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

Qian Xie (Cornell ORIE)

Joint work with Raul Astudillo, Peter Frazier, Ziv Scully, and Alexander Terenin

ECGI'24

Bayesian Optimization

Goal: optimize expensive-to-evaluate black-box function

∈ decision-making under uncertainty

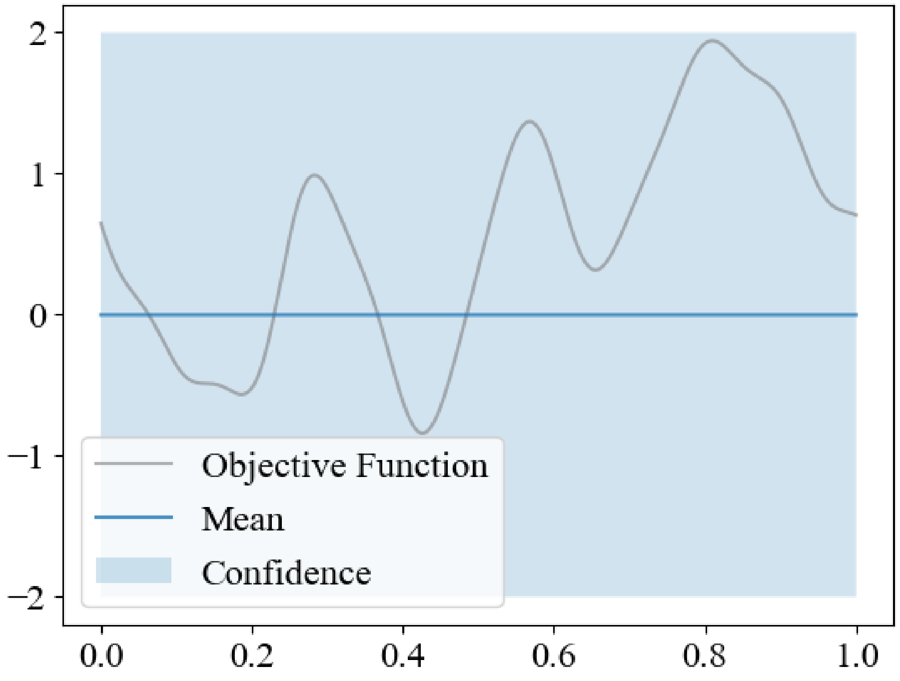

Applications:
Hyperparameter tuning
Drug/material discovery
Experiment design

Bayesian Optimization

Goal: optimize expensive-to-evaluate black-box function

∈ decision-making under uncertainty

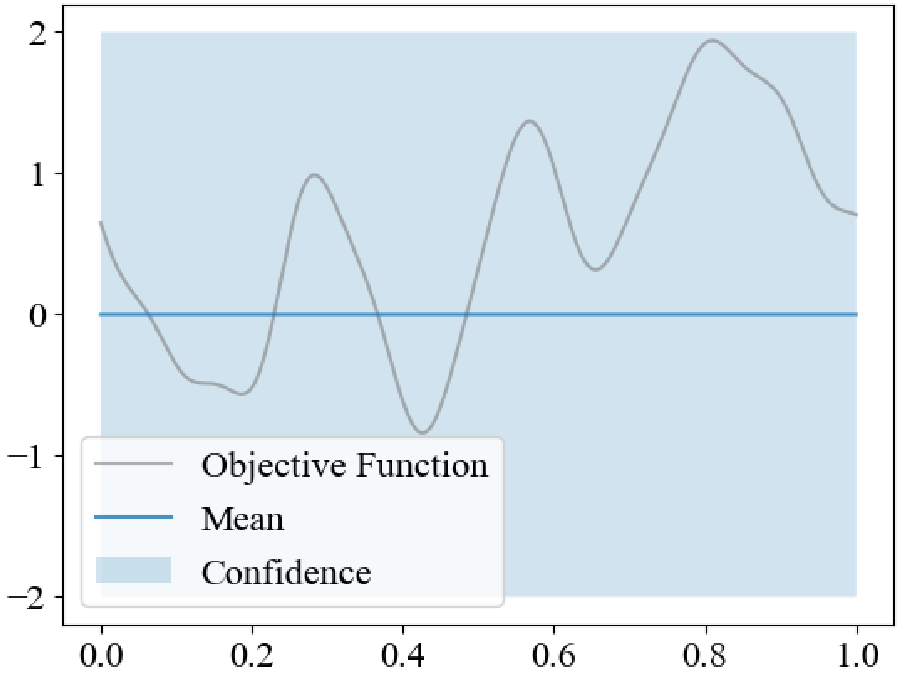

Applications:
Hyperparameter tuning
Drug/material discovery
Experiment design

Bayesian Optimization

Goal: optimize expensive-to-evaluate black-box function

$\in$ decision-making under uncertainty

Applications:
Hyperparameter tuning
Drug/material discovery
Experiment design

x: hyperparameter/configuration

An unknown random function $f: \mathcal{X} \to \mathbb{R}$ drawn from a Gaussian process prior

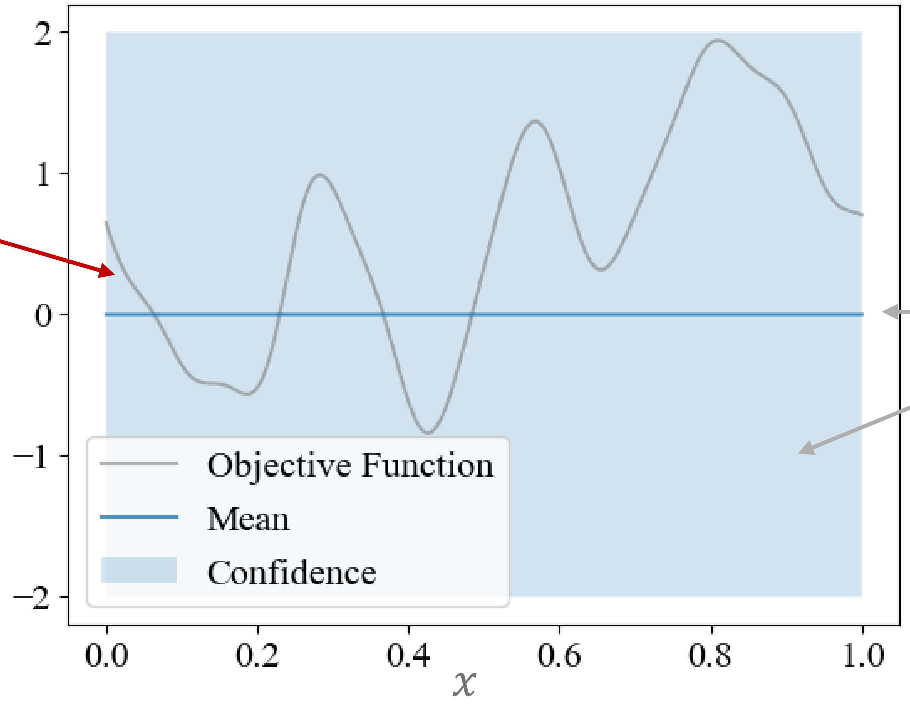

mean: prediction

variance: confidence/uncertainty

Gaussian process: infinite-dimensional generalization of multivariate normal distributions

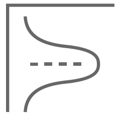

Bayesian Optimization

Goal: optimize expensive-to-evaluate black-box function

Applications:
Hyperparameter tuning
Drug/material discovery
Experiment design

x: hyperparameter/configuration

An unknown random function $f: \mathcal{X} \to \mathbb{R}$ drawn from a Gaussian process prior

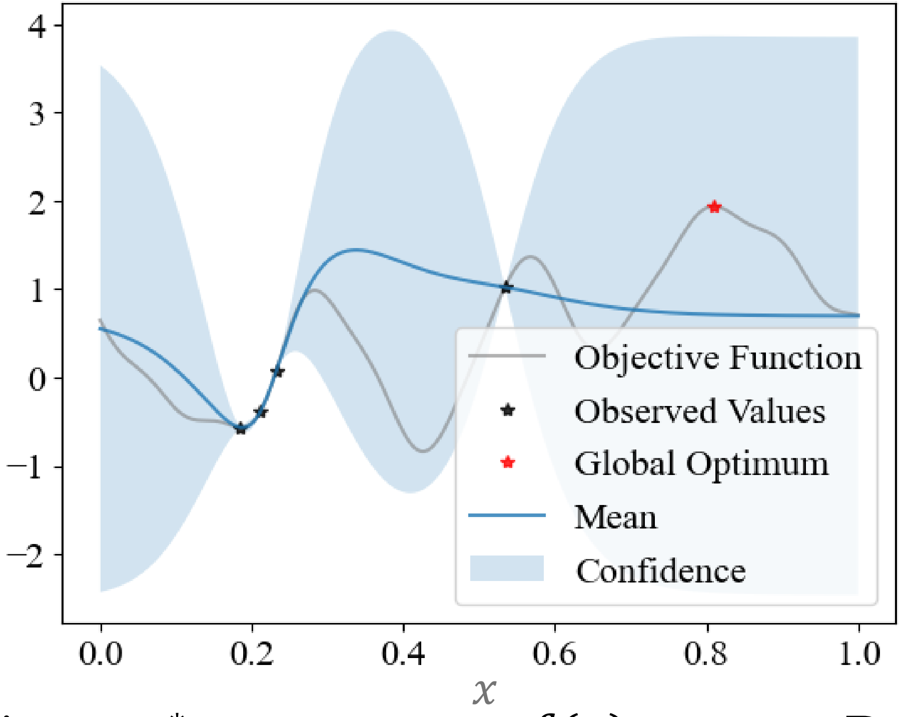

Objective: find global optimum $x^* = \text{argmax}_{x \in \mathcal{X}} f(x)$

Decision: evaluate a set of points

Bayesian Optimization

Goal: optimize expensive-to-evaluate black-box function

An unknown random function $f: \mathcal{X} \to \mathbb{R}$ drawn from a Gaussian process prior

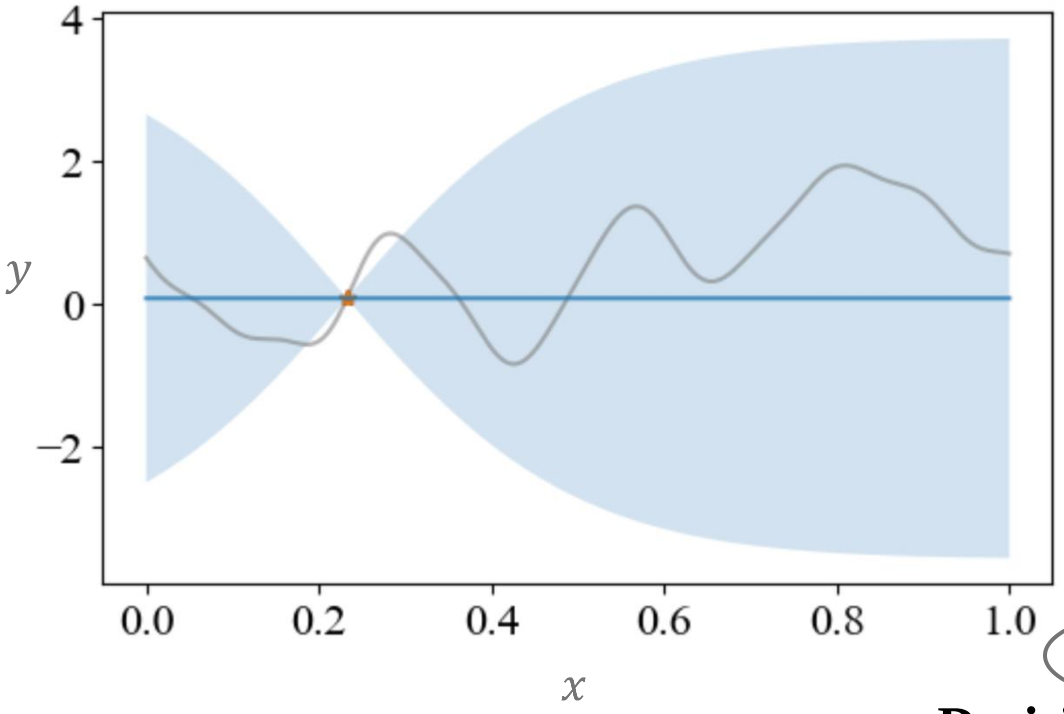

Applications:
Hyperparameter tuning
Drug/material discovery
Experiment design

x: hyperparameter/configuration

Decision: evaluate a set of points adaptively

Bayesian Optimization

Goal: optimize expensive-to-evaluate black-box function

An unknown random function $f: \mathcal{X} \to \mathbb{R}$ drawn from a Gaussian process prior

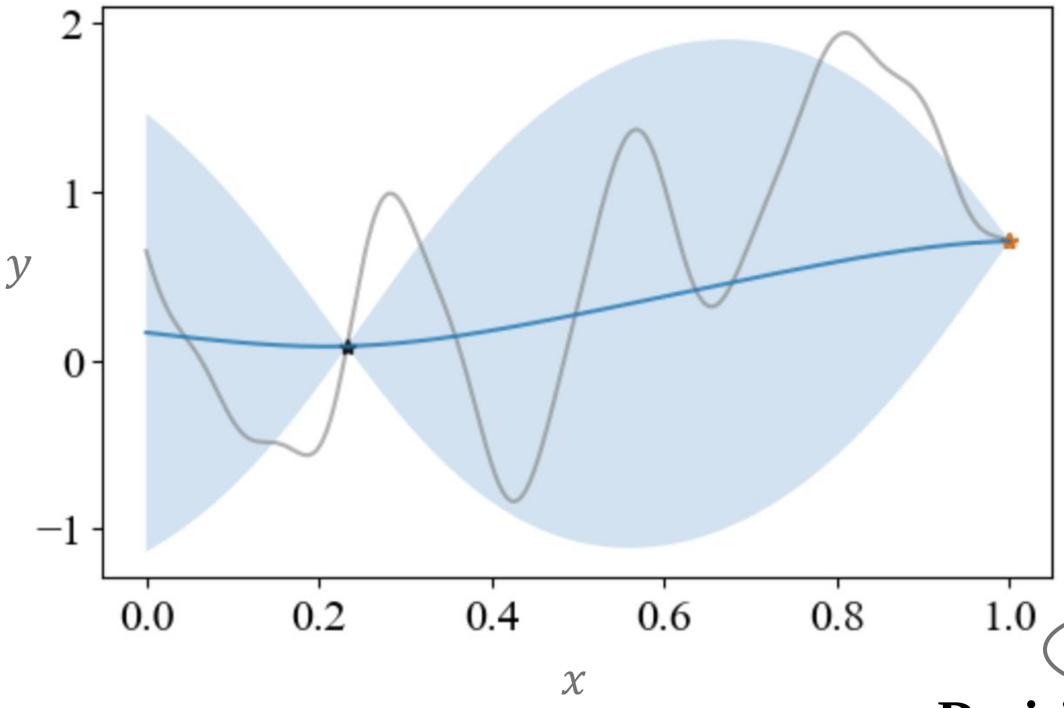

Applications:
Hyperparameter tuning
Drug/material discovery
Experiment design

x: hyperparameter/configuration

adaptively

Decision: evaluate a set of points

Bayesian Optimization

Goal: optimize expensive-to-evaluate black-box function

An unknown random function $f: \mathcal{X} \to \mathbb{R}$ drawn from a Gaussian process prior

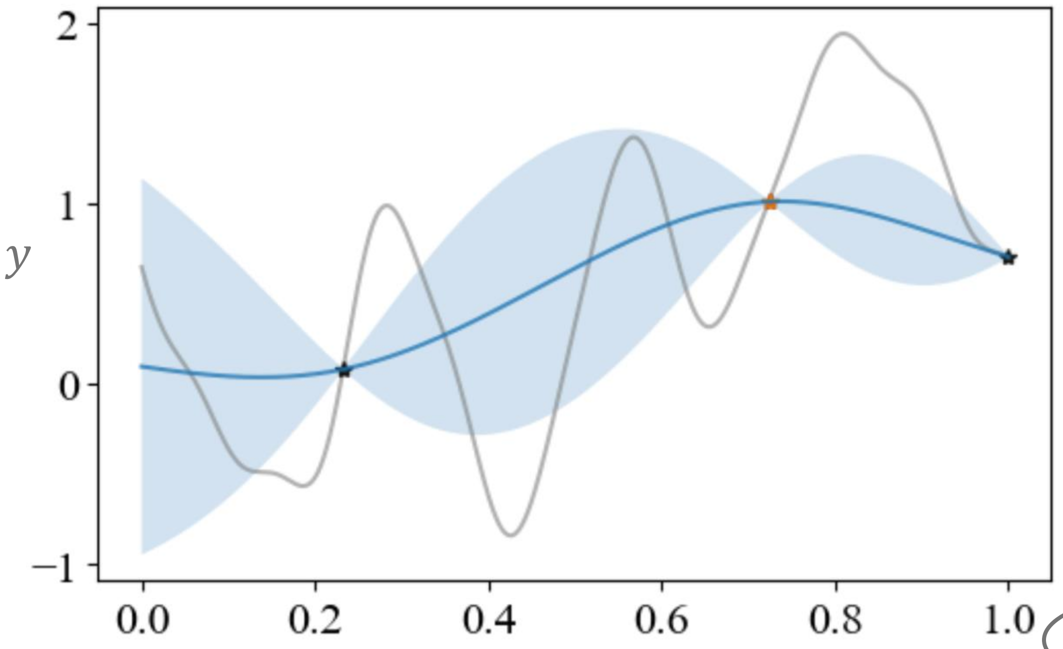

y

x

Applications:
Hyperparameter tuning
Drug/material discovery
Experiment design

x: hyperparameter/configuration

adaptively

Decision: evaluate a set of points

Bayesian Optimization

Goal: optimize expensive-to-evaluate black-box function

An unknown random function $f: \mathcal{X} \to \mathbb{R}$ drawn from a Gaussian process prior

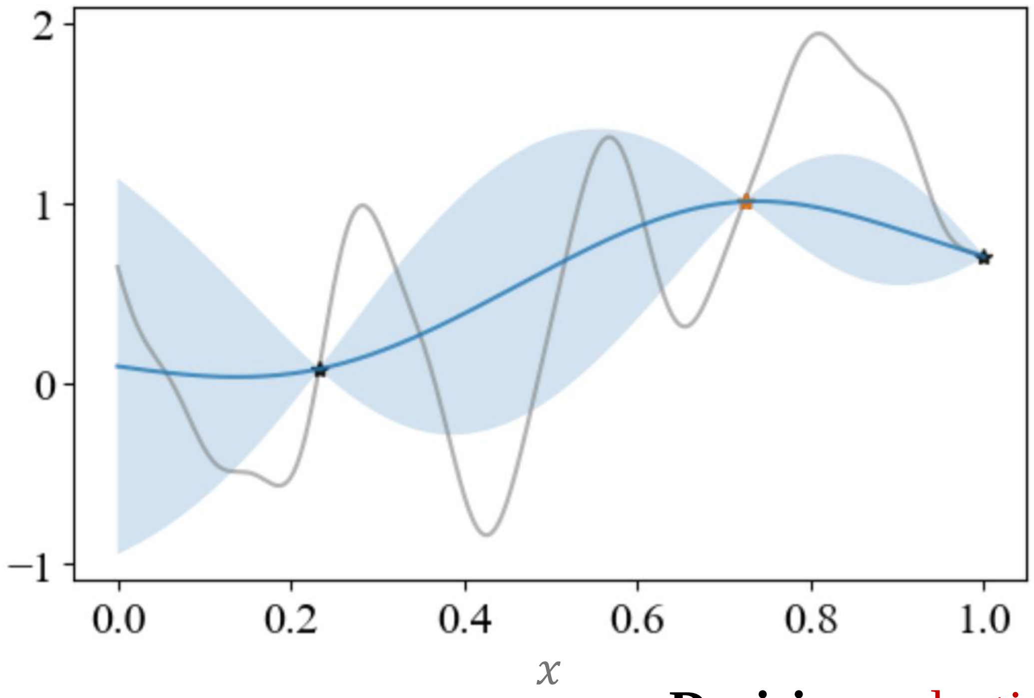

Applications:
Hyperparameter tuning
Drug/material discovery
Experiment design

x: hyperparameter/configuration

Decision: adaptively evaluate a set of points

$$x_1, x_2, \ldots, x_T \in \mathcal{X}$$

T: time budget

Bayesian Optimization

Goal: optimize expensive-to-evaluate black-box function

An unknown random function $f: \mathcal{X} \to \mathbb{R}$ drawn from a Gaussian process prior

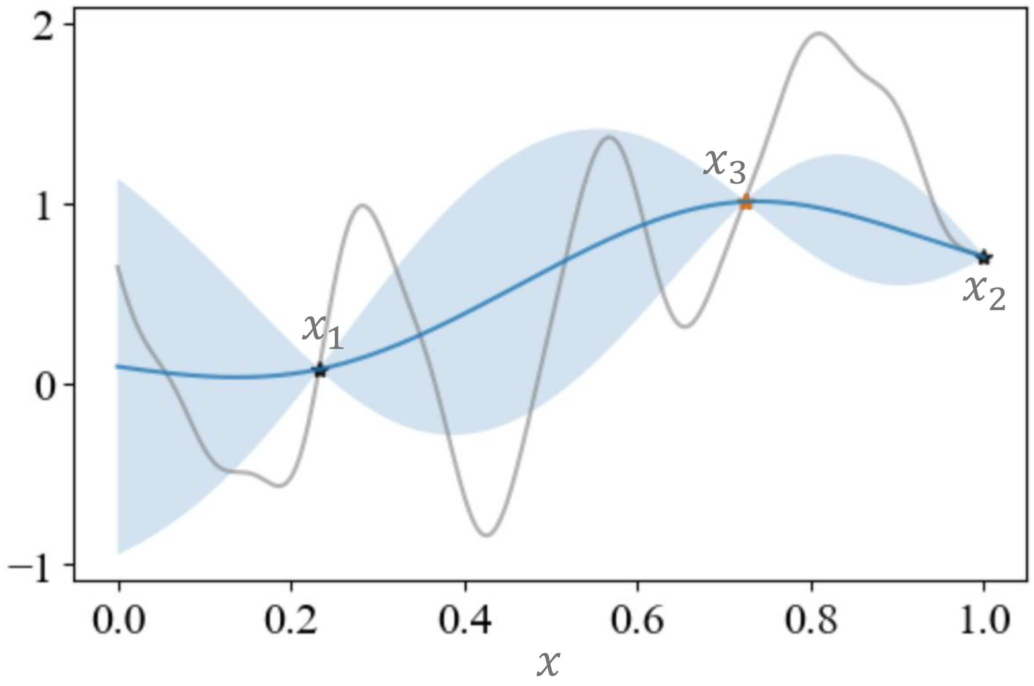

Applications:
Hyperparameter tuning
Drug/material discovery
Experiment design

x: hyperparameter/configuration

Objective: optimize best observed value at time T

$$\sup_{\text{policy}} \mathbb{E} \max_{t=1,2,\ldots,T} f(x_t)$$

Decision: adaptively evaluate a set of points

$$x_1, x_2, \ldots, x_T \in \mathcal{X}$$

T: time budget

Why is it hard?

