A new design tool for Bayesian optimization: the

Gittins index

Ziv Scully
Cornell University

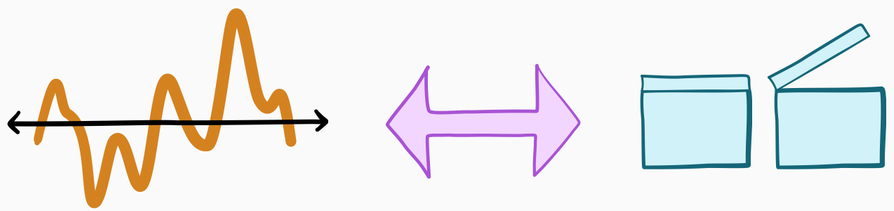

Collaborators

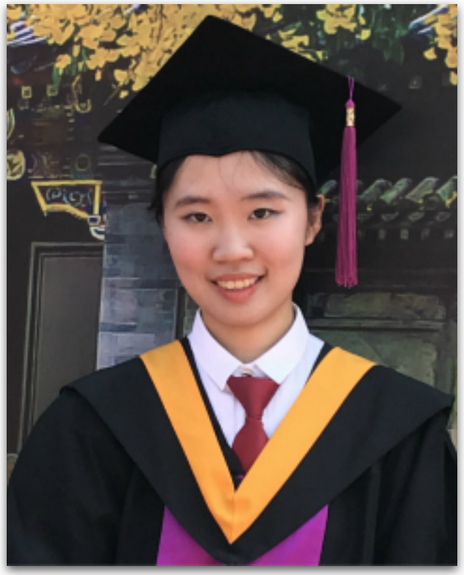

Qian Xie
Cornell

Alex Terenin
Cornell

Collaborators

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

Raul Astudillo
Caltech

Peter Frazier
Cornell

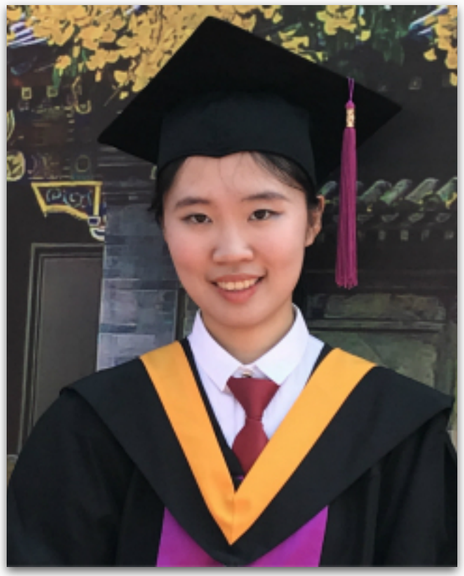

Qian Xie
Cornell

Alex Terenin
Cornell

Collaborators

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

Raul Astudillo
Caltech

Peter Frazier
Cornell

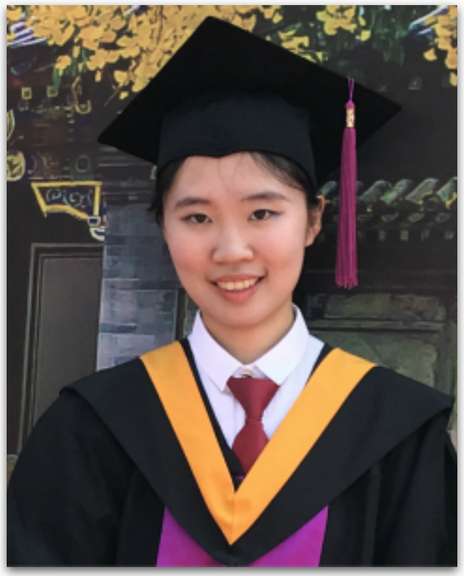

Qian Xie
Cornell

Alex Terenin
Cornell

Theo Brown
UCL/UKAE

work in progress

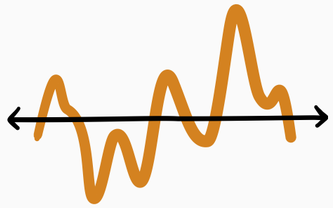

BayesOpt

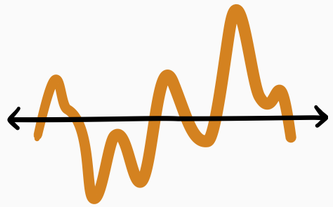

BayesOpt

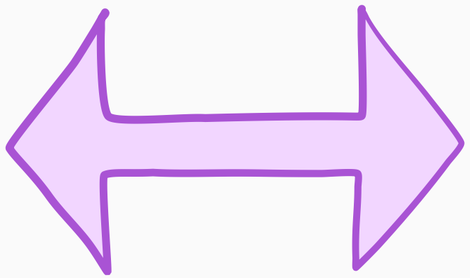

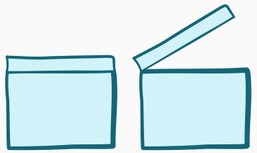

Pandora’s box

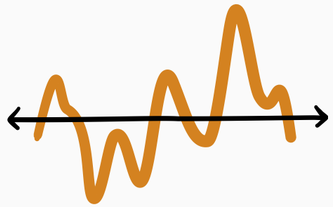

BayesOpt

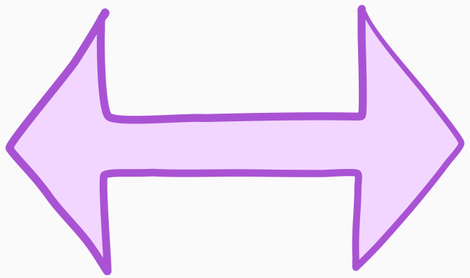

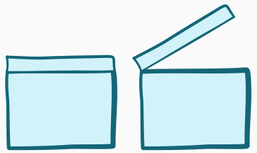

Pandora's box

Gittins is optimal

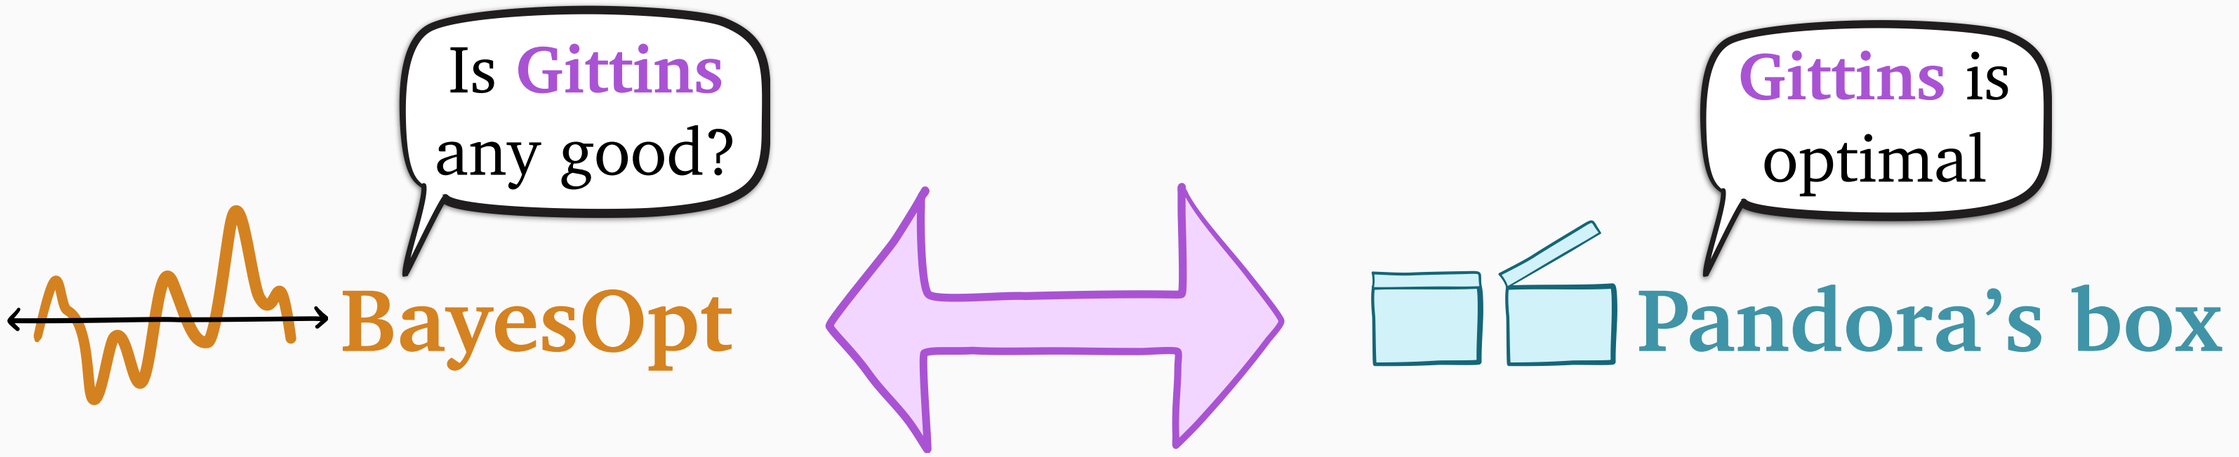
Is Gittins any good?
BayesOpt
Gittins is optimal
Pandora's box

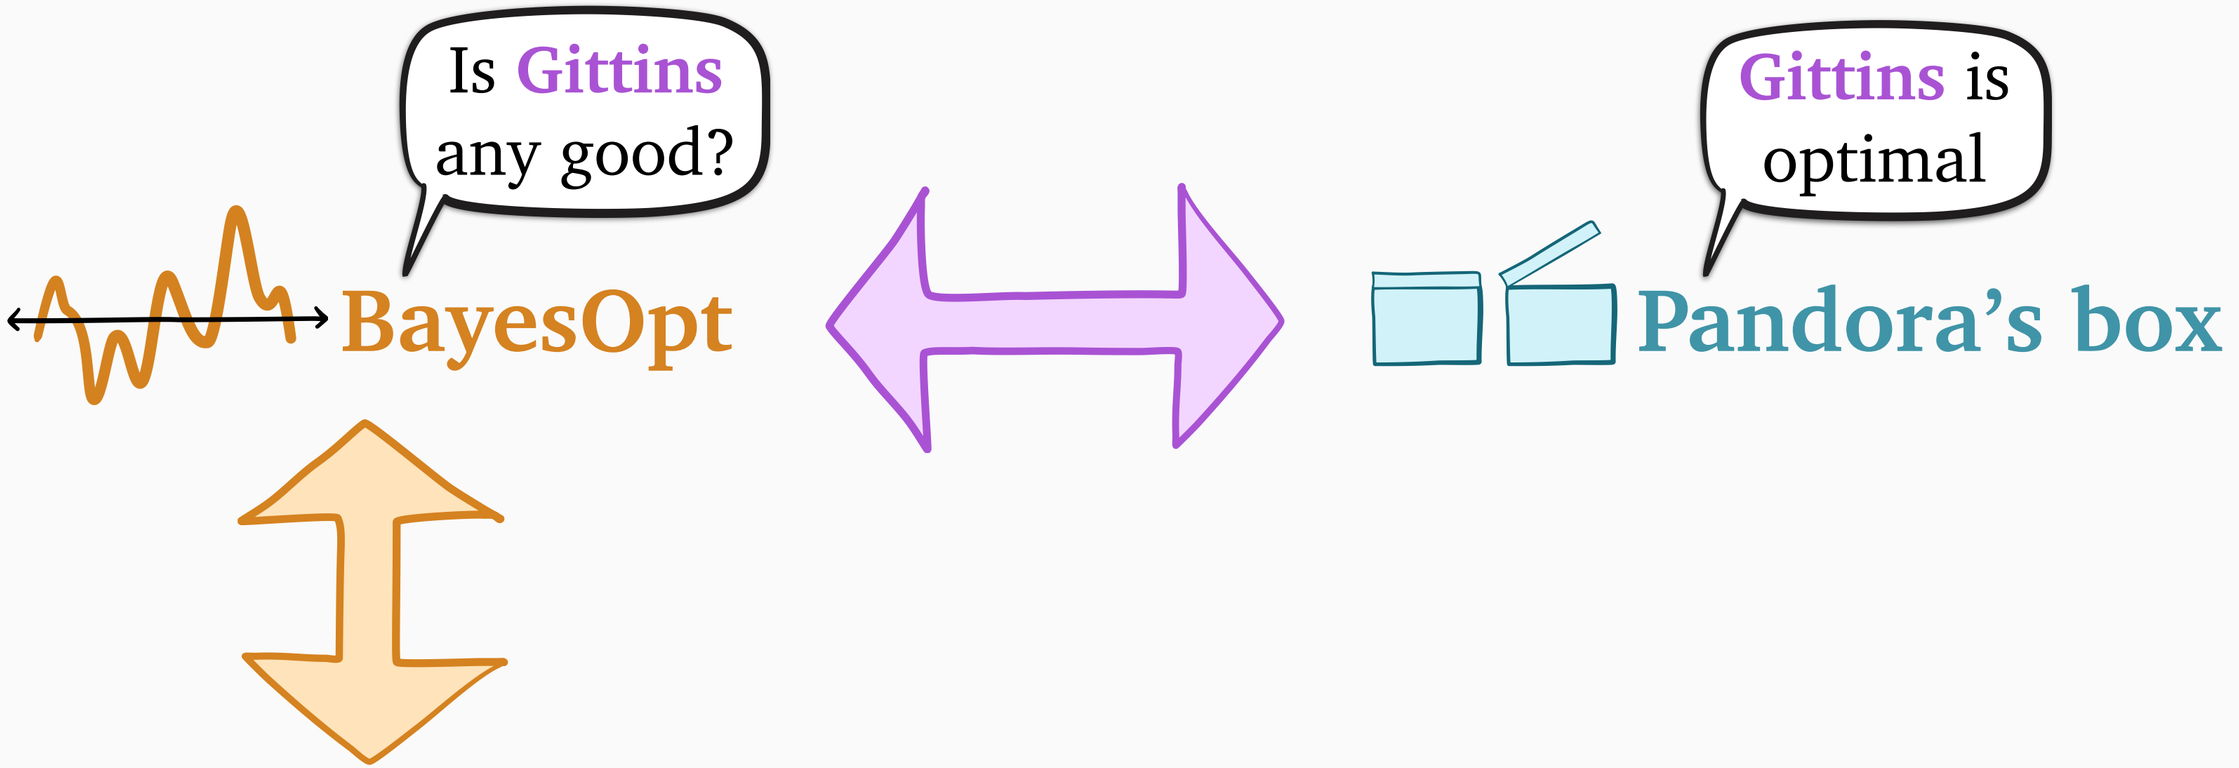
Is Gittins any good?
BayesOpt
Gittins is optimal
Pandora's box
"Fancy" BayesOpt

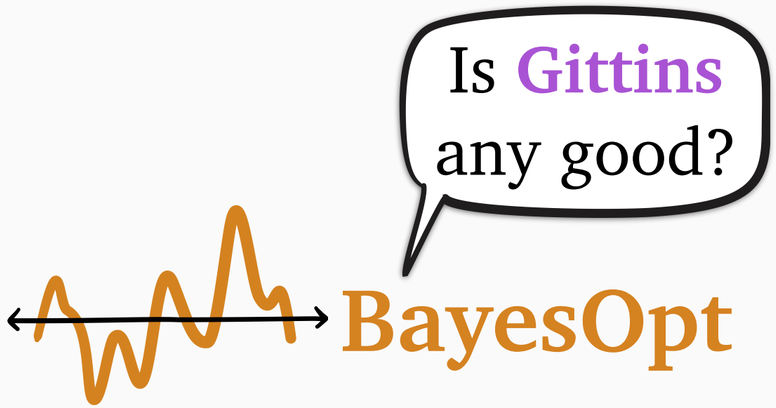

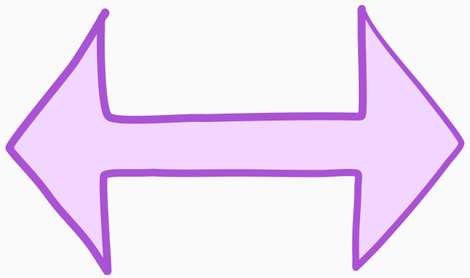

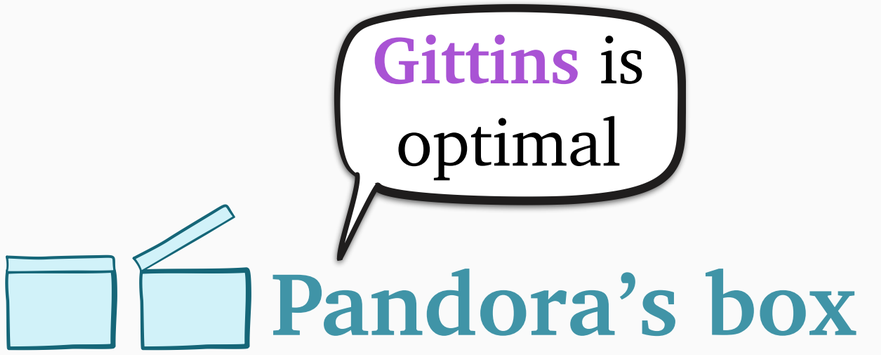

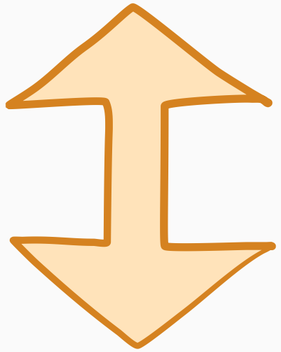

"Fancy" BayesOpt

- Many dimensions
- Heterogeneous costs
- Partial feedback
- Batch evaluations

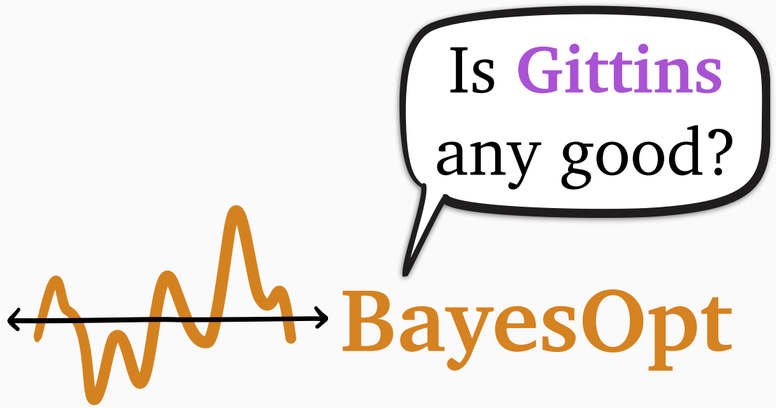

BayesOpt

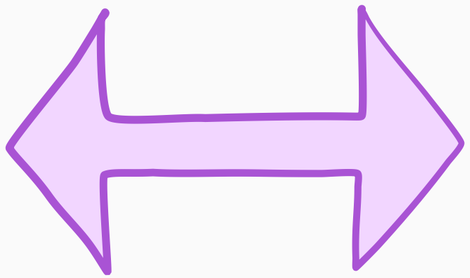

Pandora's box

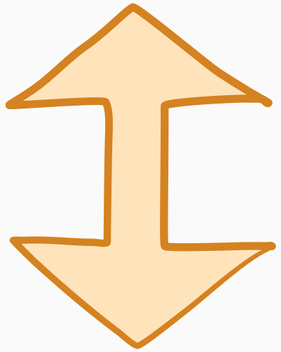

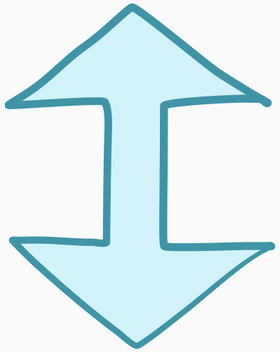

"Fancy" BayesOpt

- Many dimensions
- Heterogeneous costs
- Partial feedback
- Batch evaluations

Markovian MAB

(multi-armed bandit)

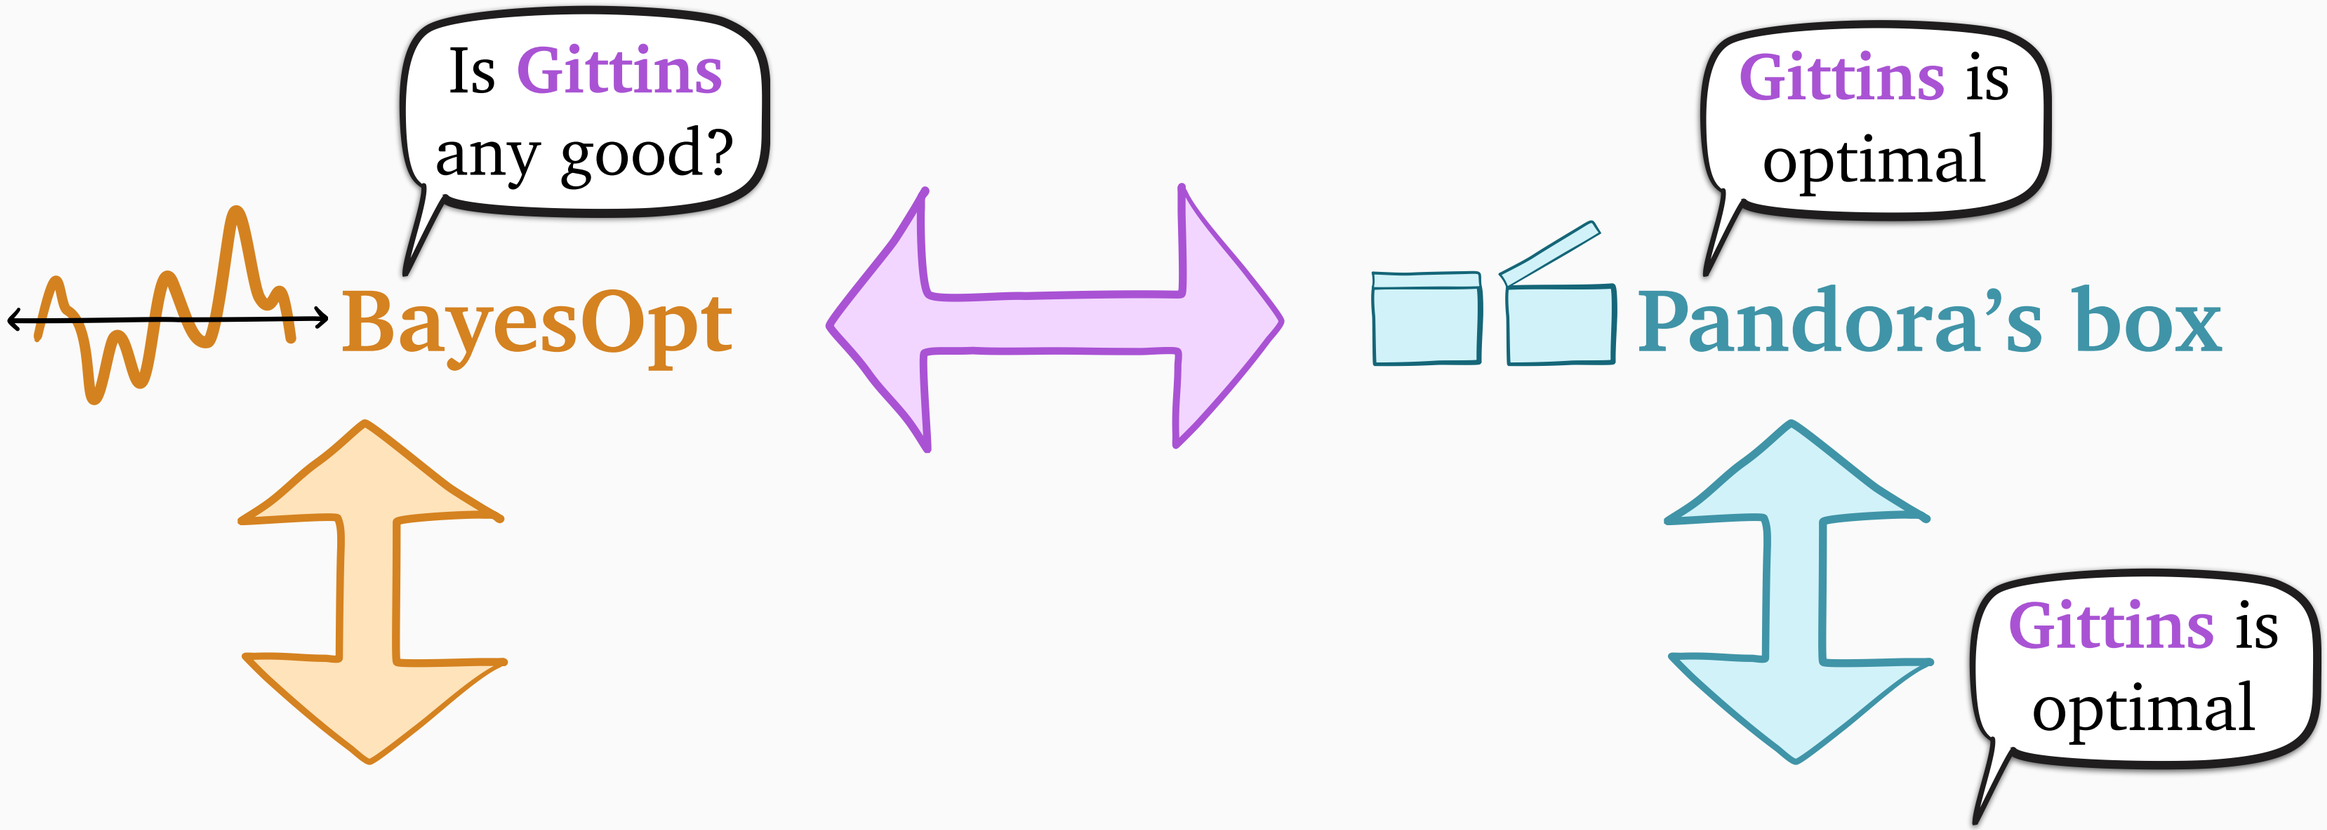

"Fancy" BayesOpt

- Many dimensions
- Heterogeneous costs
- Partial feedback
- Batch evaluations

Markovian MAB

(multi-armed bandit)

