Auctions between Regret-Minimizing Agents

Yoav Kolumbus: yoav.kolumbus@cornell.edu

Based on joint work with Noam Nisan

Algorithmic Game Theory, CS 6840
September 2024

Outline

- Online auctions, automated bidding agents, the meta-game model
- Learning agents and regret minimization
  Hedge/Multiplicative-Weights algorithm - recap
- Short detour: On the convergence of regret minimization dynamics
- Auctions with learning agents
- Open problems (time permitting)

Online auctions

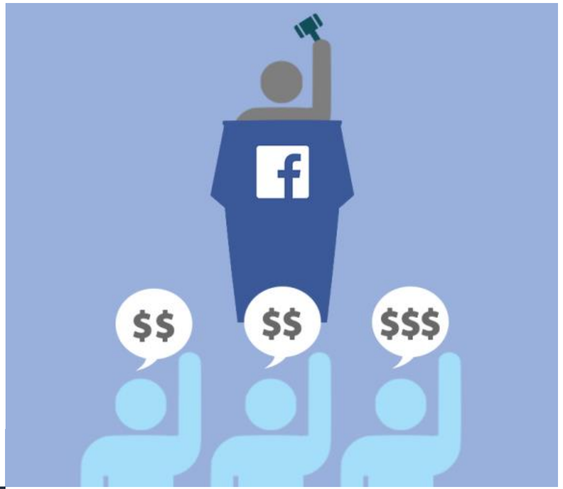

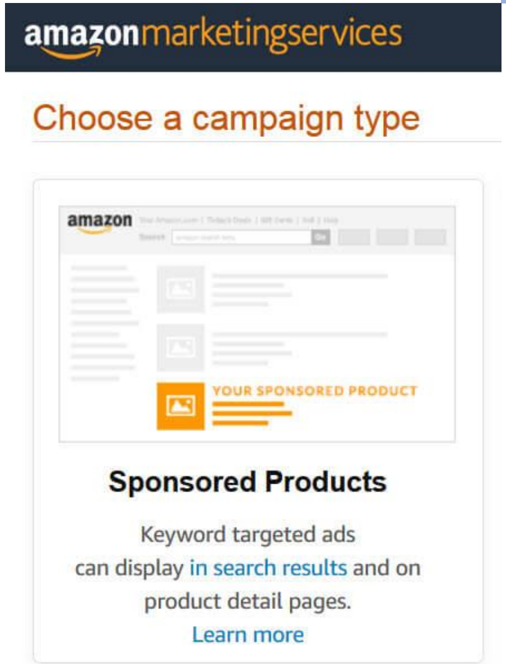

Bing ads

Example: keyword auctions

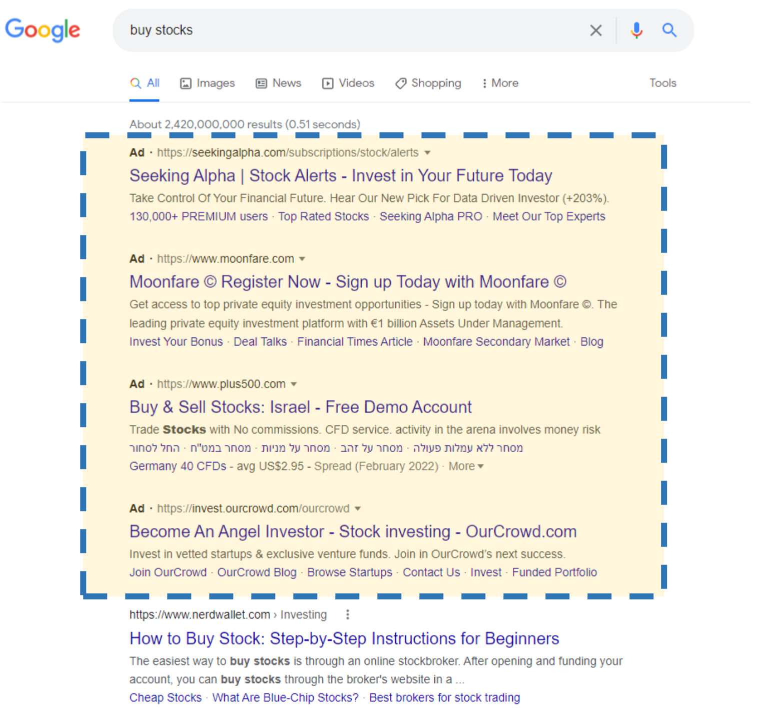

An online search query

Bidders place bids for ad positions

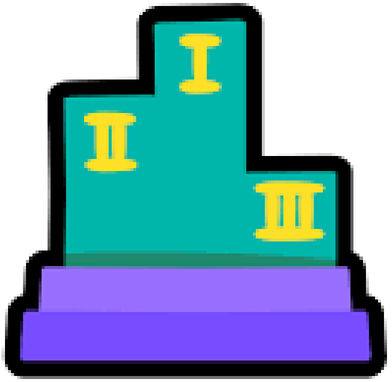

Auctioneer sets the allocations of items to bidders and their payments (per click)

Automated bidding

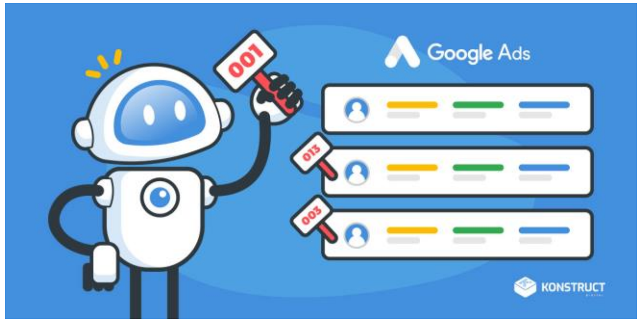

Basic facts:

- A very large industry (significant part of revenues for Google, Meta…)
- Auctions run at fast rates (thousands per second and more)
- Most of the bids are being placed by various auto-bidding tools

How does auto-bidding work? (in a nutshell)

- Users enter key parameters into the auto-bidding agent interface
- Then the agents place bids, interact, and learn
- Users observe the long-term outcomes

