Perfect Competition

I. What is a perfectly competitive market?

The remainder of the class will focus primarily on analyzing four different market structures: (1) perfect competition, (2) monopoly, (3) monopolistic competition, and (4) oligopoly. For now we will focus on the first two market structures, which are at the extremes of a continuum of market structures. The latter two market structures fill in this continuum with monopolistic competition being close to perfect competition and oligopoly being close to monopoly.

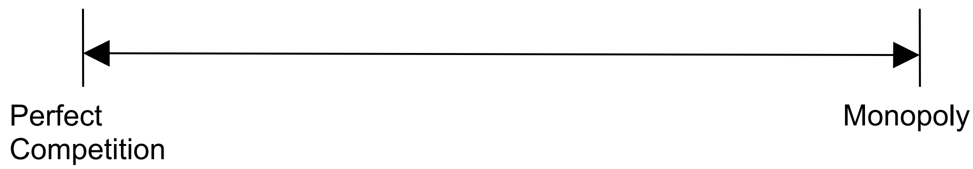

Let us begin by defining both perfect competition and monopoly:

Perfect Competition	Monopoly
1. Lots of buyers and sellers	1. Lots of buyers only one seller
2. All firms are small relative to the market.	2. Single firm is the market.
3. Homogenous product.	3. Homogenous product.
4. Free entry and exit from market.	4. Barriers to entry.

A comparison of the characteristics of both market structures illustrates that perfect competition and monopoly are, in fact, polar opposites. Now consider the implications of the characteristics for perfect competition upon both the perfectly competitive firm and market

Consider first the market, which is similar to the markets we have already discussed in previous chapters. That is, market demand has similar characteristics as any other market demand curve. Market demand represents the sum of the demands for all the individual consumers in the market. Thus, it is the individual consumers who determine their own demand and, hence, market demand. Their tastes/preferences for the good, income, the number of consumers in the market, prices of related goods, and expectations about future prices all combine to determine the exact characteristics of market demand. However, market demand is downward sloping as shown in Figure 1.

Similarly, market supply is also shown in Figure 1 and, again, supply is determined by all of the firms in the market, is upward sloping, and depends upon costs of production, prices of related goods, and expectations about future prices. As shown in Figure 1, the interaction of market demand and supply set the equilibrium market price, P_M.

The first panel of figure 1 shows the representative firm's position given the market price. Perfectly competitive firms are <u>price takers</u>. Because there are so many firms in the market all producing exactly the same product any individual firm cannot affect the market price. If the firm attempts to raise their price above the market price, then no consumer would buy their product. Hence, price takers simply take the market price as given and decide how much to produce. This means that the perfectly competitive firm has a demand curve that is perfectly elastic at the market price as shown in Figure 1.

