

UNIT 3

Nash Equilibrium

A profile of strategies in which nobody wants to change their mind

Image: Jan Collsiöö, via Wikimedia Commons (public domain)

By the end of this unit, I should be able to answer these questions:

- What is a Nash equilibrium, and in what sense is it a stable outcome?
- How do I find every pure strategy equilibrium of a matrix game?
- Why does the stag hunt have two equilibria, and what decides which one is played?
- How do two competing producers choose their quantities in Cournot's model, and why do they produce more than a monopolist and less than competitive firms?
- Why does price competition between two identical firms drive the price down to cost?
- Does every game have a Nash equilibrium, and can one always be found quickly?

In the prisoner's dilemma of Unit 2 the prediction was easy, because each suspect had a strategy that was best whatever the other did. Most games are not so convenient. Two drivers who meet on a narrow road, two firms that choose how much to produce, two friends who have lost each other in a city: none of them has a dominant strategy, and yet each of these situations settles down somewhere. This unit introduces the idea that says where. A Nash equilibrium is a combination of strategies in which nobody wants to change their mind, and almost every later unit refines, tests, or computes it.

1 The limits of dominance

Two reminders come first, because this unit is read by visitors as well as by students who have Units 1 and 2 behind them. A **payoff matrix** lists the strategies of the row player down the side and those of the column player across the top, and every cell holds two numbers, the row player's payoff first, as in "3, 3"; each player wants their own number to be large. A **best response** is a strategy that gives a player the highest payoff against a given strategy of the other player, so that each column has a best response for the row player and each row has one for the column player. Unit 2 built dominance on these ideas: a strategy is dominant when it is a best response to everything the other player might do, and when every player has one, the game is solved.

To see how rarely that happens, take the two drivers of Unit 1 once more. Two cars approach each other on a road without markings, and each driver chooses, at the same moment, to keep left or to keep right. If both choose the same side, they pass, and each gets a payoff of 1. If they choose different sides, they have to brake hard, and each gets 0. Table 1 shows the matrix, with the row player's payoff written first.

Neither driver has a dominant strategy, because in Table 1 Left is the best response to Left and Right is the best response to Right. The iterated elimination of Unit 2 removes nothing either, since no strategy is dominated. Dominance therefore says nothing about this game, although anyone can see that it has two stable outcomes, in which both keep left or both keep right, and two unstable ones,

in which the drivers choose different sides. What is missing is a precise statement of what "stable" means, one that works equally for the drivers, for firms choosing among thousands of quantities, and for games with a dozen players.

	Left	**Right**
Left	1, 1	0, 0
Right	0, 0	1, 1

Table 1. The two drivers' game. The first driver chooses the row and the second the column; the first number in each cell is the row player's payoff.

2 The definition of a Nash equilibrium

The idea is to ask the stability question of every player at once. Suppose each player has settled on a strategy and knows what the others have settled on. If no single player can do better by switching to another strategy while the others keep theirs, the combination holds. JOHN FORBES NASH JR. (1928-2015) wrote this condition down in 1950, when he was a doctoral student at Princeton, and proved that it can always be met in any game with finitely many players and finitely many strategies, once players may choose at random among their strategies. Section 6 returns to his theorem; the definition comes first.

The notation is that of Unit 1. There are n players, player i chooses a strategy s_i from a set S_i, a complete list of choices $s = (s_1, \ldots, s_n)$ is called a **strategy profile**, and s_{-i} stands for the choices of everyone except player i, so that a profile can be written (s_i, s_{-i}). The payoff of player i is $u_i(s)$.

KEY RESULT Nash equilibrium

A strategy profile $s^* = (s_1^*, \ldots, s_n^*)$ is a **Nash equilibrium** if, for every player i and every strategy s_i in S_i,

$$u_i(s_i^*, s_{-i}^*) \;\geq\; u_i(s_i, s_{-i}^*).$$

In words, every player's strategy is a best response to the strategies of the others; if $\mathrm{BR}_i(s_{-i})$ denotes the set of player i's best responses to s_{-i}, the condition reads $s_i^* \in \mathrm{BR}_i(s_{-i}^*)$ for all i. If every inequality is strict for all $s_i \neq s_i^*$, the equilibrium is called **strict**. The definition applies to any number of players and any strategy sets, finite or not.

The inequality compares one player's payoff in the equilibrium with the payoff from a change made by that player alone. Nothing is said about what happens when two or more players change together; questions of that kind need the theory of coalitions, which lies beyond this unit. The definition can be read in three ways, and each reading is useful in a different situation.

1. *No regrets.* After the game has been played and everyone sees what the others did, nobody wishes they had chosen differently. The two drivers who both kept left have no regrets; the two who chose different sides and had to brake do.
2. *A self-enforcing agreement.* If the players had agreed on the profile s^* beforehand, nobody would have a reason to break the agreement in secret, so it needs no contract and no referee. An agreement on a profile that is not an equilibrium is worthless, because at least one player gains by quietly doing something else.
3. *Consistent expectations.* Each player expects something of the others, plays a best response to it, and turns out to be right. An equilibrium is the only kind of outcome in which everyone can be right about everyone else at the same time.

Three warnings belong next to the definition. First, an equilibrium need not be good for the players.

The prisoner's dilemma of Unit 2 has exactly one Nash equilibrium, mutual defection with a payoff of 1 each, although both prefer mutual cooperation with 3 each; the reason is that a player who switches to cooperation alone only lowers their own payoff, from 1 to 0. Second, a game can have several equilibria, as the drivers' game shows, and the definition does not say which one will be played; Section 4 is about that question. Third, the condition uses "at least as good", not "better". A strategy that is weakly dominated in the sense of Unit 2 can therefore be part of a Nash equilibrium, because being tied for best is enough; a strictly dominated strategy never can be, since by definition some other strategy does strictly better against everything, including the equilibrium strategies of the others.

The connection with Unit 2 also runs the other way: a dominant strategy is a best response to every profile of the others, so a profile of dominant strategies is always a Nash equilibrium, and when the dominant strategies are strict it is the only one.

