Algorithmic Game Theory and Applications

Optimal Auctions

Single-Item Auctions: Quick Refresher

Model: Single-Item Auctions

$[0,\infty)$

- **Bidders** $i = 1,2,\ldots,n$
- Each bidder i has **value** $v_i \in \mathbb{R}_+$ for the item
  - Submits a **bid** $b_i \in \mathbb{R}_+$

- An **mechanism** $\mathscr{M} = (\boldsymbol{a}, \boldsymbol{p})$ consists of
  - **allocation** rule $\boldsymbol{a} = (a_1, a_2, \ldots, a_n) : \mathbb{R}_+^n \longrightarrow [0,1]^n$, where $\sum_{i=1}^n a_i(\boldsymbol{b}) \leq 1$,
  - **payment** rule $\boldsymbol{p} = (p_1, p_2, \ldots, p_n) : \mathbb{R}_+^n \longrightarrow \mathbb{R}_+^n$

- Bidder **utility**: $u_i(\boldsymbol{b}; v_i) := a_i(\boldsymbol{b}) \cdot v_i - p_i(\boldsymbol{b})$

- Mechanism $\mathscr{M}$ is **truthful** (DSIC) if: $u_i(v_i, \boldsymbol{b}_{-i}; v_i) \geq u_i(b_i, \boldsymbol{b}_{-i}; v_i)$ $\forall i\ \forall \boldsymbol{b}_{-i}\ \forall v_i, b_i$

  \+ Individual rationality (IR): $u_i(v_i, \boldsymbol{b}_{-i}; v_i) \geq 0$ $\forall i\ \forall \boldsymbol{b}_{-i}\ \forall v_i$

Myerson's Characterization

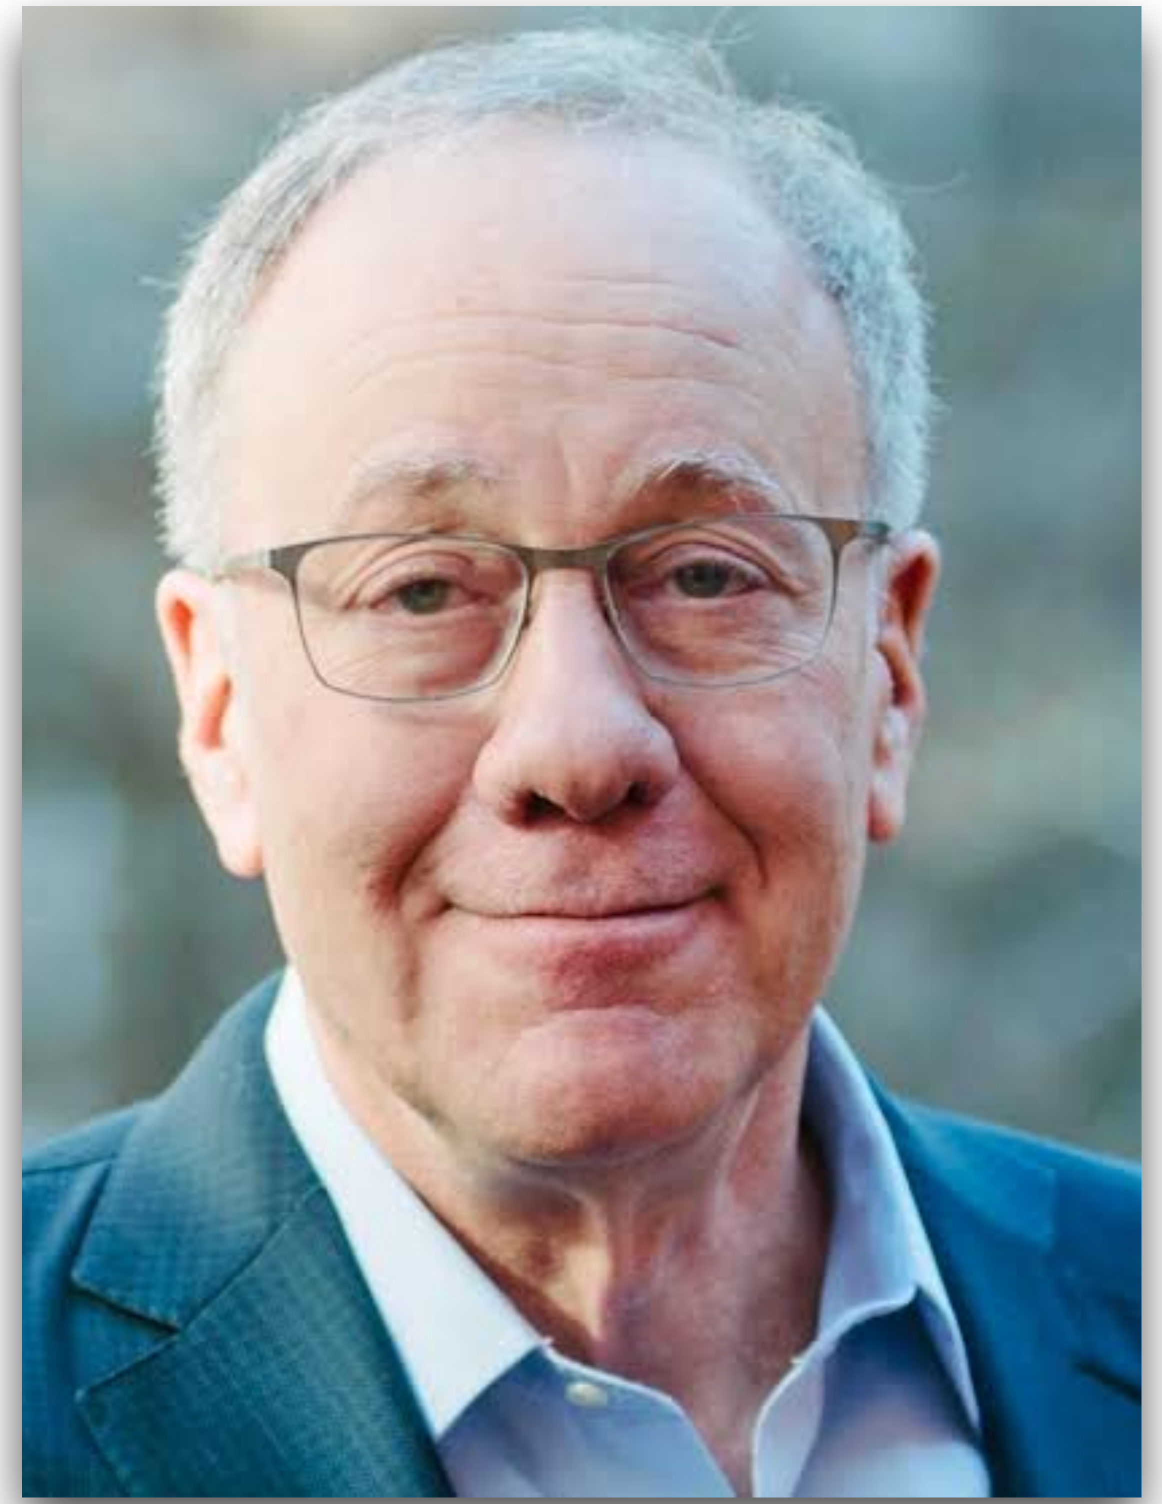

Roger Myerson (1951 -)

Nobel prize in Economics
(2007)

Myerson's Characterization [1981]

A mechanism $\mathscr{M} = (\boldsymbol{a}, \boldsymbol{b})$ is truthful if and only if, for any bidder i and all bid profiles $\boldsymbol{b}_{-i}$:

- $a_i(b_i, \boldsymbol{b}_{-i})$ is a *nondecreasing* function of b_i, and
- payments are given by $p_i(\boldsymbol{b}) = a_i(\boldsymbol{b}) \cdot b_i - \int_0^{b_i} a_i(t, \boldsymbol{b}_{-i})\, \mathrm{d}t$.

- The second-price auction (i.e., VCG) is truthful *and* welfare maximizing!
- What if our goal is to maximize the seller's *revenue* instead?
  - Shall we still *always* sell to the *highest* bidder?

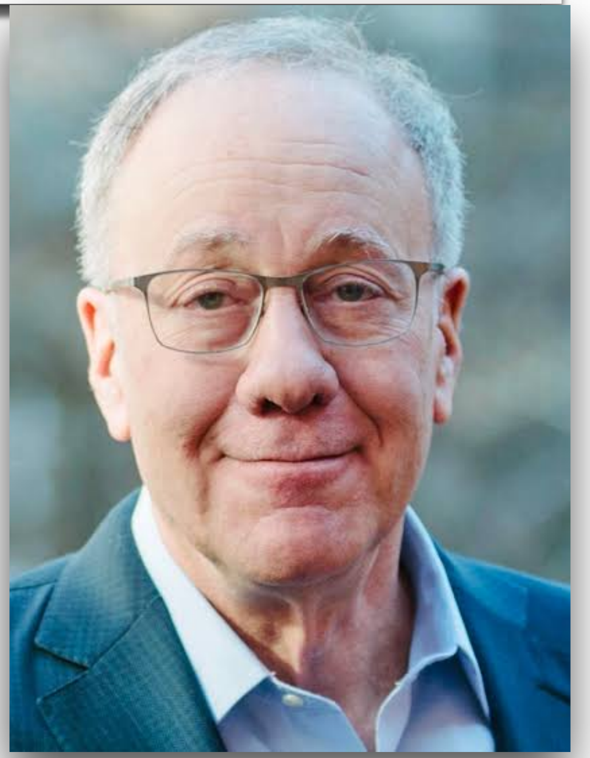

Roger Myerson

Welfare vs Revenue

Example: *Single-Bidder Deterministic* Auctions

- Always selling for free ($r = 0$) maximizes the welfare
    - However: this gives *zero revenue* to the seller!

- Where shall we set the selling price r, in order to guarantee "good" revenue?
    - *Highly* dependent on the (private) value v of the bidder.

