Transparency and Fairness in Digital Markets

Nicolò Cesa-Bianchi

Università degli Studi di Milano
Politecnico di Milano

ELiAS

Risks of AI-powered digital platforms

- Digital platforms are increasingly handling economic activities (e.g., trading, advertising, e-commerce)

Risks of AI-powered digital platforms

- Digital platforms are increasingly handling economic activities (e.g., trading, advertising, e-commerce)
- Users rely on data provided by platforms to make decisions

Risks of AI-powered digital platforms

- Digital platforms are increasingly handling economic activities (e.g., trading, advertising, e-commerce)
- Users rely on data provided by platforms to make decisions
- Lack of transparency can affect the quality of these decisions

Risks of AI-powered digital platforms

- Digital platforms are increasingly handling economic activities (e.g., trading, advertising, e-commerce)
- Users rely on data provided by platforms to make decisions
- Lack of transparency can affect the quality of these decisions
- Platforms may increase their utility in ways that are unfair to users

Risks of AI-powered digital platforms

- Digital platforms are increasingly handling economic activities (e.g., trading, advertising, e-commerce)
- Users rely on data provided by platforms to make decisions
- Lack of transparency can affect the quality of these decisions
- Platforms may increase their utility in ways that are unfair to users
- Design algorithms that incentivate transparency and fairness

Risks of AI-powered digital platforms

- Digital platforms are increasingly handling economic activities (e.g., trading, advertising, e-commerce)
- Users rely on data provided by platforms to make decisions
- Lack of transparency can affect the quality of these decisions
- Platforms may increase their utility in ways that are unfair to users
- Design algorithms that incentivate transparency and fairness
- Two case studies: first-price auctions and bilateral trade

Repeated First-Price Auctions

For $t = 1, \ldots, T$

- The environment privately generates valuation V_t and highest competing bid M_t in $[0, 1]$

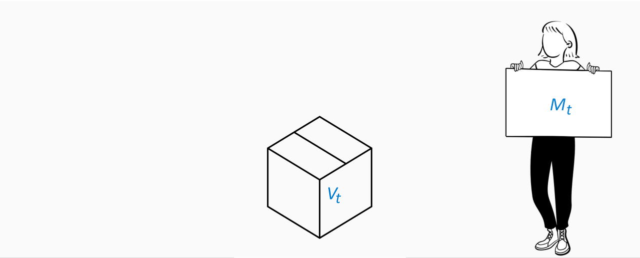

Repeated First-Price Auctions

For $t = 1, \dots, T$

- The environment privately generates valuation V_t and highest competing bid M_t in $[0, 1]$
- The bidder bids B_t without knowing V_t and M_t

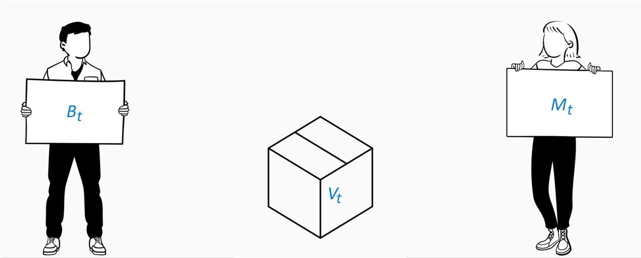

Repeated First-Price Auctions

For $t = 1, \ldots, T$

- The environment privately generates valuation V_t and highest competing bid M_t in $[0, 1]$
- The bidder bids B_t without knowing V_t and M_t
- A first-price auction takes place, and the bidder receives utility

$$U_t(B_t) = (V_t - B_t)\mathbb{I}\{B_t \geq M_t\}$$

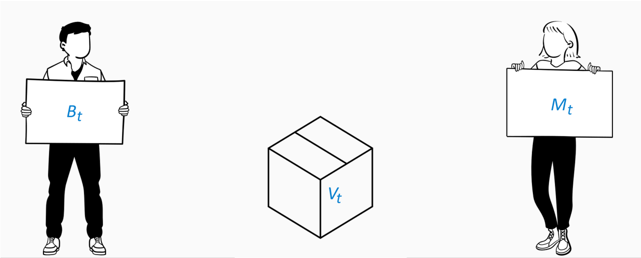

Repeated First-Price Auctions

For $t = 1, \dots, T$

- The environment privately generates valuation V_t and highest competing bid M_t in $[0, 1]$
- The bidder bids B_t without knowing V_t and M_t
- A first-price auction takes place, and the bidder receives utility

$$U_t(B_t) = (V_t - B_t)\mathbb{I}\{B_t \geq M_t\}$$

- The bidder receives some data from the platform (transparency)