3 Every pure equilibrium of a matrix game

A strategy chosen without any randomization is called a **pure strategy**, and this unit looks for equilibria in pure strategies only; Unit 4 adds strategies that mix. In a payoff matrix the definition turns into a mechanical procedure, which is best done with a pencil.

1. For each column, find the row player's largest payoff in that column and underline it. These are the row player's best responses to each strategy of the column player. If two payoffs tie for largest, underline both.
2. For each row, find the column player's largest payoff in that row and underline it. These are the column player's best responses.
3. A cell in which both payoffs are underlined is a pure strategy Nash equilibrium, because each player's strategy is then a best response to the other's. Every pure equilibrium is found this way, and nothing else is.

Figure 1 shows the procedure on four games that every game theorist knows by name, written with the payoffs this masterclass uses throughout. In the prisoner's dilemma only the cell (Defect, Defect) carries two underlines, as the dominance argument of Unit 2 already implied. The other three games have no dominant strategies, and underlining is how their equilibria are found.

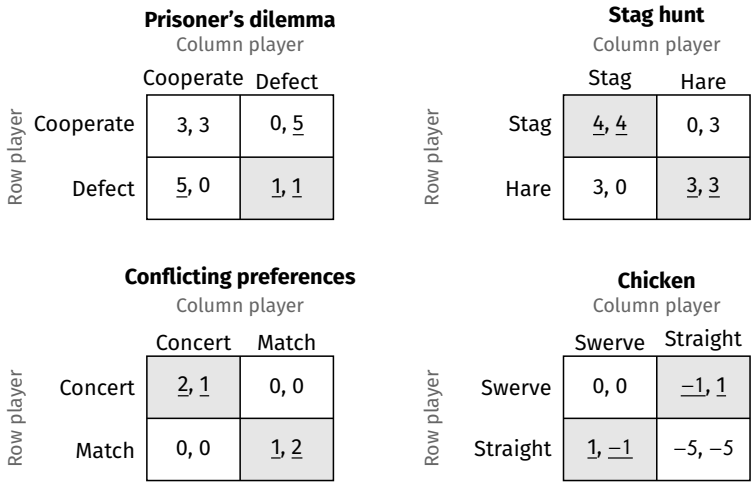


Figure 1. Four classic games solved by underlining. A line under a payoff marks a best response: under the row player's payoff, the best reply to that column; under the column player's payoff, the best reply to that row. The shaded cells carry two underlines each and are the pure strategy Nash equilibria. The row player's payoff is written first in every cell.

WORKED EXAMPLE **The stag hunt and the coordination game with conflicting preferences**

The stag hunt. Two hunters go out together. A stag can only be caught if both hunt it, and then each gets 4; a hare can be caught alone and is worth 3 to whoever catches it; a hunter who waits for the stag while the other chases a hare goes home with nothing. The matrix in Figure 1 therefore has 4, 4 for two stag hunters, 0, 3 and 3, 0 when one defects to the hare, and 3, 3 when both do. Against Stag the row player compares 4 with 3 and underlines the 4; against Hare the row player compares 0 with 3 and underlines the 3; the column player's choices are the mirror image. Two cells end up with two underlines, (Stag, Stag) and (Hare, Hare); in (Stag, Hare) the row player, who gets 0, would switch to Hare and get 3.

The coordination game with conflicting preferences. Two friends want to spend the evening together, one prefers a concert and the other a football match, and their phones are dead. Being together at the concert gives 2 to the first and 1 to the second, being together at the match gives 1 and 2, and being apart gives 0 to both. Against Concert the row player underlines the 2, against Match the 1; the column player underlines the 1 in the first row and the 2 in the second, and again two cells qualify. The game is usually called the battle of the sexes, a name from the 1950s; this masterclass uses the descriptive name instead.

The two games each have two pure strategy Nash equilibria, both on the diagonal, and neither game says which of the two will be played.

The stag hunt goes back to Jean-Jacques Rousseau (1712-1778), whose book *Discourse on the Origin and Basis of Inequality among Men* of 1755 describes hunters who agree to wait for a stag but cannot resist a passing hare. Chicken, the fourth game of Figure 1, is the game of two drivers racing towards each other. The driver who swerves is the chicken and gets -1, while the other gets 1; if both swerve, nothing happens and both get 0; if neither swerves, the cars crash, which is written as -5 for each. Its two equilibria are the cells in which exactly one driver swerves, because the driver who swerves would fall from -1 to -5 by going straight, and the driver who goes straight would fall from 1 to 0 by swerving. The same structure returns in Unit 8 as the hawk-dove game. Table 2 collects the four games with a fifth, discussed next.

Game	The Story	Pure Equilibria
Prisoner's dilemma	Confess or stay silent	(Defect, Defect)
Stag hunt	Hunt together or alone	(Stag, Stag), (Hare, Hare)
Conflicting preferences	Together, but where?	(Concert, Concert), (Match, Match)
Chicken	Who swerves first?	(Swerve, Straight), (Straight, Swerve)
Matching pennies	Match or mismatch	none

Table 2. Five classic two-player games and their equilibria in pure strategies. The first four are drawn in Figure 1; the fifth is shown in Table 4.

Underlining also works where the dominance method of Unit 2 stops, as the game of Table 3 and the example below show.

	Left	Centre	Right
Top	3, 4	1, 0	2, 5
Middle	0, 3	3, 5	3, 1
Bottom	1, 4	1, 0	4, 2

Table 3. A 3 by 3 game without dominated strategies and with a single pure strategy Nash equilibrium. The row player's payoff is written first.

WORKED EXAMPLE A 3 by 3 game that dominance cannot solve

Table 3 shows a game in which the row player chooses Top, Middle, or Bottom and the column player Left, Centre, or Right, with the row player's payoff written first. Column by column, the row player's best responses are Top against Left (3 beats 0 and 1), Middle against Centre (3 beats 1 and 1), and Bottom against Right (4 beats 2 and 3). Row by row, the column player's best responses are Right against Top (5 beats 4 and 0), Centre against Middle (5 beats 3 and 1), and Left against Bottom (4 beats 0 and 2).

