Optimization with Costly Information: a survey

SHUCHI CHAWLA

TEXAS

The University of Texas at Austin

HALG 2026

Shopping for a car...

Many possible choices...

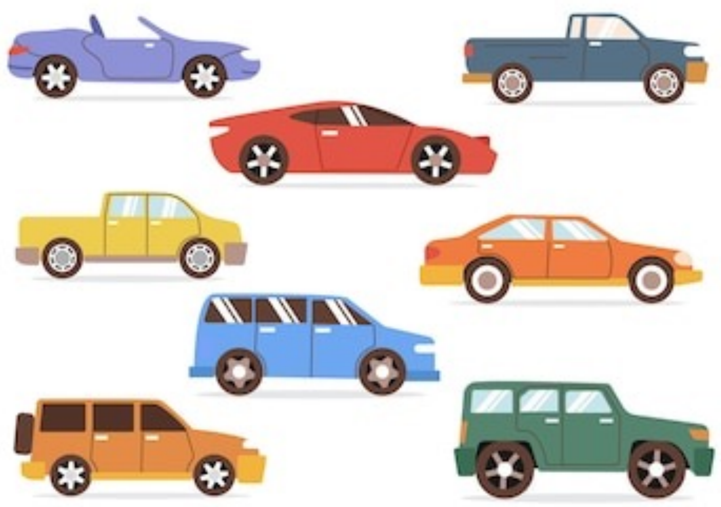

Many possible actions for each

- Read reviews
- Test drive
- Negotiate with a dealer

...

A possible "solution":

1. Pick top 5 cars based on ratings
2. Read reviews; shortlist 2 out of 5
3. Test drive the two
4. If you like one, negotiate with a dealer
   Else go back to step 1
5. Repeat until a choice is made

Goal: to select the best option without expending too much time/effort

Many settings involve costly information acquisition before optimization ...

- Hiring: read resumes and letters, interview candidates before making offers
- Procurement: evaluate suppliers before awarding contracts
- Data-driven advertising: run experiments before allocating budget
- ...

Information can be expensive

Key Question: when should an algorithm pay to learn more?

How is this different from standard stochastic optimization?

- Stochastic/online optimization: algorithm observes information for free in a fixed order; or makes decisions without observing it
- This setting: whether to obtain information is itself a decision variable

Why is this challenging?

- Task dependence: what is worth learning depends on the downstream objective
- Adaptivity: what we learn changes what and how we should inspect next

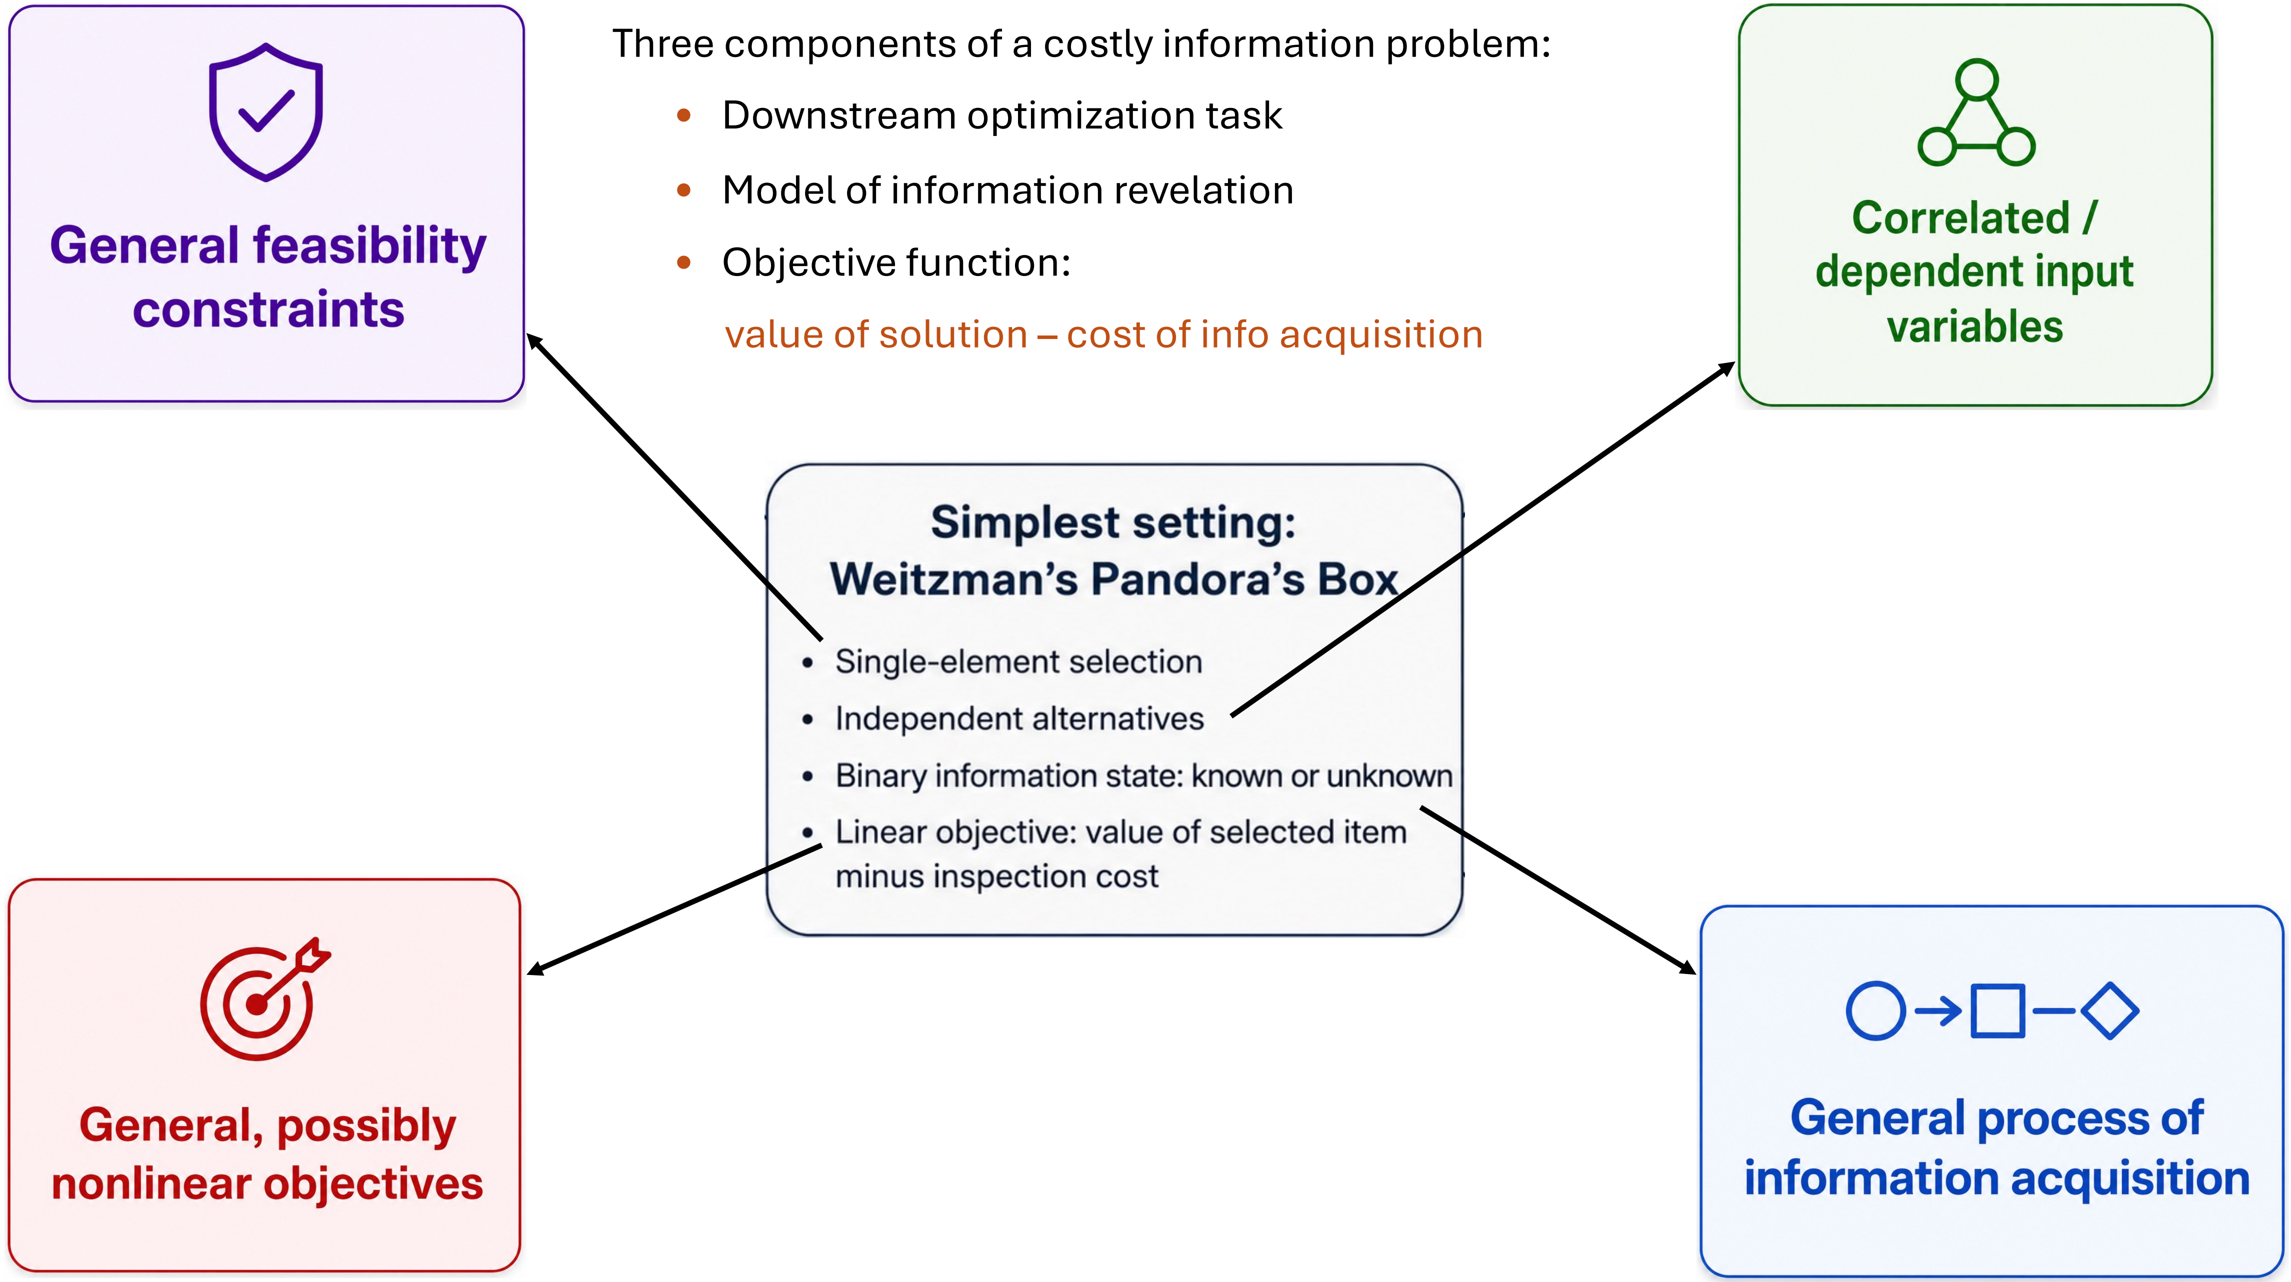
General feasibility constraints
Three components of a costly information problem:
Downstream optimization task
Model of information revelation
Objective function:
value of solution – cost of info acquisition
Correlated / dependent input variables
Simplest setting: Weitzman's Pandora's Box
Single-element selection
Independent alternatives
Binary information state: known or unknown
Linear objective: value of selected item minus inspection cost
General, possibly nonlinear objectives
General process of information acquisition