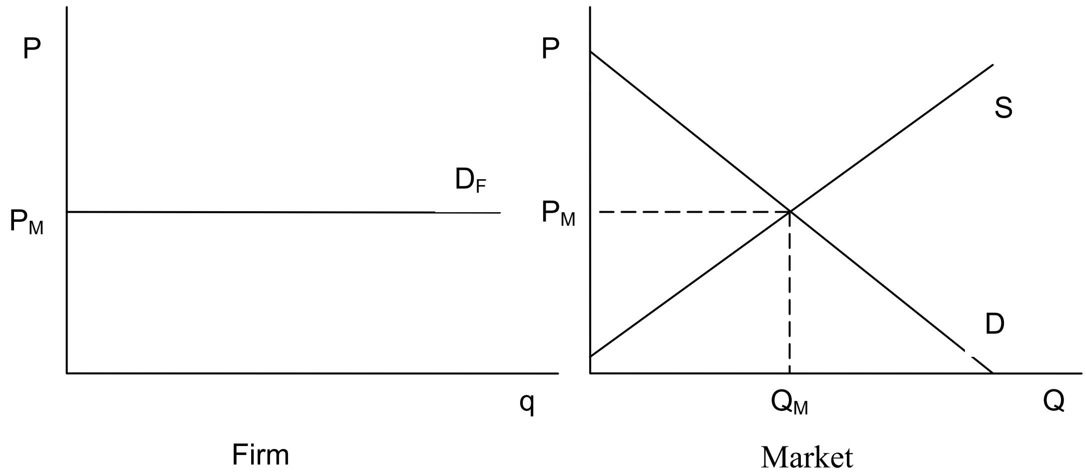

Figure 1

II. Short-run Pricing and Output

- How do Firms maximize profits (π) in the short-run?

  The most straightforward method of answering this question is to simply examine what we know about profit:

  - o π = TR – TC
    - Total Revenue (TR) = Price * Quantity = P*Q
    - Total Cost (TC) = Average Total Cost * Quantity = P*ATC
    - Hence, π = P*Q – P*ATC = (P – ATC)Q

  We will use two approaches to show perfectly competitive firms maximize profits in the short-run. The first one (TC/TR) will be intuitive but is less useful. The second one (MC/MR) is less intuitive but will, in fact, be the most useful approach.

  - o Total Cost and Total Revenue approach

    Figure 2 illustrates what we know, graphically about TR and TC. In the previous chapter, we derived TC and it is again shown in Figure 1. Recall, that TC first is relatively steep at first and then flattens out before, again, getting steeper. The shape of TC simply reflects the law of diminishing returns.

    Total Revenue is a much easier function to graph as it simply equals price times quantity. Thus, when quantity is zero, total revenue is also zero. Thereafter, with a constant price (given by the interaction of D and S in the market), total revenue rises by a constant amount (the price) for each unit of the good sold.

In Figure 2, areas A and C show where TC exceeds TR and, therefore, the firm is operating at a loss. It is only in area B where TR is greater than TC and, therefore, the firm is generating a positive profit.

However, although profit is only positive in area B, where is profit maximized? Given that profit equals the difference between TR and TC, maximum profit occurs where the distance between the two curves is at its largest, which occurs at q^*.

Figure 3 illustrates the profit function that is associated with Figure 2. Again, in areas A and C, profit is negative while only in area B is profit positive. Likewise, profit is maximized at q^*, which occurs at the same level of output in both figures.

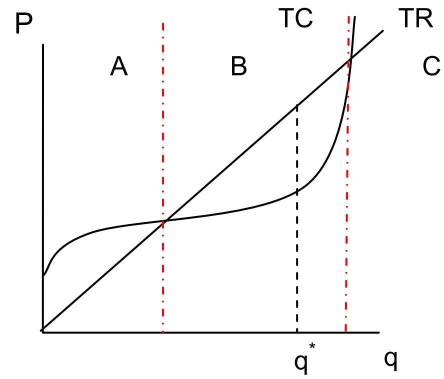

Figure 2

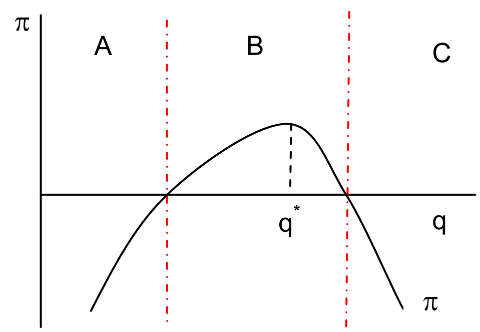

Figure 3

- o Marginal Cost and Marginal Revenue approach

  Just as we had already derived Total Cost, the same is true for marginal cost.

  Marginal Cost = the extra cost of producing one more unit of output = Δ TC / Δ q.

  Similarly:

  Marginal Revenue = the extra <u>revenue</u> of producing one more unit of output = Δ TR / Δ q.