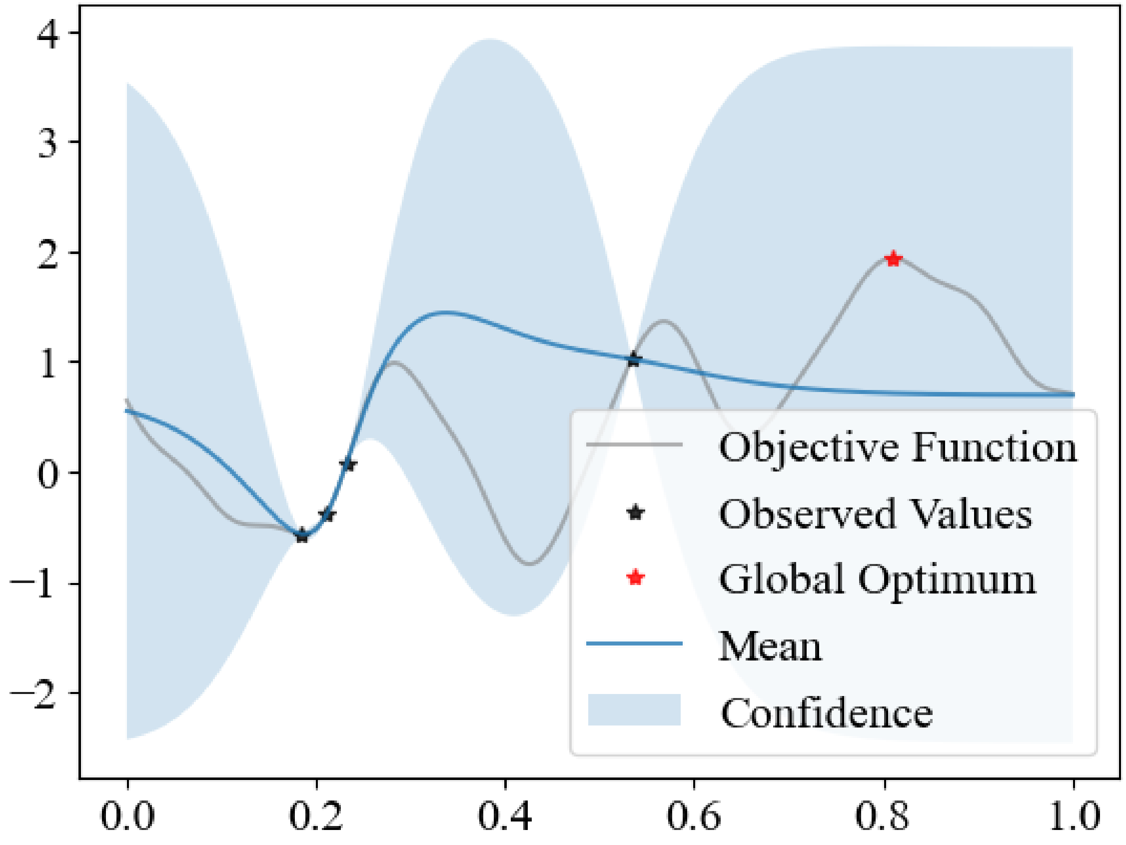

Why is it hard?

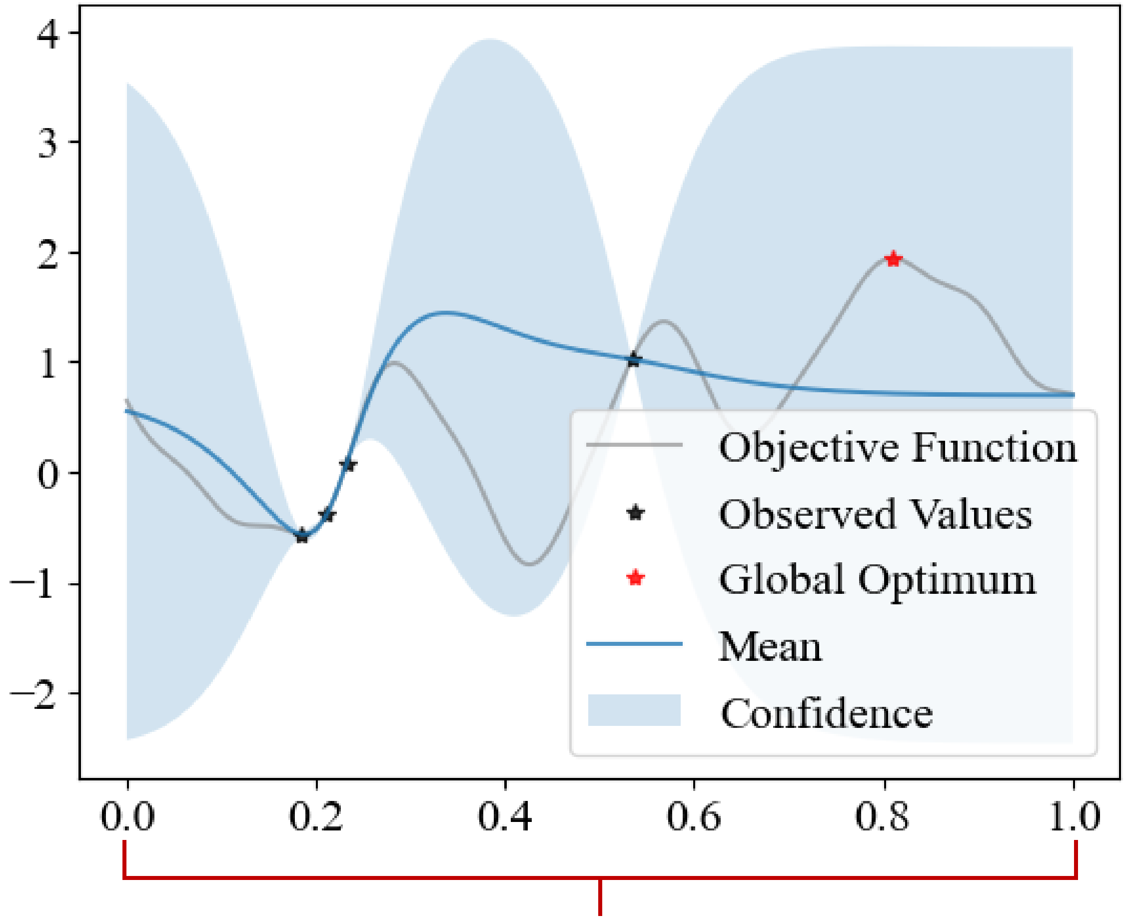

Continuous search domain

Why is it hard?

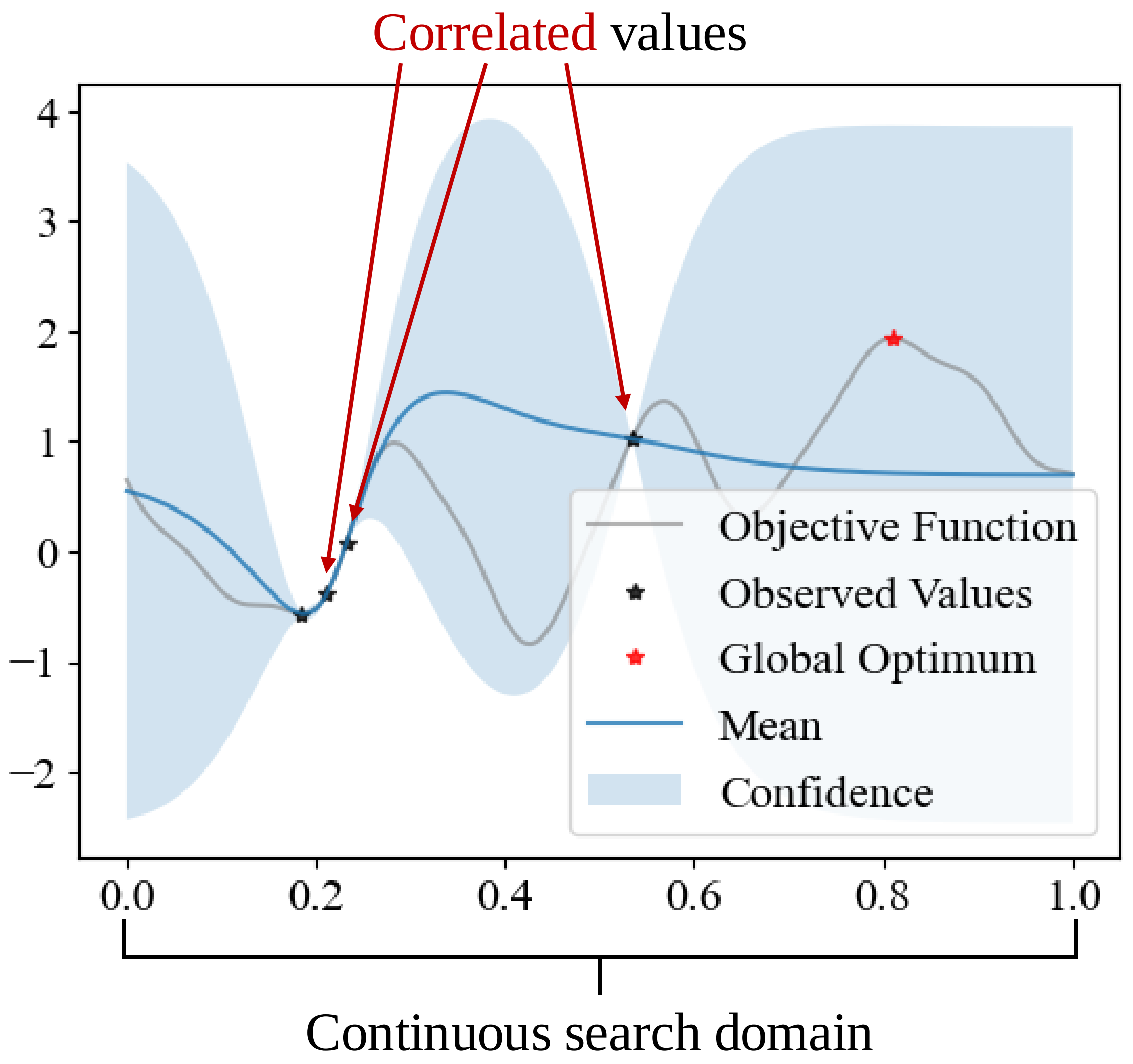

Why is it hard?

Hard budget constraint

~~$t = 1$~~

~~$t = 2$~~

~~$t = 3$~~

~~$t = 4$~~

$\vdots$

$t = T$

Correlated values

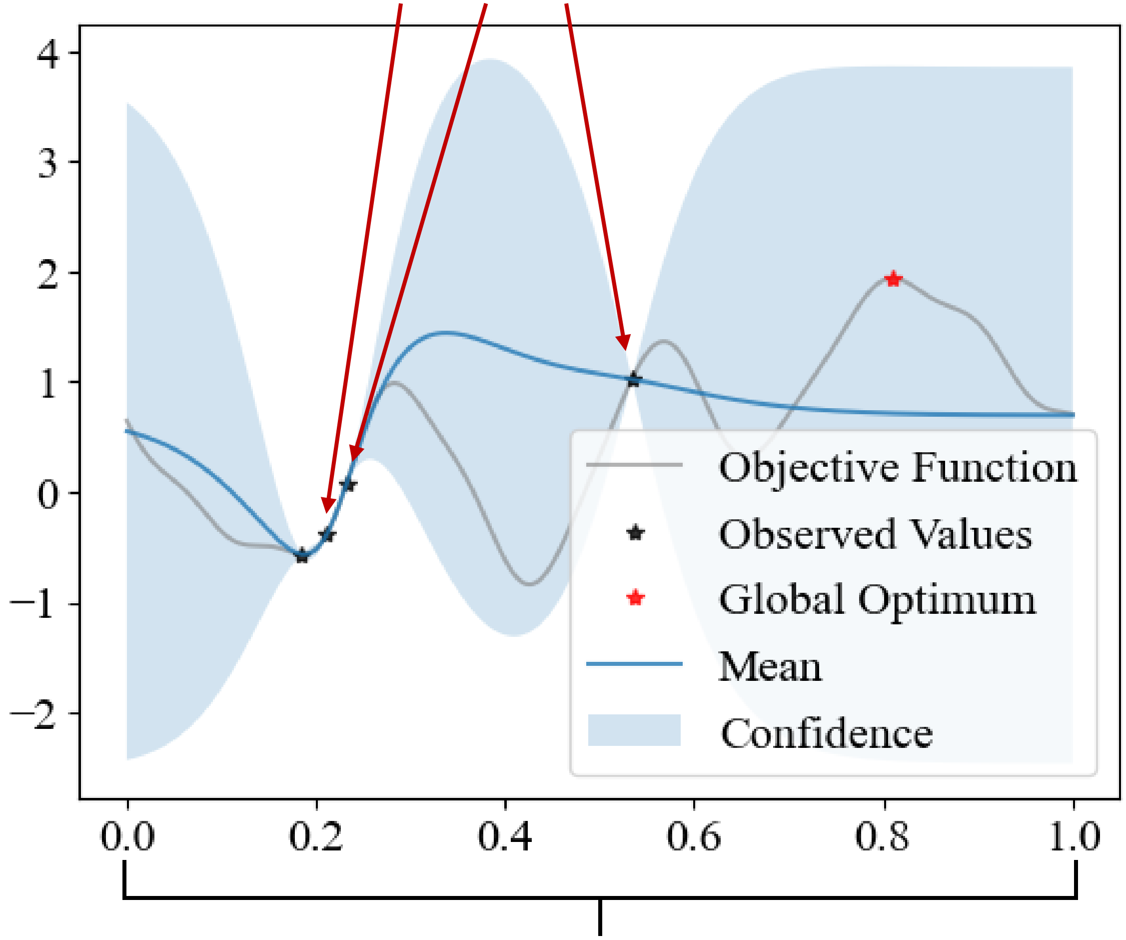

Continuous search domain

Why is it hard?

Hard budget constraint

~~$t = 1$~~

~~$t = 2$~~

~~$t = 3$~~

~~$t = 4$~~

⋮

$t = T$

Correlated values

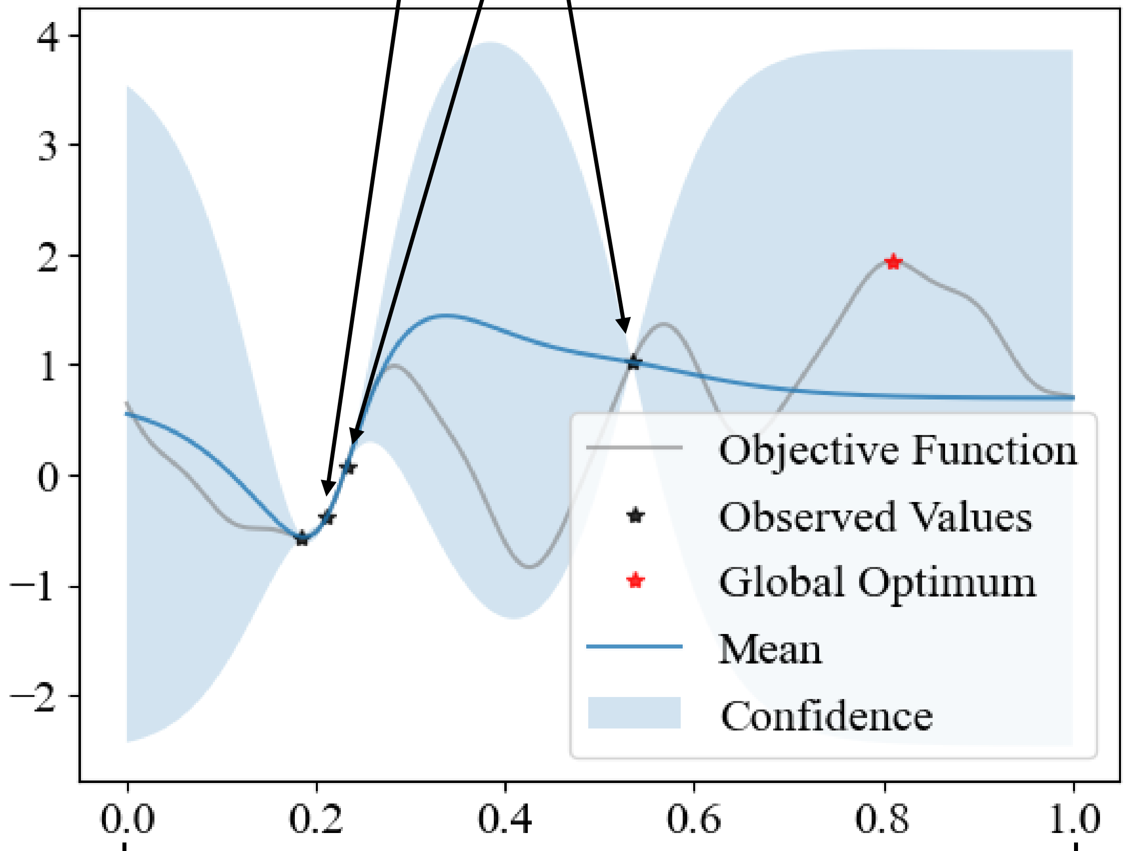

Continuous search domain

Evaluation costs handling

uniform

heterogeneous

Why is it hard?

Hard budget constraint

~~$t = 1$~~

~~$t = 2$~~

~~$t = 3$~~

~~$t = 4$~~

$\vdots$

$t = T$

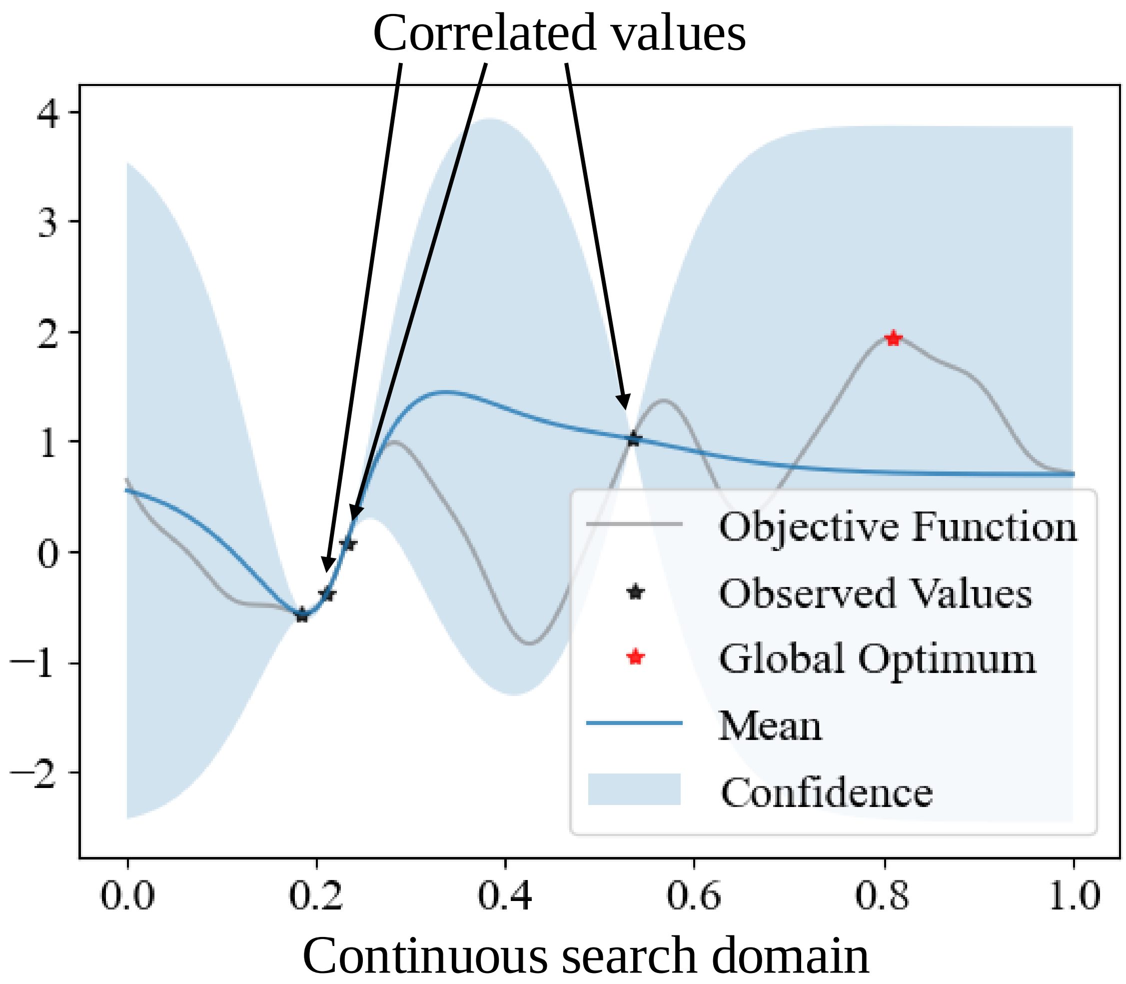

Evaluation costs handling

uniform

heterogeneous

$\Longrightarrow$ Optimal policy unknown!

Why is it hard?

Hard budget constraint

~~$t = 1$~~

~~$t = 2$~~

~~$t = 3$~~

~~$t = 4$~~

⋮

$t = T$

Correlated values

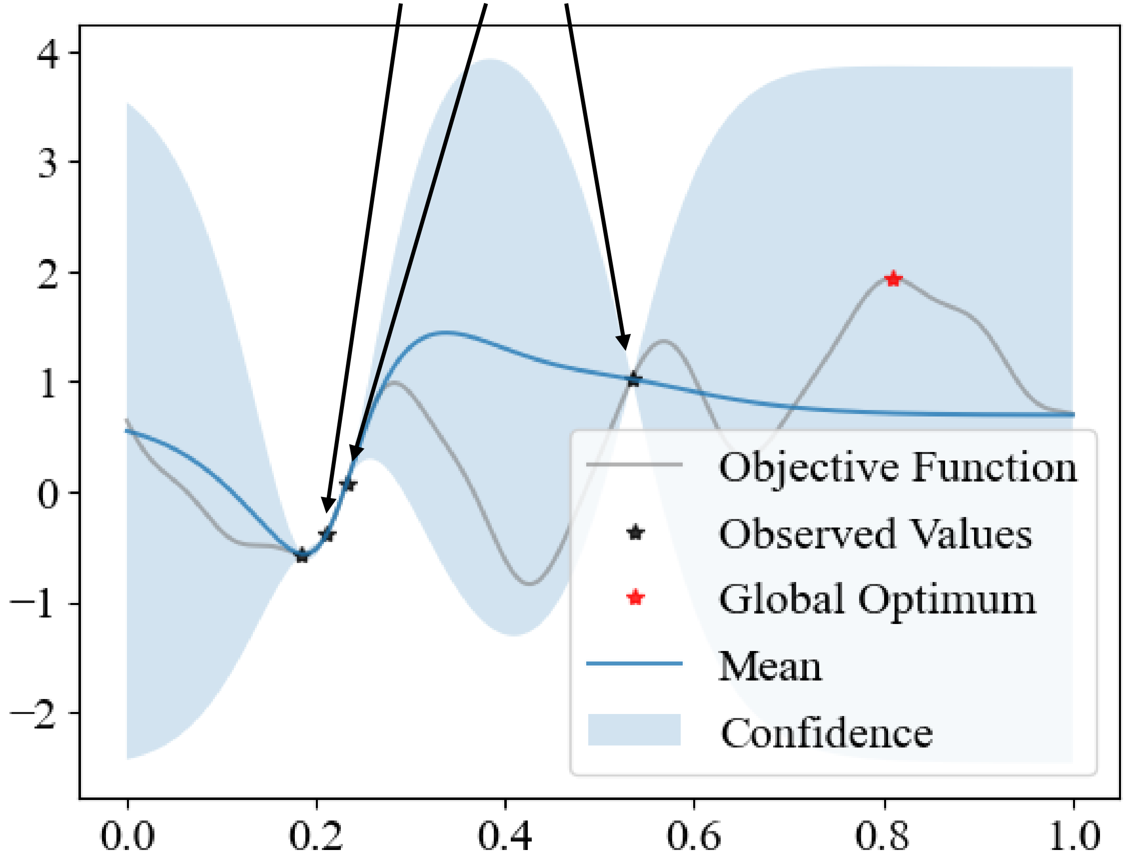

Continuous search domain

Evaluation costs handling

uniform

heterogeneous

Can we convert it to a solvable problem?

