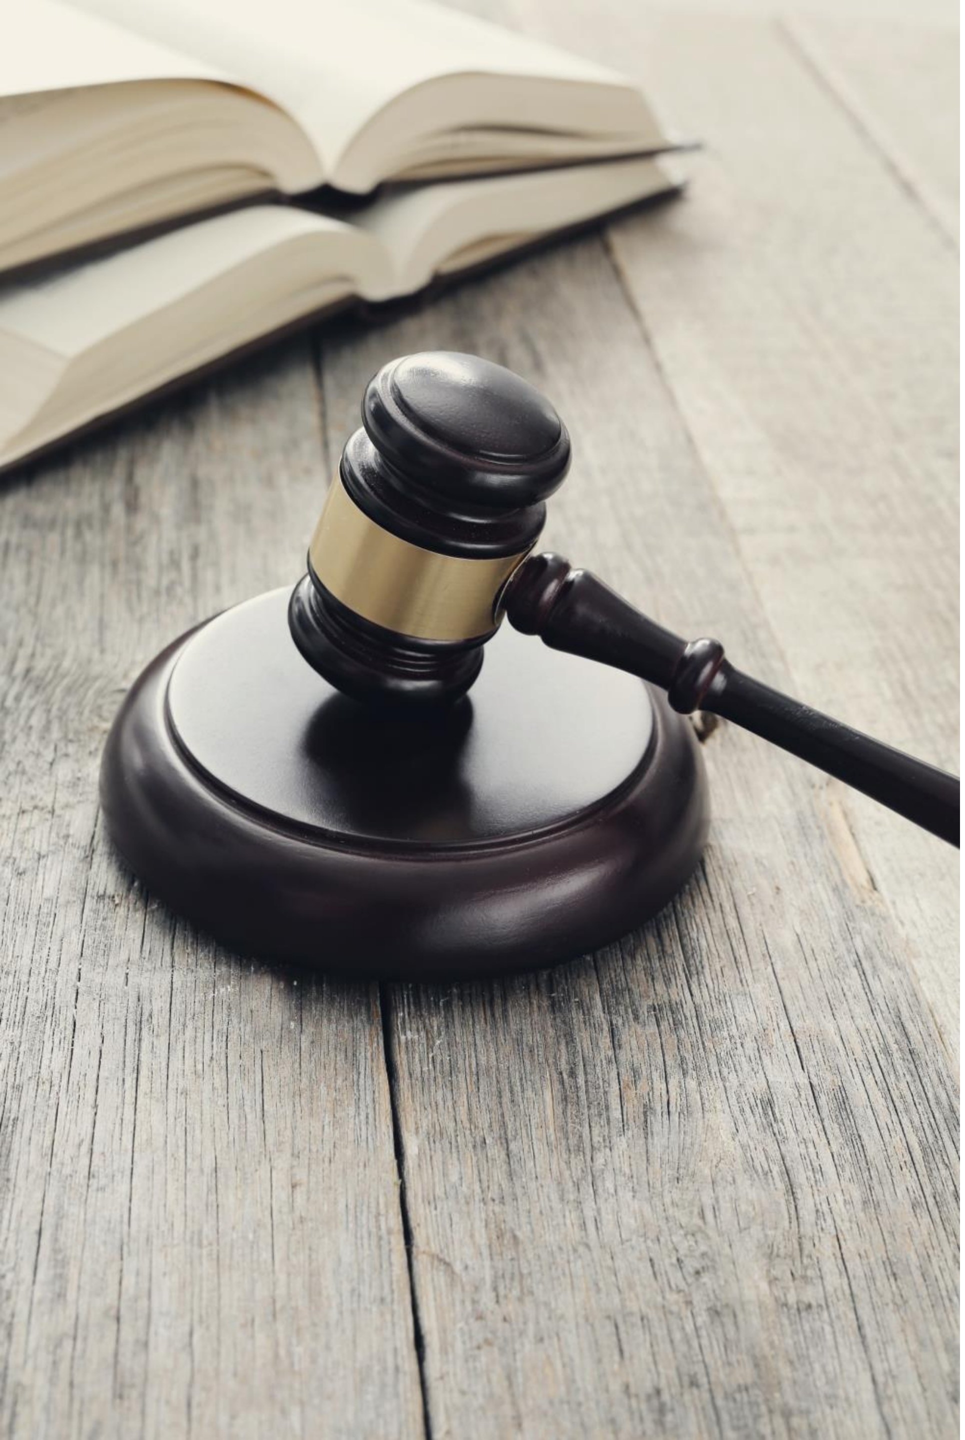

The Role of Transparency in Repeated 1st-Price Auctions with Unknown Valuations*

Nicolò Cesa-Bianchi

Tom Cesari

Roberto Colomboni

Federico Fusco

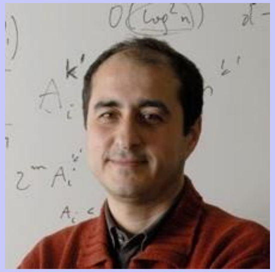

Stefano Leonardi

Workshop on Learning in Games
July 1–3, 2024

*STOC 2024

Repeated First-Price Auctions

We study repeated first-price auctions:

- An online learning framework

Repeated First-Price Auctions

We study repeated first-price auctions:

- An online learning framework
- From the bidder's perspective

Repeated First-Price Auctions

We study repeated first-price auctions:

- An online learning framework
- From the bidder's perspective
- The valuation is unknown

