Designing Markets for Daily Deals

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

WINE 2013

Motivation: Daily Deals

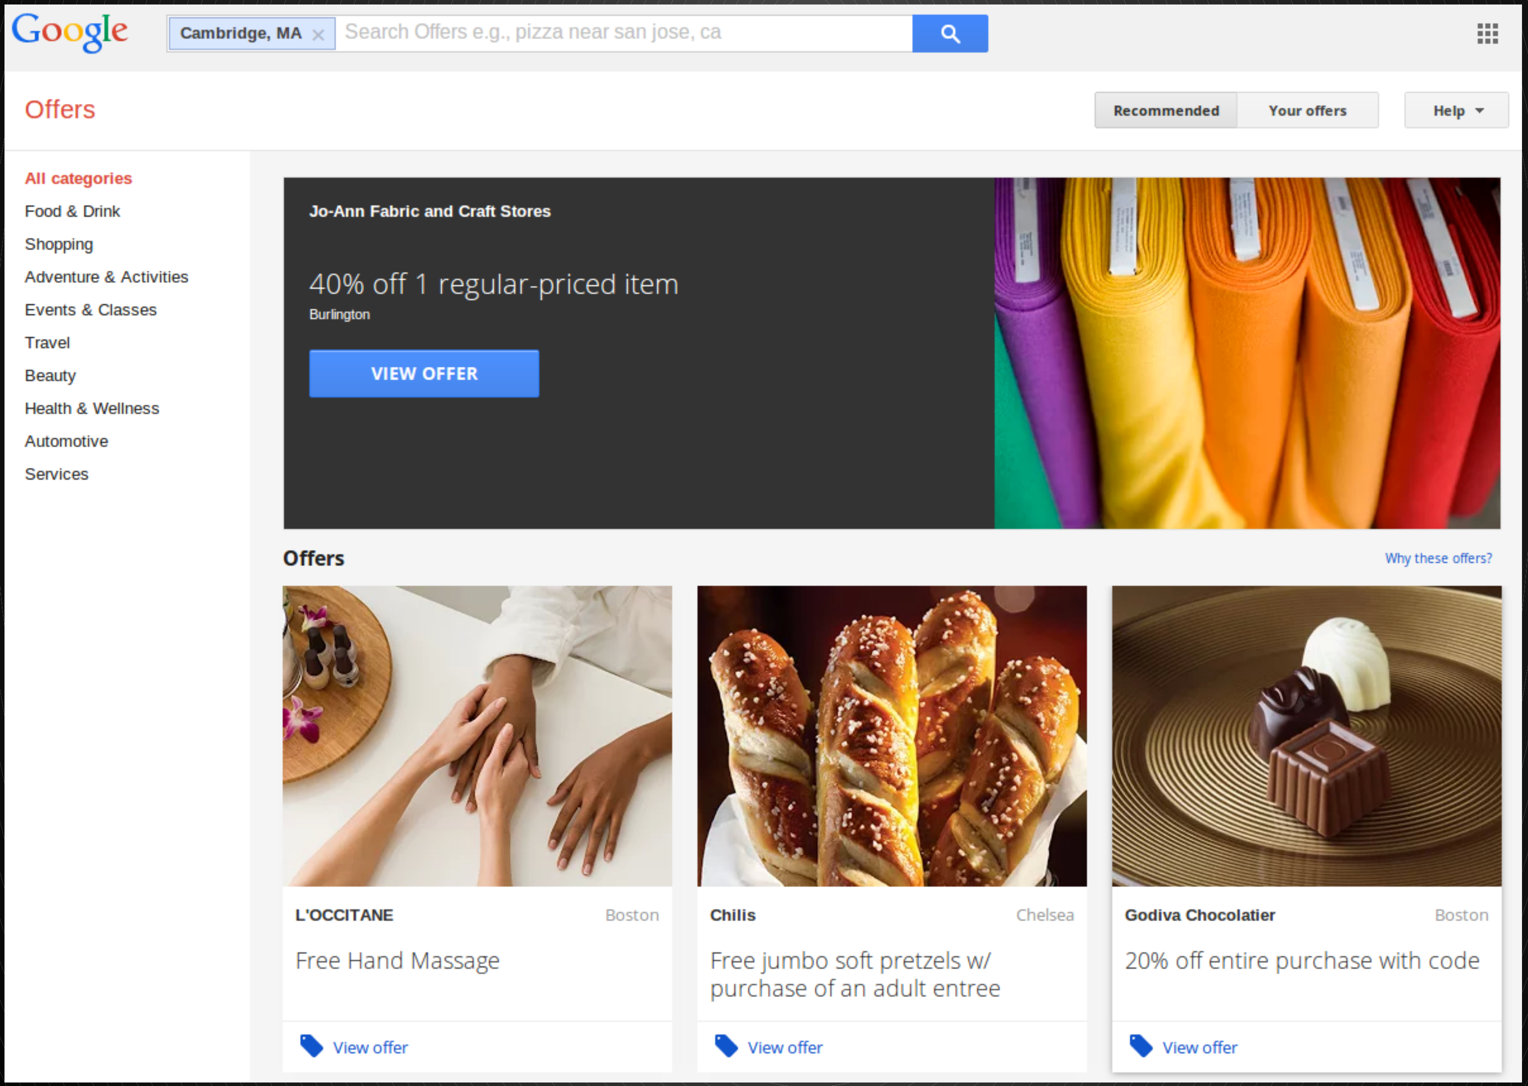

Problem statement

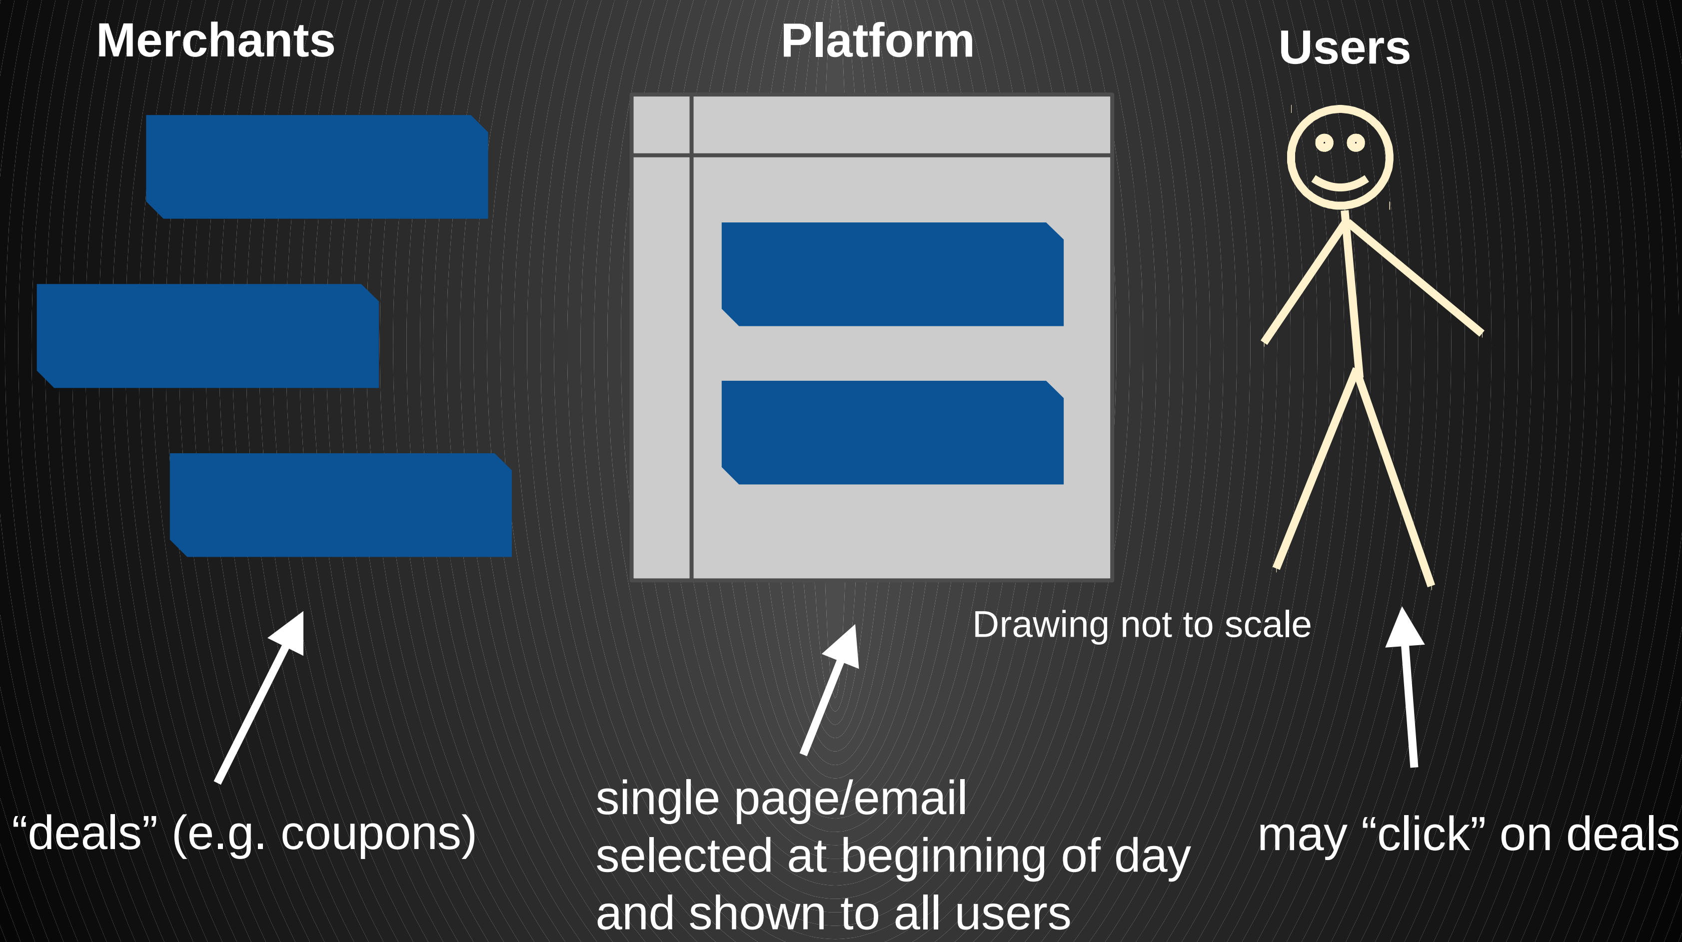

Problem statement

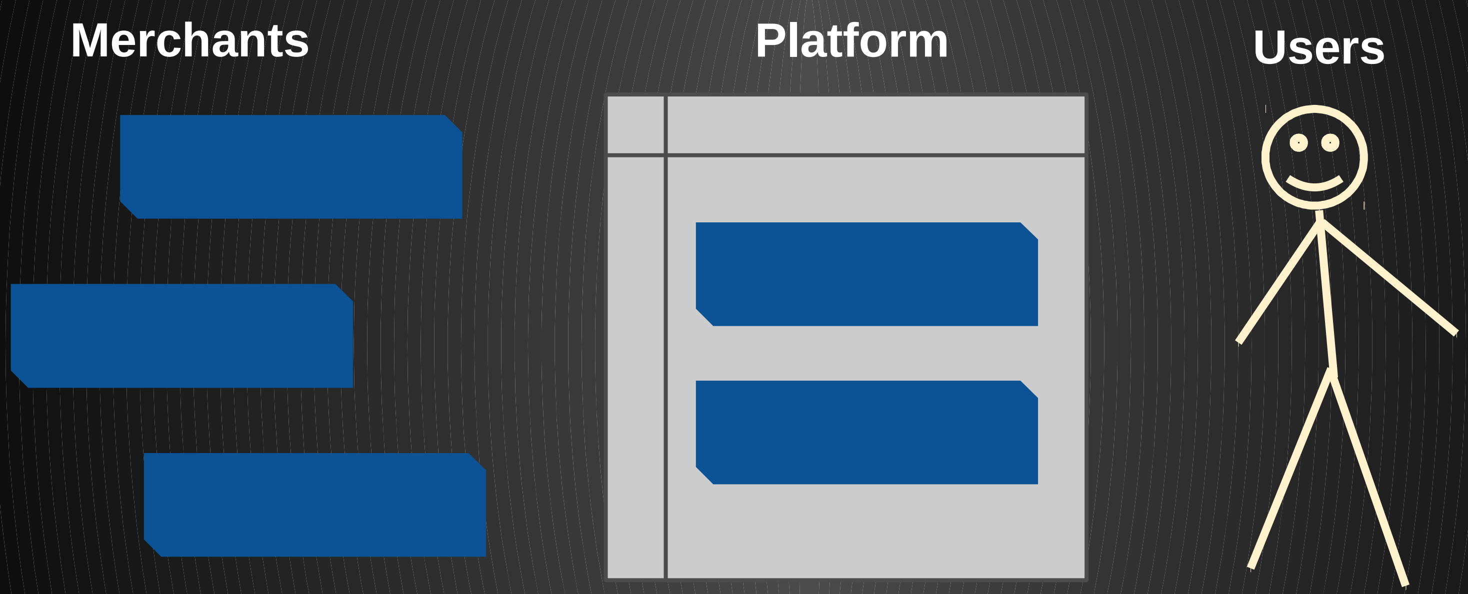

Drawing not to scale

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

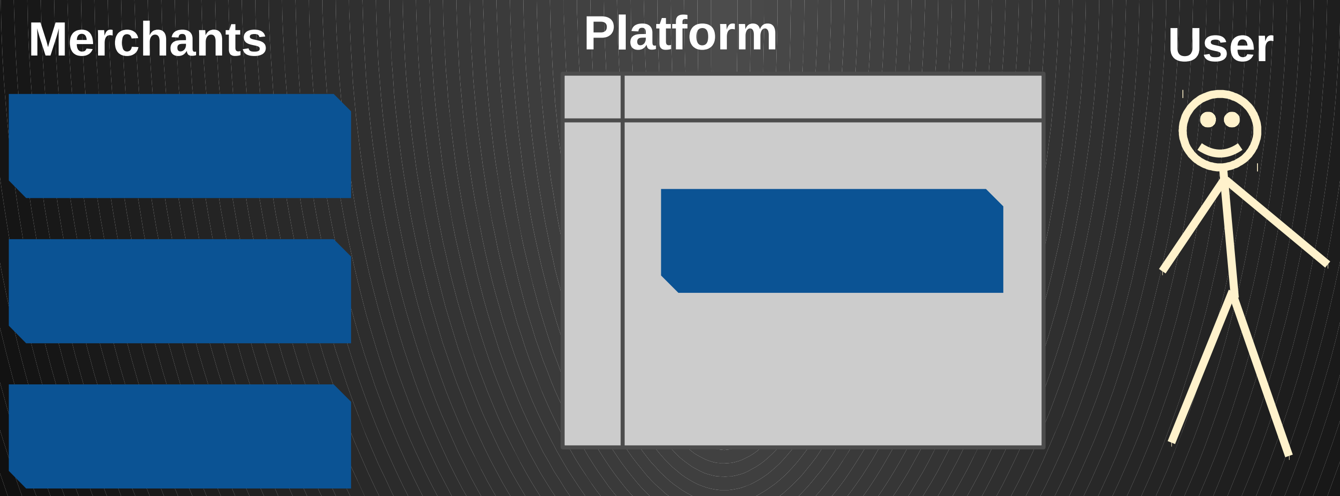

Prologue: Standard auction setting

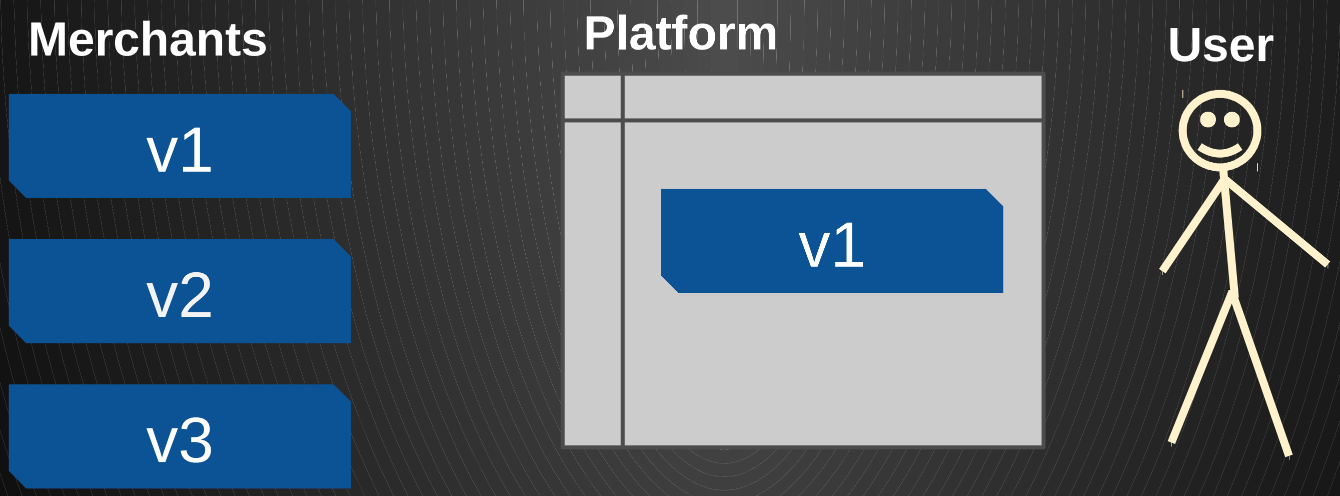

vi = value for winning

Simple model for daily deals

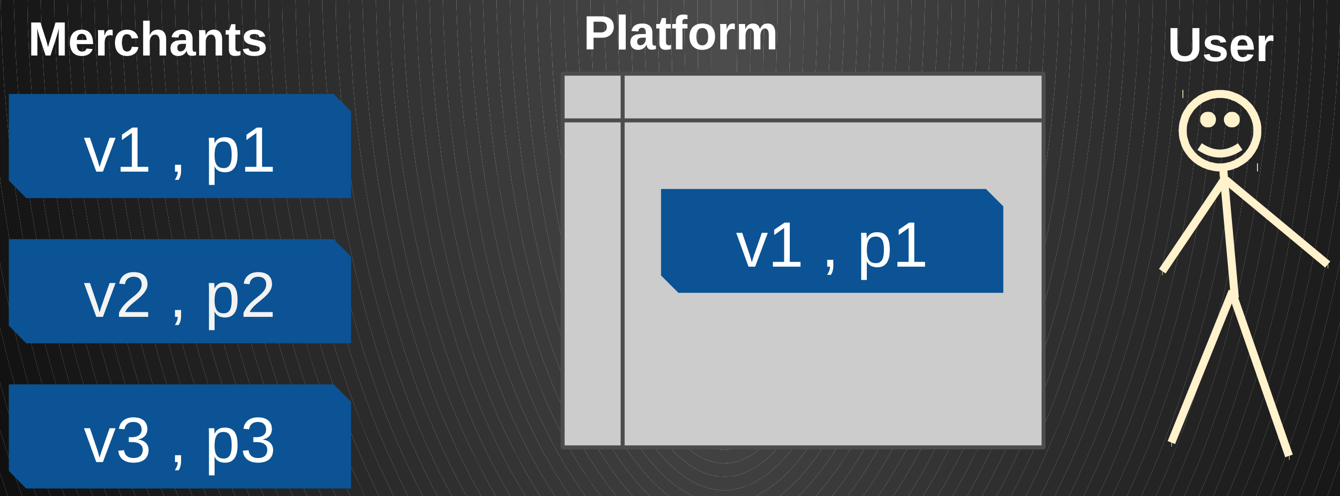

vi = value for winning
pi = probability of click

Simple model for daily deals

- User welfare is related to *pi*