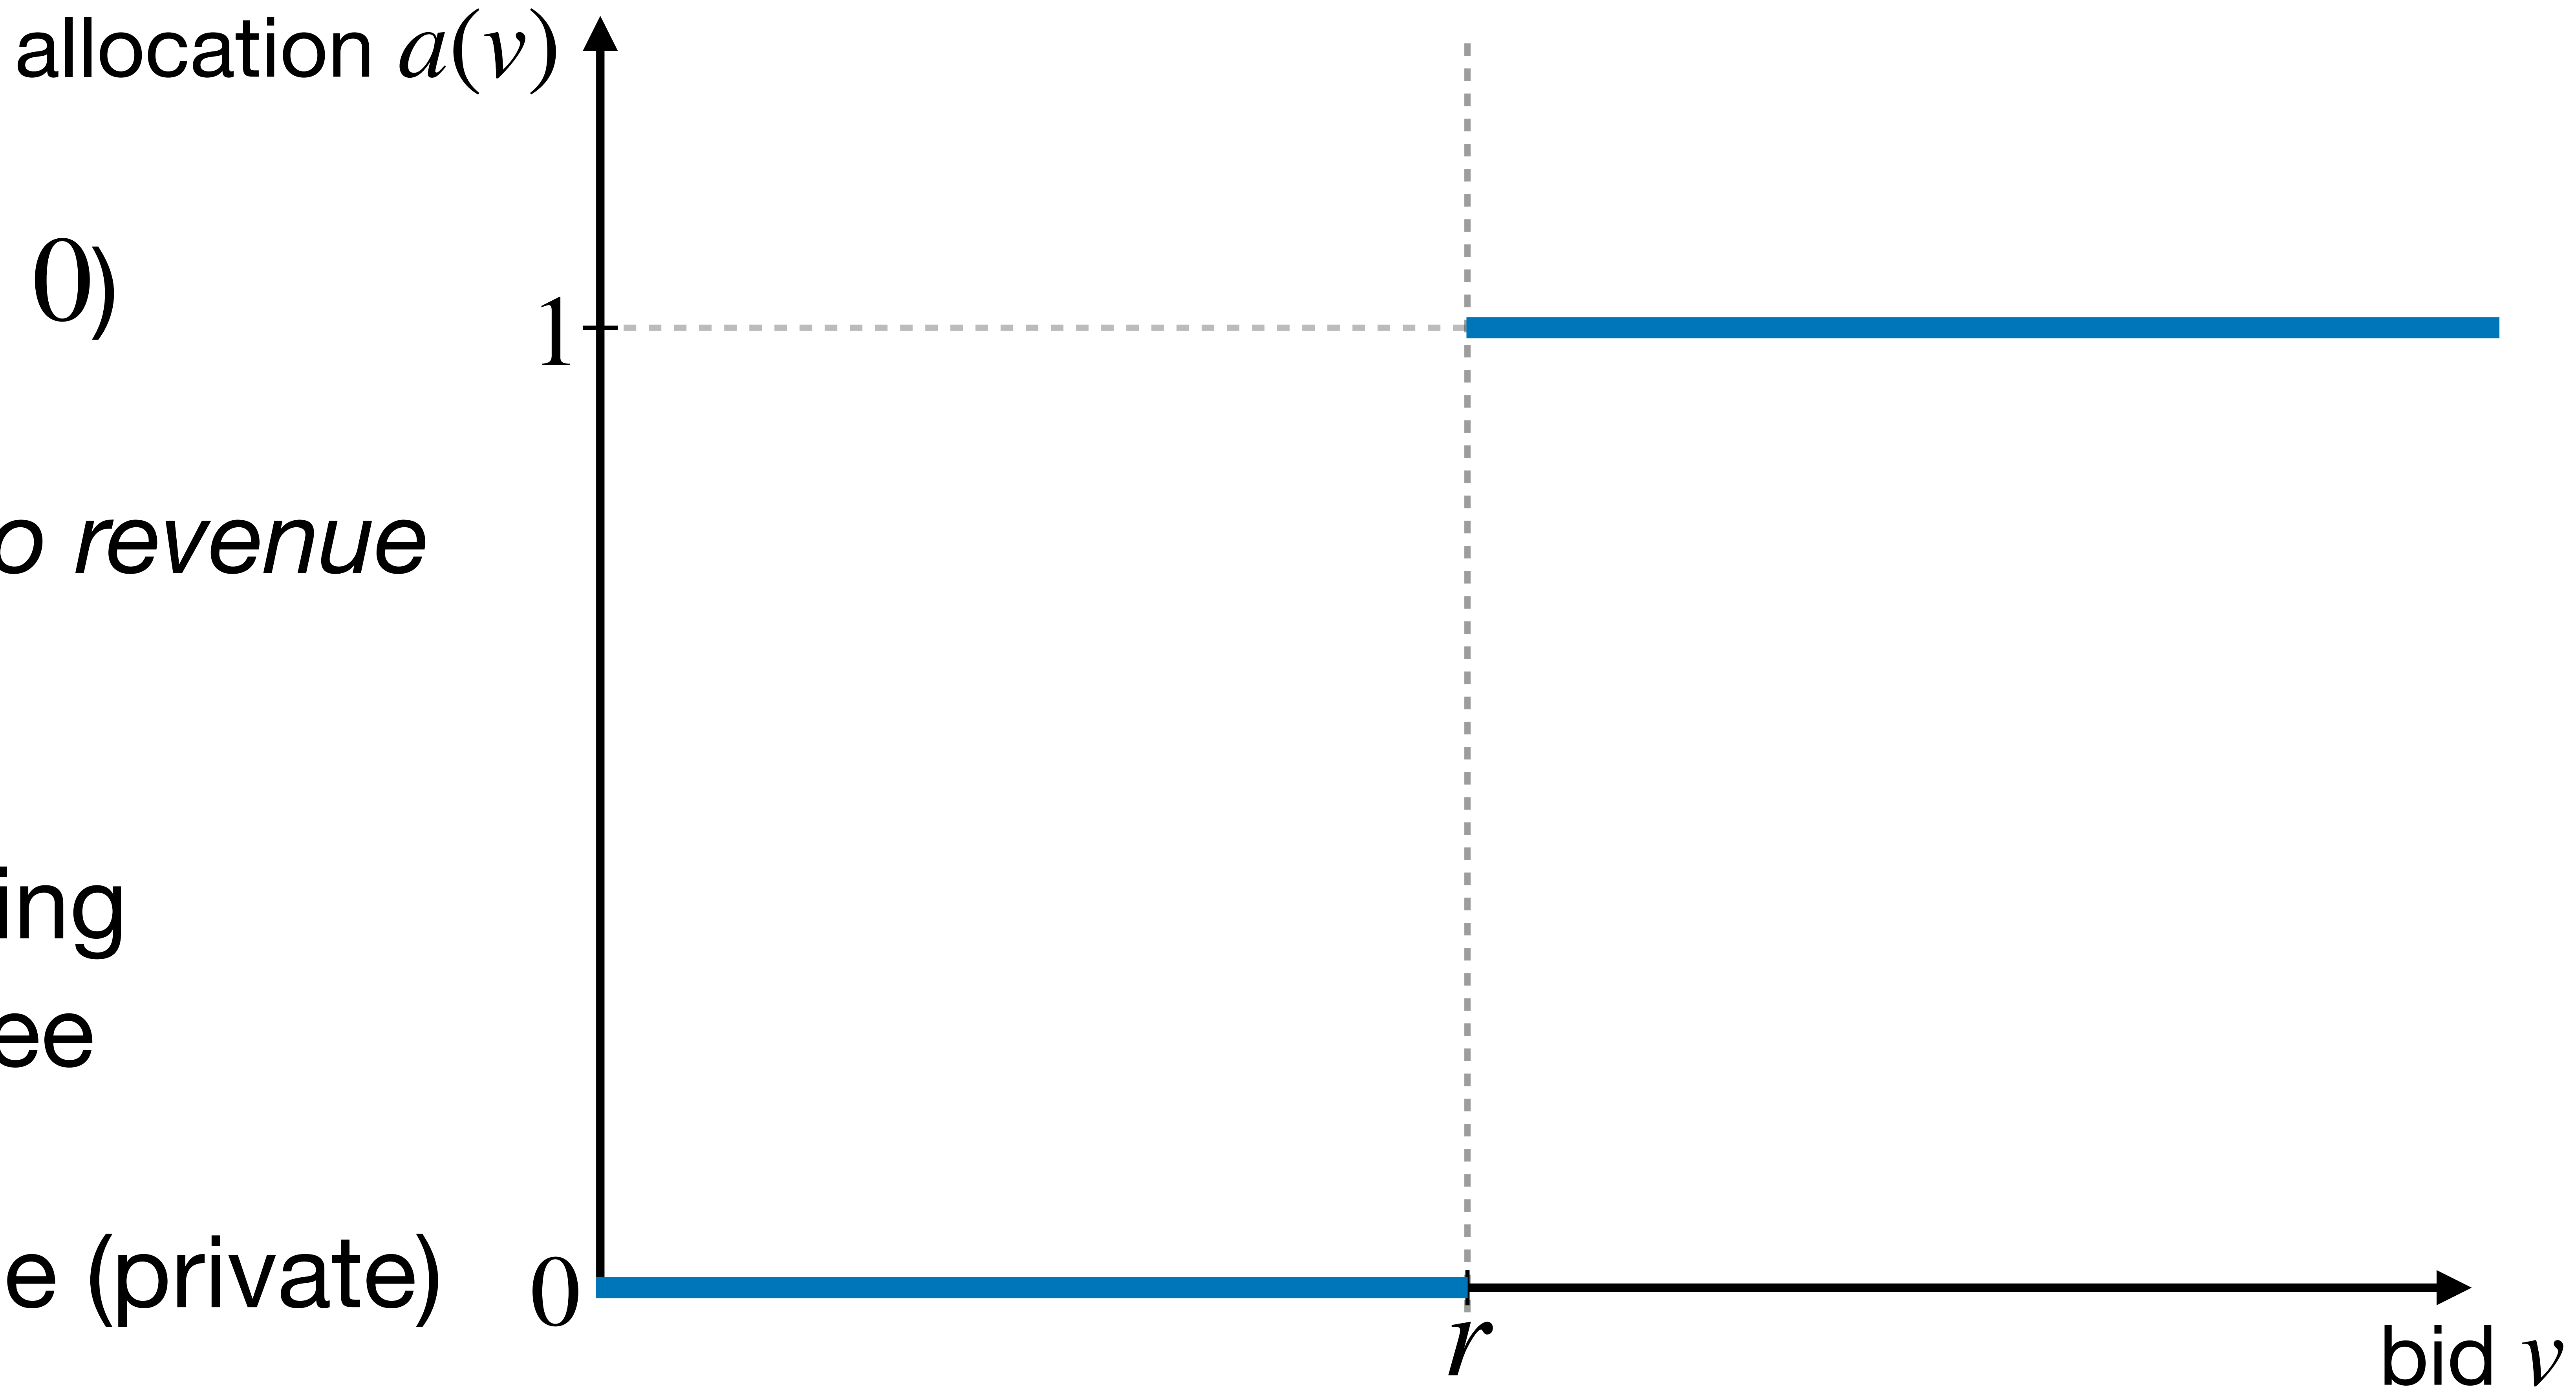

<u>Bayesian</u> (Single-Item) Auctions

Model: Bayesian Single-Item Auctions

- The seller/auction-designer has *incomplete knowledge* of the bidder values v_i:
  - $v_1, v_2, \ldots, v_n$ are *independent random variables*
  - drawn from **distributions** (“priors”) $F_1, F_2, \ldots, F_n$ supported over $[0,1]$

Continuous Probability Distributions

Quick Mathematical Detour

- Probability density function (pdf) $f : \mathbb{R}_+ \longrightarrow \mathbb{R}_+$
- Cumulative distribution function (cdf) $F : \mathbb{R}_+ \longrightarrow [0,1]$
  - The cdf is *continuous* and *nondecreasing*
  - It is given by

