1. From Extensive Form to Normal form Representation-I: Represent the extensive-form game depicted in Fig. 1 using its normal-form (matrix) representation.

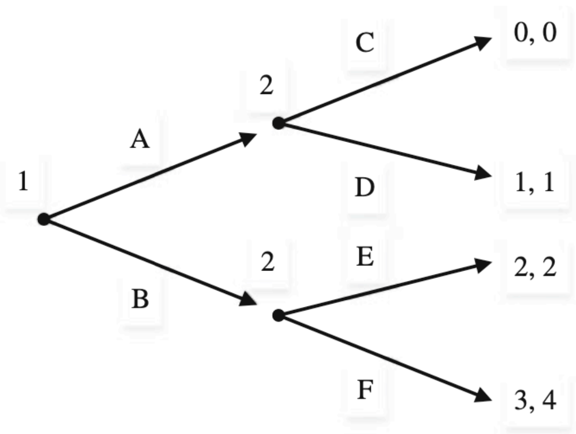

Answer

We start identifying the strategy sets of all players in the game. The cardinality of these sets (number of elements in each set) will determine the number of rows and columns in the normal-form representation of the game.

Starting from the initial node (in the root of the game tree located on the left-hand side of the figure), Player 1 must select either strategy A or B, thus implying that the strategy space for player 1, S_1, is:

$$S_1 = \{A, B\}$$

In the next stage of the game, Player 2 conditions his strategy on player 1's choice, since player 2 observes such a choice before selecting his own. We need to consider that the strategy profile of player 2 (S_2) must be a complete plan of action (complete contingent plan). Therefore, his strategy space becomes:

$$S_2 = \{CE,\ CF,\ DE,\ DF\}$$

Player 1 \ Player 2	CE	CF	DE	DF
A	0, 0	0, 0	1, 1	1, 1
B	2, 2	3, 4	2, 2	3, 4

2. From Extensive Form to Normal Form Representation-III: Consider the extensive-form game in Fig 2. Provide its normal form (matrix) representation.

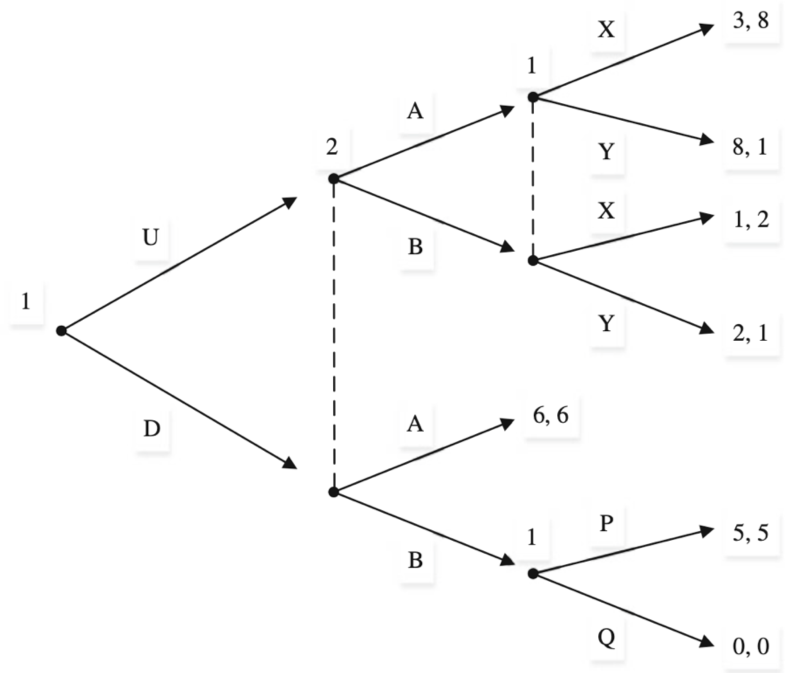

Answer

Player 2. From the extensive form game, we know player 2 only plays once and has two available choices, either A or B. The dashed line connecting the two nodes at which player 2 is called on to move indicates that player 2 cannot observe player 1's choice. Hence, he cannot condition his choice on player 1's previous action, ultimately implying that his strategy space reduces to $S_2 = \{A, B\}$.

Player 1. Player 1, however, plays twice (in the root of the game tree, and after player 2 responds) and has multiple choices:

1. First, he must select either U or D, at the initial node of the tree, i.e., left-hand side of the figure;
2. Then choose X or Y, in case that he played U at the beginning of the game. (Note that in this event he cannot condition his choice on player 2's choice, since he cannot observe whether player 2 selected A or B); and
3. Then choose P or Q, which only becomes available to player 1 in the event that player 2 responds with B after player 1 chose D.

Therefore, player 1's strategy space is composed of triplets, as follows,

$$S_1 = \{UXP, UXQ, UYP, UYQ, DXP, DXQ, DYP, DYQ\}$$

whereby the first component of every triplet describes player 1's choice at the beginning of the game (the root of the game tree), the second component represents his decision (X or Y) in the event that he chose U and afterwards player 2 responded with either A or B (something player 1 cannot observe), and the third component reflects his choice in the case that he chose D at the beginning of the game and player 2 responds with B.[1]

1 \ 2	A	B
UXP	3, 8	1, 2
UXQ	3, 8	1, 2
UYP	8, 1	2, 1
UYQ	8, 1	2, 1
DXP	6, 6	5, 5
DXQ	6, 6	0, 0
DYP	6, 6	5, 5
DYQ	6, 6	0, 0

3. Dominance Solvable Games: Two political parties simultaneously decide how much to spend on advertising, either low, medium or high, yielding the payoffs in the following matrix (Fig. 3) in which the Red party chooses rows and the Blue party chooses columns. Find the strategy profile/s that survive IDSDS.

		Blue party		
		Low	Middle	High
Red party	Low	80,80	0,50	0,100
	Middle	50,0	70,70	20,80
	High	100,0	80,20	50,50

Answer

Let us start by analyzing the Red party (row player). First, note that High is a strictly dominant strategy for the Red party, since it yields a higher payoff than both Low and Middle, regardless of the strategy chosen by the Blue party (i.e., independently of the column the Blue party selects). Indeed, $100 > 80 > 50$ when the Blue party chooses the Low column, $80 > 70 > 0$ when the Blue party selects the Middle column, and $50 > 20 > 0$ when the Blue party chooses the High column. As a consequence, both Low and Middle are strictly dominated strategies for the Red party (they are both strictly dominated by High in the bottom row), and we can thus delete the rows corresponding to Low and Middle from the payoff matrix, leaving us with the following reduced matrix (Fig. 1.11).

