Online Auctions for Dynamic Environments

Mohammad T. Hajiaghayi

University of Maryland, College Park

Example: Last-Minute Tickets

Value	$100	$80	$60
Arrival:	11am	11am	12pm
Patience:	2hrs	2hrs	1hr

How should you bid?

"Please bid your value and your patience. A decision will be made by the end of your stated patience."

Value	\$100	\$80	\$60
Arrival:	11am	11am	12pm
Patience:	2hrs	2hrs	1hr

If truthful, then:
{ <1, \$80>, <2, \$60>}

Auction: sell one ticket in each hour (given demand), to the highest bidder at second-highest bid price.

Value	~~$100~~ $65	$80	$60
Arrival:	11am	11am	12pm
Patience:	2hrs	2hrs	1hr

If truthful, then:
{ <1, $80>, <2, $60>}

However, bidder 1 could
a) reduce bid price to $65
{<2, $65>, <1, $60>}

Auction: sell one ticket in each hour (given demand), to the highest bidder at second-highest bid price.

Value	\$100	\$80	\$60
Arrival:	~~11am~~ 12pm	11am	12pm
Patience:	~~2hrs~~ 1hr	2hrs	1hr

Auction: sell one ticket in each hour (given demand), to the highest bidder at second-highest bid price.

If truthful, then:
{ <1, \$80>, <2, \$60>}

However, bidder 1 could

a) reduce bid price to \$65
{<2, \$65>, <1, \$60>}

b) delay bid until 12pm
{<2, \$0>, <1, \$60>}

Dynamic allocation problems

…are everywhere in computer science

- MoteLab (Harvard)
  - distributed sensor network testbed
  - researchers compete for the right to sense, aggregate and propagate readings
- PlanetLab (Princeton)
  - global overlay network on the Internet
  - supports network research, long-running services
- Grid computing
  - much of science research is now intensively computational
  - globally-distributed computational infrastructure
- Network resource allocation
  - e.g. dynamic negotiation for WiFi wireless port at Starbucks

Many systems are simultaneously both computational and economic systems.

…are can be found in e-commerce, elsewhere

- Sequential auctions on eBay
  - e.g. auctions for LCDs, each bidder wants one
- Expiring goods
  - e.g. auctions for last-minute air-line tickets
- Online advertisement
  - e.g. adword auction of google

Basic Set-up for Online Auctions

- Type $\theta_i = (a_i, d_i, w_i)$. Discrete time periods.
- Arrival time: a_i. Departure time: d_i .Value, w_i
- $k \geq 1$ goods to sell

- Quasi-linear utility:
  $u_i = w_i$ - price, if we assign agent i some time in $[a_i, d_i]$
  $= 0$, otherwise

- Auction: $A = \langle f, p \rangle$,
  - allocation rule, $f : \Theta^n \rightarrow$ Schedules
  - payment rule, $p : \Theta^n \rightarrow R^n$

- Truthful auction: reporting value $\langle a_i, d_i, w_i \rangle$ immediately upon arrival is a dominant strategy equilibrium (i.e. no benefit otherwise).

Setting

- Assume: agents cannot under-report a_i.
- Assume: values i.i.d. from some unknown distribution.
- Want good performance whatever the distribution is.
- Limited-supply ($k \geq 1$) of goods, sell in any period before time horizon, T.
  - single-unit
  - multi-unit
- Efficiency benchmark is the highest value in this case (in general $EFF(v) = \sum_{i \leq k} v^{(i)}$ for $k \geq 1$)
- Revenue benchmark is Vickrey price, the second highest value in this case (in general $\mathcal{F}^{(2)}(v) = \max_{2 \leq l \leq k} \{ l \cdot v^{(l)} \}$ for $k>1$, "omniscient revenue", c.f. Goldberg, Hartline et al.01)
- c-competitive if we get $1/c$ fraction of benchmark (in expectation)

Aside: The Online Selection Problem

- Remove incentives, and specialize to the case of disjoint arrival-departure intervals.

Aside: The Online Selection Problem

- Remove incentives, and specialize to the case of disjoint arrival-departure intervals.

- Reduces to the secretary problem:
  - interview n job applicants in random order, want to max prob of selecting best applicant (told n)
  - told *relative ordering* w.r.t. applicants already interviewed, must hire or pass

Aside: The Online Selection Problem

- Remove incentives, and specialize to the case of disjoint arrival-departure intervals.

- Reduces to the secretary problem:
  - interview n job applicants in random order, want to max prob of selecting best applicant (told n)
  - told *relative ordering* w.r.t. applicants already interviewed, must hire or pass

5	2	7	1,000	3	. . .

The Secretary Algorithm

- **Theorem** (Dynkin, 1962): The following stopping rule picks the maximum element with probability approaching $1/e$ as $n \to \infty$.
  - Observe the first $\lfloor n/e \rfloor$ elements. Set a threshold equal to the maximum quality seen so far.
  - Stop the next time this threshold is reached or exceeded.
- Asymptotic success probability of $1/e$ is best possible, *even if the numerical values of elements are revealed.*
  - i.e. optimal competitive ratio in the large n limit

Straw model for an Auction

- **Auction**: $p(t)=\infty$, then set $p(t \geq \tau)=\max_{i \leq j} w_i$ after $j=\lfloor n/e \rfloor$ bids received. Sell to first subsequent bid with $w_i \geq p(t)$, then set $p(t)=\infty$.
- **Not truthful**: Bidders that span transition, and with high enough values, should delay arrival.