$$F(x) = \int_0^x f(t)\,\mathrm{d}t.$$

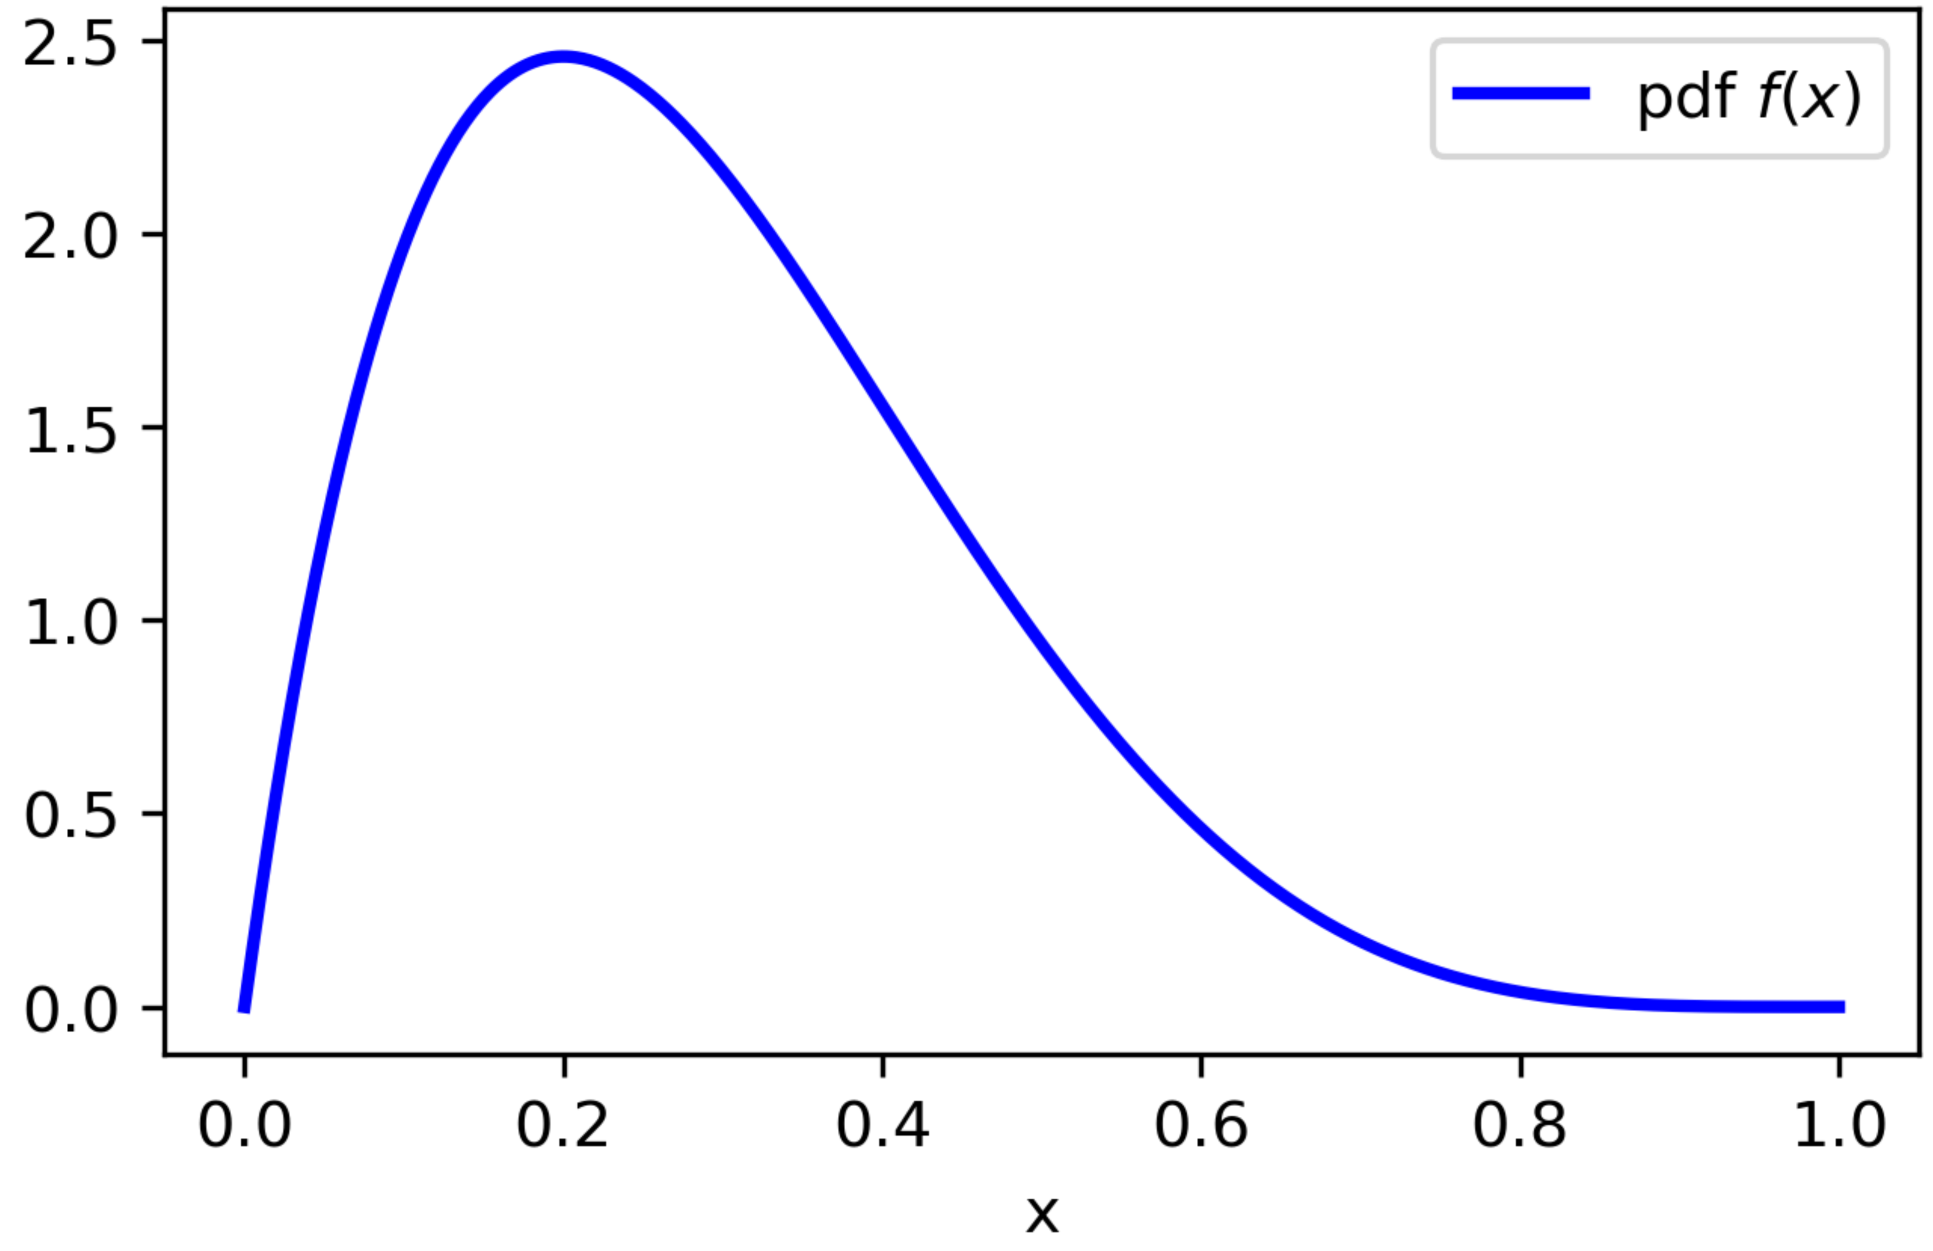

Continuous Probability Distributions

Quick Mathematical Detour

- Probability density function (pdf) $f : \mathbb{R}_+ \longrightarrow \mathbb{R}_+$
- Cumulative distribution function (cdf) $F : \mathbb{R}_+ \longrightarrow [0,1]$
  - The cdf is *continuous* and *nondecreasing*
  - It is given by

$$F(x) = \int_0^x f(t)\,\mathrm{d}t.$$

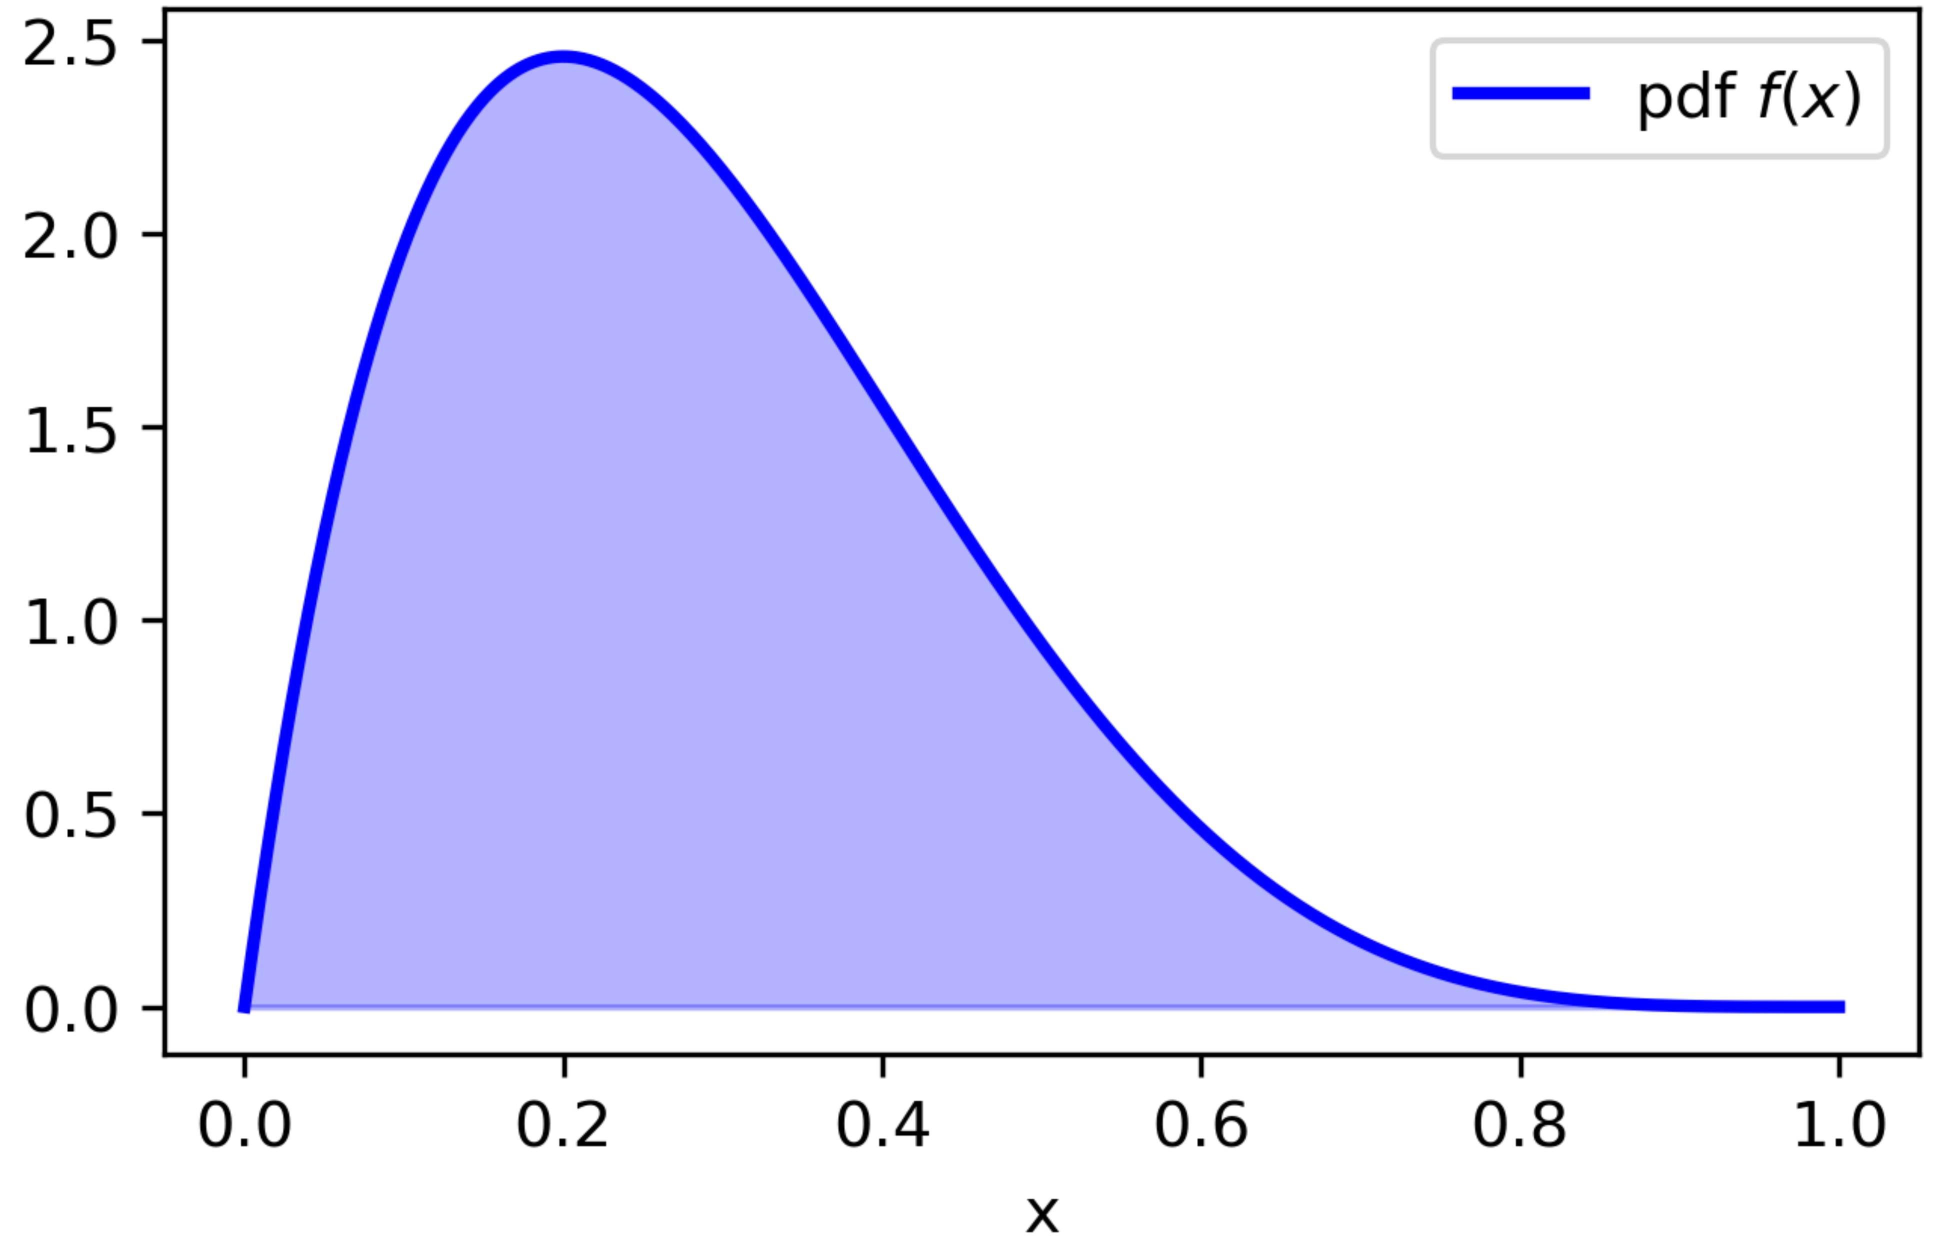

Continuous Probability Distributions

Quick Mathematical Detour

- Probability density function (pdf) $f : \mathbb{R}_+ \longrightarrow \mathbb{R}_+$
- Cumulative distribution function (cdf) $F : \mathbb{R}_+ \longrightarrow [0,1]$
  - The cdf is *continuous* and *nondecreasing*
  - It is given by