We can now check if there are any strictly dominated strategies for the Blue party (in columns). Similarly as for the Red party, High strictly dominates both Low and Middle since $50 > 20 > 0$; and we can thus delete the columns corresponding to Low and Middle from the payoff matrix, leaving us with a single cell, (High, High). Hence, (High, High) is the unique strategy surviving IDSDS.

		Blue party		
		Low	Middle	High
Red party	High	100,0	80,20	50,50

Political parties reduced normal-form game

4. Two students in the Industrial Economics course plan to take an exam tomorrow. The professor seeks to create incentives for students to study, so he tells them that the student with the highest score will receive a grade of A and the one with the lower score will receive a B. Student 1's score equals x1 + 1:5, where x1 denotes the amount of hours studying. (That is, he assume that the greater the effort, the higher her score is.) Student 2's score equals x2, where x2 is the hours she studies. Note that these score functions imply that, if both students study the same number of hours, x1 = x2, student 1 obtains a highest score, i.e., she might be the smarter of the two. Assume, for simplicity, that the hours of studying for the industrial economics class is an integer number, and that they cannot exceed 5 h, i.e., xi ∈ { 1, 2 , . . .,5}. The payoff to every student i is10–xi if she gets an A and 8–xi if she gets a B. Find which strategies survive the iterative deletion of strictly dominated strategies (IDSDS).

Answer

Part (a) Let us first show that for either player, exerting a zero effort i.e., $x_i = 0$, strictly dominates effort levels of 3, 4, and 5. If $x_i = 0$ then player i's payoff is at least 8, which occurs when she gets a B. By choosing any other effort level x_i, the highest possible payoff is $10 - x_i$, which occurs when she gets an A. Since $8 > 10 - x_i$ when $x_i > 2$, then zero effort strictly dominates efforts of 3, 4, or 5. Intuitively, we consider which is the lowest payoff that player i can obtain from exerting zero effort, and compare it with the highest payoff he could obtain from deviating to a positive effort level $x_i \neq 0$. If this holds for some effort levels (as it does here for all $x_i > 2$), it means that, regardless of what the other student $j \neq i$ does, student i is strictly better off choosing a zero effort than deviating.

Once we delete effort levels satisfying $x_i > 2$, i.e., $x_i = 3$, 4, and 5 for both player 1 (in rows) and player 2 (in columns), we obtain the 3×3 payoff matrix depicted in Fig. 1.16.

As a practice of how to construct the payoffs in this matrix, note that, for instance, when player 1 chooses $x_1 = 1$ and player 2 selects $x_2 = 2$, player 1 still gets the highest score, i.e., player 1's score is 1 + 1.5 = 2.5 thus exceeding player 2's score of 2, which implies that player 1's payoff is 10 − 1 = 9 while player 2's payoff is 8 − 2 = 6.

At this point, we can easily note that player 2 finds $x_2 = 1$ (in the center column) to be strictly dominated by $x_2 = 0$ (in the left-hand column). Indeed, regardless of which strategy player 1 uses (i.e., regardless of the particular row you look at)

Player 1 \ Player 2	0	1	2
0	10,8	10,7	8,8
1	9,8	9,7	9,6
2	8,8	8,7	8,6

Fig. 1.16 Reduced normal-form game

player 2 obtains the lowest score on the exam when he chooses an effort level of 0 or 1, since player 1 benefits from a 1.5 score advantage. Indeed, exerting a zero effort yields player 2 a payoff of 8, which is unambiguously higher than the payoff he obtains from exerting an effort of $x_2 = 1$, which is 7, regardless of the particular effort level exerted by player 1. We can thus delete the column referred to $x_2 = 1$ for player 2, which leaves us with the reduced form matrix in Fig. 1.17.

Now consider player 1. Notice that an effort of $x_1 = 1$ (in the middle row) strictly dominates $x_1 = 2$ (in the bottom row), since the former yields a payoff of 9 regardless of player 2's strategy (i.e., independently on the column), while an effort of $x_1 = 2$ only provides player 1 a payoff of 8. We can, therefore, delete the last row (corresponding to effort $x_1 = 2$) from the above matrix, which helps us further reduce the payoff matrix to the 2×2 normal-form game in Fig. 1.18.

Unfortunately, we can no longer find any strictly dominated strategy for either player. Specifically, neither of the two strategies of player 1 is dominated, since an effort of $x_1 = 0$ yields a weakly larger payoff than $x_1 = 1$, i.e., it entails a strictly larger payoff when player 2 chooses $x_2 = 0$ (in the left-hand column) but the same payoff when player 2 selects $x_2 = 2$ (in the right-hand column). Similarly, none of player 2's strategies is strictly dominated either, since $x_2 = 0$ yields the same payoff as $x_2 = 2$ when player 1 selects $x_1 = 0$ (top row), but a larger payoff when player 1 chooses $x_1 = 1$ (in the bottom row).

Therefore, the set of strategies that survive IDSDS for player 1 are $\{0, 1\}$ and for player 2 are $\{0, 2\}$. Thus, the IDSDS predicts that player 1 will exert an effort of 0 or 1, and that player 2 will exert effort of 0 or 2.

Player 1 \ Player 2	0	2
0	10,8	8,8
1	9,8	9,6
2	8,8	8,6

Fig. 1.17 Reduced normal-form game

Player 1 \ Player 2	0	2
0	10,8	8,8
1	9,8	9,6

Fig. 1.18 Reduced normal-form game

5. Cournot mergers with Efficiency Gains: Consider an industry with three identical firms each selling a homogenous good and producing at a constant cost per unit c with 1 > c > 0. Industry demand is given by $P(Q) = 1 - Q$, where $Q = q_1 + q_2 + q_3$. Competition in the marketplace is in quantities.
Part (a) Find the equilibrium quantities, price and profits.
Part (b) Consider now a merger between two of the three firms, resulting in duopolistic structure of the market (since only two firms are left: the merged firm and the remaining firm). The merger might give rise to efficiency gains, in the sense that the firm resulting from the merger produces at a cost $e.c$, with $e \leq 1$ (whereas the remaining firm still has a cost c):
    a. Find the post-merger equilibrium quantities, price and profits.
    b. Under which conditions does the merger reduce prices?
    c. Under which conditions is the merger beneficial to the merging firms?

