Algorithmic game theory

Martin Balko

11th lecture

December 13th 2024

What we learned last time

- For single-item auctions we have found an awesome auction (Vickrey's auction) with the following three properties:
  - DSIC: everybody has dominant strategy "bid truthfully" which guarantees non-negative utility,
  - strong performance: maximizing social surplus,
  - computational efficiency: running in polynomial time.
- We then generalized single-item auctions to single-parameter environments.
  - 1 seller and n bidders, the seller collects the bids $b = (b_1, \dots b_n)$,
  - the seller chooses an allocation $x(b) = (x_1(b), \dots, x_n(b))$ from X,
  - the seller sets payments $p(b) = (p_1(b), \dots, p_n(b))$.

- Is there awesome mechanism (x, p) for single-parameter environments?
- We started by looking for DSIC mechanisms and saw a powerful tool for designing them.

Myerson's lemma

- An allocation rule x is implementable if there is a payment rule p such that (x, p) is DSIC.
- An allocation rule x is monotone if, for every bidder i and all bids b_{-i}, the allocation $x_i(z; b_{-i})$ is nondecreasing in z.

Myerson's lemma

In a single-parameter environment, the following three claims hold.

(a) An allocation rule is implementable if and only if it is monotone.

(b) If an allocation rule x is monotone, then there exists a unique payment rule p such that (x, p) is DSIC (assuming $b_i = 0$ implies $p_i(b) = 0$).

(c) For every i, the payment rule p is given by the following explicit formula

$$p_i(b_i; b_{-i}) = \int_0^{b_i} z \cdot \frac{\mathrm{d}}{\mathrm{d}z} x_i(z; b_{-i}) \, \mathrm{d}z.$$

- We saw two applications: single-item and sponsored-search auctions.

Social surplus maximization

- So far, we were designing DSIC mechanisms for single-parameter environments that maximize the social surplus $\sum_{i=1}^{n} v_i x_i(b)$.
- We did not care about the revenue $\sum_{i=1}^{n} p_i(b)$.
- This makes sense in some real-world scenarios where revenue is not the first-order objective. For instance, in government auctions (e.g., to sell wireless spectrum).

Source: youtube.com

- In single-item auctions, maximizing social surplus is done by Vickrey's auction. In general single-parameter environments, we use Myerson's lemma.
- Today, we try to maximize the revenue.

Revenue maximizing auctions

Source: Reprofoto

Revenue maximization

- How to maximize the revenue $\sum_{i=1}^{n} p_i(b)$?

Source: https://merger.com/recurring-revenue/

- The situation then becomes more complicated, but we will see some nice results today.

Why is it more complicated?

- Consider single-item auction with 1 bidder with valuation v.

