Posted-Price Mechanisms: Approximating Revenue

CS 1951k/2951z

2017-03-04

We show that simple auctions can generate near-optimal revenue by using non-optimal reserve prices.

1 *Simple vs. Optimal Auctions*

While Myerson's analysis yields an elegant solution to the optimal auction design problem, his auction is not widely used in practice. There are likely multiple reasons why not, such as its complexity, and its heavy reliance on distributional knowledge. On the other hand, posted-price mechanisms *are* widely used in practice. So we might be interested to know, what is the *optimal* posted-price mechanism? We can perhaps answer this question in a single-parameter environment by spelling out the revenue function, taking derivatives, etc.,[1] but we take an alternative approach in this lecture. We analyze a randomized posted-price mechanism. Perhaps surprisingly, the guarantees we eventually derive do not rely on distributional knowledge.[2]

[1] sounds like a good homework exercise

[2] They do, however, rely on the regularity assumption.

The mechanism we will analyze is simple, namely "post a random price." We analyze this mechanism in some detail, first assuming only one bidder.[3] We then generalize to n symmetric bidders, meaning all bidders' draw their values from the same distribution.

[3] Buyer, I suppose, but we'll stick with bidder.

2 *Posted-Price Mechanism*

In **posted-price** mechanism, the center[4] announces (i.e., posts) the price π at which they are willing to sell the good,[5] after which any bidder who indicates that they are willing to pay the posted price is uniformly eligible to win the good. The winner is then charged the posted price π, and all others pay nothing.

[4] This term generalizes "auctioneer," as a posted-price mechanism is not in fact an auction. We could equivalently use the term "mechanism designer" (but that's a mouthful).

[5] Sort of like how when we go out for coffee, we pay whatever price the establishment has decided on, irrespective of possible competition from other people. (Well, aside from the long lines.)

Given a posted price mechanism with price π, we can ask:

- How should bidders behave: i.e., what should they bid?
- Assuming they behave as predicted, how good is this outcome?

"Good" in the second question implies that we are measuring something. In this lecture, that something is total expected revenue, which is the sum of the total expected payments. In the next lecture, that something will be total expected welfare.

3 Bidder Behavior

The analysis to determine what strategy a bidder should use in a posted-price mechanism is similar to that of the second-price auction. The (familiar) case analysis is described graphically in Figures 1 and 2, and summarized in Figure 3. Thus, we see that the posted-price mechanism for one good is DSIC: regardless of what any other bidder does, bidding one's true value is a dominant strategy.

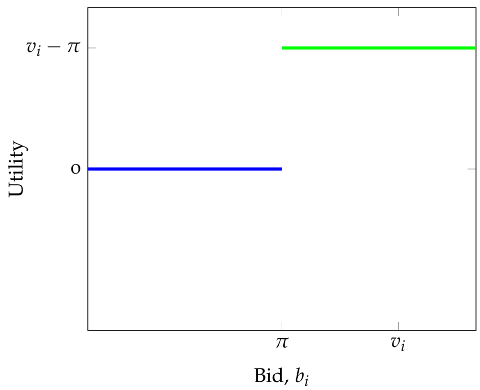

Figure 1: The utility of bidder i if the posted price is smaller than v, as a function of what she bids.

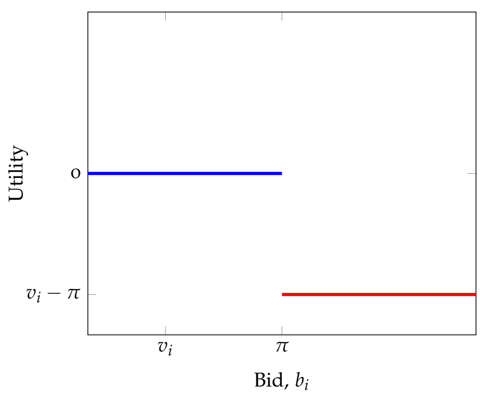

Figure 2: The utility of bidder i if the posted price is larger than v, as a function of what she bids.

