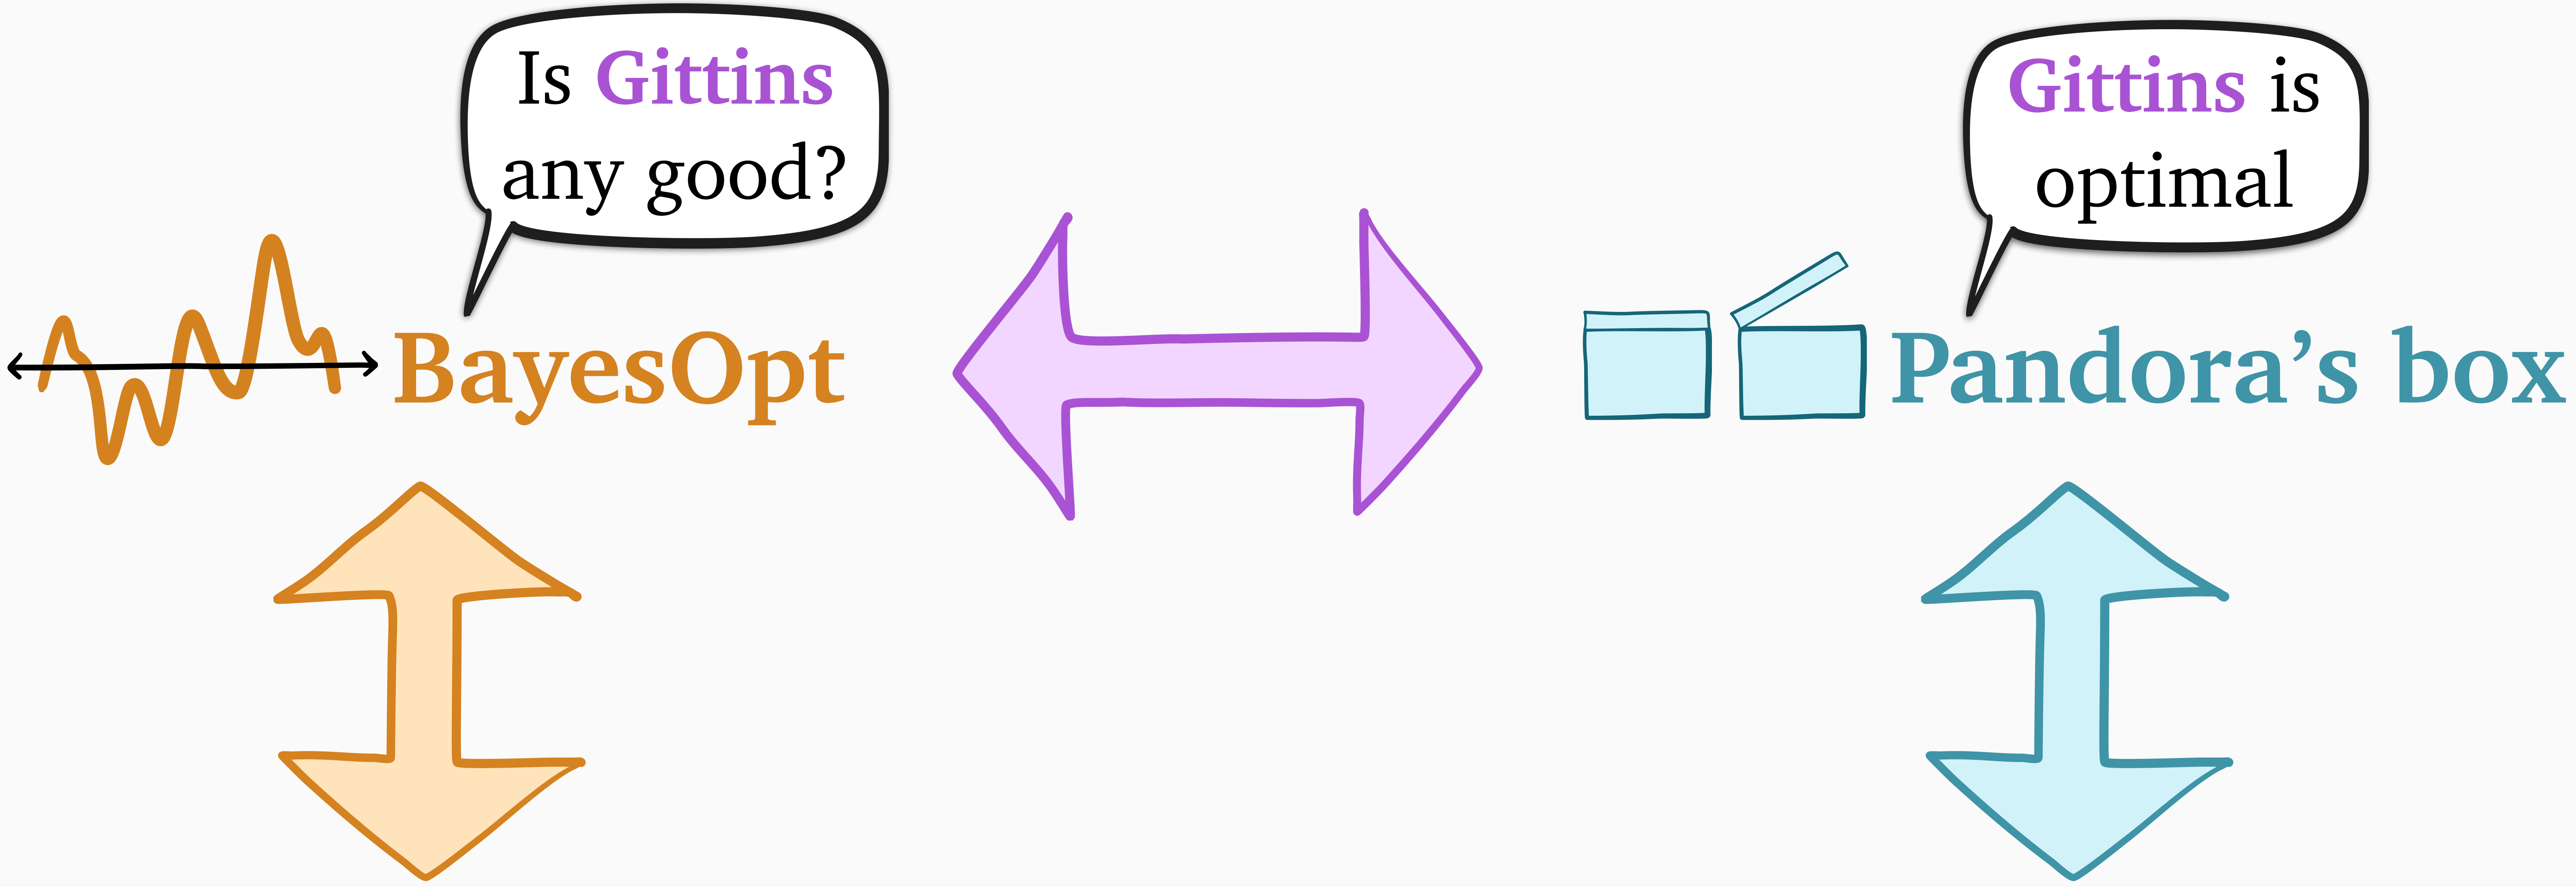

"Fancy" BayesOpt

- Multistage inspection
- Heterogeneous costs
- Partial feedback
- Automatic stopping

Markovian MAB

(multi-armed bandit)

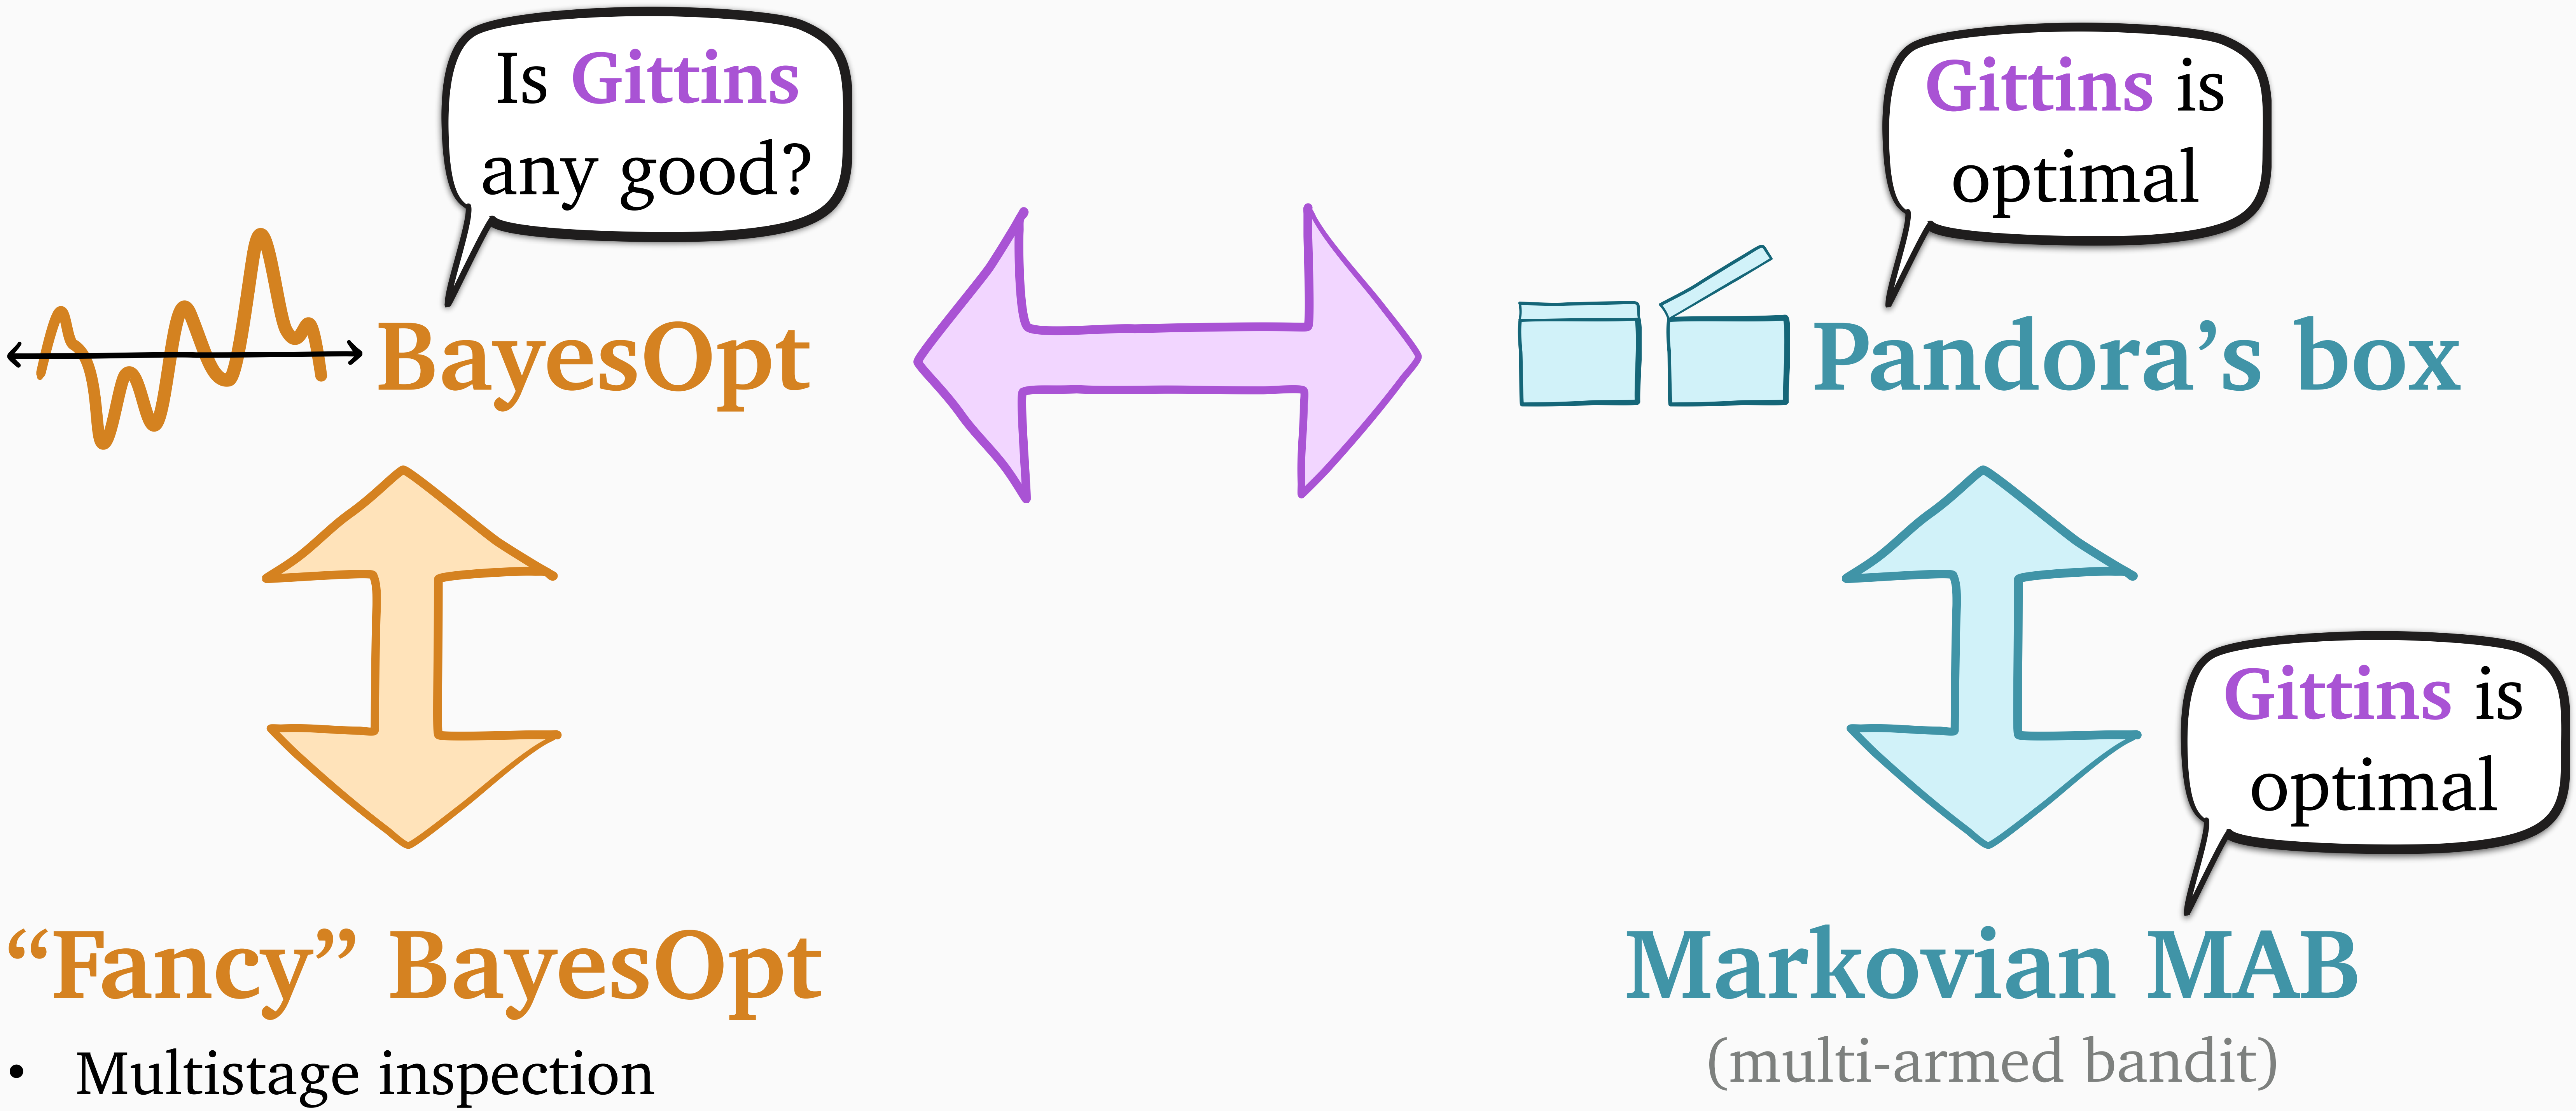

- Multistage inspection
- Heterogeneous costs
- Partial feedback
- Automatic stopping

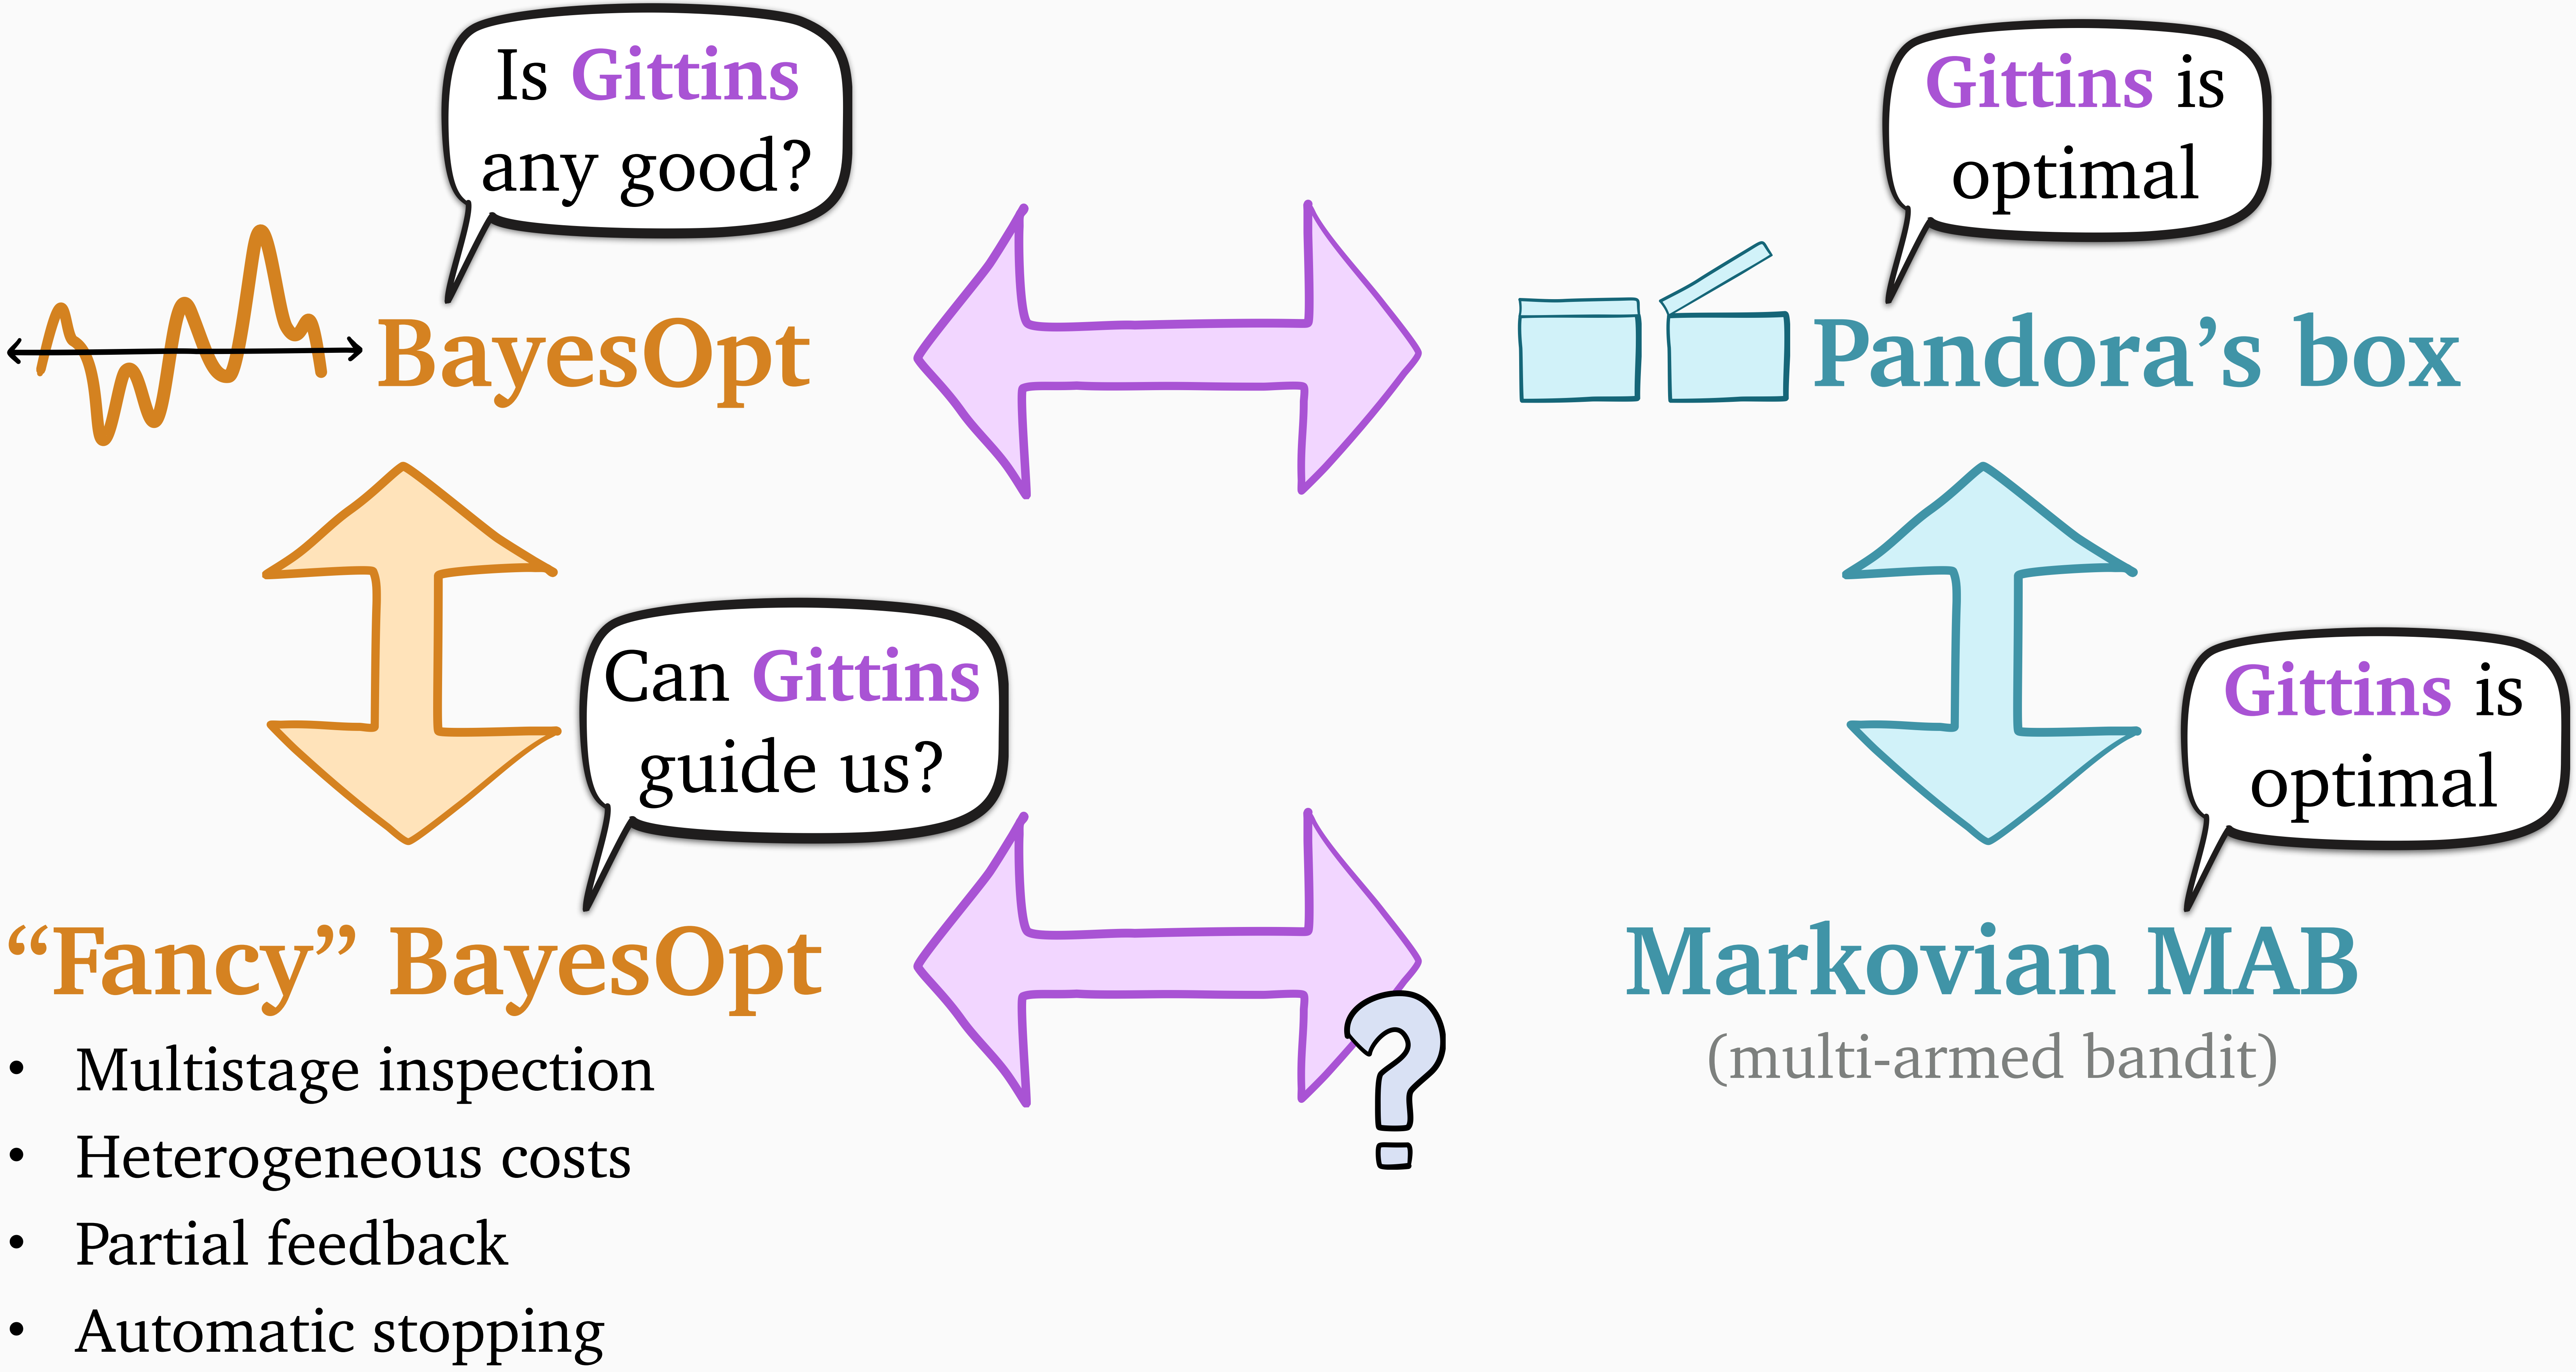
Is Gittins any good?
BayesOpt
Gittins is optimal
Pandora's box
Can Gittins guide us?
"Fancy" BayesOpt
• Multistage inspection
• Heterogeneous costs
• Partial feedback
• Automatic stopping
Gittins is optimal
Markovian MAB
(multi-armed bandit)

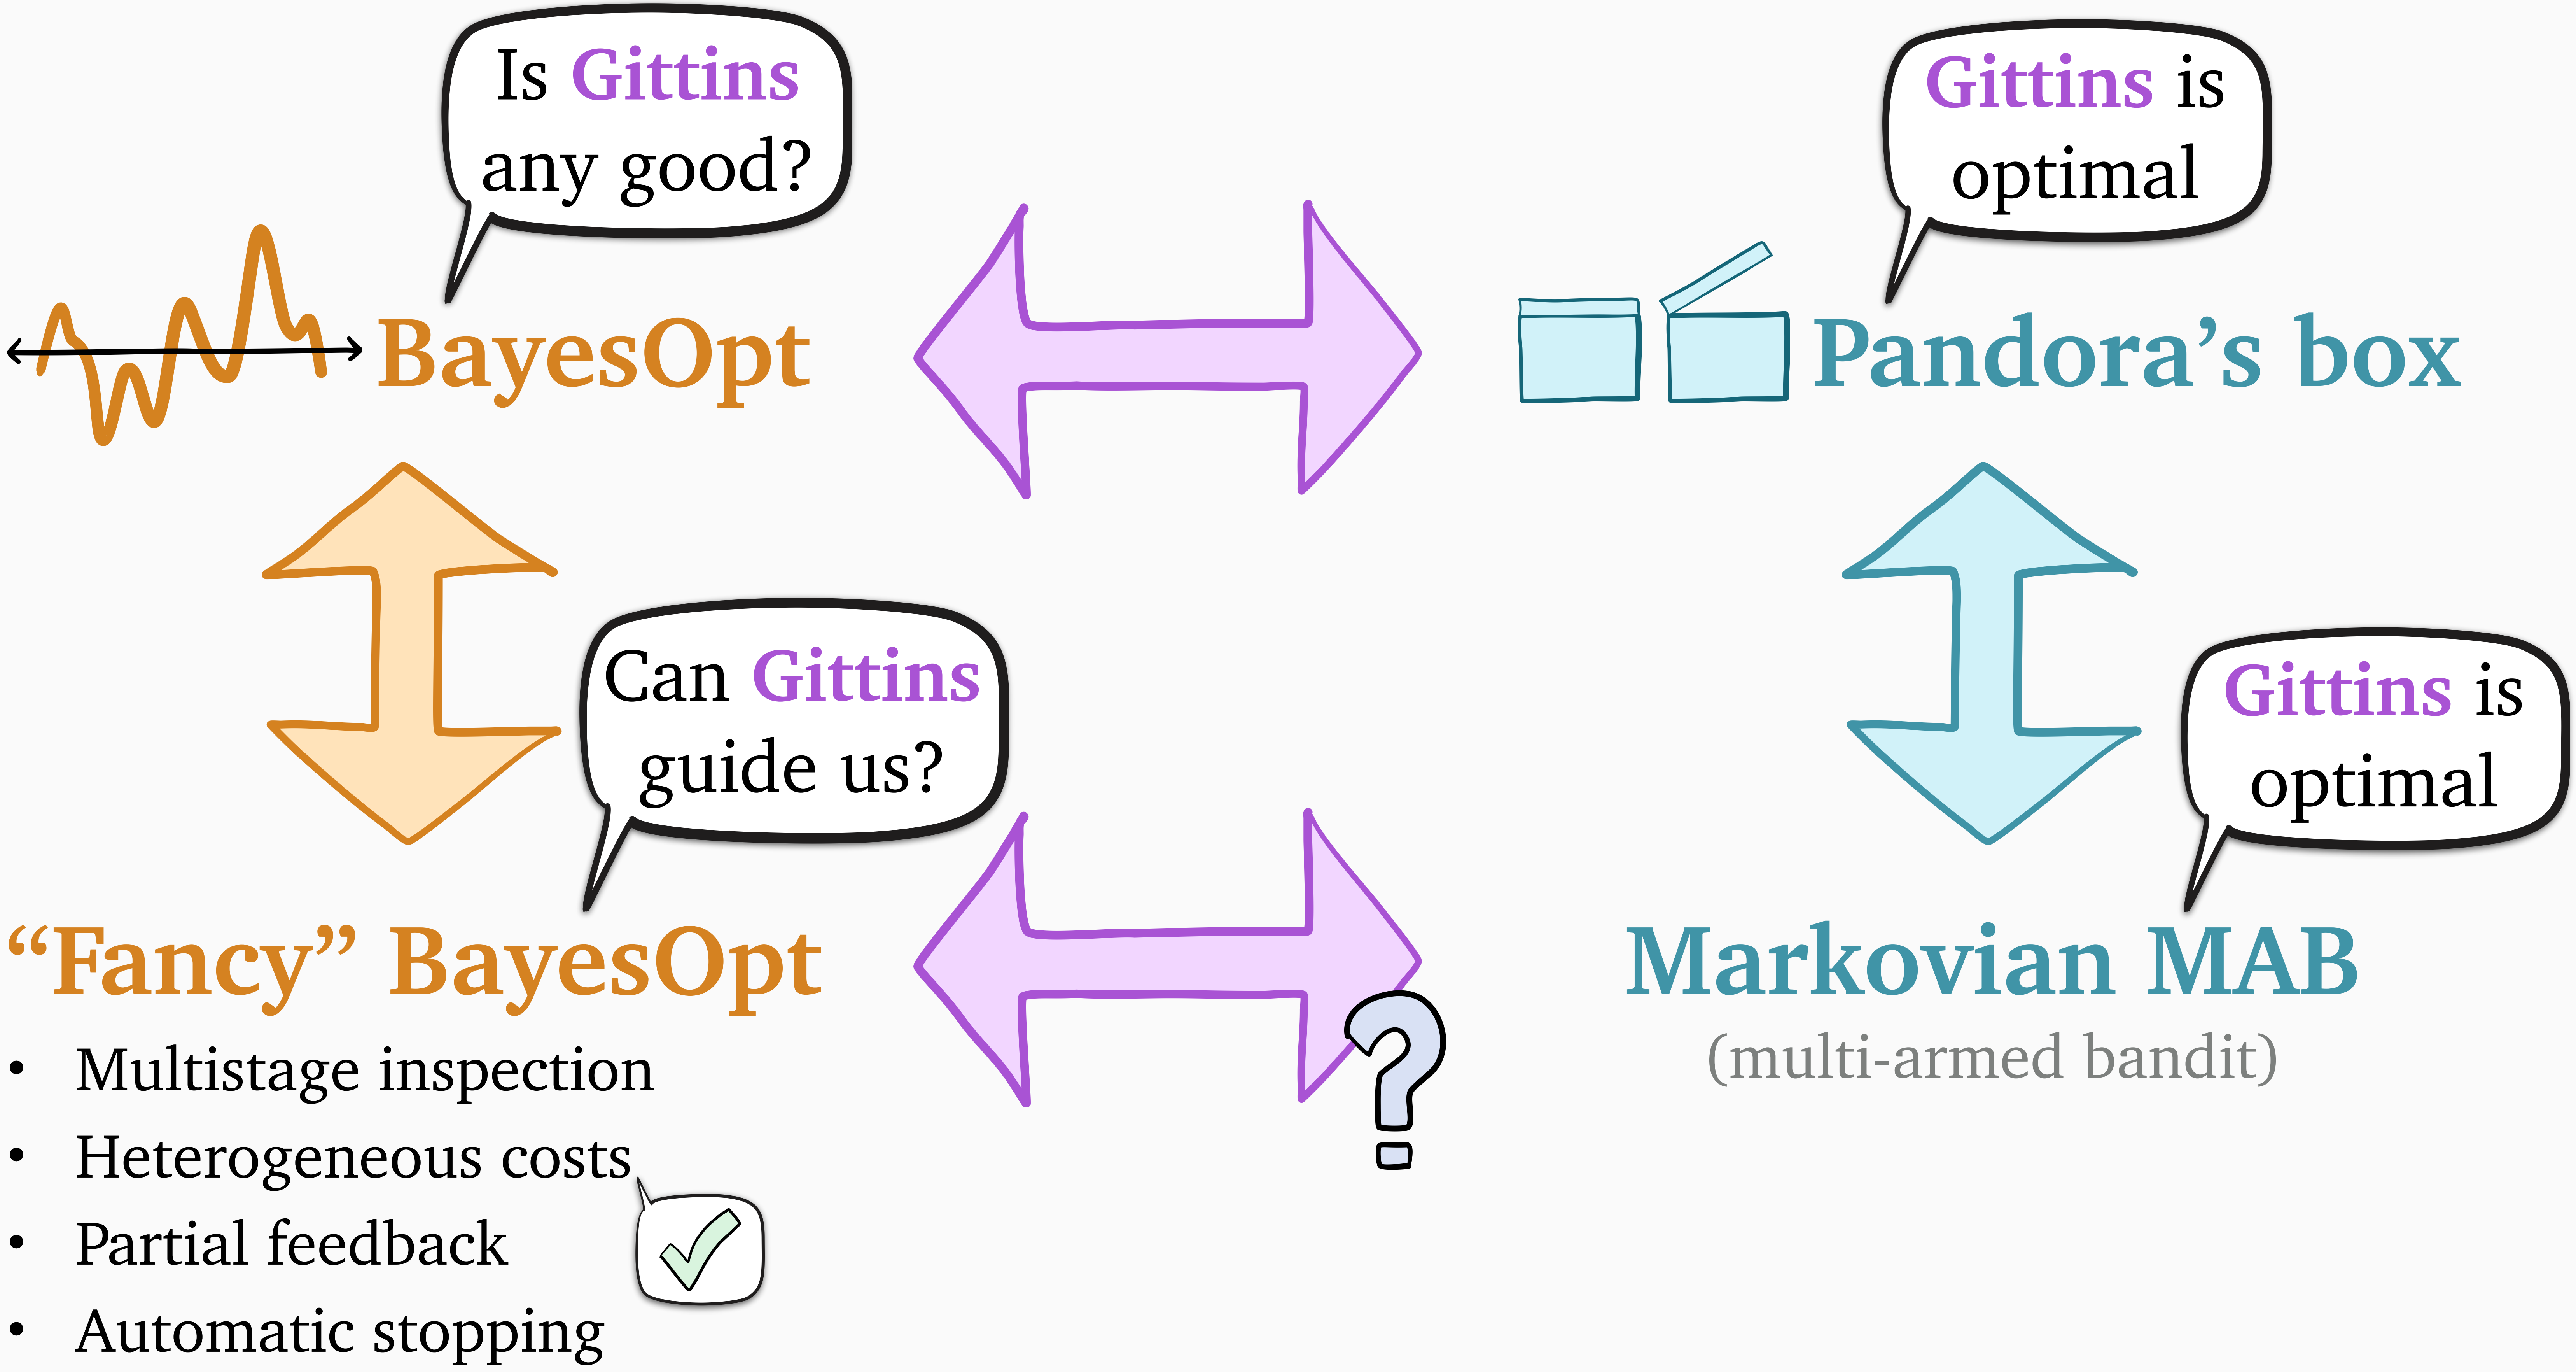
Is Gittins any good?
BayesOpt
Gittins is optimal
Pandora's box
Can Gittins guide us?
"Fancy" BayesOpt
• Multistage inspection
• Heterogeneous costs
• Partial feedback
• Automatic stopping
Gittins is optimal
Markovian MAB
(multi-armed bandit)

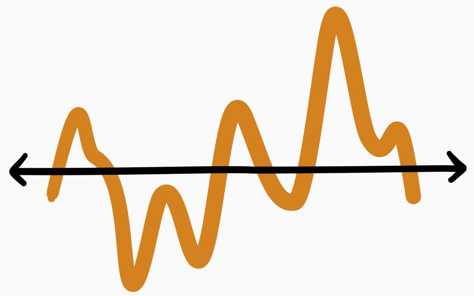

BayesOpt

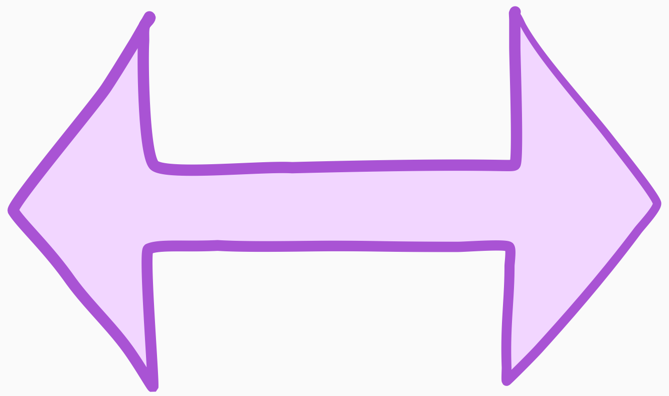

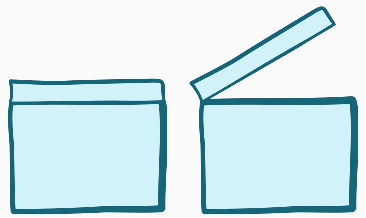

Pandora's box

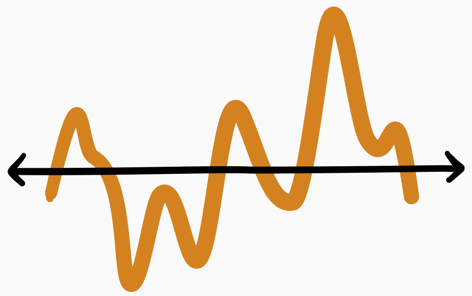
BayesOpt


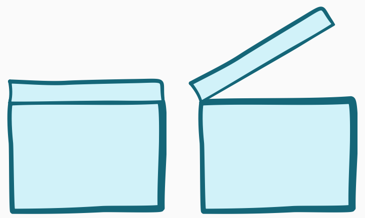
Pandora's box

Tutorial: What are Pandora's box and Gittins?

Tutorial: What are Pandora's box and Gittins?

Results: Does Gittins work for BayesOpt?

Tutorial: What are Pandora's box and Gittins?

Results: Does Gittins work for BayesOpt?

Hype: How could Gittins help practical BayesOpt?

Pandora's box problem

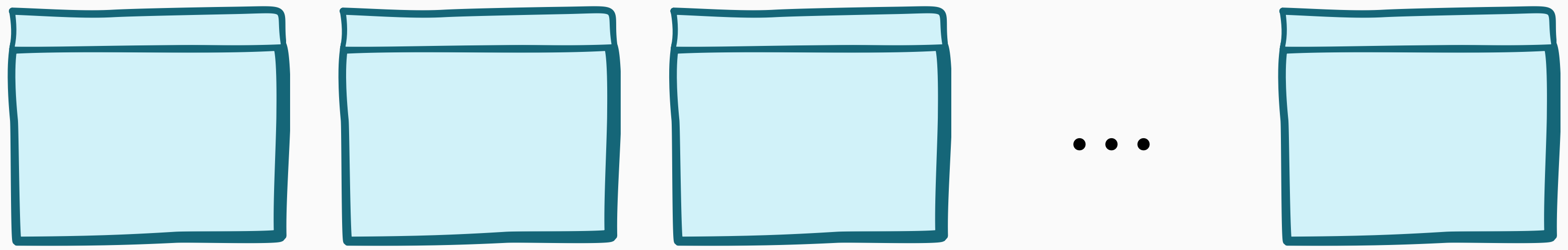

Pandora's box problem

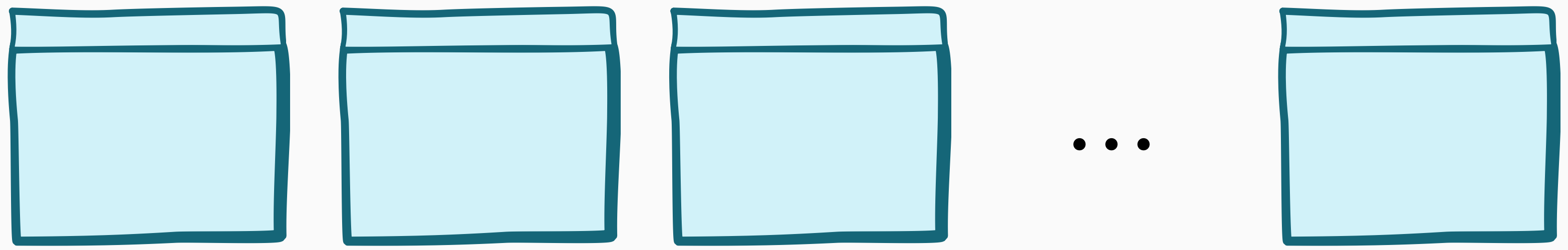

Each box:

- Opening cost c
- Hidden reward Y

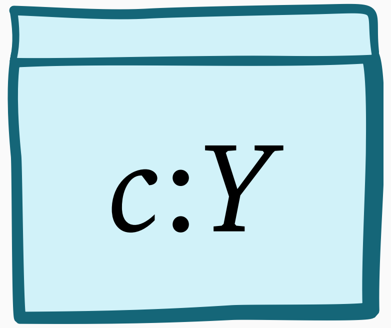

Pandora's box problem

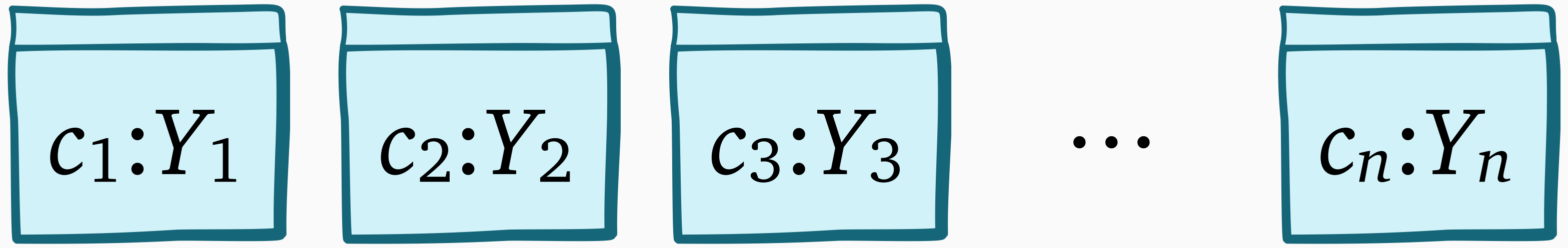

Each box:

- Opening cost c
- Hidden reward Y

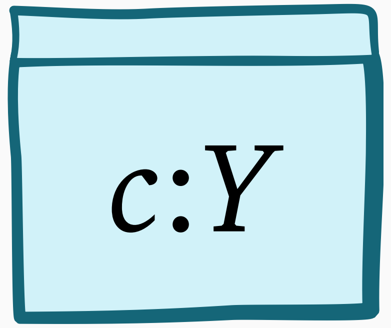

Pandora's box problem

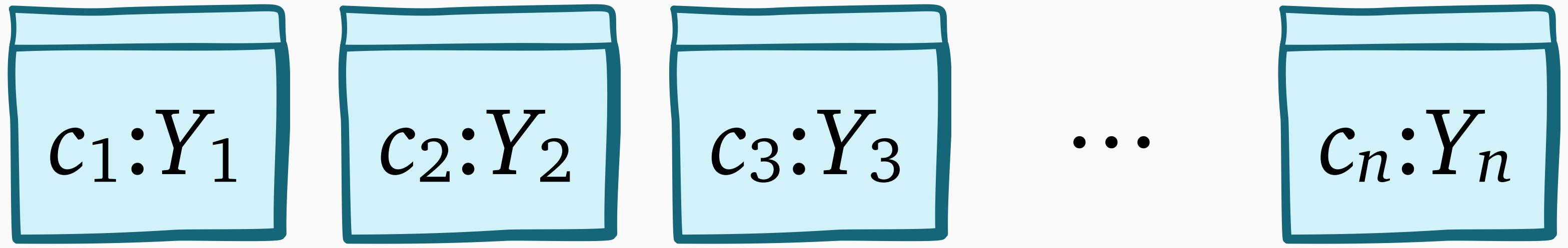

Each box:

- Opening cost c
- Hidden reward Y

independent

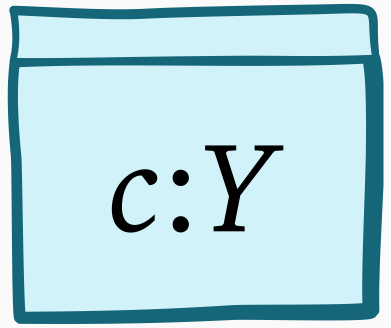

Pandora's box problem

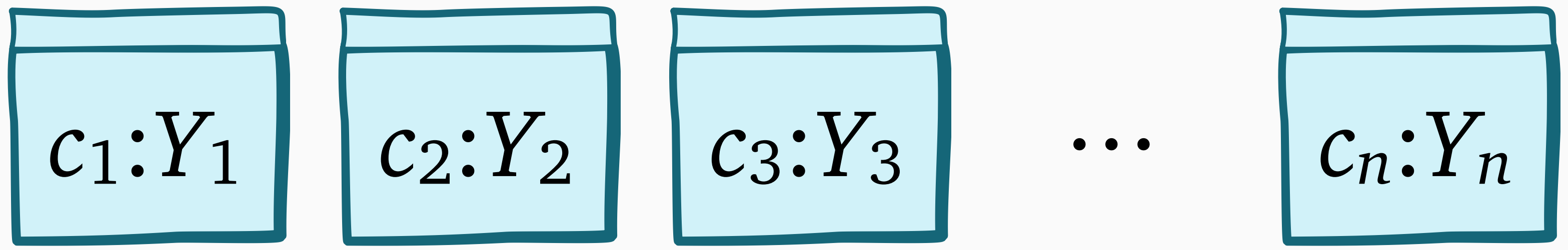

Each box:

- Opening cost c
- Hidden reward Y

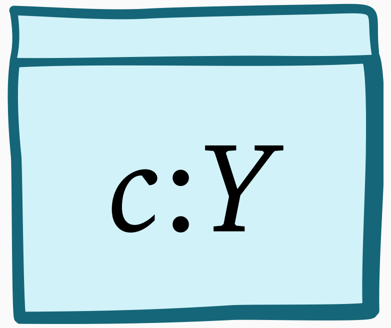

independent

Decision process:

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

Pandora's box problem

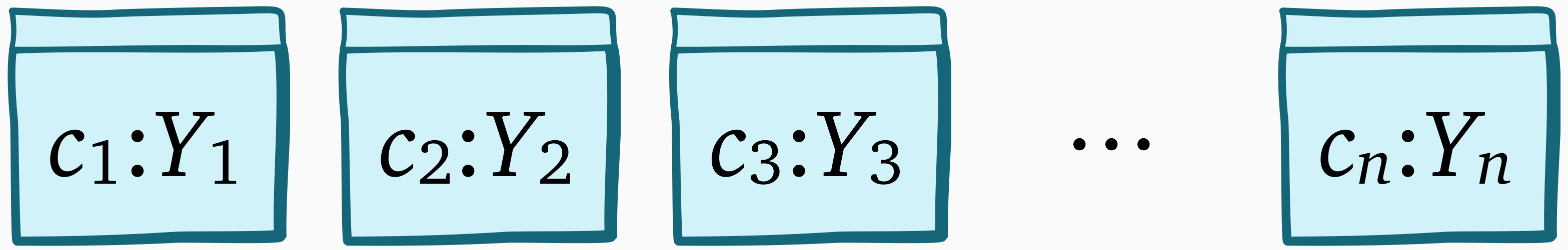

Each box:

- Opening cost c
- Hidden reward Y

independent

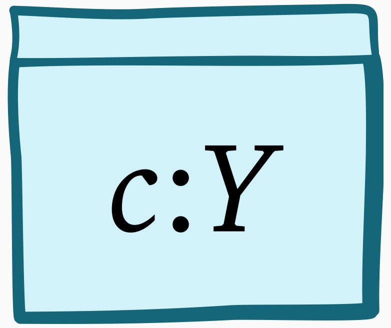

Decision process:

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

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

Pandora’s box problem

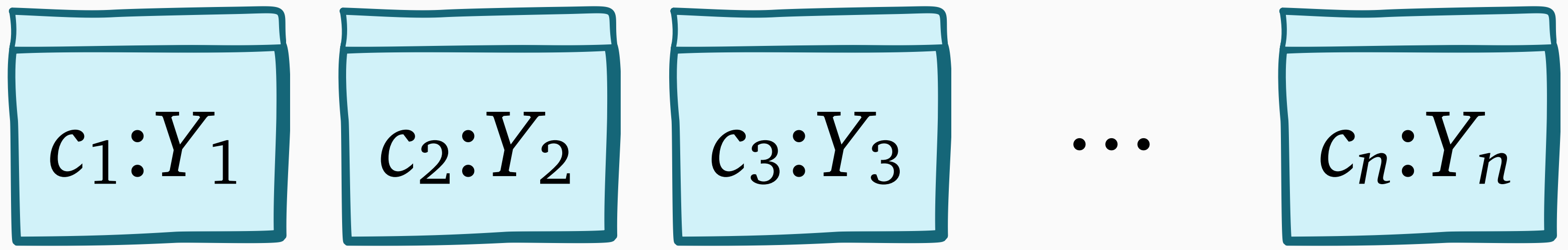

Each box:

- Opening cost c
- Hidden reward Y

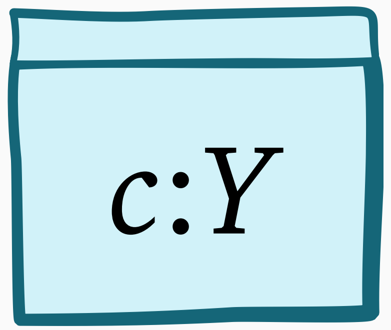

independent

Decision process:

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

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

Which box to open?

