Publicly-available solutions for

AN INTRODUCTION TO GAME THEORY

Publicly-available solutions for

AN INTRODUCTION TO GAME THEORY

MARTIN J. OSBORNE
University of Toronto

Copyright © 2012 by Martin J. Osborne

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording, or otherwise, without the prior permission of Martin J. Osborne.

This manual was typeset by the author, who is greatly indebted to Donald Knuth (TEX), Leslie Lamport (LATEX), Diego Puga (mathpazo), Christian Schenk (MiKTEX), Ed Sznyter (ppctr), Timothy van Zandt (PSTricks), and others, for generously making superlative software freely available. The main font is 10pt Palatino.

Version 6: 2012-4-7

Contents

Preface

This manual contains all publicly-available solutions to exercises in my book *An Introduction to Game Theory* (Oxford University Press, 2004). The sources of the problems are given in the section entitled "Notes" at the end of each chapter of the book. Please alert me to errors.

MARTIN J. OSBORNE
Martin.Osborne@utoronto.ca
Department of Economics, 150 St. George Street,
University of Toronto, Toronto, Canada M5S 3G7

1 Introduction

5.3 Altruistic preferences

Person 1 is indifferent between $(1, 4)$ and $(3, 0)$, and prefers both of these to $(2, 1)$. The payoff function u defined by $u(x, y) = x + \frac{1}{2}y$, where x is person 1's income and y is person 2's, represents person 1's preferences. Any function that is an increasing function of u also represents her preferences. For example, the functions $k(x + \frac{1}{2}y)$ for any positive number k, and $(x + \frac{1}{2}y)^2$, do so.

6.1 Alternative representations of preferences

The function v represents the same preferences as does u (because $u(a) < u(b) < u(c)$ and $v(a) < v(b) < v(c)$), but the function w does not represent the same preferences, because $w(a) = w(b)$ while $u(a) < u(b)$.

2 Nash Equilibrium

16.1 Working on a joint project

The game in Figure 3.1 models this situation (as does any other game with the same players and actions in which the ordering of the payoffs is the same as the ordering in Figure 3.1).

	Work hard	*Goof off*
Work hard	$3,3$	$0,2$
Goof off	$2,0$	$1,1$

Figure 3.1 Working on a joint project (alternative version).

17.1 Games equivalent to the *Prisoner's Dilemma*

The game in the left panel differs from the *Prisoner's Dilemma* in both players' preferences. Player 1 prefers (Y, X) to (X, X) to (X, Y) to (Y, Y), for example, which differs from her preference in the *Prisoner's Dilemma*, which is (F, Q) to (Q, Q) to (F, F) to (Q, F), whether we let $X = F$ or $X = Q$.

The game in the right panel is equivalent to the *Prisoner's Dilemma*. If we let $X = Q$ and $Y = F$ then player 1 prefers (F, Q) to (Q, Q) to (F, F) to (Q, F) and player 2 prefers (Q, F) to (Q, Q) to (F, F) to (F, Q), as in the *Prisoner's Dilemma*.

20.1 Games without conflict

Any two-player game in which each player has two actions and the players have the same preferences may be represented by a table of the form given in Figure 3.2, where a, b, c, and d are any numbers.

	L	R
T	a,a	b,b
B	c,c	d,d

Figure 3.2 A strategic game in which conflict is absent.

31.1 Extension of the *Stag Hunt*

Every profile $(e, \ldots, e)$, where e is an integer from 0 to K, is a Nash equilibrium. In the equilibrium $(e, \ldots, e)$, each player's payoff is e. The profile $(e, \ldots, e)$ is a Nash equilibrium since if player i chooses $e_i < e$ then her payoff is $2e_i - e_i = e_i < e$, and if she chooses $e_i > e$ then her payoff is $2e - e_i < e$.

Consider an action profile $(e_1, \ldots, e_n)$ in which not all effort levels are the same. Suppose that e_i is the minimum. Consider some player j whose effort level exceeds e_i. Her payoff is $2e_i - e_j < e_i$, while if she deviates to the effort level e_i her payoff is $2e_i - e_i = e_i$. Thus she can increase her payoff by deviating, so that $(e_1, \ldots, e_n)$ is not a Nash equilibrium.

(This game is studied experimentally by van Huyck, Battalio, and Beil (1990). See also Ochs (1995, 209–233).)

34.1 Guessing two-thirds of the average

If all three players announce the same integer $k \geq 2$ then any one of them can deviate to $k - 1$ and obtain \$1 (since her number is now closer to $\frac{2}{3}$ of the average than the other two) rather than $\$\frac{1}{3}$. Thus no such action profile is a Nash equilibrium. If all three players announce 1, then no player can deviate and increase her payoff; thus $(1, 1, 1)$ is a Nash equilibrium.

Now consider an action profile in which not all three integers are the same; denote the highest by k^*.

- Suppose only one player names k^*; denote the other integers named by k_1 and k_2, with $k_1 \geq k_2$. The average of the three integers is $\frac{1}{3}(k^* + k_1 + k_2)$, so that $\frac{2}{3}$ of the average is $\frac{2}{9}(k^* + k_1 + k_2)$. If $k_1 \geq \frac{2}{9}(k^* + k_1 + k_2)$ then k^* is further from $\frac{2}{3}$ of the average than is k_1, and hence does not win. If $k_1 < \frac{2}{9}(k^* + k_1 + k_2)$ then the difference between k^* and $\frac{2}{3}$ of the average is $k^* - \frac{2}{9}(k^* + k_1 + k_2) = \frac{7}{9}k^* - \frac{2}{9}k_1 - \frac{2}{9}k_2$, while the difference between k_1 and $\frac{2}{3}$ of the average is $\frac{2}{9}(k^* + k_1 + k_2) - k_1 = \frac{2}{9}k^* - \frac{7}{9}k_1 + \frac{2}{9}k_2$. The difference between the former and the latter is $\frac{5}{9}k^* + \frac{5}{9}k_1 - \frac{4}{9}k_2 > 0$, so k_1 is closer to $\frac{2}{3}$ of the average than is k^*. Hence the player who names k^* does not win, and is better off naming k_2, in which case she obtains a share of the prize. Thus no such action profile is a Nash equilibrium.

- Suppose two players name k^*, and the third player names $k < k^*$. The average of the three integers is then $\frac{1}{3}(2k^* + k)$, so that $\frac{2}{3}$ of the average is $\frac{4}{9}k^* + \frac{2}{9}k$. We have $\frac{4}{9}k^* + \frac{2}{9}k < \frac{1}{2}(k^* + k)$ (since $\frac{4}{9} < \frac{1}{2}$ and $\frac{2}{9} < \frac{1}{2}$), so that the player who names k is the sole winner. Thus either of the other players can switch to naming k and obtain a share of the prize rather obtaining nothing. Thus no such action profile is a Nash equilibrium.

We conclude that there is only one Nash equilibrium of this game, in which all three players announce the number 1.

(This game is studied experimentally by Nagel (1995).)

34.3 Choosing a route

A strategic game that models this situation is:

Players The four people.

Actions The set of actions of each person is $\{X, Y\}$ (the route via X and the route via Y).

Preferences Each player's payoff is the negative of her travel time.

In every Nash equilibrium, two people take each route. (In any other case, a person taking the more popular route is better off switching to the other route.) For any such action profile, each person's travel time is either 29.9 or 30 minutes (depending on the route she takes). If a person taking the route via X switches to the route via Y her travel time becomes $22 + 12 = 34$ minutes; if a person taking the route via Y switches to the route via X her travel time becomes $12 + 21.8 = 33.8$ minutes. For any other allocation of people to routes, at least one person can decrease her travel time by switching routes. Thus the set of Nash equilibria is the set of action profiles in which two people take the route via X and two people take the route via Y.

Now consider the situation after the road from X to Y is built. There is no equilibrium in which the new road is not used, by the following argument. Because the only equilibrium before the new road is built has two people taking each route, the only possibility for an equilibrium in which no one uses the new road is for two people to take the route A–X–B and two to take A–Y–B, resulting in a total travel time for each person of either 29.9 or 30 minutes. However, if a person taking A–X–B switches to the new road at X and then takes Y–B her total travel time becomes $9 + 7 + 12 = 28$ minutes.

I claim that in any Nash equilibrium, one person takes A–X–B, two people take A–X–Y–B, and one person takes A–Y–B. For this assignment, each person's travel time is 32 minutes. No person can change her route and decrease her travel time, by the following argument.

- If the person taking A–X–B switches to A–X–Y–B, her travel time increases to $12 + 9 + 15 = 36$ minutes; if she switches to A–Y–B her travel time increases to $21 + 15 = 36$ minutes.
- If one of the people taking A–X–Y–B switches to A–X–B, her travel time increases to $12 + 20.9 = 32.9$ minutes; if she switches to A–Y–B her travel time increases to $21 + 12 = 33$ minutes.
- If the person taking A–Y–B switches to A–X–B, her travel time increases to $15 + 20.9 = 35.9$ minutes; if she switches to A–X–Y–B, her travel time increases to $15 + 9 + 12 = 36$ minutes.

For every other allocation of people to routes at least one person can switch routes and reduce her travel time. For example, if one person takes A–X–B, one person takes A–X–Y–B, and two people take A–Y–B, then the travel time of those taking A–Y–B is $21 + 12 = 33$ minutes; if one of them switches to A–X–B then her travel time falls to $12 + 20.9 = 32.9$ minutes. Or if one person takes A–Y–B, one person takes A–X–Y–B, and two people take A–X–B, then the travel time of those taking A–X–B is $12 + 20.9 = 32.9$ minutes; if one of them switches to A–X–Y–B then her travel time falls to $12 + 8 + 12 = 32$ minutes.

Thus in the equilibrium with the new road every person's travel time *increases*, from either 29.9 or 30 minutes to 32 minutes.

37.1 Finding Nash equilibria using best response functions

a. The *Prisoner's Dilemma* and *BoS* are shown in Figure 6.1; *Matching Pennies* and the two-player *Stag Hunt* are shown in Figure 6.2.

	Quiet	*Fink*
Quiet	$2\ ,2$	$0\ ,3^*$
Fink	$3^*,0$	$1^*,1^*$

Prisoner's Dilemma

	Bach	*Stravinsky*
Bach	$2^*,1^*$	$0\ ,0$
Stravinsky	$0\ ,0$	$1^*,2^*$

BoS

Figure 6.1 The best response functions in the *Prisoner's Dilemma* (left) and in *BoS* (right).

	Head	*Tail*
Head	$1^*,-1$	$-1\ ,\ \ 1^*$
Tail	$-1\ ,\ \ 1^*$	$1^*,-1$

Matching Pennies

	Stag	*Hare*
Stag	$2^*,2^*$	$0\ ,1$
Hare	$1\ ,0$	$1^*,1^*$

Stag Hunt

Figure 6.2 The best response functions in *Matching Pennies* (left) and the *Stag Hunt* (right).

b. The best response functions are indicated in Figure 6.3. The Nash equilibria are (T, C), (M, L), and (B, R).

	L	C	R
T	$2\ ,2$	$1^*,3^*$	$0^*,1$
M	$3^*,1^*$	$0\ ,0$	$0^*,0$
B	$1\ ,0^*$	$0\ ,0^*$	$0^*,0^*$

Figure 6.3 The game in Exercise 37.1.

38.1 Constructing best response functions

The analogue of Figure 38.2 in the book is given in Figure 7.1.

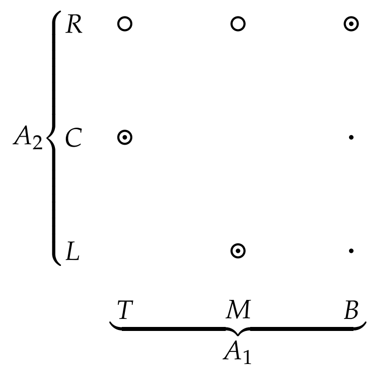

Figure 7.1 The players' best response functions for the game in Exercise 38.1b. Player 1's best responses are indicated by circles, and player 2's by dots. The action pairs for which there is both a circle and a dot are the Nash equilibria.

38.2 Dividing money

For each amount named by one of the players, the other player's best responses are given in the following table.

Other player's action	Sets of best responses
0	$\{10\}$
1	$\{9, 10\}$
2	$\{8, 9, 10\}$
3	$\{7, 8, 9, 10\}$
4	$\{6, 7, 8, 9, 10\}$
5	$\{5, 6, 7, 8, 9, 10\}$
6	$\{5, 6\}$
7	$\{6\}$
8	$\{7\}$
9	$\{8\}$
10	$\{9\}$

The best response functions are illustrated in Figure 8.1 (circles for player 1, dots for player 2). From this figure we see that the game has four Nash equilibria: $(5, 5)$, $(5, 6)$, $(6, 5)$, and $(6, 6)$.

41.1 Strict and nonstrict Nash equilibria

Only the Nash equilibrium (a_1^*, a_2^*) is strict. For each of the other equilibria, player 2's action a_2 satisfies $a_2^{***} \leq a_2 \leq a_2^{**}$, and for each such action player 1 has multiple best responses, so that her payoff is the same for a range of actions, only one of which is such that (a_1, a_2) is a Nash equilibrium.

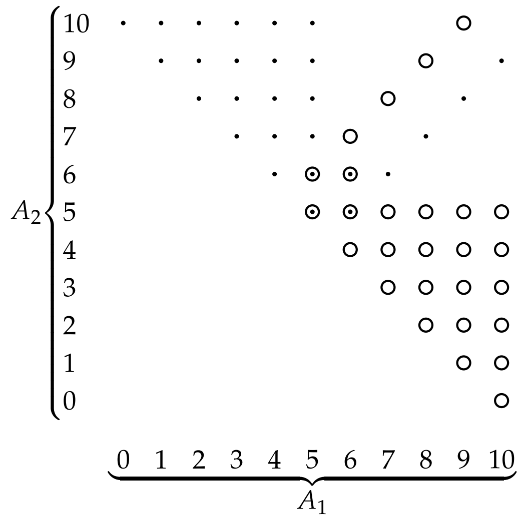

Figure 8.1 The players' best response functions for the game in Exercise 38.2.

47.1 Strict equilibria and dominated actions

For player 1, T is weakly dominated by M, and strictly dominated by B. For player 2, no action is weakly or strictly dominated. The game has a unique Nash equilibrium, (M, L). This equilibrium is not strict. (When player 2 choose L, B yields player 1 the same payoff as does M.)

47.2 Nash equilibrium and weakly dominated actions

The only Nash equilibrium of the game in Figure 8.2 is (T, L). The action T is weakly dominated by M and the action L is weakly dominated by C. (There are of course many other games that satisfy the conditions.)

	L	C	R
T	$1,1$	$0,1$	$0,0$
M	$1,0$	$2,1$	$1,2$
B	$0,0$	$1,1$	$2,0$

Figure 8.2 A game with a unique Nash equilibrium, in which both players' equilibrium actions are weakly dominated. (The unique Nash equilibrium is (T, L).)

50.1 Other Nash equilibria of the game modeling collective decision-making

Denote by i the player whose favorite policy is the median favorite policy. The set of Nash equilibria includes every action profile in which (*i*) i's action is her favorite policy x_i^*, (*ii*) every player whose favorite policy is less than x_i^* names a policy equal to at most x_i^*, and (*iii*) every player whose favorite policy is greater than x_i^* names a policy equal to at least x_i^*.

To show this, first note that the outcome is x_i^*, so player i cannot induce a better outcome for herself by changing her action. Now, if a player whose favorite position is less than x_i^* changes her action to some $x < x_i^*$, the outcome does not change; if such a player changes her action to some $x > x_i^*$ then the outcome either remains the same (if some player whose favorite position exceeds x_i^* names x_i^*) or increases, so that the player is not better off. A similar argument applies to a player whose favorite position is greater than x_i^*.

The set of Nash equilibria also includes, for any positive integer $k \leq n$, every action profile in which k players name the median favorite policy x_i^*, at most $\frac{1}{2}(n-3)$ players name policies less than x_i^*, and at most $\frac{1}{2}(n-3)$ players name policies greater than x_i^*. (In these equilibria, the favorite policy of a player who names a policy less than x_i^* may be greater than x_i^*, and vice versa. The conditions on the numbers of players who name policies less than x_i^* and greater than x_i^* ensure that no such player can, by naming instead her favorite policy, move the median policy closer to her favorite policy.)

Any action profile in which all players name the same, arbitrary, policy is also a Nash equilibrium; the outcome is the common policy named.

More generally, any profile in which at least three players name the same, arbitrary, policy x, at most $(n-3)/2$ players name a policy less than x, and at most $(n-3)/2$ players name a policy greater than x is a Nash equilibrium. (In both cases, no change in any player's action has any effect on the outcome.)

51.2 Symmetric strategic games

The games in Exercise 31.2, Example 39.1, and Figure 47.2 (both games) are symmetric. The game in Exercise 42.1 is not symmetric. The game in Section 2.8.4 is symmetric if and only if $u_1 = u_2$.

52.2 Equilibrium for pairwise interactions in a single population

The Nash equilibria are (A, A), (A, C), and (C, A). Only the equilibrium (A, A) is relevant if the game is played between the members of a single population—this equilibrium is the only *symmetric* equilibrium.

3 Nash Equilibrium: Illustrations

58.1 Cournot's duopoly game with linear inverse demand and different unit costs

Following the analysis in the text, the best response function of firm 1 is

$$b_1(q_2) = \begin{cases} \frac{1}{2}(\alpha - c_1 - q_2) & \text{if } q_2 \leq \alpha - c_1 \\ 0 & \text{otherwise} \end{cases}$$

while that of firm 2 is

$$b_2(q_1) = \begin{cases} \frac{1}{2}(\alpha - c_2 - q_1) & \text{if } q_1 \leq \alpha - c_2 \\ 0 & \text{otherwise.} \end{cases}$$

To find the Nash equilibrium, first plot these two functions. Each function has the same general form as the best response function of either firm in the case studied in the text. However, the fact that $c_1 \neq c_2$ leads to two qualitatively different cases when we combine the two functions to find a Nash equilibrium. If c_1 and c_2 do not differ very much then the functions in the analogue of Figure 59.1 intersect at a pair of outputs that are both positive. If c_1 and c_2 differ a lot, however, the functions intersect at a pair of outputs in which $q_1 = 0$.

Precisely, if $c_1 \leq \frac{1}{2}(\alpha + c_2)$ then the downward-sloping parts of the best response functions intersect (as in Figure 59.1), and the game has a unique Nash equilibrium, given by the solution of the two equations

$$\begin{aligned} q_1 &= \tfrac{1}{2}(\alpha - c_1 - q_2) \\ q_2 &= \tfrac{1}{2}(\alpha - c_2 - q_1). \end{aligned}$$

This solution is

$$(q_1^*, q_2^*) = \left(\tfrac{1}{3}(\alpha - 2c_1 + c_2), \tfrac{1}{3}(\alpha - 2c_2 + c_1)\right).$$

If $c_1 > \frac{1}{2}(\alpha + c_2)$ then the downward-sloping part of firm 1's best response function lies below the downward-sloping part of firm 2's best response function (as in Figure 12.1), and the game has a unique Nash equilibrium, $(q_1^*, q_2^*) = (0, \frac{1}{2}(\alpha - c_2))$.

In summary, the game always has a unique Nash equilibrium, defined as follows:

$$\begin{cases} \left(\frac{1}{3}(\alpha - 2c_1 + c_2), \frac{1}{3}(\alpha - 2c_2 + c_1)\right) & \text{if } c_1 \leq \frac{1}{2}(\alpha + c_2) \\ \left(0, \frac{1}{2}(\alpha - c_2)\right) & \text{if } c_1 > \frac{1}{2}(\alpha + c_2). \end{cases}$$

The output of firm 2 exceeds that of firm 1 in every equilibrium.

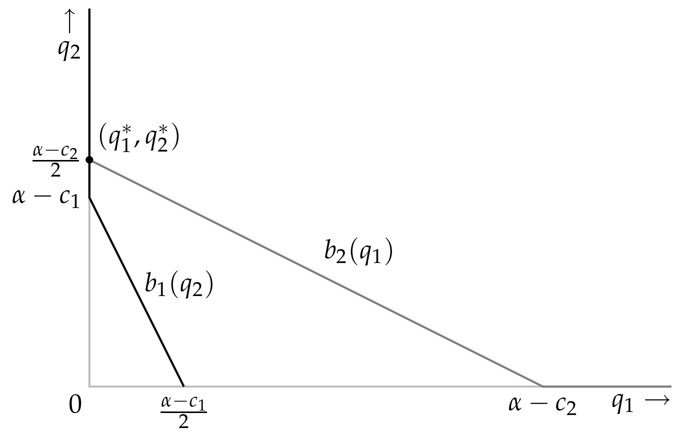

Figure 12.1 The best response functions in Cournot's duopoly game under the assumptions of Exercise 58.1 when $\alpha - c_1 < \frac{1}{2}(\alpha - c_2)$. The unique Nash equilibrium in this case is $(q_1^*, q_2^*) = (0, \frac{1}{2}(\alpha - c_2))$.

If c_2 decreases then firm 2's output increases and firm 1's output either falls, if $c_1 \leq \frac{1}{2}(\alpha + c_2)$, or remains equal to 0, if $c_1 > \frac{1}{2}(\alpha + c_2)$. The total output increases and the price falls.

