CS 580 · Algorithmic Game Theory

Mechanism Design for AI Agents

Autonomous bidding, sponsored search, and responses to public forecasts

Games and Markets with AI Agents

LECTURE MAP

Three decisions in bidding and search

Platform chooses	Agent or seller decides	Outcome to evaluate
Bid interface and spending cap	Which bid to submit or skip	User payoff and market welfare
Search/inspection subsidies	How much to sponsor and whom to inspect	Search costs and seller participation
Checks and shared records	Which services to book and prices to post	Learning, competition and buyer surplus

In every case, follow the rule through the agent's response to the user's payoff.

Autobidding

Autobidding software/AI-agent

Who chooses?	What can that choice change?
Advertiser	Objective, spending limit, delegated authority
Agent software	How observations become proposed bids
Platform	Allowed bids, feedback, payments, hard limits

$$\boldsymbol{b} \sim \Pi_c^M(\,\cdot \mid \boldsymbol{\theta}, h), \qquad (x, p) = M(\boldsymbol{b})$$

Π is the joint law of submitted bids b given users' types θ, history h, and configuration c. Mechanism M maps b to allocations x and payments p.

What does the LLM choose?

An ad impression is one opportunity to display an ad. RTBAgent estimates the click probability for impression j and computes a reference bid.

At period t, the LLM reads campaign history and selects one adjustment a_t between −0.5 and 0.5 for that period's impressions.

$$b_{tj} = \underbrace{\widehat{\Pr}(\text{click} \mid j)\,\lambda_{\text{base}}}_{\text{reference bid for impression } j}\quad (1 + a_t)$$

For a reference bid of 4, an adjustment of +0.25 produces a bid of 5.

The designer chooses which actions the bidding agent can take.

Cai et al. (WWW 2025), §3 and Eq. 11. The bid calculation is a teaching example.

What did the LLM bidder achieve in replay?

Both policies faced the same recorded market prices. The budget in each column is a fraction of historical spending. Entries count replayed clicks.

Bidding policy	Budget 1/2	Budget 1/8	Budget 1/32
DiffBid	2,275	1,235	793
RTBAgent	2,281	1,240	795

At the half-budget setting, the LLM system earned six additional clicks over DiffBid in this replay.

Rival prices stayed fixed in replay. Market welfare after rivals adapt is a different quantity.

Cai et al. (WWW 2025), §4.2 and Table 2. Counts are replayed clicks.

Ad Auction: What obj. Advertiser Delegates?

For impression t, v_t is the advertiser's value, x_t its allocation, and p_t its payment. Across T auctions, the agent chooses a bidding policy.

$$\max\ V = \sum_{t=1}^{T} v_t x_t \quad \text{subject to} \quad P = \sum_{t=1}^{T} p_t \le B, \quad V \ge \kappa P$$

B caps total spending. κ = 2 requires at least two units of acquired value per unit spent over the campaign.

Note: A surplus maximizer instead maximizes V − P. A policy that maximizes clicks can serve a yet another objective.

Campaign objective adapted from Lucier et al. (COLT 2024).

Example: Only Budget Check is not Enough

The agent knows both values and rival bids. Its total spending cap is 5.

Opportunity	Value	Price if won
First impression	8	4
Second impression	12	5

Buy first: spend 4, leave 1, miss the second impression, and get value 8.

Skip first and buy second: spend 5 and get value 12. Both plans obey the cap.

Quasilinear w/ Non-linear Price Cost

$$u_i(x_i, P_i) = v_i(x_i) - C_i(P_i), \qquad C_i \text{ increasing and convex}$$

Here P is the total payment. C(P) is the opportunity cost of that payment, measured in the same utility units as the value v(x).

Example: value 10, rival bid 6	Outcome	Utility when C(P) = 2P
Report 10	Win, pay 6	−2
Report 5	Lose, pay 0	0

With C(P) = 2P, a price of 6 costs 12 utility units. Bidding 5 avoids a loss.