Answer

Part (a) Each firm $i = \{1, 2, 3\}$ has a profit of

$$\pi_i = (1 - Q - c)q_i.$$

Hence, since $Q \equiv q_1 + q_2 + q_3$, profits can be rewritten as:

$$\pi_i = \big(1 - (q_i + q_j + q_k) - c\big)q_i,$$

The first-order conditions are given by

$$1 - 2q_i - q_j - q_k - c = 0$$

since firms are symmetric $q_i = q_j = q_k = q$ in equilibrium, that is

$$1 - 2q - q - q - c = o$$

or

$$1 - 4q - c = 0.$$

Solving for q at the symmetric equilibrium yields a Cournot output of,

$$q_c = \frac{1-c}{4}$$

Hence, equilibrium prices are $p_c = 1 - \frac{1-c}{4} - \frac{1-c}{4} - \frac{1-c}{4} = 1 - 3\left(\frac{1-c}{4}\right) = \frac{1-3c}{4}$ and equilibrium profits are $\pi_c = \left(\frac{1-3c}{4} - c\right)\frac{1-c}{4} = \frac{(1-c)^2}{16}$

Part (b)

i. After the merger, two firms are left: firm 1, with cost $e \cdot c$, and firm 3, with cost c. Hence, the two profit functions are now given by:

$$\pi_1 = (1 - Q - ec)q_1.$$

$$\pi_3 = (1 - Q - c)q_3$$

Taking first order conditions of π_1 with respect to q_1 yields

$$1 - 2q_1 - q_3 - ec = 0$$

and, solving for q_1, we obtain firm 1's best response function

$$q_1(q_3) = \frac{1 - ec}{2} - \frac{1}{2} q_3$$

Similarly taking first-order conditions of firm 3's profits, π_3, with respect to q_3 yields

$$1 - 2q_3 - q_1 - c = 0$$

which, solving for q_3, provides us with firm 3's best response function

$$q_3(q_1) = \frac{1 - c}{2} - \frac{1}{2} q_1$$

Plugging $q_3(q_1)$ into $q_1(q_3)$, yields

$$q_1^* = \frac{1 - ec}{2} - \frac{1}{2}\left(\frac{1 - c}{2} - \frac{1}{2} q_1^*\right)$$

Rearranging and solving for q_1^*, we obtain firm 1's equilibrium output

$$q_1^* = \frac{1 - c(2e - 1)}{3}$$

Plugging this output level into firm 3's best response function yields an equilibrium output of

$$q_3^* = \frac{1 - c(2 - e)}{3}$$

Note that the outsider firm can sell a positive output at equilibrium only if the merger does not give rise to strong cost savings: that is $q_3 \geq 0$ if $e \geq \frac{2c-1}{c}$ (if $c < 1/2$, then the previous payoff becomes $\frac{2c-1}{c} < 0$, implying that $e \geq \frac{2c-1}{c}$ holds for all $e \geq 0$, ultimately entailing that the outsider firm will always sell at the equilibrium. We hence concentrate on values of c that satisfy $c > 1/2$.) Figure 2.34 illustrates cutoff $e > \frac{2c-1}{c}$, where $c > 1/2$, and the region of (e, c)-combinations above this cutoff indicate parameters indicate that the merger is not sufficiently cost saving to induce the outside firm to produce positive output levels. The opposite occurs when the cost-saving parameter, c, is lower than $(2c - 1)/c$, thus indicating that the merger is so cost saving that the nonmerged firm cannot profitably compete against the merged firm.

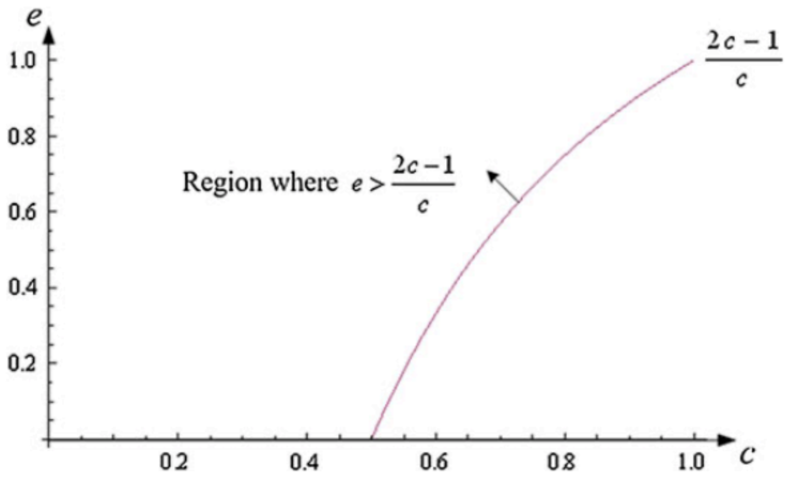

Fig. 2.34 Positive production after the merger

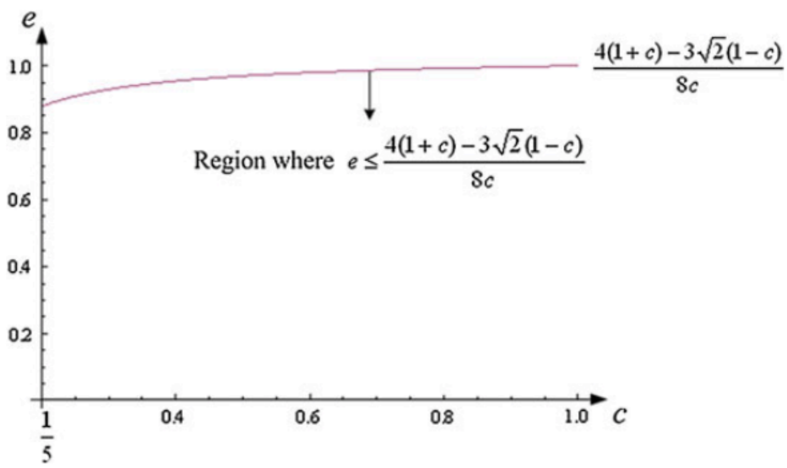

Fig. 2.35 Profitable mergers if *e* is low enough

- The equilibrium price is $p_m = 1 - \frac{1-c(2e-1)}{3} - \frac{1-c(2-e)}{3} = \frac{1+c(1+e)}{3}$, and equilibrium profits are given by $\pi_1 = \frac{(1-c(2e-1))^2}{9}$ and $\pi_3 = \frac{(1-c(2-e))^2}{9}$.

ii. Prices decrease after the merger only if there are sufficient efficiency gains: that is, $p_m \leq p_c$ can be rewritten as $e \leq \frac{5c-1}{4c}$. Note that if $c < 1/5$, then $\frac{5c-1}{4c} < 0$, implying that $e \leq \frac{5c-1}{4c}$ cannot hold for any $e \geq 0$. As a consequence, $p_m > p_c$, and prices will never fall no matter how strong efficiency gains, e, are

iii. To see if the merger is profitable, we have to study the inequality $\pi_1 \geq 2\pi_c$, which, solving for e, yields

$$e \leq \frac{4(1+c) - 3\sqrt{2}(1-c)}{8c}$$

In other words, the merger is profitable only if it gives rise to enough cost savings. Figure 2.35 depicts the cutoff of e where costs are restricted to $c \in \left[\frac{1}{5}, 1\right]$. Notice that if cost savings are sufficiently strong, i.e., parameter e is sufficiently small as depicted in the region below the cutoff, the merger is profitable.