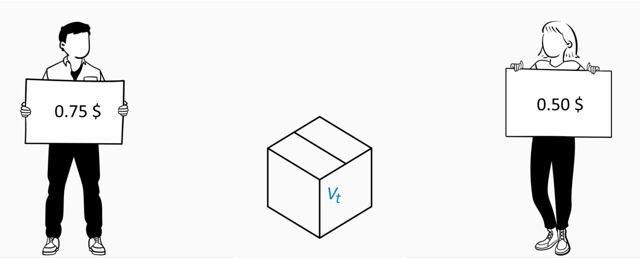

Repeated First-Price Auctions

For $t = 1, \ldots, T$

- The environment privately generates valuation V_t and highest competing bid M_t in $[0, 1]$
- The bidder bids B_t without knowing V_t and M_t
- A first-price auction takes place, and the bidder receives utility

$$U_t(B_t) = (V_t - B_t)\mathbb{I}\{B_t \geq M_t\}$$

- The bidder receives some data from the platform (transparency)

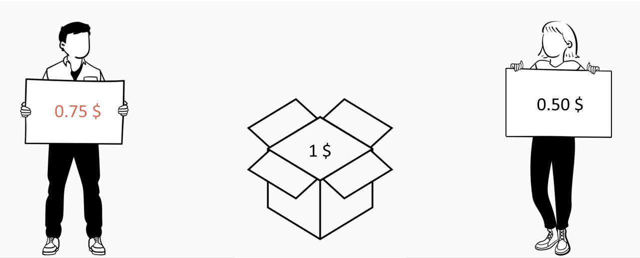

Repeated First-Price Auctions

For $t = 1, \ldots, T$

- The environment privately generates valuation V_t and highest competing bid M_t in $[0, 1]$
- The bidder bids B_t without knowing V_t and M_t
- A first-price auction takes place, and the bidder receives utility

$$U_t(B_t) = (V_t - B_t)\mathbb{I}\{B_t \geq M_t\}$$

- The bidder receives some data from the platform (transparency)

Regret

Goal: Minimize the regret with respect to the best fixed bid in hindsight

$$R_T = \max_{b \in [0,1]} \mathbb{E}\left[\sum_{t=1}^{T} U_t(b) - \sum_{t=1}^{T} U_t(B_t)\right]$$

Regret

Goal: Minimize the regret with respect to the best fixed bid in hindsight

$$R_T = \max_{b \in [0,1]} \mathbb{E}\left[\sum_{t=1}^{T} U_t(b) - \sum_{t=1}^{T} U_t(B_t)\right]$$

- The per-round utility $U_t \in [-1, 1]$

Regret

Goal: Minimize the regret with respect to the best fixed bid in hindsight

$$R_T = \max_{b \in [0,1]} \mathbb{E}\left[\sum_{t=1}^{T} U_t(b) - \sum_{t=1}^{T} U_t(B_t)\right]$$

- The per-round utility $U_t \in [-1, 1]$
- The total utility is $\mathcal{O}(T)$

Regret

Goal: Minimize the regret with respect to the best fixed bid in hindsight

$$R_T = \max_{b \in [0,1]} \mathbb{E}\left[\sum_{t=1}^{T} U_t(b) - \sum_{t=1}^{T} U_t(B_t)\right]$$

- The per-round utility $U_t \in [-1, 1]$
- The total utility is $\mathcal{O}(T)$
- We want R_T to grow sublinearly in the time horizon T

Regret

Goal: Minimize the regret with respect to the best fixed bid in hindsight

$$R_T = \max_{b \in [0,1]} \mathbb{E}\left[\sum_{t=1}^{T} U_t(b) - \sum_{t=1}^{T} U_t(B_t)\right]$$

- The per-round utility $U_t \in [-1, 1]$
- The total utility is $\mathcal{O}(T)$
- We want R_T to grow sublinearly in the time horizon T
- So that the average utility converges to that of the benchmark