60.2 Nash equilibrium of Cournot's duopoly game and the collusive outcome

The firms' total profit is $(q_1 + q_2)(\alpha - c - q_1 - q_2)$, or $Q(\alpha - c - Q)$, where Q denotes total output. This function is a quadratic in Q that is zero when $Q = 0$ and when $Q = \alpha - c$, so that its maximizer is $Q^* = \frac{1}{2}(\alpha - c)$.

If each firm produces $\frac{1}{4}(\alpha - c)$ then its profit is $\frac{1}{8}(\alpha - c)^2$. This profit exceeds its Nash equilibrium profit of $\frac{1}{9}(\alpha - c)^2$.

If one firm produces $Q^*/2$, the other firm's best response is $b_i(Q^*/2) = \frac{1}{2}(\alpha - c - \frac{1}{4}(\alpha - c)) = \frac{3}{8}(\alpha - c)$. That is, if one firm produces $Q^*/2$, the other firm wants to produce *more* than $Q^*/2$.

63.1 Interaction among resource-users

The game is given as follows.

Players The firms.

Actions Each firm's set of actions is the set of all nonnegative numbers (representing the amount of input it uses).

Preferences The payoff of each firm i is

$$\begin{cases} x_i(1 - (x_1 + \cdots + x_n)) & \text{if } x_1 + \cdots + x_n \leq 1 \\ 0 & \text{if } x_1 + \cdots + x_n > 1. \end{cases}$$

This game is the same as that in Exercise 61.1 for $c = 0$ and $\alpha = 1$. Thus it has a unique Nash equilibrium, $(x_1, \ldots, x_n) = (1/(n+1), \ldots, 1/(n+1))$.

In this Nash equilibrium, each firm's output is $(1/(n+1))(1 - n/(n+1)) = 1/(n+1)^2$. If $x_i = 1/(2n)$ for $i = 1, \ldots, n$ then each firm's output is $1/(4n)$, which exceeds $1/(n+1)^2$ for $n \geq 2$. (We have $1/(4n) - 1/(n+1)^2 = (n-1)^2/(4n(n+1)^2) > 0$ for $n \geq 2$.)

67.1 Bertrand's duopoly game with constant unit cost

The pair (c, c) of prices remains a Nash equilibrium; the argument is the same as before. Further, as before, there is no other Nash equilibrium. The argument needs only very minor modification. For an arbitrary function D there may exist no monopoly price p^m; in this case, if $p_i > c$, $p_j > c$, $p_i \geq p_j$, and $D(p_j) = 0$ then firm i can increase its profit by reducing its price slightly below $\overline{p}$ (for example).

68.1 Bertrand's oligopoly game

Consider a profile $(p_1, \ldots, p_n)$ of prices in which $p_i \geq c$ for all i and at least two prices are equal to c. Every firm's profit is zero. If any firm raises its price its profit remains zero. If a firm charging more than c lowers its price, but not below c, its profit also remains zero. If a firm lowers its price below c then its profit is negative. Thus any such profile is a Nash equilibrium.

To show that no other profile is a Nash equilibrium, we can argue as follows.

- If some price is less than c then the firm charging the lowest price can increase its profit (to zero) by increasing its price to c.
- If exactly one firm's price is equal to c then that firm can increase its profit by raising its price a little (keeping it less than the next highest price).
- If all firms' prices exceed c then the firm charging the highest price can increase its profit by lowering its price to some price between c and the lowest price being charged.

68.2 Bertrand's duopoly game with different unit costs

a. If all consumers buy from firm 1 when both firms charge the price c_2, then $(p_1, p_2) = (c_2, c_2)$ is a Nash equilibrium by the following argument. Firm 1's profit is positive, while firm 2's profit is zero (since it serves no customers).

- If firm 1 increases its price, its profit falls to zero.
- If firm 1 reduces its price, say to p, then its profit changes from $(c_2 - c_1)(\alpha - c_2)$ to $(p - c_1)(\alpha - p)$. Since c_2 is less than the maximizer of $(p - c_1)(\alpha - p)$, firm 1's profit falls.

- If firm 2 increases its price, its profit remains zero.
- If firm 2 decreases its price, its profit becomes negative (since its price is less than its unit cost).

Under this rule no other pair of prices is a Nash equilibrium, by the following argument.

- If $p_i < c_1$ for $i = 1, 2$ then the firm with the lower price (or either firm, if the prices are the same) can increase its profit (to zero) by raising its price above that of the other firm.
- If $p_1 > p_2 \geq c_2$ then firm 2 can increase its profit by raising its price a little.
- If $p_2 > p_1 \geq c_1$ then firm 1 can increase its profit by raising its price a little.
- If $p_2 \leq p_1$ and $p_2 < c_2$ then firm 2's profit is negative, so that it can increase its profit by raising its price.
- If $p_1 = p_2 > c_2$ then at least one of the firms is not receiving all of the demand, and that firm can increase its profit by lowering its price a little.

b. Now suppose that the rule for splitting up the customers when the prices are equal specifies that firm 2 receives some customers when both prices are c_2. By the argument for part *a*, the only possible Nash equilibrium is $(p_1, p_2) = (c_2, c_2)$. (The argument in part *a* that every other pair of prices is not a Nash equilibrium does not use the fact that customers are split equally when $(p_1, p_2) = (c_2, c_2)$.) But if $(p_1, p_2) = (c_2, c_2)$ and firm 2 receives some customers, firm 1 can increase its profit by reducing its price a little and capturing the entire market.

73.1 Electoral competition with asymmetric voters' preferences

The unique Nash equilibrium remains (m, m); the direct argument is exactly the same as before. (The dividing line between the supporters of two candidates with different positions changes. If $x_i < x_j$, for example, the dividing line is $\frac{1}{3}x_i + \frac{2}{3}x_j$ rather than $\frac{1}{2}(x_i + x_j)$. The resulting change in the best response functions does not affect the Nash equilibrium.)

75.3 Electoral competition for more general preferences

a. If x^* is a Condorcet winner then for any $y \neq x^*$ a majority of voters prefer x^* to y, so y is not a Condorcet winner. Thus there is no more than one Condorcet winner.

b. Suppose that one of the remaining voters prefers y to z to x, and the other prefers z to x to y. For each position there is another position preferred by a majority of voters, so no position is a Condorcet winner.

c. Now suppose that x^* is a Condorcet winner. Then the strategic game described the exercise has a unique Nash equilibrium in which both candidates choose x^*. This pair of actions is a Nash equilibrium because if either candidate chooses a different position she loses. For any other pair of actions either one candidate loses, in which case that candidate can deviate to the position x^* and at least tie, or the candidates tie at a position different from x^*, in which case either of them can deviate to x^* and win.

   If there is no Condorcet winner then for every position there is another position preferred by a majority of voters. Thus for every pair of distinct positions the loser can deviate and win, and for every pair of identical positions either candidate can deviate and win. Thus there is no Nash equilibrium.

76.1 Competition in product characteristics

Suppose there are two firms. If the products are different, then either firm increases its market share by making its product more similar to that of its rival. Thus in every possible equilibrium the products are the same. But if $x_1 = x_2 \neq m$ then each firm's market share is 50%, while if it changes its product to be closer to m then its market share rises above 50%. Thus the only possible equilibrium is $(x_1, x_2) = (m, m)$. This pair of positions is an equilibrium, since each firm's market share is 50%, and if either firm changes its product its market share falls below 50%.

Now suppose there are three firms. If all firms' products are the same, each obtains one-third of the market. If $x_1 = x_2 = x_3 = m$ then any firm, by changing its product a little, can obtain close to one-half of the market. If $x_1 = x_2 = x_3 \neq m$ then any firm, by changing its product a little, can obtain more than one-half of the market. If the firms' products are not all the same, then at least one of the extreme products is different from the other two products, and the firm that produces it can increase its market share by making it more similar to the other products. Thus when there are three firms there is no Nash equilibrium.

79.1 Direct argument for Nash equilibria of *War of Attrition*

- If $t_1 = t_2$ then either player can increase her payoff by conceding slightly later (in which case she obtains the object for sure, rather than getting it with probability $\frac{1}{2}$).
- If $0 < t_i < t_j$ then player i can increase her payoff by conceding at 0.
- If $0 = t_i < t_j < v_i$ then player i can increase her payoff (from 0 to almost $v_i - t_j > 0$) by conceding slightly after t_j.

Thus there is no Nash equilibrium in which $t_1 = t_2$, $0 < t_i < t_j$, or $0 = t_i < t_j < v_i$ (for $i = 1$ and $j = 2$, or $i = 2$ and $j = 1$). The remaining possibility is that $0 = t_i < t_j$ and $t_j \geq v_i$ for $i = 1$ and $j = 2$, or $i = 2$ and $j = 1$. In this case player i's payoff is 0, while if she concedes later her payoff is negative; player j's payoff is v_j, her highest possible payoff in the game.

85.1 Second-price sealed-bid auction with two bidders

If player 2's bid b_2 is less than v_1 then any bid of b_2 or more is a best response of player 1 (she wins and pays the price b_2). If player 2's bid is equal to v_1 then every bid of player 1 yields her the payoff zero (either she wins and pays v_1, or she loses), so every bid is a best response. If player 2's bid b_2 exceeds v_1 then any bid of less than b_2 is a best response of player 1. (If she bids b_2 or more she wins, but pays the price $b_2 > v_1$, and hence obtains a negative payoff.) In summary, player 1's best response function is

$$B_1(b_2) = \begin{cases} \{b_1 : b_1 \geq b_2\} & \text{if } b_2 < v_1 \\ \{b_1 : b_1 \geq 0\} & \text{if } b_2 = v_1 \\ \{b_1 : 0 \leq b_1 < b_2\} & \text{if } b_2 > v_1. \end{cases}$$

By similar arguments, player 2's best response function is

$$B_2(b_1) = \begin{cases} \{b_2 : b_2 > b_1\} & \text{if } b_1 < v_2 \\ \{b_2 : b_2 \geq 0\} & \text{if } b_1 = v_2. \\ \{b_2 : 0 \leq b_2 \leq b_1\} & \text{if } b_1 > v_2. \end{cases}$$

These best response functions are shown in Figure 16.1.

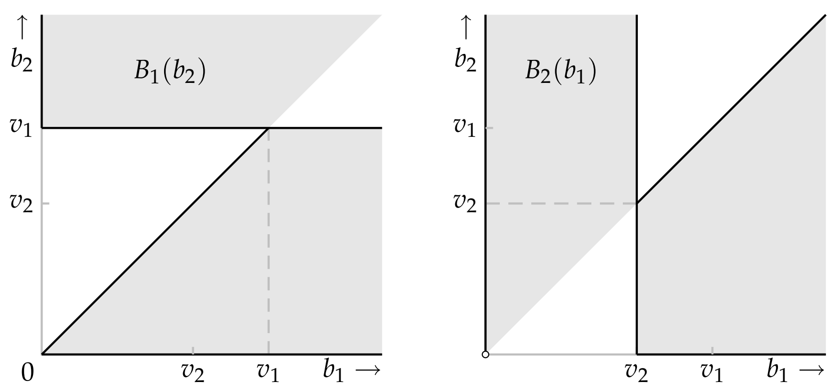

Figure 16.1 The players' best response functions in a two-player second-price sealed-bid auction (Exercise 85.1). Player 1's best response function is in the left panel; player 2's is in the right panel. (Only the edges marked by a black line are included.)

Superimposing the best response functions, we see that the set of Nash equilibria is the shaded set in Figure 17.1, namely the set of pairs (b_1, b_2) such that either

$$b_1 \leq v_2 \text{ and } b_2 \geq v_1$$

or

$$b_1 \geq v_2,\ b_1 \geq b_2, \text{ and } b_2 \leq v_1.$$

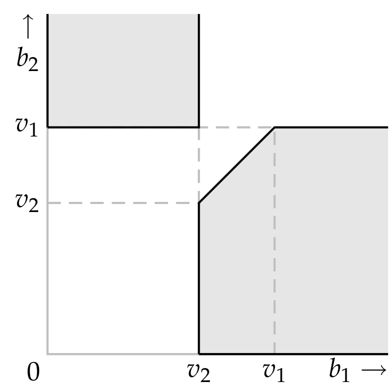

Figure 17.1 The set of Nash equilibria of a two-player second-price sealed-bid auction (Exercise 85.1).

86.2 Nash equilibrium of first-price sealed-bid auction

The profile $(b_1, \ldots, b_n) = (v_2, v_2, v_3, \ldots, v_n)$ is a Nash equilibrium by the following argument.

- If player 1 raises her bid she still wins, but pays a higher price and hence obtains a lower payoff. If player 1 lowers her bid then she loses, and obtains the payoff of 0.
- If any other player changes her bid to any price at most equal to v_2 the outcome does not change. If she raises her bid above v_2 she wins, but obtains a negative payoff.

87.1 First-price sealed-bid auction

A profile of bids in which the two highest bids are not the same is not a Nash equilibrium because the player naming the highest bid can reduce her bid slightly, continue to win, and pay a lower price.

By the argument in the text, in any equilibrium player 1 wins the object. Thus she submits one of the highest bids.

If the highest bid is less than v_2, then player 2 can increase her bid to a value between the highest bid and v_2, win, and obtain a positive payoff. Thus in an equilibrium the highest bid is at least v_2.

If the highest bid exceeds v_1, player 1's payoff is negative, and she can increase this payoff by reducing her bid. Thus in an equilibrium the highest bid is at most v_1.

Finally, any profile $(b_1, \ldots, b_n)$ of bids that satisfies the conditions in the exercise is a Nash equilibrium by the following argument.

- If player 1 increases her bid she continues to win, and reduces her payoff. If player 1 decreases her bid she loses and obtains the payoff 0, which is at most her payoff at $(b_1, \ldots, b_n)$.

- If any other player increases her bid she either does not affect the outcome, or wins and obtains a negative payoff. If any other player decreases her bid she does not affect the outcome.

89.1 All-pay auctions

Second-price all-pay auction with two bidders: The payoff function of bidder i is

$$u_i(b_1, b_2) = \begin{cases} -b_i & \text{if } b_i < b_j \\ v_i - b_j & \text{if } b_i > b_j, \end{cases}$$

with $u_1(b, b) = v_1 - b$ and $u_2(b, b) = -b$ for all b. This payoff function differs from that of player i in the *War of Attrition* only in the payoffs when the bids are equal. The set of Nash equilibria of the game is the same as that for the *War of Attrition*: the set of all pairs $(0, b_2)$ where $b_2 \geq v_1$ and $(b_1, 0)$ where $b_1 \geq v_2$. (The pair (b, b) of actions is not a Nash equilibrium for any value of b because player 2 can increase her payoff by either increasing her bid slightly or by reducing it to 0.)

First-price all-pay auction with two bidders: In any Nash equilibrium the two highest bids are equal, otherwise the player with the higher bid can increase her payoff by reducing her bid a little (keeping it larger than the other player's bid). But no profile of bids in which the two highest bids are equal is a Nash equilibrium, because the player with the higher index who submits this bid can increase her payoff by slightly increasing her bid, so that she wins rather than loses.

90.1 Multiunit auctions

Discriminatory auction To show that the action of bidding v_i and w_i is not dominant for player i, we need only find actions for the other players and alternative bids for player i such that player i's payoff is higher under the alternative bids than it is under the v_i and w_i, given the other players' actions. Suppose that each of the other players submits two bids of 0. Then if player i submits one bid between 0 and v_i and one bid between 0 and w_i she still wins two units, and pays less than when she bids v_i and w_i.

Uniform-price auction Suppose that some bidder other than i submits one bid between w_i and v_i and one bid of 0, and all the remaining bidders submit two bids of 0. Then bidder i wins one unit, and pays the price w_i. If she replaces her bid of w_i with a bid between 0 and w_i then she pays a lower price, and hence is better off.

Vickrey auction Suppose that player i bids v_i and w_i. Consider separately the cases in which the bids of the players other than i are such that player i wins 0, 1, and 2 units.

Player i wins 0 units: In this case the second highest of the other players' bids is at least v_i, so that if player i changes her bids so that she wins one or more units, for any unit she wins she pays at least v_i. Thus no change in her bids increases her payoff from its current value of 0 (and some changes lower her payoff).

Player i wins 1 unit: If player i raises her bid of v_i then she still wins one unit and the price remains the same. If she lowers this bid then either she still wins and pays the same price, or she does not win any units. If she raises her bid of w_i then either the outcome does not change, or she wins a second unit. In the latter case the price she pays is the previously-winning bid she beat, which is at least w_i, so that her payoff either remains zero or becomes negative.

Player i wins 2 units: Player i's raising either of her bids has no effect on the outcome; her lowering a bid either has no effect on the outcome or leads her to lose rather than to win, leading her to obtain the payoff of zero.

90.3 Internet pricing

The situation may be modeled as a multiunit auction in which k units are available, and each player attaches a positive value to only one unit and submits a bid for only one unit. The k highest bids win, and each winner pays the $(k+1)$st highest bid.

By a variant of the argument for a second-price auction, in which "highest of the other players' bids" is replaced by "highest rejected bid", each player's action of bidding her value is weakly dominates all her other actions.

96.2 Alternative standards of care under negligence with contributory negligence

First consider the case in which $X_1 = \hat{a}_1$ and $X_2 \leq \hat{a}_2$. The pair $(\hat{a}_1, \hat{a}_2)$ is a Nash equilibrium by the following argument.

If $a_2 = \hat{a}_2$ then the victim's level of care is sufficient (at least X_2), so that the injurer's payoff is given by (94.1) in the text. Thus the argument that the injurer's action $\hat{a}_1$ is a best response to $\hat{a}_2$ is exactly the same as the argument for the case $X_2 = \hat{a}_2$ in the text.

Since X_1 is the same as before, the victim's payoff is the same also, so that by the argument in the text the victim's best response to $\hat{a}_1$ is $\hat{a}_2$. Thus $(\hat{a}_1, \hat{a}_2)$ is a Nash equilibrium.

To show that $(\hat{a}_1, \hat{a}_2)$ is the only Nash equilibrium of the game, we study the players' best response functions. First consider the injurer's best response function. As in the text, we split the analysis into three cases.

$a_2 < X_2$: In this case the injurer does not have to pay any compensation, regardless of her level of care; her payoff is $-a_1$, so that her best response is $a_1 = 0$.

$a_2 = X_2$: In this case the injurer's best response is $\hat{a}_1$, as argued when showing that $(\hat{a}_1, \hat{a}_2)$ is a Nash equilibrium.

$a_2 > X_2$: In this case the injurer's best response is at most $\hat{a}_1$, since her payoff is equal to $-a_1$ for larger values of a_1.

Thus the injurer's best response takes a form like that shown in the left panel of Figure 20.1. (In fact, $b_1(a_2) = \hat{a}_1$ for $X_2 \leq a_2 \leq \hat{a}_2$, but the analysis depends only on the fact that $b_1(a_2) \leq \hat{a}_1$ for $a_2 > X_2$.)

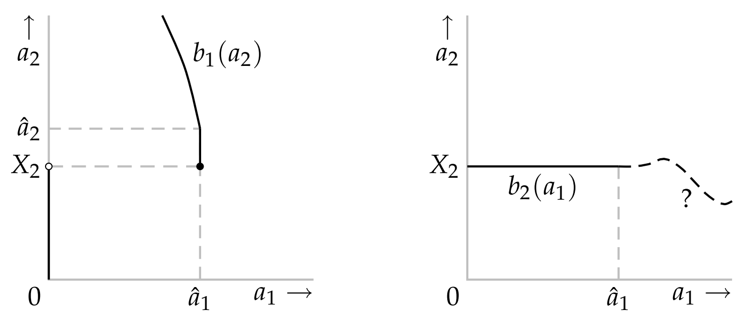

Figure 20.1 The players' best response functions under the rule of negligence with contributory negligence when $X_1 = \hat{a}_1$ and $X_2 = \hat{a}_2$. Left panel: the injurer's best response function b_1. Right panel: the victim's best response function b_2. (The position of the victim's best response function for $a_1 > \hat{a}_1$ is not significant, and is not determined in the solution.)

Now consider the victim's best response function. The victim's payoff function is

$$u_2(a_1, a_2) = \begin{cases} -a_2 & \text{if } a_1 < \hat{a}_1 \text{ and } a_2 \geq X_2 \\ -a_2 - L(a_1, a_2) & \text{if } a_1 \geq \hat{a}_1 \text{ or } a_2 < X_2. \end{cases}$$

As before, for $a_1 < \hat{a}_1$ we have $-a_2 - L(a_1, a_2) < -\hat{a}_2$ for all a_2, so that the victim's best response is X_2. As in the text, the nature of the victim's best responses to levels of care a_1 for which $a_1 > \hat{a}_1$ are not significant.

Combining the two best response functions we see that $(\hat{a}_1, \hat{a}_2)$ is the unique Nash equilibrium of the game.