Pandora's box problem

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

Each box:

- Opening cost c
- Hidden reward Y

independent

$c : Y$

Decision process:

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

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

Which box to open?

Is it time to stop?

Why is Pandora's box hard?

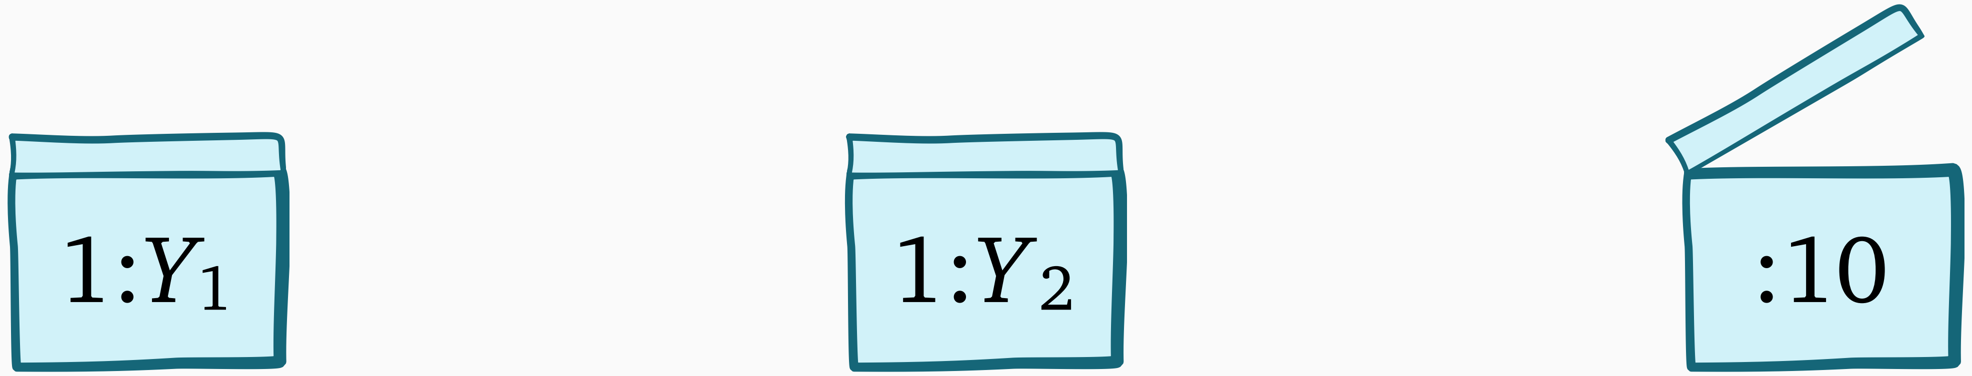

Why is Pandora's box hard?

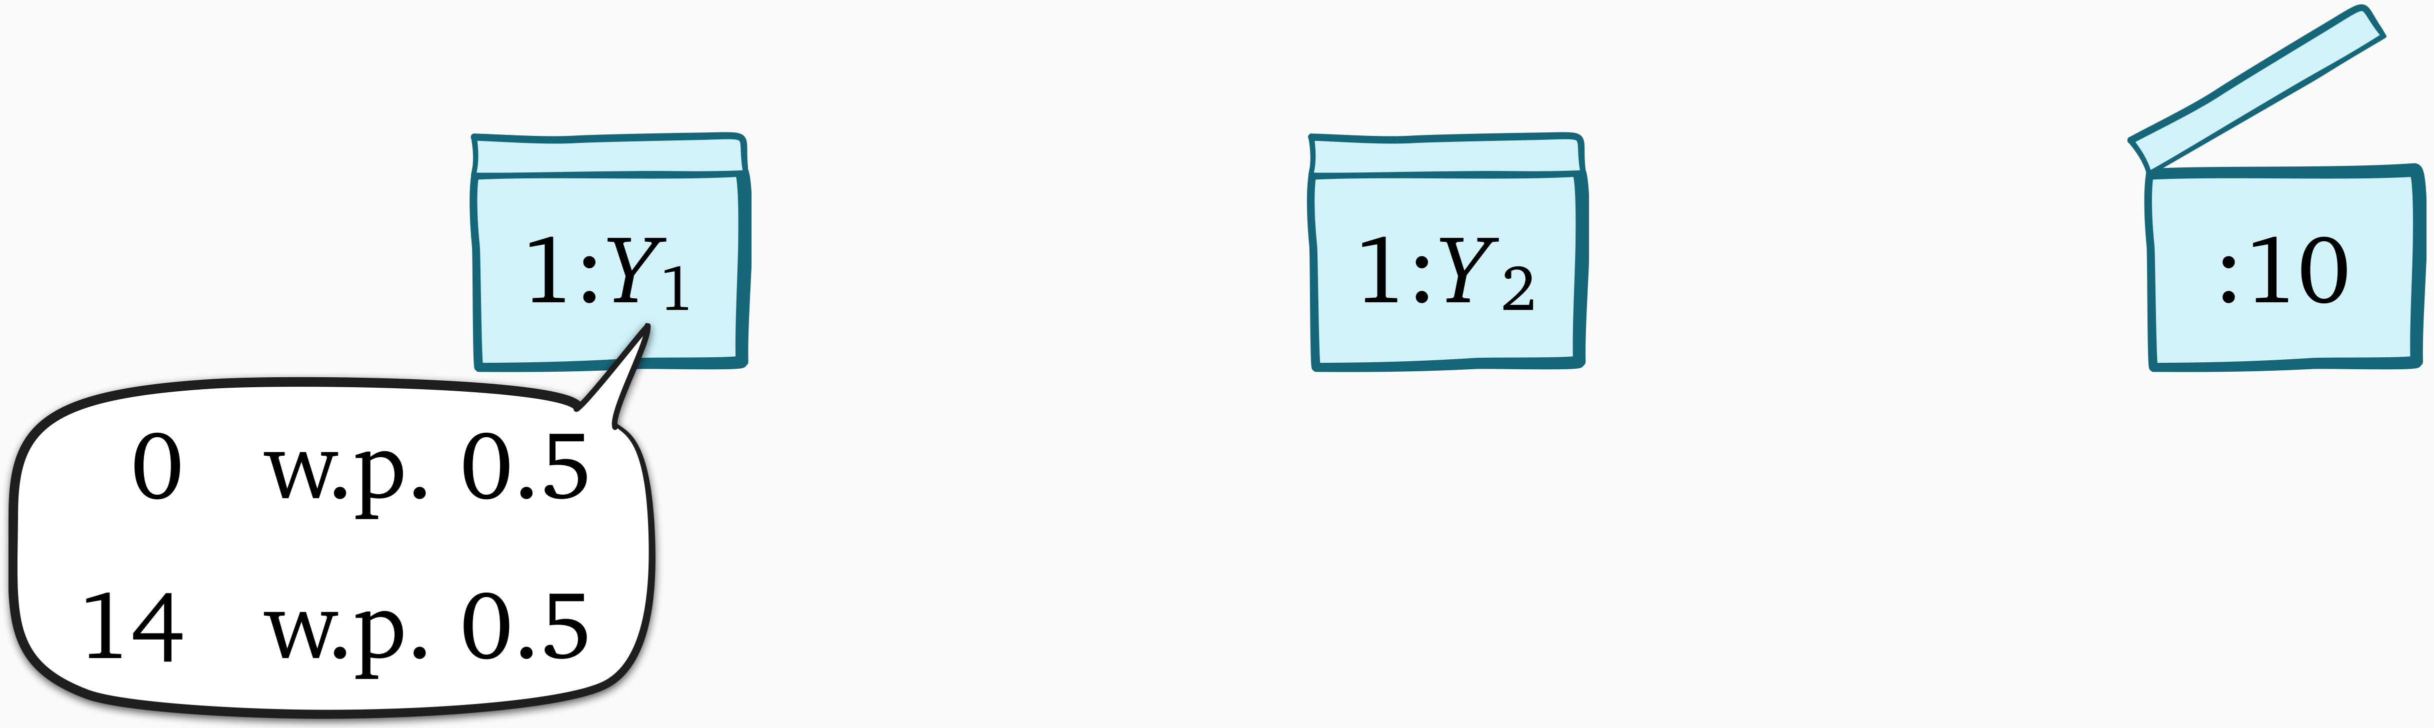

Why is Pandora's box hard?

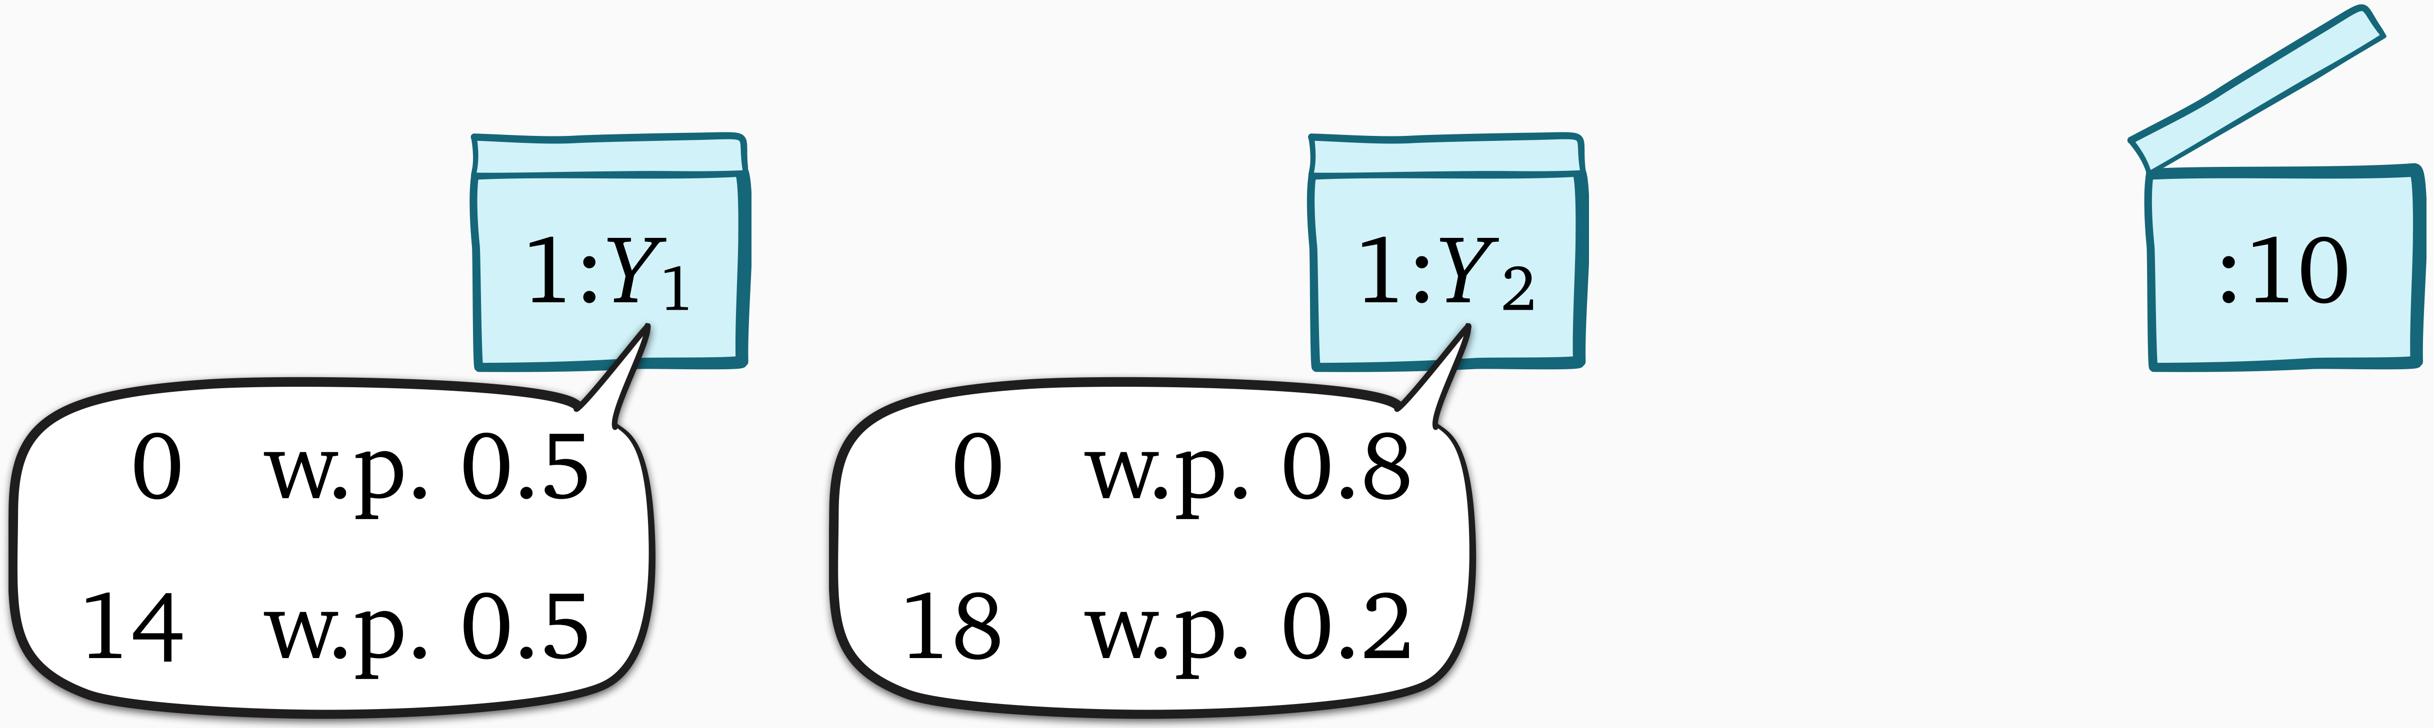

Why is Pandora's box hard?

Expected improvement of Y over r:

$$\mathrm{EI}(Y, r) = \mathbf{E}[(Y - r)^+]$$

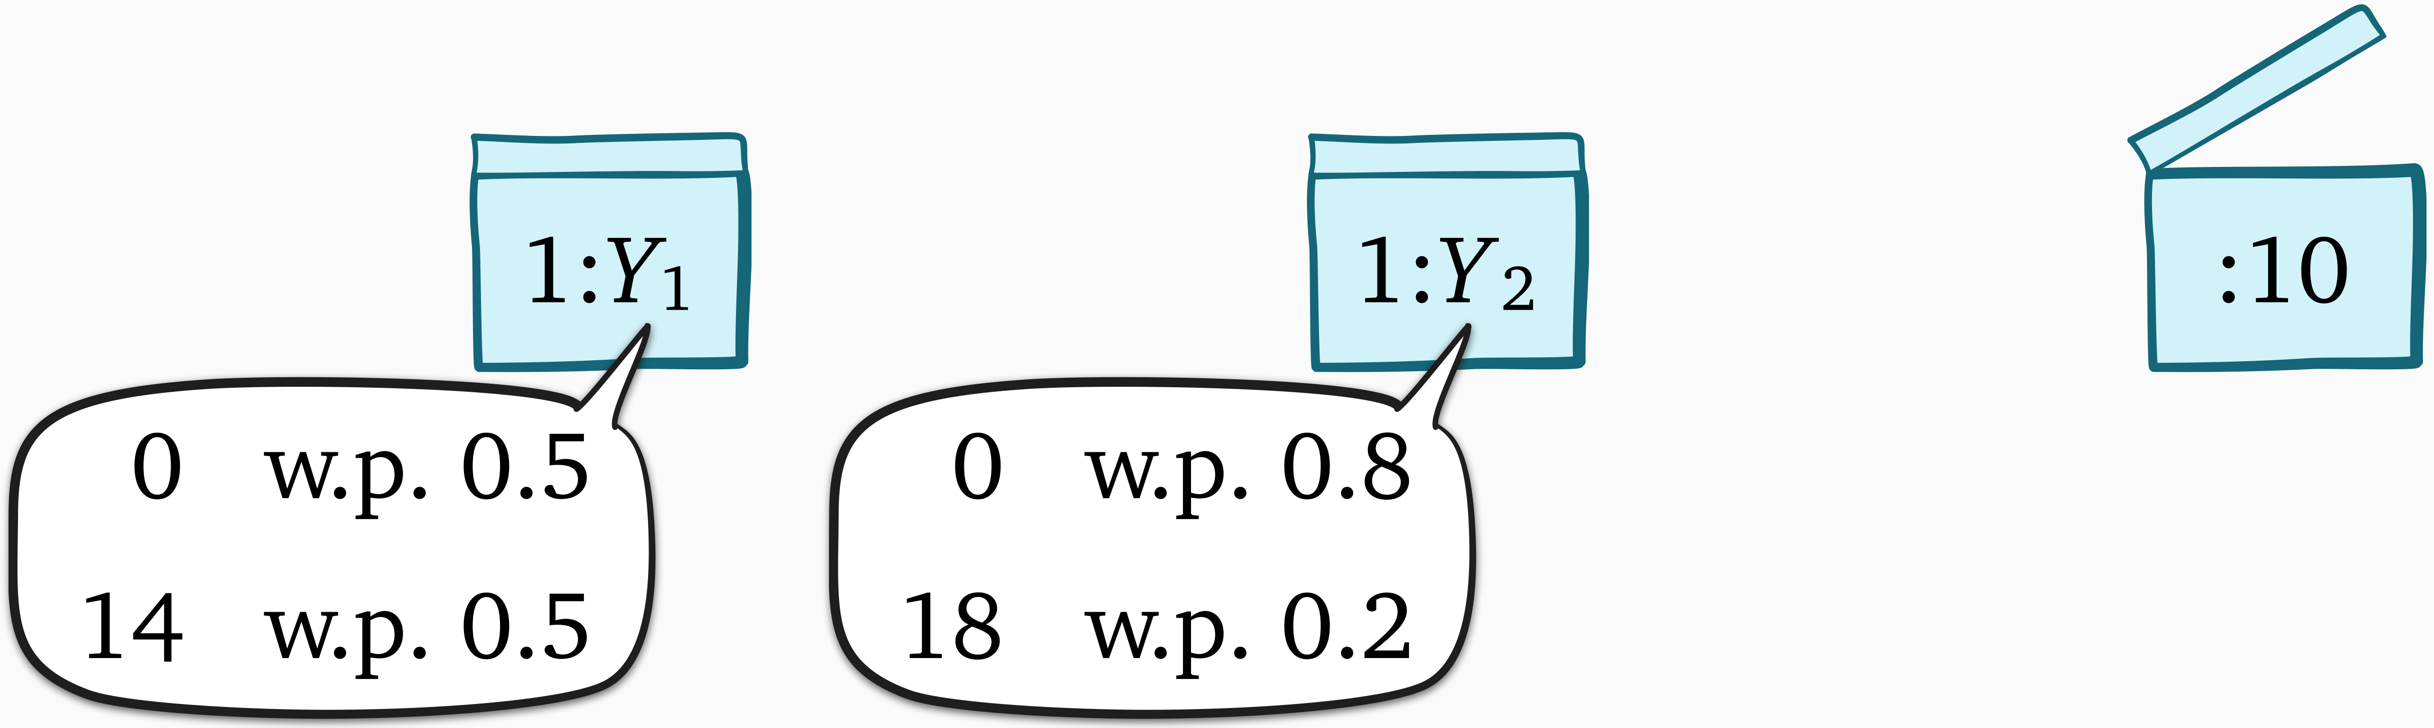

Why is Pandora's box hard?

Expected improvement of Y over r:

$$\mathrm{EI}(Y, r) = \mathbf{E}[(Y - r)^+]$$

use r = best so far

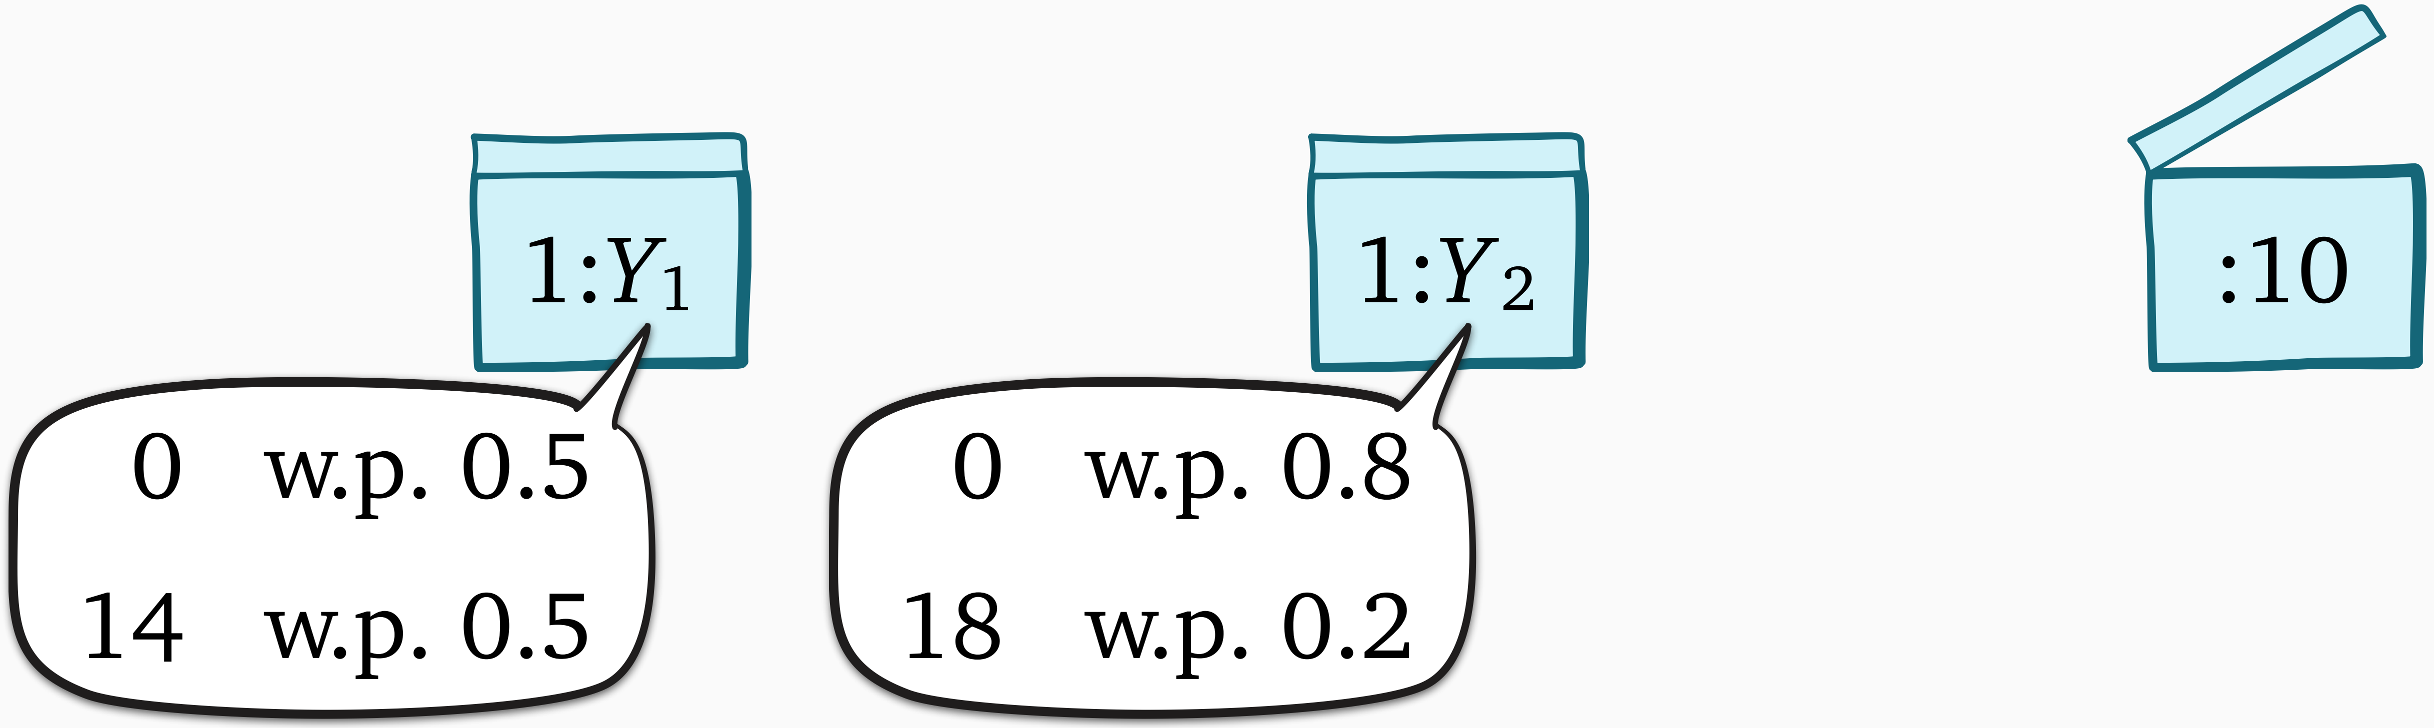

Why is Pandora's box hard?

Expected improvement of Y over r:

$$\mathrm{EI}(Y, r) = \mathbf{E}[(Y - r)^+]$$

use r = best so far

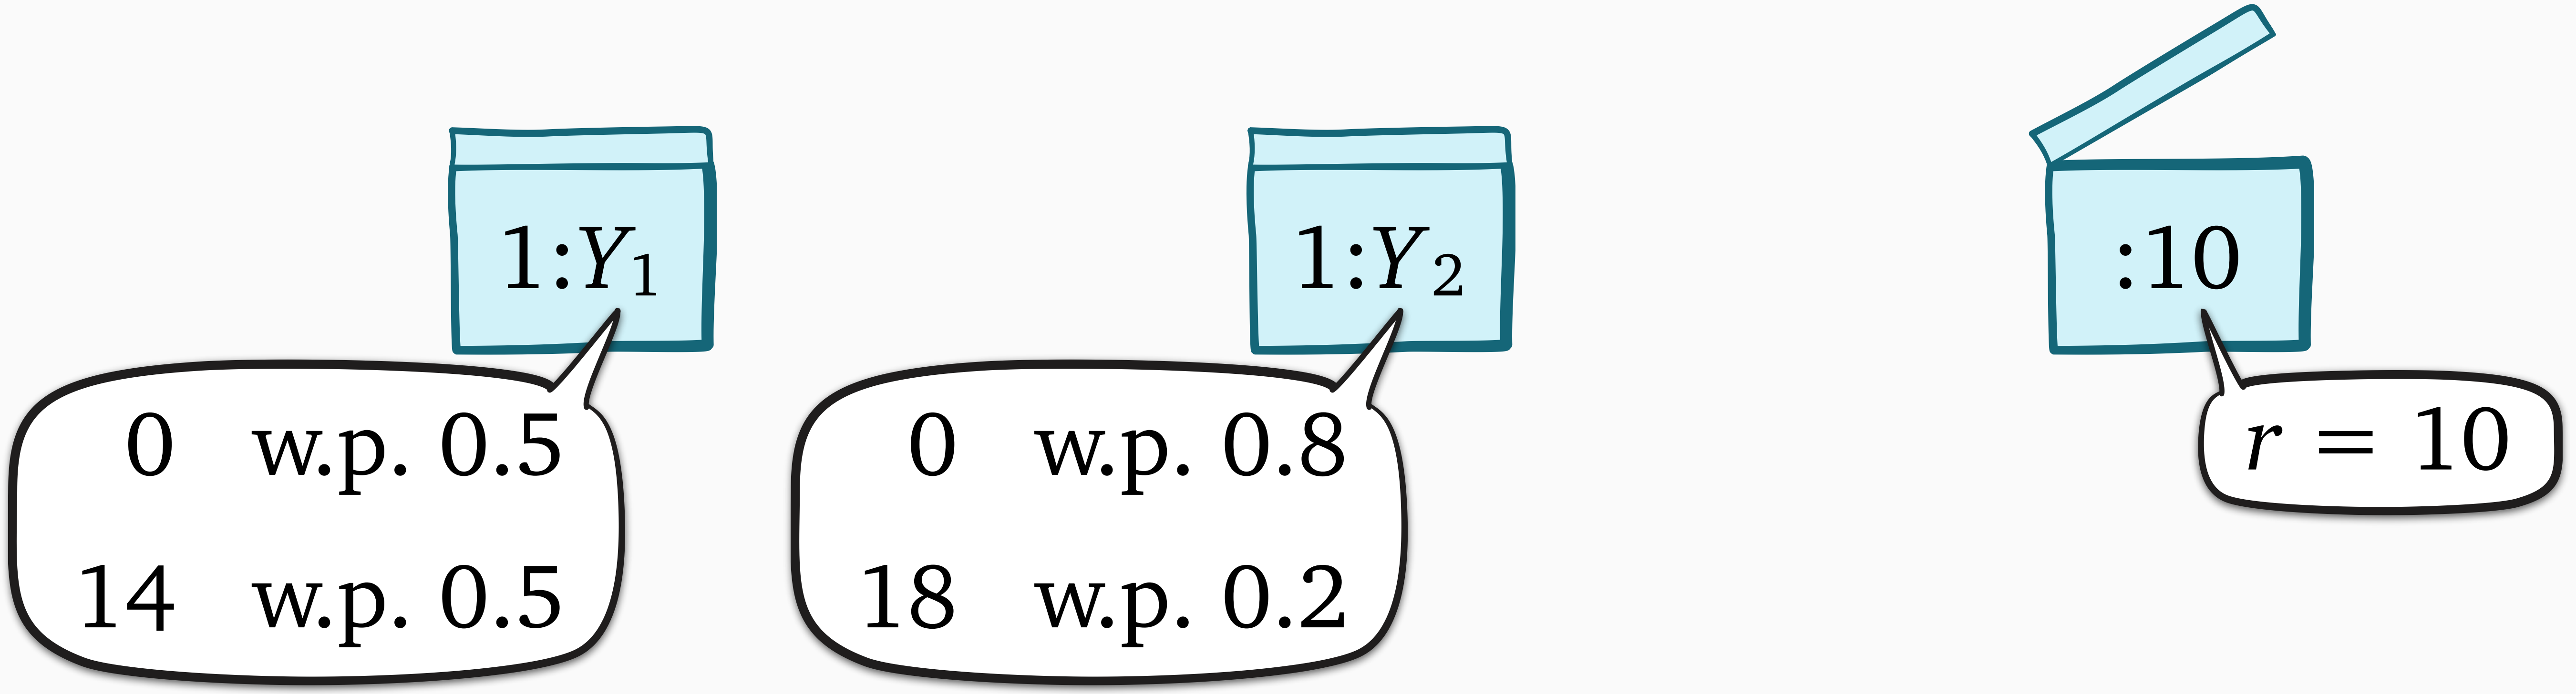

Why is Pandora's box hard?

Expected improvement of Y over r:

$$\mathrm{EI}(Y, r) = \mathbf{E}[(Y - r)^+]$$

use r = best so far

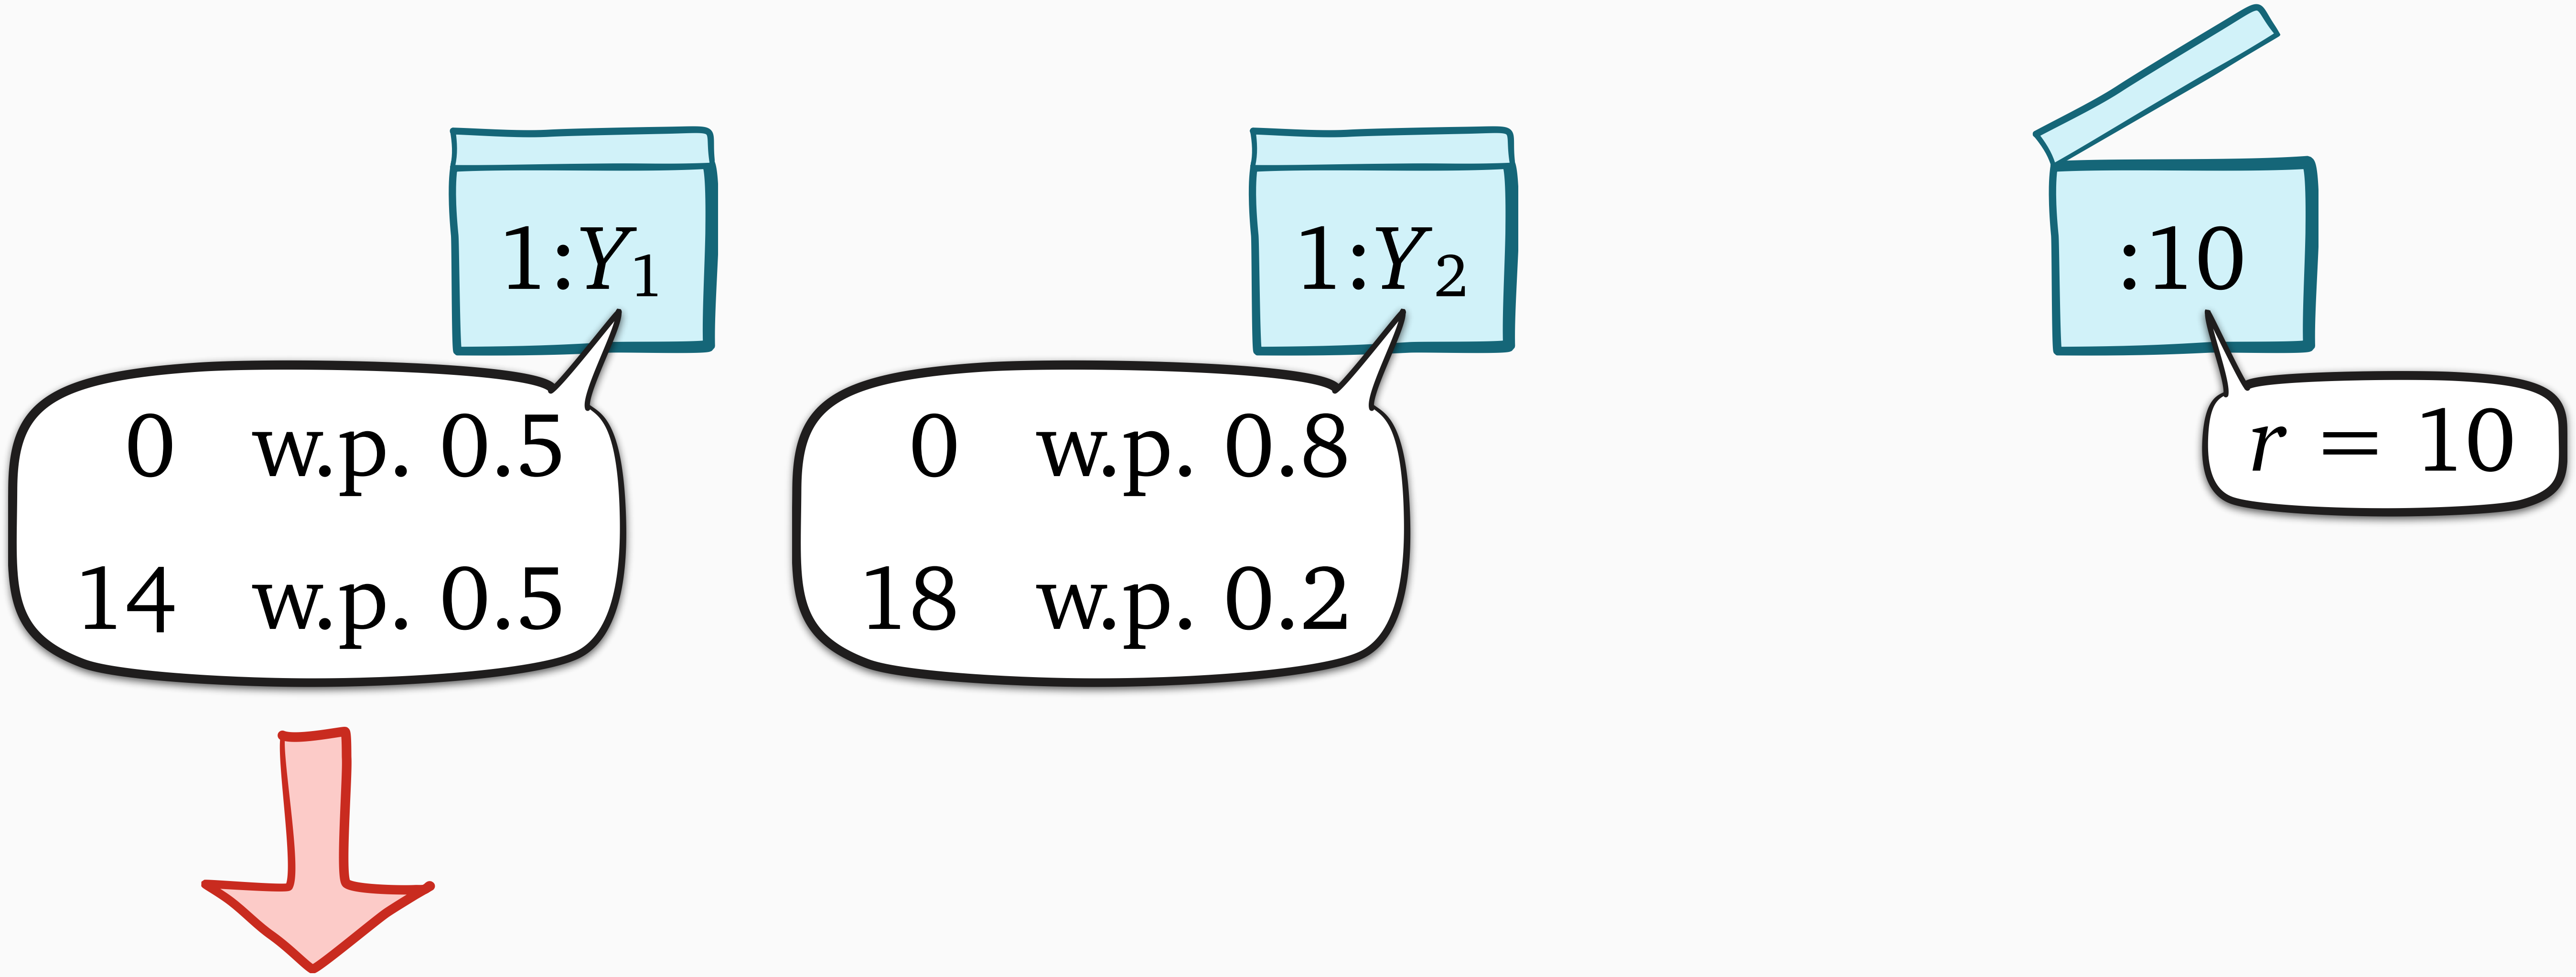

EI$(X_1, 10) = 2$

Why is Pandora's box hard?