Only one cell is a best response for both players, namely (Middle, Centre) with payoffs 3, 5. Every other cell fails for at least one player. For example, (Top, Left) gives the row player 3, but the column player would move to Right and get 5 instead of 4; from (Top, Right) the row player would move to Bottom, from (Bottom, Right) the column player to Left, and from (Bottom, Left) the row player back to Top. The best responses chase each other around the outside of the matrix and come to rest only in the centre.

No strategy is dominated here. Each row strategy is a best response to some column (Top to Left, Middle to Centre, and Bottom to Right), and each column strategy is a best response to some row, and a strategy that is a best response to something cannot be strictly dominated, because the strategy that dominated it would have to do better against that same choice too. Unit 2's method therefore cannot remove any of them. The unique pure strategy Nash equilibrium is (Middle, Centre) with payoffs 3 to the row player and 5 to the column player.

The procedure can also come back empty. In *matching pennies* each of two players secretly places a coin heads or tails up. If the coins match, the row player takes both and gets 1 while the column player gets -1; if they differ, the column player wins instead, so the matrix, shown in Table 4, has $1, -1$ in the two cells where the coins match and $-1, 1$ in the two where they differ. The row player's best responses are Heads against Heads and Tails against Tails, the column player's are Tails against Heads and Heads against Tails, and no cell collects two underlines. Whatever the two players settle on, one of them wants to switch, so the game has no equilibrium in pure strategies. Nash's theorem, stated in Section 6, says that every finite game has an equilibrium once players are allowed to randomize, and Unit 4 shows that playing Heads and Tails with probability $\frac{1}{2}$ each is the equilibrium of this game, and that the same idea explains how a penalty kicker and a goalkeeper behave.

	Heads	**Tails**
Heads	$1, -1$	$-1, 1$
Tails	$-1, 1$	$1, -1$

Table 4. Matching pennies. The row player wins when the coins match and the column player when they differ; the row player's payoff is written first. No cell holds a best response for both players.

CHECK YOUR UNDERSTANDING Question 1

In a two-player game, Up is the row player's best response to Left and Down is the best response to Right. Left is the column player's best response both to Up and to Down. Which profile is a pure strategy Nash equilibrium?

(A) (Down, Right)
(B) (Up, Left)
(C) Both (Up, Left) and (Down, Right)
(D) Neither of them

The solution is at the end of the unit.

4 Several equilibria and the choice between them

The drivers' game, the stag hunt, the coordination game with conflicting preferences, and chicken each have two pure equilibria, and the definition is silent about which one is played. Game theory has three answers.

4.1 Focal points

THOMAS SCHELLING (1921-2016) argued in his book *The Strategy of Conflict*, published in 1960, that players who must coordinate without communicating look for a **focal point**, an equilibrium that stands out from the others for some reason that both players can expect the other to notice: a tradition, a round number, a symmetry, a landmark. The reason has nothing to do with the payoffs, which are often identical across the equilibria; a focal point works because each player expects the other to be drawn to the same prominent option, and the expectation makes itself true.

WORKED EXAMPLE Schelling's questionnaire

In the section on tacit coordination of *The Strategy of Conflict* (pages 54 to 58 of the 1960 edition), Schelling reports the answers of an informal sample of about 40 people in New Haven to a list of coordination problems; each respondent was told to make the same choice as an unknown partner who was trying to do the same.

Asked to name heads or tails, 36 of 42 respondents chose heads and 6 chose tails, far more agreement than chance would give, since random choices would have split the 42 into about 21 and 21. Asked to meet a stranger somewhere in New York City on a given day, with no place and no hour agreed, an absolute majority, more than half of them, chose the information booth at Grand Central Station, and virtually all of them chose twelve noon. Asked to divide 100 dollars into two piles so that the division matched their partner's, 36 of 41 split the money fifty-fifty. In each problem the game has many equilibria with the same payoff, and the respondents picked out one of them by its prominence alone: people coordinate on focal points far more often than chance allows. Schelling shared the 2005 Nobel prize in economic sciences with ROBERT AUMANN (Hebrew University of Jerusalem).

4.2 Conventions

The second answer looks at history rather than at the matrix. A **convention** is an equilibrium that everyone plays because everyone has always played it, so that each player's expectation about the others rests on experience rather than on reasoning about the matrix. The side of the road is the standard example. Both equilibria of the drivers' game in Table 1 give the same payoffs; which one a country plays was settled long ago by habit or by a law that fixed a habit, and once in place the convention is extremely stable, because a single driver who tries the other side loses at once. It can nevertheless be changed, provided everyone changes together.

WORKED EXAMPLE Sweden's change of equilibrium on 3 September 1967

Until 1967 Sweden drove on the left, although its neighbours Norway and Finland drove on the right and most Swedish cars had their steering wheel on the left. In a national referendum on 16 October 1955, 82.9 per cent of the votes were for keeping left-hand traffic and 15.5 per cent for changing, with a turnout of 53 per cent, according to Statistics Sweden's table of national referendums. The change was unpopular, and the convention was perfectly stable, as the drivers' game predicts, since no driver and no small group could gain by moving to the right alone.

The Swedish parliament nevertheless decided on 10 May 1963 to change sides, and after four years of preparation the switch took place on Sunday 3 September 1967, known in Sweden as Dagen H, the H standing for högertrafik, right-hand traffic: all vehicles stopped at 04:50 in the morning and moved carefully to the right, and at 05:00 traffic resumed on the new side, as in the photograph at the head of this unit. Because every driver changed at the same moment, the country moved from one equilibrium of the drivers' game to the other without passing through the unstable cells in between, and since that morning keeping right has been each driver's best response. The episode shows that a convention is changed by a coordinated move of all players at once, never by individual deviation.