Now consider the case in which $X_1 = M$ and $a_2 = \hat{a}_2$, where $M \geq \hat{a}_1$. The injurer's payoff is

$$u_1(a_1, a_2) = \begin{cases} -a_1 - L(a_1, a_2) & \text{if } a_1 < M \text{ and } a_2 \geq \hat{a}_2 \\ -a_1 & \text{if } a_1 \geq M \text{ or } a_2 < \hat{a}_2. \end{cases}$$

Now, the maximizer of $-a_1 - L(a_1, \hat{a}_2)$ is $\hat{a}_1$ (see the argument following (94.1) in the text), so that if M is large enough then the injurer's best response to $\hat{a}_2$ is $\hat{a}_1$. As before, if $a_2 < \hat{a}_2$ then the injurer's best response is 0, and if $a_2 > \hat{a}_2$ then the

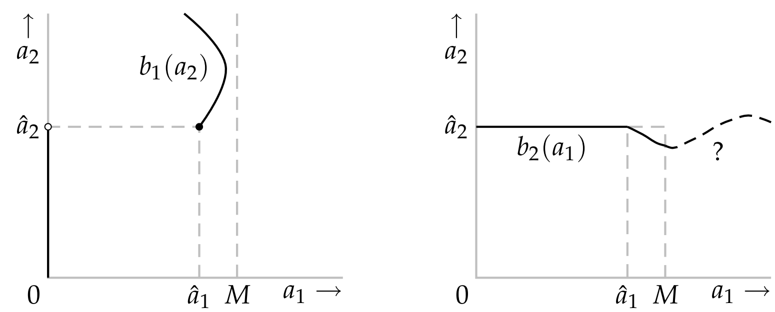

Figure 21.1 The players' best response functions under the rule of negligence with contributory negligence when $(X_1, X_2) = (M, \hat{a}_2)$, with $M \geq \hat{a}_1$. Left panel: the injurer's best response function b_1. Right panel: the victim's best response function b_2. (The position of the victim's best response function for $a_1 > M$ is not significant, and is not determined in the text.)

injurer's payoff decreases for $a_1 > M$, so that her best response is less than M. The injurer's best response function is shown in the left panel of Figure 21.1.

The victim's payoff is

$$u_2(a_1, a_2) = \begin{cases} -a_2 & \text{if } a_1 < M \text{ and } a_2 \geq \hat{a}_2 \\ -a_2 - L(a_1, a_2) & \text{if } a_1 \geq M \text{ or } a_2 < \hat{a}_2. \end{cases}$$

If $a_1 \leq \hat{a}_1$ then the victim's best response is $\hat{a}_2$ by the same argument as the one in the text. If a_1 is such that $\hat{a}_1 < a_1 < M$ then the victim's best response is at most $\hat{a}_2$ (since her payoff is decreasing for larger values of a_2). This information about the victim's best response function is recorded in the right panel of Figure 21.1; it is sufficient to deduce that $(\hat{a}_1, \hat{a}_2)$ is the unique Nash equilibrium of the game.

4 Mixed Strategy Equilibrium

101.1 Variant of *Matching Pennies*

The analysis is the same as for *Matching Pennies*. There is a unique steady state, in which each player chooses each action with probability $\frac{1}{2}$.

106.2 Extensions of *BoS* with vNM preferences

In the first case, when player 1 is indifferent between going to her less preferred concert in the company of player 2 and the lottery in which with probability $\frac{1}{2}$ she and player 2 go to different concerts and with probability $\frac{1}{2}$ they both go to her more preferred concert, the Bernoulli payoffs that represent her preferences satisfy the condition

$$u_1(S,S) = \tfrac{1}{2}u_1(S,B) + \tfrac{1}{2}u_1(B,B).$$

If we choose $u_1(S,B) = 0$ and $u_1(B,B) = 2$, then $u_1(S,S) = 1$. Similarly, for player 2 we can set $u_2(B,S) = 0$, $u_2(S,S) = 2$, and $u_2(B,B) = 1$. Thus the Bernoulli payoffs in the left panel of Figure 23.1 are consistent with the players' preferences.

In the second case, when player 1 is indifferent between going to her less preferred concert in the company of player 2 and the lottery in which with probability $\frac{3}{4}$ she and player 2 go to different concerts and with probability $\frac{1}{4}$ they both go to her more preferred concert, the Bernoulli payoffs that represent her preferences satisfy the condition

$$u_1(S,S) = \tfrac{3}{4}u_1(S,B) + \tfrac{1}{4}u_1(B,B).$$

If we choose $u_1(S,B) = 0$ and $u_1(B,B) = 2$ (as before), then $u_1(S,S) = \frac{1}{2}$. Similarly, for player 2 we can set $u_2(B,S) = 0$, $u_2(S,S) = 2$, and $u_2(B,B) = \frac{1}{2}$. Thus the Bernoulli payoffs in the right panel of Figure 23.1 are consistent with the players' preferences.

	Bach	*Stravinsky*
Bach	$2,1$	$0,0$
Stravinsky	$0,0$	$1,2$

	Bach	*Stravinsky*
Bach	$2,\frac{1}{2}$	$0,0$
Stravinsky	$0,0$	$\frac{1}{2},2$

Figure 23.1 The Bernoulli payoffs for two extensions of *BoS*.

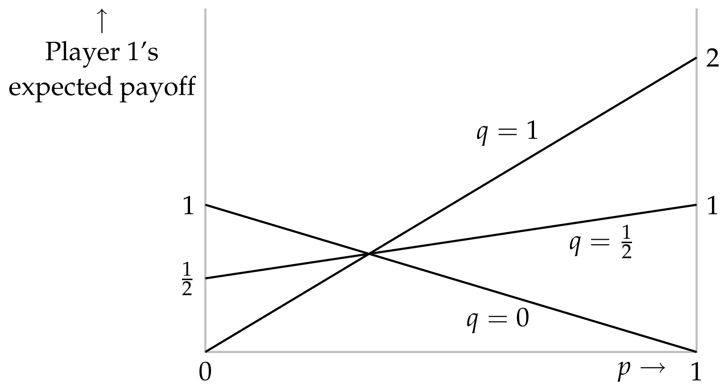

Figure 24.1 Player 1's expected payoff as a function of the probability p that she assigns to B in *BoS*, when the probability q that player 2 assigns to B is 0, $\frac{1}{2}$, and 1.

110.1 Expected payoffs

For *BoS*, player 1's expected payoff is shown in Figure 24.1.

For the game in the right panel of Figure 21.1 in the book, player 1's expected payoff is shown in Figure 24.2.

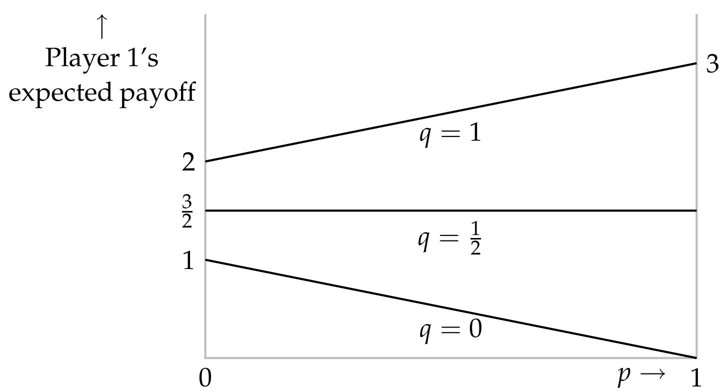

Figure 24.2 Player 1's expected payoff as a function of the probability p that she assigns to *Refrain* in the game in the right panel of Figure 21.1 in the book, when the probability q that player 2 assigns to *Refrain* is 0, $\frac{1}{2}$, and 1.

111.1 Examples of best responses

For *BoS*: for $q = 0$ player 1's unique best response is $p = 0$ and for $q = \frac{1}{2}$ and $q = 1$ her unique best response is $p = 1$. For the game in the right panel of Figure 21.1: for $q = 0$ player 1's unique best response is $p = 0$, for $q = \frac{1}{2}$ her set of best responses is the set of all her mixed strategies (all values of p), and for $q = 1$ her unique best response is $p = 1$.

114.1 Mixed strategy equilibrium of *Hawk–Dove*

Denote by u_i a payoff function whose expected value represents player i's preferences. The conditions in the problem imply that for player 1 we have

$$u_1(Passive, Passive) = \tfrac{1}{2}u_1(Aggressive, Aggressive) + \tfrac{1}{2}u_1(Aggressive, Passive)$$

and

$$u_1(Passive, Aggressive) = \tfrac{2}{3}u_1(Aggressive, Aggressive) + \tfrac{1}{3}u_1(Passive, Passive).$$

Given $u_1(Aggressive, Aggressive) = 0$ and $u_1(Passive, Aggressive = 1$, we have

$$u_1(Passive, Passive) = \tfrac{1}{2}u_1(Aggressive, Passive)$$

and

$$1 = \tfrac{1}{3}u_1(Passive, Passive),$$

so that

$$u_1(Passive, Passive) = 3 \text{ and } u_1(Aggressive, Passive) = 6.$$

Similarly,

$$u_2(Passive, Passive) = 3 \text{ and } u_2(Passive, Aggressive) = 6.$$

Thus the game is given in the left panel of Figure 25.1. The players' best response functions are shown in the right panel. The game has three mixed strategy Nash equilibria: $((0,1),(1,0))$, $((\frac{3}{4},\frac{1}{4}),(\frac{3}{4},\frac{1}{4}))$, and $((1,0),(0,1))$.

	Aggressive	*Passive*
Aggressive	0, 0	6, 1
Passive	1, 6	3, 3

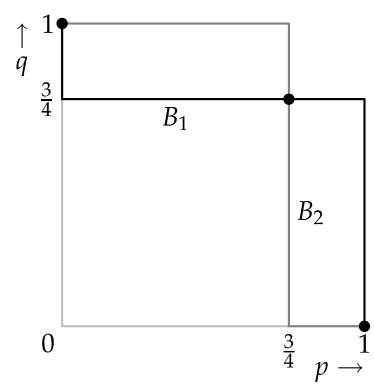

Figure 25.1 An extension of *Hawk–Dove* (left panel) and the players' best response functions when randomization is allowed in this game (right panel). The probability that player 1 assigns to *Aggressive* is p and the probability that player 2 assigns to *Aggressive* is q. The disks indicate the Nash equilibria (two pure, one mixed).

117.2 Choosing numbers

a. To show that the pair of mixed strategies in the question is a mixed strategy equilibrium, it suffices to verify the conditions in Proposition 116.2. Thus, given that each player's strategy specifies a positive probability for every action, it suffices to show that each action of each player yields the same expected payoff. Player 1's expected payoff to each pure strategy is $1/K$, because with probability $1/K$ player 2 chooses the same number, and with probability $1 - 1/K$ player 2 chooses a different number. Similarly, player 2's expected payoff to each pure strategy is $-1/K$, because with probability $1/K$ player 1 chooses the same number, and with probability $1 - 1/K$ player 2 chooses a different number. Thus the pair of strategies is a mixed strategy Nash equilibrium.

b. Let (p^*, q^*) be a mixed strategy equilibrium, where p^* and q^* are vectors, the jth components of which are the probabilities assigned to the integer j by each player. Given that player 2 uses the mixed strategy q^*, player 1's expected payoff if she chooses the number k is q_k^*. Hence if $p_k^* > 0$ then (by the first condition in Proposition 116.2) we need $q_k^* \geq q_j^*$ for all j, so that, in particular, $q_k^* > 0$ (q_j^* cannot be zero for all j!). But player 2's expected payoff if she chooses the number k is $-p_k^*$, so given $q_k^* > 0$ we need $p_k^* \leq p_j^*$ for all j (again by the first condition in Proposition 116.2), and, in particular, $p_k^* \leq 1/K$ (p_j^* cannot exceed $1/K$ for all j!). We conclude that any probability p_k^* that is positive must be at most $1/K$. The only possibility is that $p_k^* = 1/K$ for all k. A similar argument implies that $q_k^* = 1/K$ for all k.

120.2 Strictly dominating mixed strategies

Denote the probability that player 1 assigns to T by p and the probability she assigns to M by r (so that the probability she assigns to B is $1 - p - r$). A mixed strategy of player 1 strictly dominates T if and only if

$$p + 4r > 1 \quad \text{and} \quad p + 3(1 - p - r) > 1,$$

or if and only if $1 - 4r < p < 1 - \frac{3}{2}r$. For example, the mixed strategies $(\frac{1}{4}, \frac{1}{4}, \frac{1}{2})$ and $(0, \frac{1}{3}, \frac{2}{3})$ both strictly dominate T.

120.3 Strict domination for mixed strategies

(*a*) True. Suppose that the mixed strategy α_i' assigns positive probability to the action a_i', which is strictly dominated by the action a_i. Then $u_i(a_i, a_{-i}) > u_i(a_i', a_{-i})$ for all a_{-i}. Let α_i be the mixed strategy that differs from α_i' only in the weight that α_i' assigns to a_i' is transferred to a_i. That is, α_i is defined by $\alpha_i(a_i') = 0$, $\alpha_i(a_i) = \alpha_i'(a_i') + \alpha_i'(a_i)$, and $\alpha_i(b_i) = \alpha_i'(b_i)$ for every other action b_i. Then α_i strictly dominates α_i': for any a_{-i} we have $U(\alpha_i, a_{-i}) - U(\alpha_i', a_{-i}) = \alpha_i'(a_i')(u(a_i, a_{-i}) - u_i(a_i', a_{-i})) > 0$.

(*b*) False. Consider a variant of the game in Figure 120.1 in the text in which player 1's payoffs to (T, L) and to (T, R) are both $\frac{5}{2}$ instead of 1. Then player 1's mixed strategy that assigns probability $\frac{1}{2}$ to M and probability $\frac{1}{2}$ to B is strictly dominated by T, even though neither M nor B is strictly dominated.

127.1 Equilibrium in the expert diagnosis game

When $E = rE' + (1-r)I'$ the consumer is indifferent between her two actions when $p = 0$, so that her best response function has a vertical segment at $p = 0$. Referring to Figure 126.1 in the text, we see that the set of mixed strategy Nash equilibria correspond to $p = 0$ and $\pi/\pi' \leq q \leq 1$.

130.3 Bargaining

The game is given in Figure 27.1.

	0	2	4	6	8	10
0	5,5	4,6	3,7	2,8	1,9	0,10
2	6,4	5,5	4,6	3,7	2,8	0,0
4	7,3	6,4	5,5	4,6	0,0	0,0
6	8,2	7,3	6,4	0,0	0,0	0,0
8	9,1	8,2	0,0	0,0	0,0	0,0
10	10,0	0,0	0,0	0,0	0,0	0,0

Figure 27.1 A bargaining game.

By inspection it has a single symmetric pure strategy Nash equilibrium, $(10, 10)$.

Now consider situations in which the common mixed strategy assigns positive probability to two actions. Suppose that player 2 assigns positive probability only to 0 and 2. Then player 1's payoff to her action 4 exceeds her payoff to either 0 or 2. Thus there is no symmetric equilibrium in which the actions assigned positive probability are 0 and 2. By a similar argument we can rule out equilibria in which the actions assigned positive probability are any pair except 2 and 8, or 4 and 6.

If the actions to which player 2 assigns positive probability are 2 and 8 then player 1's expected payoffs to 2 and 8 are the same if the probability player 2 assigns to 2 is $\frac{2}{5}$ (and the probability she assigns to 8 is $\frac{3}{5}$). Given these probabilities, player 1's expected payoff to her actions 2 and 8 is $\frac{16}{5}$, and her expected payoff to every other action is less than $\frac{16}{5}$. Thus the pair of mixed strategies in which every player assigns probability $\frac{2}{5}$ to 2 and $\frac{3}{5}$ to 8 is a symmetric mixed strategy Nash equilibrium.

Similarly, the game has a symmetric mixed strategy equilibrium (α^*, α^*) in which α^* assigns probability $\frac{4}{5}$ to the demand of 4 and probability $\frac{1}{5}$ to the demand of 6.

In summary, the game has three symmetric mixed strategy Nash equilibria in which each player's strategy assigns positive probability to at most two actions: one in which probability 1 is assigned to 10, one in which probability $\frac{2}{5}$ is assigned to 2 and probability $\frac{3}{5}$ is assigned to 8, and one in which probability $\frac{4}{5}$ is assigned to 4 and probability $\frac{1}{5}$ is assigned to 6.

132.2 Reporting a crime when the witnesses are heterogeneous

Denote by p_i the probability with which each witness with cost c_i reports the crime, for $i = 1, 2$. For each witness with cost c_1 to report with positive probability less than one, we need

$$\begin{aligned} v - c_1 &= v \cdot \Pr\{\text{at least one other person calls}\} \\ &= v\left(1 - (1-p_1)^{n_1-1}(1-p_2)^{n_2}\right), \end{aligned}$$

or

$$c_1 = v(1-p_1)^{n_1-1}(1-p_2)^{n_2}. \tag{28.1}$$

Similarly, for each witness with cost c_2 to report with positive probability less than one, we need

$$\begin{aligned} v - c_2 &= v \cdot \Pr\{\text{at least one other person calls}\} \\ &= v\left(1 - (1-p_1)^{n_1}(1-p_2)^{n_2-1}\right), \end{aligned}$$

or

$$c_2 = v(1-p_1)^{n_1}(1-p_2)^{n_2-1}. \tag{28.2}$$

Dividing (28.1) by (28.2) we obtain

$$1 - p_2 = c_1(1-p_1)/c_2.$$

Substituting this expression for $1 - p_2$ into (28.1) we get

$$p_1 = 1 - \left(\frac{c_1}{v} \cdot \left(\frac{c_2}{c_1}\right)^{n_2}\right)^{1/(n-1)}.$$

Similarly,

$$p_2 = 1 - \left(\frac{c_2}{v} \cdot \left(\frac{c_1}{c_2}\right)^{n_1}\right)^{1/(n-1)}.$$

For these two numbers to be probabilities, we need each of them to be nonnegative and at most one, which requires

$$\left(\frac{c_2^{n_2}}{v}\right)^{1/(n_2-1)} < c_1 < \left(vc_2^{n_1-1}\right)^{1/n_1}.$$

136.1 Best response dynamics in Cournot's duopoly game

The best response functions of both firms are the same, so if the firms' outputs are initially the same, they are the same in every period: $q_1^t = q_2^t$ for every t. For each period t, we thus have

$$q_i^t = \tfrac{1}{2}(\alpha - c - q_i^t).$$

Given that $q_i^1 = 0$ for $i = 1, 2$, solving this first-order difference equation we have

$$q_i^t = \tfrac{1}{3}(\alpha - c)[1 - (-\tfrac{1}{2})^{t-1}]$$

for each period t. When t is large, q_i^t is close to $\frac{1}{3}(\alpha - c)$, a firm's equilibrium output.

In the first few periods, these outputs are 0, $\frac{1}{2}(\alpha - c)$, $\frac{1}{4}(\alpha - c)$, $\frac{3}{8}(\alpha - c)$, $\frac{5}{16}(\alpha - c)$.

139.1 Finding all mixed strategy equilibria of two-player games

Left game:

- There is no equilibrium in which each player's mixed strategy assigns positive probability to a single action (i.e. there is no pure equilibrium).

- Consider the possibility of an equilibrium in which one player assigns probability 1 to a single action while the other player assigns positive probability to both her actions. For neither action of player 1 is player 2's payoff the same for both her actions, and for neither action of player 2 is player 1's payoff the same for both her actions, so there is no mixed strategy equilibrium of this type.

- Consider the possibility of a mixed strategy equilibrium in which each player assigns positive probability to both her actions. Denote by p the probability player 1 assigns to T and by q the probability player 2 assigns to L. For player 1's expected payoff to her two actions to be the same we need

  $$6q = 3q + 6(1 - q),$$

  or $q = \frac{2}{3}$. For player 2's expected payoff to her two actions to be the same we need

  $$2(1 - p) = 6p,$$

  or $p = \frac{1}{4}$. We conclude that the game has a unique mixed strategy equilibrium, $((\frac{1}{4}, \frac{3}{4}), (\frac{2}{3}, \frac{1}{3}))$.

Right game:

- By inspection, (T, R) and (B, L) are the pure strategy equilibria.

- Consider the possibility of a mixed strategy equilibrium in which one player assigns probability 1 to a single action while the other player assigns positive probability to both her actions.
  - $\{T\}$ for player 1, $\{L, R\}$ for player 2: no equilibrium, because player 2's payoffs to (T, L) and (T, R) are not the same.
  - $\{B\}$ for player 1, $\{L, R\}$ for player 2: no equilibrium, because player 2's payoffs to (B, L) and (B, R) are not the same.
  - $\{T, B\}$ for player 1, $\{L\}$ for player 2: no equilibrium, because player 1's payoffs to (T, L) and (B, L) are not the same.
  - $\{T, B\}$ for player 1, $\{R\}$ for player 2: player 1's payoffs to (T, R) and (B, R) are the same, so there is an equilibrium in which player 1 uses T with probability p if player 2's expected payoff to R, which is $2p + 1 - p$, is at least her expected payoff to L, which is $p + 2(1 - p)$. That is, the game has equilibria in which player 1's mixed strategy is $(p, 1 - p)$, with $p \geq \frac{1}{2}$, and player 2 uses R with probability 1.
- Consider the possibility of an equilibrium in which both players assign positive probability to both their actions. Denote by q the probability that player 2 assigns to L. For player 1's expected payoffs to T and B to be the same we need $0 = 2q$, or $q = 0$, so there is no equilibrium in which both players assign positive probability to both their actions.

In summary, the mixed strategy equilibria of the game are $((0, 1), (1, 0))$ (i.e. the pure equilibrium (B, L)) and $((p, 1 - p), (0, 1))$ for $\frac{1}{2} \leq p \leq 1$ (of which one equilibrium is the pure equilibrium (T, R)).

145.1 All-pay auction with many bidders

Denote the common mixed strategy by F. Look for an equilibrium in which the largest value of z for which $F(z) = 0$ is 0 and the smallest value of z for which $F(z) = 1$ is $z = K$.