Truthful Auction:

- At time τ (for $\lfloor n/e \rfloor$ arrival) let $p \geq q$ be the top two bids yet received.
- If any agent bidding p has not yet departed, sell to that agent (breaking ties randomly) at price q.
- Else, sell to the next agent whose bid is at least p (breaking ties randomly)

Adaptive Limited-Supply Auction

- At time τ, denoting arrival $j=\lfloor n/e \rfloor$, let $p \geq q$ be the top two bids yet received.
- If any agent bidding p has not yet departed, sell to that agent (breaking ties randomly) at price q.
- Else, sell to the next agent whose bid is at least p.

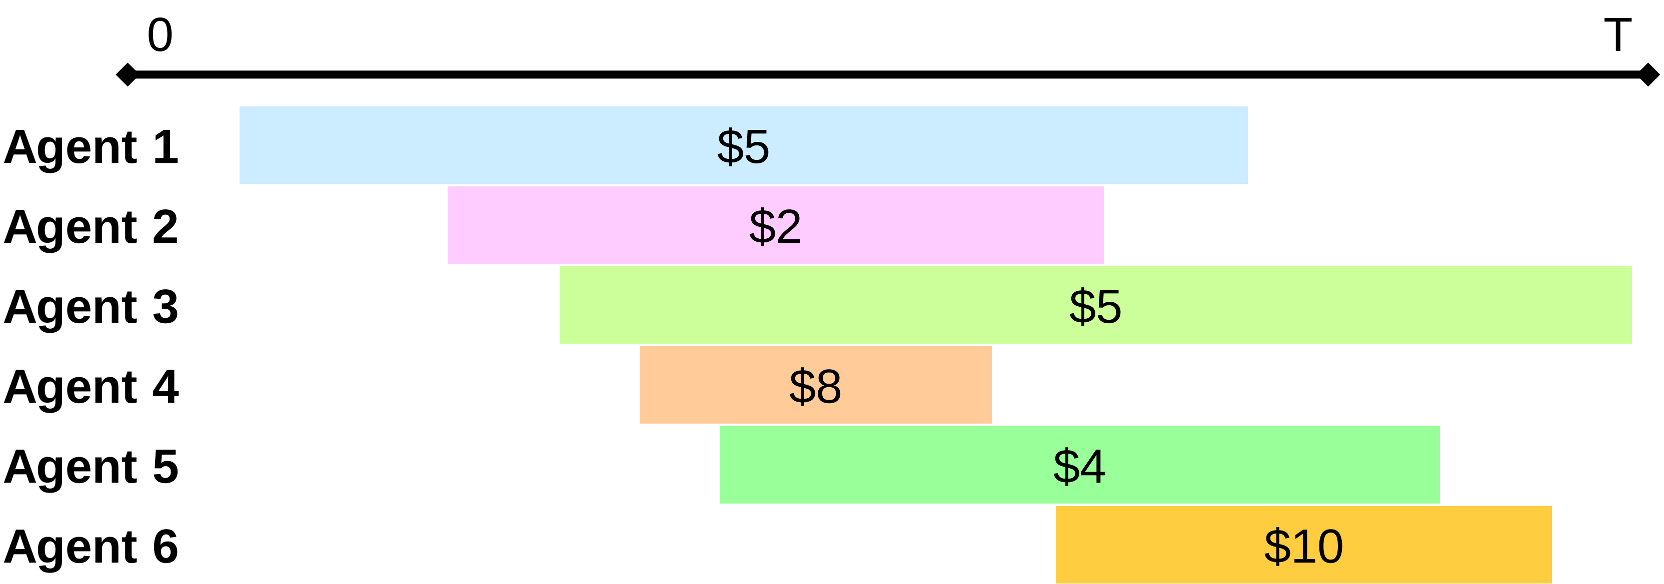

Adaptive Limited-Supply Auction

- At time τ, denoting arrival $j=\lfloor n/e \rfloor$, let $p \geq q$ be the top two bids yet received.
- If any agent bidding p has not yet departed, sell to that agent (breaking ties randomly) at price q.
- Else, sell to the next agent whose bid is at least p.