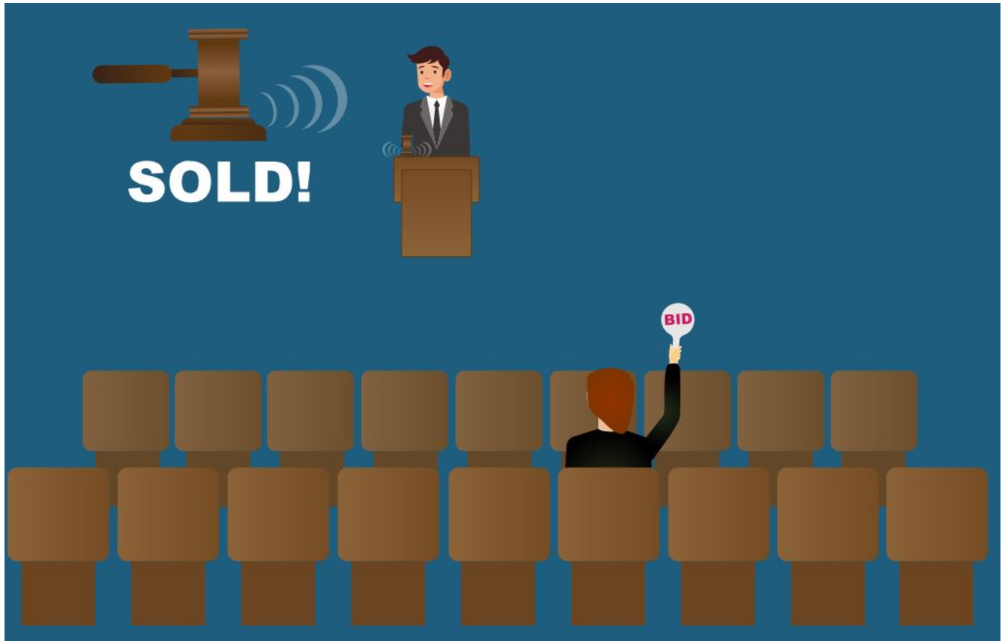

Source: https://auctions.propertysolvers.co.uk/

- the only DSIC auction: the seller posts a price r, then his revenue is either r if $v \geq r$ and 0 otherwise.
- Maximizing the social surplus is trivial by putting $r = 0$.
- However, when maximizing the revenue, it is not clear how we should set r, since we do not know the valuation v.

New model

- Thus, we need a model to reason about trade-offs between our performance on different inputs.
- We introduce some randomness outside of discrete probability spaces, so let us first recall some notions from the probability theory.
  - If F is a probability distribution drawing some value v, then $F(z)$ is the probability that v has value at most z.
  - If F is a probability distribution with density f and with support $[0, v_{max}]$, then

$$f(x) = \frac{\mathrm{d}}{\mathrm{d}x}F(x) \text{ and } F(x) = \int_0^x f(z)\,\mathrm{d}z.$$

  - For a random variable X, the expected value of X is

$$\mathbb{E}_{z\sim F}[X(z)] = \int_0^{v_{max}} X(z) \cdot f(z)\,\mathrm{d}z.$$

Bayesian model

- The Bayesian model consists of
  - a single-parameter environment (x, p),
  - for each bidder i, the private valuation v_i of i is drawn from a probability distribution F_i with density function f_i and with support contained in $[0, v_{max}]$. We assume that the distributions $F_1, \ldots, F_n$ are independent, but not necessarily the same.
- The distributions $F_1, \ldots, F_n$ are known to the mechanism designer. The bidders do not know the distributions $F_1, \ldots, F_n$. Since we discuss only DSIC auctions, the bidders do not need to know $F_1, \ldots, F_n$, as they have dominant strategies.
- The goal is to maximize, among all DSIC auctions, the expected revenue

$$\mathbb{E}_{v=(v_1,\ldots,v_n)\sim(F_1\times\cdots\times F_n)}\left[\sum_{i=1}^{n} p_i(v)\right],$$

where the expectation is taken with respect to the given distribution $F_1 \times \cdots \times F_n$ over the valuations $(v_1, \ldots, v_n)$.

Example: single-bidder single-item auction

- Consider the example with a single bidder in a single-item auction. It is now easier to handle.
- The expected revenue of a posted price r equals $r \cdot (1 - F_1(r))$, since $(1 - F_1(r))$ is the probability that v is at least r.
- Given a probability distribution F_1, it is usually a simple matter to maximize this value.
- For example, if F_1 is the uniform distribution on $[0, 1]$, where $F_1(x) = x$ for every $x \in [0, 1]$, then we achieve the maximum revenue by choosing $r = 1/2$.
- The posted price r that makes the expected revenue as high as possible is called the monopoly price.

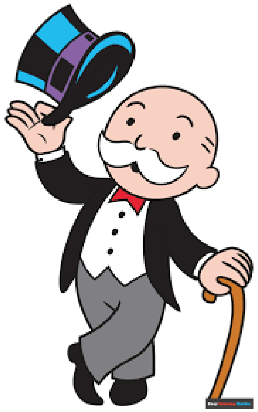

Example: two-bidders single-item auction

- Consider a single-item auction with two bidders that have their valuations v_1 and v_2 drawn uniformly from $[0, 1]$.

- Now, there are more DSIC auctions to think of, for example, the Vickrey's auction, for which the expected revenue is $1/3$ (Exercise).
- Besides Vickrey's auctions, there are other DSIC auctions that perform better with respect to the expected revenue (Exercise).
- Today, we describe such optimal auctions in more general settings using Myerson's lemma.

Maximizing expected revenue

- To characterize DSIC auctions that maximize the expected revenue, we find a more useful formula for the expected revenue.