Babaioff et al. (ITCS 2021), §2. Linear C(P) = 2P is an illustrative convex cost.

Best response to fixed rival bids

Fix the rivals' reports in a mechanism that is truthful for quasilinear bidders.

Bidder i compares every available report against those fixed rivals using its true value minus its cost of payment.

$$b_i^* \in \arg\max_{\widehat{b}_i} \left\{ v_i\Big(x_i(\widehat{b}_i, b_{-i})\Big) - C_i\Big(p_i(\widehat{b}_i, b_{-i})\Big) \right\}$$

Thm. With suitable tie-breaking, one optimal report scales the whole valuation.

$$b_i^* = \alpha_i v_i, \qquad \alpha_i = \frac{1}{C_i'(P_i^*)}, \qquad P_i^* = p_i(b_i^*, b_{-i})$$

For smooth costs and a nonbinding budget, the multiplier uses the marginal cost at the payment induced by that report.

With value 10 and C(P) = 2P ⇒ α = ½. The b_i^* =5 loses to 6 and earns 0.

Babaioff et al. (ITCS 2021), Theorem 3.1. At a kink, 1/α is a supporting slope of C at P*.

Proof. (why the scaled report is optimal)

Hold rivals fixed. The scaled report yields true value V* and payment P*. Any other report yields V and P. Let m = C′(P*) = 1/α.

Truthfulness for valuation αv bounds the extra value from any deviation.

$$\alpha V^* - P^* \geq \alpha V - P \quad \Longrightarrow \quad V - V^* \leq m(P - P^*)$$

Convexity bounds the extra cost of the deviation by the same tangent.

$$C'(P^*) \leq \frac{C(P) - C(P^*)}{P - P^*} \Rightarrow m(P - p^*) \leq C(P) - C(P^*)$$

Combining these bounds gives V − C(P) ≤ V* − C(P*).

Best Response Bidding

Keep value 10, cost 2p, and rival bid 6. The unrestricted optimal payoff is 0.

Suppose the agent must bid and its only tool adjusts a reference bid of 14 by at most 50%. The reachable bids are 7 through 21.

Action available	What happens?	User payoff
Bid between 7 and 21	Win and pay 6	$10 - 2 \times 6 = -2$
Skip the impression	Lose and pay 0	0

Adding "skip" makes the best payoff attainable again.

Constructed action-space example inspired by the bounded adjustment in RTBAgent (Cai et al., 2025).

Evaluating the Outcome

One ad impression. Both bidders get zero if they lose.

Bidder	Payoff from winning at price p	Largest acceptable price
A	$10 - 2p$	\$5 $(10 - 2 \times 5 = 0)$
B	$6 - p$	\$6 $(6 - 6 = 0)$

At bids of \$5 and \$6, B wins and pays \$5.

Transferable welfare adds each bidder's maximum acceptable price for its assigned bundle. Here A contributes \$0 and B contributes \$6.

Theorem (Welfare). At a pure equilibrium of ROI-optimal scaled reports, this welfare is at least half the largest sum achievable by any feasible allocation.

Algorithm: max (V - P) s.t. budget constraint

Maximize: V − P subject to P ≤ B.

Give scarce budget a shadow price μ.

$$L = V - P + \mu(B - P) = V - (1 + \mu)P + \mu B$$

At rival price z, winning is worthwhile when $v_t - (1 + \mu_t)z \geq 0$. The corresponding second-price bid is the largest acceptable price.

Pacing Eq. Bids:

$$b_t = \frac{v_t}{1 + \mu_t}, \qquad \mu_{t+1} = \left[\mu_t + \eta(p_t - B/T)\right]_+$$

If spending exceeds the target B/T, the update raises μ and lowers the next bid. This simple controller can also check the remaining budget before submitting a bid.

The only error signal here is spending relative to budget pace.

Algorithm: max V, s.t. Budget and RoI constraints

Maximize V subject to $P \le B$ and $V \ge \kappa P$.

Write ρ = B/T for target spend per round and $\bar{v}$ for an upper bound on values.