4.3 Payoff dominance and risk dominance

The third answer looks at the payoffs themselves. In the stag hunt, (Stag, Stag) gives both hunters 4 and (Hare, Hare) gives both 3, so both prefer the first equilibrium; an equilibrium that every player prefers to another is called **payoff dominant**. If the players could rely on each other, they would hunt the stag. The difficulty is that hunting the stag is only a good idea when the other hunter does the same, and a hunter who is not sure of that may prefer the safe hare. John Harsanyi (1920-2000) and Reinhard Selten (1930-2016) made this trade-off precise in their book *A General Theory of Equilibrium Selection in Games* of 1988 and called the safer equilibrium **risk dominant**.

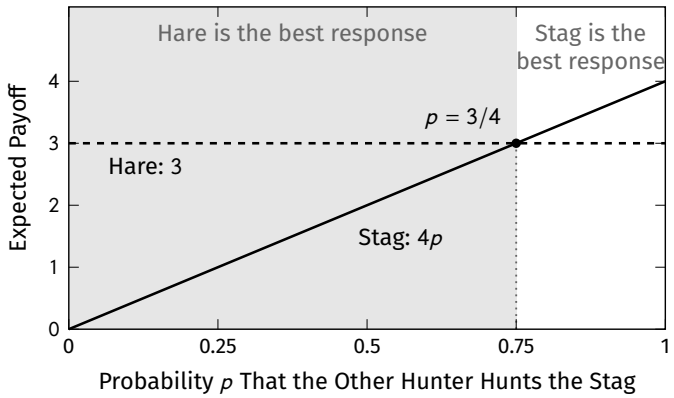


Figure 2. The stag hunt from one hunter's point of view. The expected payoff of hunting the stag, $4p$, rises with the probability p that the other hunter does the same, while the hare pays 3 regardless. The lines cross at $p = 3/4$, and the shaded part of the belief interval, three quarters of it, is where the hare is the best response. Computed from the payoffs of Figure 1.

To see what risk means here, suppose the row player believes that the other hunter will hunt the stag with probability p and the hare with probability $1 - p$. The expected payoff of Stag, which is the payoff averaged with these probabilities, is then $4p + 0\,(1 - p) = 4p$, and the expected payoff of Hare is $3p + 3\,(1 - p) = 3$ whatever p is. Stag is the better choice when $4p \geq 3$, that is, when $p \geq \frac{3}{4}$, where the two lines of Figure 2 cross. A hunter needs to be at least 75 per cent sure of the partner before the stag is worth the risk, whereas the hare is the better choice for every belief below that. The range of possible beliefs, from $p = 0$ to $p = 1$, is called the **belief interval**, and the hare's part of it is three times as large as the stag's. That asymmetry is what risk dominance measures.

CHECK YOUR UNDERSTANDING Question 2

In a stag hunt, hunting the stag together gives 5 to each hunter, hunting the stag alone gives 0, and hunting the hare gives 3 whatever the other does. What is the smallest probability of the other hunter choosing the stag that makes the stag a best response?

(A) 0.5
(B) 0.75
(C) 0.6
(D) 0.4

The solution is at the end of the unit.

KEY RESULT Risk dominance in a 2 by 2 coordination game

Let a game have two pure strategy equilibria, (A, A) and (B, B). The loss a player suffers by deviating alone from an equilibrium is called that player's **deviation loss**. The equilibrium (A, A) *risk dominates* (B, B) if the product of the two players' deviation losses at (A, A) is larger than the product at (B, B):

$$\big(u_1(A,A) - u_1(B,A)\big)\big(u_2(A,A) - u_2(A,B)\big) \;>\; \big(u_1(B,B) - u_1(A,B)\big)\big(u_2(B,B) - u_2(B,A)\big).$$

The equilibrium with the larger product is a best response against the wider range of beliefs about the other player. In a symmetric game both players have the same deviation loss at each equilibrium, so the rule compares the squares of the two losses, and with losses d_A at (A, A) and d_B at (B, B) the strategy A is the best response on a share $d_A/(d_A+d_B)$ of the belief interval; comparing the products is therefore the same as comparing the two regions of Figure 2.

For the stag hunt the arithmetic is short. At (Stag, Stag) a hunter who deviates to Hare falls from 4 to 3, a loss of 1, so the product is $1 \times 1 = 1$. At (Hare, Hare) a hunter who deviates to Stag falls from 3 to 0, a loss of 3, and the product is $3 \times 3 = 9$. Since $9 > 1$, (Hare, Hare) is the risk dominant equilibrium and (Stag, Stag) the payoff dominant one, and the two criteria point in opposite directions; the shares of the belief interval are $1/(1+3) = 1/4$ for the stag and $3/(1+3) = 3/4$ for the hare, the quarter and three quarters of Figure 2. Which one real players follow depends on how well they know each other and on how often they have played; Unit 6 returns to the stag hunt when the same players meet repeatedly.

5 Cournot competition

All the games so far had two or three strategies per player. Many of the most useful games have infinitely many, because the strategy is a number: a price, a quantity, a bid, a share. Underlining is impossible in such a game, but the definition of Section 2 applies unchanged, and the first game of this kind, the competition between the two sellers of one market, which is called a **duopoly**, was solved more than a century before Nash wrote down the general idea.

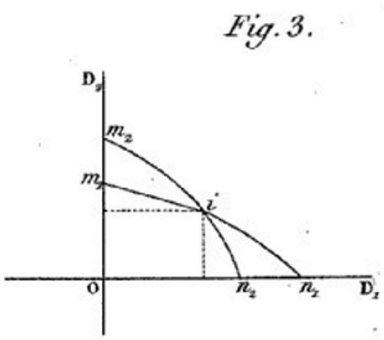


Figure 3. Cournot's figure from the part of the *Recherches* of 1838 headed "Of the competition of producers". Each curve gives one owner's best quantity against the other's, and the solution of the duopoly lies where they cross, which is the construction of Figure 4. Credit: Antoine Augustin Cournot, via Wikimedia Commons (public domain).