Because payments for the winner are pre-determined by the auctioneer, a bid in this mechanism is in fact a binary signal. Placing a bid $b_i \geq \pi$ is telling the auctioneer "I am willing to purchase the good at that price." Placing a bid $b_i < \pi$ is telling the auctioneer "I am not willing to purchase the good at that price." Thus, the following bids also comprise a dominant strategy:

$$b_i \in \begin{cases} [\pi, \infty), & \text{if } v_i \geq \pi \\ (-\infty, \pi), & \text{otherwise.} \end{cases}$$

In other words, there are multiple dominant strategies in this mechanism. For example, the following two strategies would each do just

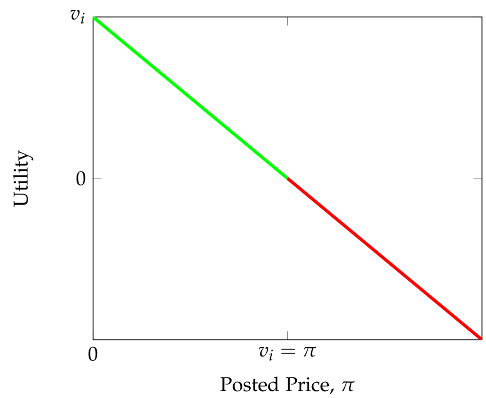

Figure 3: The utility of bidder i if she wins, as a function of the posted price.

as well as bidding truthfully in a posted-price mechanism:

$$b_i = \begin{cases} \pi, & \text{if } v_i \geq \pi \\ 0, & \text{otherwise} \end{cases}$$

$$b_i = \begin{cases} \infty, & \text{if } v_i \geq \pi \\ -\infty, & \text{otherwise.} \end{cases}$$

Because of the myriad of dominant strategies, the posted-price mechanism differs from the second-price auction in an important way. *Bidders need not bid truthfully!* In particular, bidders need not divulge their private information in order to maximize their utilty.

4 *Revenue Maximization*

Unlike the basic first- and second-price auctions (without a reserve), which always produce a winner, the posted-price mechanism provides no such guarantee. If the posted price is larger than the upper bound on bidder values, then no one will ever win. Thus, given distributional knowledge about bidders' values, the auctioneer could reason about what posted price would maximize revenue.

Assuming only one bidder, if the mechanism posts a price π, then the probability of a sale is equal to the probability that the bidder's value is at least π: $\Pr(v \geq \pi) = 1 - \Pr(v \leq \pi) = 1 - F(\pi)$.

The expected revenue function is thus given by:

$$R(\pi) = \pi (1 - F(\pi)).$$

More concretely, when values are distributed uniformly on $[0, 1]$, then the expected revenue at posted price π is given by:

$$R(\pi) = \pi (1 - \pi).$$

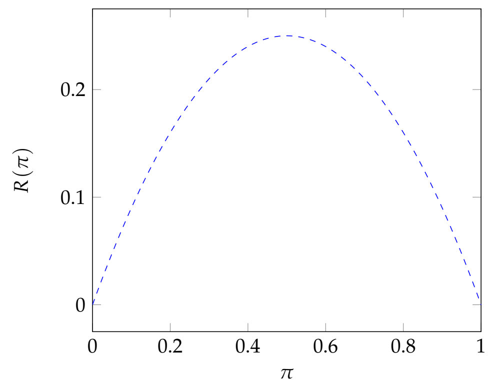

Figure 4: Expected revenue at posted price π, assuming one bidder whose values are uniformly distributed on $[0, 1]$.

This revenue curve is depicted in Figure 4.

Why does the revenue curve have this shape? Well, if the posted price is 0, then revenue is necessarily 0. Moreover, increasing the posted price by a little bit increases revenue. At the other extreme, if the posted price is 1, then revenue is again 0, as the probability of drawing a bidder whose type is 1 is 0. Again, decreasing the posted price by a little bit increases revenue.

Exercise: Solve for the optimal posted price assuming one bidder whose values are distributed uniformly on $[0, 1]$. Can you provide an interpretation for this price in terms of the optimal auction?

