Where Information and Incentives Collide

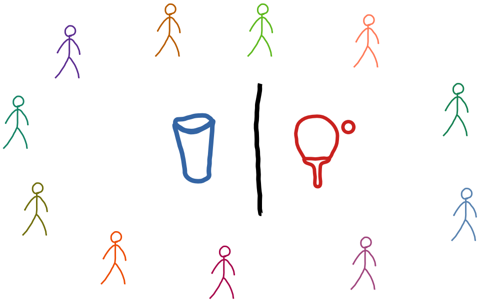

Bo Waggoner

University of Colorado, Boulder

September 14, 2023

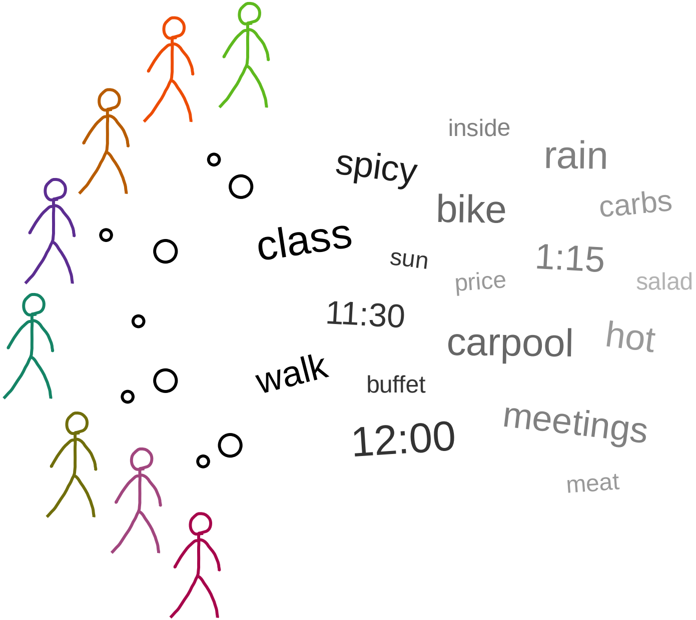

inside
spicy
rain
bike
carbs
class
sun
1:15
price
salad
11:30
carpool
hot
walk
buffet
meetings
12:00
meat


inside
rain
spicy
bike
carbs
class
sun
price
1:15
salad
11:30
carpool
hot
walk
buffet
12:00
meetings
meat
Let’s just go to C4C

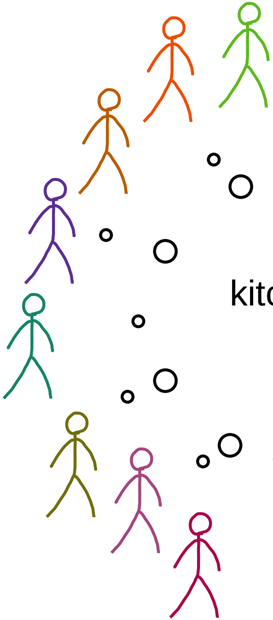

collaboration

staff

rec rooms

focus

kitchenettes

faculty

classrooms

windows

students

community

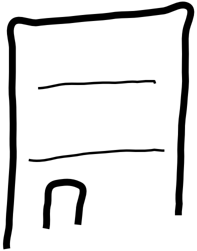

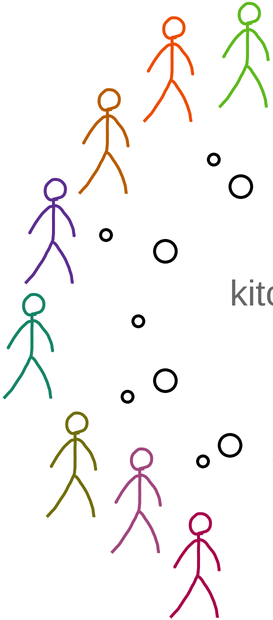

collaboration

staff

rec rooms

focus

kitchenettes

faculty

classrooms

windows

students

community

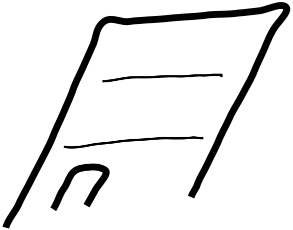

stays up?

Outline

1. Contracts
2. Public projects
3. Matching

Themes

- Gathering hidden information
- Navigating preferences and strategic behavior
- Coordinating good decisionmaking
- Solving algorithmic problems in societal contexts

Contracts

Joint work with Maneesha Papireddygari

The 2022 ACM Conference on Economics and Computation (**EC ’22**)

Eliciting predictions

First question: how to **elicit a prediction**?

Eliciting predictions

First question: how to **elicit a prediction**?

- An expert makes a **prediction** p

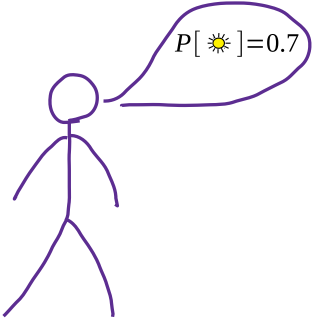

Eliciting predictions

First question: how to **elicit a prediction**?

- An expert makes a **prediction** p
- We **observe** whether the event happened, y

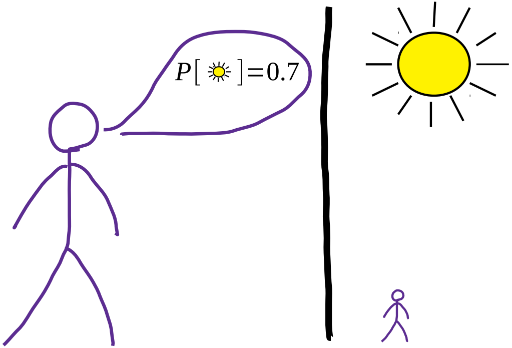

Eliciting predictions

First question: how to **elicit a prediction**?

