ORIE 4154: Pricing and Market Design **Fall 2026**

Lecture 2 — From Values to Demand

Lecturer: Sid Banerjee *Date: August 27, 2026*

2.1 Overview

In Lecture 1 we assumed that every buyer's value was known. This let us ask what prices can accomplish when information is not the issue: divide surplus, ration scarce supply, decentralize an assignment, or correct a congestion externality.

We now remove the complete-information assumption. The seller no longer knows an individual buyer's value, but instead knows a *distribution* of values. Our central question is:

How should we represent demand so that its economically useful structure becomes visible?

The main chain of ideas is

$$\boxed{\begin{array}{c}\text{value distribution} \longrightarrow \text{demand curve} \longrightarrow \text{revenue} \\ \longrightarrow \text{quantile space} \longrightarrow \text{marginal revenue}\end{array}}$$

Along the way we will connect four objects that recur throughout the course:

elasticity, hazard rate, virtual value, choice probability.

Sections marked * are optional reading. They deepen the main development but are not needed for the core chain of ideas.

2.2 From reservation values to a demand curve

Suppose there are N potential buyers. Buyer i has a reservation value drawn from a continuous distribution with CDF F:

$$V_i \sim F,$$

independently across buyers. For now, think of F as known to the seller even though each realization V_i is private to the buyer.

At a posted price p, buyer i purchases if and only if

$$V_i \geq p.$$

Thus the realized number of purchases is

$$Q(p) = \sum_{i=1}^{N} \mathbf{1}\{V_i \geq p\},$$

while expected demand is

$$D(p) = \mathbb{E}[Q(p)] = N\mathbb{P}(V_i \geq p) = N(1 - F(p)). \tag{2.1}$$

This is the simplest bridge from micro-level heterogeneity to an aggregate demand curve.

Remark 2.1 (Expected versus realized demand). *The function $D(p)$ is a model of* expected *demand. If N is finite, actual sales fluctuate around $D(p)$. Writing $q(p) = 1 - F(p)$, independence gives*

$$Q(p) \sim \mathrm{Binomial}(N, q(p)), \qquad \mathrm{sd}(Q(p)) = \sqrt{Nq(p)(1 - q(p))}.$$

For fixed $q(p) \in (0, 1)$, the relative fluctuation satisfies

$$\frac{\mathrm{sd}(Q(p))}{\mathbb{E}[Q(p)]} = \sqrt{\frac{1 - q(p)}{Nq(p)}} = O(N^{-1/2}).$$

This is our first example of statistical averaging: additive uncertainty grows like $\sqrt{N}$, while uncertainty relative to market size vanishes. Later, when F itself is unknown, observing $Q(p)$ after choosing p will be the basic feedback in our learning-to-price problem.

If there is no marginal cost, expected revenue is

$$R(p) = pD(p) = Np(1 - F(p)). \tag{2.2}$$

For much of this lecture we divide by N and work with revenue per potential customer,

$$\bar{R}(p) = p(1 - F(p)).$$

The factor N can always be restored at the end.

2.3 Price space is not always the nicest coordinate system

A natural hope is that $R(p)$ is concave in p, so that the pricing problem is a standard concave maximization problem. This fails even for extremely familiar distributions.

Example 2.2 (Exponential values). *Let*

$$V \sim \mathrm{Exp}(1), \qquad F(p) = 1 - e^{-p}, \quad p \geq 0.$$

Then

$$D(p) = Ne^{-p}$$

and

$$R(p) = Npe^{-p}.$$

Differentiating,

$$R'(p) = Ne^{-p}(1-p),$$

so the optimal price is $p^* = 1$. But

$$R''(p) = Ne^{-p}(p-2).$$

Hence R is concave for $p < 2$ and convex for $p > 2$. Even this very smooth, unimodal pricing problem is not globally concave in price.

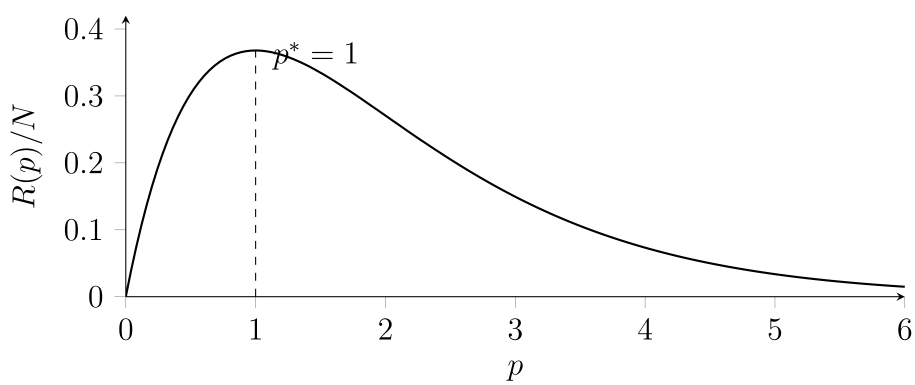

Figure 2.1. For exponential values, revenue pe^{-p} is easy to optimize but is not globally concave in the price p.