$$F(x) = \int_0^x f(t)\,\mathrm{d}t.$$

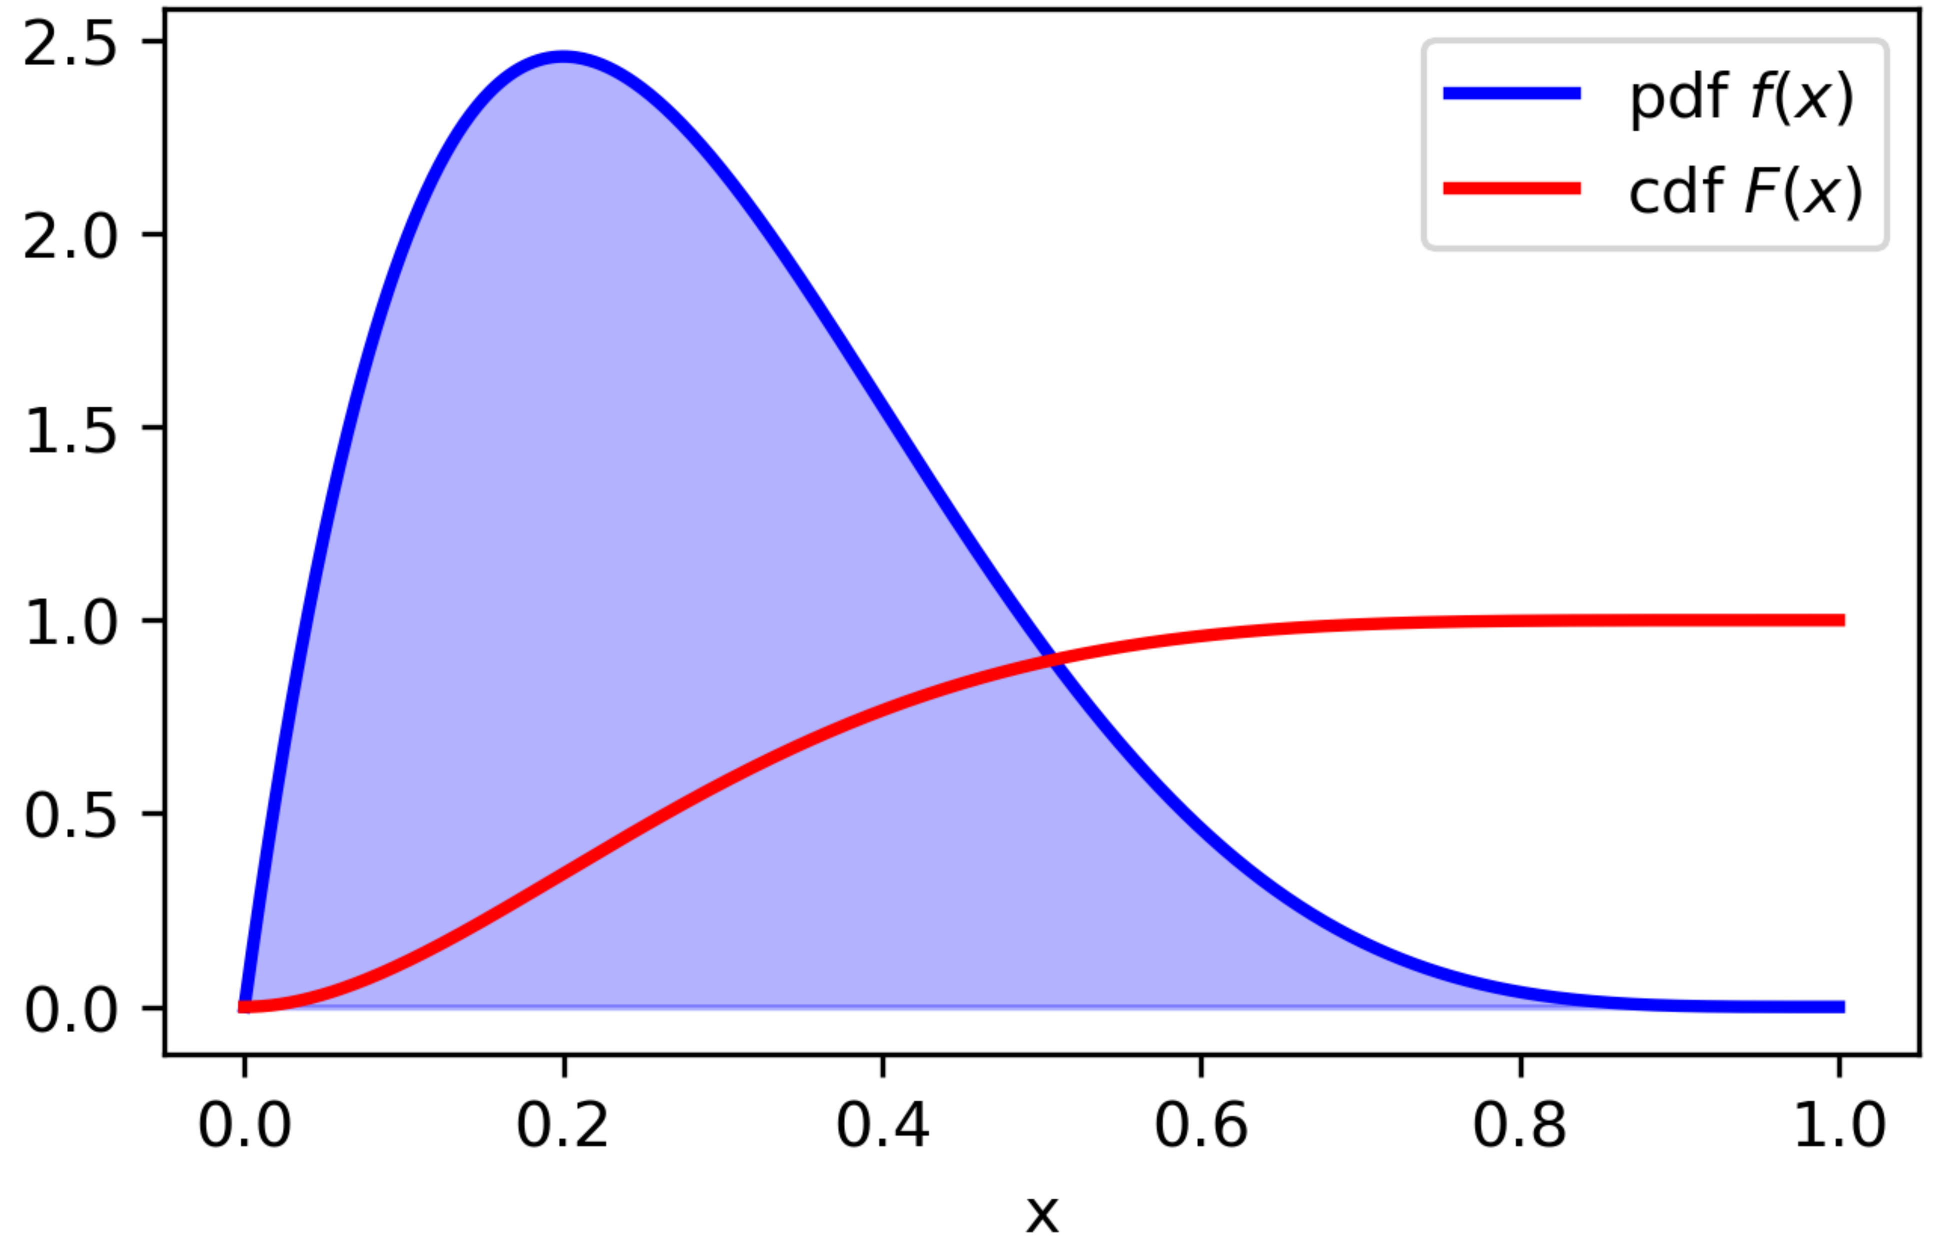

Continuous Probability Distributions

Quick Mathematical Detour

- Probability density function (pdf) $f : \mathbb{R}_+ \longrightarrow \mathbb{R}_+$
- Cumulative distribution function (cdf) $F : \mathbb{R}_+ \longrightarrow [0,1]$
  - The cdf is *continuous* and *nondecreasing*
  - It is given by

$X \sim F$

$$F(x) = \int_0^x f(t)\,\mathrm{d}t.$$

- If X is a random variable with distribution F, then:
  - $\mathrm{Prob}[a \leq X \leq b] = \int_a^b f(x)\,\mathrm{d}x = F(b) - F(a)$
  - Expectations: $\mathbb{E}[g(X)] = \int_0^\infty g(x) \cdot f(x)\,\mathrm{d}x$

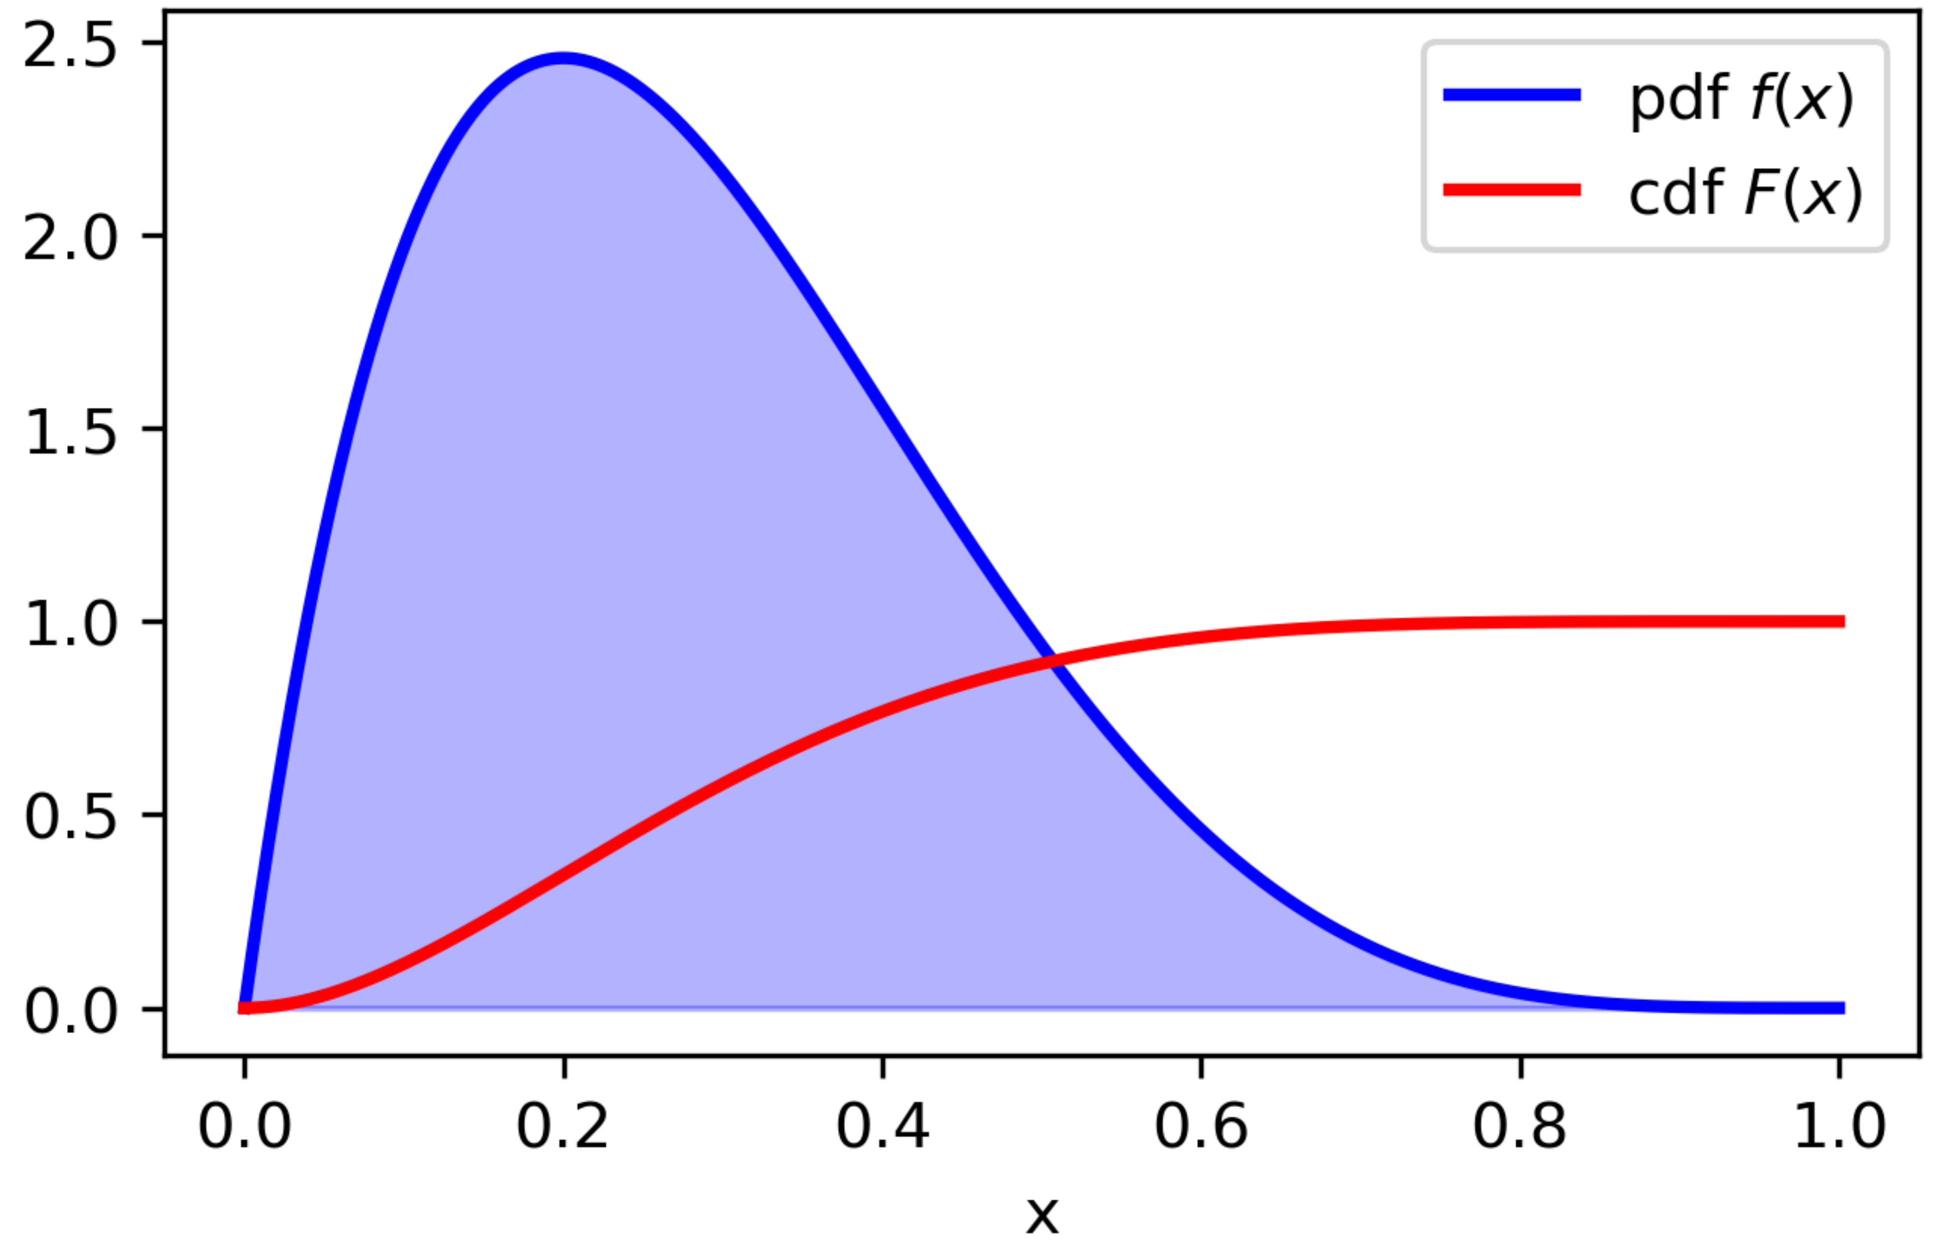

Continuous Probability Distributions

Quick Mathematical Detour

- Probability density function (pdf) $f : \mathbb{R}_+ \longrightarrow \mathbb{R}_+$
- Cumulative distribution function (cdf) $F : \mathbb{R}_+ \longrightarrow [0,1]$
  - The cdf is *continuous* and *nondecreasing*
  - It is given by

$$F(x) = \int_0^x f(t)\,\mathrm{d}t.$$

$X \sim F$

- If X is a random variable with distribution F, then:
  - $\mathrm{Prob}[a \le X \le b] = \int_a^b f(x)\,\mathrm{d}x = F(b) - F(a)$
  - Expectations: $\mathbb{E}[g(X)] = \int_0^\infty g(x) \cdot f(x)\,\mathrm{d}x$