Theorem 3.13

Let (x, p) be a DSIC mechanism in a single-parameter environment forming a Bayesian model with n bidders, probability distributions $F_1, \ldots, F_n$, and density functions $f_1, \ldots, f_n$. Let $F = F_1 \times \cdots \times F_n$. Then

$$\mathbb{E}_{v \sim F}\left[\sum_{i=1}^{n} p_i(v)\right] = \mathbb{E}_{v \sim F}\left[\sum_{i=1}^{n} \varphi_i(v_i) \cdot x_i(v)\right],$$

where

$$\varphi_i(v_i) = v_i - \frac{1 - F_i(v_i)}{f_i(v_i)}$$

is the virtual valuation of bidder i.

Maximizing expected revenue: remarks

- In other words, for every DSIC auction, the expected revenue is equal to the expected virtual social surplus $\sum_{i=1}^{n} \varphi_i(v_i) \cdot x_i(v)$. So when maximizing the expected revenue, we end up with maximizing a similar value as before.
- Note that the virtual valuation $\varphi_i(v_i) = v_i - \frac{1-F_i(v_i)}{f_i(v_i)}$ of a bidder can be negative and that it depends on v_i and F_i, and not on parameters of the others.
- Example: consider a bidder i with valuation v_i drawn from the uniform distribution on $[0,1]$. Then $F_i(z) = z$, $f_i(z) = 1$, and $\varphi_i(z) = z - (1-z)/1 = 2z - 1 \in [-1,1]$.
- A coarse intuition behind the virtual valuations: the first term v_i in this expression can be thought of as the maximum revenue obtainable from bidder i and the second term as the inevitable revenue loss caused by not knowing v_i in advance.

Maximizing expected revenue: proof I

- Since (x, p) is DSIC, we assume $b_i = v_i$ for each bidder i.
- By Myerson's lemma, we have the following payment formula

$$p_i(b_i; b_{-i}) = \int_0^{b_i} z \cdot \frac{d}{dz} x_i(z; b_{-i})\, dz.$$

- It can be shown that this formula holds for any monotone function $x_i(z; b_{-i})$, if we suitably interpret the derivation $\frac{d}{dz} x_i(z; b_{-i})$. All arguments in this proof, such as integration by parts, can be made rigorous for bounded monotone functions, but we omit the details.
- We fix a bidder i and the bids v_{-i} of others. By the payment formula, the expected payment of i for v_{-i} is determined by x_i and equals

$$\begin{aligned}
\mathbb{E}_{v_i \sim F_i}[p_i(v)] &= \int_0^{v_{max}} p_i(v) f_i(v_i)\, dv_i \\
&= \int_0^{v_{max}} \left(\int_0^{v_i} z \cdot \frac{d}{dz} x_i(z; v_{-i})\, dz \right) f_i(v_i)\, dv_i.
\end{aligned}$$

Maximizing expected revenue: proof II

- After reversing the integration order (Fubini's theorem), we obtain

$$\int_0^{v_{max}} \left(\int_0^{v_i} z \cdot \frac{dx_i(z; v_{-i})}{dz} \, dz \right) f_i(v_i) \, dv_i = \int_0^{v_{max}} \left(\int_z^{v_{max}} f_i(v_i) \, dv_i \right) z \cdot \frac{dx_i(z; v_{-i})}{dz} \, dz.$$

- We have used the fact that $z \leq v_i$. Since f_i is the density function, the inner integral can be simplified and we get

$$\int_0^{v_{max}} (1 - F_i(z)) \cdot z \cdot \frac{dx_i(z; v_{-i})}{dz} \, dz.$$

- We apply integration by parts, that is $\int fg' = fg - \int f'g$, choosing $f(z) = (1 - F_i(z)) \cdot z$ and $g'(z) = \frac{d}{dz} x_i(z; v_{-i})$, we get

$$[(1 - F_i(z)) \cdot z \cdot x_i(z; v_{-i})]_0^{v_{max}} - \int_0^{v_{max}} x_i(z; v_{-i}) \cdot (1 - F_i(z) - z f_i(z)) \, dz.$$

- The first term is zero, as it equals to 0 for $z = 0$ and also for $z = v_{max}$, as v_{max} is a finite number.