A player who bids a_i wins if and only if the other $n - 1$ players all bid less than she does, an event with probability $(F(a_i))^{n-1}$. Thus, given that the probability that she ties for the highest bid is zero, her expected payoff is

$$(K - a_i)(F(a_i))^{n-1} + (-a_i)(1 - (F(a_i))^{n-1}).$$

Given the form of F, for an equilibrium this expected payoff must be constant for all values of a_i with $0 \leq a_i \leq K$. That is, for some value of c we have

$$K(F(a_i))^{n-1} - a_i = c \text{ for all } 0 \leq a_i \leq K.$$

For $F(0) = 0$ we need $c = 0$, so that $F(a_i) = (a_i/K)^{1/(n-1)}$ is the only candidate for an equilibrium strategy.

The function F is a cumulative probability distribution on the interval from 0 to K because $F(0) = 0$, $F(K) = 1$, and F is increasing. Thus F is indeed an equilibrium strategy.

We conclude that the game has a mixed strategy Nash equilibrium in which each player randomizes over all her actions according to the probability distribution $F(a_i) = (a_i/K)^{1/(n-1)}$; each player's equilibrium expected payoff is 0.

Each player's mean bid is K/n.

147.2 Preferences over lotteries

The first piece of information about the decision-maker's preferences among lotteries is consistent with her preferences being represented by the expected value of a payoff function: set $u(a_1) = 0$, $u(a_2)$ equal to any number between $\frac{1}{2}$ and $\frac{1}{4}$, and $u(a_3) = 1$.

The second piece of information about the decision-maker's preferences is not consistent with these preferences being represented by the expected value of a payoff function, by the following argument. For consistency with the information about the decision-maker's preferences among the four lotteries, we need

$$\begin{aligned} 0.4u(a_1) + 0.6u(a_3) > 0.5u(a_2) \ + \ 0.5u(a_3) > \\ 0.3u(a_1) + 0.2u(a_2) \ + \ 0.5u(a_3) > 0.45u(a_1) + 0.55u(a_3). \end{aligned}$$

The first inequality implies $u(a_2) < 0.8u(a_1) + 0.2u(a_3)$ and the last inequality implies $u(a_2) > 0.75u(a_1) + 0.25u(a_3)$. Because $u(a_1) < u(a_3)$, we have $0.75u(a_1) + 0.25u(a_3) > 0.8u(a_1) + 0.2u(a_3)$, so that the two inequalities are incompatible.

149.2 Normalized vNM payoff functions

Let $\overline{a}$ be the best outcome according to her preferences and let $\underline{a}$ be the worse outcome. Let $\eta = -u(\underline{a})/(u(\overline{a}) - u(\underline{a}))$ and $\theta = 1/(u(\overline{a}) - u(\underline{a})) > 0$. Lemma 148.1 implies that the function v defined by $v(x) = \eta + \theta u(x)$ represents the same preferences as does u; we have $v(\underline{a}) = 0$ and $v(\overline{a}) = 1$.

5 Extensive Games with Perfect Information: Theory

163.1 Nash equilibria of extensive games

The strategic form of the game in Exercise 156.2*a* is given in Figure 33.1.

	EG	EH	FG	FH
C	$1,0$	$1,0$	$3,2$	$3,2$
D	$2,3$	$0,1$	$2,3$	$0,1$

Figure 33.1 The strategic form of the game in Exercise 156.2*a*.

The Nash equilibria of the game are (C, FG), (C, FH), and (D, EG).
The strategic form of the game in Figure 160.1 is given in Figure 33.2.

	E	F
CG	$1,2$	$3,1$
CH	$0,0$	$3,1$
DG	$2,0$	$2,0$
DH	$2,0$	$2,0$

Figure 33.2 The strategic form of the game in Figure 160.1.

The Nash equilibria of the game are (CH, F), (DG, E), and (DH, E).

164.2 Subgames

The subgames of the game in Exercise 156.2*c* are the whole game and the six games in Figure 34.1.

168.1 Checking for subgame perfect equilibria

The Nash equilibria (CH, F) and (DH, E) are not subgame perfect equilibria: in the subgame following the history (C, E), player 1's strategies CH and DH induce the strategy H, which is not optimal.

The Nash equilibrium (DG, E) is a subgame perfect equilibrium: (*a*) it is a Nash equilibrium, so player 1's strategy is optimal at the start of the game, given player 2's strategy, (*b*) in the subgame following the history C, player 2's strategy E induces the strategy E, which is optimal given player 1's strategy, and (*c*) in the subgame following the history (C, E), player 1's strategy DG induces the strategy G, which is optimal.

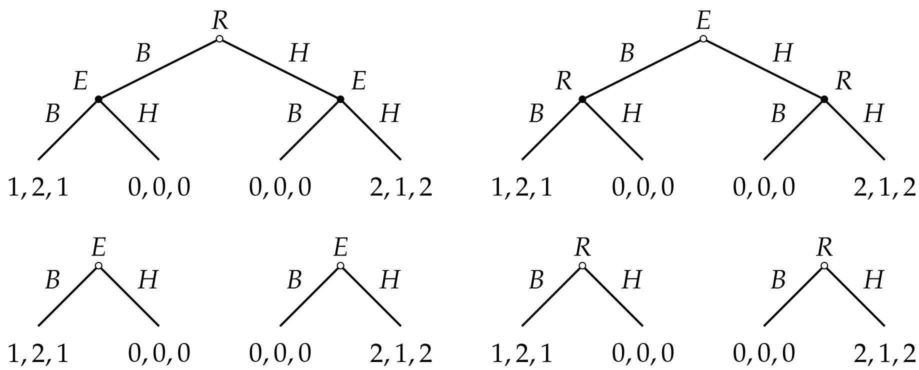

Figure 34.1 The proper subgames of the game in Exercise 156.2*c*.

174.1 Sharing heterogeneous objects

Let $n = 2$ and $k = 3$, and call the objects a, b, and c. Suppose that the values person 1 attaches to the objects are 3, 2, and 1 respectively, while the values player 2 attaches are 1, 3, 2. If player 1 chooses a on the first round, then in any subgame perfect equilibrium player 2 chooses b, leaving player 1 with c on the second round. If instead player 1 chooses b on the first round, in any subgame perfect equilibrium player 2 chooses c, leaving player 1 with a on the second round. Thus in every subgame perfect equilibrium player 1 chooses b on the first round (though she values a more highly.)

Now I argue that for any preferences of the players, $G(2,3)$ has a subgame perfect equilibrium of the type described in the exercise. For any object chosen by player 1 in round 1, in any subgame perfect equilibrium player 2 chooses her favorite among the two objects remaining in round 2. Thus player 2 never obtains the object she least prefers; in any subgame perfect equilibrium, player 1 obtains that object. Player 1 can ensure she obtains her more preferred object of the two remaining by choosing that object on the first round. That is, there is a subgame perfect equilibrium in which on the first round player 1 chooses her more preferred object out of the set of objects excluding the object player 2 least prefers, and on the last round she obtains x_3. In this equilibrium, player 2 obtains the object less preferred by player 1 out of the set of objects excluding the object player 2 least prefers. That is, player 2 obtains x_2. (Depending on the players' preferences, the game also may have a subgame perfect equilibrium in which player 1 chooses x_3 on the first round.)

177.3 Comparing simultaneous and sequential games

a. Denote by (a_1^*, a_2^*) a Nash equilibrium of the strategic game in which player 1's payoff is maximal in the set of Nash equilibria. Because (a_1^*, a_2^*) is a Nash equilibrium, a_2^* is a best response to a_1^*. By assumption, it is the only best response to a_1^*. Thus if player 1 chooses a_1^* in the extensive game, player 2 must choose a_2^* in any subgame perfect equilibrium of the extensive game. That is, by choosing a_1^*, player 1 is assured of a payoff of at least $u_1(a_1^*, a_2^*)$. Thus in any subgame perfect equilibrium player 1's payoff must be at least $u_1(a_1^*, a_2^*)$.

b. Suppose that $A_1 = \{T, B\}$, $A_2 = \{L, R\}$, and the payoffs are those given in Figure 35.1. The strategic game has a unique Nash equilibrium, (T, L), in which player 2's payoff is 1. The extensive game has a unique subgame perfect equilibrium, (B, LR) (where the first component of player 2's strategy is her action after the history T and the second component is her action after the history B). In this subgame perfect equilibrium player 2's payoff is 2.

	L	R
T	1, 1	3, 0
B	0, 0	2, 2

Figure 35.1 The payoffs for the example in Exercise 177.3*b*.

c. Suppose that $A_1 = \{T, B\}$, $A_2 = \{L, R\}$, and the payoffs are those given in Figure 35.2. The strategic game has a unique Nash equilibrium, (T, L), in which player 2's payoff is 2. A subgame perfect equilibrium of the extensive game is (B, RL) (where the first component of player 2's strategy is her action after the history T and the second component is her action after the history B). In this subgame perfect equilibrium player 1's payoff is 1. (If you read Chapter 4, you can find the mixed strategy Nash equilibria of the strategic game; in all these equilibria, as in the pure strategy Nash equilibrium, player 1's expected payoff exceeds 1.)

	L	R
T	2, 2	0, 2
B	1, 1	3, 0

Figure 35.2 The payoffs for the example in Exercise 177.3*c*.

179.3 Three Men's Morris, or Mill

Number the squares 1 through 9, starting at the top left, working across each row. The following strategy of player 1 guarantees she wins, so that the subgame perfect equilibrium outcome is that she wins. First player 1 chooses the central square (5).

- Suppose player 2 then chooses a corner; take it to be square 1. Then player 1 chooses square 6. Now player 2 must choose square 4 to avoid defeat; player 1 must choose square 7 to avoid defeat; and then player 2 must choose square

3 to avoid defeat (otherwise player 1 can move from square 6 to square 3 on her next turn). If player 1 now moves from square 6 to square 9, then whatever player 2 does she can subsequently move her counter from square 5 to square 8 and win.

- Suppose player 2 then chooses a noncorner; take it to be square 2. Then player 1 chooses square 7. Now player 2 must choose square 3 to avoid defeat; player 1 must choose square 1 to avoid defeat; and then player 2 must choose square 4 to avoid defeat (otherwise player 1 can move from square 5 to square 4 on her next turn). If player 1 now moves from square 7 to square 8, then whatever player 2 does she can subsequently move from square 8 to square 9 and win.

6 Extensive Games with Perfect Information: Illustrations

183.1 Nash equilibria of the ultimatum game

For *every* amount x there are Nash equilibria in which person 1 offers x. For example, for any value of x there is a Nash equilibrium in which person 1's strategy is to offer x and person 2's strategy is to accept x and any offer more favorable, and reject every other offer. (Given person 2's strategy, person 1 can do no better than offer x. Given person 1's strategy, person 2 should accept x; whether person 2 accepts or rejects any other offer makes no difference to her payoff, so that rejecting all less favorable offers is, in particular, optimal.)

183.2 Subgame perfect equilibria of the ultimatum game with indivisible units

In this case each player has finitely many actions, and for both possible subgame perfect equilibrium strategies of player 2 there is an optimal strategy for player 1.

If player 2 accepts all offers then player 1's best strategy is to offer 0, as before.

If player 2 accepts all offers except 0 then player 1's best strategy is to offer one cent (which player 2 accepts).

Thus the game has two subgame perfect equilibria: one in which player 1 offers 0 and player 2 accepts all offers, and one in which player 1 offers one cent and player 2 accepts all offers except 0.

186.1 Holdup game

The game is defined as follows.

Players Two people, person 1 and person 2.

Terminal histories The set of all sequences (low, x, Z), where x is a number with $0 \leq x \leq c_L$ (the amount of money that person 1 offers to person 2 when the pie is small), and $(high, x, Z)$, where x is a number with $0 \leq x \leq c_H$ (the amount of money that person 1 offers to person 2 when the pie is large) and Z is either Y ("yes, I accept") or N ("no, I reject").

Player function $P(\varnothing) = 2$, $P(low) = P(high) = 1$, and $P(low, x) = P(high, x) = 2$ for all x.

Preferences Person 1's preferences are represented by payoffs equal to the amounts of money she receives, equal to $c_L - x$ for any terminal history (low, x, Y) with $0 \leq x \leq c_L$, equal to $c_H - x$ for any terminal history

$(high, x, Y)$ with $0 \leq x \leq c_H$, and equal to 0 for any terminal history (low, x, N) with $0 \leq x \leq c_L$ and for any terminal history $(high, x, N)$ with $0 \leq x \leq c_H$. Person 2's preferences are represented by payoffs equal to $x - L$ for the terminal history (low, x, Y), $x - H$ for the terminal history $(high, x, Y)$, $-L$ for the terminal history (low, x, N), and $-H$ for the terminal history $(high, x, N)$.

189.1 Stackelberg's duopoly game with quadratic costs

From Exercise 59.1, the best response function of firm 2 is the function b_2 defined by

$$b_2(q_1) = \begin{cases} \frac{1}{4}(\alpha - q_1) & \text{if } q_1 \leq \alpha \\ 0 & \text{if } q_1 > \alpha. \end{cases}$$

Firm 1's subgame perfect equilibrium strategy is the value of q_1 that maximizes $q_1(\alpha - q_1 - b_2(q_1)) - q_1^2$, or $q_1(\alpha - q_1 - \frac{1}{4}(\alpha - q_1)) - q_1^2$, or $\frac{1}{4}q_1(3\alpha - 7q_1)$. The maximizer is $q_1 = \frac{3}{14}\alpha$.

We conclude that the game has a unique subgame perfect equilibrium, in which firm 1's strategy is the output $\frac{3}{14}\alpha$ and firm 2's strategy is its best response function b_2.

The outcome of the subgame perfect equilibrium is that firm 1 produces $q_1^* = \frac{3}{14}\alpha$ units of output and firm 2 produces $q_2^* = b_2(\frac{3}{14}\alpha) = \frac{11}{56}\alpha$ units. In a Nash equilibrium of Cournot's (simultaneous-move) game each firm produces $\frac{1}{5}\alpha$ (see Exercise 59.1). Thus firm 1 produces more in the subgame perfect equilibrium of the sequential game than it does in the Nash equilibrium of Cournot's game, and firm 2 produces less.

196.4 Sequential positioning by three political candidates

The following extensive game models the situation.

Players The candidates.

Terminal histories The set of all sequences $(x_1, \ldots, x_n)$, where x_i is either *Out* or a position of candidate i (a number) for $i = 1, \ldots, n$.

Player function $P(\varnothing) = 1$, $P(x_1) = 2$ for all x_1, $P(x_1, x_2) = 3$ for all (x_1, x_2), ..., $P(x_1, \ldots, x_{n-1}) = n$ for all $(x_1, \ldots, x_{n-1})$.

Preferences Each candidate's preferences are represented by a payoff function that assigns n to every terminal history in which she wins, k to every terminal history in which she ties for first place with $n - k$ other candidates, for $1 \leq k \leq n - 1$, 0 to every terminal history in which she stays out, and -1 to every terminal history in which she loses, where positions attract votes as in Hotelling's model of electoral competition (Section 3.3).

When there are two candidates the analysis of the subgame perfect equilibria is similar to that in the previous exercise. In every subgame perfect equilibrium candidate 1's strategy is m; candidate 2's strategy chooses m after the history m, some position between x_1 and $2m - x_1$ after the history x_1 for any position x_1, and any position after the history *Out*.

Now consider the case of three candidates when the voters' favorite positions are distributed uniformly from 0 to 1. I claim that every subgame perfect equilibrium results in the first candidate's entering at $\frac{1}{2}$, the second candidate's staying out, and the third candidate's entering at $\frac{1}{2}$.

To show this, first consider the best response of candidate 3 to each possible pair of actions of candidates 1 and 2. Figure 39.1 illustrates these optimal actions in every case that candidate 1 enters. (If candidate 1 does not enter then the subgame is exactly the two-candidate game.)

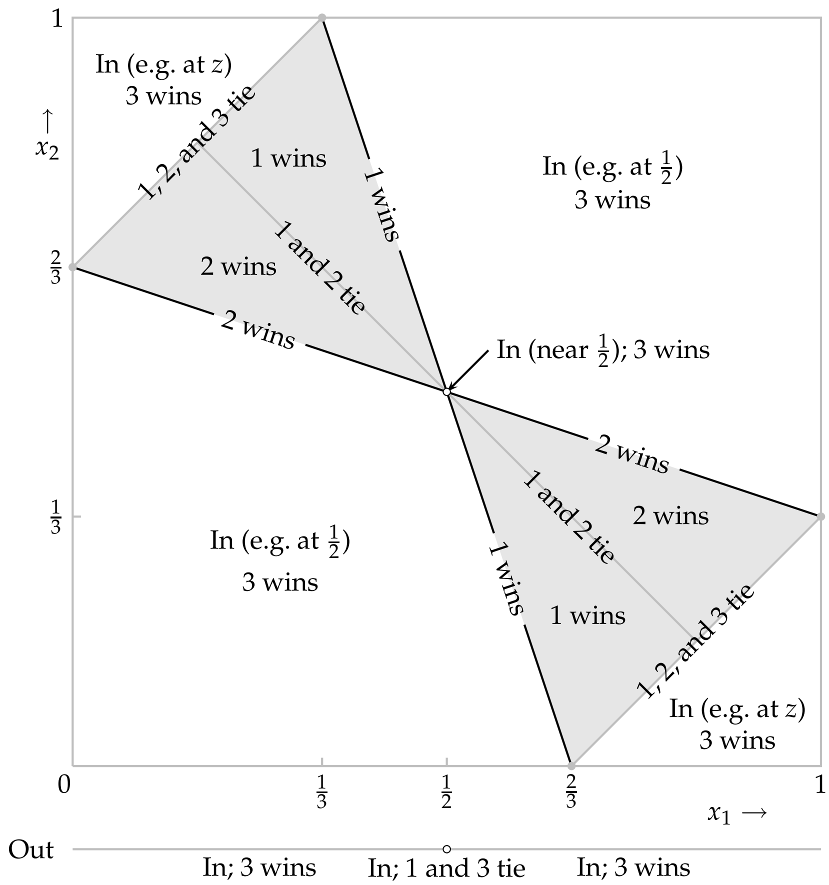

Figure 39.1 The outcome of a best response of candidate 3 to each pair of actions by candidates 1 and 2. The best response for any point in the gray shaded area (including the black boundaries of this area, but excluding the other boundaries) is *Out*. The outcome at each of the four small disks at the outer corners of the shaded area is that all three candidates tie. The value of z is $1 - \frac{1}{2}(x_1 + x_2)$.

Now consider the optimal action of candidate 2, given x_1 and the outcome of candidate 3's best response, as given in Figure 39.1. In the figure, take a value of x_1 and look at the outcomes as x_2 varies; find the value of x_2 that induces the best outcome for candidate 2. For example, for $x_1 = 0$ the only value of x_2 for which candidate 2 does not lose is $\frac{2}{3}$, at which point she ties with the other two candidates. Thus when candidate 1's strategy is $x_1 = 0$, candidate 2's best action, given candidate 3's best response, is $x_2 = \frac{2}{3}$, which leads to a three-way tie. We find that the outcome of the optimal value of x_2, for each value of x_1, is given as follows.

$$\begin{cases} 1, 2, \text{and } 3 \text{ tie } (x_2 = \frac{2}{3}) & \text{if } x_1 = 0 \\ 2 \text{ wins} & \text{if } 0 < x_1 < \frac{1}{2} \\ 1 \text{ and } 3 \text{ tie } (2 \text{ stays out}) & \text{if } x_1 = \frac{1}{2} \\ 2 \text{ wins} & \text{if } \frac{1}{2} < x_1 < 1 \\ 1, 2, \text{and } 3 \text{ tie } (x_2 = \frac{1}{3}) & \text{if } x_1 = 1. \end{cases}$$

Finally, consider candidate 1's best strategy, given the responses of candidates 2 and 3. If she stays out then candidates 2 and 3 enter at m and tie. If she enters then the best position at which to do so is $x_1 = \frac{1}{2}$, where she ties with candidate 3. (For every other position she either loses or ties with both of the other candidates.)

We conclude that in every subgame perfect equilibrium the outcome is that candidate 1 enters at $\frac{1}{2}$, candidate 2 stays out, and candidate 3 enters at $\frac{1}{2}$. (There are many subgame perfect equilibria, because after many histories candidate 3's optimal action is not unique.)

(The case in which there are many potential candidates, is discussed on the page `http://www.economics.utoronto.ca/osborne/research/CONJECT.HTM`.)

198.1 The race $G_1(2, 2)$

The consequences of player 1's actions at the start of the game are as follows.

Take two steps: Player 1 wins.

Take one step: Go to the game $G_2(1, 2)$, in which player 2 initially takes two steps and wins.

Do not move: If player 2 does not move, the game ends. If she takes one step we go to the game $G_1(2, 1)$, in which player 1 takes two steps and wins. If she takes two steps, she wins. Thus in a subgame perfect equilibrium player 2 takes two steps, and wins.

We conclude that in a subgame perfect equilibrium of $G_1(2, 2)$ player 1 initially takes two steps, and wins.

203.1 A race with a liquidity constraint

In the absence of the constraint, player 1 initially takes one step. Suppose she does so in the game with the constraint. Consider player 2's options after player 1's move.

Player 2 takes two steps: Because of the liquidity constraint, player 1 can take at most one step. If she takes one step, player 2's optimal action is to take one step, and win. Thus player 1's best action is not to move; player 2's payoff exceeds 1 (her steps cost 5, and the prize is worth more than 6).

