EconS 503 - Advanced Microeconomics - II
Midterm Exam #1 - Answer key

1. **Cournot competition between public and private firms.** Consider a market, such as oil and natural gas in several countries, with one public firm (firm 0), and one private firm (firm 1). Both firms produce a homogenous good with identical and constant marginal cost c per unit of output, where $1 > c \geq 0$, and face inverse linear demand function $p(X) = 1 - X$ where $X = x_0 + x_1$ denotes aggregate output. The private firm maximizes its profit

$$\pi_1 = (1 - X)x_1 - cx_1$$

And the public firm maximizes a combination of social welfare and profits

$$V_0 = \theta W + (1 - \theta)\pi_0$$

where social welfare (W) is given by

$$W = \int_0^X p(y)dy - cX,$$

and its profits are $\pi_2 = (1 - X)x_2 - cx_2$. Intuitively, parameter θ represents the weight that the manager of the public firm assigns to social welfare W, while $(1-\theta)$ denotes the weight that he assigns to its profits π_2. Both firms simultaneously and independently choose output (as in the Cournot model of quantity competition).

(a) Find the best-response functions of the private firm, $x_1(x_0)$, and of the public firm, $x_0(x_1)$.

(b) Depict the best response functions of each firm. For simplicity, you can restrict your analysis to the cases in which $\theta = 0$, $\theta = \frac{1}{2}$ and $\theta = 1$. Intuitively explain the rotation in the public firm's best response function as θ increases.

(c) Calculate the equilibrium quantities. How are they affected by an increase in parameter θ?

- See answer key at the end of this handout.

2. **Subgame perfection.** Find the subgame perfect equilibrium of the game tree depicted in the following figure. Player 1 chooses first whether to go In or Out. If he goes Out the game is over. If he goes In, player 2 gets to choose between A, B and C. If player 2 chooses C the game is over. If, in contrast, he chooses A or B, player 3 gets to select between X and Y, but without observing whether player 2 chose A or B. (That's why the two nodes of player 3 are connected with an information set.) Payoff vectors, as usual, denote the payoff that player 1 obtains, followed by the payoff that

player 2 obtains, and that of player 3.

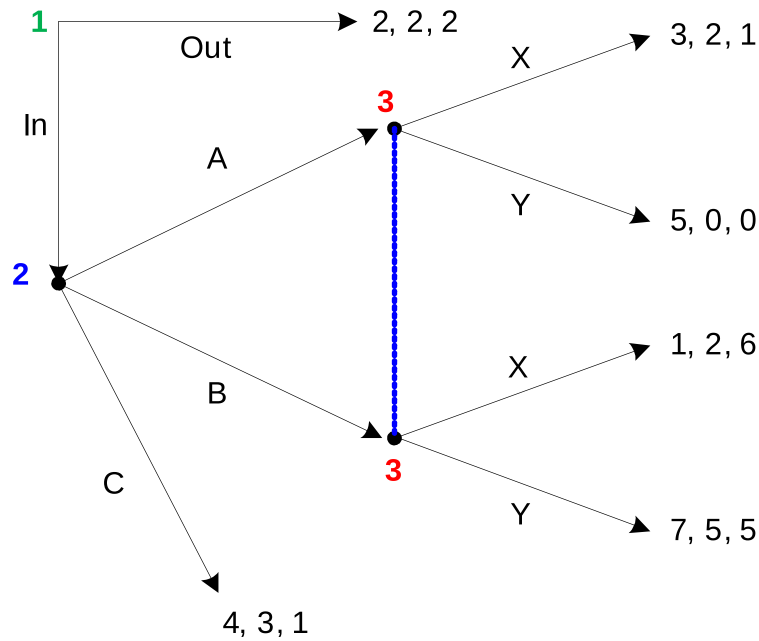

- See answer key at the end of this handout.

3. **Bertrand collusion with imperfect monitoring.** Consider two symmetric firms that consider collusion. The fully collusive price in the market is given by $p_m > 0$, and gives a profit of $\pi_m > 0$ for each firm. Firms also have a common discount factor $\delta \in (0, 1)$. They play the Bertrand game of price competition an infinite number of periods.