Motivation: Real-Time Bidding

Motivation: Real-Time Bidding

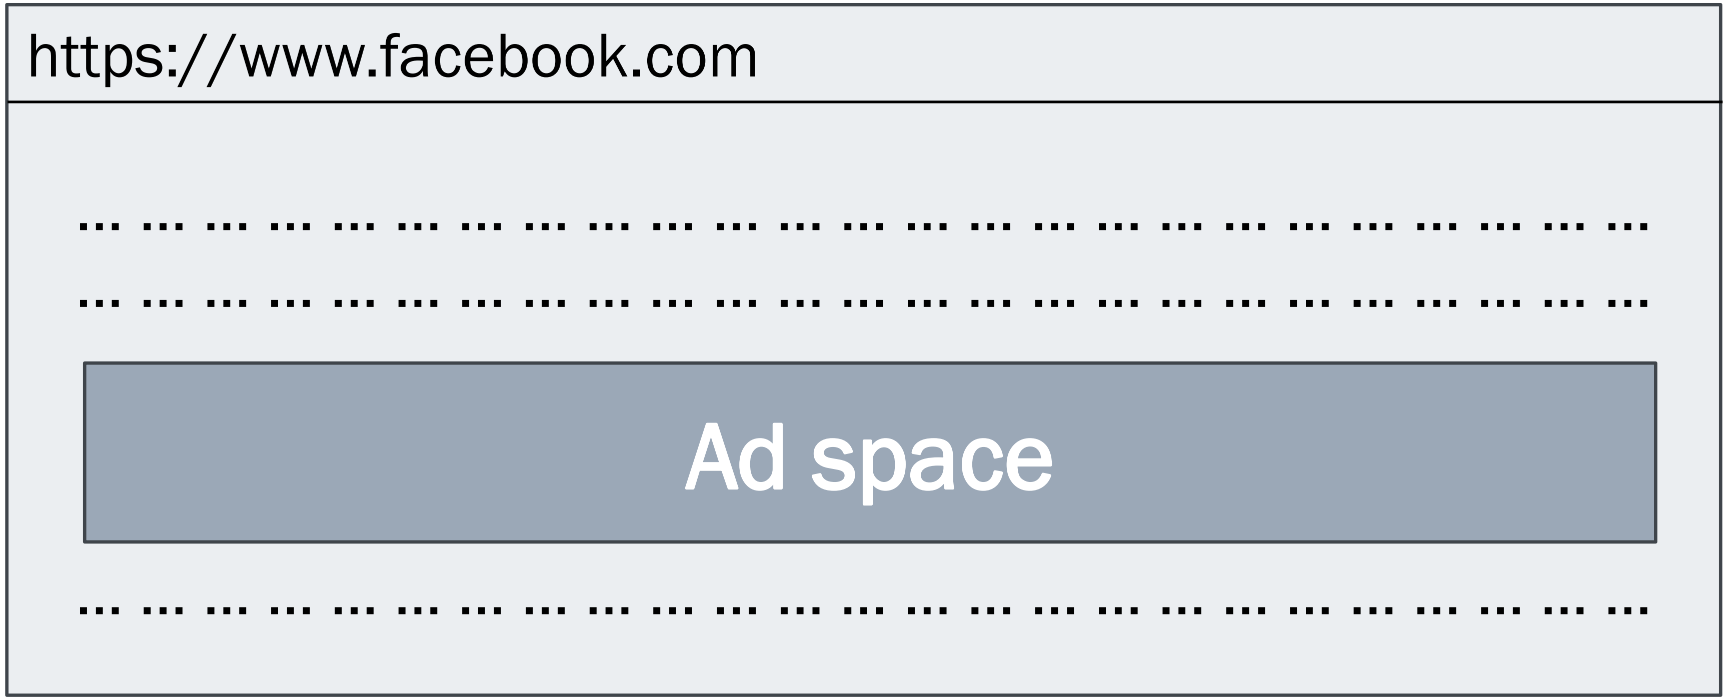

Motivation: Real-Time Bidding

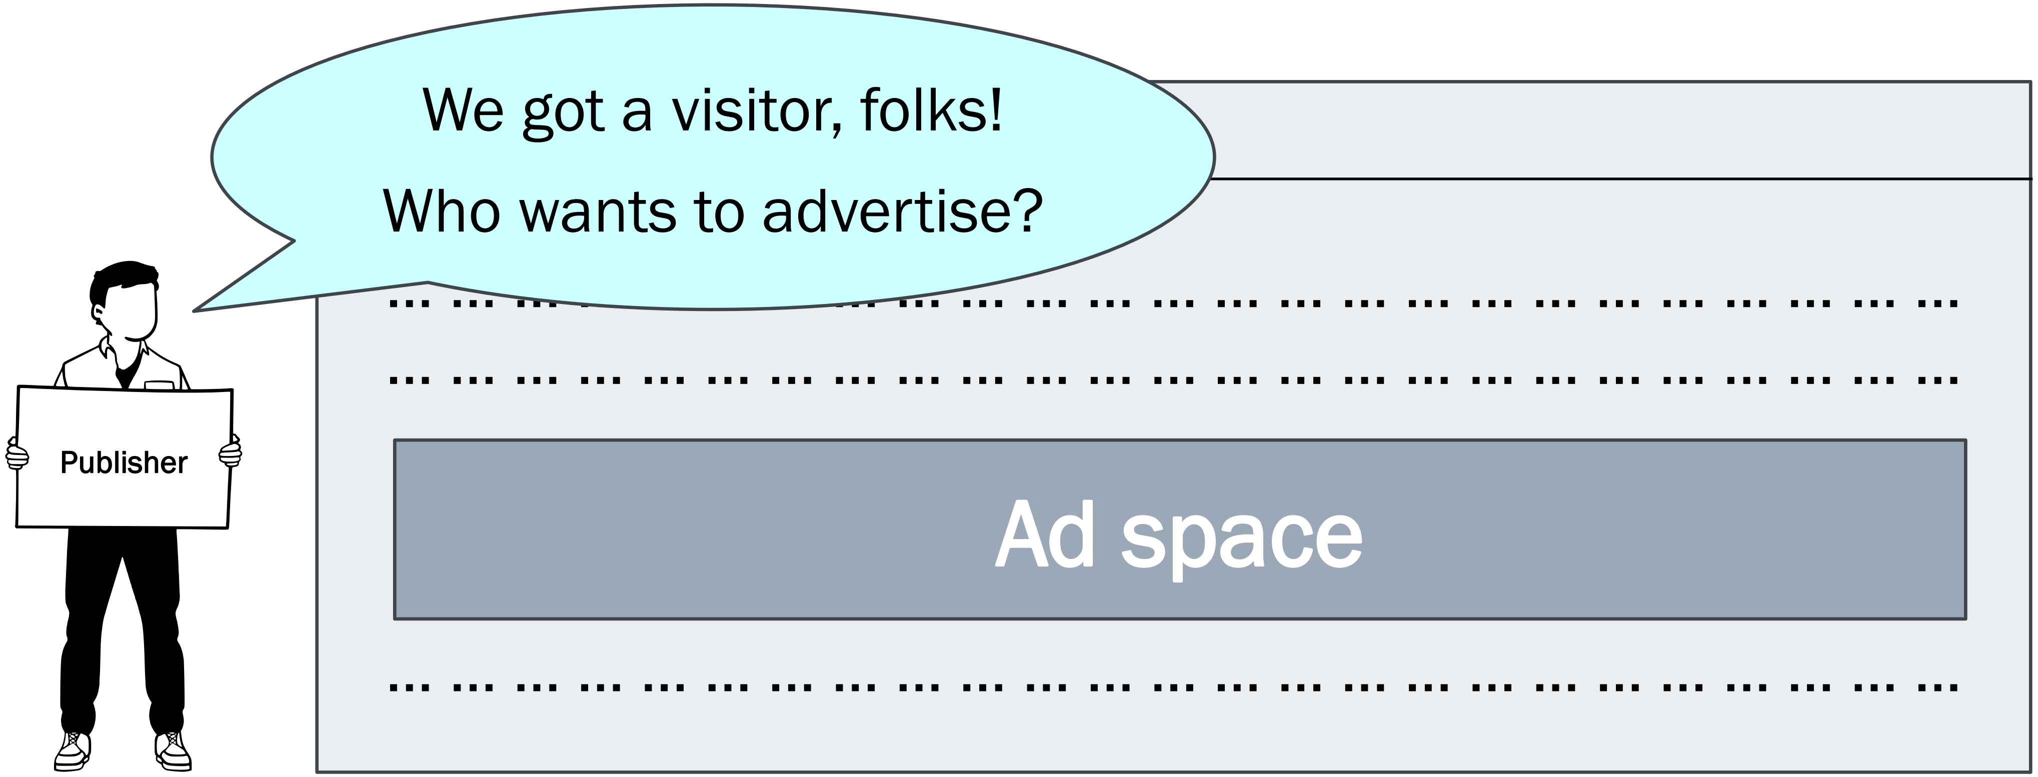

Motivation: Real-Time Bidding

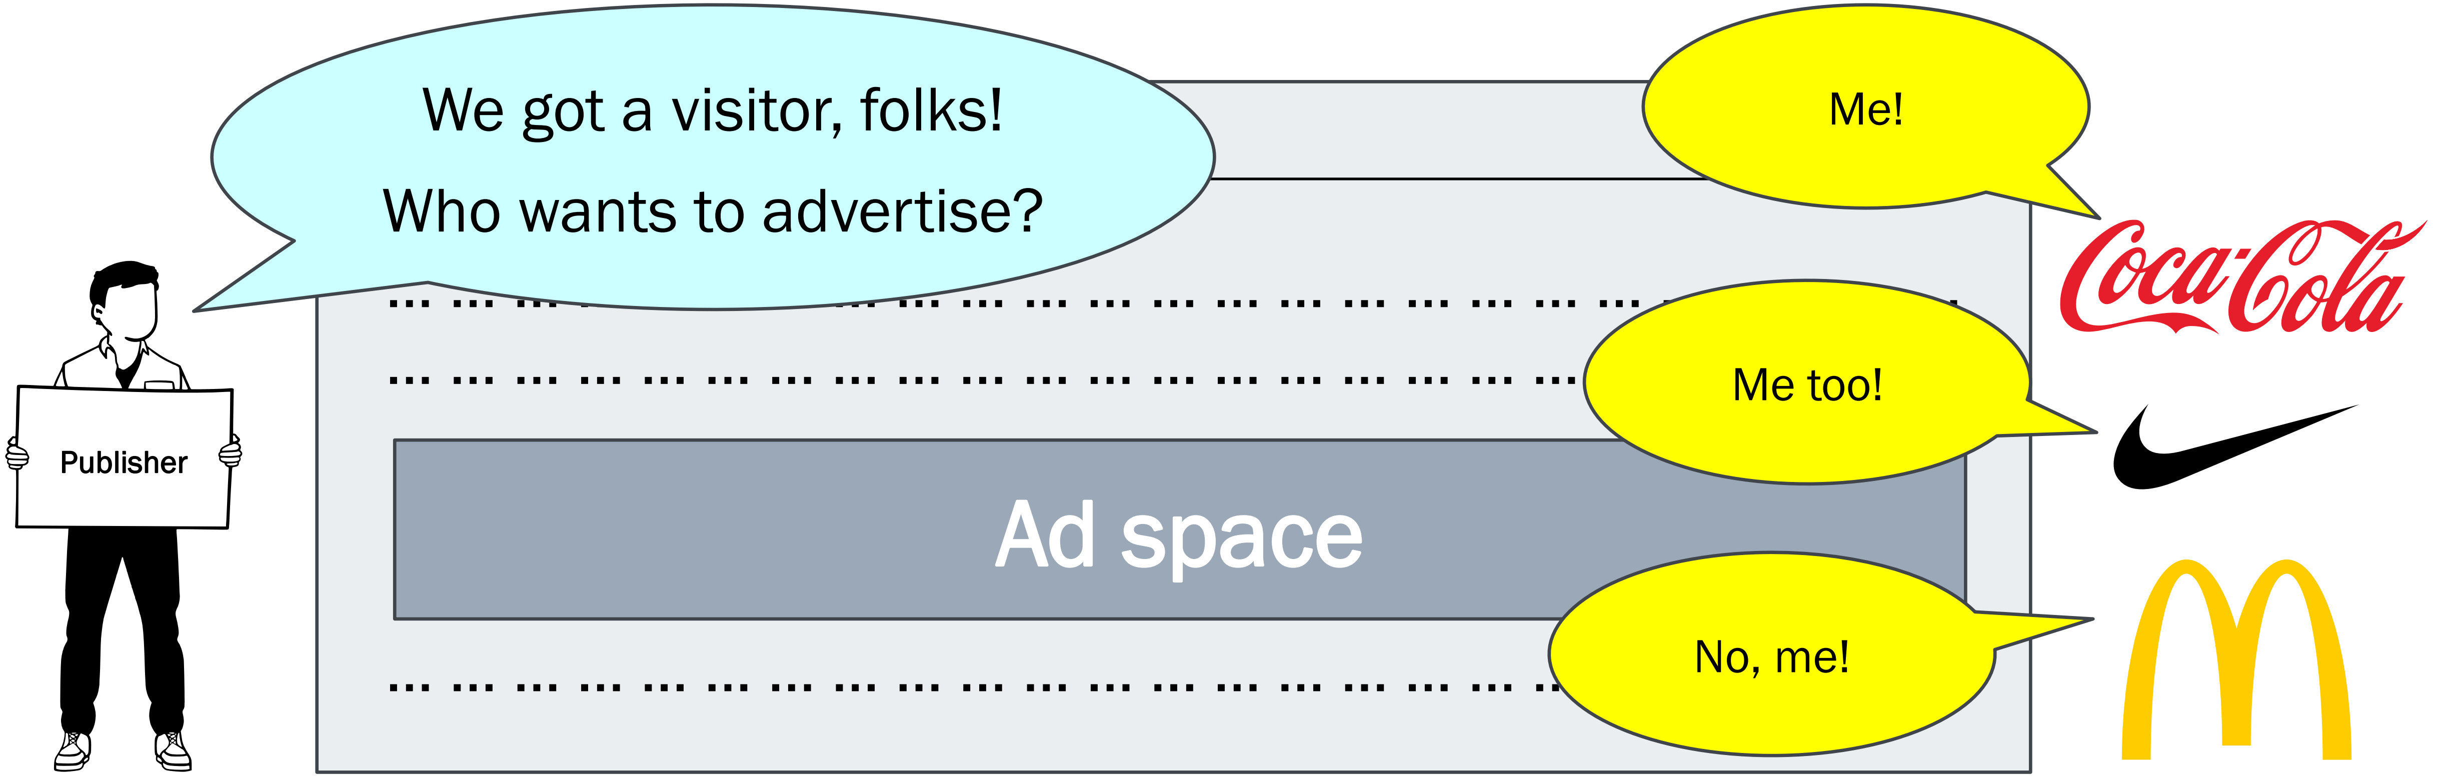

Motivation: Real-Time Bidding

