Designing Markets for Daily Deals

Yang Cai (Berkeley/McGill)
Mohammad Mahdian (Google)
Aranyak Mehta (Google)
Bo Waggoner (Harvard)

WINE 2013

Motivation: Daily Deals

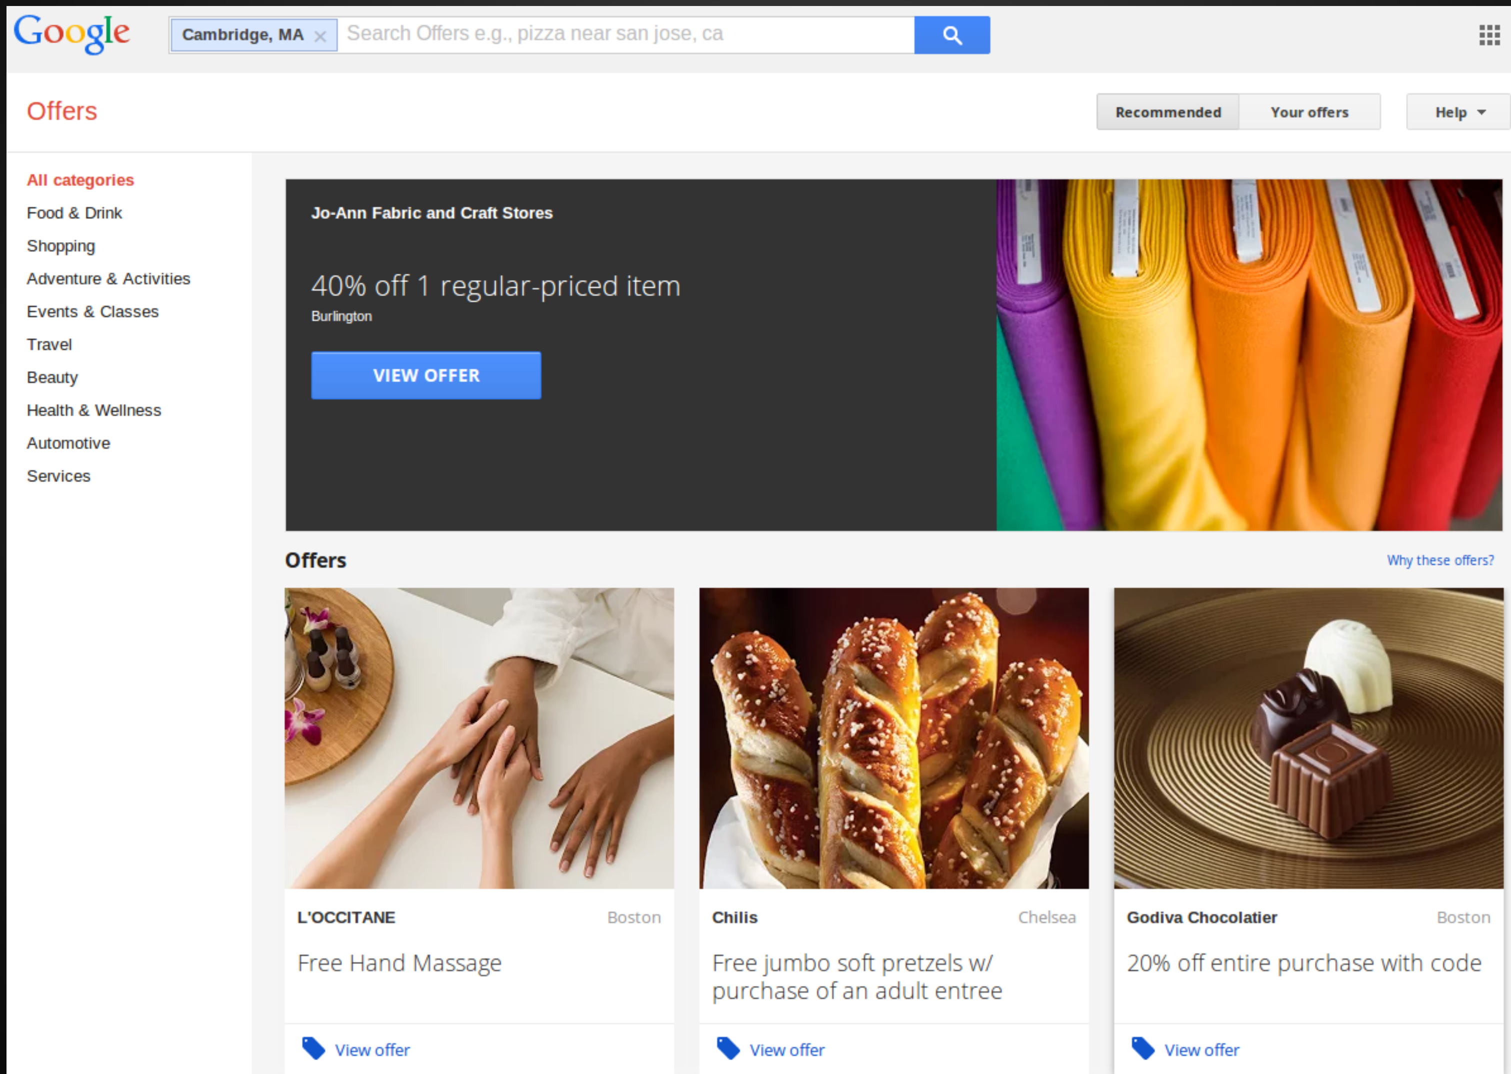

Problem statement

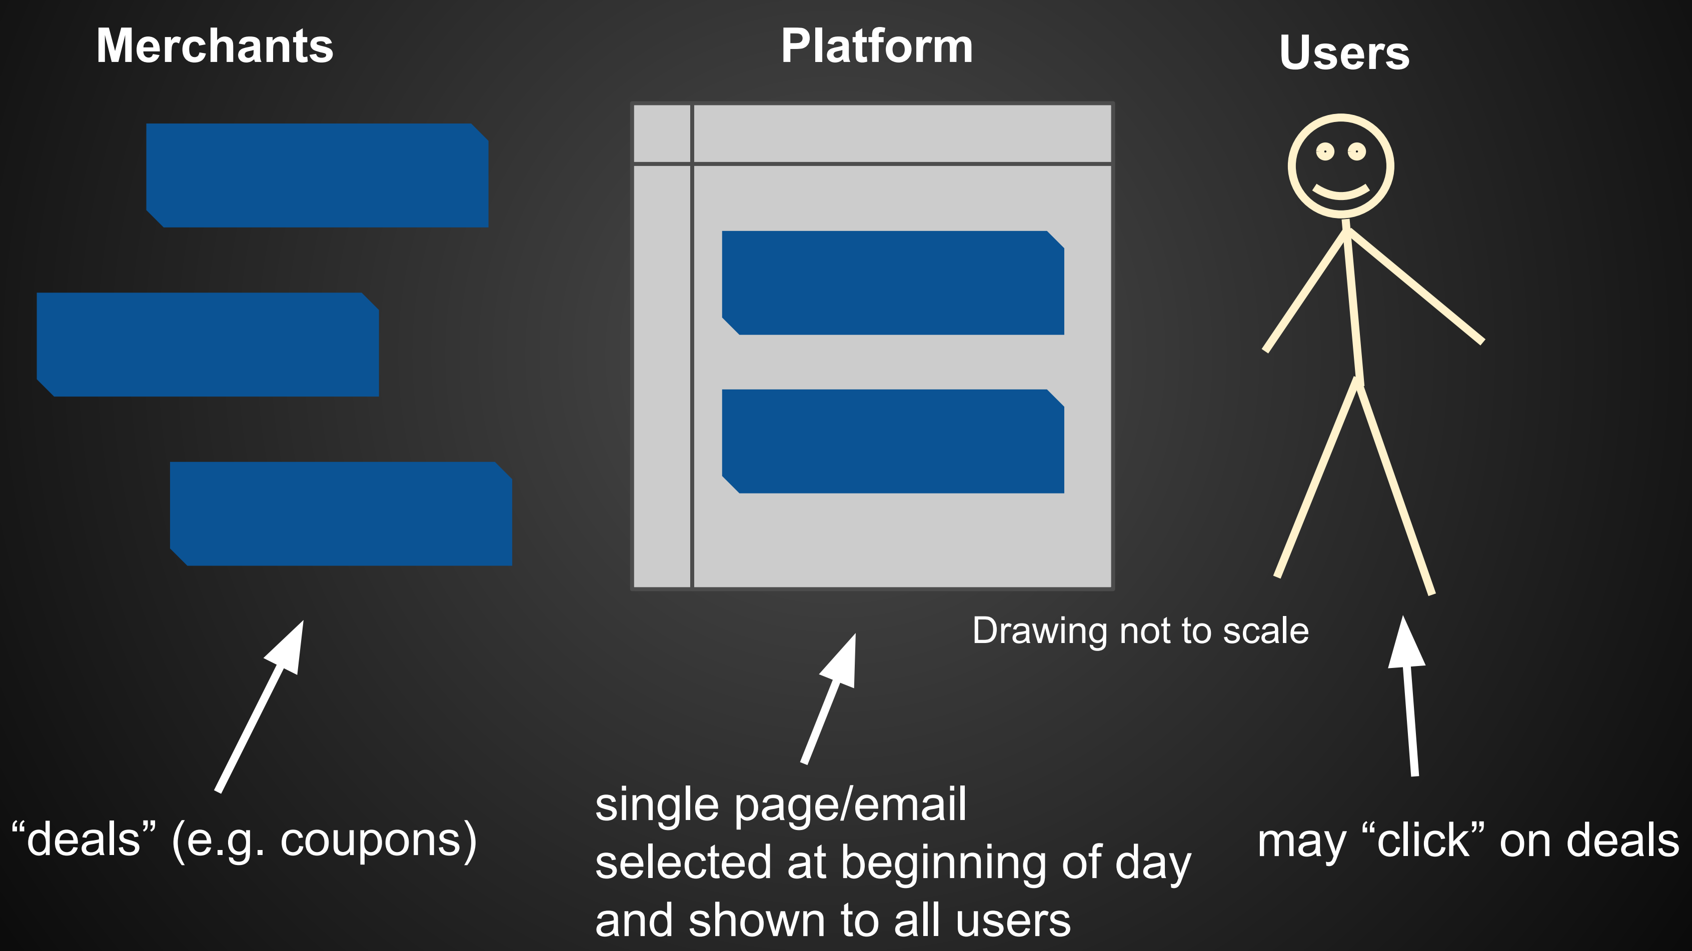

Problem statement

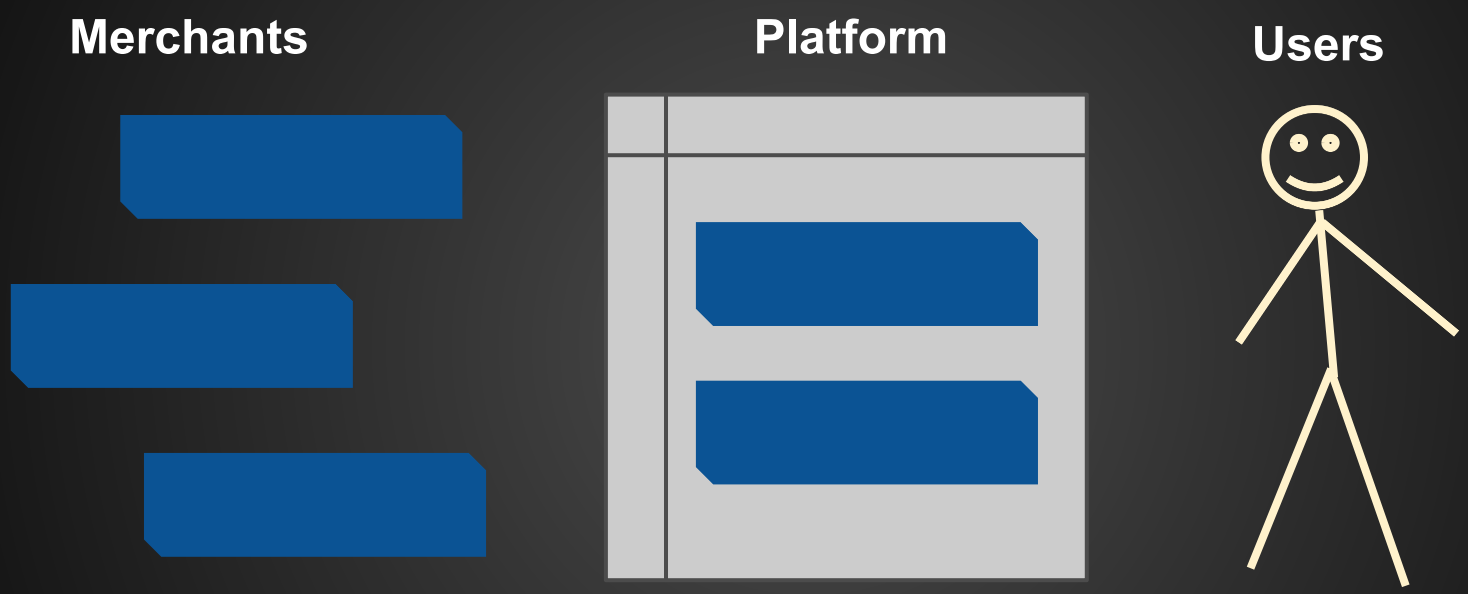

Drawing not to scale

Task: design an *auction* to pick deals
Twist: care about *users'* welfare
Challenge: merchants know value to users; platform may not

Outline

1. Really simple model for daily deals, results

2. Really general model, characterization

3. Applications and conclusion

Goals of talk: (a) state/solve daily deals problem
(b) **general** auction takeaways

Outline

1. Really simple model for daily deals, results
2. Really general model, characterization
3. Applications and conclusion

Goals of talk: (a) state/solve daily deals problem
(b) **general** auction takeaways