Classical special case: Weitzman's Pandora's Box [1979]

- Given: n boxes with unknown values $X_i \sim D_i$ and opening costs c_i.
- Optimization problem: select one alternative
- Information acquisition: opening box i costs c_i
  - Closed box ≡ no information
  - Open box ≡ full knowledge of value
- Algorithmic task: At every step decide whether to open some closed box i, or to stop and accept the best value seen so far.
- Maximize $X_{\text{selected}} - c_{\text{opening}}$

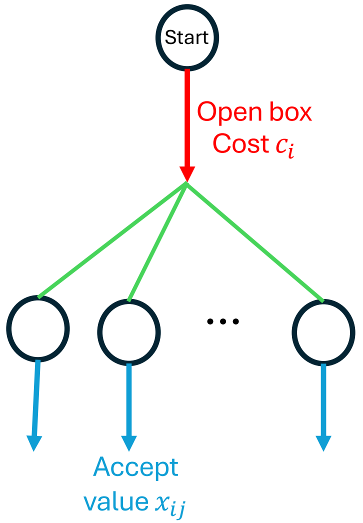

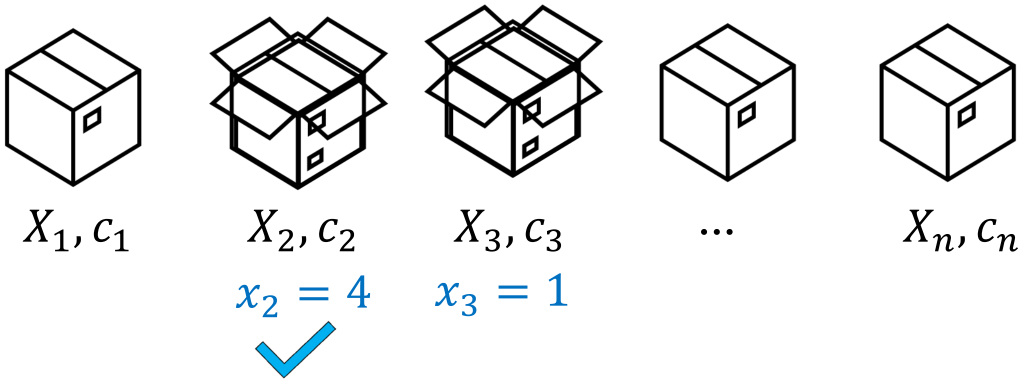

Classical special case: Weitzman's Pandora's Box [1979]

- Given: n boxes with unknown values $X_i \sim D_i$ and opening costs c_i.
- Optimization problem: select one alternative
- Information acquisition: opening box i costs c_i
  - Closed box ≡ no information
  - Open box ≡ full knowledge of value
- Algorithmic task: At every step decide whether to open some closed box i, or to stop and accept the best value seen so far.
- Maximize $X_{\text{selected}} - c_{\text{opening}}$

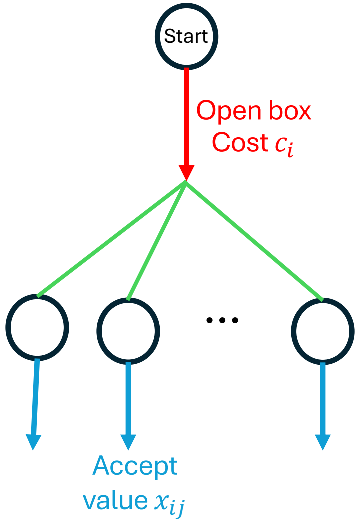

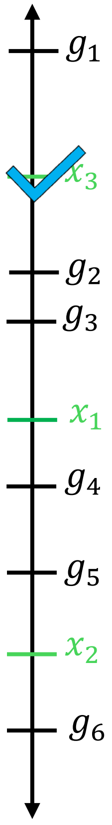

> **Weitzman's algorithm: index-based policy**
> - Compute the "Gittins index" g_i of every box
> - Open boxes in decreasing order of index
> - Until we find a value larger than all remaining indices
> - Stop and return the largest value

Weitzman's algorithm is optimal for Pandora's Box. – Let's prove it!

Technique: amortization and surrogate costs

Key idea: amortize the action costs down to the leaves.

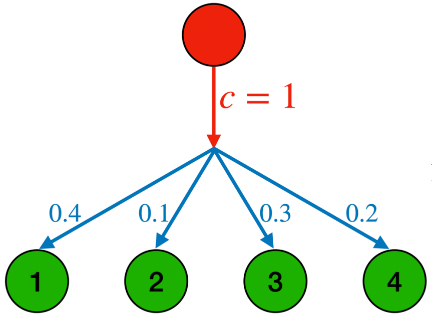

Cost share	0.5	2	1	1.5
Surrogate value	0.5	0	2	2.5

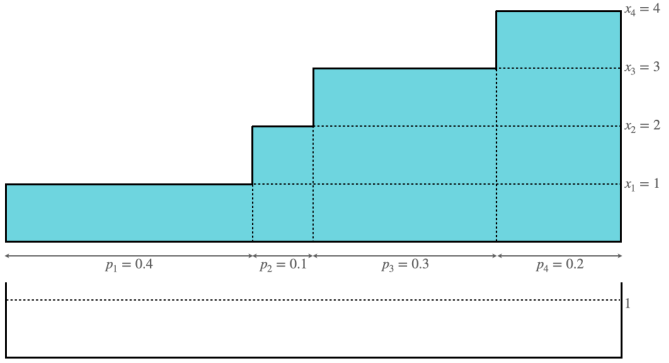

Water draining amortization: cost shares are charged to the max value nodes

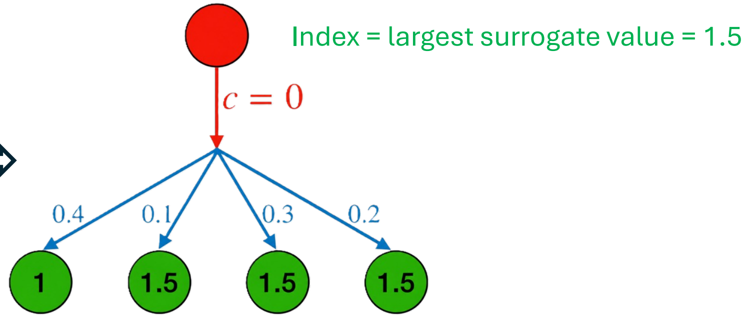

Surrogate value = value of node - amortized cost share

$\mathrm{E}[\text{cost share}] = \text{action cost } c$

Promise of payment: if a node with ~~cost share > 0 is realized~~ surrogate value = index is realized, the alg must accept at that node.

Weitzman's algorithm: index-based policy

- Open boxes in decreasing order of index
- Until we find a value larger than all remaining indices
- Stop and return the largest value

Claim 1: For any alg and any amortization, surrogate value of selection ≥ actual value.

Claim 2: For any alg satisfying promise of payment, surrogate value of selection = actual value.

Claim 3: Weitzman's algorithm satisfies promise of payment and accepts the largest realized surrogate value.

⟹ optimality

Extension to multi-stage information acquisition

General process of information acquisition

- Optimization problem: select one alternative
- Information acquisition:
  - Sequence of information revelation steps
  - Each costly step further refines distribution
- Algorithmic task: At every step decide whether to learn more about i, or to stop and accept the best value seen so far.

References:
- Weitzman '79 defines indices
- Dumitrescu Tetali Winkler '03 extend indices to Markov chains
- Amortization idea is due to Kleinberg Waggoner Weyl '16
- Water draining proof is due to C. Christou Harlev Scully '26

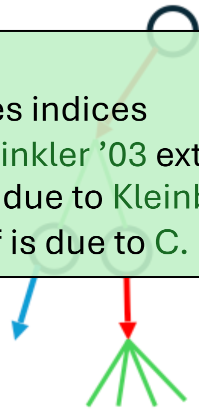

Optimal algorithm: index-based policy
- Compute the "Gittins index" g_i of every state
- Advance the state with the largest index
- If this leads to an "accept" action, stop

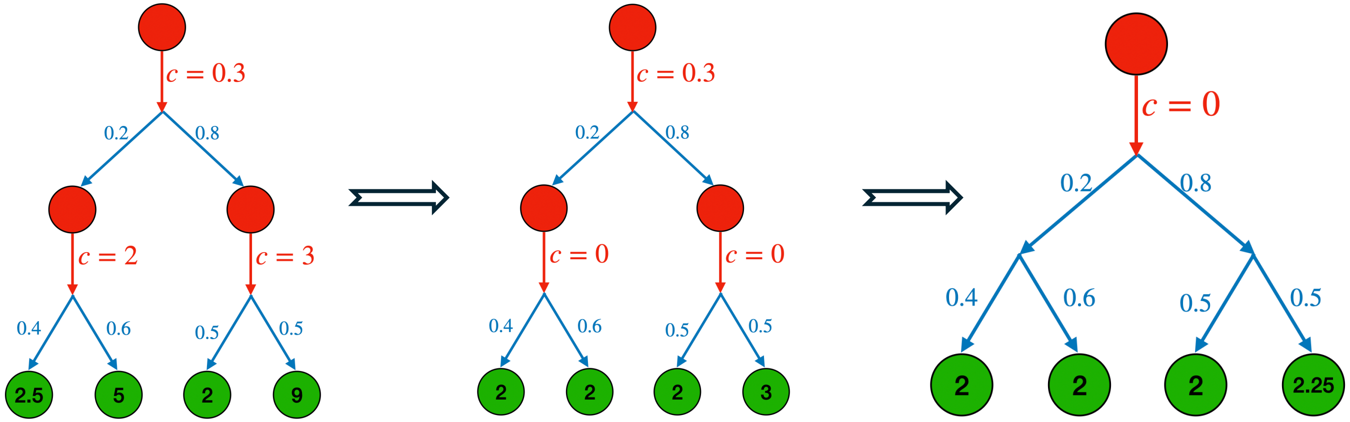

The same water draining amortization and analysis works!

Extension to general feasibility constraints

Concept: Frugality

[Singla '18, Gupta Jiang Scully Singla '20]

General feasibility constraints

- Optimization problem: select subset $S \in \mathcal{F} \longleftarrow$ feasibility constraint
- Information acquisition: each alternative lies in a box; has binary state: known or unknown
- Algorithmic task: At every step decide whether to open some closed box i, or to stop and accept the best feasible set seen so far.

Recall:

- For every algorithm, surrogate value collected $\geq$ actual value collected
- And, if the algorithm satisfies promise of payment, surrogate value collected = actual value

Question: can we maximize surrogate value collected while satisfying promise of payment?

Singla's idea: use a greedy algorithm – select a box to open based on its index and previous revealed info; must accept if value $\geq$ index.