Rather than set the price at an optimum, which would yield the optimal revenue OPT, let's consider the simple mechanism in which the center selects a price at random, simply by sampling from the distribution F. We call the expected revenue of this mechanism APX: i.e.,

$$\begin{aligned}\text{APX} &= \mathop{\mathbb{E}}_{\pi \sim F}[R(\pi)] \\ &= \int_{\underline{v}}^{\overline{v}} \pi \, (1 - F(\pi)) \, f(\pi) \, \mathrm{d}\pi\end{aligned}$$

Just how well can this mechanism do? That is, what is the ratio of APX to OPT? We will develop some machinery in this lecture that will enable us to answer this question. The machinery is a bit complex, but with it, the answer to the question will be simple.

Here is a summary of the necessary machinery, together with our plan for the rest of today's lecture:

1. We redefine revenue in terms of quantiles, which leads to a simple interpretation of APX as the area under the revenue curve.

2. Next, we show that the virtual value function is the derivative of the revenue curve. By the regularity assumption, this derivative is non-decreasing, which implies the revenue curve is concave.

3. Finally, we derive an approximation ratio by picture.

4. All of the above applies only in the single-bidder case. We conclude by showing how to extend this reasoning to a symmetric setting with multiple bidders, assuming infinite supply. The final result is an approximation ratio for a **prior-independent** posted-price mechanism with multiple bidders, assuming infinite supply.

   Aside: An analysis of a mechanism is **prior-independent** if no distributional knowledge about the participants' private information is assumed. Such an analysis would be worst-case in the sense that it holds for all distributions (but it might still involve expectations over an unspecified distribution, as does the present analysis).

   An analysis of a mechanism is **prior-free** if it is not even assumed that participants draw their private information from distributions. Such an analysis would be worst-case in the stronger sense that it holds for all realizations of the participants' private information.

5 *Revenue in Quantile Space*

Recall the formula for expected revenue at posted price π:

$$R(\pi) = \pi \left(1 - F(\pi)\right).$$

In words, this quantity is the product of the posted price and the probability of a sale. The probability of a sale is the probability that a draw from F exceeds π. But this is precisely the meaning of a quantile. So at a given quantile q, the probability of a sale is simply q. Moreover, the value of a sale at quantile q is $v(q) = F^{-1}(1-q)$: i.e.,

$$R(q) = F^{-1}(1-q)\, q,$$

Next, let's investigate expected revenue in quantile space:

$$\begin{aligned}
\mathop{\mathbb{E}}_{q \sim U[0,1]}\left[R(q)\right] &= \mathop{\mathbb{E}}_{q \sim U[0,1]}\left[F^{-1}(1-q)\, q\right] \\
&= \int_0^1 F^{-1}(1-q)\, q\, f(q)\, \mathrm{d}q \\
&= \int_0^1 F^{-1}(1-q)\, q\, \mathrm{d}q \\
&= \int_0^1 R(q)\, \mathrm{d}q
\end{aligned}$$

The first step in this derivation follows from the definition of the revenue curve. The second step follows from the definition of expectation. The third step follows from the fact that quantiles are necessarily uniformly distributed, so that $f(q) = 1$. In words, *in quantile space, the expected revenue is the area under the revenue curve.*

Sample revenue curves are plotted in value and quantile space in Figures 5 through 8. The area under the curves plotted in quantile space is the expected revenue.

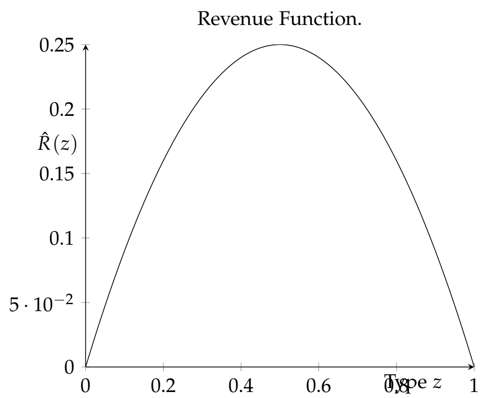

Figure 5: Revenue function of the uniform distribution, plotted in value space.

