# Part I

# Foundations

# Part II

# Bargaining

# Part III

# Auctions and Mechanism Design

# Chapter 4

# Private Value Auctions

**Definition 4.1: Private-Value Auction Setting**

A **private-value auction** consists of $n$ risk-neutral bidders competing for a single indivisible object. Each bidder $i$ independently draws a valuation $X_i \sim F$, where $F$ is a continuous cumulative distribution function on $[0, \bar{x}]$ that is common knowledge. The realization $x_i$ is bidder $i$'s *private information*—no other bidder, nor the seller, observes it. Each bidder submits a bid $b_i \in \mathbb{R}_{\geq 0}$, and the auctioneer determines the winner and the payment from the bid profile $(b_1, \ldots, b_n)$ according to the auction's rules.

## 4.1 Second Price Sealed Bid Auction (SPA)

**Definition 4.2: Second-Price Sealed-Bid Auction**

Each bidder simultaneously submits a sealed bid. The bidder with the highest bid wins the object and pays the *second*-highest bid; ties are broken uniformly at random.

**Claim**

It is a weakly dominant strategy for each bidder to bid their true value in a SPA, i.e., $b_i = x_i$.

***Proof for Claim.***

Graphical proof:

Suppose bidder 1 has value $x_1$. We show that bidding $b_1 = x_1$ is a weakly dominant strategy for bidder 1. Let $x_{-1}$ denote the highest competing bid. Bidder 1's profit is

$$\pi_1(b_1, x_{-1}) = \begin{cases} x_1 - x_{-1} & \text{if } b_1 > x_{-1}, \\ 0 & \text{if } b_1 < x_{-1}. \end{cases}$$

This corresponds to the **black** line in the figure below.

If instead bidder 1 bids $b_1^+ > x_1$, the profit becomes

$$\pi_1(b_1^+, x_{-1}) = \begin{cases} x_1 - x_{-1} & \text{if } b_1^+ > x_{-1}, \\ 0 & \text{if } b_1^+ < x_{-1}. \end{cases}$$

This corresponds to the **red** line in the figure below.

If bidder 1 bids $b_1^- < x_1$, the profit becomes

$$\pi_1(b_1^-, x_{-1}) = \begin{cases} x_1 - x_{-1} & \text{if } b_1^- > x_{-1}, \\ 0 & \text{if } b_1^- < x_{-1}. \end{cases}$$

This corresponds to the **blue** line in the figure below.

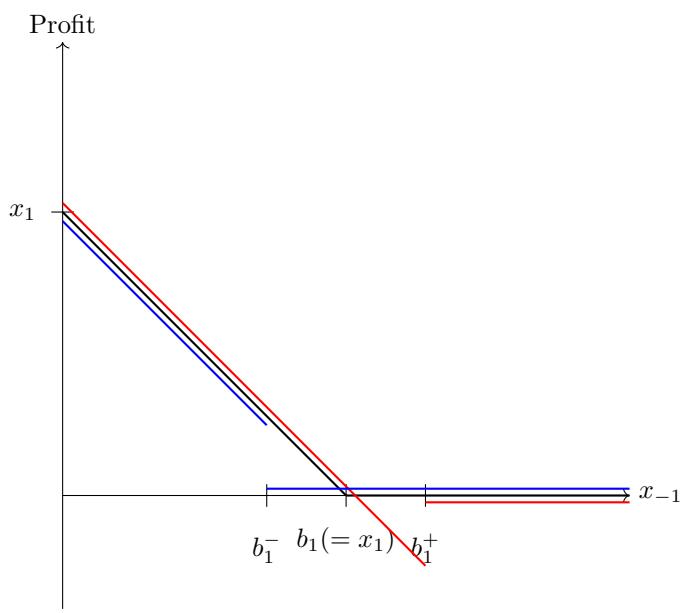


The figure shows that the black profit profile (truthful bid $b_1 = x_1$) weakly dominates both the red one ($b_1^+$, which can yield negative payoffs) and the blue one ($b_1^-$, which forfeits positive surplus when $b_1^- < x_{-1} < x_1$) at every realization of $x_{-1}$. Hence $b_1 = x_1$ is weakly dominant.

## 4.2 First Price Sealed Bid Auction (FPA)

**Definition 4.3: First-Price Sealed-Bid Auction**

Each bidder simultaneously submits a sealed bid. The highest bidder wins and pays *her own* bid; ties are broken uniformly at random.

In a first price sealed bid auction, each bidder submits a bid without knowing the bids of others. The highest bidder wins the item and pays their own bid.

**Example (Motivating Example of Two Bidders).**

© Rui Zhou 2026. Do not redistribute without permission.

Suppose the true value $X_i$ for bidder $i$ is drawn from a uniform distribution on $[0, 1]$. Player 1 believes that Player 2 will bid $B_2 = kX_2$ for some $\frac{1}{2} \leq k < 1$. If Player 1 bids $b_1 \leq k$, her expected profit is

$$\begin{aligned}\mathbb{E}\left(\pi_1(b_1)\right) &= \Pr(kX_2 < b_1)\,(x_1 - b_1) \\ &= \Pr\left(X_2 < \frac{b_1}{k}\right)(x_1 - b_1) \\ &= \frac{b_1}{k}(x_1 - b_1).\end{aligned}$$

Then the best response of Player 1 to her belief $B_2 = kX_2$ with $\frac{1}{2} \leq k < 1$ is to bid $b_1 = \frac{x_1}{2}$.

Similarly, we can show that it is a symmetric equilibrium for both players to bid $b_i = \frac{x_i}{2}$ in a FPA.

**Remark (Why $k \geq 1/2$?).**

The condition $k \geq 1/2$ is strictly necessary to ensure that the unconstrained interior maximizer $b_1^* = \frac{x_1}{2}$ is globally valid for all possible valuations $x_1 \in [0, 1]$.