  We already know that MC is u-shaped because of the law of diminishing returns. What does MR look like for a perfectly competitive firm? Recall that a perfectly competitive firm is a <u>price taker</u>, as shown above in Figure 1. The firm takes the market price as given and then decides how much output to produce. Thus, every time the firm produces another unit of output they can <u>always</u> sell it for the market price. In other words, their demand equals the market price which, in turn, equals the marginal revenue.

Figure 4 illustrates the MC/MR approach to profit maximization. The dotted line in Figure 4 occurs where MC = MR.

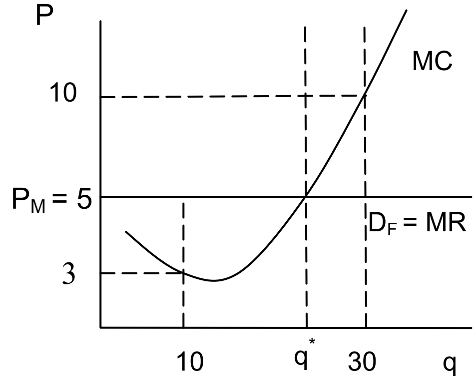

Figure 4

To the left of the dotted line MR exceeds MC. For example, consider whether the firm should produce the tenth unit of output. At this level of output, MR equals 5 (as it always does) while MC equals 3. Should the firm produce the tenth unit when it costs $3 to produce it and they get $5 in revenue? Clearly, the firm's profit will rise by $2 (5 – 3) if they do produce the tenth unit. As a result a profit maximizing unit will always produce such a unit. This gives us our first insight into profit maximization:

- A profit maximizing firm will always increase output as long as MR > MC.

Consider now a firm that is producing to the right of the dotted line where MR = MC. To the right, MC is greater than MR. For example, at the 30th unit of output MC equals $10 while MR is still $5. Clearly the firm would lose profit if it chose to produce this unit of output because it would cost $10 to produce it while only gaining $5 in revenue. Our second insight into profit maximization suggests that:

- A profit maximizing firm will always decrease output whenever MC > MR.

Notice that if the profit maximizing firm will increase output when MR > MC and decrease output whenever MC > MR that there is only one place where the firm can end up eventually:

- All profit maximizing firms will produce where MR = MC.

It is important to understand that both the TC/TR and MC/MR methods of finding the profit maximizing output will yield the same answer. In fact, neither method would be useful if it was not consistent with other methods.

- Graphical Analysis – Questions to be answered using Figure 5.

1. At what q is π maximized?

   Recall from above that all profit maximizing firms produce at the quantity where MR = MC. In Figure 5 this occurs at q_1.

2. What P does a firm charge?

   Recall that a perfectly competitive firm is a price taker that takes the market price as given. Hence, in Figure 5, the market price is P = MR = D.

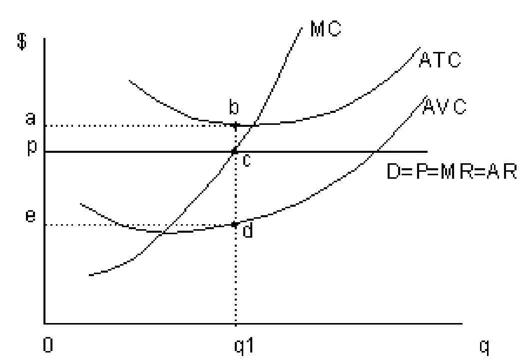

Figure 5

3. What does π equal?

   Obviously firms are primarily interested in their profits. However, in order to calculate profits in Figure 5, there are a number of other questions that must be answered before profit can be calculated. All of these are calculated at the profit maximizing output, q_1.