Expected improvement of Y over r:

$$\mathrm{EI}(Y, r) = \mathbf{E}[(Y - r)^+]$$

use r = best so far

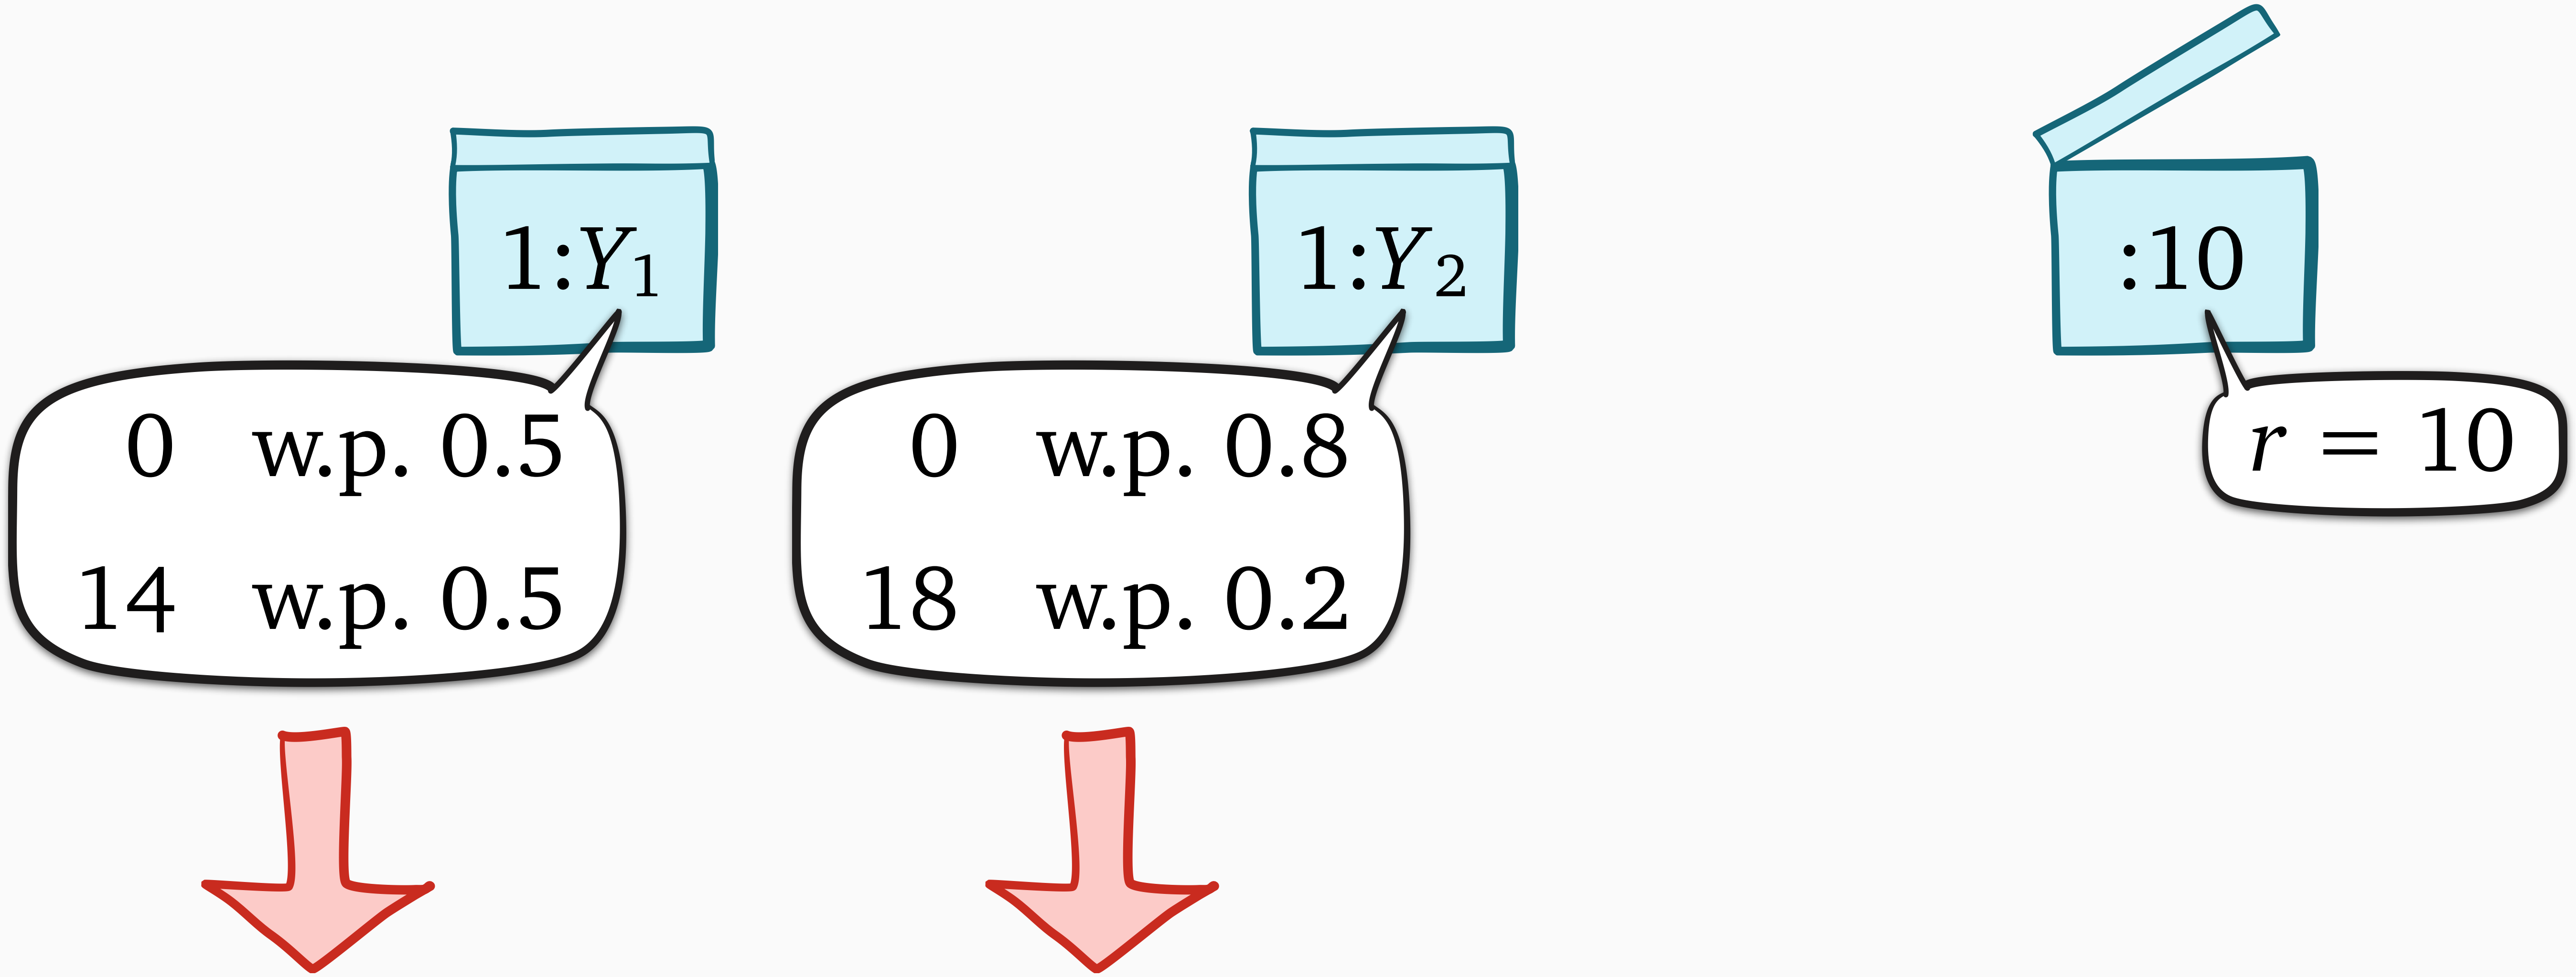

$\mathrm{EI}(X_1, 10) = 2$ $\mathrm{EI}(X_2, 10) = 1.6$

Why is Pandora's box hard?

Expected improvement of Y over r:

$$\mathrm{EI}(Y, r) = \mathbf{E}[(Y - r)^+]$$

use r = best so far

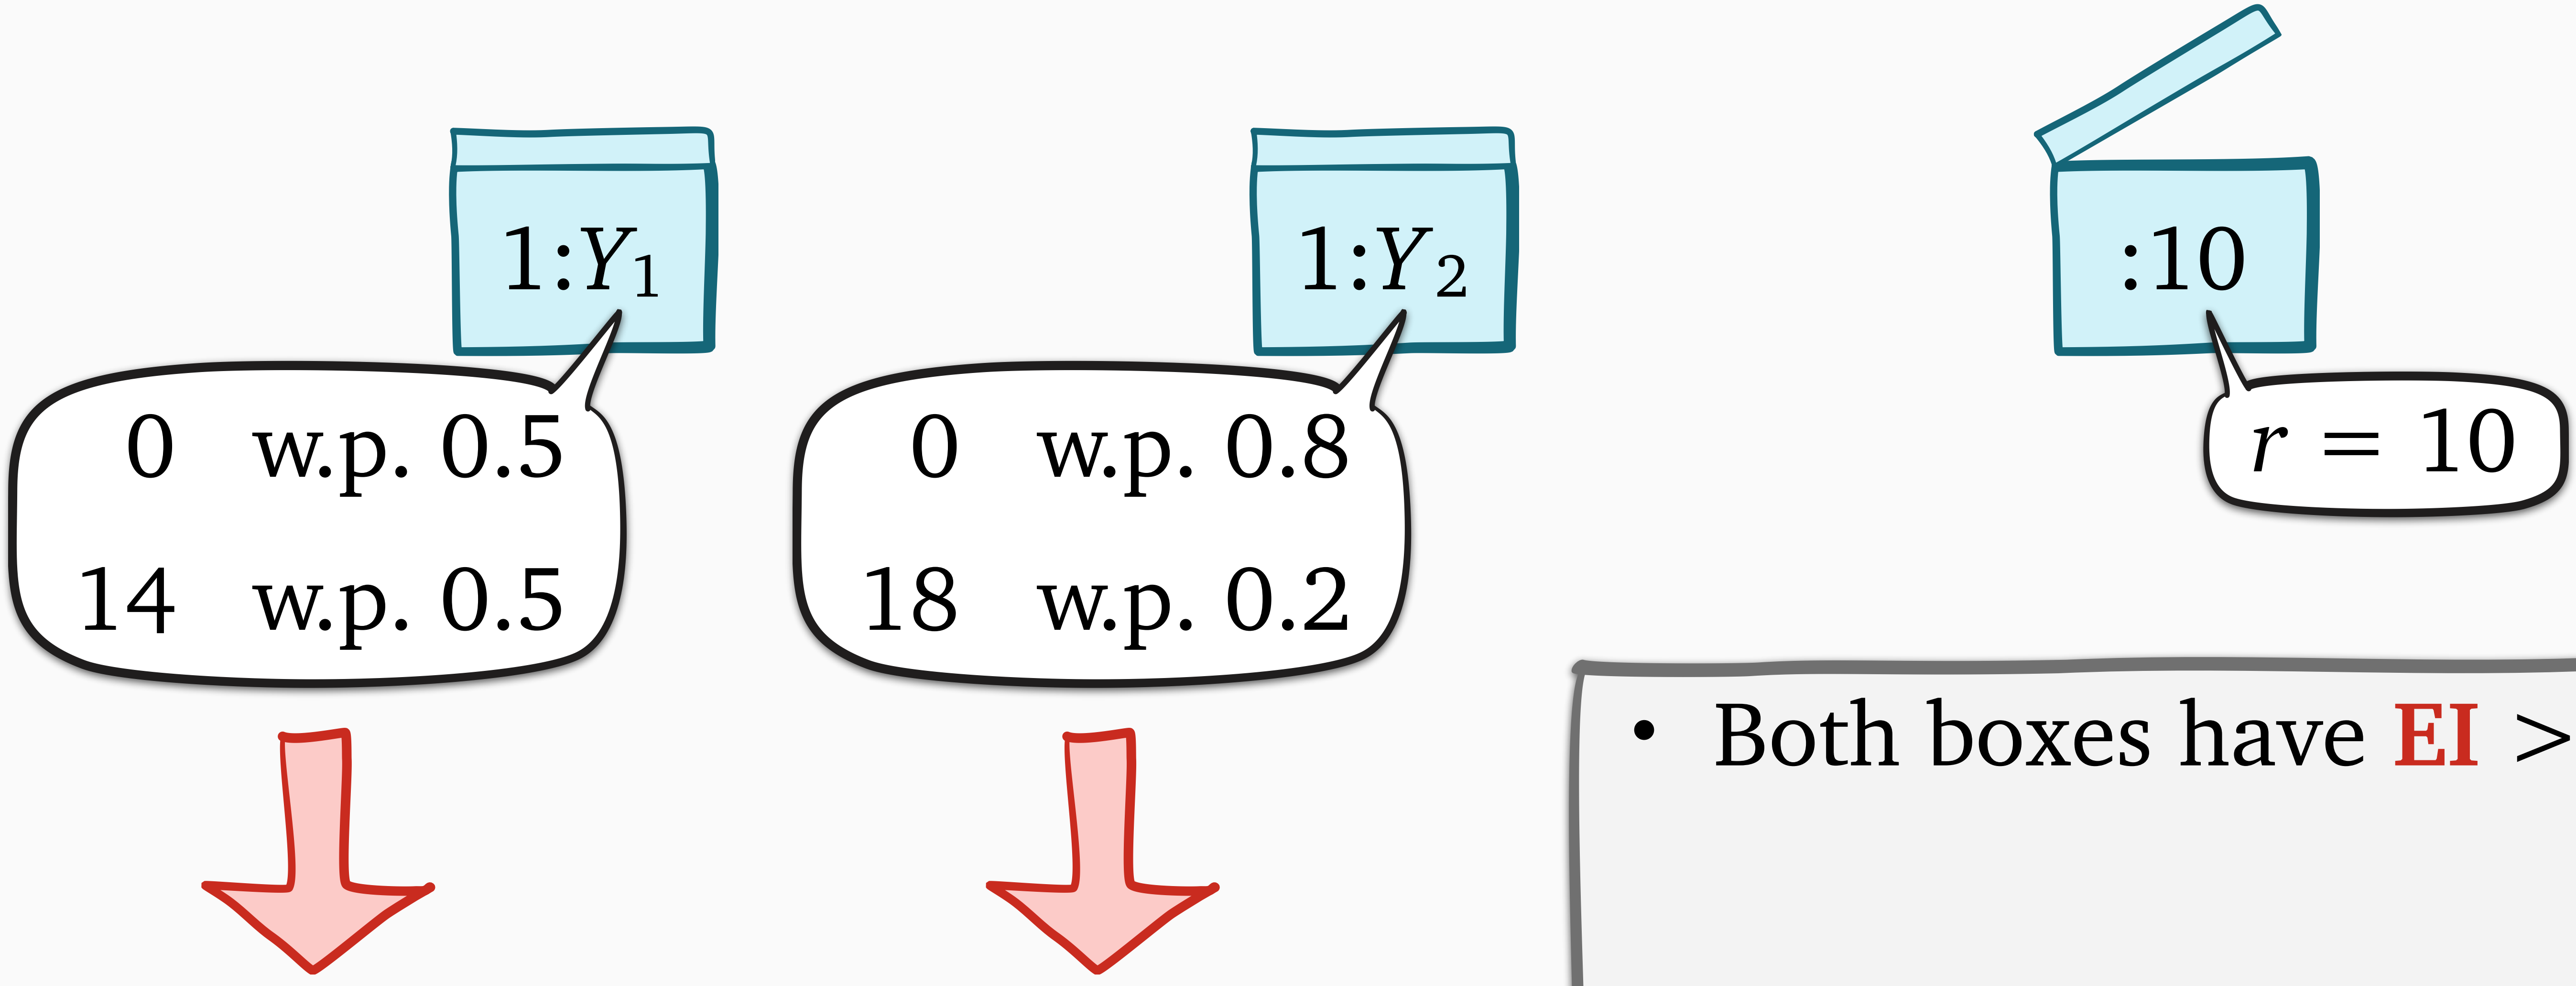

EI$(X_1, 10) = 2$ **EI**$(X_2, 10) = 1.6$

- Both boxes have **EI** > open cost

Why is Pandora's box hard?

Expected improvement of Y over r:

$$\mathrm{EI}(Y, r) = \mathbf{E}[(Y - r)^+]$$

use r = best so far

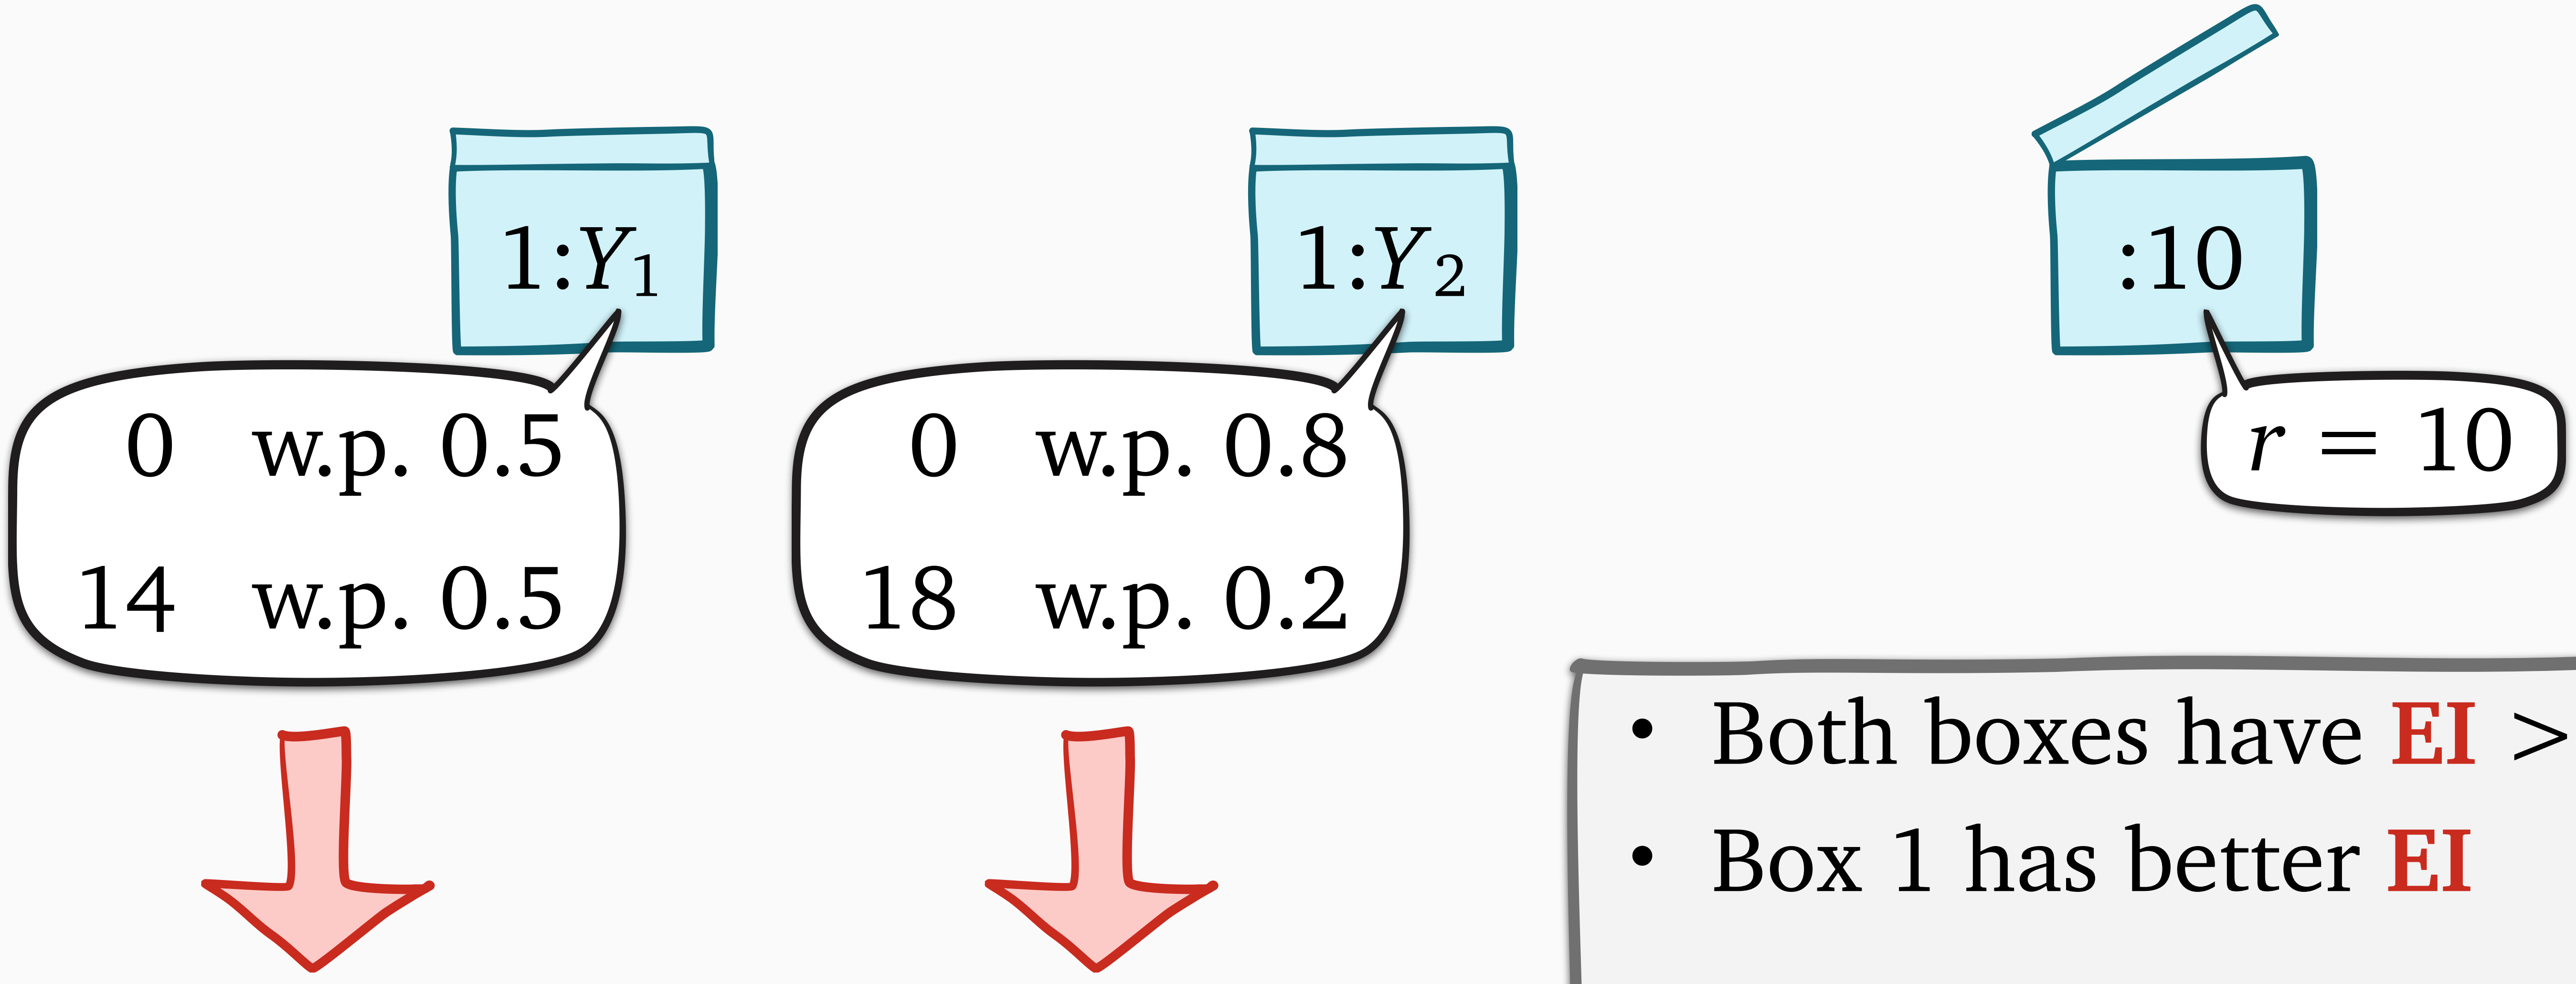

EI$(X_1, 10) = 2$ **EI**$(X_2, 10) = 1.6$

- Both boxes have **EI** > open cost
- Box 1 has better **EI**

Why is Pandora's box hard?

Expected improvement of Y over r:

$$\mathrm{EI}(Y, r) = \mathbf{E}[(Y - r)^+]$$

use r = best so far

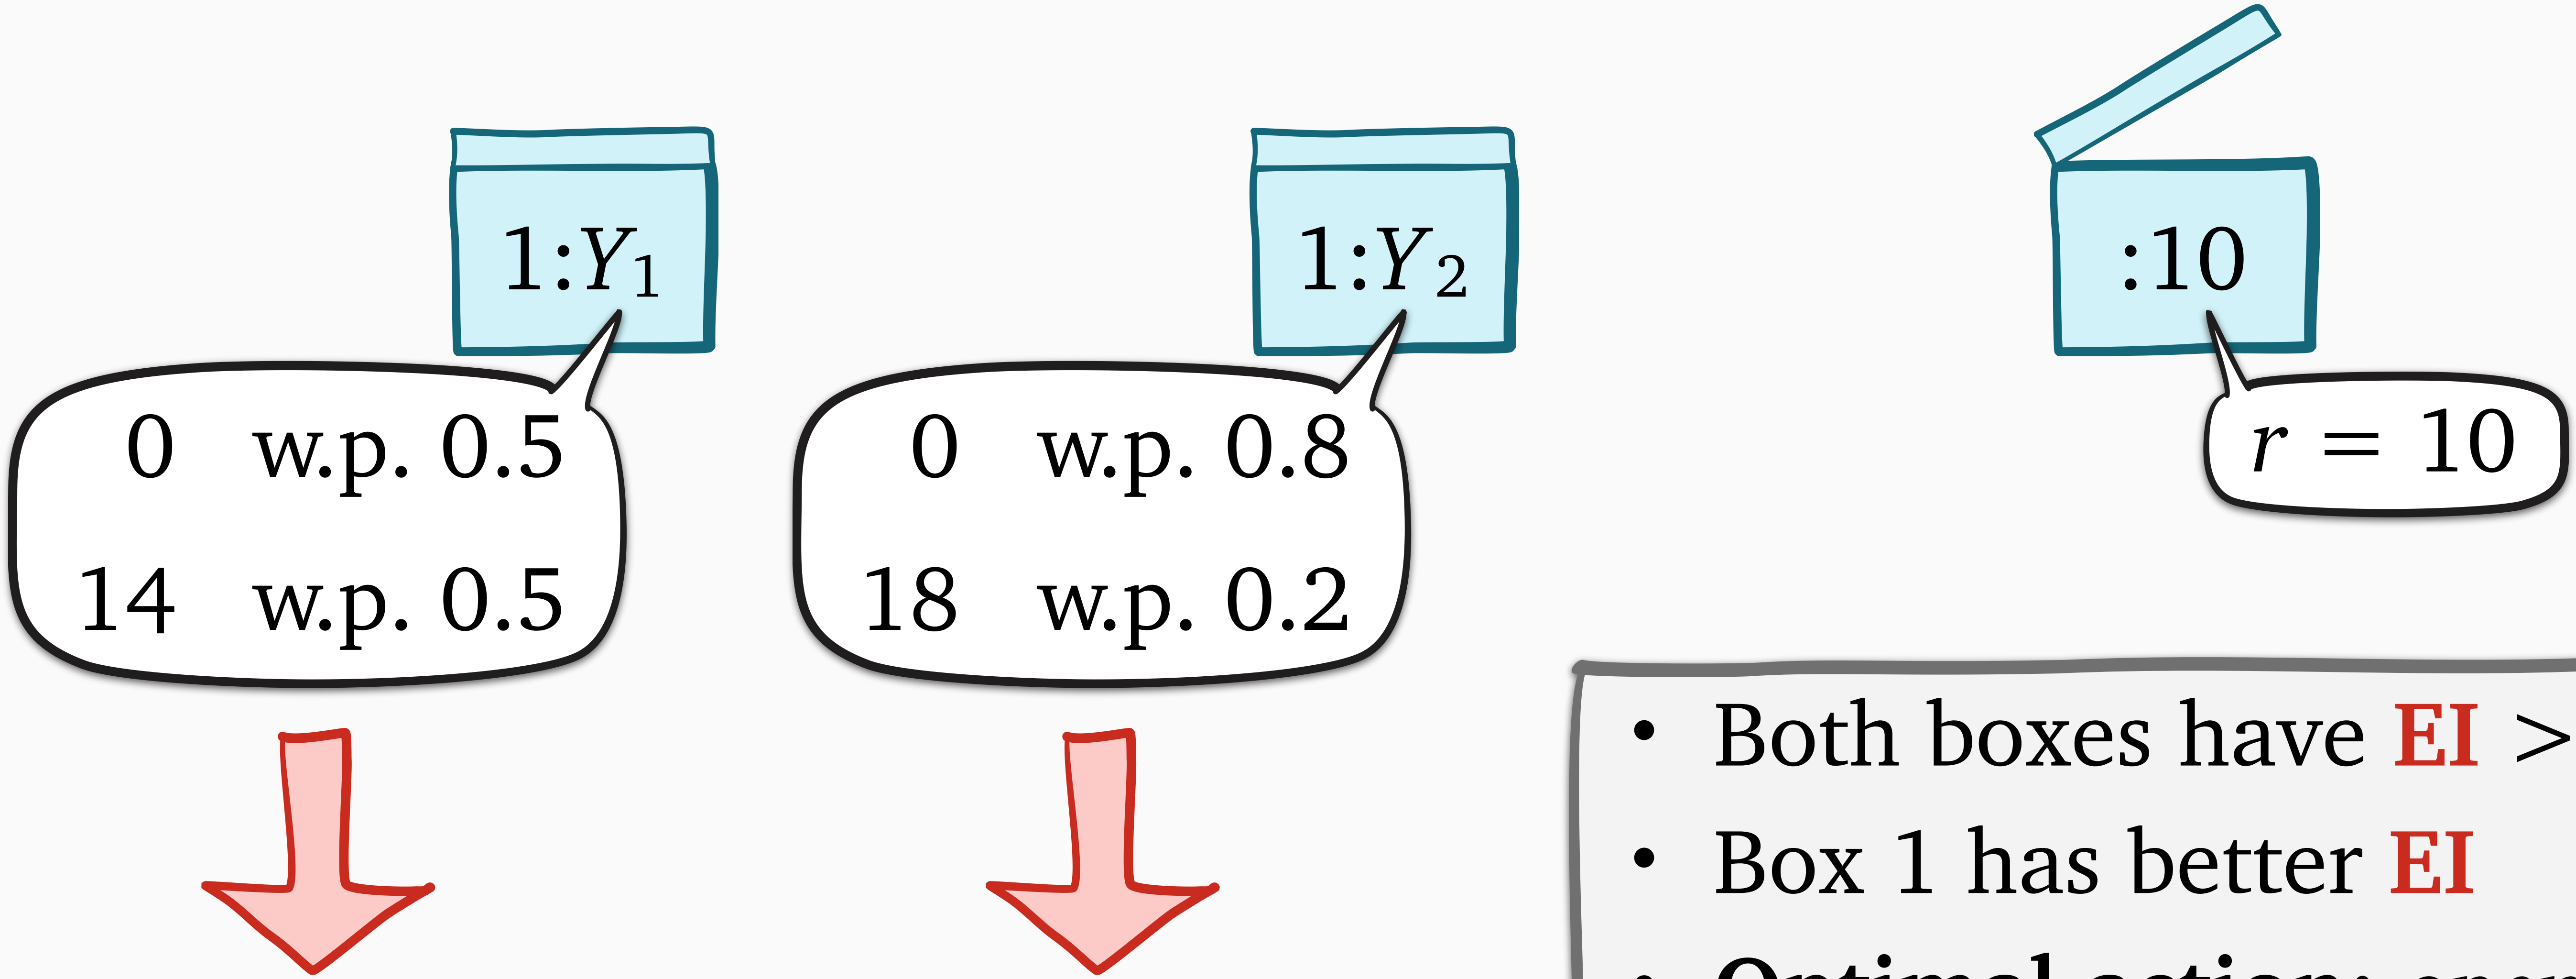

EI$(X_1, 10) = 2$ **EI**$(X_2, 10) = 1.6$

- Both boxes have **EI** > open cost
- Box 1 has better **EI**
- **Optimal action:** *open box 2!*

Why is Pandora's box hard?

⚠ Just using **EI** is suboptimal!

Expected improvement of Y over r:

$$\mathrm{EI}(Y, r) = \mathbf{E}[(Y - r)^+]$$

use r = best so far

$1{:}Y_1$	0 w.p. 0.5; 14 w.p. 0.5
$1{:}Y_2$	0 w.p. 0.8; 18 w.p. 0.2
$:10$	$r = 10$

EI$(X_1, 10) = 2$

EI$(X_2, 10) = 1.6$

- Both boxes have **EI** > open cost
- Box 1 has better **EI**
- **Optimal action:** *open box 2!*

Optimal policy: **Gittins**

Optimal policy: **Gittins**

Step 1: *rate* each box separately

Step 2: *act* on box of best rating

Optimal policy: **Gittins**

Step 1: *rate* each box separately

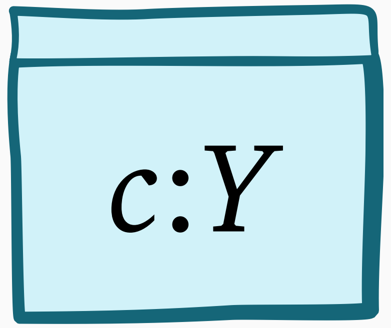

Step 2: *act* on box of best rating

Optimal policy: Gittins

Step 1: *rate* each box separately

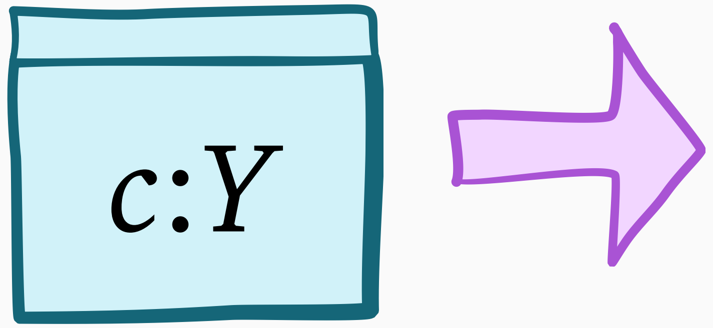

Gittins index:

$g(c{:}Y)$

Step 2: *act* on box of best rating

Optimal policy: Gittins

Step 1: *rate* each box separately

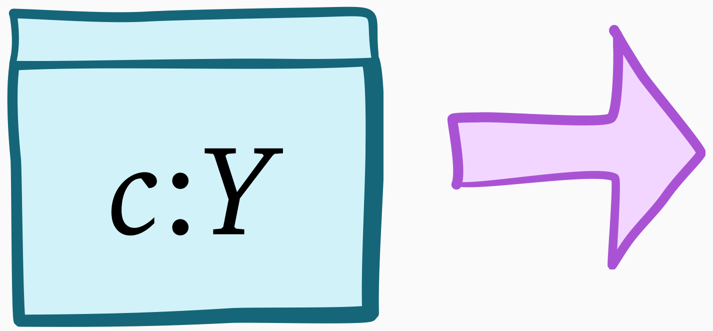

Gittins index:

$$g(c{:}Y)$$

higher is better

Step 2: *act* on box of best rating

Optimal policy: Gittins

Step 1: *rate* each box separately

$c{:}Y$ ⇒ **Gittins index:** $g(c{:}Y)$

higher is better

Step 2: *act* on box of best rating

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

Optimal policy: Gittins

Step 1: *rate* each box separately

Step 2: *act* on box of best rating

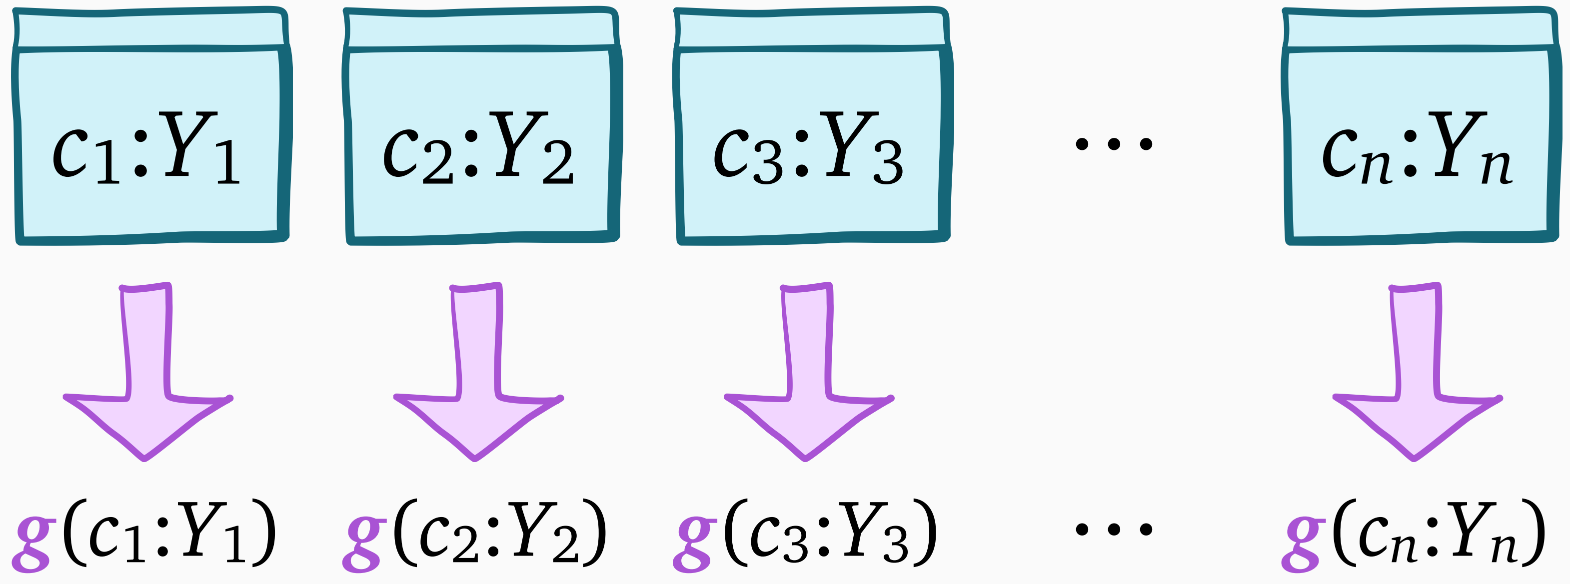

Optimal policy: Gittins

Step 1: *rate* each box separately

Step 2: *act* on box of best rating

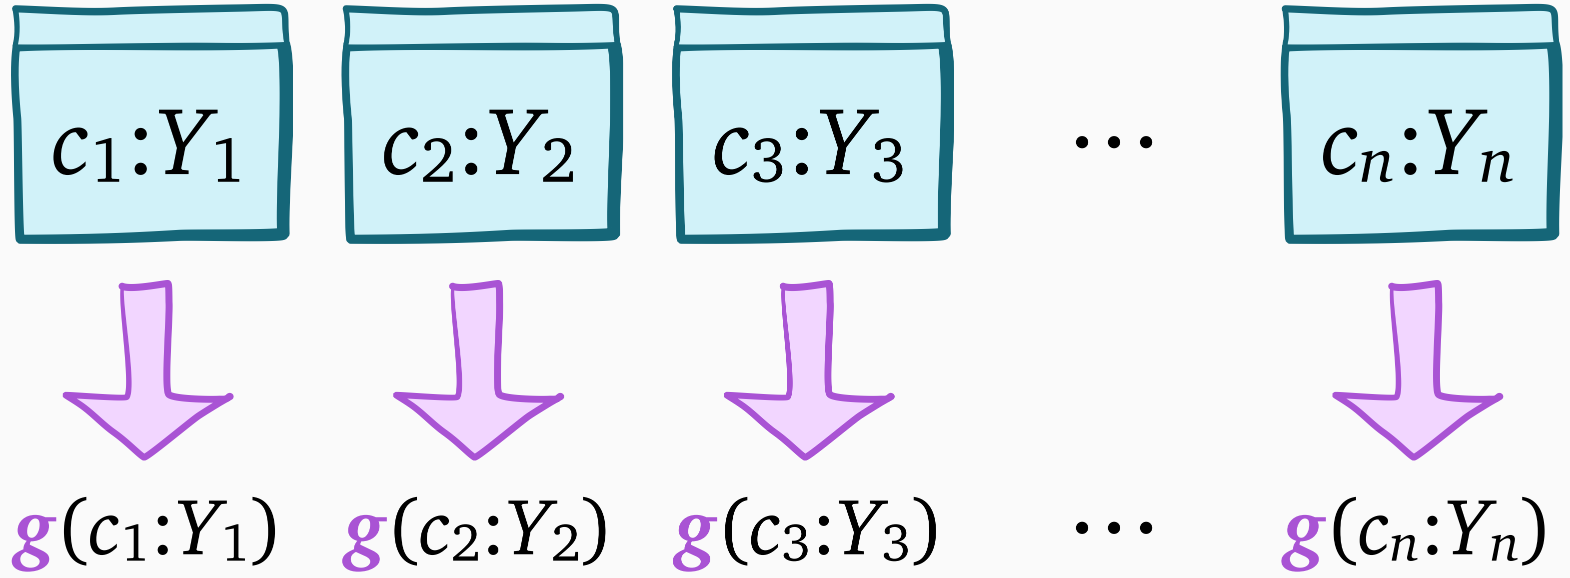

Gittins **policy:** if box of max Gittins index is…

- *closed:* open it
- *open:* select it

Optimal policy: Gittins

Step 1: *rate* each box separately

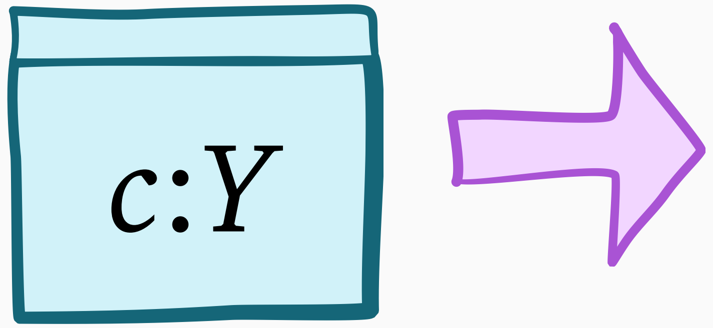

Gittins index:

$g(c{:}Y)$

higher is better

Step 2: *act* on box of best rating

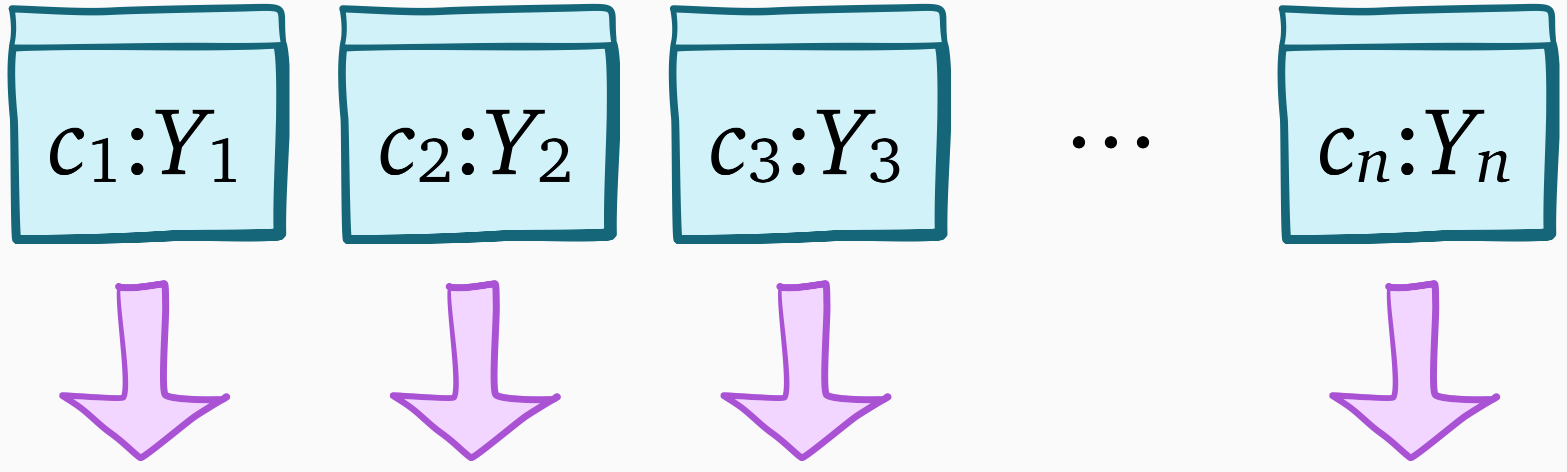

$g(c_1{:}Y_1)$ $g(c_2{:}Y_2)$ $g(c_3{:}Y_3)$ … $g(c_n{:}Y_n)$

Gittins policy: if box of max Gittins index is…

- *closed:* open it
- *open:* select it

} *act* on it

Optimal policy: Gittins

Step 1: *rate* each box separately

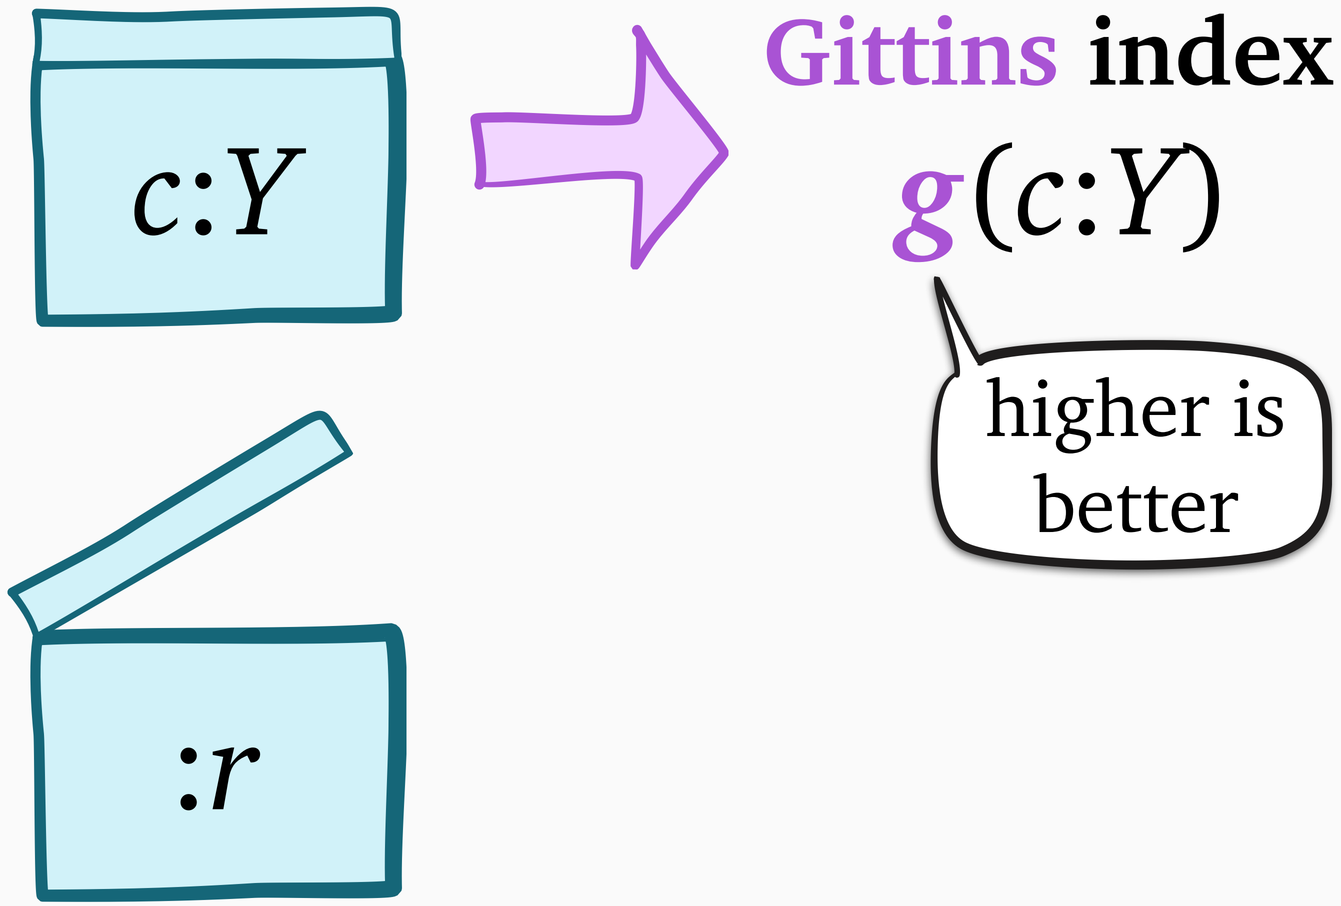

Step 2: *act* on box of best rating

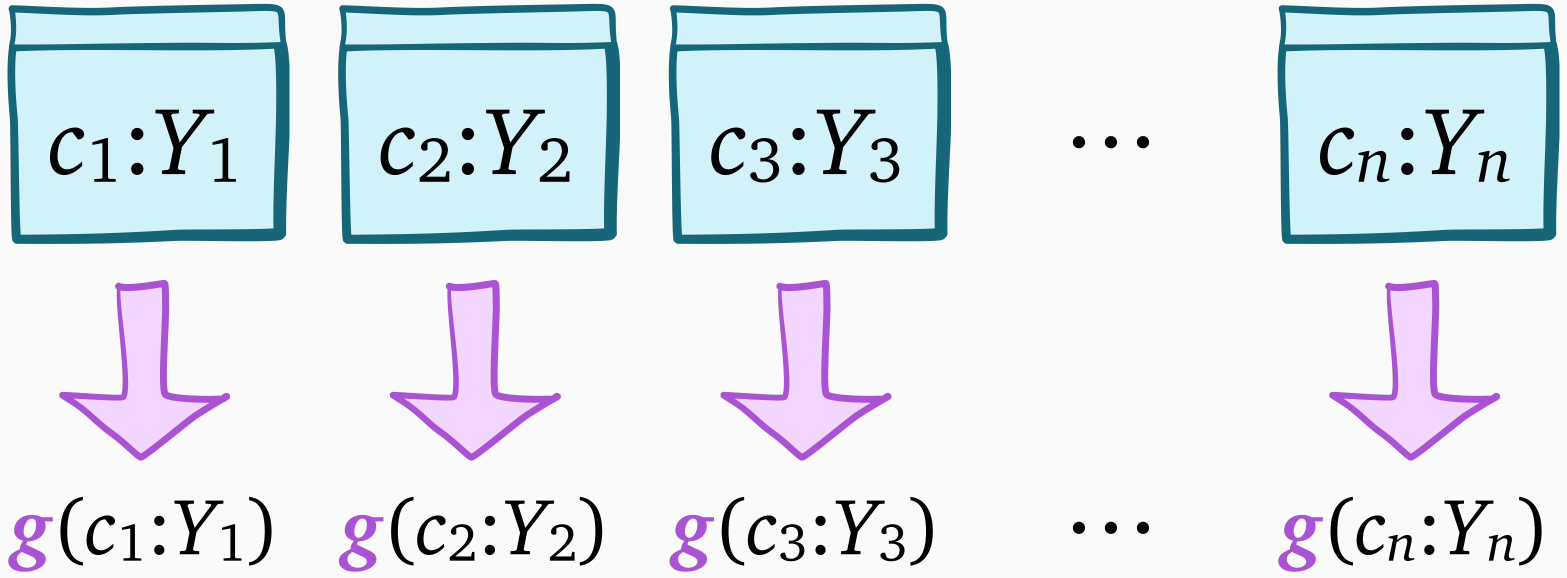

Gittins **policy:** if box of max Gittins index is…

- *closed:* open it
- *open:* select it

} *act* on it

Optimal policy: Gittins

Step 1: *rate* each box separately

$c{:}Y$ → **Gittins index:** $g(c{:}Y)$

higher is better

$:r$ → $g(:r) = r$

Step 2: *act* on box of best rating

$c_1{:}Y_1$ → $g(c_1{:}Y_1)$, $c_2{:}Y_2$ → $g(c_2{:}Y_2)$, $c_3{:}Y_3$ → $g(c_3{:}Y_3)$, $\cdots$, $c_n{:}Y_n$ → $g(c_n{:}Y_n)$