In the derivation, Player 1's expected profit function $\mathbb{E}(\pi_1(b_1)) = \frac{b_1}{k}(x_1 - b_1)$ is predicated on the assumption that $b_1 \leq k$, since Player 2's maximum possible bid is essentially $k \cdot 1 = k$. To guarantee that the optimal bid always falls within this valid domain (i.e., $\frac{x_1}{2} \leq k$), this inequality must hold for the supremum of $x_1$. Since $x_1 \leq 1$, this immediately requires $k \geq \frac{1}{2}$.

If we were to relax this condition such that $k < 1/2$, a bidder with a sufficiently high valuation ($x_1 > 2k$) would find her unconstrained best response to be $\frac{x_1}{2} > k$. However, since bidding exactly $k$ (or marginally above it) already guarantees a 100% probability of winning, bidding $\frac{x_1}{2}$ would strictly decrease her profit without any marginal gain in win rate. Consequently, the best response for high-valuation types would collapse into a corner solution $b_1 = k$. Thus, $k \geq 1/2$ rules out this corner solution and preserves the symmetric linear strategy over the entire support.

### 4.2.1 Symmetric Equilibrium of FPA

Now we extend the above example to the case with $n$ bidders whose private values are drawn from an arbitrary distribution function $F$. We start the analysis by assuming there is a symmetric equilibrium where each bidder bids $b_i = \beta(X_i)$ for some strictly increasing function $\beta$.

Without loss of generality, suppose bidder 1 has value $x_1$ and bids $b_1$, while all other

© Rui Zhou 2026. Do not redistribute without permission.

bidders follow the symmetric strategy $\beta$. The probability that bidder 1 wins is

$$\begin{aligned}
\Pr\left(b_1 > \max_{i\neq 1} b_i\right) &= \Pr\left(b_1 > \max_{i\neq 1} \beta(X_i)\right) \\
&= \Pr\left(b_1 > \beta\left(\max_{i\neq 1} X_i\right)\right) \\
&= \Pr\left(\beta^{-1}(b_1) > \max_{i\neq 1} X_i\right) \\
&= \Pr\big(X_i < \beta^{-1}(b_1), \forall i \neq 1\big) \\
&= \left[F(\beta^{-1}(b_1))\right]^{n-1}.
\end{aligned}$$

For simplicity, we define

$$Y_1 = \max_{i\neq 1} X_i.$$

Let $G$ be the distribution function of $Y_1$. Then we have

$$\begin{aligned}
G(y) &= \Pr(Y_1 \leq y) \\
&= \Pr\left(\max_{i\neq 1} X_i \leq y\right) \\
&= \Pr(X_i \leq y, \forall i \neq 1) \\
&= \left[F(y)\right]^{n-1}.
\end{aligned}$$

Then the probability that bidder 1 wins the auction can be rewritten as

$$\Pr\left(b_1 > \max_{i\neq 1} \beta(X_i)\right) = \Pr\big(Y_1 < \beta^{-1}(b_1)\big) = G(\beta^{-1}(b_1)).$$

The expected profit for bidder 1 is then given by

$$\begin{aligned}
\mathbb{E}\left(\pi_1(b_1)\right) &= \Pr\left(b_1 > \max_{i\neq 1} \beta(X_i)\right)(x_1 - b_1) \\
&= G(\beta^{-1}(b_1))(x_1 - b_1).
\end{aligned}$$

Since $b_1$ is the best response to the symmetric strategy $\beta$, we have

$$\begin{aligned}
\frac{\mathrm{d}}{\mathrm{d}b_1}\mathbb{E}\left(\pi_1(b_1)\right) &= \frac{\mathrm{d}}{\mathrm{d}b_1} G(\beta^{-1}(b_1))(x_1 - b_1) \\
&= G'(\beta^{-1}(b_1)) \cdot \frac{\mathrm{d}}{\mathrm{d}b_1}\beta^{-1}(b_1) \cdot (x_1 - b_1) - G(\beta^{-1}(b_1)) \\
&= 0.
\end{aligned}$$

Note that $\frac{\mathrm{d}}{\mathrm{d}b_1}\beta^{-1}(b_1) = \frac{1}{\beta'(\beta^{-1}(b_1))}$.[1] And by the optimality of $\beta$, we have $b_1 = \beta(x_1)$

[1] This is true because $f(f^{-1}(x)) = x$ for any strictly increasing function $f$. Differentiating both sides with respect to $x$, we have $f'(f^{-1}(x)) \cdot \frac{\mathrm{d}}{\mathrm{d}x} f^{-1}(x) = 1$. Hence, $\frac{\mathrm{d}}{\mathrm{d}x} f^{-1}(x) = \frac{1}{f'(f^{-1}(x))}$.

© Rui Zhou 2026. Do not redistribute without permission.

and thus $\beta^{-1}(b_1) = x_1$. So the FOC can be rewritten as

$$\begin{aligned}
&\frac{g(x_1)}{\beta'(x_1)}(x_1 - \beta(x_1)) - G(x_1) = 0 \\
\Longrightarrow\ &\beta'(x_1)G(x_1) + \beta(x_1)g(x_1) = x_1 g(x_1) \\
\Longrightarrow\ &\frac{\mathrm{d}}{\mathrm{d}x_1}(\beta(x_1)G(x_1)) = x_1 g(x_1) \\
\Longrightarrow\ &\beta(x_1)G(x_1) - \beta(0)G(0) = \int_0^{x_1} yg(y)\mathrm{d}y \\
\Longrightarrow\ &\beta(x_1)G(x_1) = \int_0^{x_1} yg(y)\mathrm{d}y \\
\Longrightarrow\ &\beta(x_1) = \frac{\int_0^{x_1} yg(y)\mathrm{d}y}{G(x_1)} \\
\Longrightarrow\ &\beta(x_1) = \mathbb{E}\left[Y_1 | Y_1 < x_1\right].
\end{aligned}$$

A bidder with the lowest possible valuation $x = 0$ wins with probability zero, and any positive bid would yield a strictly negative expected payoff; hence $\beta(0) = 0$, which fixes the constant of integration.