Example

Uniform distribution on $[0,1]$:

$$f(x) = 1, \;\; F(x) = x$$

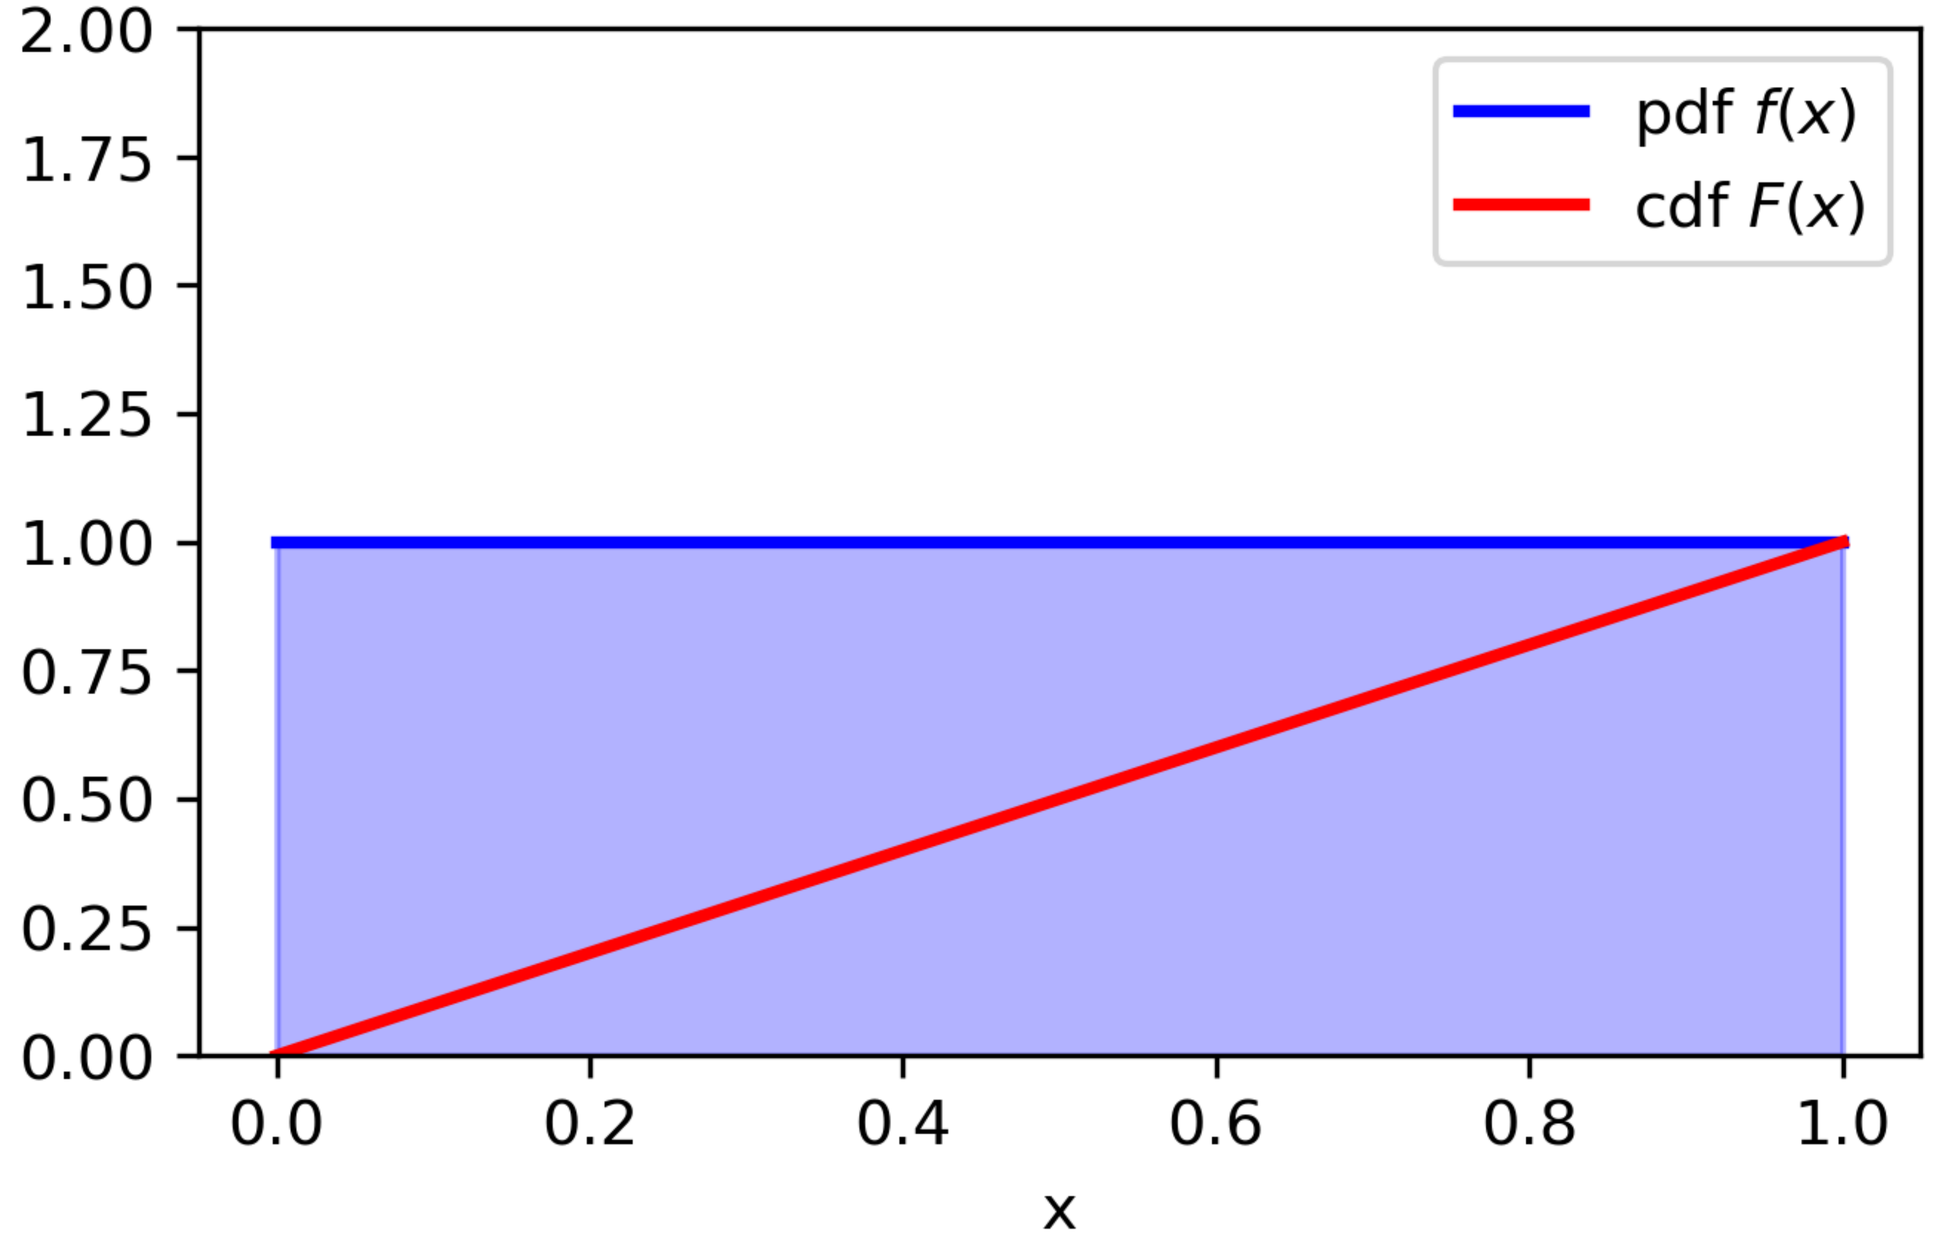

Model: Bayesian Single-Item Auctions

- The seller/auction-designer has *incomplete knowledge* of the bidder values v_i:
  - $v_1, v_2, \ldots, v_n$ are *independent random variables*
  - drawn from **distributions** ("priors") $F_1, F_2, \ldots, F_n$ supported over $[0,1]$
- **Optimization** objectives are defined **in expectation**:

$$W(\mathscr{M}) := \mathbb{E}_{\boldsymbol{v}\sim\boldsymbol{F}}\left[\sum_{i=1}^{n} a_i(\boldsymbol{v})v_i\right]$$

Welfare

$$\boxed{R(\mathscr{M}) := \mathbb{E}_{\boldsymbol{v}\sim\boldsymbol{F}}\left[\sum_{i=1}^{n} p_i(\boldsymbol{v})\right]}$$

Revenue

- Goal: find a ***revenue-maximizing*** *truthful* mechanism

$$\max_{\text{truthful } \mathscr{M}} R(\mathscr{M}) = \max_{\text{monotone } \boldsymbol{a}} \mathbb{E}\left[\sum_{i=1}^{n}\left(a_i(\boldsymbol{v})v_i - \int_0^{v_i} a_i(t, \boldsymbol{v}_{-i})\,\mathrm{d}t\right)\right]$$

"Critical" value: Myerson's lemma

Revenue Maximization

Example: *Single*-Bidder, *Deterministic* Auctions

- Bidder wins if and only if $v \geq r$
  - And then pays r
- Expected revenue:
  $r \cdot \text{Prob}[v \geq r] = r(1 - F(r))$
- Optimal **"monopoly reserve"**:
  $r^* = \arg\max_{r \in [0,1]} r(1 - F(r))$
- Example: Uniform distribution
  - $r(1 - F(r)) = r(1 - r)$
  - $r^* = 1/2$; optimal revenue $1/4$