Feedback Models

Feedback depends on the transparency of the platform

	HIGHEST COMPETING BID	VALUATION
Full	Always observed	
Transparent	Always observed	Observed if auction is won
Semi-Transparent	Observed if auction is lost	
Bandit	Never observed	

Environments

- Stochastic model: (V_t, M_t) drawn i.i.d. from a fixed but unknown distribution

Environments

- Stochastic model: (V_t, M_t) drawn i.i.d. from a fixed but unknown distribution
- Adversarial model: (V_t, M_t) generated by an oblivious adversary

Environments

- Stochastic model: (V_t, M_t) drawn i.i.d. from a fixed but unknown distribution
- Adversarial model: (V_t, M_t) generated by an oblivious adversary
- Smoothed models:
  - Smoothed adversarial: each (V_t, M_t) drawn from a different distribution with bounded density
  - Smoothed stochastic: (V_t, M_t) drawn i.i.d. from a distribution with bounded density

Results

	Stochastic i.i.d.		Adversarial	
	Smooth	General	Smooth	General
V_t always, M_t always	$\Omega(\sqrt{T})$			$\Omega(T)$
V_t if won, M_t always		$\mathcal{O}(\sqrt{T})$	$\widetilde{\mathcal{O}}(\sqrt{T})$	
V_t if won, M_t if lost	$\Omega(T^{2/3})$	$\widetilde{\mathcal{O}}(T^{2/3})$		
V_t if won, M_t never		$\Omega(T)$	$\mathcal{O}(T^{2/3})$	

- V_t is valuation and M_t is highest competing bid
- Full feedback and transparent feedback are equivalent
- Semi-transparent feedback and bandit feedback are almost equivalent
- Disclosing other bids significantly decreases the bidder's regret

Bilateral trading (e.g., Uber/Lyft app)

1. Seller and buyer arrive with private S_t and B_t

Bilateral trading (e.g., Uber/Lyft app)

1. Seller and buyer arrive with private S_t and B_t
2. The platform posts a price p_t

Bilateral trading (e.g., Uber/Lyft app)

1. Seller and buyer arrive with private S_t and B_t
2. The platform posts a price p_t
3. If $S_t \leq p_t \leq B_t$, then the trade happens

Bilateral trading (e.g., Uber/Lyft app)

1. Seller and buyer arrive with private S_t and B_t
2. The platform posts a price p_t
3. If $S_t \leq p_t \leq B_t$, then the trade happens
4. Feedback: YES if trade happened, NO otherwise

Bilateral trading (e.g., Uber/Lyft app)

1. Seller and buyer arrive with private S_t and B_t
2. The platform posts a price p_t
3. If $S_t \le p_t \le B_t$, then the trade happens
4. Feedback: YES if trade happened, NO otherwise

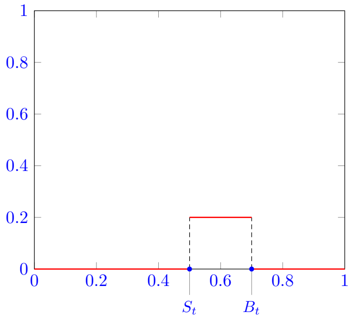

Gain from Trade

$$\mathrm{GFT}_t(p) = \Big(\underbrace{(B_t - p)}_{\text{buyer's net gain}} + \underbrace{(p - S_t)}_{\text{seller's net gain}} \Big) \underbrace{\mathbb{I}\{S_t \le p \le B_t\}}_{\text{trade happens}}$$

Bilateral trading (e.g., Uber/Lyft app)

1. Seller and buyer arrive with private S_t and B_t
2. The platform posts a price p_t
3. If $S_t \leq p_t \leq B_t$, then the trade happens
4. Feedback: YES if trade happened, NO otherwise

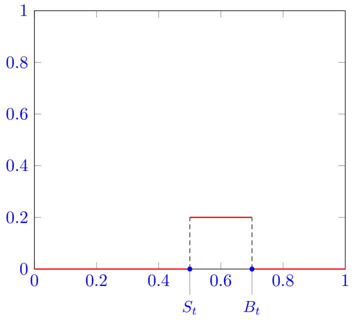

Gain from Trade

$$\mathrm{GFT}_t(p) = \Big(\underbrace{(B_t - p)}_{\text{buyer's net gain}} + \underbrace{(p - S_t)}_{\text{seller's net gain}} \Big) \underbrace{\mathbb{I}\{S_t \leq p \leq B_t\}}_{\text{trade happens}} = (B_t - S_t)\mathbb{I}\{S_t \leq p \leq B_t\}$$