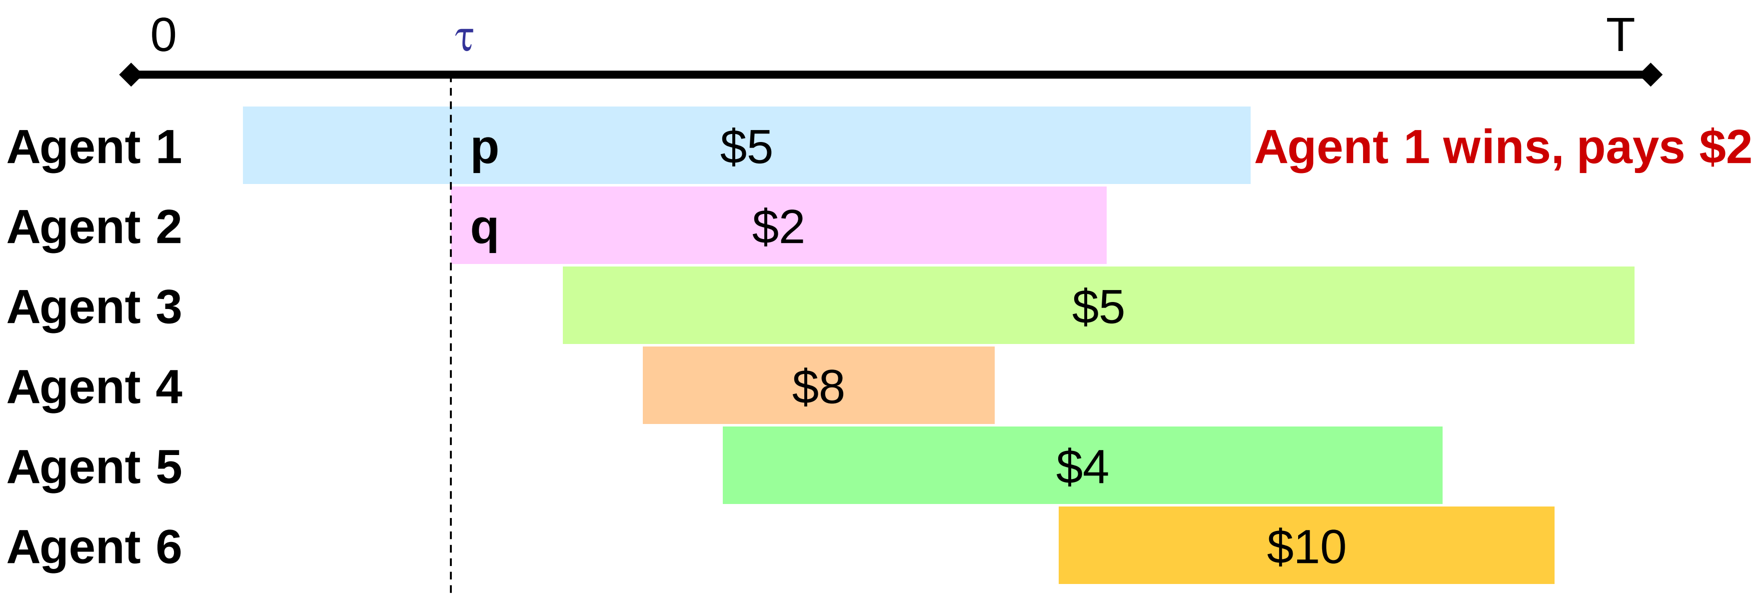

Adaptive Limited-Supply Auction

- At time τ, denoting arrival $j=\lfloor n/e \rfloor$, let $p \geq q$ be the top two bids yet received.
- If any agent bidding p has not yet departed, sell to that agent (breaking ties randomly) at price q.
- Else, sell to the next agent whose bid is at least p.

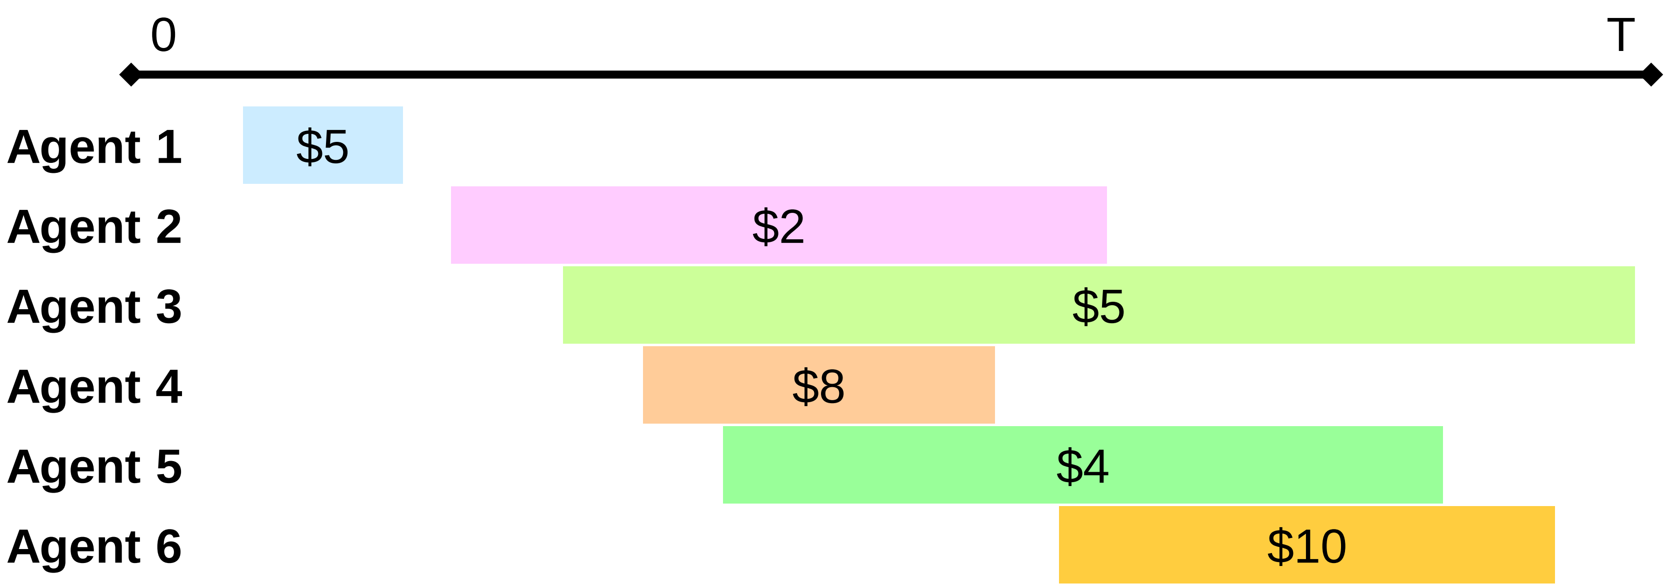

Adaptive Limited-Supply Auction

- At time τ, denoting arrival $j=\lfloor n/e \rfloor$, let $p \geq q$ be the top two bids yet received.
- If any agent bidding p has not yet departed, sell to that agent (breaking ties randomly) at price q.
- Else, sell to the next agent whose bid is at least p.