Player 2 moves one step: Again because of the liquidity constraint, player 1 can take at most one step. If she takes one step, player 2 can take two steps and win, obtaining a payoff of more than 1 (as in the previous case).

Player 2 does not move: Player 1, as before, can take one step on each turn, and win; player 2's payoff is 0.

We conclude that after player 1 moves one step, player 2 should take either one or two steps, and ultimately win; player 1's payoff is -1. A better option for player 1 is not to move, in which case player 2 can move one step at a time, and win; player 1's payoff is zero.

Thus the subgame perfect equilibrium outcome is that player 1 does not move, and player 2 takes one step at a time and wins.

7 Extensive Games with Perfect Information: Extensions and Discussion

210.2 Extensive game with simultaneous moves

The game is shown in Figure 43.1.

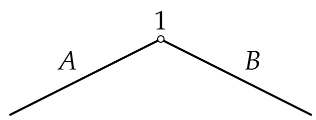

	C	D
C	$4,2$	$0,0$
D	$0,0$	$2,4$

	E	F
E	$3,1$	$0,0$
F	$0,0$	$1,3$

Figure 43.1 The game in Exercise 210.2.

The subgame following player 1's choice of A has two Nash equilibria, (C, C) and (D, D); the subgame following player 1's choice of B also has two Nash equilibria, (E, E) and (F, F). If the equilibrium reached after player 1 chooses A is (C, C), then regardless of the equilibrium reached after she chooses (E, E), she chooses A at the beginning of the game. If the equilibrium reached after player 1 chooses A is (D, D) and the equilibrium reached after she chooses B is (F, F), she chooses A at the beginning of the game. If the equilibrium reached after player 1 chooses A is (D, D) and the equilibrium reached after she chooses B is (E, E), she chooses B at the beginning of the game.

Thus the game has four subgame perfect equilibria: (ACE, CE), (ACF, CF), (ADF, DF), and (BDE, DE) (where the first component of player 1's strategy is her choice at the start of the game, the second component is her action after she chooses A, and the third component is her action after she chooses B, and the first component of player 2's strategy is her action after player 1 chooses A at the start of the game and the second component is her action after player 1 chooses B at the start of the game).

In the first two equilibria the outcome is that player 1 chooses A and then both players choose C, in the third equilibrium the outcome is that player 1 chooses A and then both players choose D, and in the last equilibrium the outcome is that player 1 chooses B and then both players choose E.

217.1 Electoral competition with strategic voters

I first argue that in any equilibrium each candidate that enters is in the set of winners. If some candidate that enters is not a winner, she can increase her payoff by deviating to *Out*.

Now consider the voting subgame in which there are more than two candidates and not all candidates' positions are the same. Suppose that the citizens' votes are equally divided among the candidates. I argue that this list of citizens' strategies is not a Nash equilibrium of the voting subgame.

For either the citizen whose favorite position is 0 or the citizen whose favorite position is 1 (or both), at least two candidates' positions are better than the position of the candidate furthest from the citizen's favorite position. Denote a citizen for whom this condition holds by i. (The claim that citizen i exists is immediate if the candidates occupy at least three distinct positions, or they occupy two distinct positions and at least two candidates occupy each position. If the candidates occupy only two positions and one position is occupied by a single candidate, then take the citizen whose favorite position is 0 if the lone candidate's position exceeds the other candidates' position; otherwise take the citizen whose favorite position is 1.)

Now, given that each candidate obtains the same number of votes, if citizen i switches her vote to one of the candidates whose position is better for her than that of the candidate whose position is furthest from her favorite position, then this candidate wins outright. (If citizen i originally votes for one of these superior candidates, she can switch her vote to the other superior candidate; if she originally votes for neither of the superior candidates, she can switch her vote to either one of them.) Citizen i's payoff increases when she thus switches her vote, so that the list of citizens' strategies is not a Nash equilibrium of the voting subgame.

We conclude that in every Nash equilibrium of every voting subgame in which there are more than two candidates and not all candidates' positions are the same at least one candidate loses. Because no candidate loses in a subgame perfect equilibrium (by the first argument in the proof), in any subgame perfect equilibrium either only two candidates enter, or all candidates' positions are the same.

If only two candidates enter, then by the argument in the text for the case $n = 2$, each candidate's position is m (the median of the citizens' favorite positions).

Now suppose that more than two candidates enter, and their common position is not equal to m. If a candidate deviates to m then in the resulting voting subgame only two positions are occupied, so that for every citizen, any strategy that is not weakly dominated votes for a candidate at the position closest to her favorite position. Thus a candidate who deviates to m wins outright. We conclude that in any subgame perfect equilibrium in which more than two candidates enter, they all choose the position m.

220.1 Top cycle set

a. The top cycle set is the set $\{x, y, z\}$ of all three alternatives because x beats y beats z beats x.

b. The top cycle set is the set $\{w, x, y, z\}$ of all four alternatives. As in the previous case, x beats y beats z beats x; also y beats w.

224.1 Exit from a declining industry

Period t_1 is the largest value of t for which $P_t(k_1) \geq c$, or $60 - t \geq 10$. Thus $t_1 = 50$. Similarly, $t_2 = 70$.

If both firms are active in period t_1, then firm 2's profit in this period is $-ck_2 = -10(20) = -200$. (Note that the price is zero, because $k_1 + k_2 > 50$.) Its profit in any period t in which it is alone in the market is $(100 - t - c - k_2)k_2 = (70 - t)(20)$. Thus its profit from period $t_1 + 1$ through period t_2 is

$$(19 + 18 + \ldots + 1)(20) = 3800.$$

Hence firm 2's loss in period t_1 when both firms are active is (much) less than the sum of its profits in periods $t_1 + 1$ through t_2 when it alone is active.

227.1 Variant of ultimatum game with equity-conscious players

The game is defined as follows.

Players The two people.

Terminal histories The set of sequences (x, β_2, Z), where x is a number with $0 \leq x \leq c$ (the amount of money that person 1 offers to person 2), β_2 is 0 or 1 (the value of β_2 selected by chance), and Z is either Y ("yes, I accept") or N ("no, I reject").

Player function $P(\varnothing) = 1$, $P(x) = c$ for all x, and $P(x, \beta_2) = 2$ for all x and all β_2.

Chance probabilities For every history x, chance chooses 0 with probability p and 1 with probability $1 - p$.

Preferences Each person's preferences are represented by the expected value of a payoff equal to the amount of money she receives. For any terminal history (x, β_2, Y) person 1 receives $c - x$ and person 2 receives x; for any terminal history (x, β_2, N) each person receives 0.

Given the result from Exercise 183.4 stated in the question, an offer x of player 1 is accepted with probability either 0 or p if $x = 0$, is accepted with probability p if $0 < x < \frac{1}{3}$, is accepted with probability either p or 1 if $x = \frac{1}{3}$, and is accepted with probability 1 if $x > \frac{1}{3}$. By an argument like that for the original ultimatum game, in any equilibrium in which player 1 makes an offer of 0, player 2 certainly accepts the offer if $\beta_2 = 0$, and in any equilibrium in which player 1 makes an offer of $\frac{1}{3}$, player 2 certainly accepts the offer if $\beta_2 = 1$. Thus player 1's expected payoff to making the offer x is

$$\begin{cases} p(1-x) & \text{if } 0 \leq x < \frac{1}{3} \\ 1 - x & \text{if } \frac{1}{3} \leq x < 1. \end{cases}$$

The maximizer of this function is $x = \frac{1}{3}$ if $p < \frac{2}{3}$ and $x = 0$ if $p > \frac{2}{3}$; if $p = \frac{2}{3}$ then both offers are optimal. (If you do not see that the maximizer takes this form, plot the expected payoff as a function of x.)

We conclude that if $p \neq \frac{2}{3}$, the subgame perfect equilibria of the game are given as follows.

$p < \frac{2}{3}$ Player 1 offers $\frac{1}{3}$. After a history in which $\beta_2 = 0$, player 2 accepts an offer x with $x > 0$ and either accepts or rejects the offer 0. After a history in which $\beta_2 = 1$, player 2 accepts an offer x with $x \geq \frac{1}{3}$ and rejects an offer x with $x < \frac{1}{3}$.

$p > \frac{2}{3}$ Player 1 offers 0. After a history in which $\beta_2 = 0$, player 2 accepts all offers. After a history in which $\beta_2 = 1$, player 2 accepts an offer x with $x > \frac{1}{3}$, rejects an offer x with $x < \frac{1}{3}$, and either accepts or rejects the offer $\frac{1}{3}$.

If $p = \frac{2}{3}$, both these strategy pairs are subgame perfect equilibria.

We see that if $p > \frac{2}{3}$ then in a subgame perfect equilibrium player 1's offers are rejected by every player 2 with for whom $\beta_2 = 1$ (that is, with probability $1 - p$).

230.1 Nash equilibria when players may make mistakes

The players' best response functions are indicated in Figure 46.1. We see that the game has two Nash equilibria, (A, A, A) and (B, A, A).

	A	B
A	$1^*, 1^*, 1^*$	$0, 0, 1^*$
B	$1^*, 1^*, 1^*$	$1^*, 0, 1^*$

A

	A	B
A	$0, 1^*, 0$	$1^*, 0, 0$
B	$1^*, 1^*, 0$	$0, 0, 0$

B

Figure 46.1 The player's best response functions in the game in Exercise 230.1.

The action A is not weakly dominated for any player. For player 1, A is better than B if players 2 and 3 both choose B; for players 2 and 3, A is better than B for all actions of the other players.

If players 2 and 3 choose A in the modified game, player 1's expected payoffs to A and B are

$$A\text{: } (1 - p_2)(1 - p_3) + p_1 p_2 (1 - p_3) + p_1 (1 - p_2) p_3 + (1 - p_1) p_2 p_3$$

$$B\text{: } (1 - p_2)(1 - p_3) + (1 - p_1) p_2 (1 - p_3) + (1 - p_1)(1 - p_2) p_3 + p_1 p_2 p_3.$$

The difference between the expected payoff to B and the expected payoff to A is

$$(1 - 2p_1)[p_2 + p_3 - 3p_2 p_3].$$

If $0 < p_i < \frac{1}{2}$ for $i = 1, 2, 3$, this difference is positive, so that (A, A, A) is not a Nash equilibrium of the modified game.

233.1 Nash equilibria of the chain-store game

Any terminal history in which the event in each period is either *Out* or (In, A) is the outcome of a Nash equilibrium. In any period in which challenger chooses *Out*, the strategy of the chain-store specifies that it choose F in the event that the challenger chooses *In*.

8 Coalitional Games and the Core

245.1 Three-player majority game

Let (x_1, x_2, x_3) be an action of the grand coalition. Every coalition consisting of two players can obtain one unit of output, so for (x_1, x_2, x_3) to be in the core we need

$$\begin{aligned} x_1 + x_2 &\geq 1 \\ x_1 + x_3 &\geq 1 \\ x_2 + x_3 &\geq 1 \\ x_1 + x_2 + x_3 &= 1. \end{aligned}$$

Adding the first three conditions we conclude that

$$2x_1 + 2x_2 + 2x_3 \geq 3,$$

or $x_1 + x_2 + x_3 \geq \frac{3}{2}$, contradicting the last condition. Thus no action of the grand coalition satisfies all the conditions, so that the core of the game is empty.

248.1 Core of landowner–worker game

Let a_N be an action of the grand coalition in which the output received by each worker is at most $f(n) - f(n-1)$. No coalition consisting solely of workers can obtain any output, so no such coalition can improve upon a_N. Let S be a coalition of the landowner and $k-1$ workers. The total output received by the members of S in a_N is at least

$$f(n) - (n-k)(f(n) - f(n-1))$$

(because the total output is $f(n)$, and every *other* worker receives at most $f(n) - f(n-1)$). Now, the output that S can obtain is $f(k)$, so for S to improve upon a_N we need

$$f(k) > f(n) - (n-k)(f(n) - f(n-1)),$$

which contradicts the inequality given in the exercise.

249.1 Unionized workers in landowner–worker game

The following game models the situation.

Players The landowner and the workers.

Actions The set of actions of the grand coalition is the set of all allocations of the output $f(n)$. Every other coalition has a single action, which yields the output 0.

Preferences Each player's preferences are represented by the amount of output she obtains.

The core of this game consists of every allocation of the output $f(n)$ among the players. The grand coalition cannot improve upon any allocation x because for every other allocation x' there is at least one player whose payoff is lower in x' than it is in x. No other coalition can improve upon any allocation because no other coalition can obtain any output.

249.2 Landowner–worker game with increasing marginal products

We need to show that no coalition can improve upon the action a_N of the grand coalition in which every player receives the output $f(n)/n$. No coalition of workers can obtain any output, so we need to consider only coalitions containing the landowner. Consider a coalition consisting of the landowner and k workers, which can obtain $f(k+1)$ units of output by itself. Under a_N this coalition obtains the output $(k+1)f(n)/n$, and we have $f(k+1)/(k+1) < f(n)/n$ because $k < n$. Thus no coalition can improve upon a_N.

254.1 Range of prices in horse market

The equality of the number of owners who sell their horses and the number of nonowners who buy horses implies that the common trading price p^*

- is not less than σ_{k^*}, otherwise at most k^*-1 owners' valuations would be less than p^* and at least k^* nonowners' valuations would be greater than p^*, so that the number of buyers would exceed the number of sellers
- is not less than β_{k^*+1}, otherwise at most k^* owners' valuations would be less than p^* and at least k^*+1 nonowners' valuations would be greater than p^*, so that the number of buyers would exceed the number of sellers
- is not greater than β_{k^*}, otherwise at least k^* owners' valuations would be less than p^* and at most k^*-1 nonowners' valuations would be greater than p^*, so that the number of sellers would exceed the number of buyers
- is not greater than σ_{k^*+1}, otherwise at least k^*+1 owners' valuations would be less than p^* and at most k^* nonowners' valuations would be greater than p^*, so that the number of sellers would exceed the number of buyers.

That is, $p^* \geq \max\{\sigma_{k^*}, \beta_{k^*+1}\}$ and $p^* \leq \min\{\beta_{k^*}, \sigma_{k^*+1}\}$.

258.1 House assignment with identical preferences

Because the players rank the houses in the same way, we can refer to the "best house", the "second best house", and so on. In any assignment in the core, the player who owns the best house is assigned this house (because she has the option of keeping it). Among the remaining players, the one who owns the second best house must be assigned this house (again, because she has the option of keeping it). Continuing to argue in the same way, we see that there is a single assignment in the core, in which every player is assigned the house she owns initially.

261.1 Median voter theorem

Denote the median favorite position by m. If $x < m$ then every player whose favorite position is m or greater—a majority of the players—prefers m to x. Similarly, if $x > m$ then every player whose favorite position is m or less—a majority of the players—prefers m to x.

267.2 Empty core in roommate problem

Notice that ℓ is at the bottom of each of the other players' preferences. Suppose that she is matched with i. Then j and k are matched, and $\{i, k\}$ can improve upon the matching. Similarly, if ℓ is matched with j then $\{i, j\}$ can improve upon the matching, and if ℓ is matched with k then $\{j, k\}$ can improve upon the matching. Thus the core is empty (ℓ has to be matched with *someone*!).

9 Bayesian Games

276.1 Equilibria of a variant of *BoS* with imperfect information

If player 1 chooses S then type 1 of player 2 chooses S and type 2 chooses B. But if the two types of player 2 make these choices then player 1 is better off choosing B (which yields her an expected payoff of 1) than choosing S (which yields her an expected payoff of $\frac{1}{2}$). Thus there is no Nash equilibrium in which player 1 chooses S.

Now consider the mixed strategy Nash equilibria. If both types of player 2 use a pure strategy then player 1's two actions yield her different payoffs. Thus there is no equilibrium in which both types of player 2 use pure strategies and player 1 randomizes.

Now consider an equilibrium in which type 1 of player 2 randomizes. Denote by p the probability that player 1's mixed strategy assigns to B. In order for type 1 of player 2 to obtain the same expected payoff to B and S we need $p = \frac{2}{3}$. For this value of p the best action of type 2 of player 2 is S. Denote by q the probability that type 1 of player 2 assigns to B. Given these strategies for the two types of player 2, player 1's expected payoff if she chooses B is

$$\tfrac{1}{2} \cdot 2q = q$$

and her expected payoff if she chooses S is

$$\tfrac{1}{2} \cdot (1 - q) + \tfrac{1}{2} \cdot 1 = 1 - \tfrac{1}{2}q.$$

These expected payoffs are equal if and only if $q = \frac{2}{3}$. Thus the game has a mixed strategy equilibrium in which the mixed strategy of player 1 is $(\frac{2}{3}, \frac{1}{3})$, that of type 1 of player 2 is $(\frac{2}{3}, \frac{1}{3})$, and that of type 2 of player 2 is $(0, 1)$ (that is, type 2 of player 2 uses the pure strategy that assigns probability 1 to S).

Similarly the game has a mixed strategy equilibrium in which the strategy of player 1 is $(\frac{1}{3}, \frac{2}{3})$, that of type 1 of player 2 is $(0, 1)$, and that of type 2 of player 2 is $(\frac{2}{3}, \frac{1}{3})$.

For no mixed strategy of player 1 are both types of player 2 indifferent between their two actions, so there is no equilibrium in which both types randomize.

277.1 Expected payoffs in a variant of *BoS* with imperfect information

The expected payoffs are given in Figure 54.1.

	(B,B)	(B,S)	(S,B)	(S,S)
B	0	1	1	2
S	1	$\frac{1}{2}$	$\frac{1}{2}$	0

Type n_1 of player 1

	(B,B)	(B,S)	(S,B)	(S,S)
B	1	$\frac{2}{3}$	$\frac{1}{3}$	0
S	0	$\frac{2}{3}$	$\frac{4}{3}$	2

Type y_2 of player 2

	(B,B)	(B,S)	(S,B)	(S,S)
B	0	$\frac{1}{3}$	$\frac{2}{3}$	1
S	2	$\frac{4}{3}$	$\frac{2}{3}$	0

Type n_2 of player 2

Figure 54.1 The expected payoffs of type n_1 of player 1 and types y_2 and n_2 of player 2 in Example 276.2.

282.2 An exchange game

The following Bayesian game models the situation.

Players The two individuals.

States The set of all pairs (s_1, s_2), where s_i is the number on player i's ticket (an integer from 1 to m).

Actions The set of actions of each player is $\{Exchange, Don't\ exchange\}$.

Signals The signal function of each player i is defined by $\tau_i(s_1, s_2) = s_i$ (each player observes her own ticket, but not that of the other player)

Beliefs Type s_i of player i assigns the probability $\Pr_j(s_j)$ to the state (s_1, s_2), where j is the other player and $\Pr_j(s_j)$ is the probability with which player j receives a ticket with the prize s_j on it.

Payoffs Player i's Bernoulli payoff function is given by $u_i((X, Y), \omega) = \omega_j$ if $X = Y = Exchange$ and $u_i((X, Y), \omega) = \omega_i$ otherwise.

Let M_i be the highest type of player i that chooses *Exchange*. If $M_i > 1$ then type 1 of player j optimally chooses *Exchange*: by exchanging her ticket, she cannot obtain a smaller prize, and may receive a bigger one. Thus if $M_i \geq M_j$ and $M_i > 1$, type M_i of player i optimally chooses *Don't exchange*, because the expected value of the prizes of the types of player j that choose *Exchange* is less than M_i. Thus in any possible Nash equilibrium $M_i = M_j = 1$: the only prizes that may be exchanged are the smallest.

287.1 Cournot's duopoly game with imperfect information

We have

$$b_1(q_L, q_H) = \begin{cases} \frac{1}{2}(\alpha - c - (\theta q_L + (1-\theta)q_H)) & \text{if } \theta q_L + (1-\theta)q_H \le \alpha - c \\ 0 & \text{otherwise.} \end{cases}$$

The best response function of each type of player 2 is similar:

$$b_I(q_1) = \begin{cases} \frac{1}{2}(\alpha - c_I - q_1) & \text{if } q_1 \le \alpha - c_I \\ 0 & \text{otherwise} \end{cases}$$

for $I = L, H$.

The three equations that define a Nash equilibrium are

$$q_1^* = b_1(q_L^*, q_H^*),\ q_L^* = b_L(q_1^*), \text{ and } q_H^* = b_H(q_1^*).$$

Solving these equations under the assumption that they have a solution in which all three outputs are positive, we obtain

$$\begin{aligned} q_1^* &= \tfrac{1}{3}(\alpha - 2c + \theta c_L + (1-\theta)c_H) \\ q_L^* &= \tfrac{1}{3}(\alpha - 2c_L + c) - \tfrac{1}{6}(1-\theta)(c_H - c_L) \\ q_H^* &= \tfrac{1}{3}(\alpha - 2c_H + c) + \tfrac{1}{6}\theta(c_H - c_L) \end{aligned}$$

If both firms know that the unit costs of the two firms are c_1 and c_2 then in a Nash equilibrium the output of firm i is $\frac{1}{3}(\alpha - 2c_i + c_j)$ (see Exercise 58.1). In the case of imperfect information considered here, firm 2's output is less than $\frac{1}{3}(\alpha - 2c_L + c)$ if its cost is c_L and is greater than $\frac{1}{3}(\alpha - 2c_H + c)$ if its cost is c_H. Intuitively, the reason is as follows. If firm 1 knew that firm 2's cost were high then it would produce a relatively large output; if it knew this cost were low then it would produce a relatively small output. Given that it does not know whether the cost is high or low it produces a moderate output, less than it would if it knew firm 2's cost were high. Thus if firm 2's cost is in fact high, firm 2 benefits from firm 1's lack of knowledge and optimally produces more than it would if firm 1 knew its cost.