- An expert makes a **prediction** p
- We **observe** whether the event happened, y
- We assign a **score** or **payment** $S(p, y)$

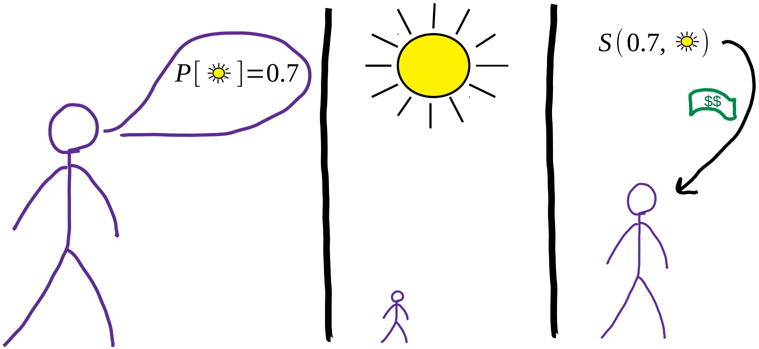

Characterization of proper scoring rules

Fact (McCarthy 1956; Savage 1971; . . .)

A scoring rule is **proper** *(meaning truthful) if and only if it is*

$$S(p, y) = G(p) + \nabla G(p) \cdot (\delta_y - p)$$

for some convex G*.*

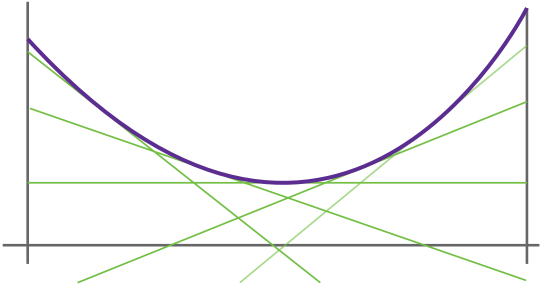

Next problem: information acquisition

What if the expert can **acquire costly information**?

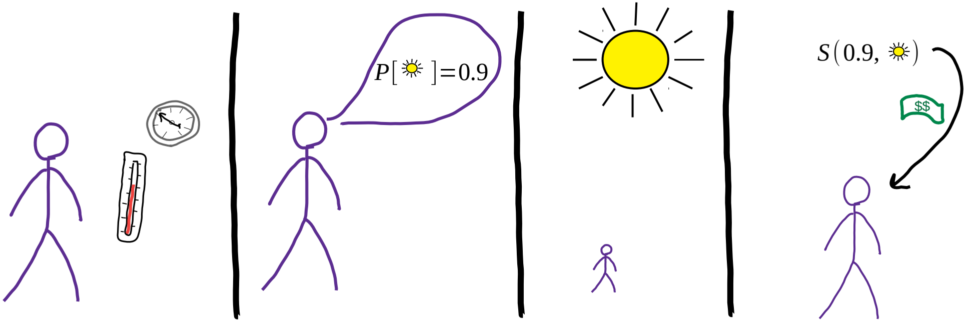

How do we incentivize truthful, accurate predictions?

Next problem: hidden action (moral hazard)

How to write **contracts** to incentivize good, unverifiable work?

- Long history in economics
- Recent algorithmic work in CS/Econ

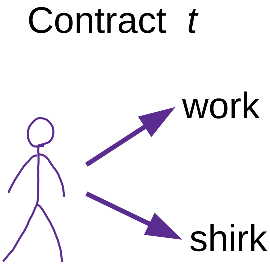

Next problem: hidden action (moral hazard)

How to write **contracts** to incentivize good, unverifiable work?

- Long history in economics
- Recent algorithmic work in CS/Econ

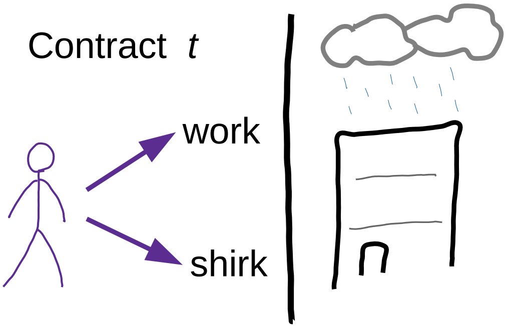

Next problem: hidden action (moral hazard)

How to write **contracts** to incentivize good, unverifiable work?

- Long history in economics
- Recent algorithmic work in CS/Econ

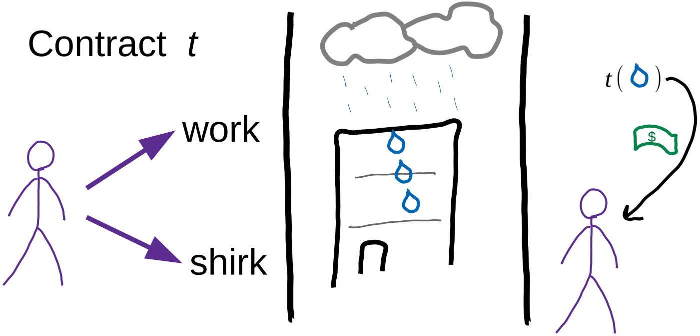

What about both at once?

Final problem: **hidden action with information acquisition**.