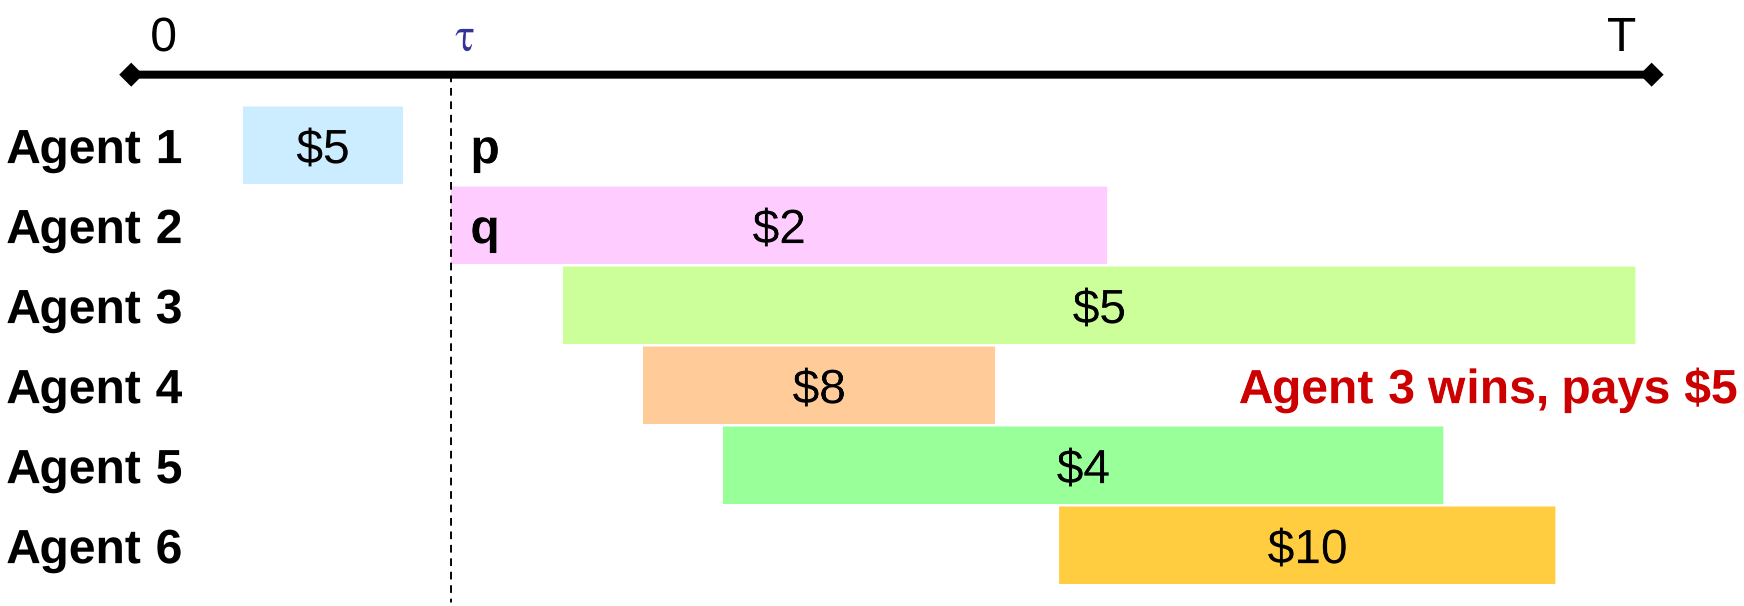

Analysis: Truthfulness

- If agent i wins, the price charged to her does not depend on her reported valuation.
- Possibility agent i wins is (weakly) increasing in w_i, hence no incentive to understate w_i.
- Reporting $w'_i > w_i$ cannot increase the possibility that agent i wins at a price $\leq w_i$, hence no incentive to overstate w_i.
- Price facing agent i is never influenced by d_i, so no incentive to misstate d_i

… just need to check effect of arrival time.

Analysis: Truthfulness

- **Claim**: Given two arrival times $a_i < a'_i$, it's always better to report a_i if possible.
- Let r, s be the $(\lfloor n/e \rfloor - 1)$-th and $\lfloor n/e \rfloor$-th arrival times excluding agent I (say $\lfloor n/e \rfloor = 3$ in this case).