288.1 Cournot's duopoly game with imperfect information

The best response $b_0(q_L, q_H)$ of type 0 of firm 1 is the solution of

$$\max_{q_0} [\theta(P(q_0 + q_L) - c)q_0 + (1-\theta)(P(q_0 + q_H) - c)q_0].$$

The best response $b_\ell(q_L, q_H)$ of type ℓ of firm 1 is the solution of

$$\max_{q_\ell} (P(q_\ell + q_L) - c)q_\ell$$

and the best response $b_h(q_L, q_H)$ of type h of firm 1 is the solution of

$$\max_{q_h}(P(q_h + q_H) - c)q_h.$$

The best response $b_L(q_0, q_\ell, q_h)$ of type L of firm 2 is the solution of

$$\max_{q_L}[(1-\pi)(P(q_0 + q_L) - c_L)q_L + \pi(P(q_\ell + q_L) - c_L)q_L]$$

and the best response $b_H(q_0, q_\ell, q_h)$ of type H of firm 2 is the solution of

$$\max_{q_H}[(1-\pi)(P(q_0 + q_H) - c_H)q_H + \pi(P(q_h + q_H) - c_H)q_H].$$

A Nash equilibrium is a profile $(q_0^*, q_\ell^*, q_h^*, q_L^*, q_H^*)$ for which q_0^*, q_ℓ^*, and q_h^* are best responses to q_L^* and q_H^*, and q_L^* and q_H^* are best responses to q_0^*, q_ℓ^*, and q_h^*. When $P(Q) = \alpha - Q$ for $Q \leq \alpha$ and $P(Q) = 0$ for $Q > \alpha$ we find, after some exciting algebra, that

$$\begin{aligned}
q_0^* &= \frac{1}{3}\left(\alpha - 2c + c_H - \theta\left(c_H - c_L\right)\right) \\
q_\ell^* &= \frac{1}{3}\left(\alpha - 2c + c_L + \frac{(1-\theta)(1-\pi)(c_H - c_L)}{4-\pi}\right) \\
q_h^* &= \frac{1}{3}\left(\alpha - 2c + c_H - \frac{\theta(1-\pi)(c_H - c_L)}{4-\pi}\right) \\
q_L^* &= \frac{1}{3}\left(\alpha - 2c_L + c - \frac{2(1-\theta)(1-\pi)(c_H - c_L)}{4-\pi}\right) \\
q_H^* &= \frac{1}{3}\left(\alpha - 2c_H + c + \frac{2\theta(1-\pi)(c_H - c_L)}{4-\pi}\right).
\end{aligned}$$

When $\pi = 0$ we have

$$\begin{aligned}
q_0^* &= \frac{1}{3}\left(\alpha - 2c + c_H - \theta\left(c_H - c_L\right)\right) \\
q_\ell^* &= \frac{1}{3}\left(\alpha - 2c + c_L + \frac{(1-\theta)(c_H - c_L)}{4}\right) \\
q_h^* &= \frac{1}{3}\left(\alpha - 2c + c_H - \frac{\theta(c_H - c_L)}{4}\right) \\
q_L^* &= \frac{1}{3}\left(\alpha - 2c_L + c - \frac{(1-\theta)(c_H - c_L)}{2}\right) \\
q_H^* &= \frac{1}{3}\left(\alpha - 2c_H + c + \frac{\theta(c_H - c_L)}{2}\right),
\end{aligned}$$

so that q_0^* is equal to the equilibrium output of firm 1 in Exercise 287.1, and q_L^* and q_H^* are the same as the equilibrium outputs of the two types of firm 2 in that exercise.

When $\pi = 1$ we have

$$\begin{aligned} q_0^* &= \frac{1}{3}\left(\alpha - 2c + c_H - \theta\left(c_H - c_L\right)\right) \\ q_\ell^* &= \frac{1}{3}\left(\alpha - 2c + c_L\right) \\ q_h^* &= \frac{1}{3}\left(\alpha - 2c + c_H\right) \\ q_L^* &= \frac{1}{3}\left(\alpha - 2c_L + c\right) \\ q_H^* &= \frac{1}{3}\left(\alpha - 2c_H + c\right), \end{aligned}$$

so that q_ℓ^* and q_L^* are the same as the equilibrium outputs when there is perfect information and the costs are c and c_L (see Exercise 58.1), and q_h^* and q_H^* are the same as the equilibrium outputs when there is perfect information and the costs are c and c_H.

Now, for an arbitrary value of π we have

$$\begin{aligned} q_L^* &= \frac{1}{3}\left(\alpha - 2c_L + c - \frac{2(1-\theta)(1-\pi)(c_H - c_L)}{4-\pi}\right) \\ q_H^* &= \frac{1}{3}\left(\alpha - 2c_H + c + \frac{2\theta(1-\pi)(c_H - c_L)}{4-\pi}\right). \end{aligned}$$

To show that for $0 < \pi < 1$ the values of these variables lie between their values when $\pi = 0$ and when $\pi = 1$, we need to show that

$$0 \le \frac{2(1-\theta)(1-\pi)(c_H - c_L)}{4-\pi} \le \frac{(1-\theta)(c_L - c_H)}{2}$$

and

$$0 \le \frac{2\theta(1-\pi)(c_H - c_L)}{4-\pi} \le \frac{\theta(c_L - c_H)}{2}.$$

These inequalities follow from $c_H \ge c_L$, $\theta \ge 0$, and $0 \le \pi \le 1$.

290.1 Nash equilibria of game of contributing to a public good

Any type v_j of any player j with $v_j < c$ obtains a negative payoff if she contributes and 0 if she does not. Thus she optimally does not contribute.

Any type $v_i \ge c$ of player i obtains the payoff $v_i - c \ge 0$ if she contributes, and the payoff 0 if she does not, so she optimally contributes.

Any type $v_j \ge c$ of any player $j \ne i$ obtains the payoff $v_j - c$ if she contributes, and the payoff $(1 - F(c))v_j$ if she does not. (If she does not contribute, the probability that player i does so is $1 - F(c)$, the probability that player i's valuation is at least c.) Thus she optimally does not contribute if $(1 - F(c))v_j \ge v_j - c$, or $F(c) \le c/v_j$. This condition must hold for all types of every player $j \ne i$, so we need $F(c) \le c/\overline{v}$ for the strategy profile to be a Nash equilibrium.

294.1 Weak domination in second-price sealed-bid action

Fix player i, and choose a bid for every type of every other player. Player i, who does not know the other players' types, is uncertain of the highest bid of the other players. Denote by $\overline{b}$ this highest bid. Consider a bid b_i of type v_i of player i for which $b_i < v_i$. The dependence of the payoff of type v_i of player i on $\overline{b}$ is shown in Figure 58.1.

		Highest of other players' bids			
		$\overline{b} < b_i$	$b_i = \overline{b}$ (m-way tie)	$b_i < \overline{b} < v_i$	$\overline{b} \geq v_i$
i's bid	$b_i < v_i$	$v_i - \overline{b}$	$(v_i - \overline{b})/m$	0	0
	v_i	$v_i - \overline{b}$	$v_i - \overline{b}$	$v_i - \overline{b}$	0

Figure 58.1 Player i's payoffs to her bids $b_i < v_i$ and v_i in a second-price sealed-bid auction as a function of the highest of the other player's bids, denoted $\overline{b}$.

Player i's expected payoffs to the bids b_i and v_i are weighted averages of the payoffs in the columns; each value of $\overline{b}$ gets the same weight when calculating the expected payoff to b_i as it does when calculating the expected payoff to v_i. The payoffs in the two rows are the same except when $b_i \leq \overline{b} < v_i$, in which case v_i yields a payoff higher than does b_i. Thus the expected payoff to v_i is at least as high as the expected payoff to b_i, and is greater than the expected payoff to b_i unless the other players' bids lead this range of values of $\overline{b}$ to get probability 0.

Now consider a bid b_i of type v_i of player i for which $b_i > v_i$. The dependence of the payoff of type v_i of player i on $\overline{b}$ is shown in Figure 58.2.

		Highest of other players' bids			
		$\overline{b} \leq v_i$	$v_i < \overline{b} < b_i$	$b_i = \overline{b}$ (m-way tie)	$\overline{b} > b_i$
i's bid	v_i	$v_i - \overline{b}$	0	0	0
	$b_i > v_i$	$v_i - \overline{b}$	$v_i - \overline{b}$	$(v_i - \overline{b})/m$	0

Figure 58.2 Player i's payoffs to her bids v_i and $b_i > v_i$ in a second-price sealed-bid auction as a function of the highest of the other player's bids, denoted $\overline{b}$.

As before, player i's expected payoffs to the bids b_i and v_i are weighted averages of the payoffs in the columns; each value of $\overline{b}$ gets the same weight when calculating the expected payoff to v_i as it does when calculating the expected payoff to b_i. The payoffs in the two rows are the same except when $v_i < \overline{b} \leq b_i$, in which case v_i yields a payoff higher than does b_i. (Note that $v_i - \overline{b} < 0$ for $\overline{b}$ in this range.) Thus the expected payoff to v_i is at least as high as the expected payoff to b_i, and is greater than the expected payoff to b_i unless the other players' bids lead this range of values of $\overline{b}$ to get probability 0.

We conclude that for type v_i of player i, every bid $b_i \neq v_i$ is weakly dominated by the bid v_i.

299.1 Asymmetric Nash equilibria of second-price sealed-bid common value auctions

Suppose that each type t_2 of player 2 bids $(1 + 1/\lambda)t_2$ and that type t_1 of player 1 bids b_1. Then by the calculations in the text, with $\alpha = 1$ and $\gamma = 1/\lambda$,

- a bid of b_1 by player 1 wins with probability $b_1/(1 + 1/\lambda)$
- the expected value of player 2's bid, *given that it is less than* b_1, is $\frac{1}{2}b_1$
- the expected value of signals that yield a bid of less than b_1 is $\frac{1}{2}b_1/(1 + 1/\lambda)$ (because of the uniformity of the distribution of t_2).

Thus player 1's expected payoff if she bids b_1 is

$$(t_1 + \tfrac{1}{2}b_1/(1 + 1/\lambda) - \tfrac{1}{2}b_1) \cdot \frac{b_1}{1 + 1/\lambda},$$

or

$$\frac{\lambda}{2(1+\lambda)^2} \cdot (2(1+\lambda)t_1 - b_1)b_1.$$

This function is maximized at $b_1 = (1 + \lambda)t_1$. That is, if each type t_2 of player 2 bids $(1 + 1/\lambda)t_2$, any type t_1 of player 1 optimally bids $(1 + \lambda)t_1$. Symmetrically, if each type t_1 of player 1 bids $(1 + \lambda)t_1$, any type t_2 of player 2 optimally bids $(1 + 1/\lambda)t_2$. Hence the game has the claimed Nash equilibrium.

299.2 First-price sealed-bid auction with common valuations

Suppose that each type t_2 of player 2 bids $\frac{1}{2}(\alpha + \gamma)t_2$ and type t_1 of player 1 bids b_1. To determine the expected payoff of type t_1 of player 1, we need to find the probability with which she wins, and the expected value of player 2's signal if player 1 wins. (The price she pays is her bid, b_1.)

Probability of player 1's winning: Given that player 2's bidding function is $\frac{1}{2}(\alpha + \gamma)t_2$, player 1's bid of b_1 wins only if $b_1 \geq \frac{1}{2}(\alpha + \gamma)t_2$, or if $t_2 \leq 2b_1/(\alpha + \gamma)$. Now, t_2 is distributed uniformly from 0 to 1, so the probability that it is at most $2b_1/(\alpha + \gamma)$ is $2b_1/(\alpha + \gamma)$. Thus a bid of b_1 by player 1 wins with probability $2b_1/(\alpha + \gamma)$.

Expected value of player 2's signal if player 1 wins: Player 2's bid, given her signal t_2, is $\frac{1}{2}(\alpha + \gamma)t_2$, so that the expected value of signals that yield a bid of less than b_1 is $b_1/(\alpha + \gamma)$ (because of the uniformity of the distribution of t_2).

Thus player 1's expected payoff if she bids b_1 is

$$2(\alpha t_1 + \gamma b_1/(\alpha + \gamma) - b_1) \cdot \frac{b_1}{\alpha + \gamma},$$

or

$$\frac{2\alpha}{(\alpha+\gamma)^2}((\alpha + \gamma)t_1 - b_1)b_1.$$

This function is maximized at $b_1 = \frac{1}{2}(\alpha+\gamma)t_1$. That is, if each type t_2 of player 2 bids $\frac{1}{2}(\alpha+\gamma)t_2$, any type t_1 of player 1 optimally bids $\frac{1}{2}(\alpha+\gamma)t_1$. Hence, as claimed, the game has a Nash equilibrium in which each type t_i of player i bids $\frac{1}{2}(\alpha+\gamma)t_i$.

309.2 Properties of the bidding function in a first-price sealed-bid auction

We have

$$\begin{aligned}
\beta^{*\prime}(v) &= 1 - \frac{(F(v))^{n-1}(F(v))^{n-1} - (n-1)(F(v))^{n-2}F'(v)\int_{\underline{v}}^{v}(F(x))^{n-1}\,dx}{(F(v))^{2n-2}} \\
&= 1 - \frac{(F(v))^{n} - (n-1)F'(v)\int_{\underline{v}}^{v}(F(x))^{n-1}\,dx}{(F(v))^{n}} \\
&= \frac{(n-1)F'(v)\int_{\underline{v}}^{v}(F(x))^{n-1}\,dx}{(F(v))^{n}} \\
&> 0 \quad \text{if } v > \underline{v}
\end{aligned}$$

because $F'(v) > 0$ (F is increasing). (The first line uses the quotient rule for derivatives and the fact that the derivative of $\int^{v} f(x)dx$ with respect to v is $f(v)$ for any function f.)

If $v > \underline{v}$ then the integral in (309.1) is positive, so that $\beta^*(v) < v$. If $v = \underline{v}$ then both the numerator and denominator of the quotient in (309.1) are zero, so we may use L'Hôpital's rule to find the value of the quotient as $v \to \underline{v}$. Taking the derivatives of the numerator and denominator we obtain

$$\frac{(F(v))^{n-1}}{(n-1)(F(v))^{n-2}F'(v)} = \frac{F(v)}{(n-1)F'(v)},$$

the numerator of which is zero and the denominator of which is positive. Thus the quotient in (309.1) is zero, and hence $\beta^*(\underline{v}) = \underline{v}$.

309.3 Example of Nash equilibrium in a first-price auction

From (309.1) we have

$$\begin{aligned}
\beta^*(v) &= v - \frac{\int_0^v x^{n-1}\,dx}{v^{n-1}} \\
&= v - \frac{\int_0^v x^{n-1}\,dx}{v^{n-1}} \\
&= v - v/n = (n-1)v/n.
\end{aligned}$$

10 Extensive Games with Imperfect Information

316.1 Variant of card game

An extensive game that models the game is shown in Figure 61.1.

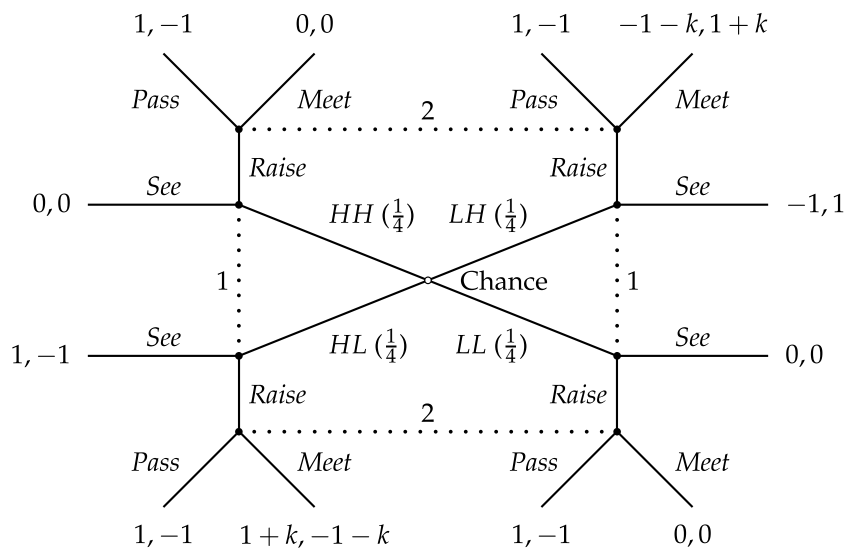

Figure 61.1 An extensive game that models the situation in Exercise 316.1.

318.2 Strategies in variants of card game and entry game

Card game: Each player has two information sets, and has two actions at each information set. Thus each player has four strategies: *SS*, *SR*, *RS*, and *RR* for player 1 (where *S* stands for *See* and *R* for *Raise*, the first letter of each strategy is player 1's action if her card is *High*, and the second letter if her action is her card is *Low*), and *PP*, *PM*, *MP*, and *MM* for player 2 (where *P* stands for *Pass* and *M* for *Meet*).

Entry game: The challenger has a single information set (the empty history) and has three actions after this history, so it has three strategies—*Ready*, *Unready*, and *Out*. The incumbent also has a single information set, at which two actions are available, so it has two strategies—*Acquiesce* and *Fight*.

331.2 Weak sequential equilibrium and Nash equilibrium in subgames

Consider the assessment in which the Challenger's strategy is (Out, R), the Incumbent's strategy is F, and the Incumbent's belief assigns probability 1 to the history (In, U) at her information set. Each player's strategy is sequentially rational. The Incumbent's belief satisfies the condition of weak consistency because her information set is not reached when the Challenger follows her strategy. Thus the assessment is a weak sequential equilibrium.

The players' actions in the subgame following the history *In* do not constitute a Nash equilibrium of the subgame because the Incumbent's action F is not optimal when the Challenger chooses R. (The Incumbent's action F *is* optimal given her belief that the history is (In, U), as it is in the weak sequential equilibrium. In a Nash equilibrium she acts as if she has a belief that coincides with the Challenger's action in the subgame.)

340.1 Pooling equilibria of game in which expenditure signals quality

We know that in the second period the high-quality firm charges the price H and the low-quality firm charges any nonnegative price, and the consumer buys the good from a high-quality firm, does not buy the good from a low-quality firm that charges a positive price, and may or may not buy from a low-quality firm that charges a price of 0.

Consider an assessment in which each type of firm chooses (p^*, E^*) in the first period, the consumer believes the firm is high-quality with probability π if it observes (p^*, E^*) and low quality if it observes any other (price, expenditure) pair, and buys the good if and only if it observes (p^*, E^*).

The payoff of a high-quality firm under this assessment is $p^* + H - E^* - 2c_H$, that of a low-quality firm is $p^* - E^*$, and that of the consumer is $\pi(H - p^*) + (1 - \pi)(-p^*) = \pi H - p^*$.

This assessment is consistent—the only first-period action of the firm observed in equilibrium is (p^*, E^*), and after observing this pair the consumer believes, correctly, that the firm is high-quality with probability π.

Under what conditions is the assessment sequentially rational?

Firm If the firm chooses a (price, expenditure) pair different from (p^*, E^*) then the consumer does not buy the good, and the firm's profit is 0. Thus for the assessment to be an equilibrium we need $p^* + H - E^* - 2c_H \geq 0$ (for the high-quality firm) and $p^* - E^* \geq 0$ (for the low-quality firm).

Consumer If the consumer does not buy the good after observing (p^*, E^*) then its payoff is 0, so for the assessment to be an equilibrium we need $\pi H - p^* \geq 0$.

In summary, the assessment is a weak sequential equilibrium if and only if

$$\max\{E^*, E^* - H + 2c_H\} \leq p^* \leq \pi H.$$

346.1 Comparing the receiver's expected payoff in two equilibria

The receiver's payoff as a function of the state t in each equilibrium is shown in Figure 63.1. The area above the black curve is smaller than the area above the gray curve: if you shift the black curve $\frac{1}{2}t_1$ to the left and move the section from 0 to $\frac{1}{2}t_1$ to the interval from $1 - \frac{1}{2}t_1$ to 1 then the area above the black curve is a subset of the area above the gray curve.

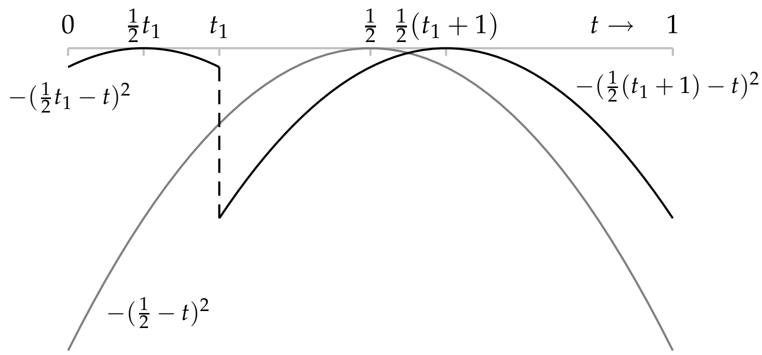

Figure 63.1 The gray curve gives the receiver's payoff in each state in the equilibrium in which no information is transferred. The black curve gives her payoff in each state in the two-report equilibrium.

350.1 Variant of model with piecewise linear payoff functions

The equilibria of the variant are exactly the same as the equilibria of the original model.