Really Simple Model

- **One** winning deal
- **One** user

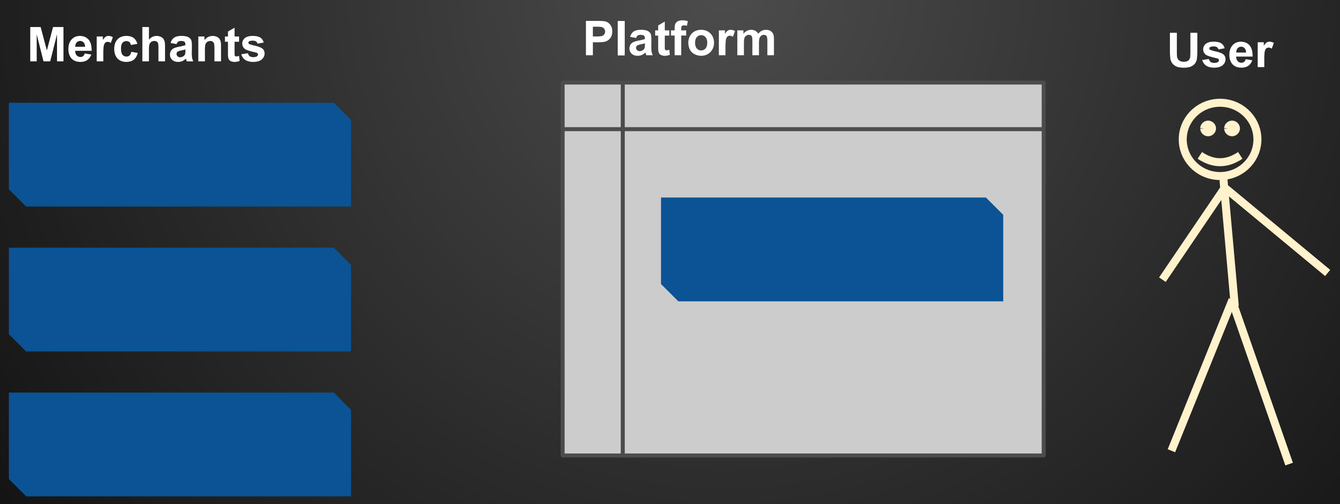

Prologue: Standard auction setting

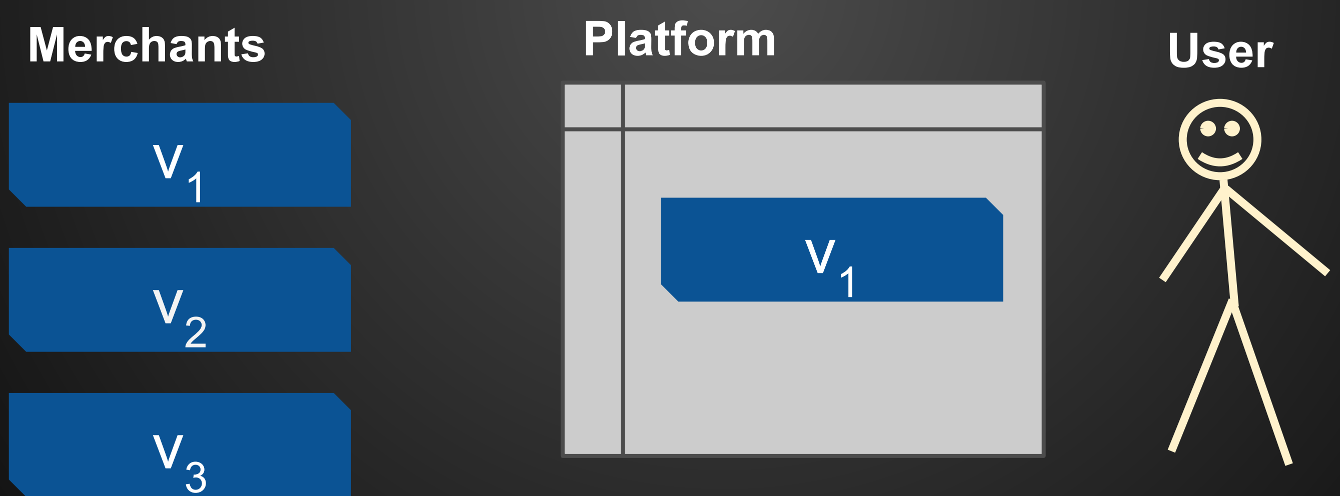

v_i = value for winning

Simple model for daily deals

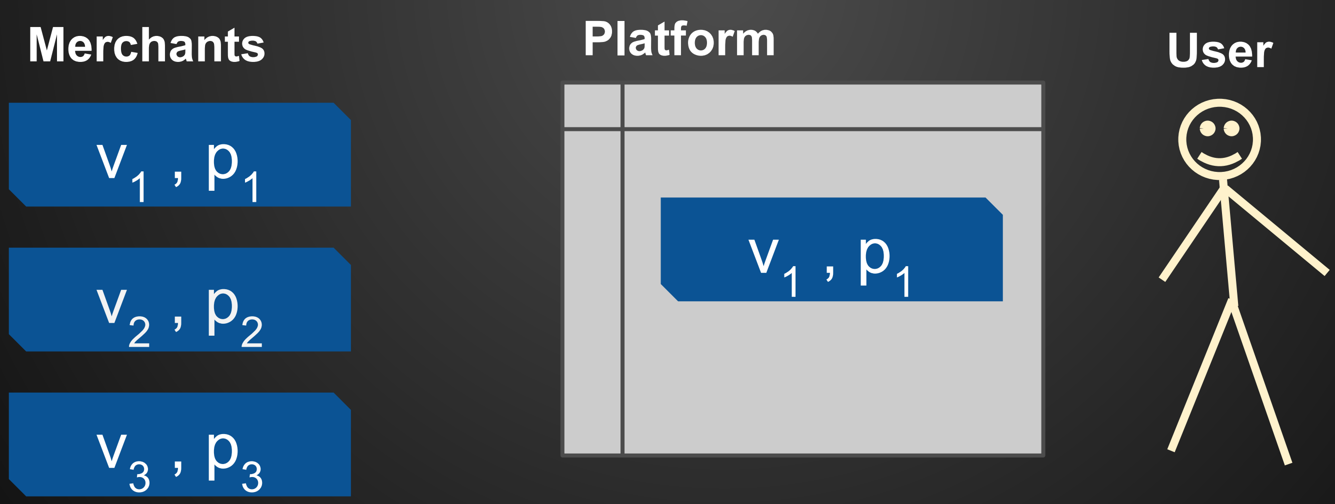

v_i = value for winning
p_i = probability of click

Simple model for daily deals

- User welfare is related to p_i
- First try: require p_i to exceed "quality" threshold

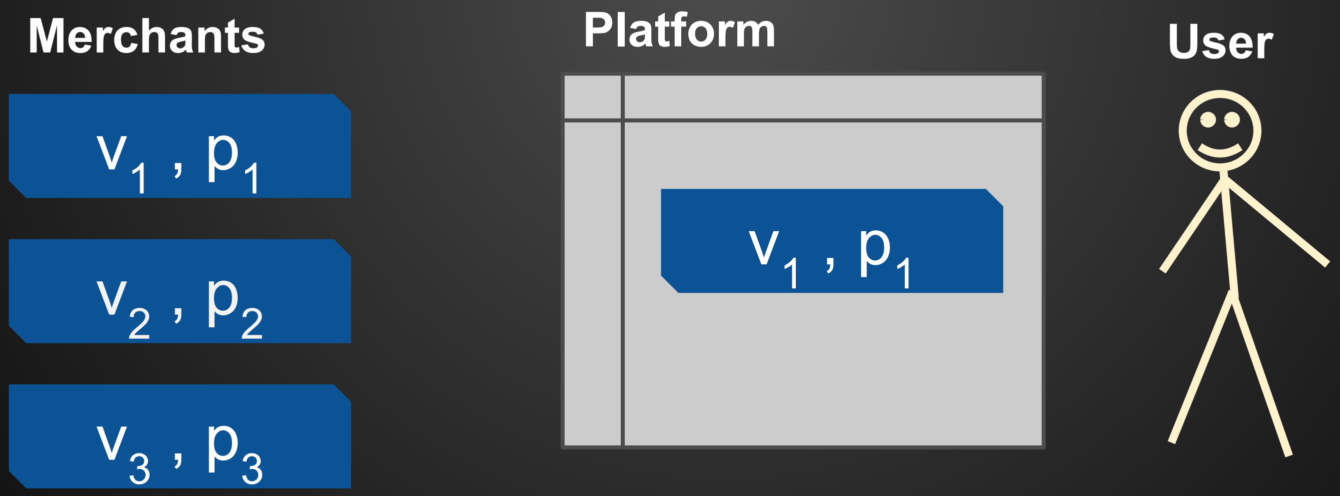

v_i = value for winning
p_i = probability of click

Simple model for daily deals

- User welfare is related to p_i
- First try: require p_i to exceed "quality" threshold
- **Fails!** (cannot even get constant factor of v_i)

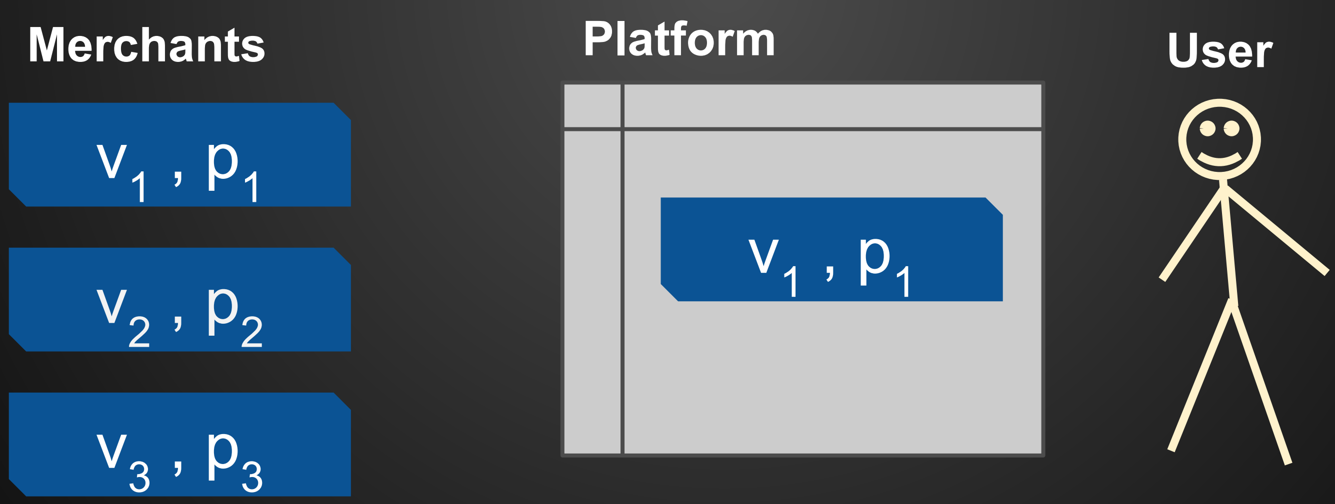

v_i = value for winning
p_i = probability of click

Maximizing total welfare

- User welfare is related to p_i
- Model relationship by a function $g(p_i)$
- Goal: maximize $\boldsymbol{v_i + g(p_i)}$

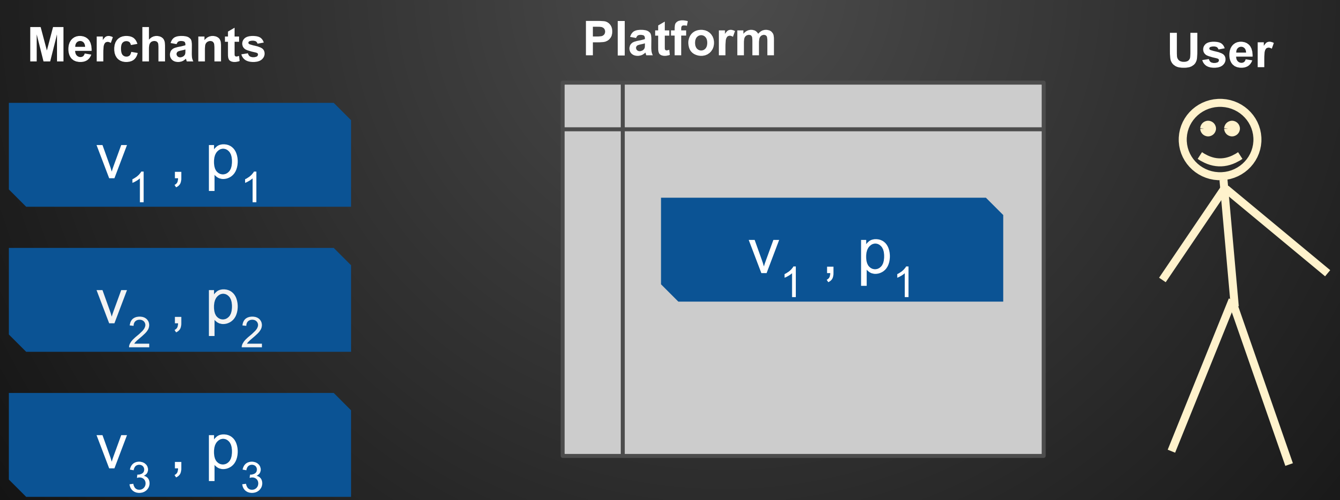

v_i = value for winning
p_i = probability of click

welfare = $g(p_i)$

Q: For what user welfare functions *g* *(p)* can we truthfully max welfare?