- First try: require *pi* to exceed “quality” threshold

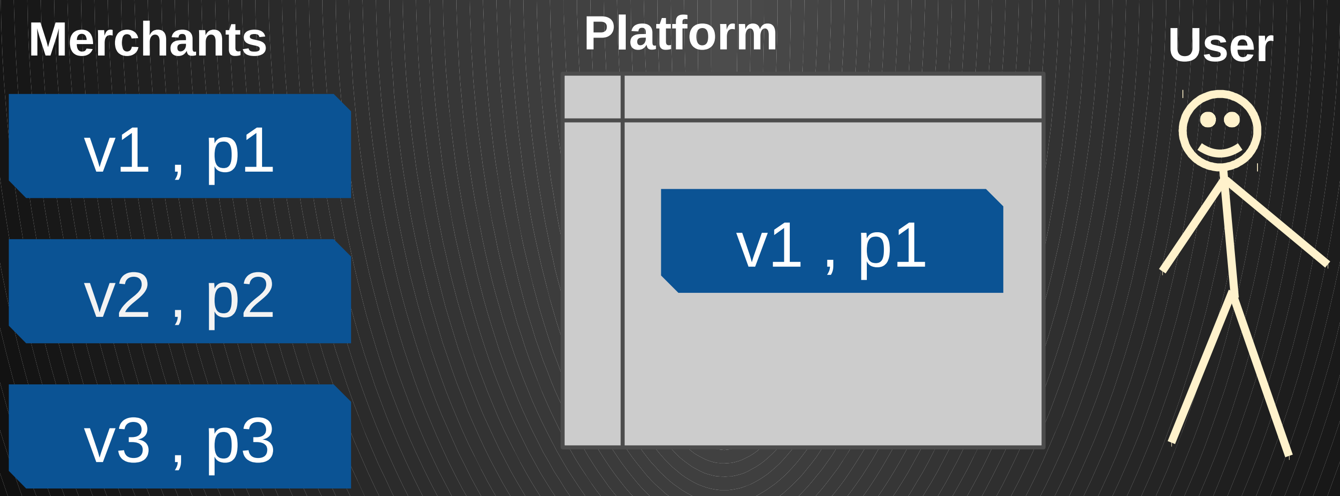

vi = value for winning
pi = probability of click

Simple model for daily deals

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

vi = value for winning
pi = probability of click

Maximizing total welfare

- User welfare is related to *pi*
- Model relationship by a function *g(pi)*
- Goal: maximize ***vi + g(pi)***

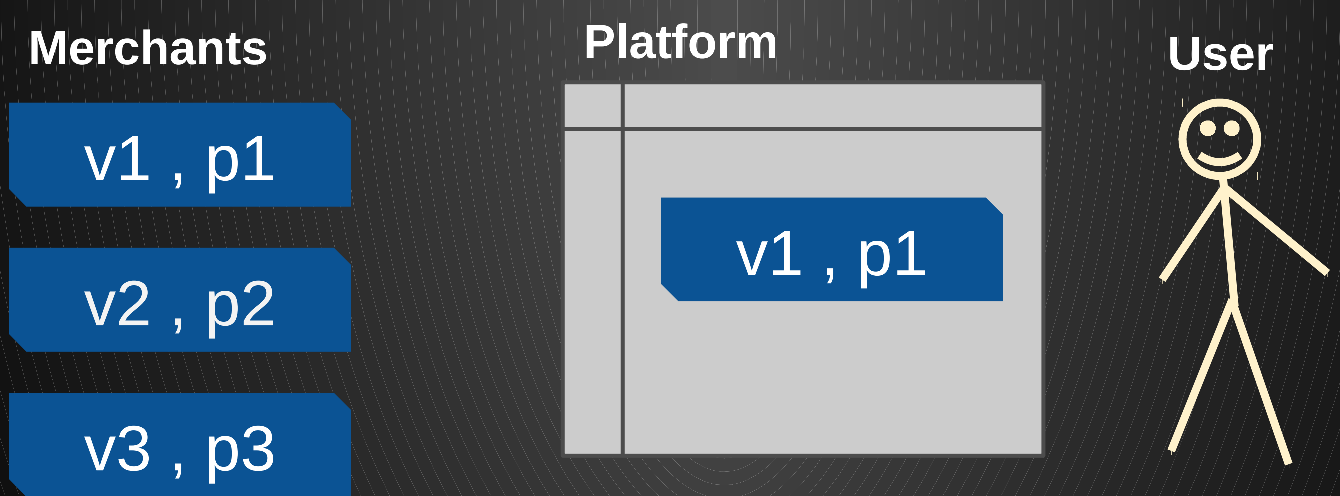

vi = value for winning
pi = probability of click

welfare = *g(pi)*

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

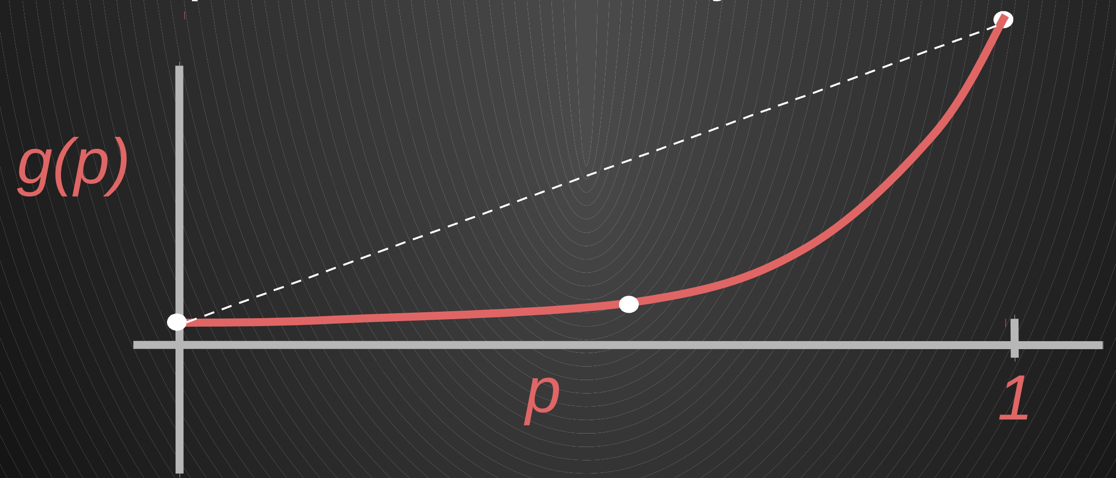

Q: For what user welfare functions *g(p)* can we truthfully max welfare?