6. Ultimatum Bargaining Game: In the ultimatum bargaining game, a proposer is given a pie, normalized to a size $1, and he is asked to make a monetary offer, x, to the responder who, upon receiving the offer, only has the option to accept or reject it (as if he received an "ultimatum" from the proposer). If the offer x is accepted, then the responder receives it while the proposer keeps the remainder of the pie 1−x. However, if he rejects it, both players receive a zero payoff. Operating by backward induction, the responder should accept any offer x from the proposer (even if it is low) since the alternative (reject the offer) yields an even lower payoff (zero). Anticipating such as a response, the proposer should then offer one cent (or the smallest monetary amount) to the responder, since by doing so the proposer guarantees acceptance and maximizes his own payoff. Therefore, according to the subgame perfect equilibrium prediction in the ultimatum bargaining game, the proposer should make a tiny offer (one cent or, if possible, an amount approaching zero), and the responder should accept it, since his alternative (reject the offer) would give him a zero payoff. However, in experimental tests of the ultimatum bargaining game, subjects who are assigned the role of proposer rarely make offers close to zero to the subject who plays as a responder. Furthermore, sometimes subjects in the role of the responder reject positive offers, which seems to contradict our equilibrium predictions. In order to explain this dissonance between theory and experiments, many scholars have suggested that players' utility function is not as selfish as that specified in standard models (where players only care about the monetary payoff they receive). Instead, the utility function should also include social preferences, measured by the difference between the payoff a player obtains and that of his opponent, which gives rise to envy (when the monetary amount he receives is lower than that of his opponent) or guilt (when his monetary payoff is lower than his opponent's). In particular, suppose that the responder's utility is given by

$$u_R(x, y) = x + \alpha(x - y)$$

where x is the responder's monetary payoff, y is the proposer's monetary payoff, and a is a positive constant. That is, the responder not only cares about how much money he receives, x, but also about the payoff inequality that emerges at the end of the game, $\alpha(x - y)$, which gives rise to either envy, if x < y, or guilt, if x > y. For simplicity, assume that the proposer is selfish, i.e., his utility function only con- siders his own monetary payoffs $u_P(x, y) = y$ as in the basic model.

   a. Use a game tree to represent this game in its extensive form, writing the payoffs in terms of m, the monetary offer of the proposer, and parameter α
   b. Find the subgame perfect equilibrium. Describe how equilibrium payoffs are affected by changes in parameter α.

Answer

Part (a) First, player 1 offers a division of the pie, m, to player 2, who either accepts or rejects it. However, payoffs are not the same as in the standard ultimatum bargaining game. While the payoff of player 1 (proposer) is just the remaining share of the pie that he does not offer to player 2, i.e., $1 - m$, the payoff of player 2 (responder) is

$$m + \alpha[m - (1 - m)] = m + \alpha(2m - 1)$$

where m is the payoff of the responder, and 1-m is the payoff of the proposer. We depict this modified ultimatum bargaining game in Fig. 4.1.

Part (b) Operating by backward induction, we first focus on the last mover (player 2, the responder). In particular, player 2 accepts any offer m from player 1 such that:

$$m + \alpha(2m - 1) \geq 0,$$

since the payoff he obtains from rejecting the offer is zero. Solving for m, this implies that player 2 accepts any offer m that satisfies $m \geq \frac{\alpha}{1+2\alpha}$. Anticipating such a response from player 2, player 1 offers the minimal m that generates acceptance, i.e., $m^* = \frac{\alpha}{1+2\alpha}$, since by doing so player 1 can maximize the share of the pie he keeps. This implies that equilibrium payoffs are:

$$(1 - m^*, m^*) = \left(\frac{1-\alpha}{1+2\alpha}, \frac{\alpha}{1+2\alpha}\right)$$

The proposer's equilibrium payoff is decreasing in α since its derivative with respect to α is

$$\frac{\partial(1 - m^*)}{\partial\alpha} = -\frac{(2\alpha^2 - \alpha + 2)}{(1+2\alpha)^2}$$

which is negative for all $\alpha > 0$. In contrast, the responder's equilibrium payoff is increasing in α given that its derivative with respect to α is

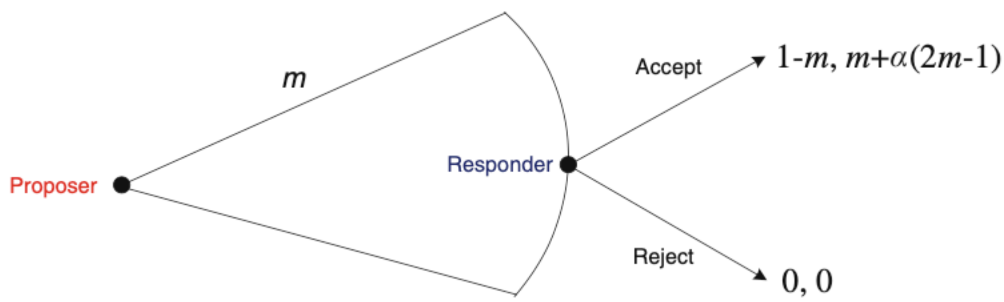

Fig. 4.1 Ultimatum bargaining game (extensive-form)

$$\frac{\partial m^*}{\partial\alpha} = \frac{1}{(1+2\alpha)^2}$$

which is positive for all $\alpha > 0$.

7. Stackelberg with Two Firms: Consider a leader and a follower in a Stackelberg game of quantity competition. Firms face an inverse demand curve p(Q) = 1 − Q, where Q = $q_1 + q_2$ denotes aggregate output. The leader faces a constant marginal cost $c_L > 0$ while the follower's constant marginal cost is $c_F > 0$, where $1 > c_F \geq c_L$, indicating that the leader has access to a more efficient technology than the follower.
    a. Find the follower's best response function.
    b. Determine each firm's output strategy in the subgame perfect equilibrium (SPNE) of this sequential-move game.
    c. Under which conditions on c_L can you guarantee that both firms produce strictly positive output levels?