Gittins policy: if box of max Gittins index is…

- *closed:* open it
- *open:* select it

} *act* on it

Optimal policy: Gittins

Step 1: *rate* each box separately

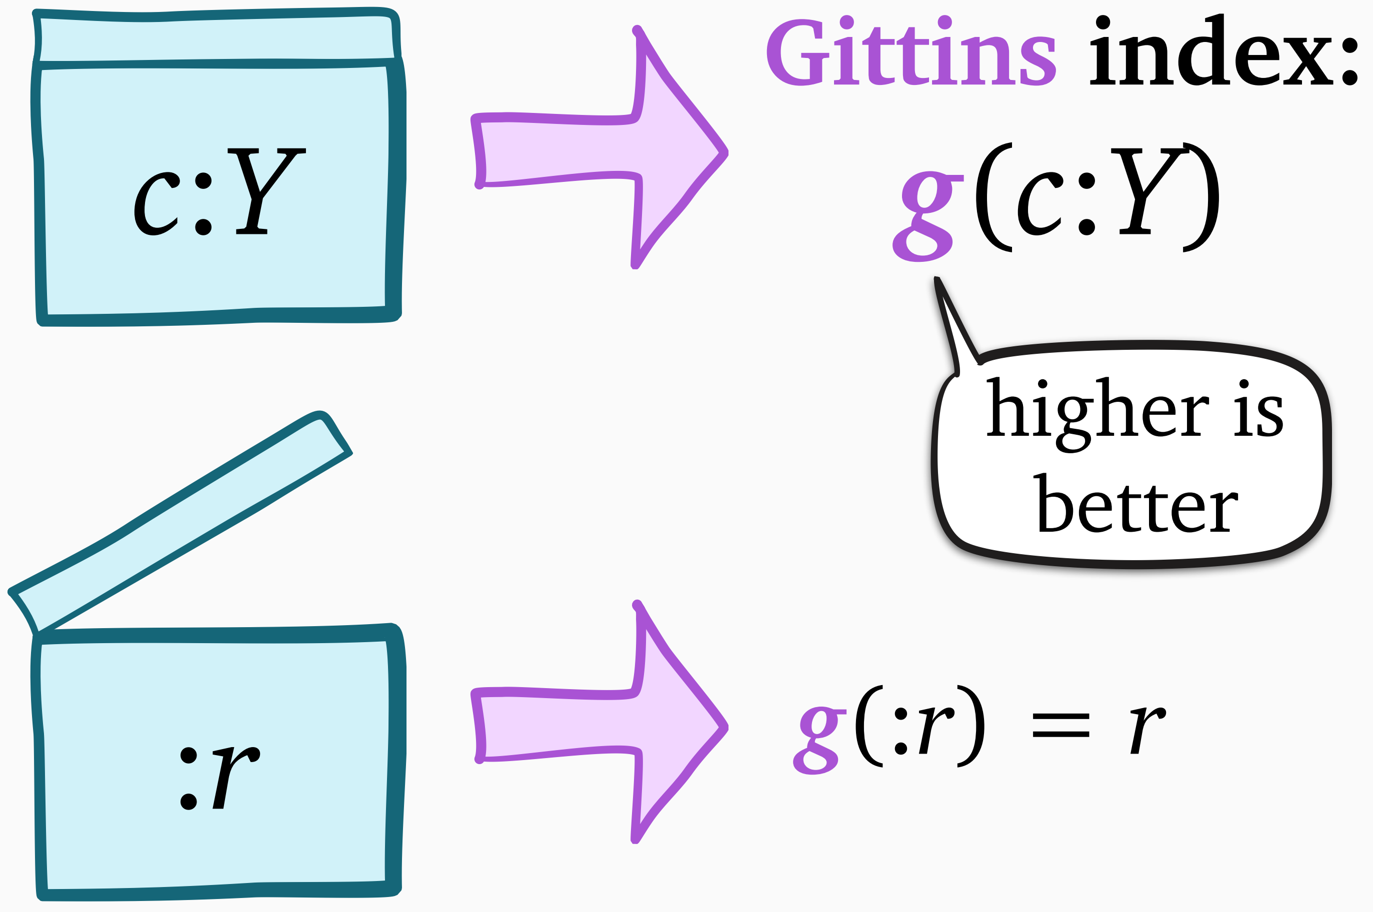

Theorem: [Weitzman, 1979]
the Gittins policy is optimal

Step 2: *act* on box of best rating

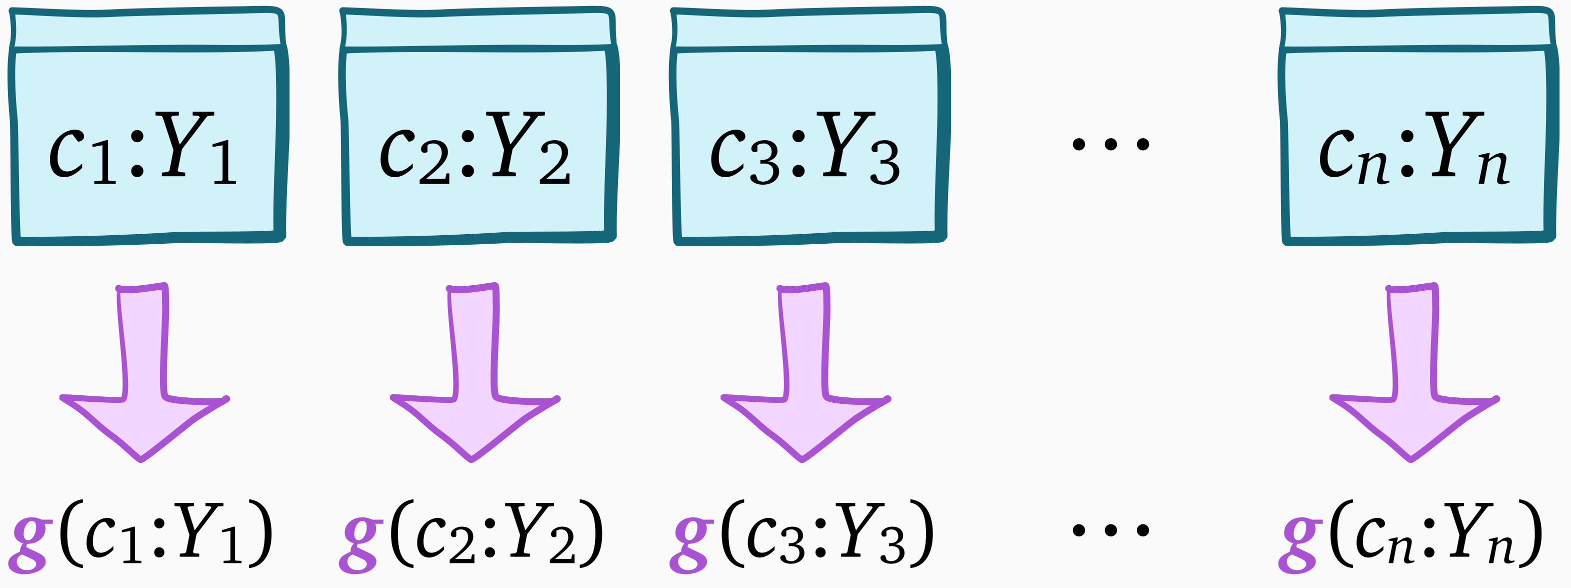

Gittins policy: if box of max Gittins index is…

- *closed:* open it
- *open:* select it

} *act* on it

Defining the **Gittins** index

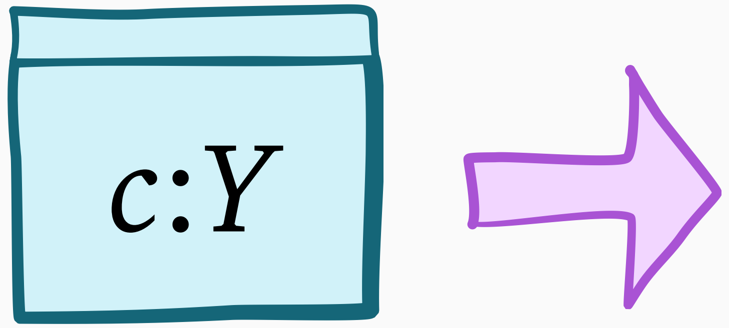

$g(c:Y)$

Defining the Gittins index

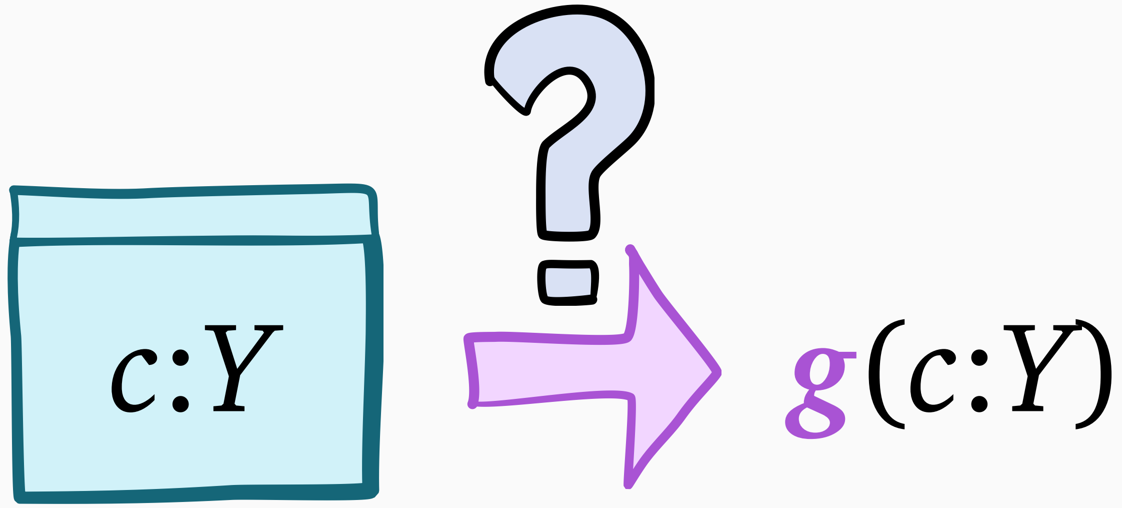

Defining the **Gittins** index

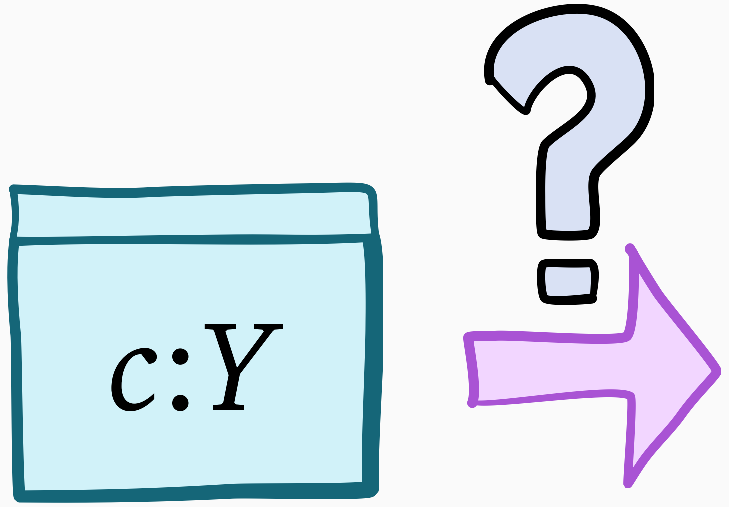

$g(c{:}Y)$

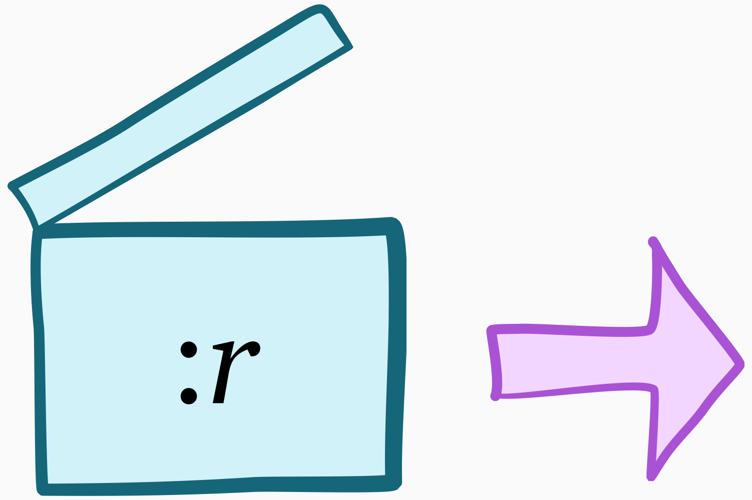

$g({:}r) = r$

Defining the **Gittins** index

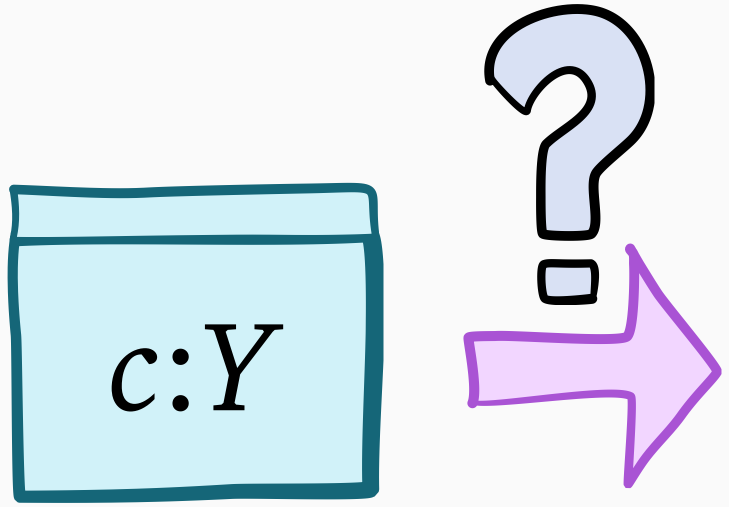

$\mathbf{g}(c{:}Y)$

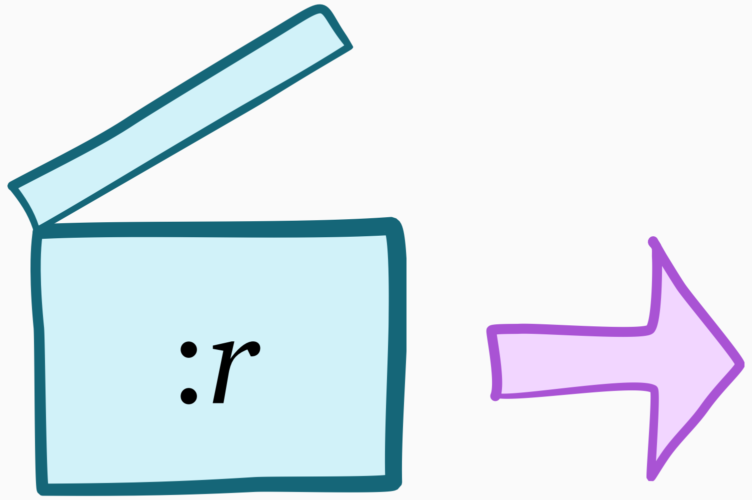

$\mathbf{g}({:}r) = r$

1.5-box problem

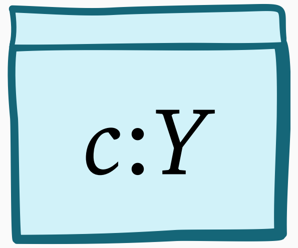

vs.

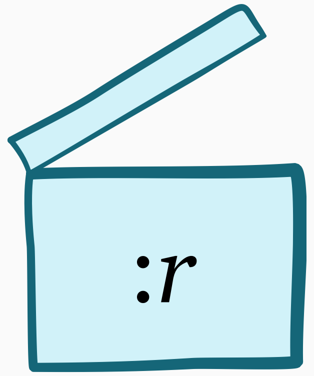

Defining the Gittins index

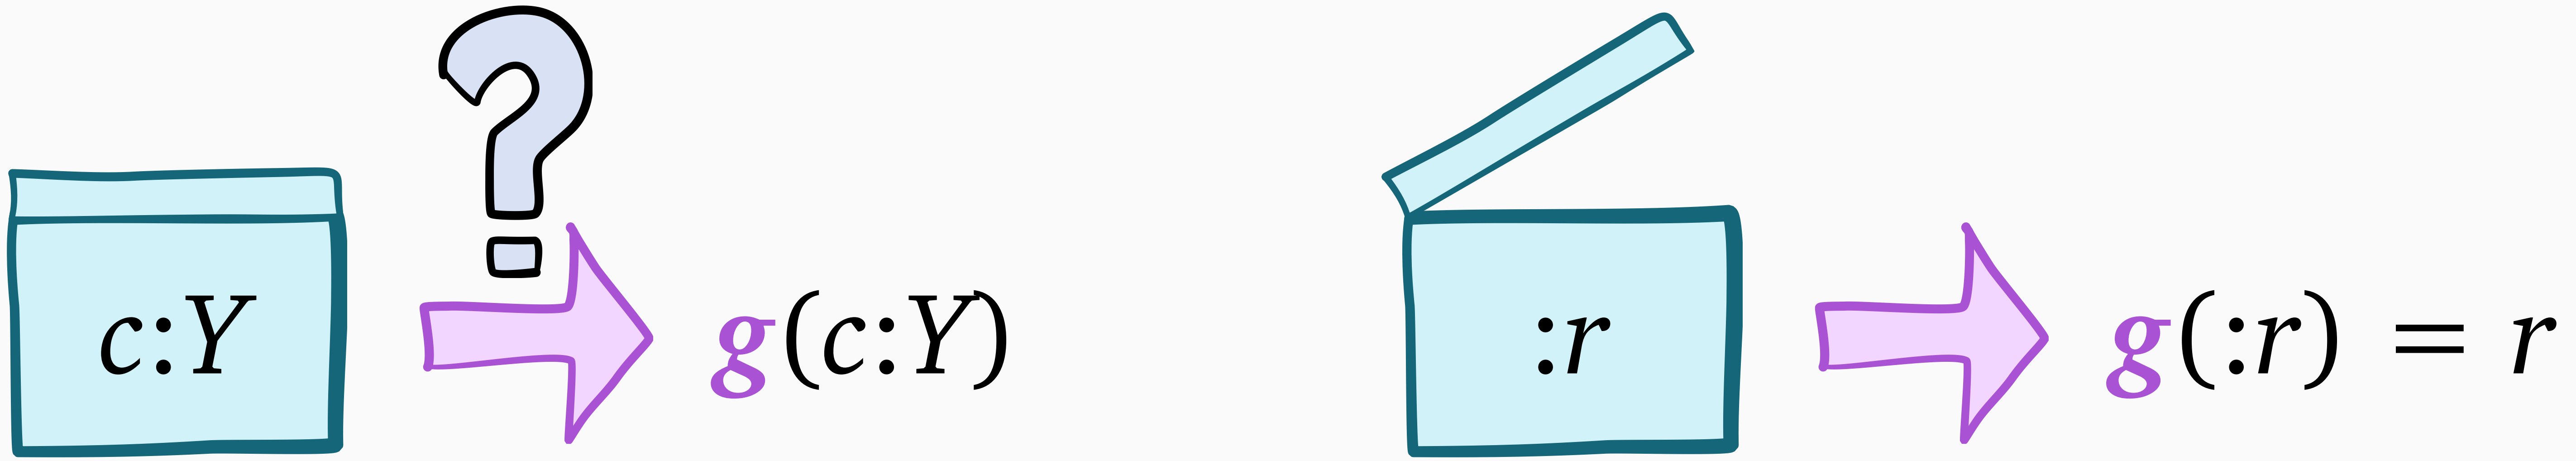

1.5-box problem

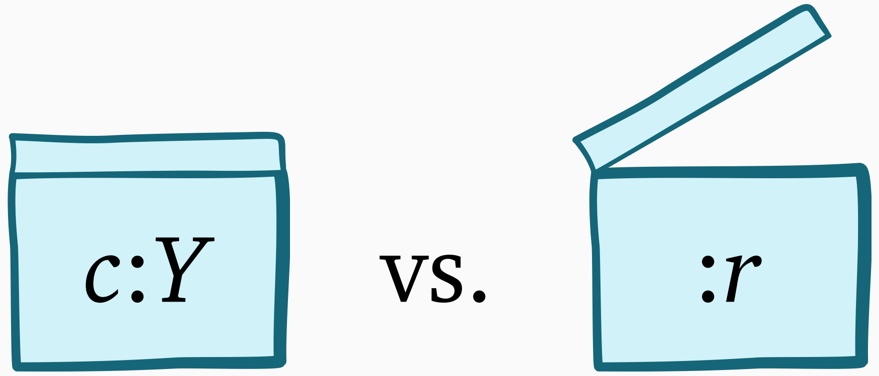

Key question: what to do in 1.5-box problem?

Defining the **Gittins** index

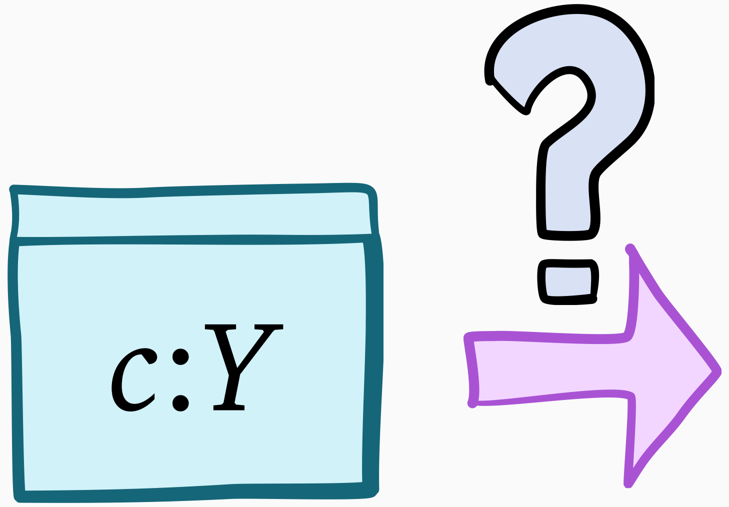

$g(c{:}Y)$

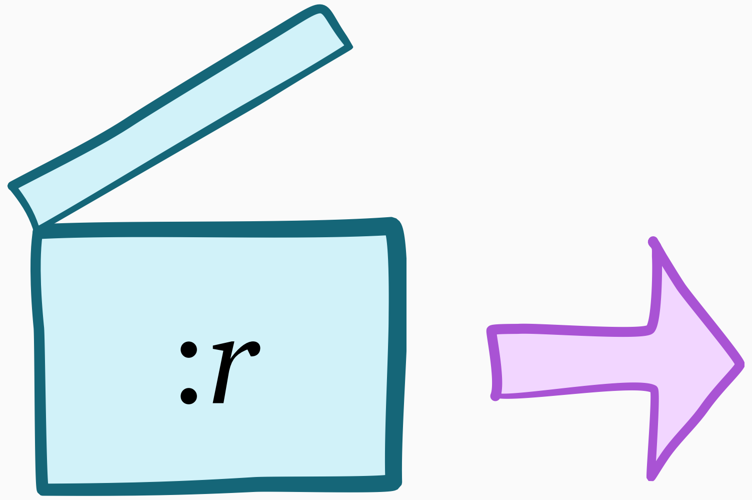

$g({:}r) = r$

1.5-box problem

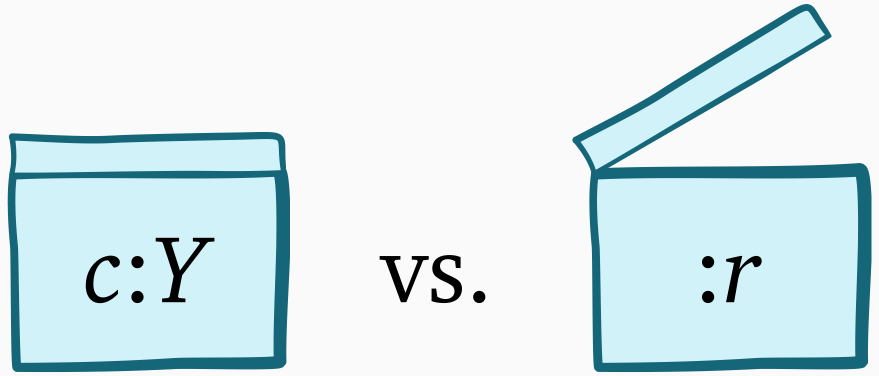

Key question: what to do in 1.5-box problem?

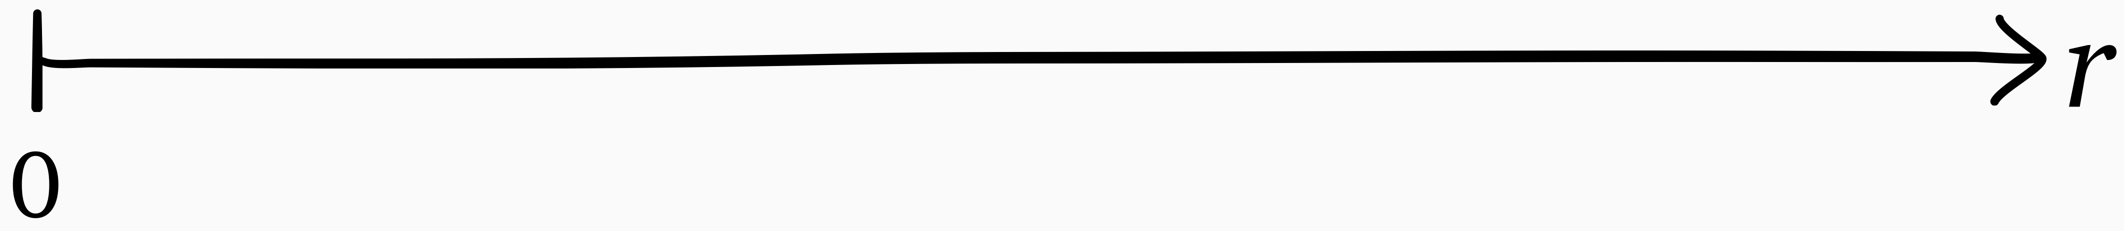

Defining the **Gittins** index

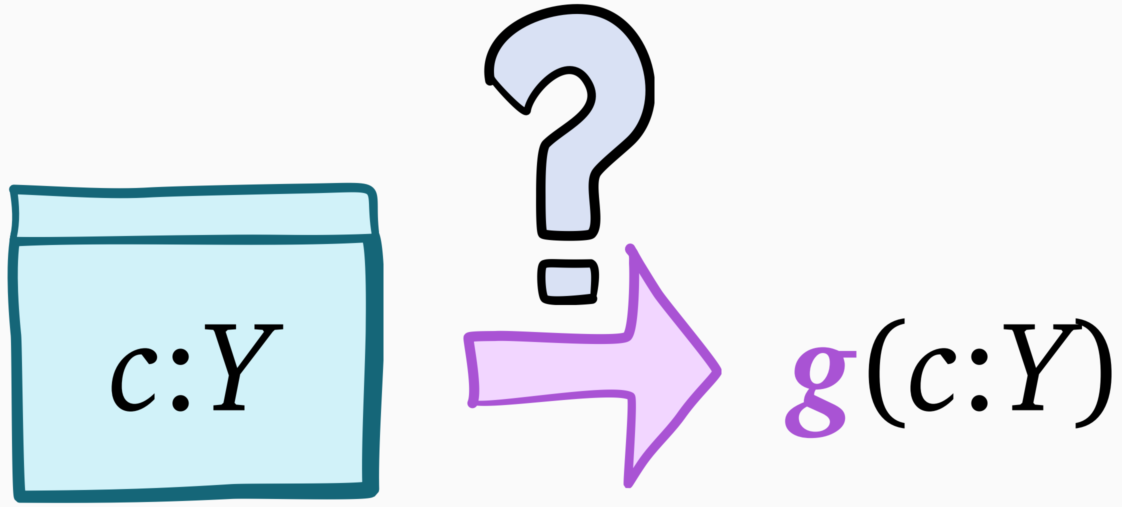

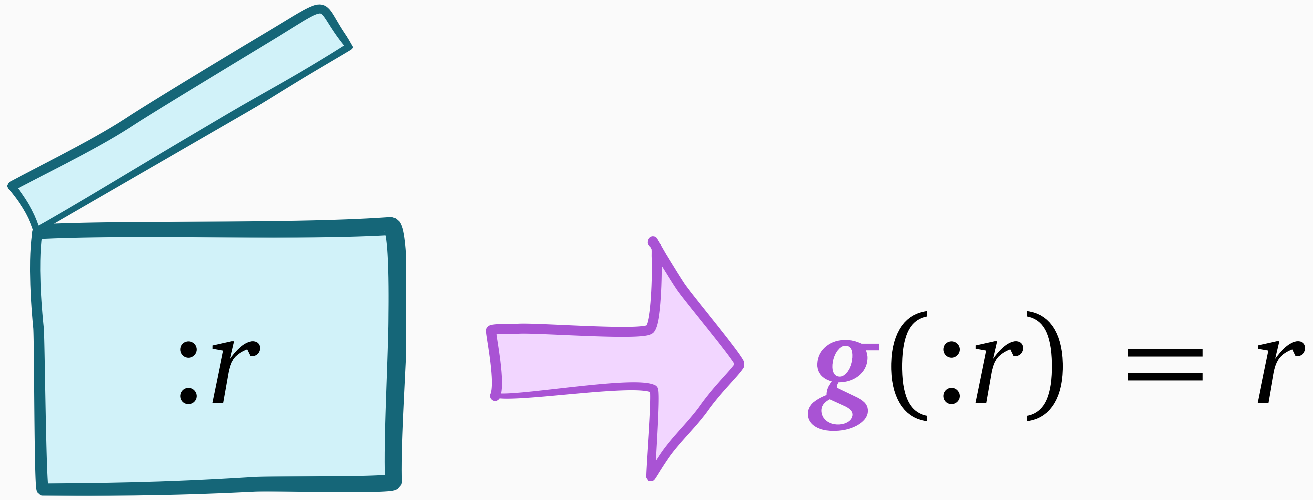

1.5-box problem

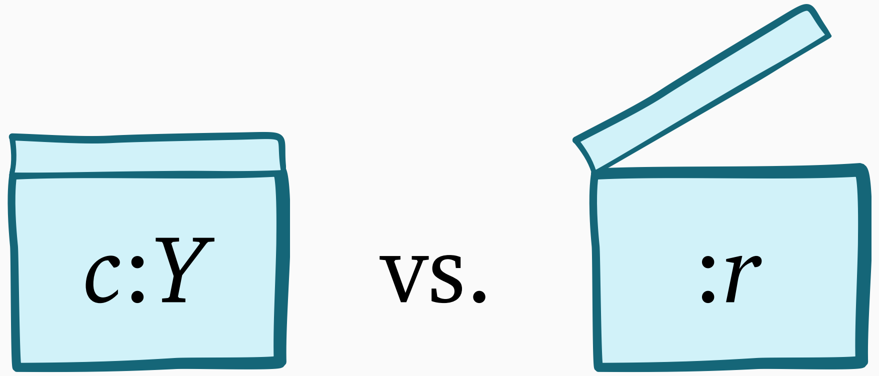

Key question: what to do in 1.5-box problem?

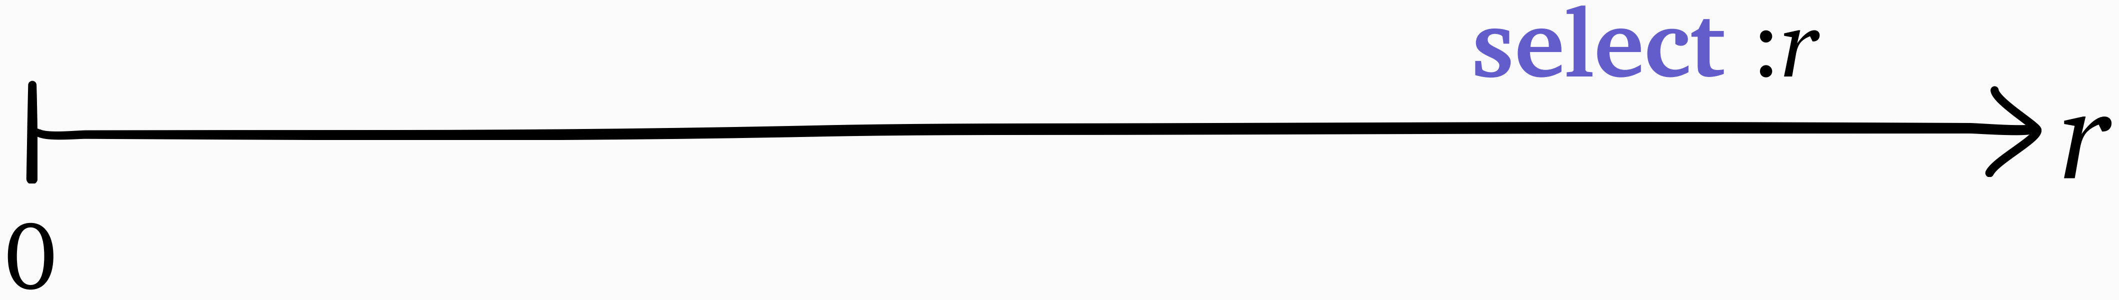

Defining the Gittins index

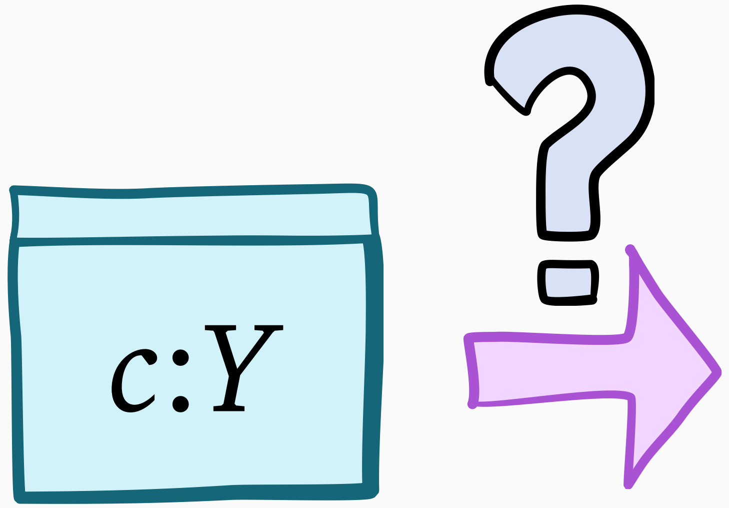

$g(c{:}Y)$

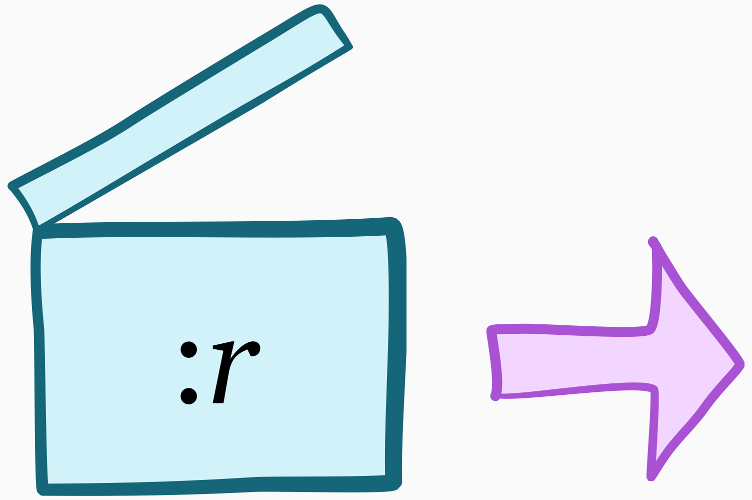

$g({:}r) = r$

1.5-box problem

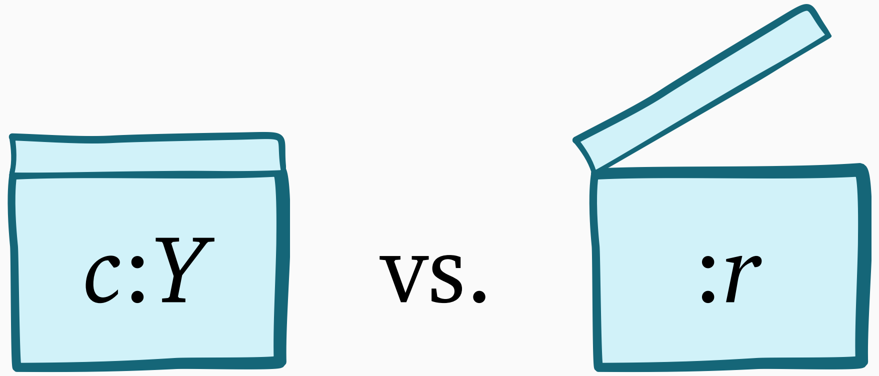

Key question: what to do in 1.5-box problem?

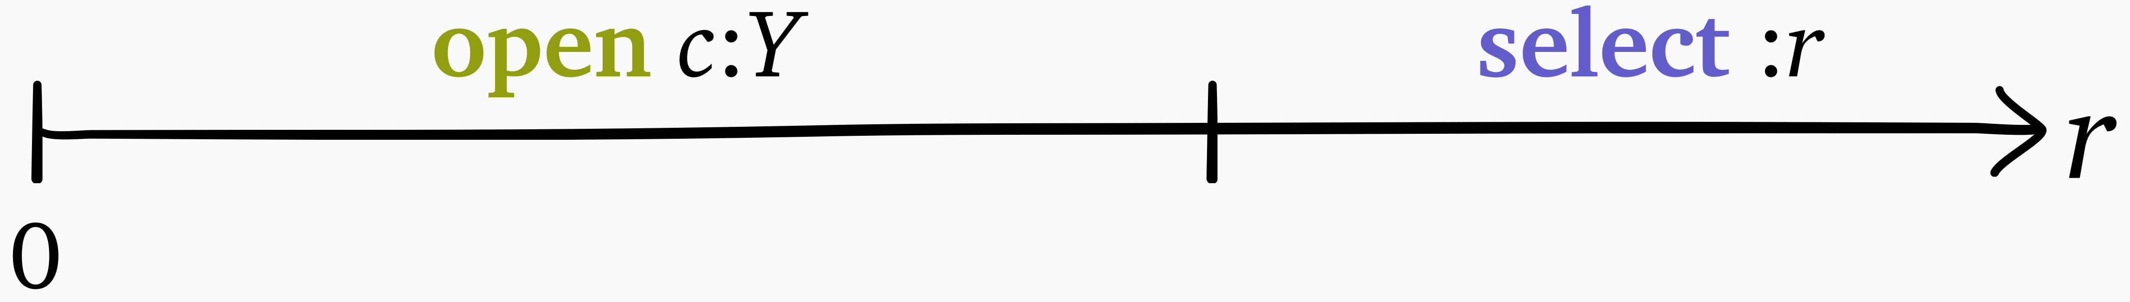

Defining the Gittins index

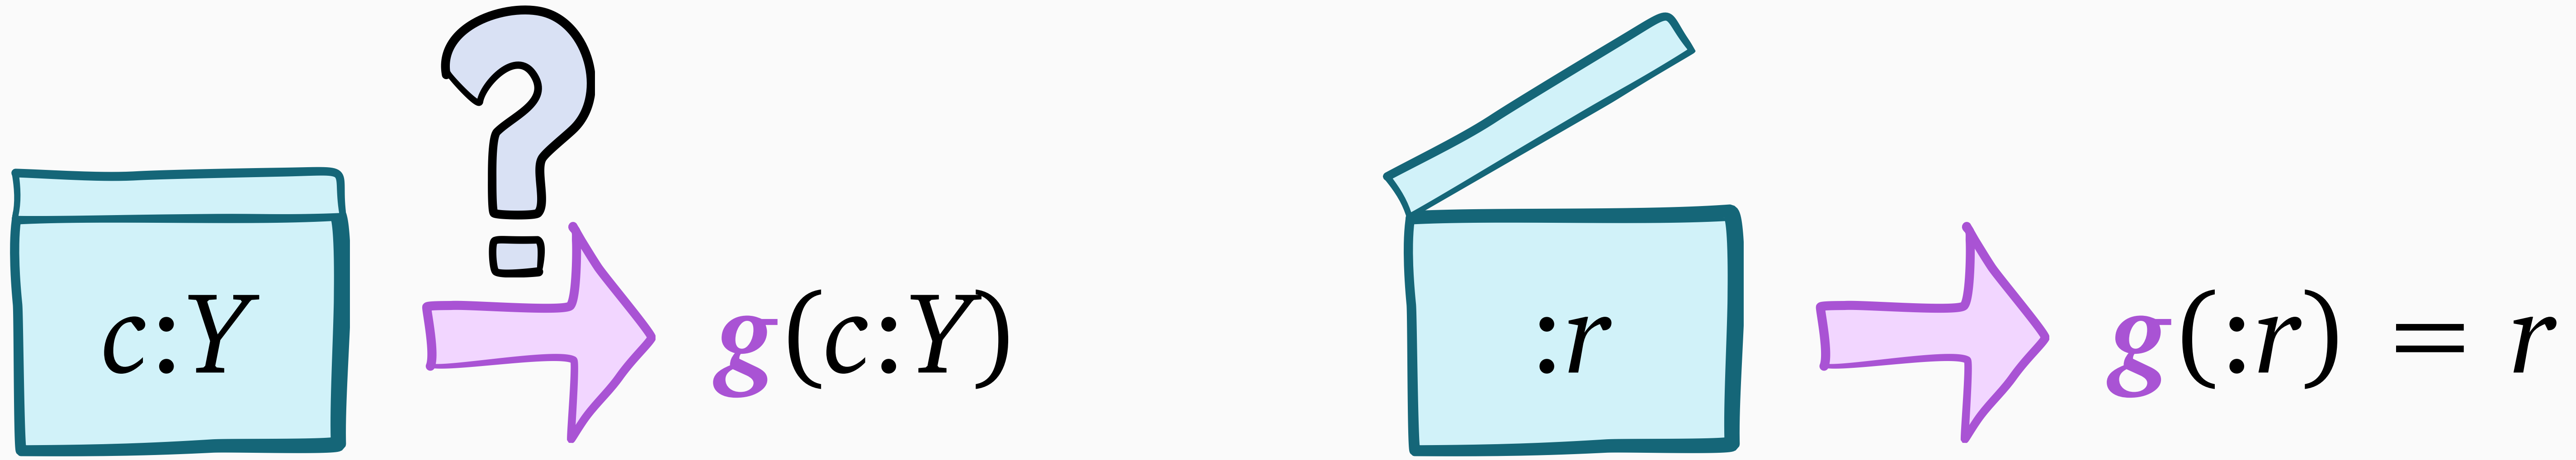

1.5-box problem

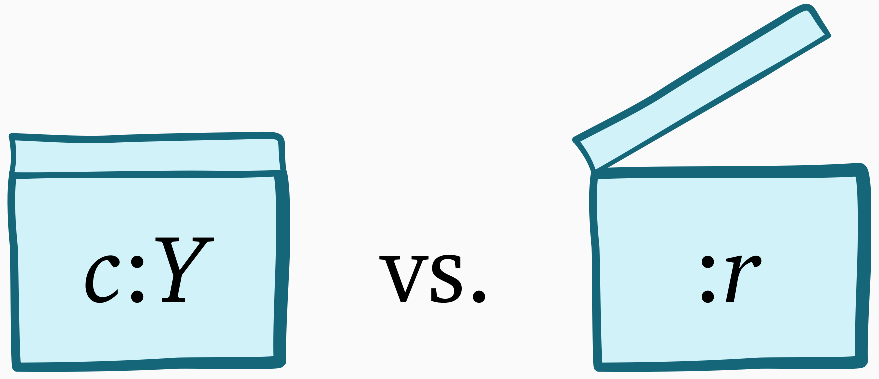

Key question: what to do in 1.5-box problem?