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] = v1 + g(p1) - (v2 + g(p2))

Outline

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

Takeaways from simple model

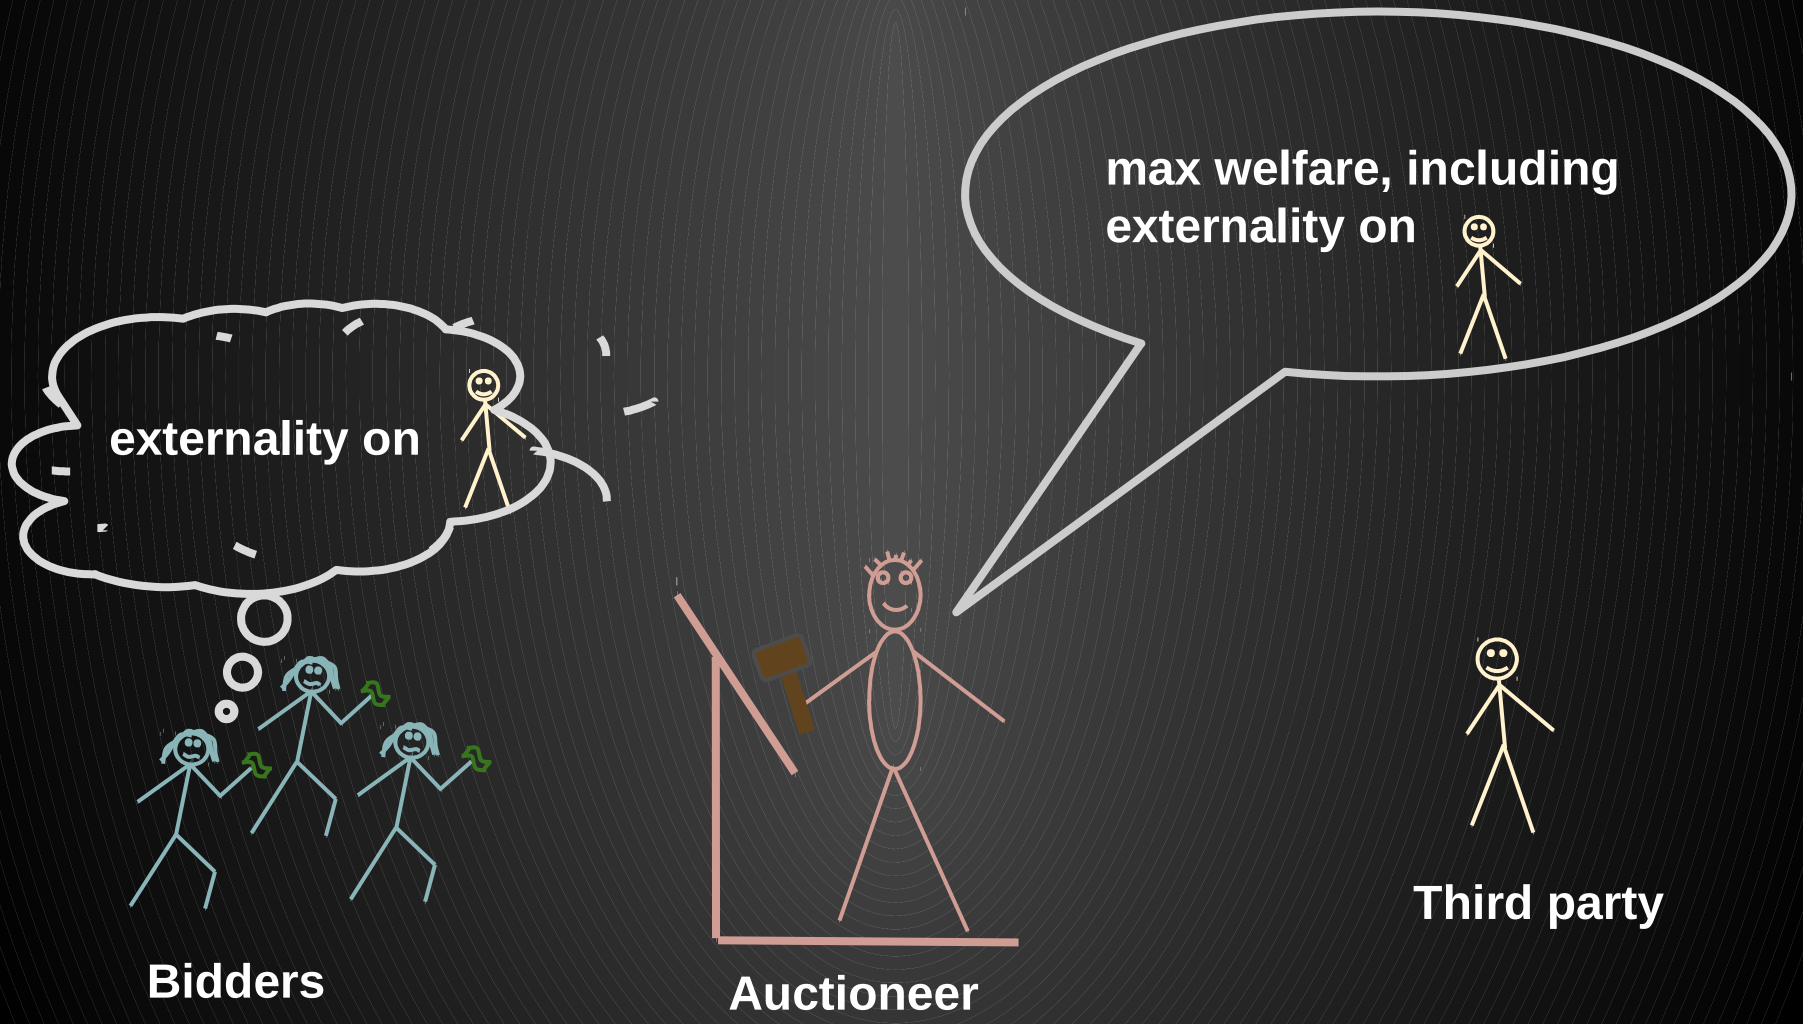

Takeaways from simple model

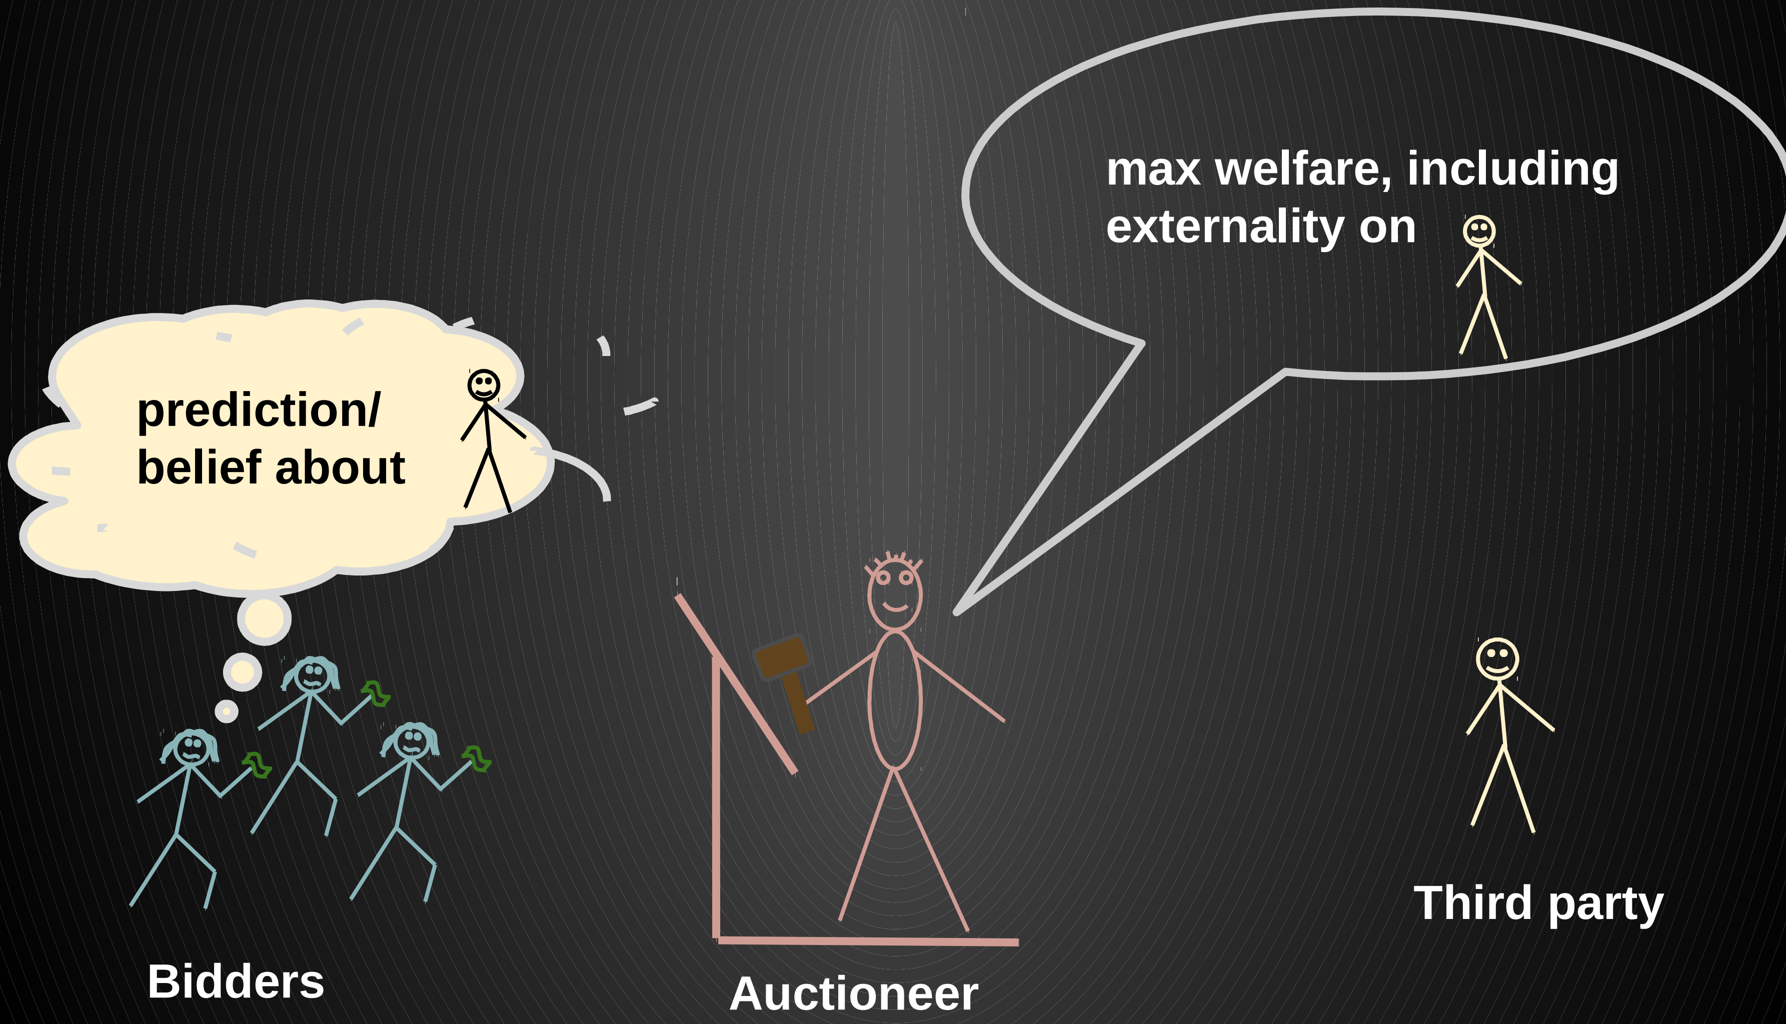

Takeaways from simple model