The classic auction setting

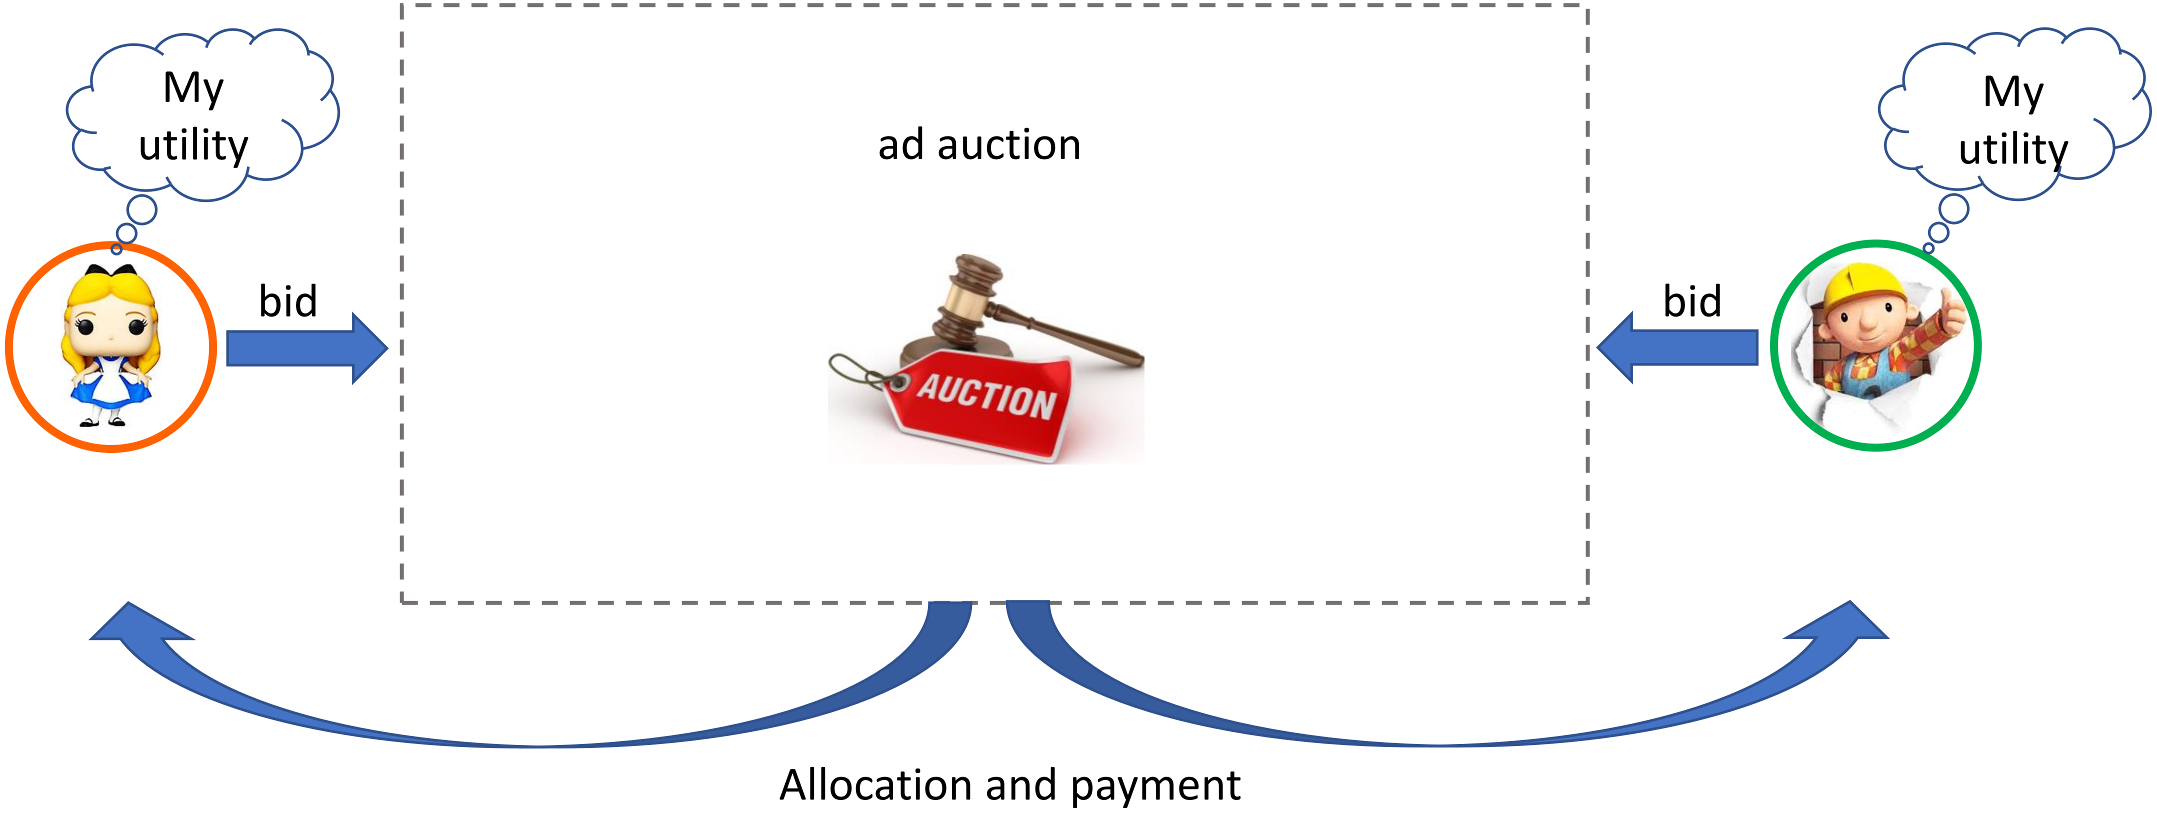

The classic auction setting

$$\text{Utility} = \begin{cases} value - payment, & \text{if I win the item} \\ \text{zero}, & \text{if I do not win} \end{cases}$$

