

Fair Information Sharing for Treasure Hunting



Harvard EconCS, Feb 2015

Yiling Chen



Kobbi Nissim



Bo Waggoner







pirates searching for treasure....




each has some
prior knowledge


and wants the treasure
for herself





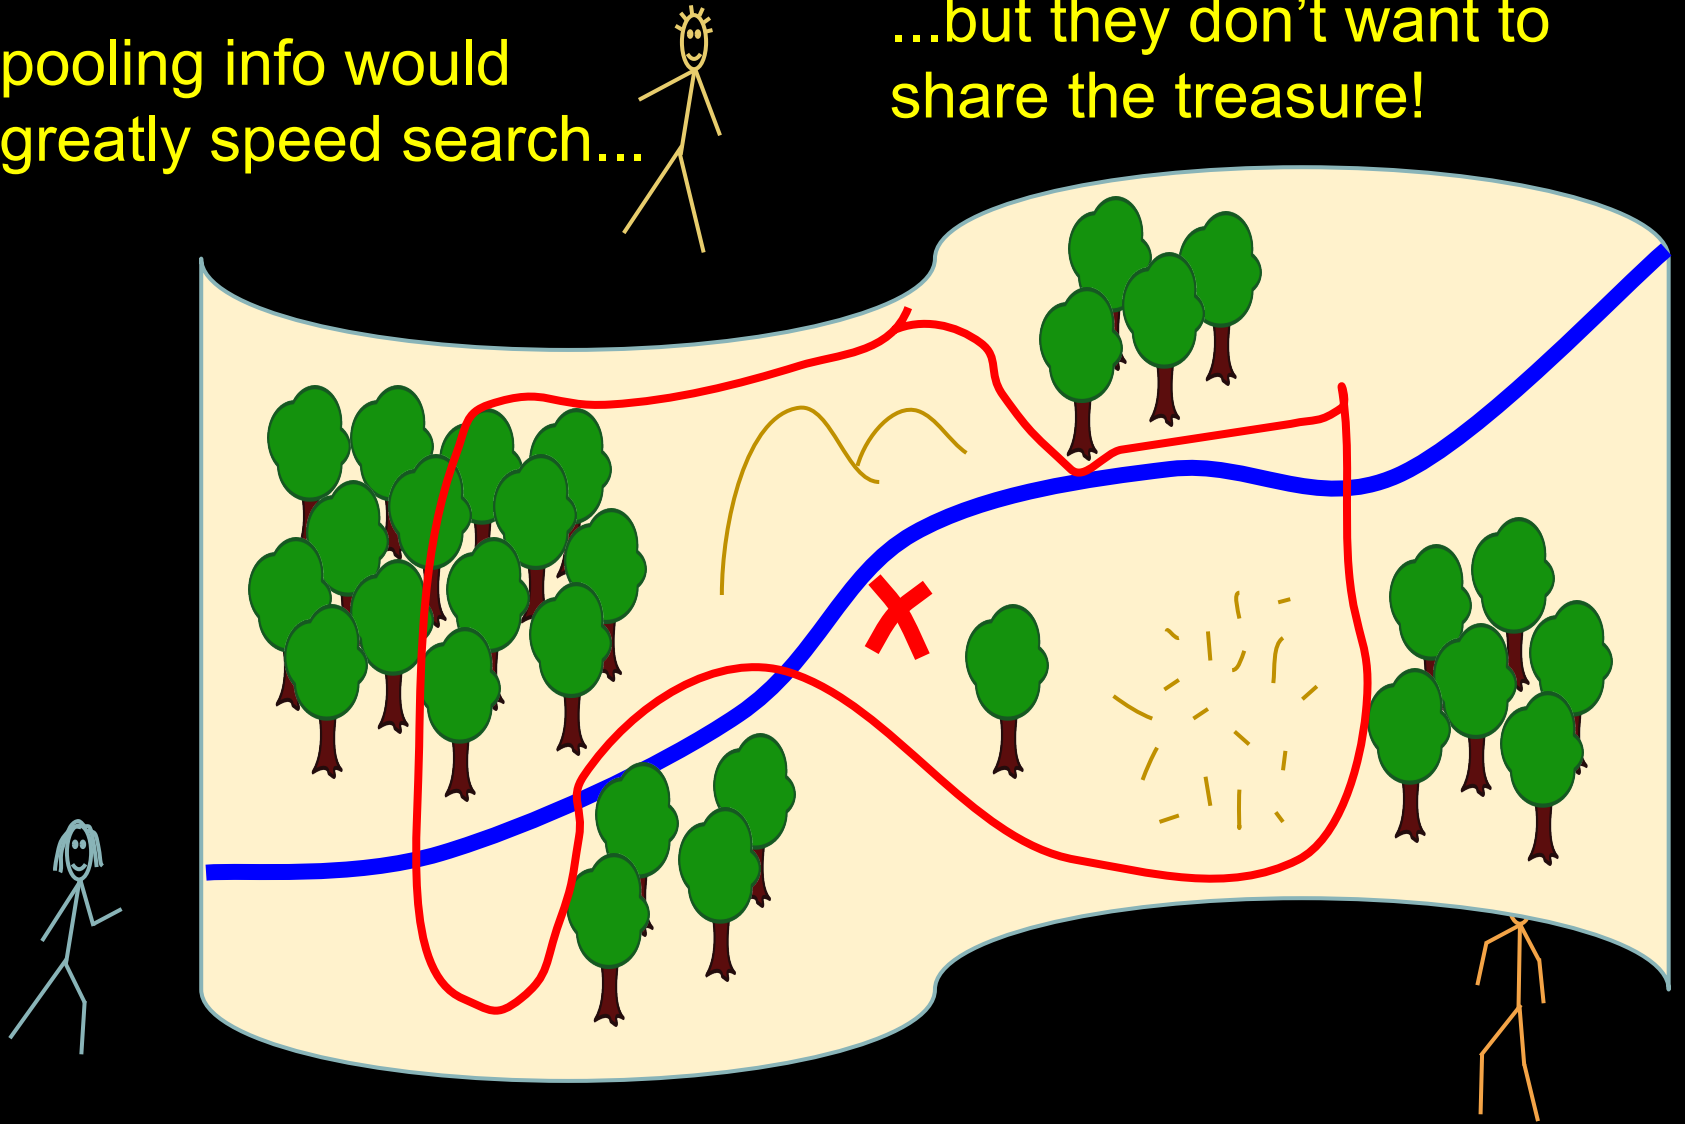


Problem: could take a long time to find the treasure!




pooling info would greatly speed search...

...but they don't want to share the treasure!



Working together?



Captain wants to convince pirates to pool info

- Goal: **design a mechanism** (without money) for *cooperation in a competitive environment*
- Examples: scientific credit, ...

Outline

1. Bo talks: summary of paper (~30min)
 - a. model and goals
 - b. proposed mechanism
 - c. results about the mechanism
 - d. extension to “composable” mechanisms
2. “Guided Discussion” (~20-30min)
 - a. approaches / solution concepts
 - b. goals / desiderata
 - c. models
3. Recap (~5min)


Outline

1. Bo talks: summary of paper
 - a. model and goals
 - b. proposed mechanism
 - c. results about the mechanism
 - d. extension to “composable” mechanisms

Model

discuss: other models, variants

- **island:** set S of locations
- **pirate knows:** subset S_i containing the treasure
(believes treasure is uniformly random in S_i)
- **beliefs about S_k :** arbitrary
(but believes treasure is uniformly random in S_i)
- each location takes one day to dig




Goals -- informally

1. "Welfare"

subject to

2. "Fairness"

3. "Truthfulness"