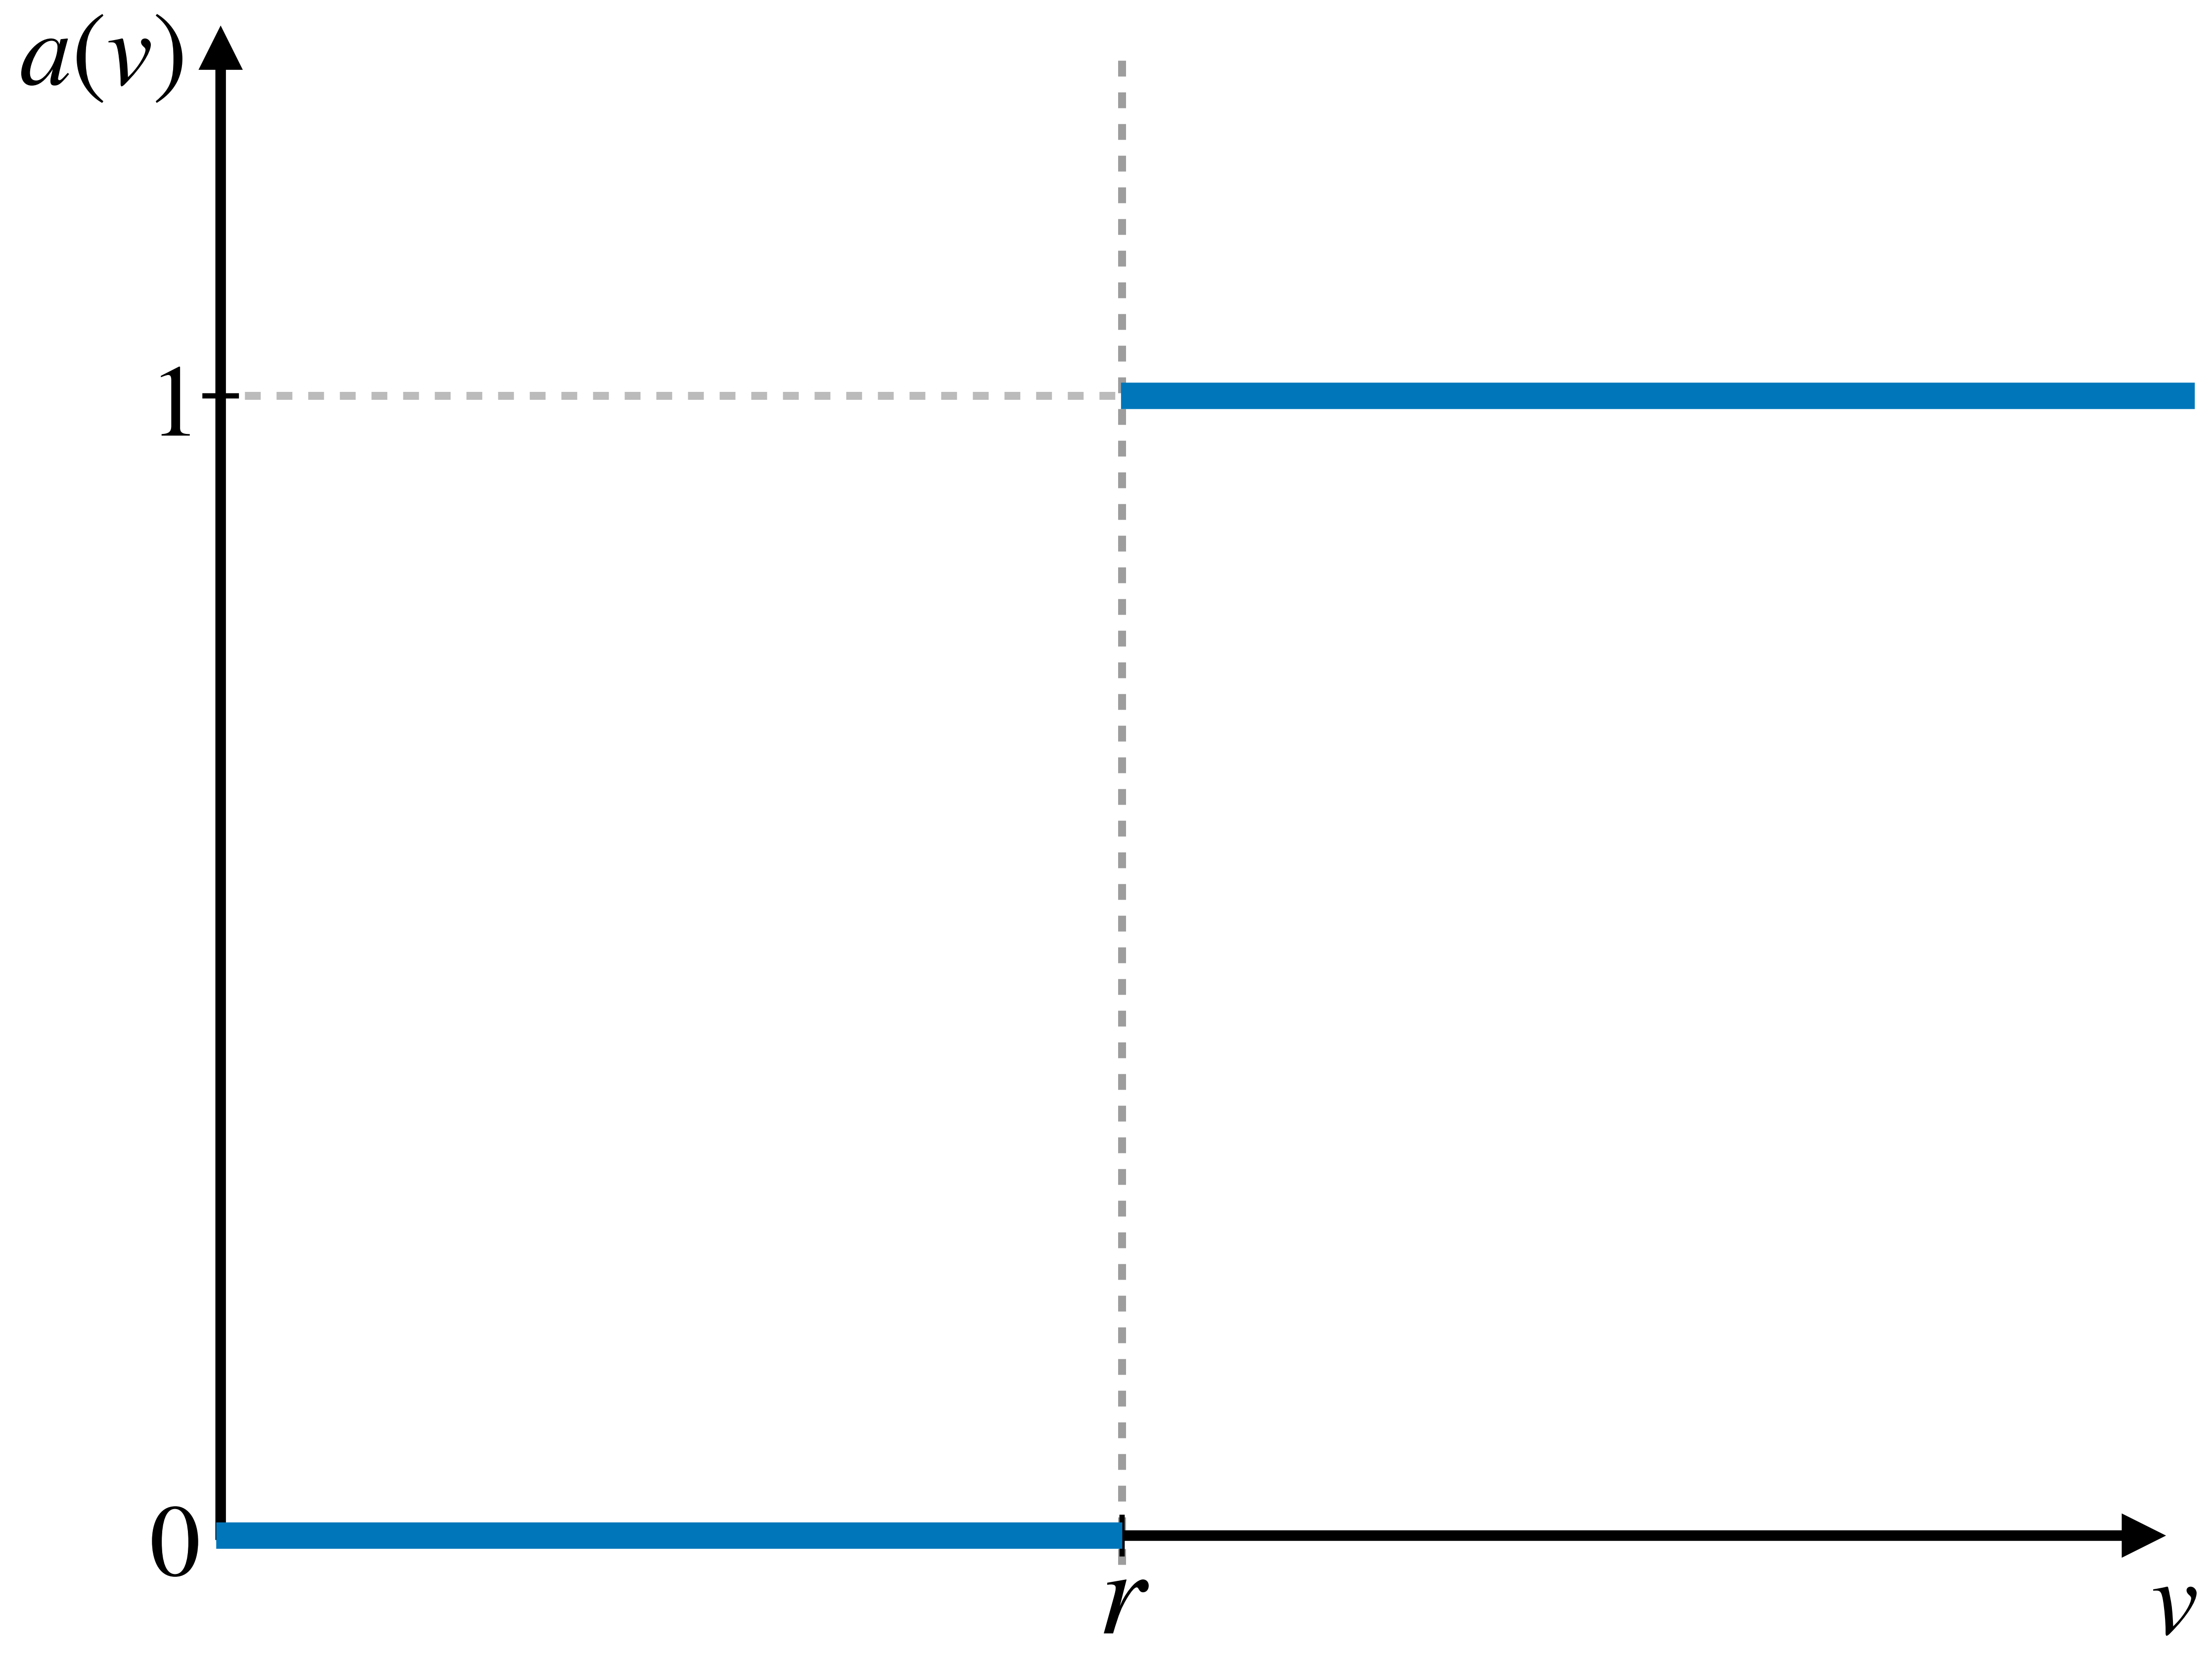

Virtual Values

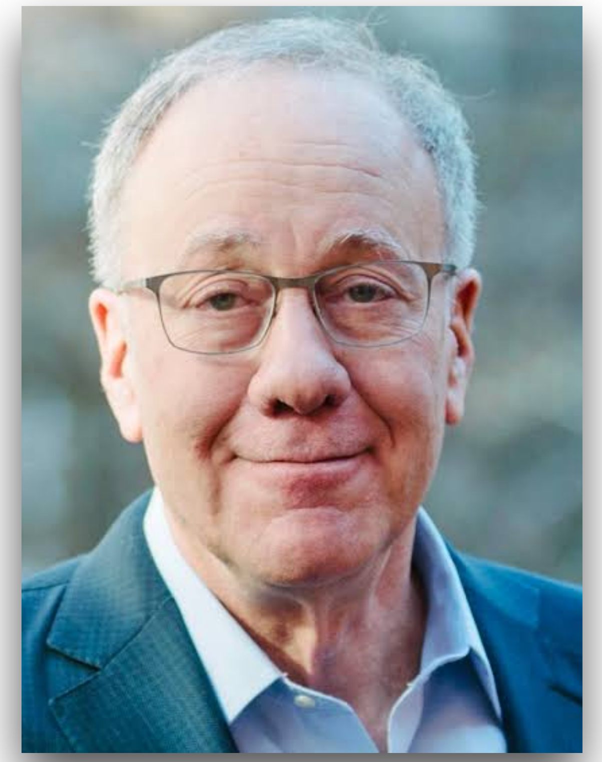

Roger Myerson

- The **virtual value** function of bidder i is defined by

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

  - If ϕ_i is *(strictly) increasing*, we will call distribution F_i **(strictly) regular**.

THEOREM (R. Myerson [1981])

For any *truthful* mechanism $(\boldsymbol{a}, \boldsymbol{p})$ it holds that

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

revenue

"virtual" social welfare

Regular Distributions
Examples

- Most “common” probability distributions in economic theory are strictly regular.
- Intuition: no “thick tails”
- Examples: uniform, (truncated) normal, exponential, …
- Virtual value function for uniform distribution
  - $$\phi(x) = x - \frac{1 - F(x)}{f(x)} = x - \frac{1 - x}{1} = 2x - 1$$