   - $Q = q_1$ = distance 0 to q_1.
   - P = MR = MC = distance 0 to p.
   - The average costs:
     - ATC = distance 0 to a.
     - AVC = distance 0 to e.
     - AFC = ATC – AVC = distance e to a.
   - Recall that Total Revenue equals price times quantity. In graph this is the multiplication of two distances making a right angle: Price – distance 0 to p and quantity – distance 0 to q_1. Graphically, multiplying two sides of a right angle together gives the area of the resultant rectangle defined by the points, 0 – p – c – q_1. Thus TR = the area of the rectangle 0–p–c–q_1.
   - Similarly for the three total costs:
     - ATC = TC/q. Hence, TC = ATC*q = the area of rectangle 0–a–b–q_1.
     - AVC = TVC/q. Hence, TVC = AVC*q = the area of rectangle 0–e–d–q_1.
     - ACF = TFC/q. Hence, TFC = AFC*q = the area of rectangle e–a–b–d.
   - And, finally, profit. There are two methods to calculate profit that both yield the same answer:
     - π = TR – TC = the area of the rectangle 0–p–c–q_1 minus the area of rectangle 0–a–b–q_1 = the area of rectangle p-a-b-c.
     - π = TR – TC = p*q – ATC*q = (p – ATC)q. This equation is the most useful because it is clear now that in Figure 5, profit is negative (because p < ATC). P – ATC = distance a – p while q = distance p – c, which again yields the area of rectangle p-a-b-c.
   - Students should know all of the above formulas and graphical calculations.

- Loss Minimization and the Short-Run Shutdown Point

  When considering what profit equals for a given firm there are clearly three possibilities:

  - $\pi > 0$, which occurs whenever P > ATC. In this case firms clearly continue to produce.
  - $\pi = 0$, which occurs whenever P = ATC. In this case firms will still continue to produce (zero economic profit implies that all costs, including opportunity costs are covered.)
  - $\pi < 0$, which occurs whenever P < ATC. Should the firm shutdown in this situation?

    Figure 5 will be used to illustrate the issue here. For illustration we will assume that:

    - $q_1 = 1000$.
    - P = 5
    - At q_1 ATC (a) = 7; AVC (e) = 3, which implies that AFC = ATC – AVC = 4.
    - π = TR – TC = p*q – ATC*q = 5000 – 7000 = - 2000.

    In this case, it turns out that it is best for the firm to continue to produce. This is true because the firm is in the short-run, with fixed costs that must still be paid when the firm shuts down. Calculate TFC:

    - TFC = AFC*q = 4*1000 = 4000.

    Thus, if the firm produces it loses $2,000 and if it shuts down it loses its TFC or $4,000. Clearly, the firm is better off to produce and not shut down.

    How low must the price go before the firm will shut down?

    - Essentially, the key concept here is for the firm to only shut down if it cannot cover those costs that it can avoid by shutting down – the total variable costs. Thus, firms should continue to produce as long as their revenue (as it does in the example above) covers the variable costs. Notice that fixed costs are irrelevant. A common mistake is to incorrectly think that one must continue to produce <u>because</u> of fixed costs.
    - To prove this algebraically consider the breakeven point where profits if the firm produces equals profits if they shutdown:
      - TR – TC = – TFC or
      - TR – TVC – TFC = – TFC or
      - TR – TVC = 0 (Notice that TFC cancels out because the firm must pay them whether producing or not) or
      - TR = TVC, thus we breakeven when total revenue equals total variable costs. Continuing to simplify the equation:
      - P*q = AVC*q or

- P = AVC, thus the firm should shutdown if price falls below AVC. This is called the short-run shutdown point.

- What is the firm's short-run supply curve?

  A supply curve shows quantity supplied given the price, ceteris paribus. Figure 6 shows the MC and AVC for a perfectly competitive firm. If the price in the market equals P_m then the firm will produce q_1, where P = MR = MC or where P = MC.

  Likewise, if the price rises to P_1 then the firm produces q_2, where the new price equals the MC curve. Notice in both cases that it is the MC curve that equates price and quantity supplied. Thus, the MC curve is the firm's supply curve.