Answer

Part (a) The follower observes the leader's output level, q_L, and chooses its own production, q_F, to solve:

$$\max_{q_F \geq 0} (1 - q_L - q_F)q_F - c_F q_F$$

Taking first order conditions with respect to q_F yields

$$1 - q_L - 2q_F - c_F = 0$$

and solving for q_F we obtain the follower's best response function

$$q_F(q_L) = \frac{1 - c_F}{2} - \frac{1}{2}q_L$$

which, as usual, is decreasing in the follower's costs, i.e., the vertical intercept decreases in c_F, indicating that, graphically, the best-response function experiences a downward shift as c_F increases. In addition, the follower's best-response function decreases in the leader's output decision (as indicated by the negative slope, $-\frac{1}{2}$).
Part (b) The leader anticipates that the follower will respond with best response function $q_F(q_L) = \frac{1-c_F}{2} - \frac{1}{2}q_L$, and plugs it into the leader's own profit maximization problem, as follows

$$\max_{q_L \geq 0} \left[1 - q_L - \underbrace{\left(\frac{1 - c_F}{2} - \frac{1}{2}q_L\right)}_{q_F}\right] q_L - c_L q_L$$

which simplifies into

$$\max_{q_L \geq 0} \frac{1}{2}[(1 - c_F) - q_L - c_L]q_L.$$

Taking first order conditions with respect to q_L yields

$$\frac{1}{2}(1 - c_F) - q_L - c_L = 0$$

and solving for q_L we find the leader's equilibrium output level

$$q_L^* = \frac{1 + c_F - 2c_L}{2}$$

which, thus, implies a follower's equilibrium output of

$$q_F^* = q_F\left(\frac{1 + c_F - 2c_L}{2}\right) = \frac{1 - 3c_F + 2c_L}{4}$$

Note that the equilibrium output of every firm $i=\{L,F\}$ is decreasing in its own cost, c_i, and increasing in its rival's cost, c_j, where $j \neq i$. Hence, the SPNE of the game is

$$\left(q_L^*, q_F(q_L)\right) = \left(\frac{1 + c_F - 2c_L}{2}, \frac{1 - c_F}{2} - \frac{1}{2}q_L\right)$$

which allows the follower to optimally respond to both the equilibrium output level from the leader, q_L^*, but also to off-the-equilibrium production decisions $q_L \neq q_L^*$.
Part (c) The follower (which operates under a cost disadvantage), produces a positive output level in equilibrium, i.e., $q_F^* > 0$, if and only if

$$\frac{1 - 3c_F + 2c_L}{4} > 0,$$

which, solving for c_L, yields

$$c_L > -\frac{1}{2} + \frac{3}{2}c_F \equiv C_A.$$

Similarly, the leader produces a positive output level, $q_L^* > 0$, if and only if

$$\frac{1 + c_F - 2c_L}{2} > 0$$

or, solving for c_L,

$$c_L < \frac{1}{2} + \frac{1}{2}c_F \equiv C_B.$$

Figure 4.10 depicts cutoffs C_A (for the follower) and C_B (for the leader), in the (c_F, c_L)—quadrant. (Note that we focus on points below the 45°-line, since the leader experiences a cost advantage relative to the follower, i.e., $c_L \leq c_F$.)

First, note that cutoff C_B is not binding since it lies above the 45°-line. Intuitively, the leader produces a positive output level for all (c_F, c_L)—pairs in the admissible region of cost pairs (below the 45°-line). However, cutoff C_A restricts the of cost pairs below the 45°-line to only that above cutoff C_A. Hence, in the shaded area of the figure, both firms produce a strictly positive output in equilibrium.

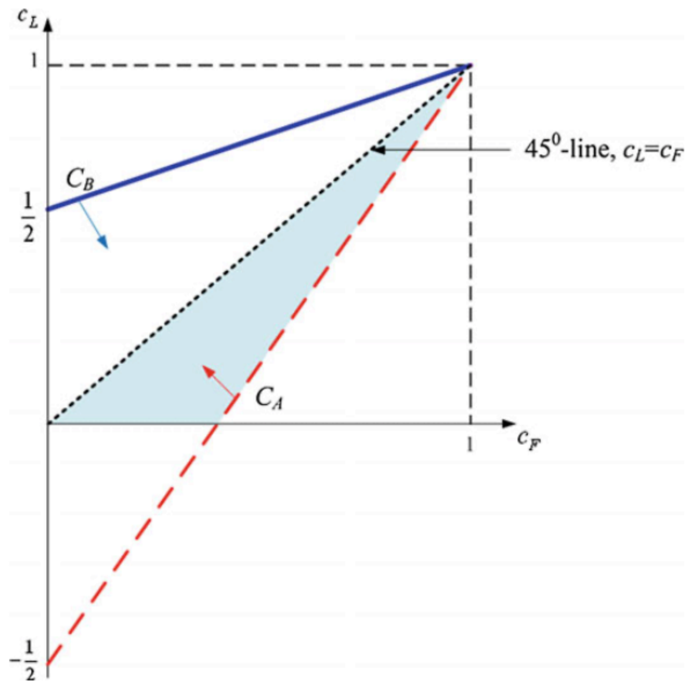

Fig. 4.10 Region of cost pairs for which both firms produce positive output

8. Backward Induction-II: Consider the sequential-move game depicted in the Figure 4.17. The game describes Apple's decision to develop the new iPhone with radically new software which allows for faster and better applications (apps). These apps are, however, still not developed by app developers (such as Rovio, the Finnish company that introduced Angry Birds). If Apple does not develop the new iPhone, then all companies make zero profit in this emerging market. If, instead, the iPhone is introduced, then company 1 (the leader in the app industry) gets to decide whether to develop apps that are compatible with the new iPhone's software. Upon observing company 1's decision, the followers (firm 2 and 3) simultaneously decide whether to develop apps (D) or not develop (ND). Find all subgame perfect Nash equilibria in this sequential-move game.