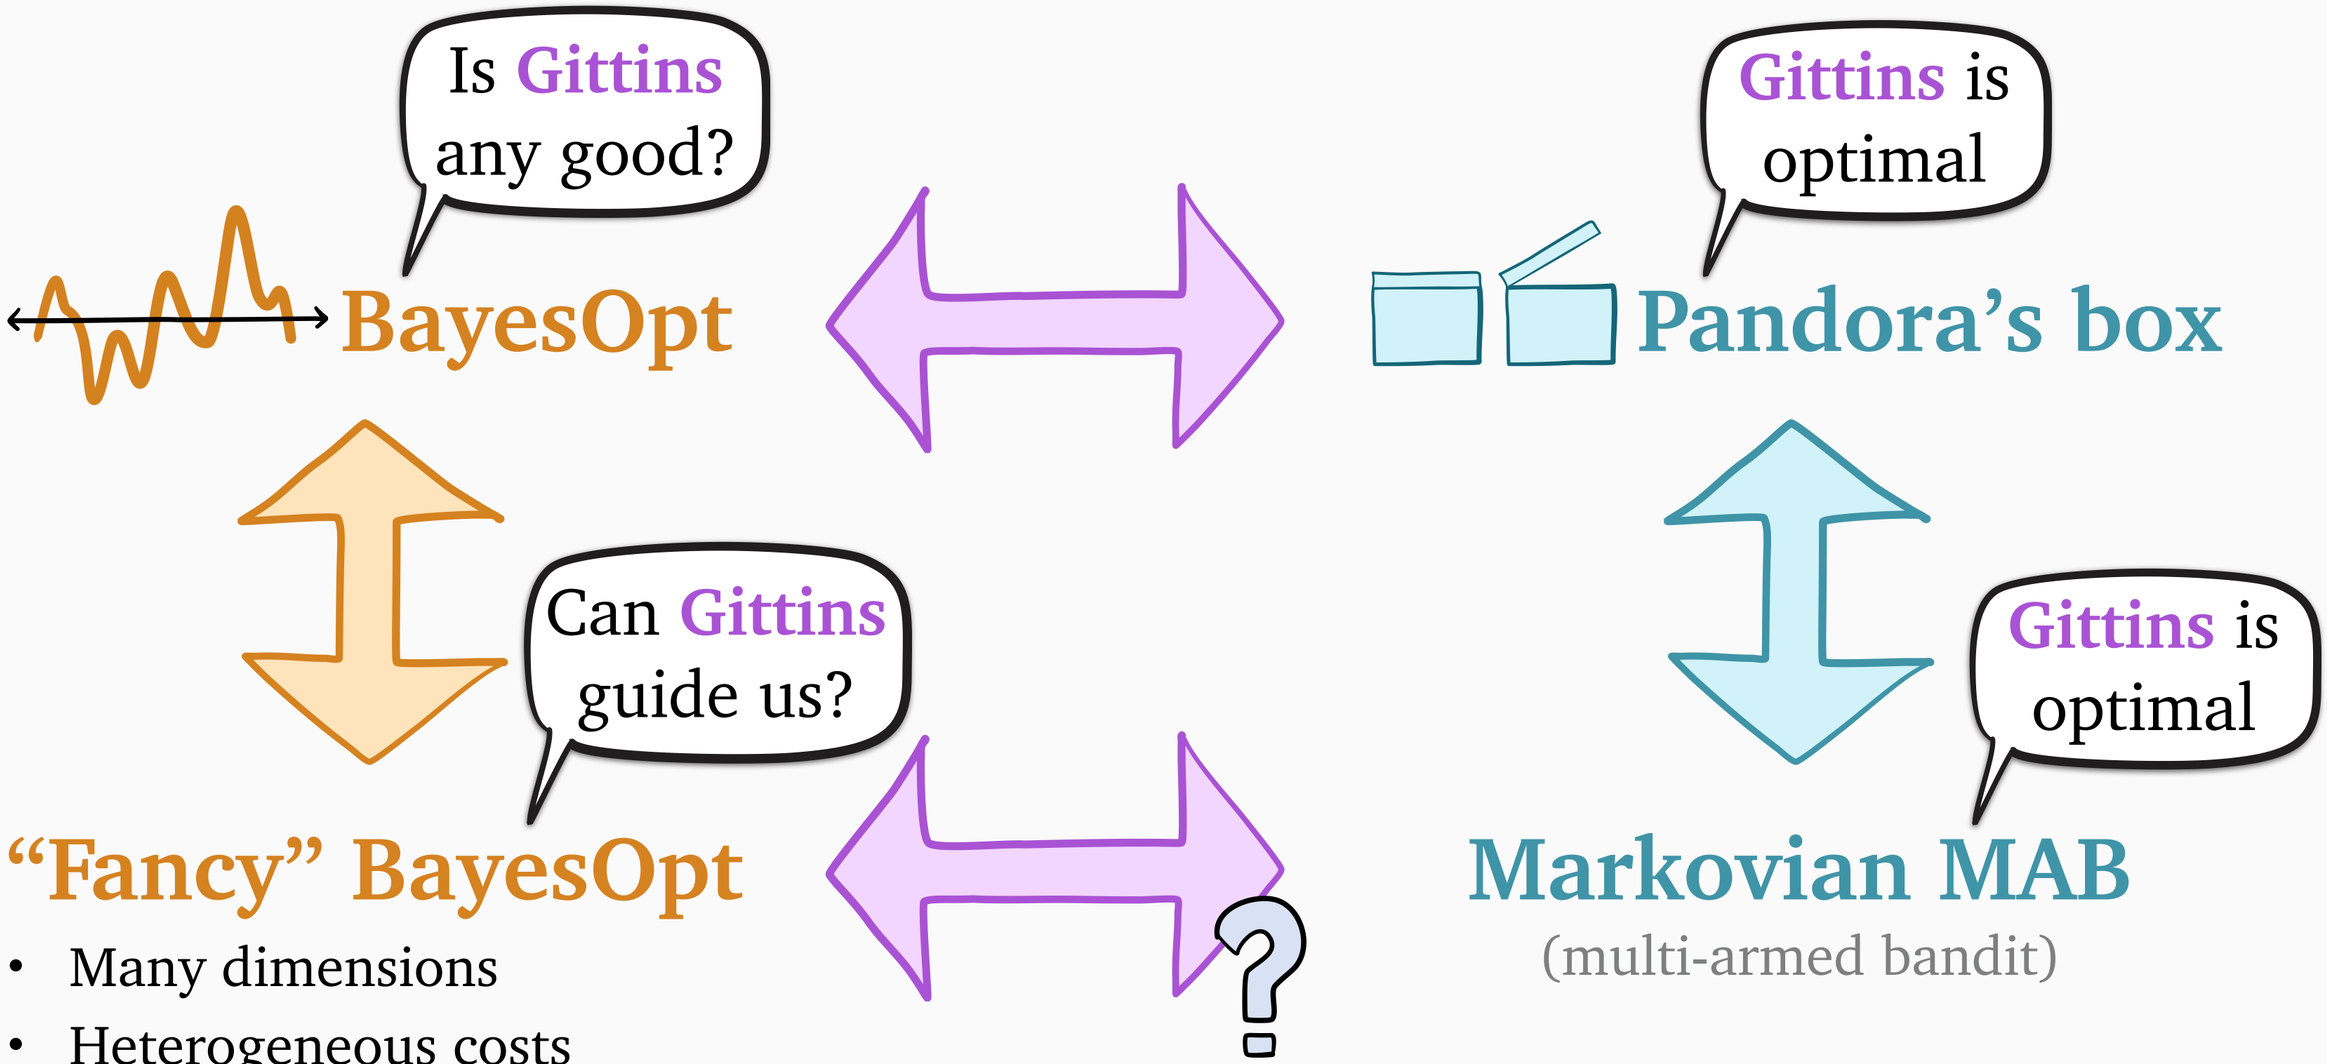

- Many dimensions
- Heterogeneous costs
- Partial feedback
- Batch evaluations

Is **Gittins** any good?

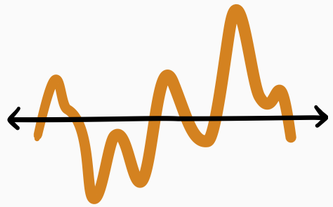

BayesOpt

Gittins is optimal

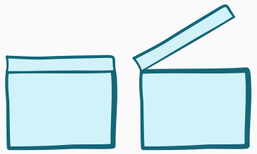

Pandora's box

Can **Gittins** guide us?

Gittins is optimal

"Fancy" BayesOpt

- Many dimensions
- Heterogeneous costs
- Partial feedback
- Batch evaluations

Markovian MAB

(multi-armed bandit)

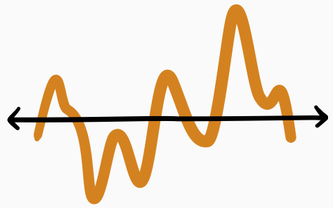

BayesOpt

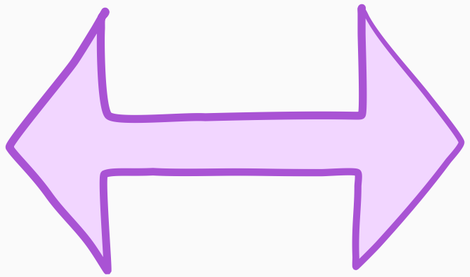

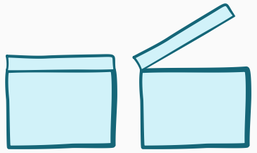

Pandora’s box

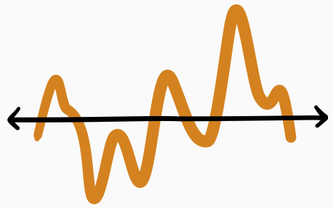

BayesOpt

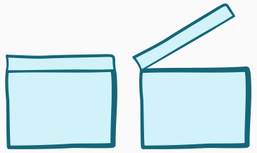

Pandora's box

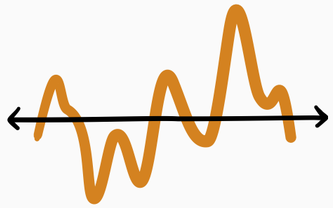

BayesOpt

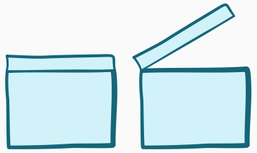

Pandora's box

Tutorial: What are **Pandora's box** and **Gittins**?

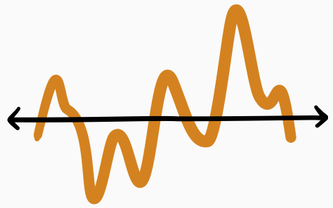

BayesOpt

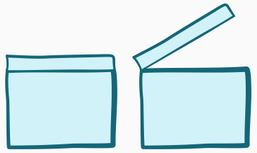

Pandora's box

Tutorial: What are **Pandora's box** and **Gittins**?

Results: Does **Gittins** work for **BayesOpt**?

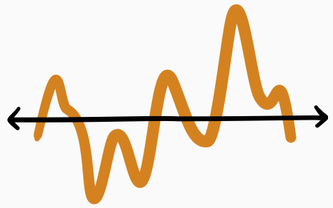

BayesOpt

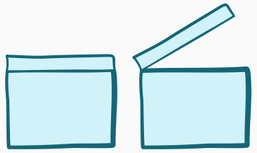

Pandora's box

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

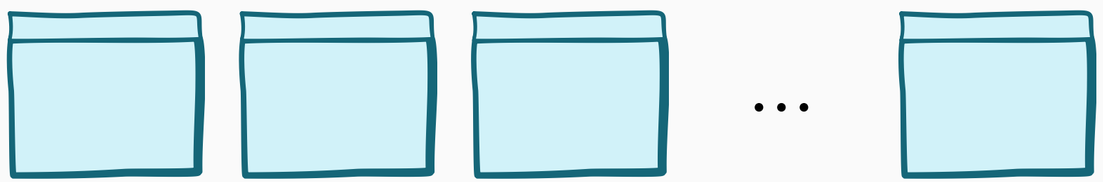

Pandora's box problem

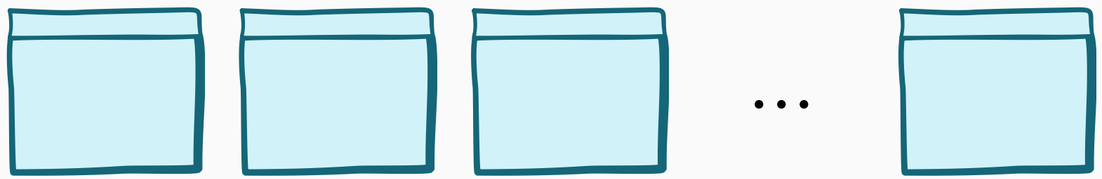

Each box:

- Opening cost c
- Hidden reward Y

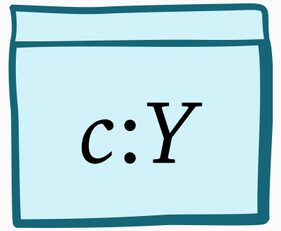

Pandora's box problem

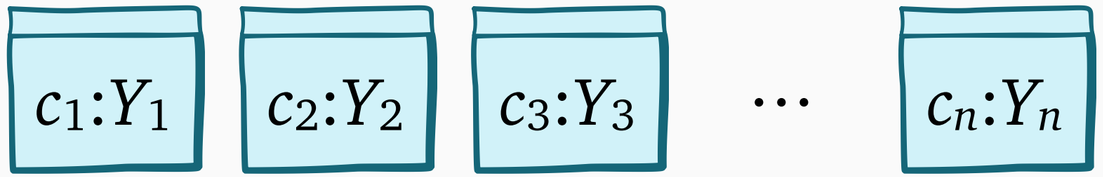

Each box:

- Opening cost c
- Hidden reward Y

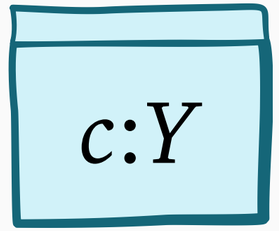

Pandora's box problem

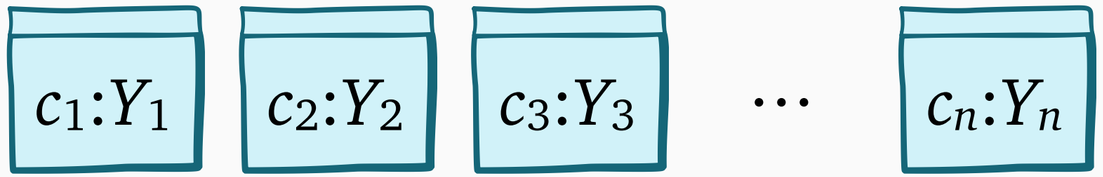

Each box:

- Opening cost c
- Hidden reward Y

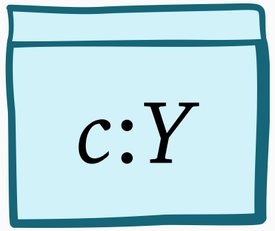

independent

Pandora's box problem

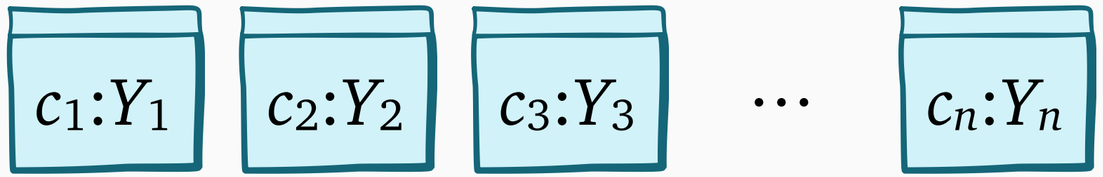

Each box:

- Opening cost c
- Hidden reward Y

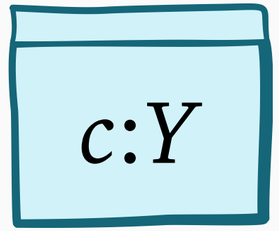

independent

Decision process:

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

Pandora's box problem

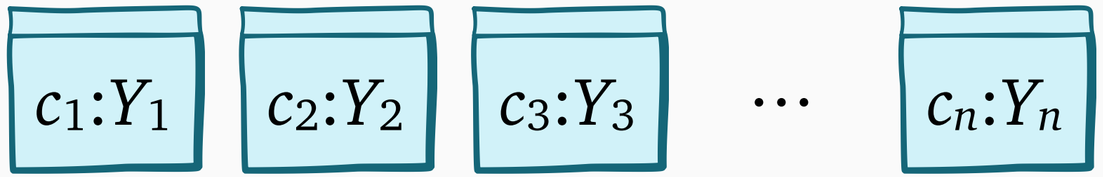

Each box:

- Opening cost c
- Hidden reward Y

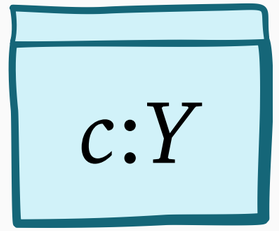

independent

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

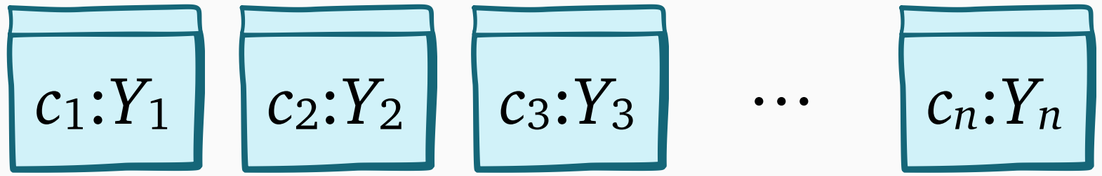

Each box:

- Opening cost c
- Hidden reward Y

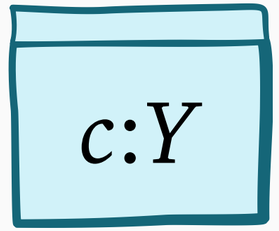

independent

Decision process:

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

Which box to open?

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

Pandora's box problem

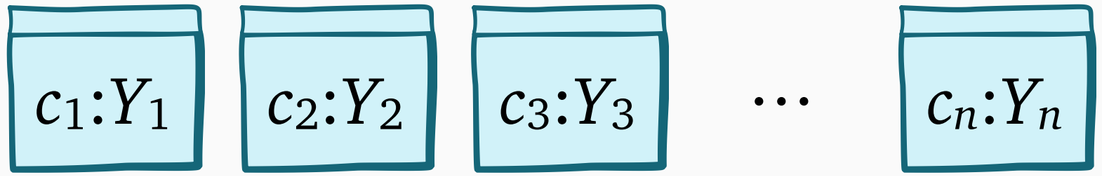

Each box:

- Opening cost c
- Hidden reward Y

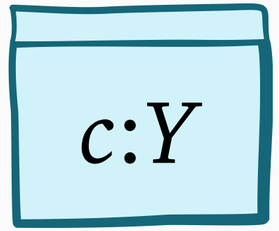

independent

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?

Is it time to stop?

Why is Pandora's box hard?

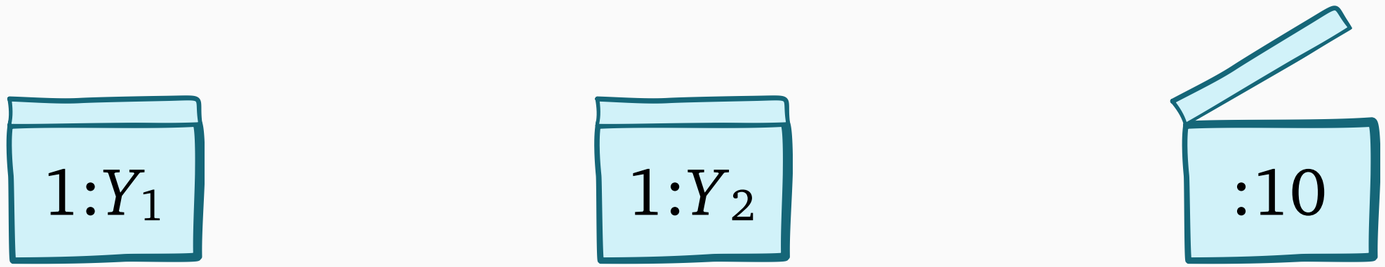

Why is Pandora's box hard?

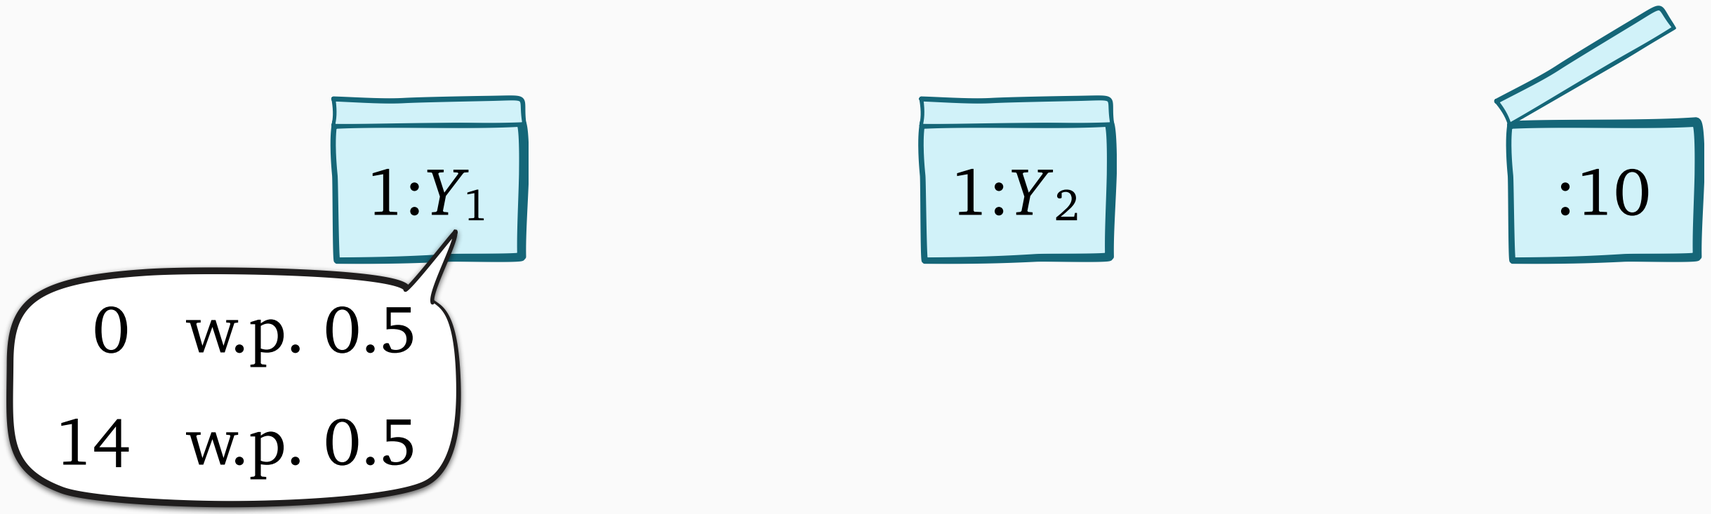

Why is Pandora's box hard?

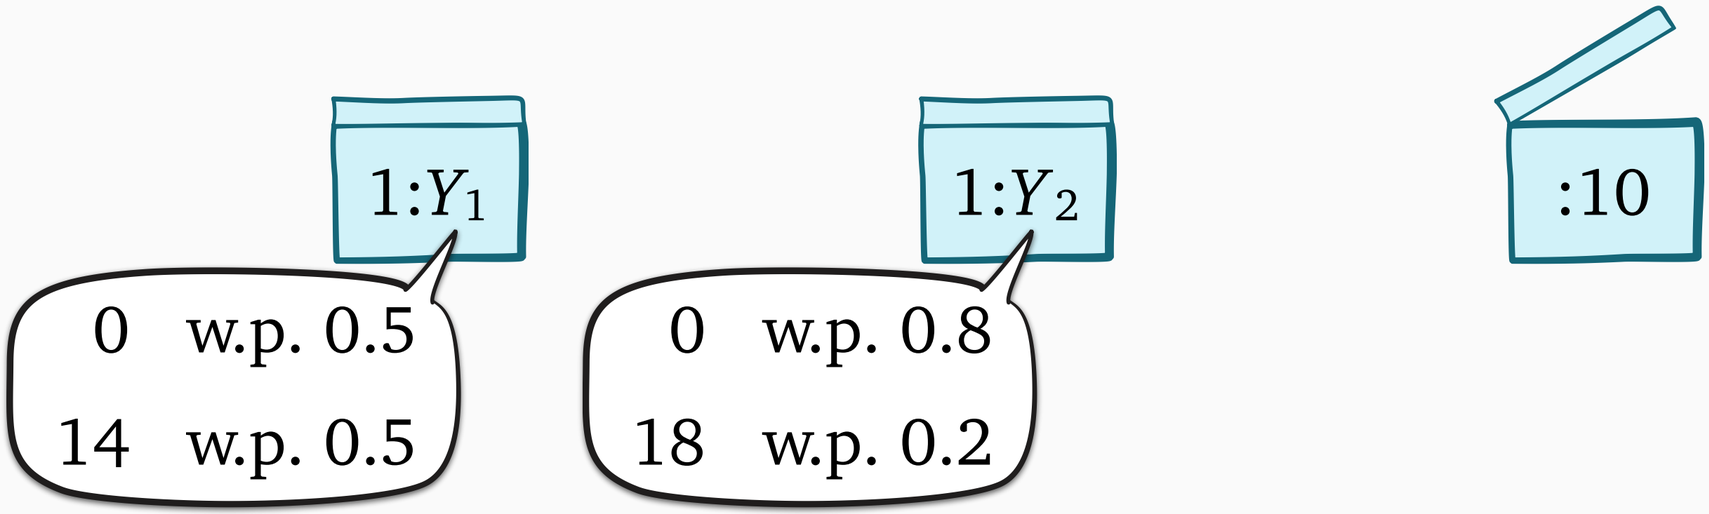

Why is Pandora's box hard?

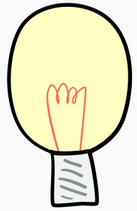

Expected improvement of X over r:

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

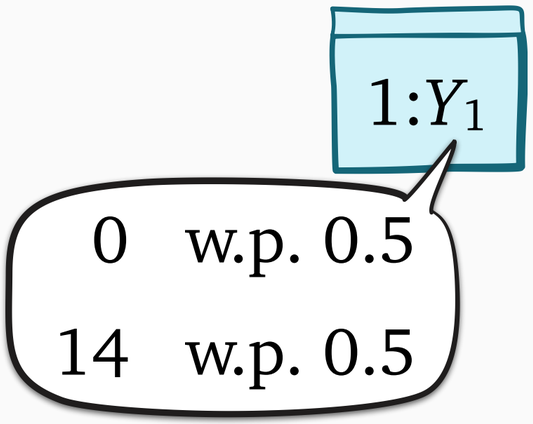

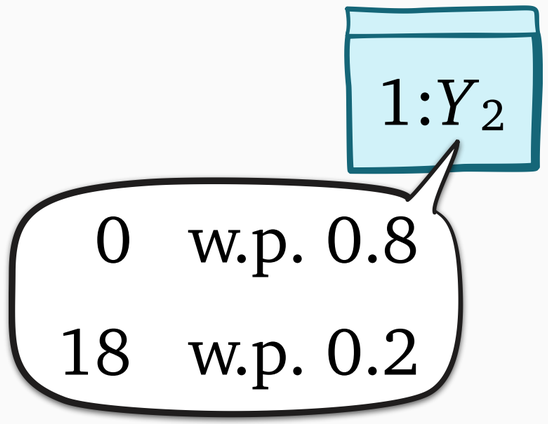

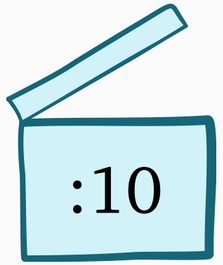

Why is Pandora's box hard?

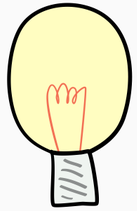

Expected improvement of X over r:

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

for now, think r = best so far

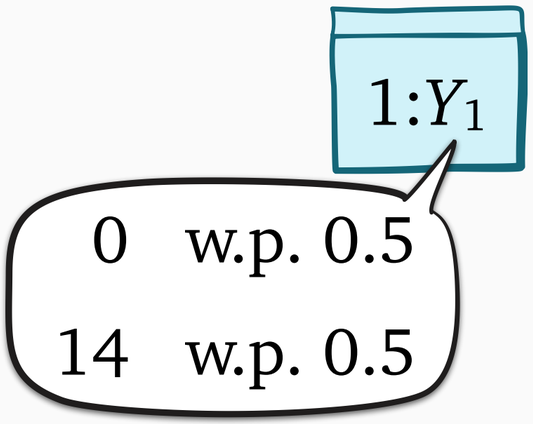

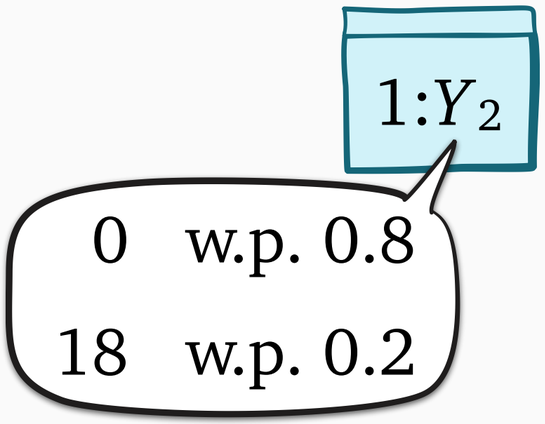

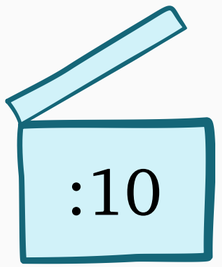

Why is Pandora's box hard?