Bayesian Optimization

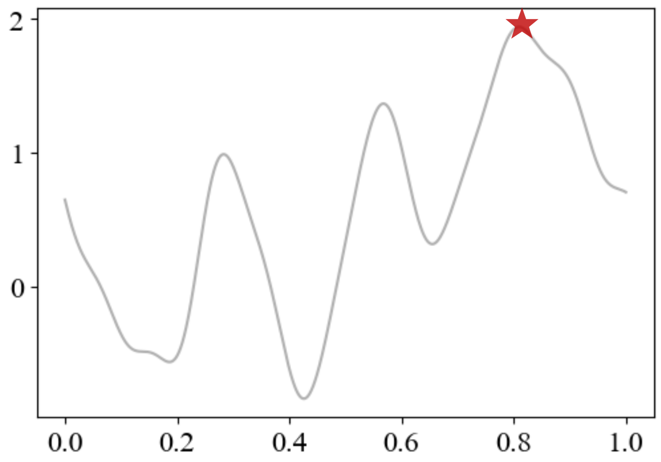

Continuous

Correlated

Hard budget constraint

Bayesian Optimization

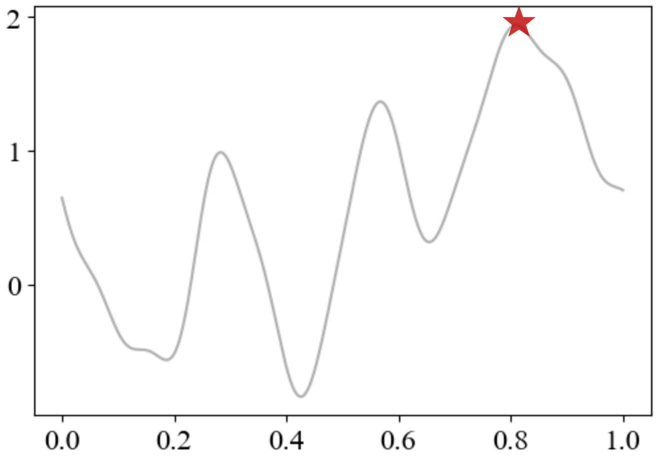

⟹

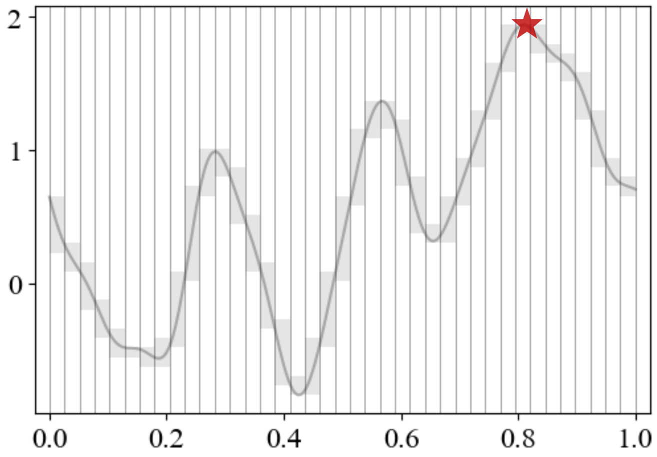

Continuous ⟹ Discrete

Correlated

Hard budget constraint

Bayesian Optimization

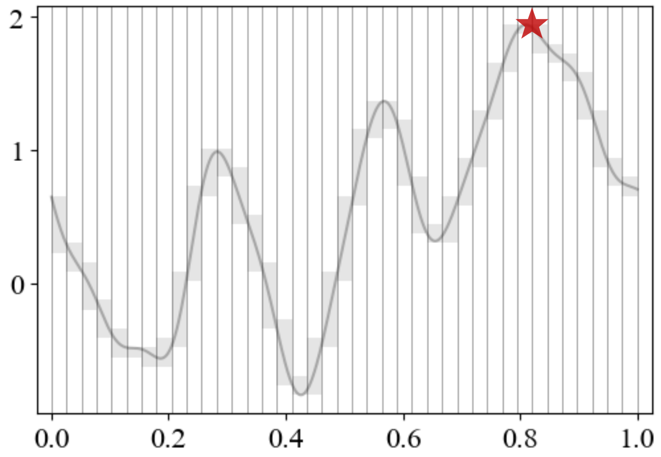

Continuous ⟹ Discrete

Correlated

Hard budget constraint

Bayesian Optimization

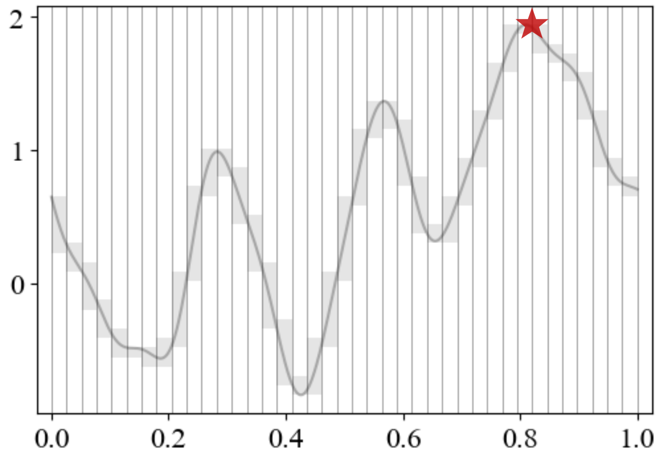

⟹

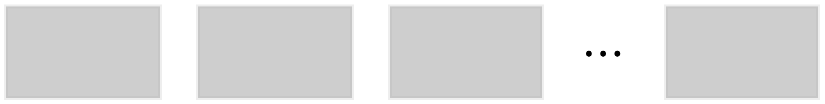

Continuous ⟹ Discrete

Correlated ⟹ Independent

Hard budget constraint

Bayesian Optimization

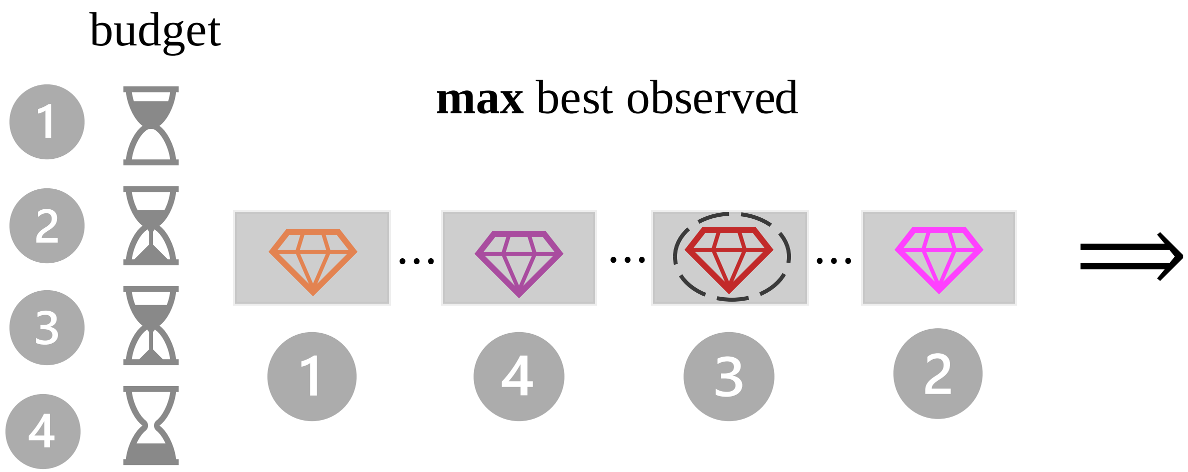

Continuous $\Longrightarrow$ Discrete

Correlated $\Longrightarrow$ Independent

Hard budget constraint

Bayesian Optimization

budget | **max** best observed | **max** (best observed – costs) | costs

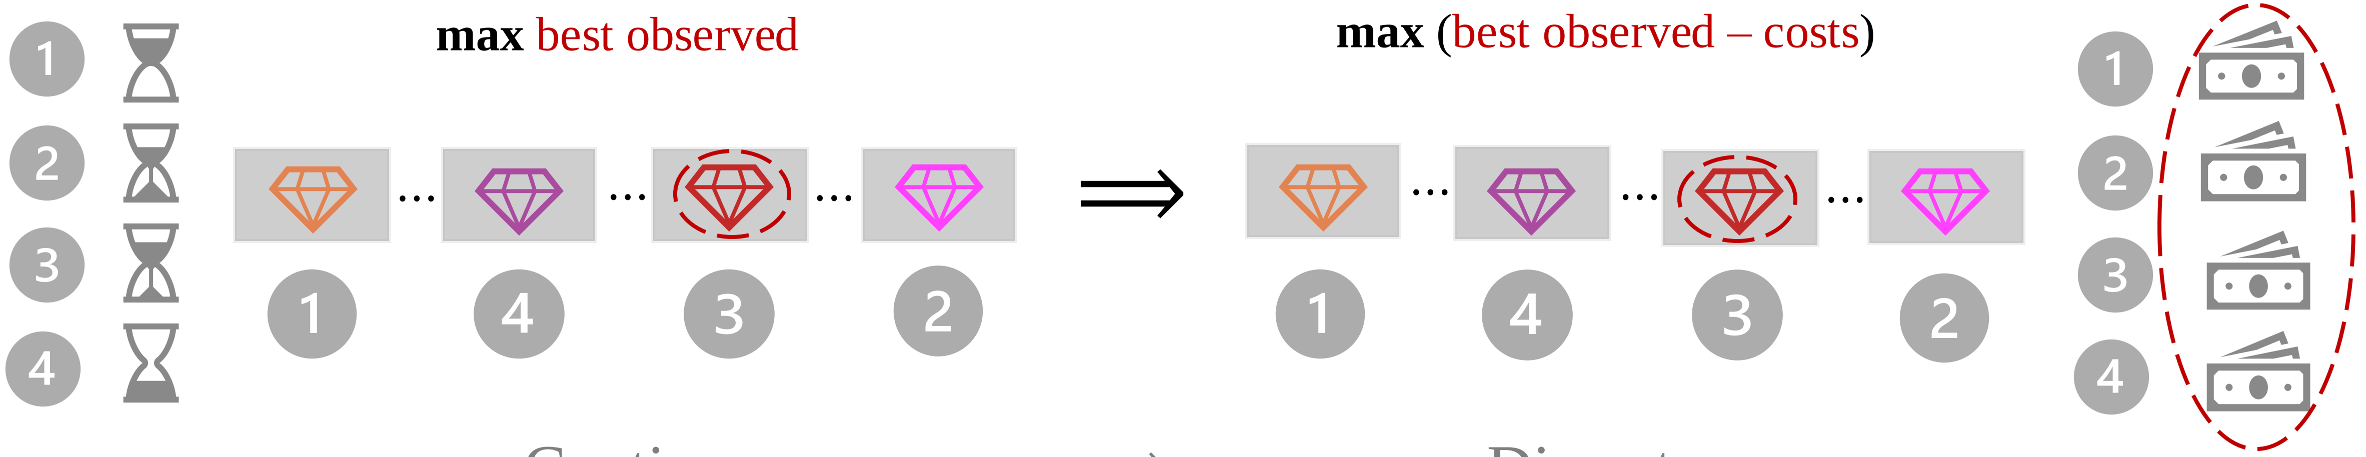

Continuous $\Longrightarrow$ Discrete

Correlated $\Longrightarrow$ Independent

Hard budget constraint $\Longrightarrow$ (Lagrangian relaxation) Cost per sample

Bayesian Optimization ⟹ Pandora's Box

[Weitzman'79]

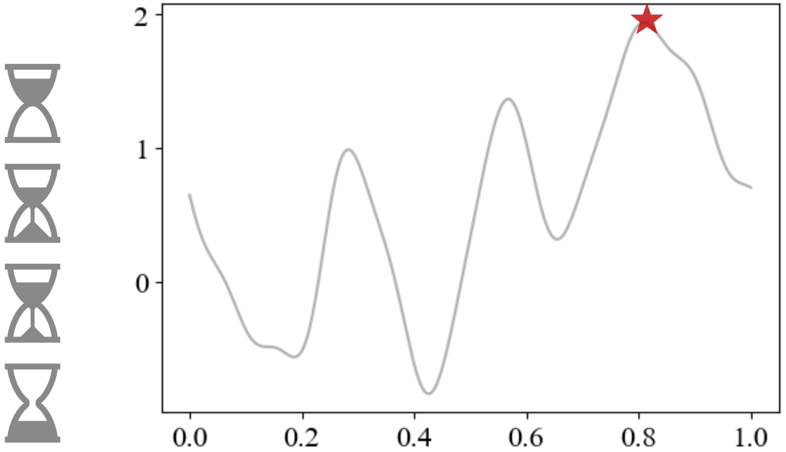

⟹

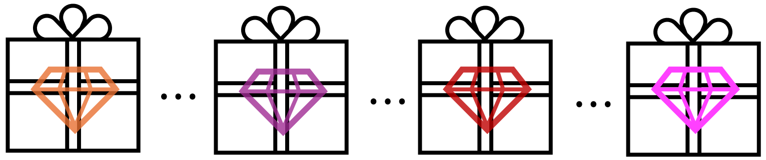

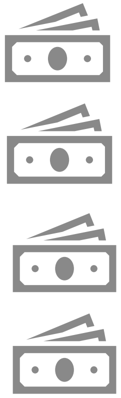

Continuous	⟹	Discrete
Correlated	⟹	Independent
Hard budget constraint	⟹	Cost per sample