EXCURSION ***Cournot and the two mineral springs***

ANTOINE AUGUSTIN COURNOT (1801-1877), a French mathematician, published his book *Recherches sur les principes mathématiques de la théorie des richesses* (Researches into the Mathematical Principles of the Theory of Wealth) in 1838. Its seventh part, headed "On the competition of producers", imagines two owners of neighbouring mineral springs who sell identical water to the same customers and each decide, independently, how much to bring to market. Cournot worked out how much each owner would sell if neither could gain by changing their own quantity given the other's, and drew the two curves whose crossing gives the answer (Figure 3). That condition is the Nash equilibrium of the quantity game, found 112 years before Nash, which is why the outcome is often called a Cournot-Nash equilibrium. The book was almost ignored until JOSEPH BERTRAND (1822-1900) reviewed it in the *Journal des Savants* in 1883; his objection that firms set prices rather than quantities became a model of its own, treated in Section 5.2.

5.1 The model and the best-response function

The game is set up as follows, in the form this masterclass uses in Units 3, 5, and 6. Two firms, and only two, produce the same good. Each firm i chooses a quantity $q_i \geq 0$, and the two choices are made at the same time, without knowing the other's. The market price depends on the total quantity through the inverse demand function

$$P = a - b\,(q_1 + q_2), \tag{1}$$

with $a > 0$ and $b > 0$, which says that every extra item brought to market lowers the price by b; the price is zero if the total reaches a/b. A demand function gives the quantity that buyers take at each price; this one is called inverse because it gives the price at each total quantity instead. Each item costs c to produce, with $c < a$, and there are no other costs. The payoff of firm i is its profit, revenue minus cost,

$$\pi_i(q_1, q_2) = \big(P - c\big)\, q_i = \big(a - c - b\, q_1 - b\, q_2\big)\, q_i. \tag{2}$$

The program's numbers are $a = 120$, $b = 1$, and $c = 30$, so that $P = 120 - (q_1 + q_2)$, which means that the price would be 120 if nothing were sold and falls by 1 for every item brought to market, while each item costs 30 to make. This kind of game, with a few firms, an identical product, and quantities chosen at the same time, is called **Cournot competition**.

To find the equilibrium, the best responses are needed first. Firm 1's profit for a given q_2 is a quadratic function of q_1, and because the coefficient of q_1^2 in Equation (2) is $-b$, a negative number, its graph is a parabola opening downwards, so the best response is the vertex of that parabola, its highest point. The vertex can be found by completing the square, a method from school algebra; the Key Result states what comes out, and the Proof box derives it in full.

KEY RESULT **Cournot best response and equilibrium**

With inverse demand $P = a - b\,(q_1 + q_2)$ and a cost of c per item for both firms,

$$\mathrm{BR}_i(q_j) = \frac{a - c - b\,q_j}{2b}, \qquad q_1^* = q_2^* = \frac{a - c}{3b}, \qquad P^* = \frac{a + 2c}{3}.$$

Each firm's best quantity falls by half an item for every item the rival adds. In equilibrium the two firms together sell two thirds of the competitive quantity $(a - c)/b$, the quantity at which the price has fallen to the cost c, and the price lies between that competitive price and the monopoly price $(a + c)/2$ that a single seller would set; the example below gives all three with numbers. The formulas hold for two firms with the same constant cost per item; the interactive element of this section also computes different costs.

PROOF **The Cournot best responses and the equilibrium by completing the square**

Fix the quantity q_2 of firm 2 and write $m = a - c - b\,q_2$ for the profit firm 1 would make per item if it sold a single item. By Equation (2), firm 1's profit is

$$\pi_1 = (m - b\,q_1)\,q_1 = m\,q_1 - b\,q_1^2.$$

Completing the square turns the right-hand side into a constant minus a square, since $(q_1 - m/(2b))^2 = q_1^2 - (m/b)\,q_1 + m^2/(4b^2)$ and the extra last term, multiplied by $-b$, is cancelled by adding $m^2/(4b)$:

$$\pi_1 = -b\left(q_1^2 - \frac{m}{b}\,q_1\right) = -b\left(q_1 - \frac{m}{2b}\right)^2 + \frac{m^2}{4b}.$$

The square is never negative, so the profit is largest when the square is zero, that is, when $q_1 = m/(2b)$, and the largest profit is $m^2/(4b)$. Writing m out again gives the best-response function of firm 1,

$$q_1 = \mathrm{BR}_1(q_2) = \frac{a - c - b\,q_2}{2b},$$

as long as $a - c - b\,q_2 \geq 0$; if firm 2 floods the market beyond that, firm 1's best response is to produce nothing. The game is symmetric, so firm 2's best response to q_1 is $\mathrm{BR}_2(q_1) = (a - c - b\,q_1)/(2b)$.

In a Nash equilibrium both equations hold at once, which gives two linear equations in two unknowns. Because the game is symmetric, the equilibrium is symmetric too, and the pair can be solved by setting $q_1 = q_2 = q^*$ in firm 1's equation:

$$q^* = \frac{a - c - b\,q^*}{2b} \quad \Longrightarrow \quad 2b\,q^* = a - c - b\,q^* \quad \Longrightarrow \quad q^* = \frac{a - c}{3b}.$$

Substituting back into the inverse demand gives the price $P^* = a - 2b\,q^* = (a + 2c)/3$, and each firm's profit is $(P^* - c)\,q^* = (a - c)^2/(9b)$. That the symmetric solution is the only one follows from subtracting the two best-response equations, which gives $q_1 - q_2 = (q_1 - q_2)/2$ and hence $q_1 = q_2$. This completes the derivation.

The graph of a best-response function is called a **reaction curve**, after Cournot's own drawing, and Figure 4 shows both for the program's numbers. Firm 1's quantity is on the horizontal axis. The curves are straight lines here because the demand is linear, and they cross exactly once, at the Nash equilibrium.