Bilateral trading (e.g., Uber/Lyft app)

1. Seller and buyer arrive with private S_t and B_t
2. The platform posts a price p_t
3. If $S_t \le p_t \le B_t$, then the trade happens
4. Feedback: YES if trade happened, NO otherwise

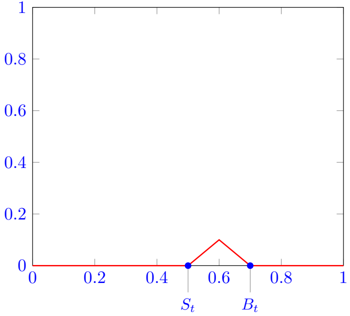

Gain from Trade

$$\mathrm{GFT}_t(p) = \Big(\underbrace{(B_t - p)}_{\text{buyer's net gain}} + \underbrace{(p - S_t)}_{\text{seller's net gain}} \Big) \underbrace{\mathbb{I}\{S_t \le p \le B_t\}}_{\text{trade happens}} = \boxed{(B_t - S_t)\mathbb{I}\{S_t \le p \le B_t\}}$$

Fair Gain from Trade

$$\mathrm{FGFT}_t(p) = \min\Big\{[B_t - p]_+, [p - S_t]_+\Big\}$$

Bilateral trading (e.g., Uber/Lyft app)

1. Seller and buyer arrive with private S_t and B_t
2. The platform posts a price p_t
3. If $S_t \leq p_t \leq B_t$, then the trade happens
4. Feedback: YES if trade happened, NO otherwise

Neither GFT_t nor FGFT_t are observed

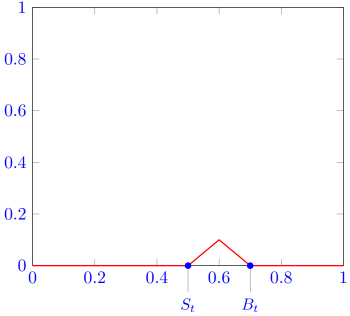

Gain from Trade

$$\mathrm{GFT}_t(p) = \Big(\underbrace{(B_t - p)}_{\text{buyer's net gain}} + \underbrace{(p - S_t)}_{\text{seller's net gain}} \Big) \underbrace{\mathbb{I}\{S_t \leq p \leq B_t\}}_{\text{trade happens}} = \boxed{(B_t - S_t)\mathbb{I}\{S_t \leq p \leq B_t\}}$$

Fair Gain from Trade

$$\mathrm{FGFT}_t(p) = \min\Big\{[B_t - p]_+, [p - S_t]_+\Big\}$$

Results

Goal: Minimize the regret with respect to the best fixed price in hindsight

Regret from GFT: $$\max_{p\in[0,1]} \mathbb{E}\left[\sum_{t=1}^{T} \mathrm{GFT}_t(p) - \sum_{t=1}^{T} \mathrm{GFT}_t(p_t)\right]$$

Regret from FGFT: $$\max_{p\in[0,1]} \mathbb{E}\left[\sum_{t=1}^{T} \mathrm{FGFT}_t(p) - \sum_{t=1}^{T} \mathrm{FGFT}_t(p_t)\right]$$

Results

Goal: Minimize the regret with respect to the best fixed price in hindsight

$$\text{Regret from GFT:} \qquad \max_{p\in[0,1]} \mathbb{E}\left[\sum_{t=1}^{T} \mathrm{GFT}_t(p) - \sum_{t=1}^{T} \mathrm{GFT}_t(p_t)\right]$$

$$\text{Regret from FGFT:} \qquad \max_{p\in[0,1]} \mathbb{E}\left[\sum_{t=1}^{T} \mathrm{FGFT}_t(p) - \sum_{t=1}^{T} \mathrm{FGFT}_t(p_t)\right]$$

- $\mathrm{FGFT} \le \frac{1}{2}\mathrm{GFT}$ and $\max_p \mathrm{FGFT}(p) = \dfrac{1}{2} \max_p \mathrm{GFT}(p)$