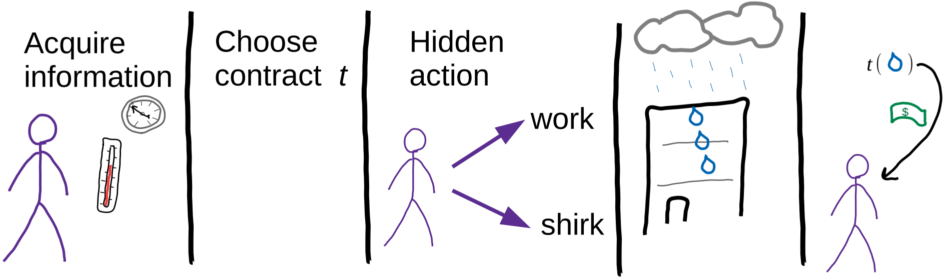

Overview of problems

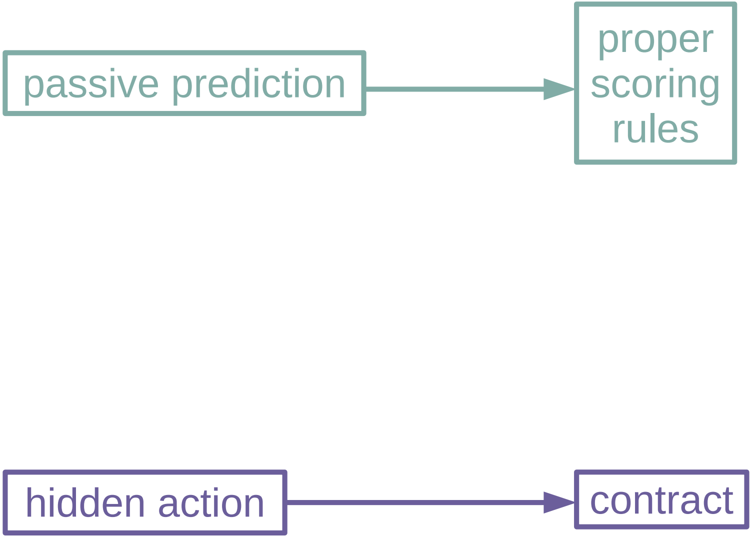

Overview of problems

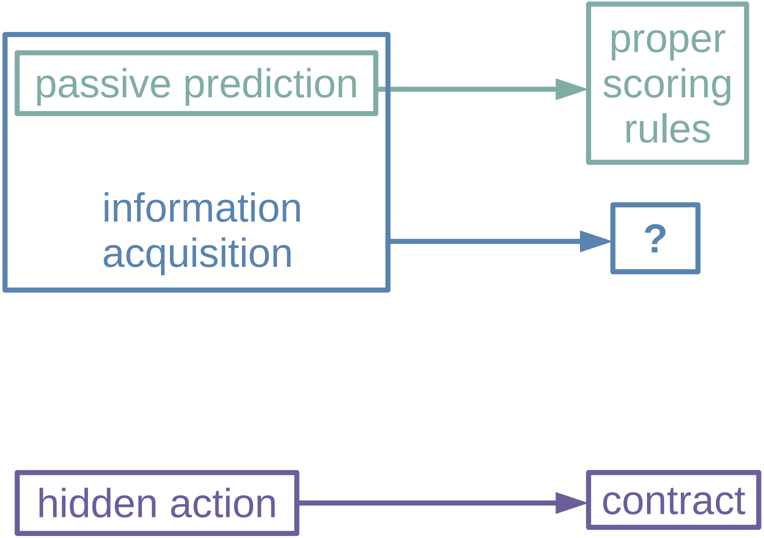

Overview of problems

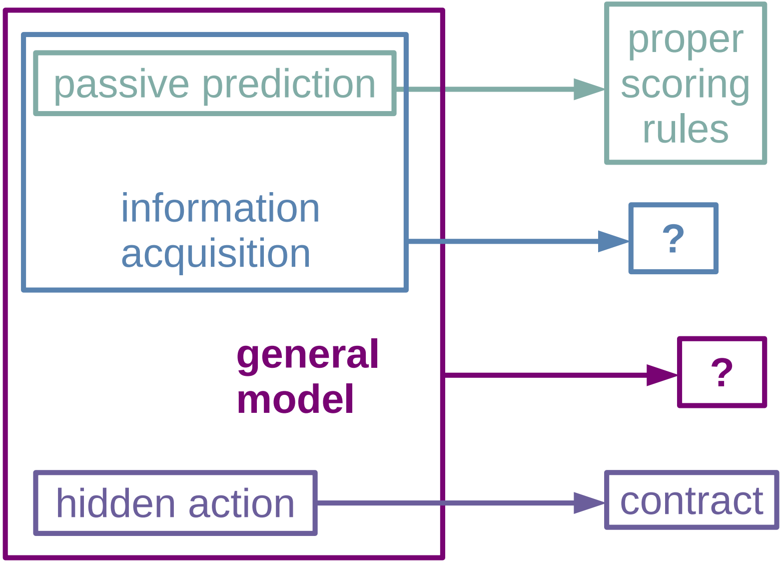

Key insight

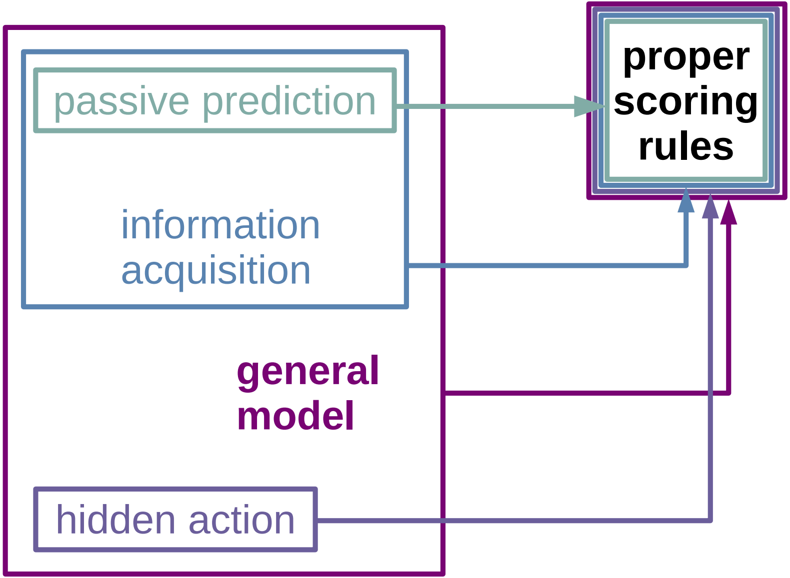
passive prediction
information acquisition
general model
hidden action
proper scoring rules

Results

Proposition

Any solution to any of the above problems is, without loss of generality, a **proper scoring rule**.

