
COSTLY EXPLORATION: GRAPH CONNECTIVITY WITH NOISY QUERIES

Evangelia Gergatsouli


University of Wisconsin-Madison

Nov 15 2022


Based on joint work with: Dimitris Fotakis¹, Charilaos Pipis¹, Miltos Stouras²,

Christos Tzamos³

¹NTUA, ²EPFL, ³UW Madison




Model



Moldgraph G

Adversary




Realized graph



Oracle


- ▶ Initial graph G
- ▶ Adversary deletes edges
- ▶ *Noisy Oracle* answers questions of the form “Does edge e exist”?

Model



Moldgraph G

Adversary



Realized graph


Does e_1 exist?



Oracle


- ▶ Initial graph G
- ▶ Adversary deletes edges
- ▶ *Noisy Oracle* answers questions of the form “Does edge e exist”?

Model



Moldgraph G

Adversary



Realized graph

Does e_1 exist?


Yes!



Oracle


- ▶ Initial graph G
- ▶ Adversary deletes edges
- ▶ *Noisy Oracle* answers questions of the form “Does edge e exist”?

Model



Moldgraph G

Adversary



Realized graph

Does e_1 exist?

Yes!



Oracle

Goal: find a spanning tree with minimal number of queries to \mathcal{O}

Error Regimes

2-sided error

Oracle's answer

Truth

Yes!

No!

Yes!

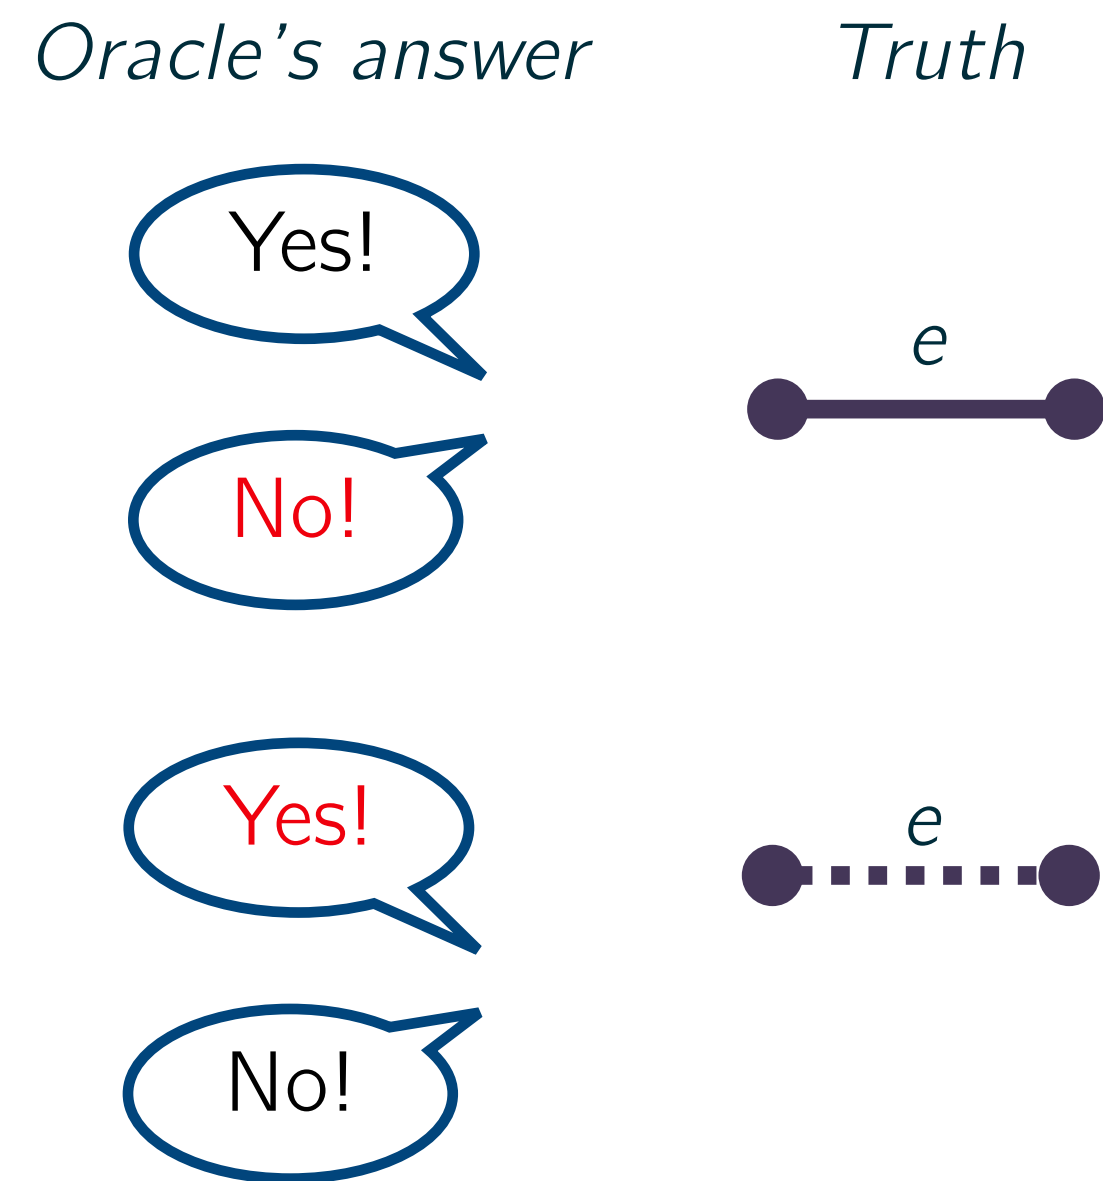
No!



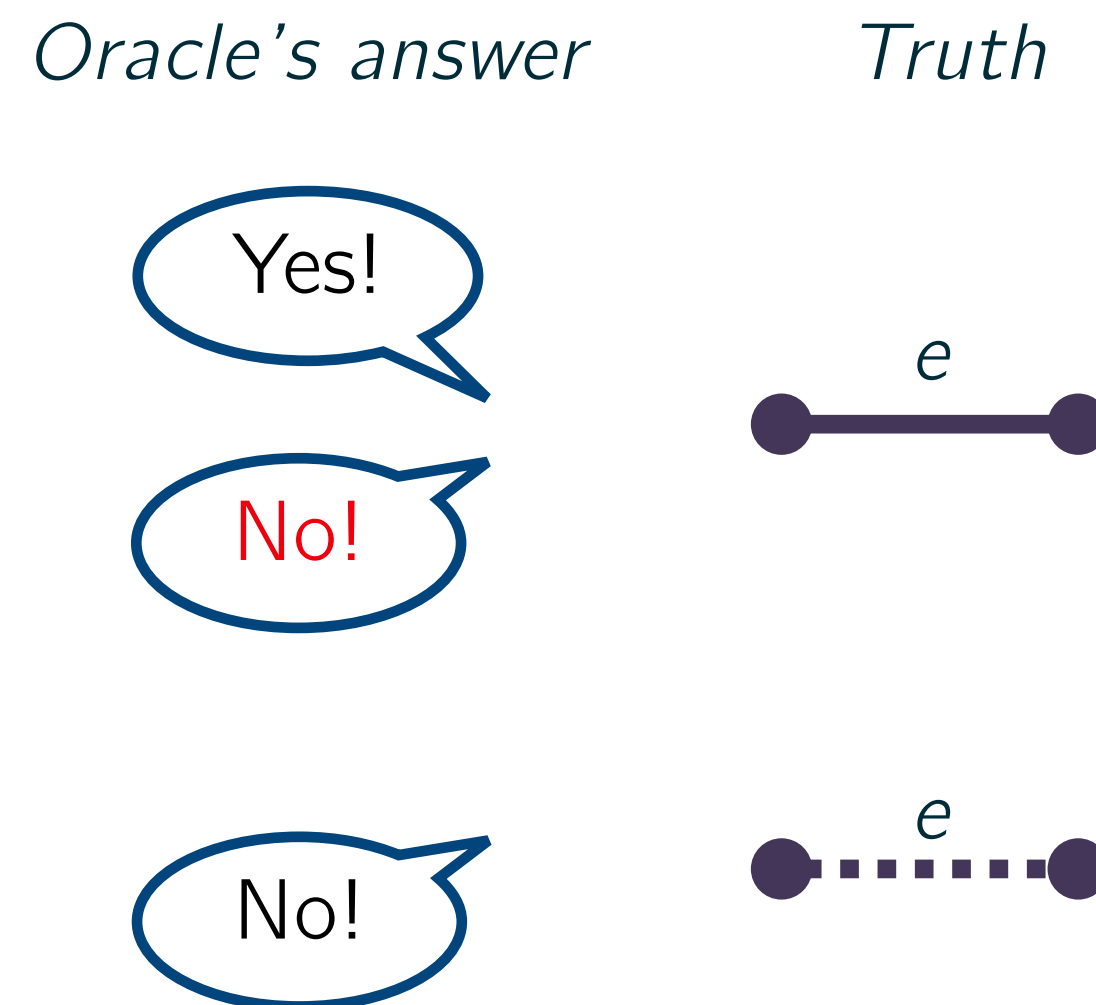
Error Regimes

1-sided errors


2-sided error



False Negatives (FN)