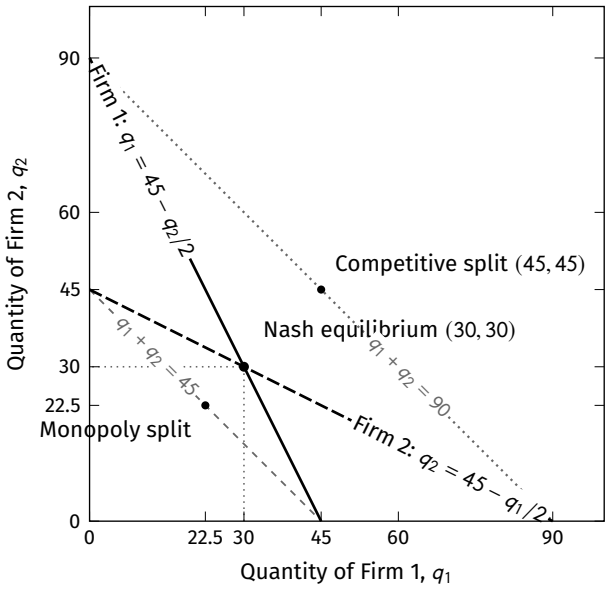


Figure 4. The two best-response lines of the Cournot game with $P = 120 - (q_1 + q_2)$ and a cost of 30 per item. They cross at the Nash equilibrium $(30, 30)$. The dashed grey line $q_1 + q_2 = 45$ is the monopoly quantity, with the equal split at $(22.5, 22.5)$, and the dotted grey line $q_1 + q_2 = 90$ is the competitive quantity. Computed from the Key Result above.

WORKED EXAMPLE **Cournot's duopoly with the program's numbers**

With $a = 120$, $b = 1$, and $c = 30$ the best-response functions are $q_1 = 45 - q_2/2$ and $q_2 = 45 - q_1/2$. The equilibrium is

$$q_1^* = q_2^* = \frac{120 - 30}{3} = 30, \qquad P^* = 120 - 60 = 60, \qquad \pi_1 = \pi_2 = (60 - 30) \times 30 = 900.$$

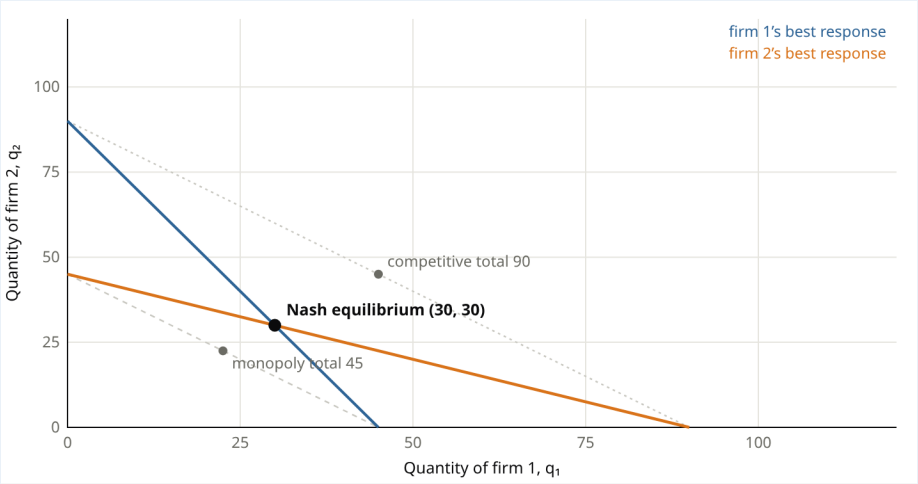
To check that nobody wants to deviate, let firm 1 try 35 while firm 2 keeps 30. The price falls to $120 - 65 = 55$, and firm 1's profit becomes $(55 - 30) \times 35 = 875$, less than 900. Trying 25 instead gives a price of $120 - 55 = 65$, a margin of 35 per item, and a profit of $35 \times 25 = 875$, again less. The quantity 30 is the best response, as the formula said.

Two other outcomes of the same market put the equilibrium in perspective. A single owner of both firms, a monopolist, would choose the total Q that maximizes $(120 - Q - 30)\,Q$; completing the square as in the Proof box, or reading off that the parabola $(90 - Q)\,Q$ peaks at half of 90, gives $Q = 45$, a price of 75, and a total profit of $45 \times 45 = 2025$, or 1012.5 per firm if the market were shared equally. If instead many firms competed and none could affect the price, the price would fall to the cost of 30, the quantity would be $120 - 30 = 90$, and profits would be zero. The duopoly sits between the two: 30 items each and a total of 60, at a price of 60, with a profit of 900 for each firm, against the monopoly's 45 at a price of 75 and the competitive 90 at a price of 30.

Both firms would be better off at the monopoly split of 22.5 each, with 1012.5 instead of 900 apiece, but that pair is not an equilibrium: against a rival who sells 22.5 the best response is $45 - 11.25 = 33.75$, each firm is tempted to expand, and the expansion of both drives the price down to 60. The duopoly is a prisoner's dilemma with infinitely many strategies instead of two.

The economic reason for the extra output is that each firm counts only its own loss when it lowers the price by selling one more item and ignores the loss it inflicts on the rival, whereas a monopolist bears both losses and stops earlier. Unit 5 shows what changes when one firm chooses its quantity first and the other observes it, the case of Stackelberg leadership, and Unit 6 asks whether two firms that meet year after year can hold the monopoly split after all.

INTERACTIVE Cournot competition between two firms



Move the demand intercept and the two firms' costs and watch the two best-response lines, their crossing at the Nash equilibrium, and the readouts of quantities, price, and profits next to the monopoly and the competitive outcome. Give one firm a lower cost and its line moves outwards, so that it sells more and earns more than its rival.

Click the picture or this link to try it in the online version of this unit.

CHECK YOUR UNDERSTANDING Question 3

In the Cournot game with inverse demand $P = 120 - (q_1 + q_2)$, both firms' cost per item falls from 30 to 0. How much does each firm produce in the new Nash equilibrium?