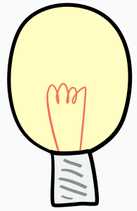

Expected improvement of X over r:

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

for now, think r = best so far

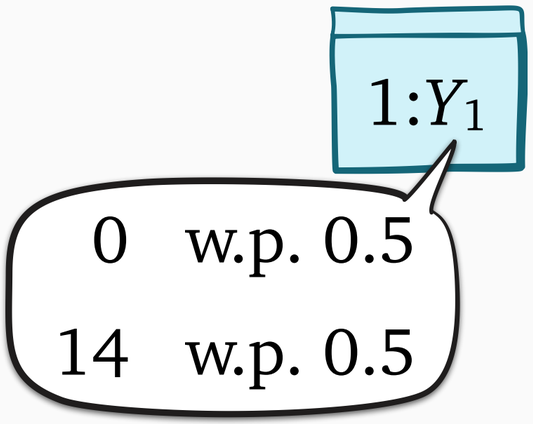

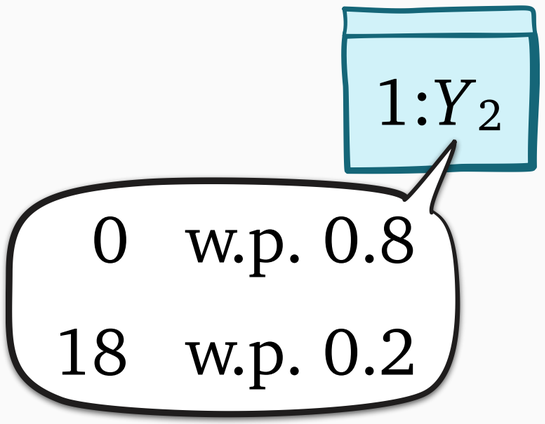

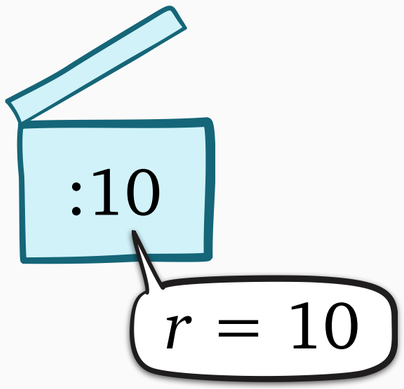

Why is Pandora's box hard?

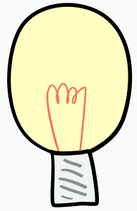

Expected improvement of X over r:

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

for now, think r = best so far

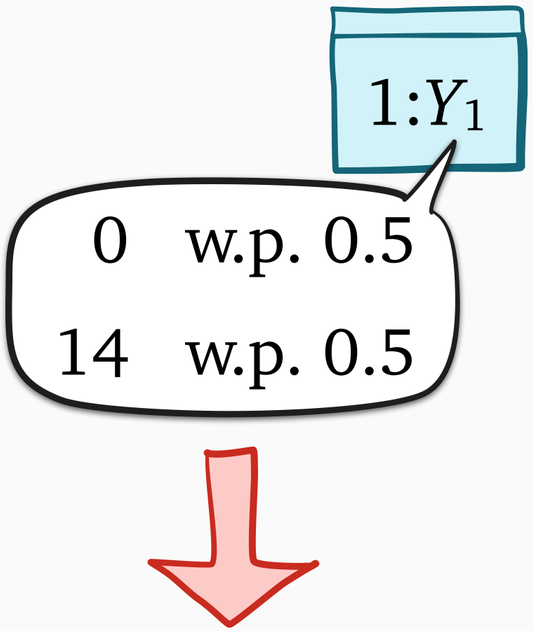

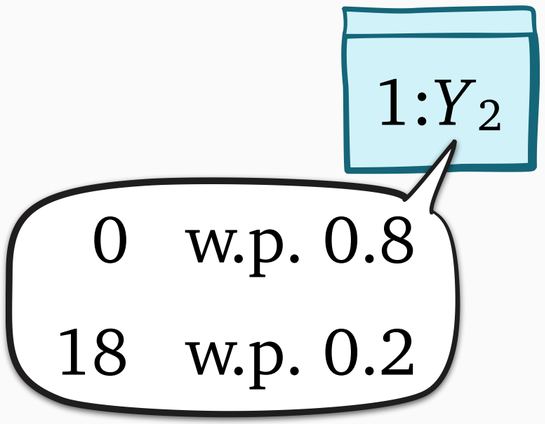

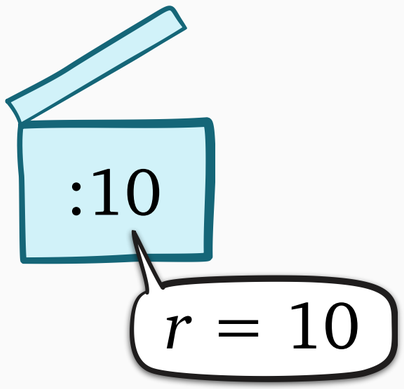

EI(X_1, 10) = 2

Why is Pandora's box hard?

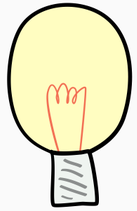

Expected improvement of X over r:

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

for now, think r = best so far

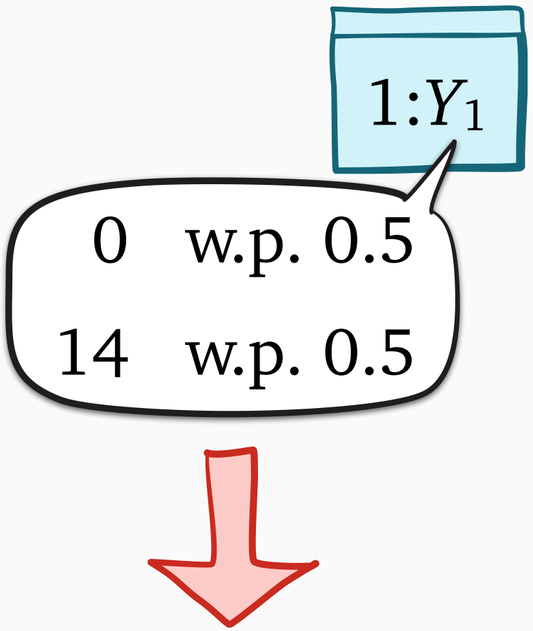

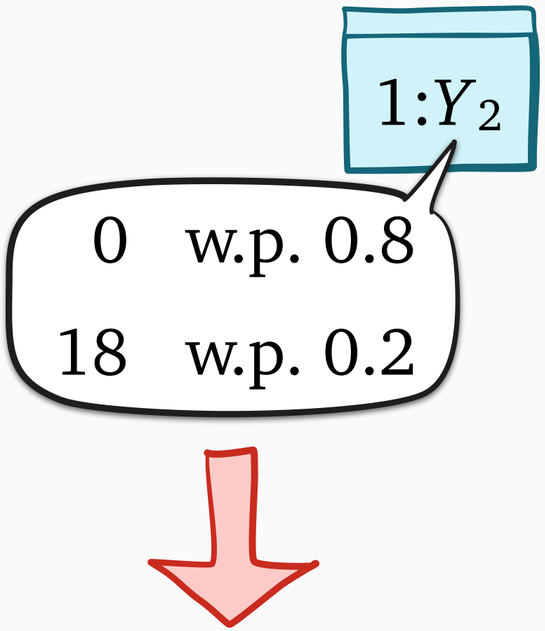

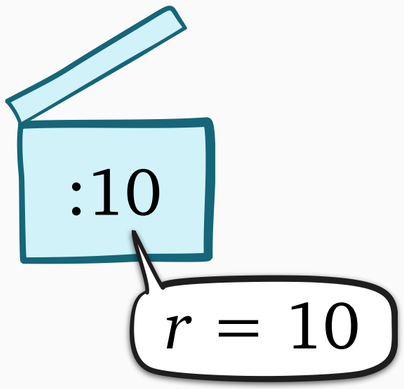

EI(X_1, 10) = 2

EI(X_2, 10) = 1.6

Why is Pandora's box hard?

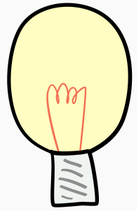

Expected improvement of X over r:

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

for now, think r = best so far

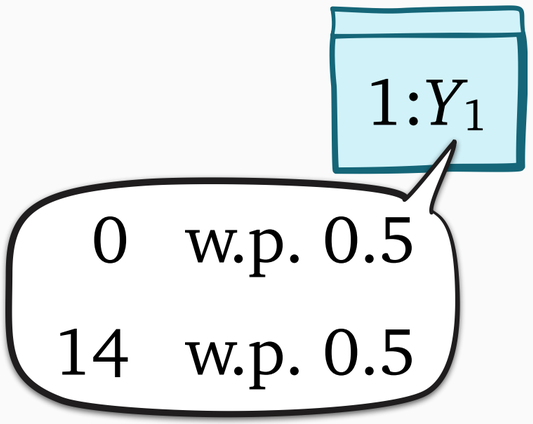

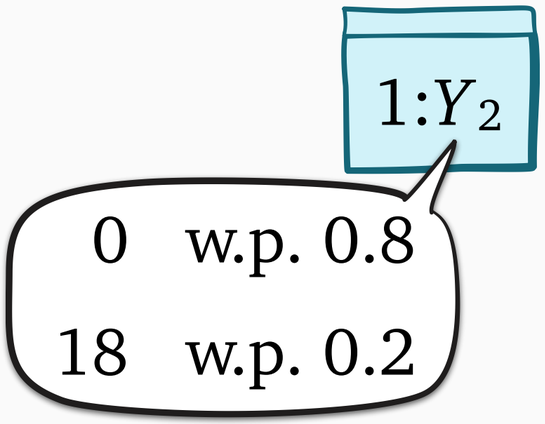

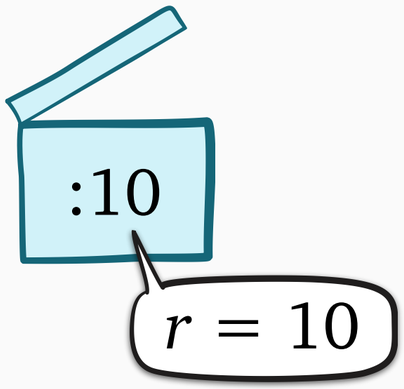

EI$(X_1, 10) = 2$

EI$(X_2, 10) = 1.6$

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

Why is Pandora's box hard?

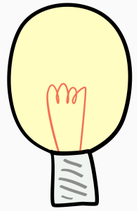

Expected improvement of X over r:

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

for now, think r = best so far

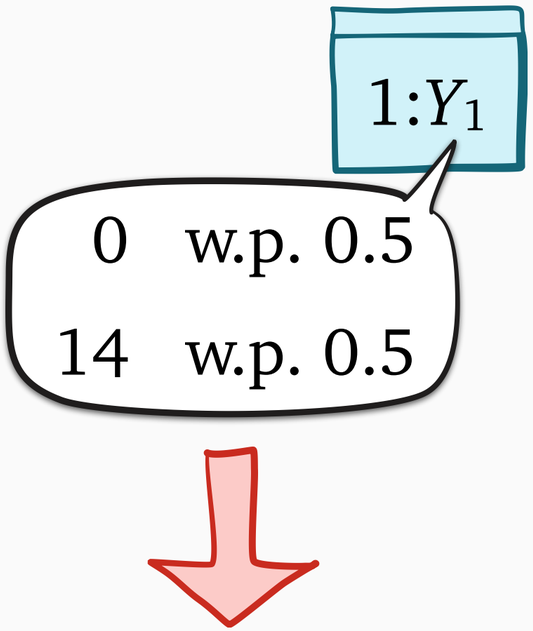

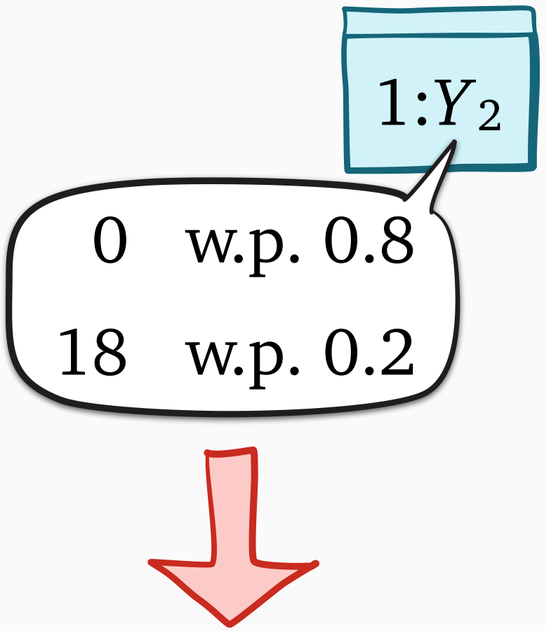

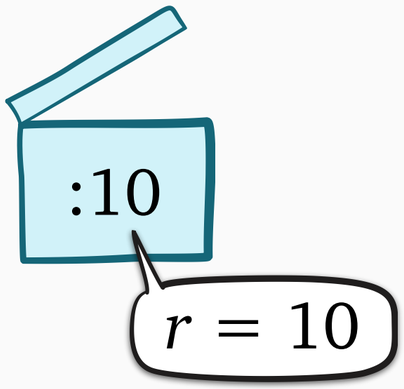

EI$(X_1, 10) = 2$

EI$(X_2, 10) = 1.6$

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

Why is Pandora's box hard?

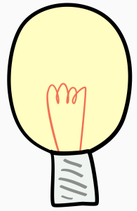

Expected improvement of X over r:

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

for now, think r = best so far

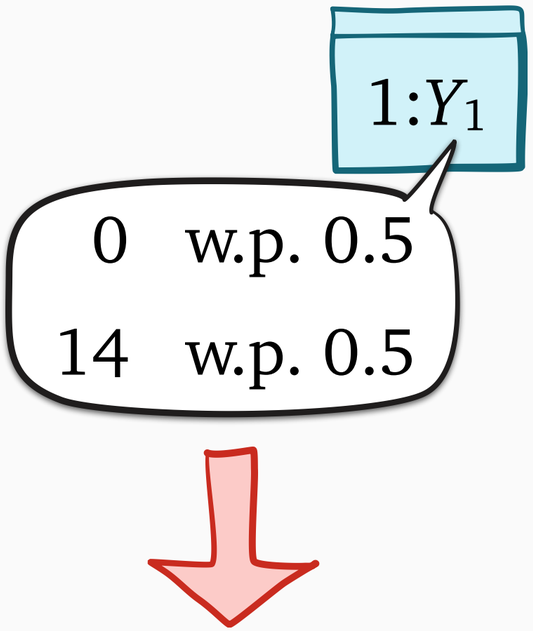

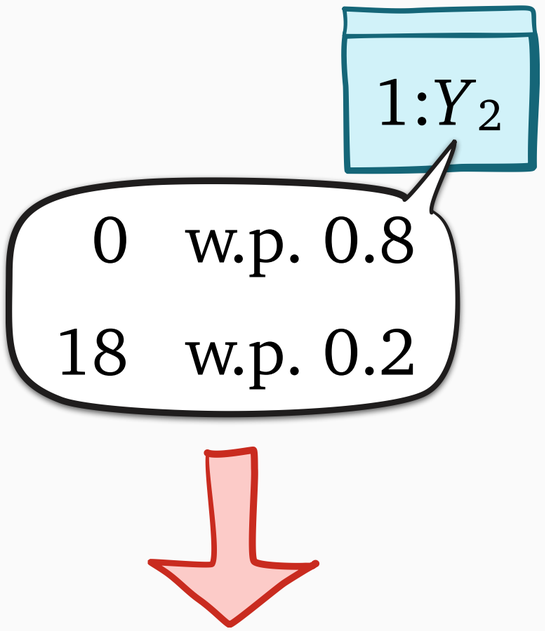

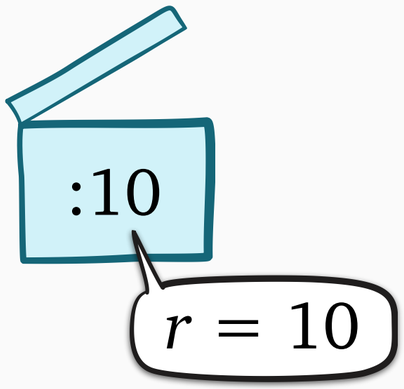

EI$(X_1, 10) = 2$

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?

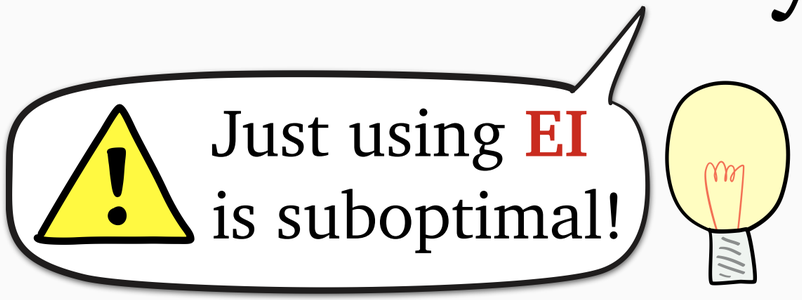

Expected improvement of X over r:

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

for now, think
r = best so far

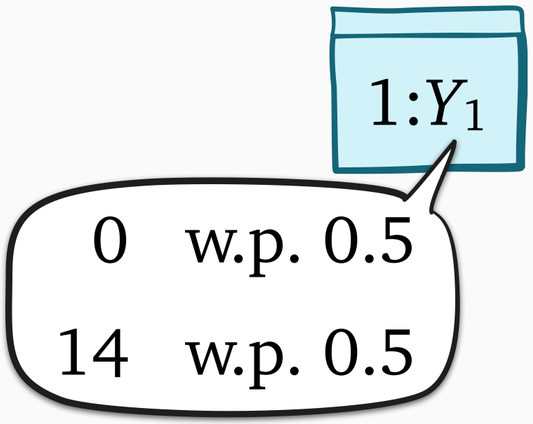

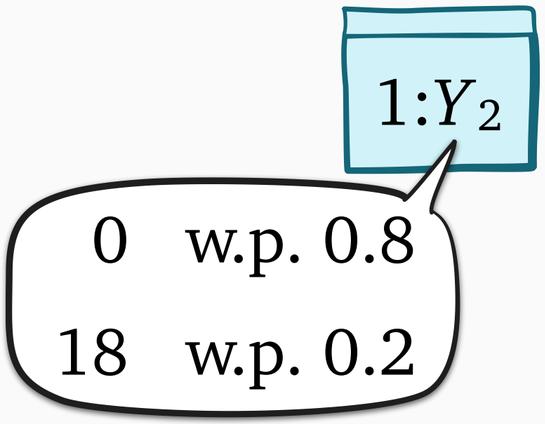

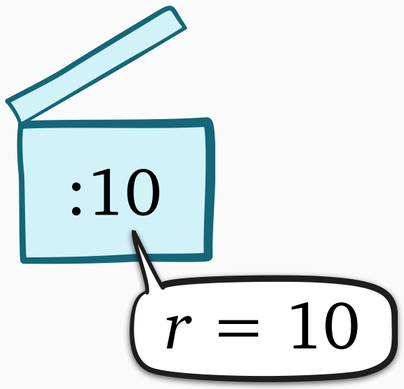

EI$(X_1, 10) = 2$

EI$(X_2, 10) = 1.6$

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

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

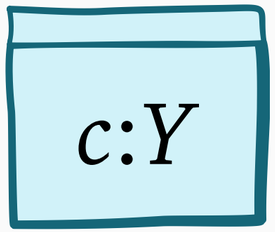

Step 2: *act* on box of best rating

Optimal policy: **Gittins**

Step 1: *rate* each box separately

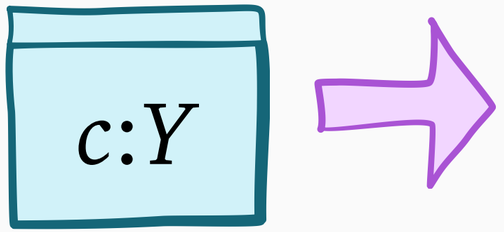

Gittins index:

$g(c{:}Y)$

Step 2: *act* on box of best rating

Optimal policy: Gittins

Step 1: *rate* each box separately

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

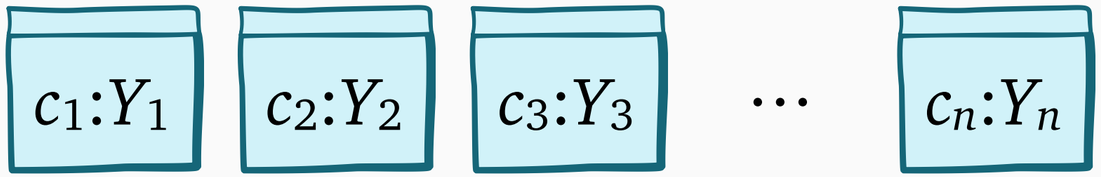

Optimal policy: Gittins

Step 1: *rate* each box separately

Step 2: *act* on box of best rating

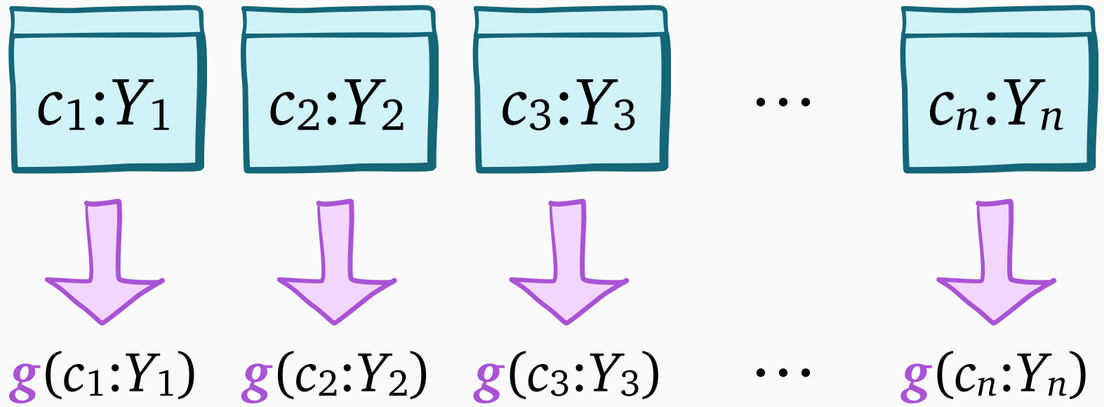

Optimal policy: Gittins

Step 1: *rate* each box separately

Step 2: *act* on box of best rating

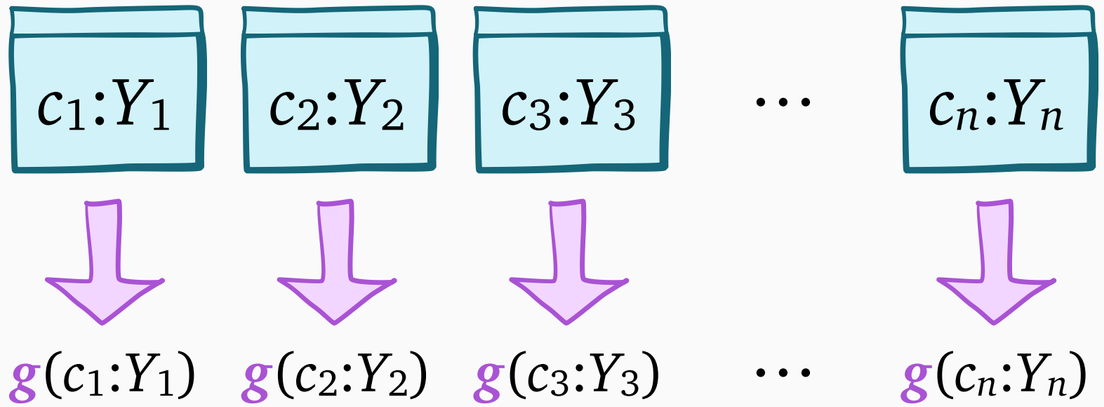

Gittins policy: if box of max Gittins index is…

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

Optimal policy: Gittins

Step 1: *rate* each box separately

Step 2: *act* on box of best rating

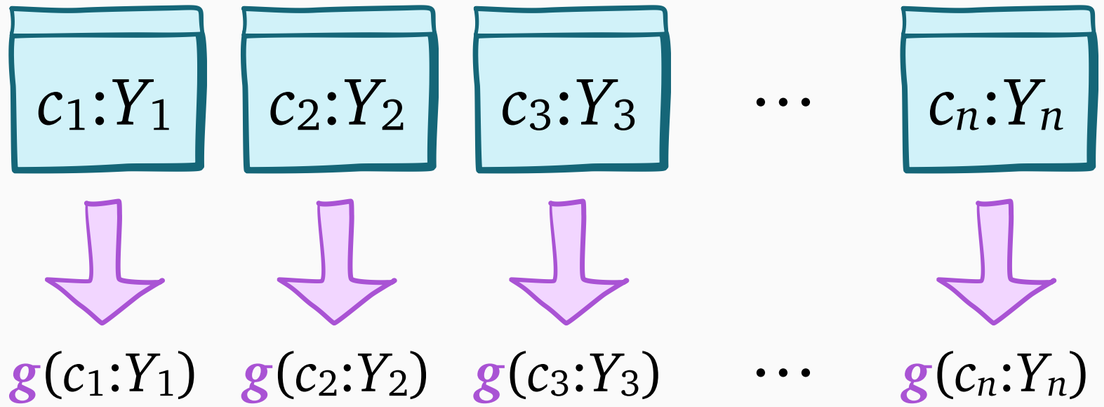

Gittins **policy:** if box of max Gittins index is…

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

} *act* on it

Optimal policy: Gittins

Step 1: *rate* each box separately

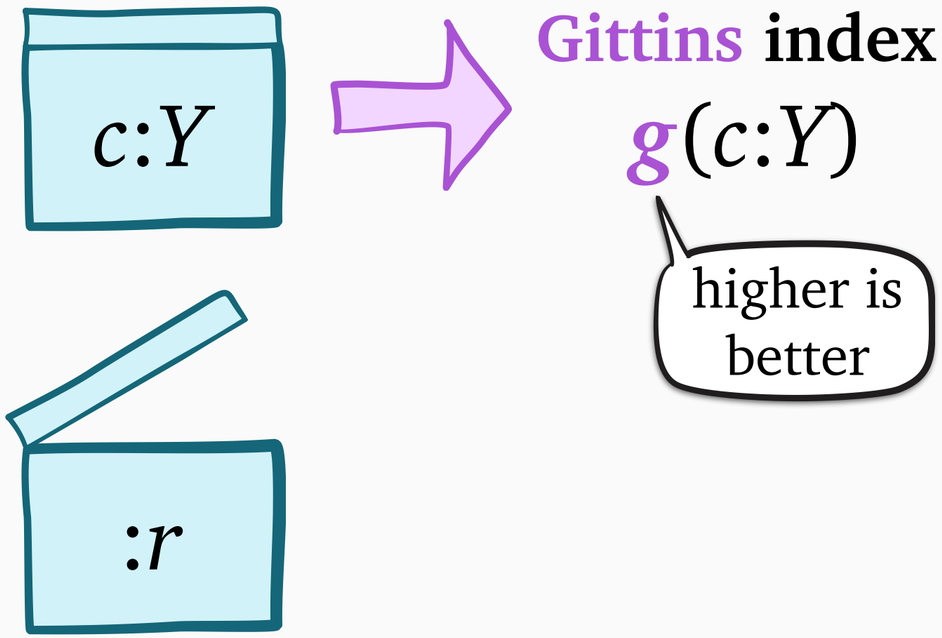

Step 2: *act* on box of best rating

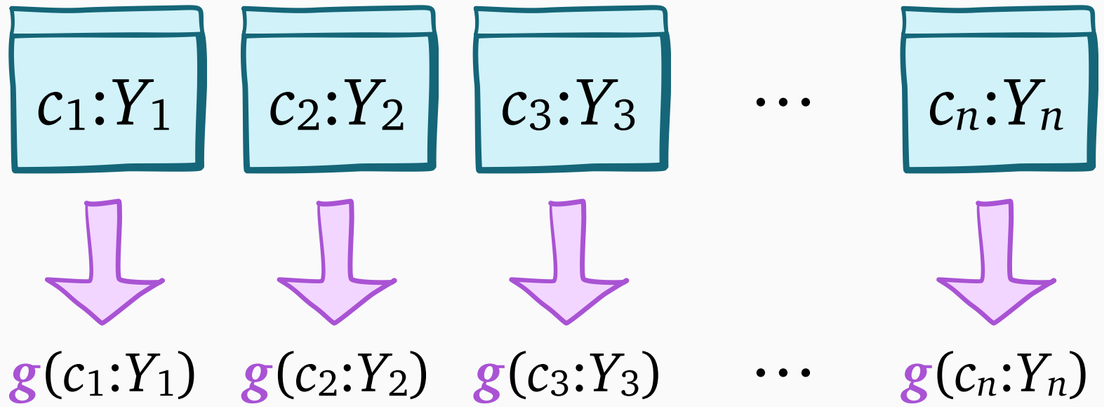

Gittins policy: if box of max Gittins index is…

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

} *act* on it

Optimal policy: Gittins

Step 1: *rate* each box separately

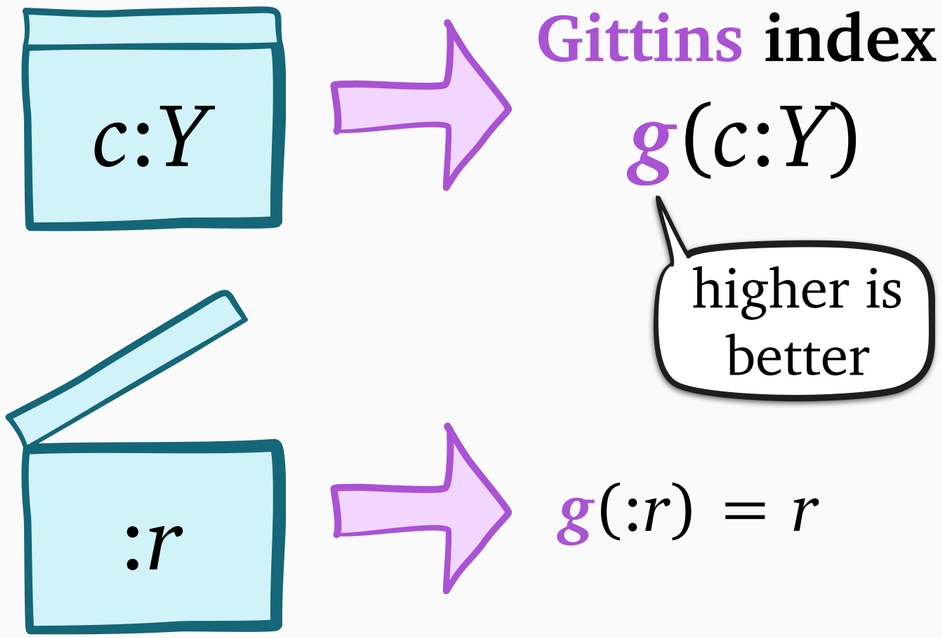

Step 2: *act* on box of best rating

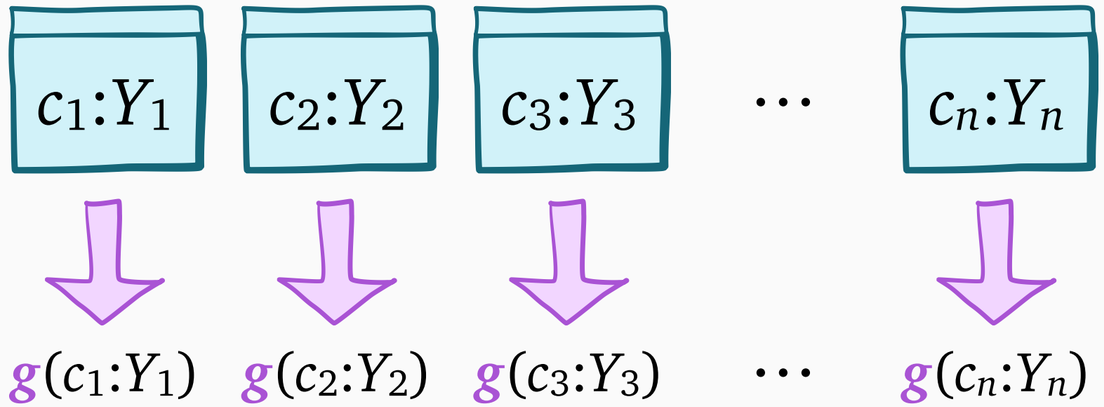

Gittins **policy:** if box of max Gittins index is…

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

} *act* on it

Optimal policy: Gittins

Step 1: *rate* each box separately

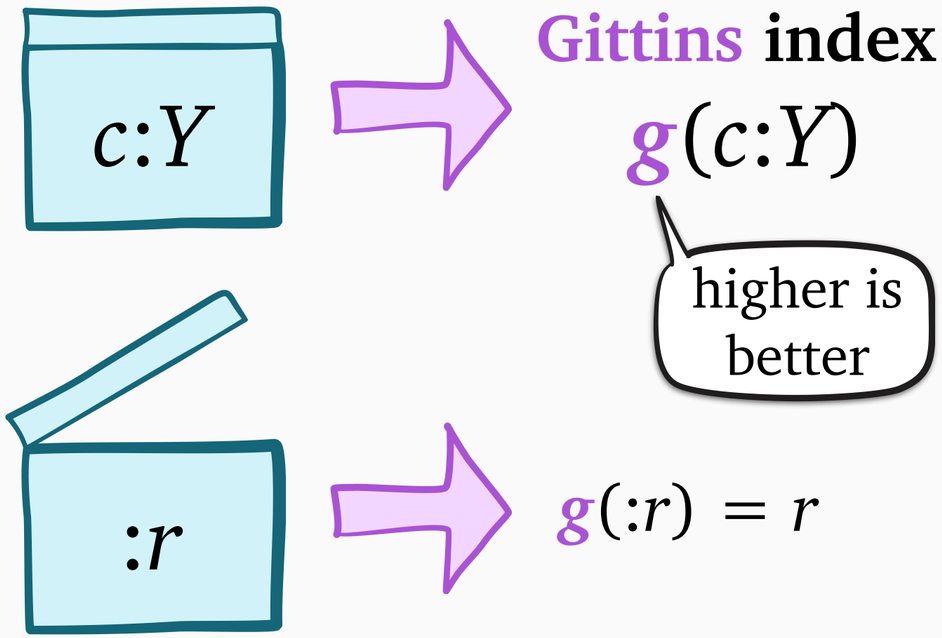

Theorem: [Weitzman, 1979]
the Gittins policy is optimal

Step 2: *act* on box of best rating

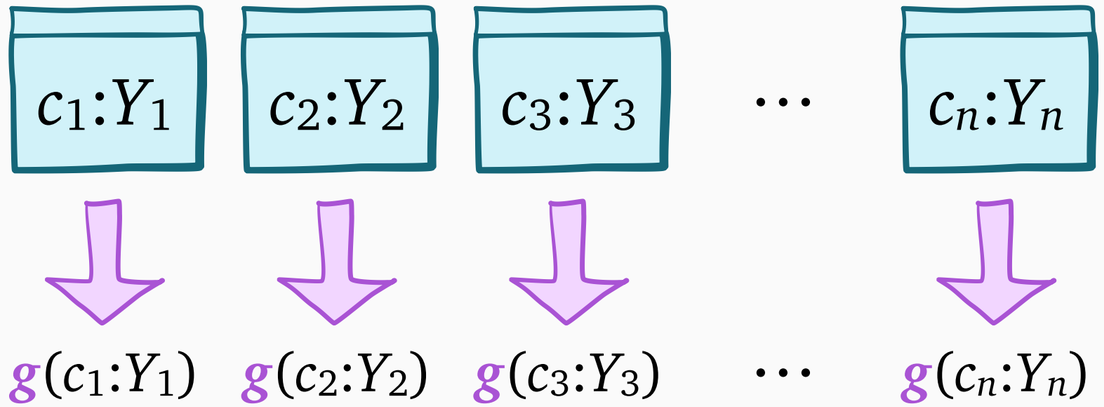

Gittins **policy:** if box of max Gittins index is…

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

} *act* on it

Defining the **Gittins** index

Defining the **Gittins** index

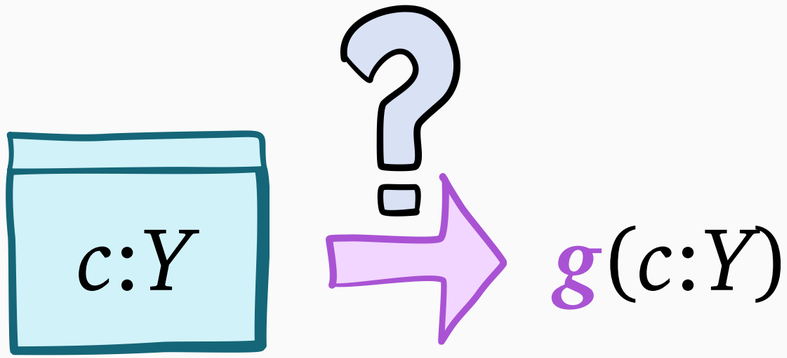

Defining the **Gittins** index

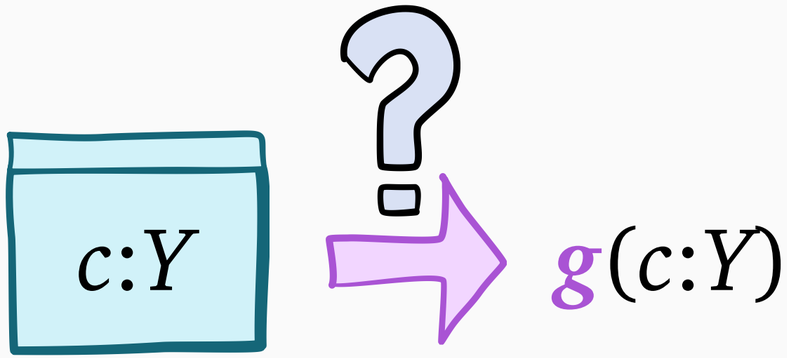

$g(c{:}Y)$

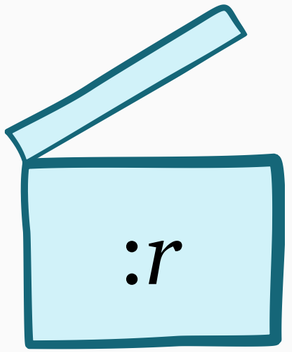

$g(:r) = r$

Defining the **Gittins** index

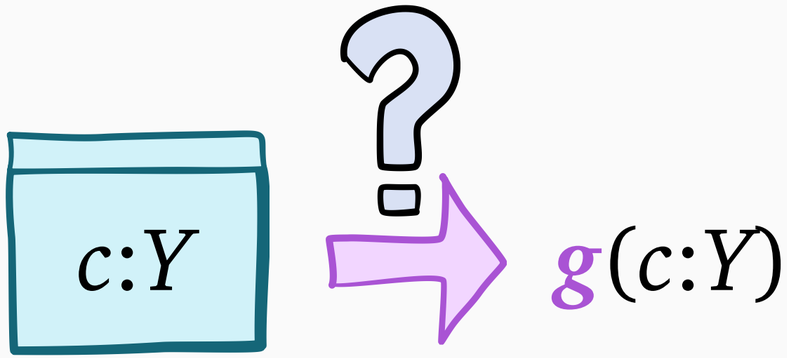

$g(c{:}Y)$

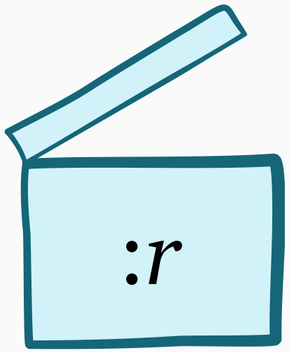

$g(:r) = r$

1.5-box problem

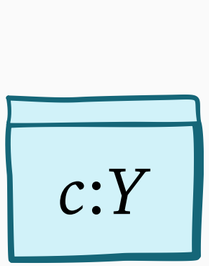

vs.

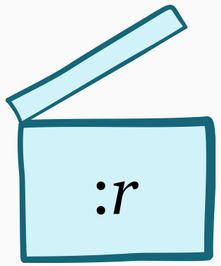

Defining the **Gittins** index

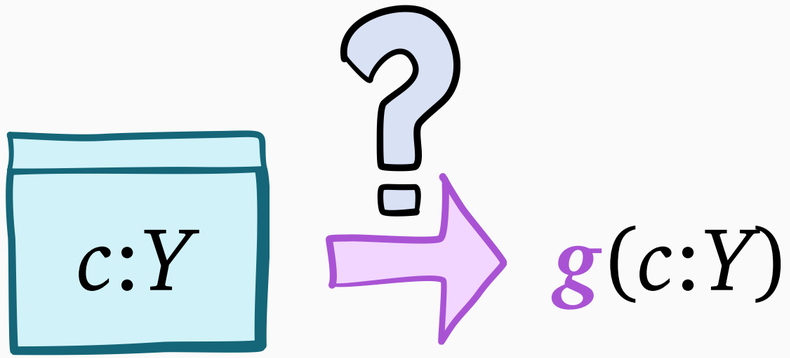

$g(c{:}Y)$

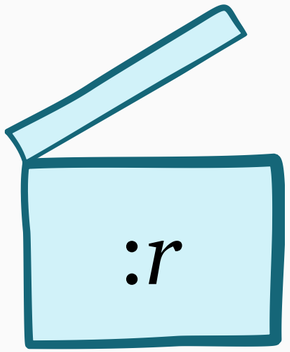

$g(:r) = r$

1.5-box problem

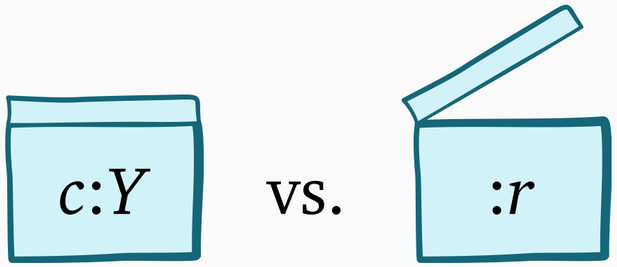

Key question: what to do in 1.5-box problem?

Defining the **Gittins** index

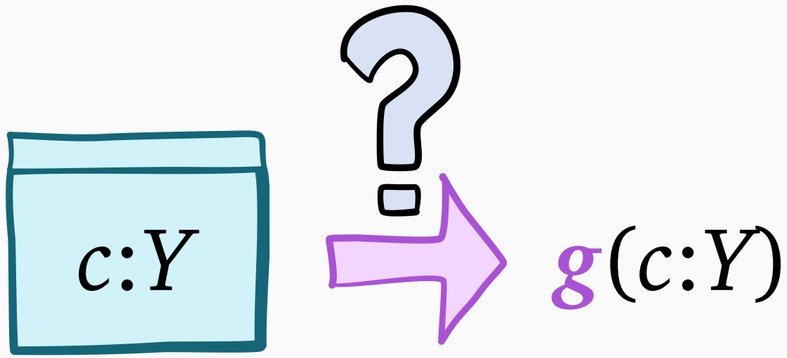

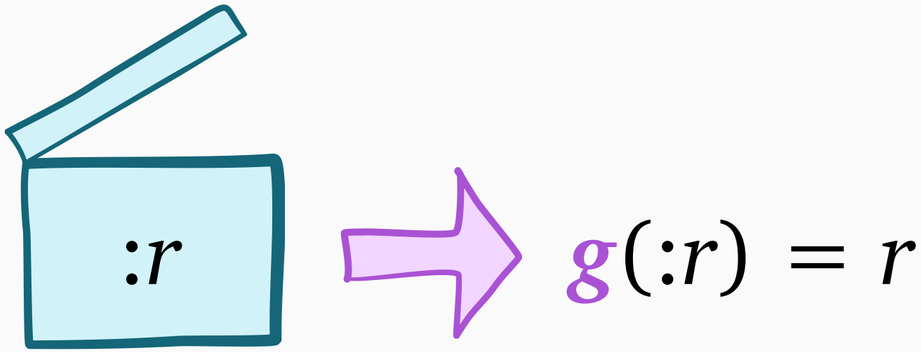

1.5-box problem

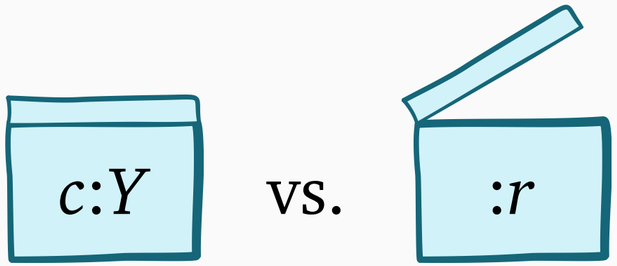

Key question: what to do in 1.5-box problem?

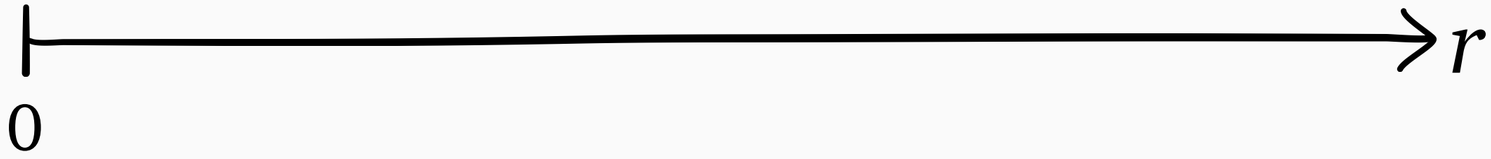

Defining the **Gittins** index

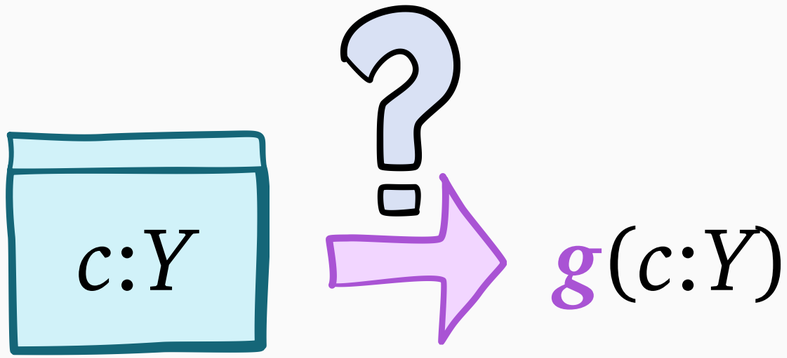

$g(c{:}Y)$

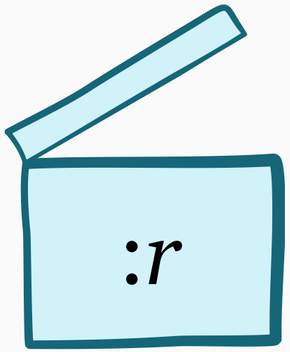

$g(:r) = r$

1.5-box problem

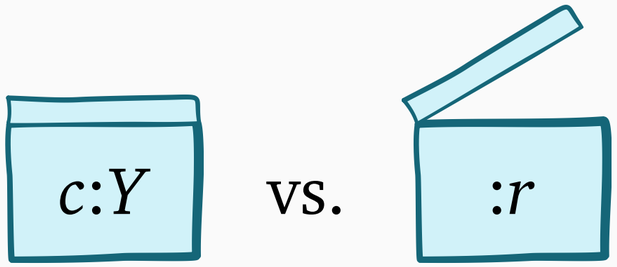

Key question: what to do in 1.5-box problem?

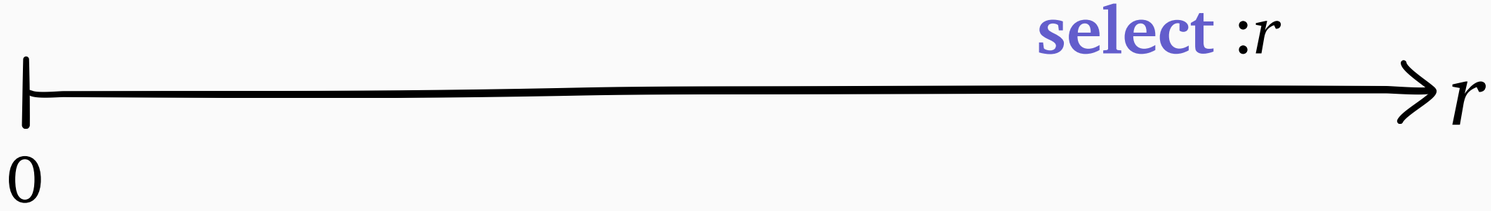

Defining the **Gittins** index

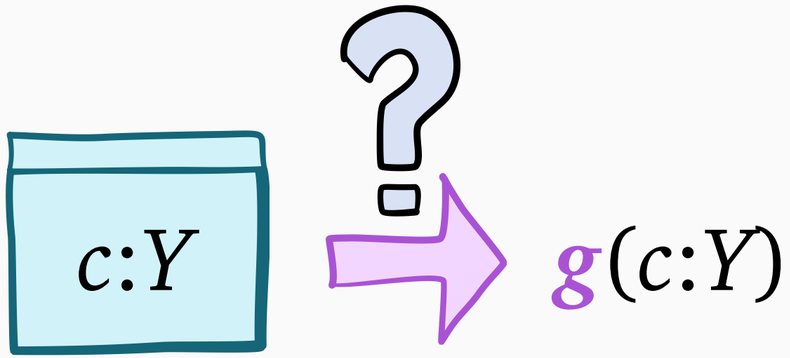

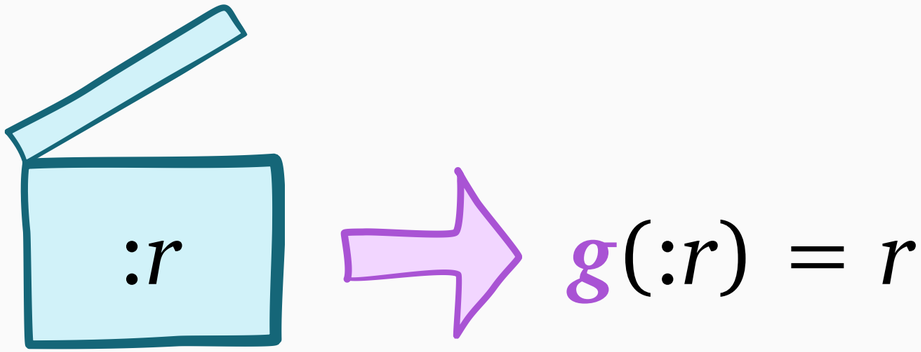

1.5-box problem

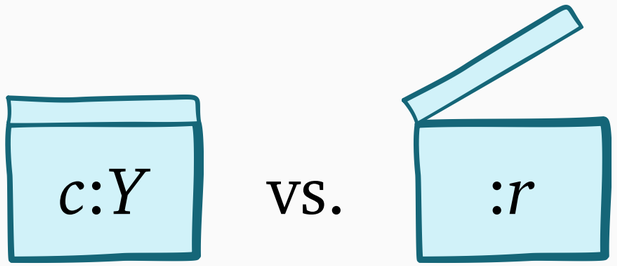

Key question: what to do in 1.5-box problem?

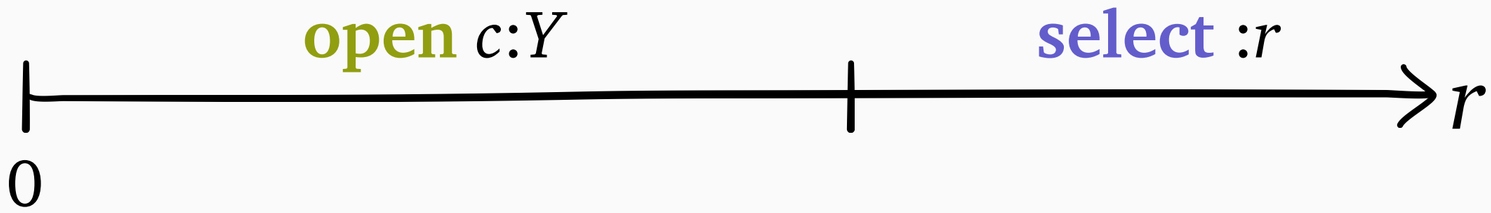

Defining the **Gittins** index

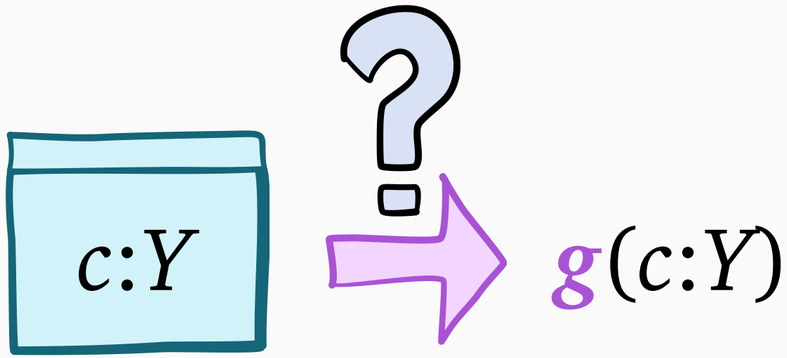

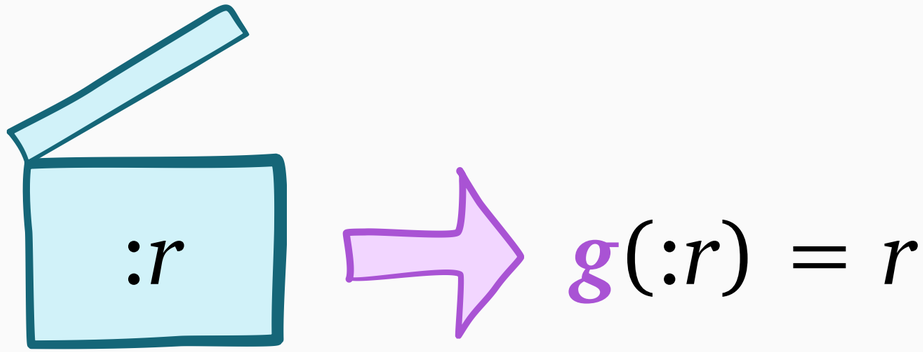

1.5-box problem

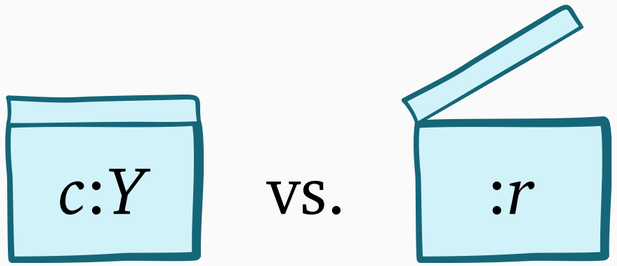

Key question: what to do in 1.5-box problem?