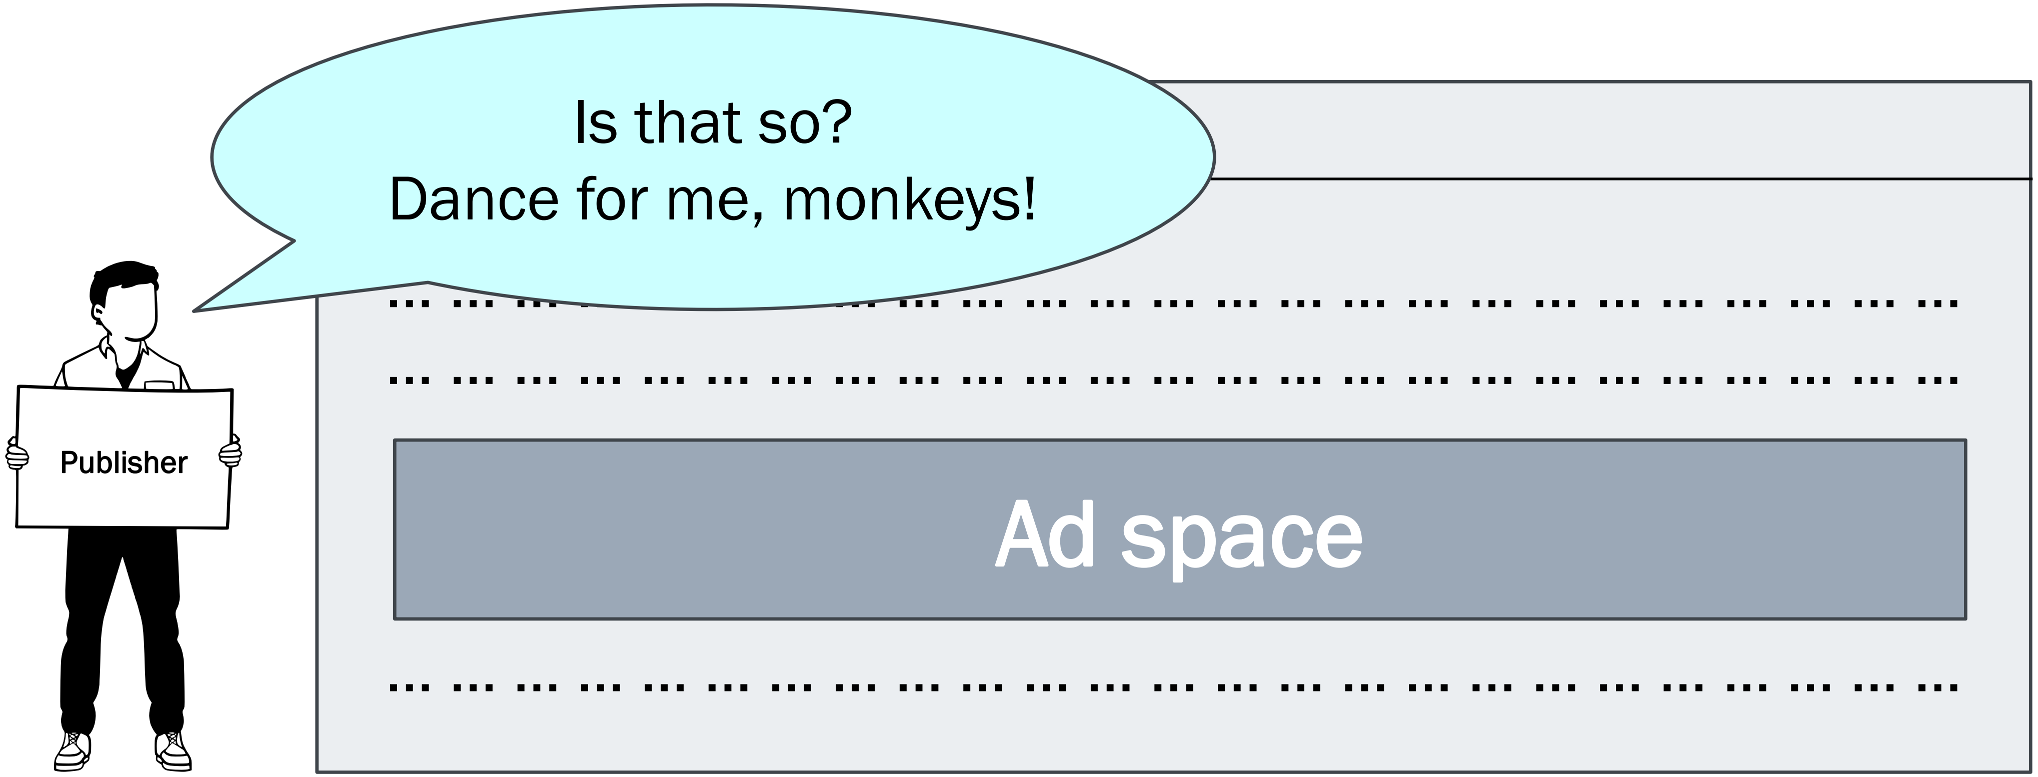

Motivation: Real-Time Bidding

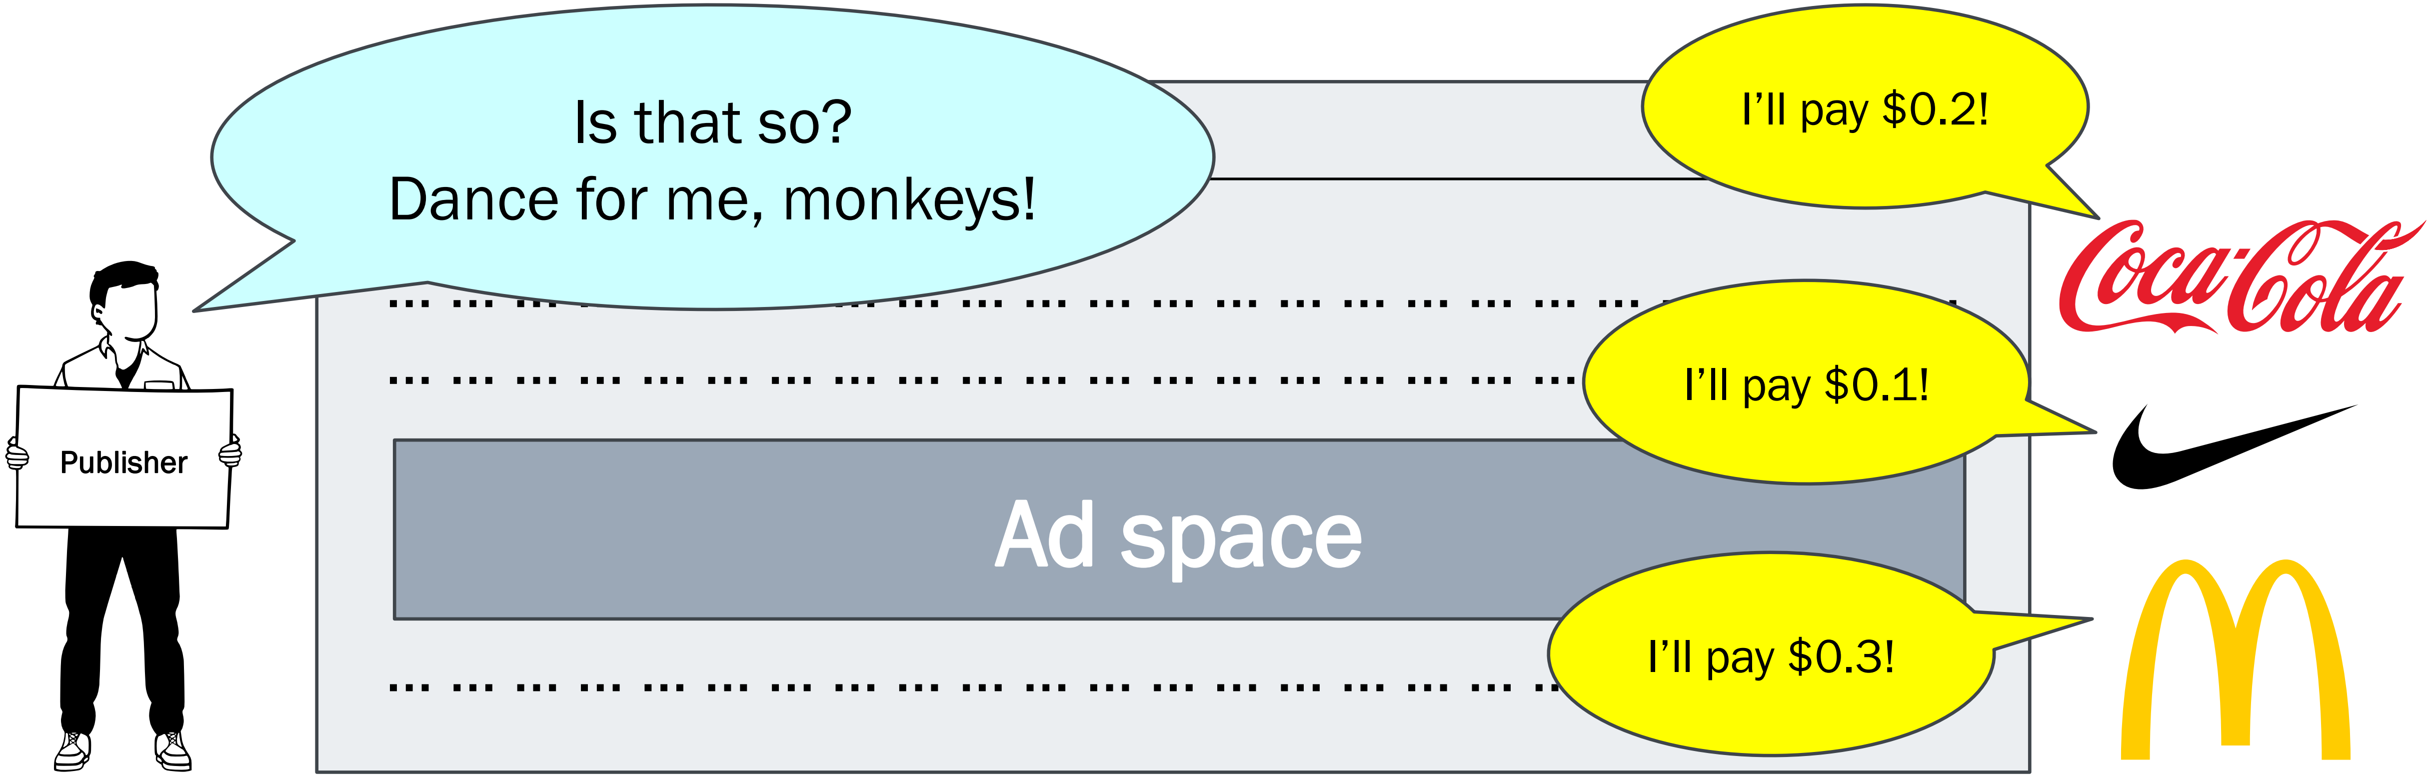

Motivation: Real-Time Bidding

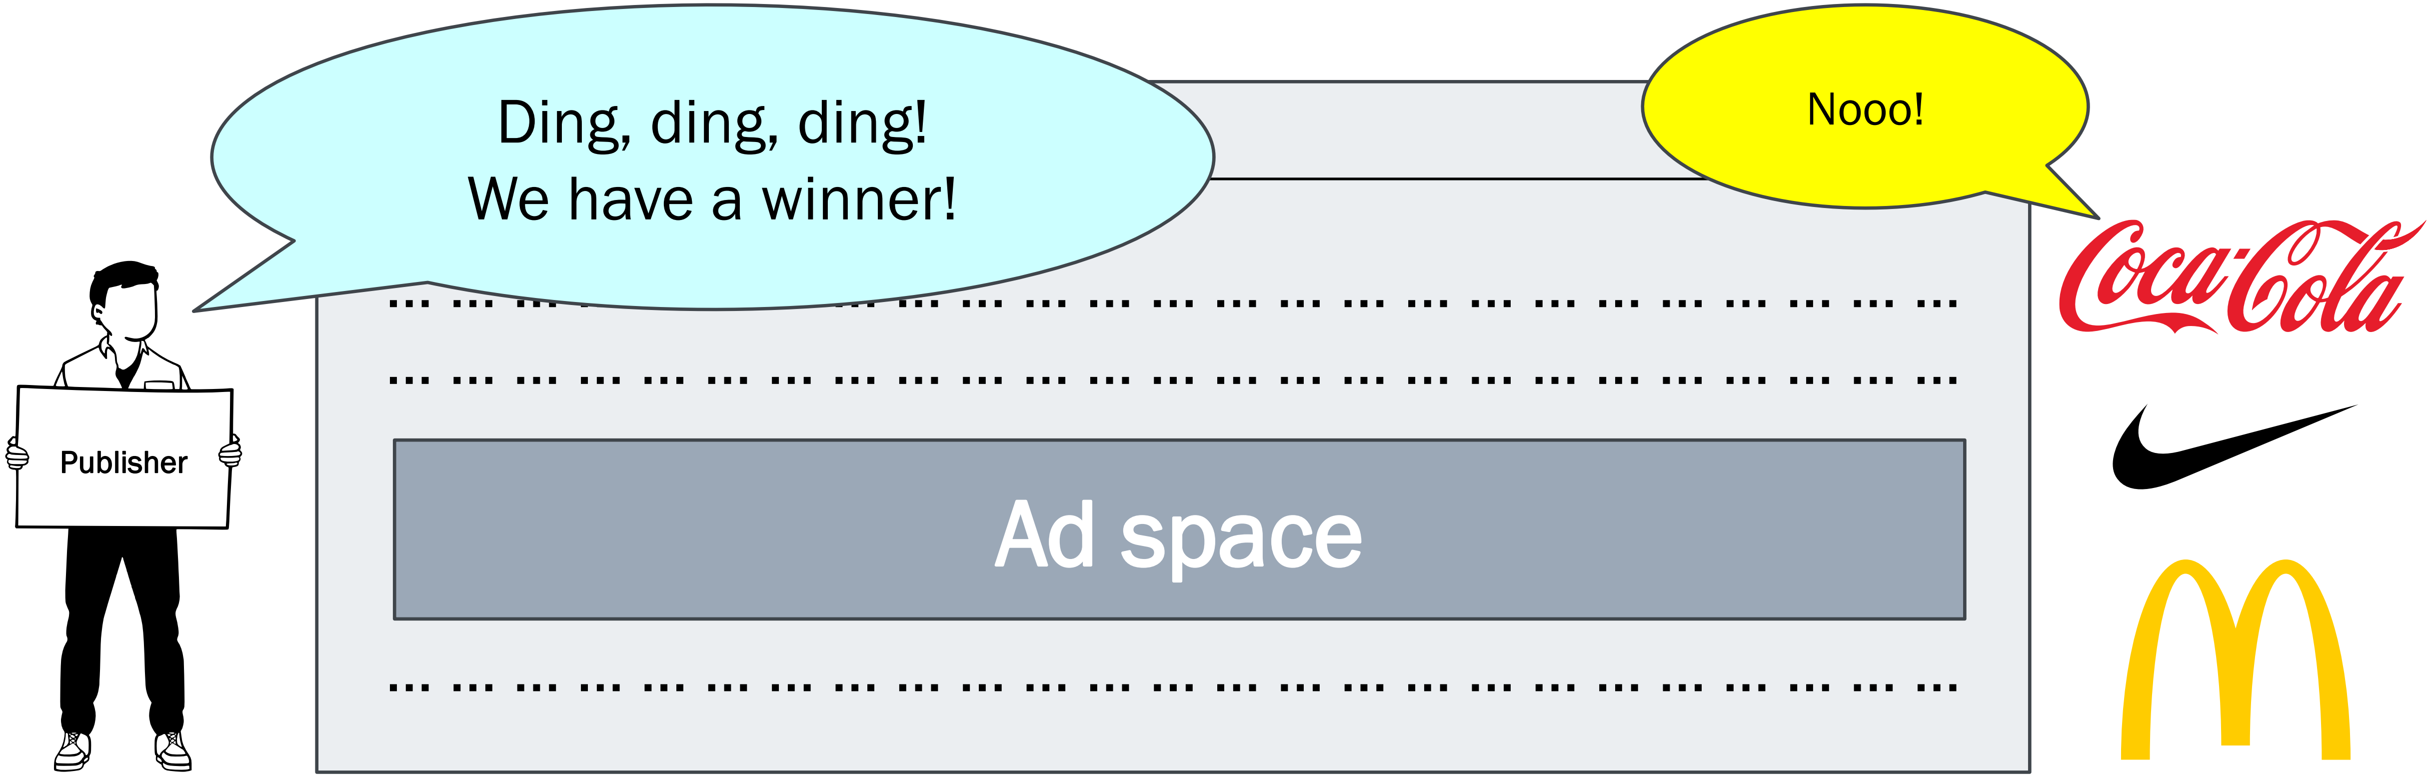

Motivation: Real-Time Bidding

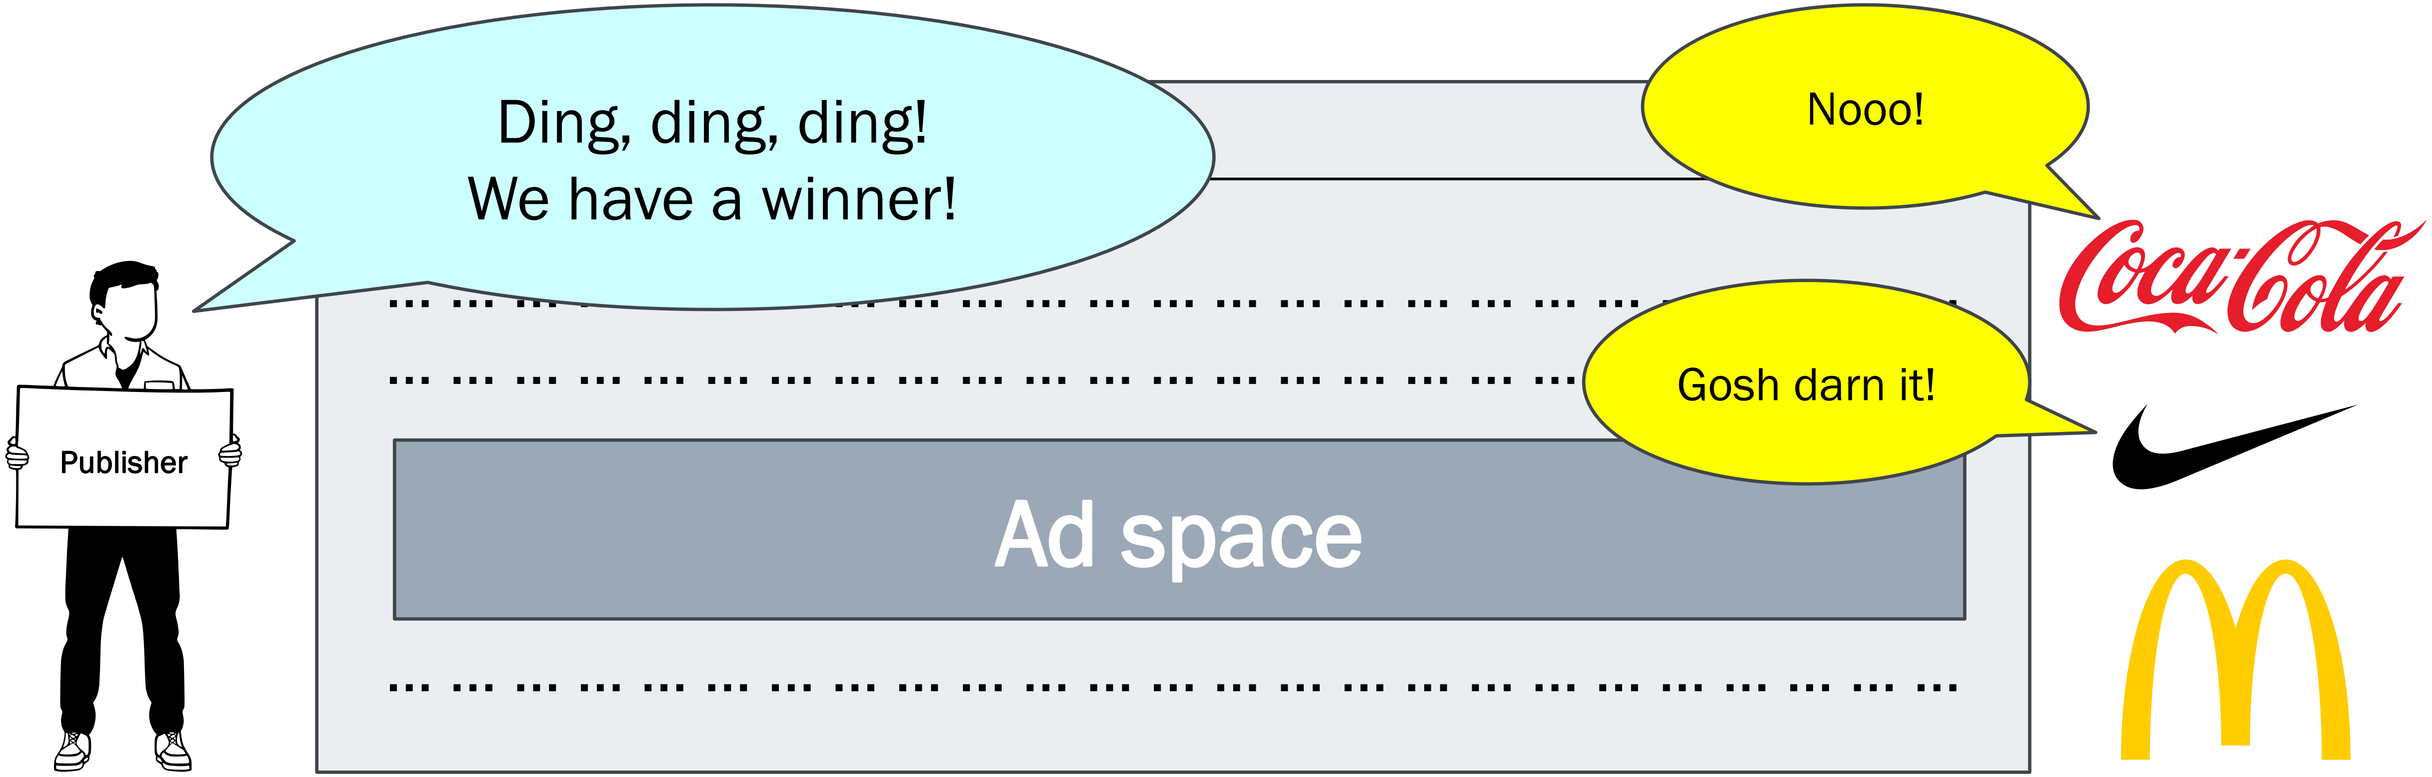

Motivation: Real-Time Bidding

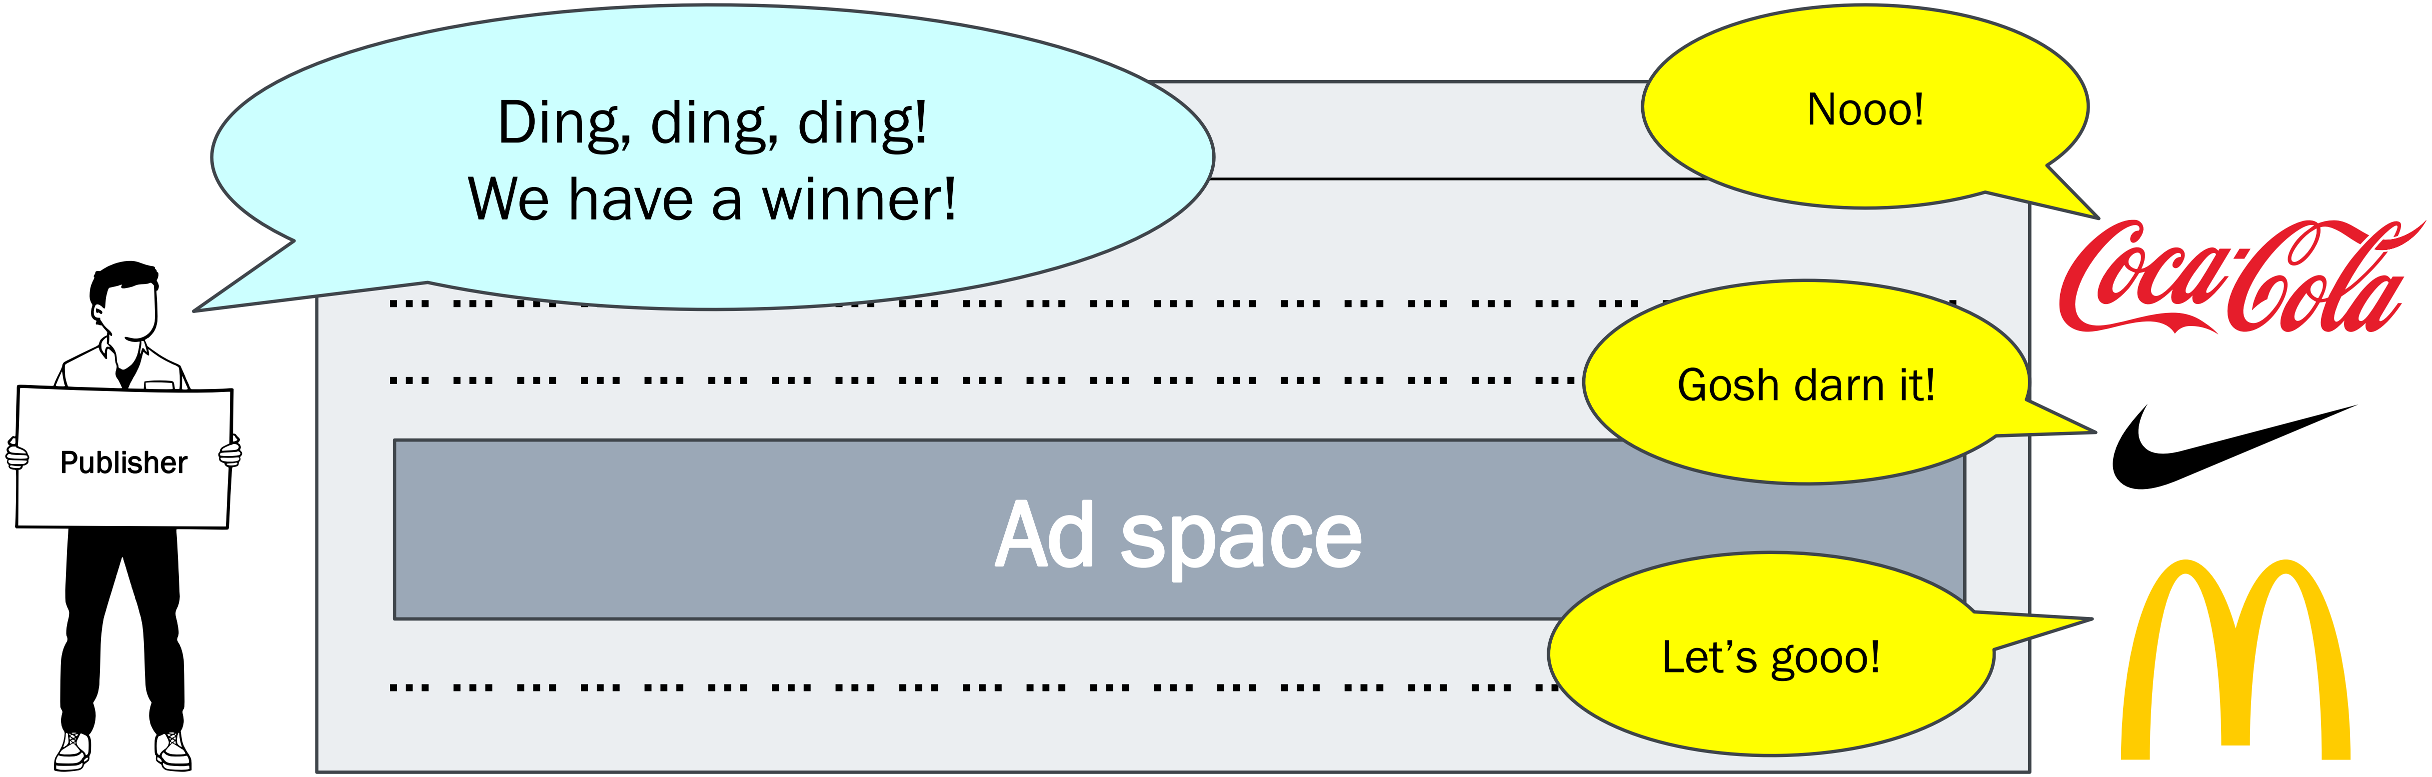