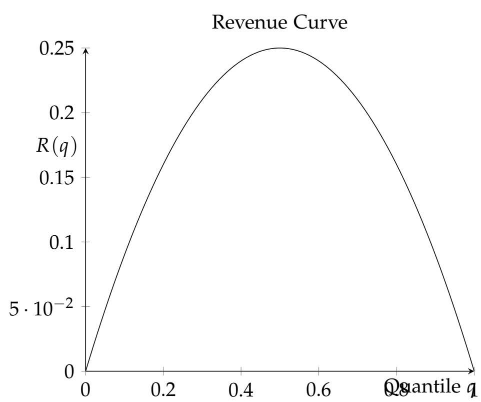

Figure 6: Revenue curve of the uniform distribution, plotted in quantile space.

Now observe the following:

$$\begin{aligned}\frac{\mathrm{d}q}{\mathrm{d}\pi} &= \frac{\mathrm{d}}{\mathrm{d}\pi}\left(1 - F(\pi)\right) \\ &= -f(\pi).\end{aligned}$$

Equivalently, $\mathrm{d}q = -f(\pi)\,\mathrm{d}\pi$. Therefore,

$$\begin{aligned}\int_0^1 R(q)\,\mathrm{d}q &= \int_0^1 F^{-1}(1-q)\,q\,\mathrm{d}q \\ &= -\int_1^0 F^{-1}(1-q)\,q\,\mathrm{d}q \\ &= \int_{\underline{v}}^{\overline{v}} \pi\,(1 - F(\pi))\,f(\pi)\,\mathrm{d}\pi \\ &= \mathop{\mathbb{E}}_{\pi \sim F}\left[R(\pi)\right] \\ &= \mathrm{APX}\end{aligned}$$

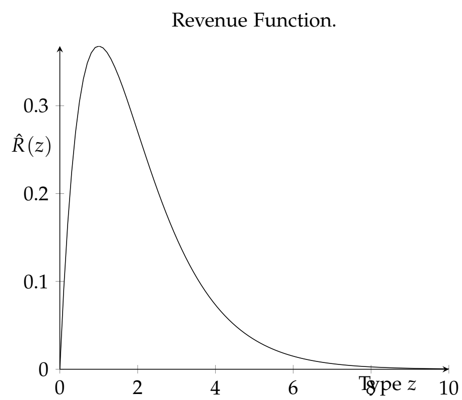

Figure 7: Revenue function of the exponential distribution, $\lambda = 1$, plotted in value space.

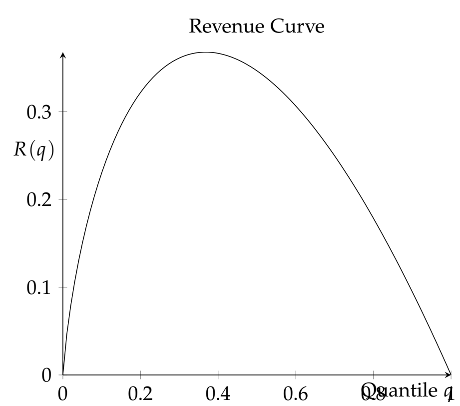

Figure 8: Revenue curve corresponding to the exponential distribution, $\lambda = 1$, plotted in quantile space.

In sum, we have expressed APX in quantile space as expected revenue. It follows that APX is the area under the revenue curve.

6 *Properties of the Revenue Curve*

We have shown that APX is the area under the revenue curve in quantile space. But if the revenue curve is arbitrarily complex, it may be difficult to compute this integral. We now set out to show that the revenue curve cannot be arbitrarily complex; on the contrary, it is always concave, assuming F is regular.

6.1 *Virtual Values*

For starters, we show how virtual values relate to the revenue curve. Specifically, we differentiate the revenue curve R w.r.t. quantile q:

$$\begin{aligned}\frac{\mathrm{d}R(q)}{\mathrm{d}q} &= \frac{\mathrm{d}\left(qF^{-1}(1-q)\right)}{\mathrm{d}q} \\ &= [q]'[F^{-1}(1-q)] + [q][F^{-1}(1-q)]' \\ &= F^{-1}(1-q) + [q][F^{-1}(1-q)]'.\end{aligned}$$

To differentiate the function inverse, we use the chain rule. For a function $f(g(x))$, let $z = f(y)$ and $y = g(x)$. Then:

$$\frac{\mathrm{d}z}{\mathrm{d}x} = \frac{\mathrm{d}z}{\mathrm{d}y}\frac{\mathrm{d}y}{\mathrm{d}x} = \frac{\mathrm{d}f(y)}{\mathrm{d}y}\frac{\mathrm{d}g(x)}{\mathrm{d}x}.$$