(A) 30
(B) 45
(C) 60
(D) 40

The solution is at the end of the unit.

5.2 Bertrand competition

Bertrand's objection of 1883 was that firms do not announce quantities and let the market find the price; they announce prices and let the buyers decide where to shop. The change seems small and turns the result upside down. In **Bertrand competition** two firms with the same cost c per item and the same product each name a price, at the same time. Buyers, who are fully informed and have no preference between the firms, all buy from the cheaper one; if the prices are equal, the buyers split evenly. With the demand of Equation (1), the quantity sold at a price P is $Q = (a - P)/b$, which for the program's numbers is $120 - P$.

The equilibrium can be found by going through the possible pairs of prices one case at a time, without any algebra.

1. At a common price above cost each firm serves half the market, and either firm can take the whole market at almost the same margin per item by pricing slightly lower, which roughly doubles its profit, so no such pair is stable.

2. If the prices differ and the cheaper price is above cost, the dearer firm sells nothing and gains by undercutting the cheaper firm, so no such pair is stable either.
3. A firm that prices below cost sells every item at a loss and gains by raising its price, so no pair with a price below cost is stable.
4. Only $p_1 = p_2 = c$ survives. Both firms earn zero there, a firm that raises its price sells nothing and still earns zero, and a firm that lowers it sells every item at a loss, so no unilateral change pays. The condition of Section 2 holds only weakly, because a firm that raises its price ends up tied with one that keeps it rather than worse off.

One detail deserves a precise statement. When a price can be any number, there is no largest price below the rival's, so a firm facing a rival price above cost has no single best response, only a whole range of better replies, each a little below the rival. For the argument this is enough, because the definition only asks whether some change pays, not which change is best; if prices move in steps of one cent, the best response is exactly one cent below the rival.

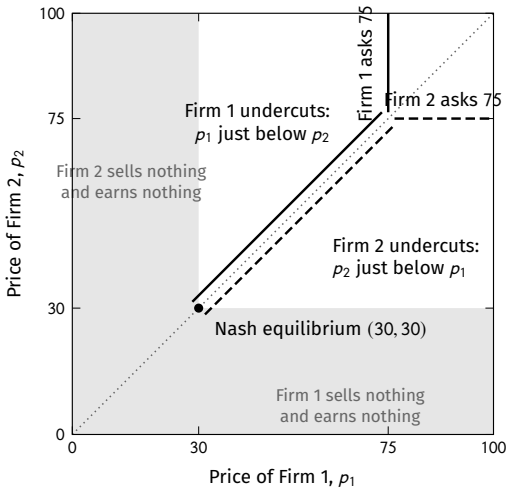


Figure 5. The best responses of the Bertrand game with $Q = 120 - P$ and a cost of 30, drawn as correspondences, the name used when a best response may be a whole set of prices rather than a single one. Against a rival price between 30 and 75 each firm's best response is a price just below it, drawn slightly off the diagonal; against a rival price above the monopoly price of 75 it is 75 itself; against a rival price at or below cost, in the shaded regions, any higher price gives zero profit and is as good as any other. The only point on both correspondences is the equilibrium $(30, 30)$.

KEY RESULT The Bertrand equilibrium

When two firms with the same constant cost c per item and an identical product set prices at the same time, and buyers all go to the cheaper firm, the unique Nash equilibrium is

$$p_1^* = p_2^* = c,$$

with zero profit for both. Two firms are enough to produce the competitive price. The result needs the stated assumptions; firms with limited capacity, with products that differ, or with different costs, and firms that meet repeatedly, do not end up at cost.

WORKED EXAMPLE Price competition with the program's numbers

With the demand $Q = 120 - P$ and a cost of 30 per item, the Bertrand equilibrium has both firms charging 30, selling 45 items each and 90 in total, and earning nothing, which is the competitive outcome of the Cournot example reached with two firms instead of many. The argument is

easiest to feel with concrete prices. Suppose both firms charge 60, the Cournot price, and each sells 30 items for a profit of 900. If firm 1 cuts its price to 59, it sells all $120 - 59 = 61$ items at a margin of 29, a profit of 1769, almost double. Firm 2 answers with 58 and takes $62 \times 28 = 1736$. The undercutting continues until a price cut no longer pays, which happens only at 30. In equilibrium the price equals the cost of 30, each firm sells 45 items, and both earn zero. If prices can only move in whole cents, a common price one cent above cost is also an equilibrium, because undercutting it would mean selling at cost; the conclusion does not change.

With the same demand and the same costs, quantity competition gives a price of 60 and a profit of 900 each, and price competition gives a price of 30 and nothing; the two games differ in what a firm can commit to, a quantity or a price. Where capacity is fixed in advance, as for a steel mill or a hotel, Cournot's model fits better, and where firms can supply any demand at the price they name, as in much online retail, Bertrand's does.

6 Nash's theorem and the cost of finding an equilibrium

The examples of this unit all had an equilibrium, except matching pennies in pure strategies. The theorem that made the concept central guarantees that this is no accident.

EXCURSION ***A thesis of 27 pages***

Nash was born on 13 June 1928 in Bluefield, West Virginia, and came to Princeton in 1948, four years after John von Neumann (1903-1957) and Oskar Morgenstern (1902-1977) had published *Theory of Games and Economic Behavior* there, a book that completely solved two-player zero-sum games, the games in which one player's gain is the other's loss. Nash asked what rational players would do in any game when each chose for themselves, without binding agreements. A note of two pages, "Equilibrium Points in n-Person Games", appeared in the *Proceedings of the National Academy of Sciences* in January 1950 and stated the existence theorem with a proof based on a fixed point theorem of Shizuo Kakutani (1911-2004). His doctoral thesis, "Non-Cooperative Games", about 27 typed pages long, was accepted in May 1950 under Albert Tucker (1905-1995), who told the prison story of Unit 2 that same month, and a fuller version appeared in the *Annals of Mathematics* in 1951 with a proof from the older fixed point theorem of L. E. J. Brouwer (1881-1966). From 1959 he was ill for many years and published almost nothing. In 1994 the Nobel prize in economic sciences was awarded to Nash, Harsanyi, and Selten for their analysis of equilibria in non-cooperative games. Nash died on 23 May 2015.