Maximizing expected revenue: proof III

- Thus, so far we derived that the expected revenue of i for v_{-i} is

$$0 - \int_0^{v_{max}} x_i(z; v_{-i}) \cdot (1 - F_i(z) - zf_i(z))\, dz.$$

- By the definition of $\varphi_i(v_i)$, the second term can be rewritten as

$$\int_0^{v_{max}} \left(z - \frac{1 - F_i(z)}{f_i(z)}\right) x_i(z; v_{-i}) f_i(z)\, dz = \int_0^{v_{max}} \varphi_i(z) x_i(z; v_{-i}) f_i(z)\, dz.$$

- This is an expected value where $z \sim F_i$ and thus, for every v_{-i},

$$\mathbb{E}_{v_i \sim F_i}\left[p_i(v_i; v_{-i})\right] = \mathbb{E}_{v_i \sim F_i}\left[\varphi_i(v_i) \cdot x_i(v_i; v_{-i})\right].$$

- Taking the expectation with respect to v_{-i}, this is true also for $v \sim F$.
- Finally, applying the linearity of expectation twice, we get the final form

$$\mathbb{E}_{v \sim F}\left[\sum_{i=1}^{n} p_i(v)\right] = \sum_{i=1}^{n} \mathbb{E}_{v \sim F}\left[p_i(v)\right] = \sum_{i=1}^{n} \mathbb{E}_{v \sim F}\left[\varphi_i(v_i) x_i(v)\right] = \mathbb{E}_{v \sim F}\left[\sum_{i=1}^{n} \varphi_i(v_i) x_i(v)\right]$$

$\square$

Maximizing expected virtual social surplus

- By Theorem 3.13, "maximizing expected revenue = maximizing expected virtual social surplus". We characterize auctions that do that.
- We start with single-item auctions. We assume that there is a probability distribution F such that $F_1 = \cdots = F_n = F$ (thus all virtual valuations $\varphi_1, \ldots, \varphi_n$ are equal to a function φ). We also assume that F is regular, that is, $\varphi(v)$ is strictly increasing in v.
- All this holds if F is the uniform probability distribution $F(v) = v$ on $[0, 1]$, as then $\varphi(v) = v - (1 - F(v))/f(v) = 2v - 1$.

- We show that for such auctions the Vickrey auction with reserved price maximizes expected revenue.
- In a Vickrey auction with reserve price r, the item is given to the highest bidder, unless all bids are less than r, in which case no one gets the item. The winner (if any) is charged the second-highest bid or r, whichever is larger.

Vickrey with reserve price is optimal

Proposition 3.14

Let F be a regular probability distribution with density f and the virtual valuation φ and let $F_1, \ldots, F_n$ be independent probability distributions on valuations of n bidders such that $F = F_1 = \cdots = F_n$ (and thus $\varphi = \varphi_1 = \cdots = \varphi_n$). Then, Vickrey auction with reserve price $\varphi^{-1}(0)$ maximizes the expected revenue.

- Proof: By Theorem 3.13, we want to maximize the expected virtual social surplus. To do so, we can choose $x(b) \in X$ for each input b, while we do not have any control over F and φ.
- In a single-item auction, the feasibility constraint is $\sum_{i=1}^{n} x_i(b) \leq 1$, so we give the item to bidder i with the highest $\varphi(b_i)$. If every bidder has a negative virtual valuation, then we do not award the item to anyone.
- It remains to show that this allocation rule x is monotone, as then, by Myerson's lemma, it can be extended to a DSIC auction (x, p). Since F is regular, the function φ is strictly increasing. Thus, x is monotone and we get Vickrey's auction with a reserve price $\varphi^{-1}(0)$. □

- The Vickrey auction with reserve price is used by eBay.