Frugality = performance ratio of greedy algorithm; Approximation ratio of Singla's algorithm = frugality of $\mathcal{F}$.

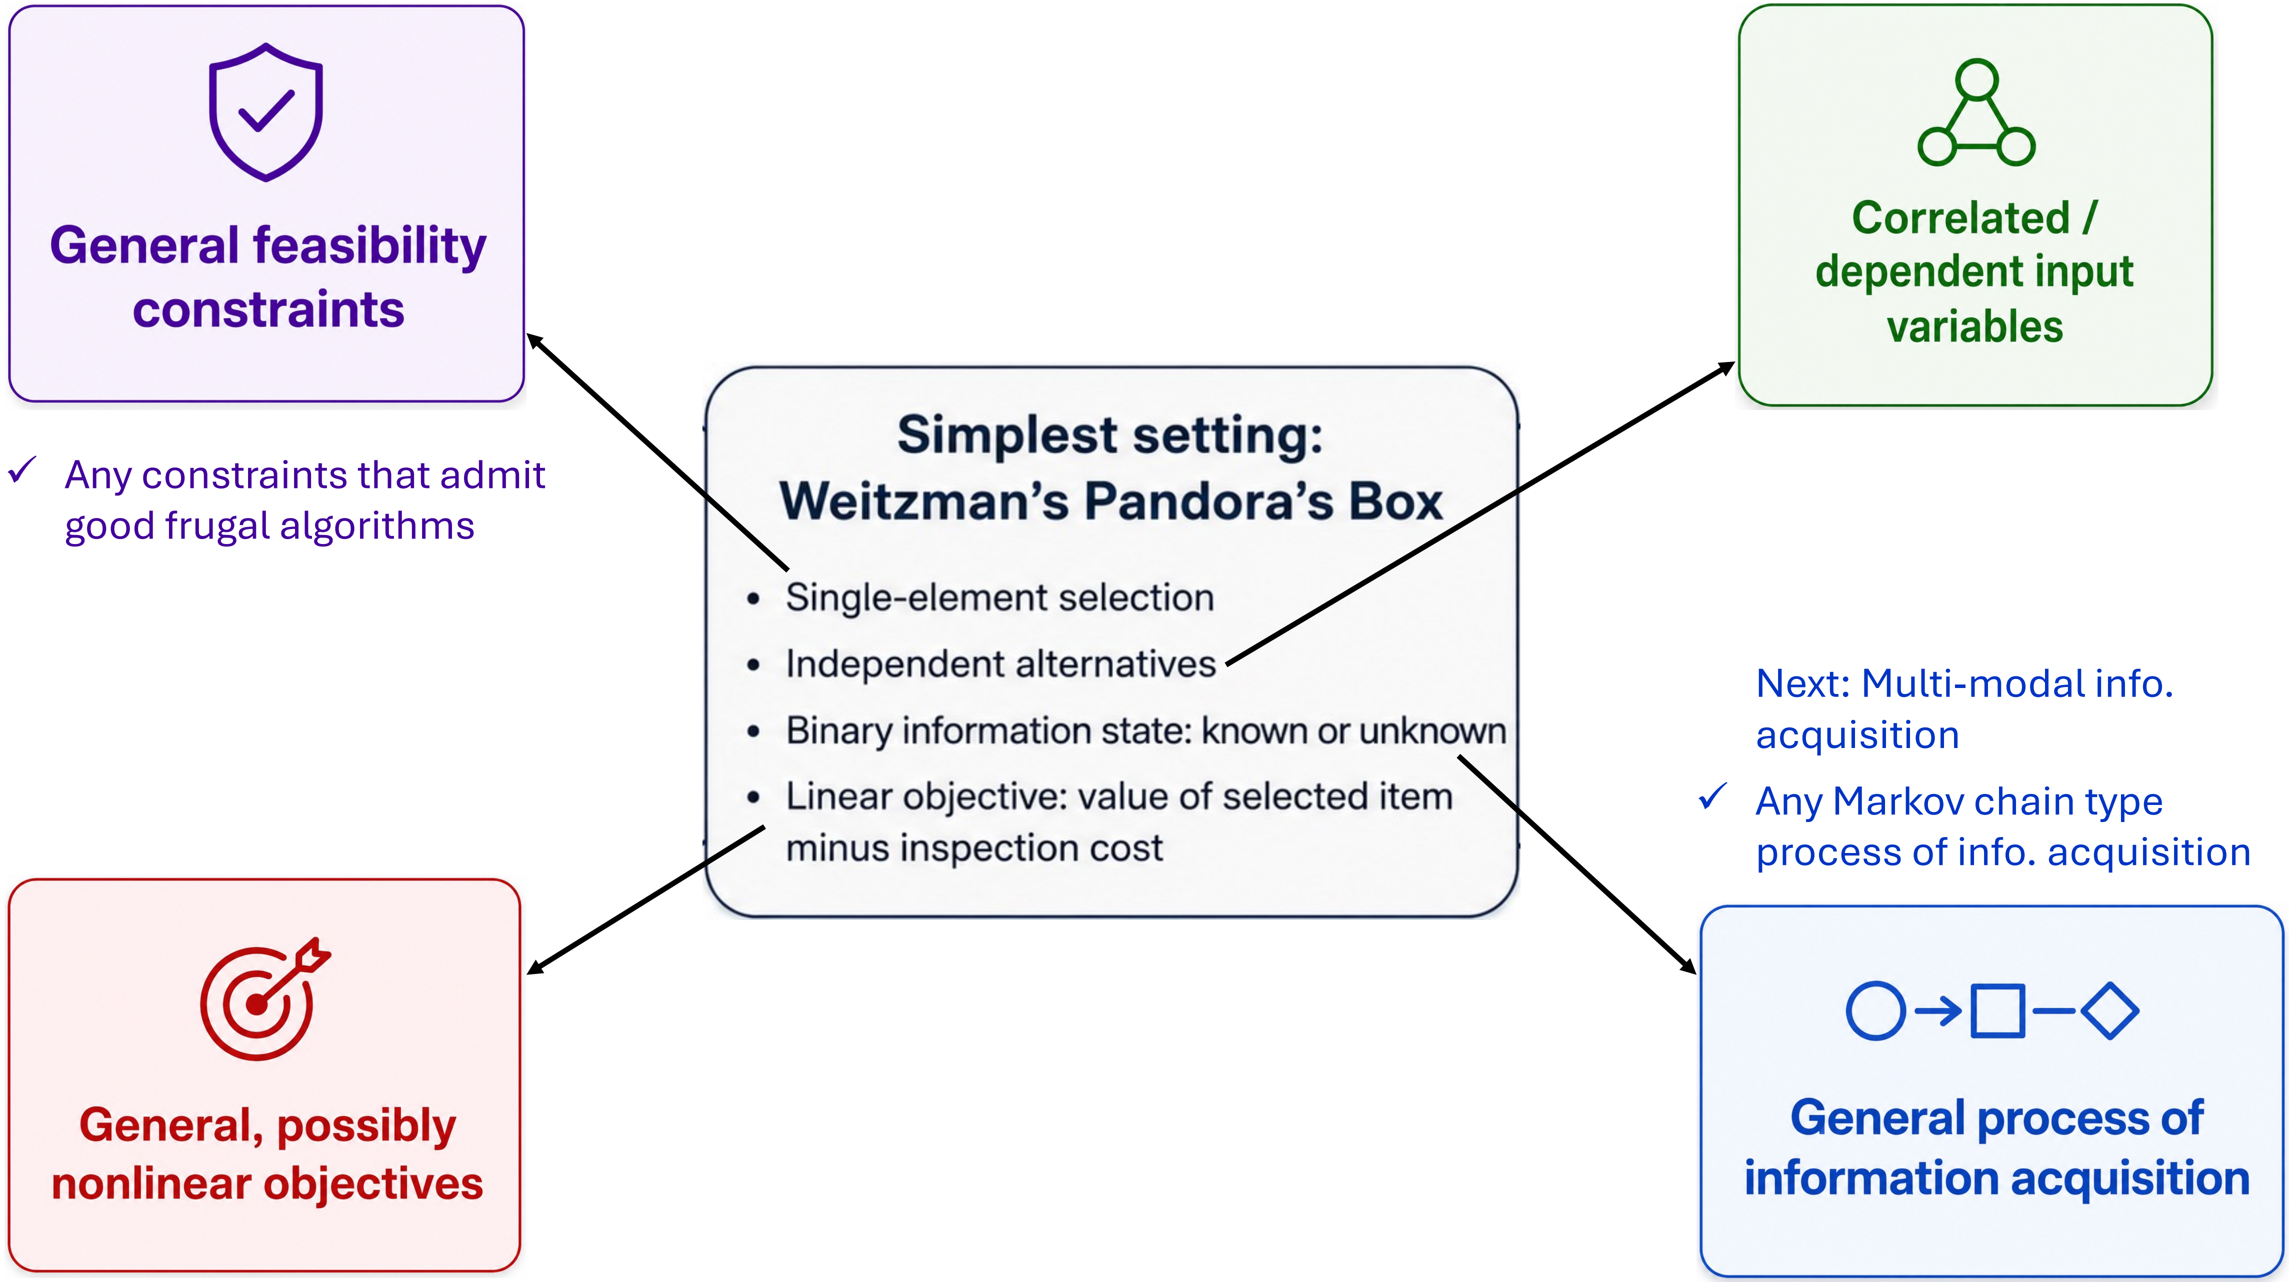

General feasibility constraints
✓ Any constraints that admit good frugal algorithms
Correlated / dependent input variables
Simplest setting: Weitzman's Pandora's Box
• Single-element selection
• Independent alternatives
• Binary information state: known or unknown
• Linear objective: value of selected item minus inspection cost
Next: Multi-modal info. acquisition
✓ Any Markov chain type process of info. acquisition
General, possibly nonlinear objectives
General process of information acquisition

Extension to multi-modal information acquisition

General process of information acquisition

Markov Decision Process $\mathcal{M}_i$:
one per alternative X_i

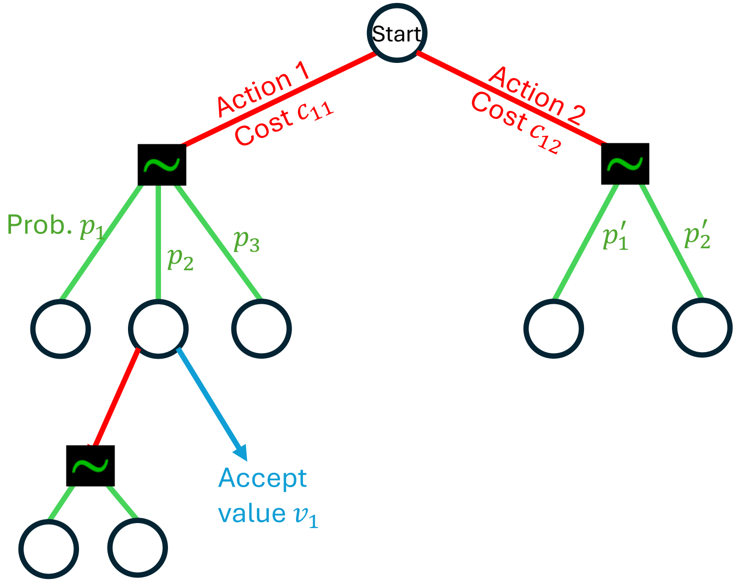

Two components in the algorithm

- Which alternative to explore?
- How to explore it?