Notice that $x = F(F^{-1}(x))$. Taking the derivatives of both sides of this equation, letting $z = F(y)$ and $y = F^{-1}(x)$, yields

$$1 = \frac{\mathrm{d}z}{\mathrm{d}x} = \frac{\mathrm{d}F(y)}{\mathrm{d}y}\frac{\mathrm{d}F^{-1}(x)}{\mathrm{d}x} = f(y)[F^{-1}(x)]'.$$

Rearranging,

$$[F^{-1}(x)]' = \frac{1}{f(y)} = \frac{1}{f(F^{-1}(x))}.$$

Thus,

$$[q][F^{-1}(1-q)]' = \frac{-q}{f(F^{-1}(1-q))}.$$

Now, continuing where we left off,

$$\frac{\mathrm{d}R(q)}{\mathrm{d}q} = F^{-1}(1-q) + \left(\frac{-q}{f(F^{-1}(1-q))}\right).$$

Since $q(v) = 1 - F(v)$ and $v = F^{-1}(1 - q(v))$, we conclude that

$$\begin{aligned}\frac{\mathrm{d}R(q)}{\mathrm{d}q} &= v - \frac{1-F(v)}{f(v)} \\ &= \varphi(v).\end{aligned}$$

Therefore, the derivative of the revenue curve, also called the **marginal revenue**, is the virtual value function!

6.2 *Concavity*

Next, we prove the revenue curve is concave, assuming the distribution F is regular. In other words, we assume the virtual value function in *value* space is non-decreasing, or equivalently, the virtual value function in *quantile* space is non-increasing.

As an example, if values are uniformly distributed on $[0, 1]$, the virtual value function in value space $\varphi(v) = 2v - 1$ is non-decreasing. Since $v = F^{-1}(1 - q) = 1 - q$, the virtual value function in quantile space $\varphi(q) = 1 - 2q$ is non-increasing.

Definition 6.1 (Concave function)**.** A function f is concave if, for any $c \in [0, 1]$,

$$f((1 - c)x + cy) \geq (1 - c)f(x) + cf(y).$$

Equivalently, for an x, y in the domain,

$$f\left(\frac{x + y}{2}\right) \geq \frac{f(x) + f(y)}{2}.$$

Remark 6.2. You can understand concavity by graphing f. Draw a line from point $(x, f(x))$ to $(y, f(y))$. The function f is concave if it lies above the line in the interval $[x, y]$, for all choices of x and y.

Proposition 6.3. *Assuming regularity, the revenue curve is concave.*

Proof. We show that the revenue curve must be concave using a bit of calculus. Let $q_1 \leq q_2$, so that $v(q_1) \geq v(q_2)$. Integrating the virtual value function from quantile q_1 to q_2 yields:

$$\int_{q_1}^{q_2} \varphi(v(q))\, \mathrm{d}q = R(q)\Big|_{q_1}^{q_2} = R(q_2) - R(q_1).$$

It follows that

$$\begin{aligned}\int_{q_1}^{\frac{q_1+q_2}{2}} \varphi(v(q))\, \mathrm{d}q &= R\left(\frac{q_1 + q_2}{2}\right) - R(q_1) \\ \int_{\frac{q_1+q_2}{2}}^{q_2} \varphi(v(q))\, \mathrm{d}q &= R(q_2) - R\left(\frac{q_1 + q_2}{2}\right).\end{aligned}$$

Since the virtual value function is non-increasing in quantile space,

$$\begin{aligned}\int_{q_1}^{\frac{q_1+q_2}{2}} \varphi(v(q))\, \mathrm{d}q &\geq \int_{\frac{q_1+q_2}{2}}^{q_2} \varphi(v(q))\, \mathrm{d}q \\ R\left(\frac{q_1 + q_2}{2}\right) - R(q_1) &\geq R(q_2) - R\left(\frac{q_1 + q_2}{2}\right) \\ R\left(\frac{q_1 + q_2}{2}\right) &\geq \frac{R(q_2) - R(q_1)}{2}.\end{aligned}$$

We conclude that the revenue curve is concave. □

We can also proceed via proof by picture to show that integrating a non-increasing function yields a concave function.

