ADDITIONAL MATERIAL: COURNOT

EEBL - Game Theory

University of Tor Vergata

2021-2022

Guillaume Pommey

This note explains more precisely why do we use the first-order condition of a problem as a way to find the solution to the maximization problem in a Cournot setting.

Assume there are two firms competing à la Cournot. Consider the following inverse demand function $P(Q) = 4 - Q$ where $Q = q_1 + q_2$ is the aggregate quantity (the sum of each firm's quantity). Further assume that the two firms have the same marginal cost $c = 1$.

Let us focus on the profit of firm 1:

$$\begin{aligned}\pi_1(q_1, q_2) &= (P(q_1 + q_2) - c)q_1 \\ &= (3 - q_1 - q_2)q_1 \\ &= 3q_1 - q_1^2 - q_1 q_2.\end{aligned}$$

Notice that we can interpret the first term $(9 - q_1 - q_2)$ as being the unit margin of firm 1. Indeed it corresponds to (Price - marginal cost). While the other term q_1 is simply the quantity sold by firm 1. So the profit is simply (Unit Margin)*Quantity.

Consider firm 1's maximization problem:

$$\max_{q_1} \left[\pi_1(q_1, q_2) = 3q_1 - q_1^2 - q_1 q_2\right],$$

that is, firm 1 must choose q_1 to maximize its profits for any given q_2. In other words, firm 1 must choose which amount q_1 to produce for each possible q_2 that firm 2 may produce (which is what we call the best-response function).

Now let us plot the profit of firm 1 on the plane (π_1, q_1) for different values of q_2, namely, for $q_2 = 0$, $q_2 = 0.5$ and $q_2 = 1$. The three functions we plot are the following:[1]

$$\pi_1(q_1, 0) = 3q_1 - q_1^2 - 0q_1,$$
$$\pi_1(q_1, 0.5) = 3q_1 - q_1^2 - 0.5q_1,$$
$$\pi_1(q_1, 1) = 3q_1 - q_1^2 - 1q_1.$$

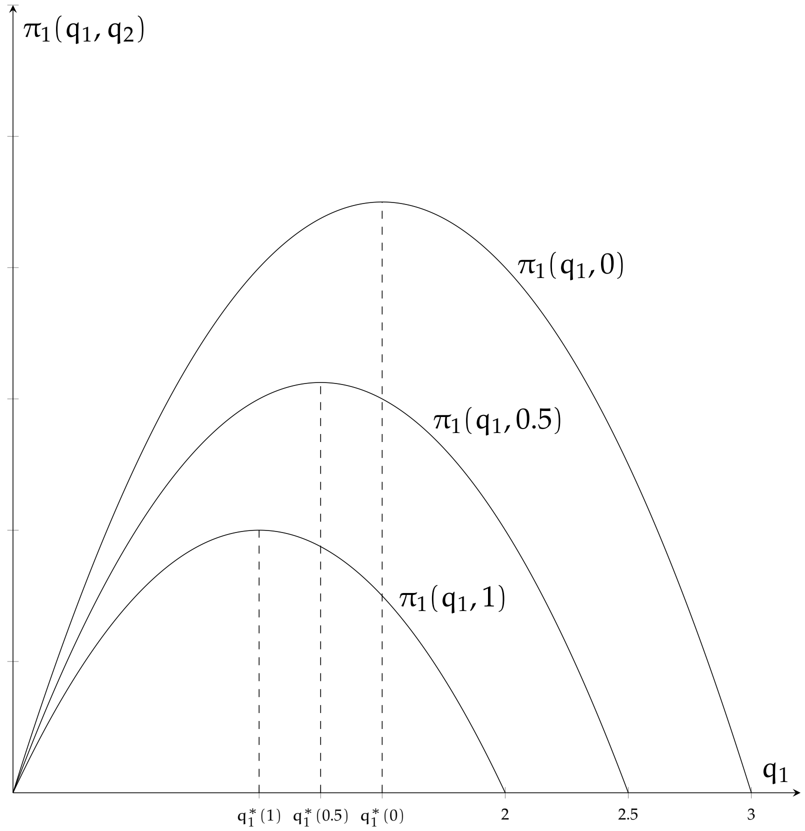

From the plot, it is quite easy to identify where is the maximum of firm 1's profits for each given value of q_2, it is at the top of each bell-shaped curve.