- At every step, choose an MDP $\mathcal{M}_i$ and take an action in it.
- Process terminates when some "accept" action is taken.

Old problem: Bandit Super Process [Nash'73, Whittle'80, Glazebrook'82, ...]

But much recent interest:
[Doval'18, Beyhaghi-Kleinberg'19, Ke-Villas-Boas'19, Aouad-Ji-Shaposhnik'20 Beyhaghi-Cai'22, Fu-Li-Liu'22, Scully-Doval'24, ...]

See survey by Beyhaghi & Cai (2023)

Extension to multi-modal information acquisition

General process of information acquisition

- Global ordering is not index-based
  - Whether to evaluate X_1 or X_2 first depends on what we know about X_3
- Optimal local decisions depend on what else we know.
  - Which action to take in an MDP depends on our state in other MDPs.
- Actions that are never locally optimal may be globally optimal!
- NP-hard, even with just two action choices per alternative [Fu Li Liu’22]

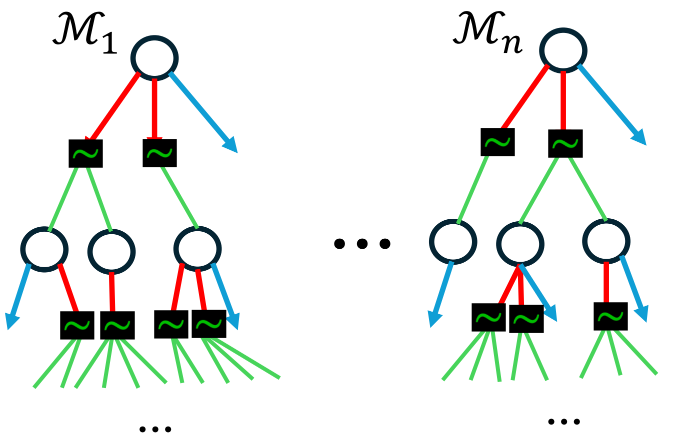

Concept: Commitment Policies and Commitment Gap

[Beyhaghi Kleinberg '19]

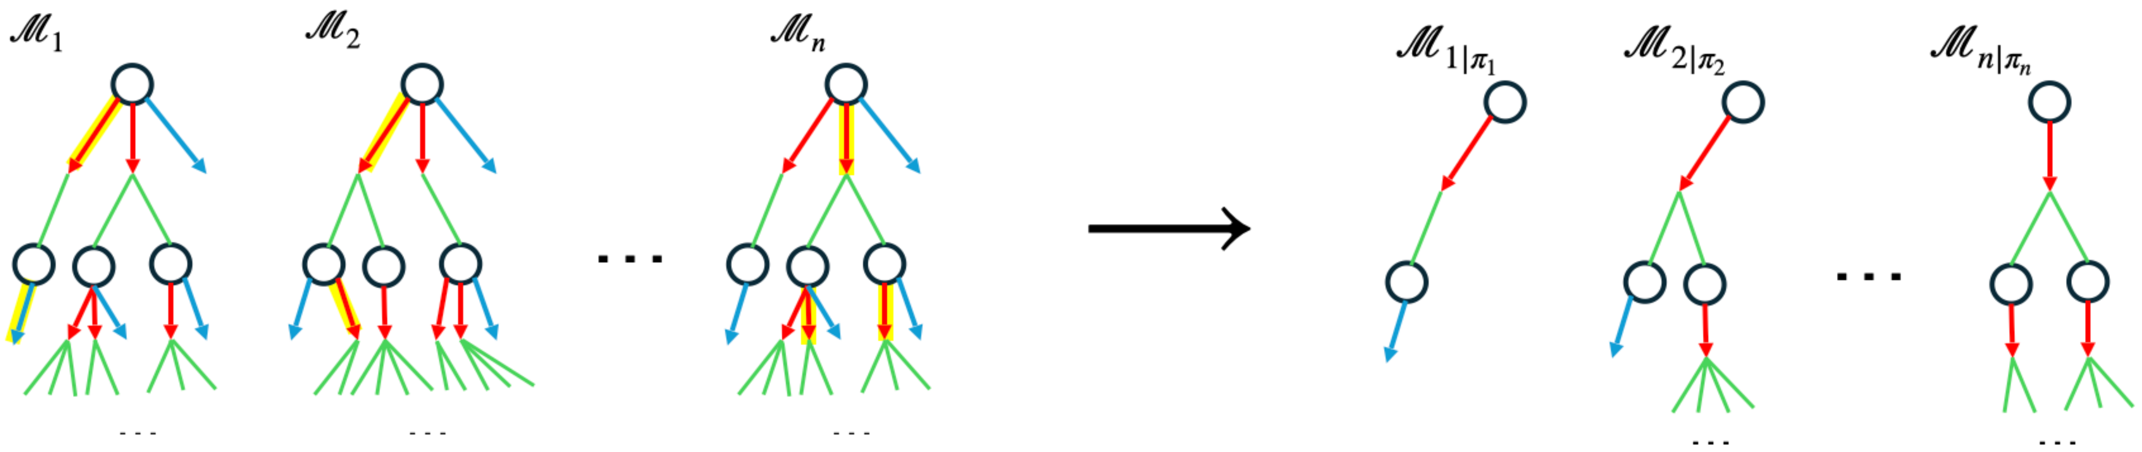

- A commitment π_i for the MDP $\mathcal{M}_i$ is a mapping from all states s_i to actions $a \in A(s_i)$.
- For any $\Pi = (\pi_1, \dots, \pi_n)$ we can efficiently find $\mathrm{OPT}_{|\Pi}$

$$\text{Commitment Gap} = \min_{\Pi} \frac{\mathrm{OPT}}{\mathrm{OPT}_{|\Pi}}$$

Technique: Ex Ante Relaxation

[Bowers Lindgren Waggoner ’26]

- Original problem: Select one car to maximize $\mathrm{E}[X_{\text{selected}} - c(\text{actions})]$
- New relaxation: Select a subset S of cars to maximize $\mathrm{E}[\Sigma_{i \in S} X_i - c(\text{actions})]$
  such that $\mathrm{E}[|S|] \leq 1$

	$X_1 \sim D_1$	$X_2 \sim D_2$	$X_3 \sim D_3$	$X_4 \sim D_4$
Prob. of acceptance:	0.1	0.3	0.5	0.1

An algorithm based on the ex ante relaxation

[C. Christou Dang '25]

- Solve the ex ante relaxation $\Longrightarrow$ obtain acceptance probability q_i for each alternative i
  - Can solve optimally using a linear program of size $\sum_i |\mathcal{M}_i|$
- For each alternative i, find the "optimal" local policy π_i that accepts with probability q_i
- Commit to the local actions $(\pi_1, \dots, \pi_n)$ and run the index policy of Weitzman and DTW03

	$X_1 \sim D_1$	$X_2 \sim D_2$	$X_3 \sim D_3$	$X_4 \sim D_4$
Prob. of acceptance:	0.1	0.3	0.5	0.1
Prescribed local actions:		Read reviews Test drive if ★ ★ ★+		Read reviews Buy if ★ ★ ★ ★ ★

Why does this work?

[C. Christou Dang ’25]

- Solve the ex ante relaxation $\Longrightarrow$ obtain acceptance probability q_i for each alternative i
- For each alternative i, find the “optimal” local policy π_i that accepts with probability q_i
- Commit to the local actions $(\pi_1, \dots, \pi_n)$ and run the index policy of Weitzman and DTW03

$$\underbrace{\text{Ex Post OPT}(\mathcal{M}_1^{\pi_1}, \dots, \mathcal{M}_n^{\pi_n})}_{\text{Our algorithm}} \;\underset{\text{What is this gap?}}{?}\; \text{Ex Ante OPT}(\mathcal{M}_1^{\pi_1}, \dots, \mathcal{M}_n^{\pi_n}) \underset{\text{By construction}}{=} \text{Ex Ante OPT}(\mathcal{M}_1, \dots, \mathcal{M}_n) \underset{\text{Relaxation}}{\geq} \underbrace{\text{Ex Post OPT}(\mathcal{M}_1, \dots, \mathcal{M}_n)}_{\text{Benchmark}}$$

Correlation gap $\Longrightarrow$ Approximation ratio

- The correlation gap of the max function is $e/(e-1)$. [Agrawal-Ding-Saberi-Ye'10, Yan'11]

Theorem: The approximation ratio of the above algorithm for single element selection over any MDPs is at most $e/(e-1)$.

$W_1 = 0 \text{ or } 1 \text{ w.p.} ½$

$W_2 = 0 \text{ or } 1 \text{ w.p.} ½$

Ex Post expected value = $\frac{3}{4}$

Ex Ante expected value = 1

$$\mathrm{E}[\max_i W_i^{\pi_i}] \xleftrightarrow{\text{Correlation gap of the max function}} \mathrm{E}[\max_i Z_i]$$

where $(Z_1, \ldots, Z_n)$ are correlated random variables such that:

- the marginal distribution of Z_i is identical to the distribution of $W_i^{\pi_i}$
- the expected maximum is maximized

$\parallel$ $\parallel$

$$\text{Ex Post OPT}(\mathcal{M}_1^{\pi_1}, \ldots, \mathcal{M}_n^{\pi_n}) \;?\; \text{Ex Ante OPT}(\mathcal{M}_1^{\pi_1}, \ldots, \mathcal{M}_n^{\pi_n}) = \text{Ex Ante OPT}(\mathcal{M}_1, \ldots, \mathcal{M}_n) \geq \text{Ex Post OPT}(\mathcal{M}_1, \ldots, \mathcal{M}_n)$$

Our algorithm

Benchmark

What is this gap?

By construction

Relaxation

Technique: Reduction to Bayesian Selection

[C. Christou Dang '25]

Fix distributions $D_1, \ldots, D_n$ and $F \subseteq 2^{[n]}$. Select $S \in F$ to maximize $\sum_{i \in S} X_i$ in expectation over $X_i \sim D_i$.

Benchmark $\longrightarrow$

- Ex Ante OPT: Maximizes sum subject to ex ante feasibility
- Ex Post OPT: Selects the optimal set after learning the X_is

Any alg satisfying promise of payment w.r.t. the $W_i^{\pi_i}$s $\longrightarrow$

- Semi-Online: Choose i; Choose threshold τ_i; Observe whether $X_i \geq \tau_i$
  - If $X_i \geq \tau_i$, must accept
  - If $X_i < \tau_i$, continue and may revisit i with a different threshold later

- Fixed-Order Online: Given i in adversarial order, see $X_i \sim D_i$, decide accept or reject
- Free-Order Online: Visit i in optimal order, see $X_i \sim D_i$, decide accept or reject

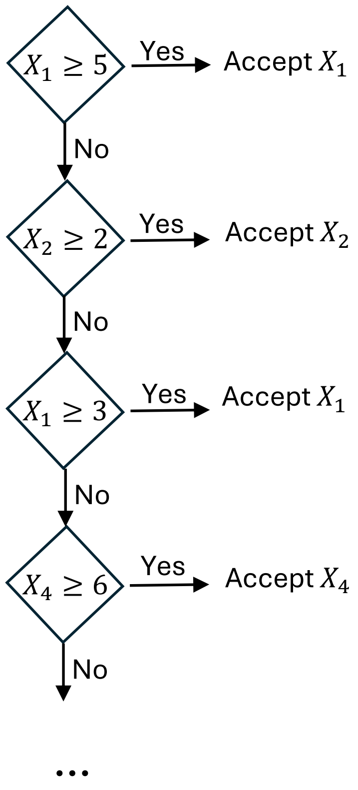

Technique: Reduction to Bayesian Selection

[C. Christou Dang '25]

Fix distributions $D_1, \ldots, D_n$ and $F \subseteq 2^{[n]}$. Select $S \in F$ to maximize $\sum_{i \in S} X_i$ in expectation over $X_i \sim D_i$.