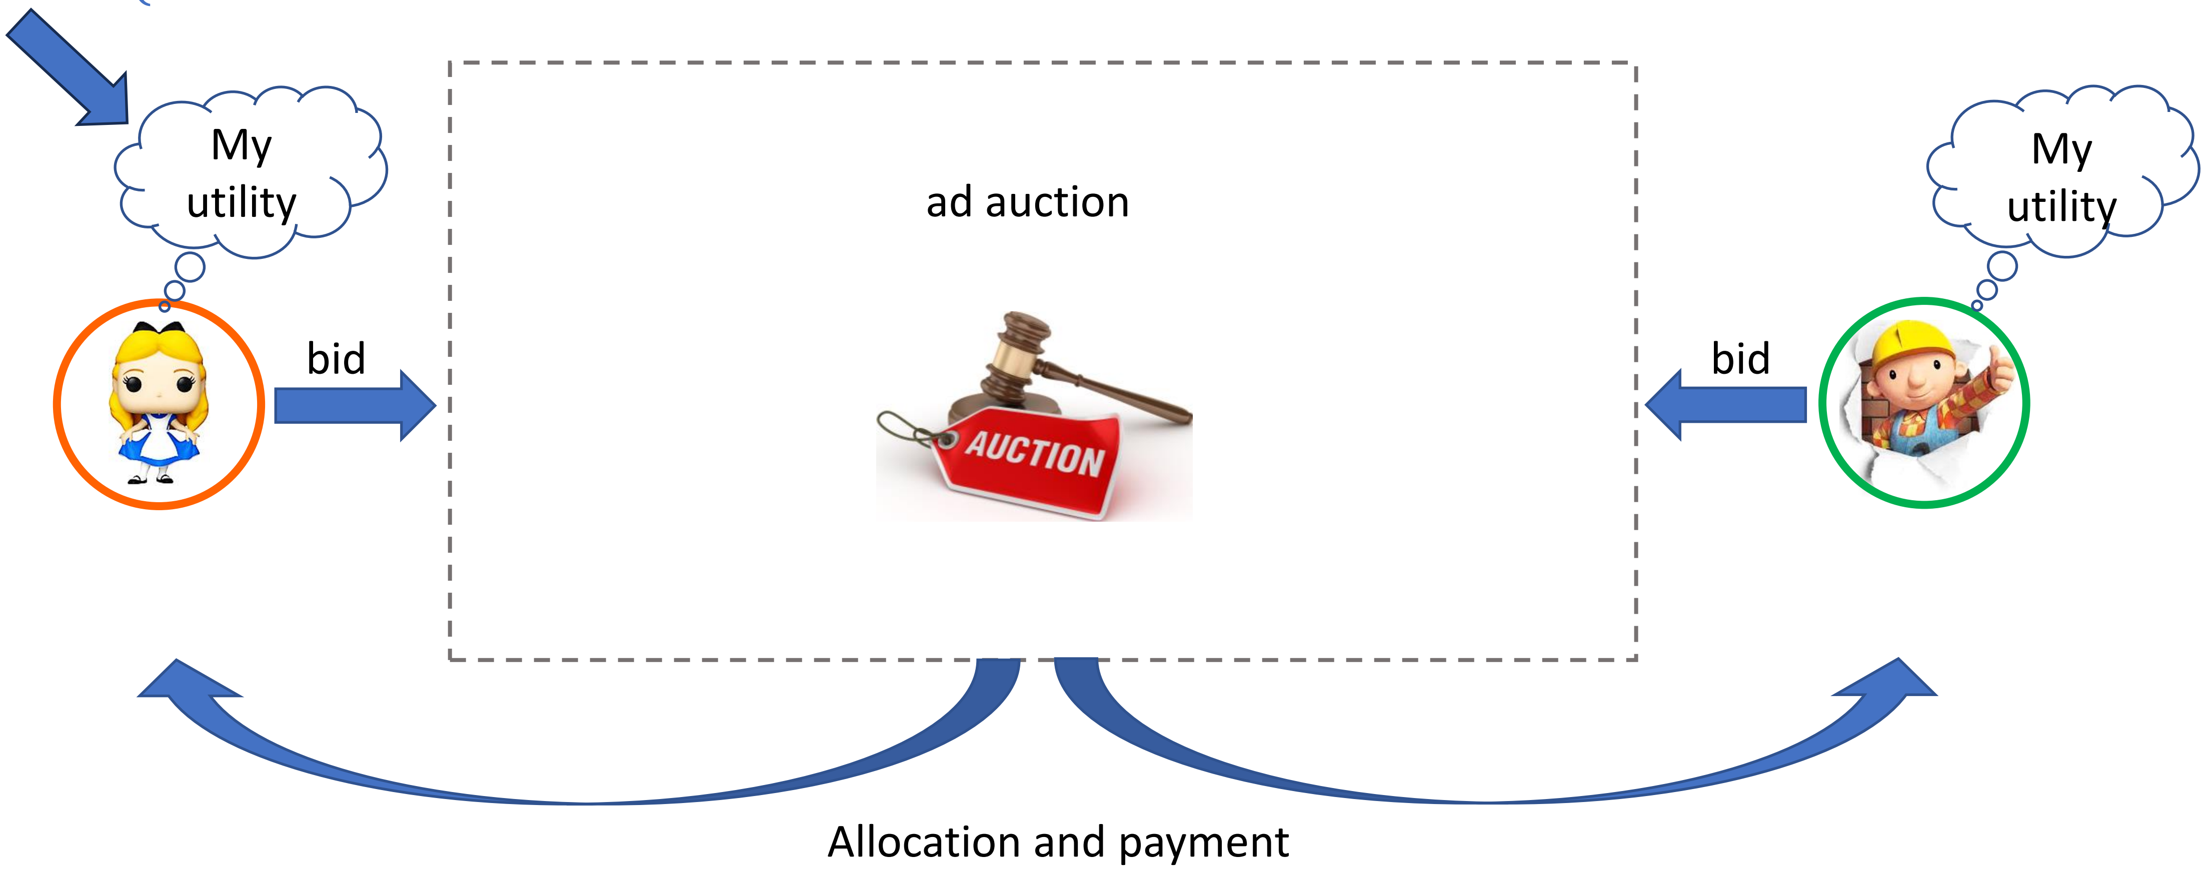

The auto-bidding setting

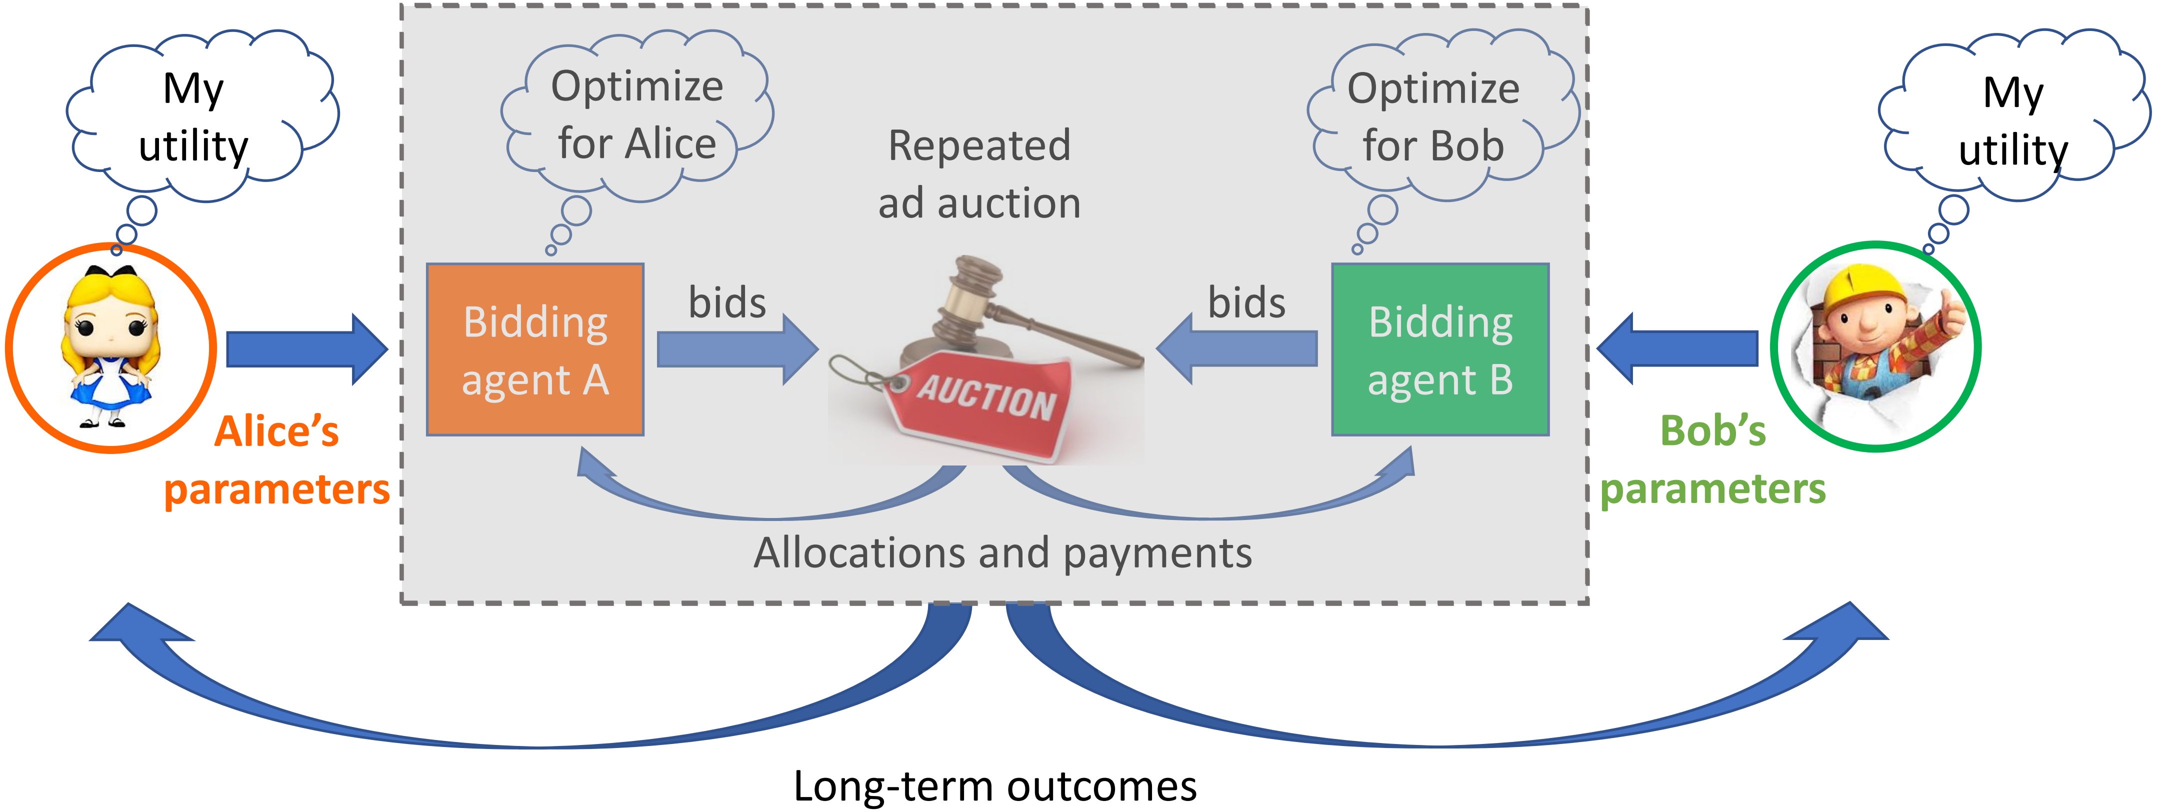