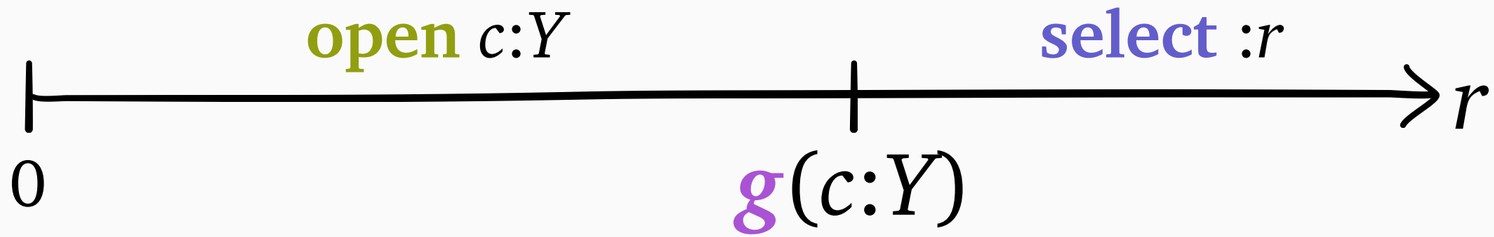

Defining the **Gittins** index

1.5-box problem

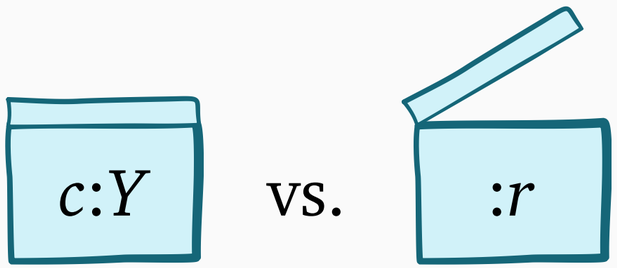

Key question: what to do in 1.5-box problem?

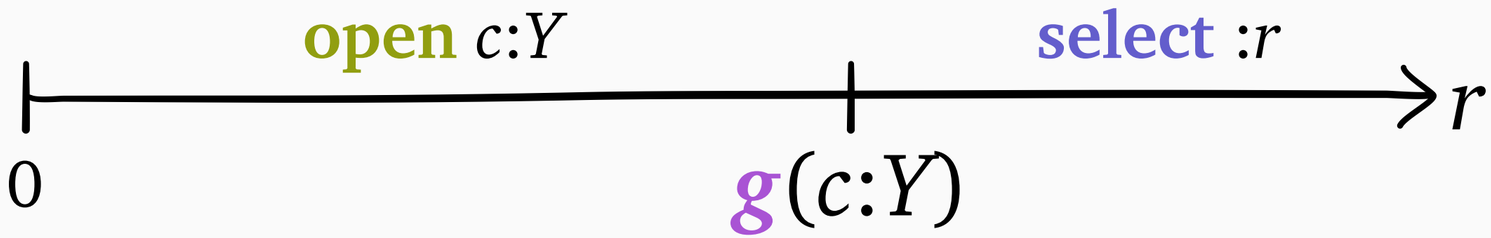

Defining the **Gittins** index

1.5-box problem

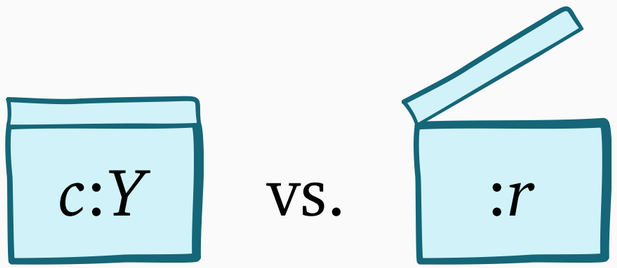

Key question: what to do in 1.5-box problem?

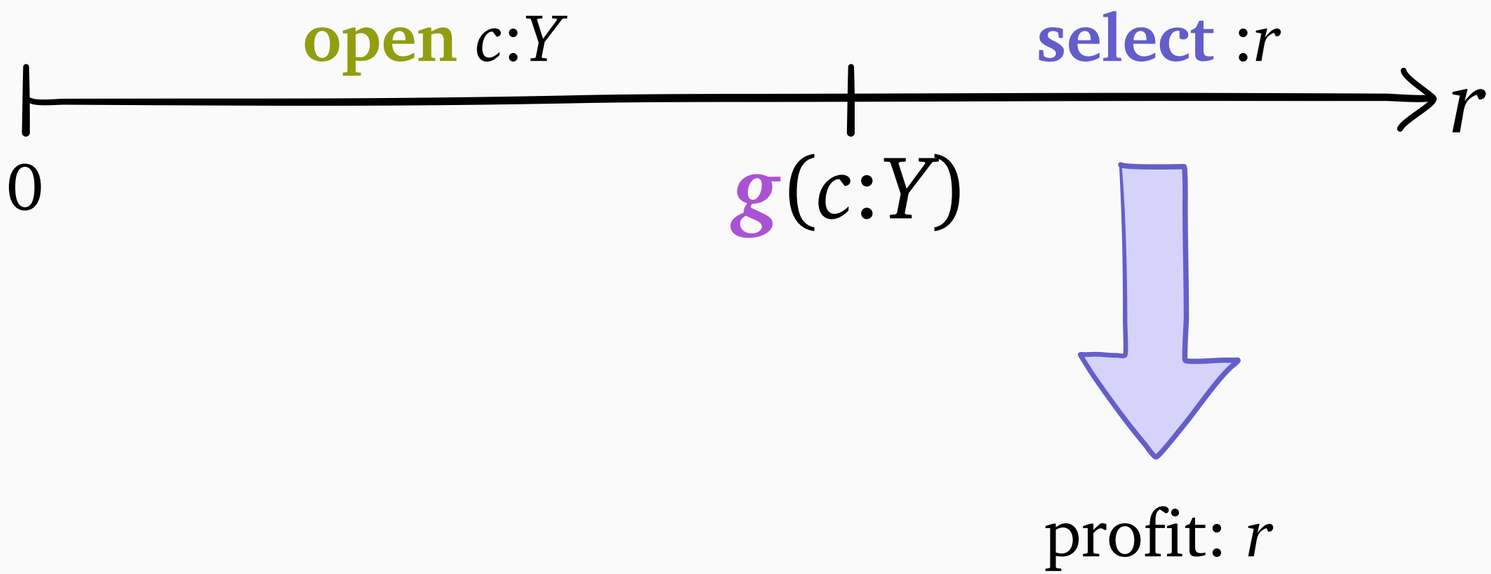

Defining the **Gittins** index

1.5-box problem

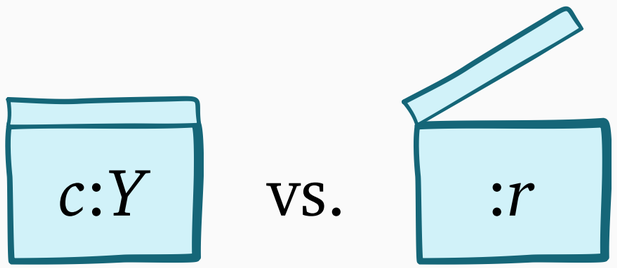

Key question: what to do in 1.5-box problem?

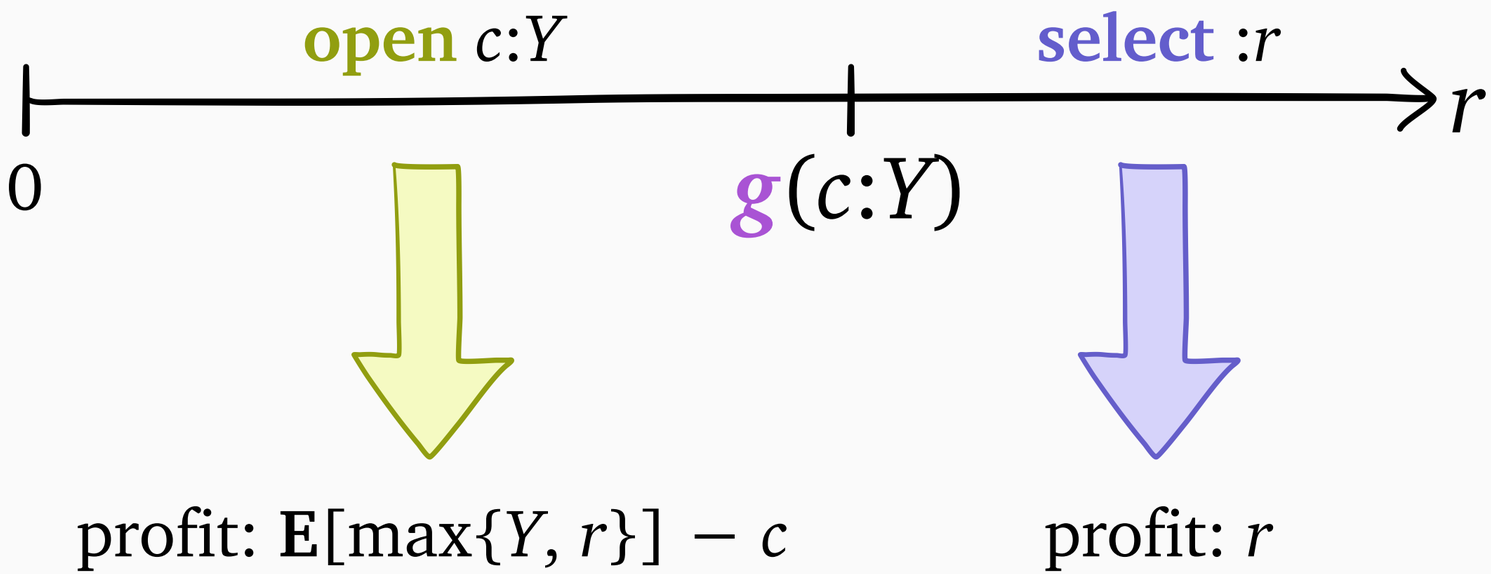

Defining the **Gittins** index

1.5-box problem

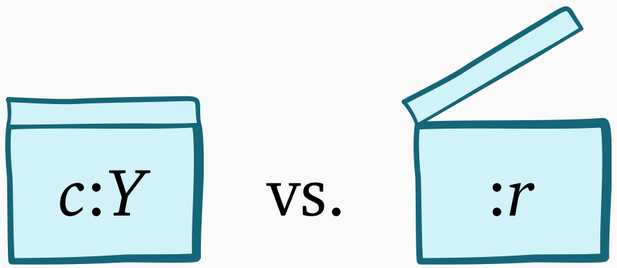

Key question: what to do in 1.5-box problem?

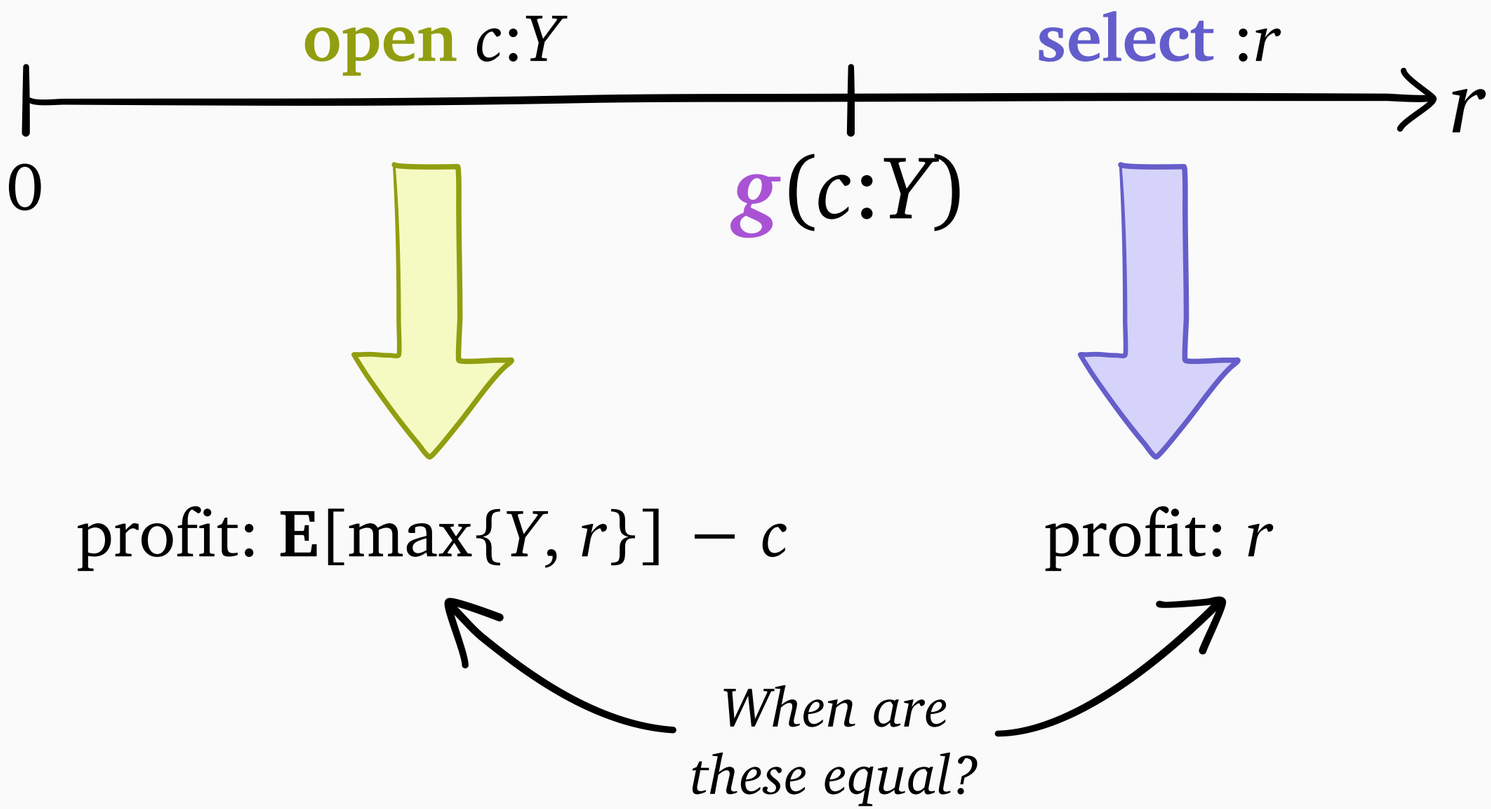

Defining the Gittins index

1.5-box problem

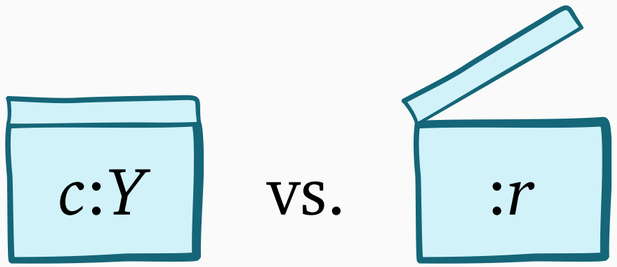

Key question: what to do in 1.5-box problem?

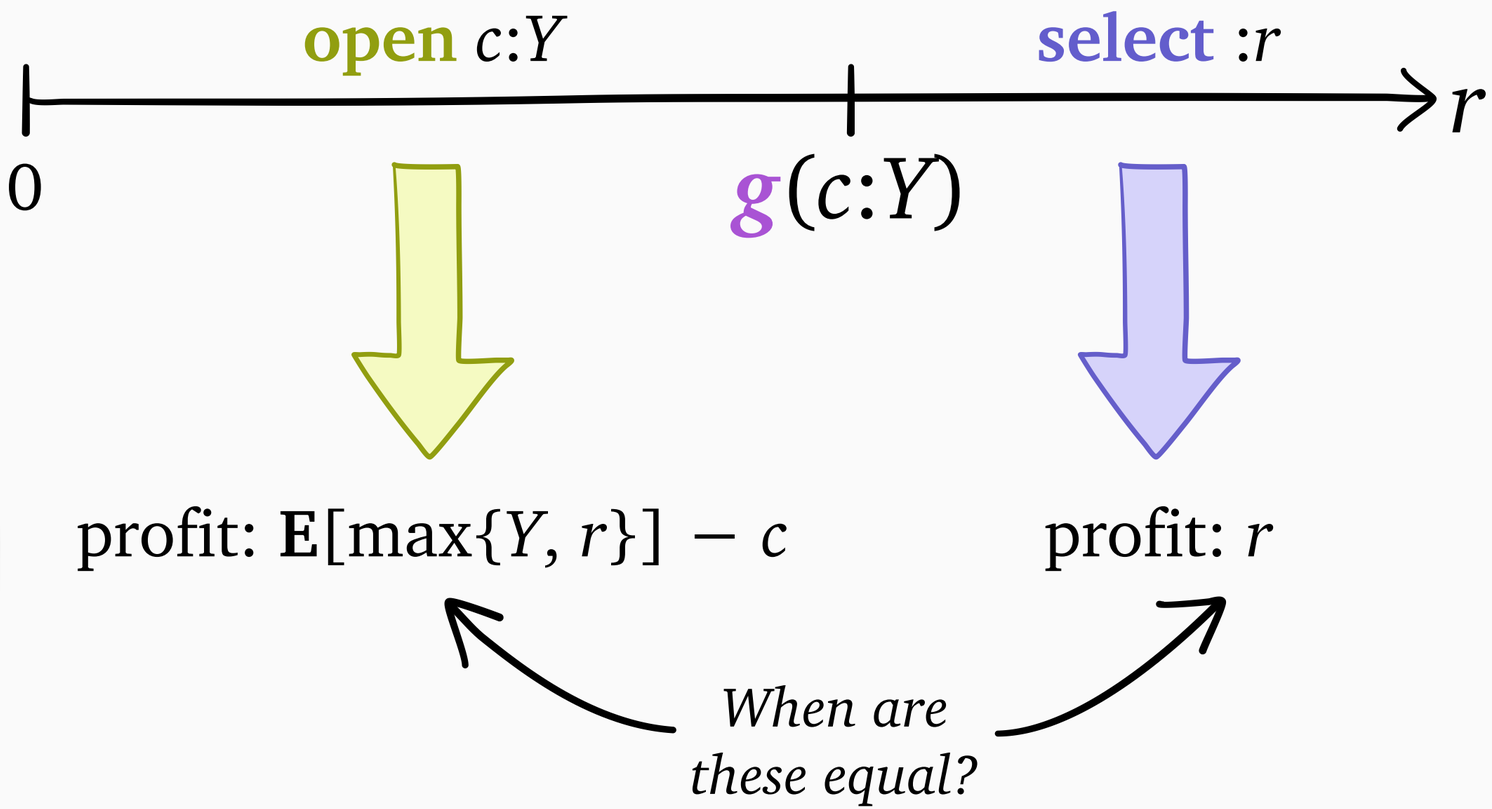

Defn: $g(c{:}Y)$ is solution r to

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

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

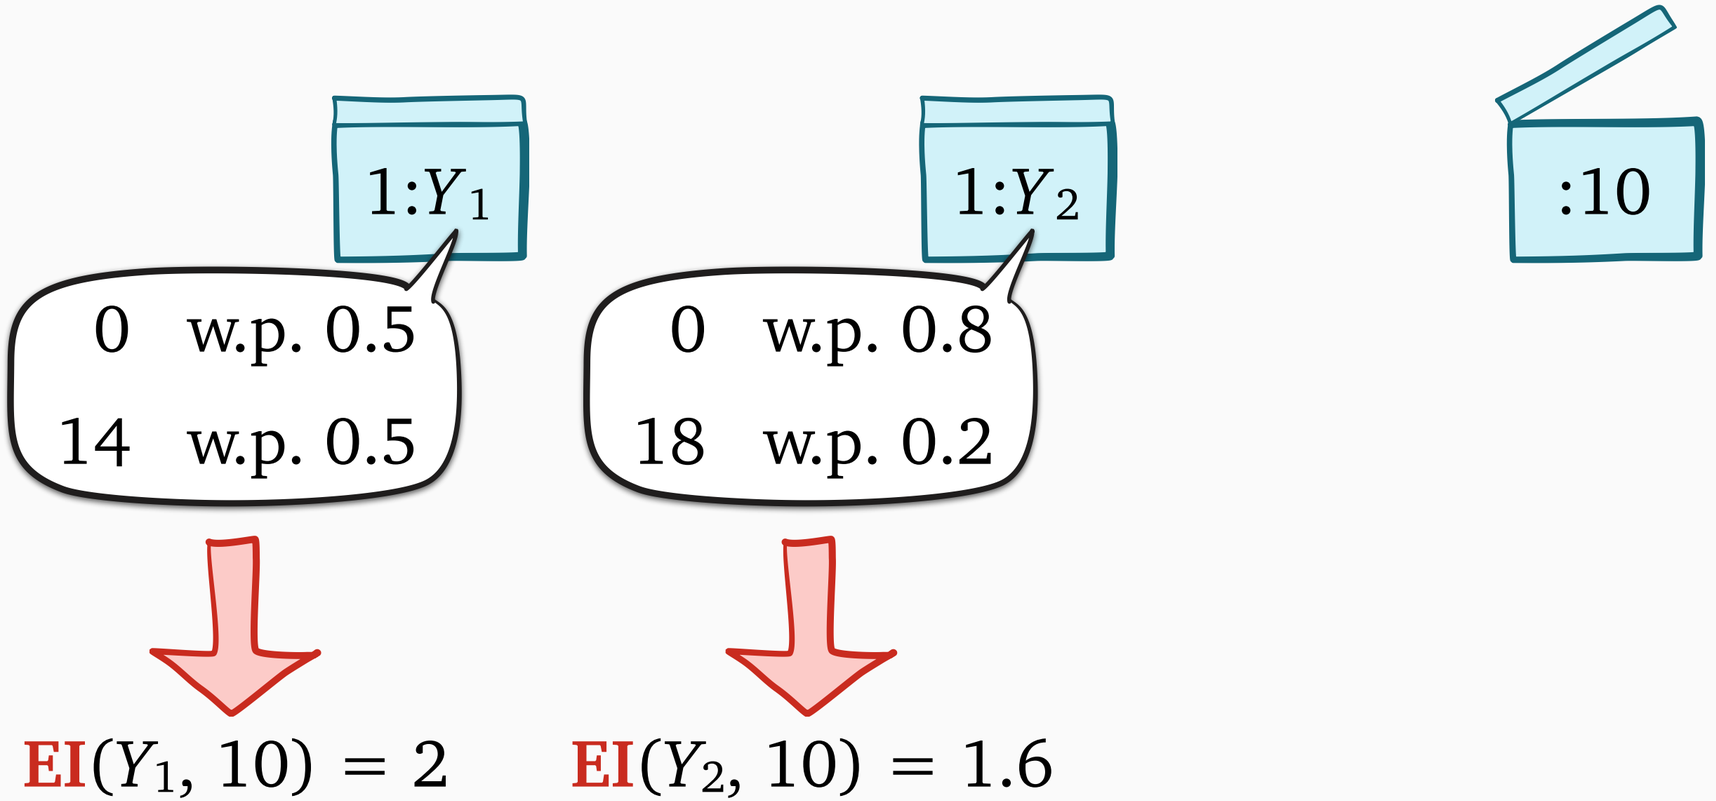

Difference between **EI** and **Gittins**

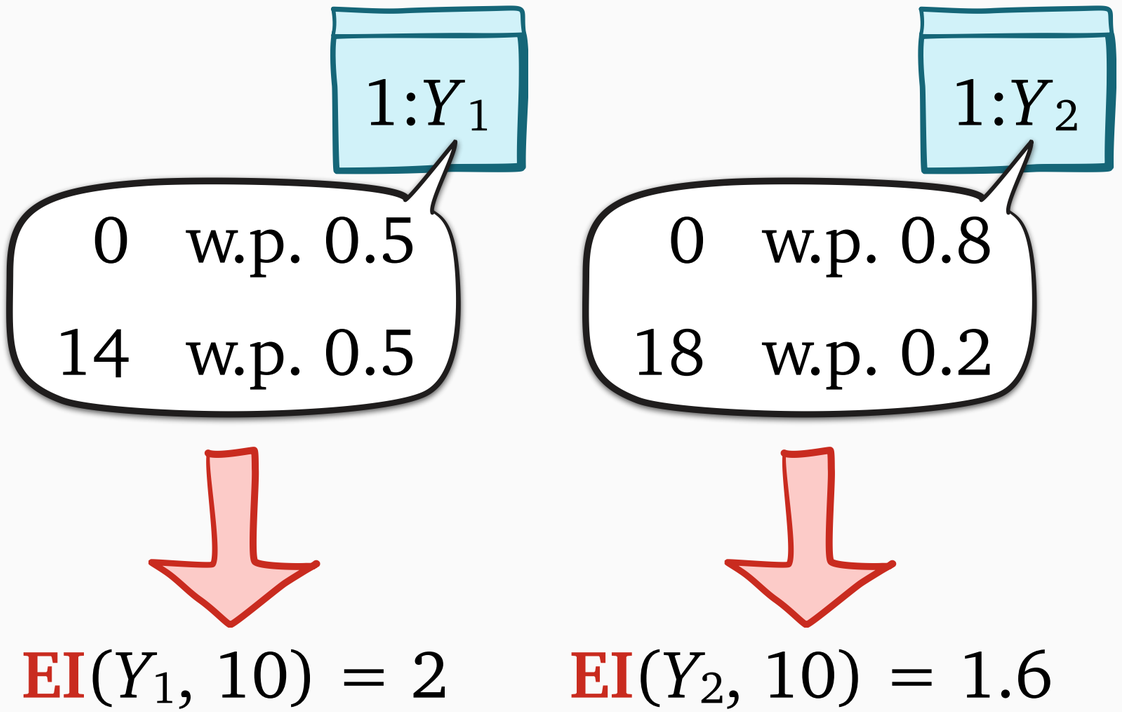

EI$(Y_1, 10) = 2$ **EI**$(Y_2, 10) = 1.6$

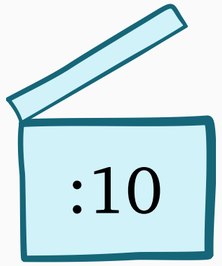

With $r = 10$, what is **EI**$(Y, r) - c$?