**Remark.**

- **Why $\frac{\int_0^{x_1} yg(y)\mathbf{d}y}{G(x_1)} = \mathbb{E}\left[Y_1 | Y_1 < x_1\right]$?**

  This equality follows directly from the definition of conditional expectation.

  First, determine the conditional density of the highest competing valuation $Y_1$ given the event that the bidder wins (i.e., $Y_1 < x_1$). Restricting $Y_1$ to $[0, x_1]$, we normalize $g(y)$ by the total probability $\Pr(Y_1 < x_1) = G(x_1)$ to obtain

  $$f(y \mid Y_1 < x_1) = \frac{g(y)}{G(x_1)} \quad \text{for } 0 \le y < x_1.$$

  By the standard definition of expected value, we integrate the variable $y$ multiplied by its conditional PDF over the restricted support:

  $$\begin{aligned}
  \mathbb{E}\left[Y_1 | Y_1 < x_1\right] &= \int_0^{x_1} y \cdot f(y \mid Y_1 < x_1)\mathrm{d}y \\
  &= \int_0^{x_1} y \left(\frac{g(y)}{G(x_1)}\right)\mathrm{d}y.
  \end{aligned}$$

  Since $G(x_1)$ only depends on $x_1$ and acts as a constant with respect to the integration variable $y$, we can factor it out of the integral:

  $$\mathbb{E}\left[Y_1 | Y_1 < x_1\right] = \frac{1}{G(x_1)} \int_0^{x_1} yg(y)\mathrm{d}y.$$

- **Economic Intuition of FPA's Bidding Strategy**: The result $\beta(x) = \mathbb{E}\left[Y_1 | Y_1 < x\right]$ reveals a profound economic intuition that elegantly connects the First-Price Auction (FPA) with the Second-Price Auction (SPA).

  In an SPA, a winning bidder pays exactly the highest competing valuation, $Y_1$. This

© Rui Zhou 2026. Do not redistribute without permission.

payment is determined *ex-post* (after the bids are submitted).

In an FPA, however, a winning bidder pays her own bid. Because the bid must be chosen *ex ante* (without knowing $Y_1$), the bidder effectively "simulates" the SPA in her mind: conditional on winning (which implies $Y_1 < x$), her best unbiased estimate of what she *would* have paid in an SPA is the expected highest competing valuation, $\mathbb{E}[Y_1 \mid Y_1 < x]$. The optimal FPA bid is therefore the bidder's expected SPA payment.

In an FPA, a bidder with value $x$ bids $\beta(x) = \mathbb{E}[Y_1 \mid Y_1 < x]$ and pays her own bid upon winning. She wins iff $Y_1 < x$, which happens with probability $G(x)$. The expected payment $m(x)$ is therefore

$$m(x) = G(x) \cdot \mathbb{E}\left[Y_1 | Y_1 < x\right] = \int_0^x y g(y) \mathrm{d}y.$$

Furthermore, by integration by parts,

$$\begin{aligned} m(x) &= \beta(x) G(x) \\ &= \int_0^x y g(y) \mathrm{d}y \\ &= \int_0^x y \mathrm{d}G(y) \\ &= \left[y G(y)\right]\Big|_0^x - \int_0^x G(y) \mathrm{d}y \\ &= x G(x) - \int_0^x G(y) \mathrm{d}y. \end{aligned}$$

This derived identity, $m(x) = xG(x) - \int_0^x G(y)\mathrm{d}y$, allows for a beautiful geometric decomposition of the total social surplus generated by the auction. This surplus-splitting perspective is visualized in the figure below:

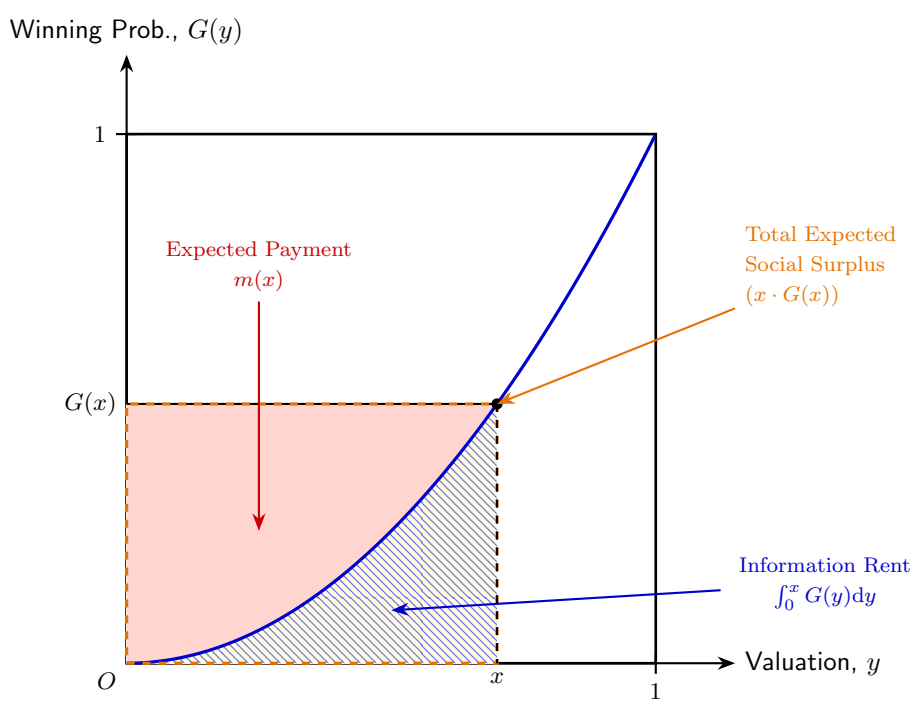


© Rui Zhou 2026. Do not redistribute without permission.