The auto-bidding setting

The "meta-auction"

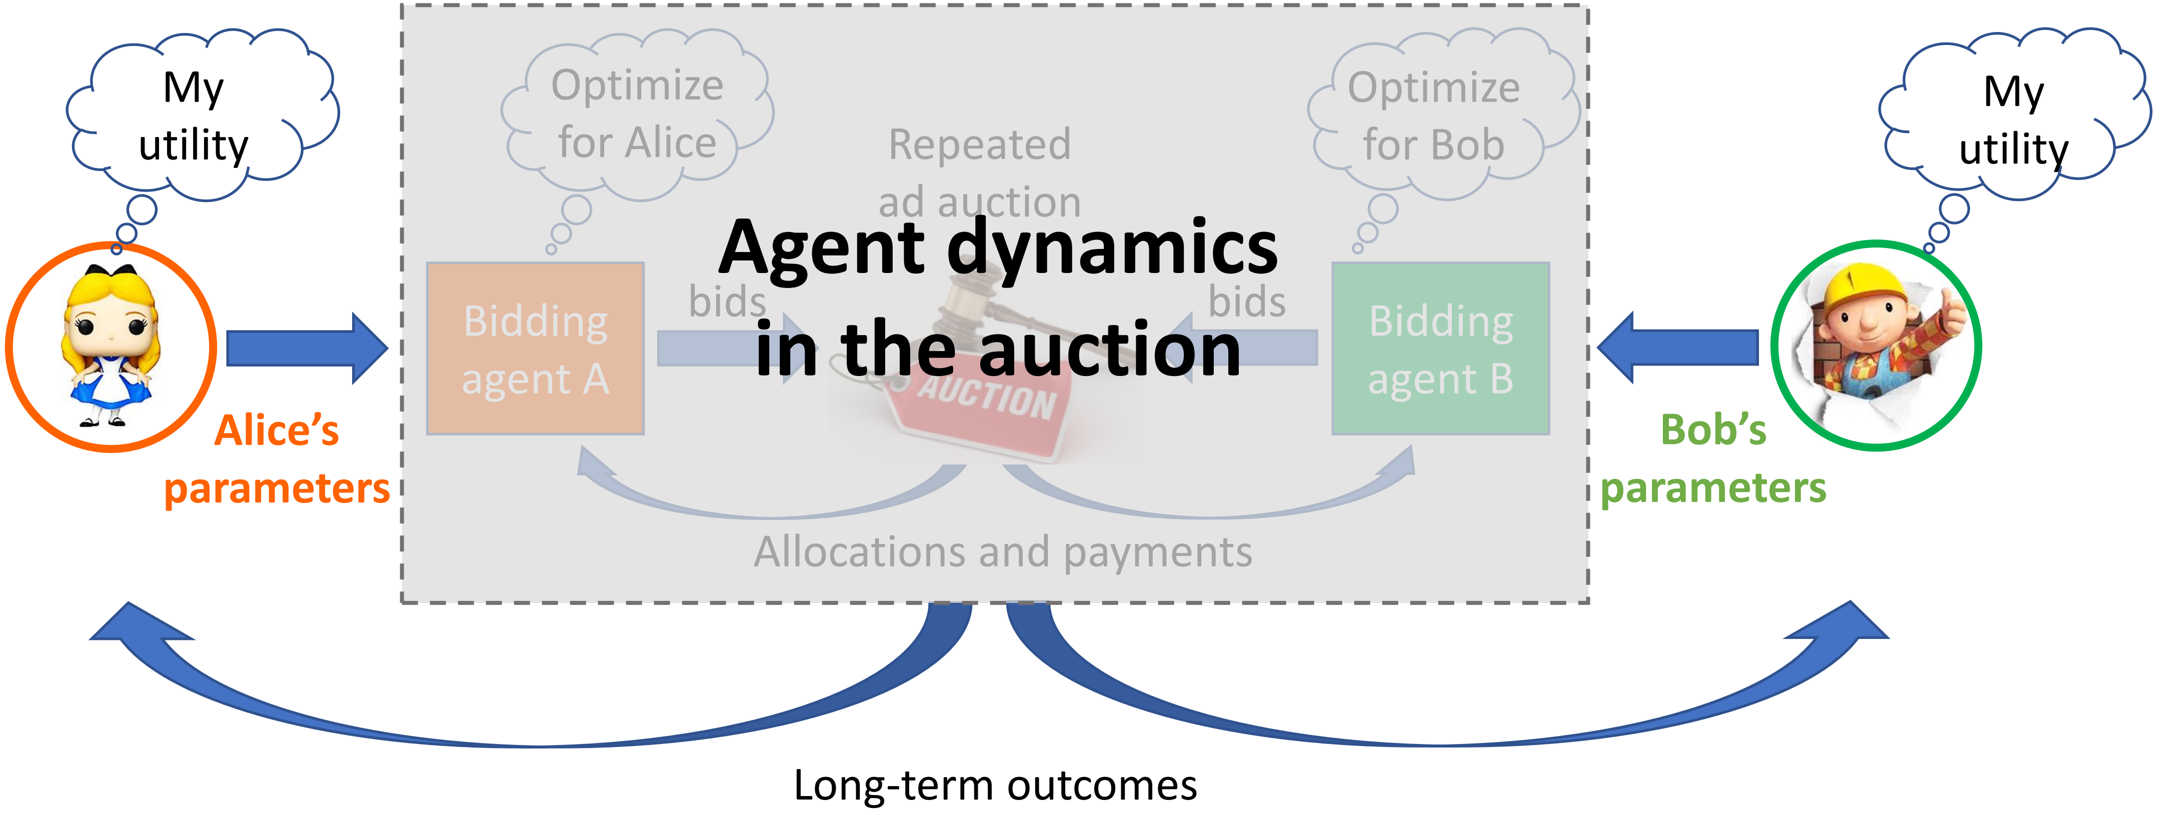

The "meta-game"