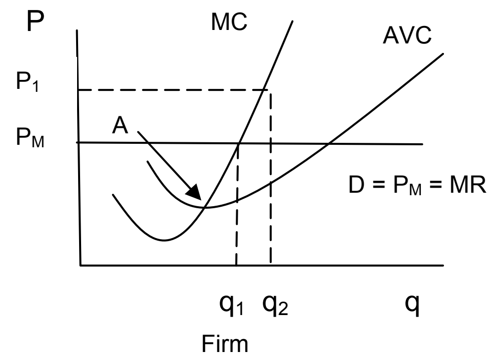

Figure 6

  The only exception would be if the price fell below AVC. In that case, recall that the firm will shut-down in the short-run. Point A in Figure 6, the point where MC crosses AVC, shows the point at which price would equal AVC and the firm would shut-down.

  Recall that the law of supply requires an upward sloping supply curve, when price increases so does quantity supplied. Since the supply curve equals the firm's MC above AVC and MC is upward sloping, supply curves are upward sloping because so is MC. And as we learned before, MC is upward sloping because of the law of diminishing returns.

- What is the industry/market short-run supply curve?

  This is a relatively straightforward concept. If at a price of $10 a perfectly competitive firm produces 4,000 units of output and there are 1,000 similar firms, then total industry output is 4 million. Thus, the industry supply curve is simply the sum of all the firm supply curves. Since the firm supply curves are their MC curves (above AVC), then the industry short-run supply curve is also a measure of industry MC.

III. Long-run Pricing and Output

- The long-run equilibrium in perfectly competitive markets:
  - if $\pi > 0$, the positive profits attract firms to the market in the long-run
    - More firms increases market supply
    - Increased supply decreases price
    - This process continues until $\pi = 0$
  - If $\pi < 0$, the negative profits cause some firms to go out of business / leave industry in the long-run
    - Fewer firms decreases supply
    - Decreased supply increases price
    - This process continues until $\pi = 0$
  - Thus in the long-run π must equal 0

- Market Shocks – What is the Long-run Market Supply?

  Now consider what happens in the long-run with a shock in the market, such as an increase or a decrease in market demand. Figure 7 shows the long-run equilibrium for a perfectly competitive market and firm. Profit is zero because price equals ATC.

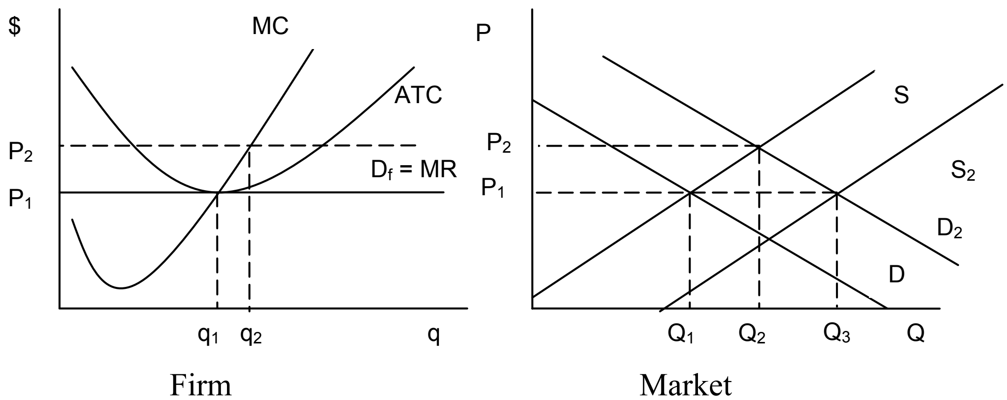

Figure 7

However, consider what happens in both the short-run and the long-run as market D increases or decreases. Initially market price equals P_1 while firm output is q_1 and market output is Q_1. Now suppose that market D increases to D_2. What happens in both the short and the long run as a result?