Goals -- our interpretation

1. Welfare -- reduce number of digs

subject to

2. Fairness -- preserve “winning chances”

3. Truthfulness -- true report maximizes $\Pr[\text{win}]$

*discuss: other
interpretations*


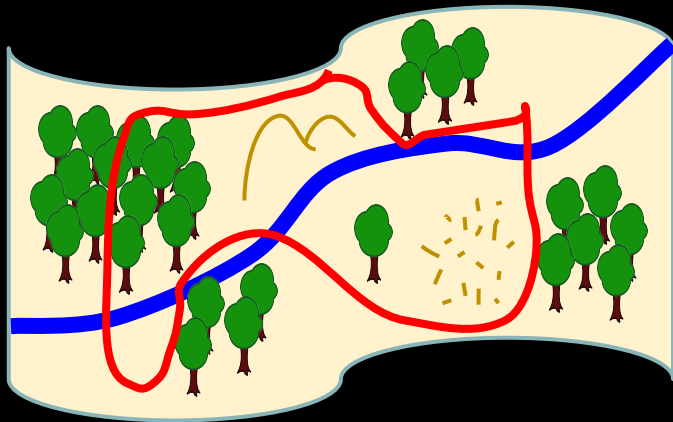
Outline

1. Bo talks: summary of paper
 - a. model and goals
 - b. proposed mechanism
 - c. results about the mechanism
 - d. extension to “composable” mechanisms

Mechanism: Framework

discuss: other frameworks


1. Each pirate reports his/her set S_i
2. Captain partitions the **intersection**
3. Pirate i may only dig in assigned area



How to Partition the Intersection?

Simplified exploration game:

- pretend i explores S_i in uniformly random order
- pretend treasure is uniformly random in intersection
- i has some probability p_i of winning the treasure
- partition according to p and assign i a p_i fraction



How to Partition the Intersection?

Simplified exploration game:

- pretend i explores S_i in uniformly random order
- pretend treasure is uniformly random in intersection
- i has some probability p_i of winning the treasure
- partition according to p and assign i a p_i fraction

Computational efficiency points:

- key obs: probabilities do not depend on set structure!
- to implement, just need to compute set intersection and partition it efficiently

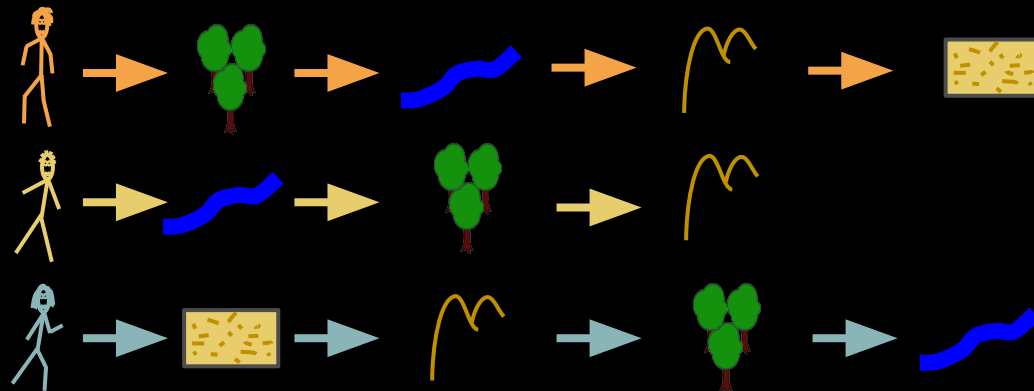
How to Partition the Intersection?

Simplified exploration game:

- pretend i explores S_i in uniformly random order
- pretend treasure is uniformly random in intersection
- i has some probability p_i of winning the treasure
- partition according to p and assign i a p_i fraction

One implementation:

- draw random order for each i
- give i all locations that i would win




How to Partition the Intersection?

Simplified exploration game:

- pretend i explores S_i in uniformly random order
- pretend treasure is uniformly random in intersection
- i has some probability p_i of winning the treasure
- partition according to p and assign i a p_i fraction

One implementation:

- draw random order for each i
- give i all locations that i would win




How to Partition the Intersection?

Simplified exploration game:

- pretend i explores S_i in uniformly random order
- pretend treasure is uniformly random in intersection
- i has some probability p_i of winning the treasure
- partition according to p and assign i a p_i fraction

One implementation:

- draw random order for each i
- give i all locations that i would win




How to Partition the Intersection?