How and Why to Manipulate your Own Agent: On the Incentives of Users of Learning Agents [Kolumbus & Nisan, NeurIPS 2022]:

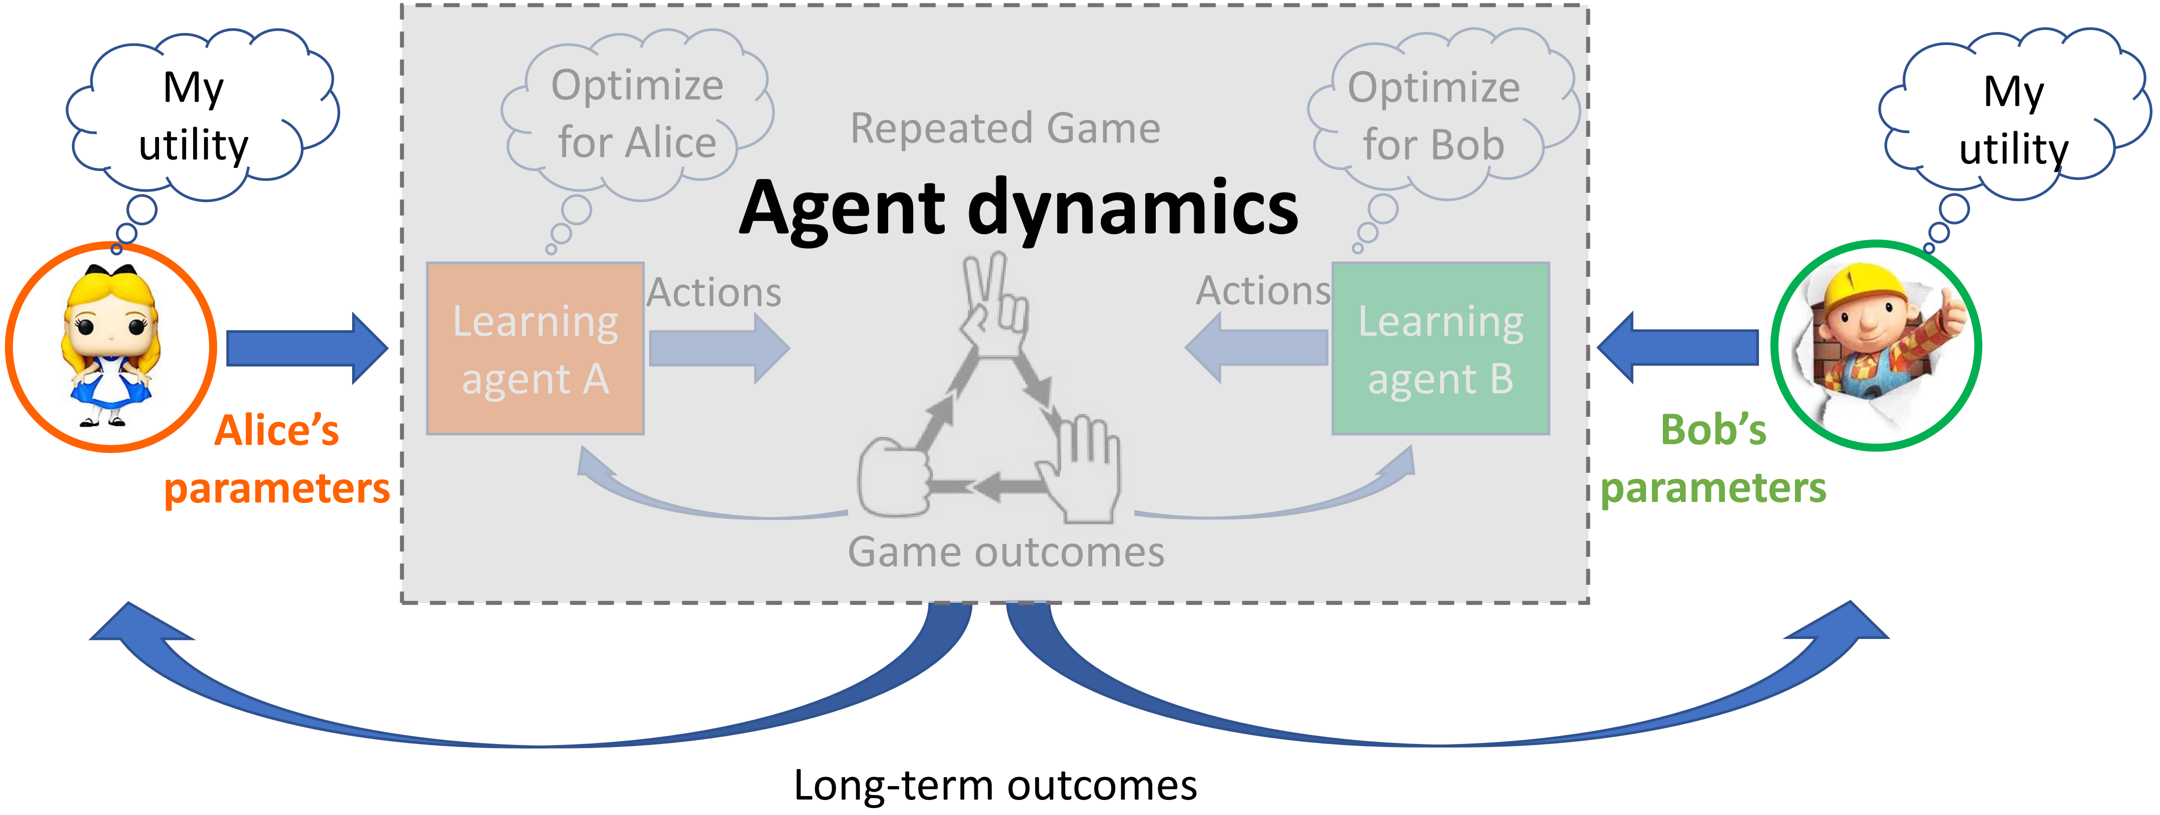

Learning agents in repeated games

- Basic idea: best-reply to learned (empirical) distribution of other bidders: "fictitious play" (a.k.a. "follow the leader")
  - But might sometimes lead to bad performance
- Improved idea: "soft best-reply" – play actions that performed better in the past with higher probability
  - Studied since the 1950s and today
- Example: The Multiplicative Weights (MW) algorithm:
  - Initialize $w_a = 1$ for every action (bid) a. Then for $t = 1,2 \dots, T$
  - Play action a with probability $w_a^t / \sum_j w_j^t$ (Note: t is an index, not an exponent)
  - After every step t, for every a, update the weights: $w_a^{t+1} = w_a^t (1 + \epsilon)^{u_a^t}$ where u_a^t is the *utility* of action a at time t