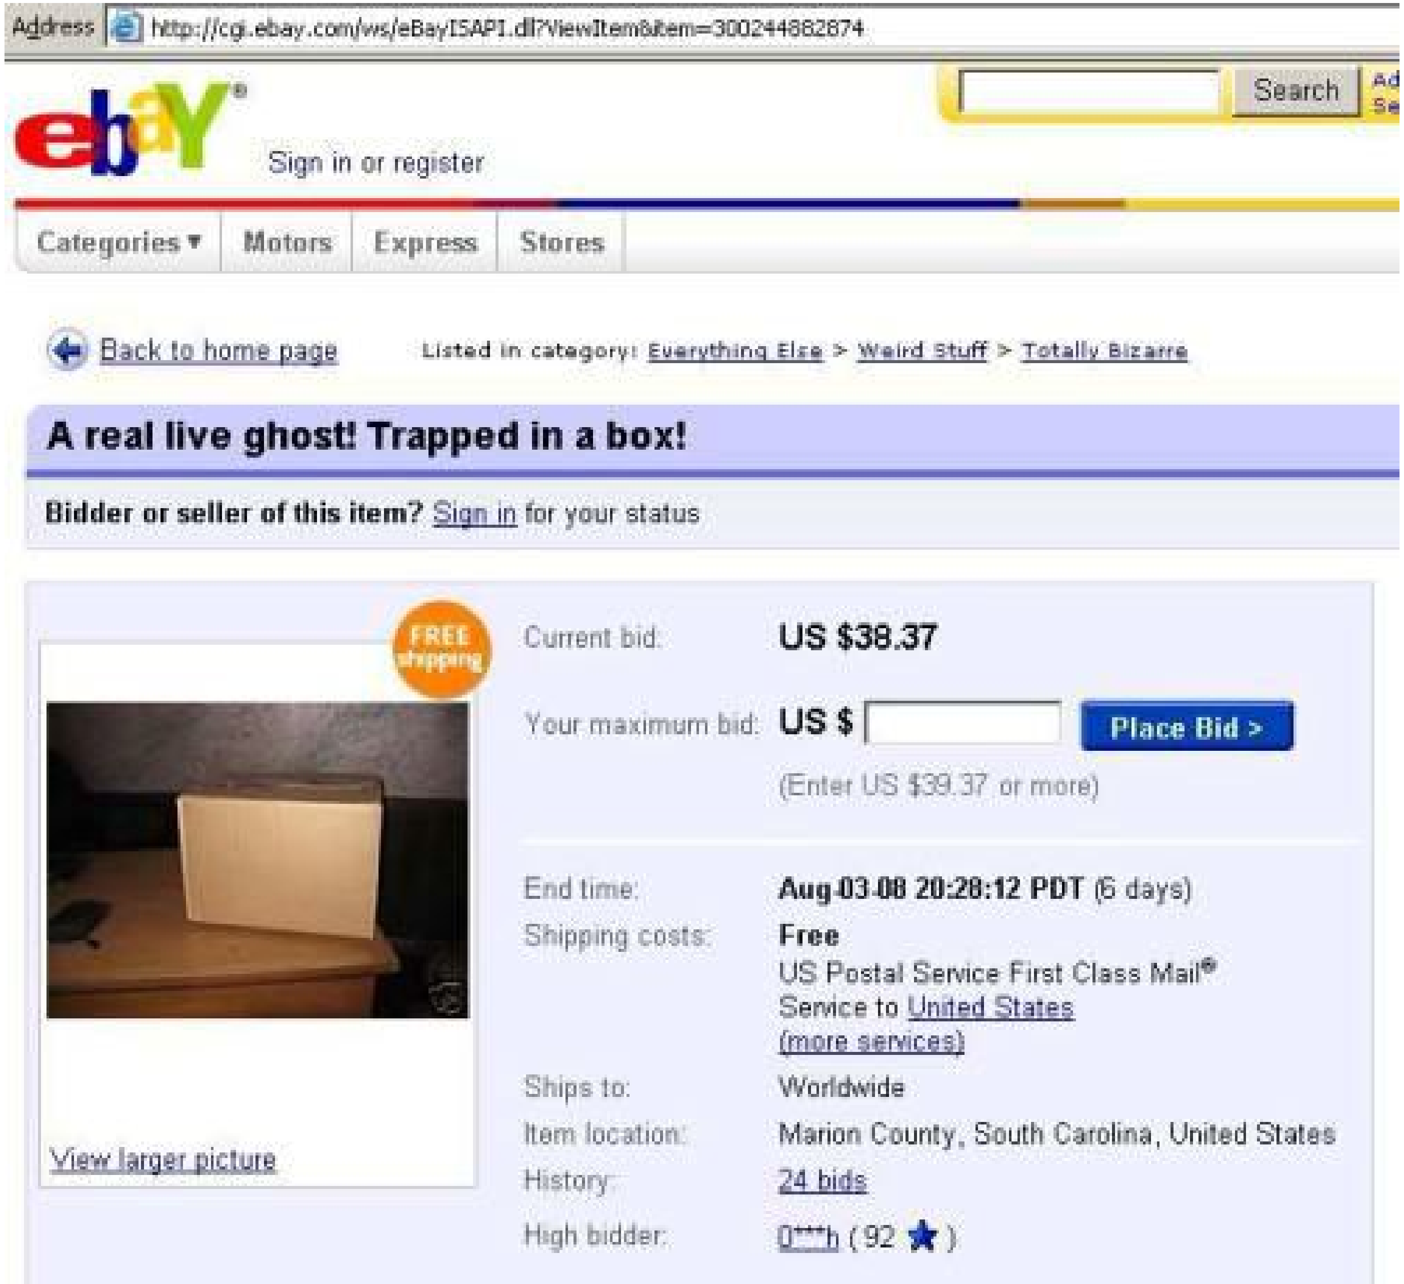

Source: https://goodgearguide.com.au

Reserve price and maximizing revenue

- The theory was developed by Roger Myerson.
- The result can be extended to single-parameter environments. With more work we can prove a version for distributions that are not regular.
- Very roughly, if the seller believes that bidders have high valuations, he should set a high reserve price accordingly.