Simplified exploration game:

- pretend i explores S_i in uniformly random order
- pretend treasure is uniformly random in intersection
- i has some probability p_i of winning the treasure
- partition according to p and assign i a p_i fraction

One implementation:

- draw random order for each i
- give i all locations that i would win



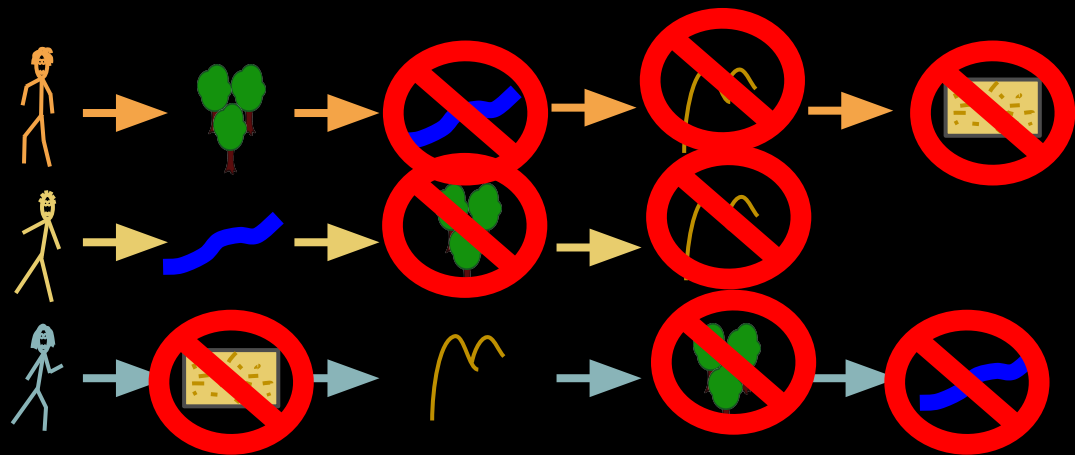
How to Partition the Intersection?

Simplified exploration game:

- pretend i explores S_i in uniformly random order
- pretend treasure is uniformly random in intersection
- i has some probability p_i of winning the treasure
- partition according to p and assign i a p_i fraction

One implementation:

- draw random order for each i
- give i all locations that i would win



Outline

1. Bo talks: summary of paper
 - a. model and goals
 - b. proposed mechanism
 - c. results about the mechanism
 - d. extension to “composable” mechanisms

Goals -- how did we do?

1. **Welfare:** reduce # of digs

Idea: compare to simplified exploration game

Result: If all sets $\geq 10 \times$ (intersection size),
number of digs is reduced by factor of 10
(as number of pirates grows, \rightarrow factor of 20).

Goals -- how did we do?

2. Fairness: preserve winning chances

Idea: compare to simplified exploration game

Result: $\Pr[\text{win}]$ is exactly the same as in simplified exploration game.

Goals -- how did we do?

3. Truthfulness: reporting truthfully maximizes $\Pr[\text{win}]$ if others are being truthful

Result: *yes*

Sidenote: ϵ -voluntary participation

- not clear how to formally define IR
- ϵ comes (in some sense) from ties and small set sizes


Goals -- how did we do?

3. Truthfulness: reporting truthfully maximizes $\Pr[\text{win}]$ if others are being truthful

Proof idea part 1:

Don't want to report a location not in S_i

- may or may not change intersection
- either way, hurts i 's chances most



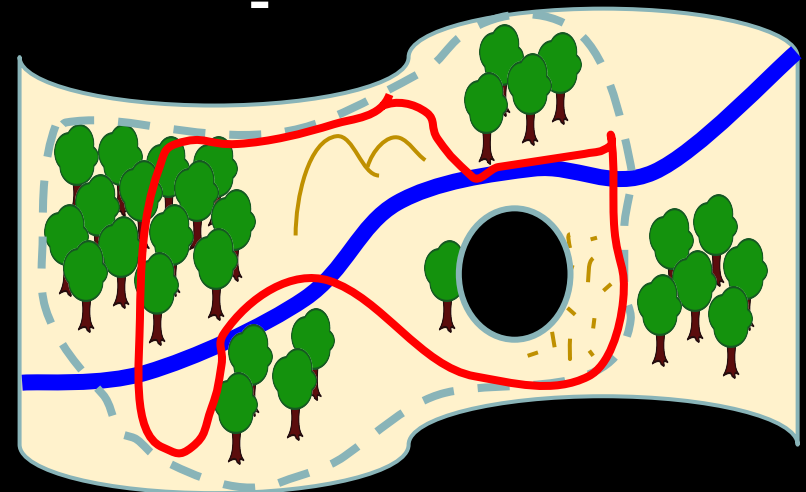
Goals -- how did we do?