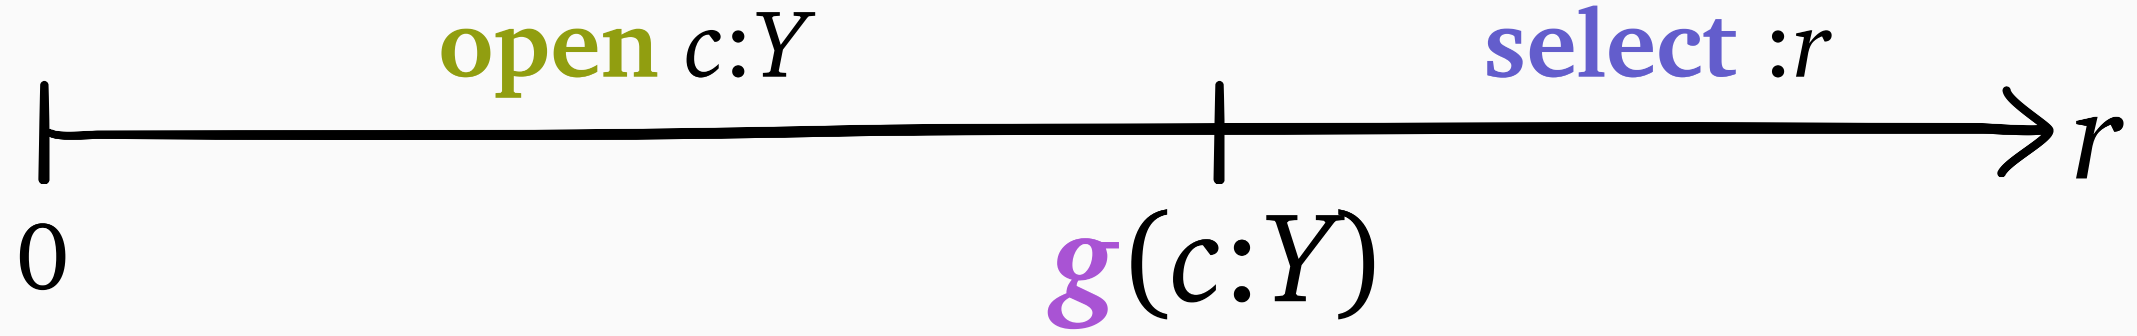

Defining the **Gittins** index

1.5-box problem

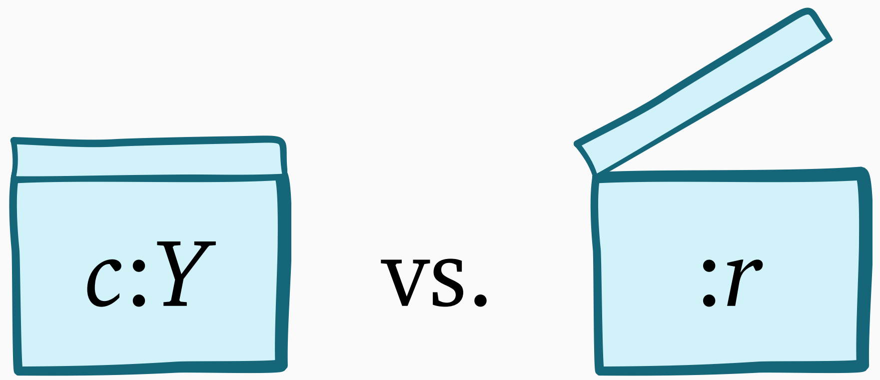

Key question: what to do in 1.5-box problem?

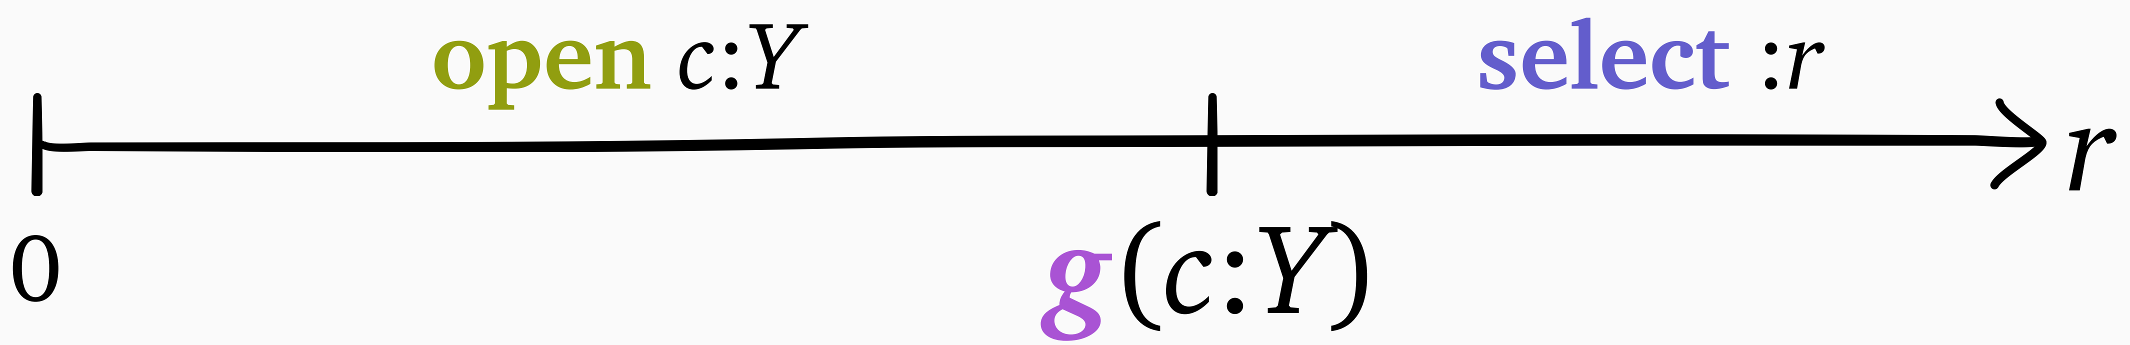

Defining the Gittins index

1.5-box problem

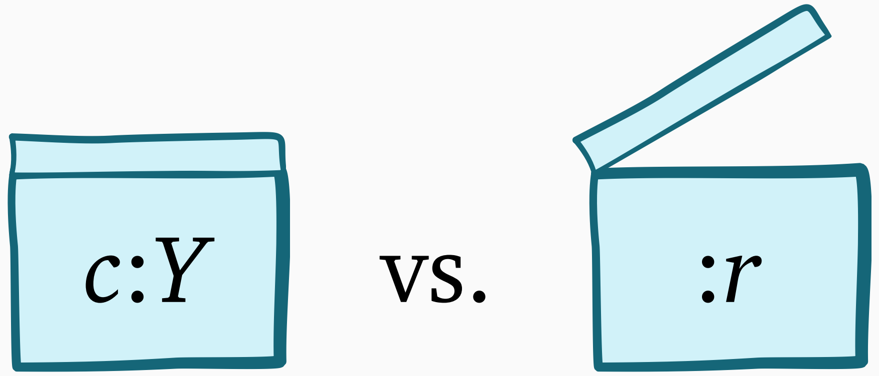

Key question: what to do in 1.5-box problem?

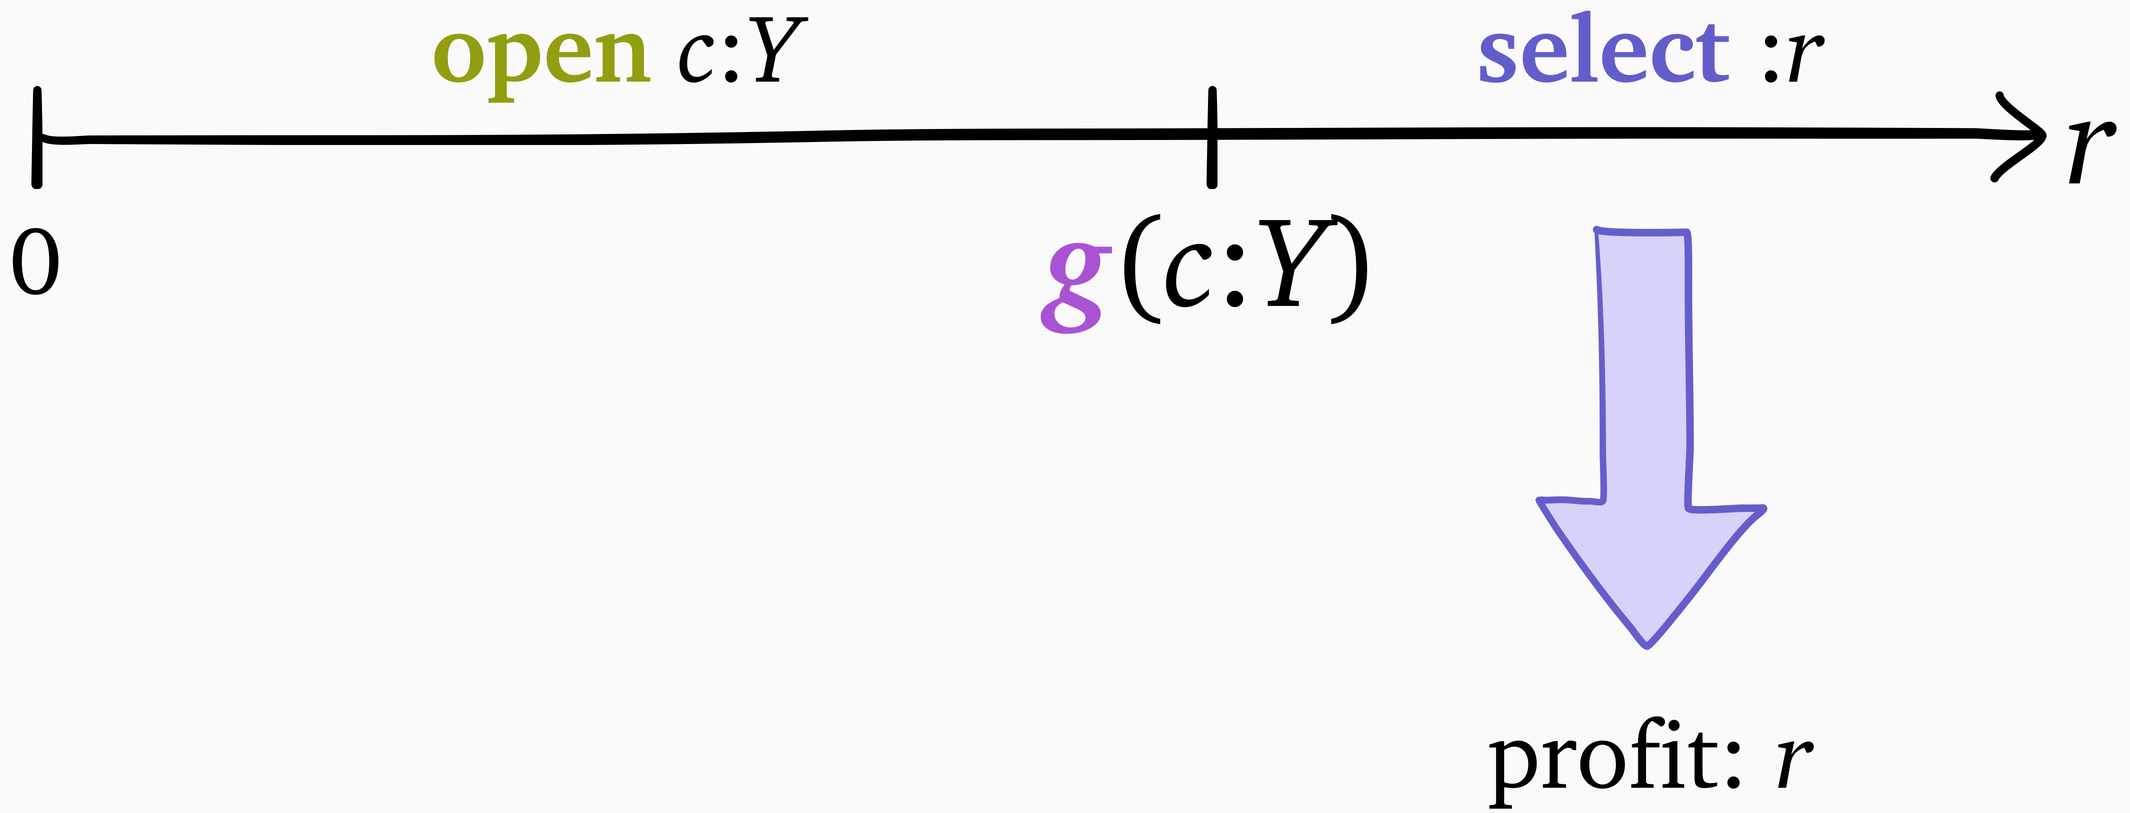

Defining the Gittins index

1.5-box problem

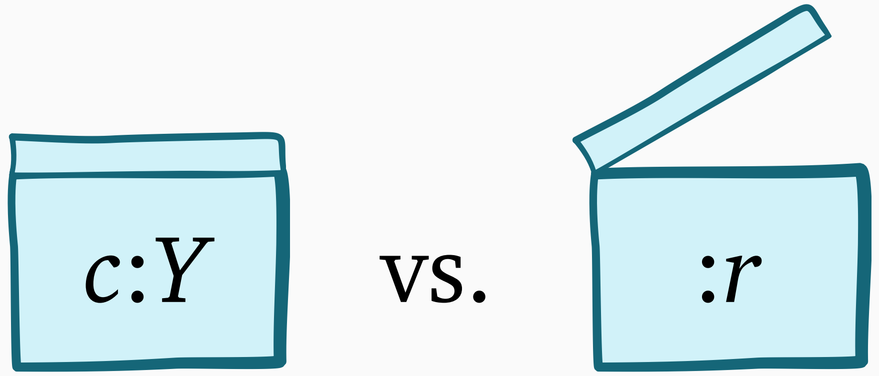

Key question: what to do in 1.5-box problem?

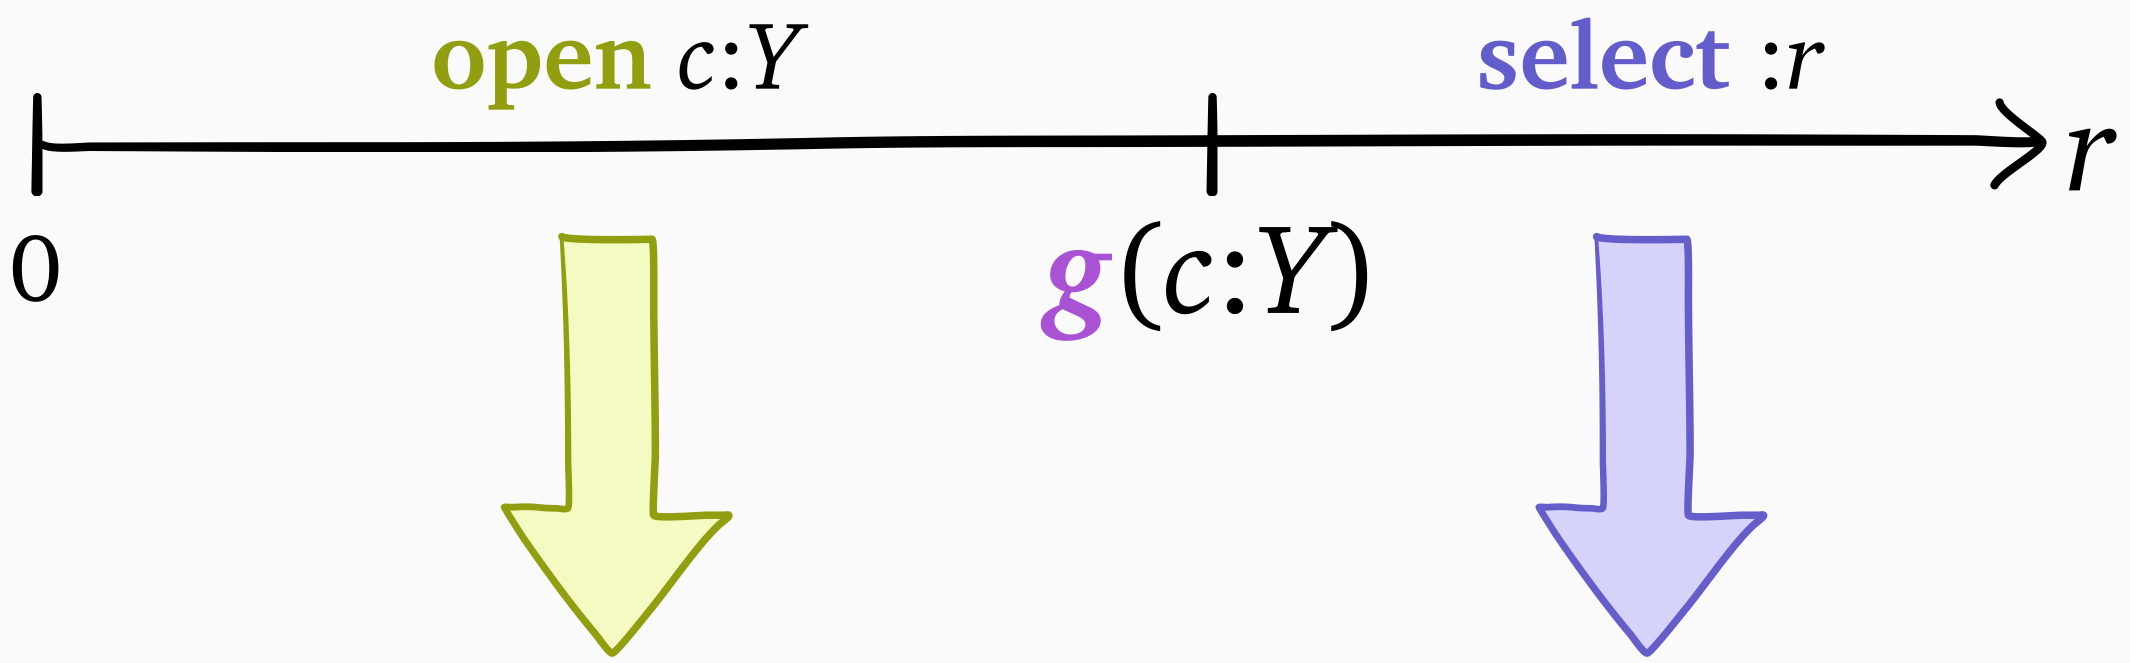

profit: $\mathbf{E}[\max\{Y, r\}] - c$

profit: r

Defining the Gittins index

1.5-box problem

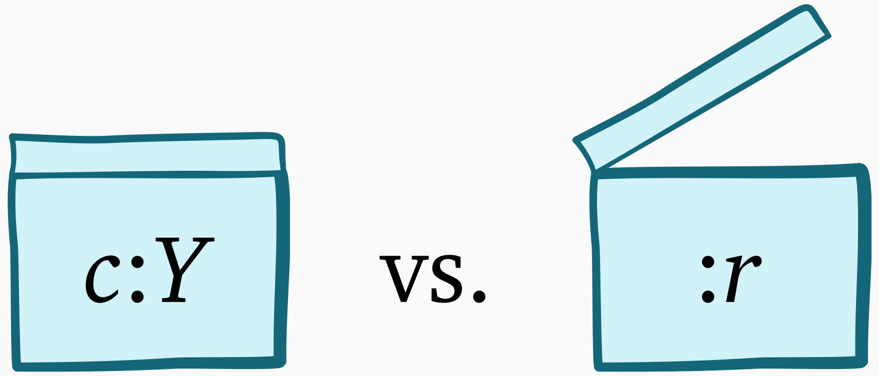

Key question: what to do in 1.5-box problem?

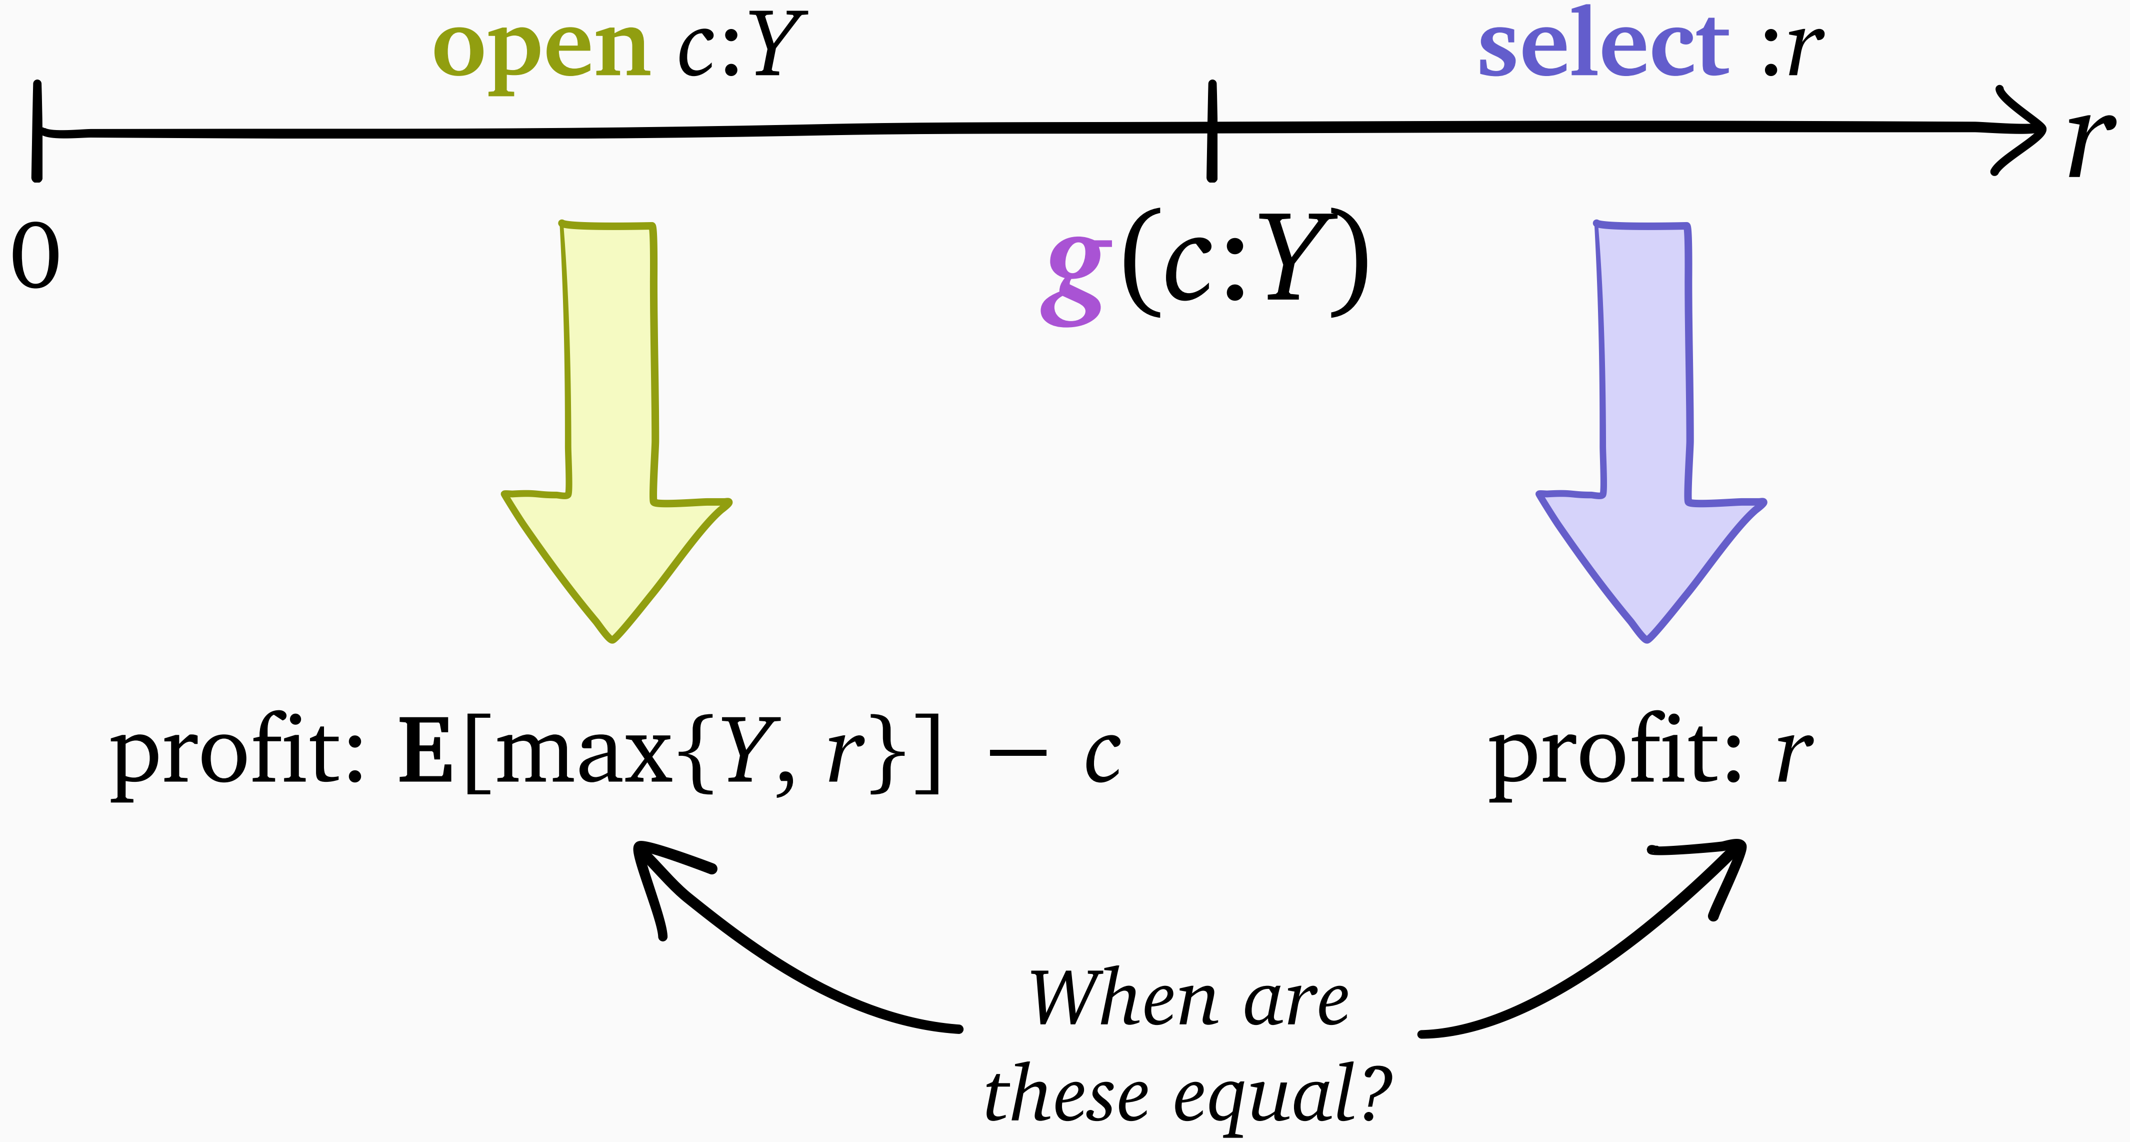

Defining the **Gittins** index

1.5-box problem

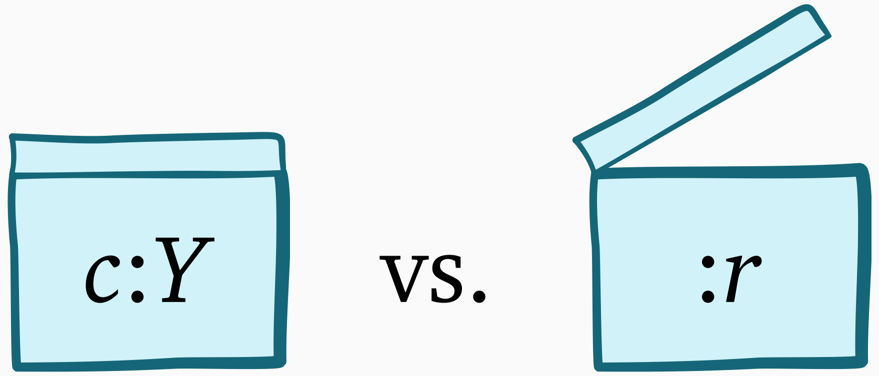

Defn: $g(c{:}Y)$ is solution r to

$$\mathrm{EI}(Y, r) = \mathbf{E}[(Y - r)^+] = c$$

Key question: what to do in 1.5-box problem?

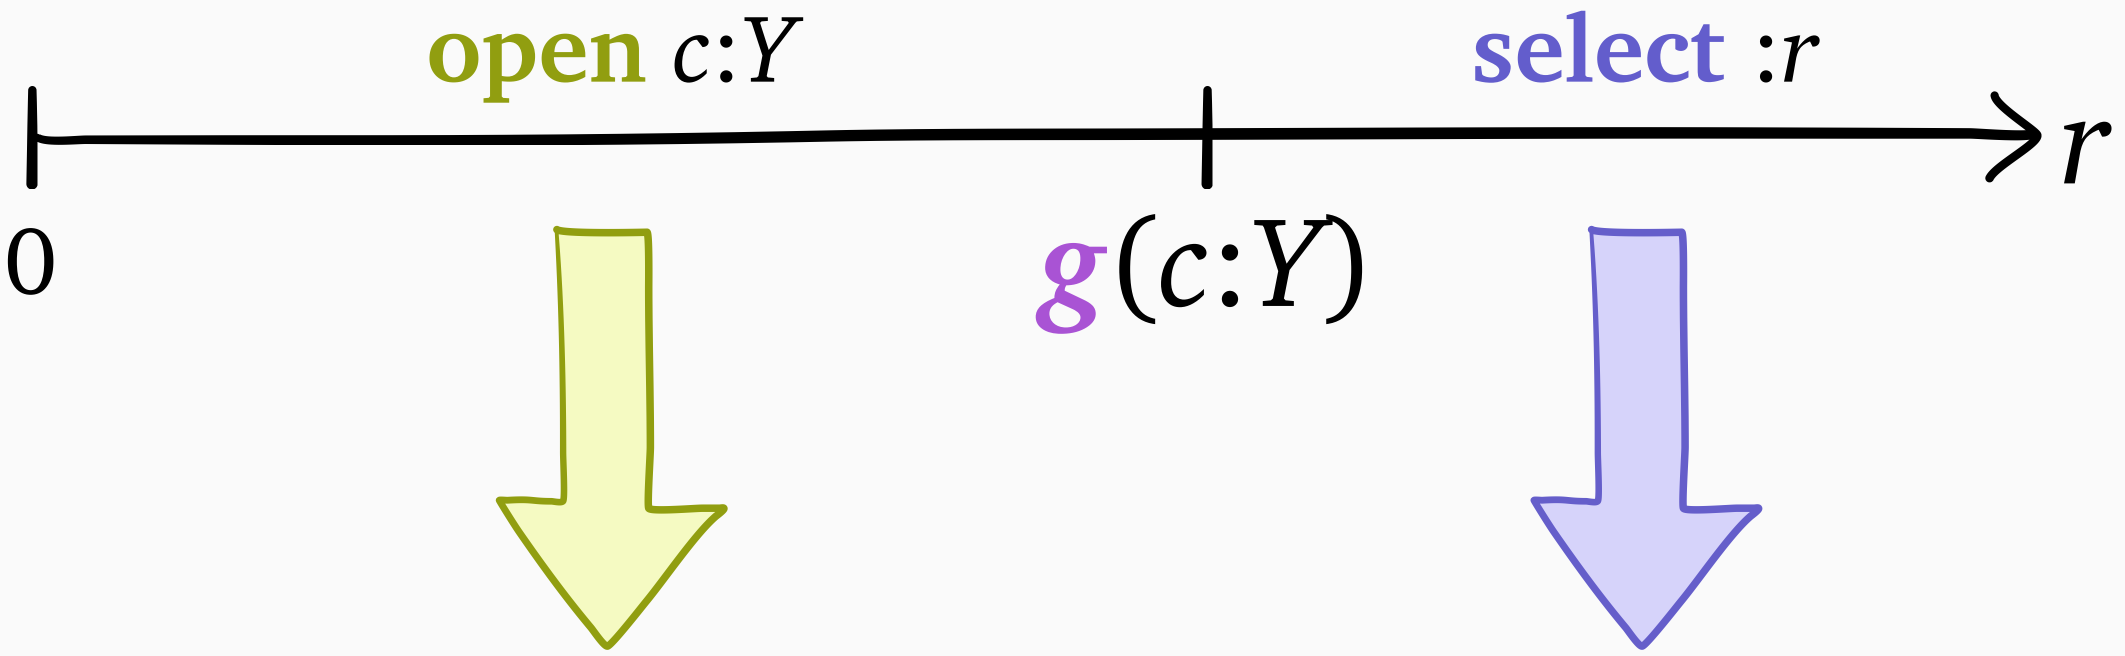

profit: $\mathbf{E}[\max\{Y, r\}] - c$

profit: r

When are these equal?

Difference between EI and Gittins

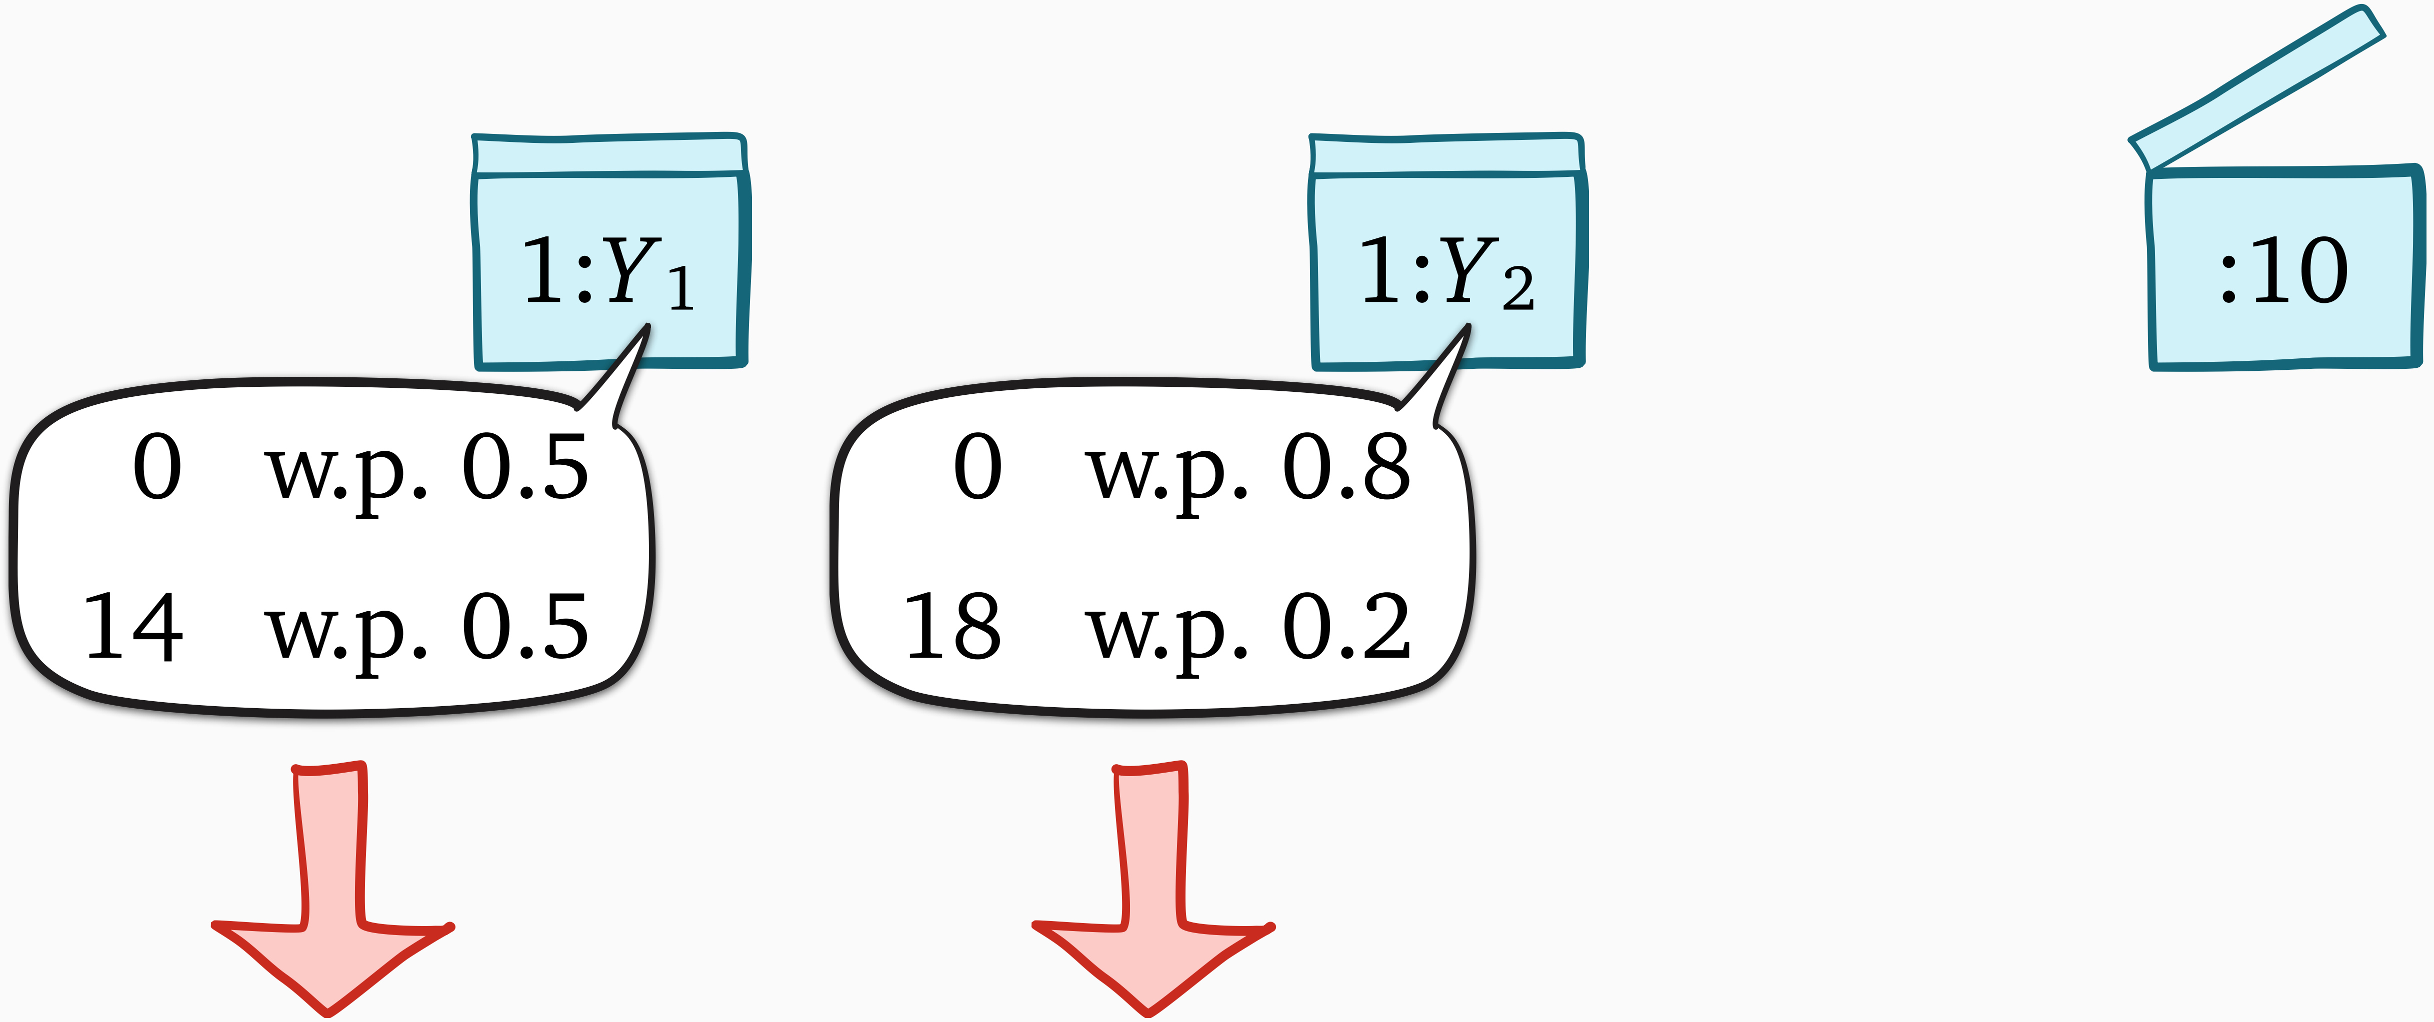

$\mathbf{EI}(Y_1, 10) = 2$ $\mathbf{EI}(Y_2, 10) = 1.6$

Difference between **EI** and **Gittins**

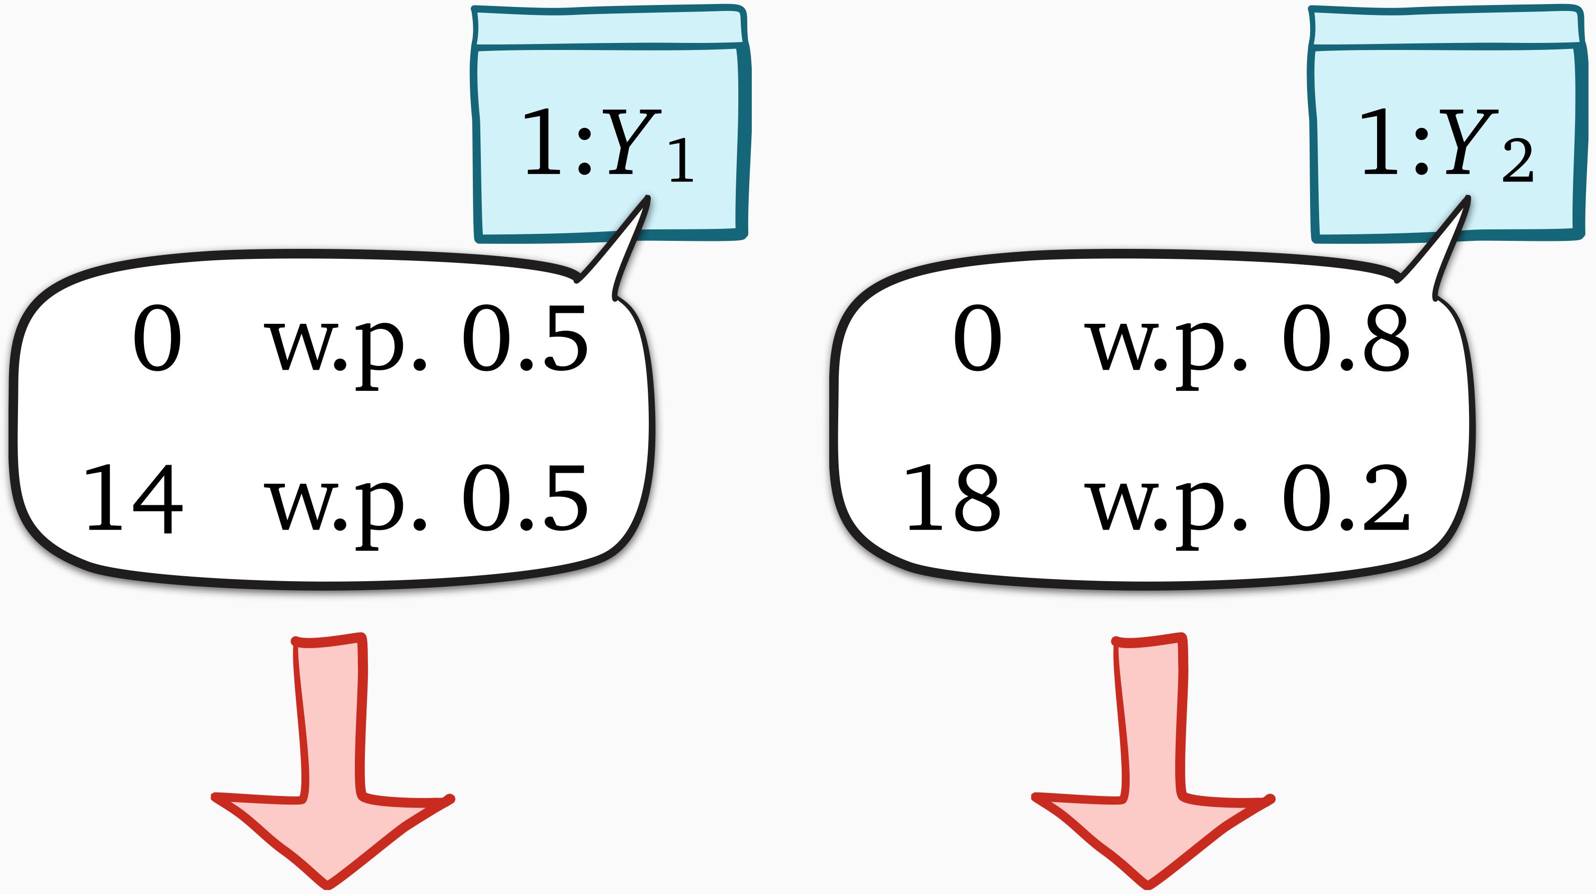

EI$(Y_1, 10) = 2$ **EI**$(Y_2, 10) = 1.6$

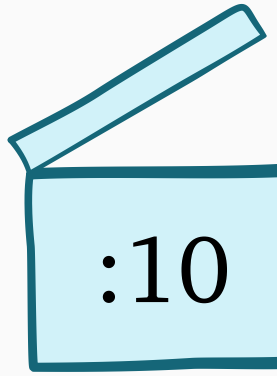

With $r = 10$, what is **EI**$(Y, r) - c$?

Difference between **EI** and **Gittins**

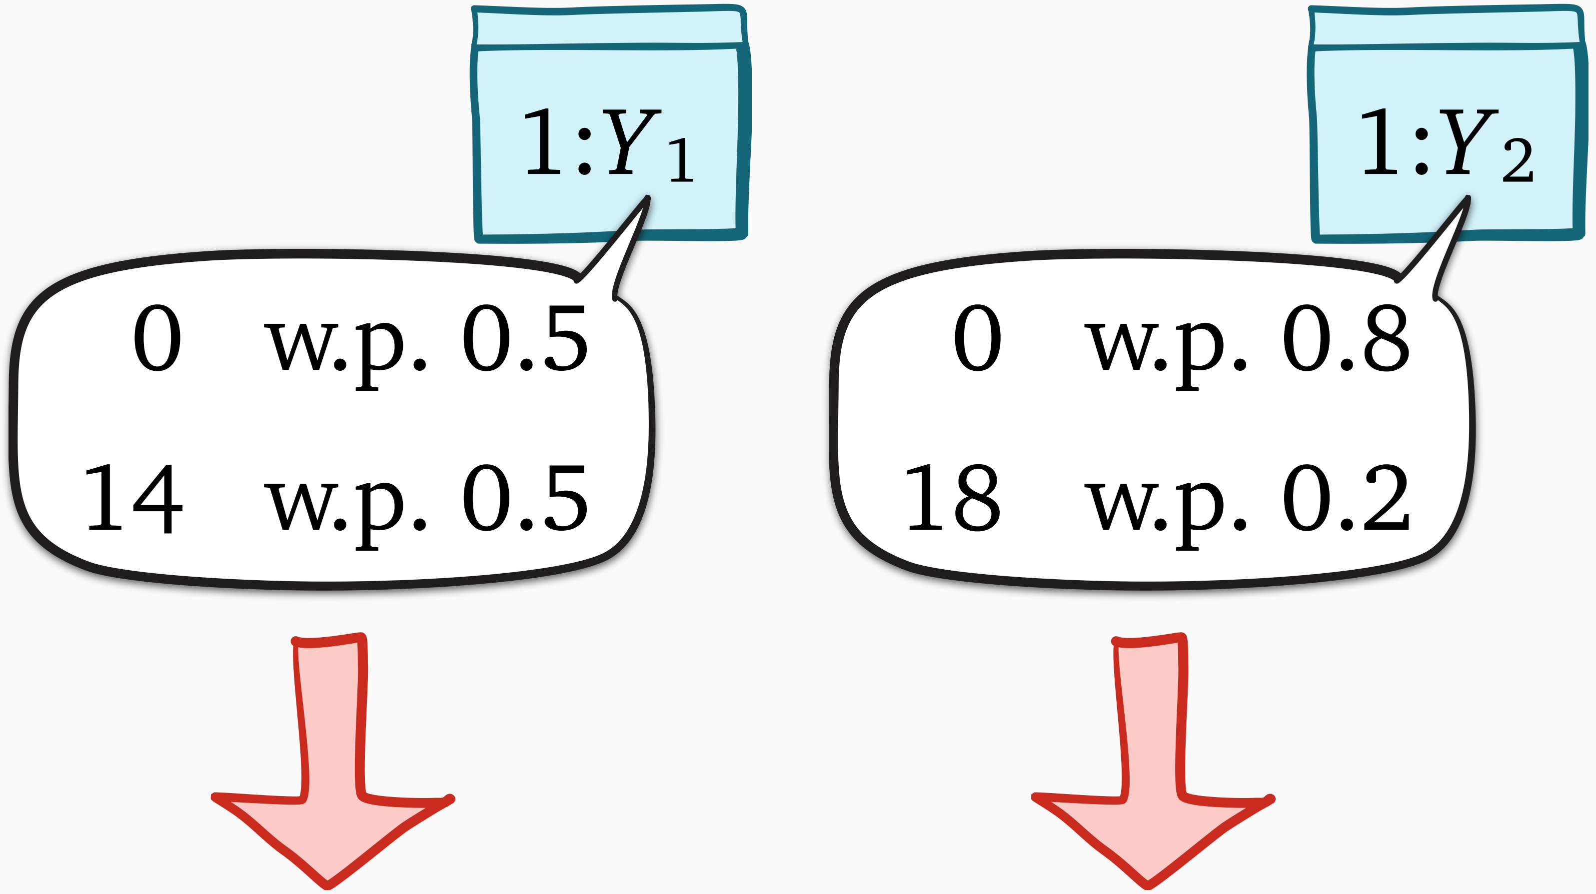

$\mathbf{EI}(Y_1, 10) = 2$ $\mathbf{EI}(Y_2, 10) = 1.6$

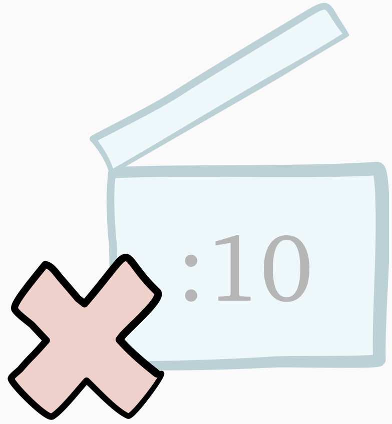

With $r = 10$, what is $\mathbf{EI}(Y, r) - c$?