Theorem 1. *g(p)* is **convex** ⇔ there exists a deterministic, truthful auction maximizing

$$v_i + g(p_i) \ .$$

Q: For what user welfare functions *g(p)* can we truthfully max welfare?

Theorem 1. *g(p)* is **convex** ⇔ there exists a deterministic, truthful auction maximizing

$$v_i + g(p_i)\ .$$

What does convex mean?
Example: $p = 0$ on first day, $p = 1$ on second day is preferred to $p = 0.5$ on both days

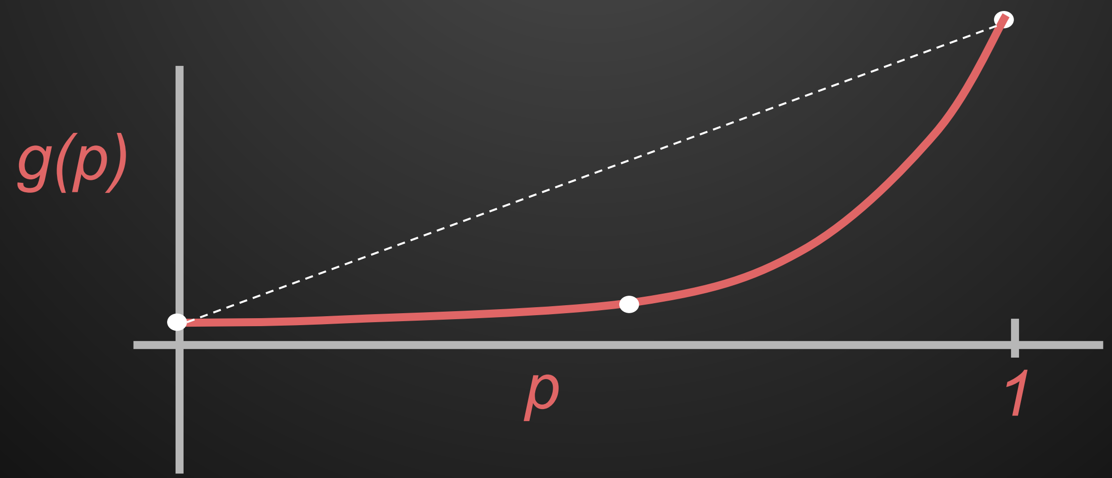

Q: For what user welfare functions *g (p)* can we truthfully max welfare?

Theorem 1. *g(p)* is **convex** ⇔ there exists a deterministic, truthful auction maximizing

$$v_i + g(p_i) .$$

Constructing the auction

Key idea: p_i = **prediction**

Q: For what user welfare functions *g(p)* can we truthfully max welfare?

Theorem 1. *g(p)* is **convex** ⇔ there exists a deterministic, truthful auction maximizing

$$v_i + g(p_i) .$$

Constructing the auction

Key idea: p_i = **prediction**

Scoring rule: Score(prediction, outcome).

Proper: truthful prediction maximizes expected score.

Q: For what user welfare functions *g(p)* can we truthfully max welfare?

Theorem 1. *g(p)* is **convex** ⇔ there exists a deterministic, truthful auction maximizing

$$v_i + g(p_i) .$$

1. Sort by $v_i + g(p_i)$ from highest to lowest.
2. Pick bidder 1.
3. Bidder 1 pays platform: $v_2 + g(p_2)$
4. Platform pays bidder 1: $Score(p_1, outcome)$

Q: For what user welfare functions *g (p)* can we truthfully max welfare?

Theorem 1. *g(p)* is **convex** ⇔ there exists a deterministic, truthful auction maximizing

$$v_i + g(p_i) \, .$$

Lemma (Savage '71). For all convex *g(p)*, there exists a proper scoring rule with expected score *g(p)* for truthfully reporting *p*.

Q: For what user welfare functions *g (p)* can we truthfully max welfare?

Theorem 1. $g(p)$ is **convex** ⇔ there exists a deterministic, truthful auction maximizing

$$v_i + g(p_i) .$$

1. Sort by $v_i + g(p_i)$ from highest to lowest.
2. Pick bidder 1.
3. Bidder 1 pays platform: $v_2 + g(p_2)$
4. Platform pays bidder 1: $Score(p_1, outcome)$