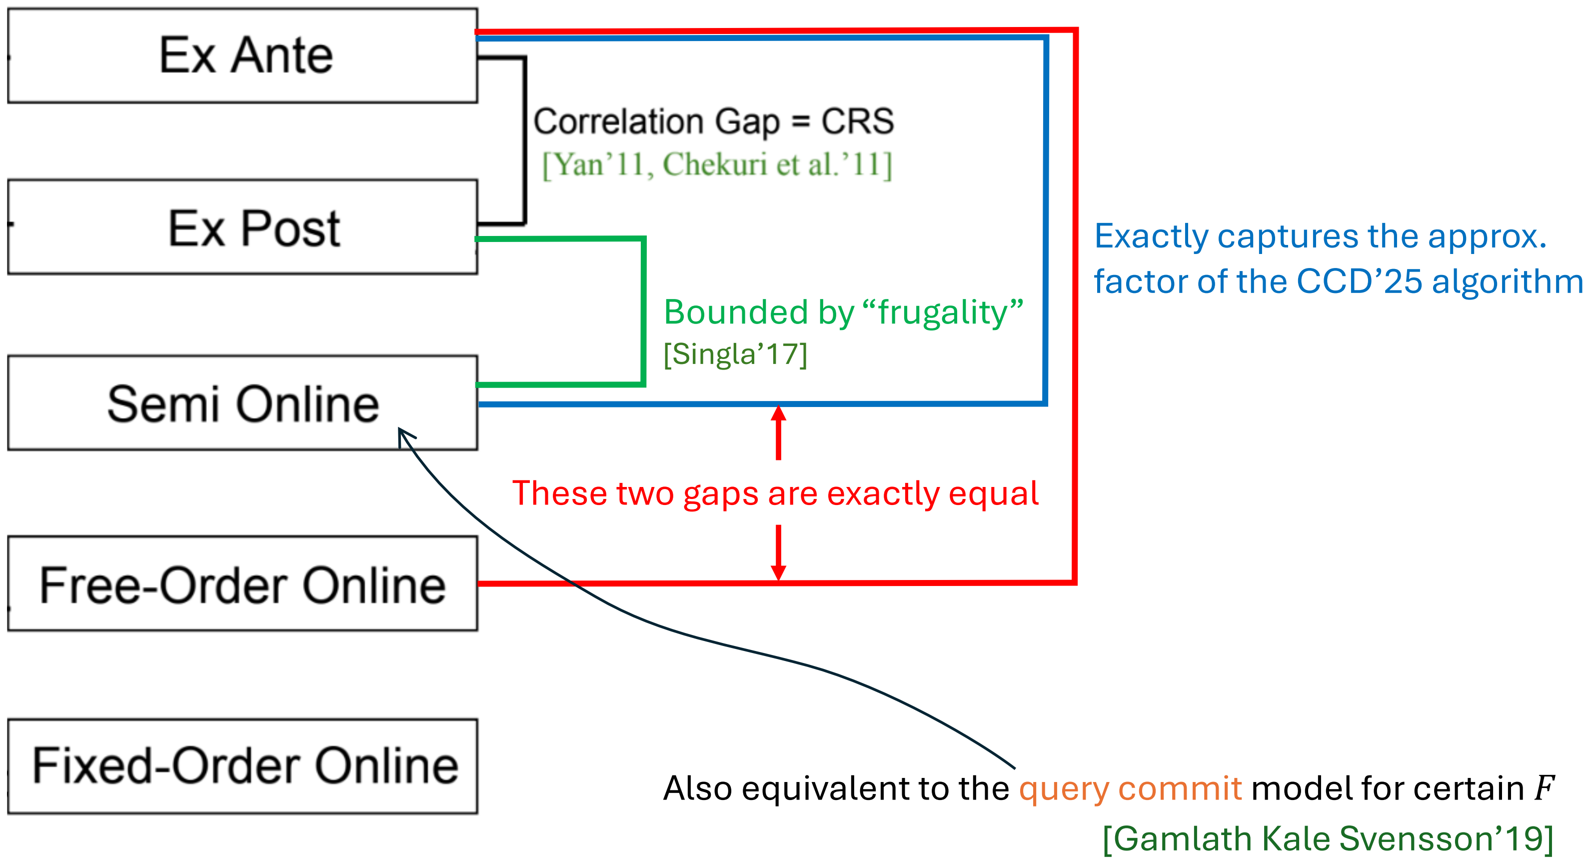

Technique: Reduction to Bayesian Selection

[C. Christou Dang '25]

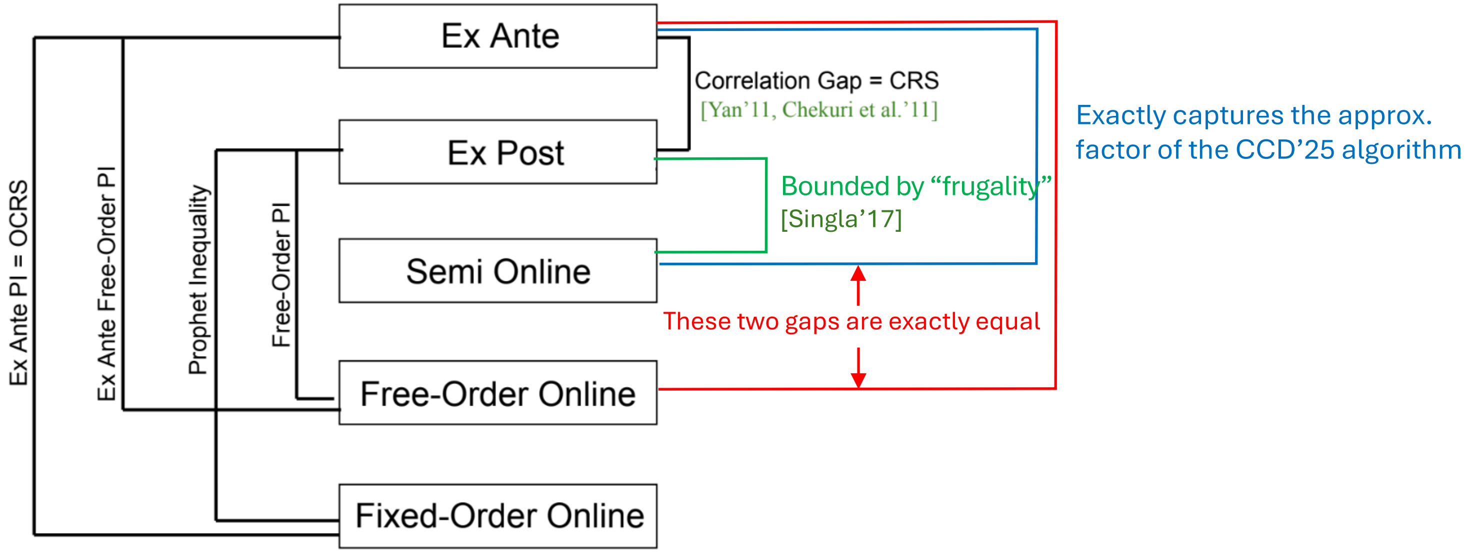

Results for optimization with costly information under constraint F

Constraint F	Matroid	k-system	Knapsack	Bipartite Matching
Correlation Gap	$e/(e-1)$ [Yan'10]	$k+1$ [Yan'10]	3.13 [Jiang Ma Zhang'21]	1.96 [Nuti-Vondrak'24]
Frugality	1	k [Singla'17]	2 [Singla'17]	1.52 [Huang et al.'25]
Approximation factor	$e/(e-1)$	$k+1$	3.13 [Jiang Ma Zhang'21]	2.98

Summary: extension to multi-modal information acquisition

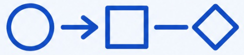

General process of information acquisition

- Maximization problems can be solved by combining
  - Ex Ante Relaxation
  - Correlation Gap + Frugality
  - Ex Ante Free Order Prophet Inequalities

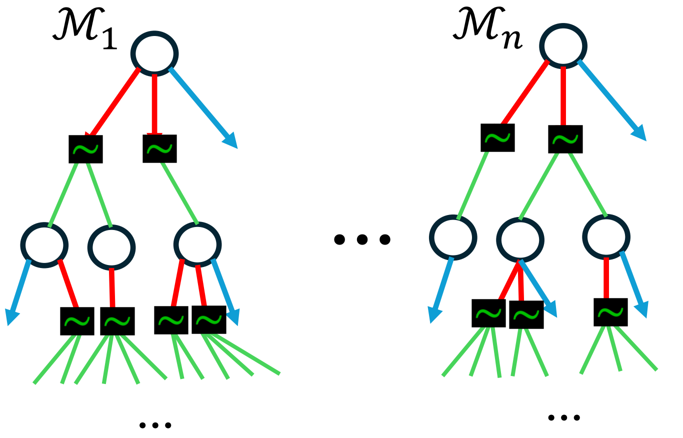

- What about minimization?

  Objective: minimize value of selection + cost of info. acquisition

Concept: Amortization for MDPs

[C. Christou Harlev Scully ’26]

Surrogate cost = value of node + amortized cost share

Index of a node = smallest surrogate cost in its subtree (best case scenario)

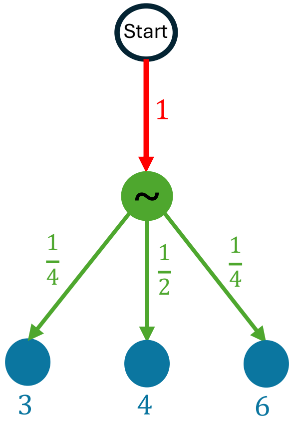

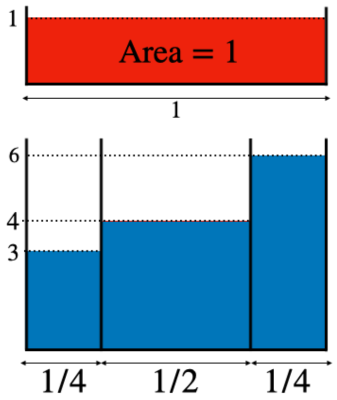

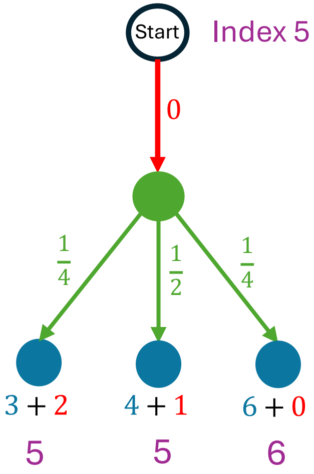

Concept: Amortization for MDPs

[C. Christou Harlev Scully '26]

Surrogate cost = value of node
+ amortized cost share

Index of a node = smallest surrogate cost
in its subtree (best case scenario)