Initialize the two multipliers and use the more conservative bid each round.

$$\mu_1^B = \frac{\bar{v}}{\rho} - 1, \qquad \mu_1^R = \kappa - 1, \qquad b_t = \frac{v_t}{1 + \max\{0, \mu_t^B, \mu_t^R\}}$$

After observing payment p_t and acquired value $v_t x_t$, update both multipliers.

$$\mu_{t+1}^B = \mu_t^B + \eta_B (p_t - \rho),$$
$$\mu_{t+1}^R = \mu_t^R + \eta_R (\gamma p_t - v_t x_t).$$

Suppose, an agent spends 3 against a budget of 5, with κ = 2 and acquired value 4. Budget error = 3 – 5 =−2, and the ROI shortfall = 3*2 – 4 = 2.

The welfare guarantees

Theorem. When every bidder uses the Algorithm with the stated learning rates, each round uses a core auction (including first- or second-price single-item auctions), value profiles are i.i.d. across rounds, then

for the liquid welfare at the realized allocation *x*,

$$\mathrm{LW}(x) = \sum_{i=1}^{n} \min\left\{ B_i, \frac{1}{\kappa_i} \sum_{t=1}^{T} v_{it} x_{it} \right\}$$

the algorithm achieves the following where OPT is the best ex-ante liquid welfare

$$\mathbb{E}[\mathrm{LW}] \;\geq\; \tfrac{1}{2}\,\mathrm{OPT}_{\text{ex ante}} - O\left(n\bar{v}\sqrt{T \log(\bar{v} n T)} \right)$$

The guarantee concerns the whole trajectory. It does not require the bids to converge.

Lucier et al. (COLT 2024), full paper §5.1, Definition 8 and Theorem 10. Fixed n and value bound as T grows.

AI Agent Inspects ⇒ Subsidize Search

Shopping Agents

User attention = Scarce resource

Shopping agent: Represents user. Can inspect a lot more, but spends *tokens*!

Paper and question	What changes in the model?
Subsidizing Sequential Search	Sellers choose what to reimburse when inspected.
Agentic Markets	Inspection becomes cheaper or reveals more. Buyers share feedback over time.

We first study who pays for one search, then what that search teaches later buyers.

Example: Inspection cost

A buyer delegates a booking. An inspection reveals whether a seller fits the buyer's needs. Quality is observed after purchase.

Single Seller: inspect when expected gains from a possible purchase exceed the inspection cost.

If fit probability is 0.6, expected value given fit is 9, price is 4, and inspection costs 2.

$$\textbf{Gain} = 0.6 \times (9 - 4) - 2 = 1.$$

If inspection cost was 4, then the agent would not inspect! -> Revenue loss!

Sellers can sponsor an inspection

One check costs c. A subsidy unit costs the seller p. Seller j privately knows its match probability t_j. A match pays the buyer u and the seller 1.

Move	Who pays or observes what?
Sellers post offers simultaneously	Each chooses a subsidy s_j between 0 and c.
Agent sees offers and inspects j	Buyer bears $c - s_j$. Seller pays $p \times s_j$ upon inspection.

After observing the offers, the agent inspects sequentially. It stops at the first match or when no remaining inspection is worthwhile.

An offer can both lower the cost to the buyer and signal seller quality.

Example: Subsidy -> Inspection

One seller remains. A match is worth u = 1, its posterior probability is 0.20, and an inspection costs c = 0.30. Stopping pays 0.

Offer if inspected	Private cost c − s	Expected gain	Decision
0	0.30	0.20 − 0.30 = −0.10	Stop
0.12	0.18	0.20 − 0.18 = +0.02	Inspect

With s = 0.12, the buyer gains 0.02 from checking. At p = 1, the seller pays 0.12 even when the check fails.

The inspection still costs 0.30 in resources. The offer changes who bears it.

Constructed example of the May paper's inspection rule. Posterior match probability is held fixed.