Implementable ⇔ externality is convex function of prediction

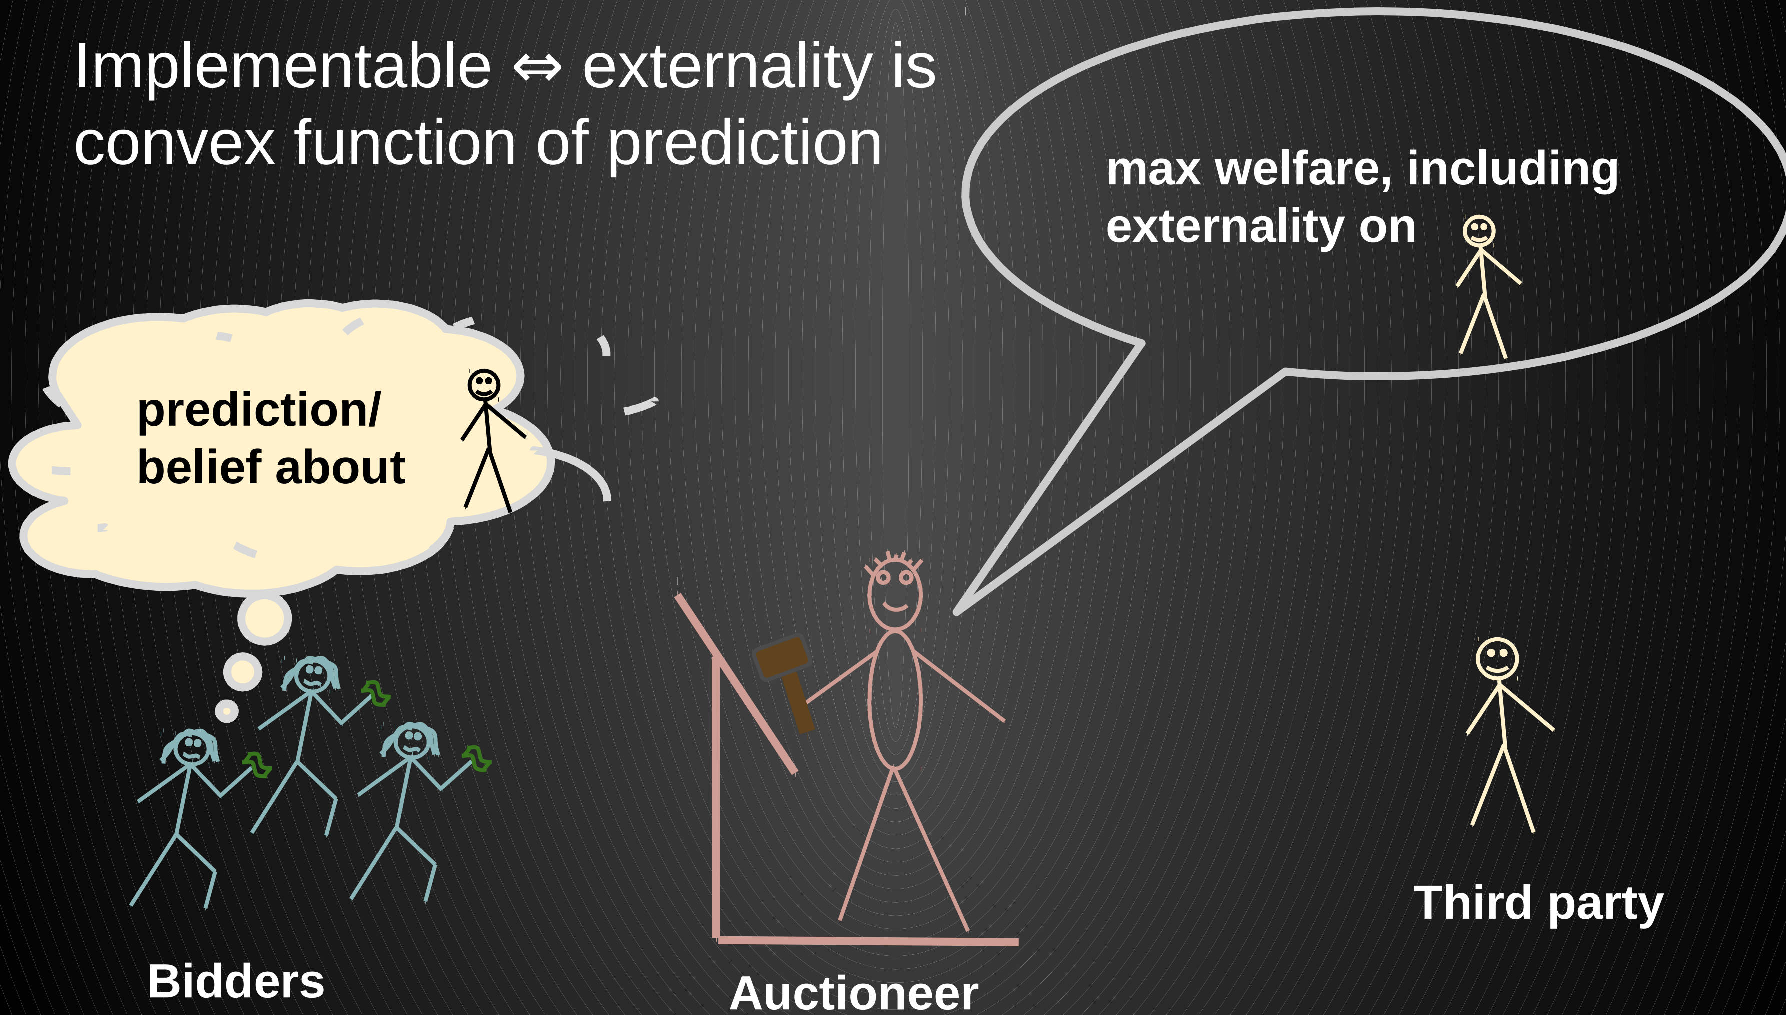

Takeaways from simple model

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

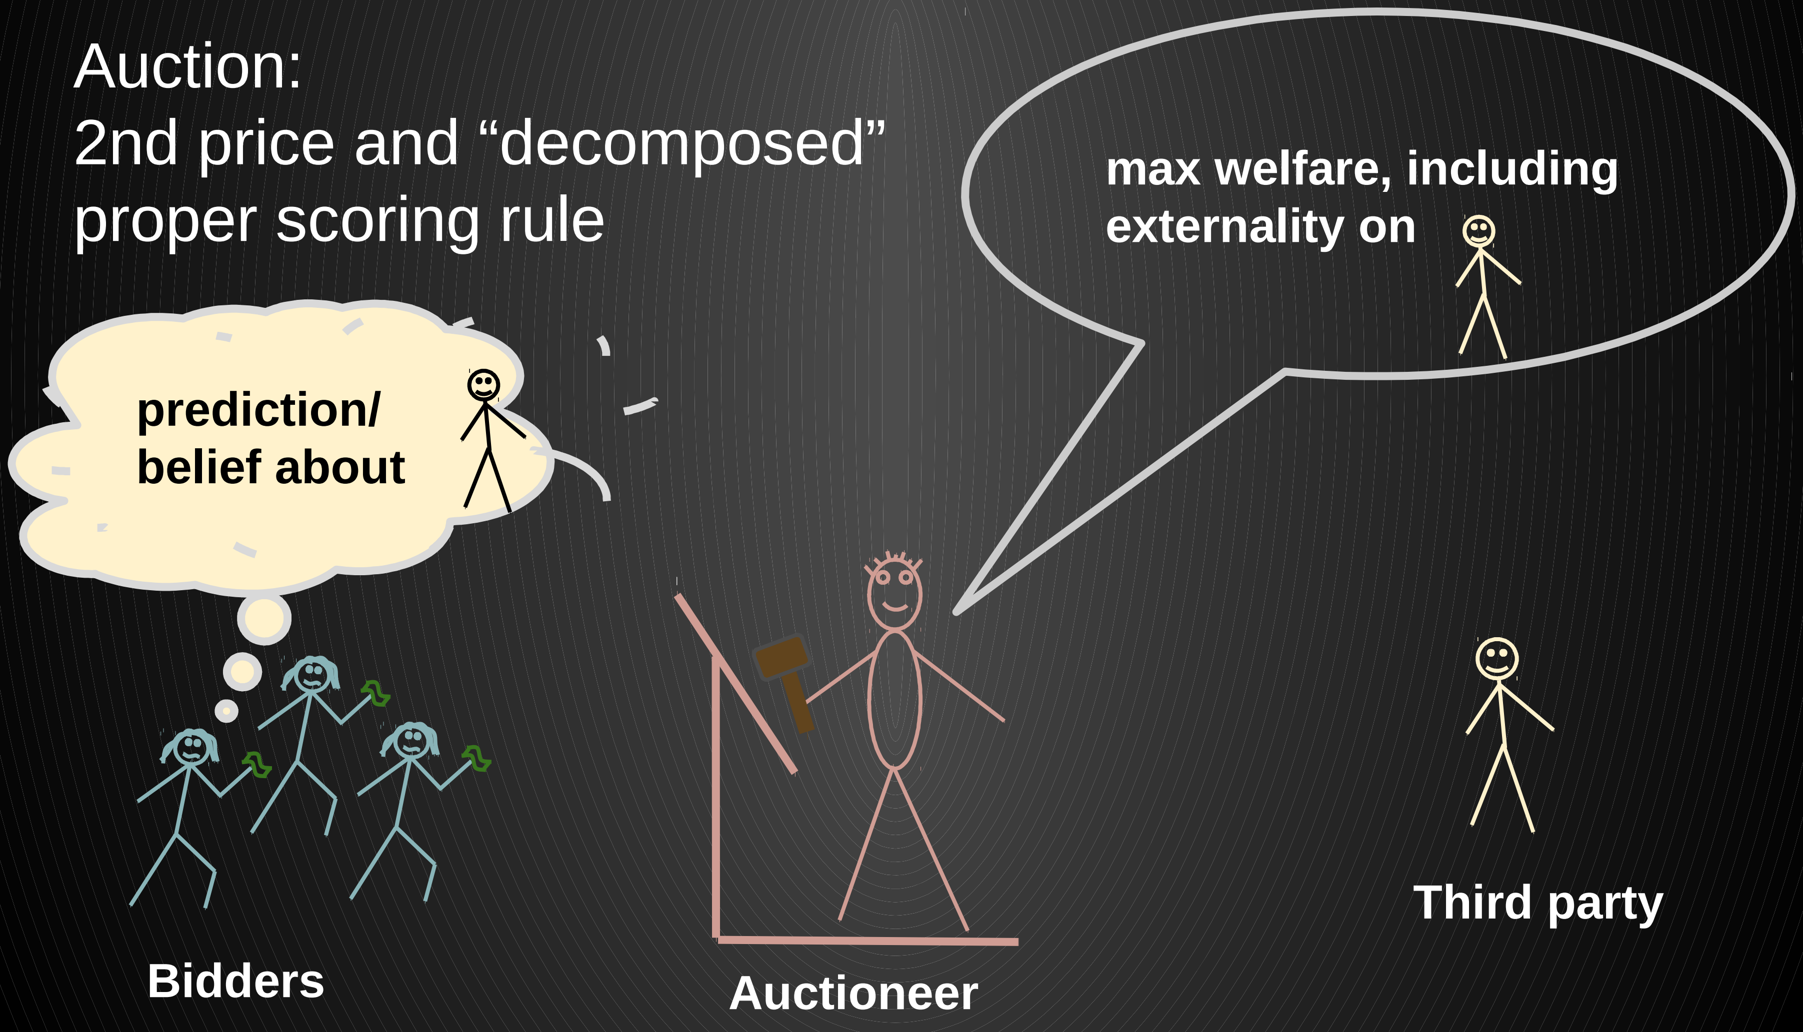

“Really General Model”

Example: “full” daily deals.

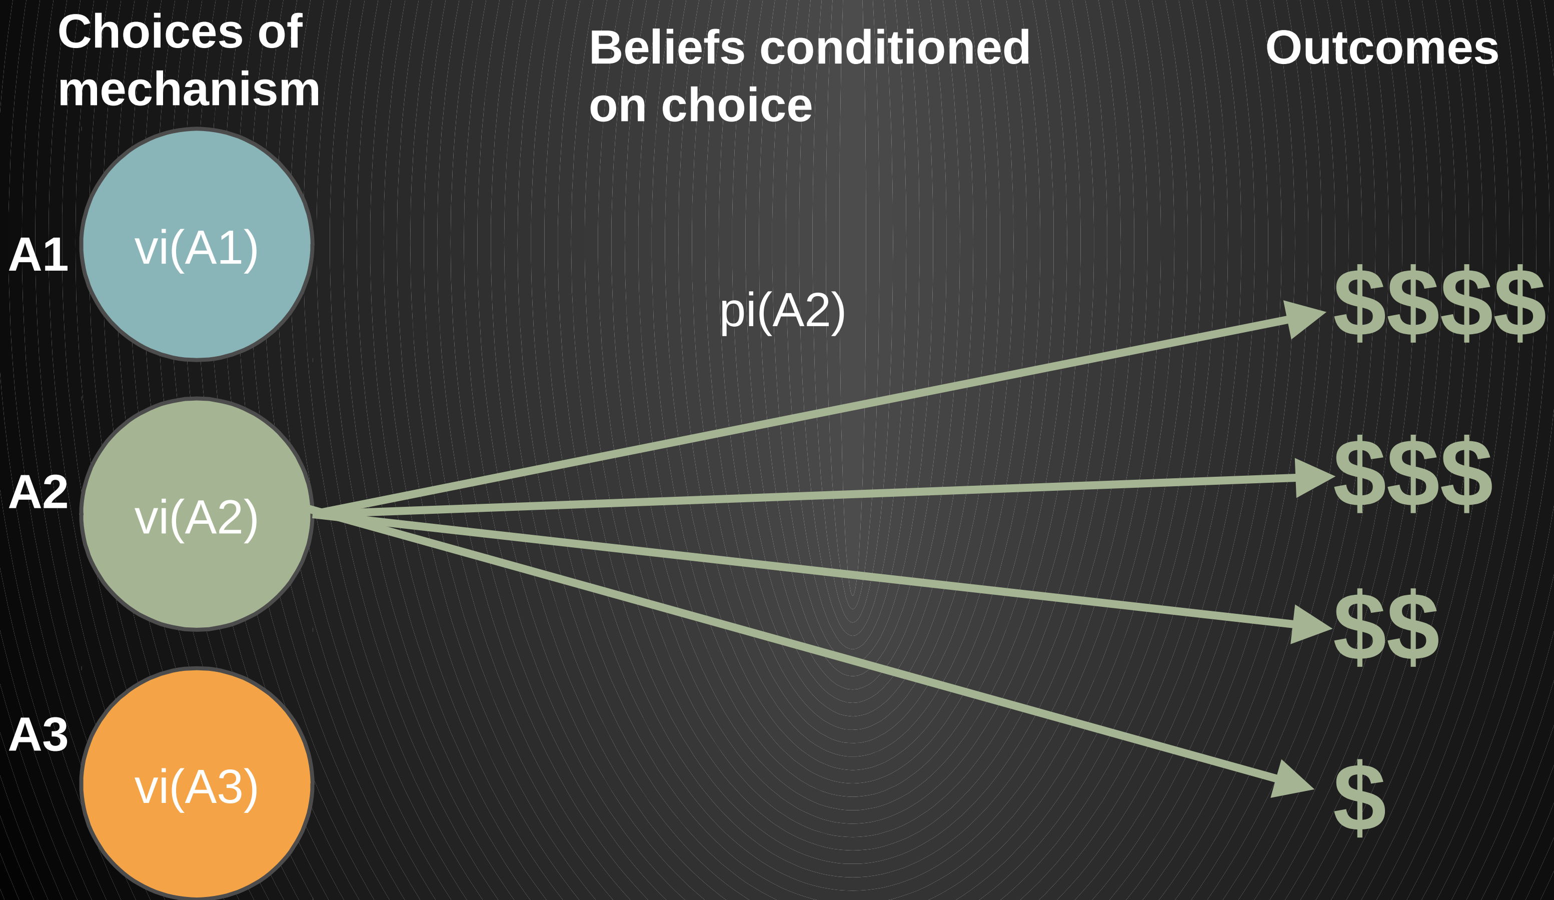

“Really General Model”

Example: “full” daily deals.