A **finite game** has finitely many players, each with finitely many pure strategies; every matrix game is finite, and so is chess, written out as a matrix of complete plans of play. A **mixed strategy** is a probability distribution over a player's pure strategies, such as heads and tails with probability $\frac{1}{2}$ each in matching pennies; Unit 4 develops them properly.

THEOREM **Nash's existence theorem**

Every finite game has at least one Nash equilibrium, provided that mixed strategies are allowed. In pure strategies alone an equilibrium may fail to exist, as matching pennies shows, but once each player may randomize over their pure strategies, a profile exists in which every player's mixed strategy is a best response to the others'. Nash stated and proved this in 1950 for any number of players.

The proof is beyond this unit, because it rests on a theorem from topology, but its idea fits in a paragraph. Think of a rule, called a map in mathematics, that takes every profile of mixed strategies and moves each player a little towards their best responses against the others, leaving a player who already plays a best response alone. A profile that the map does not move is called a **fixed point** of the map, and by construction a fixed point is exactly a Nash equilibrium. The fixed point theorems of Brouwer and Kakutani say that any such map, if it is continuous and acts on a closed and bounded set without holes, such as a disc or a solid cube, leaves at least one point unmoved, just as a paper map of a country, laid down anywhere inside that country, has one point lying exactly above the place it represents. Nash built the right map and checked that it satisfies the conditions.

6.1 Can an equilibrium be found?

The theorem says that an equilibrium exists; it does not say how to find one, because fixed point theorems show that a point exists without locating it. For a 2 by 2 game the search takes a minute. For two players with a hundred strategies each, or for many players, the question becomes serious, and not only for computers: if people are supposed to play equilibrium, some process of learning, reasoning, or trial and error must take them there, and such a process is an algorithm whether a person or a machine runs it. A method is said to run in **polynomial time** when the number of steps it needs grows no faster than some fixed power of the size of the input, here the number of payoffs in the matrix; a method whose running time grows exponentially with the size of the game, doubling, say, with every added strategy, is useless beyond toy examples. For two-player zero-sum games a polynomial-time method has existed since 1979, when linear programming, the task of making a linear quantity as large as possible under linear constraints, into which such games can be translated, was first solved in polynomial time. For general games the known methods, including the Lemke-Howson algorithm of 1964, can take exponentially many steps in the worst case, and nobody has found a better one.

THEOREM An open problem

Is there a method that finds a Nash equilibrium of every finite game, given its payoff matrix, in polynomial time? No such method is known, and no proof that none exists is known either. What is known, from the research spotlight below, is that the problem is complete for a family of problems called PPAD: a fast method for it would also solve every other problem of that family, some of which have resisted every attempt for decades.

RESEARCH SPOTLIGHT

Constantinos Daskalakis *Massachusetts Institute of Technology*

Nash proved that an equilibrium always exists, but can a computer find one?

Constantinos Daskalakis, Avanessians Professor of Electrical Engineering and Computer Science at the Massachusetts Institute of Technology, asked with Paul Goldberg (University of Oxford) and Christos Papadimitriou (Columbia University) how hard it is to compute what Nash proved to exist. The usual yardstick, NP-completeness, does not fit, because NP-complete problems draw their difficulty from instances that have no solution, and by Nash's theorem every game has one. The authors therefore placed the problem in a complexity class called PPAD, built on an existence argument of its own: in a directed graph in which every node has at most one arrow in and one out, a path that starts somewhere must end somewhere, and the task is to find the end when the graph is given only as a program and is far too large to write out. They showed that finding a Nash equilibrium, even approximately, is exactly as hard as this end-of-the-line problem, first for three or more players, which Chen and Deng (2006) extended to two. If no fast method exists for end-of-the-line problems, as most researchers

believe, none exists for Nash equilibria either, and a theory that predicts equilibrium play in every game owes an account of how the players get there. The introduction and the section on total search problems are readable with the knowledge of this unit.

Read the paper: Daskalakis, Goldberg, and Papadimitriou (2009), The Complexity of Computing a Nash Equilibrium, SIAM Journal on Computing 39, 195 to 259 (the link opens the authors' short version in Communications of the ACM 52, 89 to 97)

The open problem does not spoil the concept: for the games of this unit, and for the auctions, matchings, and repeated games of later units, equilibria can be found and checked by hand. What the result rules out is the comfortable assumption that equilibrium is the natural resting point of play in every game, however large and strange.

KEY POINTS

- A Nash equilibrium is a strategy profile in which every player's strategy is a best response to the others', so that no player gains by changing alone; read it as no regrets, as a self-enforcing agreement, or as consistent expectations.
- In a payoff matrix, underline each player's best responses; the cells with two underlines are all the pure strategy equilibria. The stag hunt, the coordination game with conflicting preferences, and chicken have two each, matching pennies none.
- Among several equilibria, focal points, conventions, and the comparison of payoff dominance with risk dominance explain which one is played. In the stag hunt the stag is payoff dominant and the hare risk dominant, and the stag is worth the risk only against a partner at least 75 per cent likely to cooperate.
- In Cournot competition with $P = a - b(q_1 + q_2)$ and cost c, the best response is $(a - c - b\,q_j)/(2b)$ and each firm sells $(a - c)/(3b)$; with the program's numbers 30 each at a price of 60, between the monopoly's 45 at 75 and the competitive 90 at 30. In Bertrand competition two identical firms that set prices end up at cost, because any common price above cost invites undercutting.
- Nash's existence theorem guarantees an equilibrium in every finite game once mixed strategies are allowed; whether one can be found in polynomial time is open, and the problem is complete for PPAD.

Solutions to the questions in the text: 1-B, 2-C, 3-D. The online version of this unit explains each answer.