For each peak, we can trace back the *best* quantity firm 1 must produce and those are $q_1^*(0) = 1.5$, $q_1^*(0.5) = 1.25$ and $q_1^*(1) = 1$. Notice that the *best* quantity q_1 changes as the quantity produced by firm 2 changes because the shape of the profit function changes and the peak is not at the same place.

We have in fact identified some elements of the best-response function of firm 1. That is, we have found what was the best-response of firm 1 to quantities $q_2 = 0$, $q_2 = 0.5$ and $q_2 = 1$. If we were to repeat the same exercise for every possible value of q_2 we would obtain the complete best-response function of firm 1.

However, as there is an infinite number of possibilities for q_2 we cannot simply find the best-response function by repeating this graphical analysis. Thankfully, firm 1's profit function

[1] If you want to play with those profit curves and see them for other values of q_2 you can click here and change the value of q_2 yourself.

has the nice mathematical property of being *differentiable* and this is what we are going to exploit.

What is the link between differentiability and the graphical approach we have used to find some elements of the best-response function?

Recall that the derivative of a function at some point x tells us by how much the function is supposed to vary if we look at a tiny increase in x. Graphically, the derivative at x is the slope of the tangent line to the graph of the function at that point.

The next figure plots some tangent lines (in green, red and blue) for some given points.

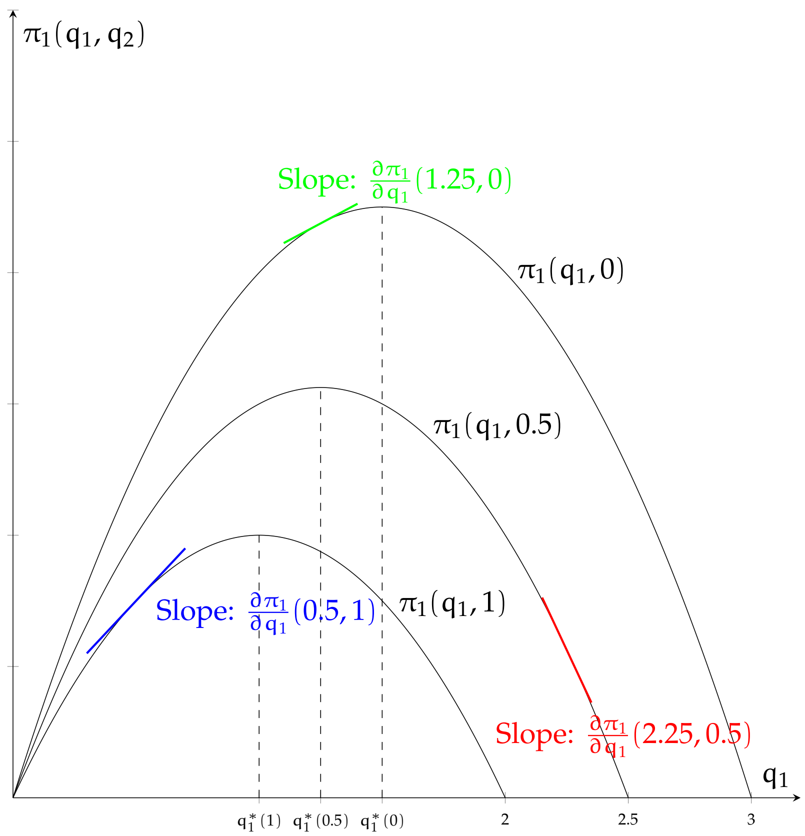

Each of those tangent lines has a slope that corresponds to the derivative of firm 1's profit with respect to q_1 and evaluated at the corresponding point. For instance, the blue tangent line slope can be computed by first taking the derivative with respect to q_1 of $\pi_1(q_1, 1)$ that is

$$\begin{aligned}\frac{\partial \pi_1}{\partial q_1}(q_1, 1) &= \frac{\partial}{\partial q_1}\left(3q_1 - q_1^2 - 1q_1\right) \\ &= 3 - 2q_1 - 1 \\ &= 2 - 2q_1.\end{aligned}$$

Then, as the blue tangent line has been computed for $q_1 = 0.5$, we evaluate the above derivative at this point and this gives us $\frac{\partial \pi_1}{\partial q_1}(0.5, 1) = 2 - 2 * 0.5 = 1$. This tells us that the slope of the tangent line at this particular point is 1, that is, the function is expected to increase in value if we go a little bit more on the right.

Notice that we expect the green tangent line to have a positive slope/derivative but we expect the red tangent line to have a negative slope/derivative.