Pandora's Box

$t = 0$

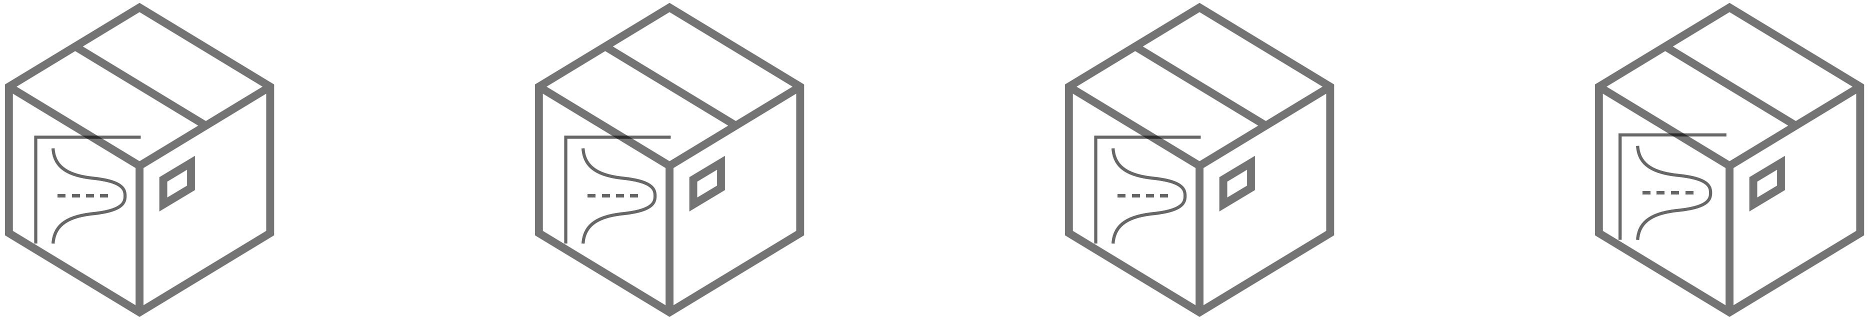

Objective: maximize net utility

Decision: adaptively evaluate a random number of boxes

Pandora’s Box

$t = 1$

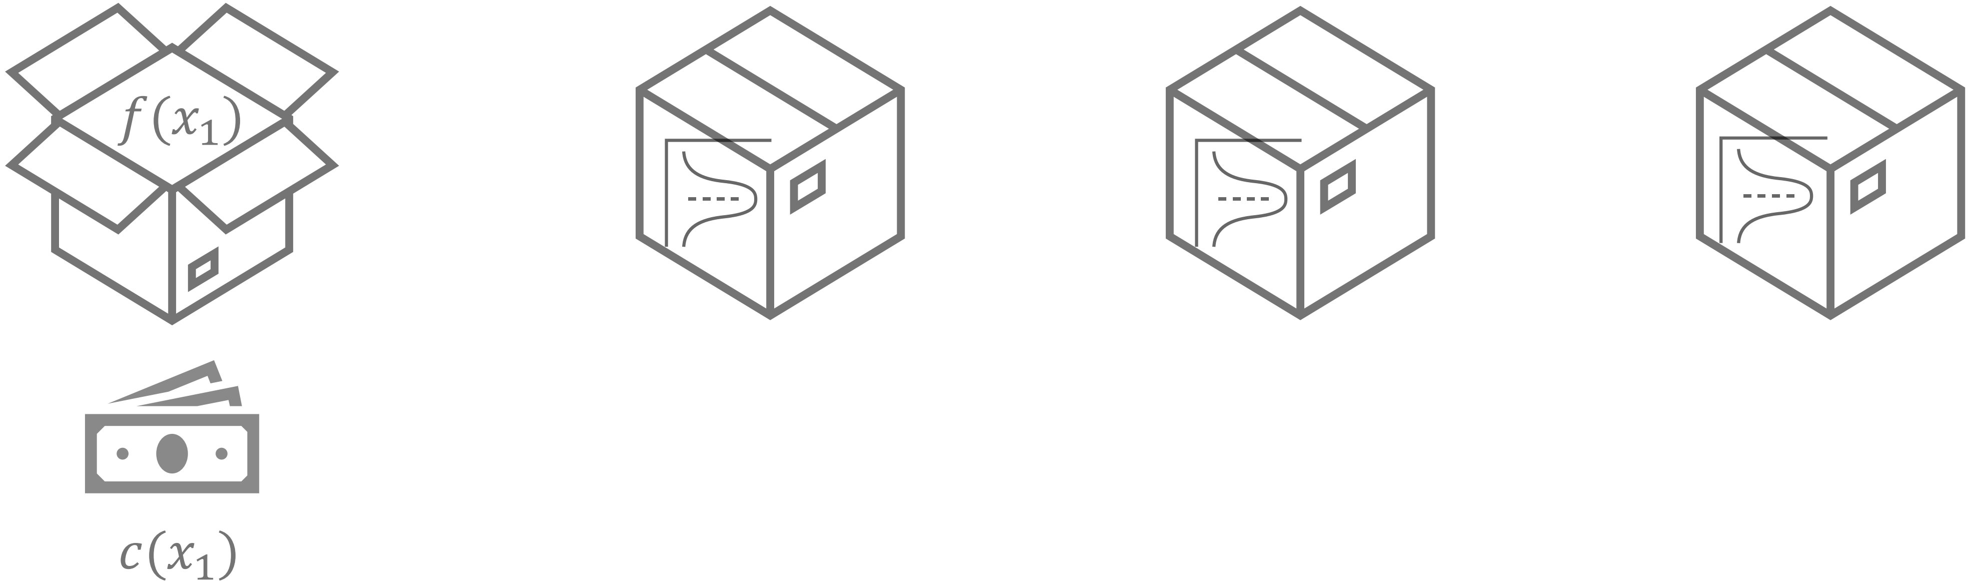

Objective: maximize net utility

Decision: adaptively evaluate a random number of boxes

Pandora's Box

$t = 2$

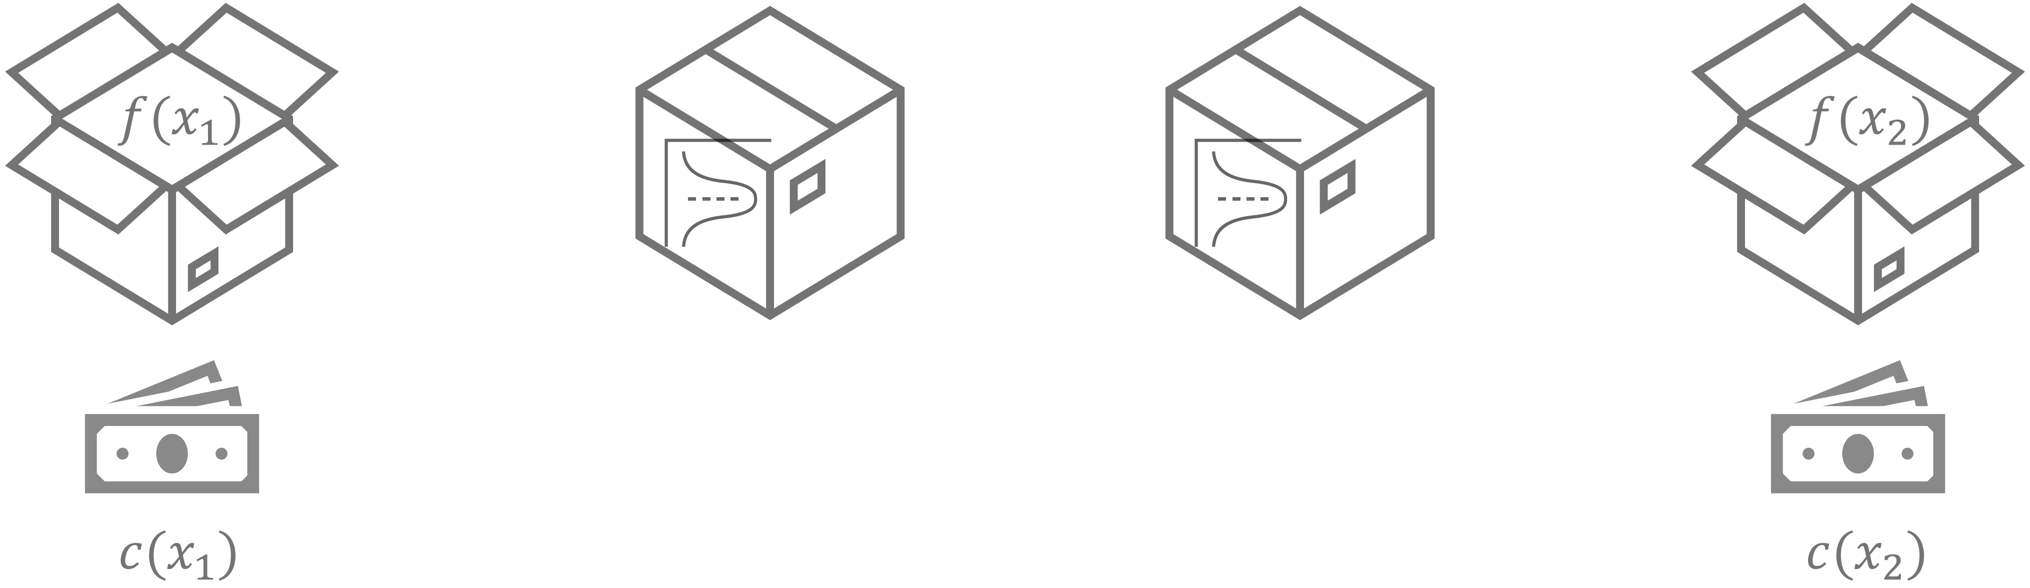

Objective: maximize net utility

Decision: adaptively evaluate a random number of boxes

Pandora’s Box

$t = 3$

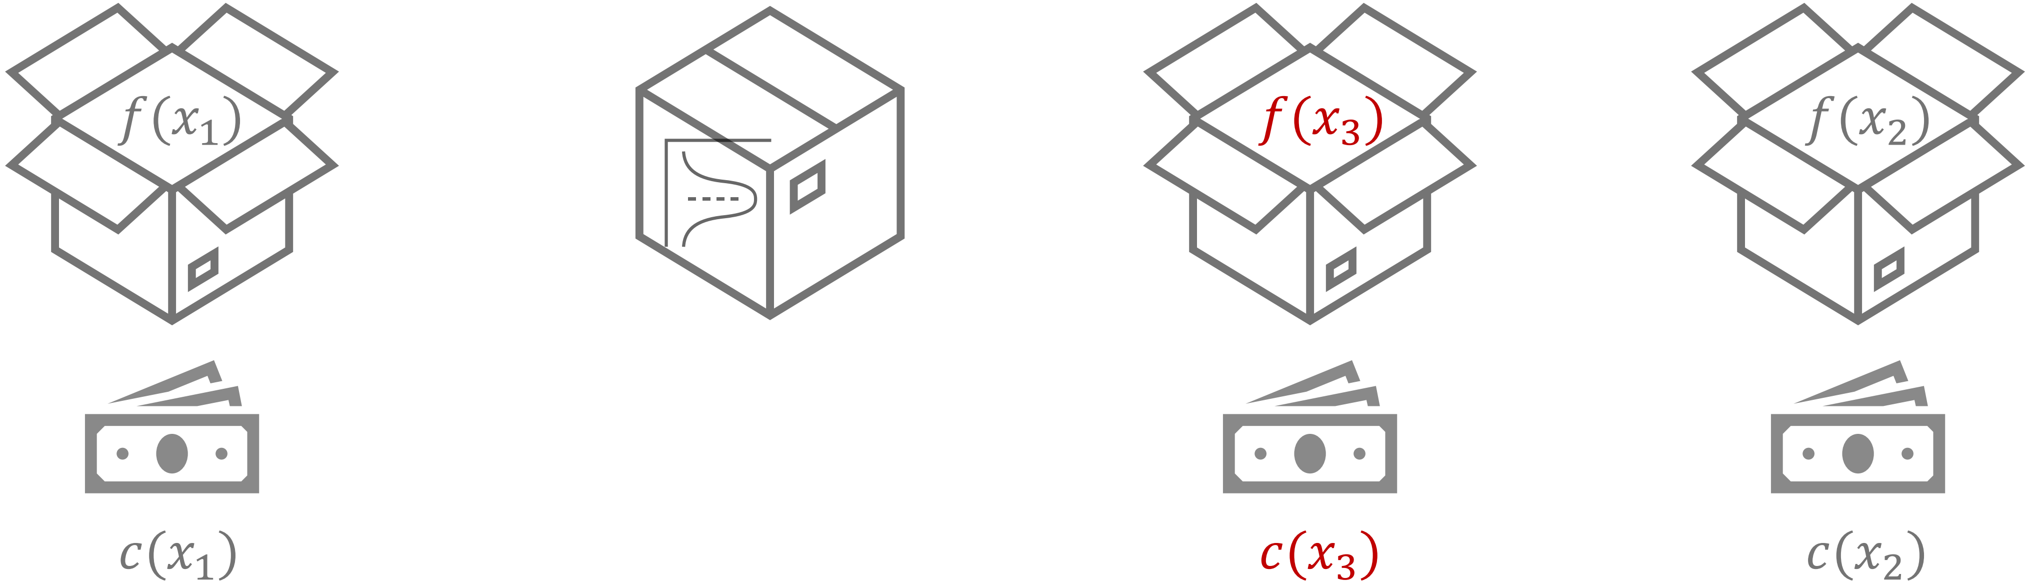

Objective: maximize net utility

Decision: adaptively evaluate a random number of boxes

Pandora's Box

$t = 3$

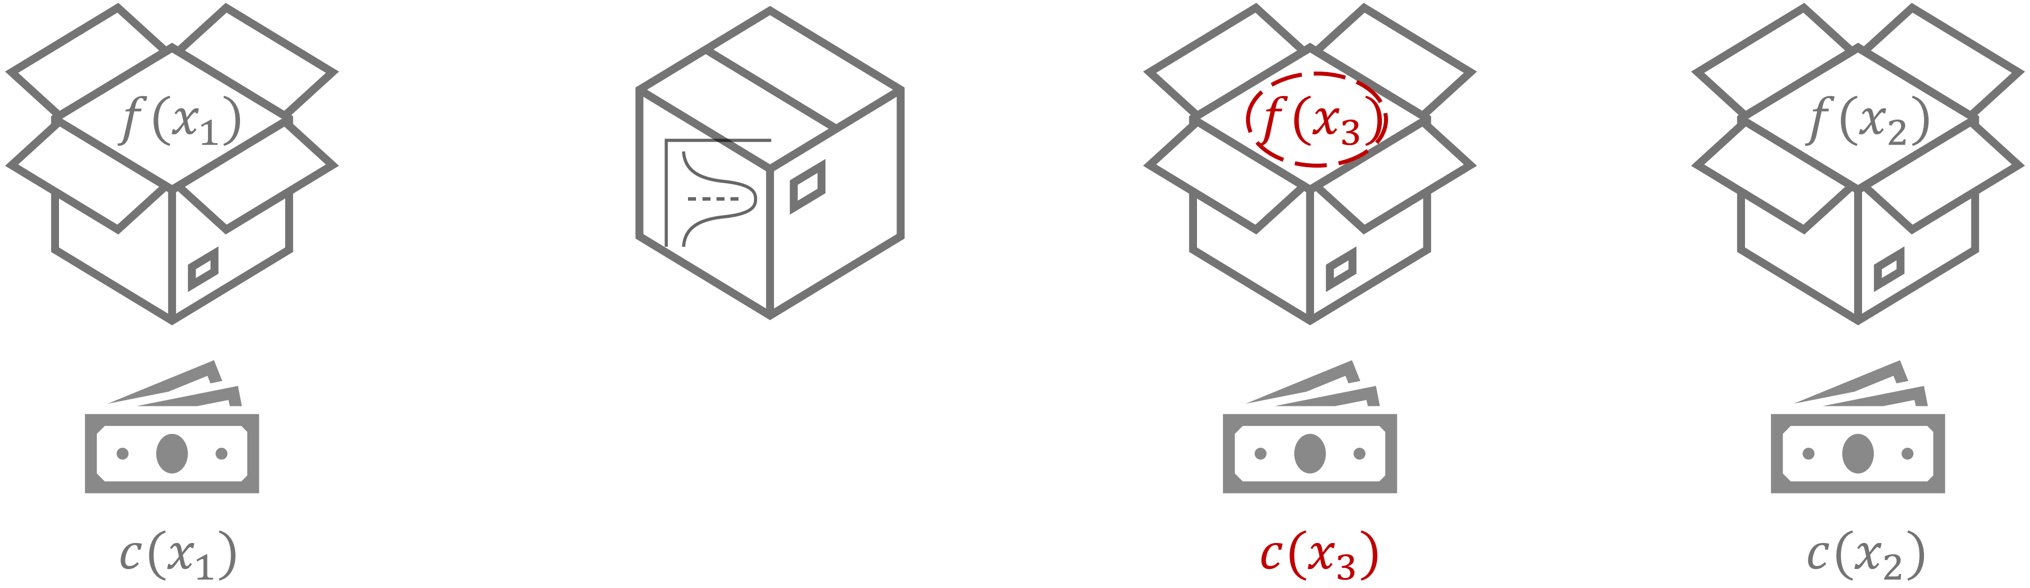

Objective: maximize net utility

Decision: adaptively evaluate a random number of boxes

max (best observed value – total costs)

Pandora’s Box

$t = 3$

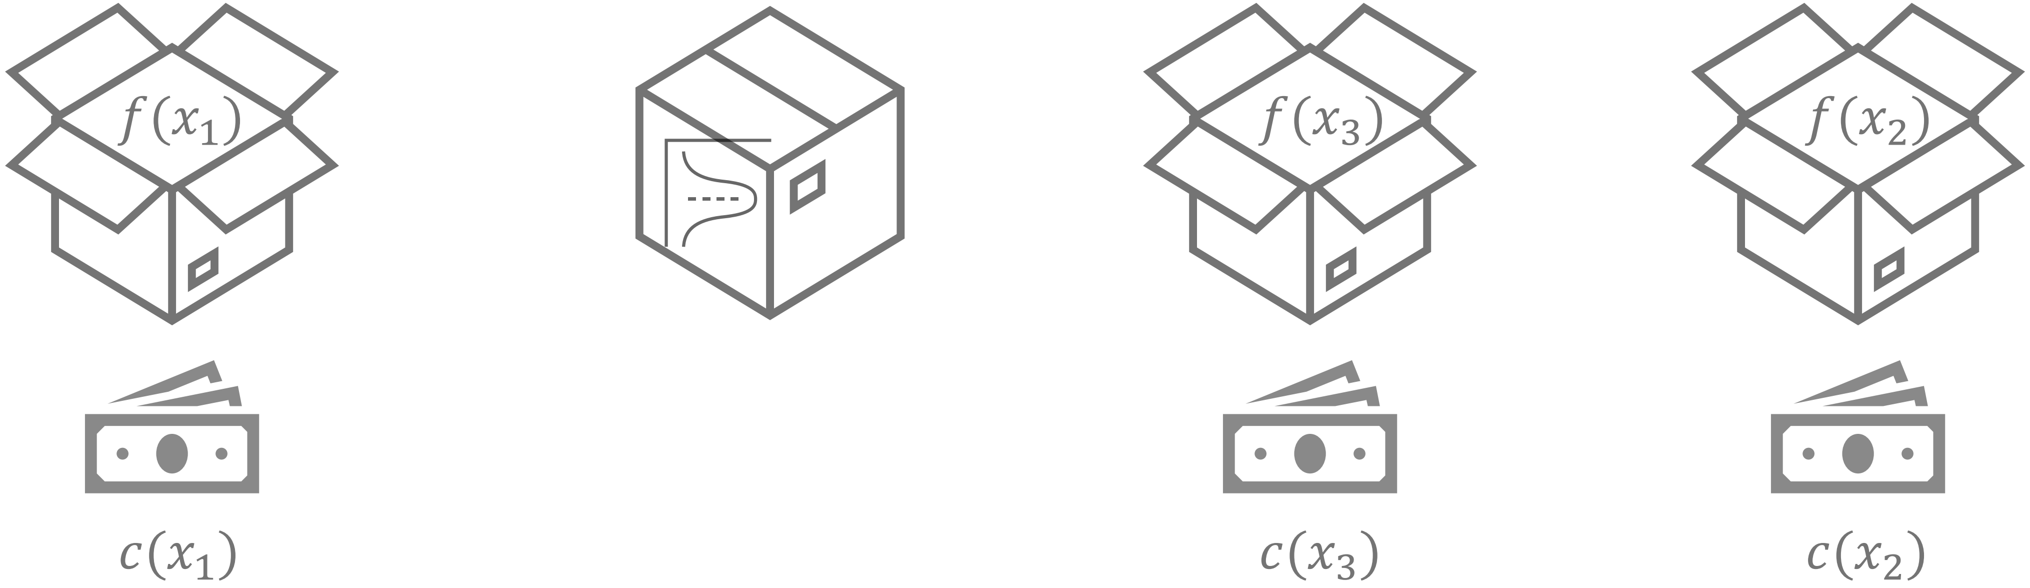

Objective: maximize net utility

$$\sup_{\text{policy}} \mathbb{E}\left(\max_{t=1,2,\dots,T} f(x_t) - \sum_{t=1}^{T} c(x_t) \right)$$

Decision: adaptively evaluate a random number of boxes

$$x_1, x_2, \dots, x_T \in \mathcal{X}$$

$\mathcal{X}$: discrete

T: random stopping time

Naïve Greedy policy can fail [Singla'18]

Naïve Greedy policy

$f(x)$

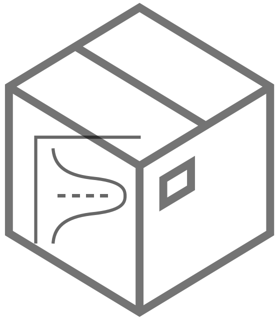

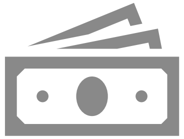

$c(x)$

vs.

Inspection rule: $\text{argmax}_x \left(\text{EI}_f(x; y_{\text{best}}) - c(x)\right)$

expected improvement - cost

Stopping rule: $\text{EI}_f(x; y_{\text{best}}) \leq c(x), \forall x \in \mathcal{X}$

expected improvement $\leq$ cost

y_{best}: current best observed value

$\text{EI}_f(x; y) := \mathbb{E}[(f(x) - y)^+]$: expected improvement of $f(x)$ over y

Naïve Greedy policy can fail [Singla'18]

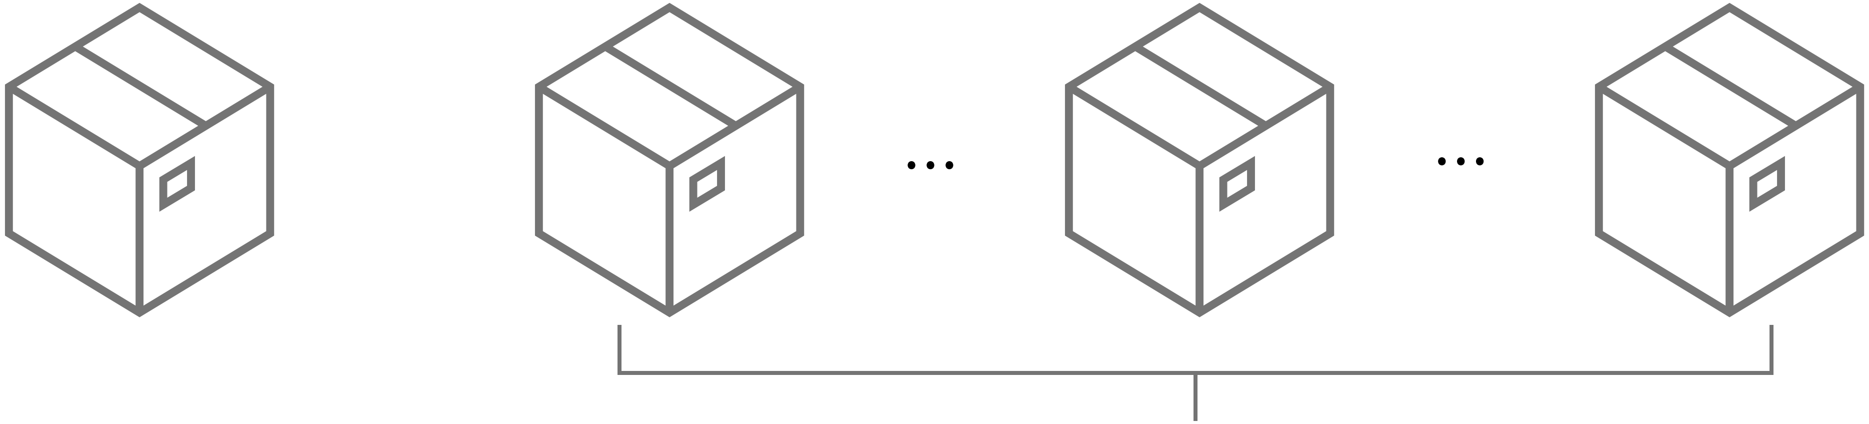

$f(1) = 200 \text{ w.p. } 1$

$c(1) = 198$

$$f(x) = \begin{cases} 200 & \text{w.p. } 0.01 \\ 0 & \text{otherwise} \end{cases}, c(x) = 1, x \in \{2, 3, \dots, 1000\}$$

Naïve Greedy policy can fail [Singla'18]

$t = 0$ $\qquad$ $y_{\text{best}} = 0$

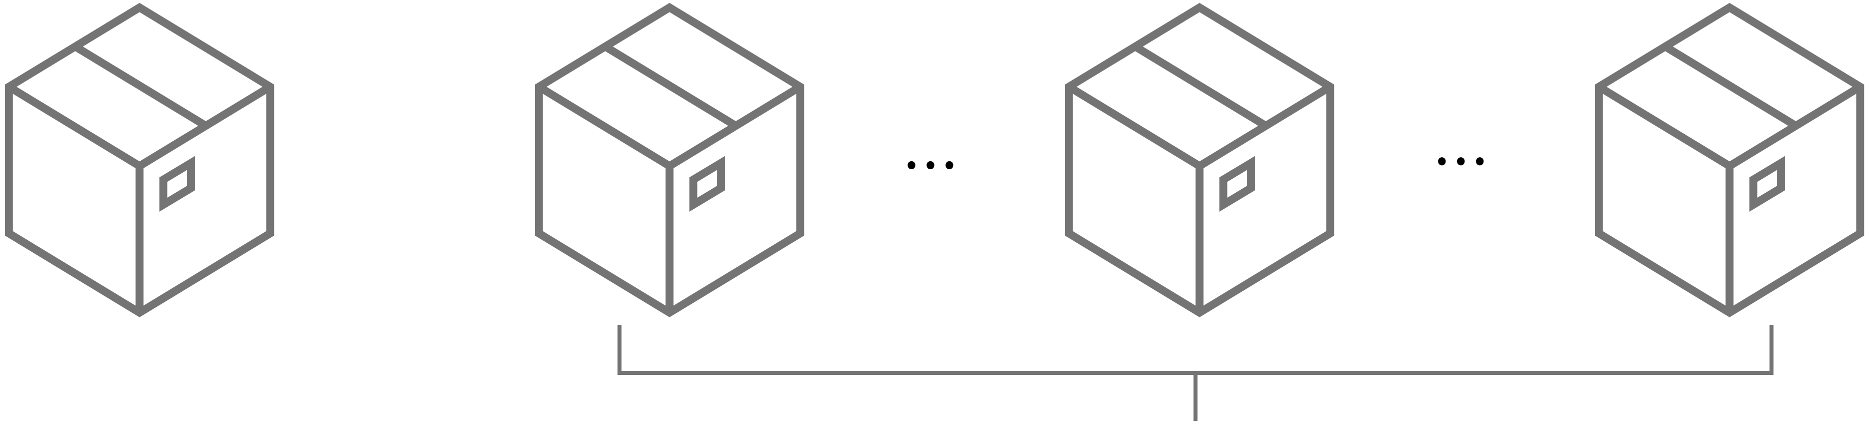

$f(1) = 200 \text{ w.p. } 1$

$c(1) = 198$

$\text{EI}_f(1; 0) - c(1)$
$= 200 - 198 = 2$

$$f(x) = \begin{cases} 200 & \text{w.p. } 0.01 \\ 0 & \text{otherwise} \end{cases}, c(x) = 1, x \in \{2, 3, \dots, 1000\}$$

$\text{EI}_f(x; 0) - c(x)$
$= 2 - 1 = 1$

Inspection rule: $\text{argmax}_x \left(\text{EI}_f(x; y_{\text{best}}) - c(x)\right)$ **Stopping rule:** $\text{EI}_f(x; y_{\text{best}}) \leq c(x), \forall x \in \mathcal{X}$

$$\text{EI}_f(x; y) := \mathbb{E}[(f(x) - y)^+]$$

Naïve Greedy policy can fail [Singla'18]

$t = 1$

$y_{\text{best}} = 200$

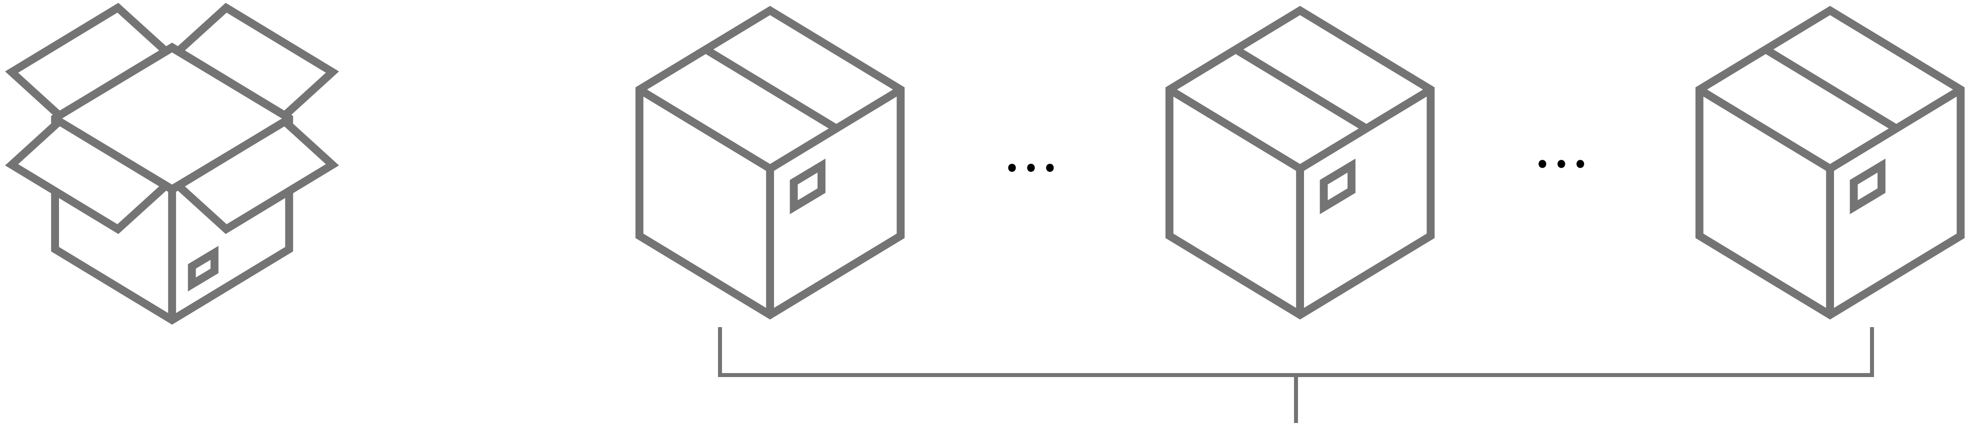

$f(1) = 200$ w.p. 1

$c(1) = 198$

$$f(x) = \begin{cases} 200 & \text{w.p. } 0.01 \\ 0 & \text{otherwise} \end{cases}, c(x) = 1, x \in \{2, 3, \dots, 1000\}$$

$\mathrm{EI}_f(x; 200) - c(x)$

$= 0 - 1 = -1 < 0$

Inspection rule: $\operatorname{argmax}_x \left(\mathrm{EI}_f(x; y_{\text{best}}) - c(x)\right)$ **Stopping rule:** $\mathrm{EI}_f(x; y_{\text{best}}) \leq c(x), \forall x \in \mathcal{X}$

$\mathrm{EI}_f(x; y) := \mathbb{E}[(f(x) - y)^+]$

Naïve Greedy policy can fail [Singla'18]

$t = 1$

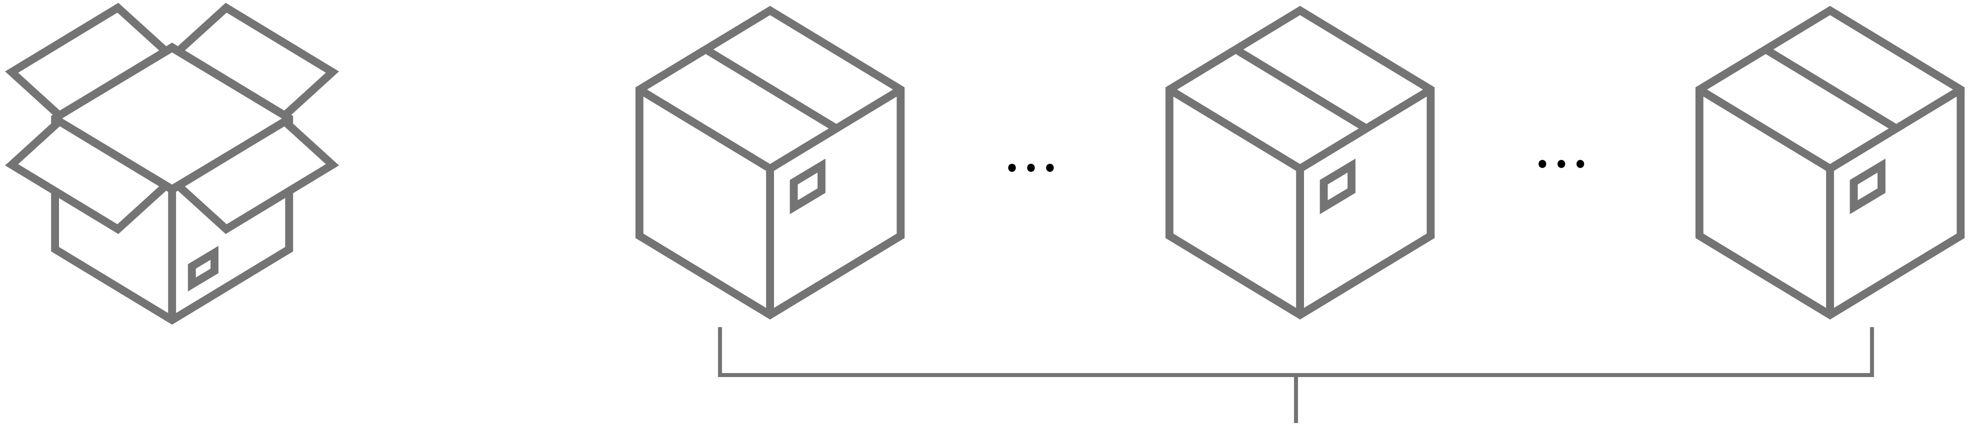

$f(1) = 200 \text{ w.p. } 1$

$c(1) = 198$

$$f(x) = \begin{cases} 200 & \text{w.p. } 0.01 \\ 0 & \text{otherwise} \end{cases}, c(x) = 1, x \in \{2, 3, \ldots, 1000\}$$

Inspection rule: $\text{argmax}_x \left(\text{EI}_f(x; y_{\text{best}}) - c(x)\right)$ **Stopping rule:** $\text{EI}_f(x; y_{\text{best}}) \leq c(x), \forall x \in \mathcal{X}$

Expected utility: $\mathbb{E}[\text{Greedy}] = 200 - 198 = 2$

Naïve Greedy policy can fail [Singla'18]