This suggests a change of variables.

2.4 Quantile space: choose market share instead of price

Define the *sale probability* or *demand quantile*

$$q = 1 - F(p). \tag{2.3}$$

If F is continuous and strictly increasing, we can invert this relation:

$$p = v(q) := F^{-1}(1 - q), \qquad q \in [0, 1]. \tag{2.4}$$

Instead of thinking of the seller as choosing a price, we can think of the seller as choosing the fraction q of the market she wants to serve. The corresponding per-customer revenue is the *revenue curve*

$$r(q) = qv(q) = qF^{-1}(1 - q). \tag{2.5}$$

Total expected revenue is $Nr(q)$.

Example 2.3 (Exponential values, revisited). *For* $V \sim \mathrm{Exp}(1)$,

$$q = e^{-p}, \qquad p = -\log q.$$

Therefore

$$r(q) = -q \log q.$$

Now

$$r''(q) = -\frac{1}{q} < 0,$$

so revenue is concave in quantile space.

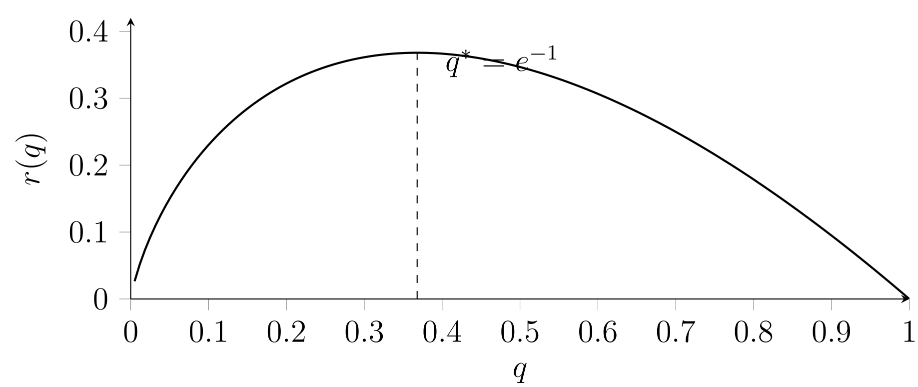

Figure 2.2. The same exponential pricing problem becomes a concave maximization problem in the demand quantile q.

The transformation does not magically make every distribution well behaved. Rather, it identifies the right structural condition under which revenue is concave.

2.5 Marginal revenue and virtual value

Assume F has a positive density f on the region of interest. Since

$$q = 1 - F(v),$$

we have

$$\frac{dq}{dv} = -f(v) \qquad \Longrightarrow \qquad \frac{dv}{dq} = -\frac{1}{f(v)}.$$

Differentiating the revenue curve gives

$$\begin{aligned} r'(q) &= v(q) + qv'(q) && (2.6)\\ &= v - \frac{1-F(v)}{f(v)}. && (2.7) \end{aligned}$$

This motivates an important definition.

Definition 1 (Virtual value / marginal revenue). *For a value distribution F with density f, define*

$$\phi(v) := v - \frac{1-F(v)}{f(v)}. \tag{2.8}$$

Then

$$\boxed{r'(q) = \phi(v(q)).} \tag{2.9}$$

For the moment, there is no auction theory hiding here: $\phi(v)$ is simply the *marginal revenue from increasing the served market quantile*. We will encounter exactly the same function much later when we study optimal auctions.

2.5.1 Revenue regularity

Because $v(q)$ is decreasing in q, the derivative $r'(q) = \phi(v(q))$ is decreasing in q exactly when $\phi(v)$ is increasing in v.

Definition 2 (Regular value distribution). *A distribution F is* regular *if its virtual value*

$$\phi(v) = v - \frac{1-F(v)}{f(v)}$$

is nondecreasing in v.