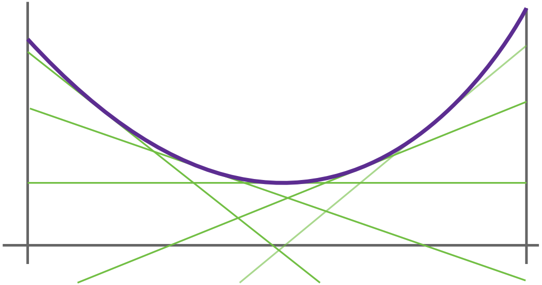

Results

Proposition

Any solution to any of the above problems is, without loss of generality, a **proper scoring rule***.*

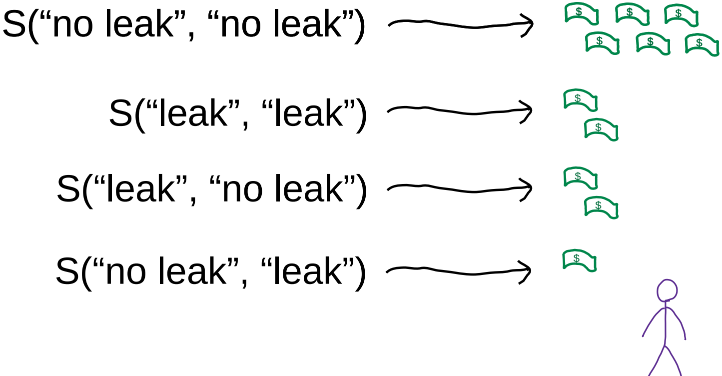

Results

Proposition

Any solution to any of the above problems is, without loss of generality, a **proper scoring rule**.

Proposition

For any of the above problems, given any target plan, we can construct an incentive-compatible, optimal scoring rule in polynomial time.

Results

Proposition

Any solution to any of the above problems is, without loss of generality, a **proper scoring rule***.*

Proposition

For any of the above problems, given any target plan, we can construct an incentive-compatible, optimal scoring rule in polynomial time.

Proposition

For the information acquisition problem, there is a closed-form solution (an inverted pyramid).

Connections and takeaways

- actions $\leftrightarrow$ predictions
- value of information
- framing **contract design** as **information elicitation**

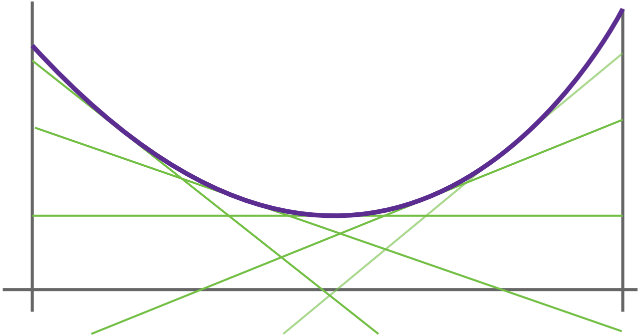

Public Projects

Joint work with Mary Monroe,

in preparation

Funding: The Ethereum Foundation (2022-)

Public Projects

A classic problem: a bunch of people want to decide what to do together.

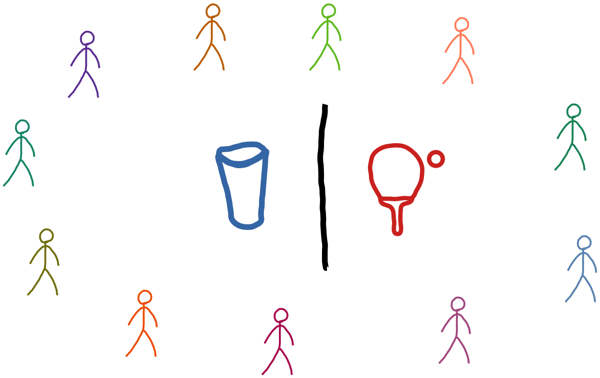

A mathematical model. . .

Each person has a **value** for each alternative

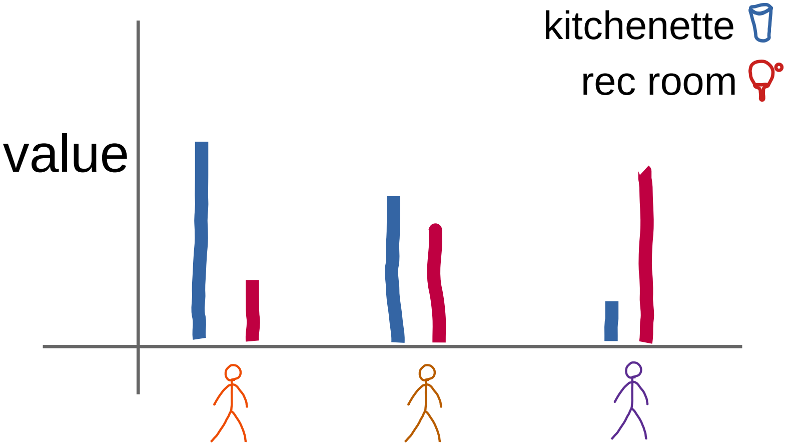

A mathematical model. . .

Each person has a **value** for each alternative
Assume: value is in units of money

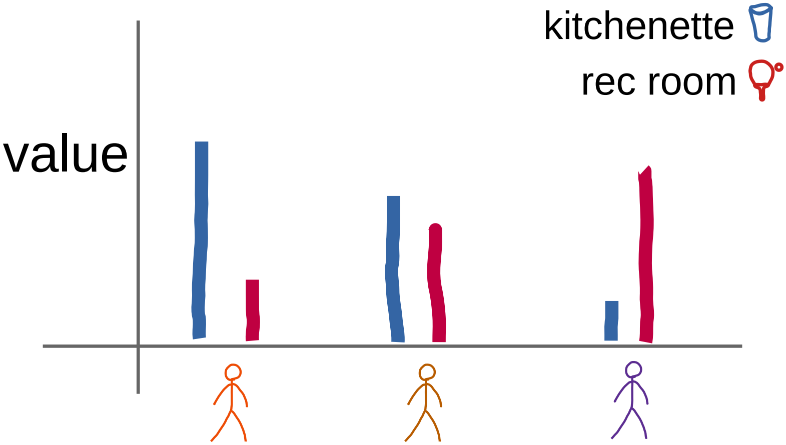

A mathematical model. . .