$t \approx 100$

$y_{\text{best}} = 200$

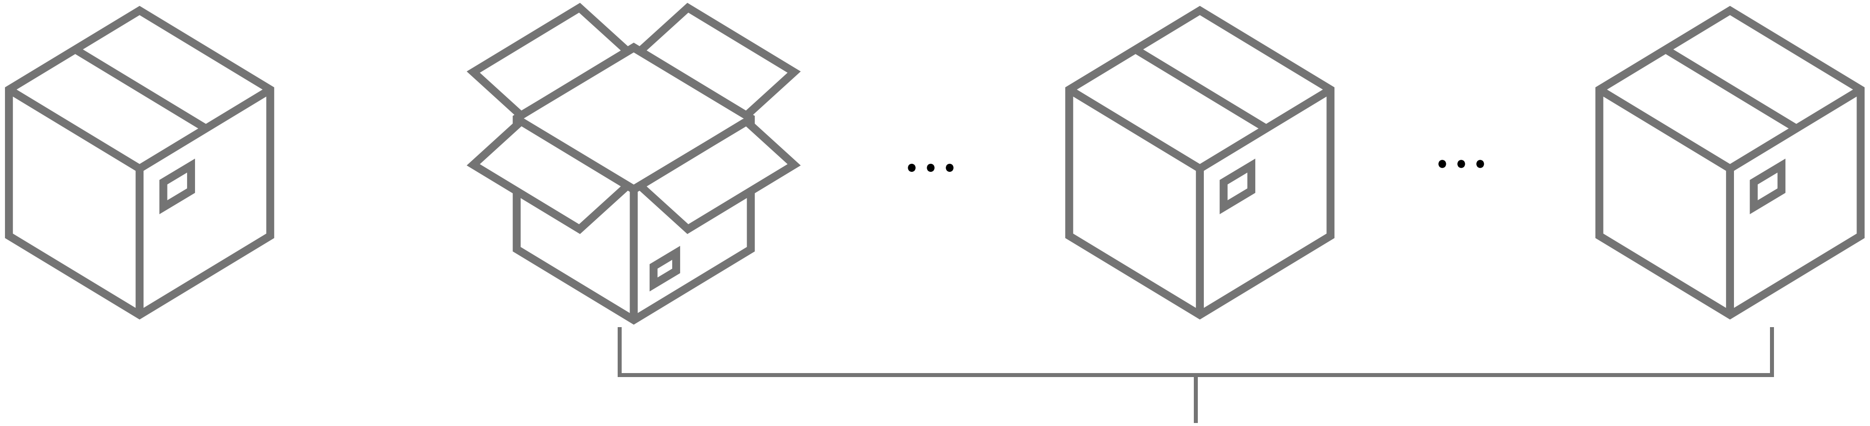

$f(1) = 200$ w.p. 1

$c(1) = 198$

$$f(x) = \begin{cases} 200 & \text{w.p. } 0.01 \\ 0 & \text{otherwise} \end{cases}, c(x) = 1, x \in \{2, 3, \dots, 1000\}$$

Inspection rule: $x \in \{2, 3, \dots, 1000\}$

Stopping rule: $y_{\text{best}} = 200$

Expected utility: $\mathbb{E}[\text{Optimal}] = 200 - 100 * 1 = 100$

1.5-Box Problem

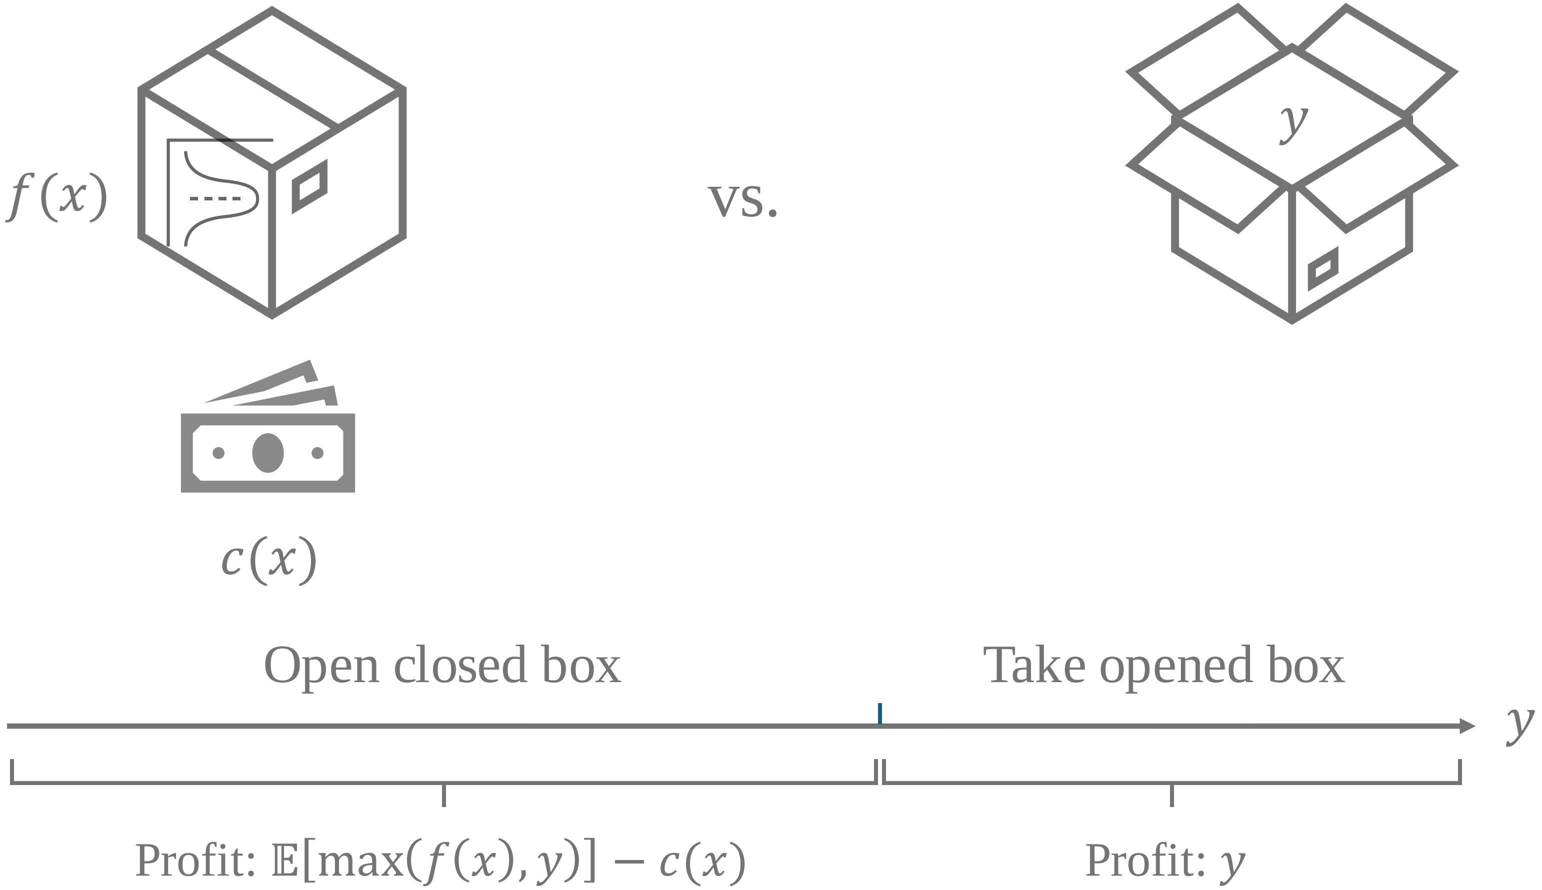

1.5-Box Problem

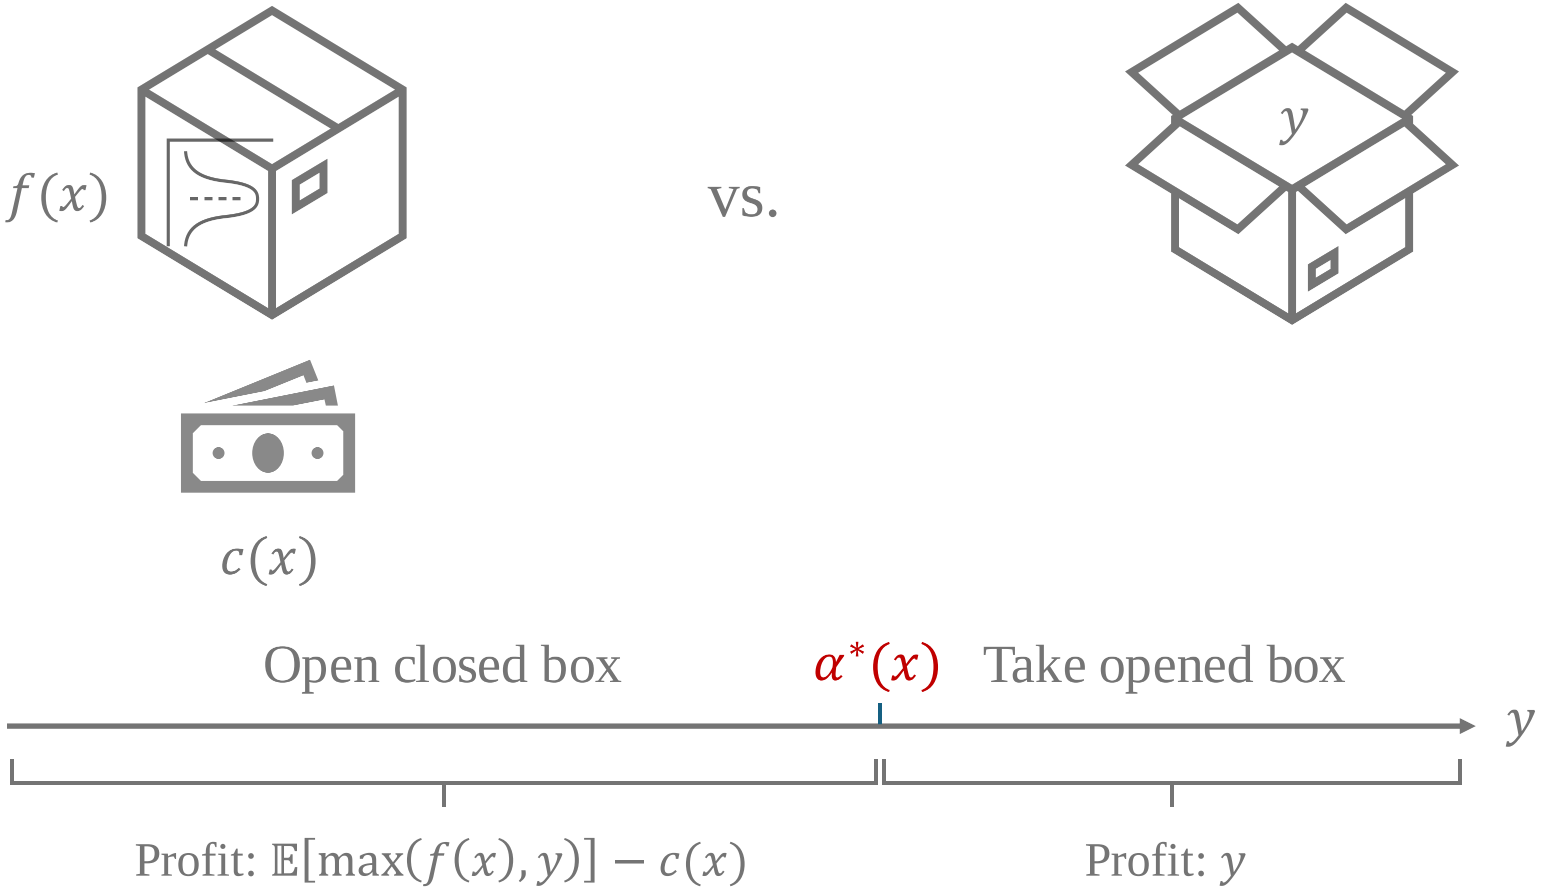

$\alpha^*(x)$: solution y to $\mathbb{E}[\max(f(x), y)] - c(x) = y$

1.5-Box Problem

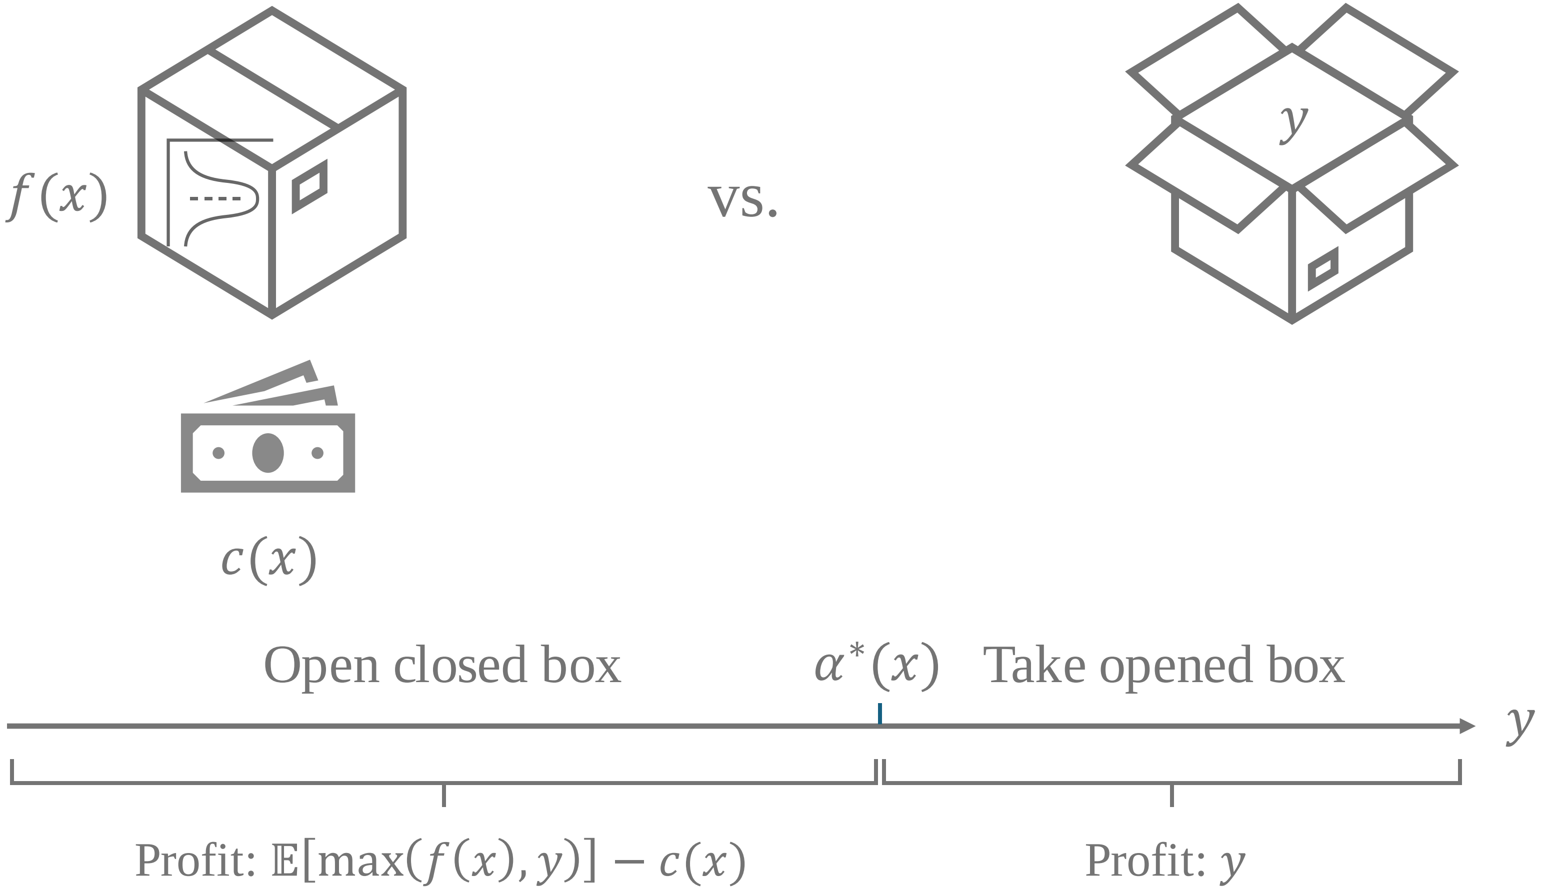

$\alpha^*(x)$: solution y to $\mathbb{E}[(f(x) - y)^+] = c(x)$

1.5-Box Problem

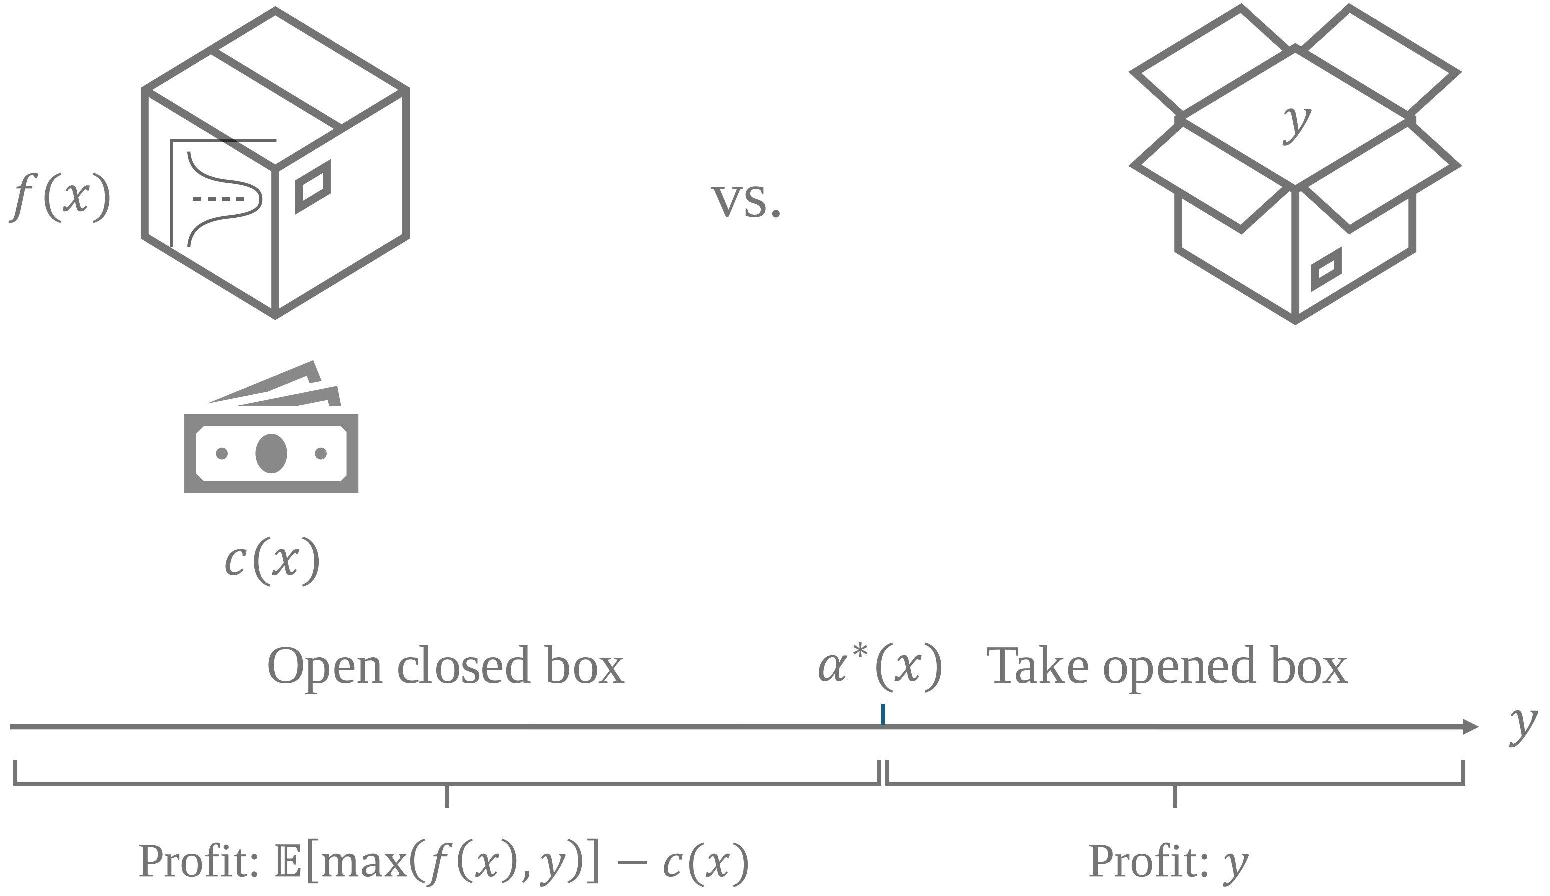

$\alpha^*(x)$: solution y to $\mathbb{E}[(f(x) - y)^+] = c(x)$ $\alpha^*(x)$: Gittins index!

Optimal policy: Gittins policy

Gittins policy

$f(x)$

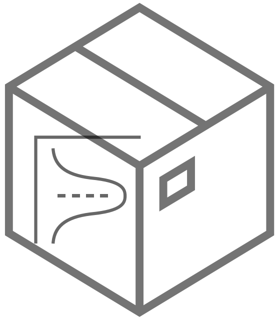

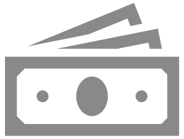

$c(x)$

vs.