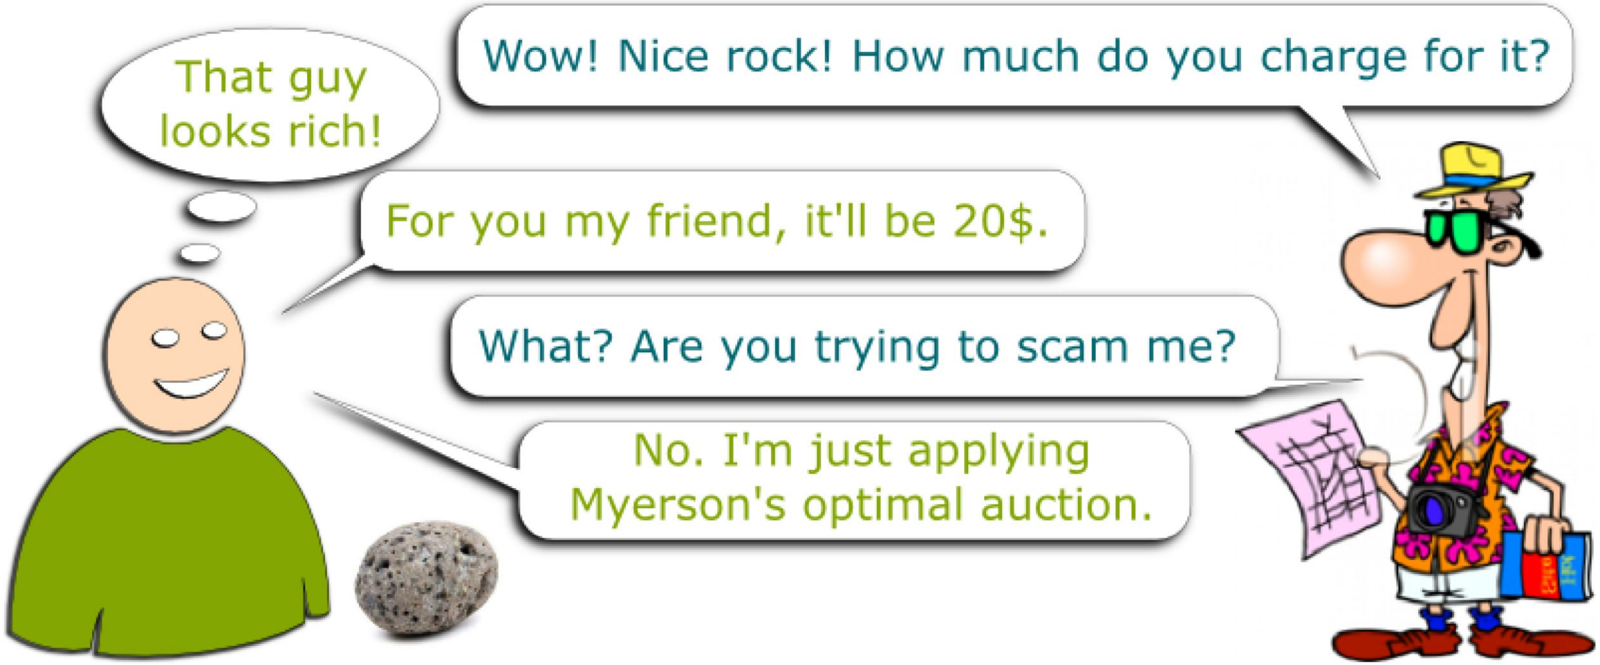

Source: https://www.science4all.org/article/auction-design/

Optimal auctions more generally

- There are also optimal DISC auctions even if we relax the conditions by not insisting on $F_1, \ldots, F_n$ being identical. However, such optimal auctions can get weird, and they does not generally resemble any auctions used in practice (Exercise).

