

5. DIFFERENTIALS AND APPLICATIONS OF DERIVATIVES

J.P. Bastos*

EPH 1100: Analytical Tools

Reading: Hoy et al., §5.2 (Differential) and §5.4 (Elasticity, Examples 5.23–5.27).

- **Elasticity:** Examples 5.23–5.27 develop arc and point elasticity; focus on the point elasticity formula and Examples 5.24 and 5.26. Skip Example 5.25: it derives constant elasticity via log-differentiation, a technique this note doesn't use (the Constant-Elasticity Demand subsection below gets the same result directly with the power rule). The book continues past Example 5.27 into L'Hôpital's Rule and an unrelated population-growth example, neither of which is assigned.

1. The Differential

The **derivative** and the **differential** are intrinsically related but distinct concepts that always appear together. Remember that we defined the derivative $f'(x)$ as a *rate*: the slope of the function at x obtained using limits, telling you how fast y changes per unit of x .

The concept of the differential is related, and useful, because it can be used to construct a **linear approximation** of a function near a point. A differential is any number measuring a change: dx is a change in x , and dy is the corresponding change in y along the tangent line at x , distinct from Δy , the actual change along the curve. In this context, you can think of dy as “rise” and dx as “run”.¹

Definition. The **differential** of $y = f(x)$, written dy , is:

$$dy = f'(x) dx$$

where dx is any real number. When dx is small, dy approximates the actual change $\Delta y = f(x + dx) - f(x)$.

The definition says: multiply the derivative $f'(x)$ by any change dx to get dy . When dx is small, dy is a good approximation of Δy ; in the limit as $dx \rightarrow 0$, the approximation becomes exact.

*Assistant Professor of Economics, University of Austin (UATX). jbastos@uaustin.org

¹Keep in mind that “rise over run” is always done along the tangent, not the curve itself. When you learn in school for linear functions, they lie above each other, so this distinction doesn't matter. For higher polynomials, it will be a rough approximation.

Alternatively, dividing both sides by dx (or equivalently, thinking of the limiting case in which dy and dx become infinitesimally small) recovers the derivative:

$$\frac{dy}{dx} = f'(x)$$

As a historical note, Leibniz and Newton each independently discovered calculus in the late 17th century, but each used very different notations. This notation of the derivative *as if* a ratio of differentials is Leibniz's notation and that's how he saw it, even though today we technically obtain the derivative using limits.²

Leibniz's notation survived in many rules because it is so expressive and intuitive (and that's why we like it), and treating dy/dx as a fraction gives correct results in the chain rule, the inverse-function theorem, and separation of variables, even though the derivative is not *technically* a quotient. For our applied purposes, either definition will work.

Why does this matter for applied work?

1. **Approximation without re-evaluating the function.** The exact change in y requires two function evaluations: $\Delta y = f(x + \Delta x) - f(x)$. When f is complicated, that can be slow or analytically messy. The differential uses only $f'(x)$ at one point (one evaluation) and gives an excellent approximation for small Δx .
2. **Keeping track of magnitudes.** In many applications, direction alone is rarely enough: you need to know the size of a response, not just its sign. Writing $dy = f'(x) dx$ keeps both the rate ($f'(x)$) and the magnitude of the input change (dx) visible and explicit in a model. This matters when comparing responses across markets, policies, or units that differ in scale. Because we can write the derivative this way, certain formulas become immediately legible and much more intuitive. Price elasticity of demand, the percentage change in Q for a one-percent change in P , reads directly in Leibniz notation as $(dQ/Q)/(dP/P)$. Section 3 develops this fully.
3. **Approximations in applied empirical work.** This notation is useful because $dy = f'(x) dx$ will give you a *quantity*: given a specific change dx in the input, it gives the approximate corresponding change dy in the output. Formally, the approximation is:

$$\Delta y \approx f'(x) \Delta x$$

It implies that if you have good data on y and good data on x , a regression of y on x will give

²Hoy et al. (2022, §5.2, p.153) specifically say to avoid this mistake, arguing that technically when we write $\frac{dy}{dx} = f'(x)$ we are writing $\frac{d}{dx}(y)$ where $\frac{d}{dx}(\cdot)$ defines an operator for “taking a limit”, whereas “ $dy = f'(x)dx$ is a procedure for linear approximation,” warning that “the two are closely related but distinct.” But it was only in the 19th century that mathematicians reconstructed differential calculus using limits and the derivative became $\lim_{h \rightarrow 0} \frac{f(x+h) - f(x)}{h}$, formally a limit rather than a ratio. So, for our purposes, taking them as interchangeable in practice is acceptable, but keep that distinction in mind.

you an approximation of the slope!³ Moreover, if you take the log of both sides and regress $\ln Y$ on $\ln X$, the coefficient of that regression, β , will approximate the elasticity!⁴

Example: A firm currently produces $q = 300$ units with $C(q) = q^2$. How much does cost rise if output increases by 1 unit?