Multiplicative weights: summary and variants

- Hedge update rule:
  - Utility: $w_a^{t+1} = w_a^t (1+\epsilon)^{u_a^t}$ or $w_a^{t+1} = w_a^t e^{\eta u_a^t}$
  - Loss: $w_a^{t+1} = w_a^t (1-\epsilon)^{l_a^t}$ or $w_a^{t+1} = w_a^t e^{-\eta l_a^t}$
- Linear multiplicative weights:
  - Utility: $w_a^{t+1} = w_a^t (1+\epsilon u_a^t)$
  - Loss: $w_a^{t+1} = w_a^t (1-\epsilon l_a^t)$

In games:

Experts → actions

The utility from an action a depends on actions of the others: $u_a^t = u_a^t(s)$.

Regret-Minimizing Agents

- Regret-minimizing agents:
  - Low regret: agents play such that the long-term average empirical payoffs approach the payoff of the best fixed strategy in hindsight
  - Examples: MW, FTPL, OGD…
  - Real-world bids in ad auctions are largely consistent with regret minimization [Nekipelov, Syrgkanis, Tardos 2015], [Noti & Syrgkanis, 2021]
- The dynamics approach the set of Coarse Correlated Equilibria (CCE)
- But the set of CCEs may be too large to analyze quantities of interest (like the utilities, revenue, etc.) → Analyze convergence of the dynamics

Detour: Convergence of no-regret dynamics (1)

An example:

MW agents playing matching pennies

- The dynamics do not converge

- The dynamics' time average does converge to the Nash equilibrium

Why the NE? In this example it is also the unique CCE [Calvó-Armengol 2006]

	Heads	Tails
Heads	1, -1	-1, 1
Tails	-1, 1	1, -1

Game

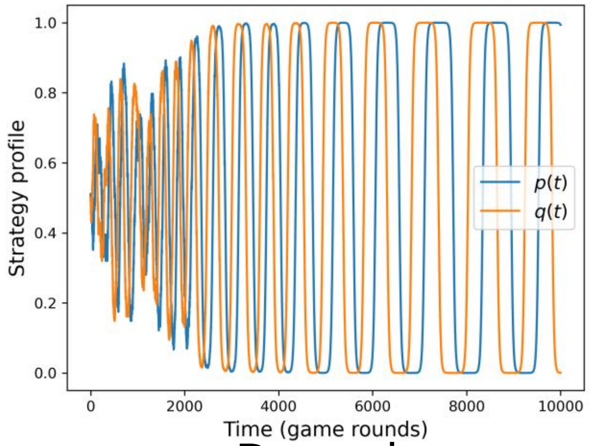

Dynamics

	0.5	0.5
0.5	0.25	0.25
0.5	0.25	0.25

Equilibrium

	0.505	0.495
0.509	0.258	0.251
0.491	0.247	0.244

Empirical Distribution

Detour: Convergence of no-regret dynamics (2)

Does the distribution of play always converge?

Theorem [Kolumbus & Nisan, NeurIPS 2022]. *For every finite game in which the set of CCEs is not a singleton there exist regret-minimizing algorithms for the players whose empirical time-average joint dynamics do not converge to any point.*

General no-regret dynamics → Convergence only with a unique CCE!

- Some potential games [Neyman 1997]
- Fully-mixed 2x2 games [Calvó-Armengol 2006]
- Dominance-solvable games
- Socially-concave games [Even-Dar, Mansour, Nadav 2009]

Analyze concrete classes of games and algorithms

Recap:

- A “meta-game” between users of bidding agents

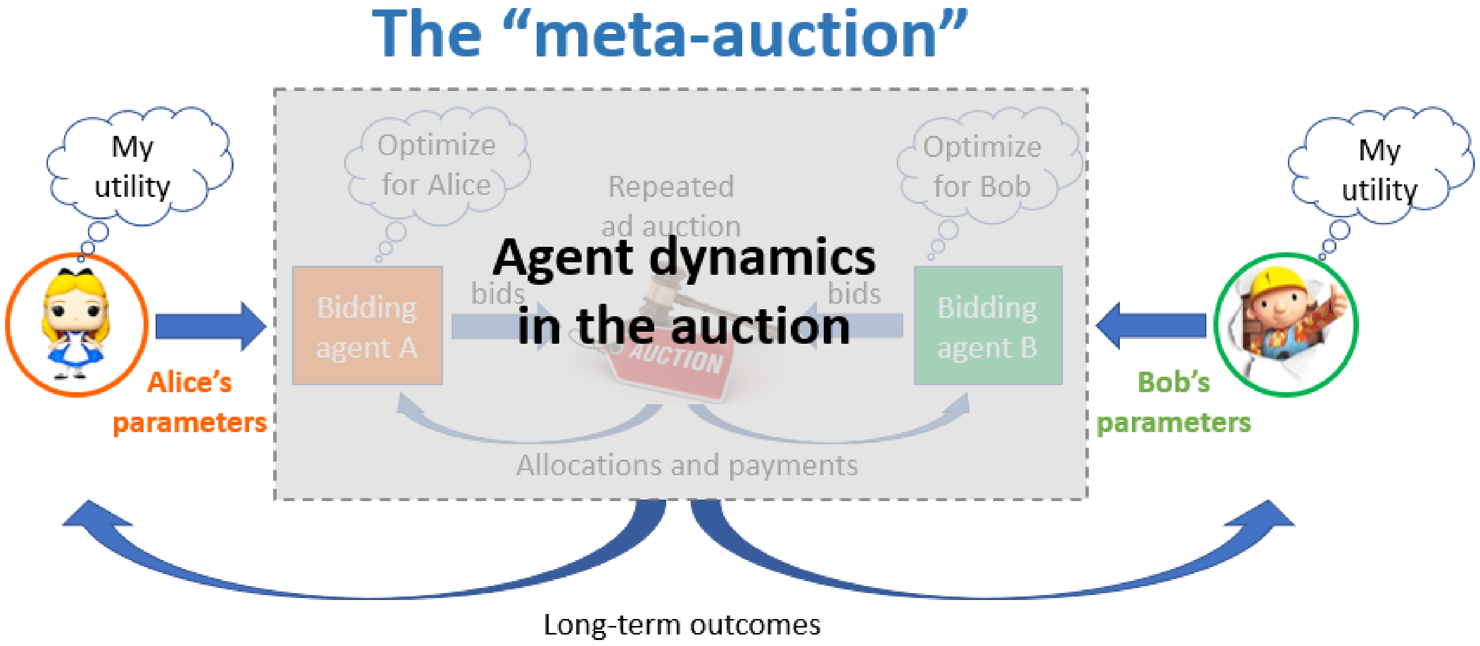

- Regret-Minimizing agents:
  - Approach the set of CCE distributions
  - Need to analyze convergence of average outcomes
  - Generally, convergence is not guaranteed
- Next:
  Auction setting, results for regret-minimizing agents

Basic Auction Formats

Second-price auction rule:

The highest bidder wins; pays to the platform an amount equal to the second highest bid

First-price auction rule:

The highest bidder wins; pays to the platform an amount equal to his own bid

The repeated auction setting:

- A single identical "item" is sold in every auction
- Each bidder has a fixed "value" v for winning the item
- Auction determines the winner and his payment p
- The utility for a bidder is $(v - p)$ if he wins, and zero otherwise
- Utilities are additive over auctions

Learning agents: A bidding agent calculates utility with the value reported by its user

Basic Auction Formats

Expect the high-value agent to win and pay the second price. Right?