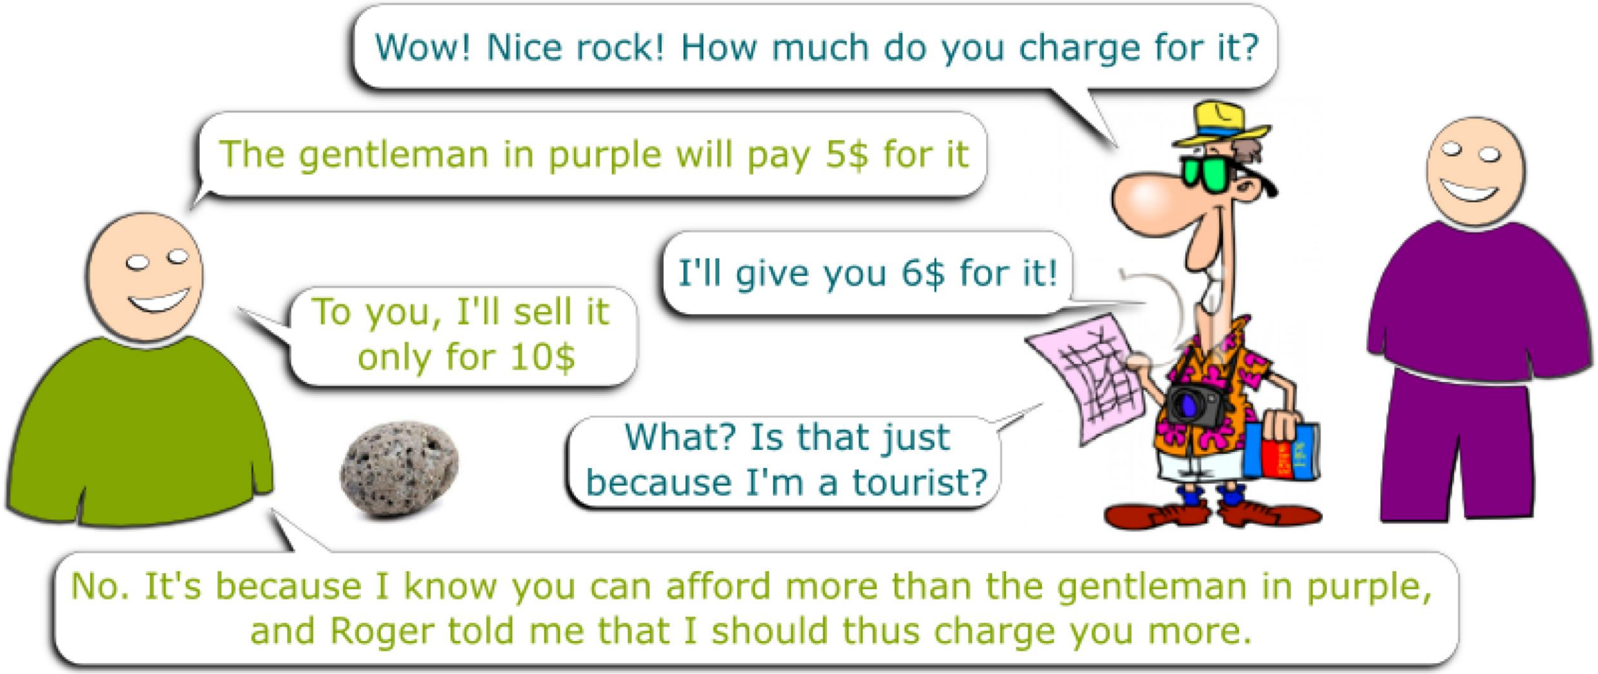

Source: https://www.science4all.org/article/auction-design/