False Positives (FP)




Error Regimes

1-sided errors


2-sided error



False Negatives (FN)



False Positives (FP)



Probability of **error** is $p < 1/2$

Previous Work

Assumes independence on
edges' existence

- Clustering with Noisy queries
[Mazumdar, Saha '17]
- Connectivity in Erdos Renyi graphs
[Fu et al. '17, Fu et al. '14]
- Computing with noisy information
[Feige et al. '94]
- MST verification, under uncertainty
[Hoffman et al. '08, Erlebach et al. '14]

Only handle consistent
answers to queries

Roadmap

| | General Graphs |
|--|-----------------------|
| 2-sided error | $\Theta(m \log m)$ |
| 1-sided error (False Positives) | $O(m \log m)$ |
| 1-sided error (False Negatives) | $\Theta(m \log m)$ |

Roadmap

| | General Graphs | |
|------------------------------------|--------------------|---|
| 2-sided error | $\Theta(m \log m)$ | <i>Lower bound even for planar graphs!</i> |
| 1-sided error (False Positives) | $O(m \log m)$ | <i>$O(m)$ if realized graph is acyclic</i> |
| 1-sided error (False Negatives) | $\Theta(m \log m)$ | <i>$O(\rho m)$ if G is ρ-sparse</i> |

Roadmap

| | General Graphs | |
|------------------------------------|--------------------|---|
| 2-sided error | $\Theta(m \log m)$ | <i>Lower bound even for planar graphs!</i> |
| 1-sided error (False Positives) | $O(m \log m)$ | <i>$O(m)$ if realized graph is acyclic</i> |
| 1-sided error (False Negatives) | $\Theta(m \log m)$ | <i>$O(\rho m)$ if G is ρ-sparse</i> |

*Special cases also **tight***

2-sided Error - Algorithm

Algorithm:

- ▶ Query every edge of moldgraph $\log m$ times
- ▶ Every $e \in E$ with more “Yes” is treated as realized
- ▶ **If** graph connected
 - ▶ Output graph
- ▶ **Else**
 - ▶ Output any spanning tree of moldgraph


2-sided Error - Algorithm

Algorithm:

- ▶ Query every edge of moldgraph $\log m$ times
- ▶ Every $e \in E$ with more “Yes” is treated as realized
- ▶ **If** graph connected
 - ▶ Output graph
- ▶ **Else**
 - ▶ Output any spanning tree of moldgraph

Algorithm finds a spanning tree **w.h.p.**
using $O(m \log m)$ queries

2-sided Error - Lower Bound




Moldgraph G

Observations:

- ▶ Each pair is independent of others (each pair is a **cut**)
- ▶ If less than $\log m$ queries per pair \rightarrow **error w.h.p.**

2-sided Error - Lower Bound



Moldgraph G

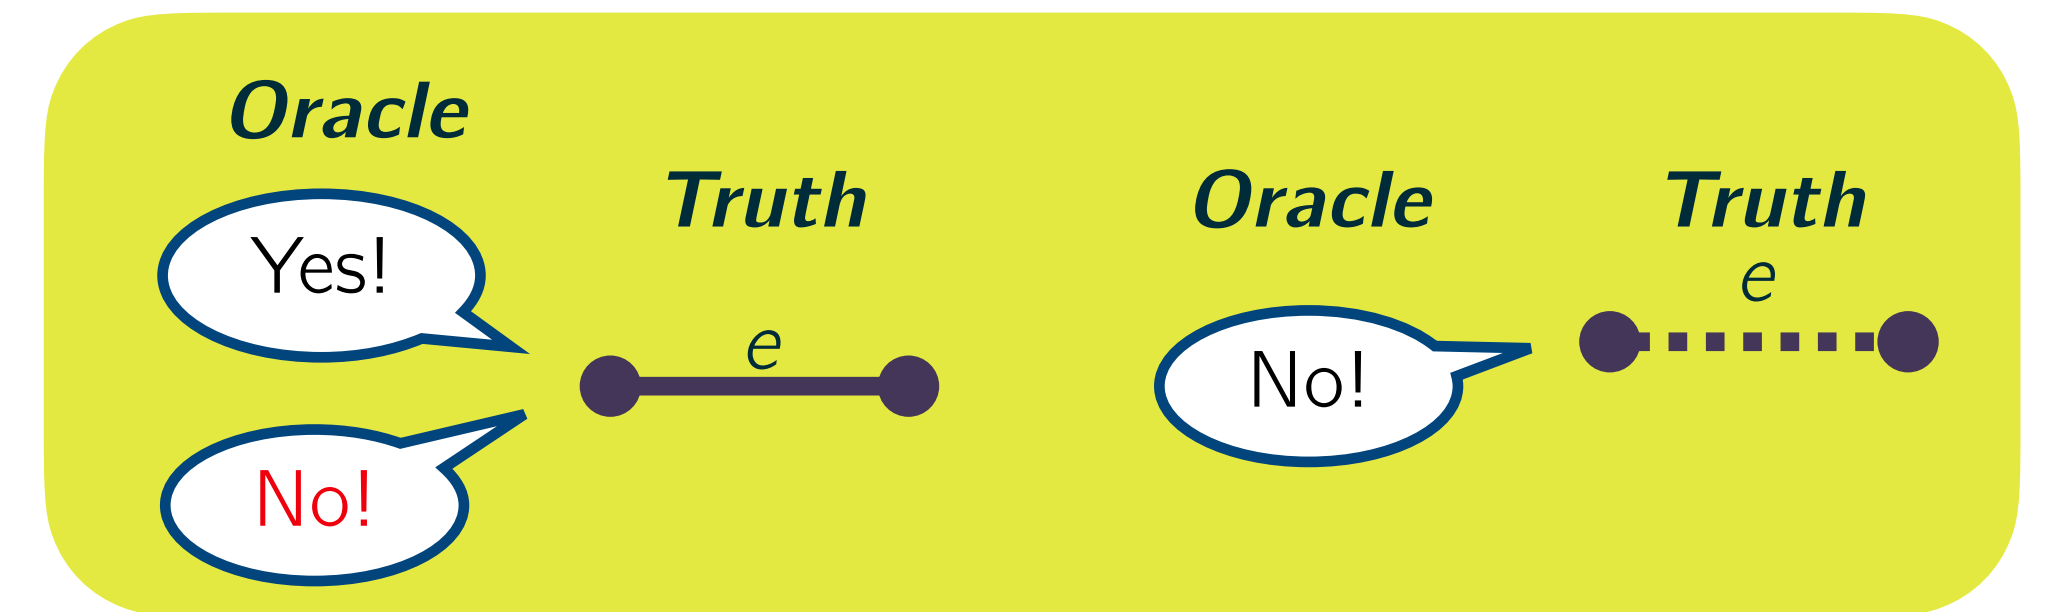
Observations:

- ▶ Each pair is independent of others (each pair is a **cut**)
- ▶ If less than $\log m$ queries per pair \rightarrow **error w.h.p.**

No algorithm can do better than
 $\Omega(m \log m)$ queries

1-sided, False Negatives

Recall:




Observations:

- ▶ Every realized edge, will eventually give a "Yes" answer
- ▶ Every cut of G contains at least one realized edge
- ▶ Cuts are small in sparse graphs

ρ -sparse:
edges $\leq \rho \cdot$ nodes

1-sided, False Negatives

Recall:



Observations:

- ▶ Every realized edge, will eventually give a "Yes" answer
- ▶ Every cut of G contains at least one realized edge
- ▶ Cuts are small in sparse graphs