Equivalently,

$$\boxed{F \text{ regular} \quad \Longleftrightarrow \quad r(q) = qF^{-1}(1-q) \text{ is concave in } q.} \tag{2.10}$$

Remark 2.4 (Terminology). *Talluri–van Ryzin also use the word "regularity" for a collection of smoothness and boundedness assumptions on an aggregate demand function. Here we use* regular distribution *in the mechanism-design sense: monotone virtual value, equivalently concavity of the quantile revenue curve.*

Example 2.5 (Uniform values). *If $V \sim \mathrm{Unif}[0,1]$, then*

$$1 - F(v) = 1 - v, \qquad f(v) = 1,$$

so

$$\phi(v) = 2v - 1.$$

Hence F is regular. In quantile space, $v(q) = 1 - q$ and

$$r(q) = q(1-q),$$

a concave quadratic.

2.6 Hazard rates, virtual values, and elasticity

The *hazard rate* of F is

$$h(v) := \frac{f(v)}{1 - F(v)}. \tag{2.11}$$

Hence virtual value can be written as

$$\phi(v) = v - \frac{1}{h(v)}. \tag{2.12}$$

If the hazard rate is nondecreasing, then $1/h(v)$ is nonincreasing, so $\phi(v)$ is strictly increasing. Thus

$$\boxed{\text{monotone hazard rate (MHR)} \quad \Longrightarrow \quad \text{regular.}} \tag{2.13}$$

The converse need not hold.

Now return to aggregate demand

$$D(p) = N(1 - F(p)).$$

Its (absolute) price elasticity is

$$\epsilon(p) := -\frac{pD'(p)}{D(p)}. \tag{2.14}$$

Since $D'(p) = -Nf(p)$,

$$\boxed{\epsilon(p) = \frac{pf(p)}{1 - F(p)} = ph(p).} \tag{2.15}$$

These objects all encode the same local tradeoff. Differentiating revenue,

$$R'(p) = N\left[(1 - F(p)) - pf(p)\right] \tag{2.16}$$
$$= D(p)\left[1 - \epsilon(p)\right]. \tag{2.17}$$

Thus an interior revenue-maximizing price p^* satisfies

$$\boxed{\epsilon(p^*) = 1.} \tag{2.18}$$

Using $\epsilon(p) = ph(p)$ and (2.12), this is equivalent to

$$\boxed{\phi(p^*) = 0.} \tag{2.19}$$

So, for an interior optimum, the same necessary first-order condition can be read three ways:

$$\boxed{\text{interior pricing FOC} \iff \text{unit elasticity} \iff p^*h(p^*) = 1 \iff \phi(p^*) = 0.}$$

Example 2.6 (Exponential values, one last time). *For $V \sim \mathrm{Exp}(1)$,*

$$h(v) = 1, \qquad \phi(v) = v - 1, \qquad \epsilon(p) = p.$$

All three views give $p^* = 1$ immediately.

Next lecture we will add marginal cost c. The zero-virtual-value condition then becomes $\phi(p) = c$, which is the distributional version of the familiar monopoly markup rule.

2.7 What can go wrong? Mixtures and nonregularity *

Regularity is useful, but it is an assumption rather than a theorem. A particularly natural way it can fail is through *mixtures*: the market may contain several latent customer segments with different value distributions.

Consider a population consisting of two equally likely segments:

$$V \sim \begin{cases} \mathrm{Exp}(1), & \text{with probability } 1/2, \\ \mathrm{Exp}(0.1), & \text{with probability } 1/2. \end{cases}$$

The first group is relatively price sensitive; the second has much larger values on average. The survival function and density are

$$1 - F(v) = \frac{1}{2}e^{-v} + \frac{1}{2}e^{-0.1v}, \tag{2.20}$$
$$f(v) = \frac{1}{2}e^{-v} + 0.05e^{-0.1v}. \tag{2.21}$$

Therefore

$$\phi(v) = v - \frac{e^{-v} + e^{-0.1v}}{e^{-v} + 0.1e^{-0.1v}}. \tag{2.22}$$

Each component distribution is MHR and regular, but the mixture is not. For example,

$$\phi(1) \approx -1.78, \qquad \phi(3) \approx -3.38,$$

so virtual value actually *falls* as value rises over this range.

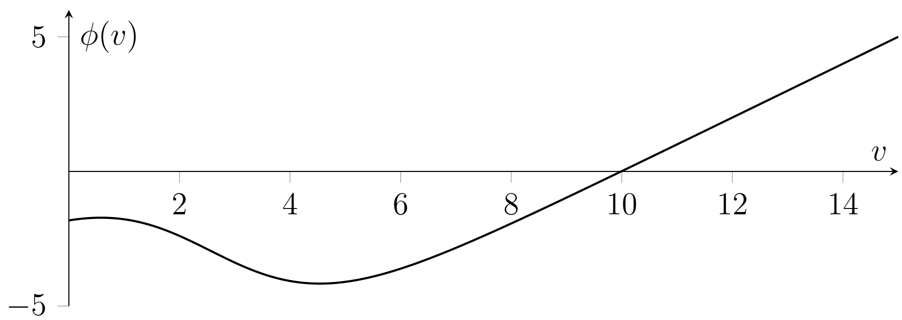

Figure 2.3. Virtual value for the 50–50 mixture of Exp(1) and Exp(0.1). The decrease over an intermediate range means the mixture is nonregular.