3. Truthfulness: reporting truthfully maximizes $\Pr[\text{win}]$ if others are being truthful

Proof idea part 2:

Don't want to omit a location in S_i

- may or may not change intersection
- will help i 's chances
- but balanced by chance it contained the treasure



Outline

1. Bo talks: summary of paper


a. model and goals

b. proposed mechanism


c. results about the mechanism

d. extension to “composable” mechanisms


Mega-Coalitions



Mega-Coalitions




Goal: create a mechanism taking in **coalitions** and outputting a **mega-coalition**



Mega-Mechanism

Idea: less-simplified exploration game

1. Each coalition (recursively) partitions its intersection (agents are coalitions of size one that give themselves their whole set)
2. Now each agent has some resulting set S_i
3. Run the simplified exploration game with these sets




Results

Fairness: **sure**

Truthfulness: **yes**

Dynamics: **a coalition ε -prefers to join earlier**



Outline

1. Bo talks: summary of paper
 - a. model and goals
 - b. proposed mechanism
 - c. results about the mechanism
 - d. extension to “composable” mechanisms

2. “Guided Discussion”
 - a. approaches / solution concepts
 - b. goals / desiderata
 - c. models


3. Recap

Outline

2. “Guided Discussion”

- a. approaches / solution concepts
- b. goals / desiderata
- c. models

Working together?



Captain wants to convince pirates to pool info


- Goal: **design a mechanism** (without money) for *cooperation in a competitive environment*
 - Examples: scientific credit, ...
- Q:** Is this a reasonable problem to solve?
a reasonable approach to solving it?

Challenges of formalizing the setting

- Knowledge of pirates?
- Power of captain?
- ...


Cooperative Game Theory?

Seemed a bad fit...



Dream framework/approach

1. Collect reports S_i
2. Give “hints” to each i
3. Pirates do whatever they want



Achievable?

Outline

2. “Guided Discussion”

- a. approaches / solution concepts
- b. goals / desiderata
- c. models

Goals / Desiderata

1. Welfare - ok, but **what is your benchmark?**

Goals / Desiderata

1. Welfare - ok, but **what is your benchmark?**
2. Fairness (**what is “fair”?**)
 - ours: preserve “spirit of competition”
 - compare: Shapley Value type solution
 - (do other notions of fairness admit truthful solutions?)

Goals / Desiderata

1. Welfare - ok, but **what is your benchmark?**
2. Fairness (**what is “fair”?**)
ours: preserve “spirit of competition”
compare: Shapley Value type solution
(do other notions of fairness admit truthful solutions?)
3. Truthfulness - **necessary?**
max Pr[win] vs max E[utility]
perhaps digging is costly
(is our mechanism is truthful in E[utility] sense?)


Outline

2. “Guided Discussion”

- a. approaches / solution concepts
- b. goals / desiderata
- c. models


Digging into our model

- **island:** set S of locations
- **pirate knows:** subset S_i containing the treasure
(believes treasure is uniformly random in S_i)
- **beliefs about S_j :** arbitrary
(but believes treasure is uniformly random in S_i)
- each location takes one day to dig




Example Bayesian game captured by our model

1. Each pirate has a **partition** of S (the island)
2. Nature picks treasure location **uniformly at random**
3. Each pirate observes $S_i =$ **element of partition**




Example Bayesian game captured by our model

1. Each pirate has a **partition** of S (the island)
2. Nature picks treasure location **uniformly at random**
3. Each pirate observes $S_i =$ **element of partition**




Example Bayesian game captured by our model

1. Each pirate has a **partition** of S (the island)
2. Nature picks treasure location **uniformly at random**
3. Each pirate observes $S_i =$ **element of partition**




Example Bayesian game captured by our model

1. Each pirate has a **partition** of S (the island)
2. Nature picks treasure location **uniformly at random**
3. Each pirate observes $S_i =$ **element of partition**




Example Bayesian game captured by our model

1. Each pirate has a **partition** of S (the island)
2. Nature picks treasure location **uniformly at random**
3. Each pirate observes $S_i =$ **element of partition**



Example Bayesian game captured by our model

1. Each pirate has a **partition** of S (the island)
2. Nature picks treasure location **uniformly at random**
3. Each pirate observes $S_i =$ **element of partition**



What about non-uniform priors?