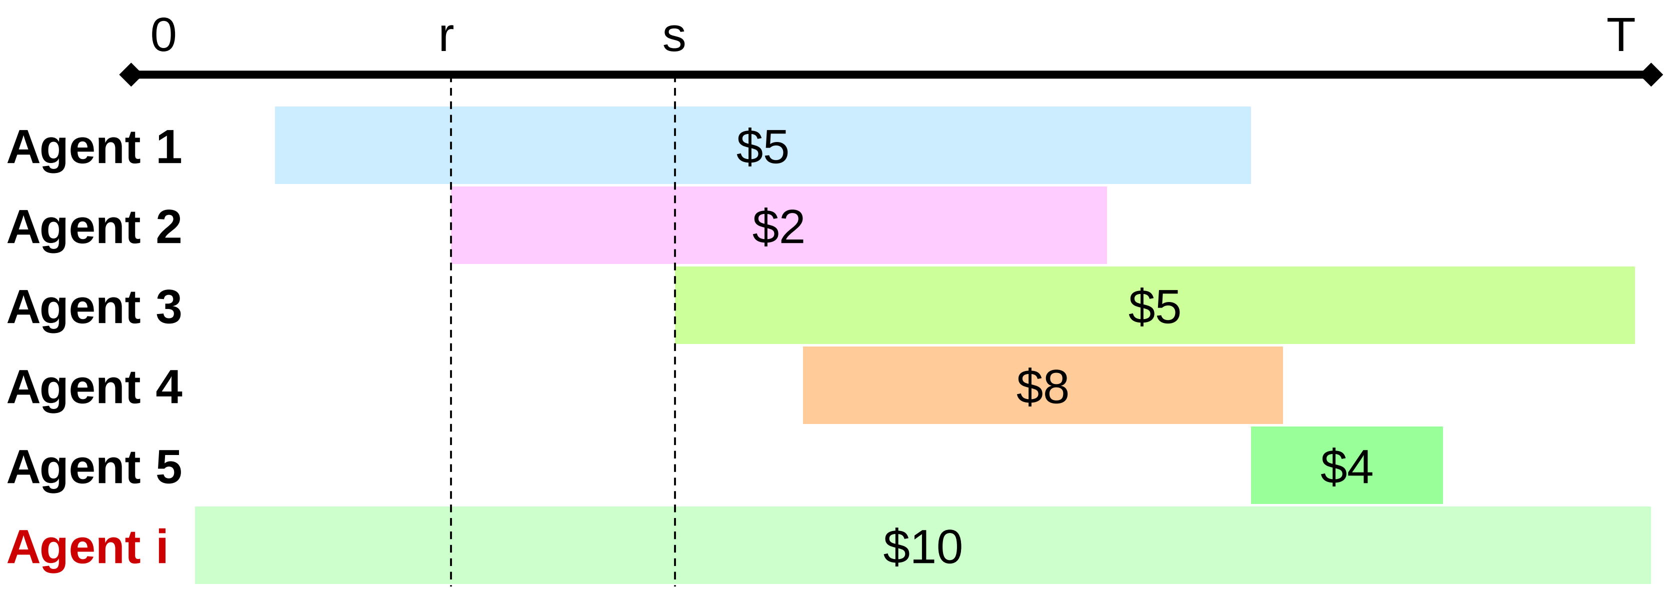

Analysis: Truthfulness

- Stating true arrival, agent 2 defines transition. Offered price $5 on transition.

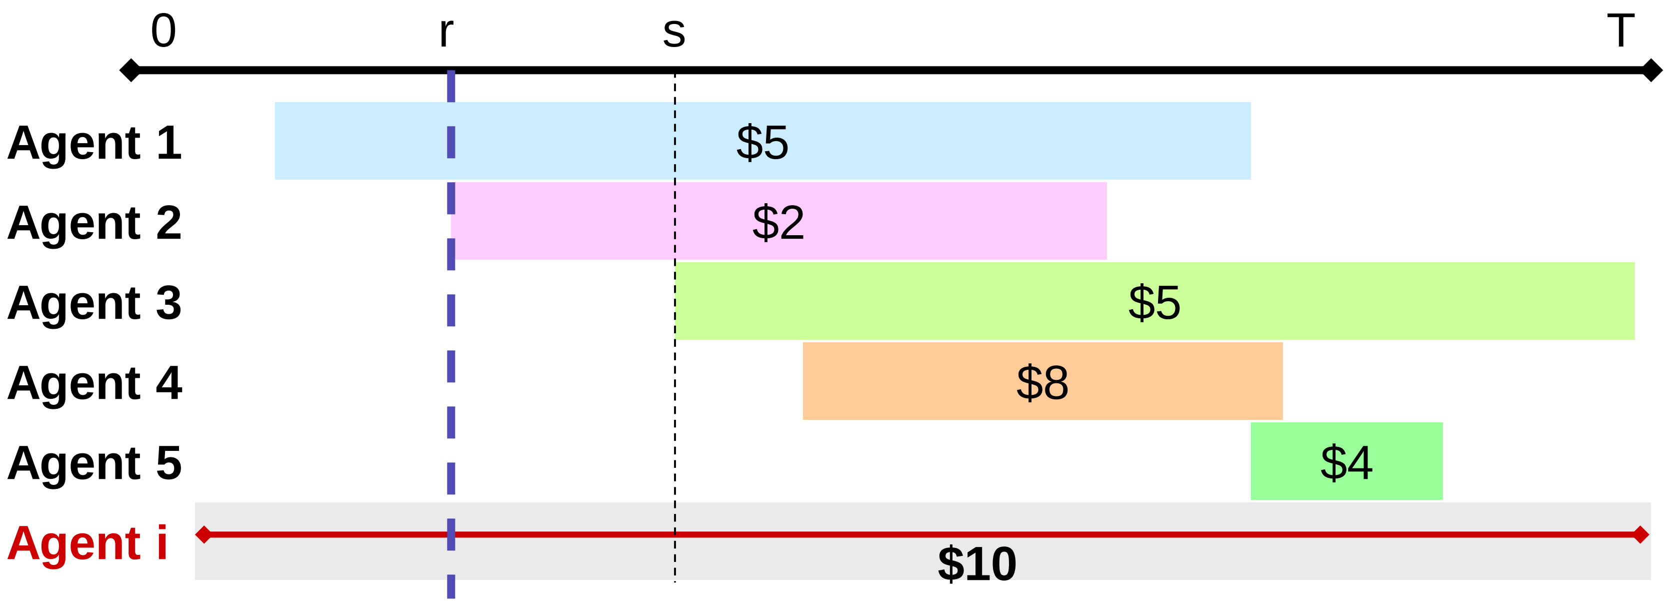

Analysis: Truthfulness

- Stating arrival time in $(a_i, r]$ changes nothing. Offered price \$5 on transition.