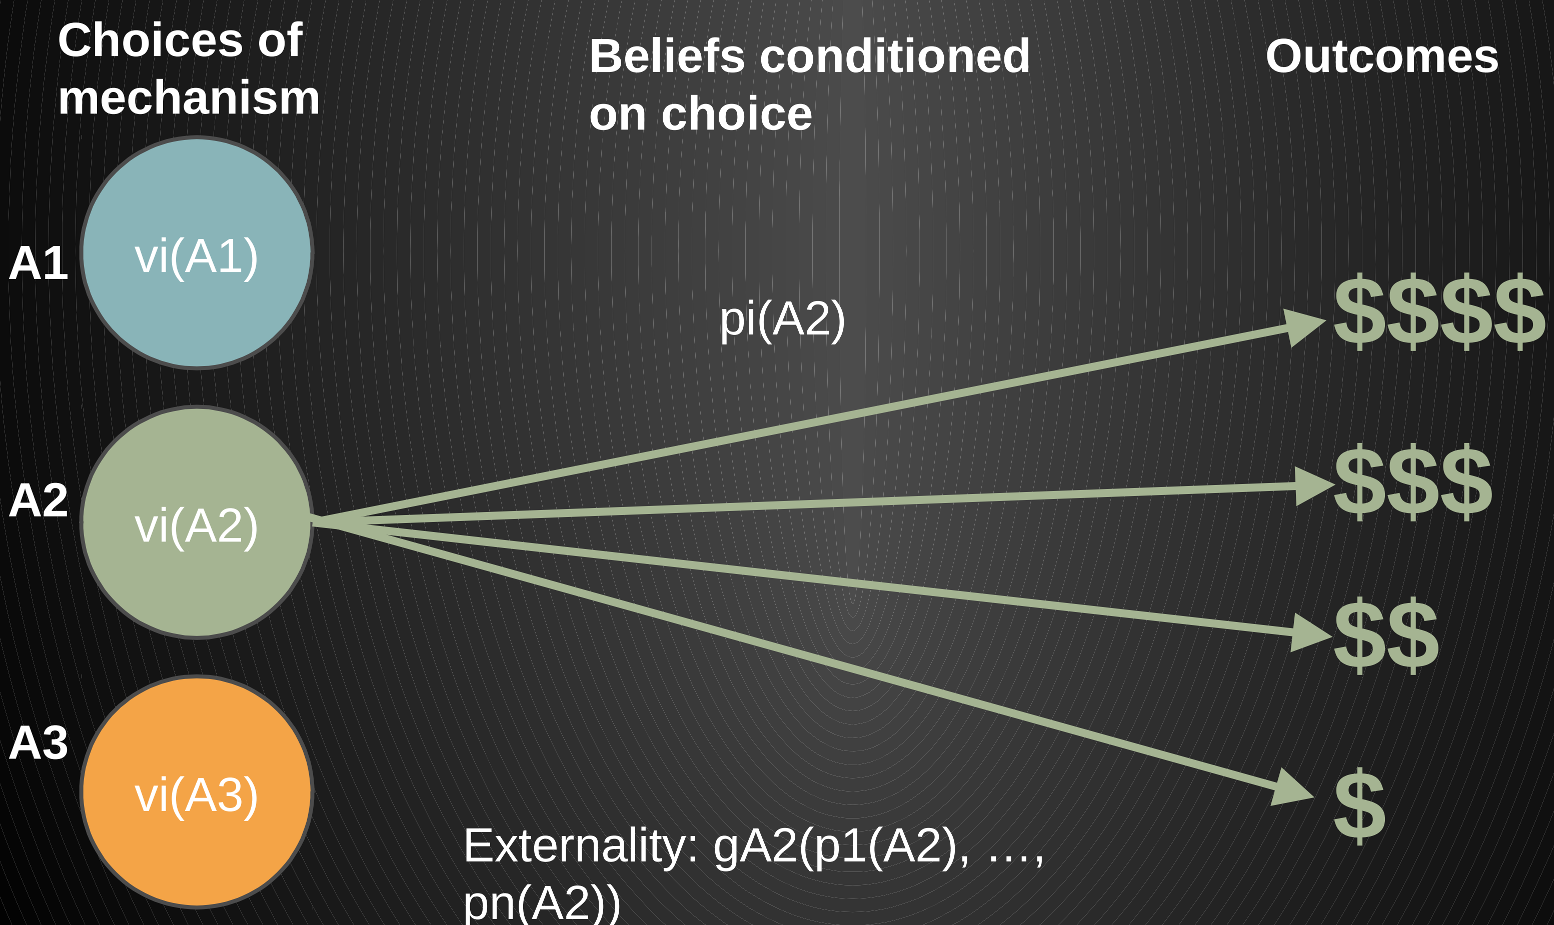

Q: For what externality functions *g* can we truthfully max welfare?

Theorem 2.
gA(p1(A),...) are **convex** in each argument ⇔ we can maximize welfare = **gA(p1(A),...) + sumi vi(A)**.

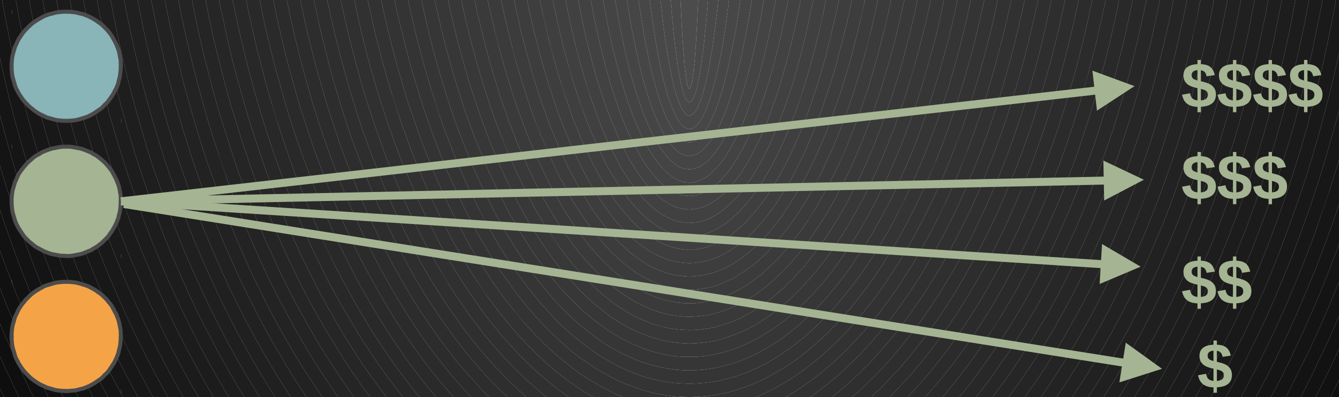

Q: For what externality functions *g* can we truthfully max welfare?

Theorem 2.
gA(p1(A),...) are **convex** in each argument ⇔ we can maximize welfare = **gA(p1(A),...) + sumi vi(A)**.

Auction: VCG and carefully constructed scoring rules.

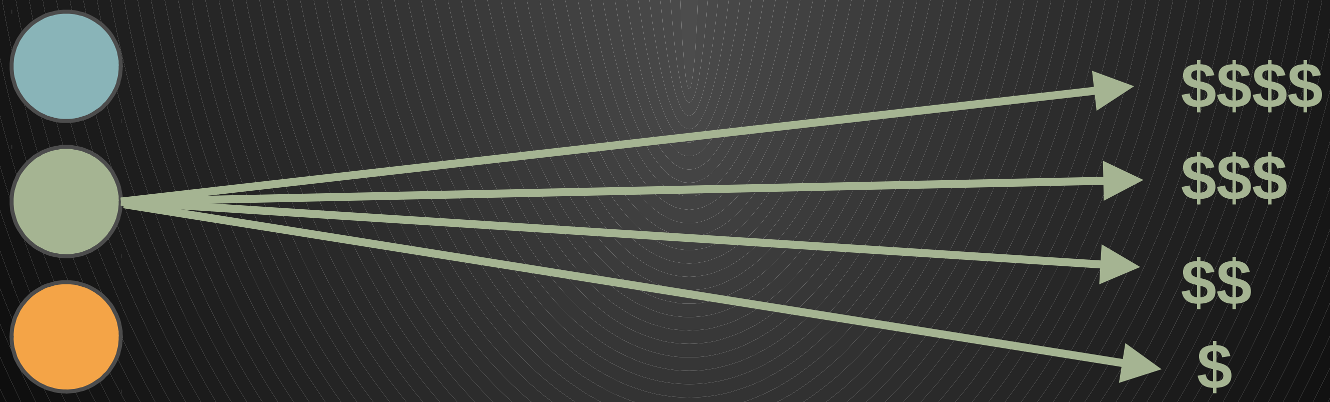

Outline

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 *vi*
  - stochastic delay ~ *pi*
- Utility of traveler: *g(p1, …, pm)* for path *1…m*
- Goal: maximize total welfare

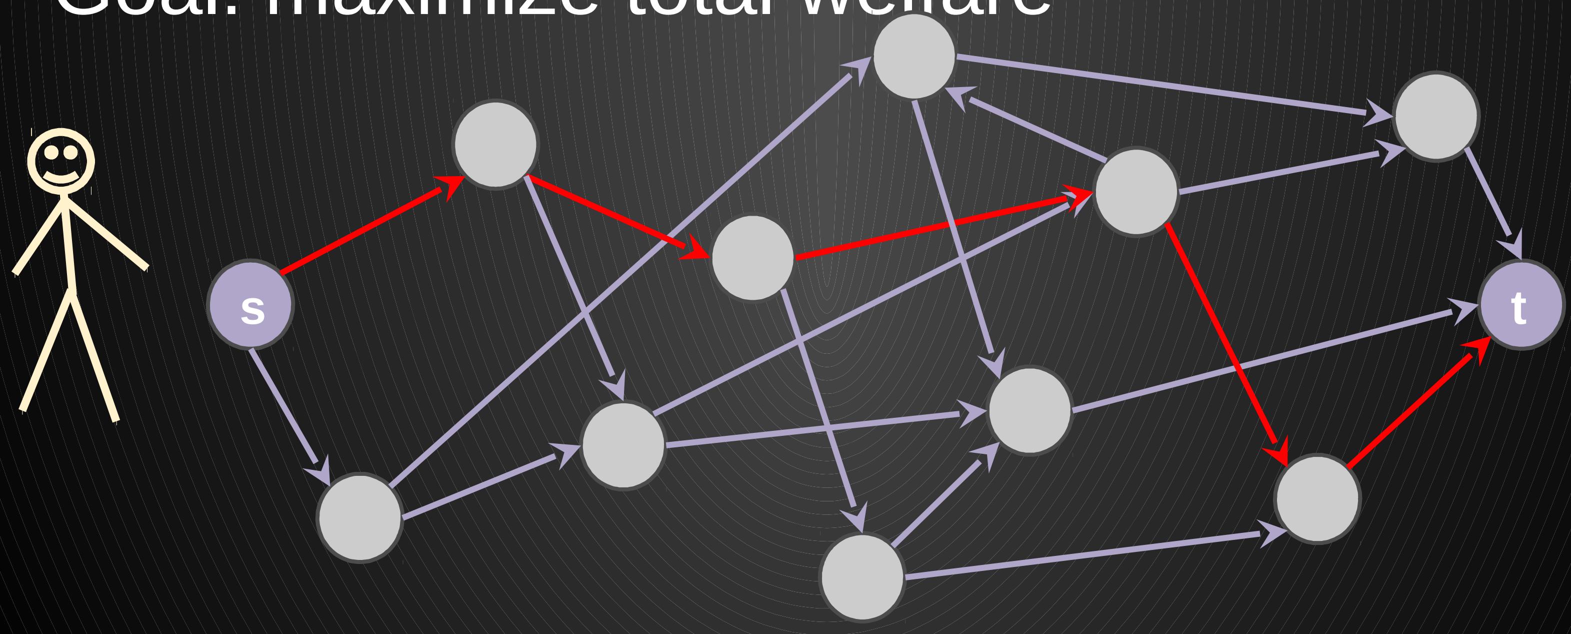

General takeaways

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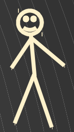
- … depending on private **predictions** of bidders
- Implementable ⇔ externality is convex function of prediction
- Auction = VCG + “decomposed” scoring rules