Motivation: Real-Time Bidding

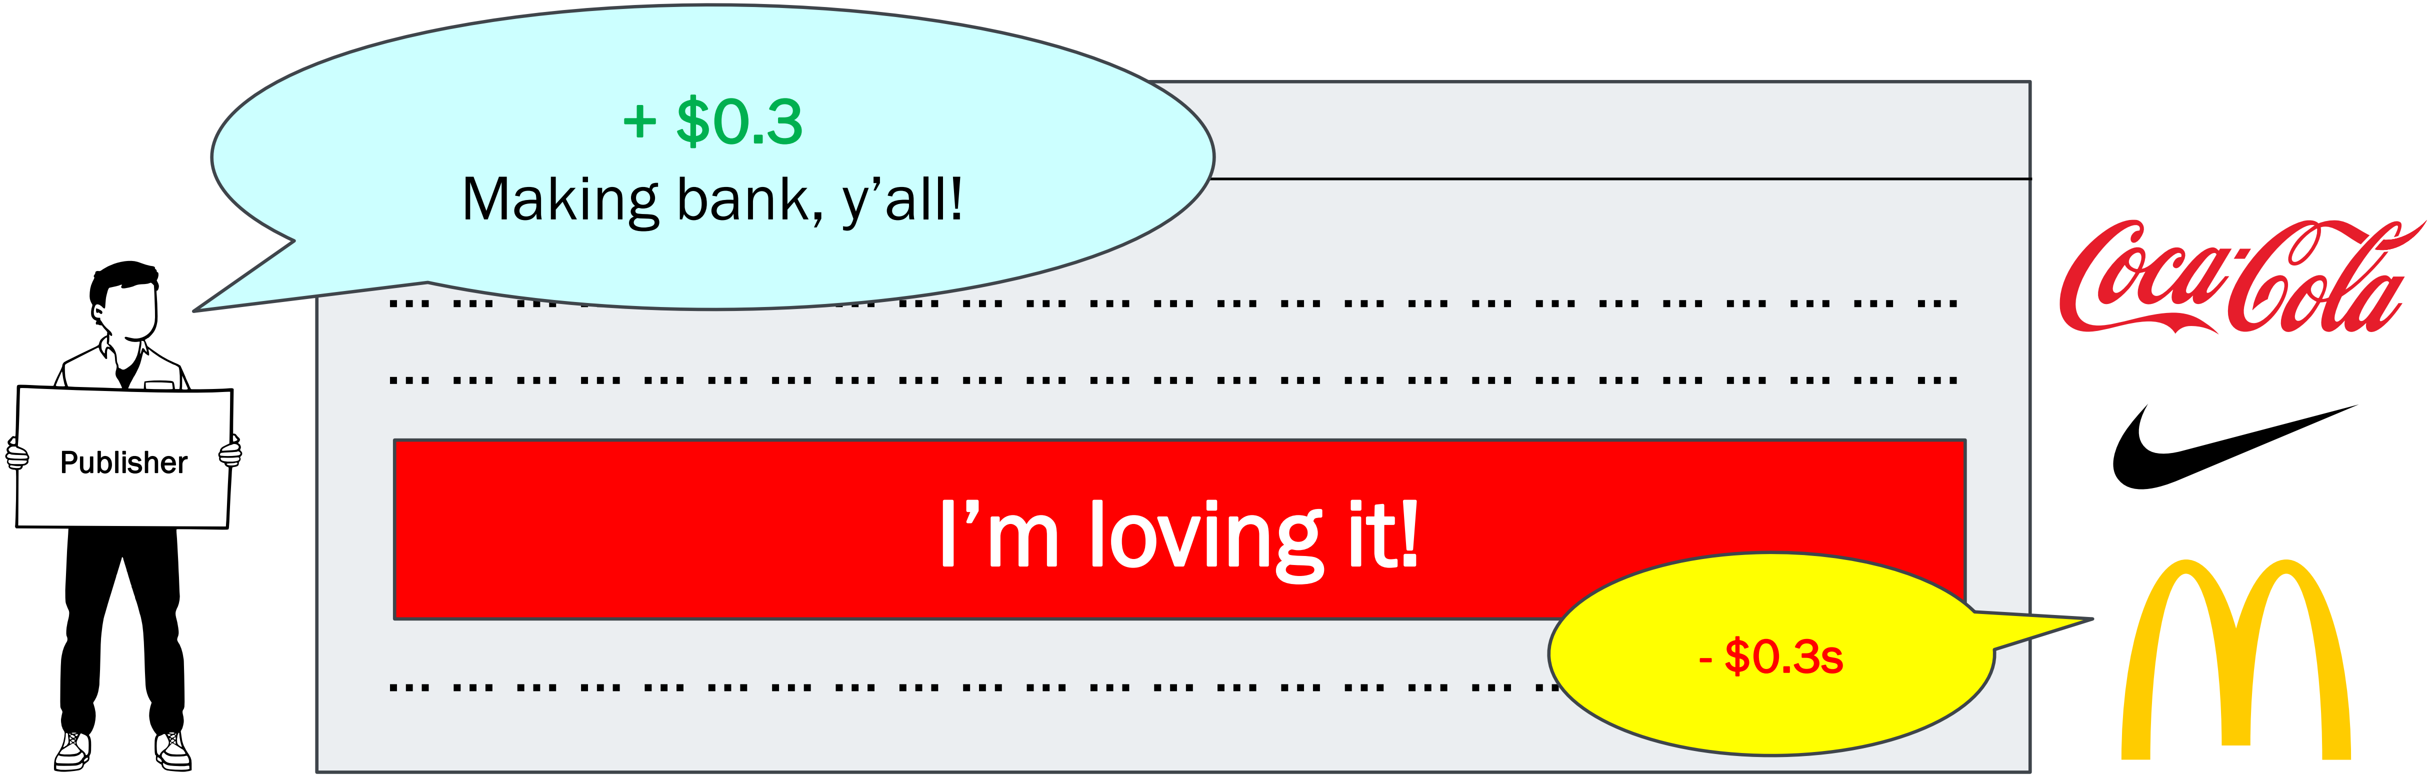

Motivation: Real-Time Bidding

How do advertisers quantify value?

- Metric 1: Click-through rate
- Metric 2: Conversion of curiosity to sales
- Metric 3: Time spent on a page

How do advertisers quantify value?