Each person has a **value** for each alternative
Assume: value is in units of money
Goal: maximize **social welfare** of project = *total value*

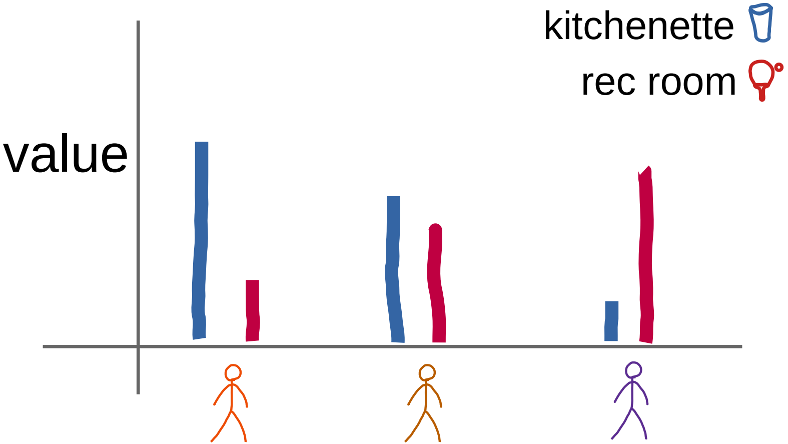

Existing solutions

Classical solution: a "VCG mechanism".

Maximizes welfare, but. . .

- *Fragile:* false-name attacks, . . .
- *Unpredictable:* payments may be zero, very high, in between. . .
- *No revenue:* often, nobody pays anything

Using quadratic voting

Proposal (Eguia et al. 2022): quadratic-voting-like approach!

- Each person casts "votes" for (or against) each option
- Pay c times the number of votes, **squared** *c a parameter*

Using quadratic voting

Proposal (Eguia et al. 2022): quadratic-voting-like approach!

- Each person casts "votes" for (or against) each option
- Pay c times the number of votes, **squared** *c a parameter*
- Pick the winner with "softmax":

$$\Pr[\text{select project } j] = \frac{e^{\text{total votes for } j}}{\sum_k e^{\text{total votes for } k}}$$

Prior work

Theorem (Eguia et al. 2022): if participants' preferences are drawn i.i.d. with bounded values, then in any symmetric Bayes-Nash equilibrium,

$$\Pr[\text{select outcome with maximum social welfare}] \to 1$$

as

$$\text{num. participants} \to \infty.$$

Our results

Our results

Proposition

With two choices, in any pure-strategy equilibrium,

$$\frac{\textit{Social Welfare}}{\textit{Optimal SW}} \geq 1 - \sqrt{\frac{2c}{U_1 - U_2}}.$$

Our results

Proposition

With two choices, in any pure-strategy equilibrium,

$$\frac{\textit{Social Welfare}}{\textit{Optimal SW}} \geq 1 - \sqrt{\frac{2c}{U_1 - U_2}}.$$

Proposition

A pure-strategy equilibrium exists if $c > \frac{3}{2} \max_{i,k} |u_k^i|$.

Proof ingredients

- analyze Hessian of utility function *related: exponential families*
- fixed-point theorem for concave utilities
- properties of xe^{-x}

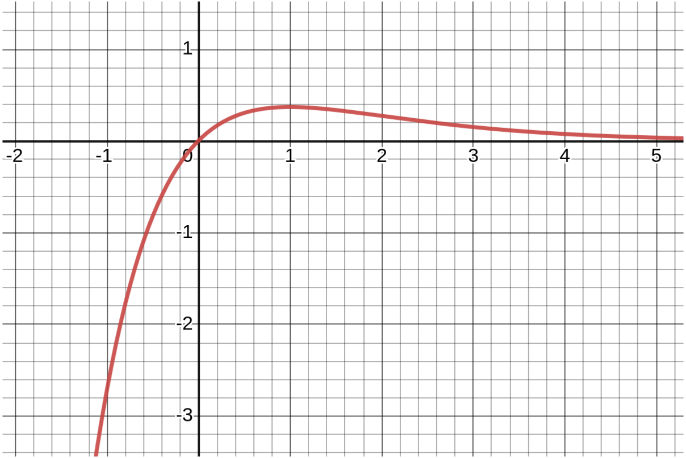

Results continued

Proposition

With m choices, in any pure-strategy equilibrium,

$$\frac{\textit{Social Welfare}}{\textit{Optimal SW}} \geq 1 - f(c, U_1, \ldots, U_m)$$

where we can write down f, but it ain't pretty.

Results continued

Proposition

With m choices, in any pure-strategy equilibrium,

$$\frac{\textit{Social Welfare}}{\textit{Optimal SW}} \geq 1 - f(c, U_1, \ldots, U_m)$$

where we can write down f, but it ain't pretty.

Conjecture

1. *If all participants agree on the ordering of the alternatives, a pure-strategy equilibrium always exists.*
2. *In mixed-strategy equil., SW remains high under many conditions.*

Future work

- Explore limits of this mechanism
- Explore connections to **prediction** and **decision markets**

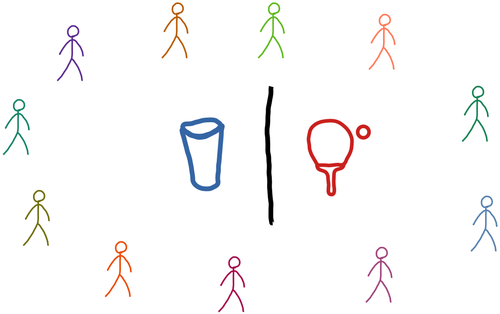

Matching

Joint work with Robin Bowers,

to appear in the 2023 Conference on Web and Internet Economics (**WINE**)

Funding: The National Science Foundation (2023-)

Matching

Classic problem: how to match e.g. workers to jobs?