Exact: $\Delta C = C(301) - C(300) = 90,601 - 90,000 = 601$.

Approximation: $dC = C'(300) \cdot 1 = 2(300) \cdot 1 = 600$.

The differential gives \$600 vs. the exact \$601, a 0.17% error. For small changes, the linear approximation is excellent.

2. Marginal Analysis: Putting the Rules to Work

The derivative is the mathematical formalization of the economist's concept of *marginal*: the change in one quantity for a small (in the limit, infinitesimal) change in another.

Marginal cost

Marginal cost is the cost of producing one additional unit. If the firm is currently making q units and wants to make one more, how much does that extra unit add to the total bill? That is what $MC(q)$ measures: not the average cost per unit, but the cost of the next unit specifically.

Why does this depend on the form of the cost function? Consider two cases.

Case 1: Linear cost. Suppose $C(q) = 10q + 100$, where 100 is a fixed cost (rent, equipment) that does not change with output. Then:

$$MC(q) = C'(q) = 10$$

Marginal cost is constant at \$10 regardless of how many units the firm is already producing. Whether you are the 1st unit or the 1,000th, the next unit always costs \$10. This reflects a situation where inputs can be acquired at constant prices with no bottleneck.

Case 2: Quadratic cost. Now suppose $C(q) = q^2 + 10q + 100$. Then:

$$MC(q) = C'(q) = 2q + 10$$

What does this tell us?

³This is a linearization, and economists linearize everything. A linear regression estimates the best-fit line through the data; when the true relationship f is nonlinear, the slope coefficient approximates the average marginal effect over the sample range. For the log-log case, the approximation is exact when the true form is $Y = \alpha X^\beta$ (constant elasticity) and approximate otherwise.

⁴Remember from Class 2 that log-log models approximate percentage changes, and notice that this is the same tangent-line approximation from Class 3: instead of following the curve, you follow the tangent. The approximation improves as $\Delta x \rightarrow 0$, because the tangent is the limit of secant lines.

- At $q = 0$: $MC = 10$. The first unit costs approximately \$10.
- At $q = 5$: $MC = 20$. The 6th unit costs approximately \$20.
- At $q = 10$: $MC = 30$. The 11th unit costs approximately \$30.⁵

Here the cost of the next unit *does* depend on the current level of production. The firm faces rising marginal cost: each additional unit is more expensive than the last. This might happen because inputs become scarcer as output grows (workers start getting in each other's way, raw materials must be sourced from more distant suppliers, and so on). We will return to what drives the shape of the cost curve when we study second derivatives in Class 6.

Marginal revenue

Marginal revenue is the additional revenue from selling one more unit. A **competitive firm** is a price taker: at the going market price, it can sell as many units as it can produce without moving that price at all. Every additional unit therefore brings in exactly P more revenue, so marginal revenue simply equals price: $MR = P$. A **monopolist** faces a different problem entirely. As the only seller in the market, it faces the entire downward-sloping demand curve: to sell more units, it must lower the price for everyone. This creates a tension. On one hand, selling at a higher price brings in more revenue per unit. On the other hand, a higher price means fewer units sold. Understanding which quantity to sell requires knowing how revenue changes as output changes, and that is what marginal revenue measures.