Theorem 6.4. *Let g be a positive, real-valued integrable function defined for all $x \geq a$. Consider a function G defined by $G(x) = \int_a^x g(t)dt$. If g is non-increasing on interval $[a, b]$, then G is concave on that interval.*

Proof by picture. Consider a non-increasing function g on interval $[a, b]$, such as the one depicted in Figure 9. Consider as well an arbitrary point $x_0 \in [a, b]$ and an arbitrary $\delta > 0$.

The value of G at x_0 is equal to the gray area in Figure 9. The black area is the incremental area corresponding to $x_0 + \delta$, and the blue area is the further incremental area corresponding to $x_0 + 2\delta$.

Since f is positive, the value of G at $x_0 + \delta$ (both the gray and the black areas), must exceed the value of G at x_0 (only the gray area); likewise for the value of G at $x_0 + \delta$ relative to the value of G at $x_0 + 2\delta$. So G is increasing. Moreover, since g is non-increasing, the blue area is no larger than the black area. These observations ensure that every line segment joining arbitrary points on G lies entirely below G. So G is concave. (See Figure 10.) □

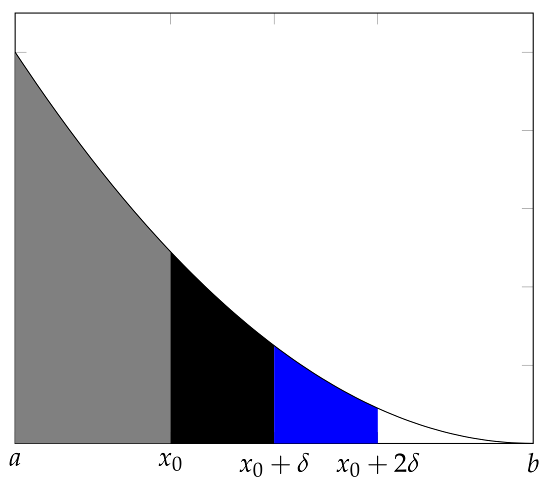

Figure 9: Decreasing function $g(x)$, where $x \in [0, 1]$.

Corollary 6.5. *The integral of the virtual value function in quantile space, namely the revenue curve, is concave.*

7 *Posted-Price Mechanisms*

We now return to our regularly scheduled program: Our goal is to derive an approximation ratio for the simple mechanism, "post a random price," in the single bidder setting, assuming F is regular.

We will analyze this mechanism not by drawing a random price from F, but equivalently, by drawing a random quantile from $U(0, 1)$,

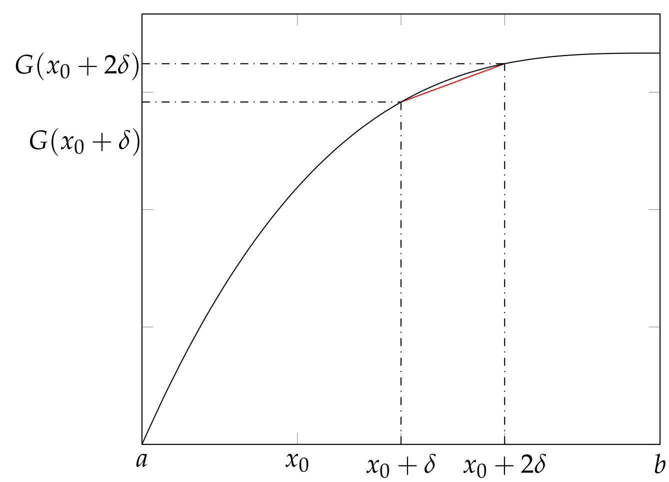

Figure 10: Concave function $G(x) = \int_0^x g(t)dt$, where $x \in [0,1]$.

and posting price $v(q)$. The expected revenue of this mechanism, APX, is the area under the revenue curve in quantile space.