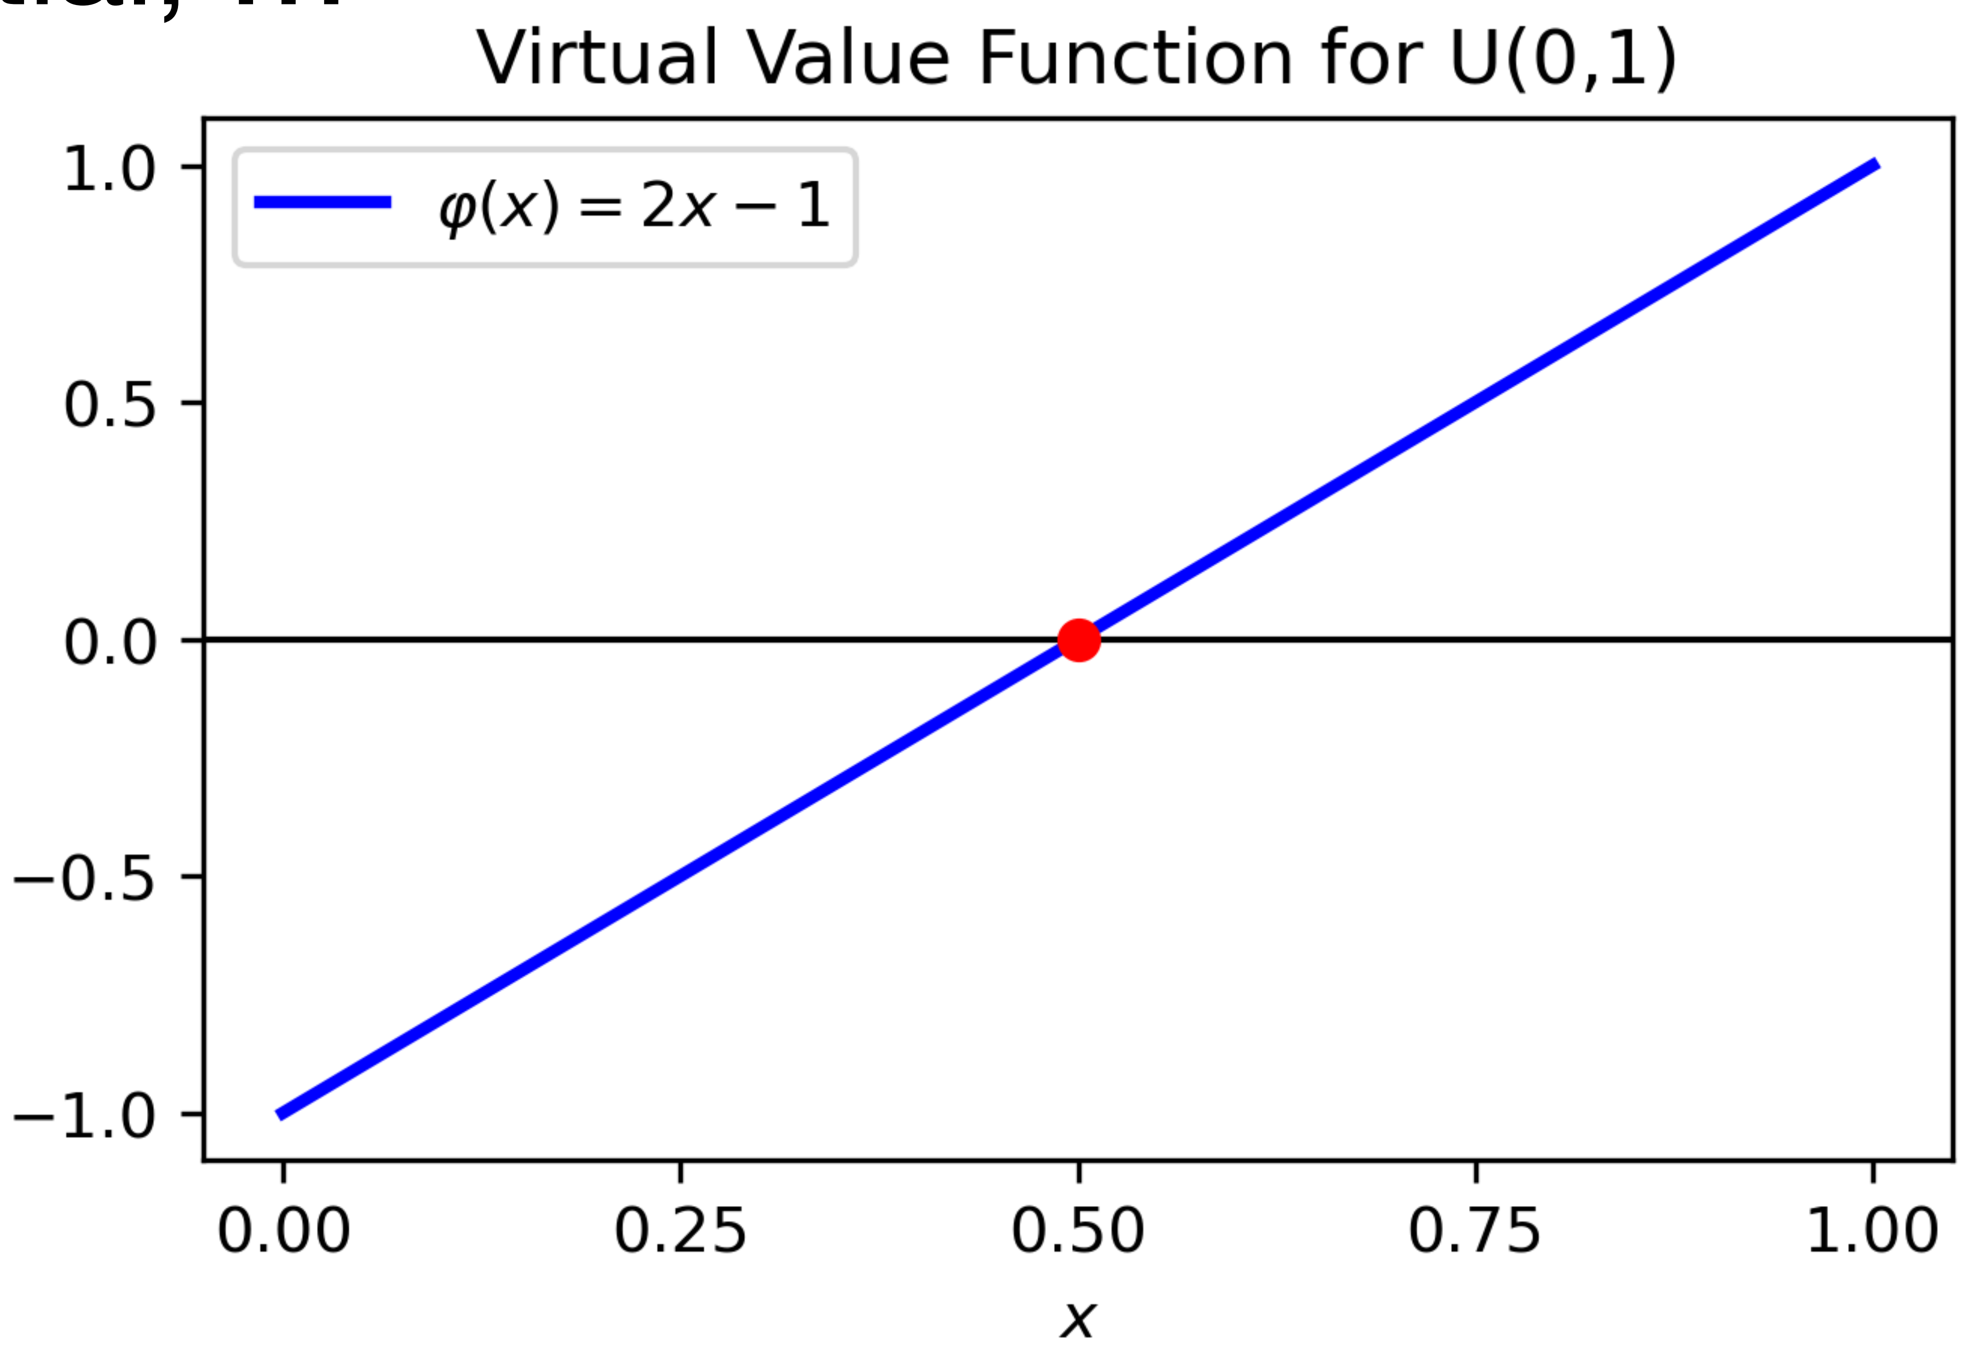

*ties occur with *zero* probability

Myerson's Optimal Auction

For any *truthful* mechanism $\mathscr{M} = (\boldsymbol{a}, \boldsymbol{p})$:

$$R(\mathscr{M}) = \mathbb{E}_{\boldsymbol{v}\sim \boldsymbol{F}}\left[\sum_{i=1}^{n} a_i(\boldsymbol{v})\phi_i(v_i)\right]$$

Optimal revenue $= \mathbb{E}_{\boldsymbol{v}\sim \boldsymbol{F}}\left[\max_{i\in[n]} \phi_i(v_i)^+\right]$

Notation: $z^+ = \max\{z, 0\}$

<u>Assuming *strictly regular* priors</u>:

- Selling the item to the bidder with the **highest* nonnegative *virtual* value**
  1. is monotone (i.e., **truthful**)
  2. maximizes the seller's revenue (i.e., is **optimal**)

Equal to monopoly reserve $r^* = \arg\max_{x\in[0,1]} x(1 - F(x))$

- If, furthermore, the priors are <u>identical</u> (aka "**iid**"), then the optimal auction is simply a **second-price auction with reserve price**

$$r := \phi^{-1}(0)$$

Proof of Myerson's Theorem

For any *truthful* mechanism $(\boldsymbol{a}, \boldsymbol{p})$: $\mathbb{E}_{\boldsymbol{v}\sim\boldsymbol{F}}\left[\sum_{i=1}^{n} p_i(\boldsymbol{v})\right] = \mathbb{E}_{\boldsymbol{v}\sim\boldsymbol{F}}\left[\sum_{i=1}^{n} a_i(\boldsymbol{v})\phi_i(v_i)\right].$

PROOF

- It is enough to prove it "per-bidder"
  - i.e., we will show that, for *any* bidder i and *all* bid/value profiles $\boldsymbol{v}_{-i}$ of the others:
  $$\mathbb{E}_{v_i\sim F_i}\left[p_i(\boldsymbol{v})\right] = \mathbb{E}_{v_i\sim F_i}\left[a_i(\boldsymbol{v})\phi_i(v_i)\right].$$
  - This is enough, due to the "linearity of expectation".
- So, from now on let's fix a bidder $i \in [n]$ and values $\boldsymbol{v}_{-i} \in [0,1]^{n-1}$.
- Simplify notation: $v := v_i$, $a(v) := a_i(v, \boldsymbol{v}_{-i})$, $p(v) := p_i(v, \boldsymbol{v}_{-i})$, $F(v) := F_i(v)$, ...

$$\mathbb{E}_{v_i \sim F_i}\left[p_i(\boldsymbol{v})\right] = \mathbb{E}_{v_i \sim F_i}\left[a_i(\boldsymbol{v})\phi_i(v_i)\right]$$

Proof of Myerson's Theorem (cont'd)

$$\mathbb{E}_{v\sim F}[p(v)] = \int_0^1 p(v)\cdot f(v)\,\mathrm{d}v$$

$$= \int_0^1 \left(v \cdot a(v) - \int_0^v a(z)\,\mathrm{d}z \right) \cdot f(v)\,\mathrm{d}v$$

$$\mathbb{E}_{v_i \sim F_i}\left[p_i(\boldsymbol{v})\right] = \mathbb{E}_{v_i \sim F_i}\left[a_i(\boldsymbol{v})\phi_i(v_i)\right]$$

Proof of Myerson's Theorem (cont'd)

$$\mathbb{E}_{v \sim F}[p(v)] = \int_0^1 p(v) \cdot f(v)\, \mathrm{d}v$$

$$= \int_0^1 \left(v \cdot a(v) - \int_0^v a(z)\, \mathrm{d}z \right) \cdot f(v)\, \mathrm{d}v$$

due to Myerson's lemma

$$\mathbb{E}_{v_i \sim F_i}\left[p_i(\boldsymbol{v})\right] = \mathbb{E}_{v_i \sim F_i}\left[a_i(\boldsymbol{v})\phi_i(v_i)\right]$$

Proof of Myerson's Theorem (cont'd)

$$\mathbb{E}_{v\sim F}[p(v)] = \int_0^1 p(v)\cdot f(v)\,\mathrm{d}v$$

$$= \int_0^1 \left(v\cdot a(v) - \int_0^v a(z)\,\mathrm{d}z \right)\cdot f(v)\,\mathrm{d}v$$

due to Myerson's lemma

$$= \int_0^1 v\cdot a(v)f(v)\,\mathrm{d}v - \int_0^1\int_0^v a(z)f(v)\,\mathrm{d}z\,\mathrm{d}v$$

$$= \int_0^1 v\cdot a(v)f(v)\,\mathrm{d}v - \int_0^1\int_z^1 a(z)f(v)\,\mathrm{d}v\,\mathrm{d}z$$

exchanging the order of integration

$$\mathbb{E}_{v_i \sim F_i}\left[p_i(\mathbf{v})\right] = \mathbb{E}_{v_i \sim F_i}\left[a_i(\mathbf{v})\phi_i(v_i)\right]$$

Proof of Myerson's Theorem (cont'd)

$$\mathbb{E}_{v \sim F}[p(v)] = \int_0^1 p(v) \cdot f(v)\,\mathrm{d}v$$

$$= \int_0^1 \left(v \cdot a(v) - \int_0^v a(z)\,\mathrm{d}z \right) \cdot f(v)\,\mathrm{d}v$$

due to Myerson's lemma

$$= \int_0^1 v \cdot a(v) f(v)\,\mathrm{d}v - \int_0^1 \int_0^v a(z) f(v)\,\mathrm{d}z\,\mathrm{d}v$$

$$= \int_0^1 v \cdot a(v) f(v)\,\mathrm{d}v - \int_0^1 \int_z^1 a(z) f(v)\,\mathrm{d}v\,\mathrm{d}z$$

exchanging the order of integration

Interchanging Sums & Integrals

Quick Mathematical Detour

$$\sum_{i=1}^{n} \sum_{j=1}^{i} a_{i,j}$$

$$\begin{bmatrix} a_{1,1} & a_{1,2} & a_{1,3} & a_{1,4} \\ a_{2,1} & a_{2,2} & a_{2,3} & a_{2,4} \\ a_{3,1} & a_{3,2} & a_{3,3} & a_{3,4} \\ a_{4,1} & a_{4,2} & a_{4,3} & a_{4,4} \end{bmatrix}$$

Interchanging Sums & Integrals

Quick Mathematical Detour

$$\sum_{i=1}^{n} \sum_{j=1}^{i} a_{i,j} = \sum_{j=1}^{n} \sum_{i=j}^{n} a_{i,j}$$

$$\begin{bmatrix} a_{1,1} & a_{1,2} & a_{1,3} & a_{1,4} \\ a_{2,1} & a_{2,2} & a_{2,3} & a_{2,4} \\ a_{3,1} & a_{3,2} & a_{3,3} & a_{3,4} \\ a_{4,1} & a_{4,2} & a_{4,3} & a_{4,4} \end{bmatrix}$$

$$\mathbb{E}_{v_i \sim F_i}\left[p_i(\boldsymbol{v})\right] = \mathbb{E}_{v_i \sim F_i}\left[a_i(\boldsymbol{v})\phi_i(v_i)\right]$$

Proof of Myerson's Theorem (cont'd)

$$\begin{aligned}
\mathbb{E}_{v \sim F}[p(v)] &= \int_0^1 v \cdot a(v) f(v)\,\mathrm{d}v - \int_0^1 \int_z^1 a(z) f(v)\,\mathrm{d}v\,\mathrm{d}z \\
&= \int_0^1 v \cdot a(v) f(v)\,\mathrm{d}v - \int_0^1 a(z) \int_z^1 f(v)\,\mathrm{d}v\,\mathrm{d}z \\
&= \int_0^1 v \cdot a(v) f(v)\,\mathrm{d}v - \int_0^1 a(z)\big(1 - F(z)\big)\,\mathrm{d}z \\
&= \int_0^1 v \cdot a(v) f(v) - a(v)\big(1 - F(v)\big)\,\mathrm{d}v
\end{aligned}$$

$$\mathbb{E}_{v_i \sim F_i}\left[p_i(\boldsymbol{v})\right] = \mathbb{E}_{v_i \sim F_i}\left[a_i(\boldsymbol{v})\phi_i(v_i)\right]$$

Proof of Myerson's Theorem (cont'd)

$$\begin{aligned}
\mathbb{E}_{v \sim F}[p(v)] &= \int_0^1 v \cdot a(v) f(v) - a(v)\big(1 - F(v)\big)\, \mathrm{d}v \\
&= \int_0^1 a(v)\left[v - \frac{1 - F(v)}{f(v)}\right] f(v)\, \mathrm{d}v \\
&= \int_0^1 a(v)\phi(v) \cdot f(v)\, \mathrm{d}v \\
&= \mathbb{E}_{v \sim F}[a(v)\phi(v)]\,.
\end{aligned}$$

“Simplicity” vs Optimality

Non-Identical Priors

- The optimal auction might be “complicated”, or even practically infeasible.
  - Higher virtual value does not always correspond to higher value!