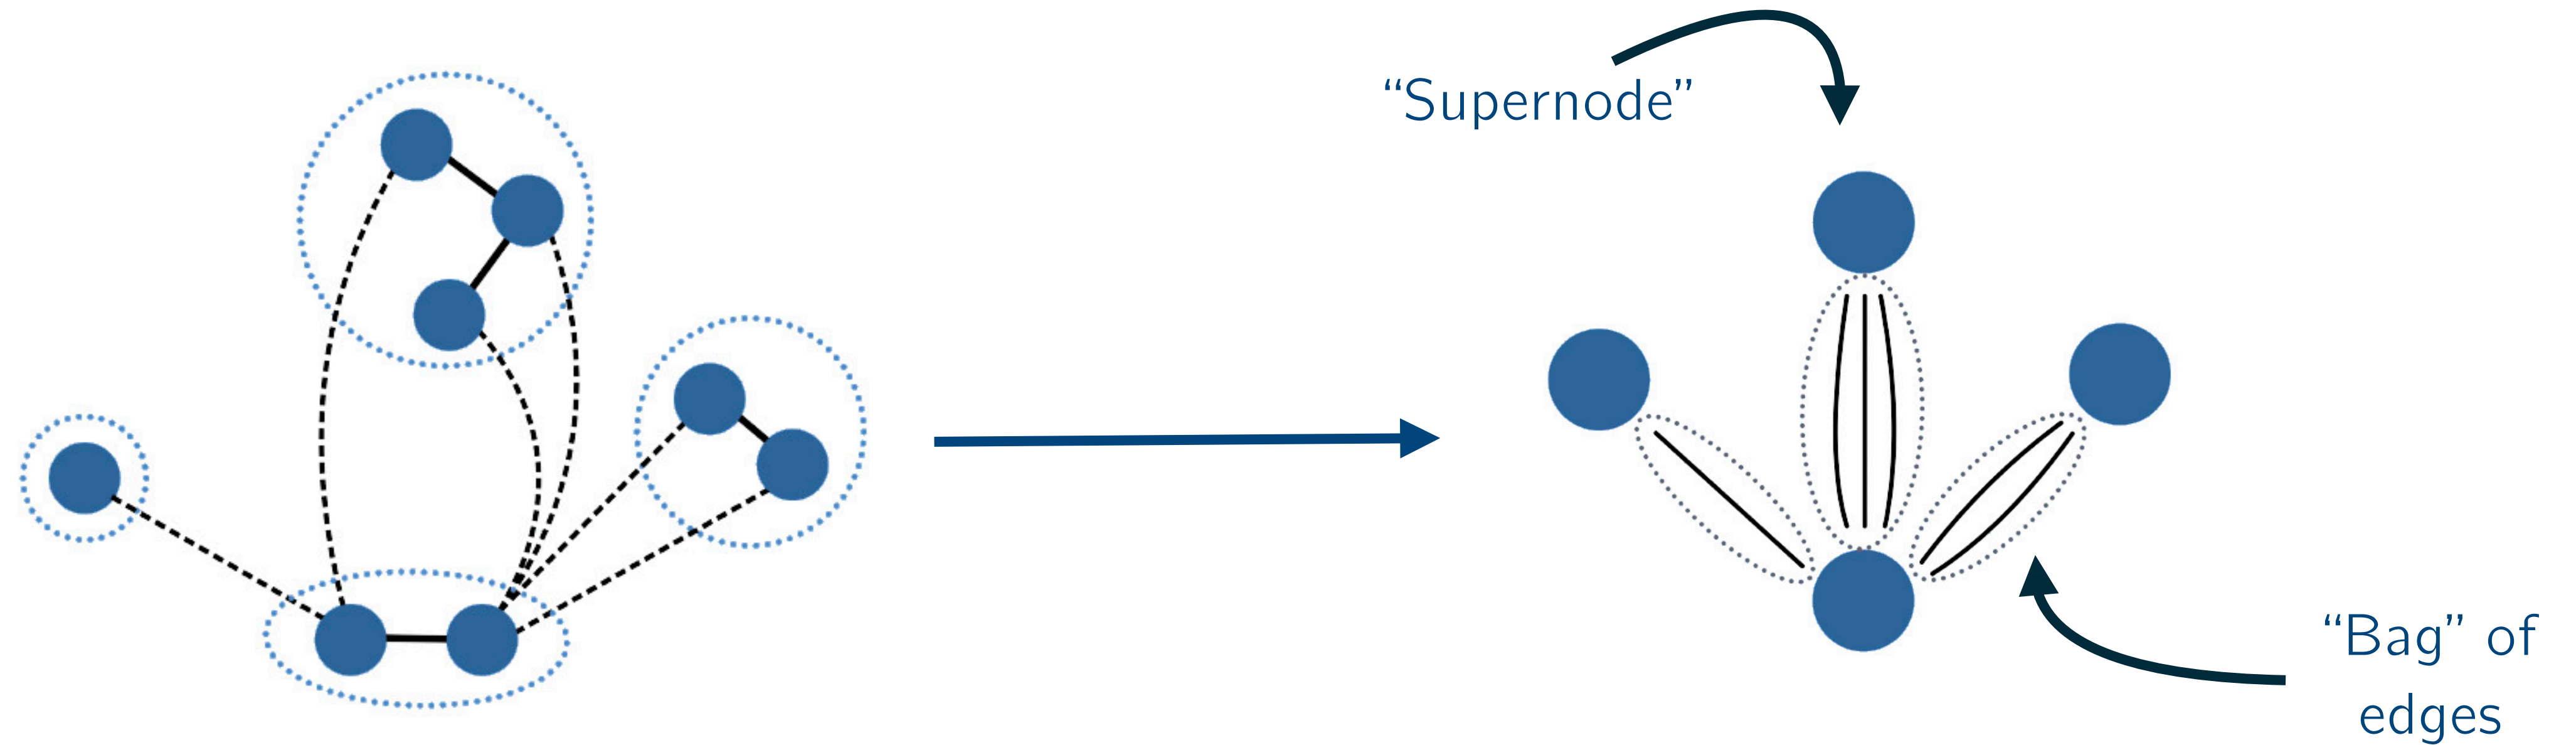
Idea 1: can naively ask about every edge

Idea 2: only need to find one edge per cut

ρ -sparse:
edges $\leq \rho \cdot$ nodes

1-sided, False Negatives

*Idea 2: only need
to find **one edge**
per cut*



1-sided, False Negatives

Naive algorithm:

- ▶ Round-robin query all edges until there is spanning tree

+

Sparse graphs algorithm:

- ▶ **While** (graph has more than 1 nodes)
 - ▶ Find min degree node $u \in V$
 - ▶ **For** every “Bag of edges” neighboring u
 - ▶ Select random edge to query
 - ▶ Find realized edge e and contract it

1-sided, False Negatives

Naive algorithm:

- ▶ Round-robin query all edges until there is spanning tree

+


Sparse graphs algorithm:

- ▶ **While** (graph has more than 1 nodes)
 - ▶ Find min degree node $u \in V$
 - ▶ **For** every “Bag of edges” neighboring u
 - ▶ Select random edge to query
 - ▶ Find realized edge e and contract it

Algorithm finds a sp. tree using $O(m \log m)$ queries.
If graph is ρ -sparse then uses $O(\rho m)$ queries

1-sided, False Positives

Recall:

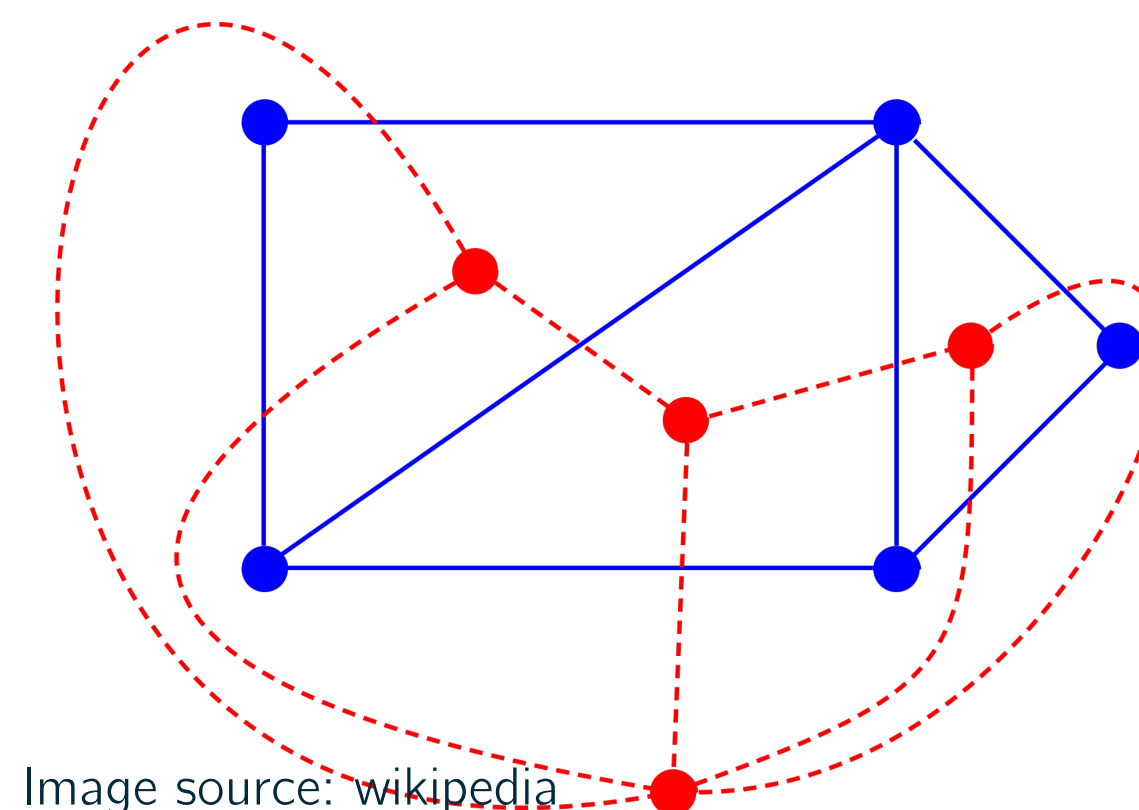


Observations:

- ▶ Every non-realized edge, will eventually give a “No” answer \longrightarrow **Idea 1:** can naively ask about every edge

Dual graph:

- ▶ Node per face of G
- ▶ Edge if faces in G are separated by edge



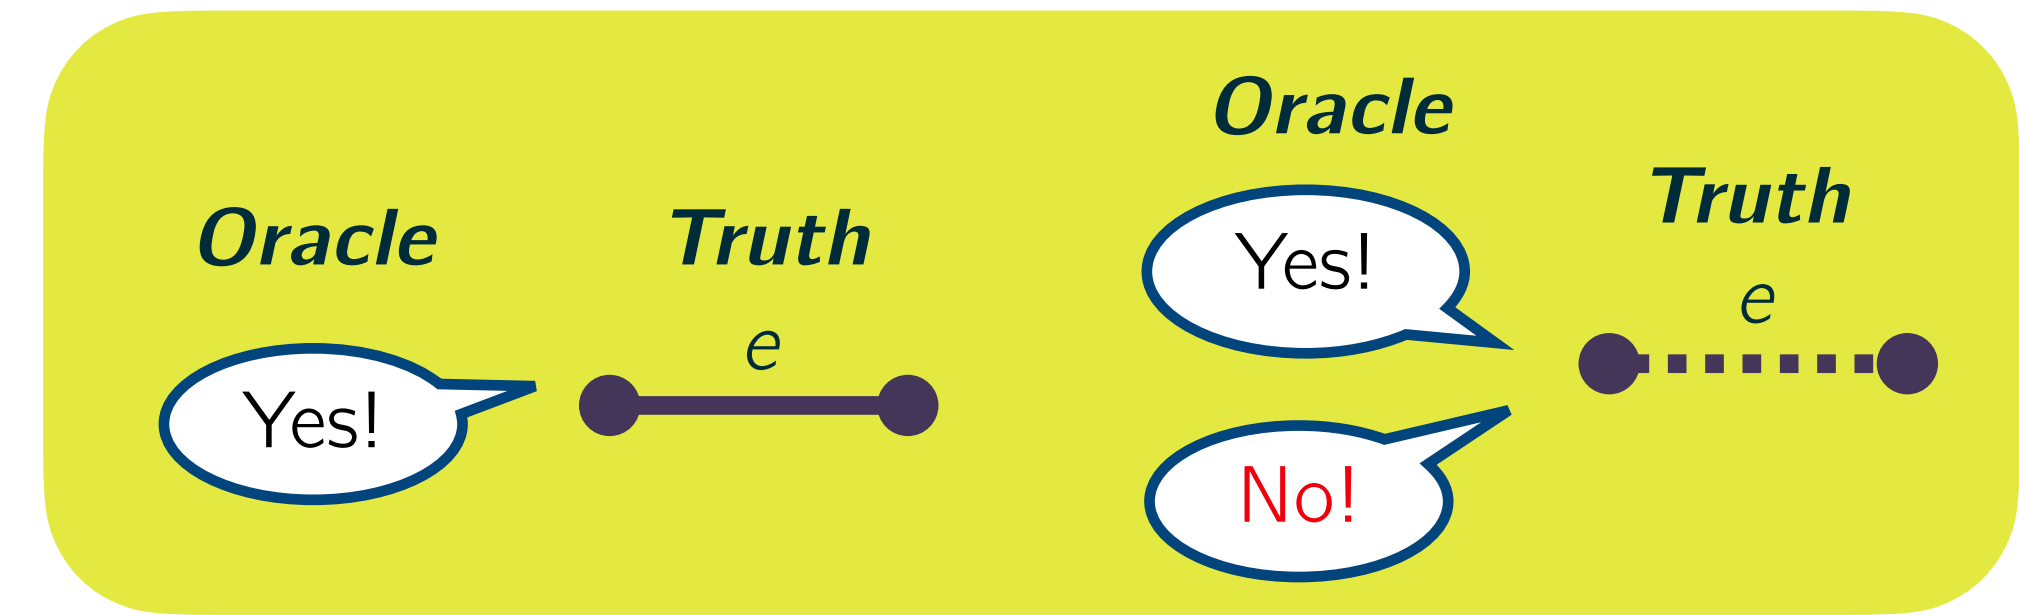
Original G

Dual G'

Image source: wikipedia

1-sided, False Positives

Recall:



Observations:

- ▶ Every non-realized edge, will eventually give a “No” answer → *Idea 1: can naively ask about every edge*
- ▶ **If** realized graph is **tree** & G is **planar**
- ▶ Every **cycle** of G contains at **least one non-realized edge** → *Idea 2: can use FN algorithm in dual graph*
- ▶ **Cycles** in initial graph \equiv **Cuts** in the dual graph

1-sided, False Positives

Naive algorithm:

▶ **While** there's no tree:

- Round robin {
- ▶ Ask about **next** edge e
 - ▶ **If** answer is "No" mark as non-realized
 - ▶ **Else** mark as realized

+

Acyclic graph algorithm:

- ▶ Construct the dual G' of initial graph
- ▶ Run FN algorithm on dual G'
- ▶ Return complement of edges

1-sided, False Positives

Naive algorithm:

- ▶ **While** there's no tree:
- ▶ Ask about **next** edge e
- ▶ **If** answer is "No" mark as non-realized
- ▶ **Else** mark as realized

+

Acyclic graph algorithm:

- ▶ Construct the dual G' of initial graph
- ▶ Run FN algorithm on dual G'
- ▶ Return complement of edges

Algorithm finds a sp. tree w.h.p. with $O(m \log m)$ queries.
If G planar and realized graph is tree then $O(m)$ queries


Summary

| | General Graphs |
|--|-----------------------|
| 2-sided error | $\Theta(m \log m)$ |
| 1-sided error (False Positives) | $O(m \log m)$ |
| 1-sided error (False Negatives) | $\Theta(m \log m)$ |

Summary

| | General Graphs | |
|------------------------------------|--------------------|---|
| 2-sided error | $\Theta(m \log m)$ | <i>Lower bound even for planar graphs!</i> |
| 1-sided error (False Positives) | $O(m \log m)$ | <i>$O(m)$ if realized graph is acyclic</i> |
| 1-sided error (False Negatives) | $\Theta(m \log m)$ | <i>$O(\rho m)$ if G is ρ-sparse</i> |



A More General Problem



- ▶ Find MST on weighted graph
- ▶ Pay to learn weights (explore alternatives)
- ▶ Stop anytime and take best so far

A More General Problem


Find the best alternative, with costly information!



- ▶ Pay to learn weights (explore alternatives)
- ▶ **Constraint:** edges should be a tree