Fig. 4.17 Developing iPhones and apps game

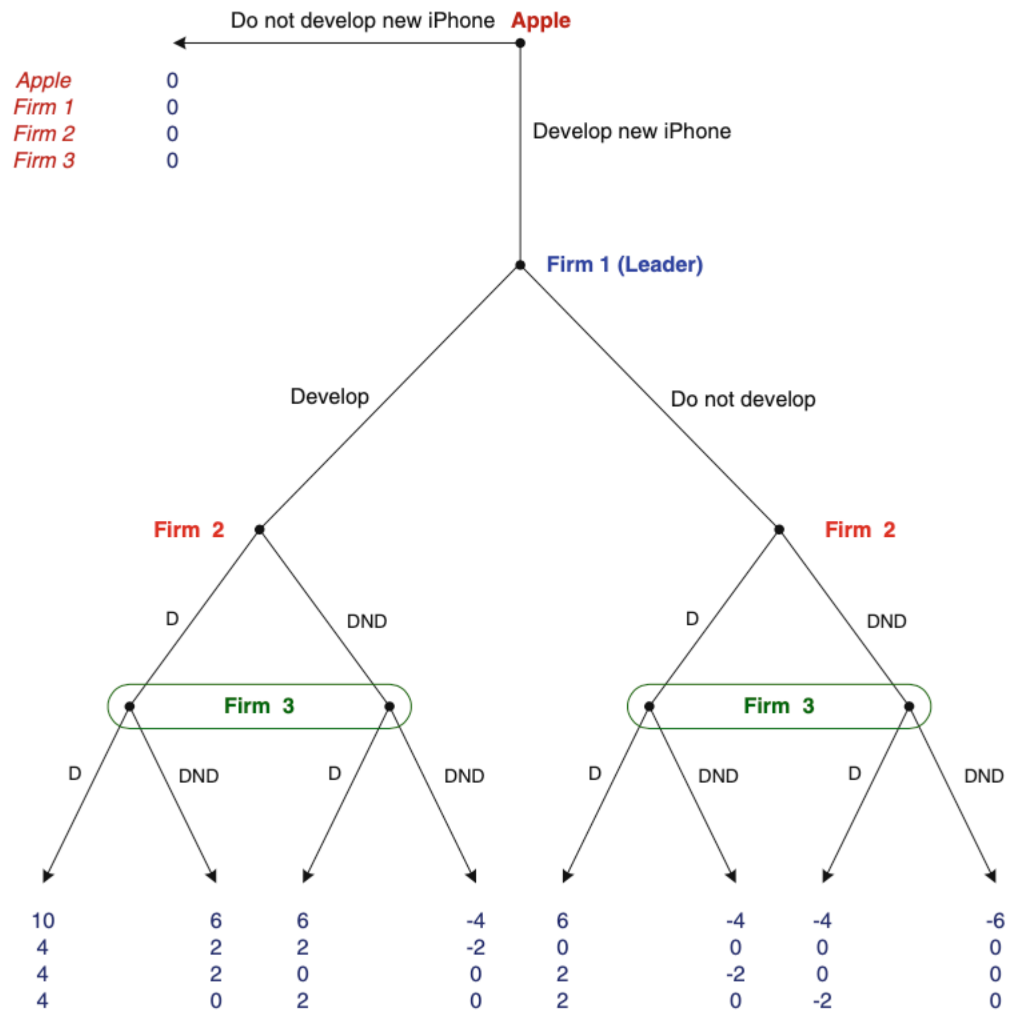

		Firm 3	
		Develop	Do not develop
Firm 2	Develop	$\underline{4}$, $\underline{4}$	2, 0
	Do not develop	0, 2	0, 0

Fig. 4.18 Smallest proper subgame (after firm 1 develops)

Answer

Company 1 develops. Consider the subgame between firms 2 and 3 which initiates after Apple develops the new iPhone and company 1 develops an application (in the left-hand side of the game tree of Fig. 4.17).

Since that subgame describes that firm 2 and 3 simultaneously choose whether or not to develop apps, we must represent it using its normal form in order to find the NEs of this subgame; as we do in the payoff matrix of Fig. 4.18.

We can identify the best responses for each player (as usual, the payoffs associated to those best responses are underlined in the payoff matrix of Fig. 4.18). In particular, $BR_2(D,D) = D$ and $BR_2(D,ND) = D$ for firm 2; and similarly for firm 3, $BR_3(D,D) = D$ and $BR_3(D,ND) = D$. Hence, *Develop* is a dominant strategy for each company, so there is a unique Nash equilibrium (*Develop, Develop*) in this subgame.

Company 1 does not develop. Next, consider the subgame associated with Apple having developed the new iPhone but company 1 not developing an application (depicted in the right-hand side of the game tree in Fig. 4.17). Since the subgame played between firms 2 and 3 is simultaneous, we represent it using its normal form in Fig. 4.19.

This subgame has two Nash equilibria: (*Develop, Develop*) and (*Do not develop, Do not develop*). (Note that, in our following discussion, we will have to separately analyze the case in which outcome (D, D) emerges as the NE of this subgame, and that in which (ND, ND) arises.)

Company 1—Case I. Let us move up the tree to the subgame initiated by Apple having developed the new iPhone. At this point, company 1 (the industry leader) has to decide whether or not to develop an application. Suppose that the Nash equilibrium for the subgame in which company 1 does not develop an application is (*Develop*, *Develop*). Replacing the two final subgames with the Nash equilibrium payoffs we found in our previous discussion, the situation is as depicted in the tree of Fig. 4.20. In particular, if firm 1 develops an application, then outcome (D, D) which entails payoffs (10, 4, 4, 4), as depicted in the terminal node at the bottom left-hand side of Fig. 4.20. If, in contrast, firm 1 does not develop an application for

		Firm 3	
		Develop	Do not develop
Firm 2	Develop	$\underline{2}$, $\underline{2}$	-2, 0
	Do not develop	0, -2	$\underline{0}$, $\underline{0}$

Fig. 4.19 Smallest proper subgame (after firm 1 does not develop)

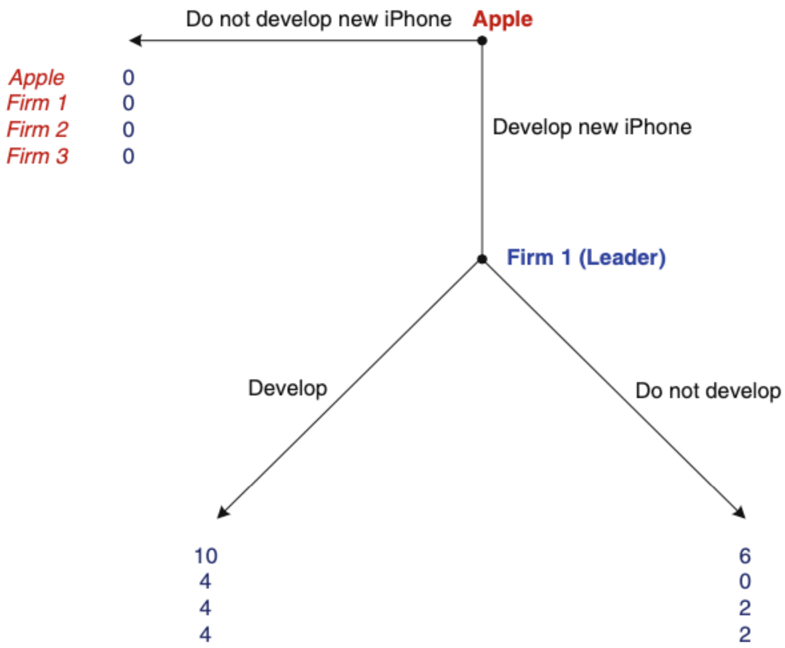

Fig. 4.20 Extensive-form subgame (Case I)

the new iPhone, then outcome (D, D) arises, yielding payoffs of (6, 0, 2, 2), as depicted in the bottom right–hand corner of Fig. 4.20.

We can now analyze firm 1's decision. If company 1 develops an application, then its payoff is 4, while its payoff is only 2 (since it anticipates the followers developing apps) from not doing so. Hence, company 1 chooses *Develop*.