   There also exists an antitrust authority, which investigates the industry in every period.

   - If firms collude, the authority will find them guilty with a probability p and will accordingly give them a fine $F > \delta\pi_n$.
   - If firms are found colluding, assume that the authority will prevent them from colluding in the future, earning profit $\pi_n > 0$ each for all subsequent periods, where the index n stands for Nash.
   - If firms do not collude, they cannot be fined.

   Consider a simple Grim Trigger Strategy (GTS) in which firms start cooperating, but punish deviation from either firm by reverting to the Nash equilibrium of the stage game forever.

   (a) Find the minimum discount factor, $\overline{\delta}$, sustaining collusion in the Subgame Perfect Equilibrium of the game.

   (b) *Comparative statics.* How do p and F affect $\overline{\delta}$? Interpret your results.

   - See answer key at the end of this handout.

4. **First-price auction with risk averse bidders.** Consider a first-price auction where every bidder i's valuation is drawn from a cumulative distribution function, $F(v_i)$, with positive density in all its support, i.e., $f(v_i) > 0$ for all $v_i \in [0, 1]$. Every bidder i's utility function is $u : \mathbb{R}_+ \to \mathbb{R}$, where $u(0) = 0$, $u' > 0$ and $u'' < 0$; thus indicating that the bidder is risk averse.

   In this exercise, we seek to find the bidding function when bidders are risk averse, risk neutral, and compare these bidding functions, but allowing for valuations to be drawn from any $F(v_i)$, not only the uniform distribution.

   We can write bidder i's expected utility maximization problem as follows.

$$\max_{b_i \geq 0} \ \Pr(win) \times u(v_i - b_i)$$

   which denotes the probability of winning the object times bidder i's net payoff from winning, $u(v_i - b_i)$. Alternatively, we can use the so-called "indirect approach" to write the above expected utility maximization problem as follows

$$\max_{z \geq 0} \ G(z) \times u(v_i - b(z))$$

   where z denotes the valuation that bidder i considers which, for generality, can coincide with his true valuation v_i or not at this point. Inserting this valuation z into the bidding function $b_i : [0, 1] \to \mathbb{R}_+$, he submits bid $b(z)$, and wins the auction with probability $G(z) \equiv F^{N-1}(z)$.

   (a) Differentiate bidder i's expected utility with respect to z to characterize the implicit function that describes bidder i's bidding function. [Hint: You can solve for $b'(v_i)$ on the left hand of the equation.]

   - Taking first-order conditions with respect to z, we obtain

$$g(z)u(v_i - b(z)) - G(z)u'(v_i - b(z))b'(z) = 0$$

     In equilibrium, it must be optimal to choose a valuation $z = v_i$, so bidder i does not have incentives to bid according to a valuation $z \neq v_i$. Inserting $z = v_i$ in the above expression, yields

$$g(v_i)u(v_i - b(v_i)) - G(v_i)u'(v_i - b(v_i))b'(v_i) = 0$$

     which we can rearrange as

$$b'_{RA}(v_i) = \frac{u(v_i - b(v_i))}{u'(v_i - b(v_i))} \frac{g(v_i)}{G(v_i)}$$

     where the subscript RA denotes "risk-averse" bidder.

   (b) Evaluate your results in a setting where all bidders are risk neutral, so that $u(x) = x$, yielding $u'(x) = 1$.

   - Since $u(x) = x$ and $u'(x) = 1$, the above result reduces to

$$b'_{RN}(v_i) = (v_i - b(v_i)) \frac{g(v_i)}{G(v_i)}$$

     where the subscript RN denotes "risk-neutral" bidder.

(c) Compare your results in parts (b) and (c). Are the bidders bidding more or less aggressive than when they are risk neutral? Interpret your results.