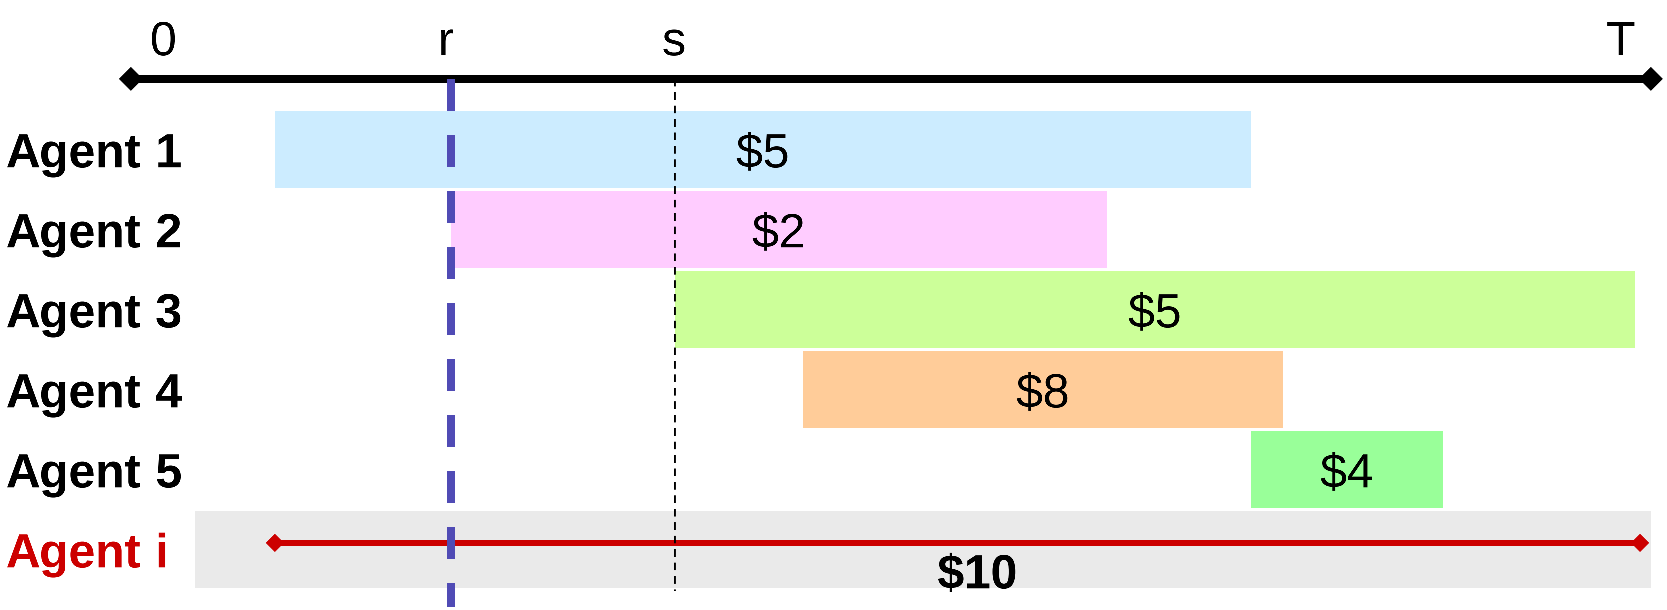

Analysis: Truthfulness

- Stating arrival time in $(a_i, r]$ changes nothing.
- Stating arrival time in (r,s) influences the transition time τ but not the pricing. Still offered price $5.

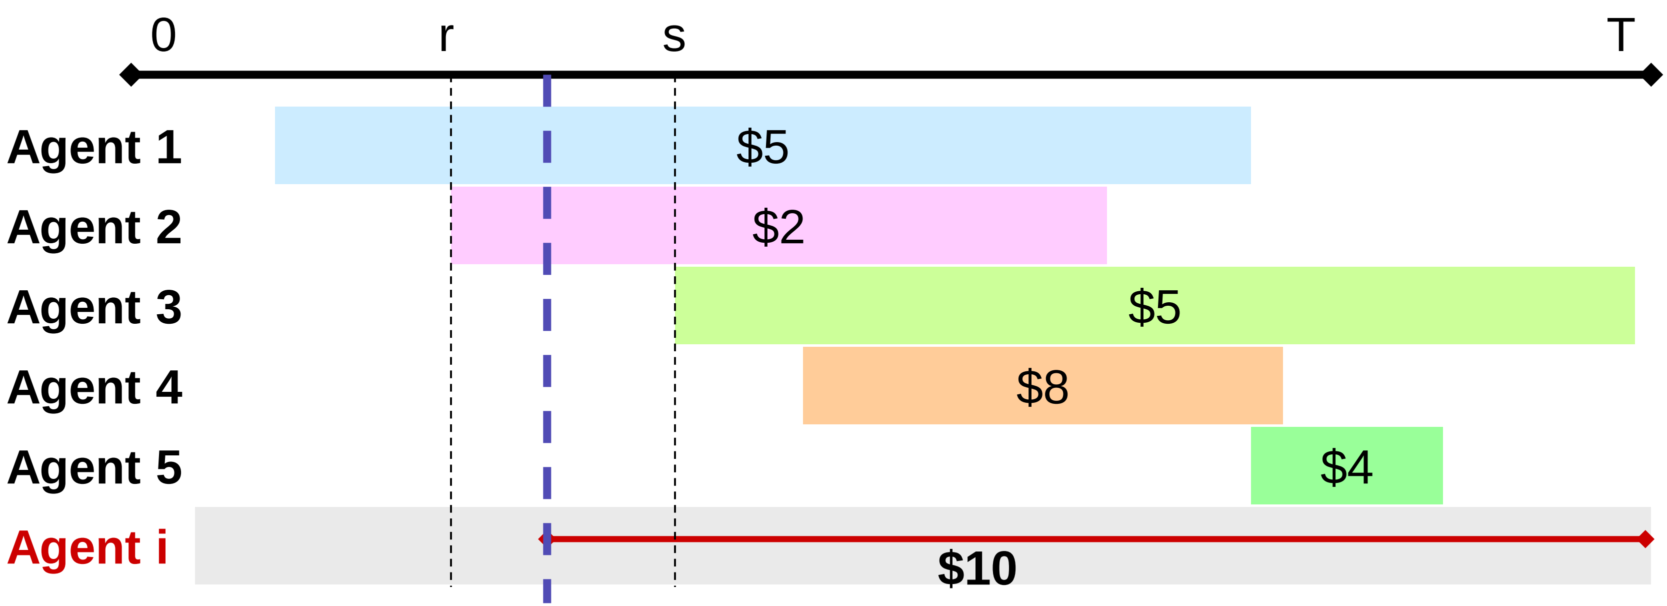

Analysis: Truthfulness

- Stating arrival time in $(a_i,r]$ changes nothing.
- Stating arrival time in (r,s) influences the transition time τ but not the pricing.
- Stating arrival time $\geq s$ influences the transition, but price not improved.