**Remark (Surplus Splitting and Information Rent).**

The identity $m(x) = xG(x) - \int_0^x G(y)\mathrm{d}y$ provides a powerful geometric and economic interpretation of the auction's outcome, as illustrated in the figure above:

- **Total Expected Surplus ($xG(x)$):** Represented by the entire rectangle defined by $(x, G(x))$, this is the total social value created when an object is allocated to a bidder with valuation $x$ (who wins with probability $G(x)$).
- **Information Rent ($\int_0^x G(y)\mathrm{d}y$):** This corresponds to the area *under* the curve $G(y)$. Because the seller cannot observe the bidder's private valuation $x$, the bidder extracts a surplus from the mechanism. This "rent" is the cost the seller must pay to induce truthful revelation. It increases with valuation, reflecting the greater bargaining power of high-value types.
- **Expected Payment ($m(x)$):** This is the area *above* the curve but within the $x \times G(x)$ rectangle. It is the portion of the total surplus that the seller successfully captures as revenue.

From this perspective, the Revenue Equivalence Theorem (shown later) holds because, in any efficient auction, the *information rent* is uniquely determined by the allocation rule $G(x)$.

**Remark (Intuition and Derivation of the Information Rent).**

- **Why the Information Rent is Equal to $\int_0^x G(y)\mathrm{d}y$?**
  In any auction, a high-value bidder wants to under-report her valuation to secure a lower price. Consider a bidder of true type $y + \mathrm{d}y$ who instead reports type $y$: she wins with probability $G(y)$, and because her actual valuation is higher by $\mathrm{d}y$, she gains a marginal utility of $G(y)\mathrm{d}y$ from the misreport.

  To achieve *Incentive Compatibility* (IC), the mechanism must compensate her for not lying. To prevent a type-$x$ bidder from mimicking *any* lower type all the way down to 0, the mechanism must compensate her for all these accumulated marginal temptations. Integrating from 0 to $x$ yields the total required information rent $\int_0^x G(y)\mathrm{d}y$.

© Rui Zhou 2026. Do not redistribute without permission.

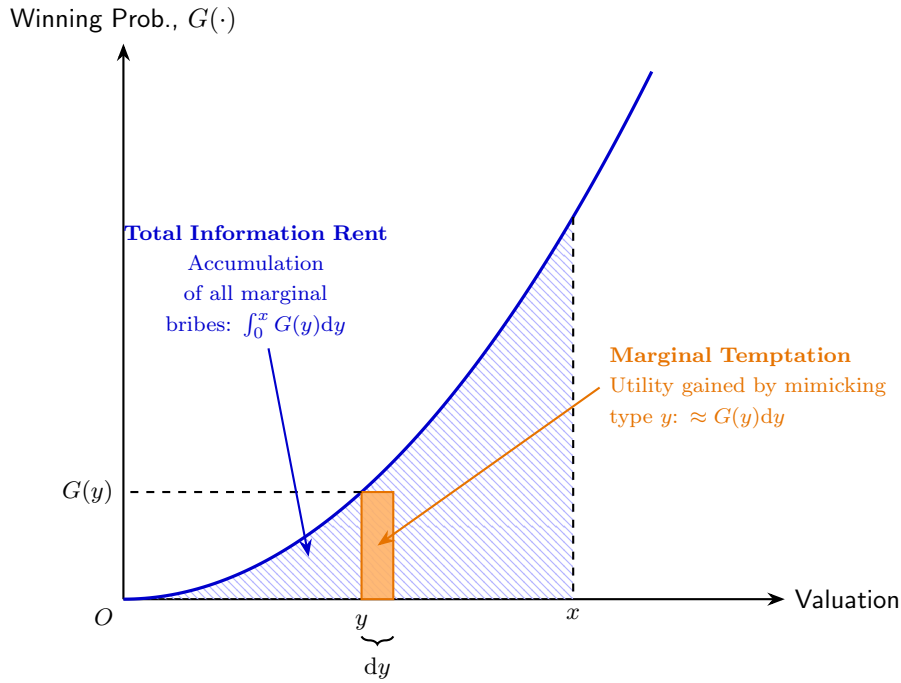


- **Mathematical Derivation (The Envelope Theorem):**
  Let $U(z|x) = xG(z) - m(z)$ be the expected utility of a bidder with true value $x$ reporting value $z$. Under truth-telling (which is optimal by design), the equilibrium utility—which is exactly the information rent—is defined as $V(x) = U(x|x)$.

  By the *Envelope Theorem*, the total derivative of $V(x)$ with respect to $x$ is simply the partial derivative of the objective function $U(z|x)$ with respect to the parameter $x$, evaluated at the optimal choice $z = x$:

  $$V'(x) = \left.\frac{\partial U(z|x)}{\partial x}\right|_{z=x} = G(x)$$

  This differential equation reveals that the marginal growth rate of the information rent is exactly the winning probability. Integrating both sides from 0 to $x$, and assuming the lowest type gets zero rent to satisfy Individual Rationality ($V(0) = 0$), we immediately obtain:

  $$V(x) = \int_0^x G(y)\mathrm{d}y$$

**Remark (An Alternative, Intuitive Form of the Bidding Strategy in a FPA).**

Recall that for a bidder with true value $x$, the bidding strategy is

$$\beta(x) = \frac{\int_0^x yg(y)\mathrm{d}y}{G(x)}.$$

© Rui Zhou 2026. Do not redistribute without permission.

By integration by parts,

$$\int_0^x yg(y)\mathrm{d}y = xG(x) - 0 \cdot G(0) - \int_0^x G(y)\mathrm{d}y.$$

We can rewrite the bidding strategy in a more intuitive form:

$$\beta(x) = x - \frac{\int_0^x G(y)\mathrm{d}y}{G(x)}.$$

This formulation perfectly captures the essence of strategic behavior in a FPA. A rational bidder will never bid their true valuation $x$ (which would yield zero surplus); instead, they "shade" their bid downwards. This formula tells us exactly *how much* they shade:

- $x$ is the bidder's true valuation (their absolute maximum willingness to pay).
- $\int_0^x G(y)\mathrm{d}y$ is the expected *information rent* the bidder is theoretically entitled to because their valuation is private.
- Dividing this expected rent by $G(x)$ converts it into a *conditional rent*—the actual rent the bidder secures *conditional on winning the auction.*

  In an FPA, a winning bidder pays exactly her bid $\beta(x)$, so her realized profit conditional on winning is the deterministic amount $x - \beta(x)$. Since she wins only with probability $G(x)$, achieving the unconditional expected information rent $\int_0^x G(t)\mathrm{d}t$ requires the realized profit to be scaled up by $1/G(x)$. Equating realized profit to the conditional rent yields $x - \beta(x) = \int_0^x G(t)\mathrm{d}t/G(x)$, the standard bid-shading identity.