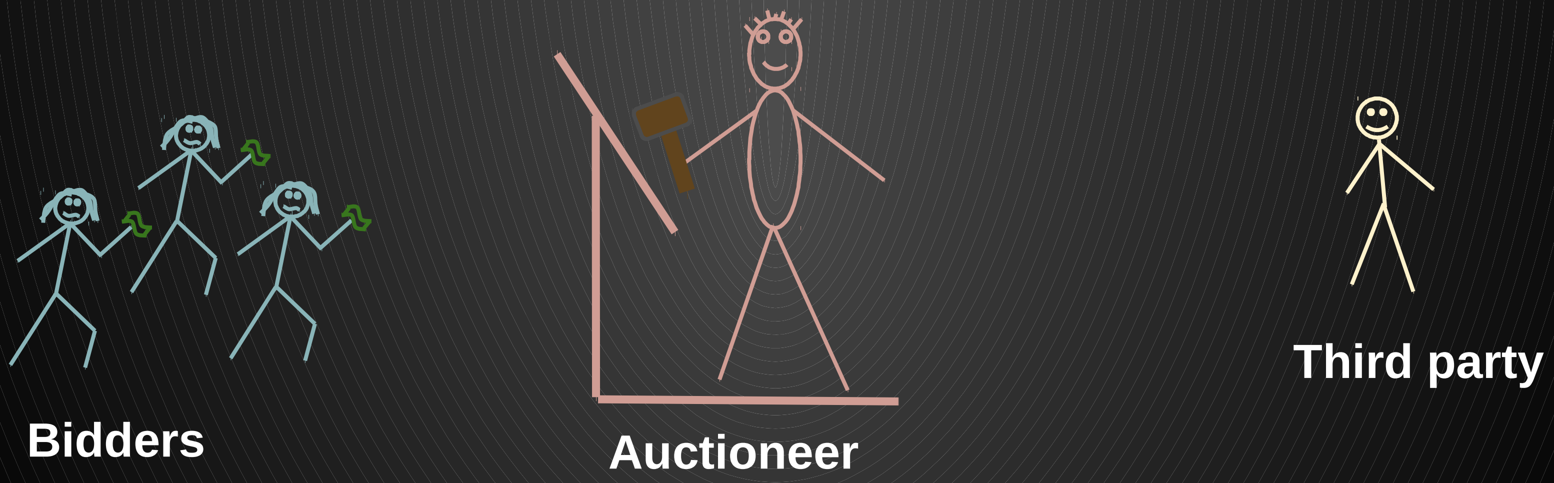

Future work

- 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

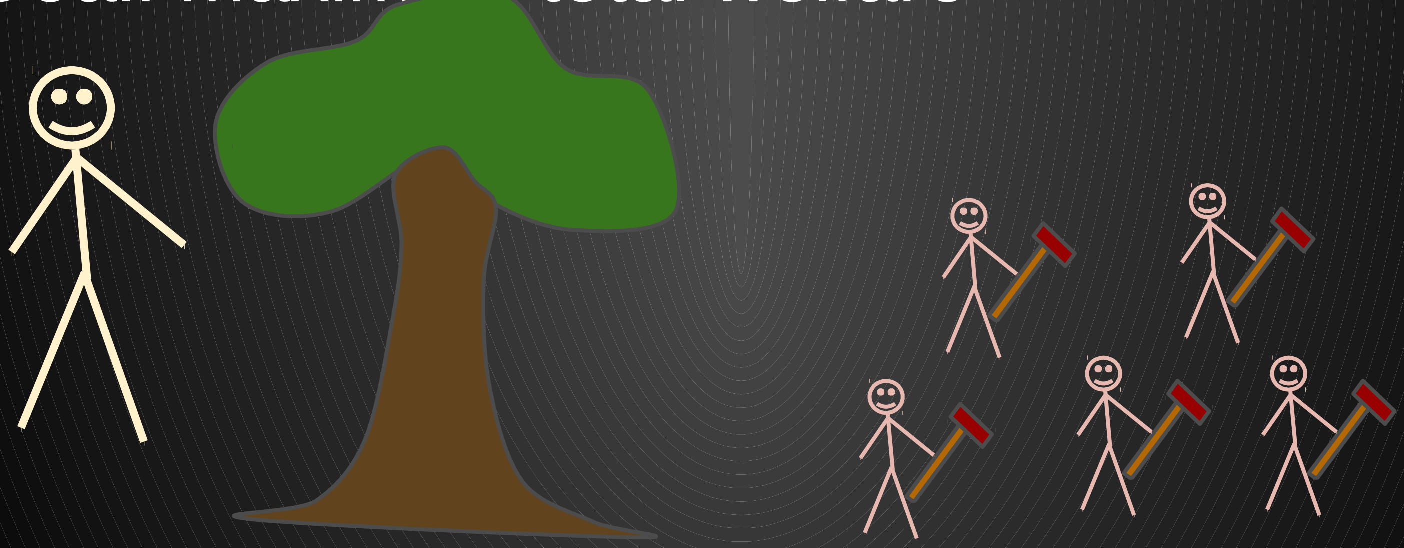

Designing Markets for Daily Deals

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

WINE 2013

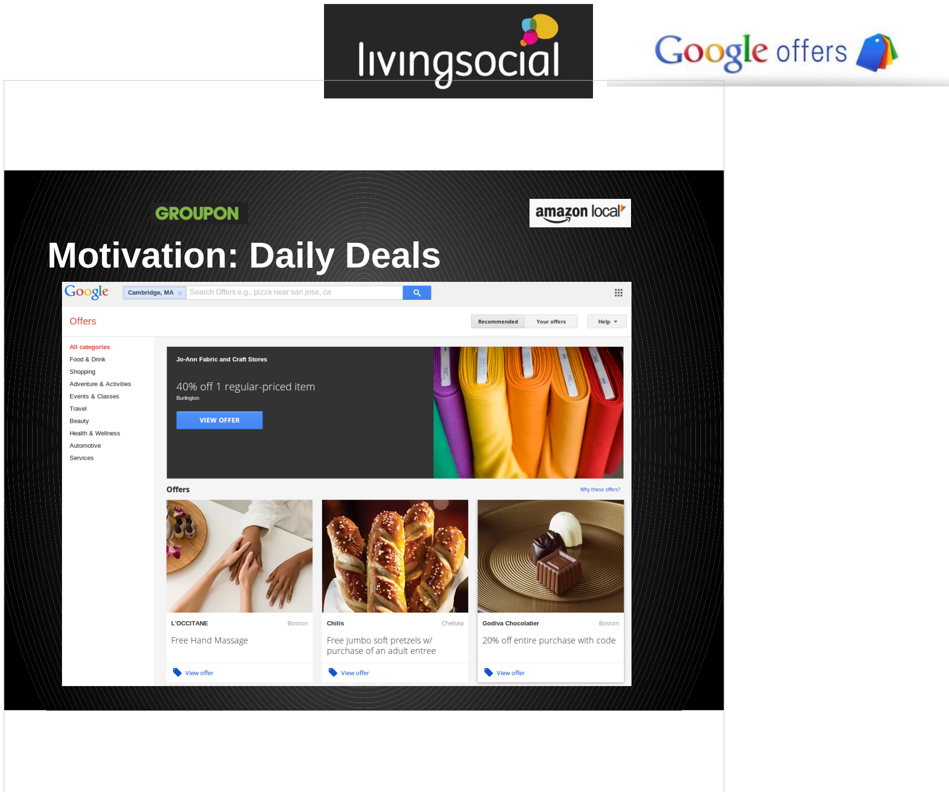

Daily deals sites are coupon aggregators.

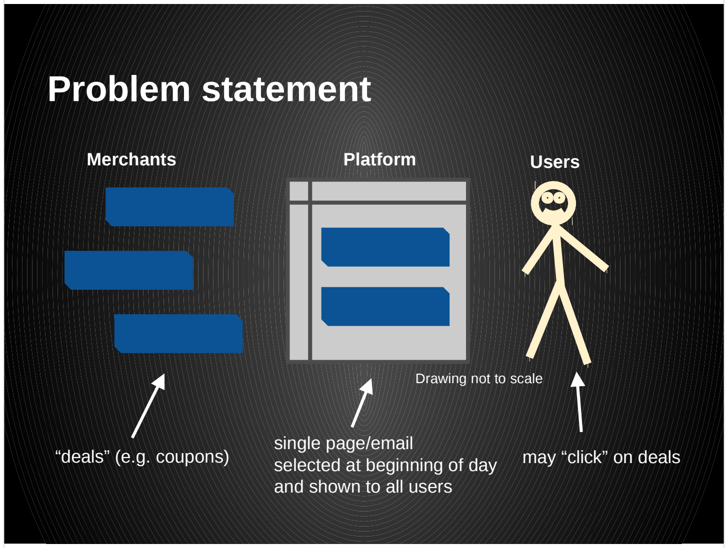

Merchants have deals (coupons) they want to offer users.
The platform, at beginning of the day, must select a single subset of deals to display on a website or send in a mass email to all users.
Users arrive, and obtain the deal in some fashion, and we just call this a click. Crucially, the platform observes this decision.
The goal is to automate the procedure of selecting deals with some sort of mechanism.

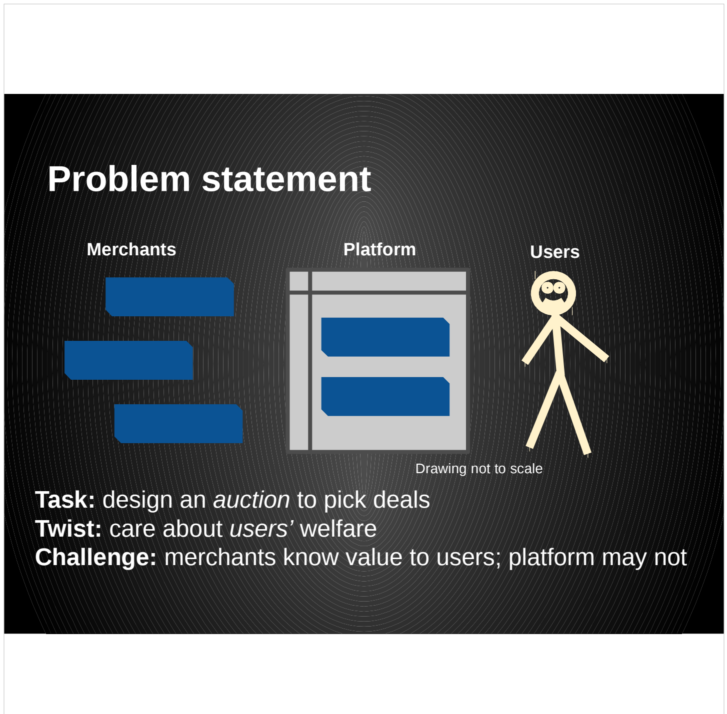

Whereas a standard auction only would concern the merchants and platform, here we want to pick deals that have good utility for the users; if the merchants have high values but the deals are terrible, users won't come back and the platform will fail.

Challenge in our setting: care about users, but impact on users is private information of merchants (many reasons, for example we've never seen this deal before today and have to make a final decision before we see its performance).

For instance, if you're familiar with ad auctions, in those models the platform has all information about the "quality" of the bidders, but here we want to understand what is possible with the opposite information asymmetry: bidders are completely informed and auctioneer is 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

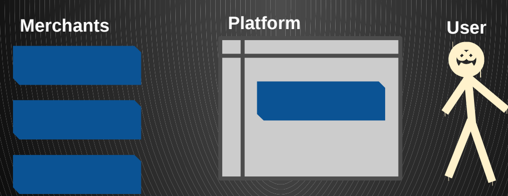

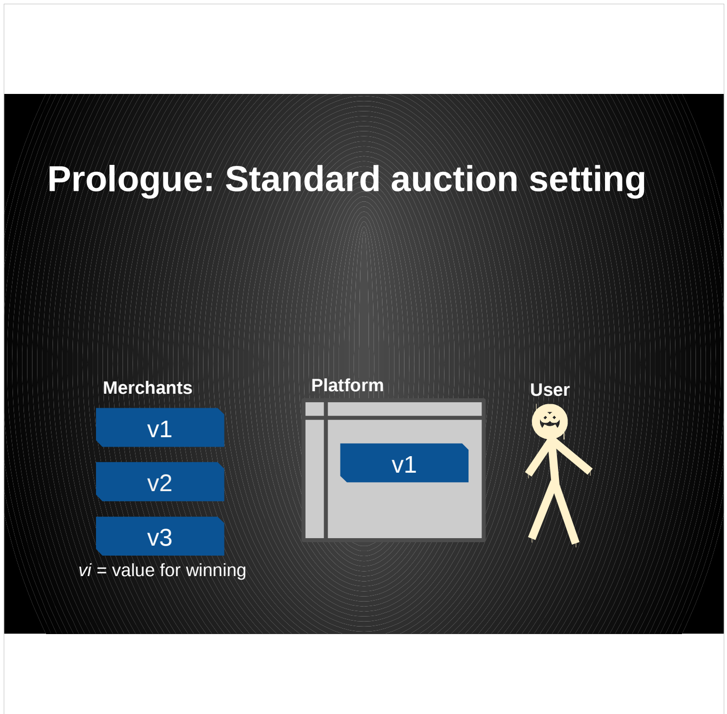

Standard auction model, each merchant has as private information a scalar vi, the value for winning the slot.