Short-run Impact	Long-run Impact
P↑ (to P_2)	With $\pi > 0$ => entry occurs and the short run S shifts right (to S_2)
q ↑ (to q_2)	P↓ (back to P_1)
Q ↑ (to Q_2)	q ↓ (back to q_1) while Q ↑ (to Q_3). How does total Q increase while firm q decreases?
π ↑ (was 0 now is > 0)	π ↓ (back to zero)

Just the opposite in both the short-run and long-run would happen if D were to decrease rather than increase. Thus, in the long-run profit is always equal to zero which is caused by either entry (if D increases) or exit (if D decreases). Notice that, in the long-run price remains constant at P_1, which must happen in order to keep profit equal to zero. The only change in the long-run with an increase or decrease in market D is that total output rises or falls, respectively.

This implies that the long-run market supply curve is perfectly elastic at the market price that yields a zero profit (P_1 in figure 7).

- Long Run Market Supply Curve

- o Assumes that $\pi = 0$
- o Shows market Q given market P
- o Three possible types of long-run supply curves exist: Constant cost, Increasing cost, and Decreasing cost industry.
  - Constant Cost – Figures 7 and 8 illustrate a constant cost industry. Essentially, if scarcity does not apply, then with entry the demand for resources ↑ but the price of resources remains constant and the cost of inputs is also constant. Therefore price of the output remains constant in the long-run to keep profit equal to zero.

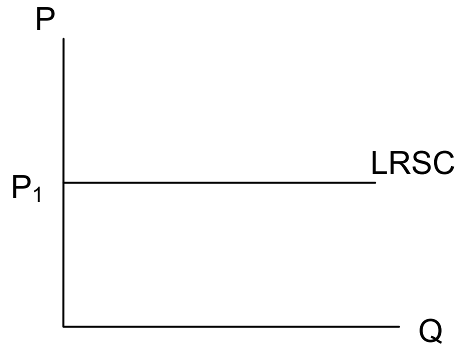

Figure 8

  - o Assumes neither market entry nor exit affects production costs.
    - During entry input supply ↑
    - During exit input supply ↓

- Increasing Cost – Figure 9 illustrates an increasing cost industry. If scarcity does apply then with entry demand for resources ↑, price of resources ↑, and costs of production ↑.

  Likewise with exit demand for resources ↓, price of resources ↓, and costs of production ↓.

  Therefore the price must rise with increasing output (entry) in the long-run to keep profit equal to zero. In other words, the long-run supply curve must have a positive slope.

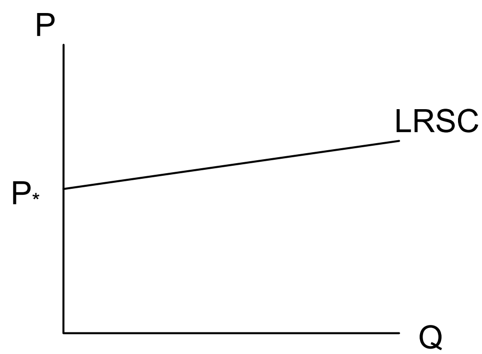

Figure 9

- Decreasing Cost – Figure 10 illustrates a decreasing cost industry. A decreasing cost industry can only occur if entry causes costs to fall. This can only happen for relatively short periods of time and occurs for two primary reasons: (1) economies of scale in a new industry and (2) new technology that decreases costs.

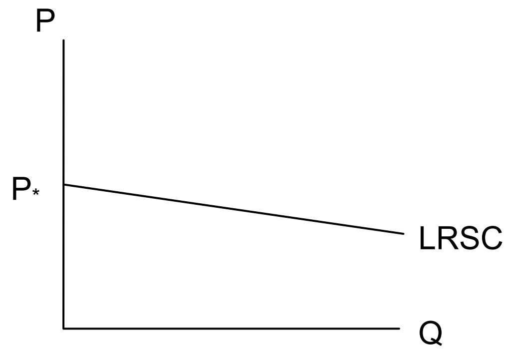

Figure 10

Therefore, the price must fall with increasing output (entry) in the long-run to keep profit equal to zero and, hence, the long-run supply curve will have a negative slope.