Will the agents reach the second-price outcome? (Or not?)

- Second-price auctions
  - Dominant strategy incentive compatible
- First-price auctions
  - Not incentive compatible
  - Equilibrium bids yield the second-price outcome
  - But coarse equilibria can lead to different outcomes [Feldman, Lucier, Nisan, 2016]

Q 1: What are the outcomes when learning agents play these auctions?

Q 2: What value should the users report to their own agents?

Answers?
Need to analyze the dynamics to see what is the meta-game that the users actually play.

Second-Price Auction

Theorem 1 In the limit empirical distribution of MW agents in a second-price auction with values $v > w$, the high agent bids uniformly in $(w, v]$. The low agent bids with full support on $[0, w]$ with monotone density. Thus, the high agent always wins and pays strictly less than the second price.

Player 1 value $= 1$

Player 2 value $= 1/2$

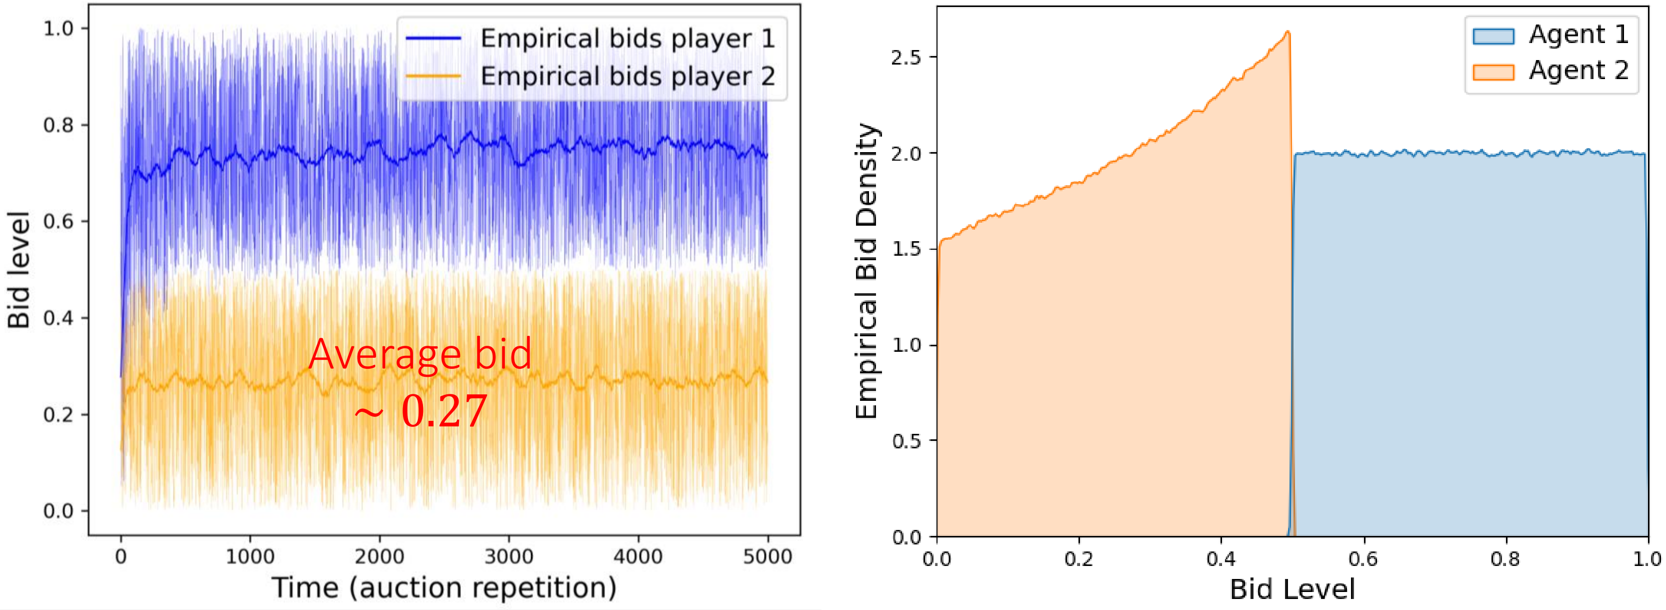

Proof Idea

- For the low agent (with value w):
  - For bids $i < j \leq w$, j dominates i. Hence:
    - Bid distribution is monotone increasing
    - Probability of bidding exactly w remains strictly positive
- For the high agent (with value v, where $v > w$):
  - Bids $\leq w$ yield on expectation strictly lower utility than bids $> w$
  - → Bids $\leq w$ appear only finitely many times (w.h.p.)
- → After a finite number of auctions, the low agent always loses
- → From this point on, the low agent's bid dist. stops changing
- → Low-agent bids $< w$ remain with strictly positive probability

Second-Price Auction

Theorem 1 In the limit empirical distribution of MW agents in a second-price auction with values $v > w$, the high agent bids uniformly in $(w, v]$. The low agent bids with full support on $[0, w]$ with monotone density. Thus, the high agent always wins and pays strictly less than the second price.