## 4.3 Reserve Prices

**Definition 4.4: Reserve Price**

A **reserve price** $r \geq 0$ is a minimum acceptable bid: the auction sells the object only if at least one submitted bid satisfies $b_i \geq r$. If all bids fall below $r$, the seller retains the object and no payment is made.

### 4.3.1 SPA with $n = 2$

In a SPA, the seller can set a reserve price $r$, meaning the item is only sold if the highest bid is no less than $r$. With the reserve price, the payment is then the maximum between the payment determined by the auction rule (namely the second-highest bid) and the reserve price. A natural question arises: *Is introducing a reserve price strictly profitable for the seller?*

To answer this intuitively, let's assume there are $n = 2$ bidders with valuations drawn from a standard Uniform distribution, $F \sim U[0, 1]$. We can perform a marginal analysis by comparing the seller's revenue with no reserve price ($r = 0$) to a scenario with a very small positive reserve price ($r \approx 0$). We can divide the possible valuation realizations into three

© Rui Zhou 2026. Do not redistribute without permission.

distinct cases:

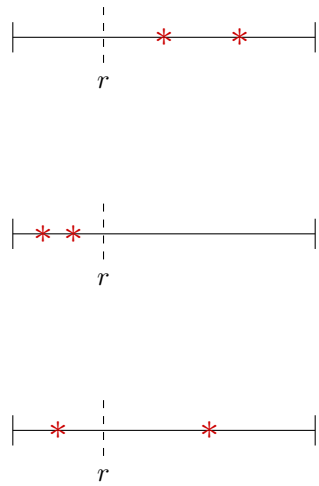


- **Case 1: Both valuations are greater than $r$.** The reserve price has no effect on the outcome. The bidder with the highest valuation wins and pays the second-highest valuation. The expected revenue is exactly the same as if there were no reserve price.
- **Case 2: Both valuations are below the reserve price $r$.** In this case, the item is not sold, yielding a revenue of zero. If there were no reserve price, the seller would have received the lower of the two valuations. Since both are drawn from $U[0, r]$, the conditional expected payment would be $r/3$. This case happens with probability $r^2$. Therefore, the unconditional expected loss for the seller is:

$$\text{Expected Loss} = r^2 \cdot \frac{r}{3} = \frac{r^3}{3}$$

- **Case 3: One valuation is above $r$, and the other is below $r$.** The bidder with the higher valuation wins. Without a reserve price, they would only pay the losing bidder's valuation (expected to be $r/2$). With the reserve price, the winner is forced to pay $r$. The probability of exactly one bidder drawing a value above $r$ and one below $r$ is $2r(1-r)$. Therefore, the expected gain for the seller is:

$$\text{Expected Gain} = 2r(1-r) \cdot \left(r - \frac{r}{2}\right) = 2r(1-r) \cdot \frac{r}{2} = r^2 - r^3$$

## Claim

$r = 0$ is never optimal for a revenue-maximizing seller.

***Proof for Claim.***

We evaluate the net change in expected revenue when introducing a small positive reserve price ($r > 0$) compared to not having one ($r = 0$). The net change is the expected gain from Case 3 minus the expected loss from Case 2:

$$\Delta\text{Revenue} = \text{Expected Gain} - \text{Expected Loss} = (r^2 - r^3) - \frac{r^3}{3} = r^2 - \frac{4}{3}r^3$$

© Rui Zhou 2026. Do not redistribute without permission.

For a sufficiently small $r$ (specifically, $0 < r < \frac{3}{4}$), the $r^2$ term strictly dominates the $r^3$ term, meaning $r^2 > \frac{4}{3}r^3$. Consequently, $\Delta\text{Revenue} > 0$. Since setting a tiny positive reserve price strictly increases expected revenue compared to setting $r = 0$, a reserve price of zero is never optimal.

### 4.3.2 Optimal Reserve Price in a FPA

**Claim**

In an FPA with reserve price $r$, the symmetric bidding strategy for a bidder with valuation $x$ is

$$\beta^r(x) = \begin{cases} \frac{rG(r)+\int_r^x tg(t)\mathrm{d}t}{G(x)} & \text{if } x > r, \\ 0 & \text{if } x \leq r. \end{cases}$$

where $G(x) = F(x)^{n-1}$ is the distribution function of the highest valuation among the $n-1$ opposing bidders, and $g(x)$ is its corresponding probability density function.

***Proof for Claim.***

- **The Conditional Expectation (Intuitive Approach)**

  Without a reserve price, a bidder's optimal strategy in a standard FPA is to bid the expected value of the highest competing valuation, conditional on winning: $\beta(x) = \mathbb{E}[Y_1|Y_1 < x]$, where $Y_1$ is the highest valuation among opponents.

  When a reserve price $r$ is introduced, the seller acts as a "ghost bidder" who steadfastly bids $r$. To win, a bidder must defeat both the actual opponents ($Y_1 < x$) and the ghost bidder. Thus, the *effective* highest competing valuation becomes $\max(Y_1, r)$. The optimal bid is the conditional expectation of this effective opponent:

  $$\beta^r(x) = \mathbb{E}[\max(Y_1, r) \mid Y_1 < x]$$

  By definition of conditional expectation, we integrate over the distribution of $Y_1$ up to $x$:

  $$\beta^r(x) = \int_0^x \max(y, r) \left(\frac{g(y)}{G(x)}\right) \mathrm{d}y$$

  We can split this integral at $t = r$:

  $$\beta^r(x) = \frac{1}{G(x)} \left(\int_0^r rg(y)\mathrm{d}y + \int_r^x yg(y)\mathrm{d}y\right)$$