A More General Problem

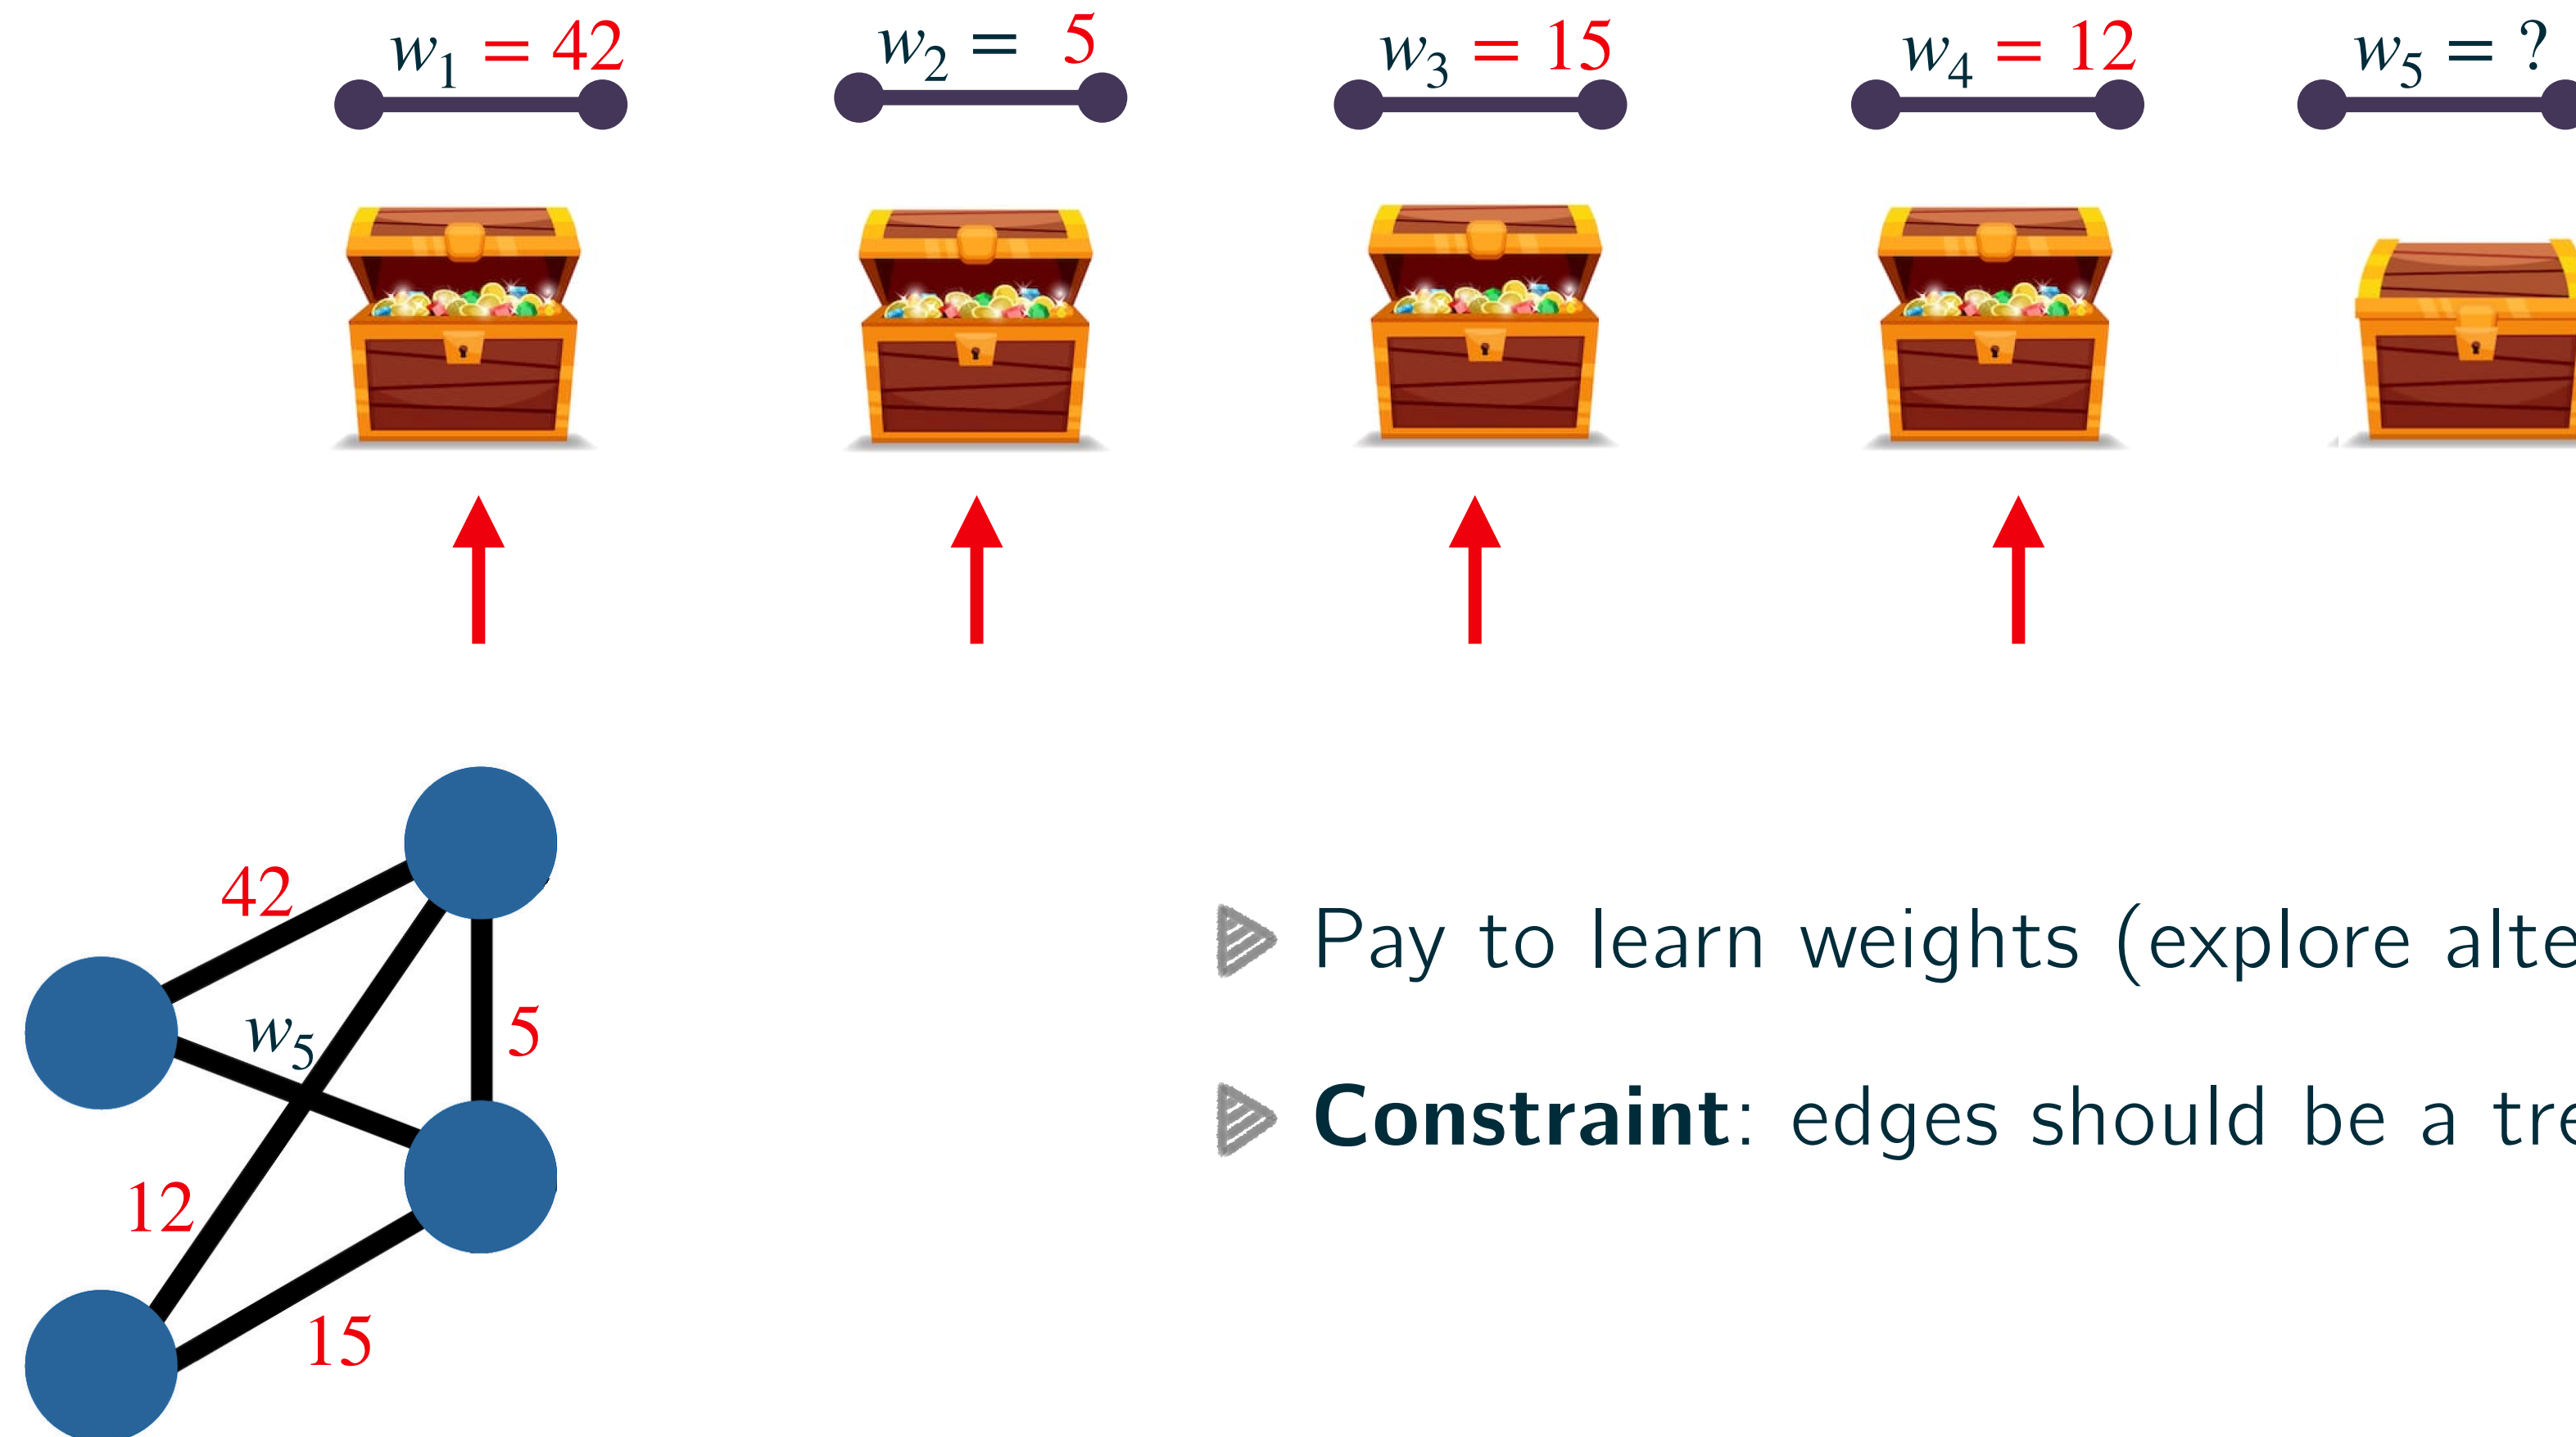
Find the best alternative, with costly information!