Inspection order

Let τ_j be the posterior match probability after seeing seller j's offer, and let $k_j = c - s_j$ be the net inspection cost.

Example: Sellers A and B. Both give utility *u* if inspection passes.
Inspecting A first costs $k_a + (1 - \tau_a)k_\beta$ in expectation.
Comparing the two orders gives A first when $k_a/\tau_a \le k_\beta/\tau_\beta$.

Equivalently, inspect in descending reservation index.

$$r_j = u_j - \frac{c - s_j}{\tau_j}, \qquad \text{where } \tau_j = E\left[t_j \middle| s_j\right]$$

Stop after a match or when the next index is nonpositive (tie convention).

At Eq.: Higher types offer larger subsidies

Let q(s) be the chance a firm offering s is inspected. Its expected profit at type t is q(s)(t − ps). An extra chance of inspection is worth more to a higher type.

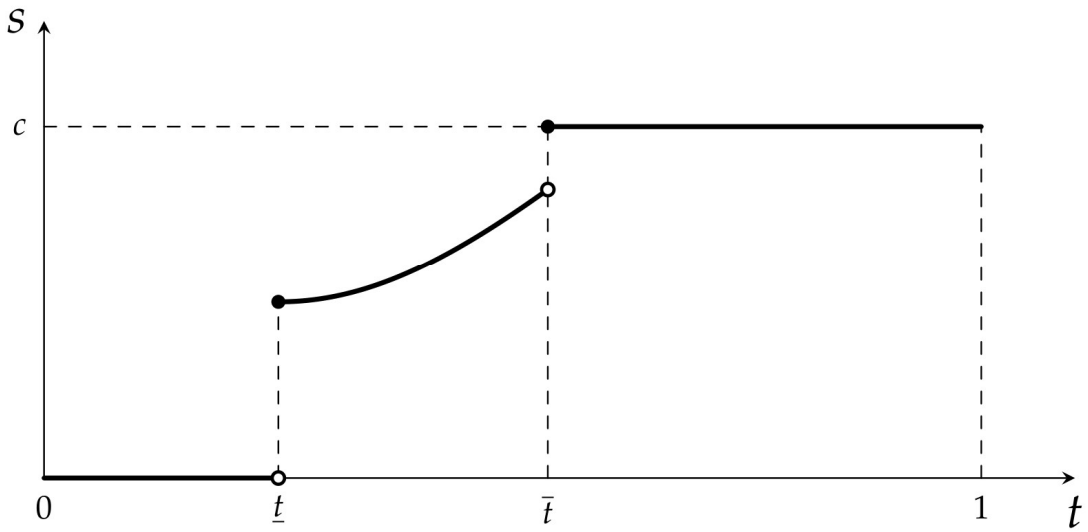
In symmetric equilibrium, higher types offer weakly larger subsidies. Larger on-path offers lower net cost and signal weakly better matches. Search proceeds from larger to smaller offers.

Seller types in the refined outcome	Subsidy and inspection
Low	Never inspected
Intermediate	Subsidy rises strictly with type
High	Pool at the full subsidy s = c

A full subsidy attracts attention but pools the best sellers.

Seller type and the equilibrium subsidy

The horizontal axis gives the seller's match probability t. The vertical axis gives its subsidy s per inspection in the refined equilibrium.



Low types go uninspected

Below the first cutoff, the seller offers zero in the schedule shown.

Middle types separate

Between the cutoffs, each higher type offers a strictly larger subsidy.

High types pool

All types at or above the upper cutoff offer s = c. The buyer can inspect each for free.

The dashed vertical lines mark where inspection begins and where the subsidy reaches its cap.

The platform can price inspection tokens

A verifiable agent query can trigger a payment. The platform sets p per unit of subsidy. Firms buy subsidy units to fund inspections of their products.

For n symmetric firms, s(t;p) is the subsidy offered by type t and q(t;p) is its inspection probability. D(p) counts expected subsidy units per firm.

$$D(p) = \mathbb{E}[s(t;p)\, q(t;p)]\,, \qquad R(p) = n\, p\, D(p)$$

At the revenue-maximizing price, some types with negative virtual value still receive inspections. The authors call this excess search relative to their benchmark.