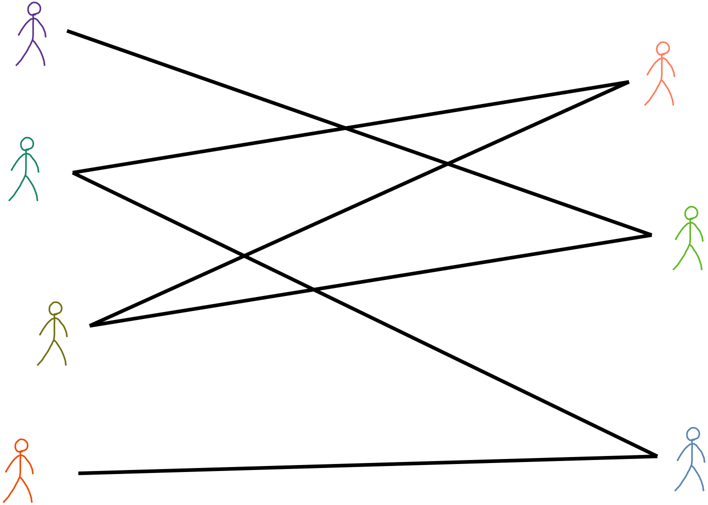

Max-Weight Matching

One formulation: maximize **total value** of the matching.

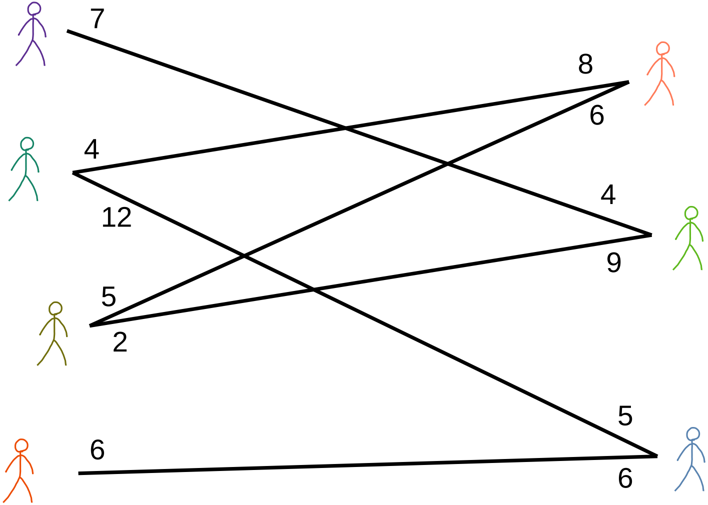

Problem: unknown values

Typically, we initially **don't know** our preferences.

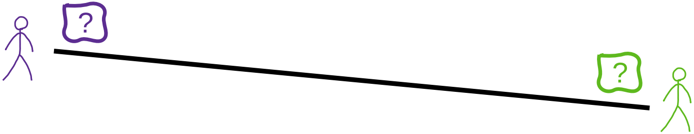

Problem: unknown values

Typically, we initially **don't know** our preferences.

We need to **spend time, effort, and money** to find out.
Reading résumés, market research, interviews, . . .

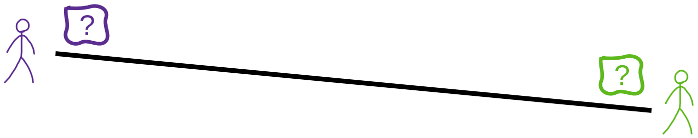

Problem: unknown values

Typically, we initially **don't know** our preferences.

We need to **spend time, effort, and money** to find out.

Reading résumés, market research, interviews, . . .

Model: each person has a **distribution** over possible values for each job, and a **cost** for finding out.

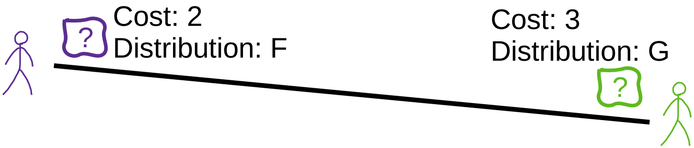

Simplified version of the problem

- Matching people to items

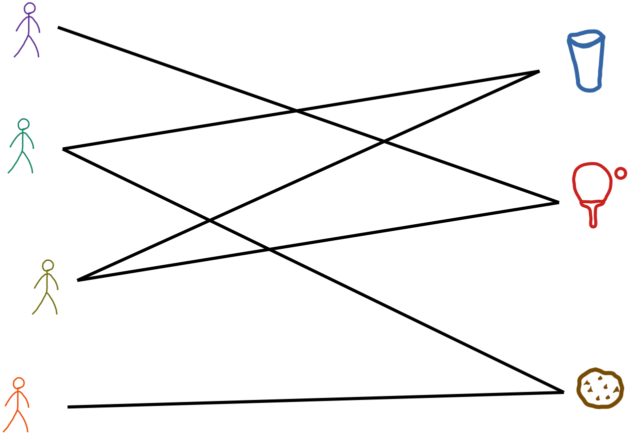

Simplified version of the problem

- Matching people to items
- Selling **one item**

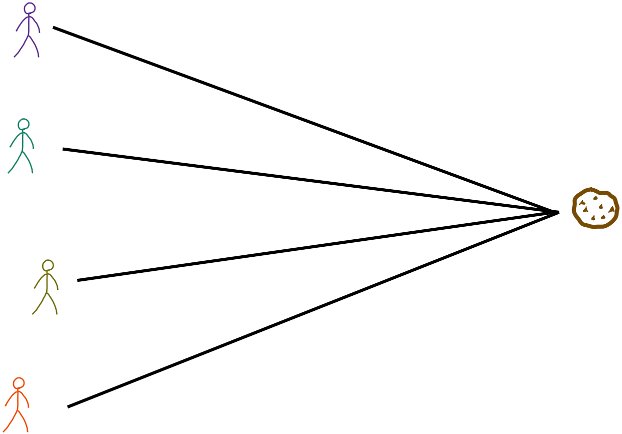

The Pandora's Box Problem

Due to Weitzman (1979) *also a case of Gittins index thm.*

The Pandora's Box Problem

Due to Weitzman (1979) *also a case of Gittins index thm.*

Optimal "descending policy":

- Compute an **index** for each alternative.
- Inspect from **highest index down** until we find a large value.