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

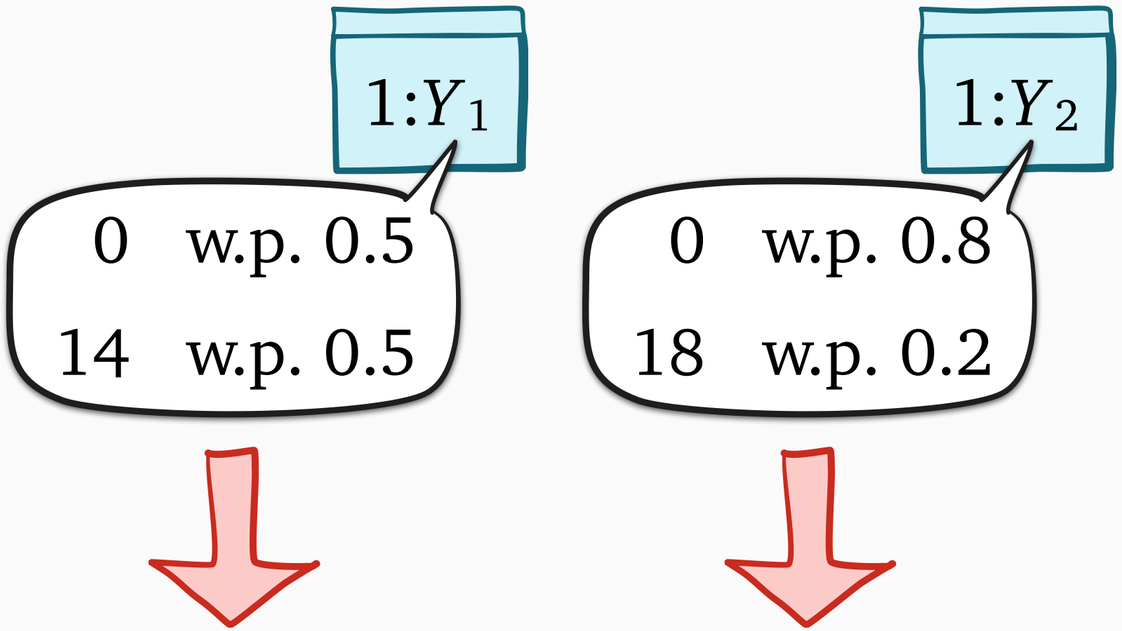

EI$(Y_1, 10) = 2$ **EI**$(Y_2, 10) = 1.6$

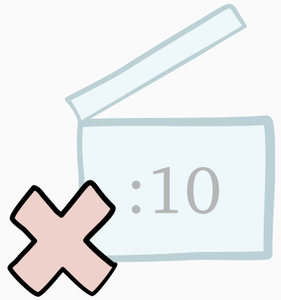

With $r = 10$, what is **EI**$(Y, r) - c$?

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

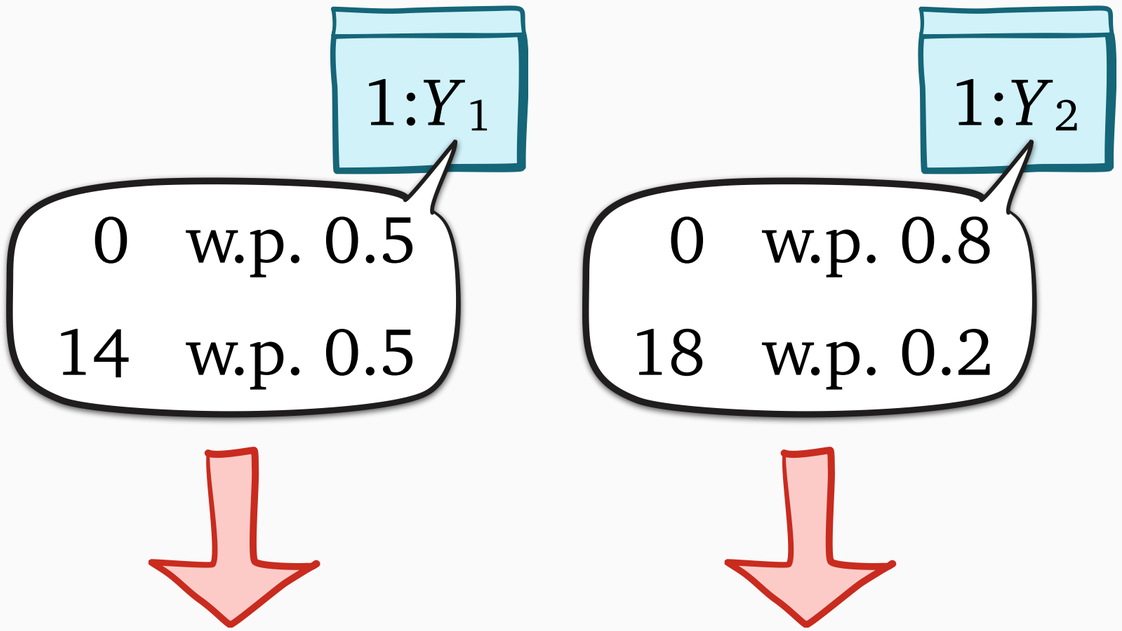

EI$(Y_1, 10) = 2$ **EI**$(Y_2, 10) = 1.6$

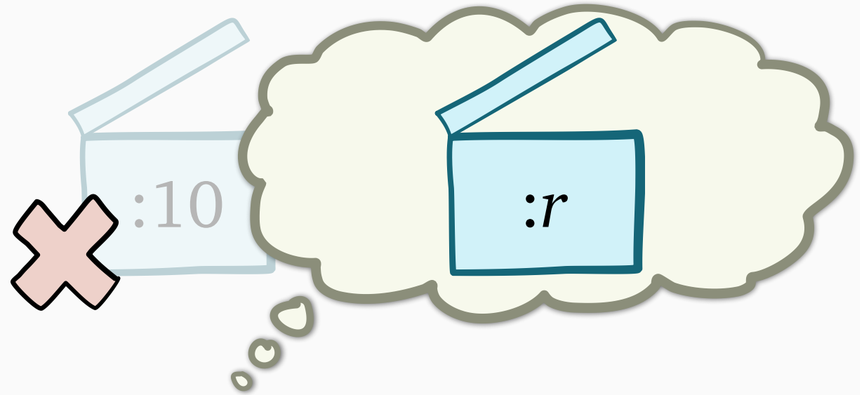

With $r = 10$, what is **EI**$(Y, r) - c$?

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

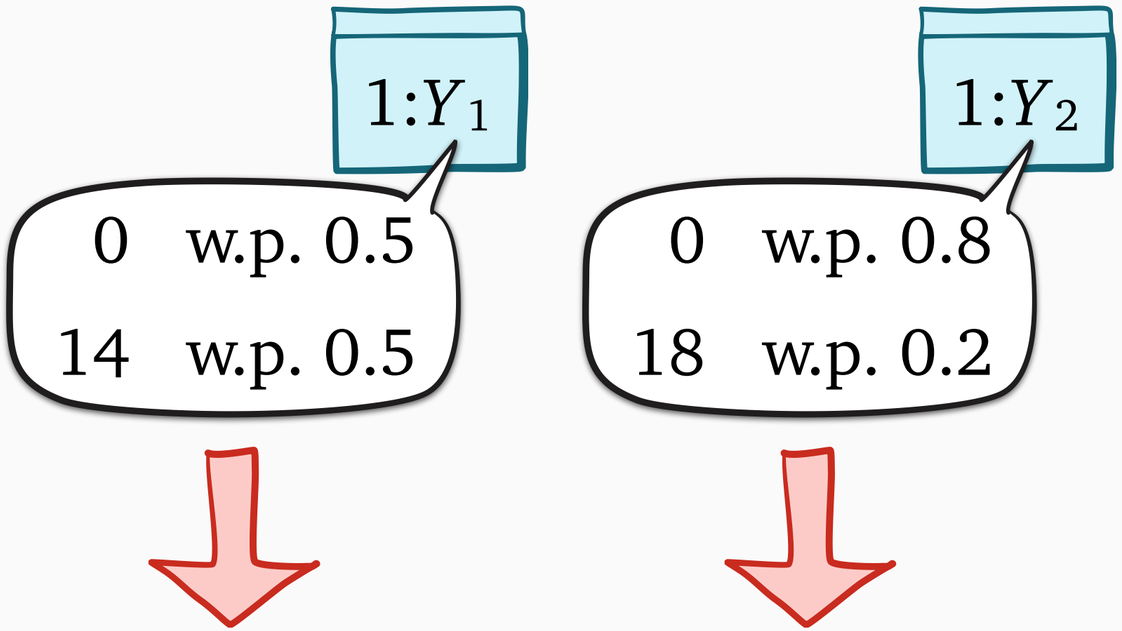

EI$(Y_1, 10) = 2$ **EI**$(Y_2, 10) = 1.6$

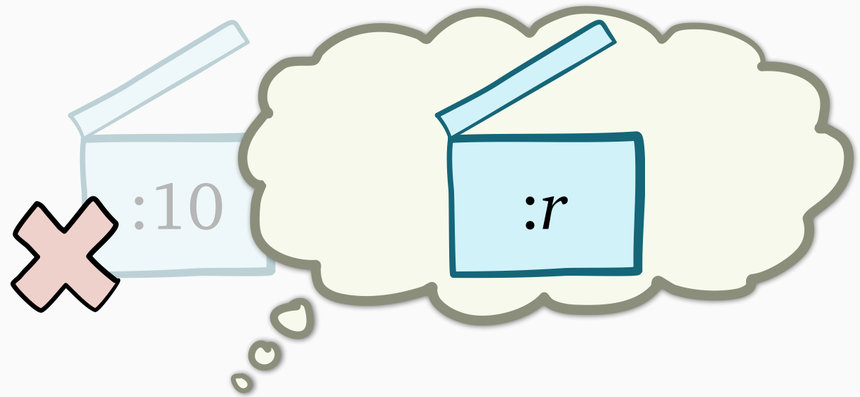

With $r = 10$, what is **EI**$(Y, r) - c$?

For what r does **EI**$(Y, r) - c = 0$?

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

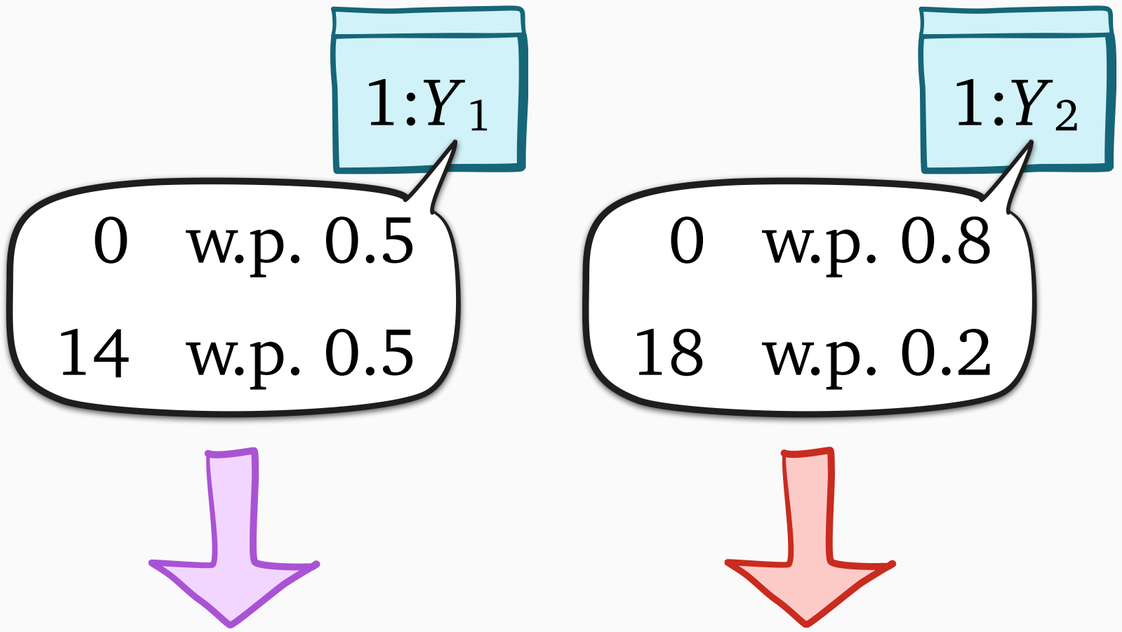

EI$(Y_1, 10) = 2$

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

EI$(Y_2, 10) = 1.6$

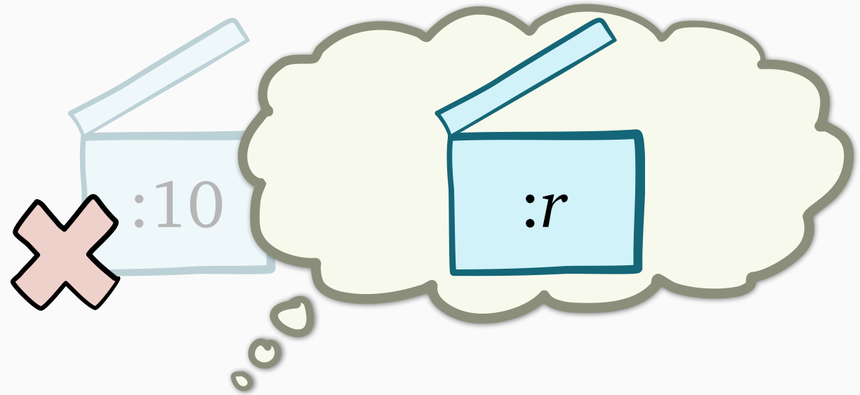

With $r = 10$, what is **EI**$(Y, r) - c$?

For what r does **EI**$(Y, r) - c = 0$?

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

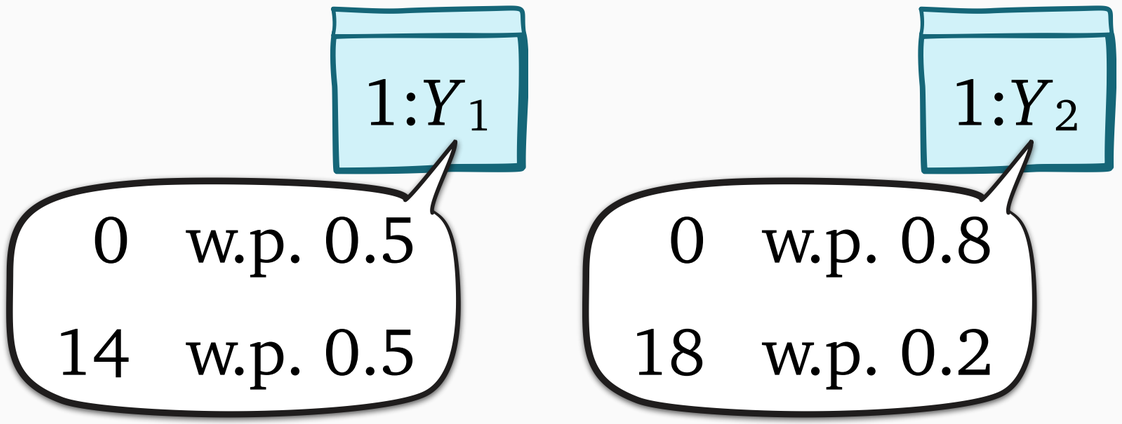

EI$(Y_1, 10) = 2$

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

EI$(Y_2, 10) = 1.6$

$g(1{:}Y_2) = 13$

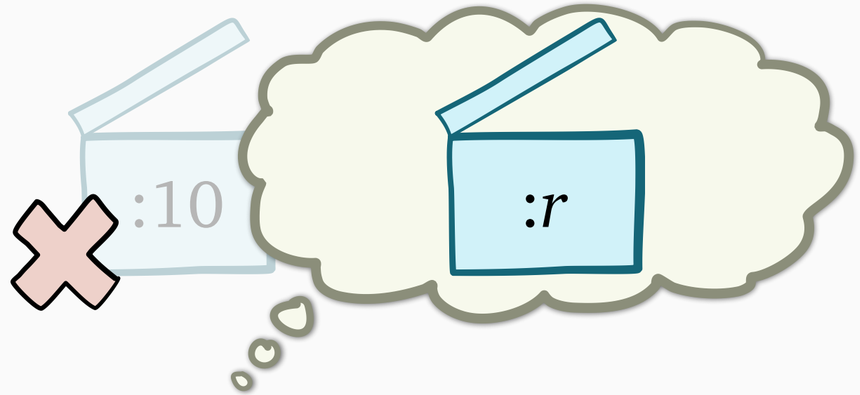

With $r = 10$, what is **EI**$(Y, r) - c$?

For what r does **EI**$(Y, r) - c = 0$?

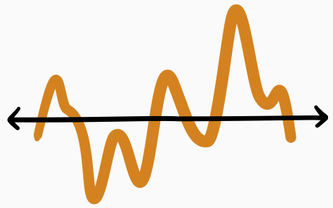

BayesOpt

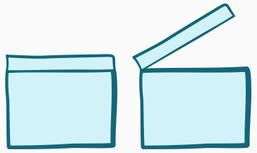

Pandora's box

Tutorial: What are **Pandora's box** and **Gittins**?

Results: Does **Gittins** work for **BayesOpt**?

Hype: How could **Gittins** help practical **BayesOpt**?

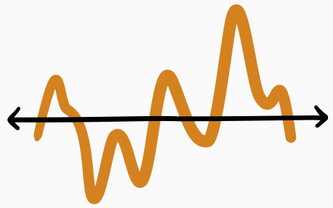

BayesOpt

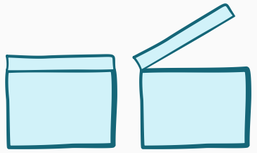

Pandora's box

Tutorial: What are **Pandora's box** and **Gittins**?

Results: Does **Gittins** work for **BayesOpt**?

Hype: How could **Gittins** help practical **BayesOpt**?

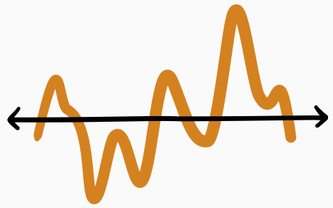

BayesOpt

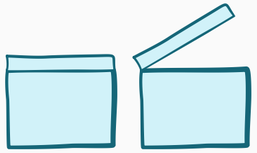

Pandora's box

Tutorial: What are **Pandora's box** and **Gittins**?

Results: Does **Gittins** work for **BayesOpt**?

Hype: How could **Gittins** help practical **BayesOpt**?

What is **BayesOpt**?

Unknown function:

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

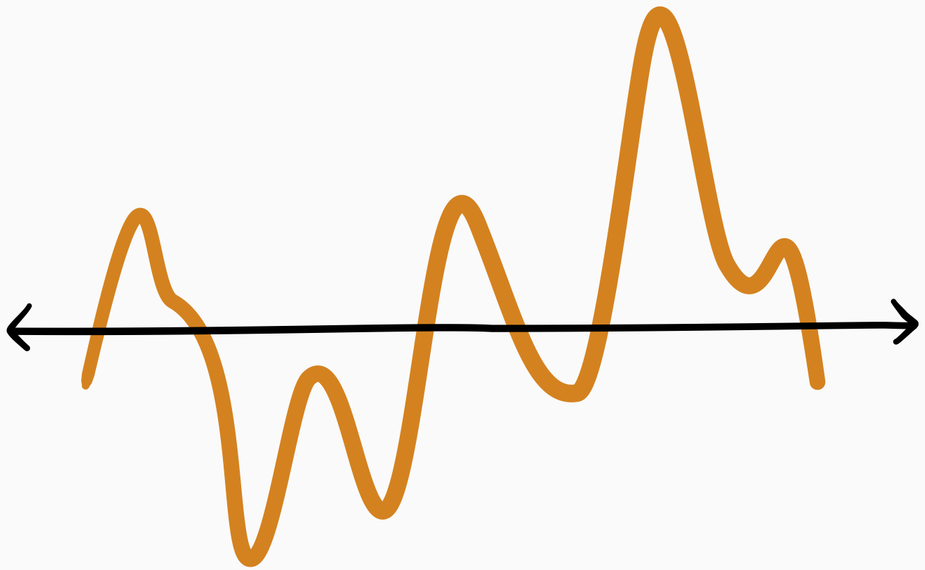

What is **BayesOpt**?

Unknown function:

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

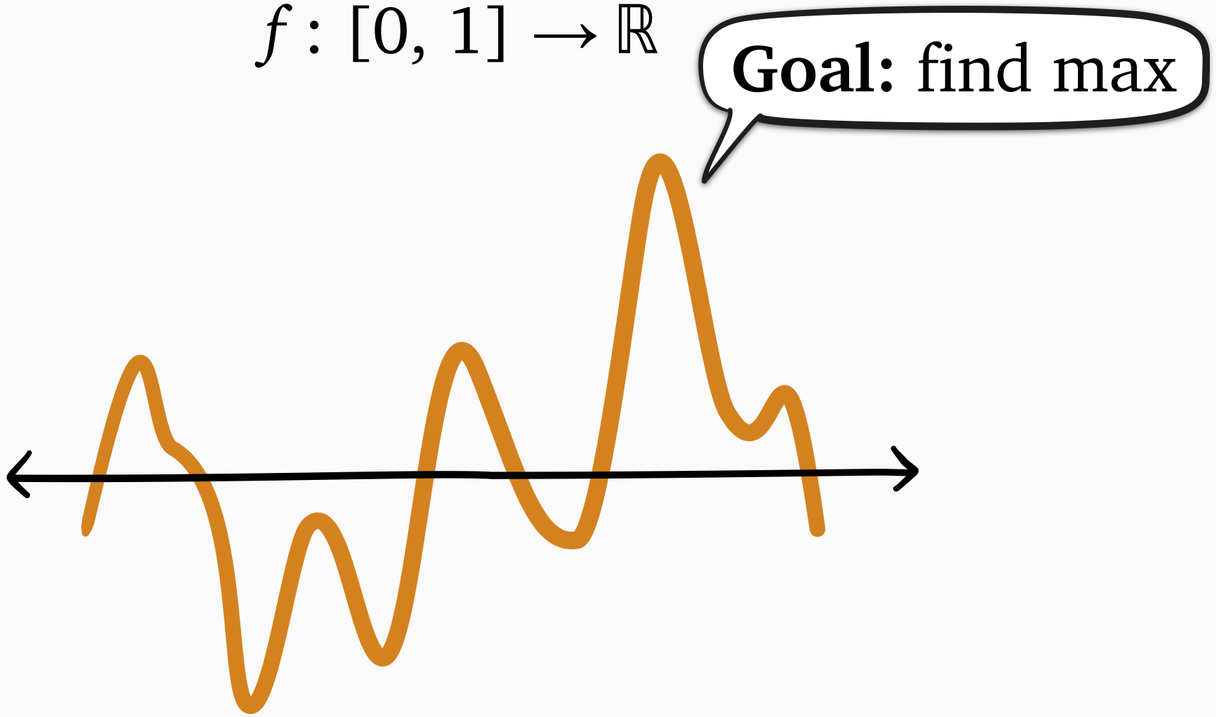

What is **BayesOpt**?

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

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

Goal: find max

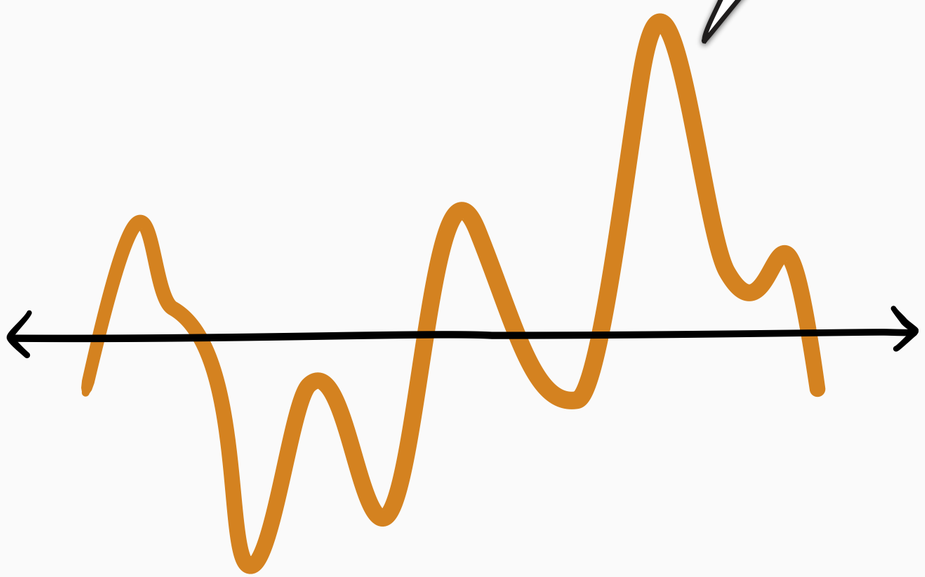

What is **BayesOpt**?

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

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

Goal: find max

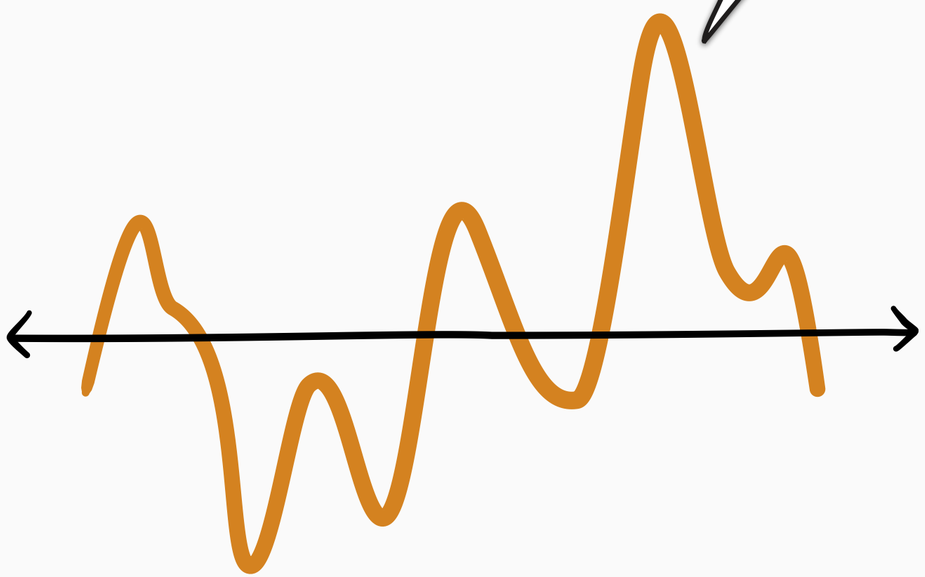

Fixed horizon: maximize

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

What is **BayesOpt**?

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

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

Goal: find max

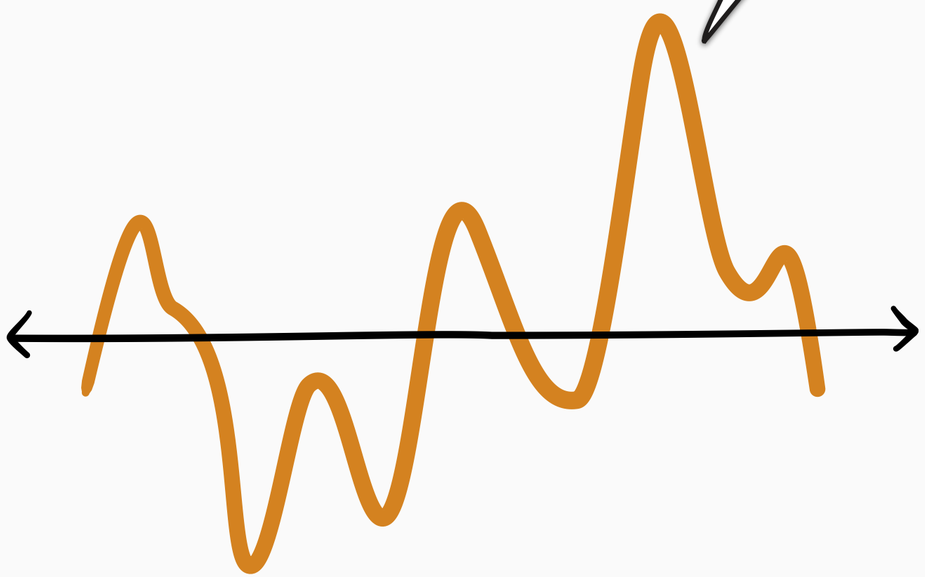

Fixed horizon: maximize

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

T fixed

x_t chosen

What is **BayesOpt**?