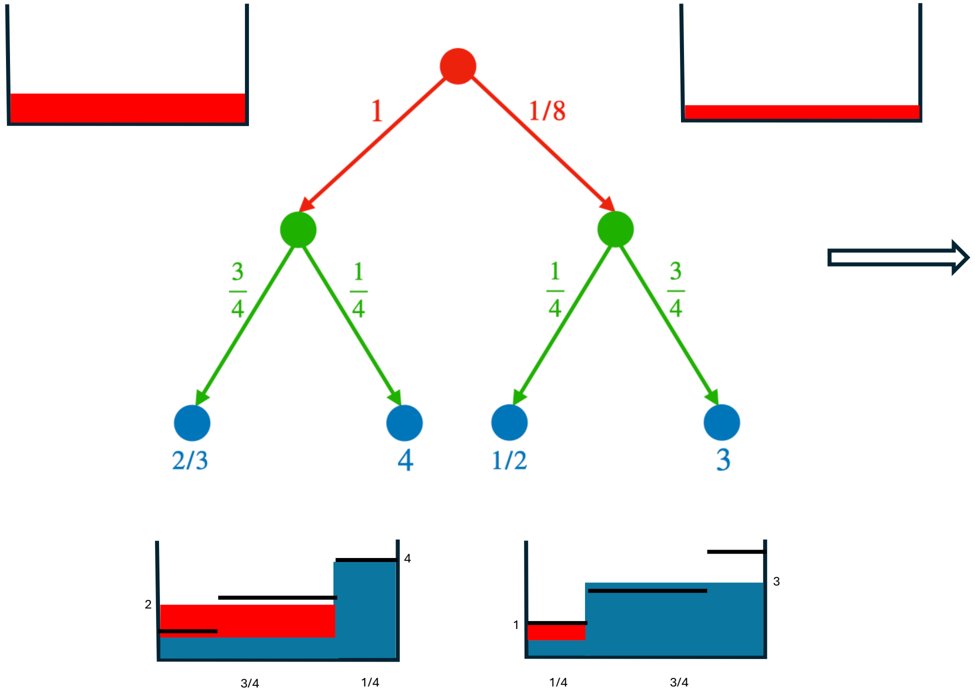

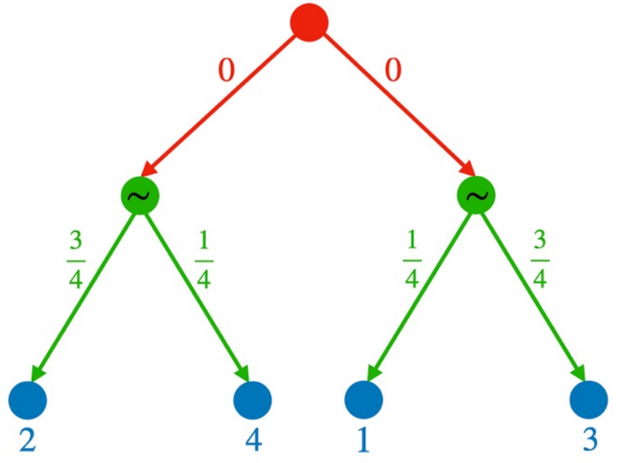

$$W_l = \begin{cases} 2 & \text{w.p. } 3/4 \\ 4 & \text{w.p. } 1/4 \end{cases} \quad W_r = \begin{cases} 1 & \text{w.p. } 1/4 \\ 3 & \text{w.p. } 3/4 \end{cases}$$

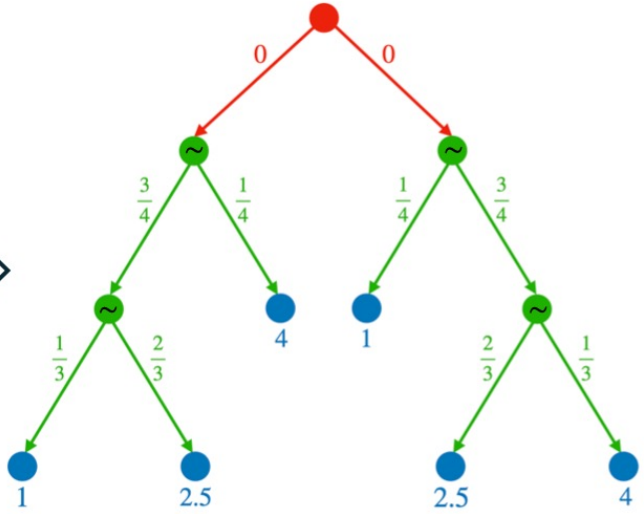

$$W^* = \begin{cases} 1 & \text{w.p.} 1/4 \\ 2.5 & \text{w.p.} 1/2 \\ 4 & \text{w.p.} 1/4 \end{cases}$$

Can we define surrogate costs W_i so that $\text{OPT} \geq \mathrm{E}[\min_i W_i]$?

- Problem: surrogate costs depend on which action is taken.
- Solution: define a common "lower bound".

We define W^* such that for **each** action a,

$$W^* \preceq_{\text{SOSD}} W^a.$$

Replacing W^a by W^* only lowers charged cost.

Concept: Amortization for MDPs

[C. Christou Harlev Scully '26]

Surrogate cost = value of node
+ amortized cost share

Index of a node = smallest surrogate cost
in its subtree (best case scenario)

Theorem: For each MDP $\mathscr{M}_i$, we can define a surrogate cost W_i^* such that:

1. For the full game $(\mathscr{M}_1, \mathscr{M}_2, \cdots, \mathscr{M}_n)$ we have $\mathrm{OPT} \geq \mathbb{E}[\min_i W_i^*]$.
2. In each local game $(\mathscr{M}_i, y)$ we have $\mathrm{OPT} = \mathbb{E}[\min(y, W_i^*)]$.

Local game $(\mathcal{M}_i, y)$

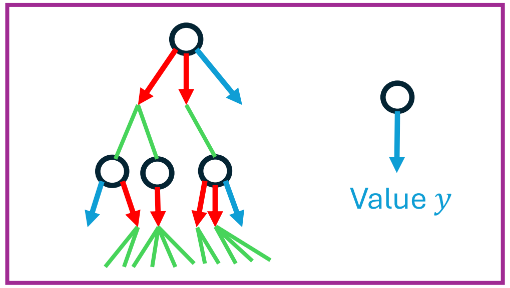

Can we define surrogate costs W_i so that OPT $\geq \mathrm{E}[\min_i W_i]$?

- Problem: surrogate costs depend on which action is taken.
- Solution: define a common "lower bound".

We define W^* such that for **each** action a,

$$W^* \preceq_{\mathrm{SOSD}} W^a.$$

Replacing W^a by W^* only lowers charged cost.

Technique: Commitment Gap via Local Approximation

[Doval Scully '24]

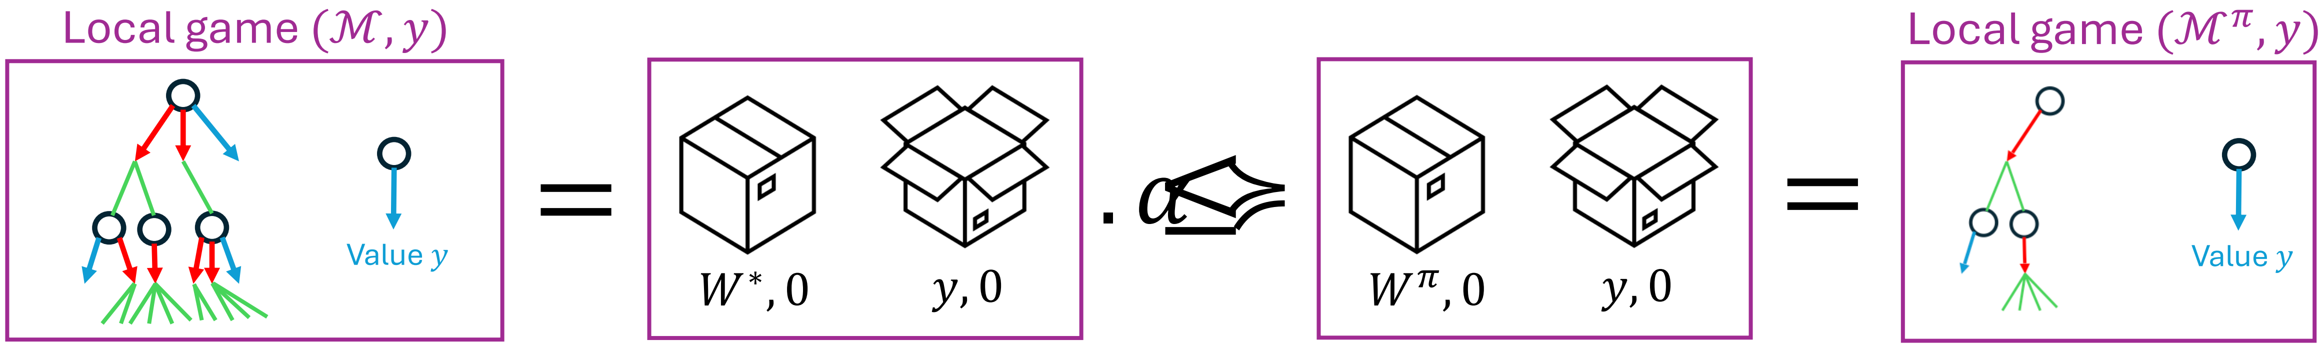

- We say that a commitment π is an α local approximation to $\mathcal{M}$ if

  αW^* second order stochastically dominates W^π

- If for all i, π_i obtains an α local approximation to $\mathcal{M}_i$,

  then $(\pi_1, \dots, \pi_n)$ obtains an α global approximation to $(\mathcal{M}_1, \dots, \mathcal{M}_n)$

- Extends automatically to objectives that admit good frugal algorithms

Local Approximation	Ex Ante Relaxation
$\mathrm{ComGap} \leq \max_i \mathrm{LocApp}(\mathscr{M}_i) \cdot \mathrm{Frugal}(\mathrm{F})$	$\mathrm{ComGap} \leq \mathrm{CorGap}(\mathrm{F}) \cdot \mathrm{Frugal}(\mathrm{F})$
• Depends on the structure of the MDPs	• Only depends on the selection constraint F
• Applies to both minimization and maximization	• Only meaningful for maximization
• Might not hold for complex MDPs	• Holds independently of the MDP structure
• Potentially overcomes the correlation gap	• Tied to the correlation gap of the constraint
• Requires computation of the commitment	• Computed via the ex-ante relaxation