The economics are intuitive. As price rises, the composition of the marginal customer changes: low-value customers disappear quickly, leaving a population increasingly dominated by the high-value segment. Mixtures can therefore create nonmonotone marginal-revenue behavior even when each segment is individually well behaved.

In quantile space, this means $r(q)$ is not concave. Later, in optimal mechanism design, we will see a standard repair: replace the revenue curve by its concave envelope, producing *ironed* marginal revenues. For now the main lesson is simply that regularity can fail for economically natural reasons.

2.8 An alternative route to demand: random utility *

The reservation-value model starts from a random willingness to pay V. A second important approach starts from random *utility*.

For a single product, suppose the buyer chooses between buying and an outside option. Let

$$U_1 = a - \beta p + \epsilon_1,$$
$$U_0 = \epsilon_0,$$

where $\beta > 0$. The buyer purchases iff

$$U_1 \geq U_0.$$

If ϵ_0 and ϵ_1 are independent Gumbel random variables with the standard common scale normalization, then

$$\mathbb{P}(\text{buy at } p) = \frac{e^{a-\beta p}}{1 + e^{a-\beta p}}. \tag{2.23}$$

Thus binary logit generates the familiar logistic demand curve

$$D(p) = N \frac{e^{a-\beta p}}{1 + e^{a-\beta p}}. \tag{2.24}$$

2.8.1 Binary random utility is a reservation-value model

The buy decision can be rewritten as

$$a - \beta p + \epsilon_1 \geq \epsilon_0$$

or

$$\frac{a + \epsilon_1 - \epsilon_0}{\beta} \geq p.$$

Define the random variable

$$V := \frac{a + \epsilon_1 - \epsilon_0}{\beta}. \tag{2.25}$$

Then buying is again exactly the event $V \geq p$, and

$$\mathbb{P}(\text{buy at } p) = 1 - F_V(p).$$

Thus, with one product and an outside option, the random-reservation-value and random-utility viewpoints are simply two representations of the same binary choice problem.

Remark 2.7 (Why random utility matters later). *For a single product, random utility adds little beyond a flexible way to generate F. Its real power appears with multiple products. If*

$$U_j = a_j - \beta p_j + \epsilon_j,$$

then the model predicts not only whether a buyer purchases but which product *she switches to when prices or availability change. We will return to this when studying substitution, MNL, and assortment optimization.*

2.9 What to remember

The main mathematical objects from this lecture fit together tightly:

Object	Interpretation
$D(p) = N(1 - F(p))$	Aggregate expected demand generated by heterogeneous reservation values.
$q = 1 - F(p)$	Fraction of the market served at price p.
$r(q) = qF^{-1}(1 - q)$	Revenue when parameterized by market share rather than price.
$\phi(v) = v - (1 - F(v))/f(v)$	Marginal revenue in quantile space: $r'(q) = \phi(v(q))$.
$h(v) = f(v)/(1 - F(v))$	Hazard rate; $\phi(v) = v - 1/h(v)$.
$\epsilon(p) = ph(p)$	Absolute price elasticity of demand.

Under regularity, $r(q)$ is concave. With zero marginal cost, its interior optimum is characterized by the equivalent conditions

$$\boxed{\epsilon(p^*) = 1 \quad \Longleftrightarrow \quad p^* h(p^*) = 1 \quad \Longleftrightarrow \quad \phi(p^*) = 0.}$$

The conceptual progression from Lectures 1–4 is now:

Lecture 1	:	know every buyer's value;
Lecture 2	:	know the distribution of values;
Lecture 3	:	optimize with known demand, then formulate learning;
Lecture 4	:	use optimism to learn while pricing.

Suggested reading

For the reservation-value and aggregate-demand viewpoint, see Talluri and van Ryzin, Chapter 7, especially Sections 7.2.1 and 7.3.1. Their Section 7.2.2 develops the random-utility, binary-logit, and multinomial-logit models. Vohra, Chapter 4, gives complementary pricing and elasticity intuition. The virtual-value interpretation developed here will reappear later in our mechanism-design lectures.