Company 1—Case II. Now suppose the Nash equilibrium of the game that arises after firm 1 does not develop an app has neither firm 2 or 3 developing an app. Replacing the two final subgames with the Nash equilibrium payoffs we found in our previous discussion, the situation is as depicted in Fig. 4.21. Specifically, if firm 1 develops, outcome (D, D) emerges, which entails payoffs (10, 4, 4, 4); while if firm 1 does not develop firm 2 and 3 respond not developing apps either, ultimately yielding a payoff vector of (−6, 0, 0, 0). In this setting, if firm 1 develops an application, its payoff is 2; while its payoff is only 0 from not doing so. Hence, firm 1 chooses *Develop*.

Thus, regardless of which Nash equilibrium is used in the subgame initiated after firm 1 chooses *Do not develop* (in the right-hand side of the game in Figs. 4.20 and 4.21), firm 1 (the leader) optimally chooses to *Develop*.

First mover (Apple). Operating by backward induction, we now consider the first mover in this game (Apple). If Apple chooses to develop the new iPhone, then, as previously derived, firm 1 develops an application and this induces all followers 2 and 3 to do so as well. Hence, Apple's payoff is 10 from introducing the new iPhone. It is then optimal for Apple to develop the new iPhone, since its payoffs from so doing, 10, is larger than from not developing it, 0. Intuitively, since Apple anticipates all app developers will react introducing new apps, it finds the initial introduction of

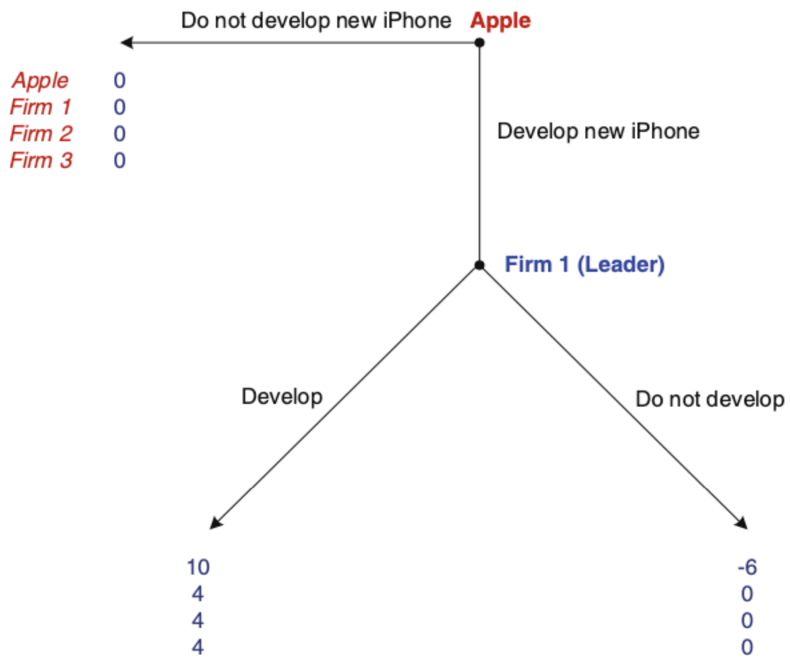

Fig. 4.21 Extensive-form subgame (Case II)

the iPhone to be very profitable. We can then identify two subgame perfect Nash equilibria (where a strategy for firm 2, as well as for firm 3, specifies a response to firm 1 choosing *Develop* and a response to company 1 choosing *Do not develop*[3]):

(*Develop iPhone, Develop, Develop/Develop, Develop/Develop*), and

(*Develop iPhone, Develop, Develop/Do not develop, Develop/Do not develop*).

Note that both SPNE result in the same equilibrium path, whereby, first, Apple introduces the new iPhone, the industry leader (firm 1) subsequently chooses to develop applications for the new iPhone, and finally firms 2 and 3 (observing firm 1's apps development) simultaneously decide to develop apps as well.

9. Cournot Competition with Asymmetric Costs: Consider a duopoly game in which firms compete a la Cournot, i.e., simultaneously and independently selecting their output levels. Assume that firm 1 and 2's constant marginal costs of production differ, i.e., $c1 > c2$. Assume also that the inverse demand function is $p = a - bQ$, with $a > c1$ and $Q = q1+q2$
    a. Find the pure strategy Nash equilibrium of this game. Under what conditions does it yield corner solutions (only one firm producing in equilibrium)?
    b. In interior equilibria (both firms producing positive amounts), examine how do equilibrium outputs vary when firm 1's costs change.

Answer

Part (a) In a Nash equilibrium (q_1^*, q_2^*), firm 1 maximizes its profits by selecting the output level q_1 that solves

$$\max_{q_1 \geq 0, q_2^*} \pi_1(q_1, q_2^*) = (a - b(q_1 + q_2^*))q_1 - c_1 q_1 = aq_1 - bq_1^2 - bq_2^* q_1 - c_1 q_1$$

where firm 1 takes the equilibrium output of firm 2, q_2^*, as given. Similarly for firm 2, which maximizes

$$\max_{q_2 \geq 0, q_1^*} \pi_2(q_1^*, q_2) = (a - b(q_1^* + q_2))q_2 - c_2 q_2 = aq_2 - bq_1^* q_2 - bq_2^2 - c_2 q_2$$

which takes the equilibrium output of firm 1, q_1^*, as given. Taking first order conditions with respect to q_1 and q_2 in the above profit-maximization problem, we obtain:

$$\frac{\partial \pi_1(q_1, q_2^*)}{\partial q_1} = a - 2bq_1 - bq_2^* - c_1 \tag{5.1}$$

$$\frac{\partial \pi_2(q_1^*, q_2)}{\partial q_2} = a - bq_1^* - 2bq_2 - c_2. \tag{5.2}$$

Solving for q_1 in (5.1) we obtain firm 1's best response function:

$$q_1(q_2) = \frac{a - bq_2 - c_1}{2b} = \frac{a - c_1}{2b} - \frac{q_2}{2}$$

(Note that, as usual, we rearranged the expression of the best response function to obtain two parts: the vertical intercept, $\frac{a-c_1}{2b}$, and the slope, $-\frac{1}{2}$). Similarly, solving for q_2 in (5.2) we find firm 2's best response function

$$q_2(q_1) = \frac{a - bq_1 - c_2}{2b} = \frac{a - c_2}{2b} - \frac{q_1}{2}.$$

Plugging firm 1's best response function into that of firm 2, we obtain:

$$q_2^* = \frac{a - c_2}{2b} - \frac{1}{2}\left(\frac{a - c_1}{2b} - \frac{q_2}{2}\right)$$