- ▶ Pay to learn weights (explore alternatives)
- ▶ **Constraint:** edges should be a tree

A More General Problem

Find the best alternative, with costly information!



- ▶ Pay to learn weights (explore alternatives)
- ▶ **Constraint:** edges should be a tree

Pandora's Box

Find the best alternative, with costly information!



- ▶ Information is not free
 - ▶ Explore alternatives (open boxes)
 - ▶ Stop anytime and take best so far
- } **Strategy**

Pandora's Box

Find the best alternative, with costly information!



\$42



??



\$15



\$12



??

Opening cost: 3
Final option: **box 4**
Total cost: 12+3

- ▶ Information is not free
 - ▶ Explore alternatives (open boxes)
 - ▶ Stop anytime and take best so far
- } **Strategy**

Pandora's Box

Find the best alternative, with costly information!



\$42



??



\$15



\$12



??


Opening cost: 3
Final option: **box 4**
Total cost: 12+3

- ▶ Information is not free
 - ▶ Explore alternatives (open boxes)
 - ▶ Stop anytime and take best so far
- } **Strategy**

Goal:
Find minimum cost strategy

Pandora's Box

Find the best alternative, with costly information!



Opening cost: 3
Final option: box 4
Total cost: 12+3

- ▶ Information is not free
 - ▶ Explore alternatives (open boxes)
 - ▶ Stop anytime and take best so far
- } **Strategy**

Goal:
Find minimum cost strategy

Previous work

Weitzman's algorithm gives the optimal! [Weitz 1979]

Algorithm:

- ▶ Assign an *expected gain* index¹ to every box
- ▶ Search boxes in order of index until:
 - ▶ Current price better than index of next box

¹Gittins index/reservation value

Previous work

Weitzman's algorithm gives the optimal! [Weitz 1979]

Algorithm:

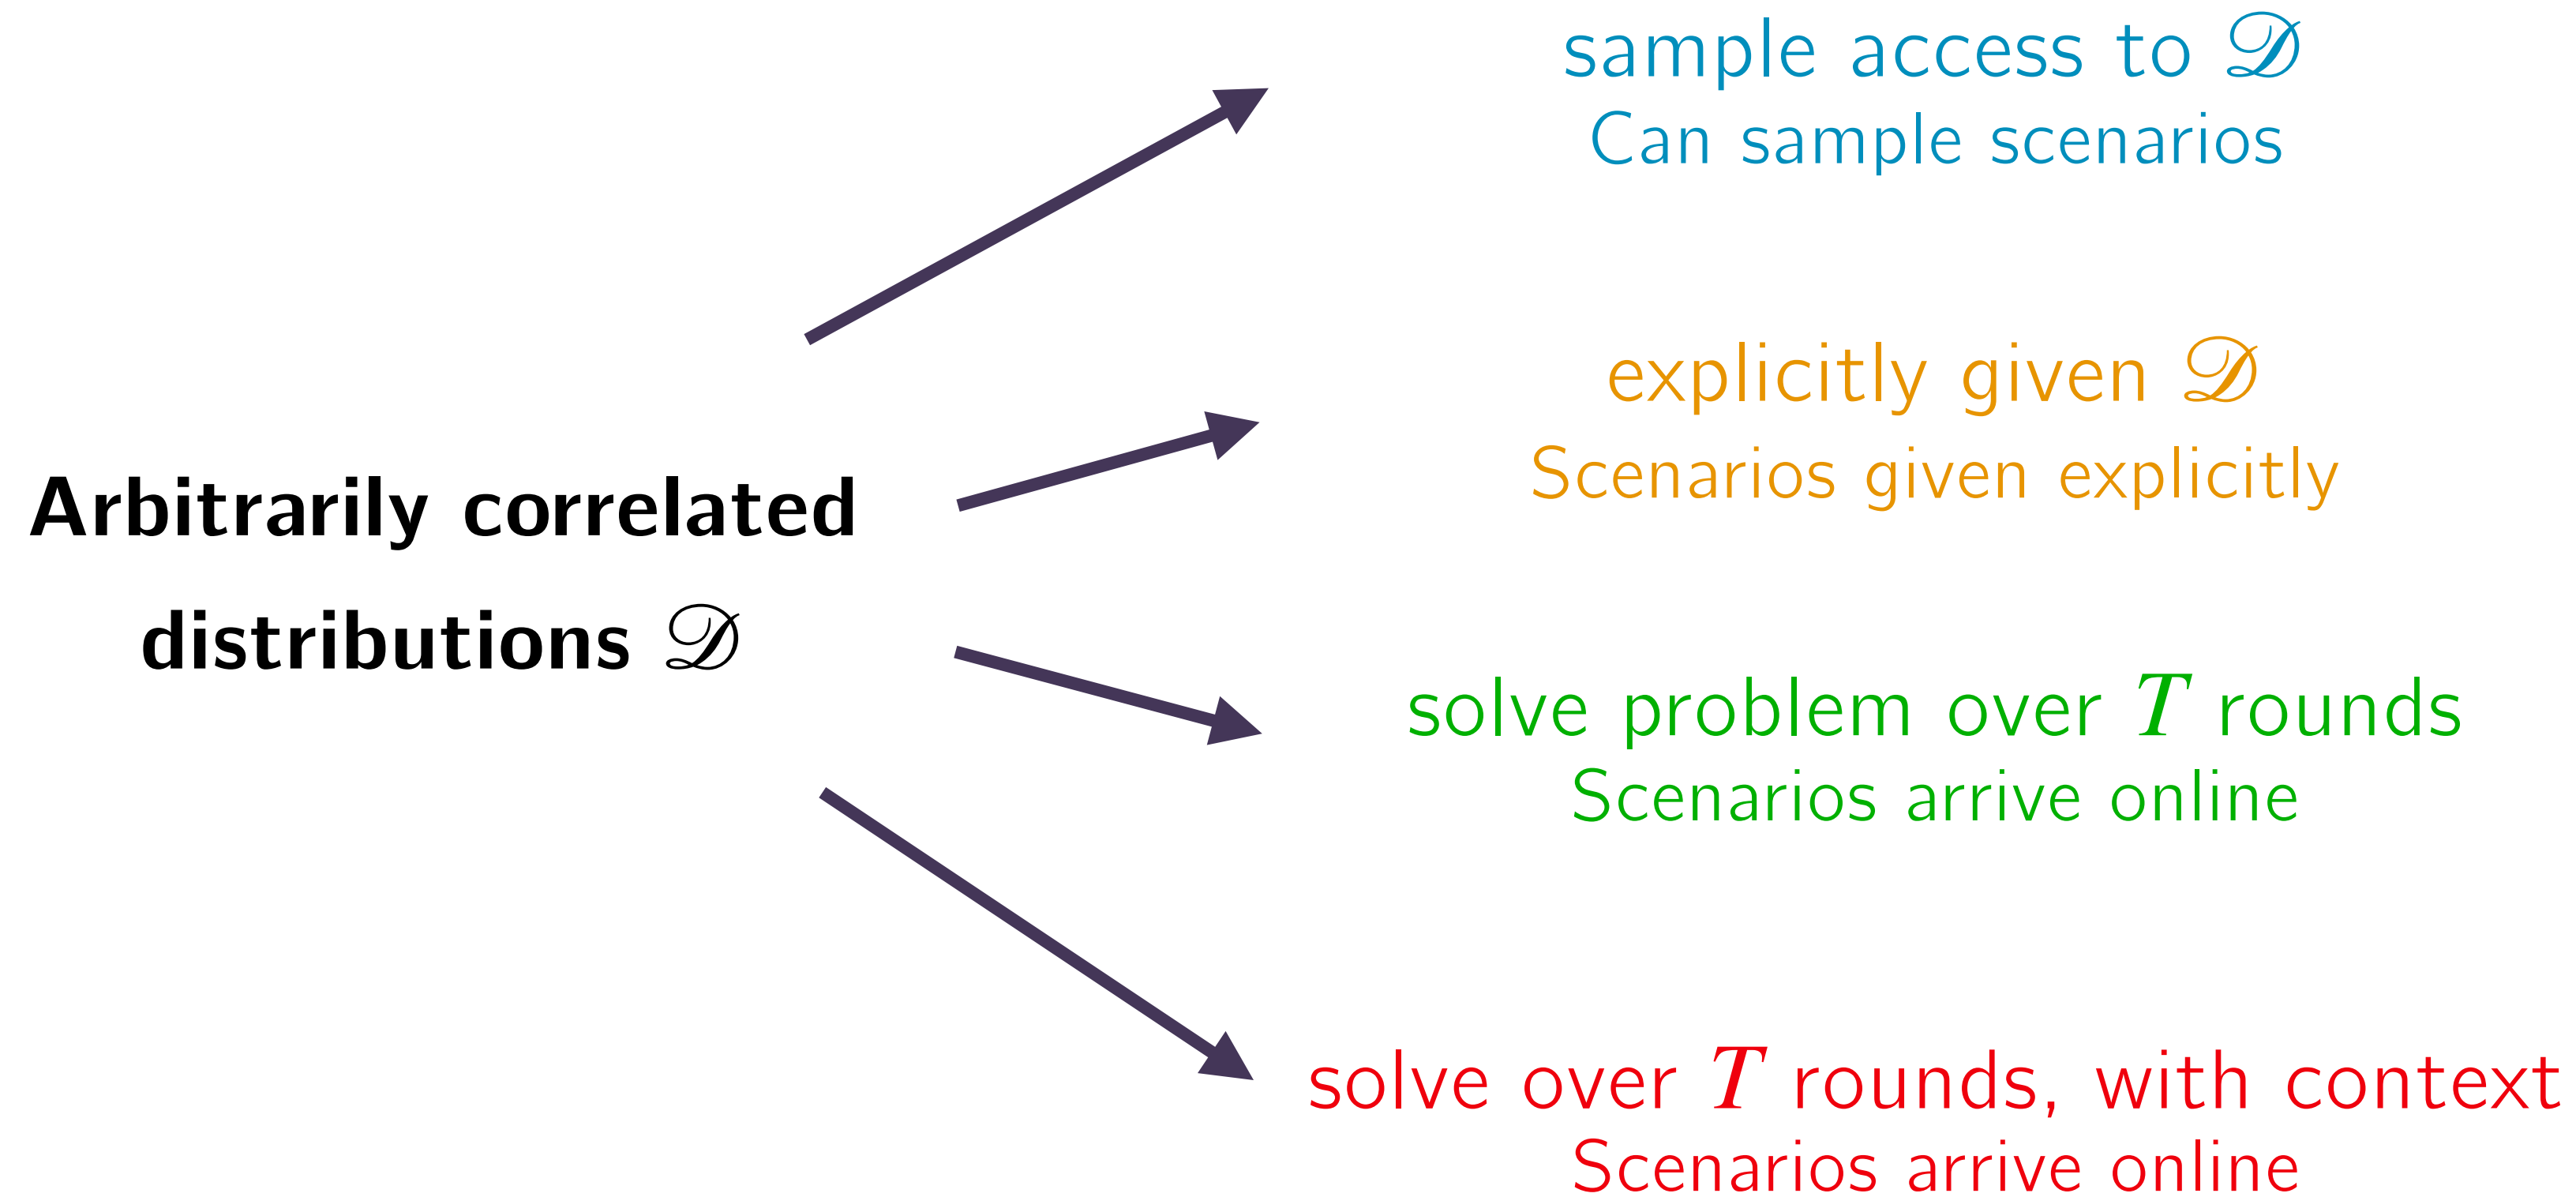
- ▶ Assign an *expected gain* index¹ to every box
- ▶ Search boxes in order of index until:
 - ▶ Current price better than index of next box