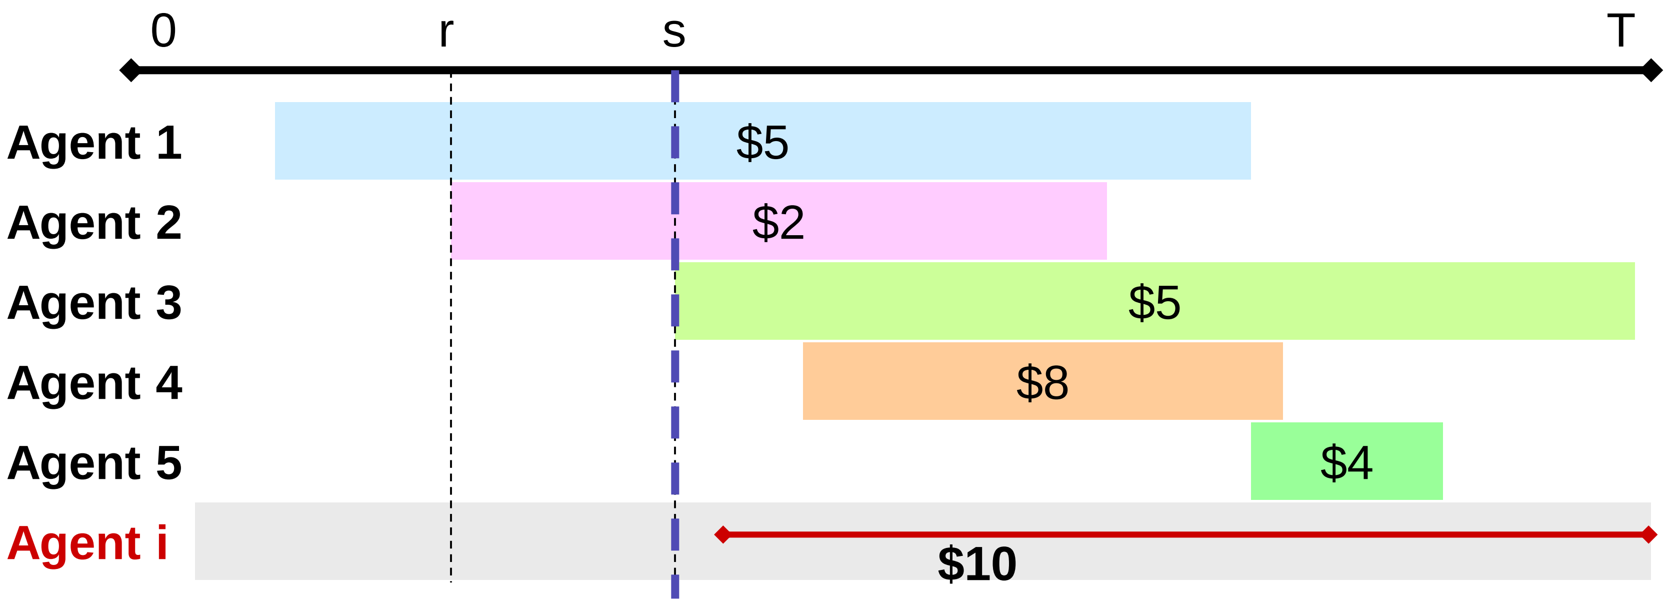

Analysis: Competitive Ratio

- **Claim**: Competitive ratio for efficiency is *e+o(1)*, assuming all valuations are distinct.
- Case 1: Item sells at time τ. Winner is highest bidder among first $\lfloor n/e \rfloor$. With probability ~1/e, this is also the highest bidder among all n agents.
- Case 2: Otherwise, the auction picks the same outcome as the secretary algorithm, whose success probability is ~1/e.

Analysis: Competitive Ratio

- **Claim**: Competitive ratio for revenue (wrt Vickrey) is $e^2+o(1)$, assuming all valuations are distinct.
- Estimate probability of selling to highest bidder at second-highest price. Use same two cases as before.
- Case 1: Probability ~(1/e)(1/e).
  - (prob 1/e that second highest also in first half)
- Case 2: Probability ~(1/e)(1/e).
  - (prob. that highest in first-half is the second-highest overall is 1/e conditioned on highest in second-half, prob. that choose highest in case 2 is 1/e)

- 4+o(1)-competitive for revenue (and also efficiency), by setting transition time at n/2.
- Lower-bounds of ~~2-competitive~~ e-competitive for efficiency, 1.5-competitive for revenue (in our model).

General approach -- Two phase

- "Learning phase"
  - use a sequence of bids to set price for rest of auction

Transition:

  - be sure that remains truthful for agents on transition

- "Accepting phase"
  - exploit information, retain truthfulness

Multi-Item Online Auction (k>1)

- (Learning) Choose pivotal bidder, $j \sim \text{Binom}(n, \frac{1}{2})$.
- (Transition) Sell up to $s = \lceil k/3 \rceil$ items at time τ, to agents present and bidding above $(s+1)$-st bid price so far.
- (Accepting) After τ, set $p =$ s-th bid and sell item to $\text{bid} \geq p$ while supply.