Difference between **EI** and **Gittins**

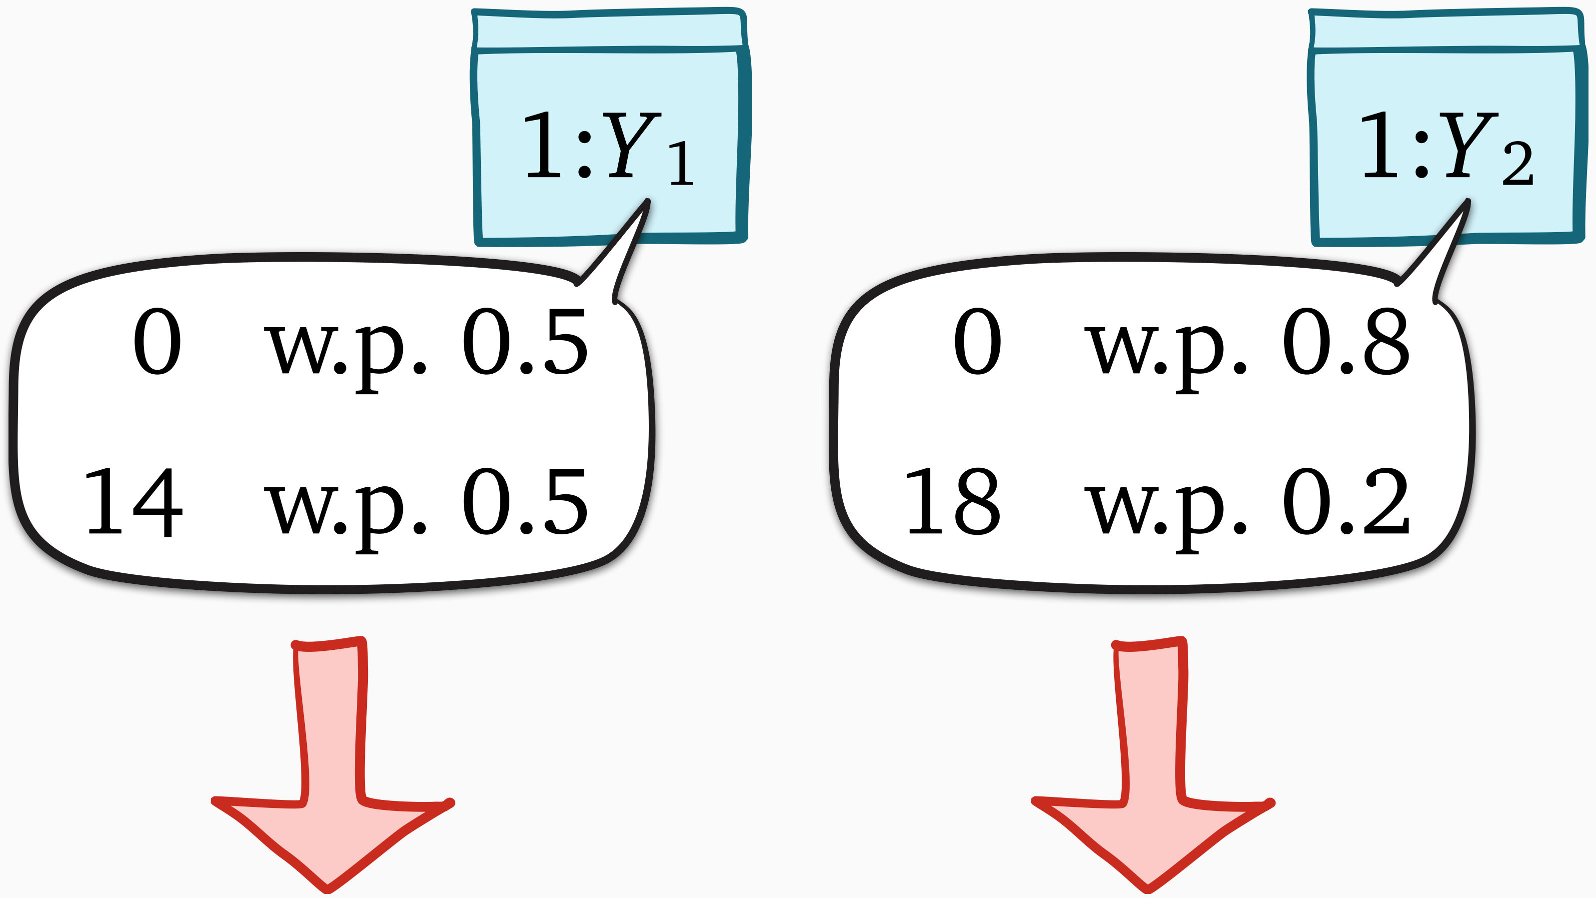

EI$(Y_1, 10) = 2$ **EI**$(Y_2, 10) = 1.6$

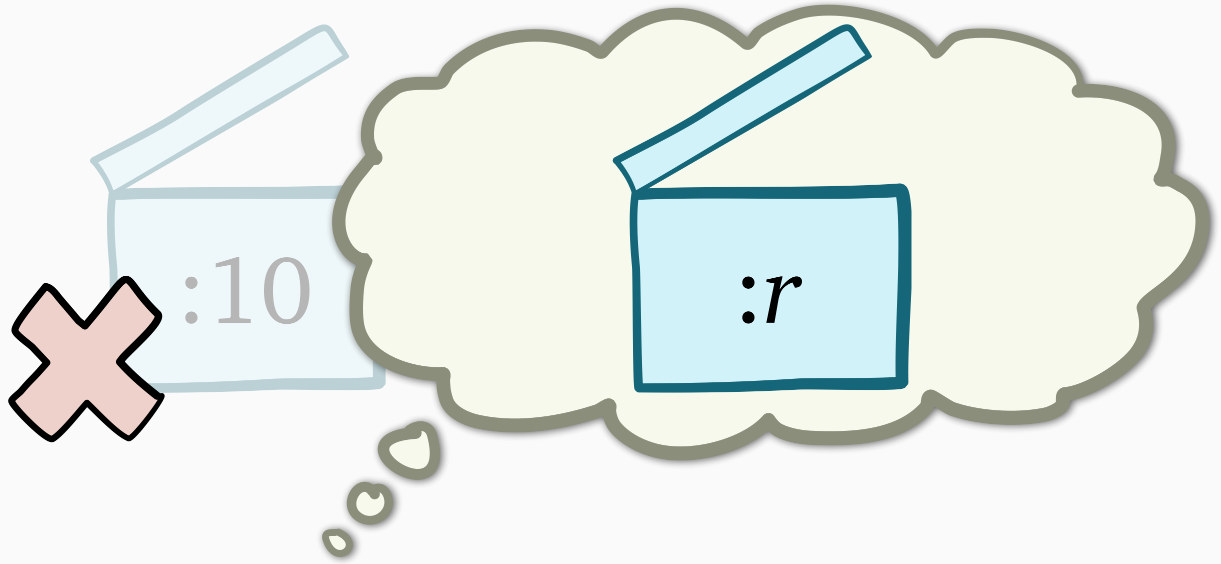

With $r = 10$, what is **EI**$(Y, r) - c$?

Difference between EI and Gittins

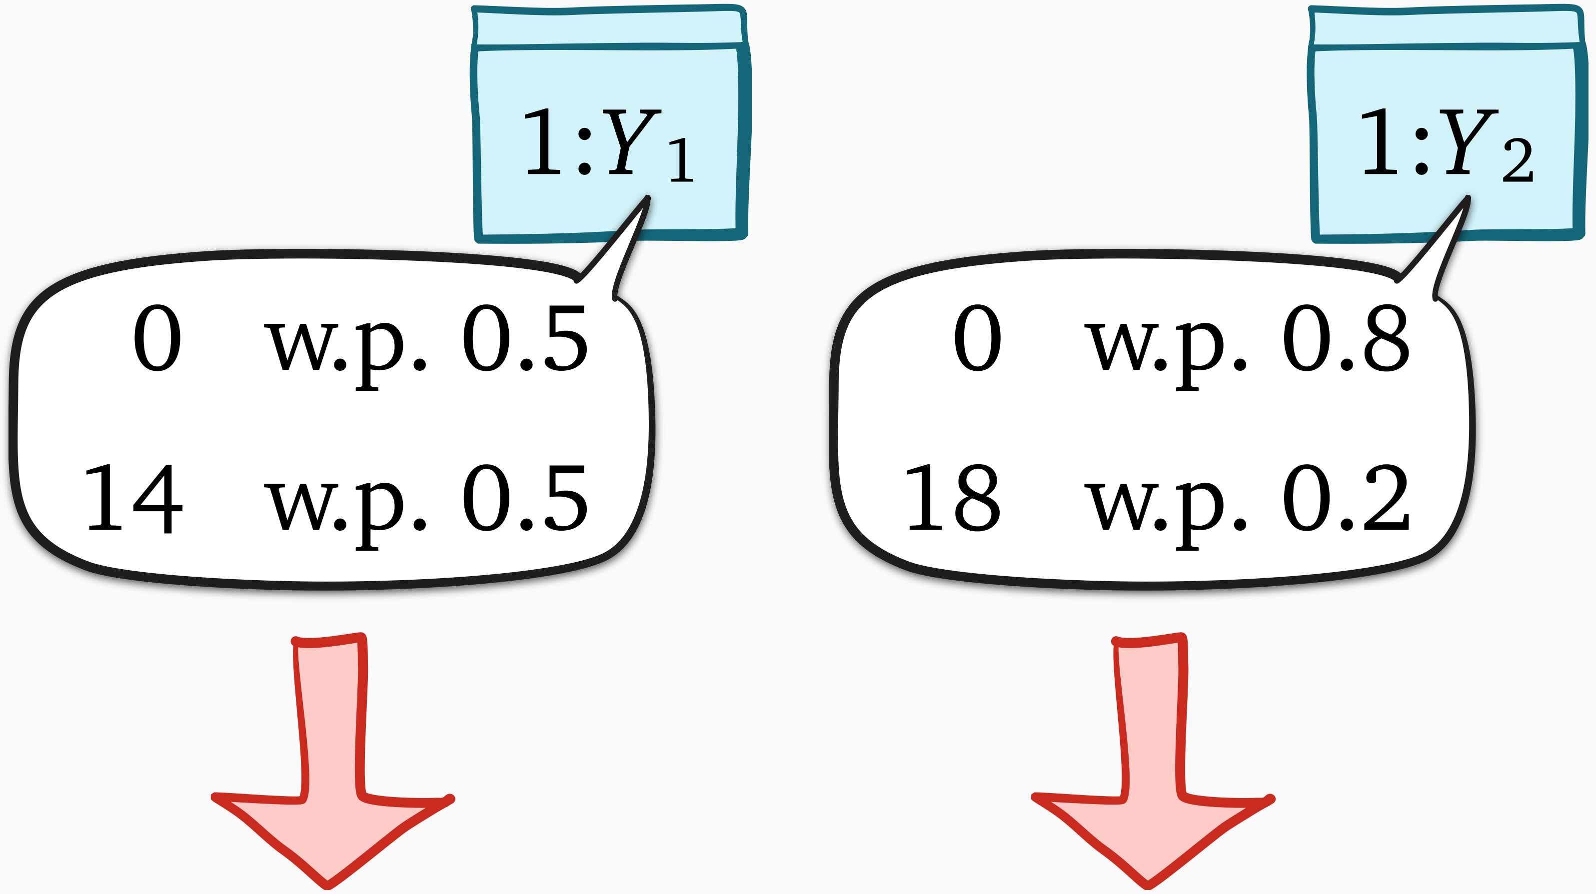

$\mathbf{EI}(Y_1, 10) = 2$ $\mathbf{EI}(Y_2, 10) = 1.6$

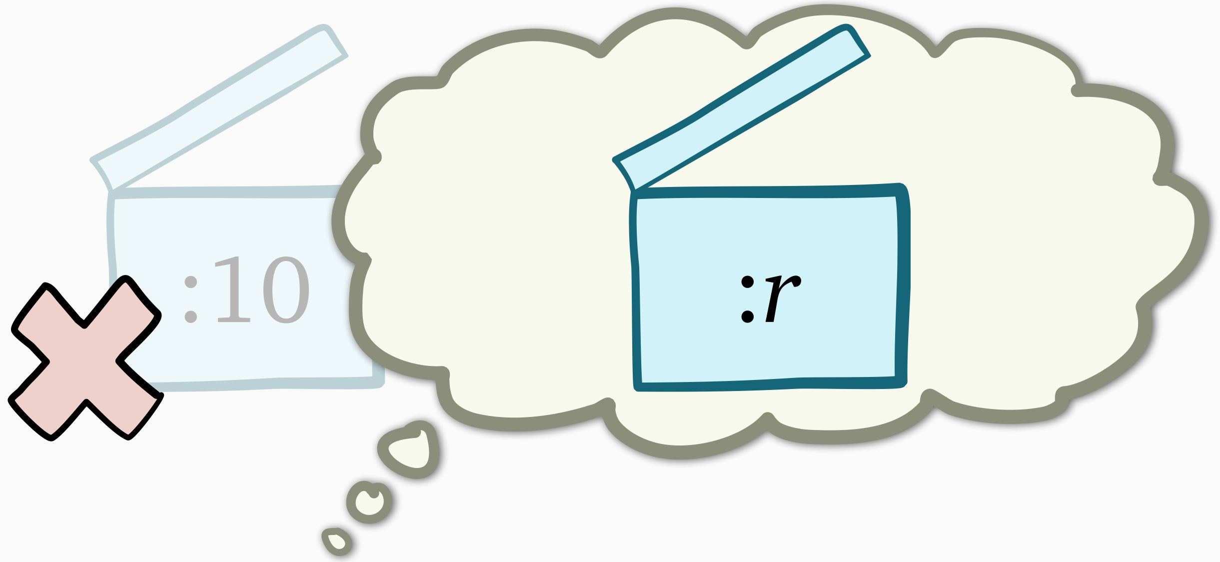

✗ With $r = 10$, what is $\mathbf{EI}(Y, r) - c$?

✓ For what r does $\mathbf{EI}(Y, r) - c = 0$?

Difference between EI and Gittins

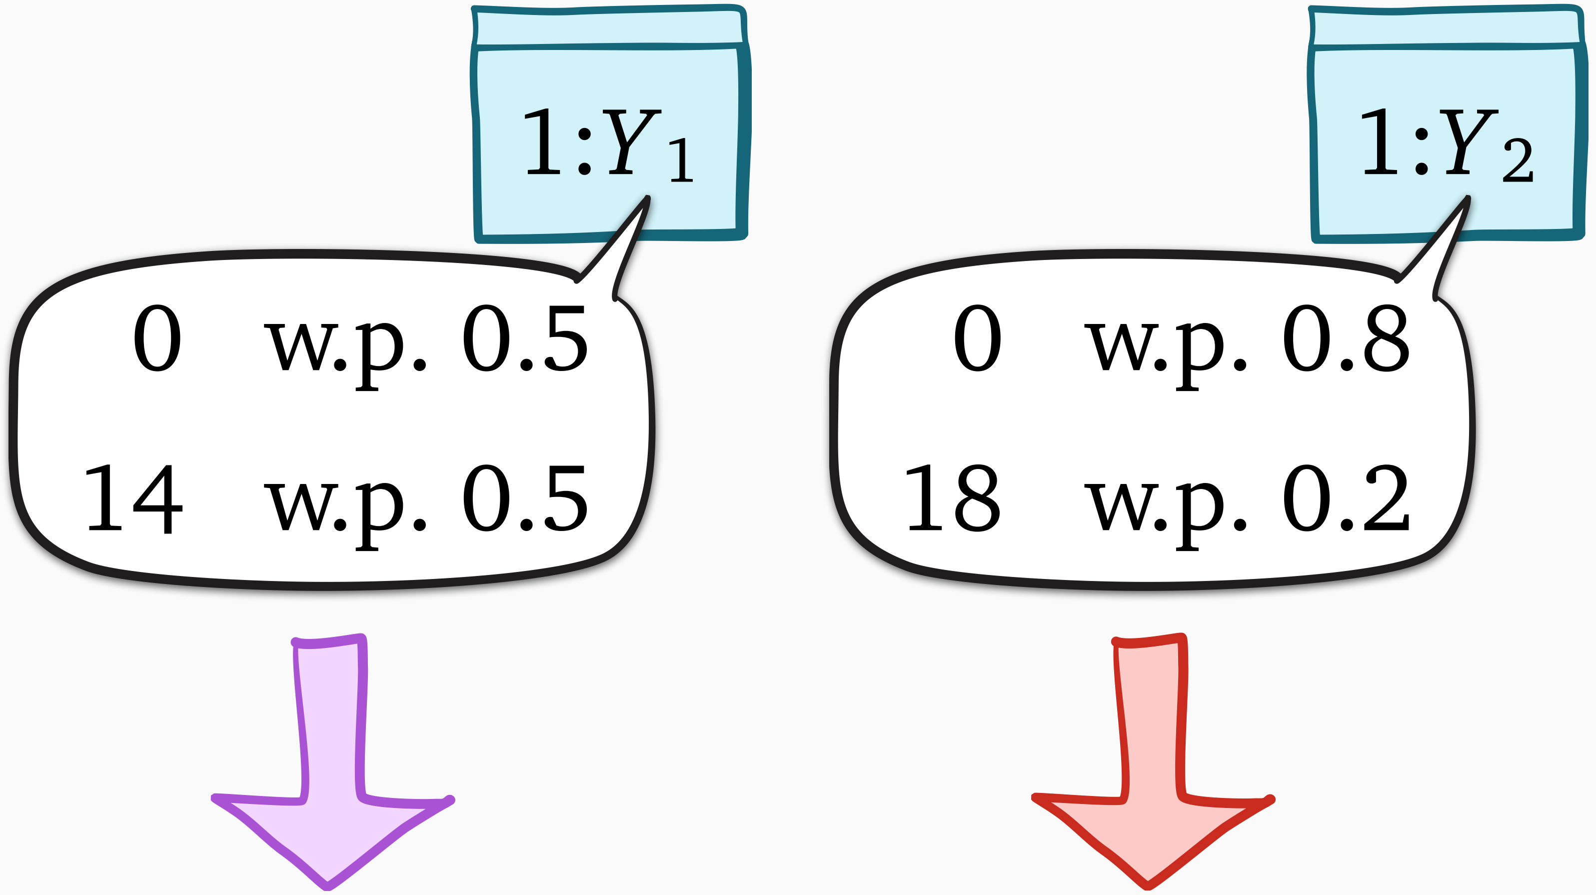

$\mathbf{EI}(Y_1, 10) = 2$

$g(1{:}Y_1) = 12$

$\mathbf{EI}(Y_2, 10) = 1.6$

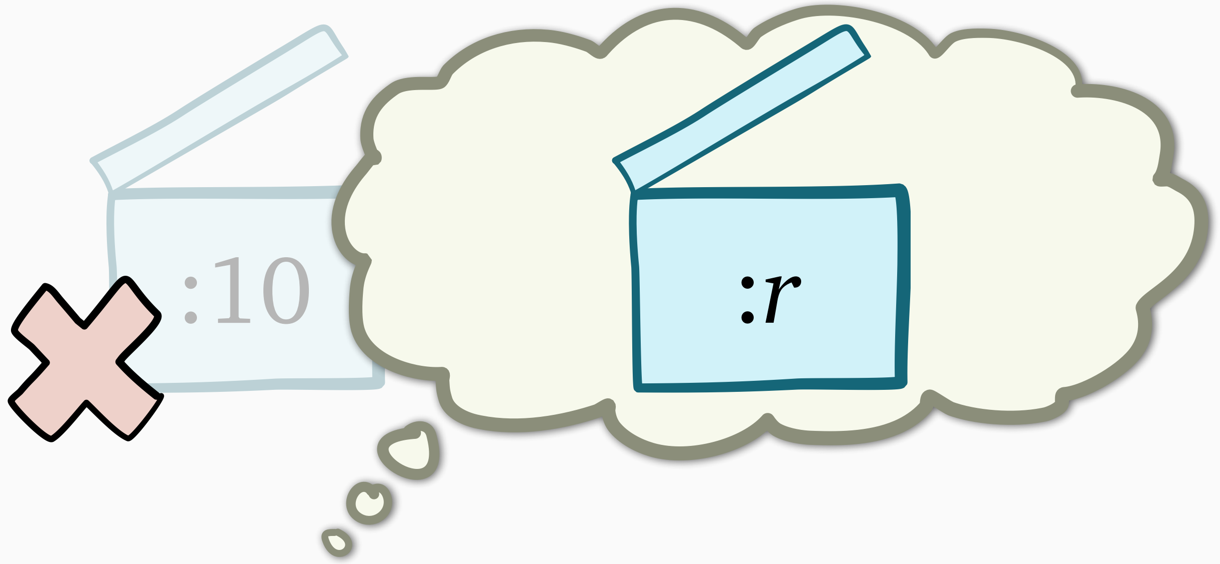

✗ With $r = 10$, what is $\mathbf{EI}(Y, r) - c$?

✓ For what r does $\mathbf{EI}(Y, r) - c = 0$?

Difference between EI and Gittins

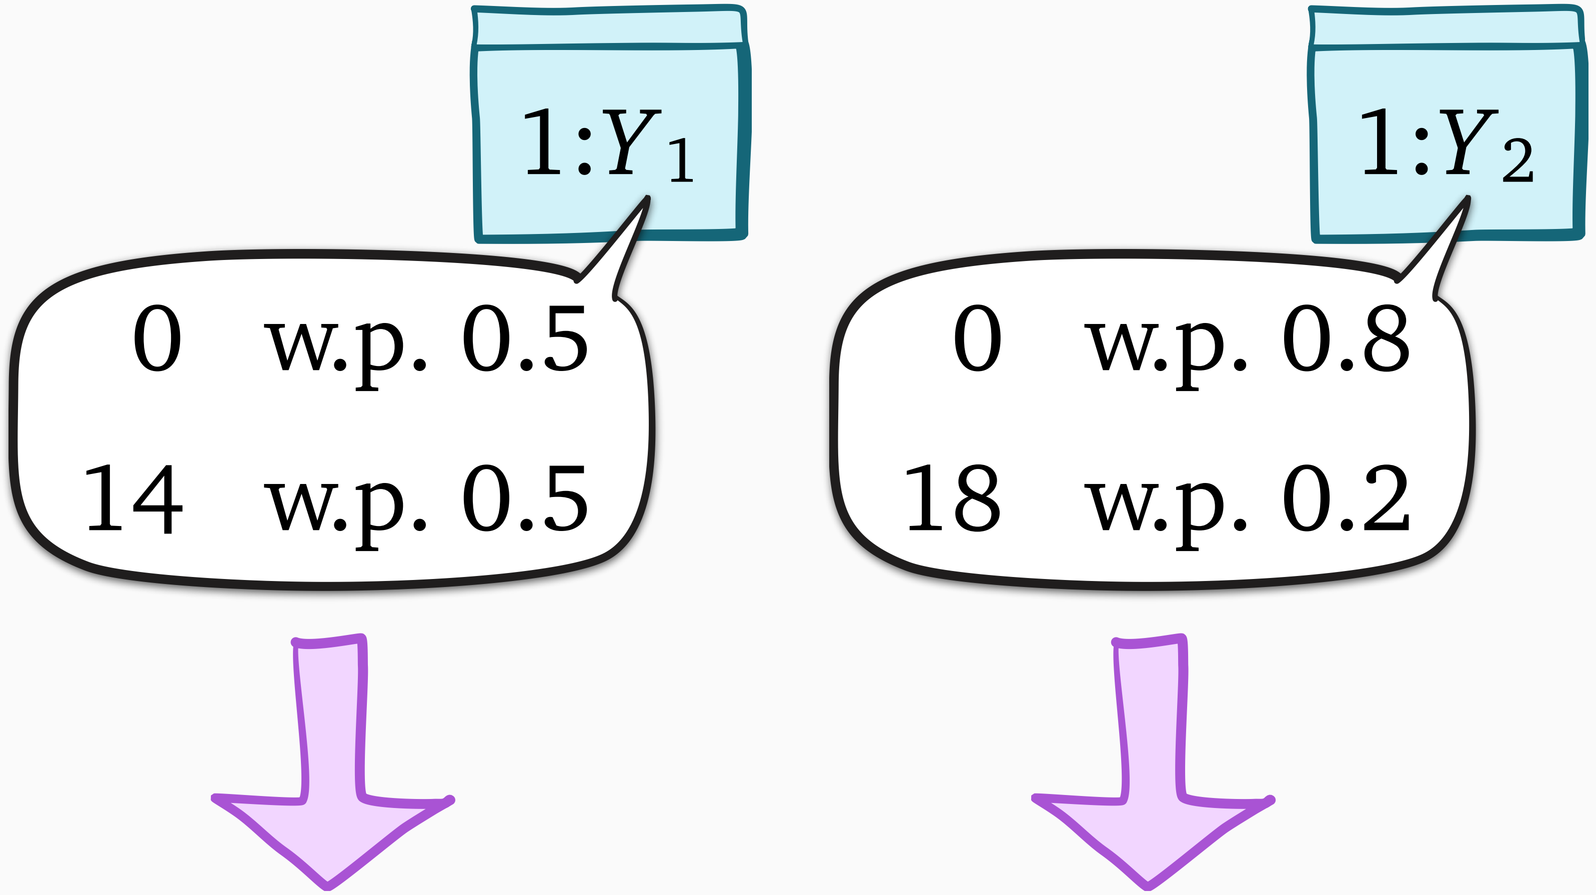

$\mathbf{EI}(Y_1, 10) = 2$
$g(1{:}Y_1) = 12$

$\mathbf{EI}(Y_2, 10) = 1.6$
$g(1{:}Y_2) = 13$

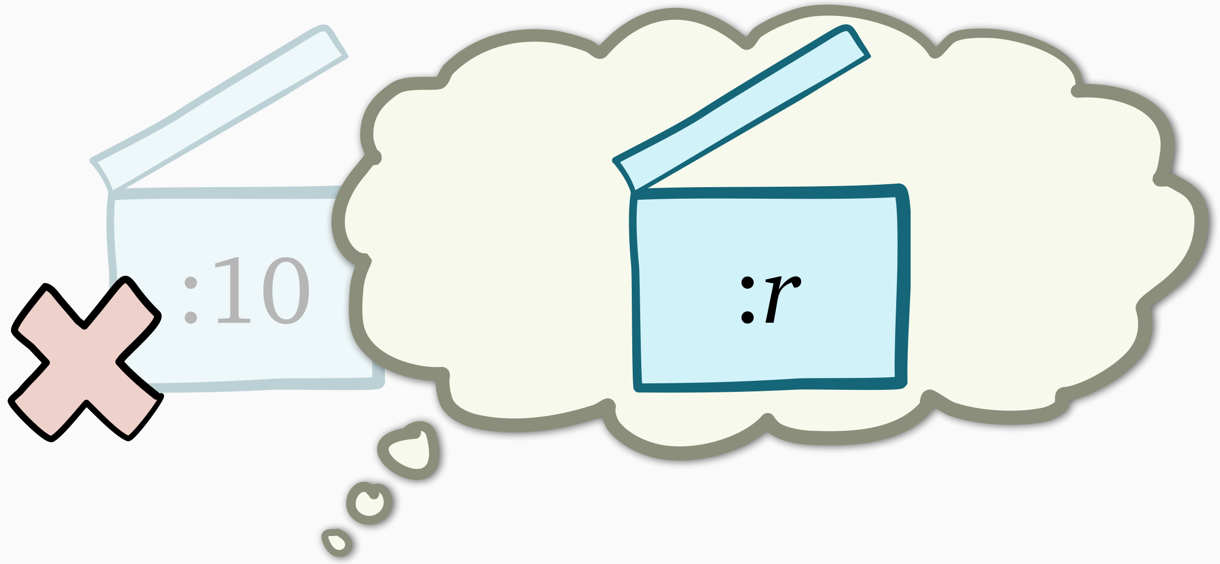

With $r = 10$, what is $\mathbf{EI}(Y, r) - c$?

For what r does $\mathbf{EI}(Y, r) - c = 0$?

Tutorial: What are Pandora's box and Gittins?

Results: Does Gittins work for BayesOpt?

Hype: How could Gittins help practical BayesOpt?

Tutorial: What are Pandora's box and Gittins?

Results: Does Gittins work for BayesOpt?

Hype: How could Gittins help practical BayesOpt?

Tutorial: What are Pandora’s box and Gittins?

Results: Does Gittins work for BayesOpt?

Hype: How could Gittins help practical BayesOpt?