Inspection rule: $\text{argmax}_x \, \alpha^*(x)$ s.t. $\text{EI}_f(x; \alpha^*(x)) = c(x)$

solution to expected improvement = cost

Stopping rule: $\alpha^*(x) \leq y_{\text{best}}, \forall x \in \mathcal{X}$

Gittins index ≤ current best

y_{best}: current best observed value

$\text{EI}_f(x; y) := \mathbb{E}[(f(x) - y)^+]$: expected improvement of $f(x)$ over y

Optimal policy: Gittins policy

$t = 0$

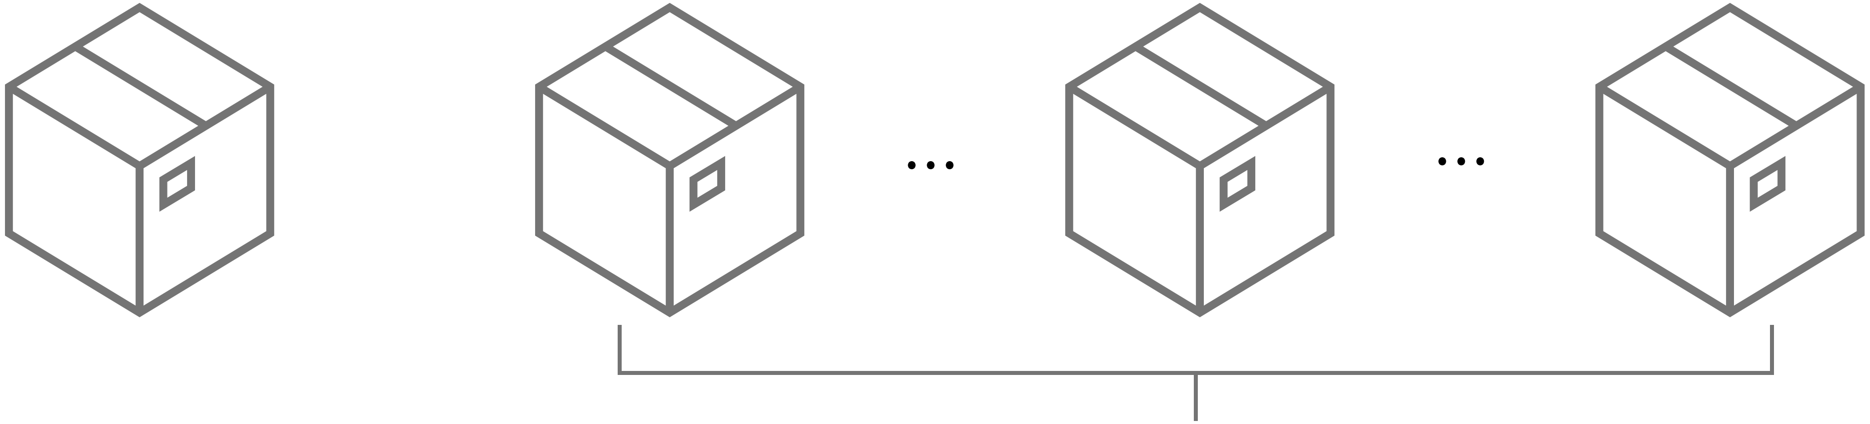

$f(1) = 200 \text{ w.p. } 1$

$c(1) = 198$

$200 - ? = 198$

$f(x) = \begin{cases} 200 & \text{w.p. } 0.01 \\ 0 & \text{otherwise} \end{cases}, c(x) = 1, x \in \{2, 3, \dots, 1000\}$

$(200 - ?) * 0.01 = 1$

Inspection rule: $\text{argmax}_x \; \alpha^*(x)$ s.t. $\text{EI}_f(x; \alpha^*(x)) = c(x)$ **Stopping rule:** $\alpha^*(x) \leq y_{\text{best}}, \forall x \in \mathcal{X}$

$\text{EI}_f(x; y) := \mathbb{E}[(f(x) - y)^+]$

Optimal policy: Gittins policy

$t = 0$

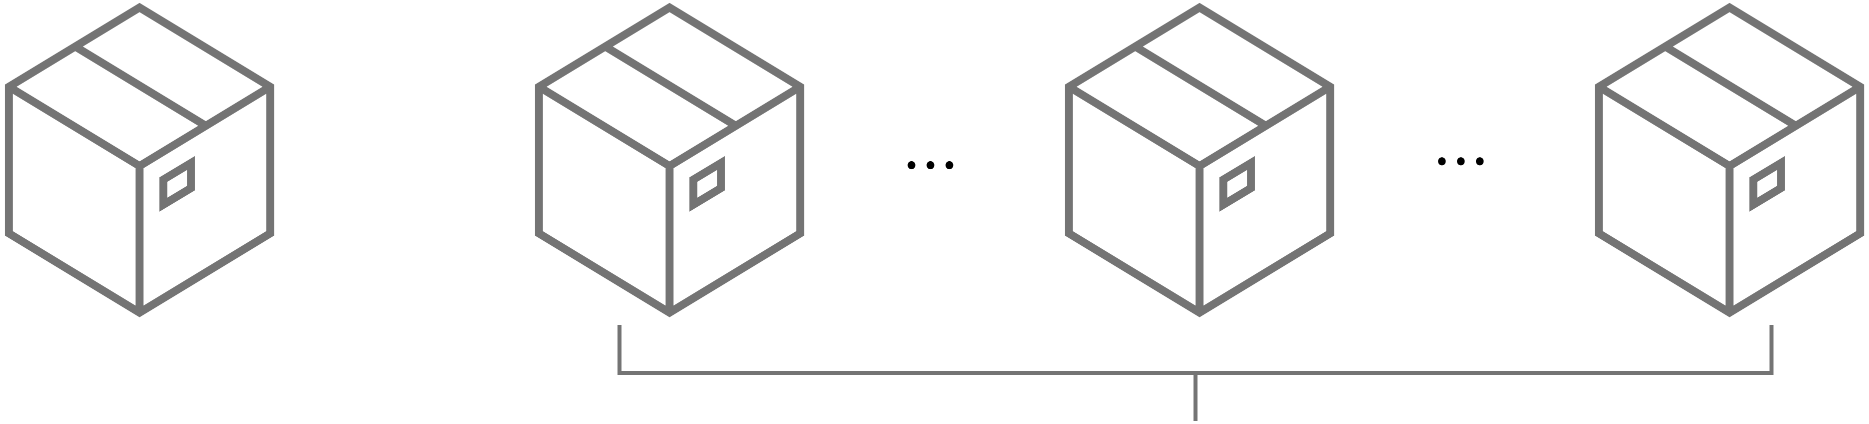

$f(1) = 200 \text{ w.p. } 1$

$c(1) = 198$

$\alpha^*(1) = 2$

$f(x) = \begin{cases} 200 & \text{w.p. } 0.01 \\ 0 & \text{otherwise} \end{cases}, c(x) = 1, x \in \{2, 3, \dots, 1000\}$

$\alpha^*(x) = 100$

Inspection rule: $\text{argmax}_x\, \alpha^*(x)$ s.t. $\text{EI}_f(x; \alpha^*(x)) = c(x)$ **Stopping rule:** $\alpha^*(x) \le y_{\text{best}}, \forall x \in \mathcal{X}$

$\text{EI}_f(x; y) := \mathbb{E}[(f(x) - y)^+]$

Optimal policy: Gittins policy

$t = 1$

$y_{\text{best}} = 200 \text{ or } 0$

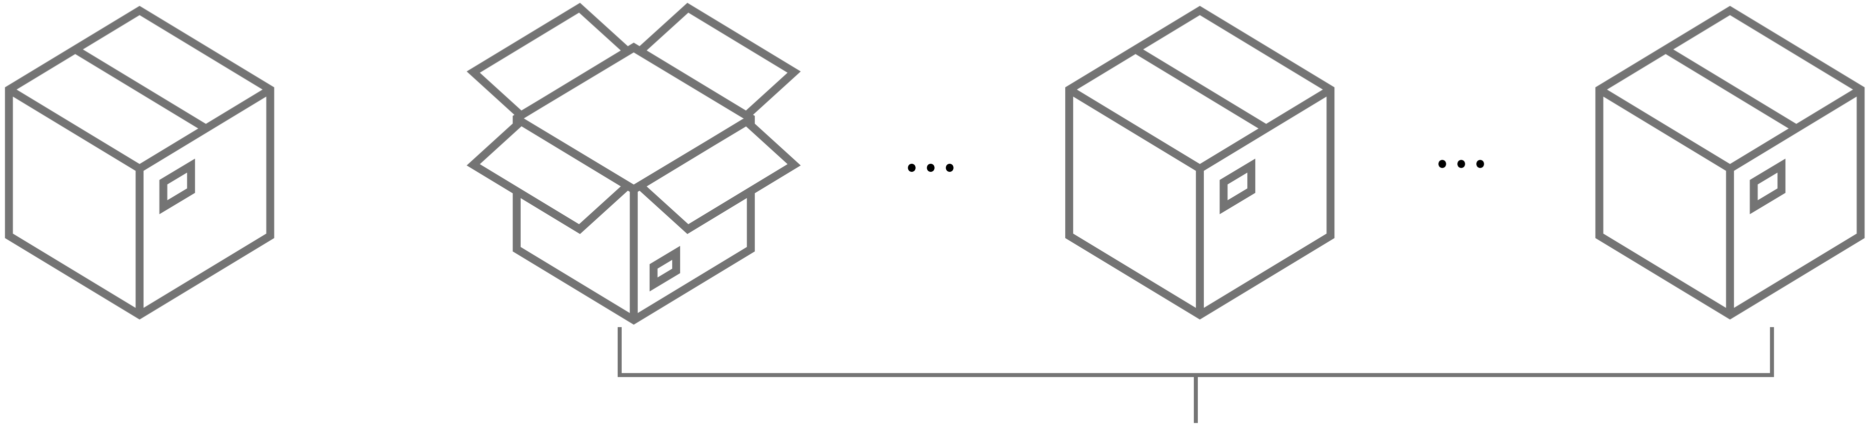

$f(1) = 200 \text{ w.p. } 1$

$c(1) = 198$

$\alpha^*(1) = 2$

$f(x) = \begin{cases} 200 & \text{w.p. } 0.01 \\ 0 & \text{otherwise} \end{cases}, c(x) = 1, x \in \{2, 3, \dots, 1000\}$

$\alpha^*(x) = 100$

Inspection rule: $\text{argmax}_x \, \alpha^*(x)$ s.t. $\text{EI}_f(x; \alpha^*(x)) = c(x)$ **Stopping rule:** $\alpha^*(x) \le y_{\text{best}}, \forall x \in \mathcal{X}$

$\text{EI}_f(x; y) := \mathbb{E}[(f(x) - y)^+]$

Optimal policy: Gittins policy

$t \approx 100$

$y_{\text{best}} = 200$

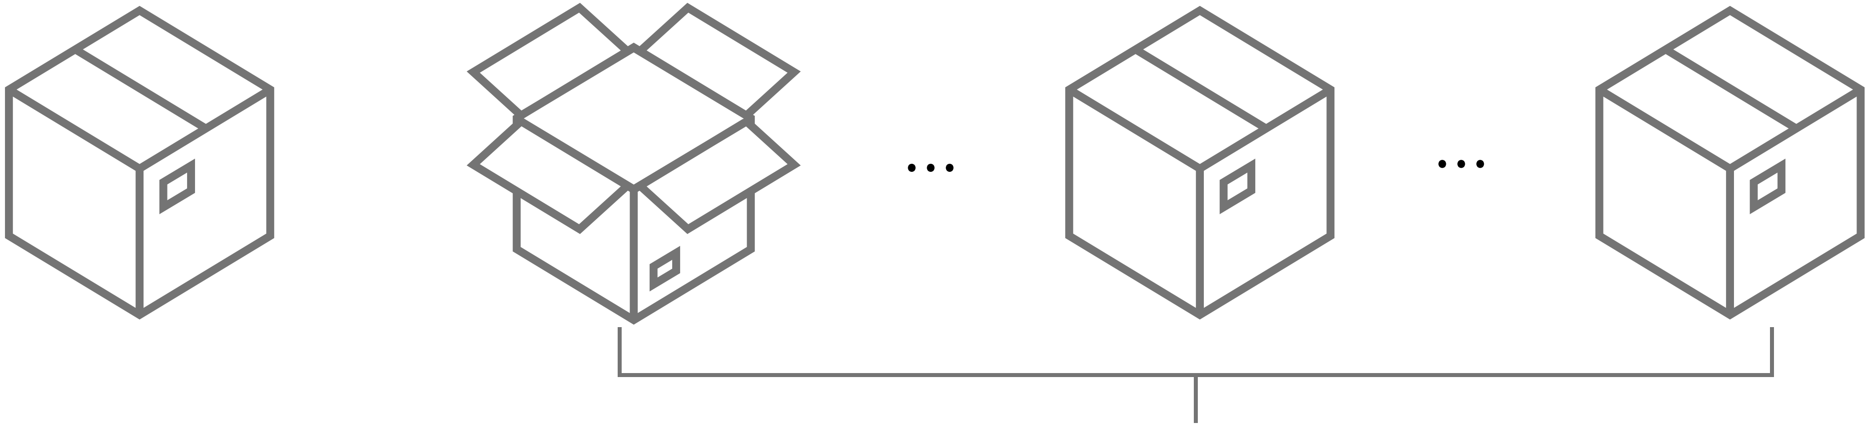

$f(1) = 200 \text{ w.p. } 1$

$c(1) = 198$

$\alpha^*(1) = 2$

$f(x) = \begin{cases} 200 & \text{w.p. } 0.01 \\ 0 & \text{otherwise} \end{cases}, c(x) = 1, x \in \{2, 3, \dots, 1000\}$

$\alpha^*(x) = 100$

Inspection rule: $\text{argmax}_x\ \alpha^*(x)$ s.t. $\text{EI}_f(x; \alpha^*(x)) = c(x)$ **Stopping rule:** $\alpha^*(x) \leq y_{\text{best}}, \forall x \in \mathcal{X}$

Expected utility: $\mathbb{E}[\text{Gittins}] = 200 - 100 * 1 = 100$

Bayesian Optimization ⟹ Pandora’s Box

Special case of Markovian/
Bayesian multi-armed bandits

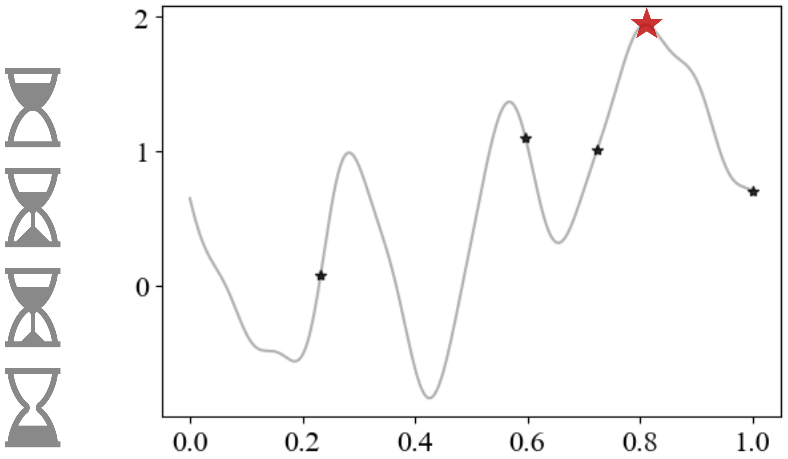

⟹

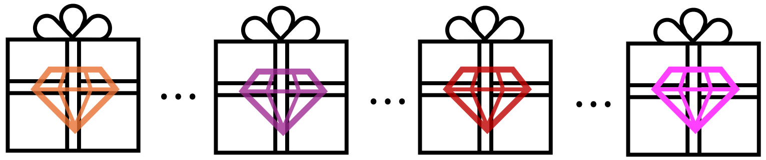

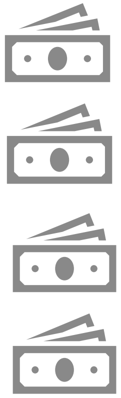

Continuous	⟹	Discrete
Correlated	⟹	Independent
Hard budget constraint	⟹	Cost per sample

Optimal policy: Gittins index [Weitzman’79]

Bayesian Optimization ⟹ Pandora's Box

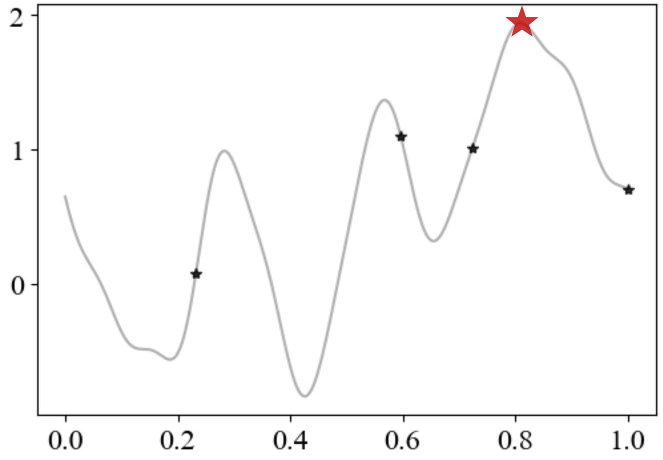

⟹

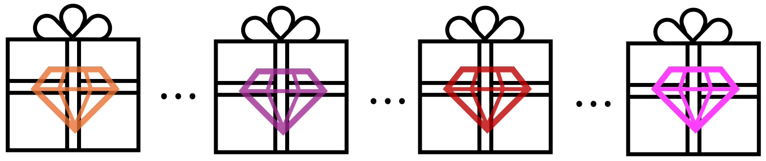

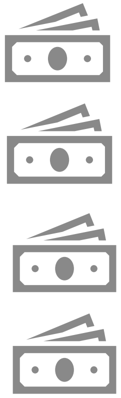

Continuous	⟹	Discrete
Correlated	⟹	Independent
Hard budget constraint	⟹	Cost per sample
Is Gittins index good?	⟸	Optimal policy: Gittins index

Bayesian Optimization $\Longrightarrow$ Pandora’s Box

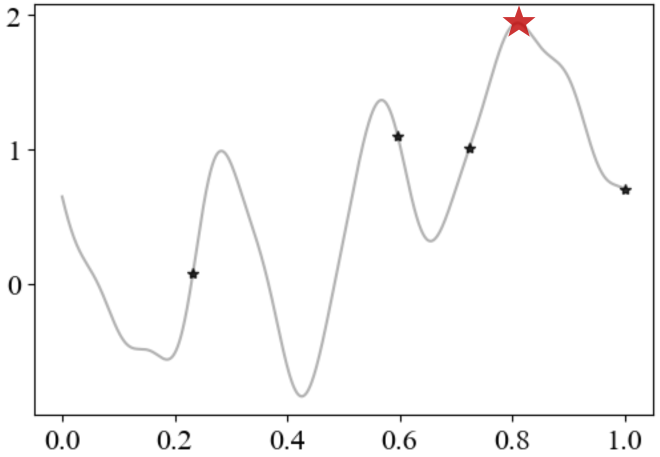

$\Longrightarrow$

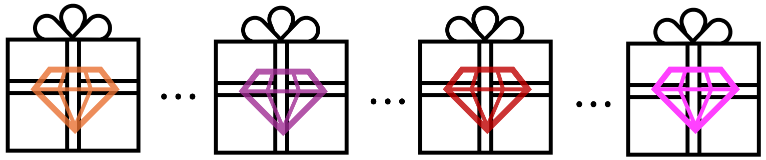

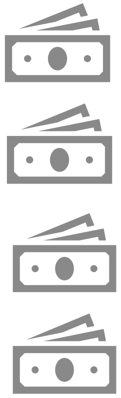

Continuous $\Longrightarrow$ Discrete

Correlated $\Longrightarrow$ Independent

Hard budget constraint $\Longrightarrow$ Cost per sample

How to translate?

Is Gittins index good? $\Longleftarrow$ Optimal policy: Gittins index

Bayesian Optimization ⟹ Pandora's Box

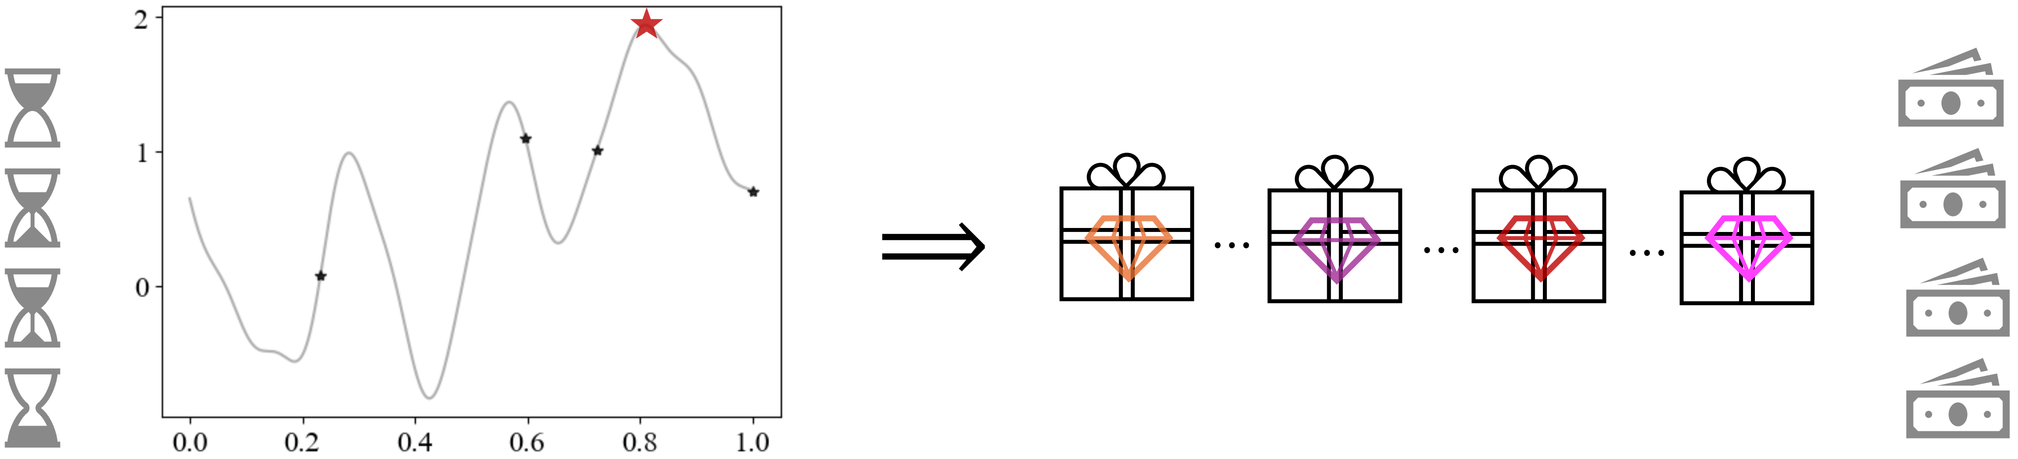

Continuous ⟹ Discrete

Correlated ⟹ Independent

Hard budget constraint ⟹ Cost per sample

How to translate?

Is Gittins index good? ⟸ Optimal policy: Gittins index

Our contributions!

How to translate?

budget

max best observed
s.t. budget constraint

Pandora's box

costs

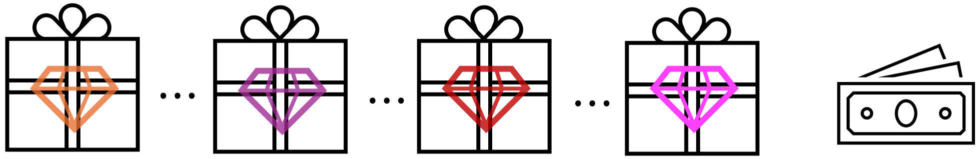

max (best observed – costs)

Expected budget constraint $\Leftrightarrow$ Cost per sample

Optimal policy? $\Leftarrow$ Optimal policy: Gittins index

Expected budget constraint ⟺ Cost per sample

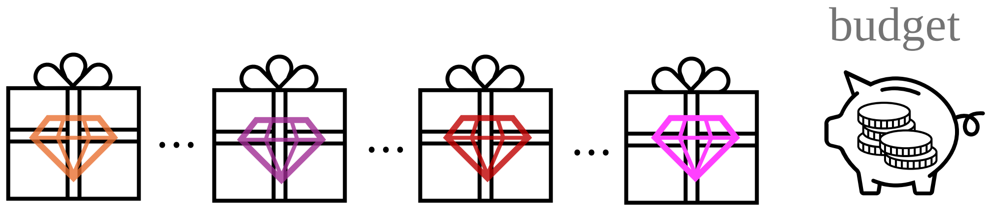

max best observed
s.t. budget constraint

Pandora's box

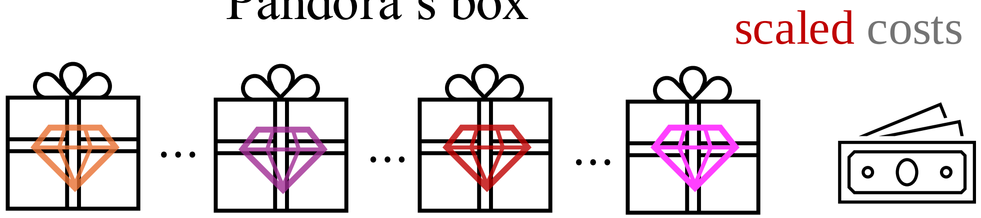

max (best observed – λ_B^* costs))

λ_B^*: budget-dependent scaling factor

Optimal policy: Gittins solution to Pandora's box with scaled costs $\Longleftarrow$ Optimal policy: Gittins index

Expected budget constraint ⟺ Cost per sample

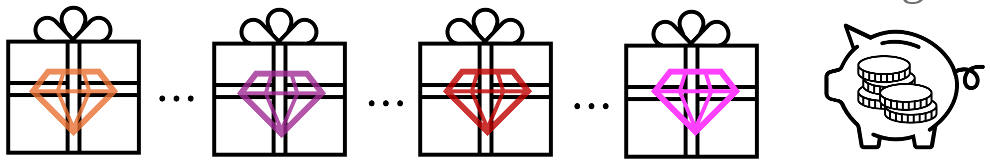

max best observed
s.t. budget constraint

Pandora's box

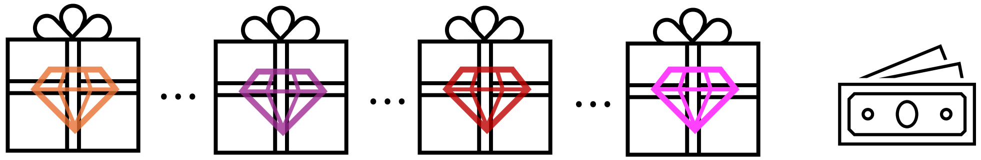

max (best observed – λ_B^* costs))

λ_B^*: budget-dependent scaling factor

Reward distribution	Reference
finite support	[Aminian, Manshadi, Niazadeh'24]
general support	our work