E[utility for winning] = $v_1 + g(p_1) - (v_2 + g(p_2))$

Outline

1. Really simple model for daily deals, results
2. Really general model, characterization
3. Applications and conclusion

Takeaways from simple model

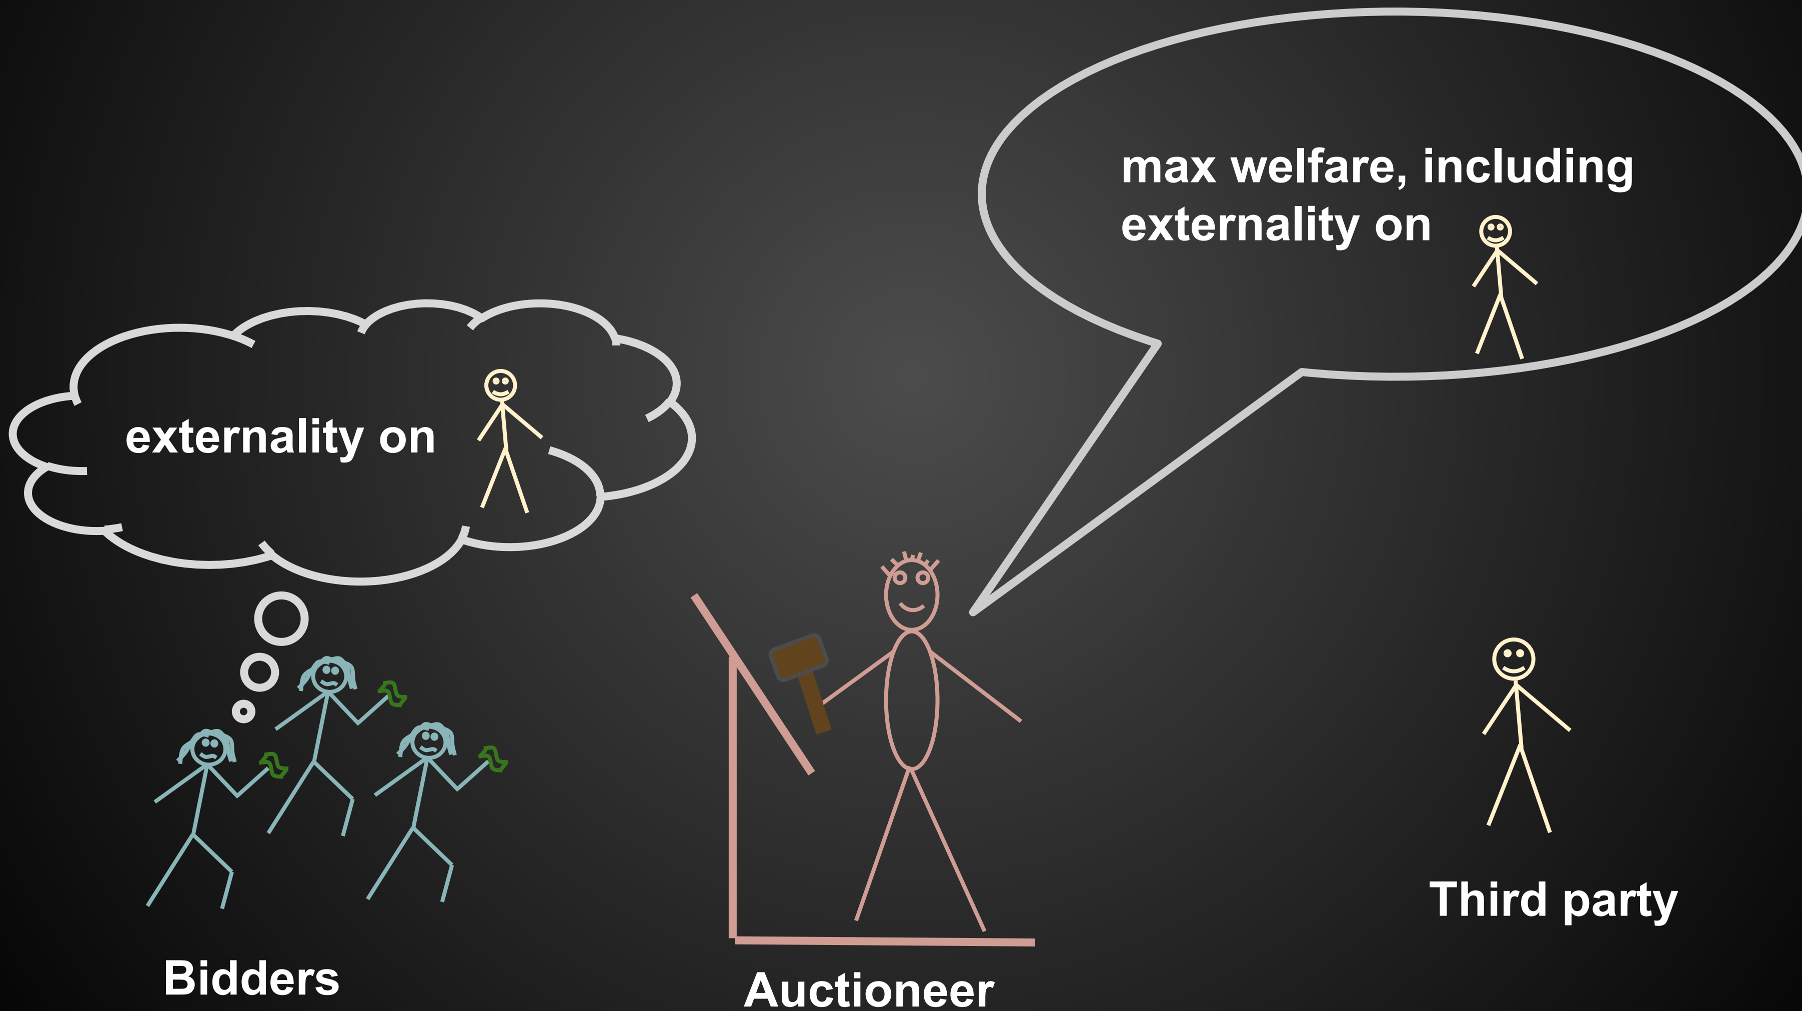

Takeaways from simple model

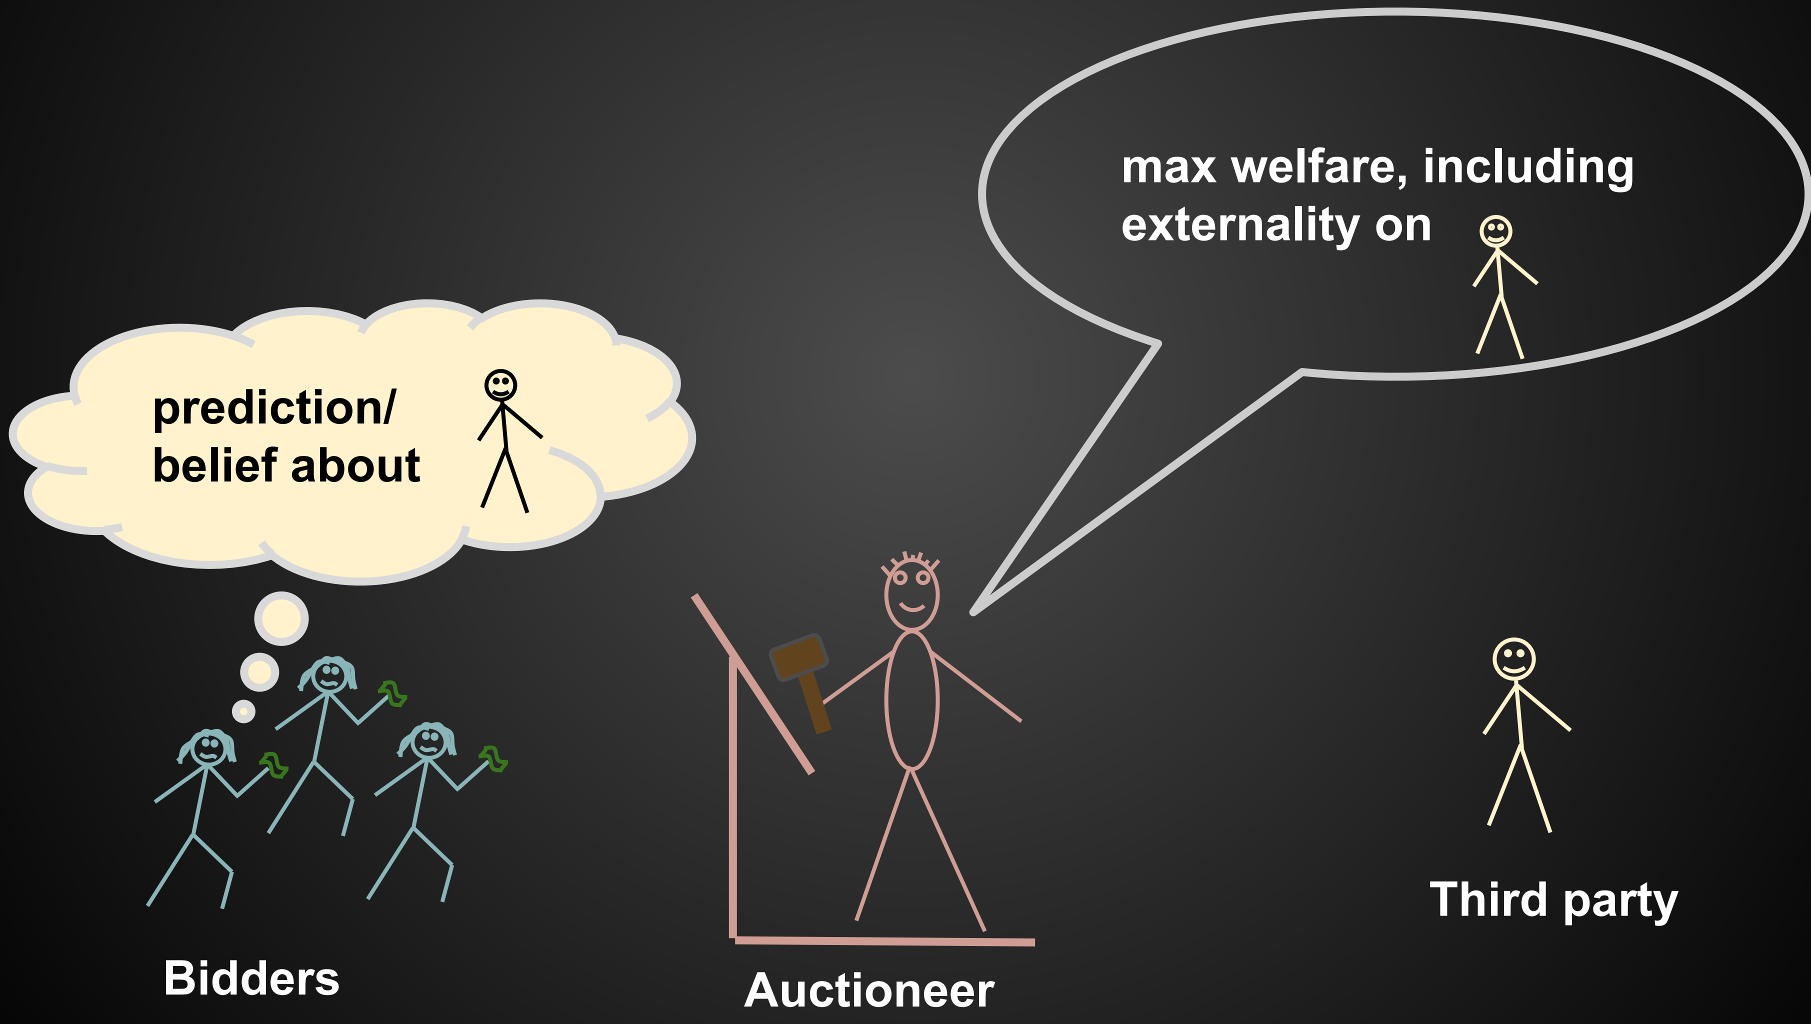