Multi-Item Online Auction (k>1)

- (Learning) Choose pivotal bidder, $j \sim \text{Binom}(n, \frac{1}{2})$.
- (Transition) Sell up to $s = \lceil k/3 \rceil$ items at time τ, to agents present and bidding above (s+1)-st bid price so far.
- (Accepting) After τ, set p= s-th bid price and sell item to bid $\geq$ p while supply.
- Truthfulness: similar, but more involved.
- Constant-competitive for efficiency.
- Constant-competitive with $\mathcal{F}^{(2)}$ for revenue, by setting $s = \lceil k/2 \rceil$, and adopting p to be the revenue-optimizing fixed price bid in first half, in accepting phase.
- Proofs are more involved.
- Tossing a fair coin gives a constant-competitive truthful algorithm for both efficiency and revenue.

Characterization of Truthful auctions

(H., Kleinberg, and Parkes., ACM-EC04)

- **Definition**. Allocation rule $f: \Theta^n \rightarrow \{0,1\}^n$ is **monotonic** if for every agent i and every $(\theta,\theta')\in\Theta^n$ with $[a'_i,d'_i]\subseteq[a_i,d_i]$, and $w_i > w'_i$, we have $f_i(\theta)\geq f_i(\theta')$.
- **Definition**. The "critical value" price is:

$$ps_i(a_i,d_i,\theta_{-i}) = \begin{cases} \min w'_i \text{ s.t. } f_i(\langle a_i,d_i,w'_i\rangle, \theta_{-i})=1 \\ \infty, \quad \text{if no such } w'_i \text{ exists} \end{cases}$$

- **Definition**. The "critical period" is the first $t\in[a_i,d_i]$ with minimal $ps_i(a_i,t,\theta_{-i})$.

> **Theorem**. An online auction is truthful if and only if the allocation rule, f, is monotonic, sets payment equal to critical value, and assigns item after the critical period.

- The only if proof is involved and uses an agent-independent price scheduling technique.

Extension: Multi-choice Secretary Problem

(Kleinberg, SODA05,
Immorlica, Kleingberg,, Mahdian, WINE06,
Babaioff, Immorlica, Kleingberg, SODA07,
Babaioff, Immorlica, Kempe, Kleingberg,
SIGecom Exch08,
Babaioff, Dinitz, Gupta, Immorlica, Talwas,
SODA09,
Bateni, Hajiaghayi, ZadiMoghaddam, TALG'13,
Etc.)

- choose k secretaries to maximize their joint performance
- e.g. secretaries should form a feasible set in a matroid
- e.g. their joint performance function is submodular
- Improves and generalize several bounds in optimal stopping theory
- constant-competitive for the submodular secretary problem
- log(n)-competitive for the matroid secretary problem
- $\log^2(n)$-competitive for the submodular matroid secretary problem

Extension: Reusable goods (Grid scheduling)

(H., Kleinberg, Mahdian, and Parkes, ACM-EC05)

Value	$100	$80	$60
Arrival:	11am	11am	12pm
Patience:	2hrs	2hrs	1hr
Duration:	1hr	1hr	1hr

- k goods in each time slot .
- Agent value $\langle a_i, d_i, w_i \rangle$. Value for one time slot in $[a_i, d_i]$.

Allocation rule for $k = 1$: In each period, t, allocate the good to the highest unassigned bid.

Payment rule for $k = 1$: Pay smallest amount could have bid and still received good (in some period).

Extension: Reusable goods (Grid scheduling)

(H., Kleinberg, Mahdian, and Parkes, ACM-EC05)

Value	\$100	\$80	\$60
Arrival:	11am	11am	12pm
Patience:	2hrs	2hrs	1hr
Duration:	1hr	1hr	1hr

- k goods in each time slot .
- Agent value $\langle a_i, d_i, w_i \rangle$. Value for one time slot in $[a_i, d_i]$.
- 2-competitive for efficiency (tight)
- log(n)-competitive for revenue with a randomized scheme (almost tight).
- Characterize truthful auctions with monotonic allocation rules

Extension: Model-based Online Auctions (v.s. Prior-free)

(H., Kleinberg, Sandholm, AAAI07)

- seller has distributional knowledge of the bid values (e.g., via the history of past transactions in the market)
- combine automated mechanism design and Prophet inequalities (a technique from optimal stopping theory)
- optimal efficiency/revenue assuming price sequence is non-decreasing (e.g. for airline tickets)
- Improve bounds for Prophet inequalityies with k stopping rules

Extension: Adwords Auction

(Mehta, Saberi, Vazirani, Vazirani, JACM)

- Internet search engine companies, such as Google, Yahoo and MSN
- Adwords Market:businesses place bids for individual keywords
- Online auction when assigning each search query to a bidder
- 1-1/e competitive for revenue (tight)

Other Extensions

- Auctions with expiring items [LN05]
- Auctions with unknown number of agents [HKS07]
- Practical implementations, e.g. in Tycoon (a market based distributed resource allocation system) [NBCSV05]
- Expressive online auctions [LKDP09]
- Multi-unit auctions with budgets [BCIMS05]
- Fair online equilibrium v.s. dominant online equilibrium (envy-freeness vs truthfulness) [GHKKKM05, DFHS08]
- etc.

Future Directions

- Real-World testing:
  - when is a prior-free approach preferable to a model-based approach? (noisy prior, prior-free on non-adversarial world.)
  - currently testing on an eBay problem, how useful is it to remove dynamic problem?
- Richer models:
  - current models insufficiently expressive, e.g. for grid computing
  - e.g., richer patience models, choices (A vs. B), bundles of resources (A and B).
- Better understanding of social networks
  - e.g. twitter, facebook, [GHISR, WINE'09]
  - Applications in online auctions and ad auctions

Thank you!
Merci!
謝謝
Todah!
Danke!
Dank Je Wel!
Gracias!
تشكر
どうもありがとう
どうもありがとうございました
Thank you!
感謝
Todah!
THANK YOU!
Merci!
GRACIAS!
MERCI!