Player 1 value $= 1$

Player 2 value $= 1/2$

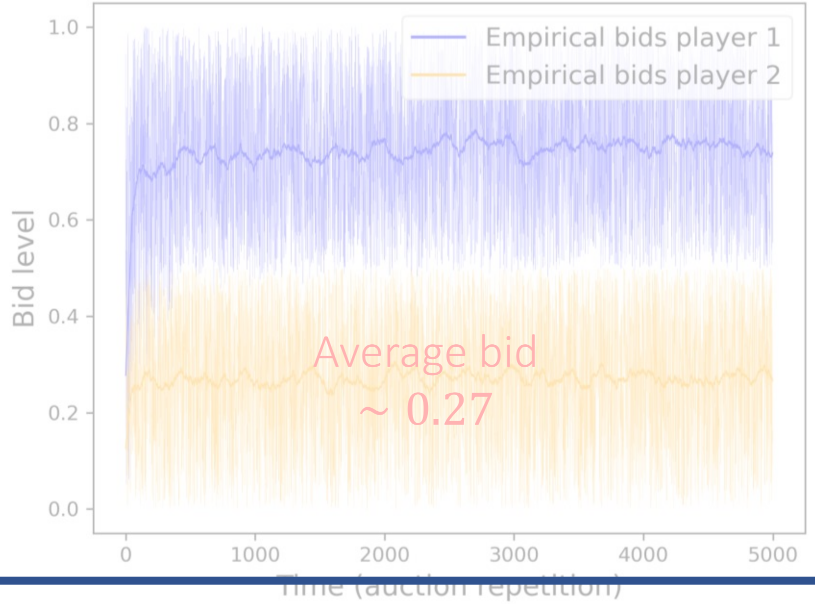

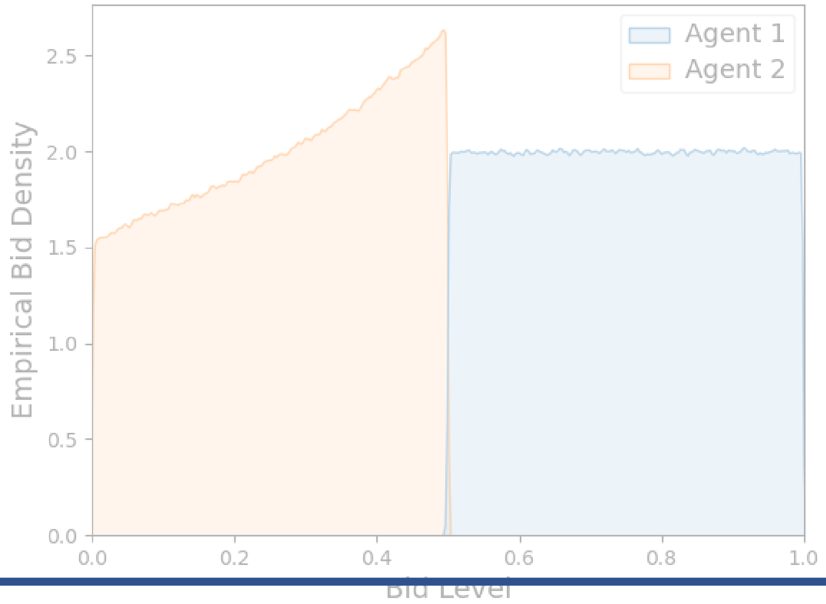

Corollary: From the perspective of the users: The second-price auction with multiplicative-weights agents is not incentive compatible.

The second-price auction with MW agents is not incentive compatible

Example:

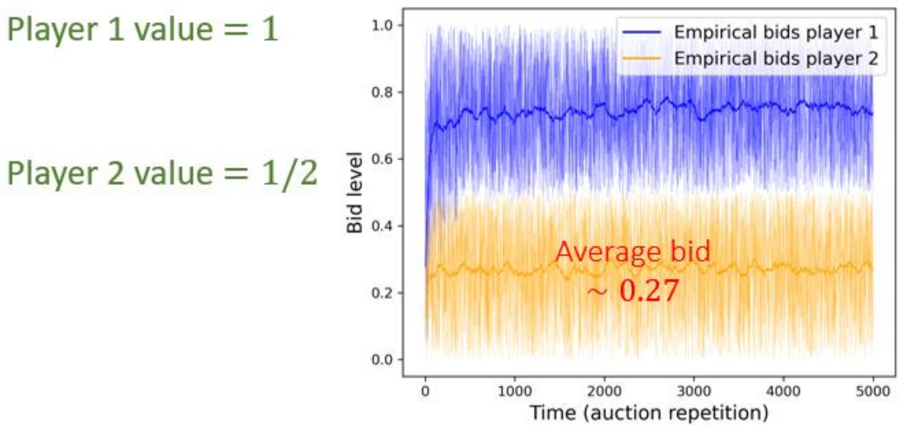

- Alice has a (true) value of 0.4
- Bob has a (true) value of 0.5
- Assume that Bob bids the truth to his agent
- If Alice reports her true value 0.4 by Thm-1 she always looses (has zero utility)
- If Alice manipulates her own agent by misreporting her value as 1, she always wins and has on average a positive utility of $0.4 - 0.27 = 0.13$
- Truthful declarations are not a dominant strategy and not even an equilibrium

Value manipulation in second price auctions

- Alice (player 1): true value = 0.4
- Bob (player 2): true value = 0.5
- Bob declares the truth
- Alice declares x to her agent

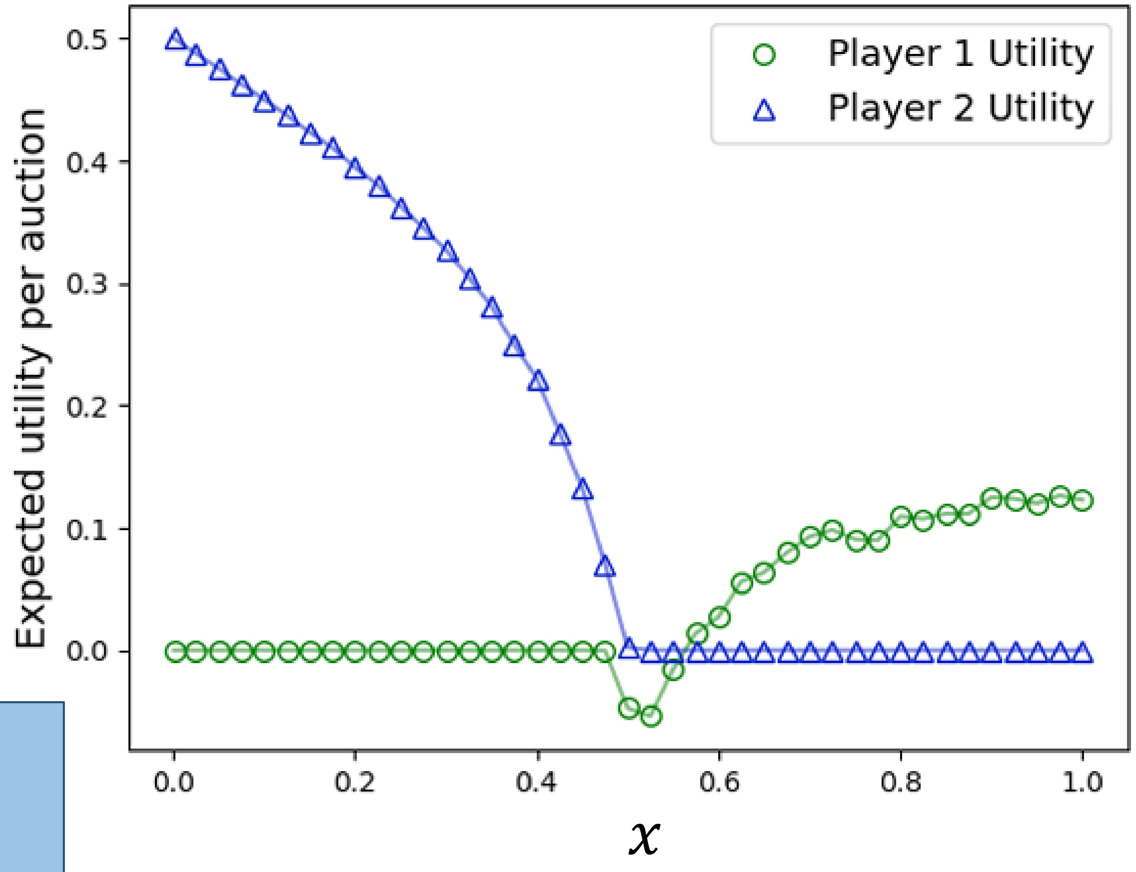

What is the meta-auction equilibrium between the users of these agents?

Meta-game equilibrium

Multiple equilibria:

- One user declares a high value (max allowed)
- Other user declares anything < (2-ε) times the other's true value

→ There exist inefficient equilibria, and

→ Revenue between zero and the first price