  Since $\int_0^r g(y)\mathrm{d}y = G(r)$, the first term evaluates to $rG(r)$. This leads directly to our familiar result:

  $$\beta^r(x) = \frac{rG(r) + \int_r^x yg(y)\mathrm{d}y}{G(x)}$$

- **First Order Approach**

  Suppose a bidder has a true valuation $x > r$, but considers reporting a valuation $z$ (and bidding $\beta(z)$). The bidder's expected utility is the probability of winning multiplied

© Rui Zhou 2026. Do not redistribute without permission.

by the surplus:

$$U(z, x) = (x - \beta(z))G(z)$$

To find the symmetric equilibrium, truth-telling must be optimal, meaning the first-order derivative with respect to $z$ evaluated at $z = x$ must be zero:

$$\left.\frac{\partial U(z, x)}{\partial z}\right|_{z=x} = -\beta'(x)G(x) + (x - \beta(x))g(x) = 0$$

Rearranging terms yields a linear ordinary differential equation (ODE):

$$\beta'(x)G(x) + \beta(x)g(x) = xg(x) \implies \frac{d}{dx}[\beta(x)G(x)] = xg(x)$$

Integrating both sides from the reserve price $r$ to $x$:

$$\beta(x)G(x) - \beta(r)G(r) = \int_r^x yg(y)\mathrm{d}y$$

Consider the marginal buyer whose valuation is exactly $x = r$. To acquire the item without strictly losing money, this bidder must bid exactly the reserve price. Thus, $\beta(r) = r$. Substituting this boundary condition into the equation gives:

$$\beta(x)G(x) - rG(r) = \int_r^x yg(y)\mathrm{d}y \implies \beta^r(x) = \frac{rG(r) + \int_r^x yg(y)\mathrm{d}y}{G(x)}$$

**Proposition 4.5: Optimal Reserve Price For a FPA**

The optimal reserve price $r^*$ for a revenue-maximizing seller in a First-Price Auction with $n$ bidders and valuations drawn from $F$ satisfies

$$r^* = \frac{1 - F(r^*)}{f(r^*)}.$$

***Proof for Proposition.***

For a FPA with a reserve price $r$, the symmetric bidding strategy is then

$$\beta^r(x) = \begin{cases} \frac{rG(r)+\int_r^x yg(y)\mathrm{d}y}{G(x)} & \text{if } x > r, \\ 0 & \text{if } x \leq r. \end{cases}$$

The expected payment for a bidder with $x > r$ is therefore

$$m^r(x) = rG(r) + \int_r^x tg(t)\mathrm{d}t.$$

And the expected payment for a bidder with value $x \leq r$ is zero: $m^r(x) = 0$.

© Rui Zhou 2026. Do not redistribute without permission.

The *ex-ante* expected payment for a bidder with valuation drawn from $F$ is then

$$m^r = \int_0^1 m^r(x)f(x)dx = r\left[1 - F(r)\right]G(r) + \int_r^1 t\left[1 - F(t)\right]g(t)\mathrm{d}t.$$

Hence, the optimal reserve price satisfies

$$\frac{\mathrm{d}m^r}{\mathrm{d}r} = \left[1 - F(r)\right]G(r) - rG(r)f(r) + r\left[1 - F(r)\right]g(r) - r\left[1 - F(r)\right]g(r) = 0.$$

Therefore, the optimal reserve price $r^*$ is

$$r^* = \frac{1 - F(r^*)}{f(r^*)}.$$

Interestingly, the optimal reserve price does not depend on the number of bidders $n$.

## 4.4 Revenue Equivalence Theorem

We make a comparison of the expected payment in a FPA and a SPA with private value $x$.

In an SPA, the bidder bids $x$ and, upon winning, pays the second-highest bid $Y_1$. Conditional on winning, expected payment is $\mathbb{E}[Y_1 \mid Y_1 < x]$; the unconditional expected payment, weighting by the win probability $G(x)$, is $m(x) = G(x) \cdot \mathbb{E}[Y_1 \mid Y_1 < x]$.

In an FPA, the bidder bids $\beta(x) = \mathbb{E}[Y_1 \mid Y_1 < x]$ and pays her own bid upon winning, which occurs iff $Y_1 < x$ (probability $G(x)$).

The expected payment in both mechanisms for a bidder with valuation $x$ is:

$$m(x) = G(x) \cdot \mathbb{E}\left[Y_1 | Y_1 < x\right] = \int_0^x yg(y)\mathrm{d}y.$$

Ex ante, the expected payment for each player is

$$m = \mathbb{E}_X\left[m(X)\right] = \int_0^1 m(x)f(x)\mathrm{d}x = \int_0^1 \left(\int_0^x yg(y)\mathrm{d}y\right) f(x)\mathrm{d}x.$$

Hence, the expected revenue (oftentimes called the expected selling price) is the same in a FPA and a SPA, which is just $n \times m$.

This result is more general than just comparing the FPA and SPA. It actually applies to any two auction mechanisms that satisfy certain conditions, as stated in the following theorems.

© Rui Zhou 2026. Do not redistribute without permission.

### Theorem 4.6: Revenue Equivalence Theorem (RET)

Assume there are $n$ risk-neutral bidders with independent private values drawn from a strictly increasing cumulative distribution function $F$ on an interval $[0, \bar{x}]$. Any auction mechanism that satisfies the following two conditions:

1. **Allocative Efficiency:** The object is always allocated to the bidder with the highest valuation.
2. **Zero Payment at the Bottom:** Any bidder with the lowest possible valuation ($x = 0$) has an expected payment of zero.

will yield the exact same expected payment $m(x)$ for every bidder type $x$, and consequently, generate the exact same expected revenue for the seller.

### Corollary 4.7

Standard auction formats such as the FPA and the SPA all satisfy these conditions and are therefore revenue equivalent.