- Metric 1: Click-through rate
- Metric 2: Conversion of curiosity to sales
- Metric 3: Time spent on a page

All happening only if the auction is won!

Setting

Setting

- The environment privately generates V_t and M_t in [0,1]

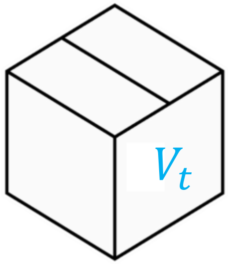

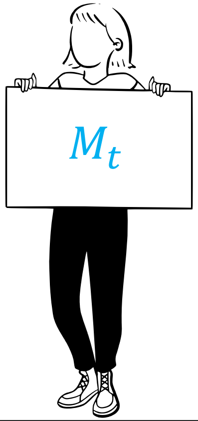

Setting

- The environment privately generates V_t and M_t in $[0,1]$
- The bidder bids $B_t \in [0,1]$

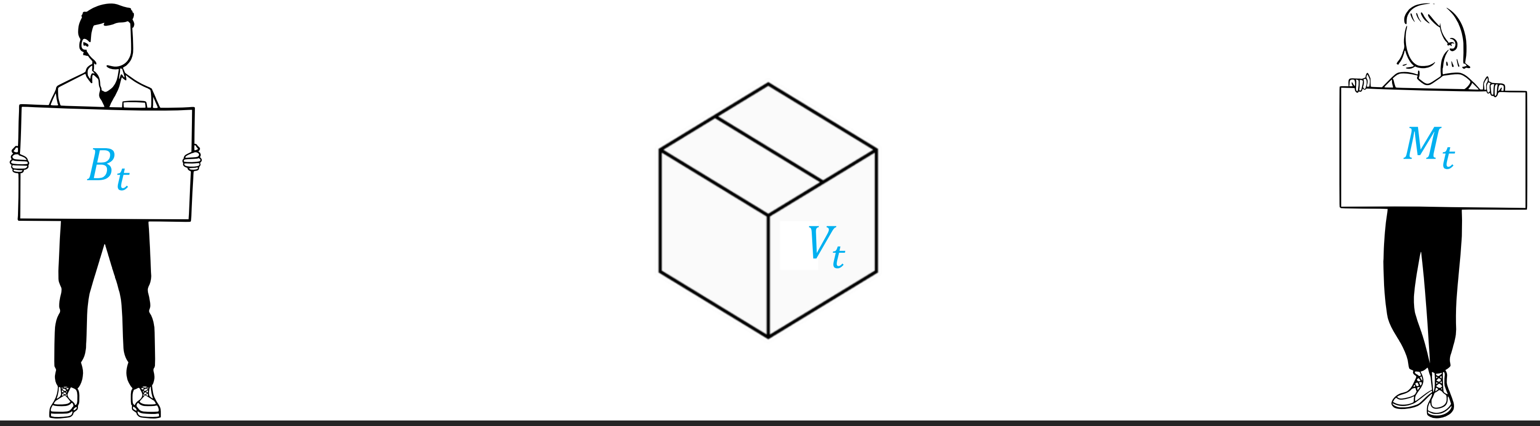

Setting

- The environment privately generates V_t and M_t in $[0,1]$
- The bidder bids $B_t \in [0,1]$ (without knowing V_t and M_t)