11 Strictly Competitive Games and Maxminimization

363.1 Maxminimizers in a bargaining game

If a player demands any amount x up to \$5 then her payoff is x regardless of the other player's action. If she demands \$6 then she may get as little as \$5 (if the other player demands \$5 or \$6). If she demands $x \geq \$7$ then she may get as little as $\$(11 - x)$ (if the other player demands $x - 1$). For each amount that a player demands, the smallest amount that you may get is given in Figure 65.1. We see that each player's maxminimizing pure strategies are \$5 and \$6 (for both of which the worst possible outcome is that the player receives \$5).

Amount demanded	0	1	2	3	4	5	6	7	8	9	10
Smallest amount obtained	0	1	2	3	4	5	5	4	3	2	1

Figure 65.1 The lowest payoffs that a player receives in the game in Exercise 38.2 for each of her possible actions, as the other player's action varies.

363.3 Finding a maxminimizer

The analog of Figure 364.1 in the text is Figure 65.2. From this figure we see that the maxminimizer for player 2 is the strategy that assigns probability $\frac{2}{5}$ to L. Player 2's maxminimized payoff is $-\frac{1}{5}$.

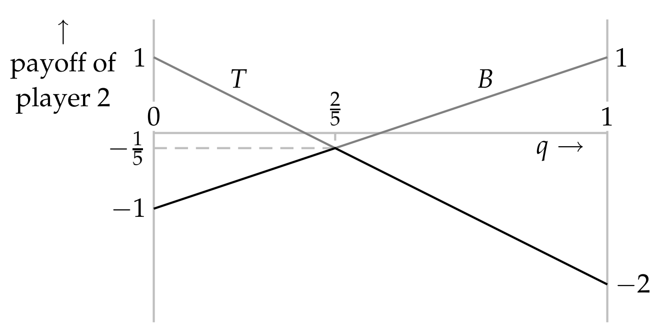

Figure 65.2 The expected payoff of player 2 in the game in Figure 363.1 for each of player 1's actions, as a function of the probability q that player 2 assigns to L.

366.2 Determining strictly competitiveness

Game in Exercise 365.1: Strictly competitive in pure strategies (because player 1's ranking of the four outcomes is the reverse of player 2's ranking). Not strictly competitive in mixed strategies (there exist no values of π and $\theta > 0$ such that $-u_1(a) = \pi + \theta u_2(a)$ for every outcome a; or, alternatively, player 1 is indifferent between (B, L) and the lottery that yields (T, L) with probability $\frac{1}{2}$ and (T, R) with probability $\frac{1}{2}$, whereas player 2 is not indifferent between these two outcomes).

Game in Figure 367.1: Strictly competitive both in pure and in mixed strategies. (Player 2's preferences are represented by the expected value of the Bernoulli payoff function $-u_1$ because $-u_1(a) = -\frac{1}{2} + \frac{1}{2}u_2(a)$ for every pure outcome a.)

370.2 Maxminimizing in BoS

Player 1's maxminimizer is $(\frac{1}{3}, \frac{2}{3})$ while player 2's is $(\frac{2}{3}, \frac{1}{3})$. Clearly neither pure equilibrium strategy of either player guarantees her equilibrium payoff. In the mixed strategy equilibrium, player 1's expected payoff is $\frac{2}{3}$. But if, for example, player 2 choose S instead of her equilibrium strategy, then player 1's expected payoff is $\frac{1}{3}$. Similarly for player 2.

372.2 Equilibrium in strictly competitive game

The claim is false. In the strictly competitive game in Figure 66.1 the action pair (T, L) is a Nash equilibrium, so that player 1's unique equilibrium payoff in the game is 0. But (B, R), which also yields player 1 a payoff of 0, is not a Nash equilibrium.

	L	R
T	$0, 0$	$1, -1$
B	$-1, 1$	$0, \ 0$

Figure 66.1 The game in Exercise 372.2.

372.4 O'Neill's game

a. Denote the probability with which player 1 chooses each of her actions 1, 2, and 3, by p, and the probability with which player 2 chooses each of these actions by q. Then all four of player 1's actions yield the same expected payoff if and only if $4q - 1 = 1 - 6q$, or $q = \frac{1}{5}$, and similarly all four of player 2's actions yield the same expected payoff if and only if $p = \frac{1}{5}$. Thus $((\frac{1}{5}, \frac{1}{5}, \frac{1}{5}, \frac{2}{5}), (\frac{1}{5}, \frac{1}{5}, \frac{1}{5}, \frac{2}{5}))$ is a Nash equilibrium of the game. The players' payoffs in this equilibrium are $(-\frac{1}{5}, \frac{1}{5})$.

b. Let (p_1, p_2, p_3, p_4) be an equilibrium strategy of player 1. In order that it guarantee her the payoff of $-\frac{1}{5}$, we need

$$\begin{aligned} -p_1 + p_2 + p_3 - p_4 &\geq -\tfrac{1}{5} \\ p_1 - p_2 + p_3 - p_4 &\geq -\tfrac{1}{5} \\ p_1 + p_2 - p_3 - p_4 &\geq -\tfrac{1}{5} \\ -p_1 - p_2 - p_3 + p_4 &\geq -\tfrac{1}{5}. \end{aligned}$$

Adding these four inequalities, we deduce that $p_4 \leq \frac{2}{5}$. Adding each pair of the first three inequalities, we deduce that $p_1 \leq \frac{1}{5}$, $p_2 \leq \frac{1}{5}$, and $p_3 \leq \frac{1}{5}$. We have $p_1 + p_2 + p_3 + p_4 = 1$, so we deduce that $(p_1, p_2, p_3, p_4) = (\frac{1}{5}, \frac{1}{5}, \frac{1}{5}, \frac{2}{5})$. A similar analysis of the conditions for player 2's strategy to guarantee her the payoff of $\frac{1}{5}$ leads to the conclusion that $(q_1, q_2, q_3, q_4) = (\frac{1}{5}, \frac{1}{5}, \frac{1}{5}, \frac{2}{5})$.

12 Rationalizability

379.2 Best responses to beliefs

Consider a two-player game in which player 1's payoffs are given in Figure 69.1. The action B of player 1 is a best response to the belief that assigns probability $\frac{1}{2}$ to both L and R, but is not a best response to any belief that assigns probability 1 to either action.

	L	R
T	3	0
M	0	3
B	2	2

Figure 69.1 The action B is a best response to a belief that assigns probability $\frac{1}{2}$ to L and to R, but is not a best response to any belief that assigns probability 1 to either L or R.

384.1 Mixed strategy equilibria of game in Figure 384.1

The game has no equilibrium in which player 2 assigns positive probability only to L and C, because if she does so then only M and B are possible best responses for player 1, but if player 1 assigns positive probability only to these actions then L is not optimal for player 2.

Similarly, the game has no equilibrium in which player 2 assigns positive probability only to C and R, because if she does so then only T and M are possible best responses for player 1, but if player 1 assigns positive probability only to these actions then R is not optimal for player 2.

Now assume that player 2 assigns positive probability only to L and R. There are no probabilities for L and R under which player 1 is indifferent between all three of her actions, so player 1 must assign positive probability to at most two actions. If these two actions are T and M then player 2 prefers L to R, while if the two actions are M and B then player 2 prefers R to L. The only possibility is thus that the two actions are T and B. In this case we need player 2 to assign probability $\frac{1}{2}$ to L and R (in order that player 1 be indifferent between T and B); but then M is better for player 1. Thus there is no equilibrium in which player 2 assigns positive probability only to L and R.

Finally, if player 2 assigns positive probability to all three of her actions then player 1's mixed strategy must be such that each of these three actions yields the same payoff. A calculation shows that there is no mixed strategy of player 1 with this property.

We conclude that the game has no mixed strategy equilibrium in which either player assigns positive probability to more than one action.

387.2 Finding rationalizable actions

I claim that the action R of player 2 is strictly dominated. Consider a mixed strategy of player 2 that assigns probability p to L and probability $1 - p$ to C. Such a mixed strategy strictly dominates R if $p + 4(1 - p) > 3$ and $8p + 2(1 - p) > 3$, or if $\frac{1}{6} < p < \frac{1}{3}$. Now eliminate R from the game. In the reduced game, B is dominated by T. In the game obtained by eliminating B, L is dominated by C. Thus the only rationalizable action of player 1 is T and the only rationalizable action of player 2 is C.

387.5 Hotelling's model of electoral competition

The positions 0 and ℓ are strictly dominated by the position m:

- if her opponent chooses m, a player who chooses m ties whereas a player who chooses 0 loses

- if her opponent chooses 0 or ℓ, a player who chooses m wins whereas a player who chooses 0 or ℓ either loses or ties

- if her opponent chooses any other position, a player who chooses m wins whereas a player who chooses 0 or ℓ loses.

In the game obtained by eliminating the two positions 0 and ℓ, the positions 1 and $\ell - 1$ are similarly strictly dominated. Continuing in the same way, we are left with the position m.

388.2 Cournot's duopoly game

From Figure 58.1 we see that firm 1's payoff to any output greater than $\frac{1}{2}(\alpha - c)$ is less than its payoff to the output $\frac{1}{2}(\alpha - c)$ for any output q_2 of firm 2. Thus any output greater than $\frac{1}{2}(\alpha - c)$ is strictly dominated by the output $\frac{1}{2}(\alpha - c)$ for firm 1; the same argument applies to firm 2.

Now eliminate all outputs greater than $\frac{1}{2}(\alpha - c)$ for each firm. The maximizer of firm 1's payoff function for $q_2 = \frac{1}{2}(\alpha - c)$ is $\frac{1}{4}(\alpha - c)$, so from Figure 58.1 we see that firm 1's payoff to any output less than $\frac{1}{4}(\alpha - c)$ is less than its payoff to the output $\frac{1}{4}(\alpha - c)$ for any output $q_2 \leq \frac{1}{2}(\alpha - c)$ of firm 2. Thus any output less than $\frac{1}{4}(\alpha - c)$ is strictly dominated by the output $\frac{1}{4}(\alpha - c)$ for firm 1; the same argument applies to firm 2.

Now eliminate all outputs less than $\frac{1}{4}(\alpha - c)$ for each firm. Then by another similar argument, any output greater than $\frac{3}{8}(\alpha - c)$ is strictly dominated by $\frac{3}{8}(\alpha - c)$. Continuing in this way, we see from Figure 59.1 that in a finite number of rounds (given the finite number of possible outputs for each firm) we reach the Nash equilibrium output $\frac{1}{3}(\alpha - c)$.

391.1 Example of dominance-solvable game

The Nash equilibria of the game are (T, L), any $((0, 0, 1), (0, q, 1 - q))$ with $0 \leq q \leq 1$, and any $((0, p, 1 - p), (0, 0, 1))$ with $0 \leq p \leq 1$.

The game is dominance solvable, because T and L are the only weakly dominated actions, and when they are eliminated the only weakly dominated actions are M and C, leaving (B, R), with payoffs $(0, 0)$.

If T is eliminated, then L and C, no remaining action is weakly dominated; (M, R) and (B, R) both remain.

391.2 Dividing money

In the first round every action $a_i \leq 5$ of each player i is weakly dominated by 6. No other action is weakly dominated, because 100 is a strict best response to 0 and every other action $a_i \geq 6$ is a strict best response to $a_i + 1$. In the second round, 10 is weakly dominated by 6 for each player, and each other remaining action a_i of player i is a strict best response to $a_1 + 1$, so no other action is weakly dominated. Similarly, in the third round, 9 is weakly dominated by 6, and no other action is weakly dominated. In the fourth and fifth rounds 8 and 7 are eliminated, leaving the single action pair $(6, 6)$, with payoffs $(5, 5)$.

392.2 Strictly competitive extensive games with perfect information

Every finite extensive game with perfect information has a (pure strategy) subgame perfect equilibrium (Proposition 173.1). This equilibrium is a pure strategy Nash equilibrium of the strategic form of the game. Because the game has only two possible outcomes, one of the players prefers the Nash equilibrium outcome to the other possible outcome. By Proposition 368.1, this player's equilibrium strategy guarantees her equilibrium payoff, so this strategy weakly dominates all her nonequilibrium strategies. After all dominated strategies are eliminated, every remaining pair of strategies generates the same outcome.

13 Evolutionary Equilibrium

400.1 Evolutionary stability and weak domination

The ESS a^* does not necessarily weakly dominate every other action in the game. For example, in the game in Figure 395.1 of the text, X is an ESS but does not weakly dominate Y.

No action can weakly dominate an ESS. To see why, let a^* be an ESS and let b be another action. Because a^* is an ESS, (a^*, a^*) is a Nash equilibrium, so that $u(b, a^*) \leq u(a^*, a^*)$. Now, if $u(b, a^*) < u(a^*, a^*)$, certainly b does not weakly dominate a^*, so suppose that $u(b, a^*) = u(a^*, a^*)$. Then by the second condition for an ESS we have $u(b, b) < u(a^*, b)$. We conclude that b does not weakly dominate a^*.

405.1 *Hawk–Dove–Retaliator*

First suppose that $v \geq c$. In this case the game has two pure symmetric Nash equilibria, (A, A) and (R, R). However, A is not an ESS, because R is a best response to A and $u(R, R) > u(A, R)$. The action pair (R, R) is a strict equilibrium, so R is an ESS. Now consider the possibility that the game has a mixed strategy equilibrium (α, α). If α assigns positive probability to either P or R (or both) then R yields a payoff higher than does P, so only A and R may be assigned positive probability in a mixed strategy equilibrium. But if a strategy α assigns positive probability to A and R and probability 0 to P, then R yields a payoff higher than does A against an opponent who uses α. Thus the game has no symmetric mixed strategy equilibrium in this case.

Now suppose that $v < c$. Then the only symmetric pure strategy equilibrium is (R, R). This equilibrium is strict, so that R is an ESS. Now consider the possibility that the game has a mixed strategy equilibrium (α, α). If α assigns probability 0 to A then R yields a payoff higher than does P against an opponent who uses α; if α assigns probability 0 to P then R yields a payoff higher than does A against an opponent who uses α. Thus in any mixed strategy equilibrium (α, α), the strategy α must assign positive probability to both A and P. If α assigns probability 0 to R then we need $\alpha = (v/c, 1 - v/c)$ (the calculation is the same as for *Hawk–Dove*). Because R yields a lower payoff against this strategy than do A and P, and the strategy is an ESS in *Hawk–Dove*, it is an ESS in the present game. The remaining possibility is that the game has a mixed strategy equilibrium (α, α) in which α assigns positive probability to all three actions. If so, then the expected payoff to this strategy is less than $\frac{1}{2}v$, because the pure strategy P yields an expected payoff less than $\frac{1}{2}v$ against any such strategy. But then $U(R,R) = \frac{1}{2}v > U(\alpha, R)$, violating the second condition in the definition of an ESS.

In summary:

- If $v \geq c$ then R is the unique ESS of the game.
- If $v < c$ then both R and the mixed strategy that assigns probability v/c to A and $1 - v/c$ to P are ESSs.

405.3 Bargaining

The game is given in Figure 27.1.

The pure strategy of demanding 10 is not an ESS because 2 is a best response to 10 and $u(2,2) > u(10,2)$.

Now let α be the mixed strategy that assigns probability $\frac{2}{5}$ to 2 and $\frac{3}{5}$ to 8. Each player's payoff at the strategy pair (α, α) is $\frac{16}{5}$. Thus the only actions a that are best responses to α are 2 and 8, so that the only mixed strategies that are best responses to α assign positive probability only to the actions 2 and 8. Let β be the mixed strategy that assigns probability p to 2 and probability $1-p$ to 8. We have

$$U(\beta, \beta) = 5p(2-p)$$

and

$$U(\alpha, \beta) = 6p + \tfrac{4}{5}.$$

We find that $U(\alpha, \beta) - U(\beta, \beta) = 5(p - \frac{2}{5})^2$, which is positive if $p \neq \frac{2}{5}$. Hence α is an ESS.

Finally let α be the mixed strategy that assigns probability $\frac{4}{5}$ to 4 and $\frac{1}{5}$ to 6. Each player's payoff at the strategy pair (α, α) is $\frac{24}{5}$. Thus the only actions a that are best responses to α are 4 and 6, so that the only mixed strategies that are best responses assign positive probability only to the actions 4 and 6. Let β be the mixed strategy that assigns probability p to 4 and probability $1-p$ to 6. We have

$$U(\beta, \beta) = 5p(2-p)$$

and

$$U(\alpha^*, \beta) = 2p + \tfrac{16}{5}.$$

We find that $U(\alpha, \beta) - U(\beta, \beta) = 5(p - \frac{4}{5})^2$, which is positive if $p \neq \frac{4}{5}$. Hence α^* is an ESS.

408.1 Equilibria of C and of G

First suppose that (α_1, α_2) is a mixed strategy Nash equilibrium of C. Then for all mixed strategies β_1 of player 1 and all mixed strategies β_2 of player 2 we have

$$U_1(\alpha_1, \alpha_2) \geq U_1(\beta_1, \alpha_2) \text{ and } U_2(\alpha_1, \alpha_2) \geq U_2(\alpha_1, \beta_2).$$

Thus

$$\begin{aligned} u((\alpha_1,\alpha_2),(\alpha_1,\alpha_2)) &= \tfrac{1}{2}U_1(\alpha_1,\alpha_2) + \tfrac{1}{2}U_2(\alpha_1,\alpha_2) \\ &\geq \tfrac{1}{2}U_1(\beta_1,\alpha_2) + \tfrac{1}{2}U_2(\alpha_1,\beta_2) \\ &= u((\beta_1,\beta_2),(\alpha_1,\alpha_2)), \end{aligned}$$

so that $((\alpha_1,\alpha_2),(\alpha_1,\alpha_2))$ is a Nash equilibrium of G. If (α_1,α_2) is a strict Nash equilibrium of C then the inequalities are strict, and $((\alpha_1,\alpha_2),(\alpha_1,\alpha_2))$ is a strict Nash equilibrium of G.

Now assume that $((\alpha_1,\alpha_2),(\alpha_1,\alpha_2))$ is a Nash equilibrium of G. Then

$$u((\alpha_1,\alpha_2),(\alpha_1,\alpha_2)) \geq u((\beta_1,\beta_2),(\alpha_1,\alpha_2)),$$

or

$$\tfrac{1}{2}U_1(\alpha_1,\alpha_2) + \tfrac{1}{2}U_2(\alpha_1,\alpha_2) \geq \tfrac{1}{2}U_1(\beta_1,\alpha_2) + \tfrac{1}{2}U_2(\alpha_1,\beta_2),$$

for all conditional strategies (β_1,β_2). Taking $\beta_2 = \alpha_2$ we see that α_1 is a best response to α_2 in C, and taking $\beta_1 = \alpha_1$ we see that α_2 is a best response to α_1 in C. Thus (α_1,α_2) is a Nash equilibrium of G.

414.1 A coordination game between siblings

The game with payoff function v is shown in Figure 75.1. If $x < 2$ then (Y,Y) is a strict Nash equilibrium of the games, so Y is an evolutionarily stable action in the game between siblings. If $x > 2$ then the only Nash equilibrium of the game is (X,X), and this equilibrium is strict. Thus the range of values of x for which the only evolutionarily stable action is X is $x > 2$.

	X	Y
X	x, x	$\frac{1}{2}x, \frac{1}{2}$
Y	$\frac{1}{2}, \frac{1}{2}x$	$1, 1$

v

Figure 75.1 The game with payoff function v derived from the game in Exercise 414.1.

414.2 Assortative mating

Under assortative mating, all siblings take the same action, so the analysis is the same as that for asexual reproduction. (A difficulty with the assumption of assortative mating is that a rare mutant will have to go to great lengths to find a mate that is also a mutant.)

416.1 Darwin's theory of the sex ratio

A normal organism produces pn male offspring and $(1-p)n$ female offspring (ignoring the small probability that the partner of a normal organism is a mutant). Thus it has $pn \cdot ((1-p)/p)n + (1-p)n \cdot n = 2(1-p)n^2$ grandchildren.

A mutant has $\frac{1}{2}n$ male offspring and $\frac{1}{2}n$ female offspring, and hence $\frac{1}{2}n \cdot ((1-p)/p)n + \frac{1}{2}n \cdot n = \frac{1}{2}n^2/p$ grandchildren.

Thus the difference between the number of grandchildren produced by mutant and normal organisms is

$$\tfrac{1}{2}n^2/p - 2(1-p)n^2 = n^2\left(\frac{1}{2p}\right)(1-2p)^2,$$

which is positive if $p \neq \frac{1}{2}$. (The point is that if $p > \frac{1}{2}$ then the fraction of a mutant's offspring that are males is higher than the fraction of a normal organism's offspring that are males, and males each bear more offspring than females. Similarly, if $p < \frac{1}{2}$ then the fraction of a mutant's offspring that are females is higher than the fraction of a normal organism's offspring that are females, and females each bear more offspring than males.)

Thus any mutant with $p \neq \frac{1}{2}$ invades the population; only $p = \frac{1}{2}$ is evolutionarily stable.