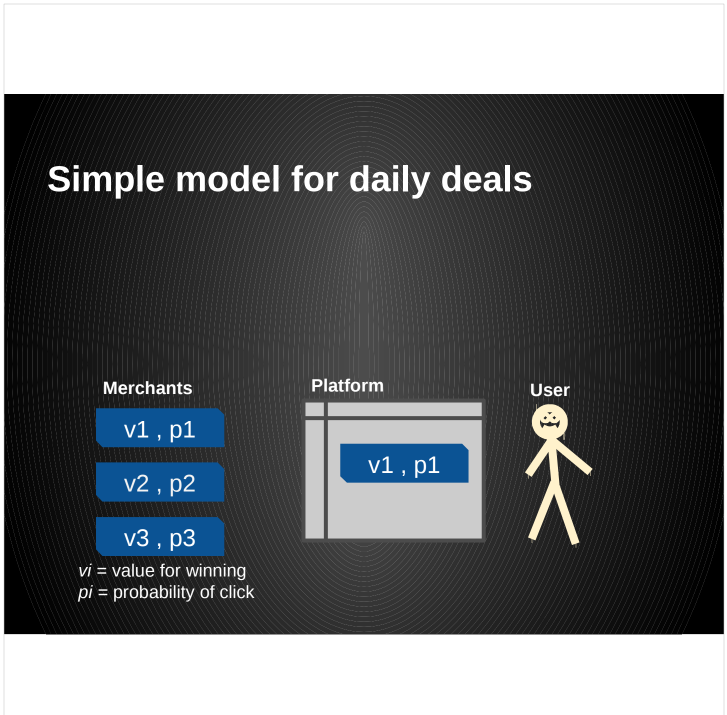

Now, merchants also have a parameter p, also private info, which is the probability that the user clicks on the deal.

v still = valuation for winning slot (hasn't changed!).

Of course, v and p can be related in some arbitrary way; doesn't matter to us, just two fixed parameters known to the merchant with these well-defined meanings.

Simple model for daily deals

- User welfare is related to *pi*
- First try: require *pi* to exceed “quality” threshold

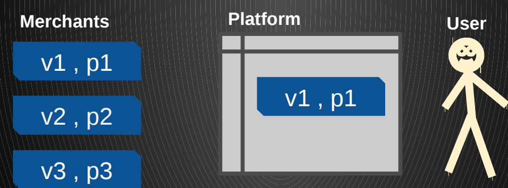

vi = value for winning
pi = probability of click

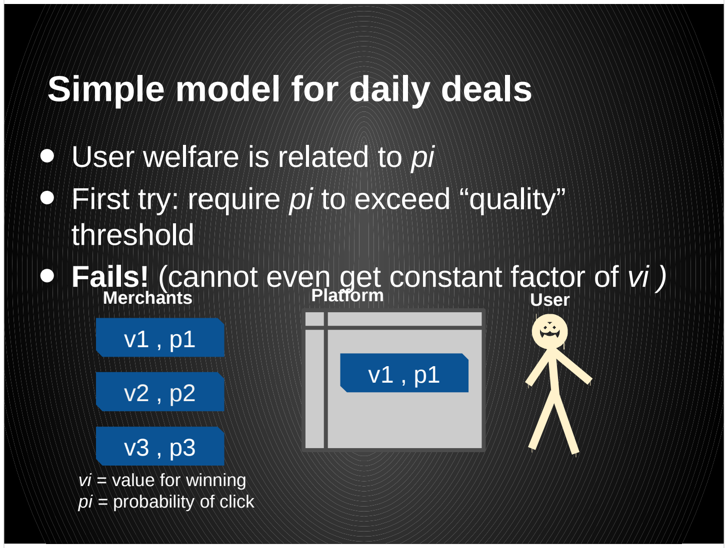

See the paper for the exact result, but the takeaway is that it seems hard to run a good auction by imposing any sort of threshold on pi (even a relatively soft one).

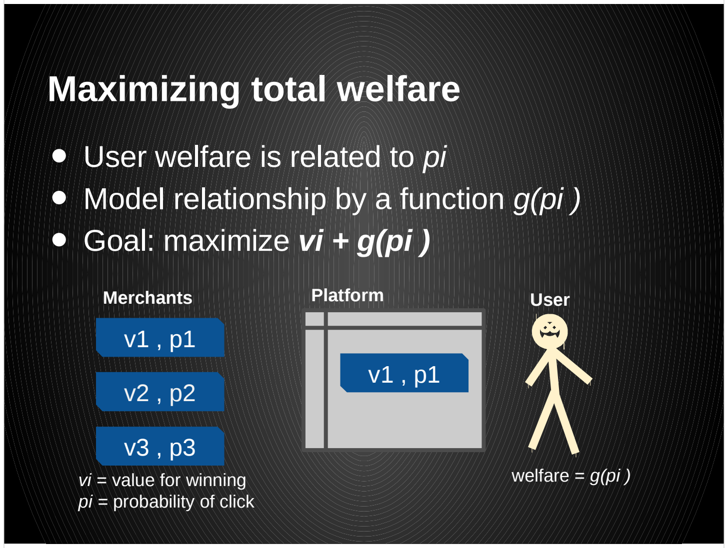

Bidders impact social welfare by both value *and* by their impact on the user; let's model this impact with a function g, thought of as "externality" on user, and now total welfare is v + g(p).

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

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

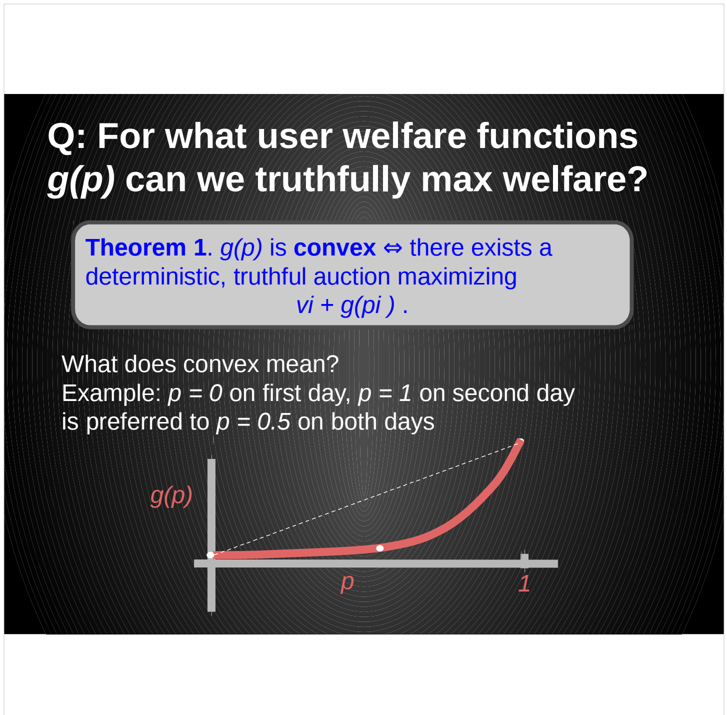

I claim that convex functions g(p) intuitively correspond to risk aversion (even though we are used to associating risk-aversion with concave functions, which are functions of wealth!).

The reason: We'd rather take something for certain than a lottery with same expectation. In this case, the user would prefer to click on exactly one deal out of the two days than face some lotteries with one click in expectation.

In general, you might think of g(p) as the "certainty equivalent" function for a lottery p; it is known that a convex g(p) corresponds to a concave utility function, i.e. risk-aversion.

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.

Example scoring rule: The log scoring rule.
Score(p, yes) = log(p).
Score(p, no) = log(1-p).
(Can be shifted/multiplied by a constant, of course!)
Of course, we can't simply apply scoring rules and auctions separately; we need to carefully combine them.

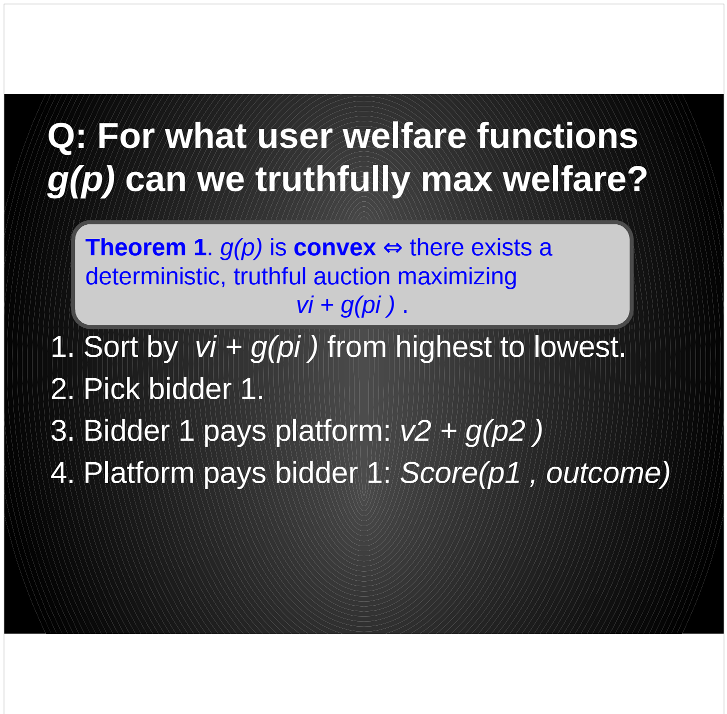

Here's the mechanism; the question is, which scoring rule do we pick for Step 4? It must be chosen carefully!

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*.

Example: For the log scoring rule,
g(p) = p log(p) + (1-p)log(1-p), i.e. the entropy function.
Note that this is convex in p and is equal to the expected score of an agent with belief p who reports truthfully.

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] = v1 + g(p1) - (v2 + g(p2))

You can easily check that bidder 1 wants to win only if v1 + g(p1) is higher than v2 + g(p2), and furthermore, prefers telling the truth to lying when she does win. This implies truthfulness of the auction.

Outline

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

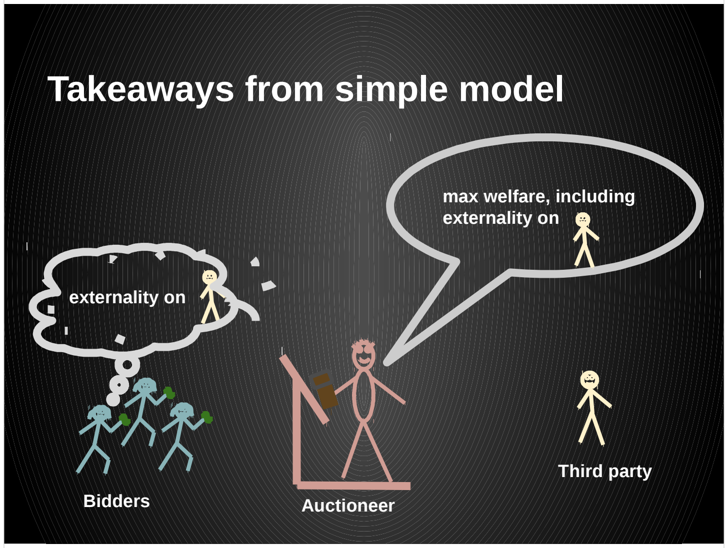
Takeaways from simple model
max welfare, including
externality on
externality on
Bidders
Auctioneer
Third party

Takeaways from simple model

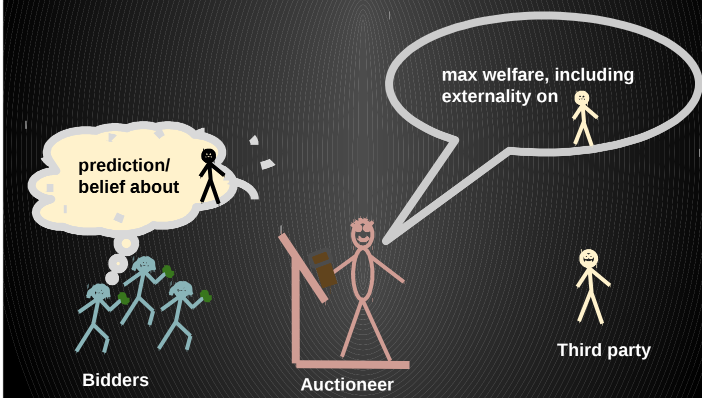

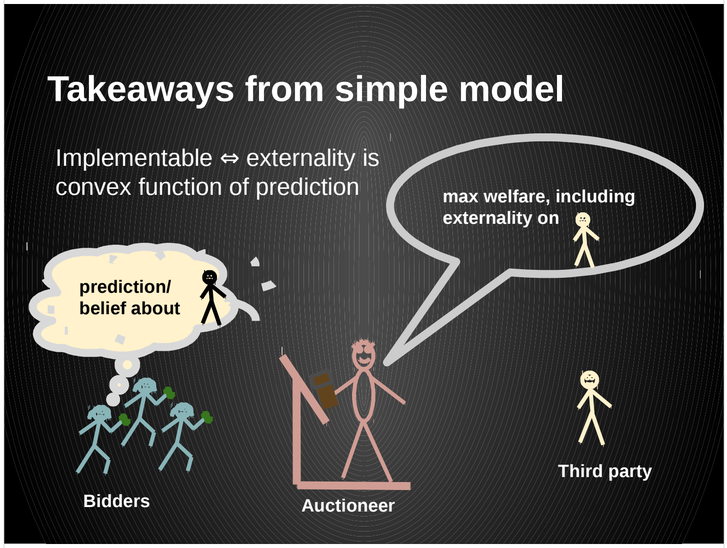
Takeaways from simple model
Implementable ⇔ externality is convex function of prediction
max welfare, including externality on
prediction/ belief about
Bidders
Auctioneer
Third party