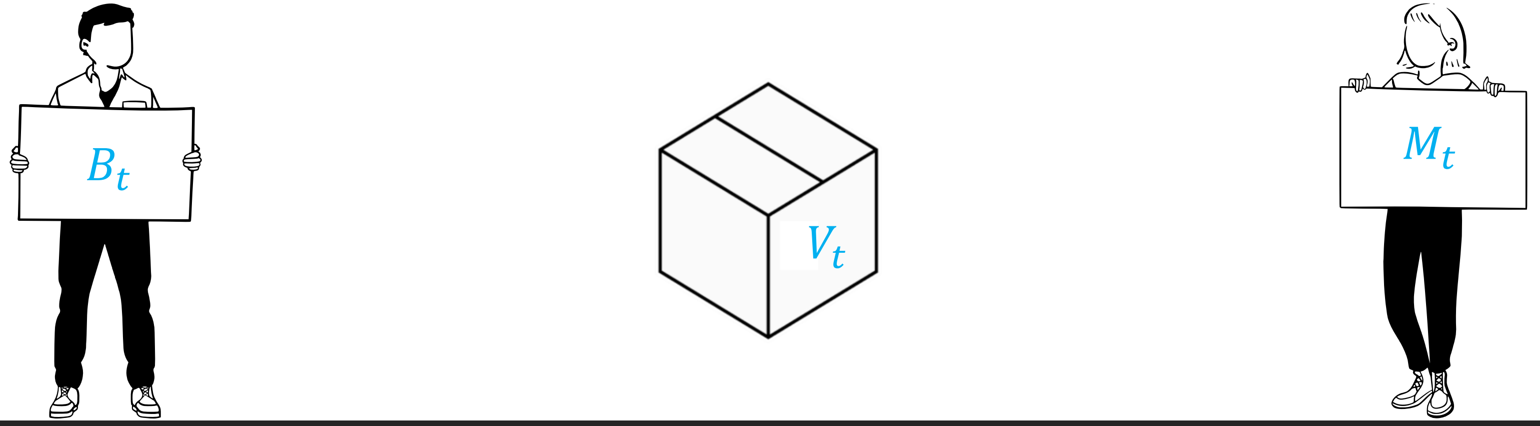

Setting

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

$$\mathrm{Util}_t(B_t) \coloneqq (V_t - B_t) \cdot \mathbb{I}\{B_t \geq M_t\}$$

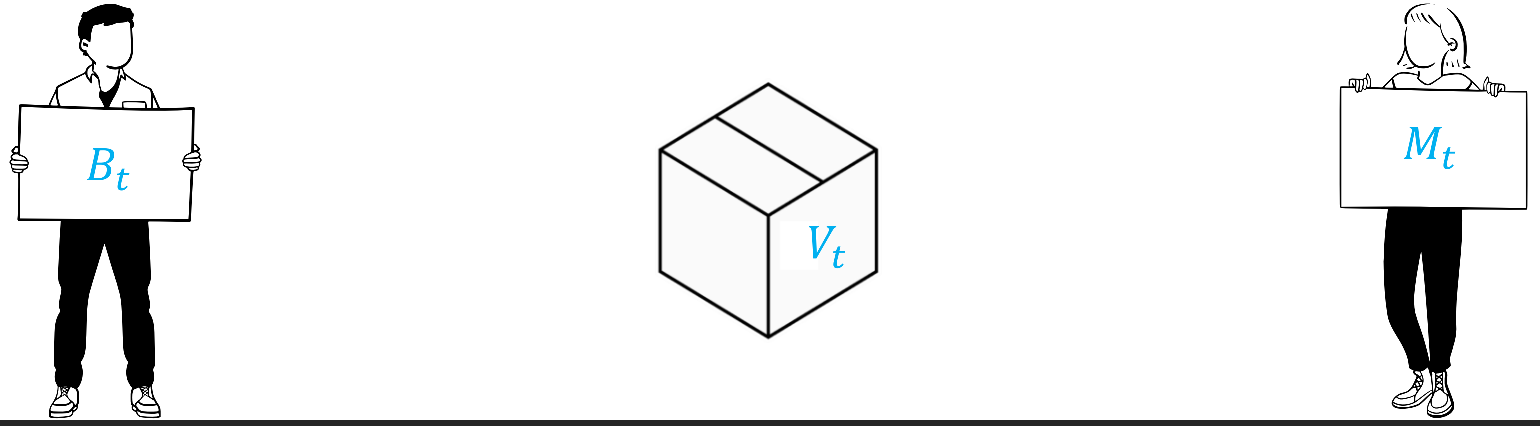

Setting

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

$$\mathrm{Util}_t(B_t) \coloneqq (V_t - B_t) \cdot \mathbb{I}\{B_t \geq M_t\}$$

- The bidder receive some feedback

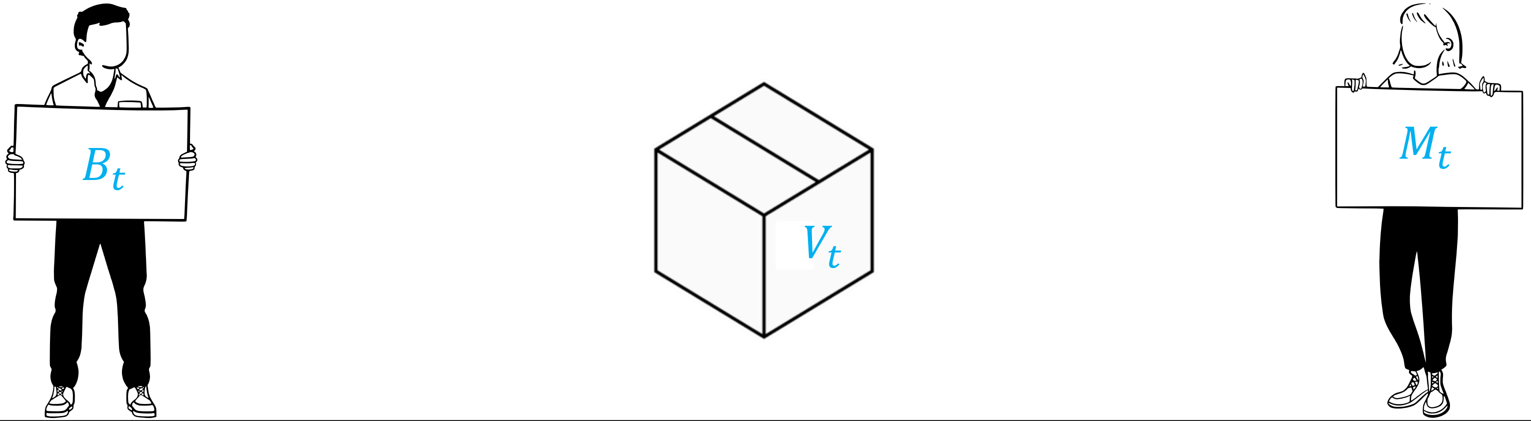

Setting

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

$$\mathrm{Util}_t(B_t) \coloneqq (V_t - B_t) \cdot \mathbb{I}\{B_t \geq M_t\}$$

- The bidder receive some feedback (transparency)

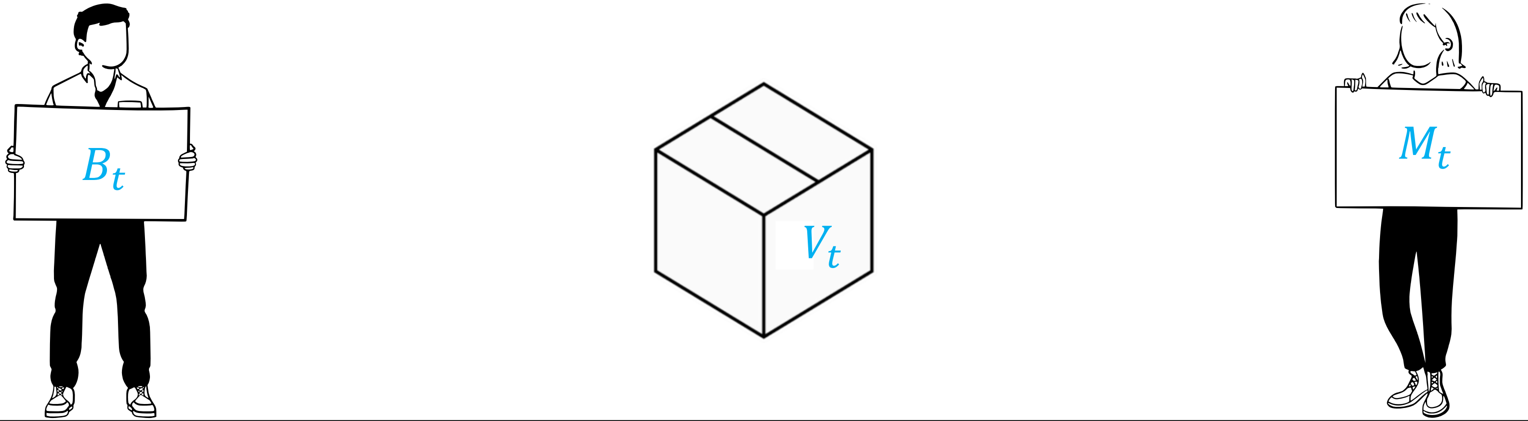

Setting

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

$$\mathrm{Util}_t(B_t) \coloneqq (V_t - B_t) \cdot \mathbb{I}\{B_t \geq M_t\}$$

- The bidder receive some feedback (transparency)

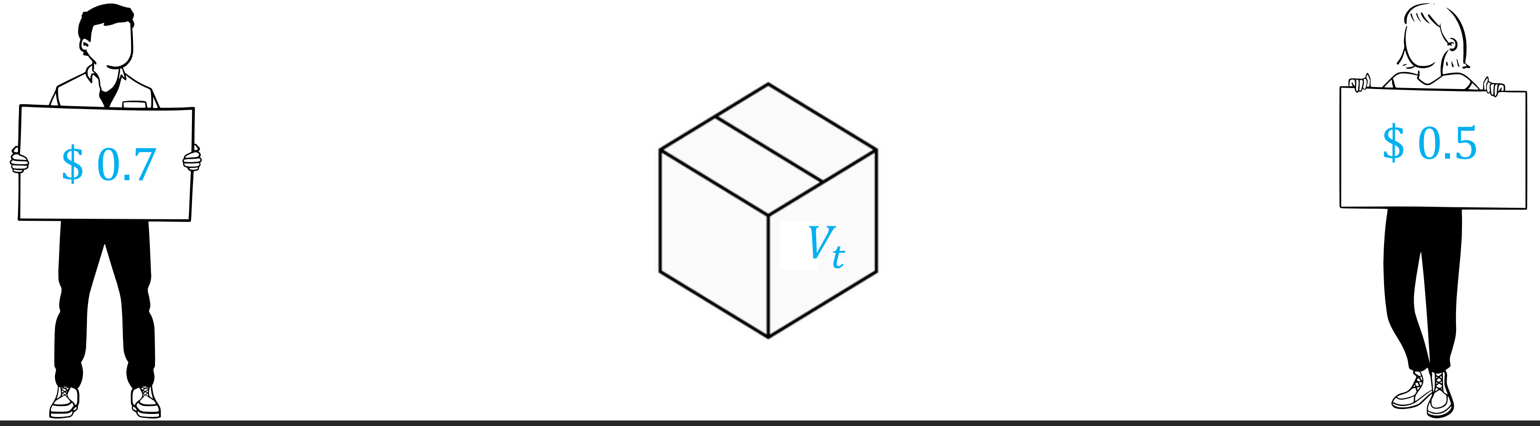

Setting

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

$$\mathrm{Util}_t(B_t) \coloneqq (V_t - B_t) \cdot \mathbb{I}\{B_t \geq M_t\}$$

- The bidder receive some feedback (transparency)

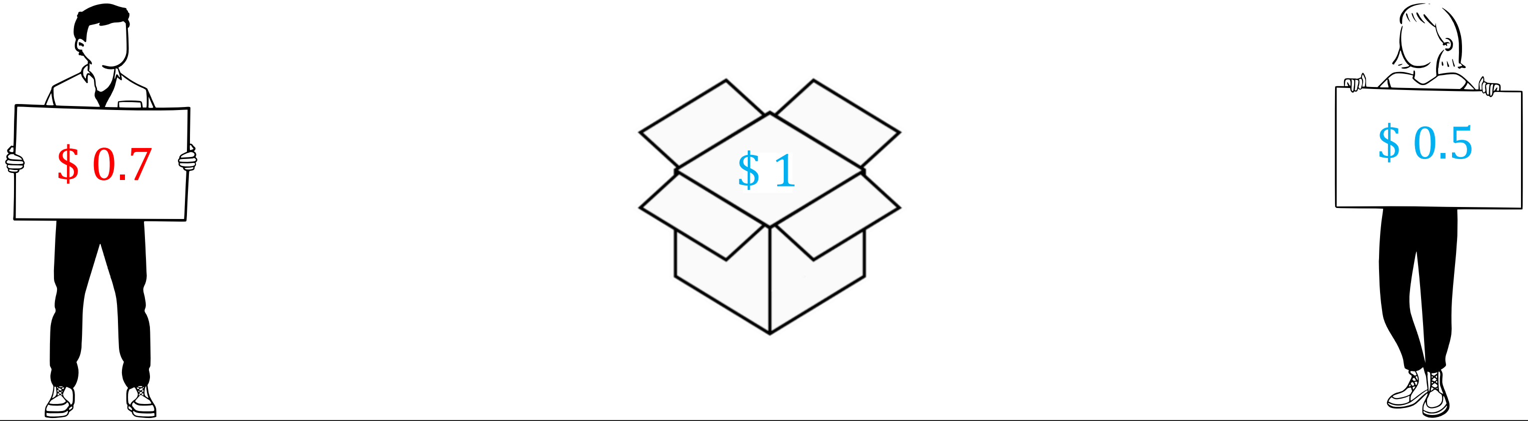

Setting

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

$$\mathrm{Util}_t(B_t) \coloneqq (V_t - B_t) \cdot \mathbb{I}\{B_t \geq M_t\}$$

- The bidder receive some feedback (transparency)