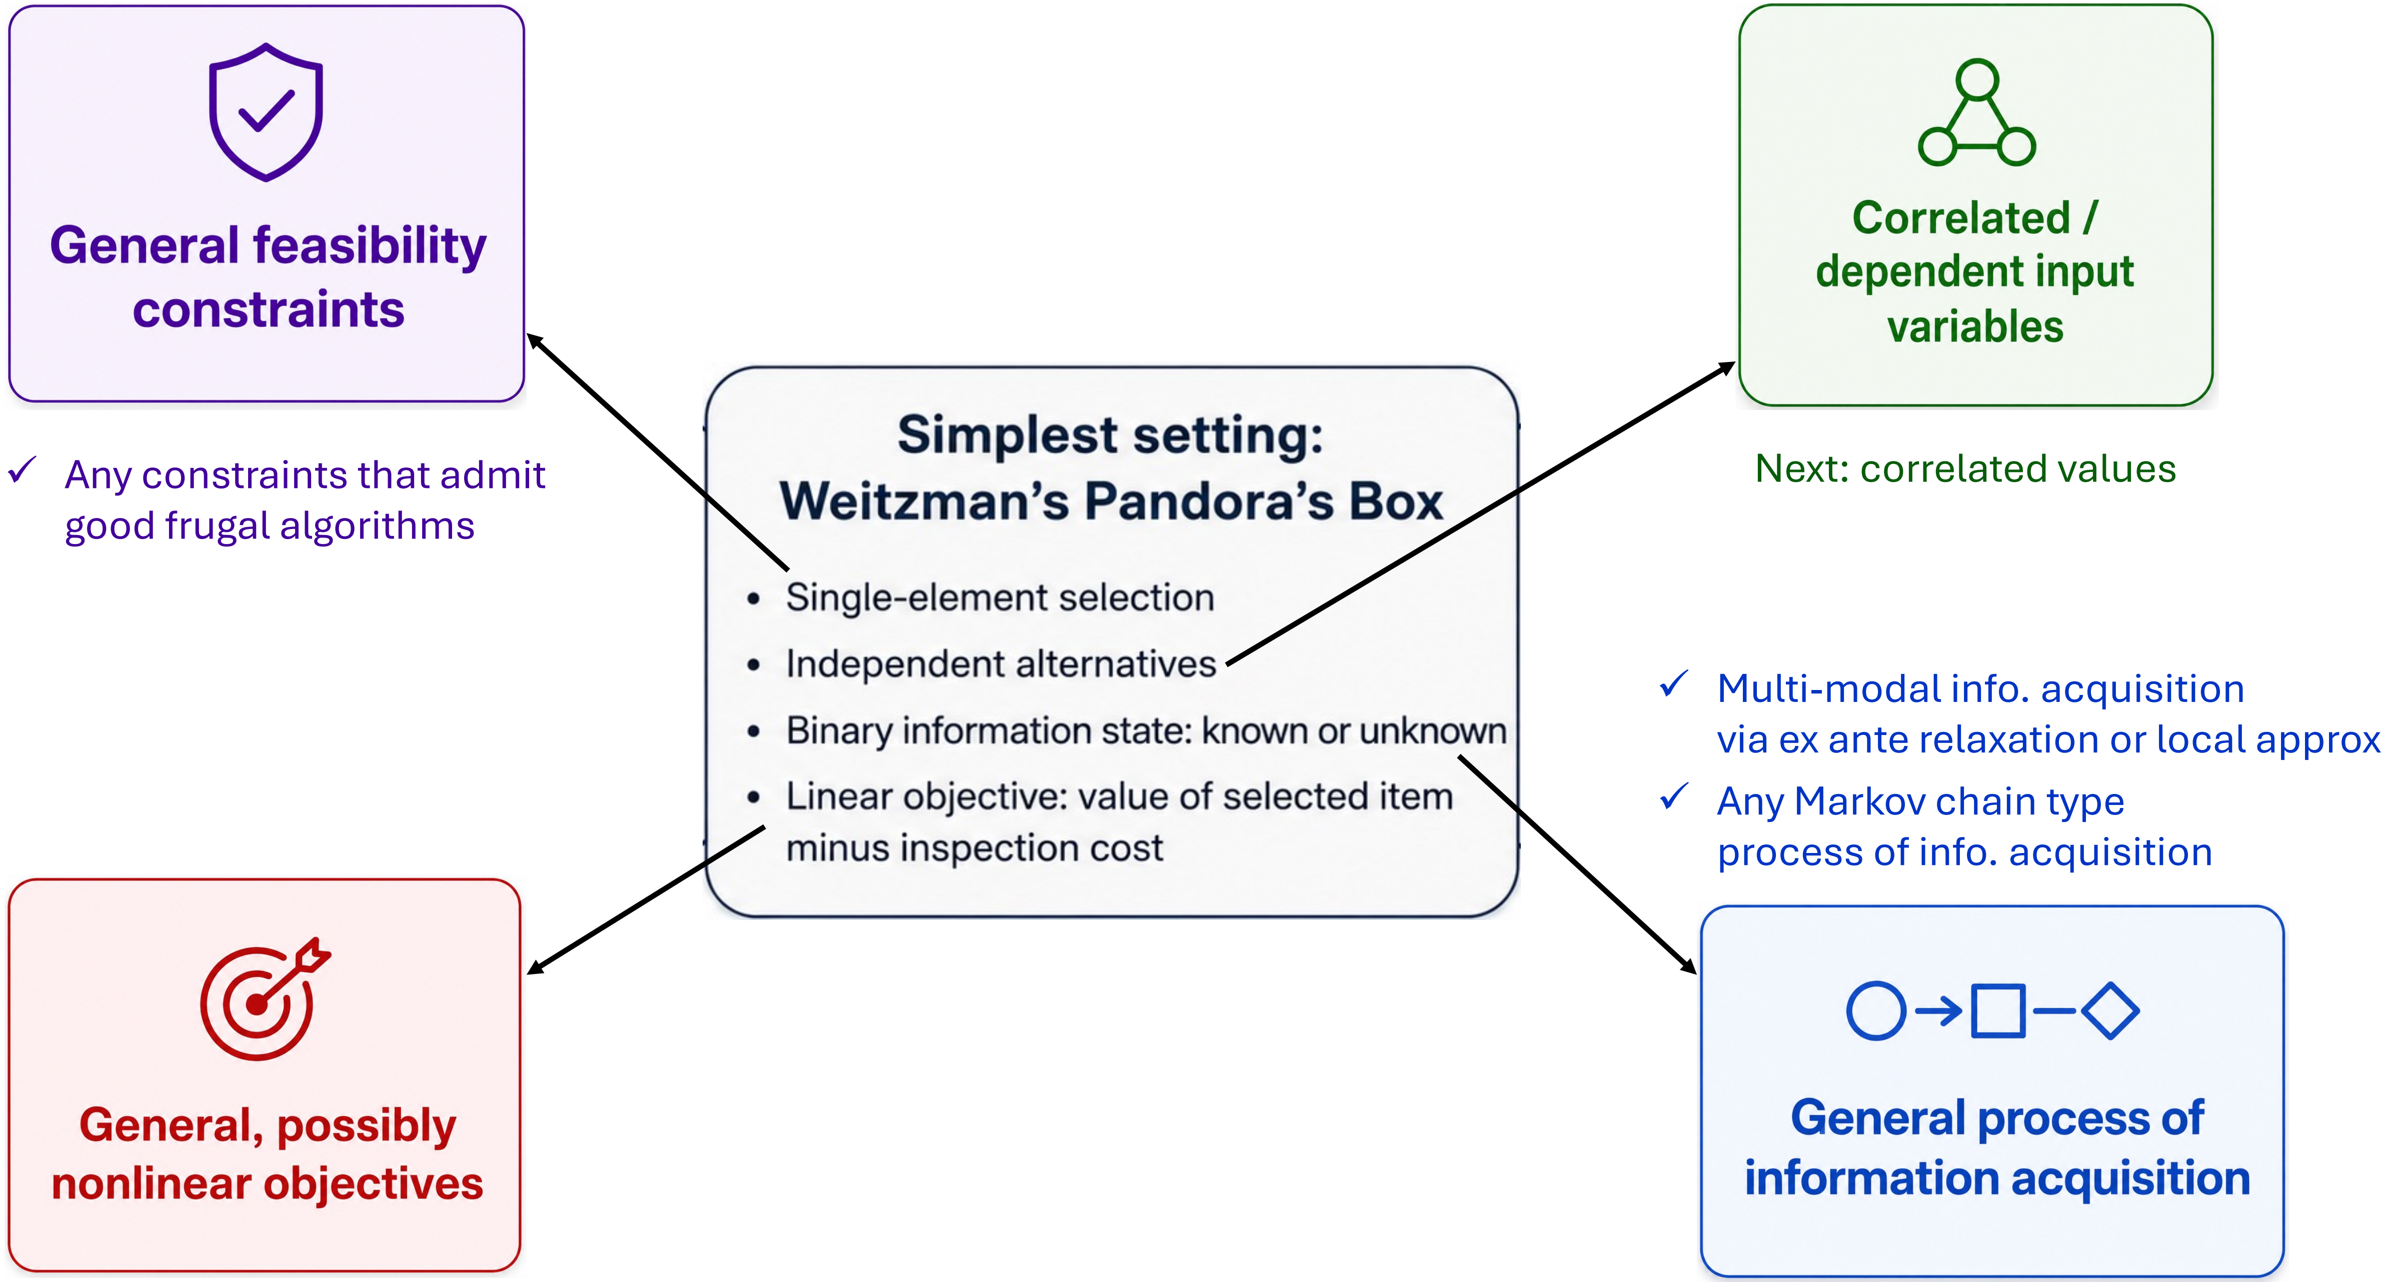
General feasibility constraints
✓ Any constraints that admit good frugal algorithms
Correlated / dependent input variables
Next: correlated values
Simplest setting:
Weitzman's Pandora's Box
• Single-element selection
• Independent alternatives
• Binary information state: known or unknown
• Linear objective: value of selected item minus inspection cost
✓ Multi-modal info. acquisition via ex ante relaxation or local approx
✓ Any Markov chain type process of info. acquisition
General, possibly nonlinear objectives
General process of information acquisition

Extension to correlated values

Correlated / dependent input variables

- Optimization problem: select one alternative
- Information acquisition: Pandora's Box
- Minimize value + opening cost
- Joint distribution over box values: $(v_1, \dots, v_n) \sim D$
  m "scenarios" representing the support of D

Example: may need to open "useless" boxes to find "useful" ones

- Let f be a random mapping from $[n]$ to a large domain.
- $S^{(i)} = \begin{cases} v_1 = f(i) \\ v_i = 0 \\ v_{i'} = \infty \quad \text{for } i' \neq 1, i \end{cases}$
- OPT = 2

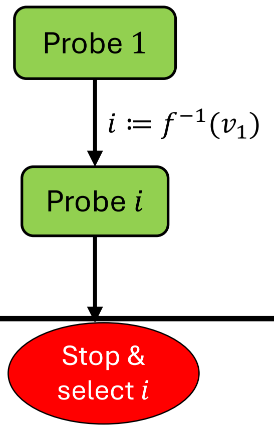

A special case: the Min Sum Set Cover problem

- All costs are 0 or ∞. I.e. boxes are "acceptable" or "unacceptable".
- For simplicity: all probing penalties are 1; Uniform distribution over m scenarios.
- Think of the scenarios as elements and boxes as sets. If an "acceptable" box is opened in a scenario, the scenario gets covered.
- Objective: minimize total penalty a.k.a. covering time.