Suppose the monopolist faces inverse demand $P(q) = 60 - q$. Total revenue is:

$$TR(q) = P(q) \cdot q = (60 - q)q = 60q - q^2$$

Applying the differentiation rules from Class 4:

$$MR(q) = TR'(q) = 60 - 2q$$

Notice that $MR < P$ for any $q > 0$. Concrete numbers make the reason exact. Suppose the monopolist raises its price from $P = \$30$ (selling $q = 30$ units) to $P = \$45$ (selling $q = 15$ units); panel (a) of Figure 1 shows this move along the demand curve. Two things happen simultaneously: (1) on the 15 units it still sells, it now collects \$15 more per unit, a gain of $15 \times 15 = \$225$; but (2) it no longer sells the 15 units between $q = 15$ and $q = 30$, units that used to bring in \$30 each, a loss of $15 \times 30 = \$450$. The net change in revenue is $225 - 450 = -\$225$. This gain-versus-loss tension is exactly why marginal revenue lies below price at every quantity: part of any gain on one margin is always offset by a loss on the other. Panel (b) shows the same logic in its continuous, calculus form; notice that panel (a)'s two rectangles meet exactly at the point $(q, \$) = (15, 30) = (q^*, MR(q^*))$, the height where the gain and loss margins balance is precisely marginal revenue.

⁵Remember that to get exact values for the q -th unit you have to calculate $C(q) - C(q - 1)$.




Figure 1: Marginal revenue: the price-quantity trade-off and its calculus.

Notes: Demand is $P = 60 - q$. Panel (a): raising price from A ($q = 30, P = 30$) to B ($q = 15, P = 45$) gains \$225 on the 15 units still sold (green) but loses \$450 from the 15 units no longer sold (red); the two rectangles meet exactly at $(q, \$) = (15, 30)$, which is exactly where the MR curve crosses that height. Panel (b): marginal revenue is $MR = 60 - 2q$, twice as steep as demand with the same vertical intercept; at any quantity $q^* > 0$, MR lies strictly below demand. The point where panel (a)'s rectangles touch is exactly the height of the MR curve at q^* : $MR(15) = 30$.

3. Elasticity

Having the derivative in hand, we can now define one of the most important concepts in applied economics: **elasticity**. Elasticity measures how responsive one variable is to a change in another, in percentage terms. Economists use several versions of it: how much quantity demanded responds to a price change (price elasticity of demand), how much it responds to a change in income (income elasticity), how much quantity supplied responds to a price change (elasticity of supply), and so on. The version we focus on here, and the most commonly used, is **price elasticity of demand**: the percentage change in quantity demanded for a percentage change in price.

In your principles course you computed this between two points on a demand curve. One way to do it uses a base price and quantity in the denominator:

$$\varepsilon_{arc} = \frac{\Delta Q/Q_0}{\Delta P/P_0}$$

where P_0, Q_0 are the price and quantity you are measuring the change from. The problem: this calculation gives a different answer depending on which point you treat as the base, because the denominator changes. Take a linear demand $Q = 60 - 2P$, moving between a price of \$10 and a price of \$20:

- **From A to B** (low price to high), using A as the base: $P_A = 10, Q_A = 40; P_B = 20, Q_B = 20$, so $\Delta P = 10, \Delta Q = -20$.

$$\varepsilon_{arc} = \frac{-20/40}{10/10} = -0.5$$

- **From B to A** (high price to low), using B as the base: $\Delta P = -10$, $\Delta Q = 20$.

$$\varepsilon_{arc} = \frac{20/20}{-10/20} = -2$$

Same two points, very different numbers. One way around this, the **midpoint rule**, is to stop picking either endpoint and instead average the two to get $\bar{P} = 15$, $\bar{Q} = 30$:

$$\varepsilon_{mid} = \frac{\Delta Q / \bar{Q}}{\Delta P / \bar{P}} = \frac{-20/30}{10/15} = -1$$

This gives a single number, but it is still only an approximation over an interval: it just averages away the ambiguity rather than removing it.




Figure 2: The arc-elasticity problem on a linear demand curve.

Notes: Demand is $Q = 60 - 2P$ (plotted with P on the vertical axis). Points A and B are connected by a right-angle triangle showing the slope $\Delta Q / \Delta P = -2$. The slope is the same everywhere on this line, but the price-to-quantity ratio P/Q differs: at A , $P/Q = 10/40 = 0.25$; at B , $P/Q = 20/20 = 1$. Arc elasticity therefore gives -0.5 or -2 depending on the baseline. Point M is the midpoint ($\bar{P} = 15$, $\bar{Q} = 30$); the midpoint rule gives $\varepsilon_{mid} = \frac{-20/30}{10/15} = -1$.