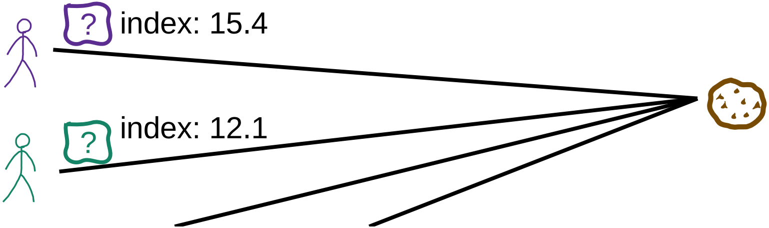

The Pandora's Box Problem

Due to Weitzman (1979) *also a case of Gittins index thm.*

Optimal "descending policy":

- Compute an **index** for each alternative.
- Inspect from **highest index down** until we find a large value.

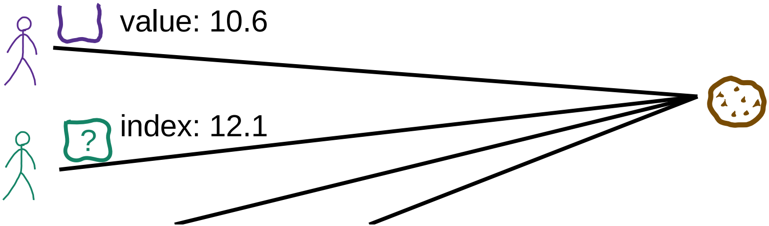

The Pandora's Box Problem

Due to Weitzman (1979) *also a case of Gittins index thm.*

Optimal "descending policy":

- Compute an **index** for each alternative.
- Inspect from **highest index down** until we find a large value.

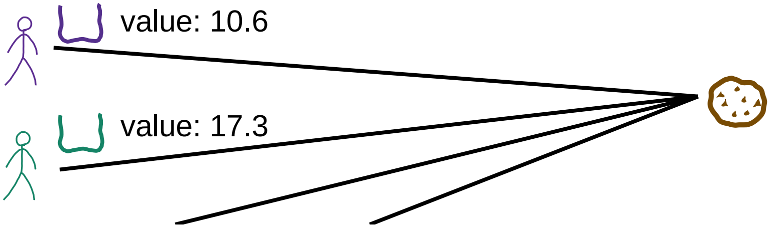

The Pandora's Box Problem

Due to Weitzman (1979) *also a case of Gittins index thm.*

Optimal "descending policy":

- Compute an **index** for each alternative.
- Inspect from **highest index down** until we find a large value.

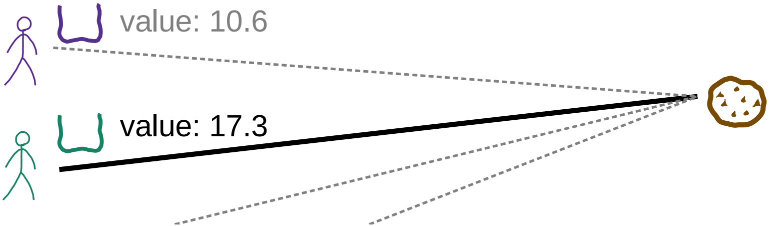

Pandora's Box for auctions

Application: selling one item [*Kleinberg, Waggoner, Weyl (EC 2016).*]

Pandora's Box for auctions

Application: selling one item [*Kleinberg, Waggoner, Weyl (EC 2016).*]

Idea: mimick the optimal policy with a **descending-price auction**.

Pandora's Box for auctions

Application: selling one item [*Kleinberg, Waggoner, Weyl (EC 2016)*.]

Idea: mimick the optimal policy with a **descending-price auction**.

Result: constant-factor approximation to **optimal social welfare**.

Pandora's Box for auctions

Application: selling one item [*Kleinberg, Waggoner, Weyl (EC 2016).*]

Idea: mimick the optimal policy with a **descending-price auction**.

Result: constant-factor approximation to **optimal social welfare**.

Observation: failure of ascending-price; and of any sealed-bid auction.

Back to our problem: two-sided matching

Proposal: the **Marshallian Match**.[1]

[1]Proposed, but not analyzed, in *Waggoner, Weyl (2019).*

Back to our problem: two-sided matching

Proposal: the **Marshallian Match**.[1]

- Participants maintain bids on all potential partners.

[1]Proposed, but not analyzed, in *Waggoner, Weyl (2019).*

Back to our problem: two-sided matching

Proposal: the **Marshallian Match**.[1]

- Participants maintain bids on all potential partners.
- A global price descends over time.

[1] Proposed, but not analyzed, in *Waggoner, Weyl (2019).*

Back to our problem: two-sided matching

Proposal: the **Marshallian Match**.[1]

- Participants maintain bids on all potential partners.
- A global price descends over time.
- When the price reaches the total bid on an edge, it matches.

[1] Proposed, but not analyzed, in *Waggoner, Weyl (2019).*

Back to our problem: two-sided matching

Proposal: the **Marshallian Match**.[1]

- Participants maintain bids on all potential partners.
- A global price descends over time.
- When the price reaches the total bid on an edge, it matches.
- Both sides **pay their bids**.