Since this expression only depends on q_2, we can now solve for q_2 to obtain the equilibrium output level for firm 2, $q_2^* = \frac{a+c_1-2c_2}{3b}$. Plugging the output we found q_2^* into firm 1's best response function $q_1 = \frac{a-2bq_2-c_2}{b}$, we obtain firm 1's equilibrium output level:

$$q_1^* = \frac{a - 2b\left[\frac{a+c_1-2c_2}{3b}\right] - c_2}{b} = \frac{a + c_2 - 2c_1}{3b}$$

Hence, for $q_1^* = \frac{a+c_2-2c_1}{3b} \leq 0$ we need firm 1's costs to be sufficiently high, i.e., solving for c_1 we obtain $\frac{a+c_2}{2} \leq c_1$. A symmetric condition holds for firm 2's output. In particular, $q_2^* = \frac{a+c_1-2c_2}{3b} \leq 0$ arises if $\frac{a+c_1}{2} \leq c_2$. Therefore, we can identify three different cases (two giving rise to corner solutions, and one providing an interior solution):

(1) if $c_i \geq \frac{a+c_j}{2}$ for all firm $i = \{1,2\}$, then both firms costs are so high that no firm produces a positive output level in equilibrium;
(2) if $c_i \geq \frac{a+c_j}{2}$ but $c_j < \frac{a+c_i}{2}$, then only firm j produces a positive output (intuitively, firm j would be the most efficient company in this setting, thus leading it to be the only producer); and
(3) if $c_i \leq \frac{a+c_j}{2}$ for all firm $i = \{1,2\}$, then both firms produce positive output levels.

Part (b) In order to know how the (q_1^*, q_2^*) varies when (c_1, c_2) change we separately differentiate each (interior) equilibrium output level with respect to c_1 and c_2, as follows:

$$\frac{\partial q_1^*}{\partial c_1} = -\frac{2}{3b} < 0 \text{ and } \frac{\partial q_1^*}{\partial c_2} = \frac{1}{3b} > 0$$

Hence, firm 1's equilibrium output decreases as its own costs go up, but increases as its rival's costs increase. A similar intuition applies to firm 2's equilibrium output:

$$\frac{\partial q_2^*}{\partial c_2} = -\frac{2}{3b} < 0 \text{ and } \frac{\partial q_2^*}{\partial c_1} = \frac{1}{3b} > 0.$$

10. Collusion when N firms compete in quantities: Consider n firms producing homogenous goods and choosing quantities in each period for an infinite number of periods. Demand in the industry is given by p = 1 − Q, Q being the sum of individual outputs. All firms in the industry are identical: they have the same constant marginal costs c < 1, and the same discount factor δ. Consider the following trigger strategy:

- Each firm sets the output q^m that maximizes joint profits at the beginning of the game, and continues to do so unless one or more firms deviate.
- After a deviation, each firm sets the quantity q, which is the Nash equilibrium of the unrepeated Cournot game.

a. Find the condition on the discount factor that allows for collusion to be sustained in this industry.
b. Indicate whether a larger number of firms in the industry facilitates or hinders the possibility of reaching a collusive outcome.

Answer

Part (a) ***Cooperation***: First, we need to find the quantities that maximize joint profits $\pi = (1 - Q)Q - cQ$. This output level coincides with that under monopoly, i.e., $Q = \frac{1-c}{2}$, yielding aggregate profits of

$$\pi = \left(1 - \frac{1-c}{2}\right)\frac{1-c}{2} - c\frac{1-c}{2} = \frac{(1-c)^2}{4}$$

for the cartel. Therefore, at the symmetric equilibrium, individual quantities are $q^m = \frac{1}{n}Q = \frac{1}{n}\frac{1-c}{2}$ and individual profits under the collusive strategy are $\pi^m = \frac{(1-c)^2}{4n}$.

Deviation: As for the deviation profits, the optimal deviation by firm i, q^{dev}, is given by solving

$$q^d(q^m) = \text{argmax}_q[1 - (n-1)q^m - q - c]q.$$

Intuitively, the above expression represents that firm i selects the output level that maximizes its profits given that all other $(n - 1)$ firms are still respecting the collusive agreement, thus producing q^m. In particular, the value of q that maximizes the above expression is obtained through the first order conditions:

$$[1 - (n-1)q^m - 2q - c] = 0$$

$$1 - (n-1)\left(\frac{1}{n}\frac{1-c}{2}\right) - 2q - c = 0$$

which, solving for q, yields,

$$q^d = (n+1)\frac{(1-c)}{4n}.$$

The profits that a firm obtains by deviating from the collusive output are hence

$$\pi^d = \left[1-(n-1)q^m - q^d(q^m)\right] \times q^d(q^m) - c \times q^d(q^m)$$
$$= \left[1-(n-1)\left(\frac{1}{n}\frac{1-c}{2}\right) - (n+1)\frac{1-c}{4n}\right](n+1)\frac{1-c}{4n} - c(n+1)\frac{1-c}{4n},$$

which simplifies to $\pi^d = \frac{(1-c)^2(n+1)^2}{16n^2}$.

Incentives to collude: Given the above profits from colluding and from deviating, every firm i chooses to collude as long as

$$\frac{1}{1-\delta}\pi^m \geq \pi^d + \frac{\delta}{1-\delta}\pi^c$$

which, solving for δ, yields

$$\delta \leq \frac{\pi^m - \pi^d}{\pi^c - \pi^d}$$

and multiplying both sides by -1, we obtain

$$\delta \geq \frac{\pi^d - \pi^m}{\pi^d - \pi^c}$$

Intuitively, the numerator represents the profit gain that a firm experiences when it unilaterally deviates from the cartel agreement, $\pi^d - \pi^m$. When such profit gain increases, cartel becomes more difficult to sustain, i.e., it is only supported for higher discount factors. Plugging the expressions of profits π^d, π^m, and π^c we previously found, yields a cutoff discount factor of

$$\delta \geq \frac{(1+n)^2}{1+6n+n^2}$$

For compactness, we denote this ratio as

$$\frac{(1+n)^2}{1+6n+n^2} \equiv \delta^c$$

Hence, under punishment strategies that involve a reversion to Cournot equilibrium forever after a deviation takes place, tacit collusion arises if and only if firms are sufficiently patient. Figure 6.4 depicts cutoff δ^c, as a function of the number of firms, n, and shades the region of δ that exceeds such a cutoff.

Differentiating the critical threshold of the discount factor, δ^c, with respect to the number of firms, n, we find that

$$\frac{\partial \delta}{\partial n} = \frac{4(n^2-1)}{(1+6n+n^2)} > 0$$

Intuition: Other things being equal, as the number of firms in the agreement increases, the more difficult it is to reach and sustain collusion in the cartel agreement, i.e., as the market becomes less concentrated collusion becomes less likely.

Fig. 6.4 Critical threshold of $\delta^c(n)$