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

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

Goal: find max

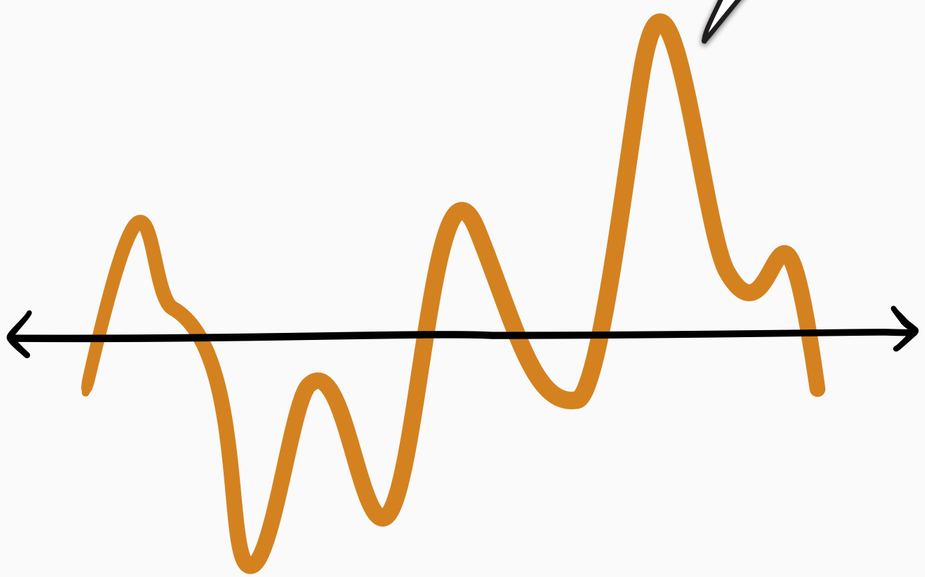

Fixed horizon: maximize

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

T fixed

x_t chosen

Cost-per-sample: maximize

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

What is **BayesOpt**?

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

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

Goal: find max

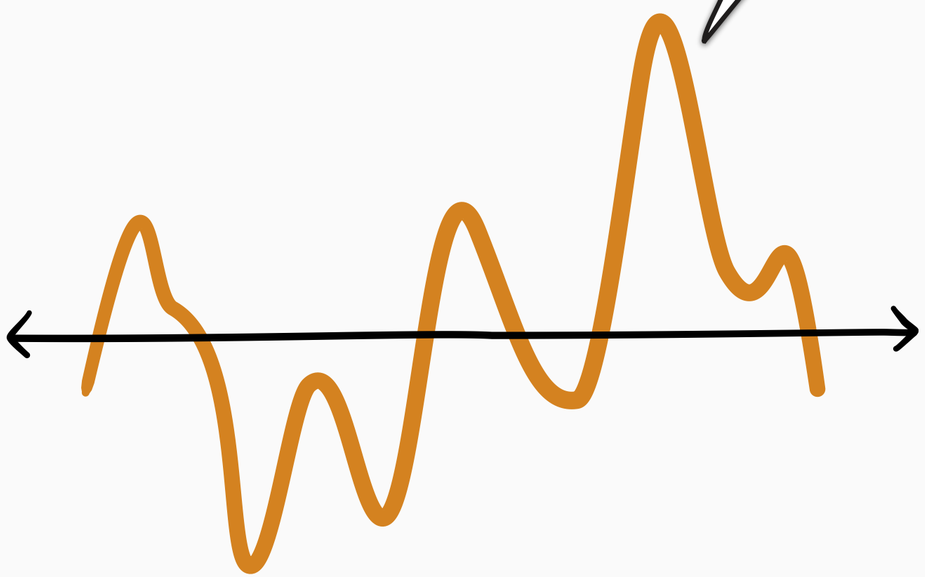

Fixed horizon: maximize

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

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

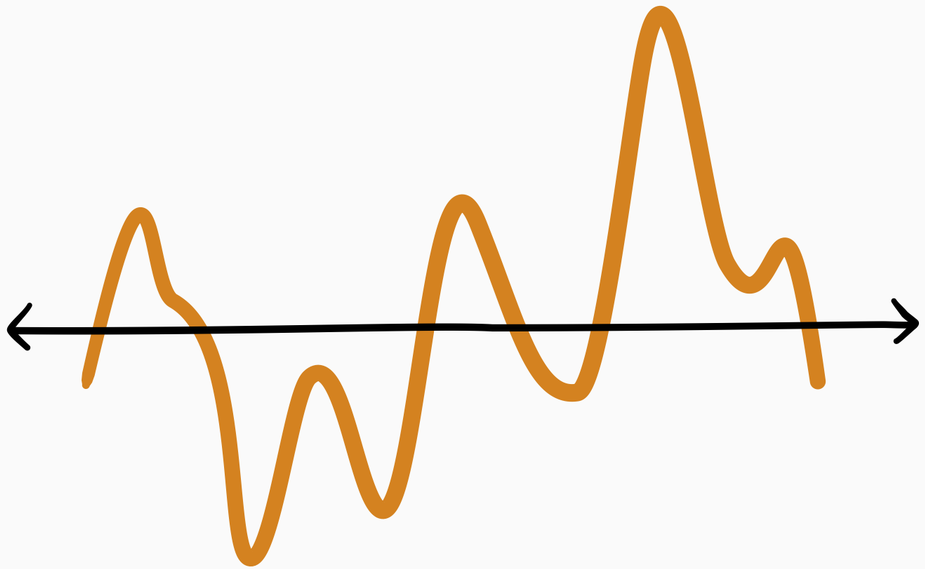

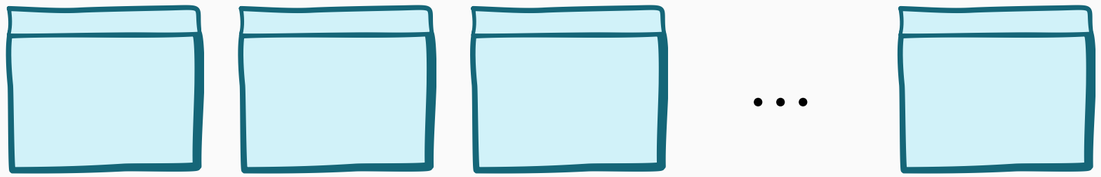

BayesOpt *vs.* Pandora's box

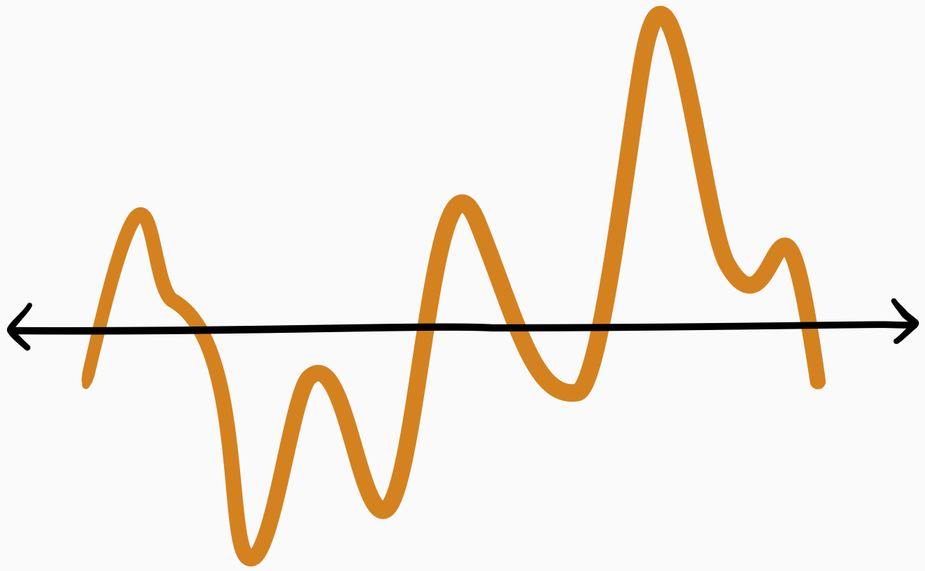

- *Continuous* domain

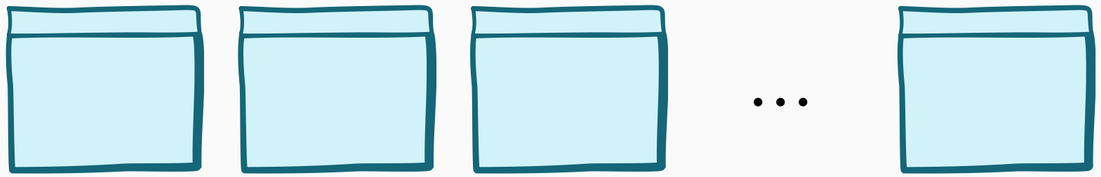

BayesOpt *vs.* Pandora's box

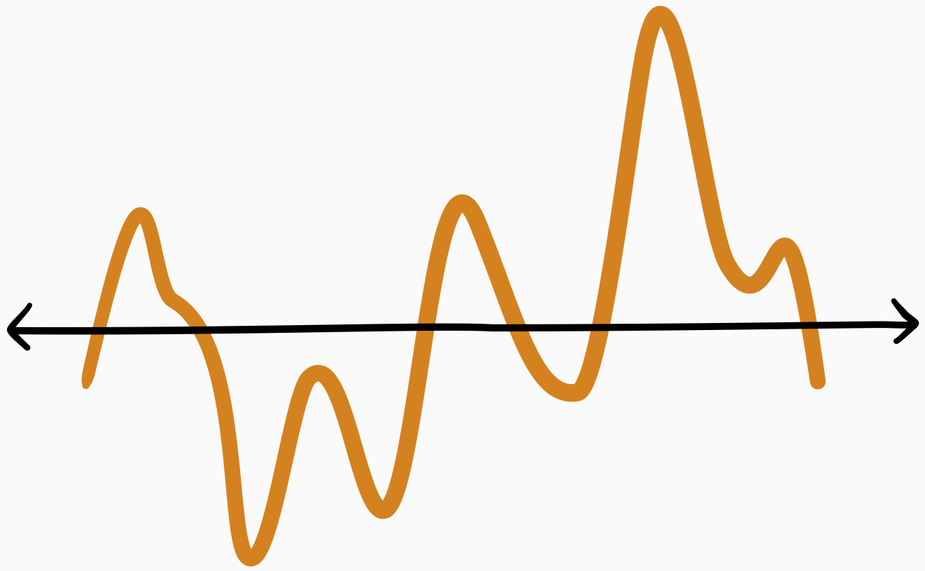

- *Continuous* domain

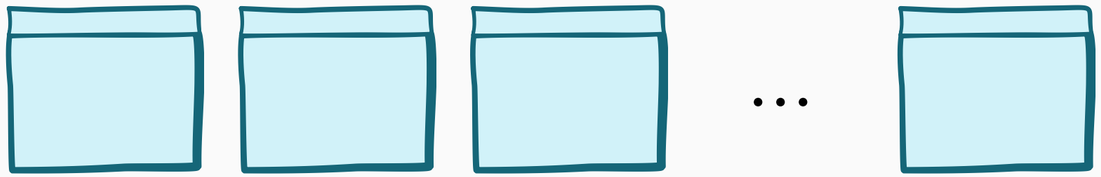

- *Discrete* domain

BayesOpt *vs.* Pandora's box

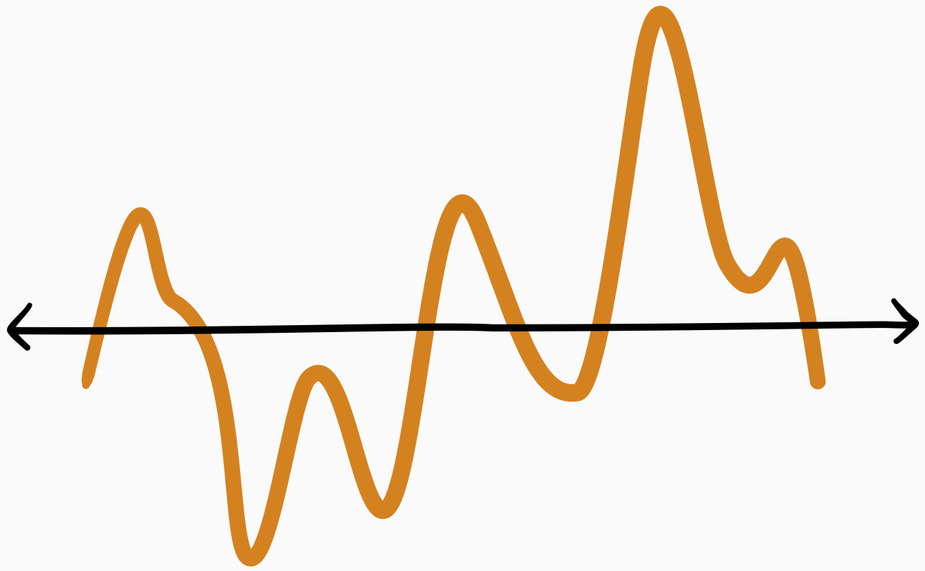

- *Continuous* domain
- *Correlated* values

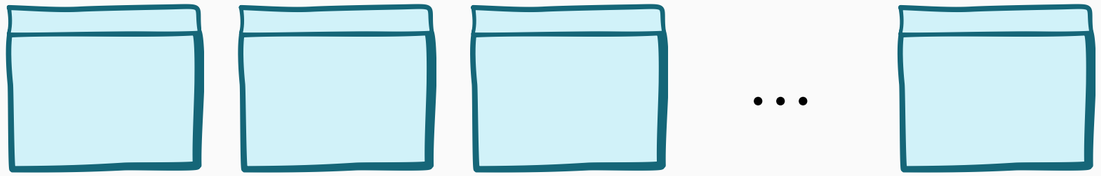

- *Discrete* domain

BayesOpt *vs.* Pandora's box

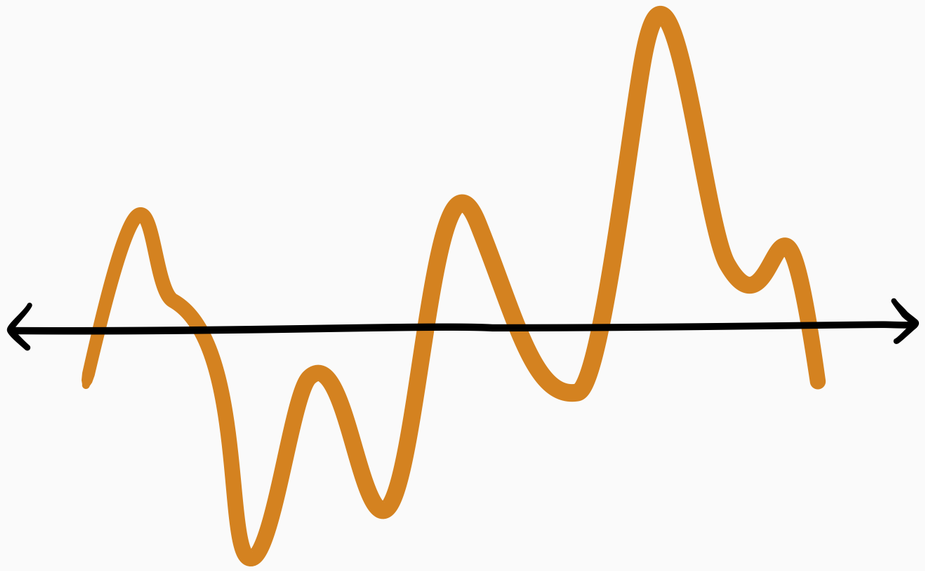

- *Continuous* domain
- *Correlated* values

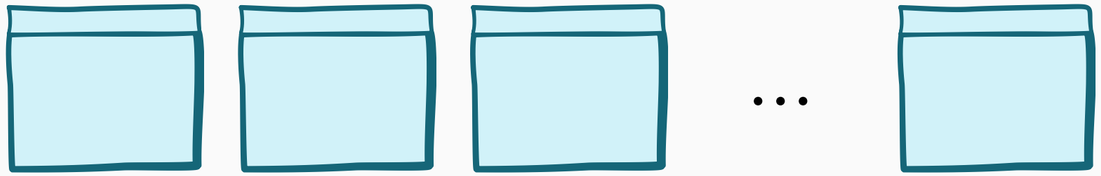

- *Discrete* domain
- *Independent* values

BayesOpt *vs.* Pandora's box

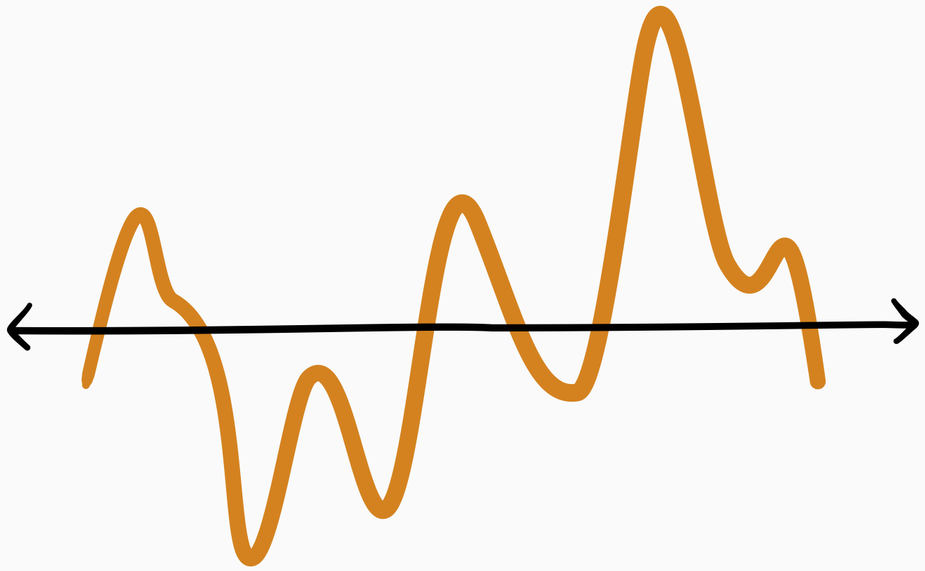

- *Continuous* domain
- *Correlated* values
- *Heuristic* acquisition fns:
  EI, UCB, TS, KG, …

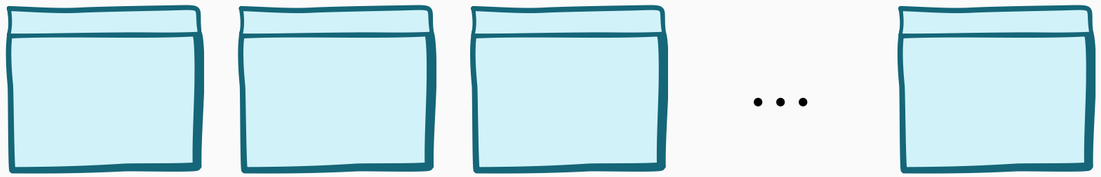

- *Discrete* domain
- *Independent* values

BayesOpt *vs.* Pandora's box

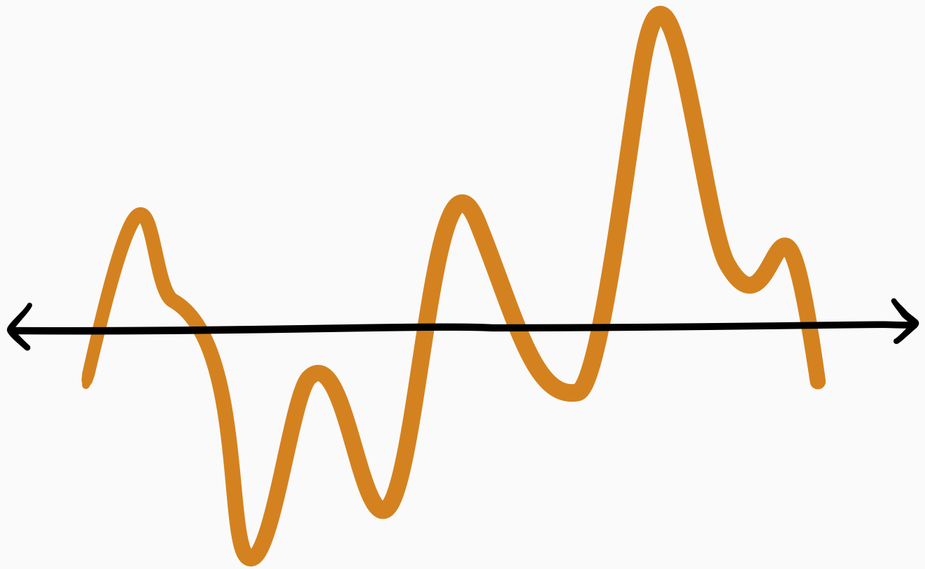

- *Continuous* domain
- *Correlated* values
- *Heuristic* acquisition fns: **EI**, UCB, TS, KG, …

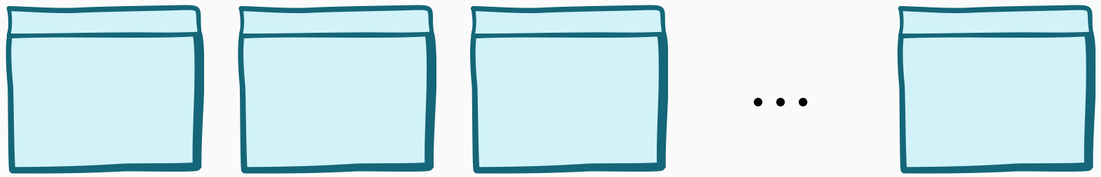

- *Discrete* domain
- *Independent* values
- *Optimal* acquisition fn: **Gittins** index

BayesOpt *vs.* Pandora's box

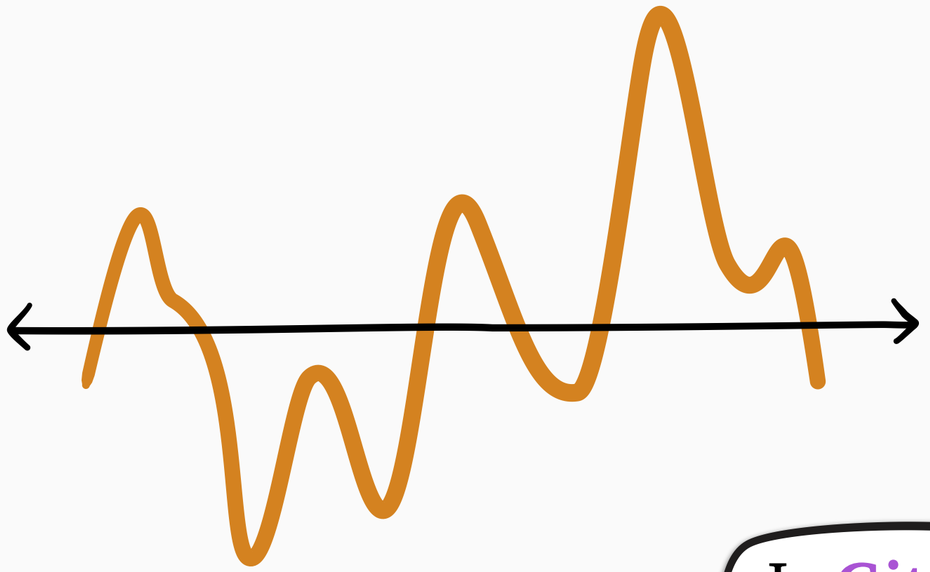

- *Continuous* domain
- *Correlated* values
- *Heuristic* acquisition fns: **EI**, UCB, TS, KG, …

Is **Gittins** any good?

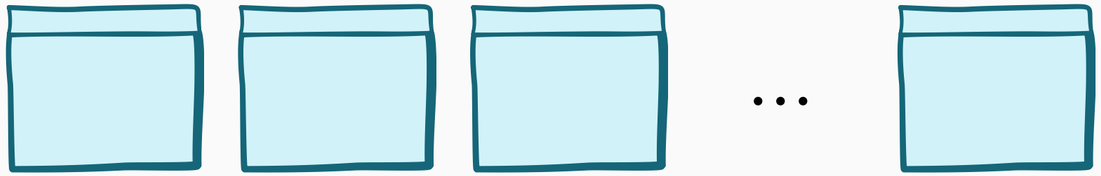

- *Discrete* domain
- *Independent* values
- *Optimal* acquisition fn: **Gittins** index