[1] Proposed, but not analyzed, in *Waggoner, Weyl (2019).*

Results (1)

Theorem

If all participants' values are positive,
the Marshallian Match guarantees, in any Bayes-Nash equilibrium,

$$\frac{\textit{Social Welfare}}{\textit{Optimal SW}} \geq \frac{1}{8}.$$

Results (1)

Theorem

If all participants' values are positive,
the Marshallian Match guarantees, in any Bayes-Nash equilibrium,

$$\frac{\textit{Social Welfare}}{\textit{Optimal SW}} \geq \frac{1}{8}.$$

- Holds for model with inspection costs
- Also holds for matchings on hypergraphs (group formation)
  factor depends on maximum group size

Some proof ingredients

Ingredient 1: greedy max-weight matching

Some proof ingredients

Ingredient 1: greedy max-weight matching

Ingredient 2: **smoothness** framework of algorithmic mechanism design

Some proof ingredients

Ingredient 1: greedy max-weight matching

Ingredient 2: **smoothness** framework of algorithmic mechanism design

Ingredient 3: Pandora's analysis ideas from KWW16

Some proof ingredients

Ingredient 1: greedy max-weight matching

Ingredient 2: **smoothness** framework of algorithmic mechanism design

Ingredient 3: Pandora's analysis ideas from KWW16

Ingredient 4: Rebate variant of Match: align incentives with early matching

Some proof ingredients

Ingredient 1: greedy max-weight matching

Ingredient 2: **smoothness** framework of algorithmic mechanism design

Ingredient 3: Pandora's analysis ideas from KWW16

Ingredient 4: Rebate variant of Match: align incentives with early matching

Ingredient 5: Information-hiding allows counterfactual analysis

Results (2)

Results (2)

Theorem

In general settings with common-knowledge values, if player strategies are 2*-ex-ante stable, the Marshallian Match guarantees*

$$\frac{\textit{Social Welfare}}{\textit{Optimal SW}} \geq \frac{1}{8}.$$

Results (2)

Theorem

In general settings with common-knowledge values,
if player strategies are 2-ex-ante stable,
the Marshallian Match guarantees

$$\frac{\textit{Social Welfare}}{\textit{Optimal SW}} \geq \frac{1}{8}.$$

- *Ex-ante stability:* extension of equilibrium to pairs of players
- Proven for model without inspection costs; may extend
- Unknown: extends to Bayes-Nash setting? (**main open problem**)

Outro

Outro

Other exciting work going on in **algorithmic economics and theoretical ML group**[2]: Prof. Raf Frongillo; Ph.D. students Dhamma, Anish, Maneesha, Rick, Robin, Mary, Melody, Elias; **theory group**: Prof. Josh Grochow, Prof. Huck Bennett, students, ...

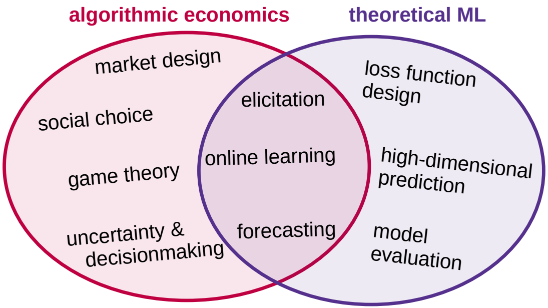

[2][JF, RF, BW (COLT 2020, NeurIPS 2021, JMLR 2023)] [RF, BW (NeurIPS 2021)] [RF, AG, AT, BW (EC 2021)] [DK, RF, BW (ICML 2023)].

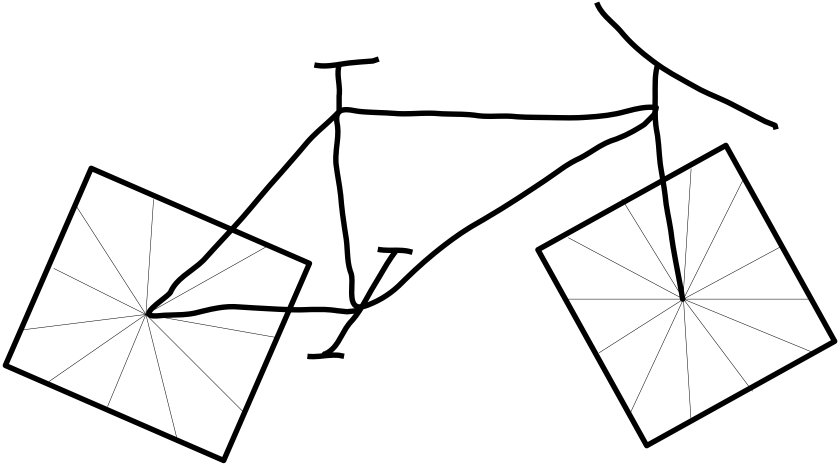

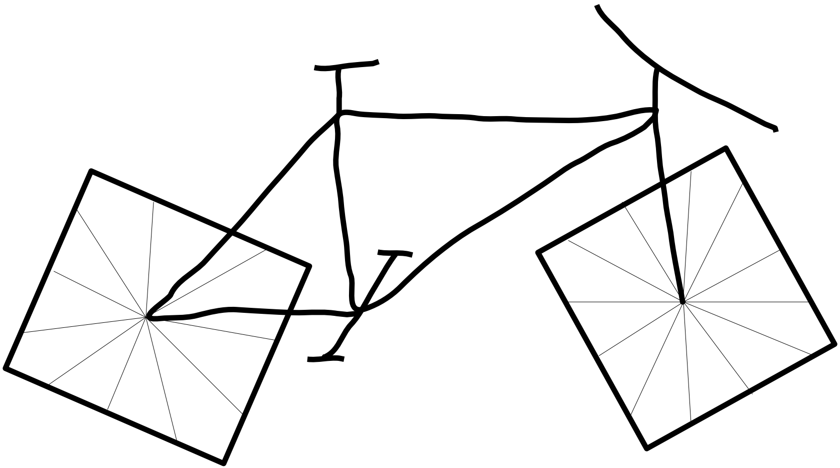

Thanks!