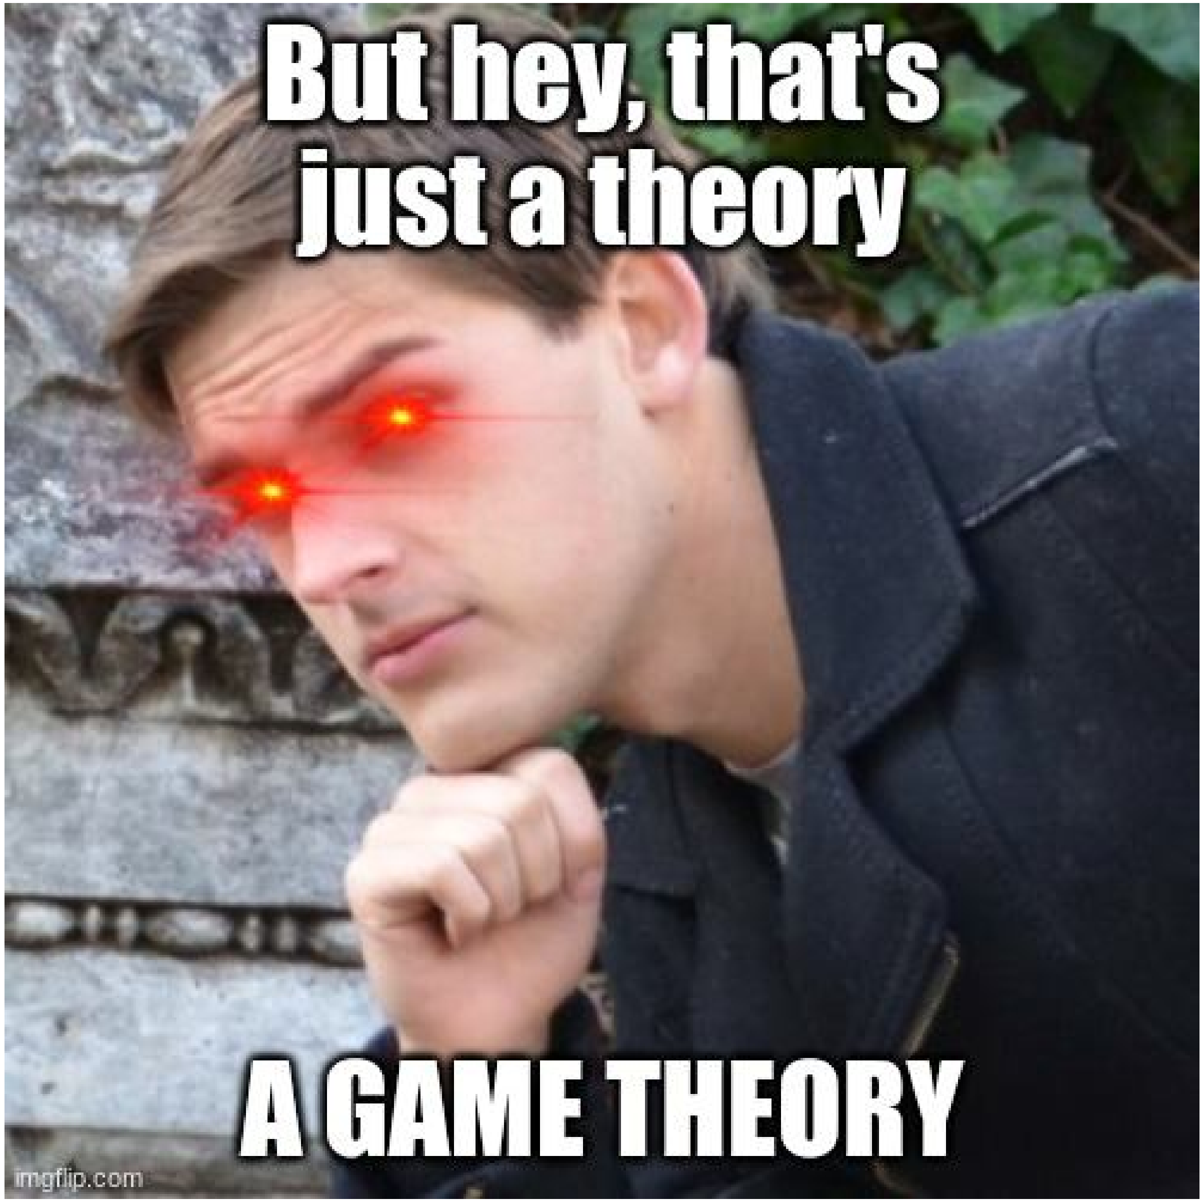

Source: https://www.reddit.com/

Thank you for your attention.