Takeaways from simple model

Implementable ⇔ externality is convex function of prediction

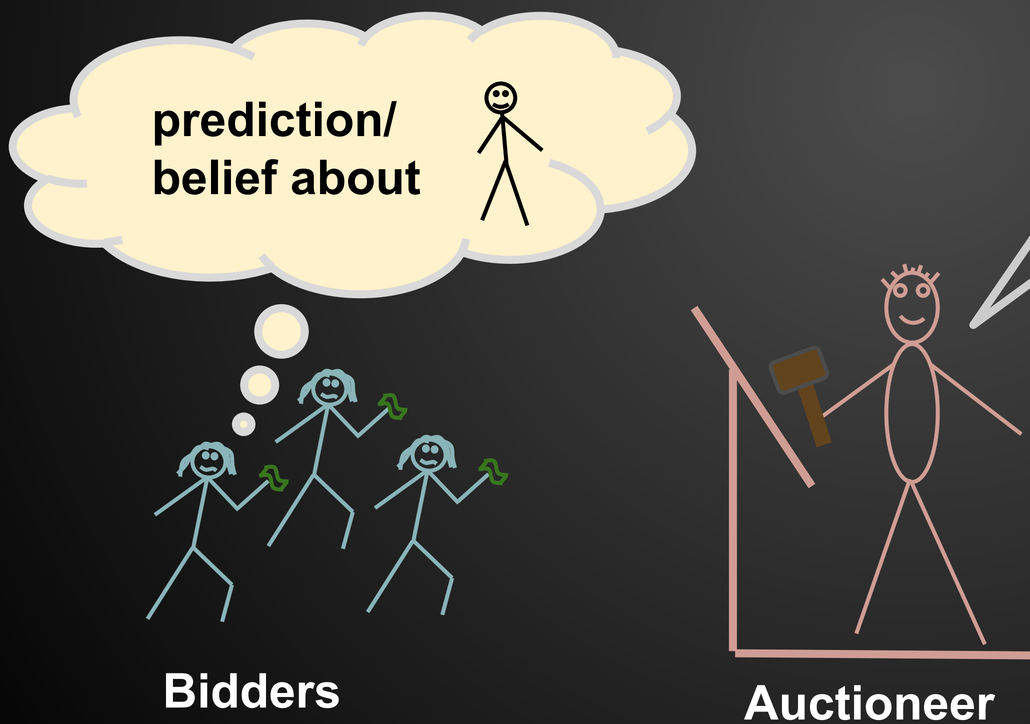

Takeaways from simple model

Auction:
2nd price and “decomposed” proper scoring rule

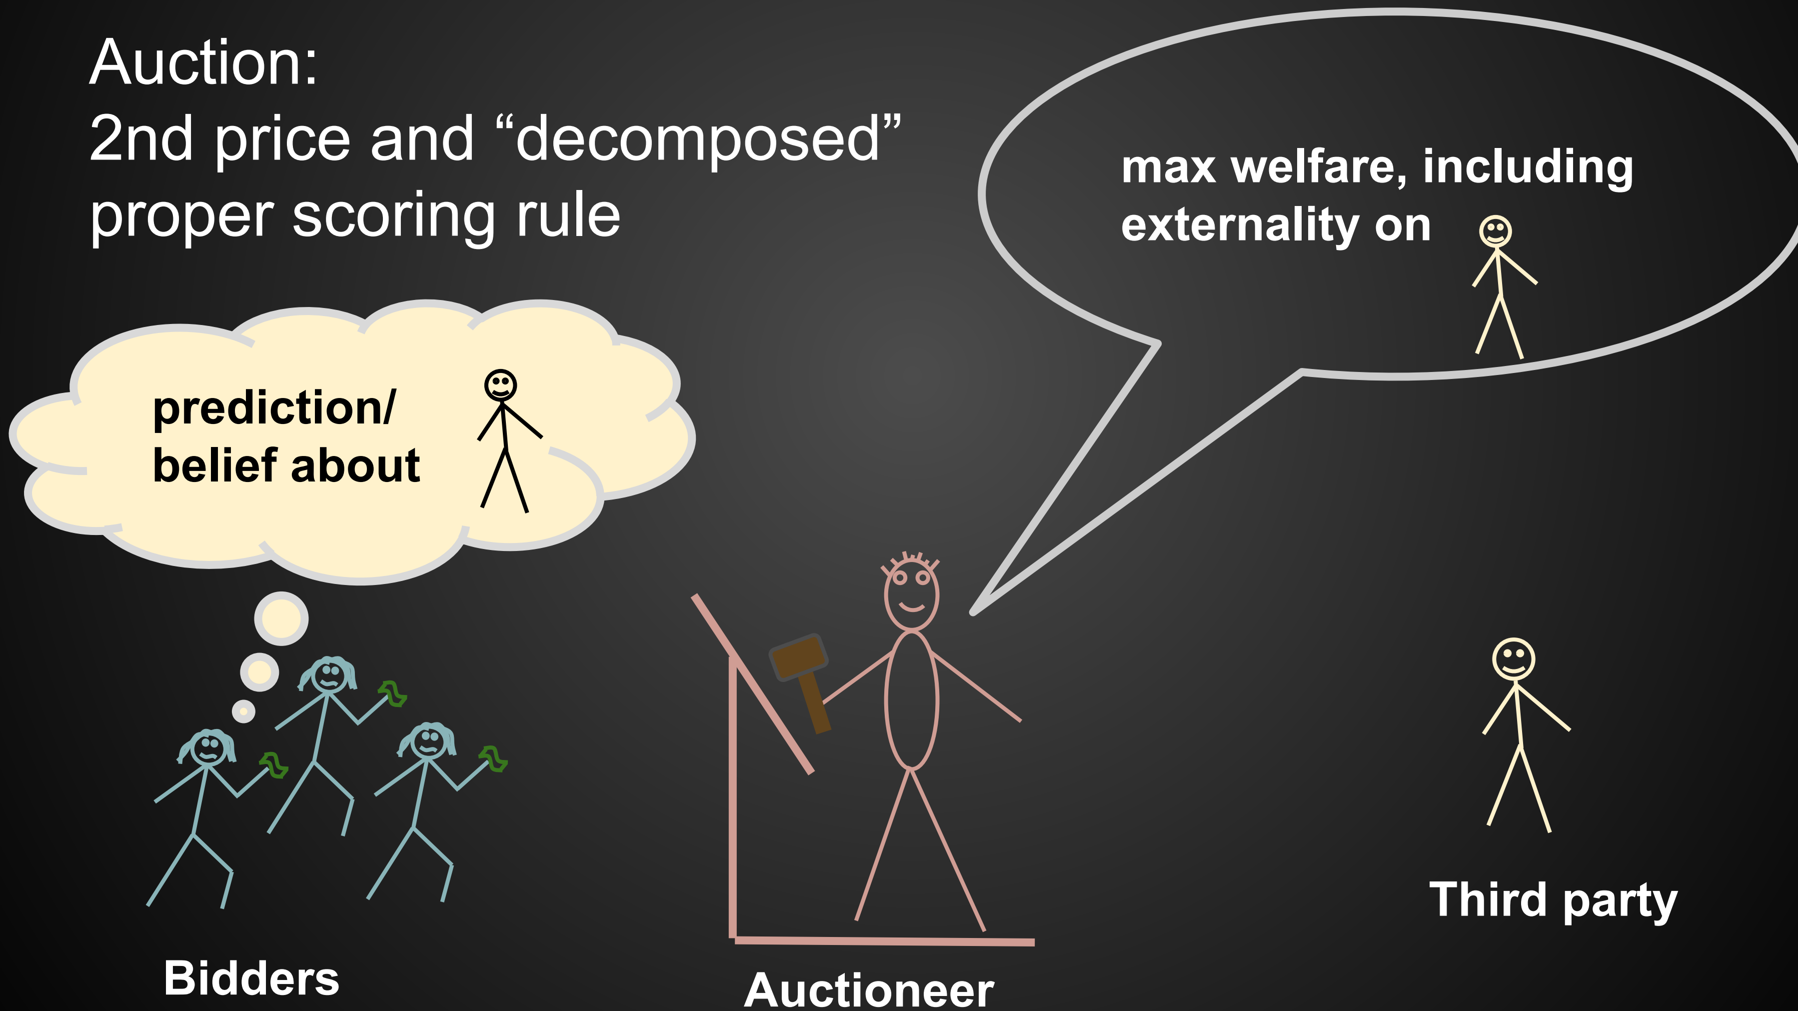

“Really General Model”

Example: “full” daily deals.