Let q^* be the quantile corresponding to the optimal posted price, π^*: i.e., $\pi^* = F^{-1}(1-q^*)$. The expected revenue in quantile space generated by posting price π^* is OPT $= R(q^*)$. We can depict this quantity by drawing a box of height $R(q^*)$ and width 1, as shown in Figure 11.

OPT upper bounds APX. To lower bound APX, observe that the area under the revenue curve is at least the area of the triangle with vertices $(0,1)$, $(1,0)$, and $(q, R(q^*))$ (see Figure 11). The area of this triangle is half the area of the box, and hence half the value of OPT.

Therefore, posting a price that is simply a random draw from F, yields, in expectation, at least half the total expected revenue of the optimal posted-price mechanism: i.e., APX $\geq \frac{1}{2}$OPT.

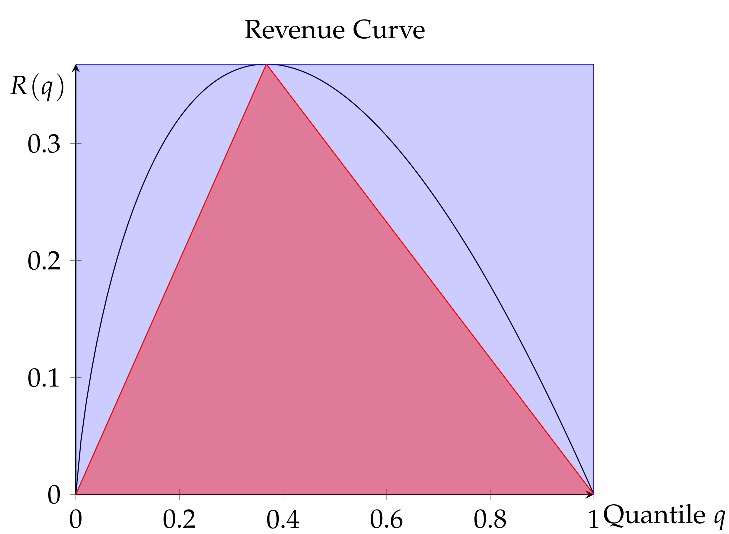

Figure 11: Revenue curve of the exponential distribution, $\lambda = 1$. The blue box represents the expected revenue generated by the optimal auction. The red triangle represents a lower bound on the expected revenue generated by posting a random price drawn from F.

7.1 *Infinite Supply*

Assume an infinite supply of copies of some good: e.g., a digital good, such as an audio or video recording.

Assume further that there are multiple potential bidders for this good, each of whom draws its value from the distribution F (i.e., buyers' values are i.i.d. draws from F). What is the total expected revenue of posting a price randomly drawn from F?

Using our earlier analysis, we can expect to generate at least half the optimal revenue from each individual bidder, and since values are i.i.d., we conclude that this mechanism, in the infinite-supply setting, yields an approximation ratio of 1/2.

7.2 *Prior-Independent Mechanisms*

We end this lecture with a simple modification that yields a prior-independent mechanism: i.e., one that is independent of F.

The modified mechanism is as follows:

1. Collect sealed bids from each bidder.
2. Select a bidder j uniformly at random.
3. Remove bidder j from the mechanism.
4. Set a reserve price to v_j for each bidder $N \setminus \{j\}$.
5. Allocate to every bidder that meets reserve, and charge them v_j.

How well does this mechanism do? Let APX denote the total expected revenue of this mechanism, and let OPT denote the total expected revenue of the optimal mechanism.

The way this mechanism selects a reserve price v_j is equivalent to drawing a random value from the distribution F. Indeed, from the point of view of each bidder in $N \setminus \{j\}$, it remains the case that the reserve price is some randomly sampled value from F. But now only $n-1$ bidders can pay, so the approximation ratio is:

$$\frac{APX}{OPT} \geq \frac{1}{2}\left(\frac{n-1}{n}\right).$$

We can improve the approximation ratio by changing the way we set reserve prices. For example, we can offer bidder j a reserve price equal to the value some other bidder $i \neq j$ submits. Again, from the point of view of bidder j, this reserve price is some randomly sampled value from F. With this modification, bidder j's contribution to total expected revenue is the same as all the other bidders, and we recover the original approximation ratio of 1/2.