Optimal policy: Gittins solution to Pandora's box with scaled costs ⟸ Optimal policy: Gittins index

Expected budget constraint ⟺ Cost per sample

Constrained

budget

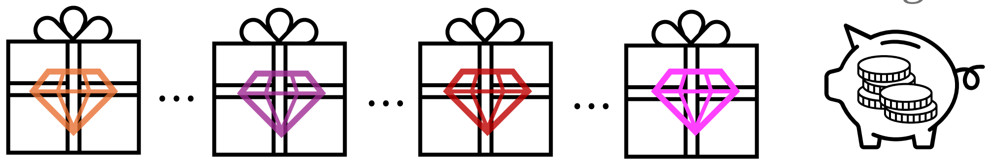

$$\sup_{\text{policy}} \mathbb{E} \max_{t=1,2,\dots,T} f(x_t)$$

$$\text{s.t. } \mathbb{E} \sum_{t=1}^{T} c(x_t) \le B$$

Lagrangian relaxation

⟺

Unconstrained

costs

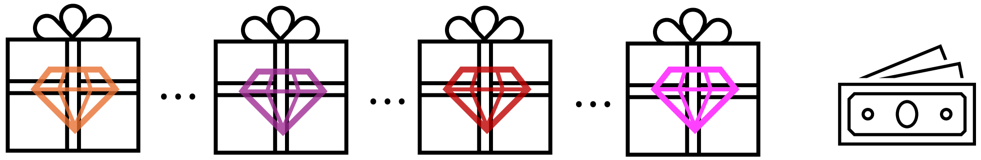

$$\sup_{\text{policy}} \inf_{\lambda \ge 0} \mathbb{E}\left(\max_{t=1,2,\dots,T} f(x_t) + \lambda \left(B - \mathbb{E} \sum_{t=1}^{T} c(x_t) \right) \right)$$

λ: Lagrange multiplier

Optimal policy: Gittins solution to Pandora's box with scaled costs ⟸ Optimal policy: Gittins index

Extension to [Aminian, Manshadi, Niazadeh'24]

Expected budget constraint ⟺ Cost per sample

Constrained

budget

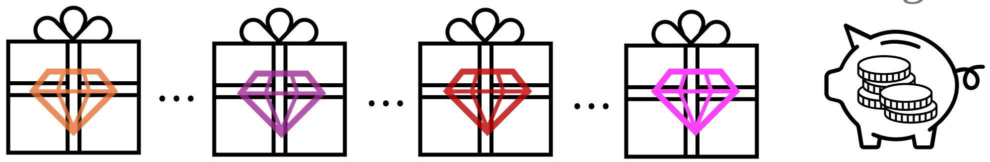

$$\sup_{\text{policy}} \mathbb{E} \max_{t=1,2,\dots,T} f(x_t)$$

$$\text{s.t. } \mathbb{E} \sum_{t=1}^{T} c(x_t) \le B$$

Lagrangian relaxation

⟺

Lagrange multiplier theorem

⟺

Unconstrained

costs

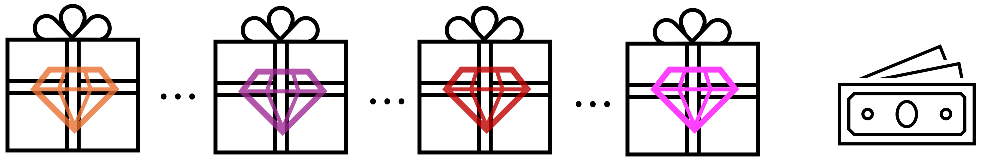

$$\sup_{\text{policy}} \inf_{\lambda\ge 0} \mathbb{E}\left(\max_{t=1,2,\dots,T} f(x_t)\right.$$

$$\inf_{\lambda\ge 0} \sup_{\text{policy}} \mathbb{E}\left(\max_{t=1,2,\dots,T} f(x_t) + \lambda\left(B - \mathbb{E}\sum_{t=1}^{T} c(x_t)\right)\right)$$

Optimal policy: Gittins solution to Pandora's box with scaled costs ⟸ Optimal policy: Gittins index

Extension to [Aminian, Manshadi, Niazadeh'24]

Expected budget constraint $\Longleftrightarrow$ Cost per sample

Constrained

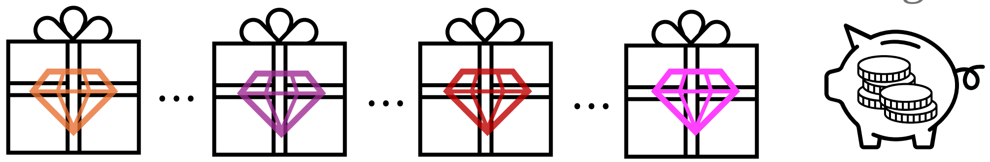

$$\sup_{\text{policy}} \mathbb{E} \max_{t=1,2,\ldots,T} f(x_t)$$

$$\text{s.t. } \mathbb{E} \sum_{t=1}^{T} c(x_t) \le B$$

Unconstrained

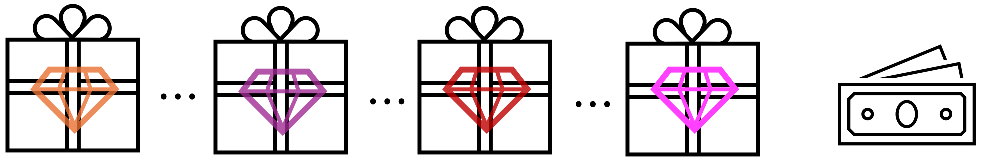

Lagrangian relaxation $\Longleftrightarrow$

$$\sup_{\text{policy}} \inf_{\lambda \ge 0} \mathbb{E} \left(\max_{t=1,2,\ldots,T} f(x_t) + \lambda \left(B - \mathbb{E} \sum_{t=1}^{T} c(x_t) \right) \right)$$

Lagrange multiplier theorem $\Longleftrightarrow$

$$\inf_{\lambda \ge 0} \underbrace{\sup_{\text{policy}} \mathbb{E} \left(\max_{t=1,2,\ldots,T} f(x_t) + \lambda \left(B - \mathbb{E} \sum_{t=1}^{T} c(x_t) \right) \right)}_{\mathcal{A}(\lambda)}$$

λ_B^*: nonzero infimum of $\mathcal{A}(\lambda)$ exists $\overset{?}{\Longleftarrow}$

Optimal policy: Gittins solution to Pandora's box with scaled costs $\Longleftarrow$ Optimal policy: Gittins index

Extension to [Aminian, Manshadi, Niazadeh'24]

Expected budget constraint $\Longleftrightarrow$ Cost per sample

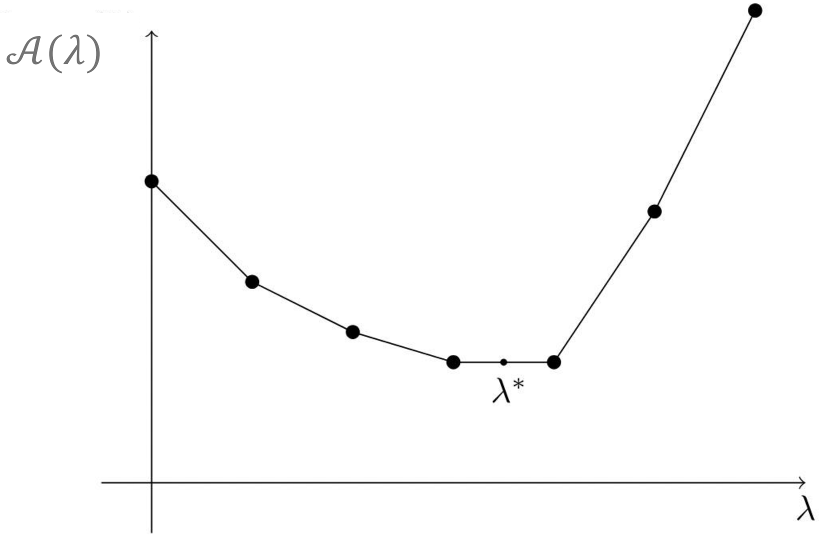

(a) Degenerate case, differentiable at λ^*.

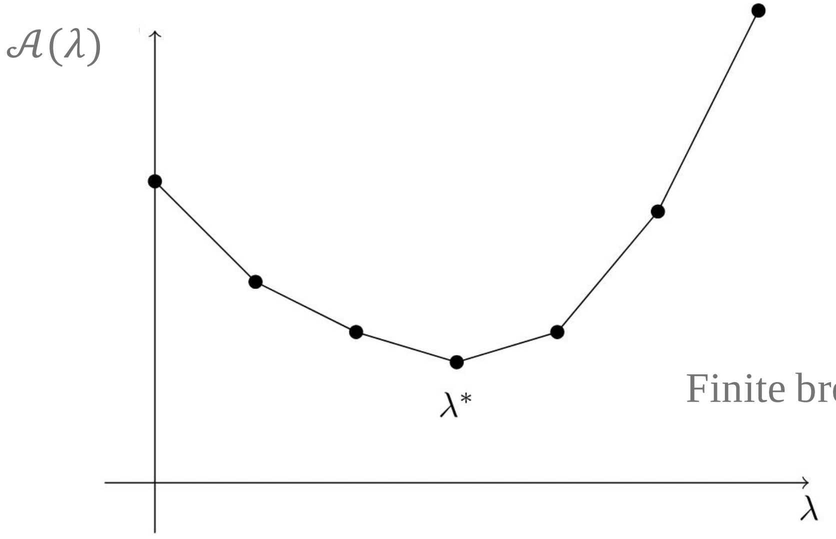

(b) Non-degenerate case, breakpoint at λ^*.

Finite breakpoints

$$\inf_{\lambda \geq 0} \sup_{\text{policy}} \mathbb{E}\left(\max_{t=1,2,\ldots,T} f(x_t) + \lambda \left(B - \mathbb{E} \sum_{t=1}^{T} c(x_t) \right) \right)$$

Envelope Theorem

λ_B^*: nonzero infimum of $\mathcal{A}(\lambda)$ exists $\Longleftarrow$ $\mathcal{A}(\lambda)$: convex (possibly non-differentiable) in λ

Optimal policy: Gittins solution to Pandora's box with scaled costs $\Longleftarrow$ Optimal policy: Gittins index

Figure from [Aminian, Manshadi, Niazadeh'24]

Expected budget constraint $\Leftrightarrow$ Cost per sample

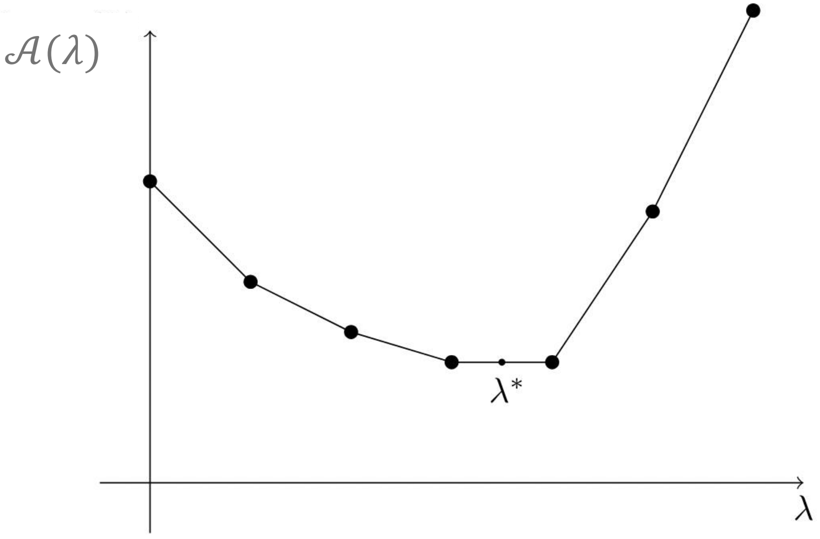

(a) Degenerate case, differentiable at λ^*.

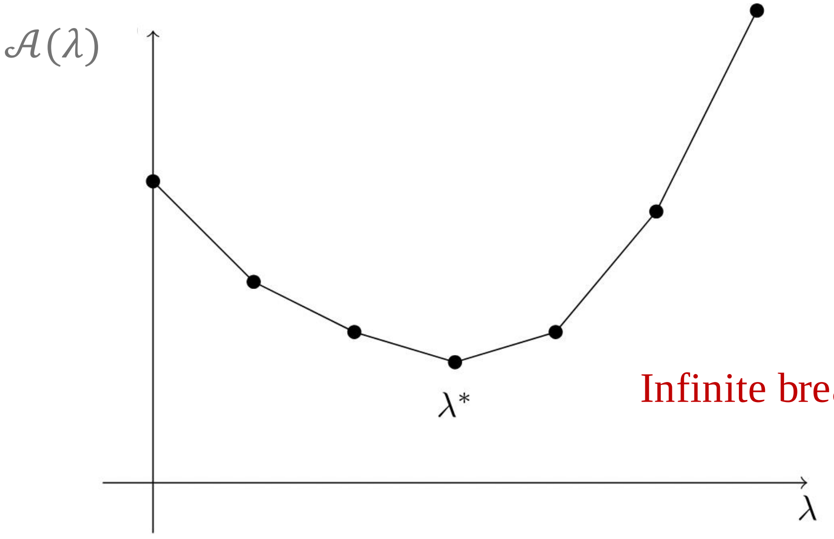

(b) Non-degenerate case, breakpoint at λ^*.

Infinite breakpoints?

$$\inf_{\lambda \geq 0} \sup_{\text{policy}} \mathbb{E}\left(\max_{t=1,2,\ldots,T} f(x_t) + \lambda \left(B - \mathbb{E} \sum_{t=1}^{T} c(x_t) \right) \right)$$

λ_B^*: nonzero infimum of $\mathcal{A}(\lambda)$ exists $\Leftarrow$ $\mathcal{A}(\lambda)$: convex (possibly non-differentiable) in λ

Optimal policy: Gittins solution to Pandora's box with scaled costs $\Leftarrow$ Optimal policy: Gittins index

Figure from [Aminian, Manshadi, Niazadeh'24]

Expected budget constraint $\Longleftrightarrow$ Cost per sample

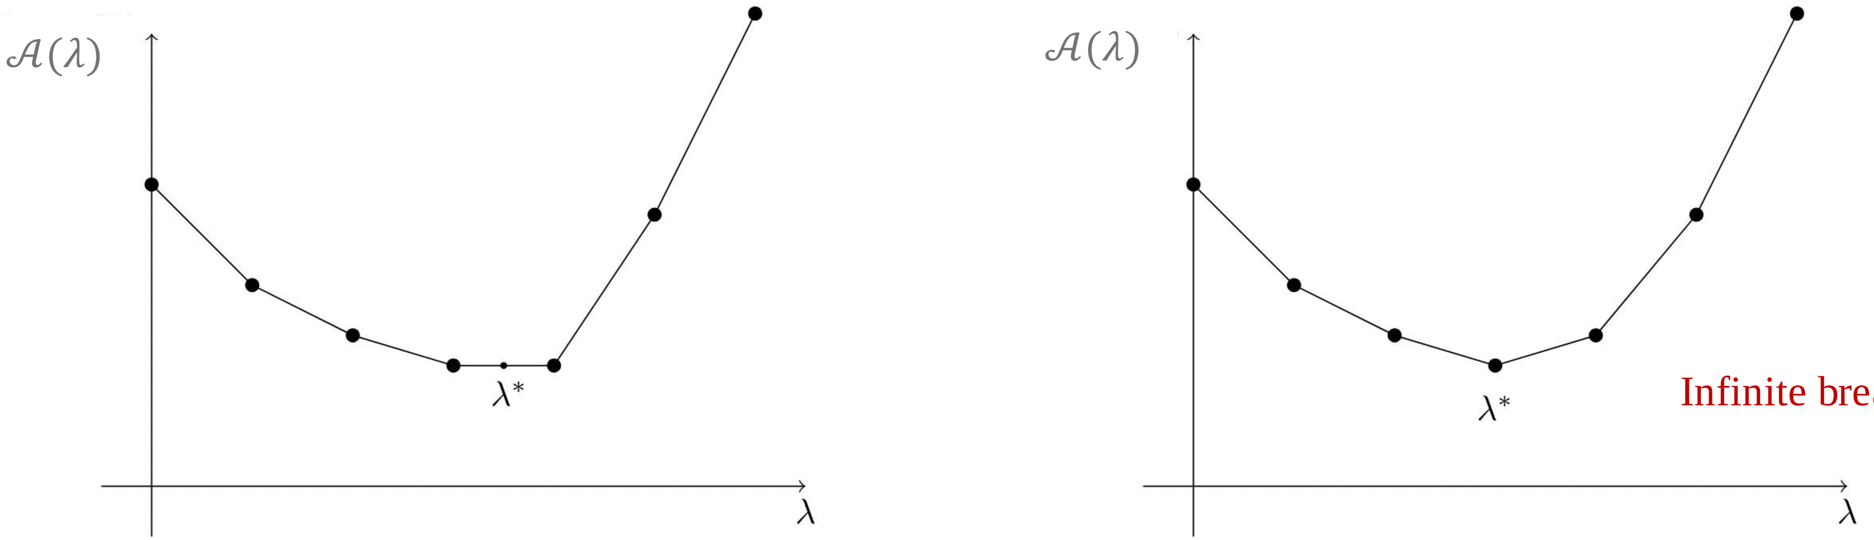

(a) Degenerate case, differentiable at λ^*.

(b) Non-degenerate case, breakpoint at λ^*.

Infinite breakpoints?

$$\inf_{\lambda \geq 0} \sup_{\text{policy}} \mathbb{E}\left(\max_{t=1,2,\ldots,T} f(x_t) + \lambda \left(B - \mathbb{E} \sum_{t=1}^{T} c(x_t) \right) \right)$$

Our work: sharp Envelope Theorem

λ_B^*: nonzero infimum of $\mathcal{A}(\lambda)$ exists $\Longleftarrow$ $\mathcal{A}(\lambda)$: convex (possibly non-differentiable) in λ

Optimal policy: Gittins solution to Pandora's box with scaled costs $\Longleftarrow$ Optimal policy: Gittins index

Figure from [Aminian, Manshadi, Niazadeh'24]

How to translate?

budget

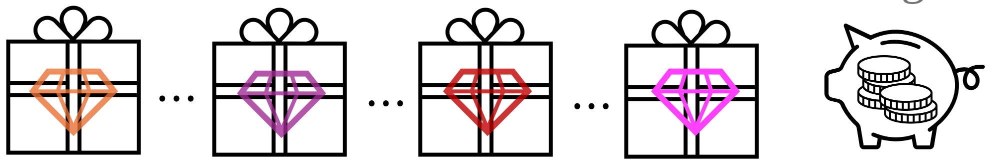

max best observed
s.t. budget constraint

Pandora's box

costs

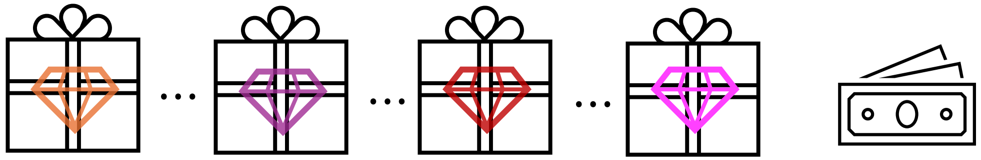

max (best observed – costs)

Hard budget constraint $\Leftarrow$ Cost per sample

$$\text{argmax}_x\ \alpha^*(x) \text{ s.t. } \text{EI}_f(x; \alpha^*(x)) = \lambda_B^* c(x) \Leftarrow \text{argmax}_x\ \alpha^*(x) \text{ s.t. } \text{EI}_f(x; \alpha^*(x)) = c(x)$$

How to translate?

Bayesian optimization

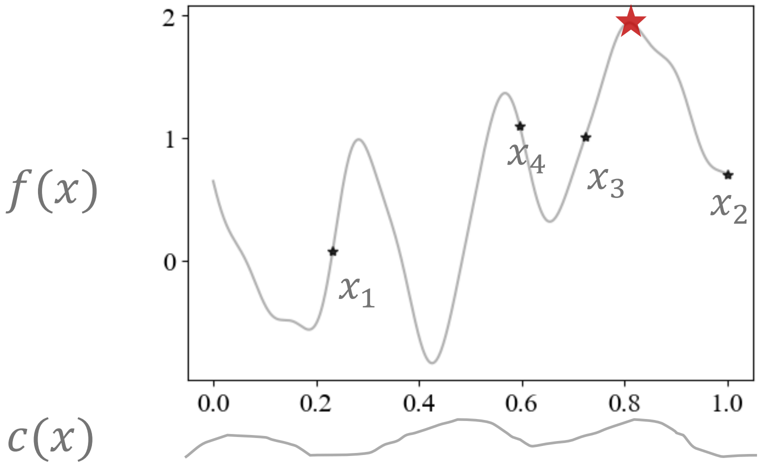

Budget-constrained
Pandora's box

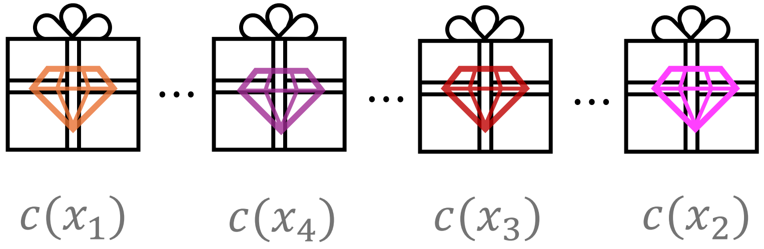

Continuous ⟸ Discrete

Correlated ⟸ Independent

Hard budget constraint ⟸ Cost per sample

How to incorporate Gaussian process? ⟸ Optimal policy: Gittins solution to Pandora's box with scaled costs

How to translate?

Bayesian optimization

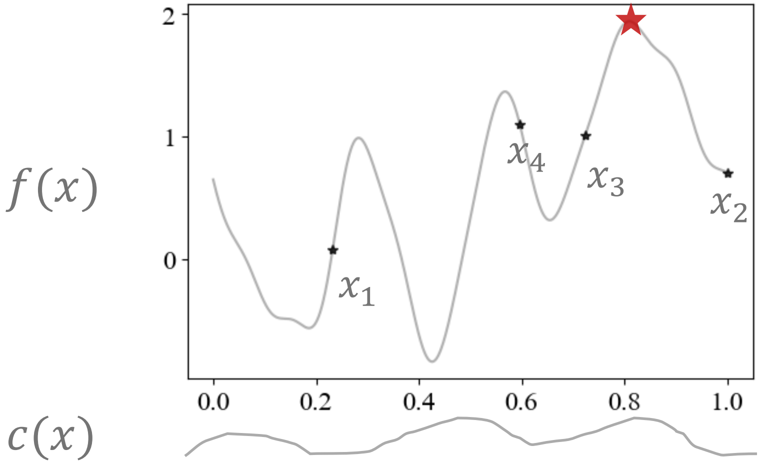

Budget-constrained
Pandora's box

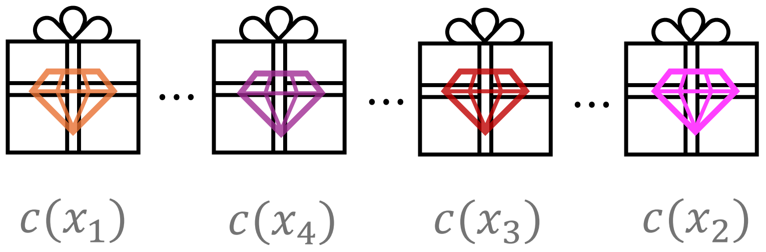

Continuous	$\Longleftarrow$	Discrete
Correlated	$\Longleftarrow$	Independent
Hard budget constraint	$\Longleftarrow$	Cost per sample

$$\text{argmax}_x\ \alpha^*(x) \text{ s.t. } \text{EI}_{f|D}(x; \alpha^*(x)) = \lambda_B^* c(x) \quad \Longleftarrow \quad \text{argmax}_x\ \alpha^*(x) \text{ s.t. } \text{EI}_f(x; \alpha^*(x)) = \lambda_B^* c(x)$$

D: observed data

How to translate?