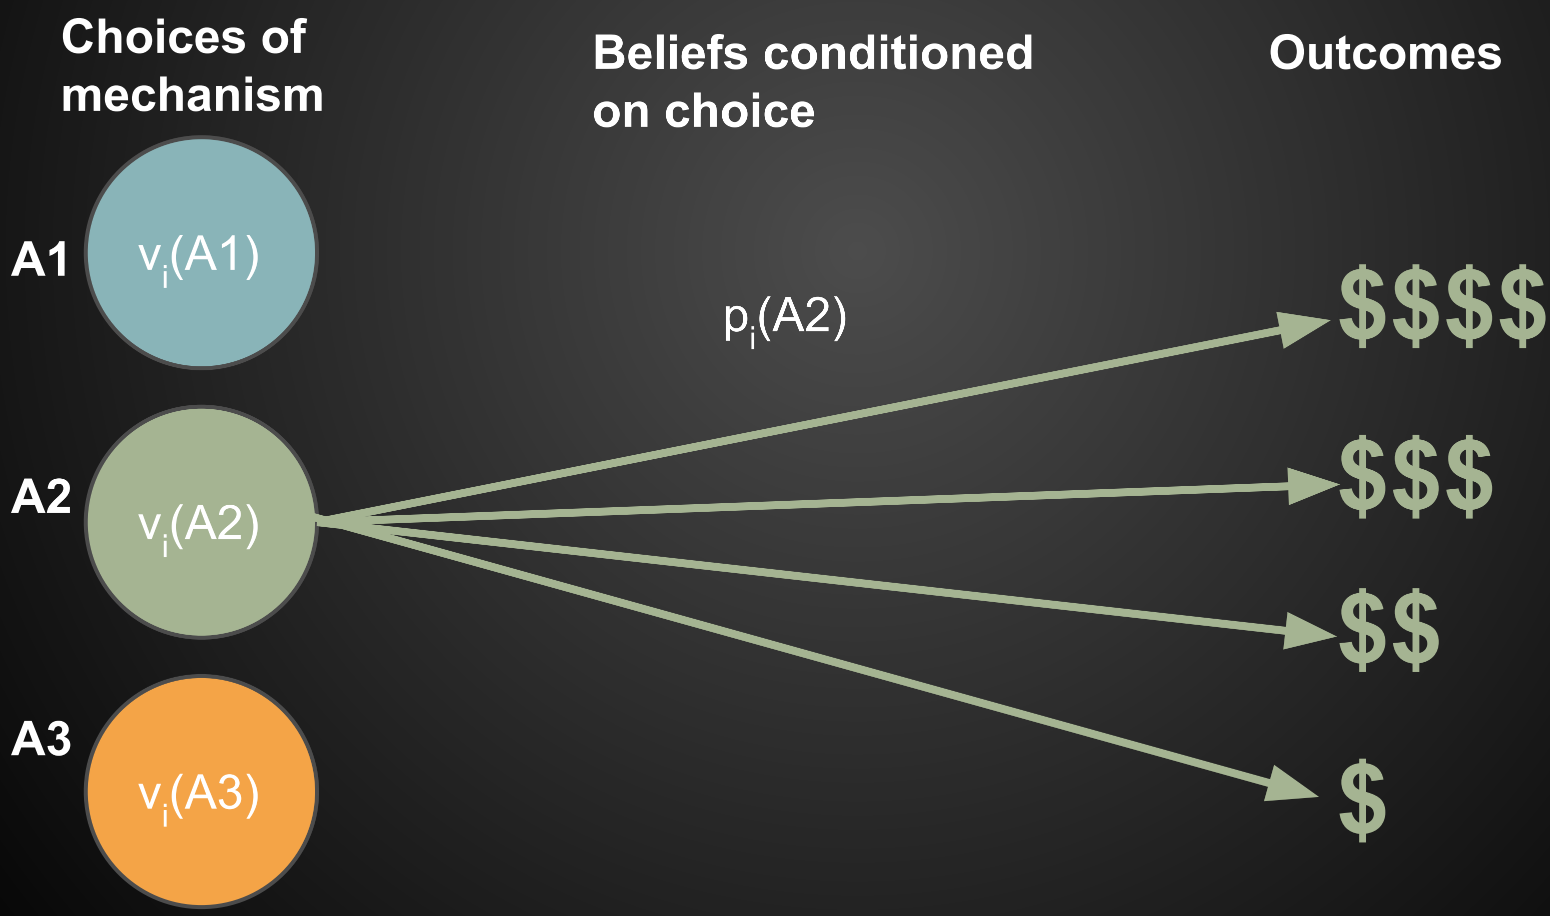

“Really General Model”

Example: “full” daily deals.

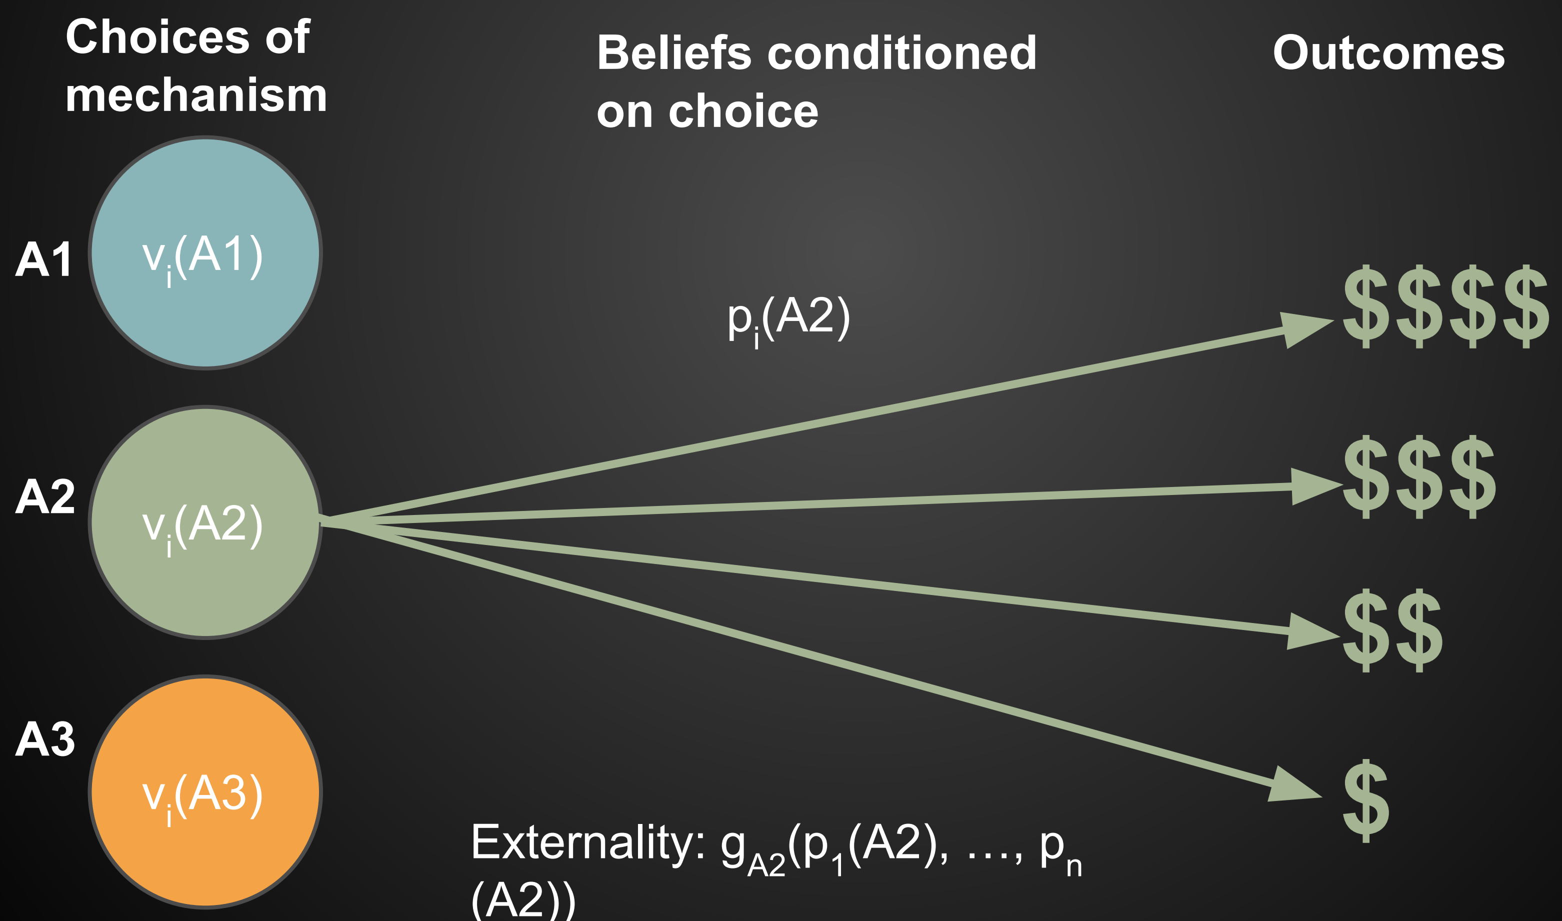

Q: For what externality functions *g* can we truthfully max welfare?

Theorem 2.
$g_A(p_1(A),\ldots)$ are **convex** in each argument ⇔ we can maximize welfare = $\mathbf{g_A(p_1(A),\ldots) + sum_i\ v_i(A)}$.

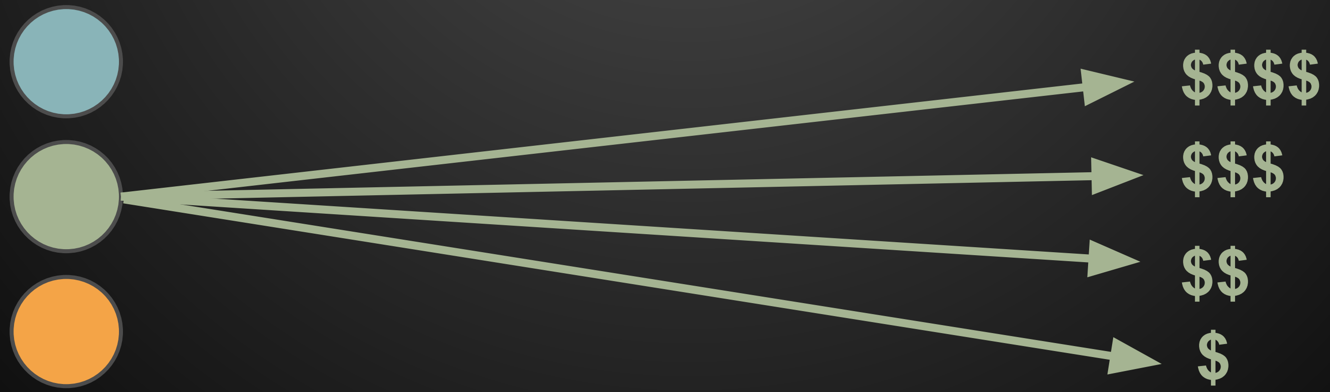

Q: For what externality functions *g* can we truthfully max welfare?

Theorem 2.
$g_A(p_1(A),...)$ are **convex** in each argument ⇔ we can maximize welfare = $\mathbf{g_A(p_1(A),...) + sum_i\ v_i(A)}$.

Auction: VCG and carefully constructed scoring rules.