Crucial assumption: distributions are **independent!**

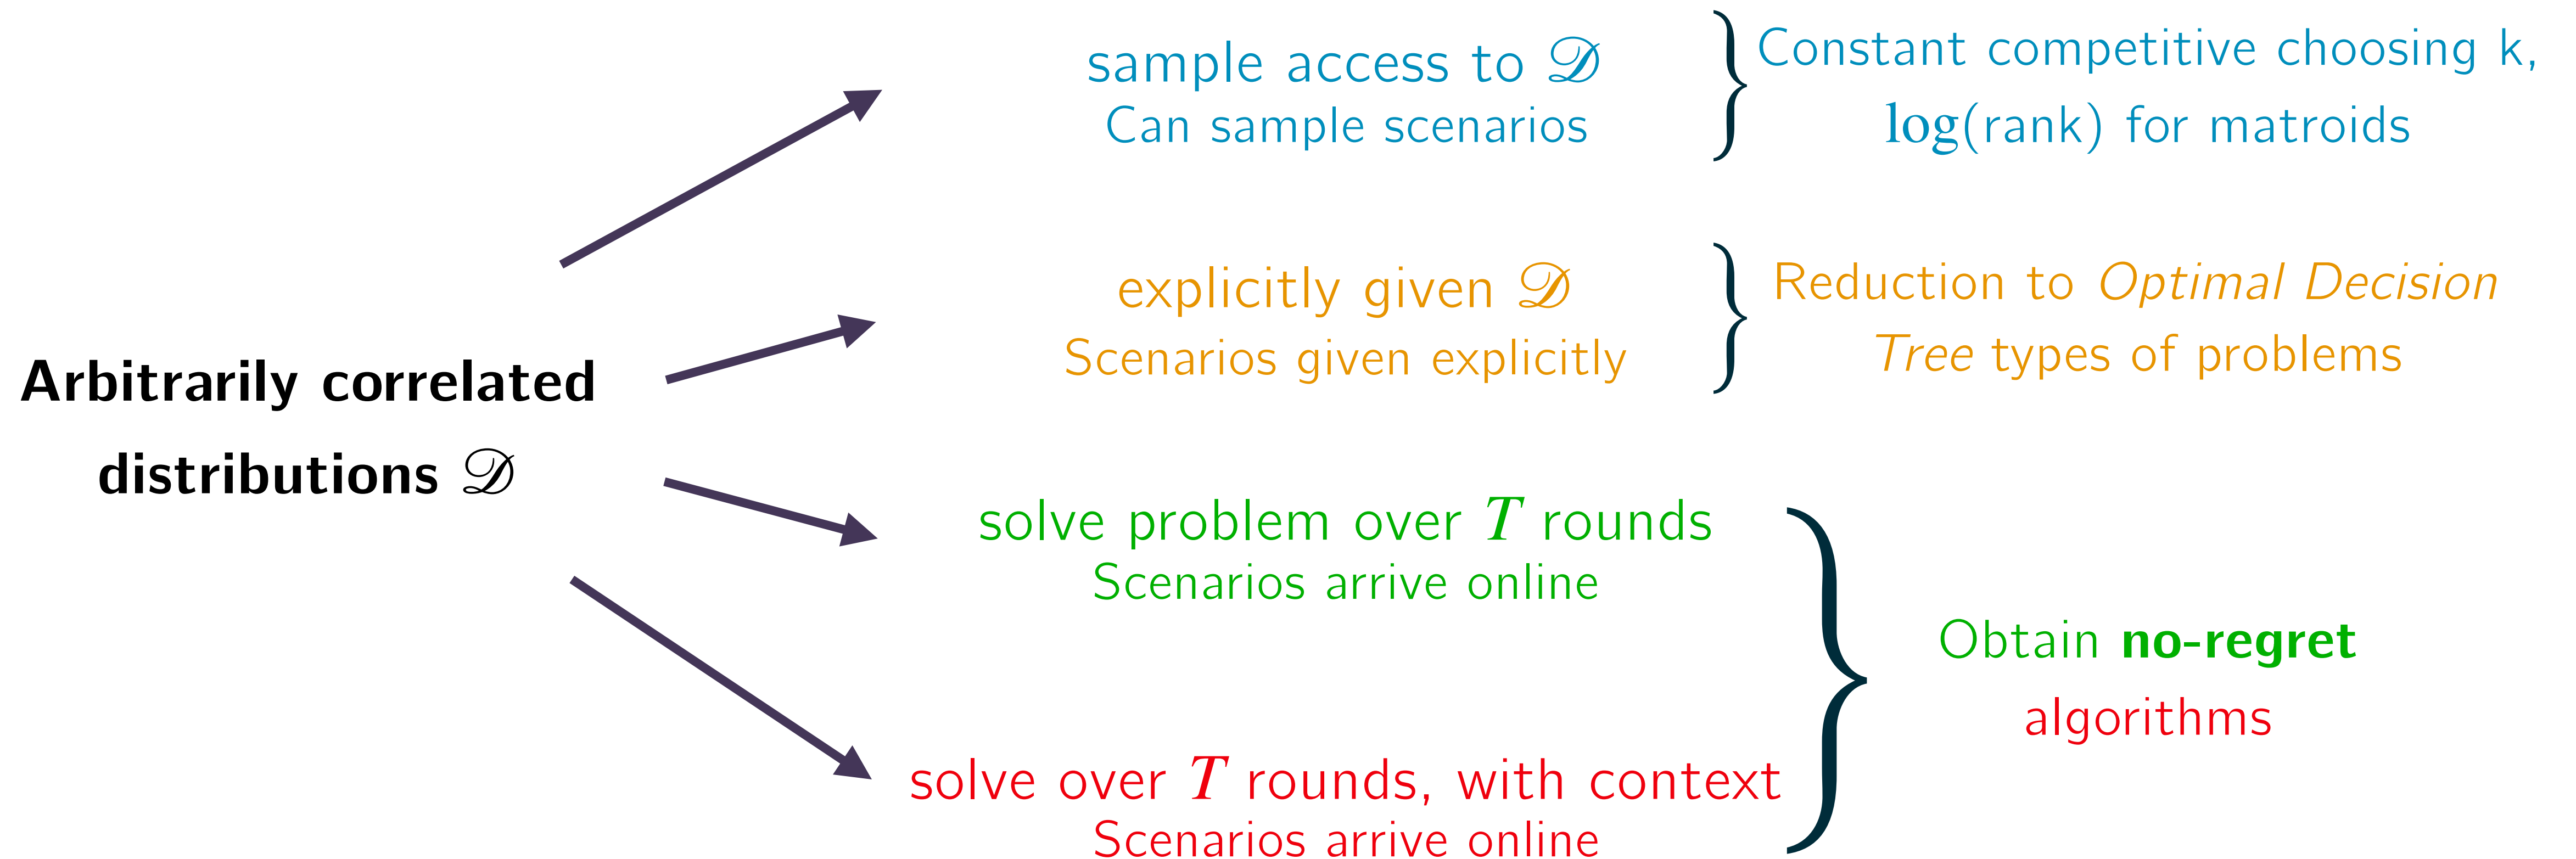
What about **correlation?**