Bayesian optimization

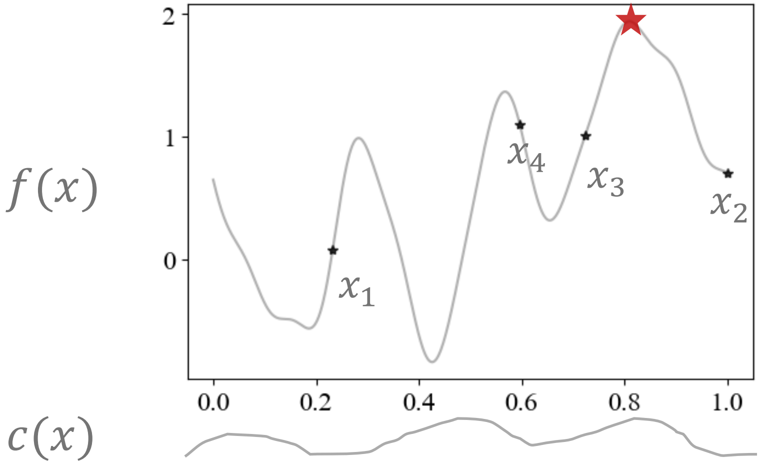

Budget-constrained
Pandora's box

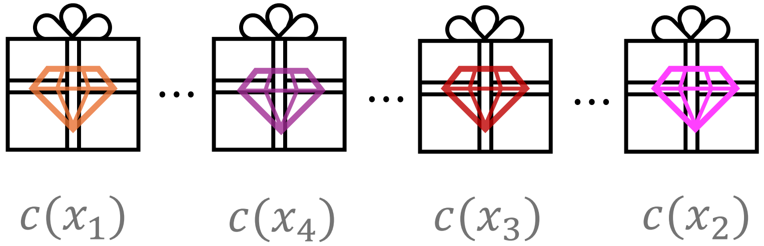

Continuous	$\Longleftarrow$	Discrete
Correlated	$\Longleftarrow$	Independent
Hard budget constraint	$\Longleftarrow$	Cost per sample

$\mathrm{argmax}_x\ \alpha^*(x)$ s.t. $\mathrm{EI}_{f|D}(x;\alpha^*(x)) = \lambda_B^* c(x)$ $\Longleftarrow$ $\mathrm{argmax}_x\ \alpha^*(x)$ s.t. $\mathrm{EI}_f(x;\alpha^*(x)) = \lambda_B^* c(x)$

popular one-step
heuristic: EI policy

How to translate?

Bayesian optimization

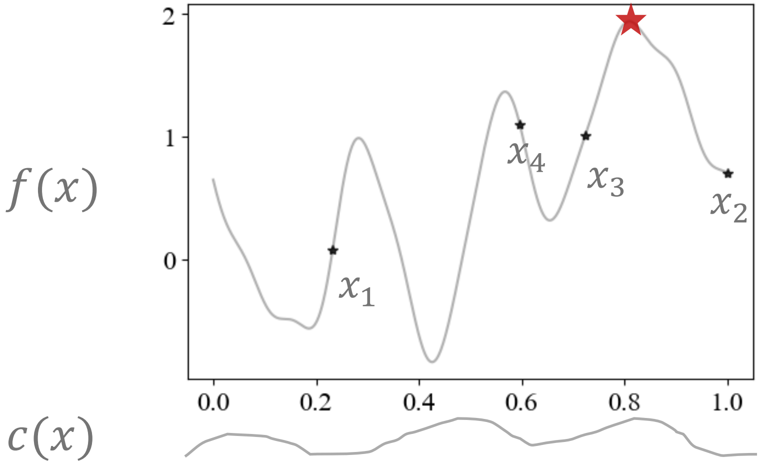

Budget-constrained
Pandora's box

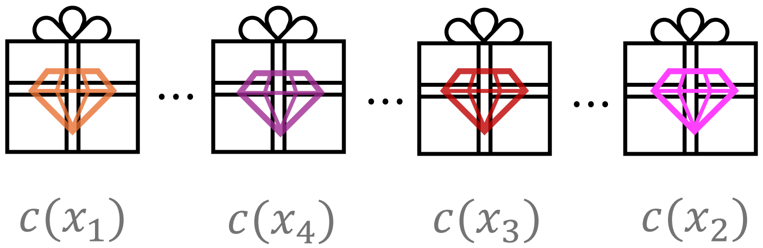

Continuous	$\Leftarrow$	Discrete
Correlated	$\Leftarrow$	Independent
Hard budget constraint	$\Leftarrow$	Cost per sample

$\text{argmax}_x\, \alpha^*(x)$ s.t. $\text{EI}_{f|D}(x; \alpha^*(x)) = \lambda_B^* c(x)$ $\Leftarrow$ $\text{argmax}_x\, \alpha^*(x)$ s.t. $\text{EI}_f(x; \alpha^*(x)) = \lambda_B^* c(x)$

ratio of EI and cost: EIPC policy

How to translate?

Bayesian optimization

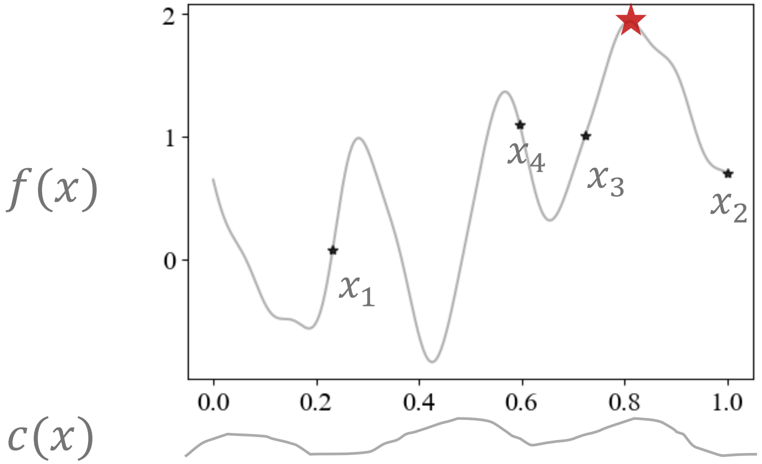

Budget-constrained
Pandora's box

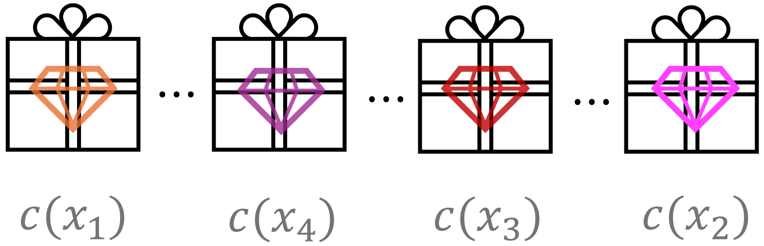

Continuous	$\Leftarrow$	Discrete
Correlated	$\Leftarrow$	Independent
Hard budget constraint	$\Leftarrow$	Cost per sample

$\text{argmax}_x\, \alpha^*(x)$ s.t. $\text{EI}_{f|D}(x; \alpha^*(x)) = \lambda_B^* c(x)$ $\Leftarrow$ $\text{argmax}_x\, \alpha^*(x)$ s.t. $\text{EI}_f(x; \alpha^*(x)) = \lambda_B^* c(x)$

EI and EIPC policy can be arbitrarily worse [Astudillo, Jiang, Balandat, Bakshy, Frazier'21]

Our Contributions

- How to translate?
- Is Pandora's Box Gittins index (PBGI) good?

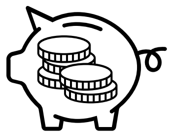

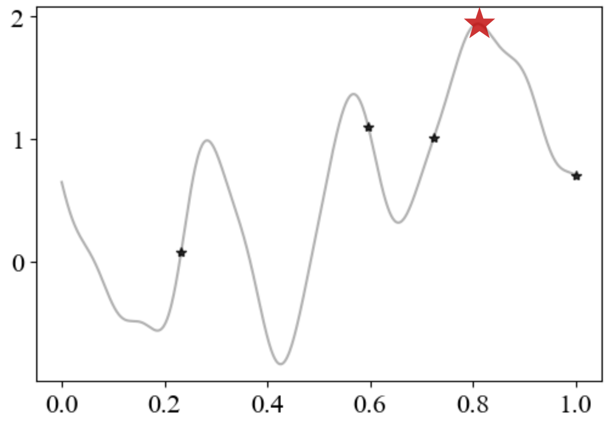

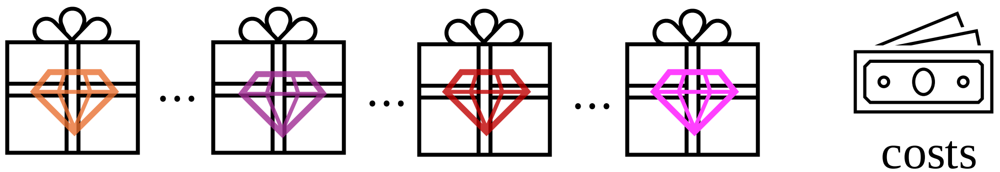

? ⟸ Pandora's Box Gittins index

Our Contributions

- Develop PBGI policy for cost-aware Bayesian optimization
- Is Pandora's Box Gittins index (PBGI) good?

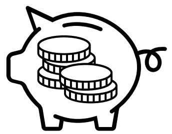

budget

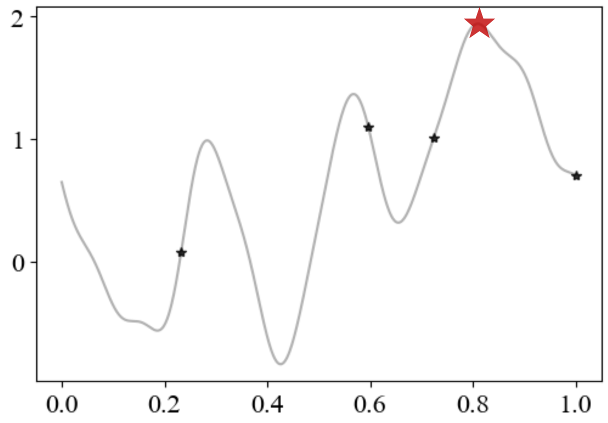

Our work

⟸

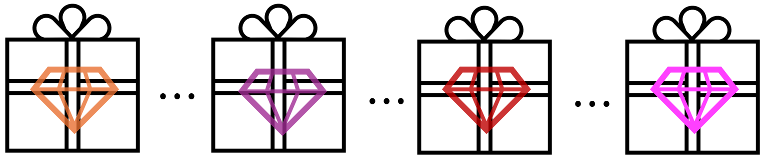

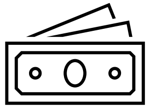

costs

Pandora's Box Gittins index

Our Contributions

- Develop PBGI policy for Bayesian optimization
- Show performance against baselines on synthetic & empirical experiments

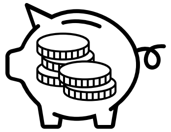

budget

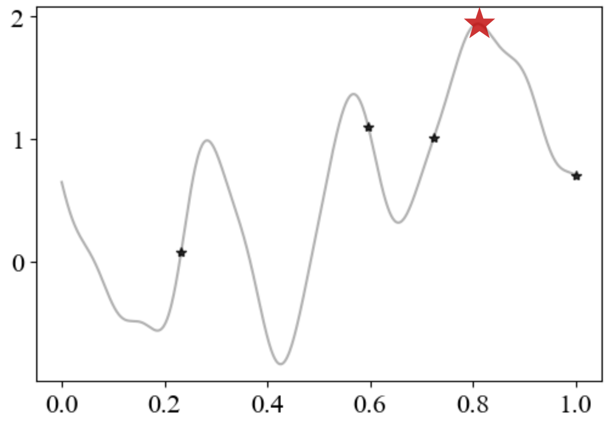

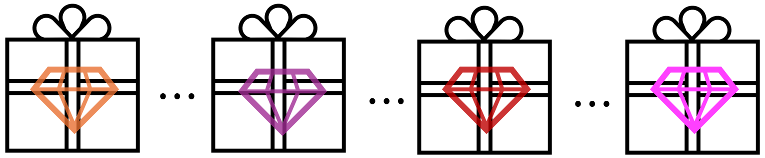

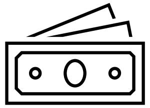

costs

Our work ⟸ Pandora’s Box Gittins index

Experiment Results: PBGI vs Baselines

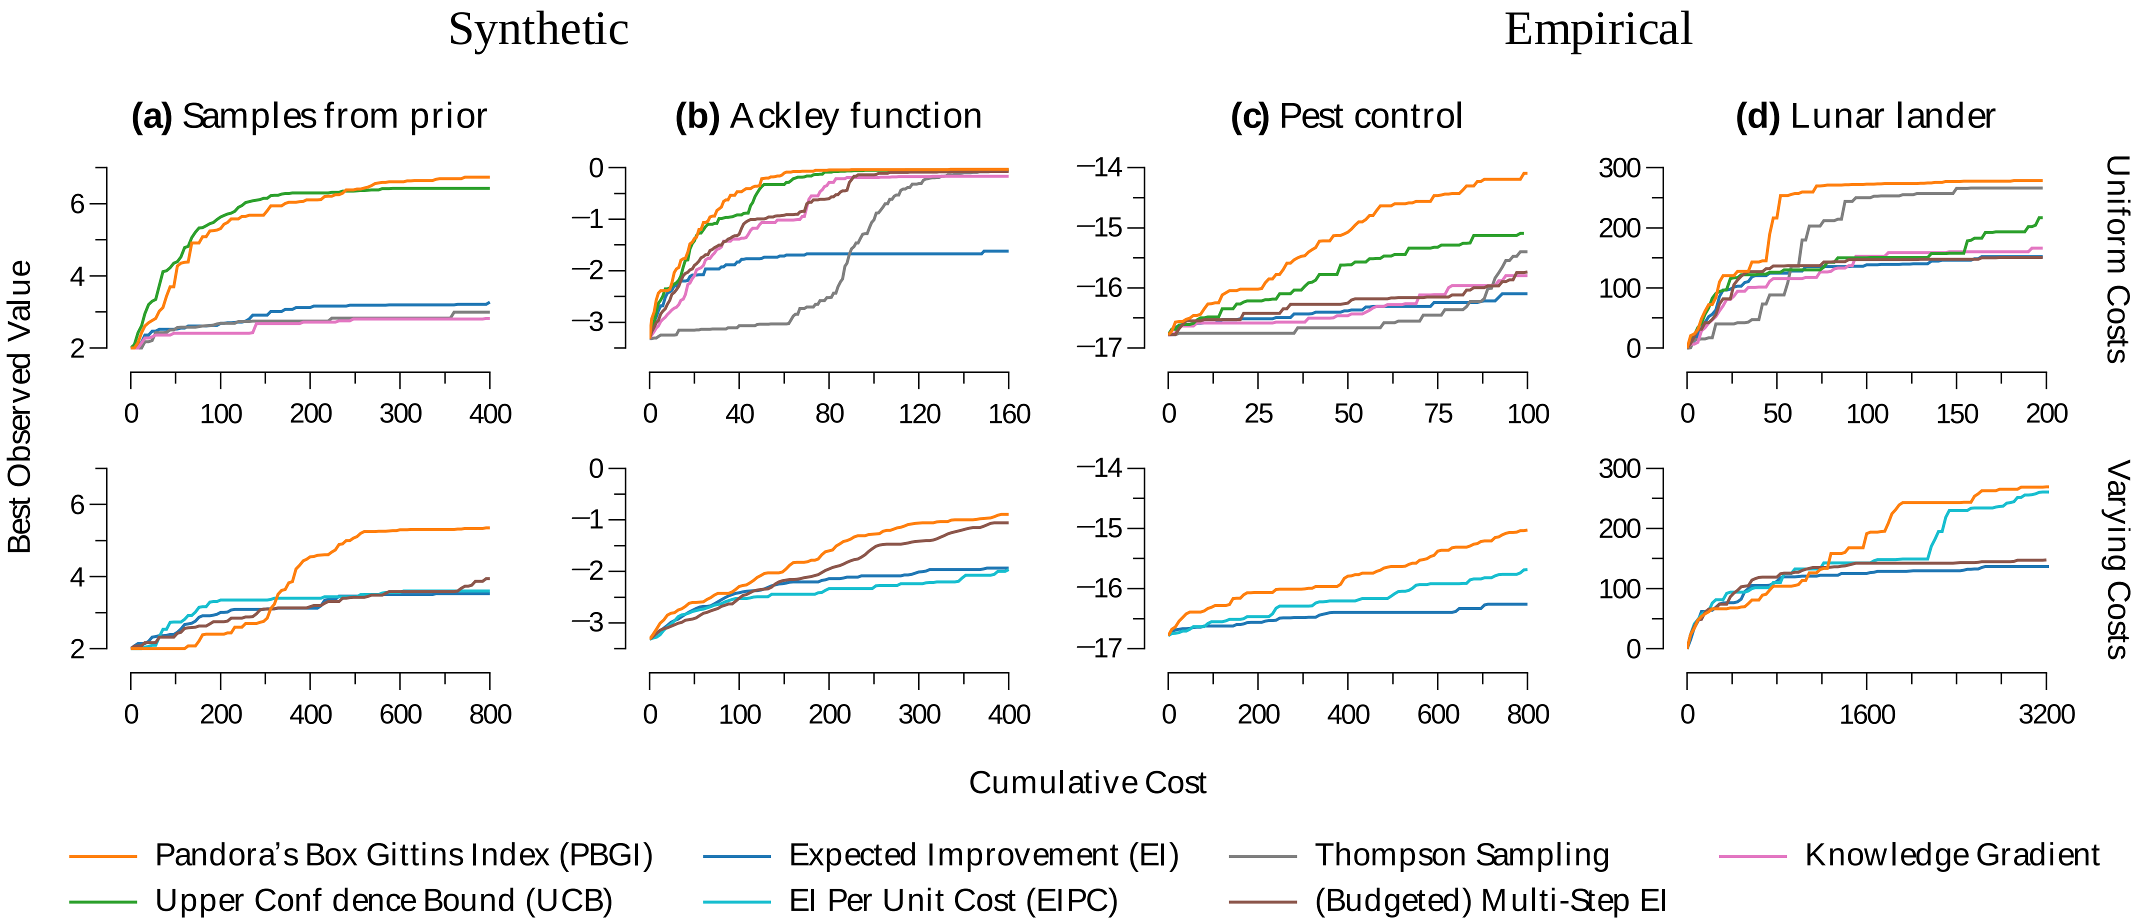

EI and EIPC policy can be arbitrarily worse [Astudillo, Jiang, Balandat, Bakshy, Frazier'21]

Conclusions

- Propose easy-to-compute PBGI policy for Bayesian optimization

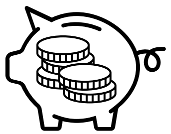

budget

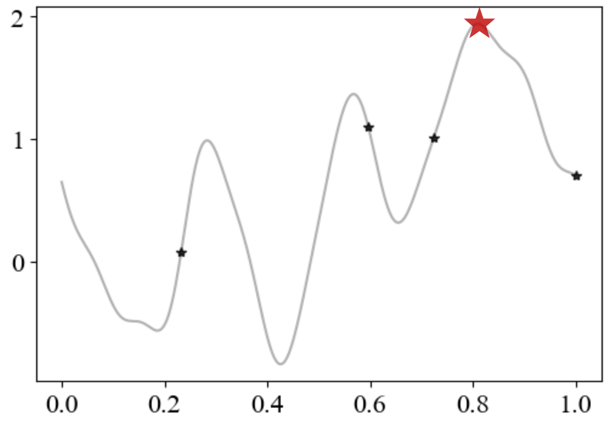

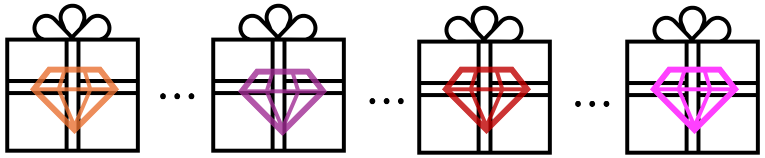

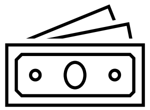

costs

Our work ⟸ Pandora's Box Gittins index

Check our preprint on arXiv!

Conclusions

- Propose easy-to-compute PBGI policy for Bayesian optimization
- Show the effectiveness of PBGI on synthetic & empirical experiments

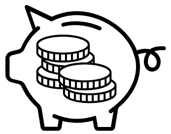

budget

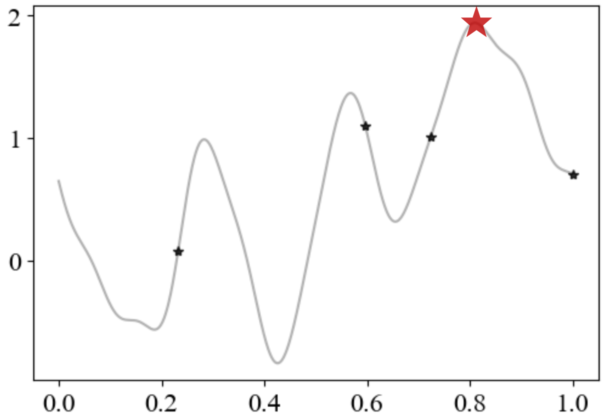

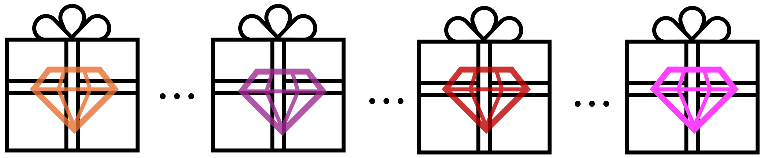

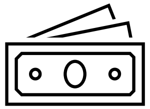

costs

Our work $\Longleftarrow$ Pandora's Box Gittins index

Check our preprint on arXiv!

Conclusions

- Propose easy-to-compute PBGI policy for Bayesian optimization
- Show the effectiveness of PBGI on synthetic & empirical experiments
- Open door for more-complex BO (partial feedback, multi-fidelity, function network, etc.) via Gittins variants (Pandora's nested boxes, "golf"-style Markovian MAB, optional inspection, etc.)

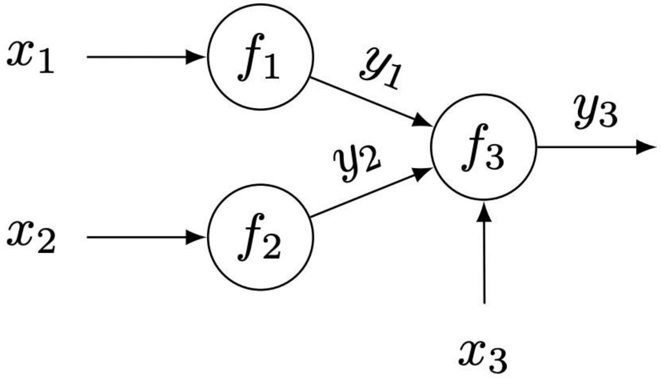

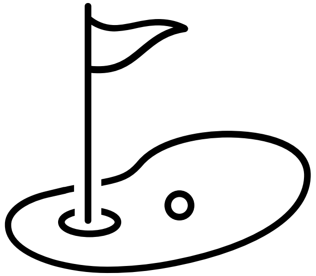

? $\Longleftarrow$ "Golf" Gittins indices

[Dumitriu, Tetali, Winkler'03]

References

- Xie, Qian, Raul Astudillo, Peter Frazier, Ziv Scully, and Alexander Terenin. "Cost-aware Bayesian optimization via the Pandora's Box Gittins index." *arXiv preprint arXiv:2406.20062* (2024).
- Frazier, Peter I. "A tutorial on Bayesian optimization." *arXiv preprint arXiv:1807.02811* (2018).
- Weitzman, Martin. *Optimal search for the best alternative*. Vol. 78, no. 8. Department of Energy, 1978.
- Aminian, Mohammad Reza, Vahideh Manshadi, and Rad Niazadeh. "Markovian search with socially aware constraints." *Available at SSRN 4347447* (2023).
- Milgrom, Paul, and Ilya Segal. "Envelope theorems for arbitrary choice sets." *Econometrica* 70, no. 2 (2002): 583-601.
- Astudillo, Raul, Daniel Jiang, Maximilian Balandat, Eytan Bakshy, and Peter Frazier. "Multi-step budgeted bayesian optimization with unknown evaluation costs." *Advances in Neural Information Processing Systems* 34 (2021): 20197-20209.