**Remark (The "Pie-Splitting" Economic Intuition).**

- **Baking the Pie (Total Social Surplus):** The allocation rule decides who ultimately receives the object, which dictates the total value created by the mechanism (the size of the pie). Because standard auctions always allocate the object to the bidder with the highest valuation, they are all efficient and thus "bake" the exact same size of pie.
- **The Bidder's Slice (Information Rent):** Because valuations are private, the seller must leave some surplus to the bidders to induce truth-telling. By the Envelope Theorem, this information rent is dictated *entirely* by the allocation probability ($V(x) = \int_0^x G(y)\mathrm{d}y$). Since the allocation rules are identical across these standard auctions, the winning probabilities are identical, meaning the bidders carve out the exact same slice of the pie.
- **The Seller's Slice (Expected Revenue):** The fundamental accounting identity of any mechanism is:

  $$\text{Expected Revenue} = \text{Total Social Surplus} - \text{Information Rent}$$

  If two auctions bake the exact same total pie, and the bidders are guaranteed to take the exact same size slice in both formats, then the remaining portion—the revenue left on the table for the seller—must logically be absolutely identical.

Having established the revenue equivalence of standard auctions, a natural question arises for a profit-maximizing seller: *Are standard auctions the best we can do?* To answer this, we must step out of the rigid, predefined rules of specific auction formats (like FPA or SPA) and adopt the "God's-eye view" of a mechanism designer. Instead of asking "what happens

© Rui Zhou 2026. Do not redistribute without permission.

in an auction," we ask: "over the space of *all possible mechanisms*, which one maximizes the seller's expected revenue?" This requires a more fundamental, generalized mathematical law that governs *any* allocation rule, not just the efficient ones. The general framework—which we develop in the next chapter via the Revelation Principle and virtual values—reveals that the optimal mechanism is itself a slightly modified auction with a carefully chosen reserve price, and that VCG (covered in the second half of the next chapter) is the corresponding object on the *efficiency*-maximizing side rather than the revenue-maximizing one.

### Theorem 4.8: General RET (Payoff Equivalence)

Assume there are $n$ risk-neutral bidders with independent private values drawn from a strictly increasing cumulative distribution function $F$ on an interval $[0, \bar{x}]$. Consider any two incentive-compatible mechanisms. If both mechanisms satisfy the following two conditions:

1. **Identical Allocation Rule:** The interim probability of receiving the object for a bidder with valuation $x$, denoted by $q(x)$, is exactly the same across both mechanisms.
2. **Identical Payment at the Bottom:** The expected payment of a bidder with the lowest possible valuation ($x = 0$) is the same across both mechanisms (typically zero).

Then, both mechanisms yield the exact same expected payment $m(x)$ for every bidder type $x$, and consequently, they generate the exact same expected revenue for the seller.

**Remark (Why is this theorem more fundamental than the Standard RET?).**

The Standard RET rigorously requires **Allocative Efficiency**—meaning the object *must* always go to the bidder with the highest valuation. The General RET, however, removes this straitjacket. It mathematically proves that for *any* arbitrary, monotonic allocation rule $q(x)$ (even inefficient ones), the payment function is completely and uniquely pinned down by the Envelope Theorem:

$$m(x) = m(0) + xq(x) - \int_0^x q(t)\mathrm{d}t$$

This "Payoff Equivalence" is the ultimate theoretical weapon for the seller. It implies that the seller's revenue maximization problem is no longer a daunting search through complex, multidimensional payment rules, but can be *entirely reduced to designing the optimal one-dimensional allocation rule* $q(x)$.

Armed with the General RET, the seller arrives at a somewhat counterintuitive economic truth: *To strictly maximize revenue, the seller must be willing to destroy social efficiency.* If the seller simply uses a standard efficient auction, they leave money on the table in the form of excessive information rents conceded to high-value bidders. To extract more surplus, for example, the seller can strategically introduce a *reserve price* ($r$). With a reserve price,

© Rui Zhou 2026. Do not redistribute without permission.

if the highest bidder's valuation is positive but falls below $r$, the seller retains the object. This explicitly violates the "*Allocative Efficiency*" condition of the Standard RET (because mutually beneficial trade is artificially blocked). However, the *General RET still applies perfectly*! We merely apply the generalized math to a newly shaped allocation rule $q^r(x)$, where the probability of winning is squashed to zero for all types below $r$. This strategic manipulation of the allocation rule to squeeze information rents leads directly to the seminal framework of Optimal Mechanism Design.

## 4.5 Risk Attitude and the Breakdown of Revenue Equivalence

All the analysis before was built on the assumption that bidders are risk-neutral. What happens when we relax this assumption? Suppose now that the bidders are strictly *risk averse*, with a utility function $u(\cdot)$ such that $u' > 0$, $u'' < 0$, and $u(0) = 0$.

First, for the SPA, risk aversion *has no effect* on the equilibrium of an SPA. Truth-telling ($\beta(x) = x$) remains a weakly dominant strategy.

However, the equilibrium in an FPA changes dramatically. Suppose the other $n-1$ bidders follow a strictly increasing and differentiable symmetric strategy $\gamma(x)$ with $\gamma(0) = 0$. A bidder with true value $x$ who submits bid $b$ has expected utility

$$\begin{aligned} U(b,x) &= u(x-b) \cdot \Pr\left(b > \max_{j \neq i} \gamma(X_j)\right) \\ &= u(x-b) \cdot \Pr\left(\gamma^{-1}(b) > \max_{j \neq i} X_j\right) \\ &= u(x-b) \cdot G(\gamma^{-1}(b)), \end{aligned}$$

where $G(\cdot)$ is the distribution function of the highest opposing valuation.