BayesOpt acquisition functions

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

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

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

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

BayesOpt acquisition functions

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

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

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

BayesOpt acquisition functions

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

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

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

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

BayesOpt acquisition functions

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

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

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

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

assume independent to compute Gittins index

BayesOpt acquisition functions

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

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

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

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

assume independent to compute Gittins index

use correlation in Bayesian update

Results: Bayesian regret

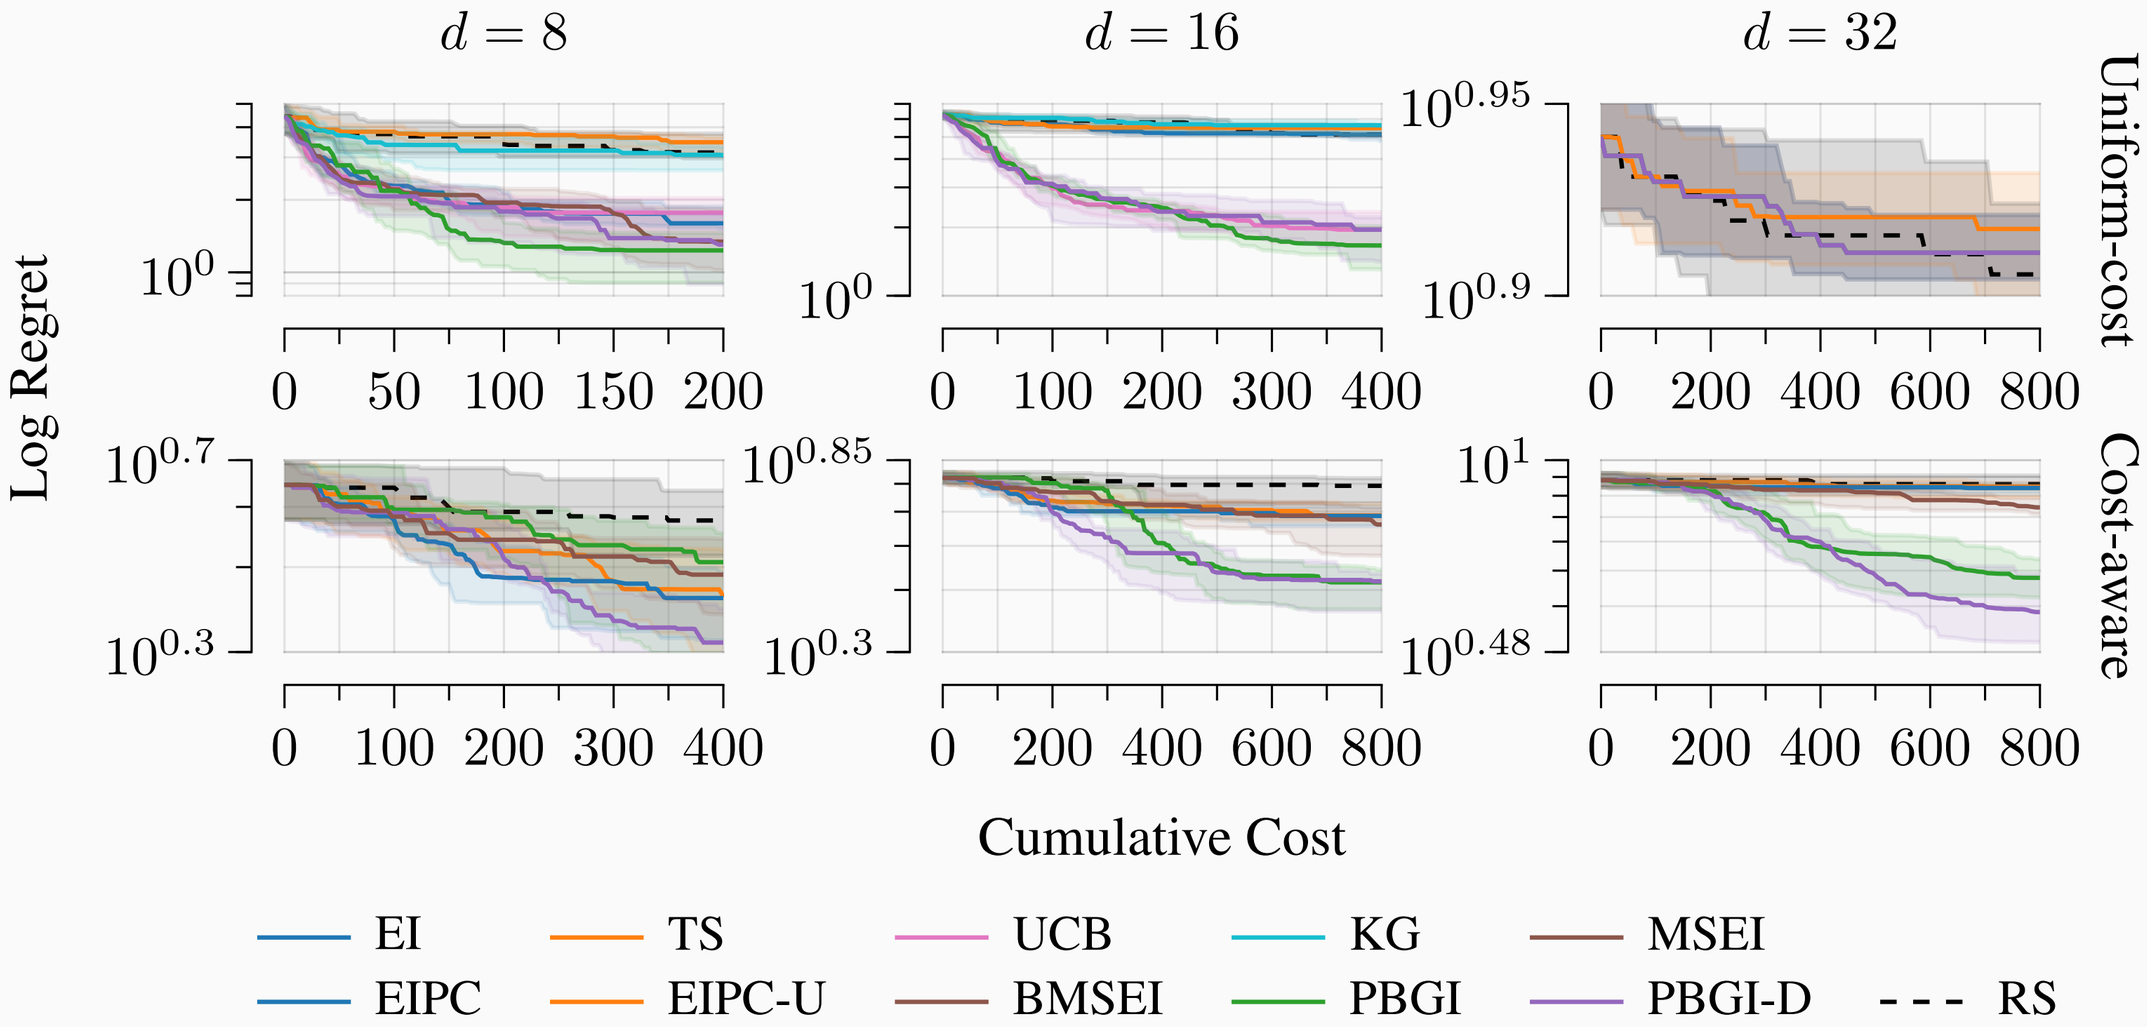

Results: classic benchmarks

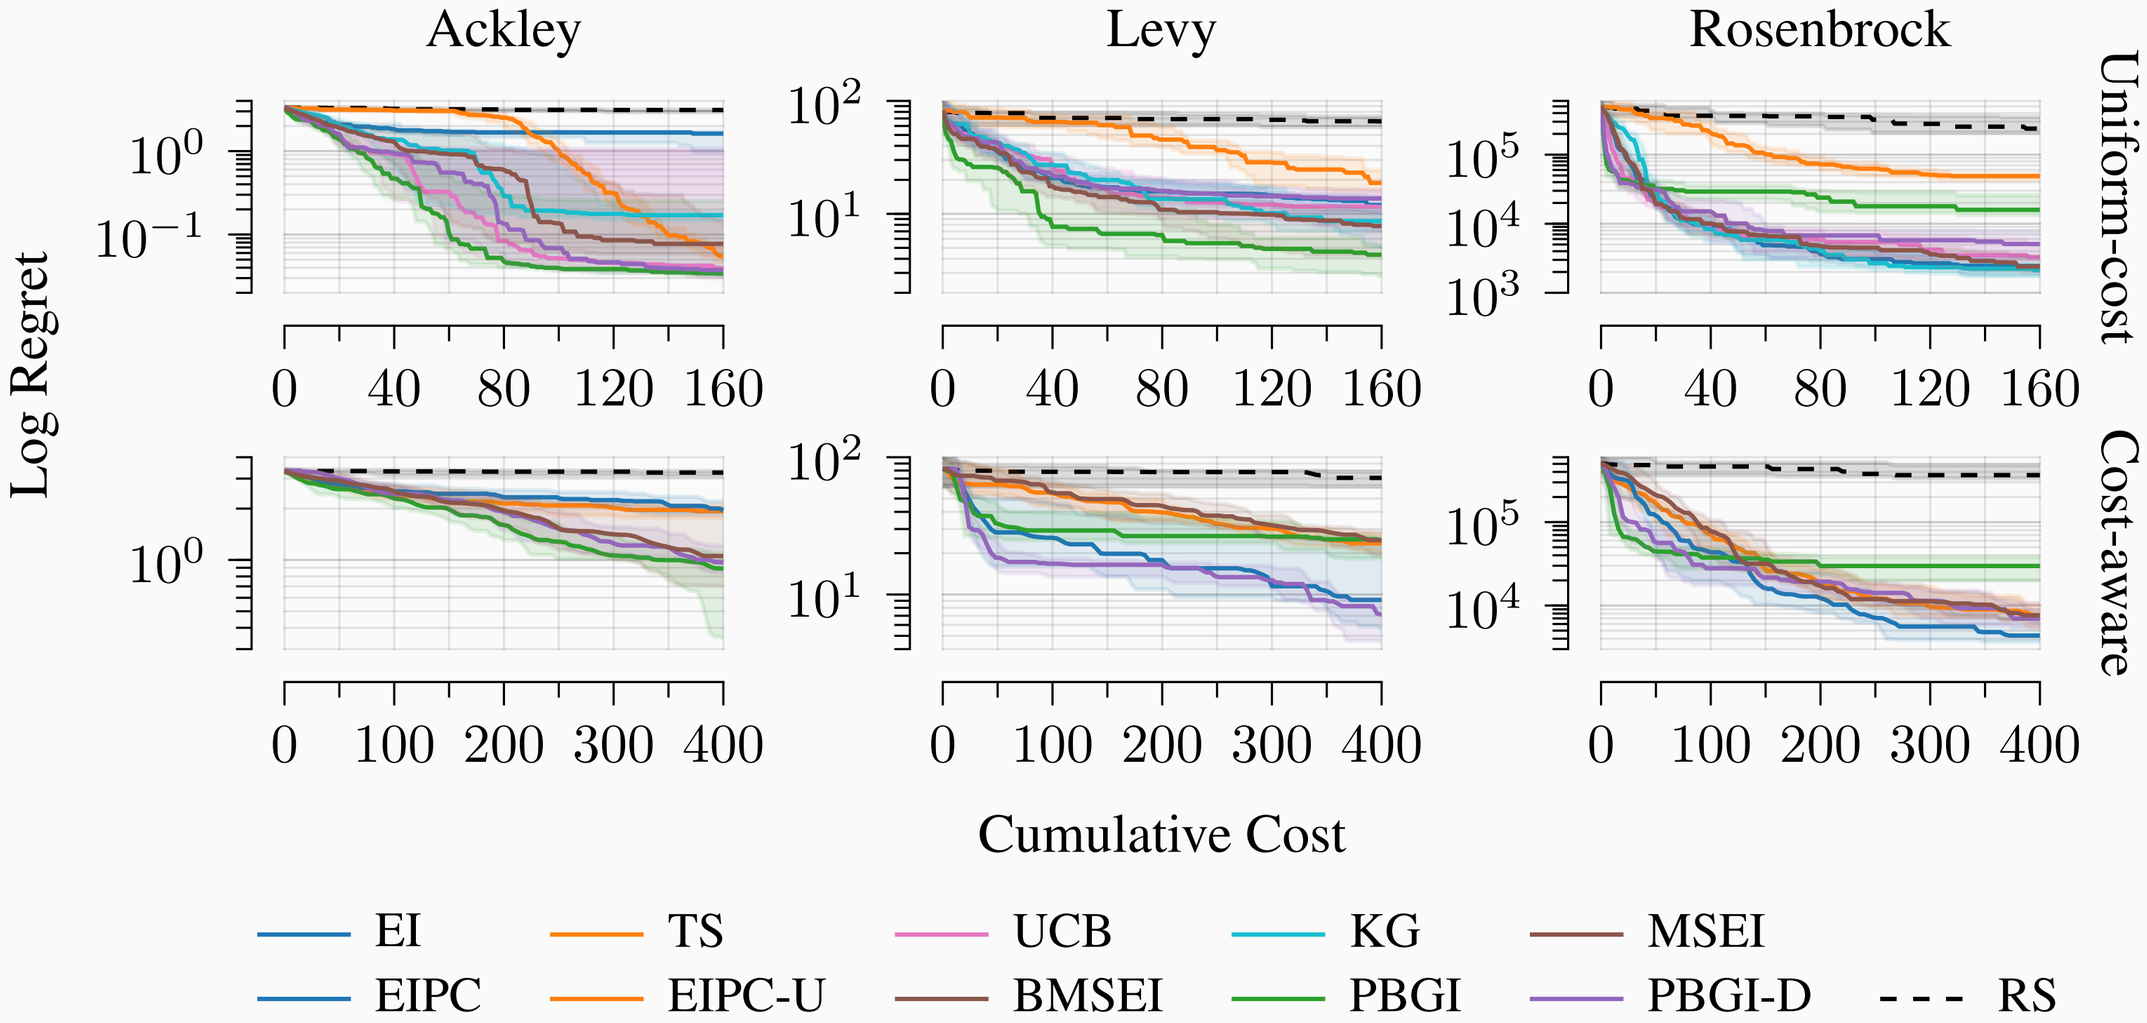

Results: empirical benchmarks

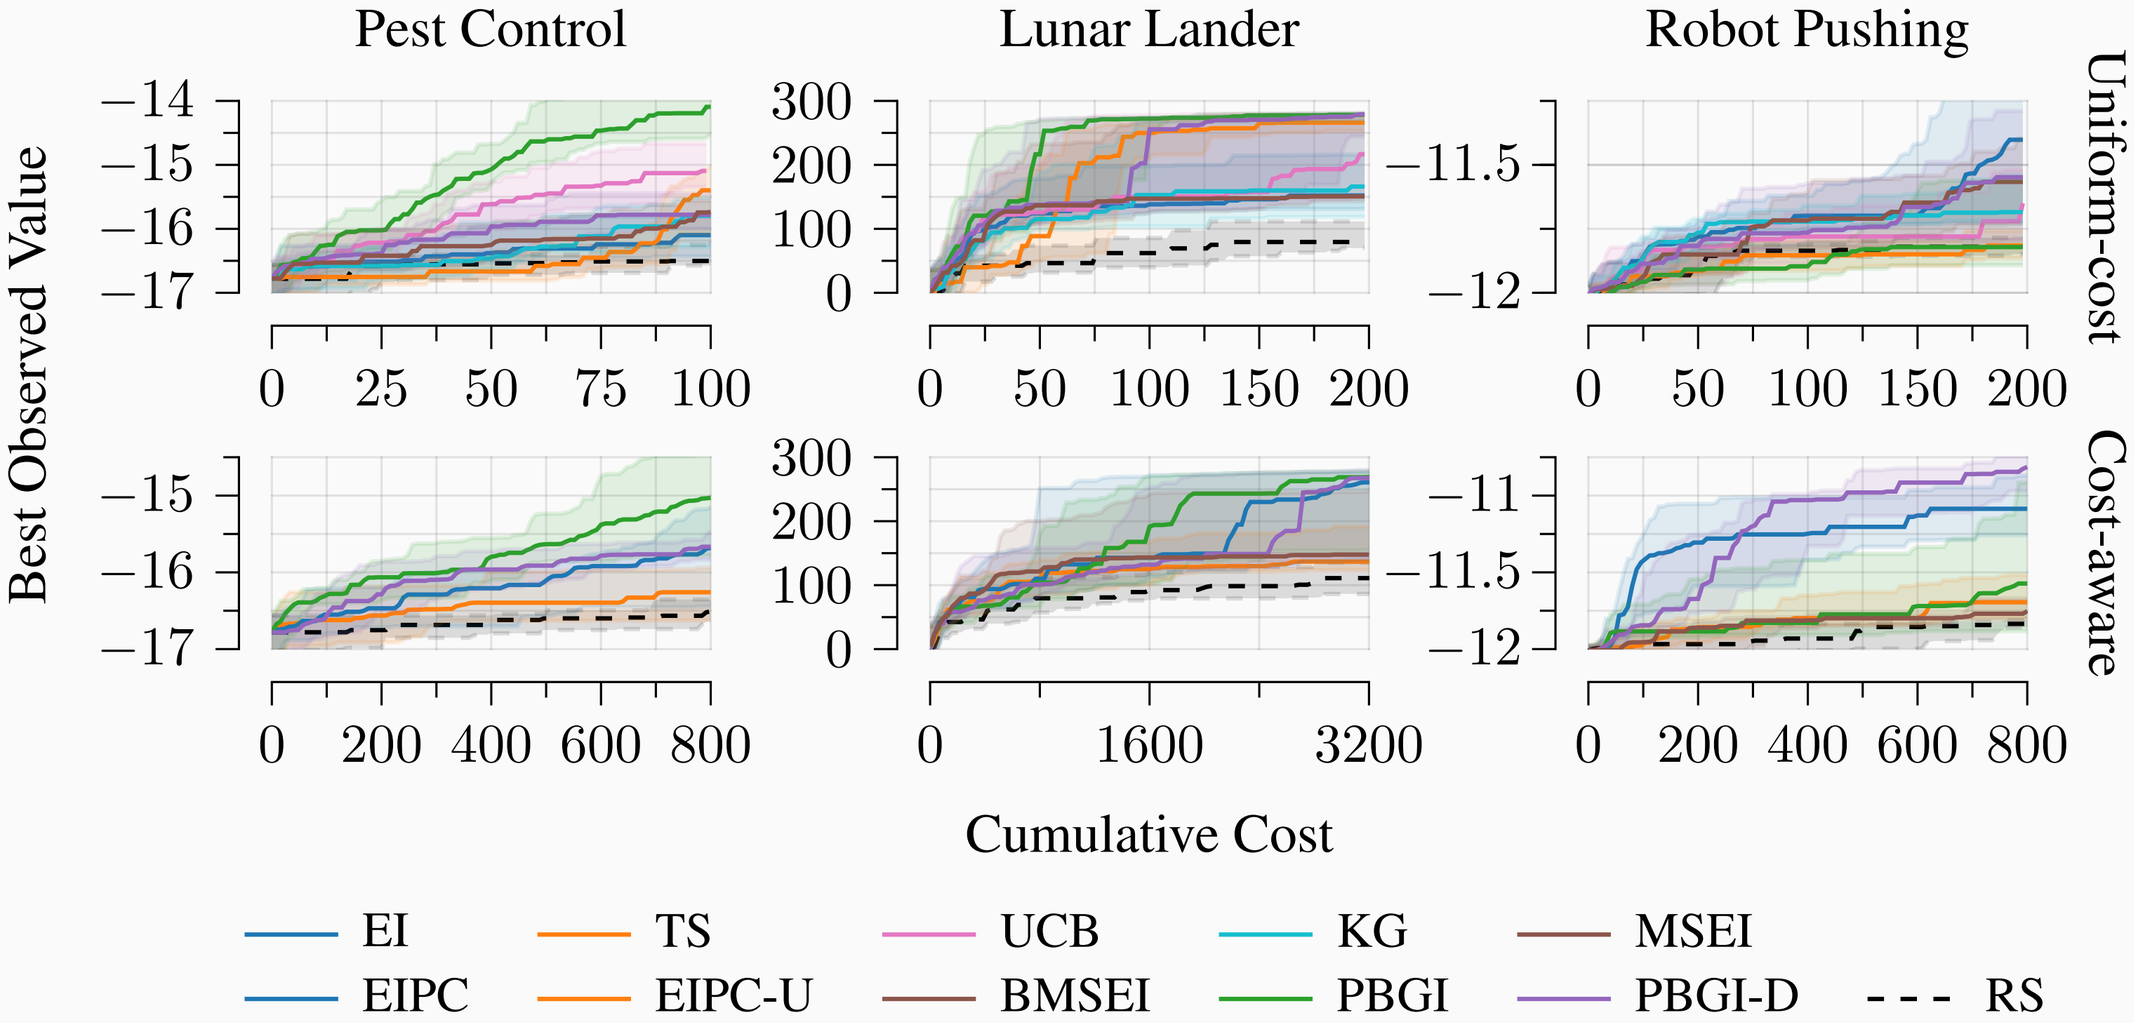

Generalization: **Markovian MAB**

Pandora's box: one stage

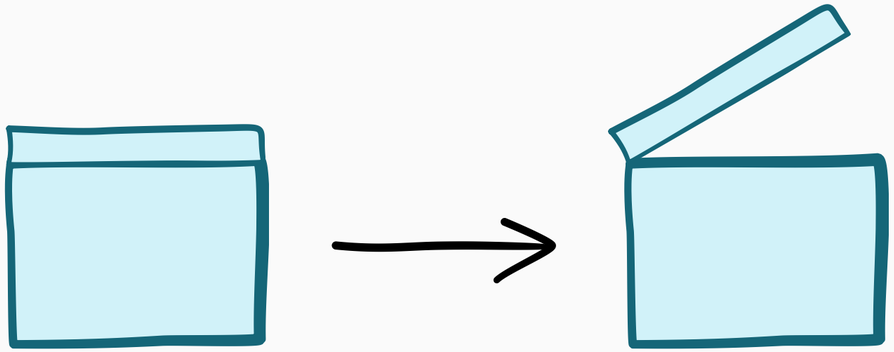

Generalization: **Markovian MAB**

Pandora's box: one stage

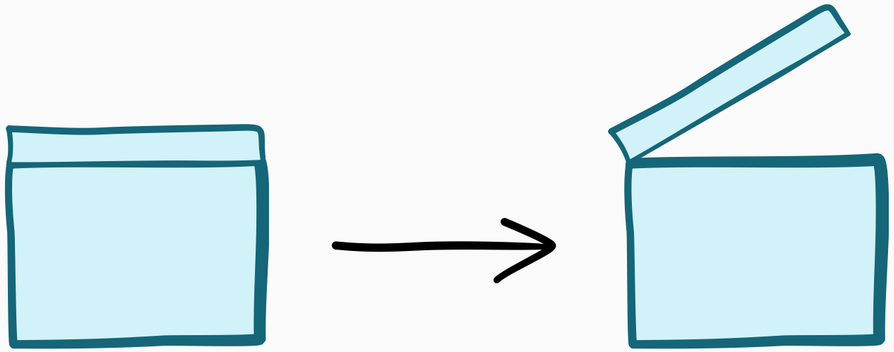

Markovian MAB: general Markov chain

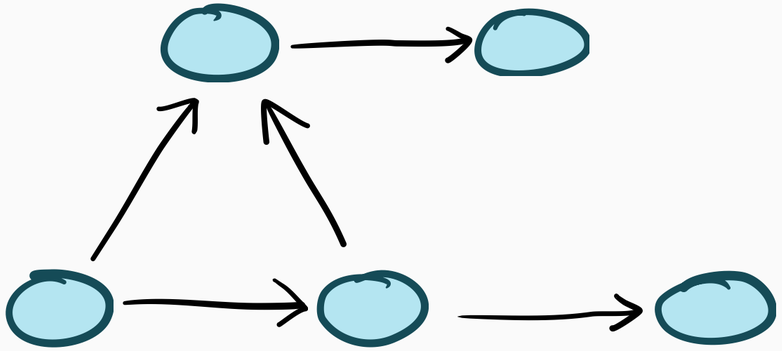

Generalization: Markovian MAB

Pandora's box: one stage

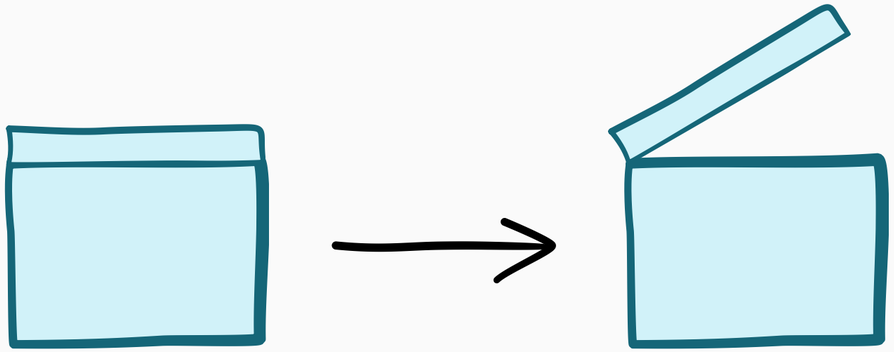

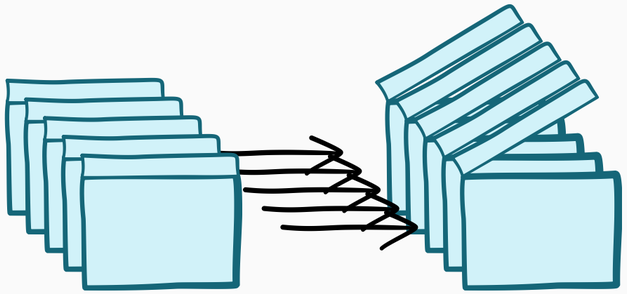

Markovian MAB: general Markov chain

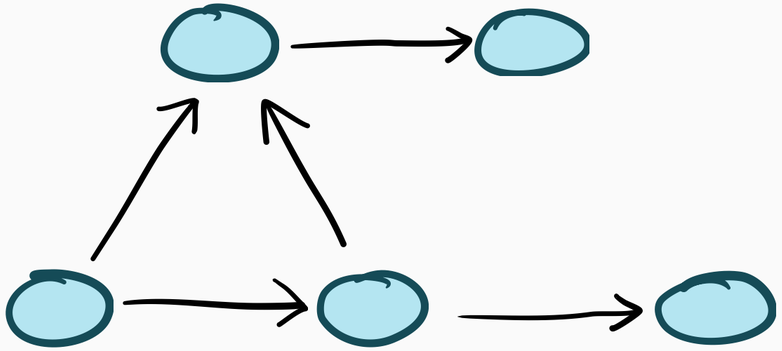

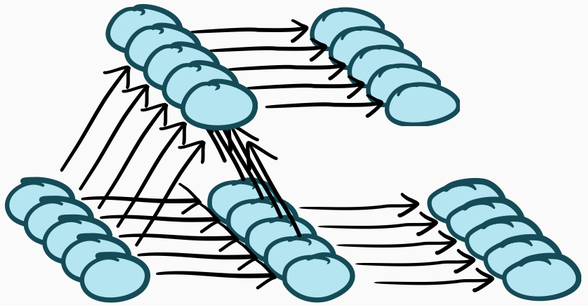

Summary

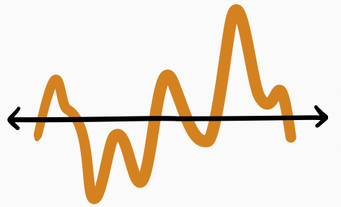

BayesOpt

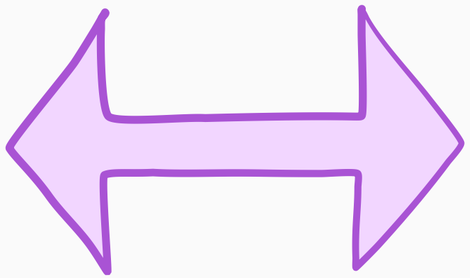

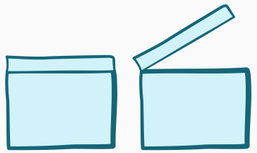

Pandora's box

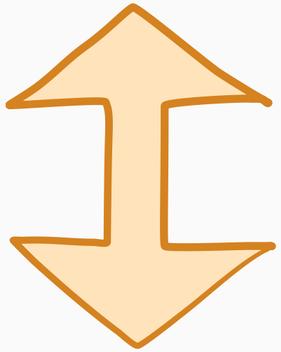

"Fancy" BayesOpt

- Many dimensions
- Heterogeneous costs
- Partial feedback
- Batch evaluations

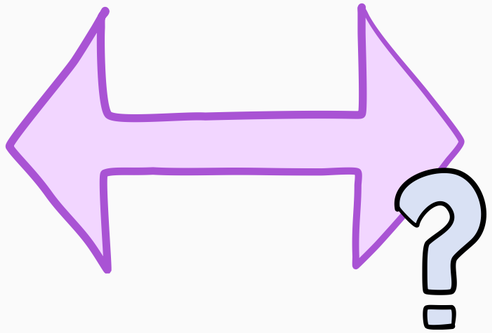

Markovian MAB

general multistage processes

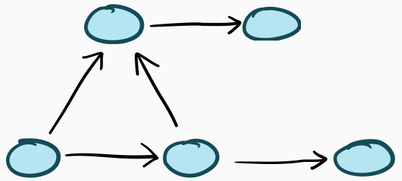

Summary

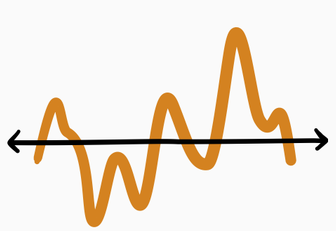

BayesOpt

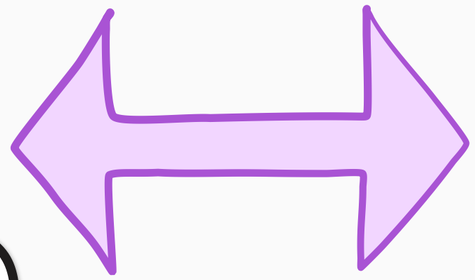

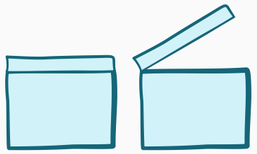

Pandora's box

"Fancy" BayesOpt

- Many dimensions
- Heterogeneous costs
- Partial feedback
- Batch evaluations

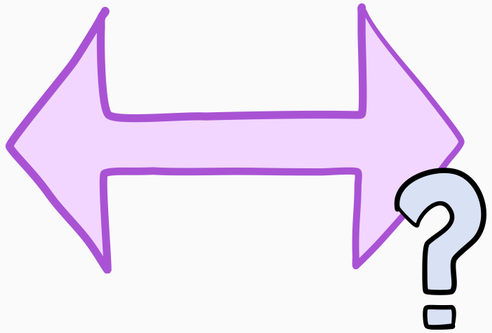

Markovian MAB

general multistage processes

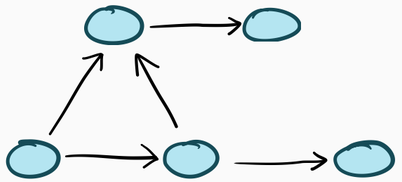

Summary

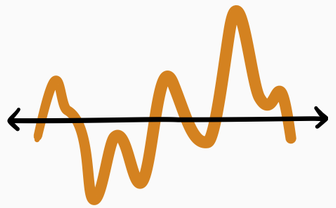

BayesOpt

Gittins is promising!

Pandora's box

Gittins is optimal

"Fancy" BayesOpt

- Many dimensions
- Heterogeneous costs
- Partial feedback
- Batch evaluations

Markovian MAB

general multistage processes