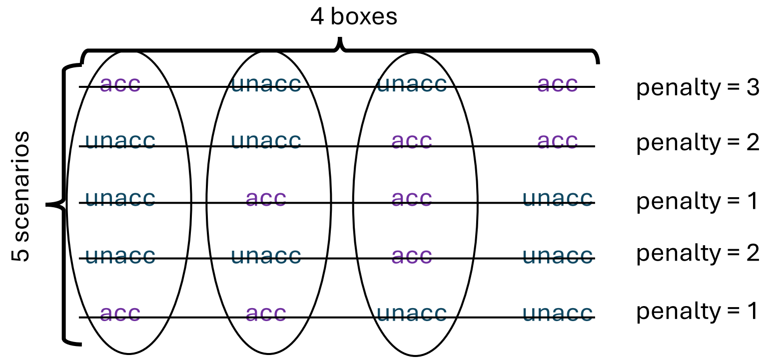

Correlated / dependent input variables

Extension to correlated values

- Optimization problem: select one alternative
- Information acquisition: Pandora's Box
- Minimize value + opening cost
- Joint distribution over box values: $(v_1, \dots, v_n) \sim D$

  m "scenarios" representing the support of D

Example: may need to open "useless" boxes to find "useful" ones

- Let f be a random mapping from $[n]$ to a large domain.
- $S^{(i)} = \begin{cases} v_1 = f(i) \\ v_i = 0 \\ v_{i'} = \infty \quad \text{for } i' \neq 1, i \end{cases}$
- OPT = 2

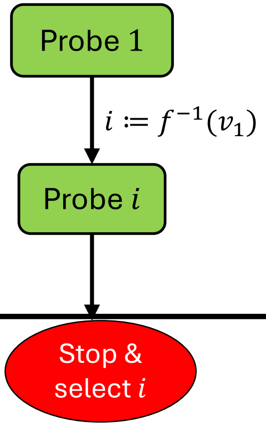

A special case: the Min Sum Set Cover problem

- All costs are 0 or ∞. I.e. boxes are "acceptable" or "unacceptable".
- For simplicity: all probing penalties are 1; Uniform distribution over m scenarios.
- Think of the scenarios as elements and boxes as sets. If an "acceptable" box is opened in a scenario, the scenario gets covered.
- Objective: minimize total penalty a.k.a. covering time.

Another special case: Optimal Decision Tree

- Use "tests" to discover which of m given hypotheses is true
- Objective: minimize probing cost
- m states of the world $\Longrightarrow O(\log m)$-approximation

  [Gupta Nagarajan Ravi'18, Li Liang Mussmann'20]
- Tight in general [Chakravarthy et al.'07]

Technique: Reduction to Optimal Decision Tree

[C. Gergatsouli McMahan Tzamos ’21]

Theorem: $\tilde{O}(\log m)$-approximation where m = number of scenarios.

[C. Gergatsouli McMahan Tzamos’22]

Open: Can we obtain an $O(1)$ approximation?

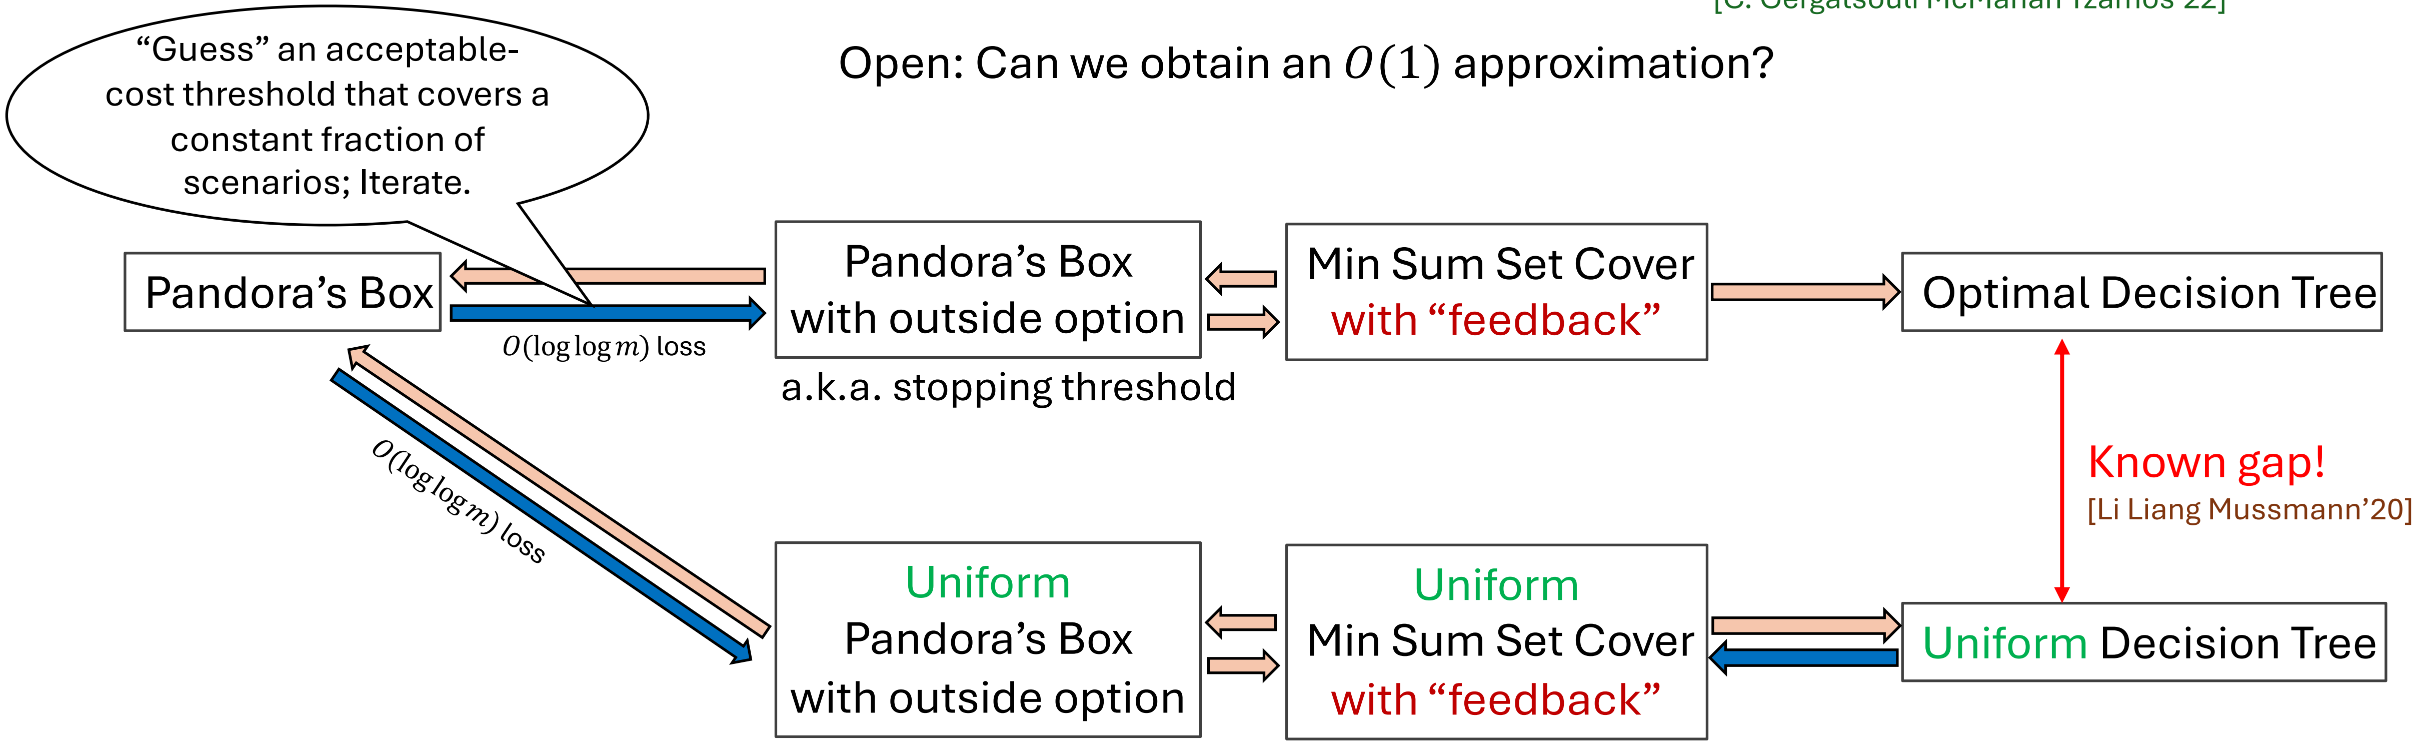

Can we say anything interesting when m is very large?

Extension to large support correlated distributions

Concept: Approximating the Partially Adaptive Optimum

Correlated / dependent input variables

Inspired by Weitzman's solution

Fully Adaptive Opt (FA-OPT)

- Probing order and stopping rule depend adaptively on obtained info.

Hard to approximate better than $\Omega(\log m)$

Partially Adaptive Opt. (PA-OPT)

- Probing order is non-adaptive
- Stopping rule depends adaptively on obtained info.

Can approximate and learn from data!
[C. Gergatsouli Teng Tzamos Zhang '20]

Non Adaptive Optimum (NA-OPT)

- Selects a subset of boxes
- Opens all simultaneously
- Chooses the best box

Hard to approximate better than $\Omega(\log m)$
(captures set cover)

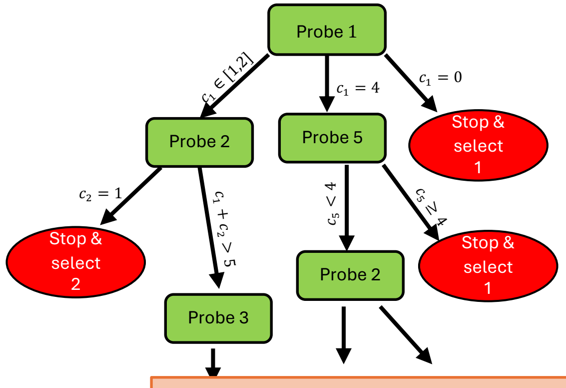

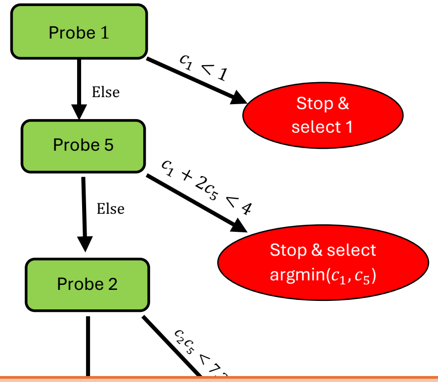

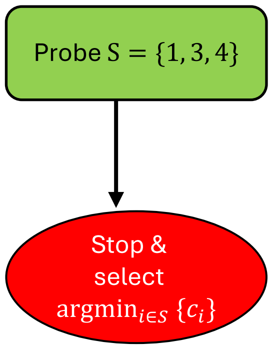

Theorem: Weitzman's algorithm, suitably adapted, is a 4.4-approximation to the PA-OPT.

[Gergatsouli Tzamos'24]

General feasibility constraints

✓ Any constraints that admit good frugal algorithms

× Maximization?
Hard to approximate NA-OPT within any constant factor via fully adaptive algos

Correlated / dependent input variables

✓ Log approx. for small support
✓ Const approx. to PA-OPT

Simplest setting: Weitzman's Pandora's Box

- Single-element selection
- Independent alternatives
- Binary information state: known or unknown
- Linear objective: value of selected item minus inspection cost

✓ Multi-modal info. acquisition via ex ante relaxation or local approx
✓ Any Markov chain type process of info. acquisition

General, possibly nonlinear objectives

Preliminary investigations by
[Olszewski Weber'15, Berger Ezra Feldman Fusco'23]

General process of information acquisition

Some open directions

- Minimization in the MDP setting
  - Few and weak results
- Maximization in the correlated setting
  - Can we make suitable assumptions to simplify the problem?
- Short range correlations?
- General objectives
  - Going beyond frugality?
  - Non-linear objectives, e.g. Submodular over selected variables? Two-dimensional coverage?

Thanks for your attention!

Questions?