To find the optimal bid, we take the FOC with respect to $b$:

$$\frac{\partial U(b,x)}{\partial b} = -u'(x-b) \cdot G(\gamma^{-1}(b)) + u(x-b) \cdot \frac{g(\gamma^{-1}(b))}{\gamma'(\gamma^{-1}(b))} = 0$$

In a symmetric equilibrium, the optimal bid must be $b = \gamma(x)$, which implies $\gamma^{-1}(b) = x$. Substituting this into the FOC gives:

$$-u'(x-\gamma(x)) \cdot G(x) + u(x-\gamma(x)) \cdot \frac{g(x)}{\gamma'(x)} = 0$$

Rearranging this equation yields the differential equation for the bidding strategy under risk aversion:

$$\gamma'(x) = \frac{u(x-\gamma(x))}{u'(x-\gamma(x))} \cdot \frac{g(x)}{G(x)}$$

To see how this compares to the risk-neutral case, recall that the risk-neutral bidding strategy $\beta(x)$ satisfies:

$$\beta'(x) = (x - \beta(x)) \cdot \frac{g(x)}{G(x)}$$

Note that because $u(\cdot)$ is a strictly concave function with $u(0) = 0$, its graph lies strictly

© Rui Zhou 2026. Do not redistribute without permission.

*below* its tangent lines. Geometrically, for any point $t > 0$, the slope of the secant line connecting the origin to $(t, u(t))$ is strictly steeper than the slope of the tangent line at that exact point. Therefore, we directly have the slope inequality:

$$\frac{u(t)}{t} > u'(t) \implies \frac{u(t)}{u'(t)} > t$$

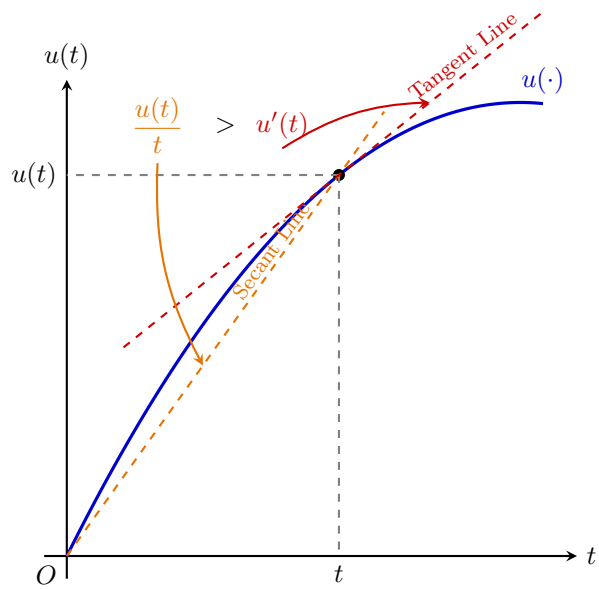


**Proposition 4.9: Overbidding under Risk Aversion**

For any strictly positive valuation $x > 0$, a risk-averse bidder bids strictly higher than a risk-neutral bidder:

$$\gamma(x) > \beta(x).$$

***Proof for Proposition.***

Note that a bidder with the lowest possible valuation ($x = 0$) bids zero regardless of their risk attitude, so we have the boundary condition $\gamma(0) = \beta(0) = 0$.

Suppose, for the sake of contradiction, that there exists some $x > 0$ where $\gamma(x) \leq \beta(x)$. Let $t = x - \gamma(x) > 0$. By the concavity property discussed above, $\frac{u(t)}{u'(t)} > t$, meaning:

$$\frac{u(x - \gamma(x))}{u'(x - \gamma(x))} > x - \gamma(x) \geq x - \beta(x).$$

Substituting this strict inequality into our differential equation for $\gamma'(x)$:

$$\gamma'(x) = \frac{u(x - \gamma(x))}{u'(x - \gamma(x))} \cdot \frac{g(x)}{G(x)} > (x - \beta(x)) \cdot \frac{g(x)}{G(x)} = \beta'(x).$$

This implies that at any point where $\gamma(x)$ is equal to or below $\beta(x)$, the slope of $\gamma$ is strictly steeper than the slope of $\beta$. Since both functions start at the same origin ($\gamma(0) = \beta(0) = 0$), $\gamma$ must immediately rise above $\beta$. Furthermore, $\gamma$ can never cross $\beta$ from above, because at any intersection, $\gamma$ must be steeper. Thus, we conclude $\gamma(x) > \beta(x)$

© Rui Zhou 2026. Do not redistribute without permission.

for all $x > 0$.

**Remark (The "Insurance Premium" Intuition).**

Why does a risk-averse bidder bid higher? In a First-Price Auction, a bidder faces a trade-off: bidding lower increases the profit *if* they win, but increases the probability of losing and getting exactly zero. A risk-averse individual deeply dislikes the downside risk of losing the item. To protect themselves against this zero-payoff outcome, they are willing to sacrifice some potential profit by bidding more aggressively. In essence, the difference $\gamma(x) - \beta(x)$ is an **insurance premium** the bidder pays to increase their probability of winning.

Because risk aversion causes bidders to bid more aggressively in the FPA, while leaving the dominant strategy in the SPA unchanged, the expected payments in the FPA are strictly higher. Therefore, under risk aversion, the First-Price Auction generates strictly higher expected revenue for the seller than the Second-Price Auction.

**Remark (Chapter Summary).**

Single-object private-value auction theory rests on three pillars. *Strategic equivalence within the standard family.* The second-price (Vickrey) auction has a weakly dominant truthful strategy, and the first-price auction has a unique symmetric equilibrium $\beta(x) = \mathbb{E}[Y_1 \mid Y_1 < x]$ in which bidders shade their bids by exactly their conditional information rent. *Revenue equivalence* (Theorem 4.4). All standard auctions that allocate the object to the highest-value bidder and give zero rent to the lowest type generate the same expected revenue—the auction format is irrelevant for revenue under risk neutrality, independent private values, and symmetric bidders. The three assumptions matter: risk aversion breaks revenue equivalence in favor of the FPA, asymmetric distributions break it ambiguously, and correlated values can be exploited by full-surplus extraction mechanisms. *Optimal reserves.* Setting $r^* = \varphi^{-1}(0)$, the root of the virtual value, strictly improves revenue at the cost of efficiency—the seller deliberately sometimes refuses to sell to extract higher payments from high-value bidders.

© Rui Zhou 2026. Do not redistribute without permission.

# Part IV

# Matching

# Part V

# Information and Dynamic Games

# Part VI

# Problem Sets and Solutions

# Part VII

# Exams and Solutions