Problem 1: what if pirate beliefs are inconsistent?

What about non-uniform priors?

Problem 1: what if pirate beliefs are inconsistent?

→ Ok, suppose they are consistent....

Problem 2: how does the mechanism aggregate reports?

What about non-uniform priors?

Problem 1: what if pirate beliefs are inconsistent?

→ Ok, suppose they are consistent....

Problem 2: how does the mechanism aggregate reports?

→ OK, suppose it knows the prior or something....

Q for audience: why can't we re-cut the island so that the prior is now uniform, then run our mechanism?

What about non-uniform priors?

Problem 1: what if pirate beliefs are inconsistent?

→ Ok, suppose they are consistent....

Problem 2: how does the mechanism aggregate reports?

→ OK, suppose it knows the prior or something....

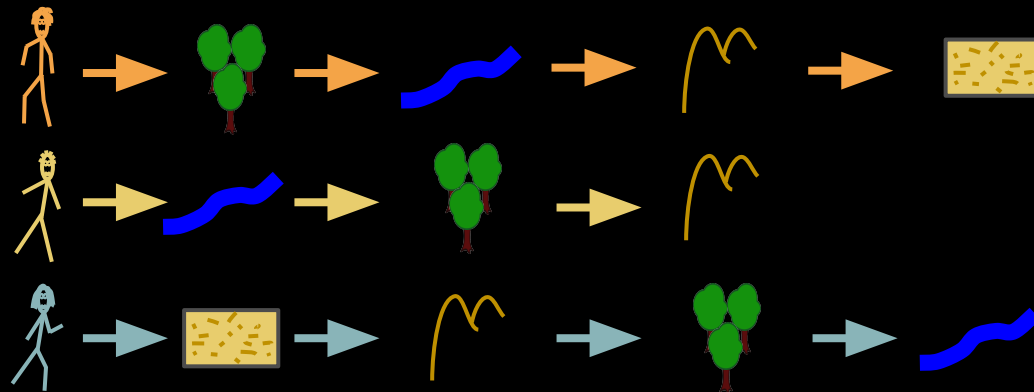
Q for audience: why can't we re-cut the island so that the prior is now uniform, then run our mechanism?

Problem 3: how to get “fairness” and truthfulness??

Ideally: robust to beliefs about other agents

A “perfect” mechanism?


Q for audience: how to change “simulated exploration” mechanism to be truthful with more general agent beliefs?



A “perfect” mechanism?

Q for audience: how to change “simulated exploration” mechanism to be truthful with more general agent beliefs?

1. Each pirate submits an exploration strategy
2. The mechanism simulates everyone’s strategy
3. Give each location to the pirate that explores it first in simulation




A “more perfect” mechanism

1. Each pirate submits their “signal” / information
2. The mechanism simulates an “equilibrium” (like what?)
3. Give each location to the pirate that explores it first in simulation

1. Satisfying?

2. How to compute “equilibrium”?



Outline

1. Bo talks: summary of paper
 - a. model and goals
 - b. proposed mechanism
 - c. results about the mechanism
 - d. extension to “composable” mechanisms

2. “Guided Discussion”
 - a. approaches / solution concepts
 - b. goals / desiderata
 - c. models

3. Recap


What did we do?

- Model of **competitive search problem**
- Mechanism for **cooperation**
- Welfare, fairness, and truthfulness properties




Future Work

- **Building up:** Extensions, variants, dynamics of coalition formation...
- **Digging down:** assumptions, model, alternative frameworks, bargaining with information sharing, alternative solution concepts...



Future Work

- **Building up:** Extensions, variants, dynamics of coalition formation...
- **Digging down:** assumptions, model, alternative frameworks, bargaining with information sharing, alternative solution concepts...



Thanks!