- Comparing the two expressions, we obtain

$$b'_{RA}(v_i) = \frac{u(v_i - b(v_i))}{u'(v_i - b(v_i))} \frac{g(v_i)}{G(v_i)} > (v_i - b(v_i)) \frac{g(v_i)}{G(v_i)} = b'_{RN}(v_i) \qquad (1)$$

  Since $u(0) = 0$ and $u''(x) < 0$, we find that $\frac{u(x)-u(0)}{x-0} > u'(x)$ for all x. Graphically, the slope of the utility function at point x is less steep than the ray connecting the origin and x. Setting $x = v_i - b(v_i)$ for $\frac{u(x)}{u'(x)} > x$, yields

$$\frac{u(v_i - b(v_i))}{u'(v_i - b(v_i))} > v_i - b(v_i).$$

  Therefore, a given increase in valuation v_i increases more significantly the bid of a risk-averse bidder, $b'_{RA}(v_i)$, than that of a risk-neutral bidder, $b'_{RN}(v_i)$.

- Since, in addition, at $v_i = 0$, bids are both zero, $b_{RN}(0) = b_{RA}(0) = 0$, we must have that

$$b_{RA}(v_i) > b_{RN}(v_i)$$

  for all $0 < v_i \leq 1$. Intuitively, bidders submit more aggressive bids when they are risk averse than risk neutral.

EconS 503 – Midterm #1, Spring 2022

Exercise #1

a) The private firm maximizes

$$\pi_1 = [1-(x_0+x_1)]x_1 - cx_1,$$

Taking the first-order conditions with respect to x_1, we find the best-response function of the private firm:

$$x_1(x_0) = \frac{1-c}{2b} - \frac{x_0}{2}.$$

The public firm maximizes $V_0 = \theta W + (1-\theta)\pi_0$. That is,

$$V = (1-\theta)([1-(x_0+x_1)]x_0 - cx_0) + \theta\left[\int_0^{x_0+x_1}(1-x)dx - c(x_0+x_1)\right],$$

Taking the first-order conditions with respect to x_0, we have

$$\frac{\partial V}{\partial x_0} = (1-\theta)(1-x_1-c-2x_0) + \theta[1-(x_0+x_1)-c)] = 0,$$

And solving for x_0, we find the best-response function of the public firm:

$$x_0 = \frac{1-c}{2\left(1-\frac{\theta}{2}\right)} - \frac{x_1}{2\left(1-\frac{\theta}{2}\right)}.$$

b)

In order to better understand the effect of θ on the public firm's best response function and, as a consequence, on equilibrium behavior, let us briefly examine some comparative statistics of parameter θ. In particular, when θ increases from 0 to 1 the best-response function of the public firm pivots outward, with its vertical intercept being unaffected. Let us analyze some extreme cases (the figure below depicts the public firm's best response function evaluated at different values of θ):

- At $\theta = 0$, the best response function of the public firm becomes analog to that of private firm: the public firm behaves exactly as the private firm and the equilibrium is symmetric, with both firms producing the same amount of output. Intuitively, this is not surprising: when $\theta = 0$ the manager of the public company assigns no importance to social welfare, and only cares about profits.
- As θ increases, the slope the best-response function of the public firm, $2\left(1-\frac{\theta}{2}\right)$, becomes flatter, and the public firm produces a larger share of industry output, i.e. the crossing point between both best response functions happens more to the southeast.
- At $\theta = 1$, only the public company produces in equilibrium, while the private firm does not produce (corner solution).