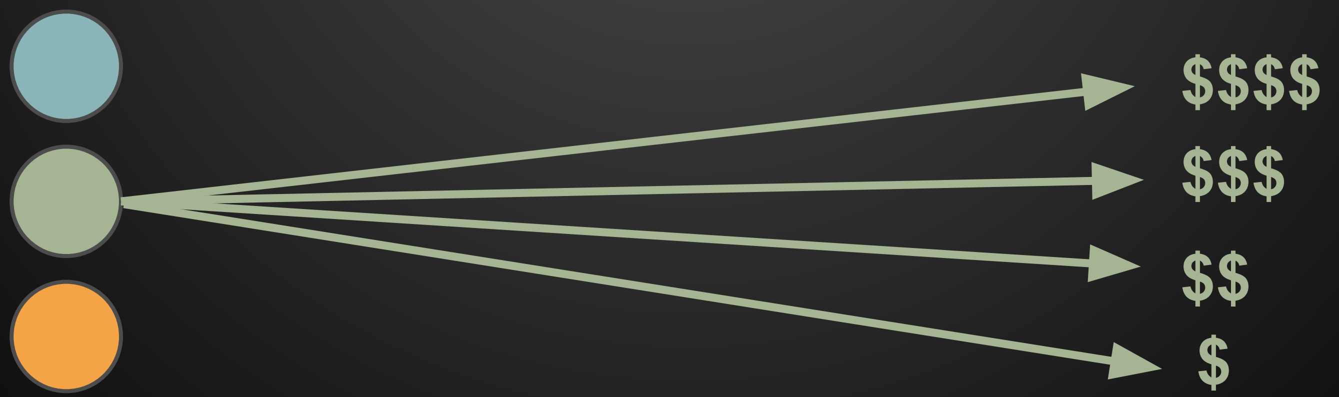

Outline

1. Really simple model for daily deals, results
2. Really general model, characterization
3. Applications and conclusion

Application of Characterization: Network Problems

- Each edge has:
  - cost v_i
  - stochastic delay ~ p_i
- Utility of traveler: $g(p_1, \ldots, p_m)$ for path $1 \ldots m$
- Goal: maximize total welfare

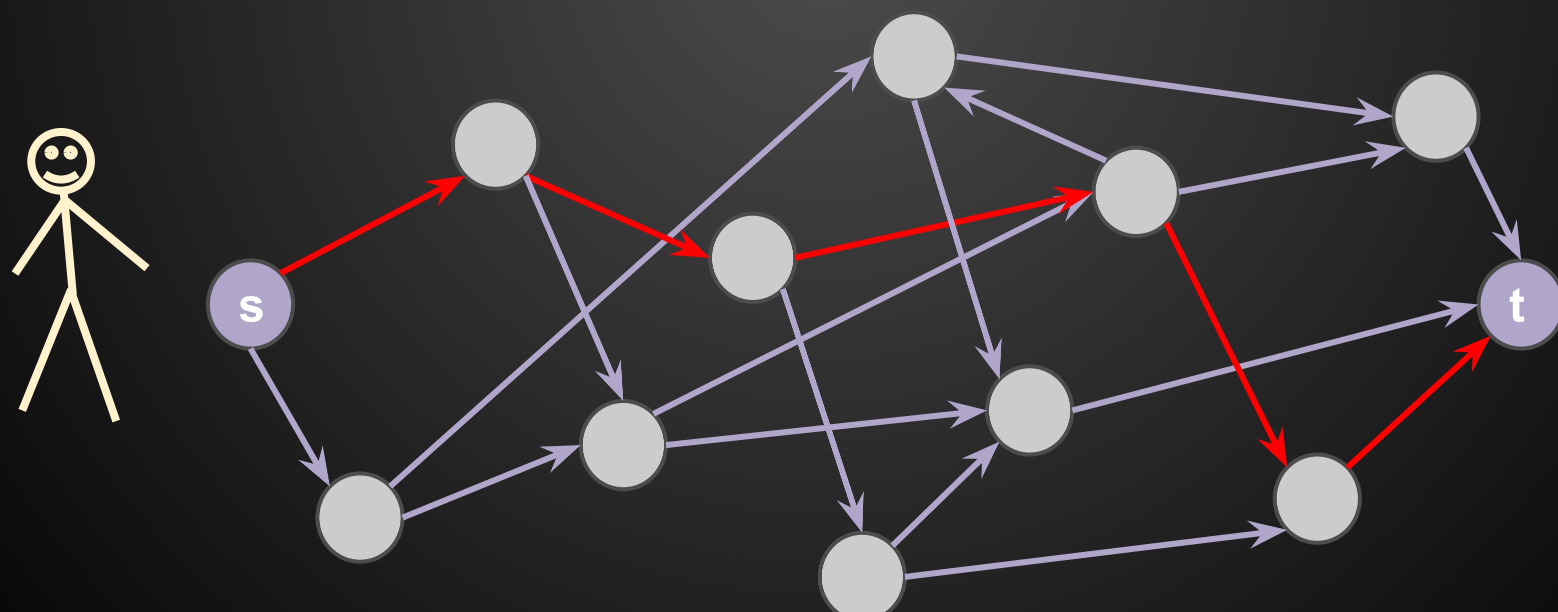

General takeaways

- Welfare includes **externality** on 
- … depending on private **predictions** of bidders
- Implementable ⇔ externality is convex function of prediction
- Auction = VCG + “decomposed” scoring rules

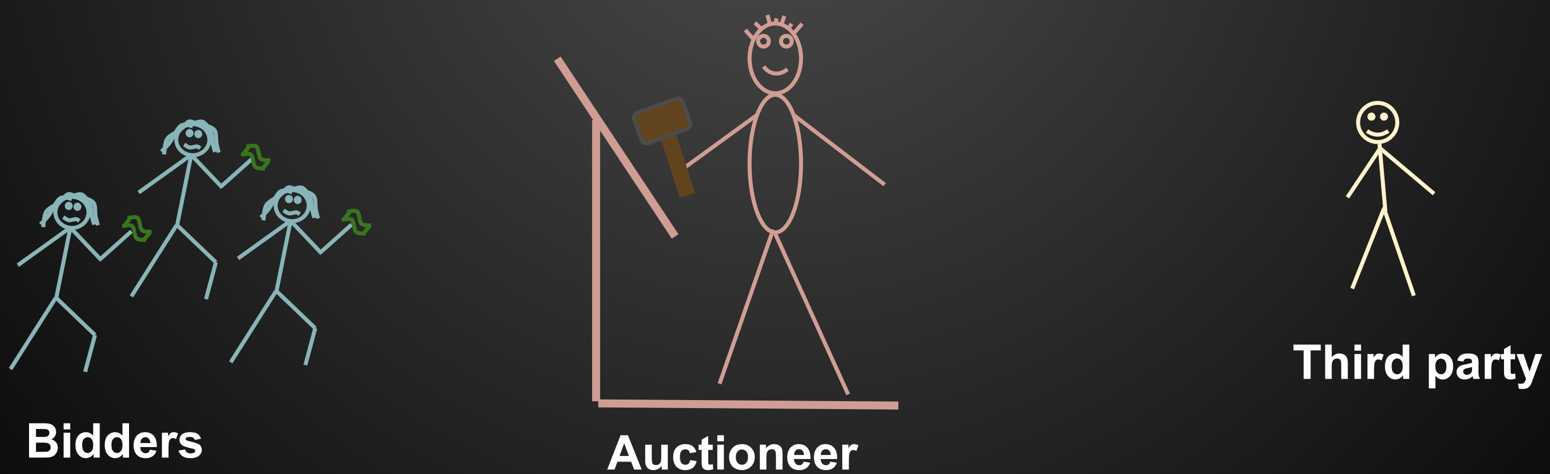

Future work

- Practicality
- Assumptions to avoid negative results
- Applications
- Revenue maximization

- Explore: convexity, implementable allocation functions, and implementable objective functions. c.f. Frongillo and Kash, *General Truthfulness Characterizations via Convex Analysis*

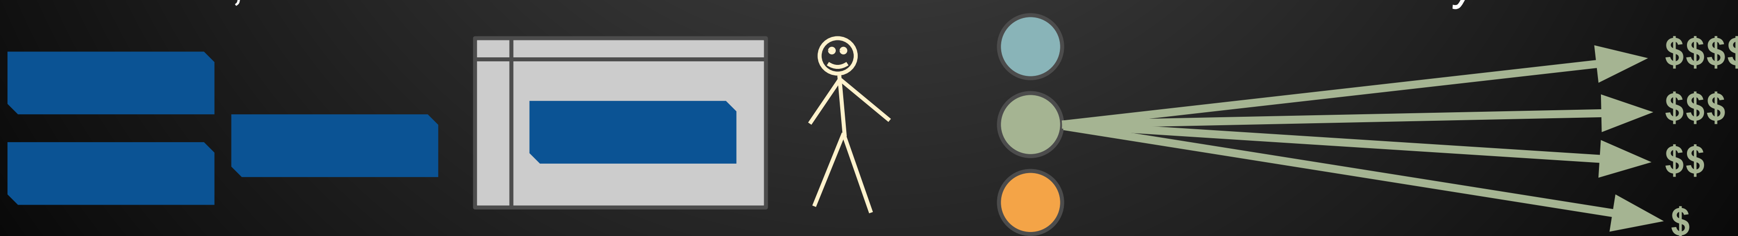

Extension: Principal-agent problems

- Each worker has a set of efforts, each with:
  - cost
  - stochastic quality
- Externality: observed quality of work
- Goal: maximize total welfare