Compare platform revenue, buyer surplus, and the real cost of inspection separately.

Candelas, Immorlica & Lucier, §6 and Theorem 4. Token-price result assumes the paper’s regularity condition.

One buyer’s inspection teaches the next buyer

Each new buyer sees a shared history, pays c per inspection, and buys at most one product. Inspections reveal fit. Purchases also reveal quality.

A seller is learned if repeated observations identify its fit and quality. It can be lost if pessimistic beliefs stop future buyers from inspecting it.

Theorem 6.1 compares markets identical except for a lower inspection cost. Under a coupling of histories, the learned set weakly expands and long-run consumer surplus weakly rises.

This result changes the cost of inspection itself. The May model holds that cost fixed and lets firms pay a share.

Lucier et al. (March 2026), §§3–6 and Theorem 6.1. Prices are fixed in this comparison.

Information sharing: Descriptive

Types H and L are equally likely. A common screen passes with probability 0.9 for H and 0.3 for L. An independent niche screen passes with probability 0.5 for both.

Observed record	Likelihood for H / L	Posterior Pr(H)
Any screens failed	$1 - (0.9 \times 0.5) / 1 - (0.3 \times 0.5)$	$0.55/(0.55+0.85) \approx 39\%$
Common passed, niche failed	$0.9 \times 0.5 / 0.3 \times 0.5$	$0.45/(0.45+0.15) = 75\%$

The same rejection raises belief in H when the buyer can see which screen failed.

Constructed Bayesian example motivated by Lucier et al. (March 2026), §6.2.

Pay less for a check, or learn more from it?

Hold catering fees fixed. Initially, a $3 check tells the buyer whether a caterer serves the venue. The custom menu is confirmed only after booking.

Change to inspection	Why the change matters
Cheaper: pay $1 for the same venue check.	More sellers are worth trying. Long-run learning and buyer surplus weakly improve.
Deeper: pay $3 to check both venue and menu before booking.	A menu mismatch can explain a rejection. A bare "rejected" record can deter later buyers, reducing learning and surplus.

If the record says "venue passed, menu failed," later agents can use each fact. With such records, deeper checks weakly improve long-run learning and buyer surplus.

Lucier et al. (March 2026), Thm. 6.1, Prop. 6.5 and Thm. 6.6. Dollar costs are illustrative.

Example: Shallow vs. Deeper Checks

Caterers A and B, each posts a booking fee before the agent checks fit.
Checks are free. If one caterer passes, the buyer pays its fee. If both pass, the buyer books the cheaper one.

Each can serve the venue with probability 0.8, and can supply menu to half these. Fits are independent across caterers. A fully suitable booking is worth $100.

Known before booking	Chance A qualifies	Expected value if A qualifies	Pr(B fails | A qualifies)
Venue only	0.80	$50	0.20
Venue and menu	0.40	$100	0.60

In the symmetric mixed equilibrium, each caterer independently draws its fee from $10–$50 (venue only) or $60–$100 (venue and menu).

Worked example using Lucier et al. (2026), Prop. 7.16 and App. D.7. Match value scaled to $100.

Best Response: $50 and $100

Hold B’s equilibrium fee distribution fixed. A earns the same expected revenue at every fee in its range. Compare A’s two endpoint fees in each market.

Per arriving buyer	Venue only	Venue and menu
Lowest fee wins whenever A passes	$10 × 0.8 = $8	$60 × 0.4 = $24
Highest fee wins only when A alone passes	$50 × 0.8 × 0.2 = $8	$100 × 0.4 × 0.6 = $24
Total expected booking revenue	2 × $8 = $16	2 × $24 = $48

When both pass, the cheaper quote wins.

Lucier et al. (2026), Thm. 7.12 and Apps. D.4, D.7. Dollar-scaled example with two symmetric caterers and free checks.