Cost-per-sample **BayesOpt**

Assumption:
$f \sim \mathrm{GP}$

Unknown function:
$f : [0, 1] \to \mathbb{R}$

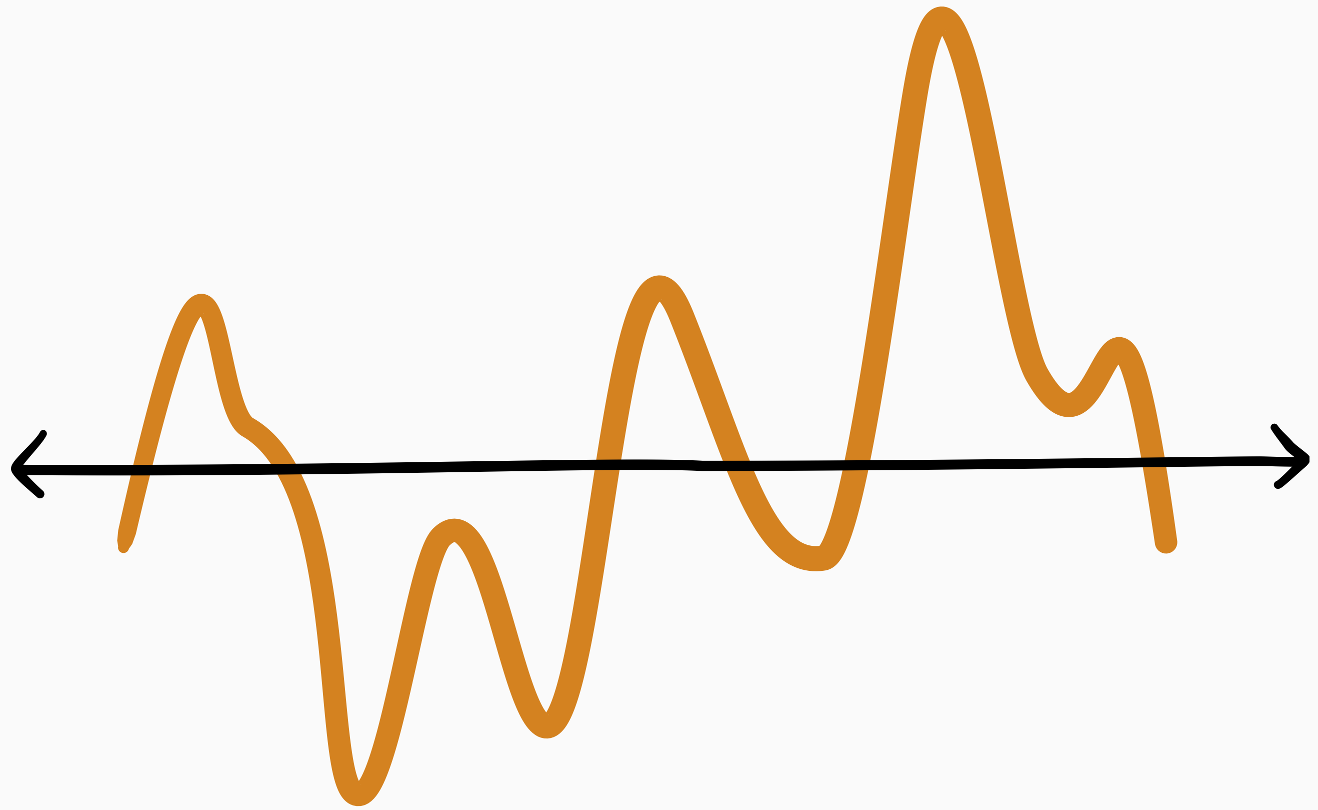

Fixed horizon: maximize

$$\mathbf{E}\Big[\max_{1 \le t \le T} f(x_t)\Big]$$

T fixed

x_t chosen

Cost-per-sample BayesOpt

Assumption:
$f \sim \text{GP}$

Unknown function:
$f : [0, 1] \to \mathbb{R}$

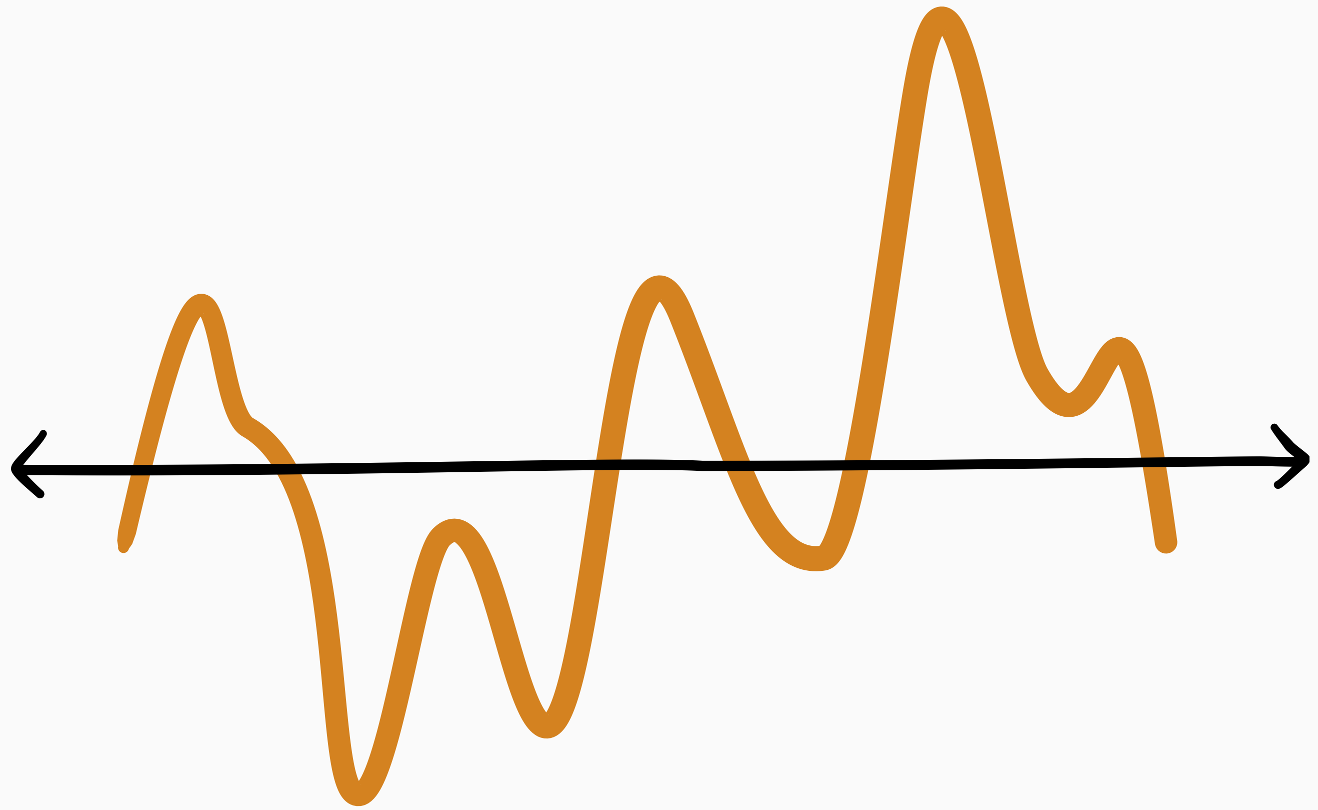

Fixed horizon: maximize

$$\mathbf{E}\left[\max_{1 \leq t \leq T} f(x_t)\right]$$

T fixed

x_t chosen

Cost-per-sample: maximize

$$\mathbf{E}\left[\max_{1 \leq t \leq T} f(x_t)\right] - cT$$

Cost-per-sample **BayesOpt**

Assumption: $f \sim \mathrm{GP}$

Unknown function: $f : [0, 1] \to \mathbb{R}$

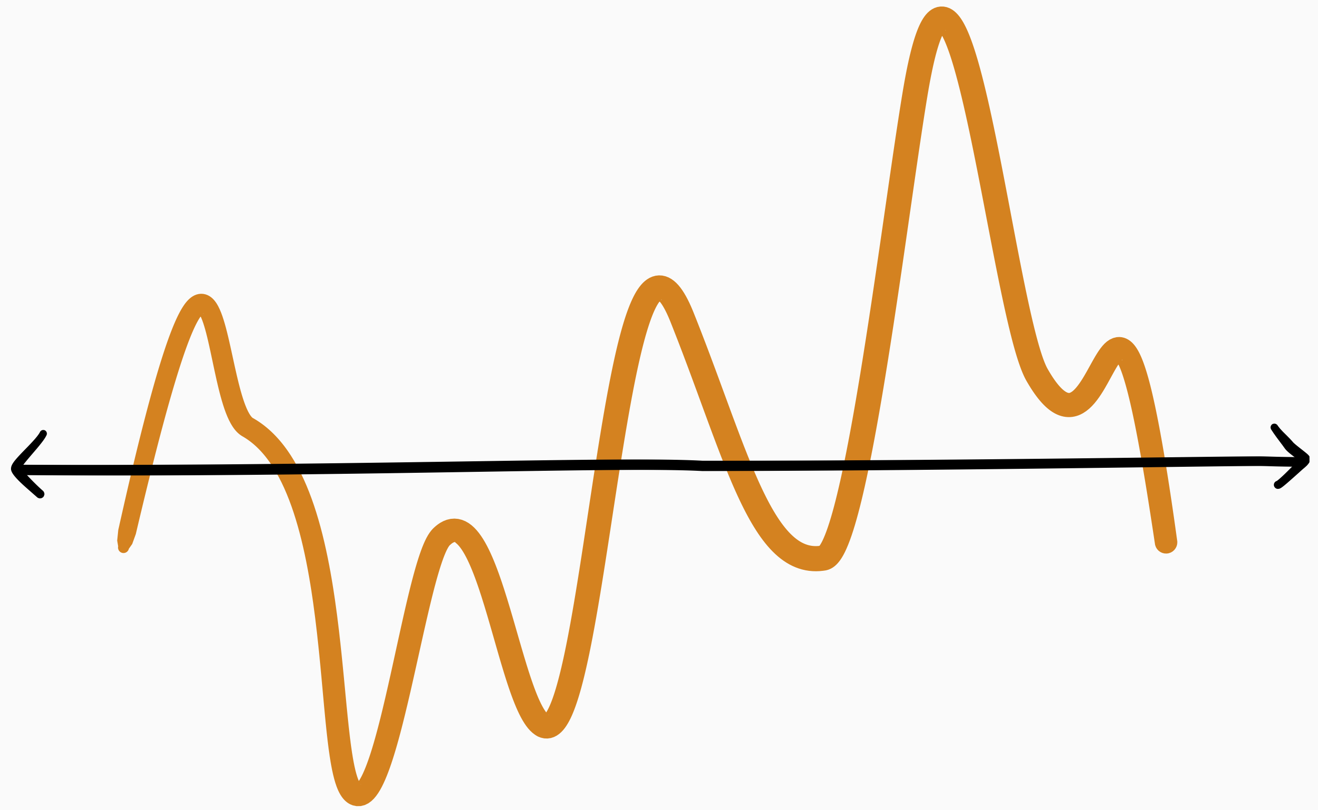

Fixed horizon: maximize

$$\mathbf{E}\Big[\max_{1\le t\le T} f(x_t)\Big]$$

T fixed

x_t chosen

Cost-per-sample: maximize

$$\mathbf{E}\Big[\max_{1\le t\le T} f(x_t)\Big] - cT$$

T chosen

x_t chosen

BayesOpt *vs.* **Pandora's box**

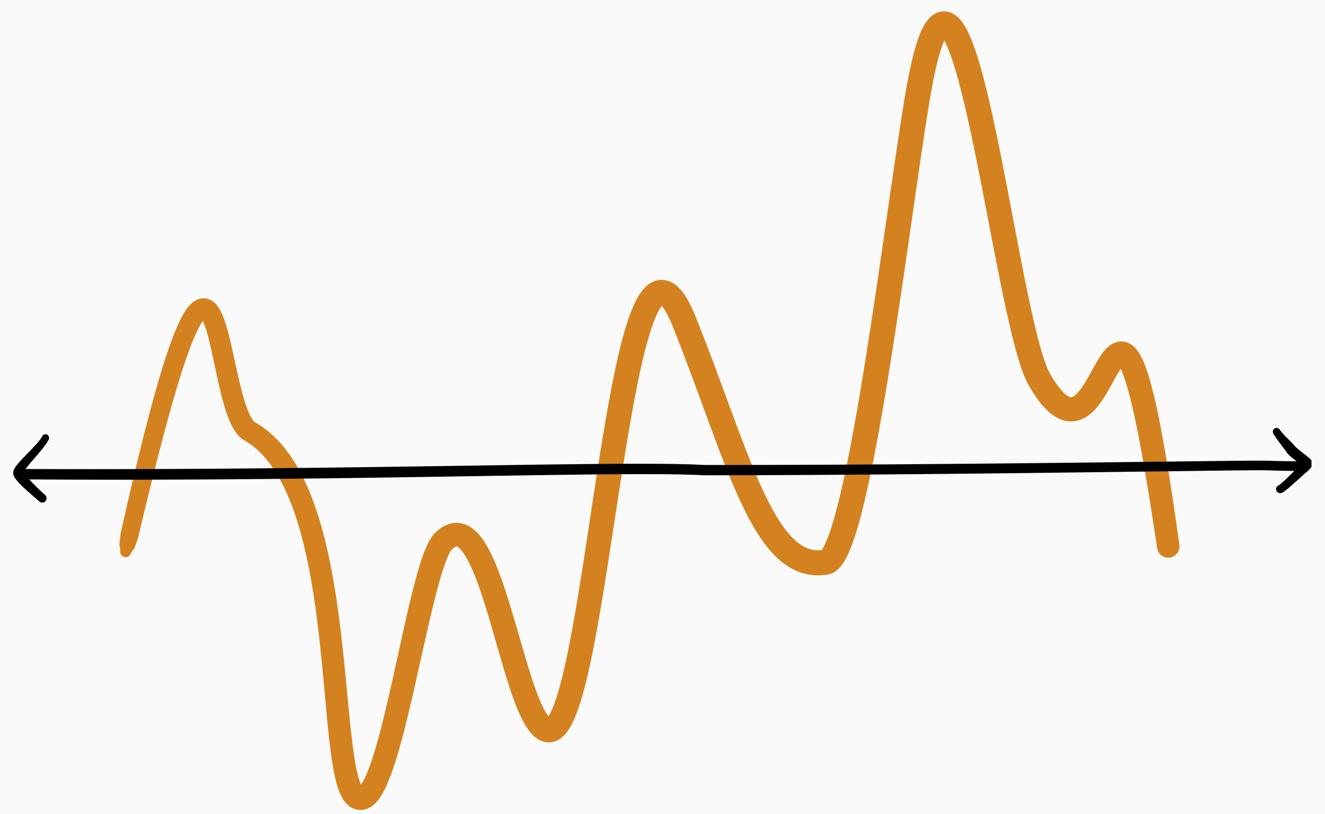

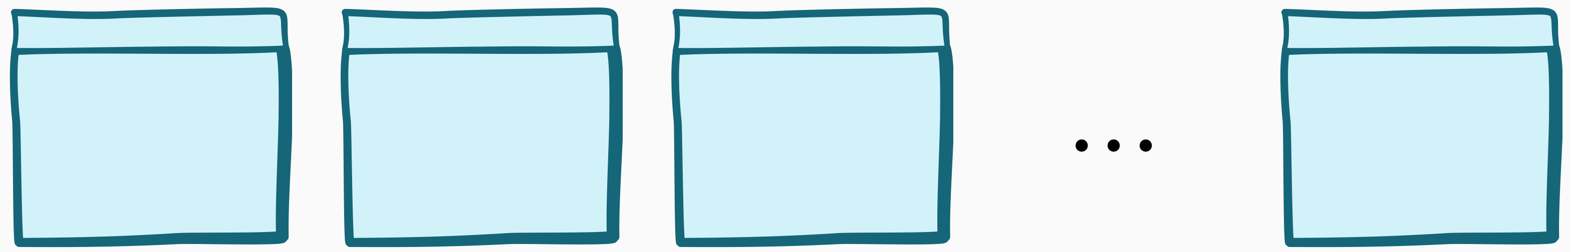

BayesOpt *vs.* **Pandora's box**

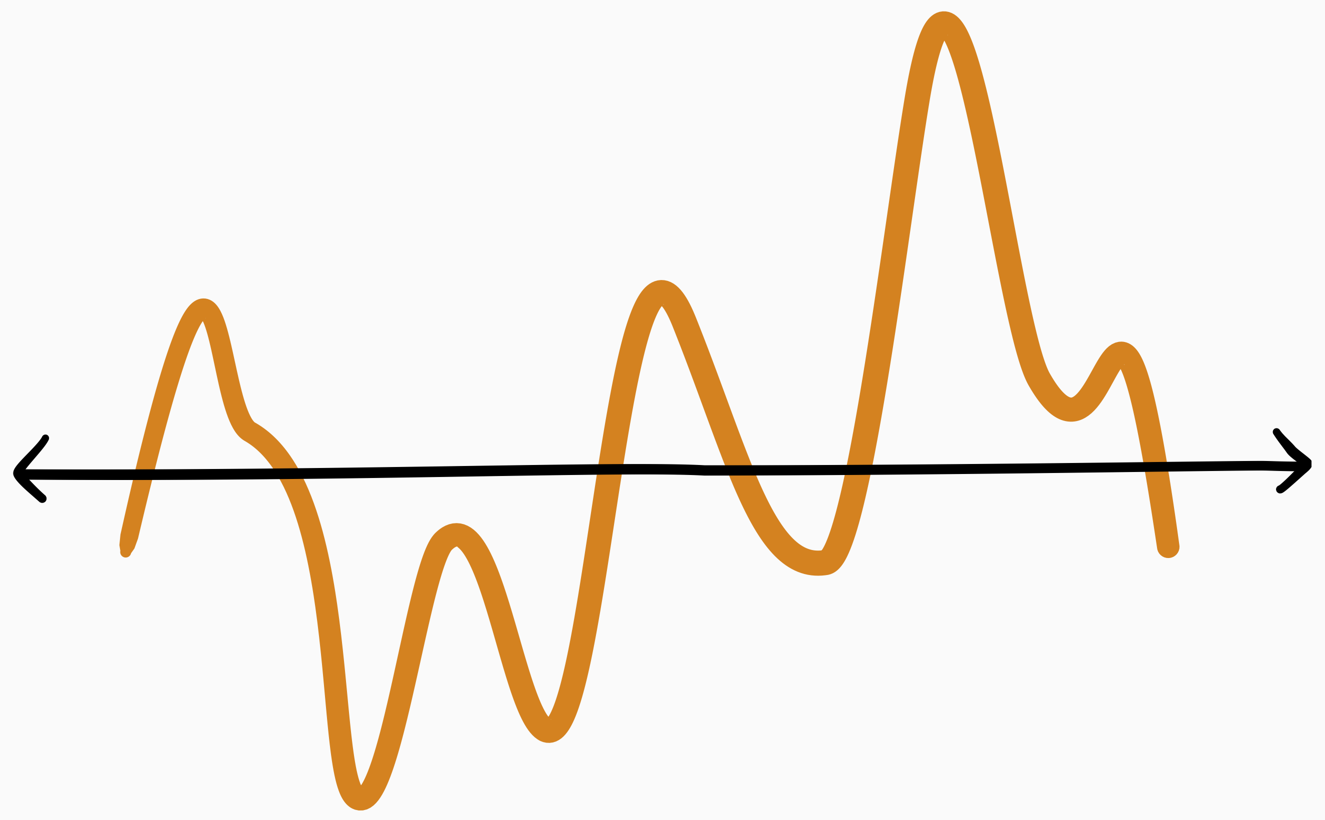

- *Continuous* domain

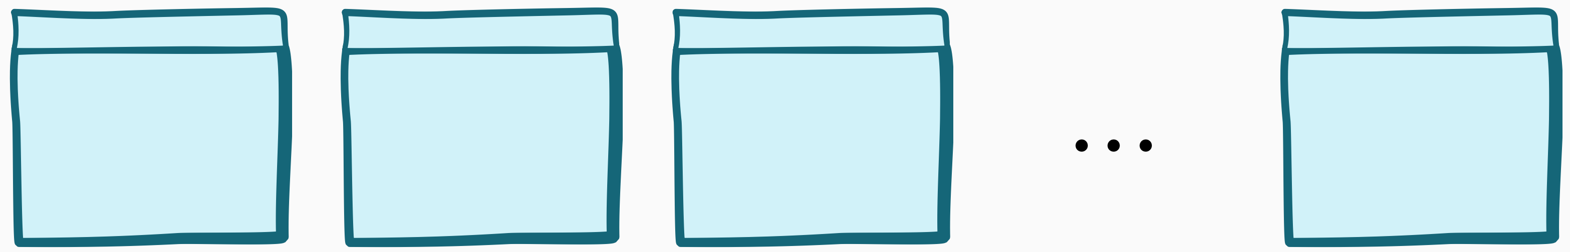

BayesOpt *vs.* Pandora's box

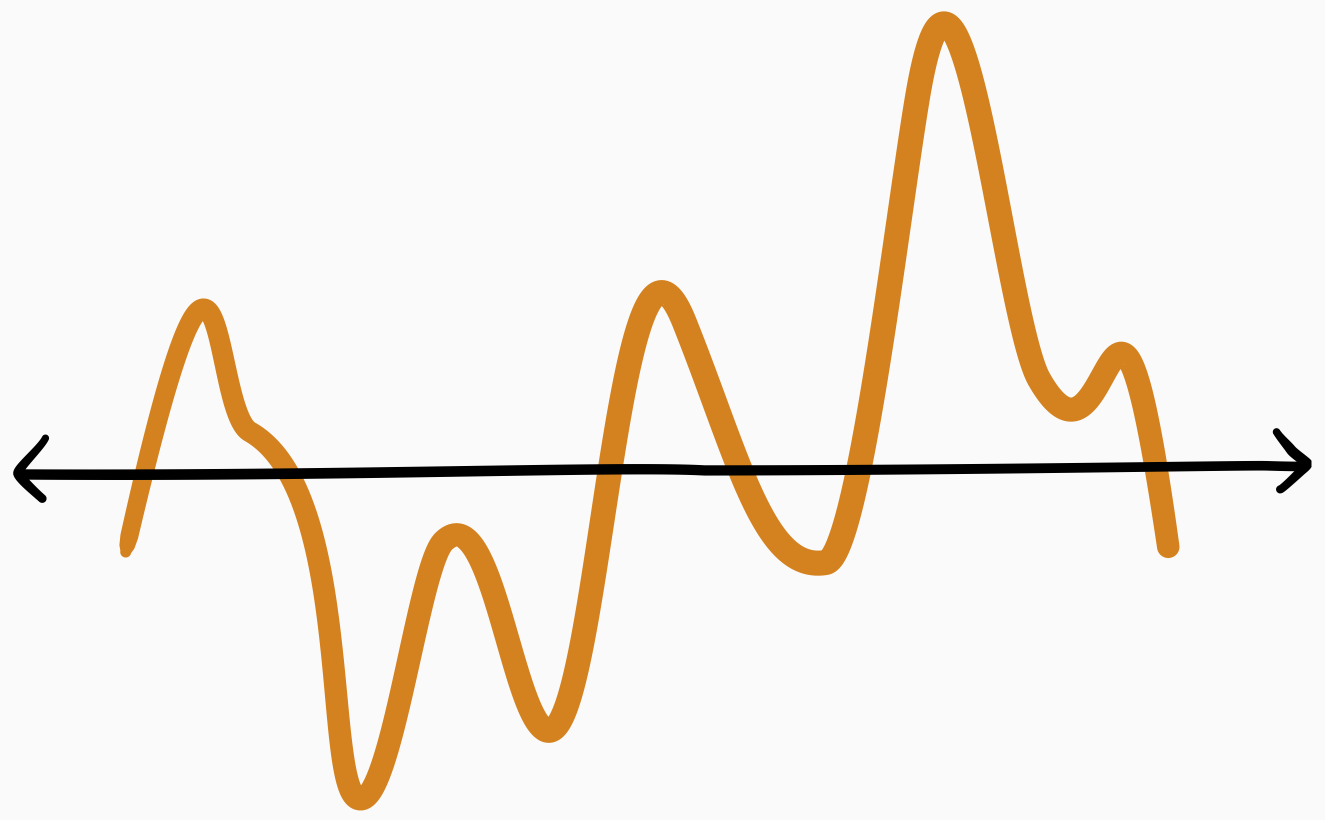

- *Continuous* domain

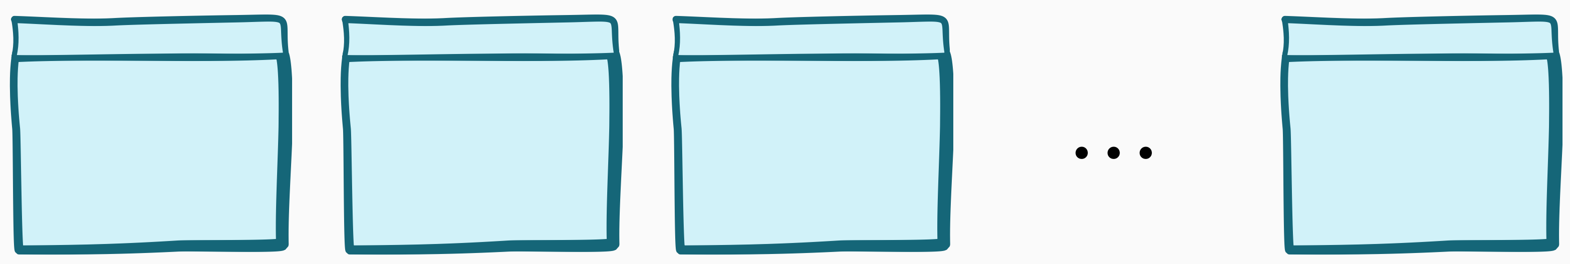

- *Discrete* domain

BayesOpt *vs.* **Pandora's box**

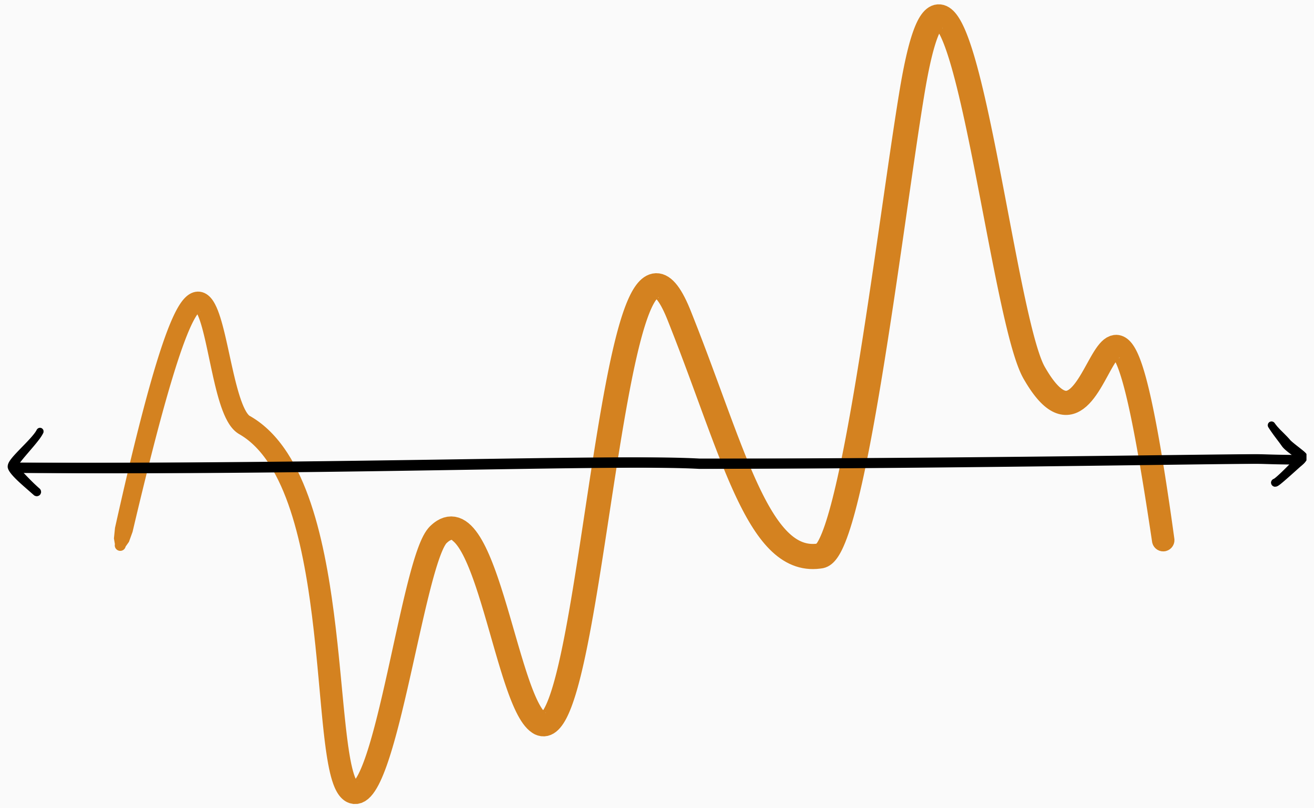

- *Continuous* domain
- *Correlated* values

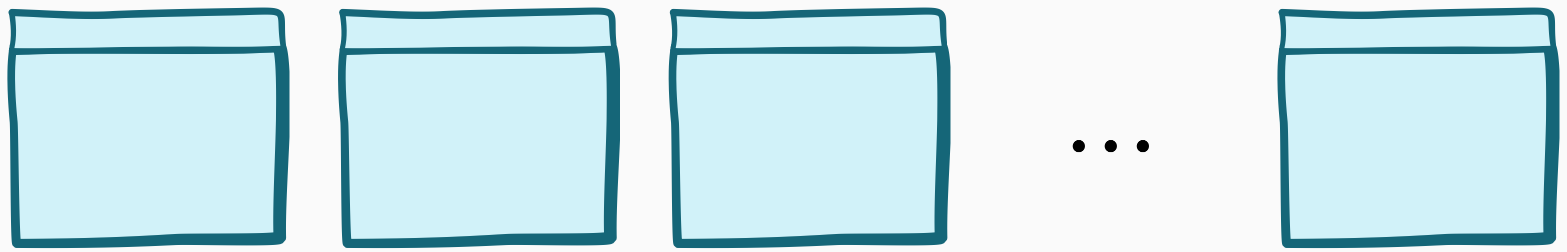

- *Discrete* domain

BayesOpt *vs.* Pandora's box

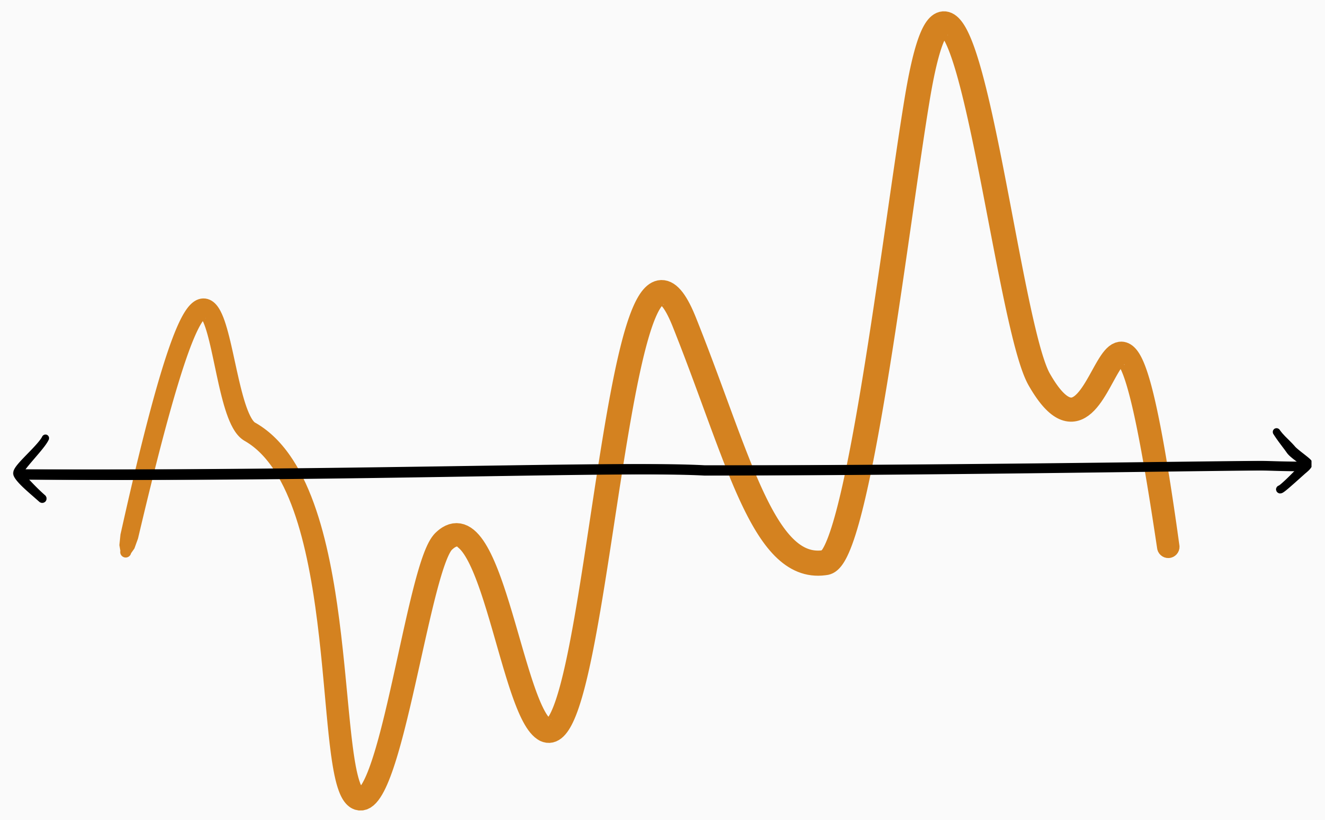

- *Continuous* domain
- *Correlated* values

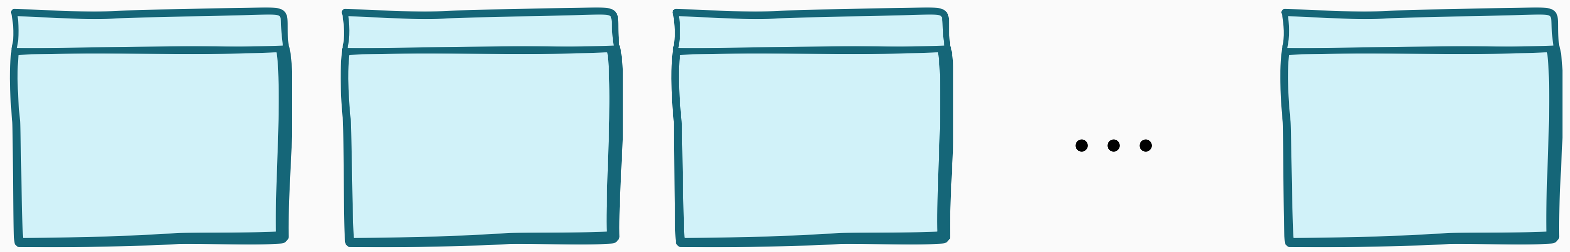

- *Discrete* domain
- *Independent* values

BayesOpt *vs.* Pandora's box

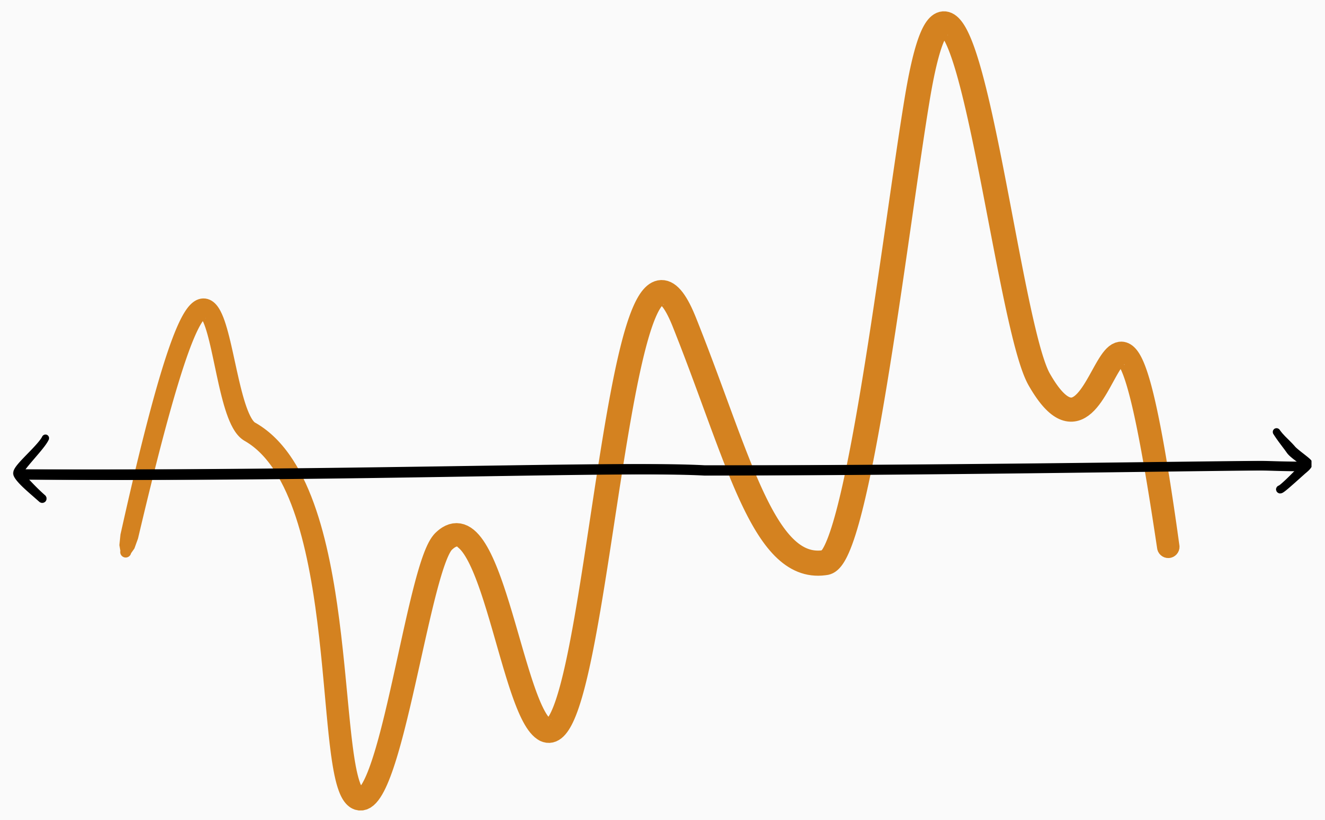

- *Continuous* domain
- *Correlated* values
- *Heuristic* acquisition fns:
  EI, UCB, TS, KG, …

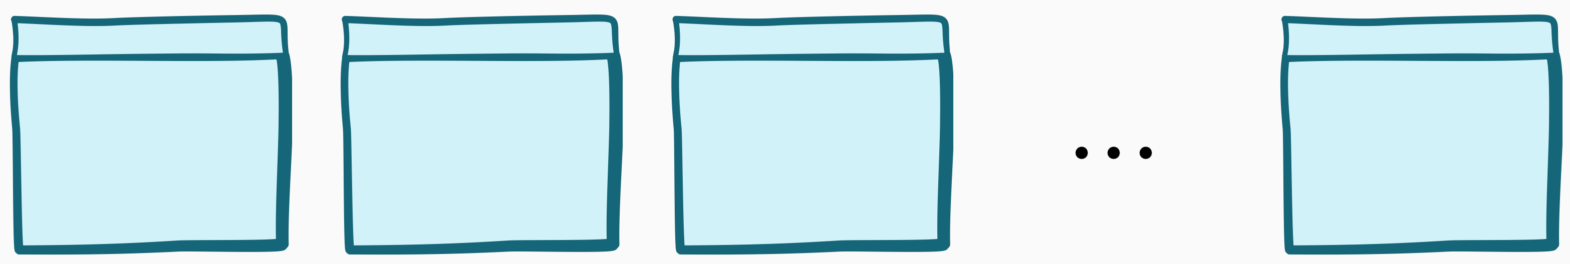

- *Discrete* domain
- *Independent* values

BayesOpt *vs.* Pandora's box

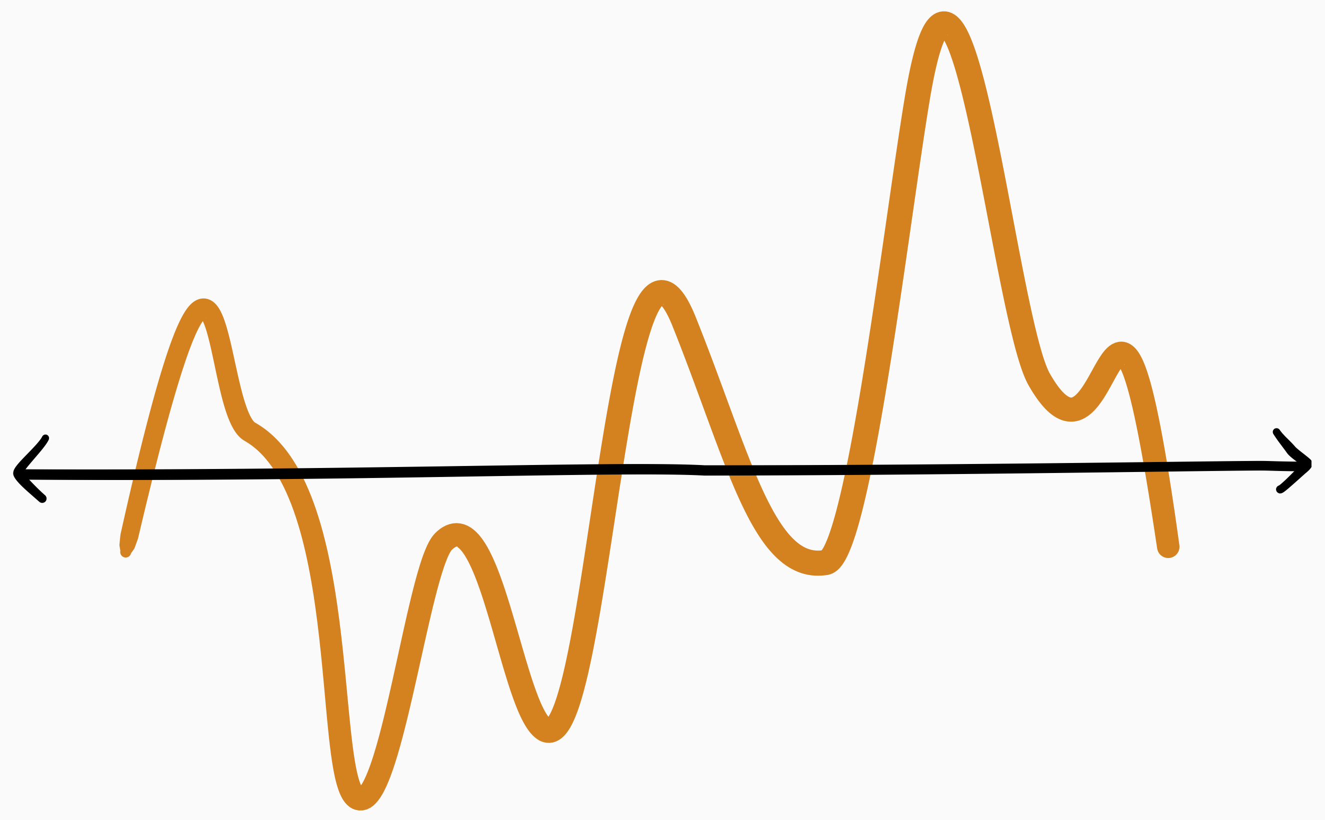

- *Continuous* domain
- *Correlated* values
- *Heuristic* acquisition fns: EI, UCB, TS, KG, …

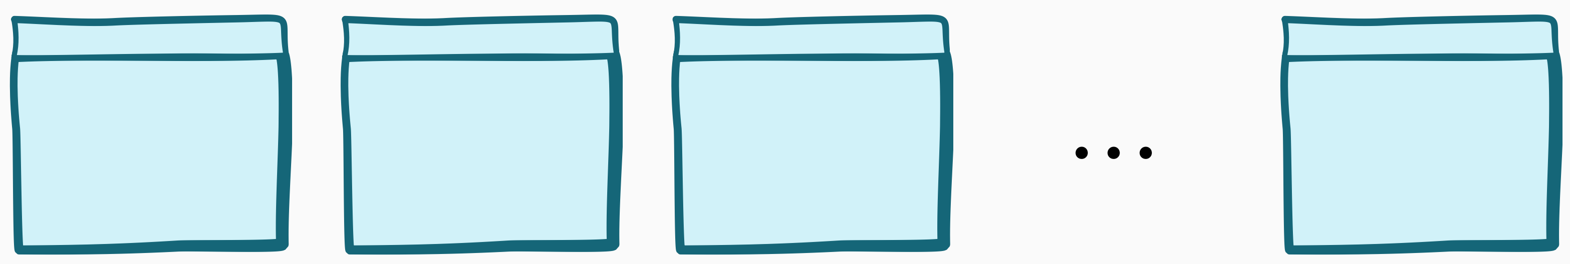

- *Discrete* domain
- *Independent* values
- *Optimal* acquisition fn: Gittins index

BayesOpt *vs.* Pandora's box

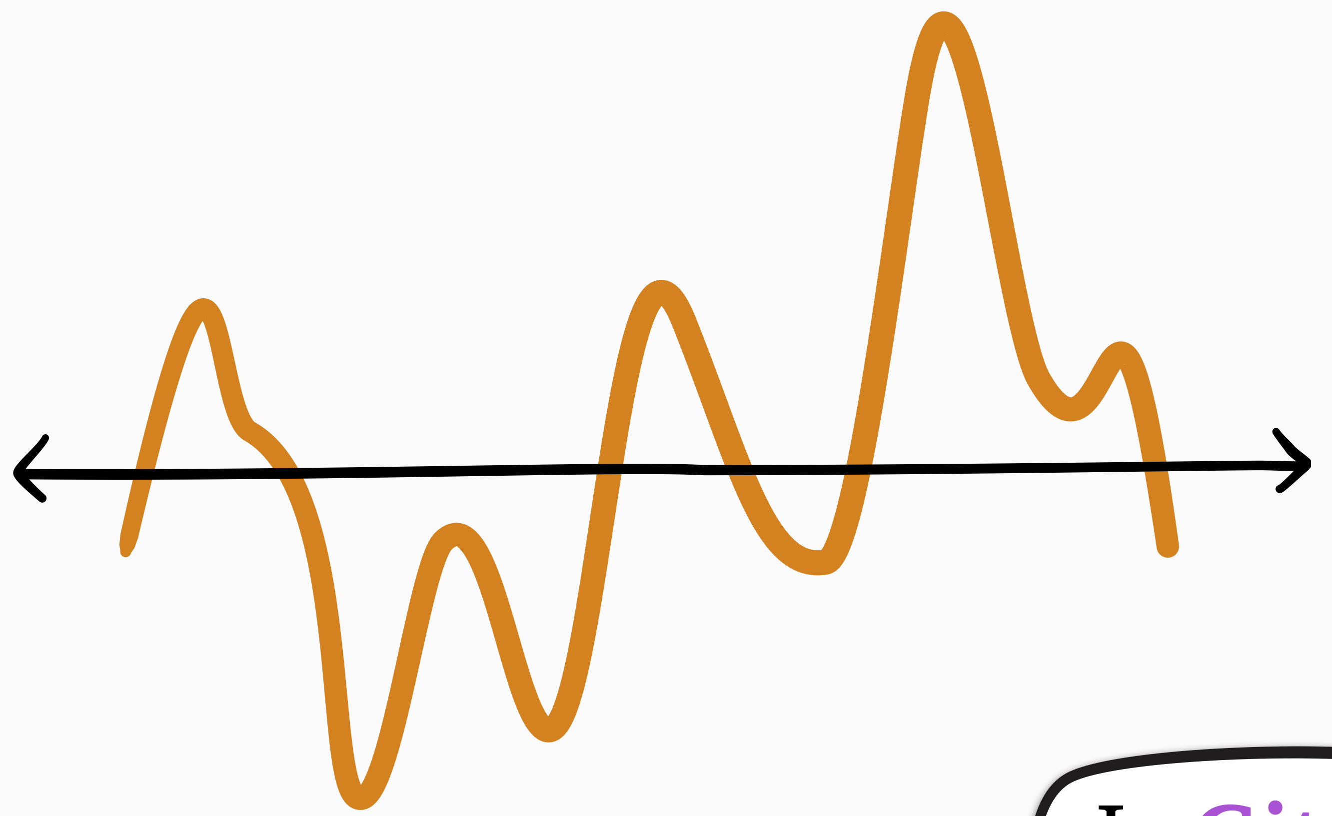

- *Continuous* domain
- *Correlated* values
- *Heuristic* acquisition fns: EI, UCB, TS, KG, …

Is Gittins any good?

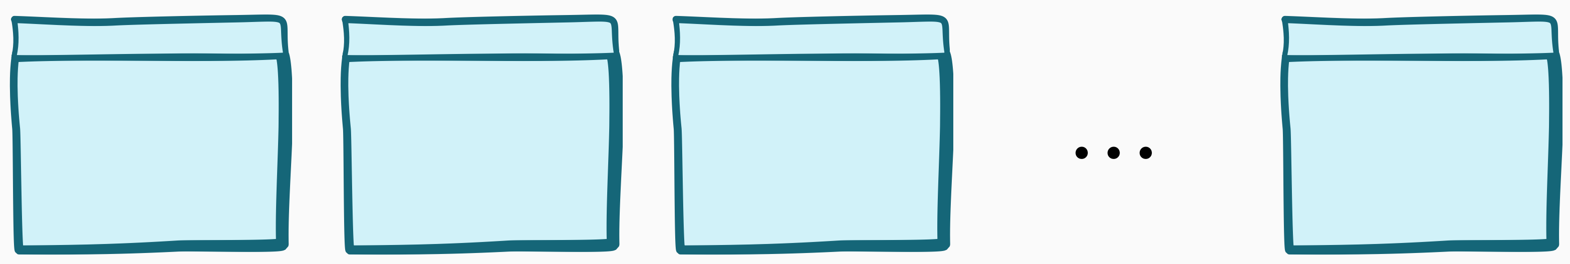

- *Discrete* domain
- *Independent* values
- *Optimal* acquisition fn: Gittins index

BayesOpt acquisition functions

$$a_{\mathrm{EI}}(x) = \mathbf{E}\Big[\big(f(x) - f_t^*\big)^+ \;\Big|\; \mathrm{data}_t\Big]$$

BayesOpt acquisition functions

$$
\begin{aligned}
a_{\mathrm{EI}}(x) &= \mathbf{E}\Big[(f(x) - f_t^*)^+ \;\Big|\; \mathrm{data}_t\Big] \\
&= \mathrm{EI}\Big([f(x) \mid \mathrm{data}_t], f_t^*\Big)
\end{aligned}
$$

BayesOpt acquisition functions

$\max\{f(x_1), \dots, f(x_t)\}$

$$a_{\mathrm{EI}}(x) = \mathbf{E}\big[(f(x) - f_t^*)^+ \mid \mathrm{data}_t\big]$$
$$= \mathrm{EI}\big([f(x) \mid \mathrm{data}_t], f_t^*\big)$$

BayesOpt acquisition functions

$\max\{f(x_1), \dots, f(x_t)\}$

$f(x_1), \dots, f(x_t)$

$$a_{\mathrm{EI}}(x) = \mathbf{E}\big[(f(x) - f_t^*)^+ \mid \mathrm{data}_t\big]$$
$$= \mathrm{EI}\big([f(x) \mid \mathrm{data}_t], f_t^*\big)$$

BayesOpt acquisition functions

$\max\{f(x_1), \ldots, f(x_t)\}$

$f(x_1), \ldots, f(x_t)$

$$a_{\text{EI}}(x) = \mathbf{E}\big[(f(x) - f_t^*)^+ \mid \text{data}_t\big]$$
$$= \text{EI}\big([f(x) \mid \text{data}_t], f_t^*\big)$$

$$a_{\text{Gittins-c}}(x) = g\big(c{:}[f(x) \mid \text{data}_t]\big)$$

BayesOpt acquisition functions

$\max\{f(x_1), \ldots, f(x_t)\}$

$f(x_1), \ldots, f(x_t)$

$$a_{\mathrm{EI}}(x) = \mathbf{E}\big[(f(x) - f_t^*)^+ \mid \mathrm{data}_t\big]$$
$$= \mathrm{EI}\big([f(x) \mid \mathrm{data}_t], f_t^*\big)$$

$$a_{\text{Gittins-}c}(x) = g\big(c{:}[f(x) \mid \mathrm{data}_t]\big)$$

assume independent to compute Gittins index

BayesOpt acquisition functions

$\max\{f(x_1), \ldots, f(x_t)\}$

$f(x_1), \ldots, f(x_t)$

$$a_{\mathrm{EI}}(x) = \mathbf{E}\left[(f(x) - f_t^*)^+ \mid \mathrm{data}_t\right]$$
$$= \mathrm{EI}\left([f(x) \mid \mathrm{data}_t], f_t^*\right)$$