Coming back to our problem of finding the best-response function thanks to the use of derivative, notice that we can draw the tangent lines at the top point of each bell-shaped curve which, visually, are obviously the maximum points of each profit curve.

The next figure represents those tangent lines at each maximum point.

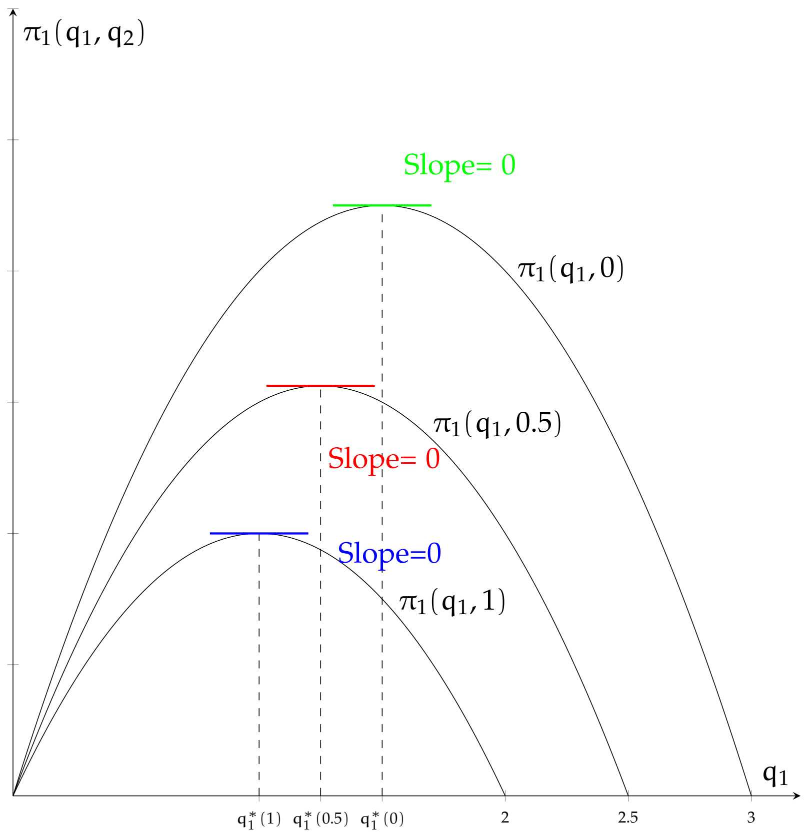

At each maximum point, the slope of the tangent line is zero and so is the derivative at this point.

We can deduce from this graphical example that if we try to find where does the derivative of firm 1's profit is zero we will find the exact point at which the profit is maximum.[2]

Therefore, if we try to solve the following equation:

$$\frac{\partial \pi_1}{\partial q_1}(q_1, 1) = 0,$$

we will obtain the value of q_1 such that the derivative of the profit is zero and this value

[2] Obviously, this idea relies on more sophisticated maths than a mere graphical analysis.

corresponds to $q_1^*(1)$ and the corresponding blue horizontal tangent line. In practice:

$$\begin{aligned} \frac{\partial \pi_1}{\partial q_1}(q_1, 1) &= 0 \\ \Leftrightarrow 2 - 2q_1 &= 0 \\ \Leftrightarrow q_1 &= 1. \end{aligned}$$

Again, if we were to proceed like this we would have make those computations for each possible value of q_2 (an infinite number of them).

We can simply let q_2 undefined in firm 1's profit function and solve for "derivative w.r.t q_1 is equal to zero" for any possible value of q_2. Formally,

$$\begin{aligned} \frac{\partial \pi_1}{\partial q_1}(q_1, q_2) &= 0 \\ \Leftrightarrow 3 - 2q_1 - q_2 &= 0 \\ \Leftrightarrow q_1 &= \frac{3 - q_2}{2}. \end{aligned}$$

We like to define $q_1^*(q_2) = \frac{3-q_2}{2}$ (the best-response function) to better indicate that it is a function q_2. Notice that now, with only one computation, we have figured out all the possible best value of q_1 for any possible q_2.

If you want to compute what is the best q_1 when $q_2 = 0.5$ (red tangent) then you simply have to plug $q_2 = 0.5$ into $q_1^*(q_2)$ and you get $q_1^*(0.5) = 1.25$. But now, thanks to the formula you recover all possible best value of q_1 for any possible value of q_2.

You should now have understood why equating the derivative of firm 1's profit to zero and solving the equation for q_1 is equivalent to looking for the value of q_1 that maximizes firm 1's profit function.