Regret

Goal: minimize the regret

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

Regret

Goal: minimize the regret

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

- We want R_T to grow sublinearly in the time horizon T

Regret

Goal: minimize the regret

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

- We want R_T to grow sublinearly in the time horizon T
- The per-round utility $\mathrm{Util}_t \in [-1,1]$

Regret

Goal: minimize the regret

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

- We want R_T to grow sublinearly in the time horizon T
- The per-round utility $\mathrm{Util}_t \in [-1,1]$

Our contribution: We fully characterize the minimax regret rate for various feedback and data generation models

Feedback Models

N.B. The feedback on M_t depends on the platform's transparency

Feedback Models

N.B. The feedback on M_t depends on the platform's transparency

	M_t	V_t
Full	Always observed	Always observed

Feedback Models

N.B. The feedback on M_t depends on the platform's transparency

	M_t	V_t
Full	Always observed	Always observed
Transparent	Always observed	Observed if auction is won

Transparent Feedback

The following bidders participated to the auction:

- Anonimous 1 bid $ 0.79 and won the auction
- Anonimous 2 bid $ 0.75
- Anonimous 3 bid $ 0.73
- Anonimous 4 bid $ 0.34
- Anonimous 5 bid $ 0.12

Feedback Models

N.B. The feedback on M_t depends on the platform's transparency

	$\boldsymbol{M_t}$	$\boldsymbol{V_t}$
Full	Always observed	Always observed
Transparent	Always observed	Observed if auction is won
Semi-Transparent	Observed if auction is lost	Observed if auction is won

Semi-Transparent Feedback

Anonimous bid $ 0.79 and won the auction

Feedback Models

N.B. The feedback on M_t depends on the platform's transparency

	M_t	V_t
Full	Always observed	Always observed
Transparent	Always observed	Observed if auction is won
Semi-Transparent	Observed if auction is lost	Observed if auction is won
Bandit	Never observed	Observed if auction is won

Semi-Transparent Feedback

I ain’t no snitch, fool. Buzz off!

Feedback Models

N.B. The feedback on M_t depends on the platform's transparency

	M_t	V_t
Full	Always observed	Always observed
Transparent	Always observed	Observed if auction is won
Semi-Transparent	Observed if auction is lost	Observed if auction is won
Bandit	Never observed	Observed if auction is won

Environments

Environments

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

We might want to avoid “atoms”

Definition (σ-smoothness)

A measure μ on $[0,1]^2$ is σ-smooth if it admits a density (w.r.t. Lebesgue) bounded by $1/\sigma$

Technical Challenges

Technical Challenges

The quality of the feedback

- Transparency regulates the ability to reconstruct counterfactual information

Technical Challenges

The quality of the feedback

- Transparency regulates the ability to reconstruct counterfactual information

The size and structure of the action space

- We typically know how to handle finite action spaces

Technical Challenges

The quality of the feedback

- Transparency regulates the ability to reconstruct counterfactual information

The size and structure of the action space

- We typically know how to handle finite action spaces
- We typically know how to handle regular objectives

The Utility Function

Recall that the utility as a function of the bid b is

$$\text{Util}_t(b) := (V_t - b) \cdot \mathbb{I}\{b \geq M_t\}$$

The Utility Function

Recall that the utility as a function of the bid b is

$$\mathrm{Util}_t(b) \coloneqq (V_t - b) \cdot \mathbb{I}\{b \geq M_t\}$$

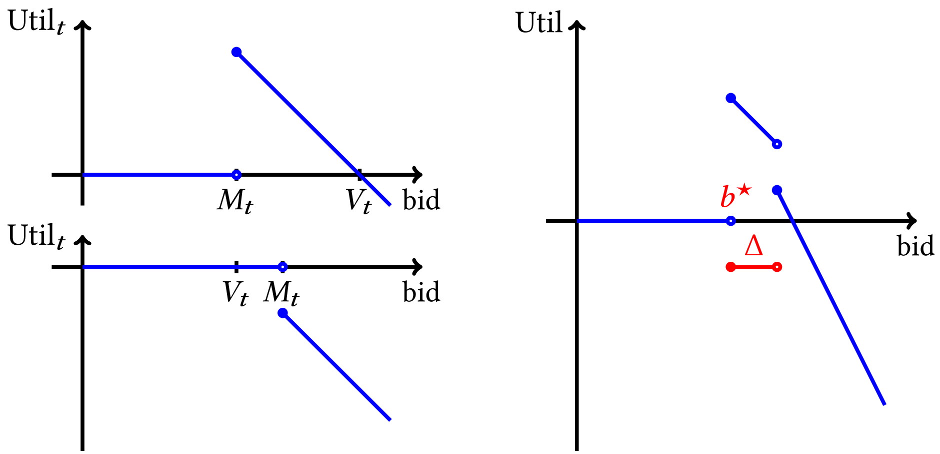

The Utility Function

Recall that the utility as a function of the bid b is

$$\text{Util}_t(b) \coloneqq (V_t - b) \cdot \mathbb{I}\{b \geq M_t\}$$

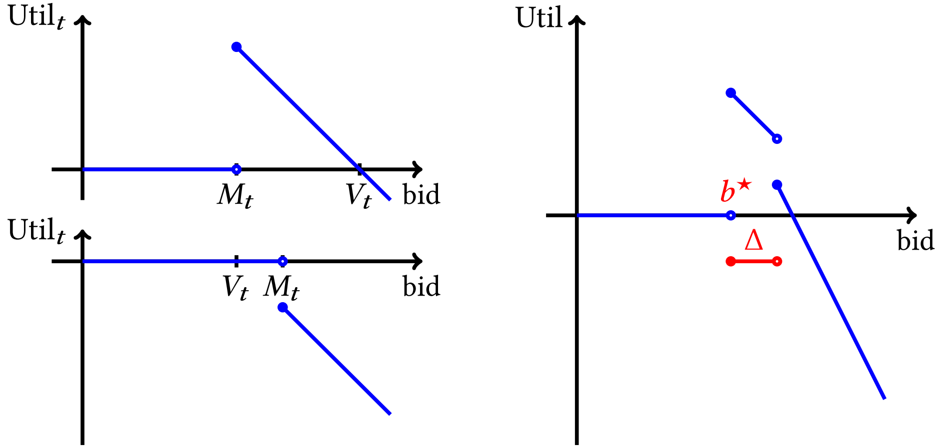

It is not (one-sided) Lipschitz nor (semi) continuous!

Our Results

	Stochastic i.i.d.		Adversarial	
	Smooth	General	Smooth	General
Full				
Transparent				
Semi-Transparent				
Bandit				

Our Results

	Stochastic i.i.d.		Adversarial	
	Smooth	General	Smooth	General
Full				$\Omega(T)$
Transparent				
Semi-Transparent				
Bandit				

Our Results

	Stochastic i.i.d.		Adversarial	
	Smooth	General	Smooth	General
Full				$\Omega(T)$
Transparent			$\tilde{\mathcal{O}}(\sqrt{T})$	
Semi-Transparent				
Bandit			$\mathcal{O}(T^{2/3})$	

Our Results

	Stochastic i.i.d.		Adversarial	
	Smooth	General	Smooth	General
Full				$\Omega(T)$
Transparent		$\mathcal{O}(\sqrt{T})$	$\tilde{\mathcal{O}}(\sqrt{T})$	
Semi-Transparent		$\tilde{\mathcal{O}}(T^{2/3})$		
Bandit			$\mathcal{O}(T^{2/3})$	

Our Results

	Stochastic i.i.d.		Adversarial	
	Smooth	General	Smooth	General
Full	$\Omega(\sqrt{T})$			$\Omega(T)$
Transparent		$\mathcal{O}(\sqrt{T})$	$\tilde{\mathcal{O}}(\sqrt{T})$	
Semi-Transparent	$\Omega(T^{2/3})$	$\tilde{\mathcal{O}}(T^{2/3})$		
Bandit			$\mathcal{O}(T^{2/3})$	

Our Results

	Stochastic i.i.d.		Adversarial	
	Smooth	General	Smooth	General
Full	$\Omega(\sqrt{T})$			$\Omega(T)$
Transparent		$\mathcal{O}(\sqrt{T})$	$\tilde{\mathcal{O}}(\sqrt{T})$	
Semi-Transparent	$\Omega(T^{2/3})$	$\tilde{\mathcal{O}}(T^{2/3})$		
Bandit			$\mathcal{O}(T^{2/3})$	

Our Results

	Stochastic i.i.d.		Adversarial	
	Smooth	General	Smooth	General
Full	$\Omega(\sqrt{T})$			$\Omega(T)$
Transparent		$\mathcal{O}(\sqrt{T})$	$\tilde{\mathcal{O}}(\sqrt{T})$	
Semi-Transparent	$\Omega(T^{2/3})$	$\tilde{\mathcal{O}}(T^{2/3})$		
Bandit		$\Omega(T)$	$\mathcal{O}(T^{2/3})$	

Our Results

	Stochastic i.i.d.		Adversarial	
	Smooth	General	Smooth	General
Full	$\Omega(\sqrt{T})$			$\Omega(T)$
Transparent		$\mathcal{O}(\sqrt{T})$	$\tilde{\mathcal{O}}(\sqrt{T})$	
Semi-Transparent	$\Omega(T^{2/3})$	$\tilde{\mathcal{O}}(T^{2/3})$		
Bandit		$\Omega(T)$	$\mathcal{O}(T^{2/3})$	

- **THM 1.** Adversarial and bandit settings are unlearnable without smoothness
- **THM 2.** Beyond that, revealing the winning bid avoids pathologies
- **THM 3.** In particular, revealing all bids drastically improves learnability (to full-info levels)

Our Results

Thank you!

	Stochastic i.i.d.		Adversarial	
	Smooth	General	Smooth	General
Full	$\Omega(\sqrt{T})$			$\Omega(T)$
Transparent		$\mathcal{O}(\sqrt{T})$	$\tilde{\mathcal{O}}(\sqrt{T})$	
Semi-Transparent	$\Omega(T^{2/3})$	$\tilde{\mathcal{O}}(T^{2/3})$		
Bandit		$\Omega(T)$	$\mathcal{O}(T^{2/3})$	

- **THM 1.** Adversarial and bandit settings are unlearnable without smoothness
- **THM 2.** Beyond that, revealing the winning bid avoids pathologies
- **THM 3.** In particular, revealing all bids drastically improves learnability (to full-info levels)