$$a_{\text{Gittins-}c}(x) = g\left(c{:}[f(x) \mid \mathrm{data}_t]\right)$$

assume independent to compute Gittins index

use correlation in Bayesian update

Results: Bayesian regret

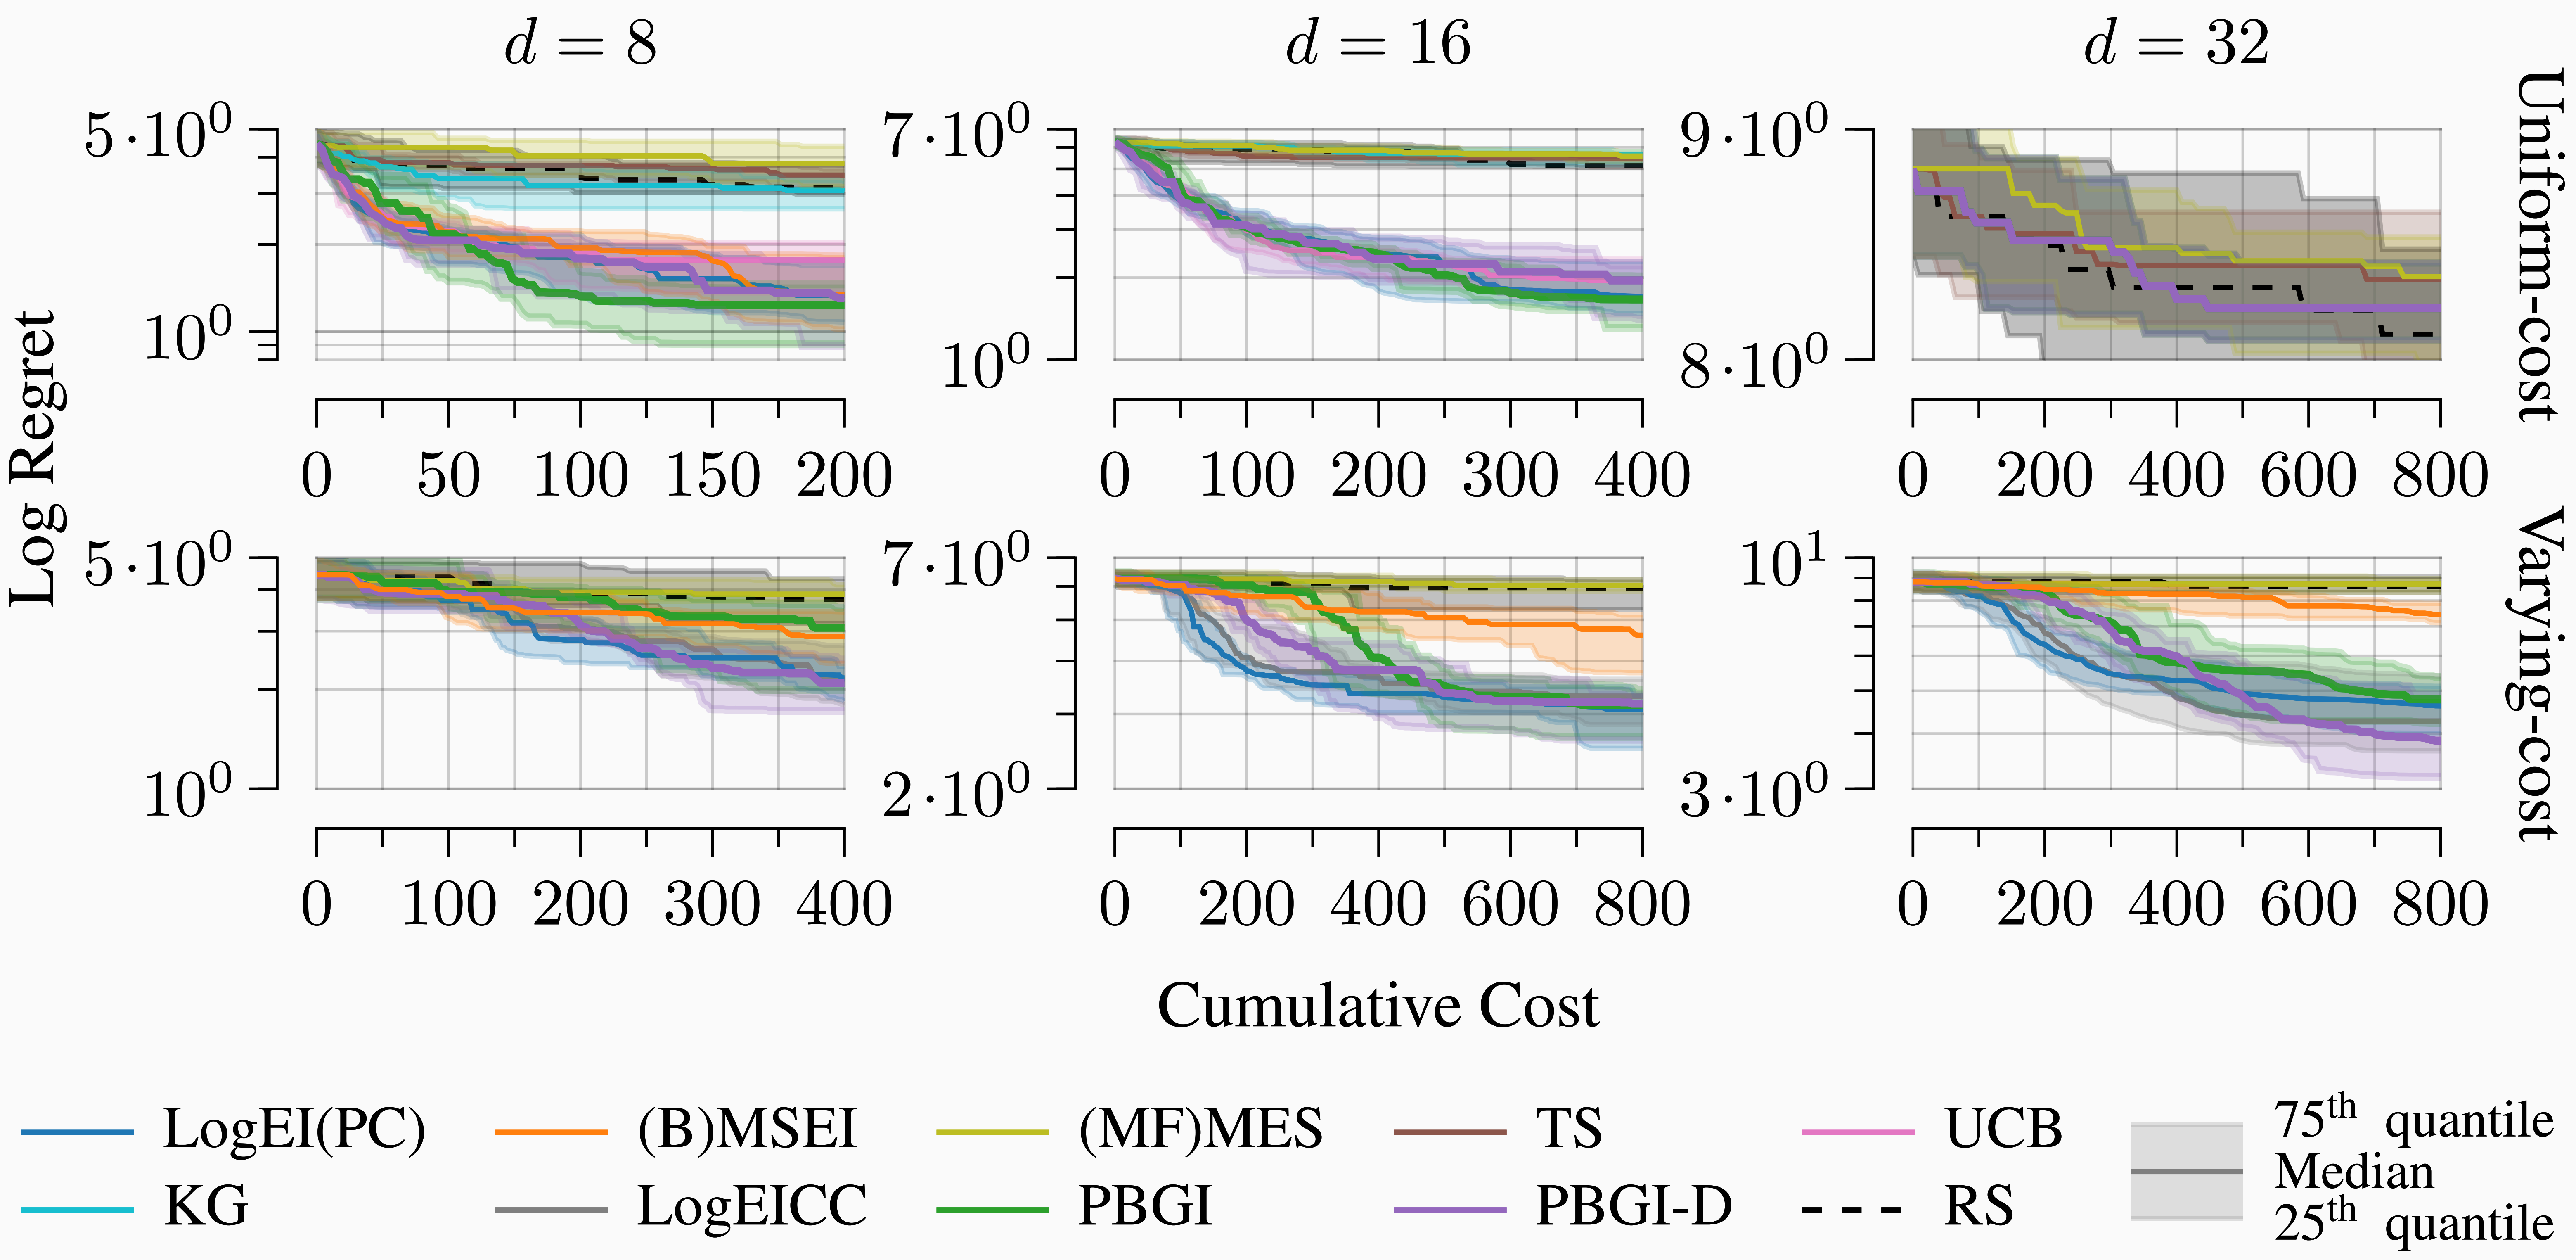

Results: classic benchmarks

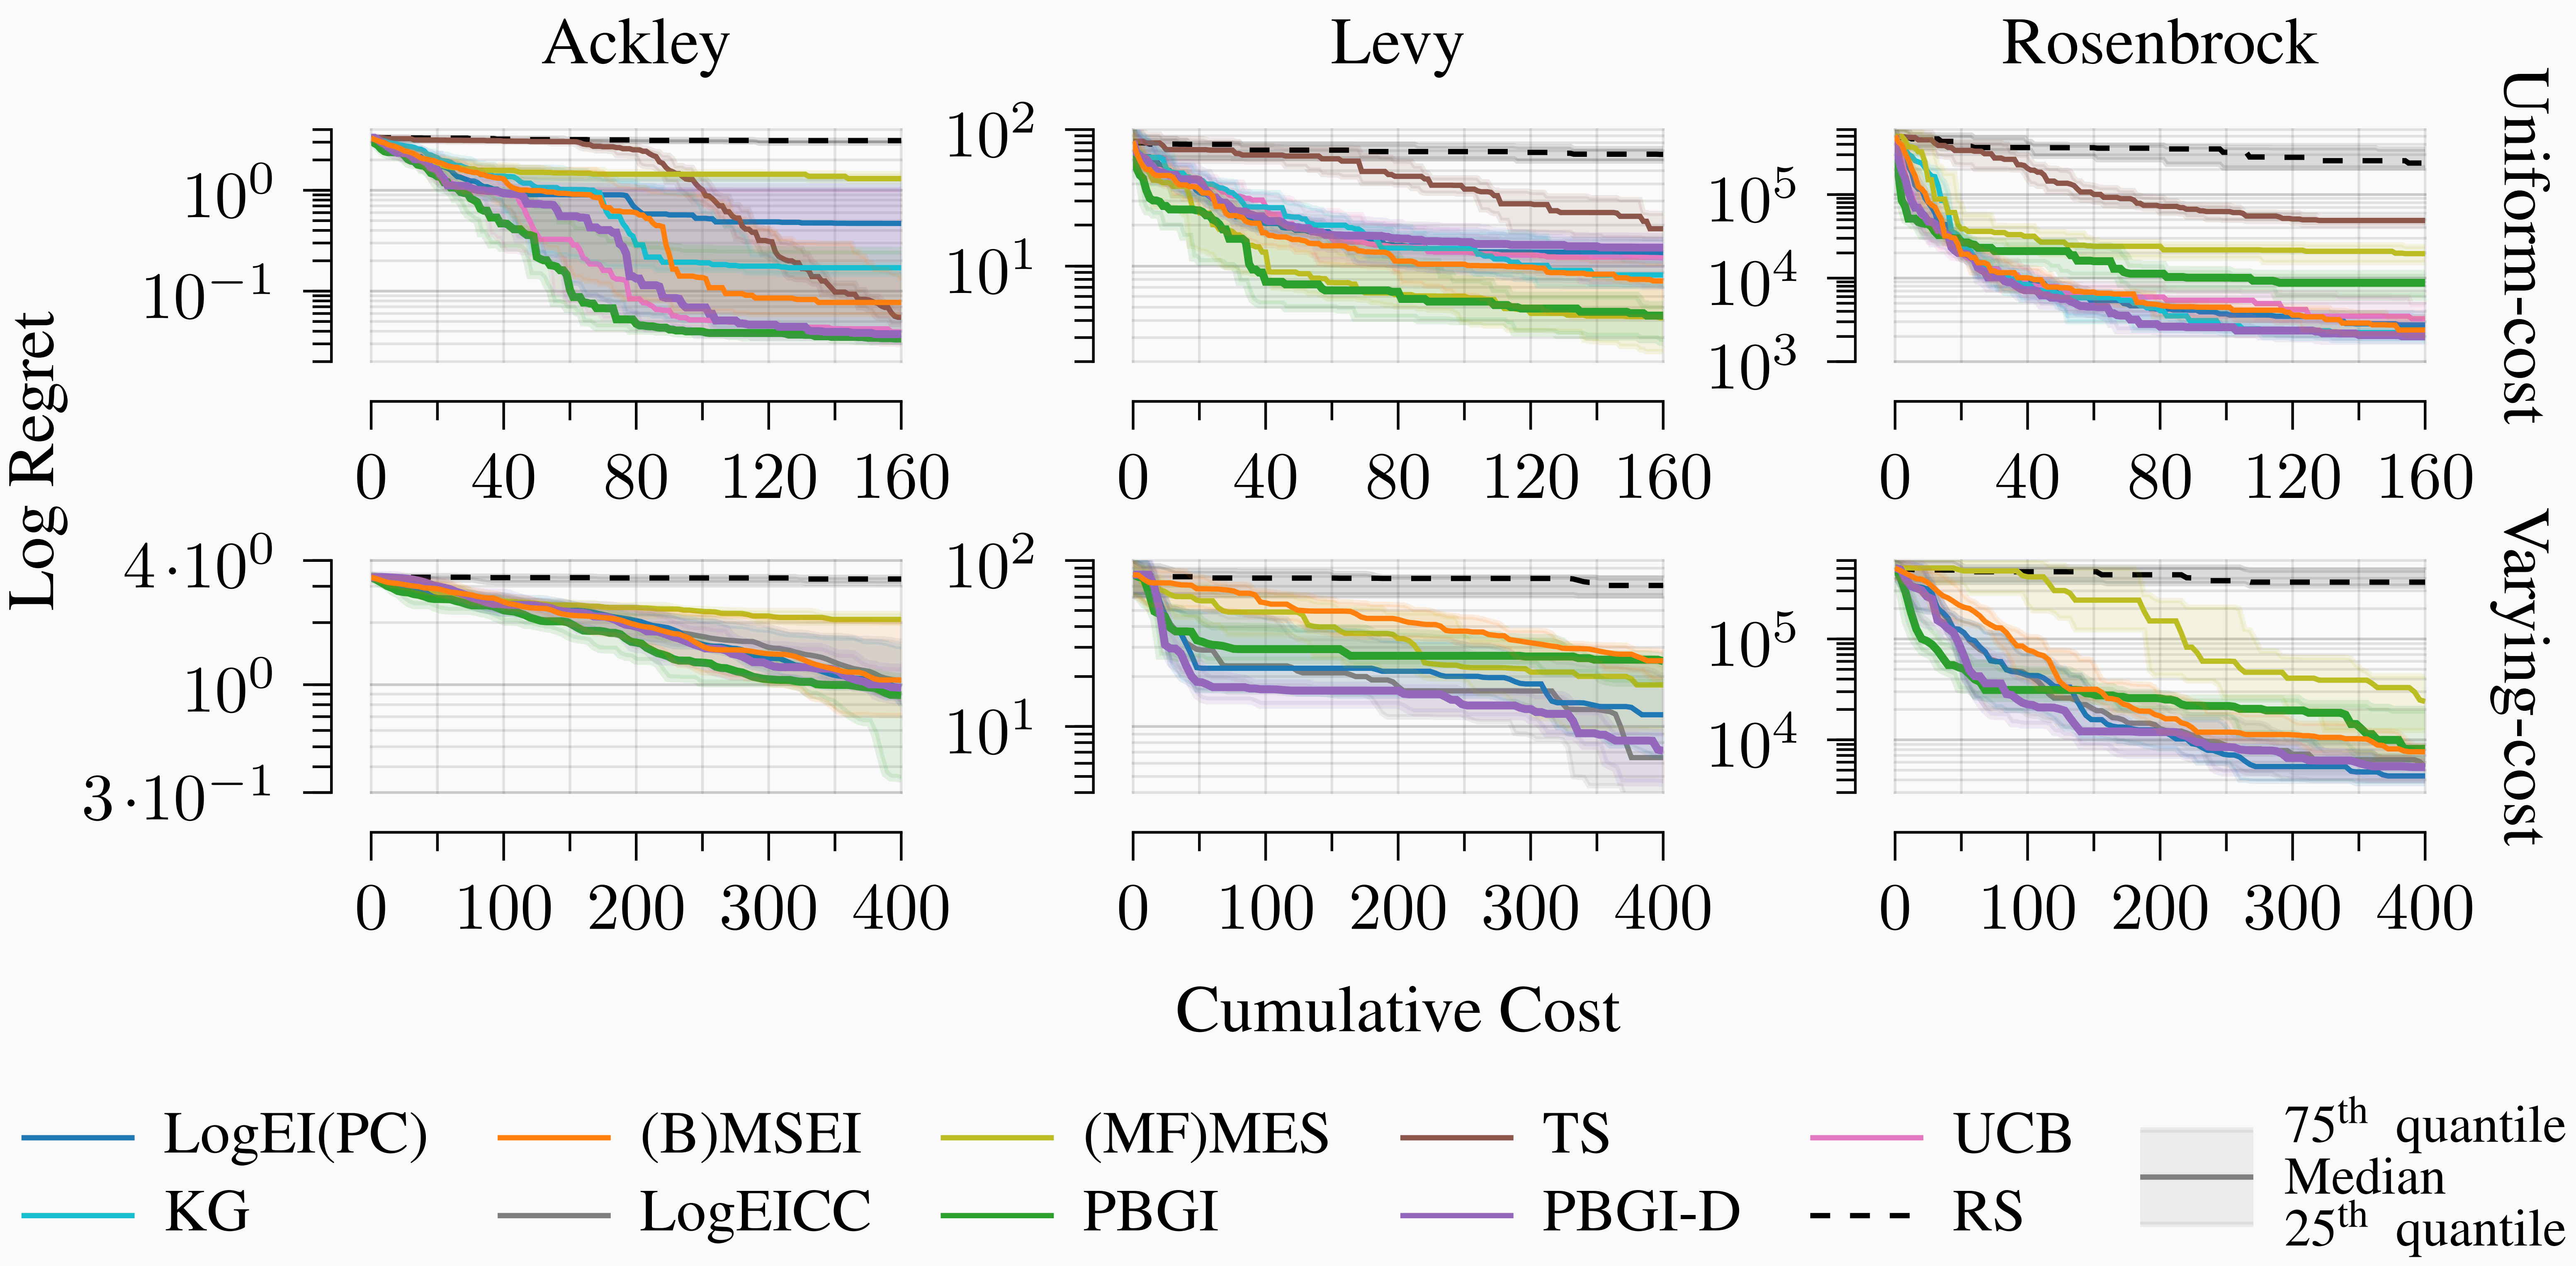

Results: empirical benchmarks

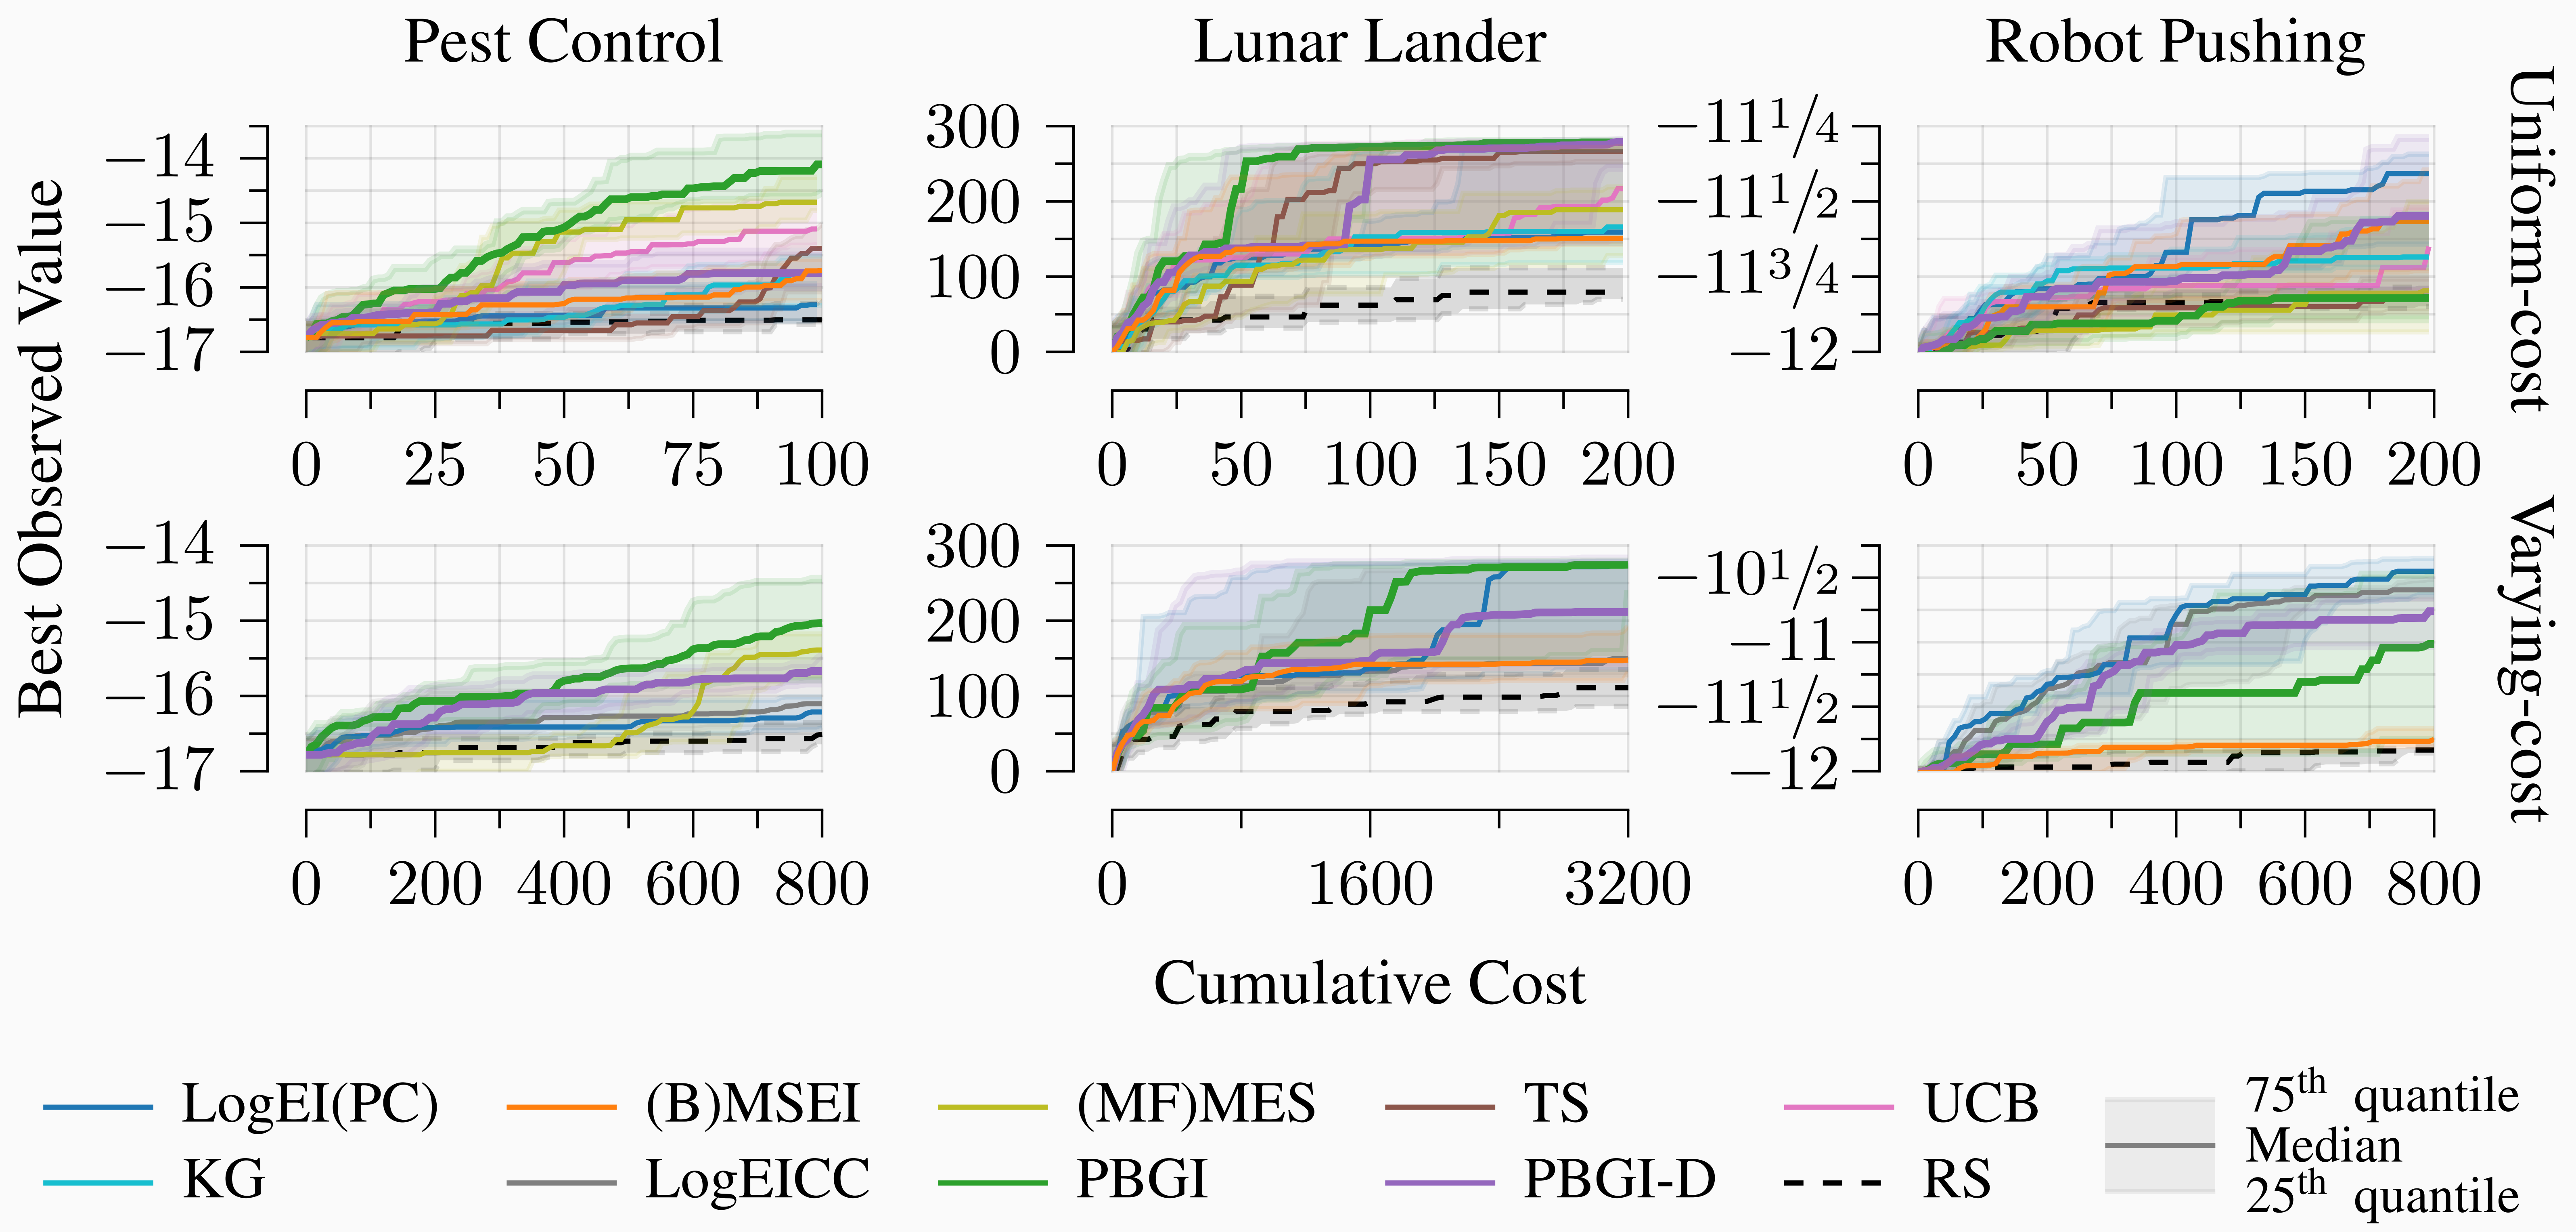

Generalization: Markovian MAB

Pandora's box: one stage

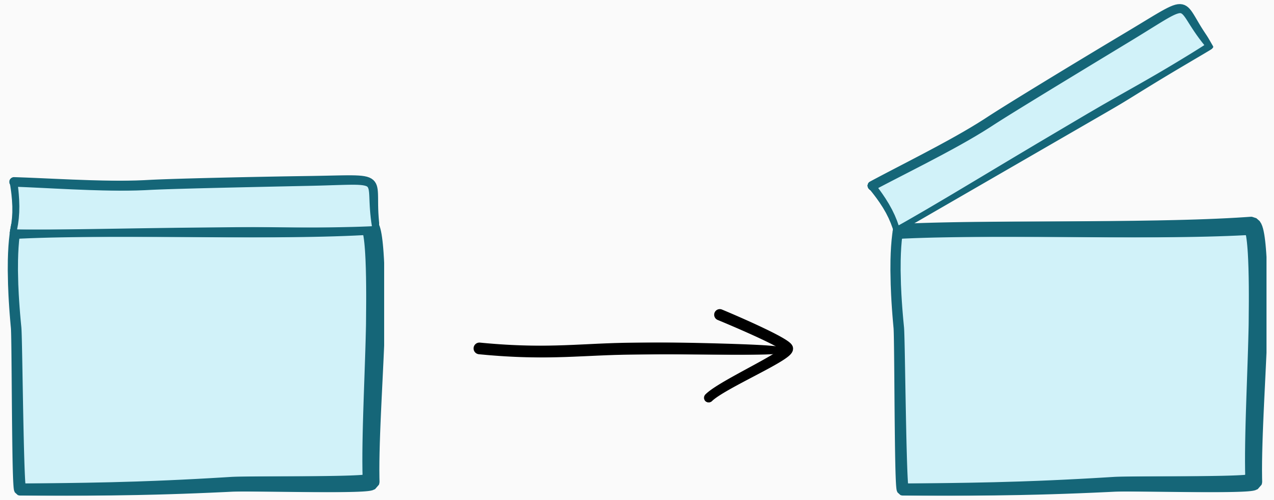

Generalization: Markovian MAB

Pandora's box: one stage

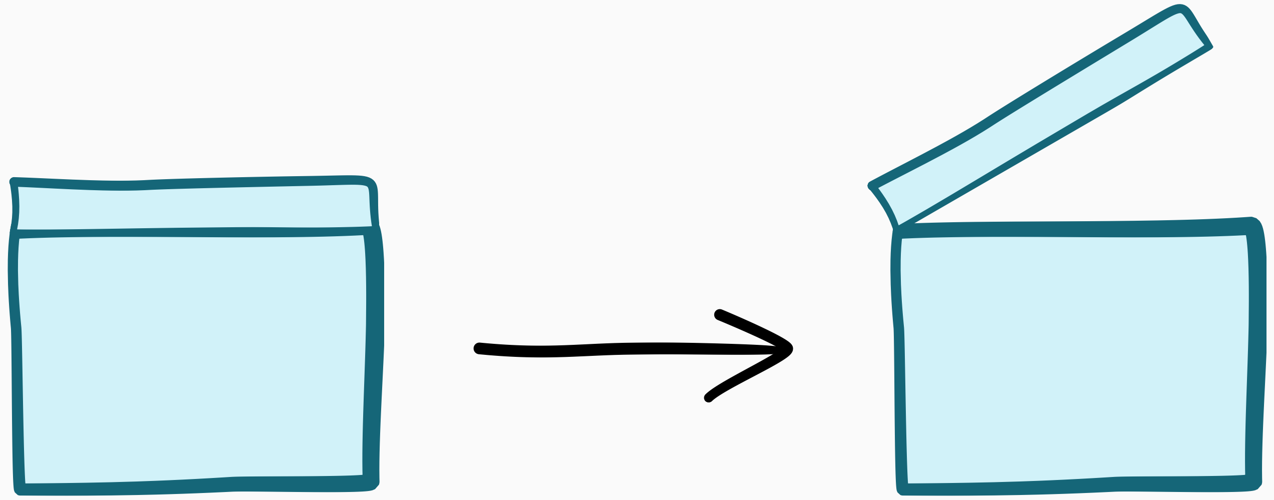

Markovian MAB: general Markov chain

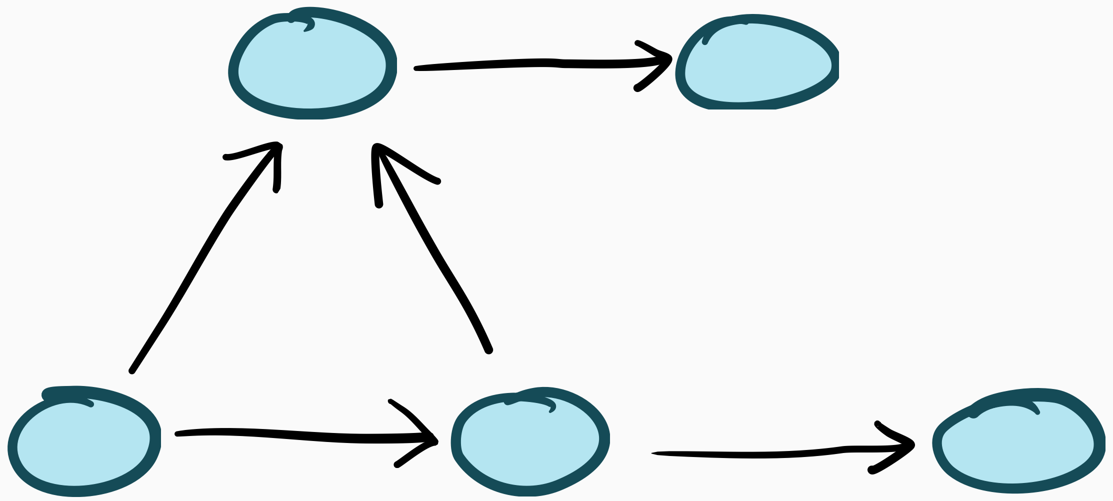

Generalization: Markovian MAB

Pandora's box: one stage

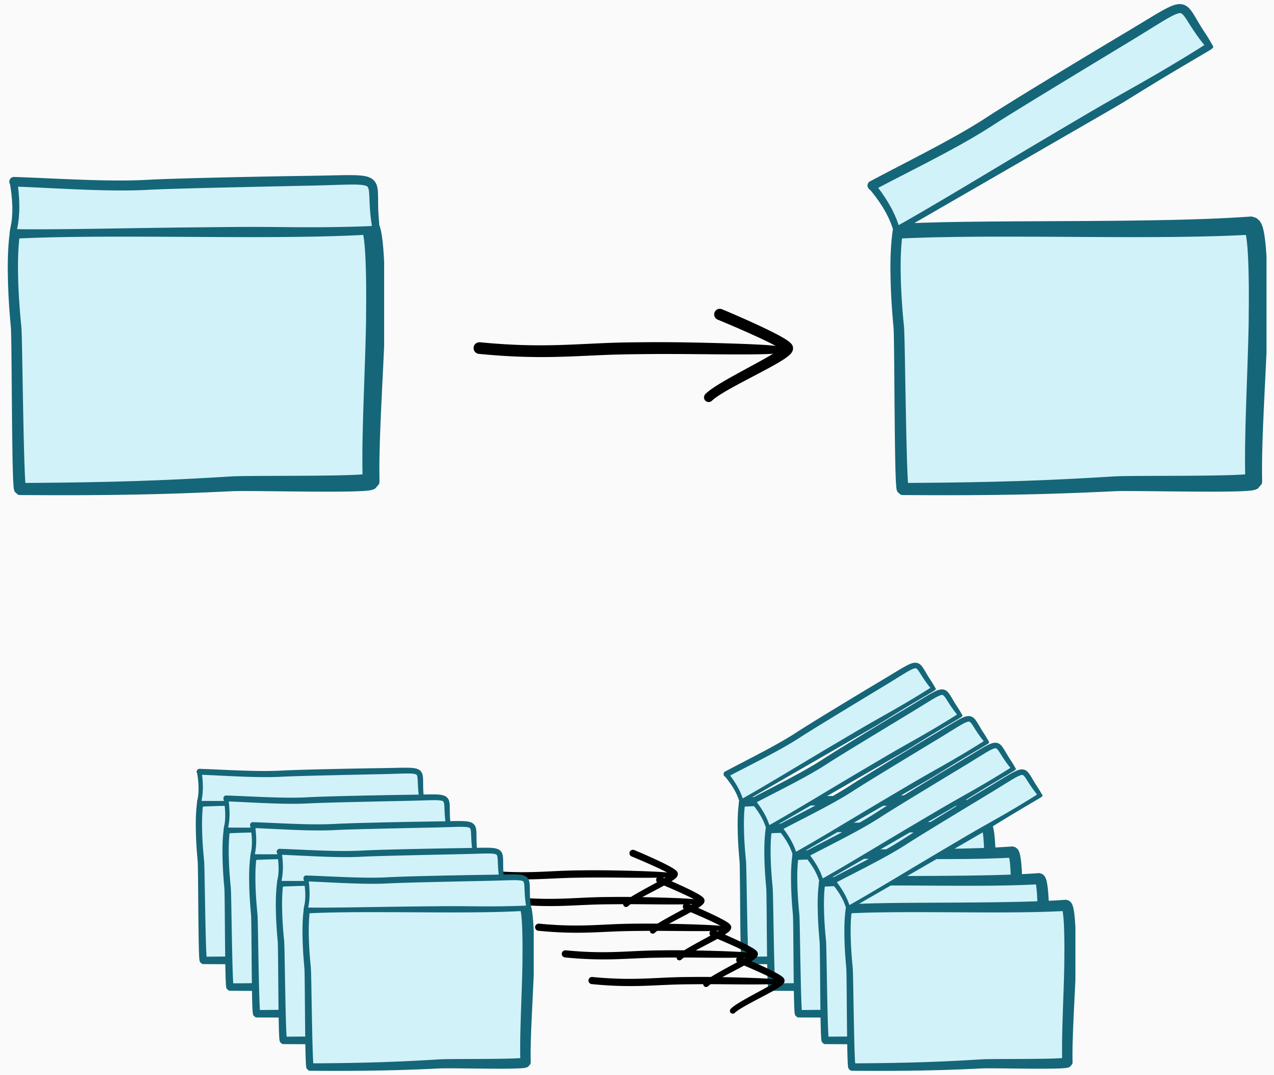

Markovian MAB: general Markov chain

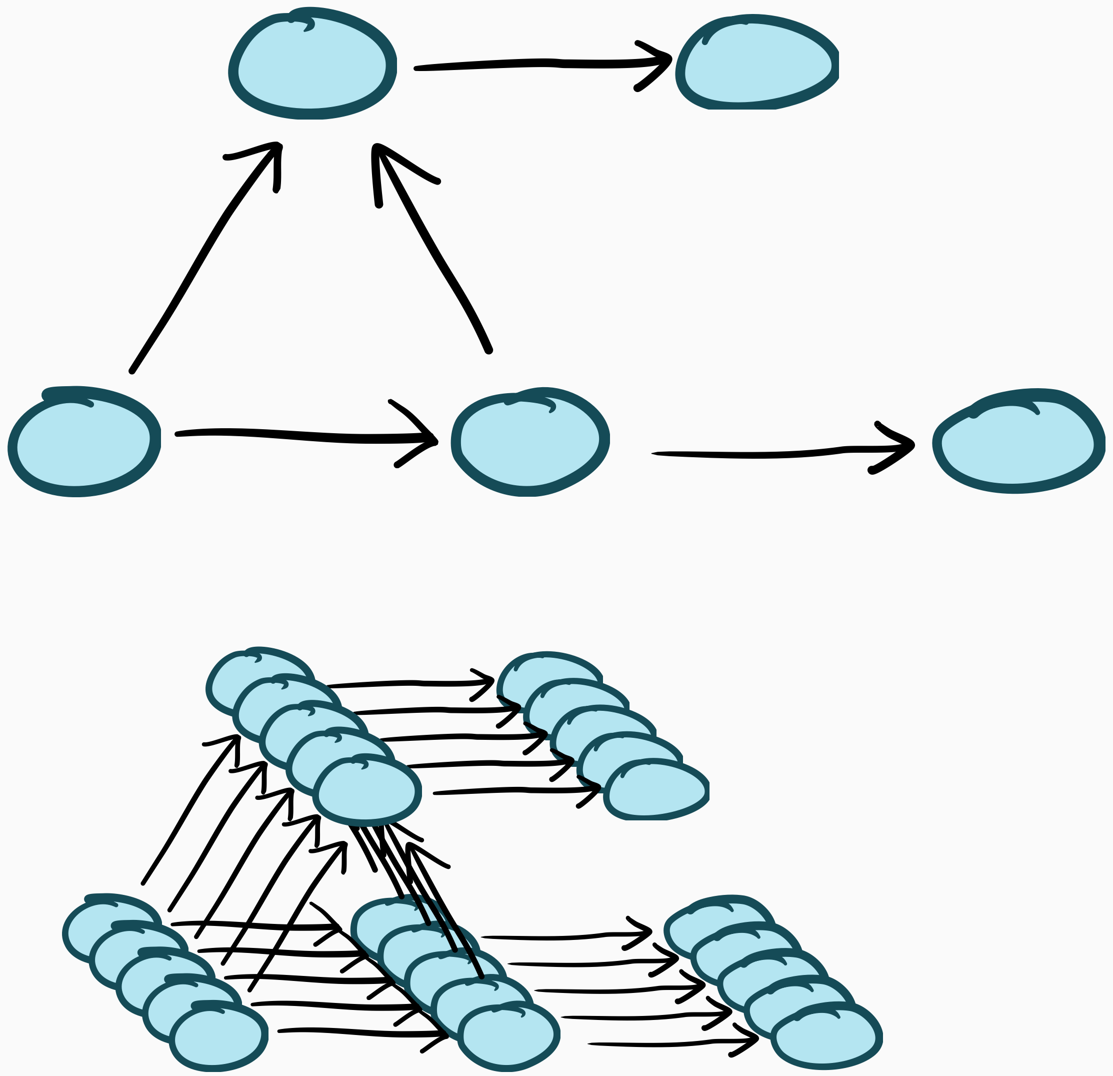

Summary

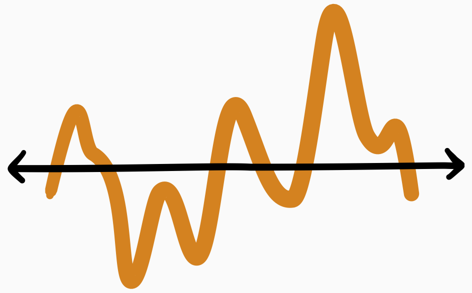 BayesOpt

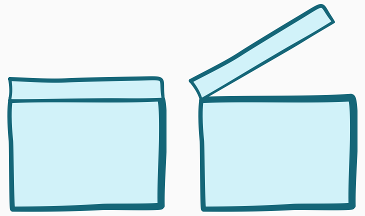 Pandora's box

Gittins is optimal

“Fancy” BayesOpt

- Multistage inspection
- Heterogeneous costs
- Partial feedback
- Automatic stopping

Markovian MAB

general multistage processes

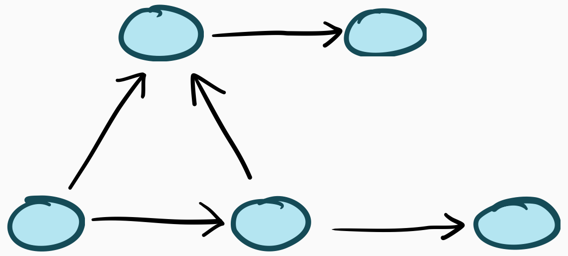

Summary

BayesOpt

Pandora's box

Gittins is promising!

Gittins is optimal

"Fancy" BayesOpt

- Multistage inspection
- Heterogeneous costs
- Partial feedback
- Automatic stopping

Markovian MAB

general multistage processes

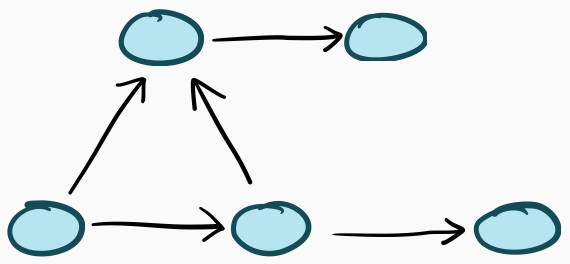

Summary

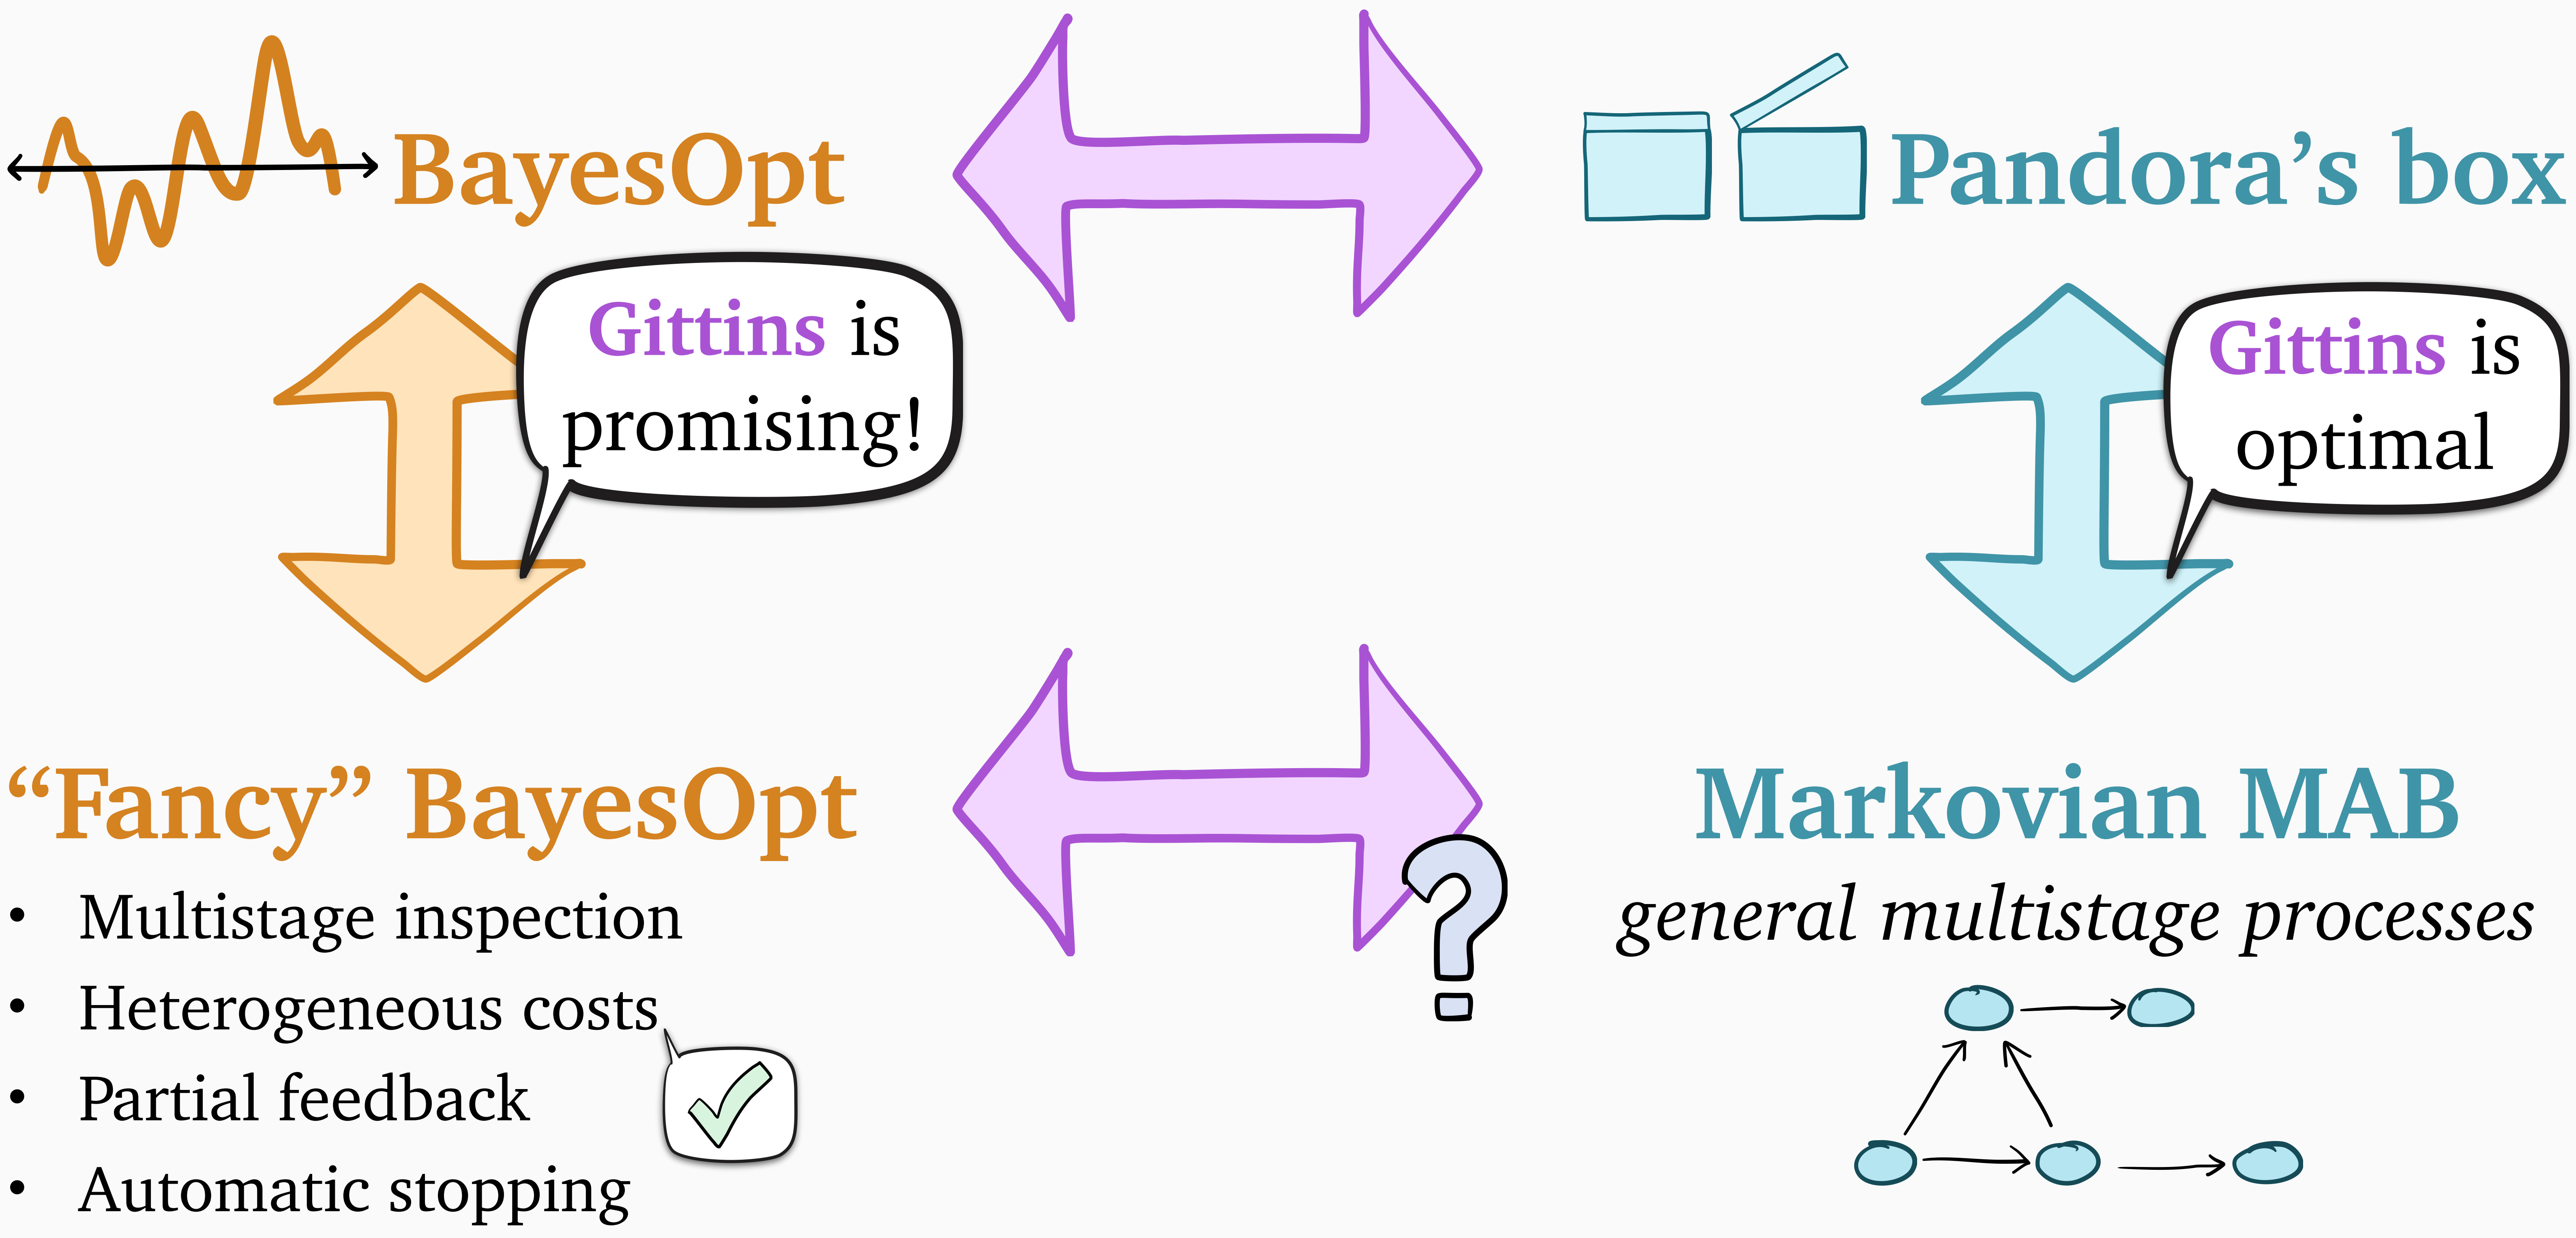
BayesOpt
Pandora's box
Gittins is promising!
Gittins is optimal
"Fancy" BayesOpt
• Multistage inspection
• Heterogeneous costs
• Partial feedback
• Automatic stopping
?
Markovian MAB
general multistage processes

Summary

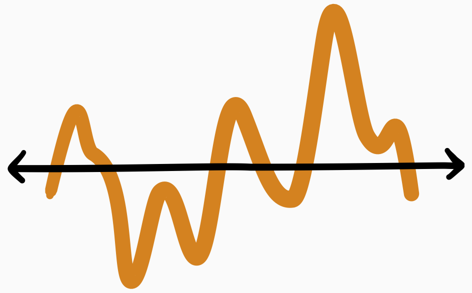

Ba

ora's box

Gittins is optimal

New tutorial

The Gittins Index: A Design Principle for Decision-Making Under Uncertainty
[arXiv:2506.10872]

Alex Terenin
Cornell

“Fancy” B

n MAB

general multistage processes

- Multistage inspection
- Heterogeneous costs
- Partial feedback
- Automatic stopping

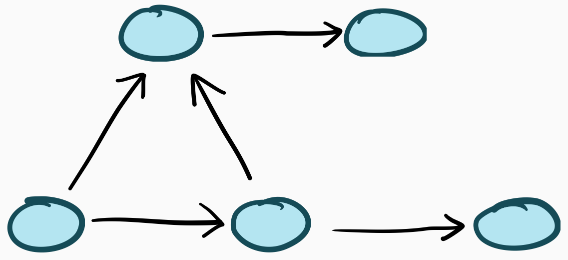

Summary

New tutorial with Alex Terenin
The Gittins Index: A Design Principle for Decision-Making Under Uncertainty
[arXiv:2506.10872]

BayesOpt

Pandora's box

Gittins is promising!

Gittins is optimal

"Fancy" BayesOpt

- Multistage inspection
- Heterogeneous costs
- Partial feedback
- Automatic stopping

Markovian MAB

general multistage processes