Note that the point elasticity evaluated at M also equals -1 : $(dQ/dP) \cdot P/Q = -2 \times (15/30) = -1$. The two coincide at the midpoint, which is why the midpoint rule is a useful approximation.

From the arc formula to point elasticity

The baseline problem has a direct fix: replace the discrete ratio $\Delta Q / \Delta P$, which needs two points and therefore a baseline, with the derivative dQ/dP , evaluated at a single point. There is no interval left, so there is nothing left to serve as a baseline; we formalize this below as **point elasticity**.

But a second, separate problem remains: the slope itself, whether written as $\Delta Q / \Delta P$ or dQ/dP , depends on the unit of measurement. For example, demand for steel measured in tonnes has a very different slope than the same demand measured in kilograms, even though the underlying

behavior is identical. To compare across goods, markets, or time periods, we need a *unit-free* measure. Multiplying the slope dQ/dP by P/Q fixes this second problem: P shares units with the denominator of dQ/dP (price) and Q shares units with the numerator (quantity), so the ratio P/Q cancels all units out.

Point elasticity

The **price elasticity of demand** at a point is:

$$\varepsilon = \frac{dQ}{dP} \cdot \frac{P}{Q}$$

This is the percentage change in quantity demanded for a 1% increase in price.⁶

The formula multiplies the slope dQ/dP by the price-to-quantity ratio P/Q , which adjusts for units. The result is dimensionless.

- $\varepsilon < -1$: **elastic** demand. A 1% price increase reduces quantity by more than 1%. Total revenue falls when price rises.
- $-1 < \varepsilon < 0$: **inelastic** demand. A 1% price increase reduces quantity by less than 1%. Total revenue rises when price rises.
- $\varepsilon = -1$: **unit elastic**. Revenue is maximized at this point (where $MR = 0$).

Example: Suppose $Q = 50 - 2P$. Then $dQ/dP = -2$.

$$\varepsilon = \frac{dQ}{dP} \cdot \frac{P}{Q} = -2 \cdot \frac{P}{Q} = \frac{-2P}{50 - 2P}$$

Note that ε varies along the demand curve even though the slope is constant. At $P = 10$: $Q = 30$, $\varepsilon = -20/30 = -2/3$ (between -1 and 0 : inelastic). At $P = 20$: $Q = 10$, $\varepsilon = -40/10 = -4$ (below -1 : elastic). At $P = 12.5$: $Q = 25$, $\varepsilon = -25/25 = -1$ (unit elastic, midpoint).

This is a key result: for linear demand, elasticity is not constant. It equals -1 at the midpoint of the demand curve, sits between -1 and 0 (inelastic) in the lower-price half, and falls below -1 (elastic) in the upper-price half.

⁶Since demand curves are downward-sloping, $dQ/dP < 0$, so ε itself comes out negative. This reflects the fact that price and quantity are inversely related: a price *increase* is associated with a quantity *decrease*, and vice versa. This is not a sign we add to the formula, but rather the formula reporting the inverse relationship directly. Be aware that some textbooks instead report elasticity as a positive number by placing a minus sign in front of the formula. I prefer to avoid that, because the sign is economically meaningful. When someone says “elasticity is -2 ,” they mean a 1% price increase reduces quantity demanded by 2%. Thus, for our purposes, if you get a positive sign for the elasticity of demand, check your algebra!




Figure 3: Elastic and inelastic ranges of a linear demand curve.

Notes: Demand is $Q = 60 - 2P$, the same curve as Figure 2. Above the midpoint M ($P = 15$), demand is elastic ($\varepsilon < -1$); below M , demand is inelastic ($-1 < \varepsilon < 0$); at M itself, demand is unit elastic ($\varepsilon = -1$). This is a general property of linear demand: elasticity falls (becomes more negative) continuously as price rises, crossing -1 exactly at the midpoint.

Constant-elasticity demand

A different functional form, $Q = \alpha P^{-\beta}$ ($\alpha, \beta > 0$), has the same elasticity at every point on the curve. Let's derive that step by step.