Results

Goal: Minimize the regret with respect to the best fixed price in hindsight

$$\text{Regret from GFT:} \qquad \max_{p\in[0,1]} \mathbb{E}\left[\sum_{t=1}^{T} \mathrm{GFT}_t(p) - \sum_{t=1}^{T} \mathrm{GFT}_t(p_t)\right]$$

$$\text{Regret from FGFT:} \qquad \max_{p\in[0,1]} \mathbb{E}\left[\sum_{t=1}^{T} \mathrm{FGFT}_t(p) - \sum_{t=1}^{T} \mathrm{FGFT}_t(p_t)\right]$$

- $\mathrm{FGFT} \le \frac{1}{2}\mathrm{GFT}$ and $\max_p \mathrm{FGFT}(p) = \dfrac{1}{2} \max_p \mathrm{GFT}(p)$
- Regret-minimizing algorithms for GFT may suffer linear regret on FGFT

Results

Goal: Minimize the regret with respect to the best fixed price in hindsight

$$\text{Regret from GFT:} \qquad \max_{p\in[0,1]} \mathbb{E}\left[\sum_{t=1}^{T} \mathrm{GFT}_t(p) - \sum_{t=1}^{T} \mathrm{GFT}_t(p_t)\right]$$

$$\text{Regret from FGFT:} \qquad \max_{p\in[0,1]} \mathbb{E}\left[\sum_{t=1}^{T} \mathrm{FGFT}_t(p) - \sum_{t=1}^{T} \mathrm{FGFT}_t(p_t)\right]$$

- $\mathrm{FGFT} \le \frac{1}{2}\mathrm{GFT}$ and $\max_p \mathrm{FGFT}(p) = \dfrac{1}{2}\max_p \mathrm{GFT}(p)$
- Regret-minimizing algorithms for GFT may suffer linear regret on FGFT
- We show that optimal regret rate is $\widetilde{\Theta}(T^{2/3})$ in both cases

Results

Goal: Minimize the regret with respect to the best fixed price in hindsight

$$\text{Regret from GFT:} \qquad \max_{p\in[0,1]} \mathbb{E}\left[\sum_{t=1}^{T} \mathrm{GFT}_t(p) - \sum_{t=1}^{T} \mathrm{GFT}_t(p_t)\right]$$

$$\text{Regret from FGFT:} \qquad \max_{p\in[0,1]} \mathbb{E}\left[\sum_{t=1}^{T} \mathrm{FGFT}_t(p) - \sum_{t=1}^{T} \mathrm{FGFT}_t(p_t)\right]$$

- $\mathrm{FGFT} \le \frac{1}{2}\mathrm{GFT}$ and $\max_p \mathrm{FGFT}(p) = \dfrac{1}{2} \max_p \mathrm{GFT}(p)$
- Regret-minimizing algorithms for GFT may suffer linear regret on FGFT
- We show that optimal regret rate is $\widetilde{\Theta}(T^{2/3})$ in both cases
- Achieved (by different algorithms) when (S_t, B_t) are independent and i.i.d.

Results

Goal: Minimize the regret with respect to the best fixed price in hindsight

$$\text{Regret from GFT:} \qquad \max_{p\in[0,1]} \mathbb{E}\left[\sum_{t=1}^{T} \mathrm{GFT}_t(p) - \sum_{t=1}^{T} \mathrm{GFT}_t(p_t)\right]$$

$$\text{Regret from FGFT:} \qquad \max_{p\in[0,1]} \mathbb{E}\left[\sum_{t=1}^{T} \mathrm{FGFT}_t(p) - \sum_{t=1}^{T} \mathrm{FGFT}_t(p_t)\right]$$

- $\mathrm{FGFT} \le \frac{1}{2}\mathrm{GFT}$ and $\max_p \mathrm{FGFT}(p) = \dfrac{1}{2}\max_p \mathrm{GFT}(p)$
- Regret-minimizing algorithms for GFT may suffer linear regret on FGFT
- We show that optimal regret rate is $\widetilde{\Theta}(T^{2/3})$ in both cases
- Achieved (by different algorithms) when (S_t, B_t) are independent and i.i.d.
- GFT is not Lipschitz: this requires smoothness of the seller/buyer distributions

Additional results

- Posted price auctions: Maximize social welfare (seller's revenue + unobserved buyer's utility)
- Financial markets: Maximize utility of liquidity providers (market makers)