Example (Contd.): Buyer value vs surplus

In equilibrium, a seller earns as much as it would by quoting the highest fee buyers accept ($50 or $100). At that fee it wins when it qualifies and its rival fails.

Expected dollars per buyer	Venue only	Venue and menu
Gross service value	$50 \times (1 - 0.2^2) = 48$	$100 \times (1 - 0.6^2) = 64$
Booking fees paid to sellers	$2 \times 50 \times 0.8 \times 0.2 = 16$	$2 \times 100 \times 0.4 \times 0.6 = 48$
Buyer surplus	$48 - 16 = 32$	$64 - 48 = 16$

Earlier menu checks create $16 more service value per buyer, but raise expected booking payments by $32. In this example, buyer surplus falls by $16.

Dollar-scaled two-seller example from Lucier et al. (2026), Prop. 7.16 and App. D.7. Amounts per arriving buyer.

Experiments and Evaluation

Evaluation: Process and Criteria

Evaluation	What changes?
Replay historical actions	Allocation and payments on the original actions
Rerun the same agent	Its actions under the new observations or tools
Let rivals or sellers adapt	Bidding updates or prices after deployment
Repeat across initial histories	Which sellers continue to be inspected and booked

Record buyer values, payments, search costs, seller participation and public observations. Report buyer surplus separately from seller profit and platform revenue.

Exercise: Did deeper screening help buyers?

The same number of shoppers arrived in each period. The platform added a deeper pre-purchase screen between periods.

Period	Bookings	Successful bookings
Before	80	60
After	50	45

1. Which statistic improved, and which fell?
2. Can you infer a change in buyer surplus?
3. How would you isolate a change in booking fees?

Use the table to separate better matches among bookings from better outcomes for all buyers.

Solution: Success rates do not identify surplus

Conditional booking success rises from 60/80 = 75% to 45/50 = 90%. The number of successful bookings falls from 60 to 45.

Consumer surplus also depends on the value of each match, prices paid, inspection costs, and the outside options of buyers who left. The table does not identify its change.

First hold the posted booking fees and screening depth fixed and vary the public record. Next vary screening depth. Finally let sellers choose new booking fees.

Compare matched arrivals and initial histories across all conditions.

Worked solution to the preceding exercise. All counts are illustrative.

What makes this mechanism design for AI agents?

Design lever	Consequence demonstrated in the lecture
Actions and authority	A skip action can restore an excluded best response.
Inspection sponsorship	Offers direct search and signal a seller’s match probability.
Shared search records	A stage-specific failure changes future agents’ beliefs.
Sellers’ pricing incentives	Better verified matches can fetch higher booking fees and leave buyers less surplus.

Evaluate the actions, information, and payments that the deployed rules induce.

Readings · Autonomous bidding

Executed policy and replay

Cai et al. (WWW 2025), RTBAgent: A LLM-based Agent System for Real-Time Bidding

Read Eq. 11 and Table 2. arXiv:2502.00792

Best responses and transferable welfare

Babaioff et al. (ITCS 2021), Non-quasi-linear Agents in Quasi-linear Mechanisms

Read §2 and Theorems 3.1 and 4.1. arXiv:2012.02893

Pacing in repeated auctions

Lucier et al. (COLT 2024), Autobidders with Budget and ROI Constraints: Efficiency, Regret, and Pacing Dynamics

Read Algorithm 2 and Theorem 10 in the full paper. arXiv:2301.13306v6

READINGS

Readings · Inspection and market learning

Who pays for inspection?

Candelas, Immorlica & Lucier (May 2026), Subsidizing Sequential Search

Read §§2–4 and Theorem 4. arXiv:2605.28985

What changes when search technology improves?

Lucier et al. (March 2026), Agentic Markets: Equilibrium Effects of Improving Consumer Search

Read §6 for search and feedback, and Proposition 7.16 with Appendix D.7 for booking prices. arXiv:2603.25893

Each entry links to the paper in the student PDF.