¹Gittins index/reservation value

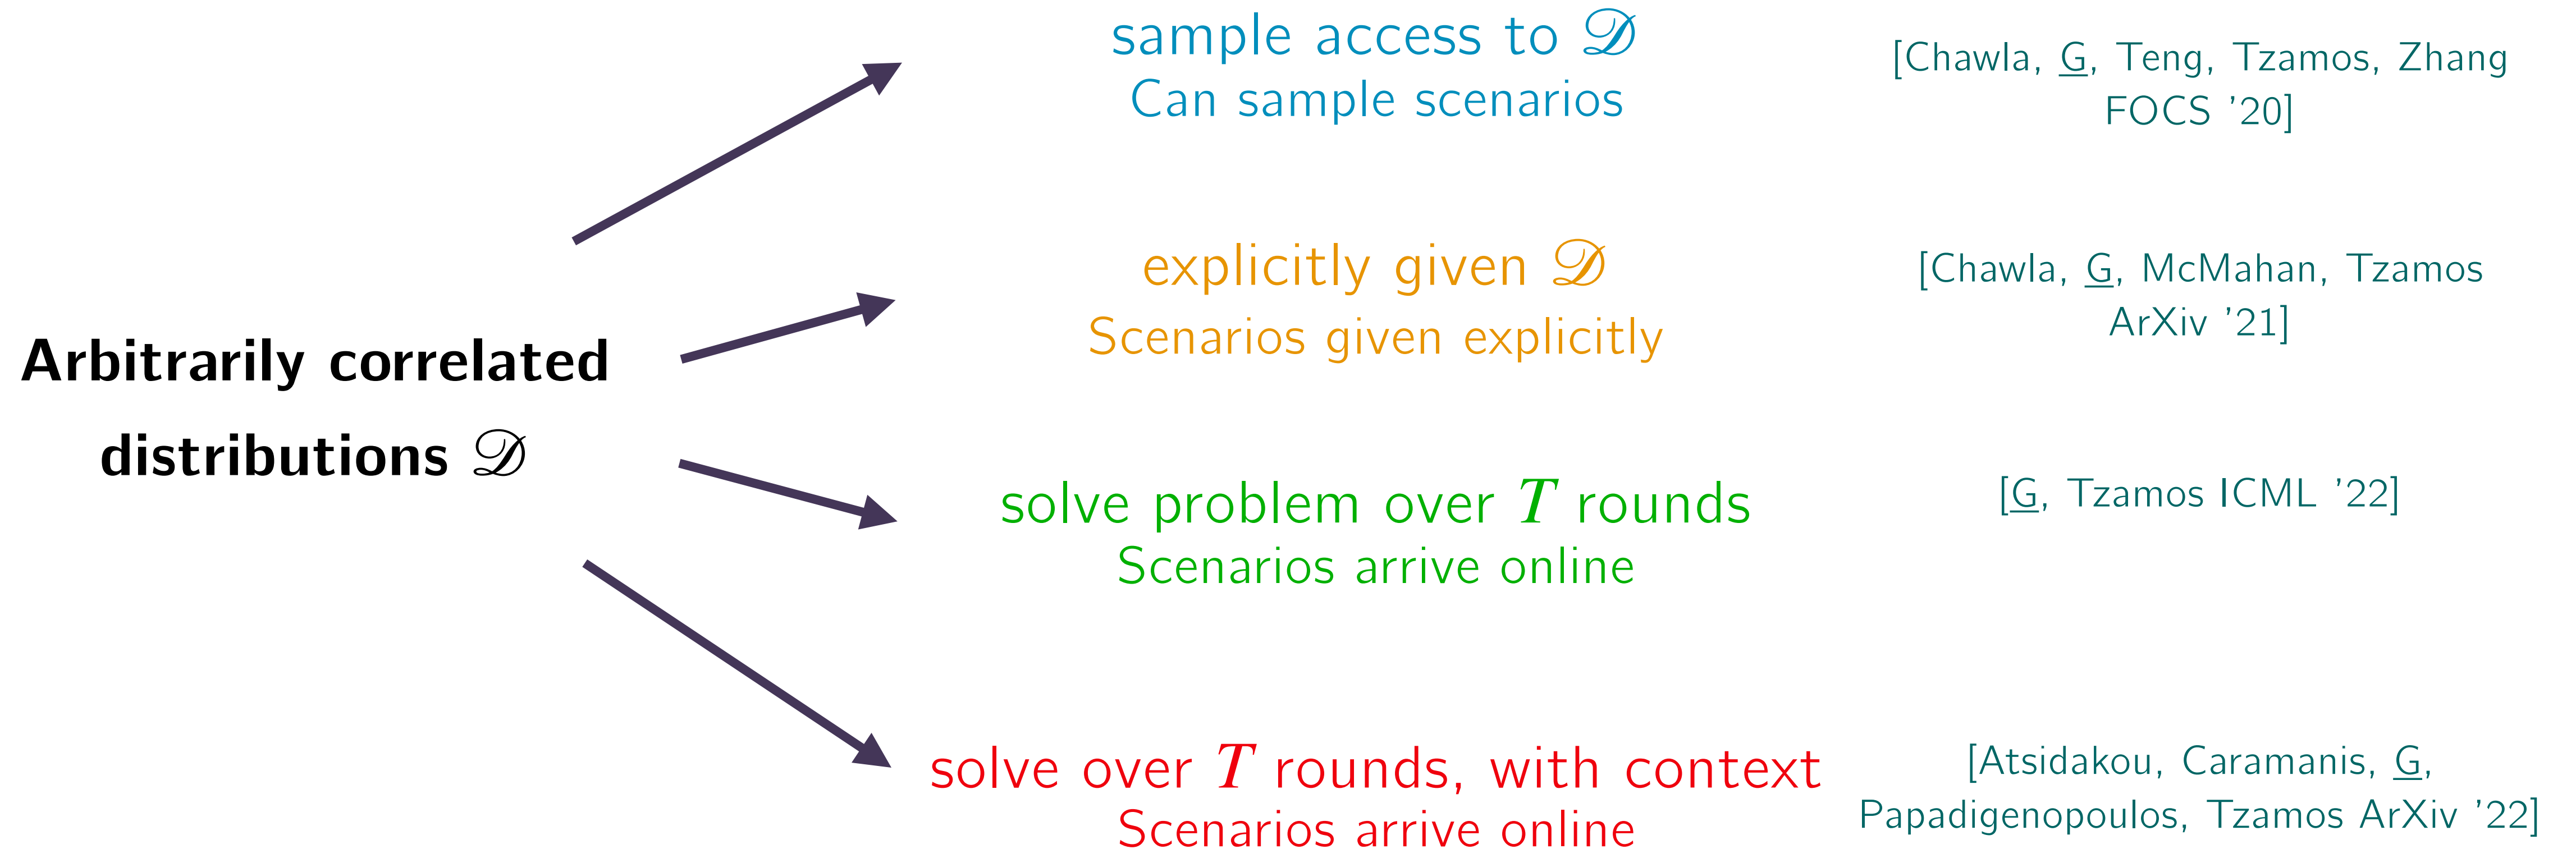
Our work



Our work



Our work



Conclusion

Thank you!