Step 1: Take the derivative. Using the power rule on $Q = \alpha P^{-\beta}$:

$$\frac{dQ}{dP} = \alpha \cdot (-\beta) P^{-\beta-1} = -\alpha\beta P^{-\beta-1}$$

Step 2: Plug into the elasticity formula. Substitute $Q = \alpha P^{-\beta}$ directly into the denominator:

$$\varepsilon = \frac{dQ}{dP} \cdot \frac{P}{Q} = \left(-\alpha\beta P^{-\beta-1}\right) \cdot \frac{P}{\alpha P^{-\beta}}$$

Step 3: Cancel. With a bit of rearranging, we can rewrite the expression $-\alpha\beta P^{-\beta-1}$ to isolate $\alpha P^{-\beta}$ (notice this is Q). Since $P^{-\beta-1} = P^{-\beta} \cdot P^{-1}$:

$$-\alpha\beta P^{-\beta-1} = -\beta \cdot \left(\alpha P^{-\beta}\right) \cdot P^{-1}$$

Substituting this into Step 2's expression:

$$\varepsilon = \left[-\beta \cdot \left(\alpha P^{-\beta}\right) \cdot P^{-1}\right] \cdot \frac{P}{\alpha P^{-\beta}} = -\beta \cdot \frac{\alpha P^{-\beta}}{P} \cdot \frac{P}{\alpha P^{-\beta}} = -\beta$$

The $\alpha P^{-\beta}$ terms cancel against each other (numerator against denominator), and so do the two

lone powers of P , leaving:

$$\varepsilon = -\beta$$

The elasticity is $-\beta$ at every price level: constant, and negative, as expected for a downward-sloping demand curve. This functional form is useful precisely because $|\varepsilon| = \beta$ is a clean parameter rather than a function of the point on the curve, unlike the linear-demand case.

Practice Problems

1. A firm faces inverse demand $P(q) = 80 - 2q$.
 - (a) Find total revenue $TR(q)$ and marginal revenue $MR(q)$.
 - (b) Verify that $MR(q) < P(q)$ for all $q > 0$, and explain why in words using the gain-vs-loss logic from Figure 1.
 - (c) Plot D and MR on the same axes, and mark the quantity where $MR = 0$.
2. A firm's total cost is $C(q) = (2q + 5)^2$ dollars.
 - (a) Use the chain rule to find $MC(q)$.
 - (b) Evaluate $MC(10)$.
 - (c) Use the differential to approximate the cost of raising output from $q = 10$ to $q = 11$. Compute the exact cost change and the percentage error.
3. A monopolist faces demand $P(q) = 140 - 4q$ and has cost $C(q) = 20q + 60$.
 - (a) Find $MR(q)$ and $MC(q)$.
 - (b) Find the profit-maximizing quantity q^* by setting $MR(q^*) = MC(q^*)$.
 - (c) What price does the firm charge at q^* ?
4. A firm's total cost is $C(q) = q^2 + 10q + 25$.
 - (a) Find average cost $AC(q) = C(q)/q$.
 - (b) Use the quotient rule to find $AC'(q)$.
 - (c) Solve $AC'(q) = 0$ for q . (This turns out to be the quantity where average cost is as low as possible.) Verify that $MC(q) = AC(q)$ at this quantity.
5. A firm currently produces 100 units with cost function $C(q) = 0.5q^2 + 10q + 200$.
 - (a) Find $MC(q)$ and evaluate $MC(100)$.
 - (b) Use the differential to estimate the increase in cost from producing 5 additional units.
 - (c) Compute the exact increase. What is the percentage error?
6. The demand function is $Q = 120 - 3P$.

- (a) Compute the price elasticity of demand ε as a function of P .
 - (b) At what price is $\varepsilon = -1$? What is happening to total revenue at this price?
 - (c) At $P = 30$, is demand elastic or inelastic? What happens to total revenue if price rises slightly?
7. The **income elasticity of demand** measures how demand responds to a change in income Y : $\eta = \frac{dQ}{dY} \cdot \frac{Y}{Q}$. Suppose demand for a good is $Q = 0.02 Y^{1.5}$, where Y is income in thousands of dollars.
- (a) Find dQ/dY and the income elasticity η as a function of Y , and interpret it.
 - (b) Is η constant, or does it vary with Y ? Compare the structure of this problem to the constant-elasticity demand subsection above.
 - (c) Economists call a good **normal** if $\eta > 0$ (demand rises with income) and **inferior** if $\eta < 0$ (demand falls as income rises). Which category does this good fall into?

References

Hoy, M., Livernois, J., McKenna, C., Rees, R., and Stengos, T. (2022). *Mathematics for Economics*. MIT Press, 4th edition.