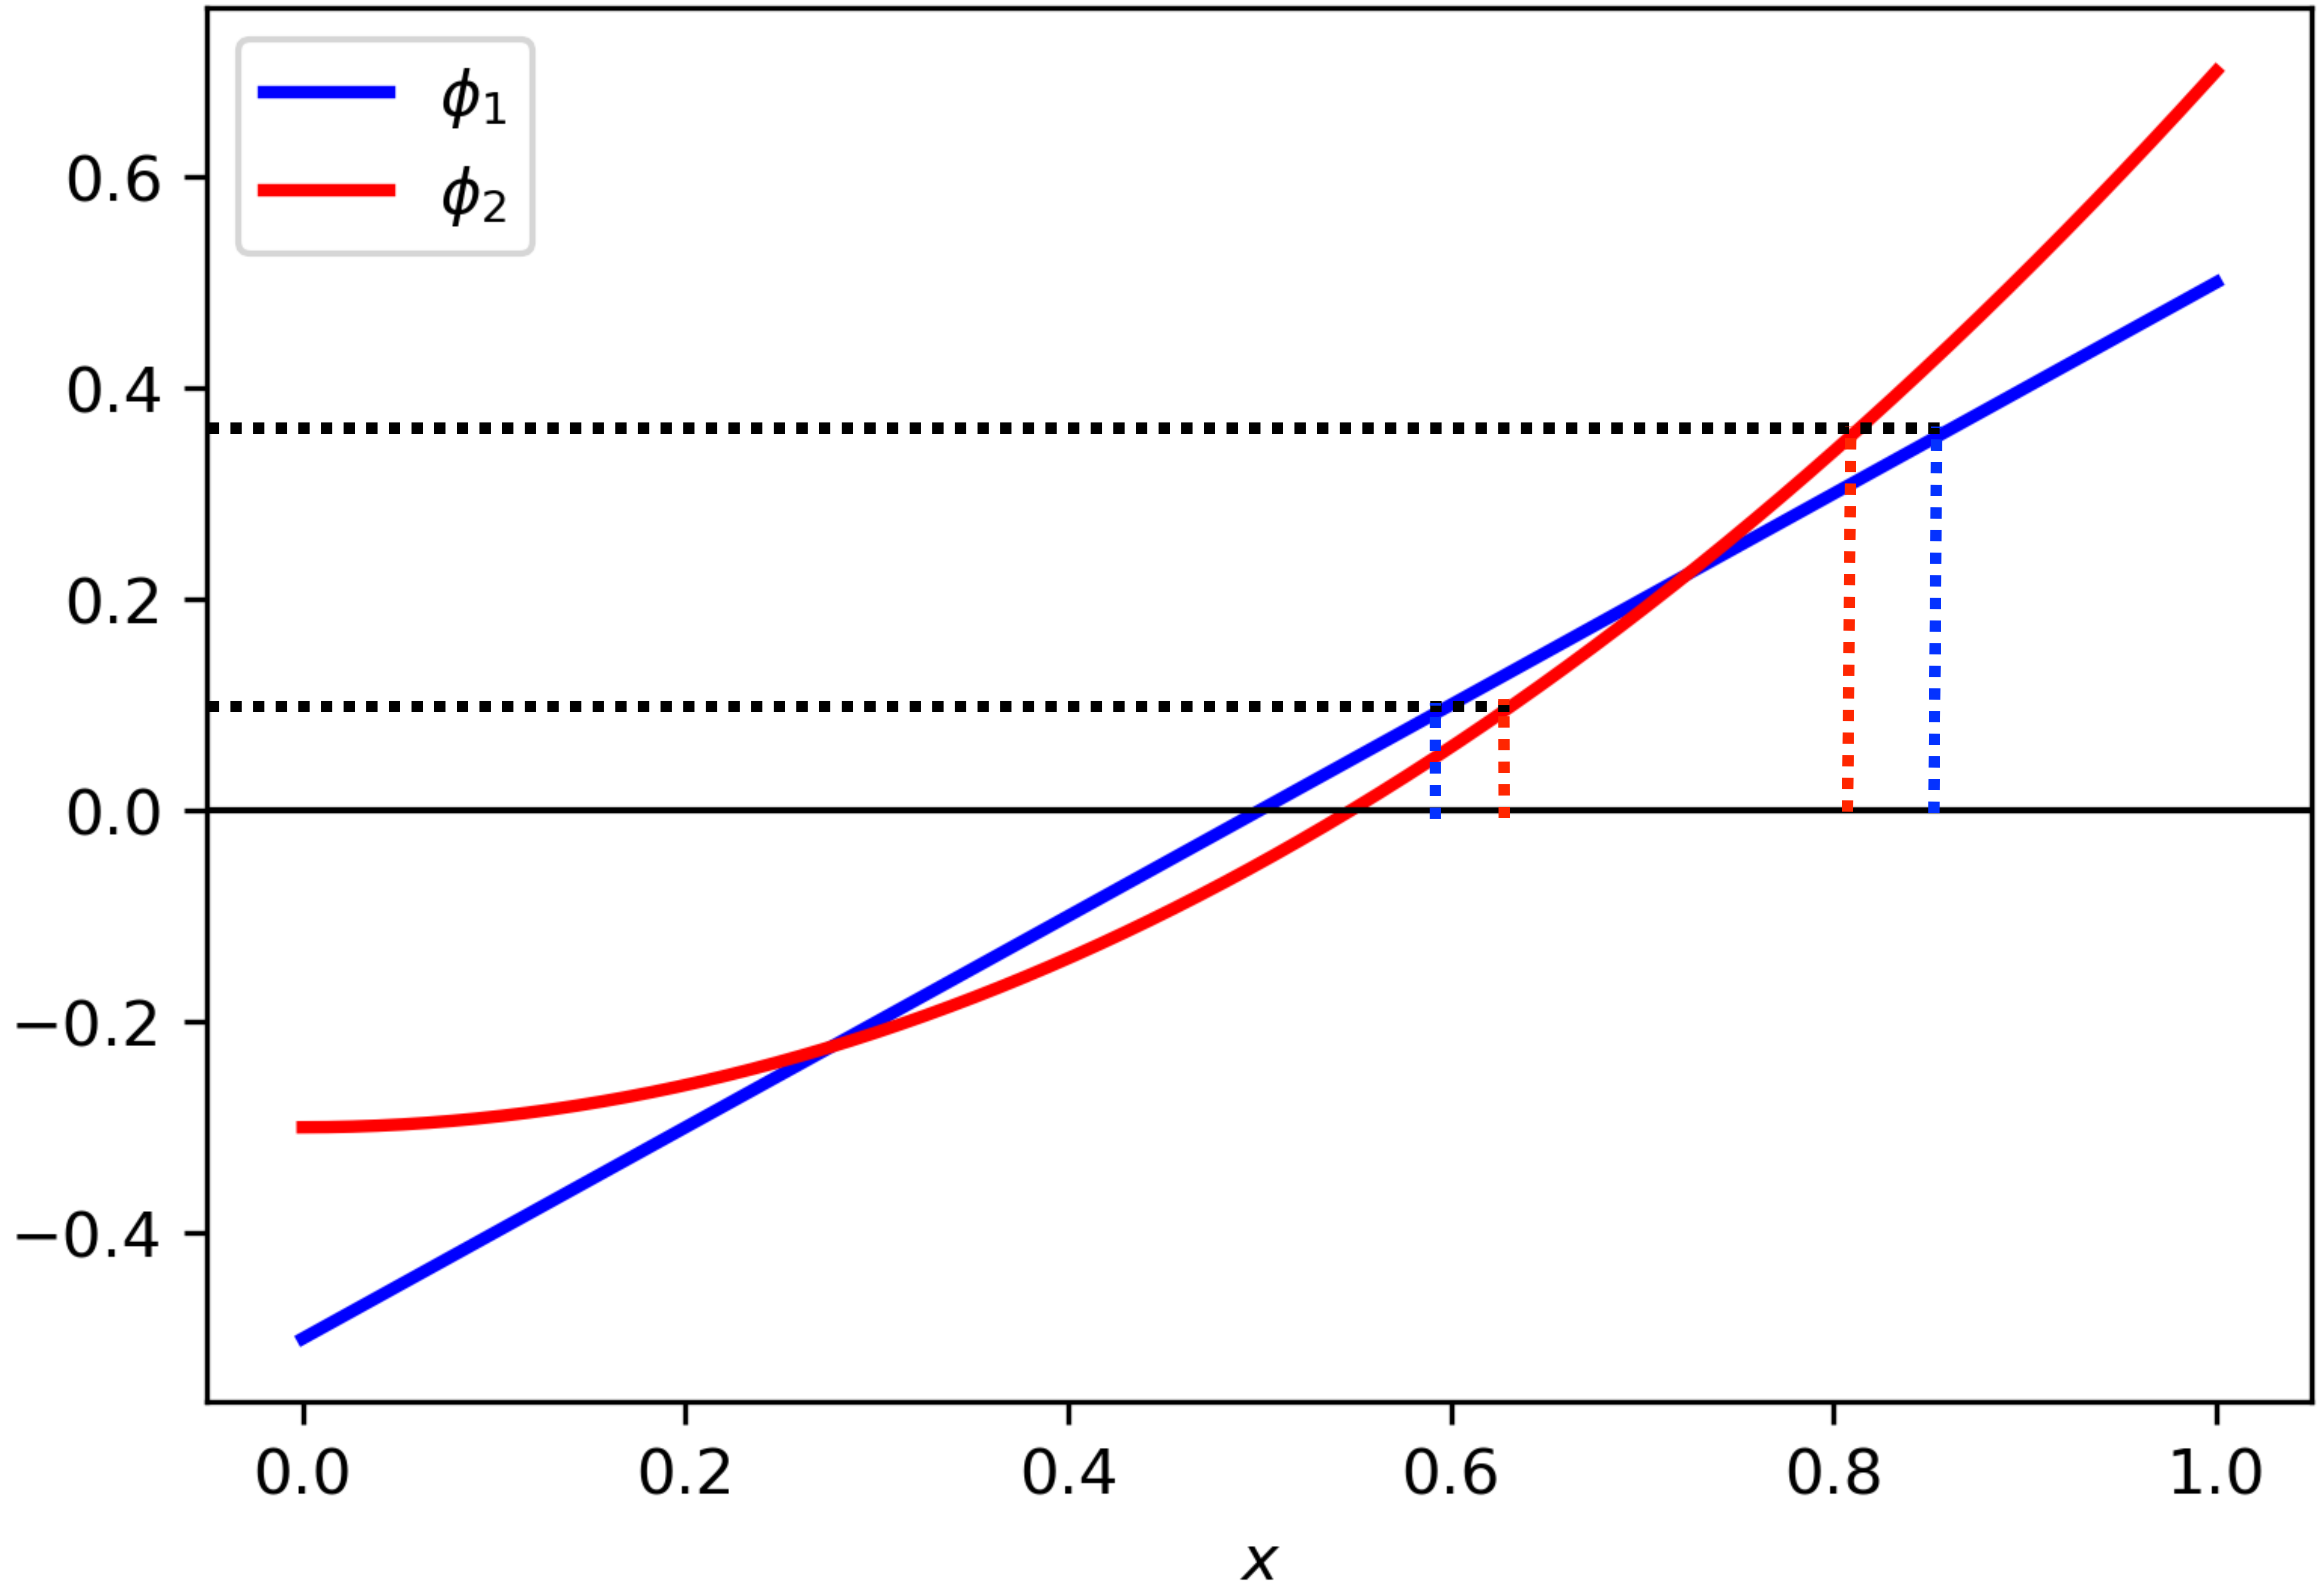

Non-Identical Priors

- The optimal auction might be “complicated”, or even practically infeasible
  - Higher virtual value does not always correspond to higher value!

- How much revenue do we loose by restricting to **simple** auction formats? For example:
  - Second-price auction; with reserves or not
  - Posted pricing (“take-it-or-leave-it”)

The Bullow-Klemperer Approximation

Identical Bidders

THEOREM (J. Bulow & P. Klemperer [1996])

For regular iid priors, the expected revenue of the second-price auction (with no reserve) with $n+1$ bidders is at least the expected revenue of the optimal auction with n bidders.

COROLLARY

For n bidders with regular iid priors, the second-price auction achieves at least a $\frac{n-1}{n}$-fraction of the optimal expected revenue.

Second-Price with Reserves & Pricing

Optimal Revenue Approximation for Non-Identical Regular Bidders

- Selling the item by posting the *same* **price** to all bidders, has an approximation ratio of (exactly)

$$2.6202$$

- The approximation ratio of the best **second-price** auction with the same (aka **"anonymous"**) **reserve** price for all bidders lies in

$$[2.1596, 2.6202]$$

- Selling via a second-price auction with **bidder-specific reserves** achieves an approximation ratio of (exactly)

$$2$$

A Small Glimpse Beyond: Multi-item Auctions

Multi-Dimensional Revenue Maximization

Complications

- Fundamental technical obstacles, even for a *single* bidder!
- Randomization is required, in general, for optimality.
- Computational hardness barriers.
- Large constant approximations only (e.g., 8).
- Generally: the exact structure, and key computational properties, of the optimal auctions **still elude us**!
    - Resolved only for a single-bidder, small number of items, and very specific distributions (most notably, uniform).