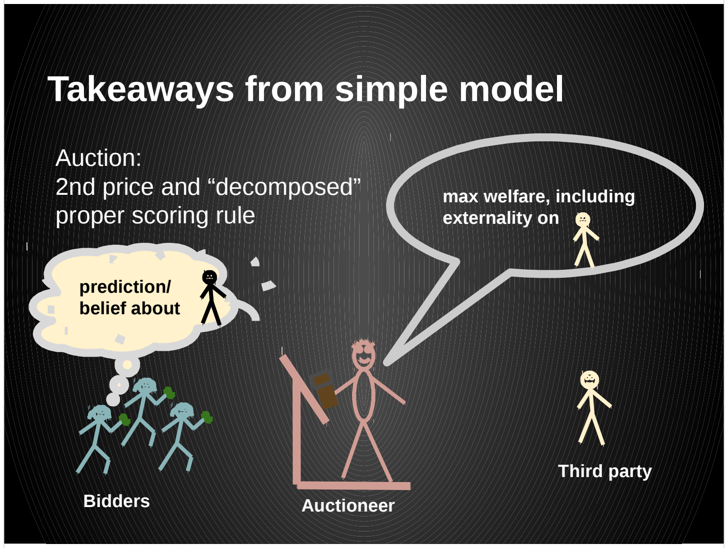
Takeaways from simple model
Auction:
2nd price and “decomposed” proper scoring rule
max welfare, including externality on
prediction/ belief about
Bidders
Auctioneer
Third party

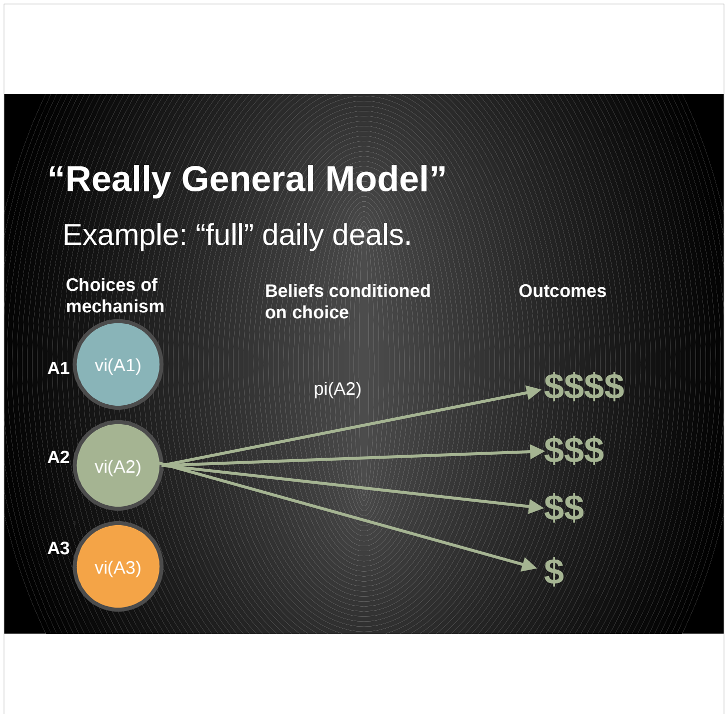

Example: a daily deals setting with multiple slots on the page for sale and many arriving users.
The choices of the mechanism are the possible page layouts: An assignment of slots to merchants.
The outcomes for bidder i could be, for instance, total number of clicks on i's advertisements (if any).
i has beliefs conditional on the mechanism's choice: Given that the mechanism chooses layout A2, i has a distribution over the total number of clicks she will get.

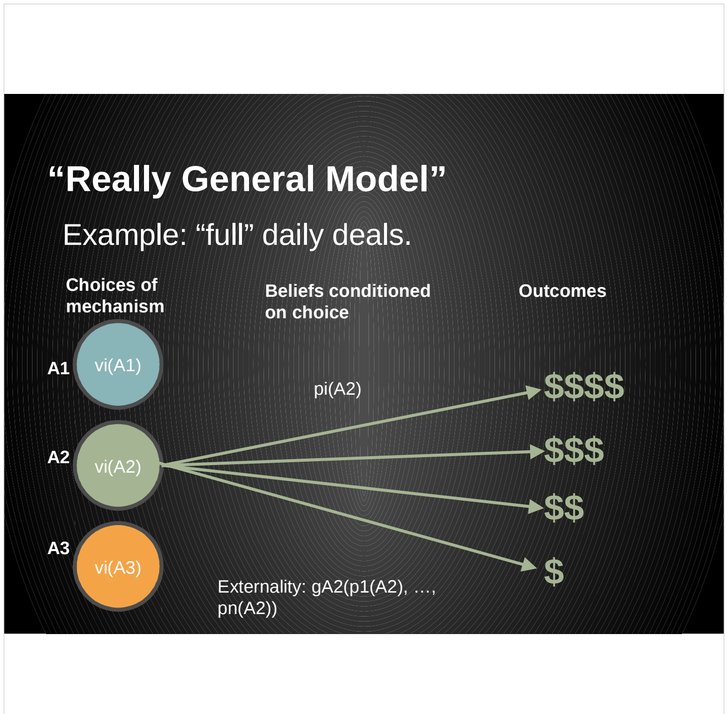

The “third party”’s welfare is modeled by a function of all the predictions.
There could be a different function for each possible choice of the mechanism.

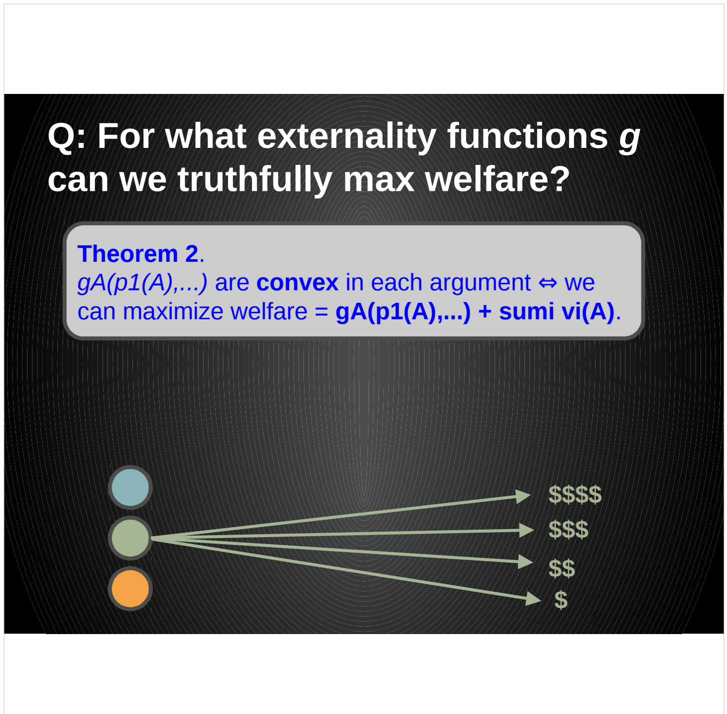

There exists a deterministic, truthful, welfare-maximizing auction if and only if all g_A(p1(A),...) are convex in each argument.

Note that this is a weaker condition than convexity: The function f(x,y) = x*y is convex in each argument, but is not a convex function.

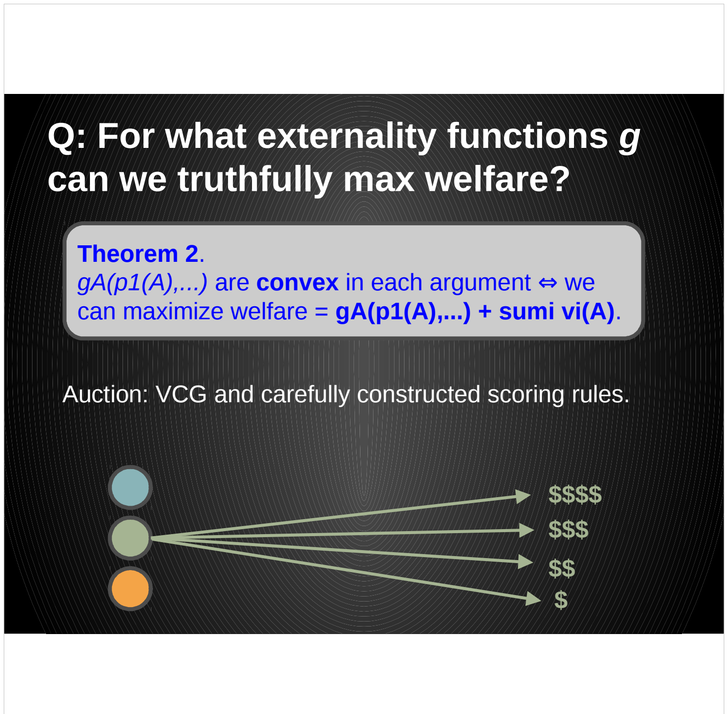

The auction relies on, for each agent, carefully decomposing gA(p1(A),...) into a scoring rule for that particular agent, after plugging in the predictions of all other agents. Then, the form of the auction looks like VCG.

Outline

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

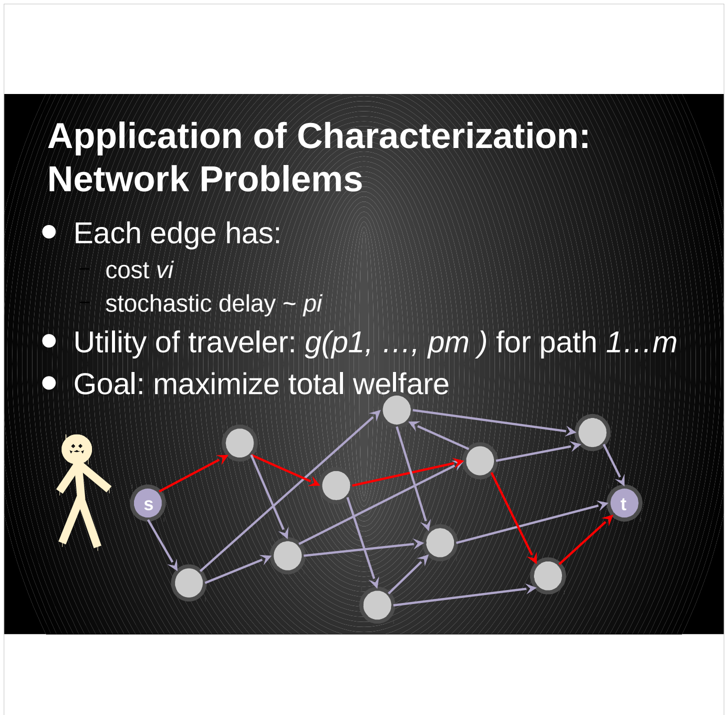

This is an unrelated problem. However, once we write it down formally, we can notice that it fits our general model, and therefore our results apply (both positive and negative).

Note that a convex function g corresponds to a risk-averse traveler: One who prefers a more certain travel time to a more uncertain one, even with a higher expected travel time.

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*

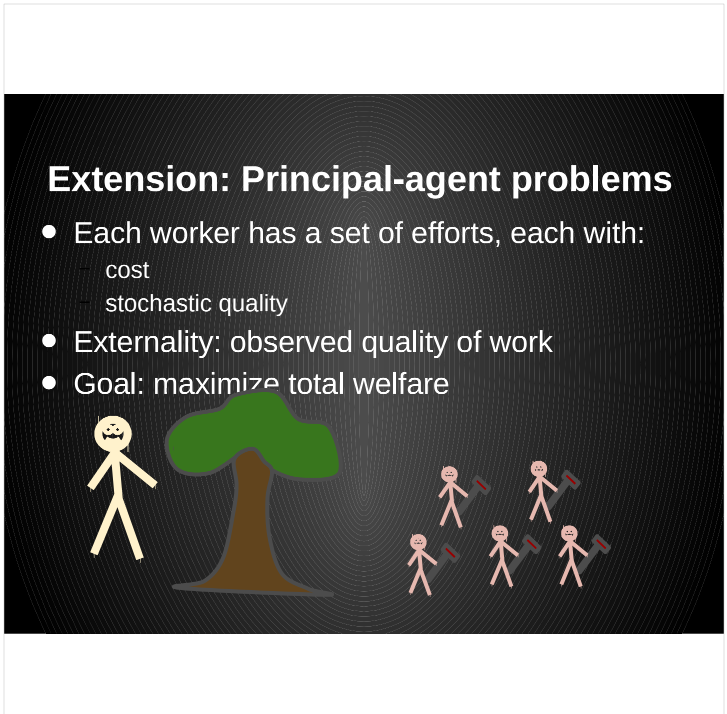

Another quick example of unrelated areas where our results apply. Take a principal-agent problem where the agents hold private information about their quality, i.e. the quality is picked from a distribution known to the agents but not the principal. The externality here is on the principal (whom we might also think of as the auctioneer, but this is no problem for the model). If the principal is risk-averse with respect to quality, our auction applies.