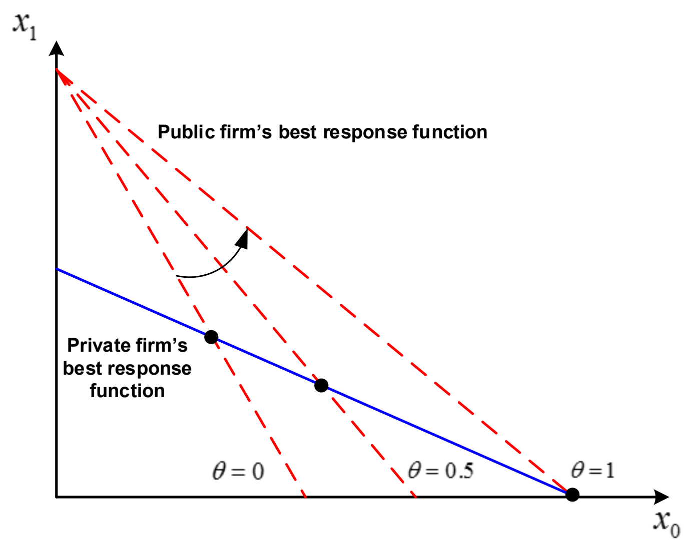

Best response function of the public and private firm

c) The equilibrium quantities solve the system of two equations:

$$x_1 = \frac{1-c}{2} - \frac{x_0}{2}, \text{ and}$$

$$x_0 = \frac{1-c}{2\left(1-\frac{\theta}{2}\right)} - \frac{x_1}{2\left(1-\frac{\theta}{2}\right)}.$$

Simultaneously solving for x_0 and x_1, we find the individual output levels:

$$x_0 = \frac{1-c}{3-2\theta}, \quad x_1 = \frac{(1-\theta)(1-c)}{3-2\theta}.$$

Note that for $\theta = 0$ the outcome is the same as that for a Cournot duopoly, with both firms producing $x_0 = x_1 = \frac{1-c}{3}$, and for $\theta = 1$ the public firm produces the competitive outcome, $x_0 = 1 - c$, and the private firm produces nothing, $x_1 = 0$. (This result was already anticipated in our discussion of the pivoting effect in the best response function in the previous question.)

The aggregate output as a function of θ is, therefore,

$$X(\theta) = x_0 + x_1 = \frac{(2-\theta)(1-c)}{3-2\theta}.$$

Finally, note that differentiating the aggregate output with respect to θ yields

$$\frac{dX(\theta)}{d\theta} = \frac{1-c}{(3-2\theta)^2} > 0.$$

Hence, an increase in θ results in an increase in output, and an increase in social welfare.

d) The socially optimal output level corresponds to $p(X) = c$, which implies $1 - X = c$, that is, $X = 1 - c$. The equilibrium output of the mixed duopoly, $\frac{(2-\theta)(1-c)}{3-2\theta}$, is below the socially optimal level, $1 - c$, for any θ satisfying $\frac{2-\theta}{3-2\theta} < 1$.That is, $2 - \theta < 3 - 2\theta$ or $\theta = 1$. However, it exactly attains this output level at $\theta = 1$.

Exercise #2

Starting from the terminal nodes, the smallest proper subgame we can identify is depicted in the figure below (which initiates after player 1 chooses In):

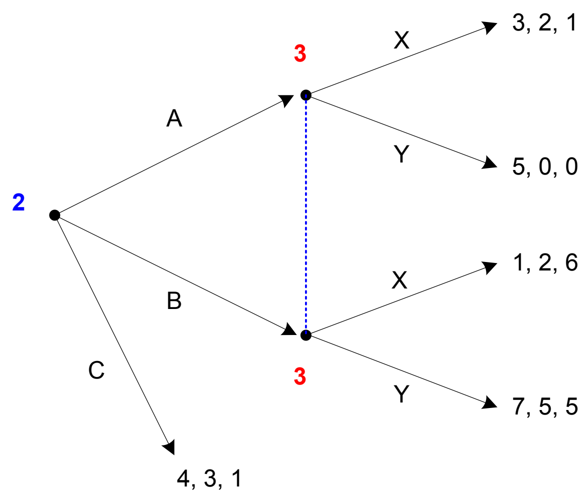

Smallest proper subgame

In this subgame in which only two players interact (players 2 and 3), player 3 chooses his action without observing player 2's choice. As a consequence, the interaction between players 2 and 3 in this subgame can be modeled as a simultaneous-move game. In order to find the Nash equilibrium of the subgame depicted in the above figure, we must first represent it in its normal (matrix) form, as depicted in the next payoff matrix:

		Player 3	
		X	*Y*
Player 2	*A*	3,2,1	5,0,0
	B	1,2,6	7,5,5
	C	4,3,1	4,3,1

Smallest proper subgame (in its Normal-form)

In addition, the payoff matrix also underlines the payoffs that arise when each player selects his best responses. In outcome (C, X) the payoffs of player 2 and 3 are underlined, thus indicating that they are playing a mutual best response to each other's strategies. Furthermore, we do not need to examine player 1, since only player 2 and 3 are called on to move in the subgame. Hence, the Nash equilibrium of this subgame predicts that players 2 and 3 choose strategy profile (C, X). We can now plug the payoff triple resulting from the Nash equilibrium of this subgame, (4,3,1), at the end of the branch indicating that player 1 chooses In (recall that this was the node initiating the smallest proper subgame), as we illustrate in the next figure.

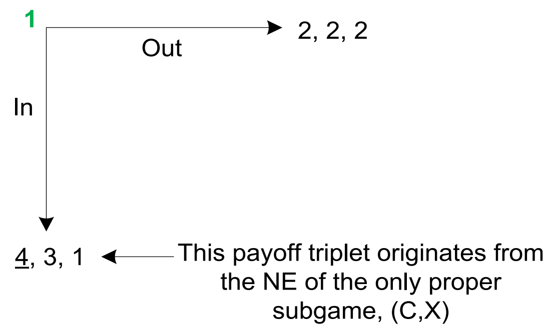

Extensive-form game (First Stage)

By inspecting the above game tree in which player 1 chooses In or Out, we can see that his payoff from In (4) is larger than from Out (2). Then, the SPNE of this game is:

(In, C, X).

Exercise #3

If the antitrust authority investigates the sector in every period, the present discounted value of collusion is given by

$$V^c = p\left(\pi_m - F + \frac{\delta}{1-\delta}\pi_n\right) + (1-p)(\pi_m + \delta V^c)$$

In words, in the current period two things can happen. Either:

- The firm is *caught* and thus obtains collusion profits π_m but pays the fine *F*. In that case, the firm can no longer collude and thus earns Nash profits π_n thereafter.
- The firm is *not caught* obtaining collusion profits π_m. In this case, the firm's discounted continuation payoff is V^c, since it will collude in the next period (having the probability of being caught once again).

Solving for the continuation payoff V^c in the above expression yields

$$V^c = \frac{(1-\delta)\pi_m - Fp + \delta(F + \pi_n)p}{(1-\delta)(1-\delta+\delta p)}$$

Therefore, we can express the incentive constraint that each firm faces as follows

$$V^c \geq \pi_d + \frac{\delta}{1-\delta}\pi_n$$

Or more explicitly,

$$\frac{(1-\delta)\pi_m - Fp + \delta(F+\pi_n)p}{(1-\delta)(1-\delta+\delta p)} \geq \pi_d + \frac{\delta}{1-\delta}\pi_n$$

Solving for the minimal discount factor δ sustaining collusion, we obtain

$$\delta \geq \bar{\delta} \equiv \frac{\pi_d - \pi_m + Fp}{\pi_d - \pi_n - \pi_d p + \pi_n p}$$

b) *Comparative statics.* Differentiating the above cutoff with respect to *p*, yields

$$\frac{\partial \bar{\delta}}{\partial p} = \frac{\pi_d + F - \pi_m}{(\pi_d - \pi_n)(1-p)^2}$$

which is positive since deviation profits π_d are higher than those under collusion and under Nash reversion, $\pi_d > \pi_m > \pi_n$.

In addition, differentiating the cutoff with respect to *F*, we obtain

$$\frac{\partial \bar{\delta}}{\partial F} = \frac{p}{(1-p)(\pi_d - \pi_n)}$$

which is also positive given that deviation profits satisfy $\pi_d > \pi_n$ by definition. Intuitively, the higher the probability of being caught, *p*, and the fine if caught, *F*, the less likely that collusion will be sustained in equilibrium.