14 Repeated Games: The *Prisoner's Dilemma*

423.1 Equivalence of payoff functions

Suppose that a person's preferences are represented by the discounted sum of payoffs with payoff function u and discount factor δ. Then if the two sequences of outcomes $(x^1, x^2, \ldots)$ and $(y^1, y^2, \ldots)$ are indifferent, we have

$$\sum_{t=0}^{\infty} \delta^{t-1} u(x^t) = \sum_{t=0}^{\infty} \delta^{t-1} u(y^t).$$

Now let $v(x) = \alpha + \beta u(x)$ for all x, with $\beta > 0$. Then

$$\sum_{t=0}^{\infty} \delta^{t-1} v(x^t) = \sum_{t=0}^{\infty} \delta^{t-1} [\alpha + \beta u(x^t)] = \sum_{t=0}^{\infty} \delta^{t-1} \alpha + \beta \sum_{t=0}^{\infty} \delta^{t-1} u(x^t)$$

and similarly

$$\sum_{t=0}^{\infty} \delta^{t-1} v(y^t) = \sum_{t=0}^{\infty} \delta^{t-1} [\alpha + \beta u(y^t)] = \sum_{t=0}^{\infty} \delta^{t-1} \alpha + \beta \sum_{t=0}^{\infty} \delta^{t-1} u(y^t),$$

so that

$$\sum_{t=0}^{\infty} \delta^{t-1} v(x^t) = \sum_{t=0}^{\infty} \delta^{t-1} v(y^t).$$

Thus the person's preferences are represented also by the discounted sum of payoffs with payoff function v and discount factor δ.

426.1 Subgame perfect equilibrium of finitely repeated *Prisoner's Dilemma*

Use backward induction. In the last period, the action C is strictly dominated for each player, so each player chooses D, regardless of history. Now consider period $T-1$. Each player's action in this period affects only the outcome in this period—it has no effect on the outcome in period T, which is (D, D). Thus in choosing her action in period $T-1$, a player considers only her payoff in that period. As in period T, her action D strictly dominates her action C, so that in any subgame perfect equilibrium she chooses D. A similar argument applies to all previous periods, leading to the conclusion that in every subgame perfect equilibrium each player chooses D in every period, regardless of history.

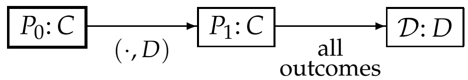

Figure 78.1 The strategy in Exercise 428.1a.

428.1 Strategies in an infinitely repeated *Prisoner's Dilemma*

a. The strategy is shown in Figure 78.1.

b. The strategy is shown in Figure 78.2.

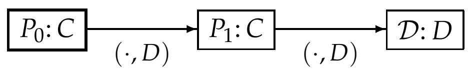

Figure 78.2 The strategy in Exercise 428.1b.

c. The strategy is shown in Figure 78.3.

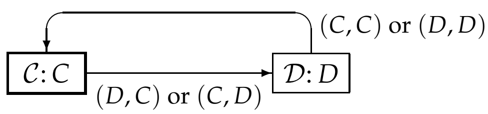

Figure 78.3 The strategy in Exercise 428.1c.

439.1 Finitely repeated *Prisoner's Dilemma* with switching cost

a. Consider deviations by player 1, given that player 2 adheres to her strategy, in the subgames following histories that end in each of the four outcomes of the game.

(C, C): If player 1 adheres to her strategy, her payoff is 3 in every period. If she deviates in the first period of the subgame, but otherwise follows her strategy, her payoff is $4 - \epsilon$ in the first period of the subgame, and 2 in every subsequent period. Given $\epsilon > 1$, player 1's deviation is not profitable, even if it occurs in the last period of the game.

(D, C) or (D, D): If player 1 adheres to her strategy, her payoff is 2 in every period. If she deviates in the first period of the subgame, but otherwise follows her strategy, her payoff is $-\epsilon$ in the first period of the subgame, $2 - \epsilon$ in the next period, and 2 subsequently. Thus adhering to her strategy is optimal for player 1.

(C, D): If player 1 adheres to her strategy, her payoff is $2 - \epsilon$ in the first period of the subgame, and 2 subsequently. If she deviates in the first period of the subgame, but otherwise follows her strategy, her payoff is 0 in the first period of the subgame, $2 - \epsilon$ in the next period, and 2 subsequently. Given $\epsilon < 2$, player 1's deviation is not optimal even if it occurs in the last period of the game.

b. Given $\epsilon > 2$, a player does not gain from deviating from (C, C) in the next-to-last or last periods, even if she is not punished, and does not optimally punish such a deviation by her opponent. Consider the strategy that chooses C at the start of the game and after any history that ends with (C, C), chooses D after any other history that has length at most $T - 2$, and chooses the action it chose in period $T - 1$ after any history of length $T - 1$ (where T is the length of the game). I claim that the strategy pair in which both players use this strategy is a subgame perfect equilibrium. Consider deviations by player 1, given that player 2 adheres to her strategy, in the subgames following the various possible histories.

History ending in (C, C), length $\leq T - 3$: If player 1 adheres to her strategy, her payoff is 3 in every period of the subgame. If she deviates in the first period of the subgame, but otherwise follows her strategy, her payoff is $4 - \epsilon$ in the first period of the subgame, and 2 in every subsequent period (her opponent switches to D). Given $\epsilon > 1$, player 1's deviation is not profitable.

History ending in (C, C), length $\geq T - 2$: If player 1 adheres to her strategy, her payoff is 3 in each period of the subgame. If she deviates to D in the first period of the subgame, her payoff is $4 - \epsilon$ in that period, and 4 subsequently (her deviation is not punished). The length of the subgame is at most 2, so given $\epsilon > 2$, her deviation is not profitable.

History ending in (D, C) or (D, D): If player 1 adheres to her strategy, her payoff is 2 in every period. If she deviates in the first period of the subgame, but otherwise follows her strategy, her payoff is $-\epsilon$ in the first period of the subgame, $2 - \epsilon$ in the next period, and 2 subsequently. Thus adhering to her strategy is optimal for player 1.

History ending in (C, D), length $\leq T - 2$: If player 1 adheres to her strategy, her payoff is $2 - \epsilon$ in the first period of the subgame (she switches to D), and 2 subsequently. If she deviates in the first period of the subgame, but otherwise follows her strategy, her payoff is 0 in the first period of the subgame, $2 - \epsilon$ in the next period, and 2 subsequently.

History ending in (C, D), length $T - 1$: If player 1 adheres to her strategy, her payoff is 0 in period T (the outcome is (C, D)). If she deviates to D, her payoff is $2 - \epsilon$ in period T. Given $\epsilon > 2$, adhering to her strategy is thus optimal.

442.1 Deviations from grim trigger strategy

- If player 1 adheres to the strategy, she subsequently chooses D (because player 2 chose D in the first period). Player 2 chooses C in the first period of the subgame (player 1 chose C in the first period of the game), and then chooses D (because player 1 chooses D in the first period of the subgame). Thus the sequence of outcomes in the subgame is $((D, C), (D, D), (D, D), \ldots)$, yielding player 1 a discounted average payoff in the subgame of

$$(1-\delta)(3+\delta+\delta^2+\delta^3+\cdots) = (1-\delta)\left(3+\frac{\delta}{1-\delta}\right) = 3-2\delta.$$

- If player 1 refrains from punishing player 2 for her lapse, and simply chooses C in every subsequent period, then the outcome in period 2 and subsequently is (C, C), so that the sequence of outcomes in the subgame yields player 1 a discounted average payoff of 2.

If $\delta > \frac{1}{2}$ then $2 > 3 - 2\delta$, so that player 1 prefers to ignore player 2's deviation rather than to adhere to her strategy and punish player 2 by choosing D. (Note that the theory does not consider the possibility that player 1 takes player 2's play of D as a signal that she is using a strategy different from the grim trigger strategy.)

443.2 Different punishment lengths in subgame perfect equilibrium

Yes, an infinitely repeated *Prisoner's Dilemma* has such subgame perfect equilibria. As for the modified grim trigger strategy, each player's strategy has to switch to D not only if the other player chooses D but also if the player herself chooses D. The only subtlety is that the number of periods for which a player chooses D after a history in which not all the outcomes were (C, C) must depend on the identity of the player who first deviated. If, for example, player 1 punishes for two periods while player 2 punishes for three periods, then the outcome (C, D) induces player 1 to choose D for two periods (to punish player 2 for her deviation) while the outcome (D, C) induces her to choose D for three periods (while she is being punished by player 2). The strategy of each player in this case is shown in Figure 81.1. Viewed as a strategy of player 1, the top part of the figure entails punishment of player 2 and the bottom part entails player 1's reaction to her own deviation. Similarly, viewed as a strategy of player 2, the bottom part of the figure entails punishment of player 1 and the top part entails player 2's reaction to her own deviation.

To find the values of δ for which the strategy pair in which each player uses the strategy in Figure 81.1 is a subgame perfect equilibrium, consider the result of each player's deviating at the start of a subgame.

First consider player 1. If she deviates when both players are in state P_0, she induces the outcome (D, C) followed by three periods of (D, D), and then (C, C) subsequently. This outcome path is worse for her than (C, C) in every period if

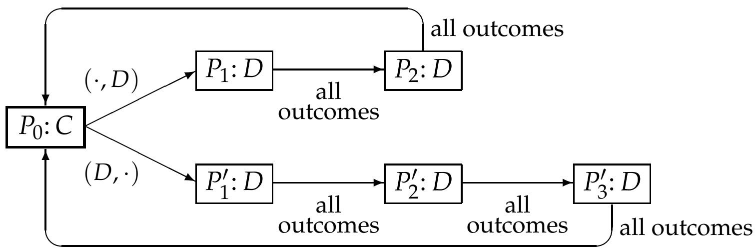

Figure 81.1 A strategy in an infinitely repeated *Prisoner's Dilemma* that punishes deviations for two periods and reacts to punishment by choosing D for three periods.

and only if $\delta^3 - 2\delta + 1 \leq 0$, or if and only if δ is at least around 0.62 (as we found in Section 14.7.2). If she deviates when both players are in one of the other states then she is worse off in the period of her deviation and her deviation does not affect the subsequent outcomes. Thus player 1 cannot profitably deviate in the first period of any subgame if δ is at least around 0.62.

The same argument applies to player 2, except that a deviation when both players are in state P_0 induces (C, D) followed by three, rather than two periods of (D, D). This outcome path is worse for player 2 than (C, C) in every period if and only if $\delta^4 - 2\delta + 1 \leq 0$, or if and only if δ is at least around 0.55 (as we found in Section 14.7.2).

We conclude that the strategy pair in which each player uses the strategy in Figure 81.1 is a subgame perfect equilibrium if and only if $\delta^3 - 2\delta + 1 \leq 0$, or if and only if δ is at least around 0.62.

445.1 *Tit-for-tat* as a subgame perfect equilibrium

Suppose that player 2 adheres to *tit-for-tat*. Consider player 1's behavior in subgames following histories that end in each of the following outcomes.

(C, C) If player 1 adheres to *tit-for-tat* the outcome is (C, C) in every period, so that her discounted average payoff in the subgame is x. If she chooses D in the first period of the subgame, then adheres to *tit-for-tat*, the outcome alternates between (D, C) and (C, D), and her discounted average payoff is $y/(1+\delta)$. Thus we need $x \geq y/(1+\delta)$, or $\delta \geq (y-x)/x$, for a one-period deviation from *tit-for-tat* not to be profitable for player 1.

(C, D) If player 1 adheres to *tit-for-tat* the outcome alternates between (D, C) and (C, D), so that her discounted average payoff is $y/(1+\delta)$. If she deviates to C in the first period of the subgame, then adheres to *tit-for-tat*, the outcome is (C, C) in every period, and her discounted average payoff is x. Thus we need $y/(1+\delta) \geq x$, or $\delta \leq (y-x)/x$, for a one-period deviation from *tit-for-tat* not to be profitable for player 1.

(D, C) If player 1 adheres to *tit-for-tat* the outcome alternates between (C, D) and (D, C), so that her discounted average payoff is $\delta y/(1+\delta)$. If she deviates to D in the first period of the subgame, then adheres to *tit-for-tat*, the outcome is (D, D) in every period, and her discounted average payoff is 1. Thus we need $\delta y/(1+\delta) \geq 1$, or $\delta \geq 1/(y-1)$, for a one-period deviation from *tit-for-tat* not to be profitable for player 1.

(D, D) If player 1 adheres to *tit-for-tat* the outcome is (D, D) in every period, so that her discounted average payoff is 1. If she deviates to C in the first period of the subgame, then adheres to *tit-for-tat*, the outcome alternates between (C, D) and (D, C), and her discounted average payoff is $\delta y/(1+\delta)$. Thus we need $1 \geq \delta y/(1+\delta)$, or $\delta \leq 1/(y-1)$, for a one-period deviation from *tit-for-tat* not to be profitable for player 1.

The same arguments apply to deviations by player 2, so we conclude that (*tit-for-tat*, *tit-for-tat*) is a subgame perfect equilibrium if and only if $\delta = (y-x)/x$ and $\delta = 1/(y-1)$, or $y - x = 1$ and $\delta = 1/x$.

15 Repeated Games: General Results

454.3 Repeated Bertrand duopoly

a. Suppose that firm i uses the strategy s_i. If the other firm, j, uses s_j, then its discounted average payoff is

$$(1-\delta)\left(\tfrac{1}{2}\pi(p^m) + \tfrac{1}{2}\delta\pi(p^m) + \cdots\right) = \tfrac{1}{2}\pi(p^m).$$

If, on the other hand, firm j deviates to a price p then the closer this price is to p^m, the higher is j's profit, because the punishment does not depend on p. Thus by choosing p close enough to p^m the firm can obtain a profit as close as it wishes to $\pi(p^m)$ in the period of its deviation. Its profit during its punishment in the following k periods is zero. Once its punishment is complete, it can either revert to p^m or deviate once again. If it can profit from deviating initially then it can profit by deviating once its punishment is complete, so its maximal profit from deviating is

$$(1-\delta)\left(\pi(p^m) + \delta^{k+1}\pi(p^m) + \delta^{2k+2}\pi(p^m) + \cdots\right) = \frac{(1-\delta)\pi(p^m)}{1-\delta^{k+1}}.$$

Thus for (s_1, s_2) to be a Nash equilibrium we need

$$\frac{1-\delta}{1-\delta^{k+1}} \leq \tfrac{1}{2},$$

or

$$\delta^{k+1} - 2\delta + 1 \leq 0.$$

(This condition is the same as the one we found for a pair of k-period punishment strategies to be a Nash equilibrium in the *Prisoner's Dilemma* (Section 14.7.2).)

b. Suppose that firm i uses the strategy s_i. If the other firm does so then its discounted average payoff is $\frac{1}{2}\pi(p^m)$, as in part *a*. If the other firm deviates to some price p with $c < p < p^m$ in the first period, and maintains this price subsequently, then it obtains $\pi(p)$ in the first period and shares $\pi(p)$ in each subsequent period, so that its discounted average payoff is

$$(1-\delta)\left(\pi(p) + \tfrac{1}{2}\delta\pi(p) + \tfrac{1}{2}\delta^2\pi(p) + \cdots\right) = \tfrac{1}{2}(2-\delta)\pi(p).$$

If p is close to p^m then $\pi(p)$ is close to $\pi(p^m)$ (because π is continuous). In fact, for any $\delta < 1$ we have $2-\delta > 1$, so that we can find $p < p^m$ such that $(2-\delta)\pi(p) > \pi(p^m)$. Hence the strategy pair is not a Nash equilibrium of the infinitely repeated game for any value of δ.

459.2 Detection lags

a. The best deviations involve prices slightly less than p^*. Such a deviation by firm i yields a discounted average payoff close to

$$(1-\delta)\left(\pi(p^*)+\delta\pi(p^*)+\cdots+\delta^{k_i-1}\pi(p^*)\right)=(1-\delta^{k_i})\pi(p^*),$$

whereas compliance with the strategy yields the discounted average payoff $\frac{1}{2}\pi(p^*)$. Thus the strategy pair is a subgame perfect equilibrium for any value of p^* if $\delta^{k_1} \geq \frac{1}{2}$ and $\delta^{k_2} \geq \frac{1}{2}$, and is not a subgame perfect equilibrium for any value of p^* if $\delta^{k_1} < \frac{1}{2}$ or $\delta^{k_2} < \frac{1}{2}$. That is, the most profitable price for which the strategy pair is a subgame perfect equilibrium is p^m if $\delta^{k_1} \geq \frac{1}{2}$ and $\delta^{k_2} \geq \frac{1}{2}$ and is c if $\delta^{k_1} < \frac{1}{2}$ or $\delta^{k_2} < \frac{1}{2}$.

b. Denote by k_i^* the critical value of k_i found in part *a*. (That is, $\delta^{k_i^*} \geq \frac{1}{2}$ and $\delta^{k_i^*+1} < \frac{1}{2}$.)

If $k_i > k_i^*$ then no change in k_j affects the outcome of the price-setting subgame, so j's best action at the start of the game is θ, in which case i's best action is the same. Thus in one subgame perfect equilibrium both firms choose θ at the start of the game, and c regardless of history in the rest of the game.

If $k_i \leq k_i^*$ then j's best action is k_j^* if the cost of choosing k_j^* is at most $\frac{1}{2}\pi(p^m)$. Thus if the cost of choosing k_i^* is at most $\frac{1}{2}\pi(p^m)$ for each firm then the game has another subgame perfect equilibrium, in which each firm i chooses k_i^* at the start of the game and the strategy s_i in the price-setting subgame.

A promise by firm i to beat another firm's price is an inducement for consumers to inform firm i of deviations by other firms, and thus reduce its detection time. To this extent, such a promise tends to promote collusion.

16 Bargaining

468.1 Two-period bargaining with constant cost of delay

In the second period, player 1 accepts any proposal that gives a positive amount of the pie. Thus in any subgame perfect equilibrium player 2 proposes $(0, 1)$ in period 2, which player 1 accepts, obtaining the payoff $-c_1$.

Now consider the first period. Given the second period outcome of any subgame perfect equilibrium, player 2 accepts any proposal that gives her more than $1 - c_2$ and rejects any proposal that gives her less than $1 - c_2$. Thus in any subgame perfect equilibrium player 1 proposes $(c_2, 1 - c_2)$, which player 2 accepts.

In summary, the game has a unique subgame perfect equilibrium, in which

- player 1 proposes $(c_2, 1 - c_2)$ in period 1, and accepts all proposals in period 2
- player 2 accepts a proposal in period 1 if and only if it gives her at least $1 - c_2$, and proposes $(0, 1)$ in period 2 after any history.

The outcome of the equilibrium is that the proposal $(c_2, 1 - c_2)$ is made by player 1 and immediately accepted by player 2.

468.2 Three-period bargaining with constant cost of delay

The subgame following a rejection by player 2 in period 1 is a two-period game in which player 2 makes the first proposal. Thus by the result of Exercise 468.1, the subgame has a unique subgame perfect equilibrium, in which player 2 proposes $(1 - c_1, c_1)$, which player 1 immediately accepts.

Now consider the first period.

- If $c_1 \geq c_2$, player 2 rejects any offer of less than $c_1 - c_2$ (which she obtains if she rejects an offer), and accepts any offer of more than $c_1 - c_2$. Thus in an equilibrium player 1 offers her $c_1 - c_2$, which she accepts.
- If $c_1 < c_2$, player 2 accepts all offers, so that player 1 proposes $(1, 0)$, which player 2 accepts.

In summary, the game has a unique subgame perfect equilibrium, in which

- player 1 proposes $(1 - (c_1 - c_2), c_1 - c_2)$ if $c_1 \geq c_2$ and $(1, 0)$ otherwise in period 1, accepts any proposal that gives her at least $1 - c_1$ in period 2, and proposes $(1, 0)$ in period 3

- player 2 accepts any proposal that gives her at least $c_1 - c_2$ if $c_1 \geq c_2$ and accepts all proposals otherwise in period 1, proposes $(1 - c_1, c_1)$ in period 2, and accepts all proposals in period 3.

17 Appendix: Mathematics

497.1 Maximizer of quadratic function

We can write the function as $-x(x-\alpha)$. Thus $r_1 = 0$ and $r_2 = \alpha$, and hence the maximizer is $\alpha/2$.

499.3 Sums of sequences

In the first case set $r = \delta^2$ to transform the sum into $1 + r + r^2 + \cdots$, which is equal to $1/(1-r) = 1/(1-\delta^2)$.

In the second case split the sum into $(1 + \delta^2 + \delta^4 + \cdots) + (2\delta + 2\delta^3 + 2\delta^5 + \cdots)$; the first part is equal to $1/(1-\delta^2)$ and the second part is equal to $2\delta(1 + \delta^2 + \delta^4 + \cdots)$, or $2\delta/(1-\delta^2)$. Thus the complete sum is

$$\frac{1+2\delta}{1-\delta^2}.$$

504.2 Bayes' law

Your posterior probability of carrying X given that you test positive is

$$\frac{\Pr(\text{positive test}|X)\Pr(X)}{\Pr(\text{positive test}|X)\Pr(X) + \Pr(\text{positive test}|\neg X)\Pr(\neg X)}$$

where $\neg X$ means "not X". This probability is equal to $0.9p/(0.9p + 0.2(1-p)) = 0.9p/(0.2 + 0.7p)$, which is increasing in p (i.e. a smaller value of p gives a smaller value of the probability). If $p = 0.001$ then the probability is approximately 0.004. (That is, if 1 in 1,000 people carry the gene then if you test positive on a test that is 90% accurate for people who carry the gene and 80% accurate for people who do not carry the gene, then you should assign probability 0.004 to your carrying the gene.) If the test is 99% accurate in both cases then the posterior probability is $(0.99 \cdot 0.001)/[0.99 \cdot 0.001 + 0.01 \cdot 0.999] \approx 0.09$.

References

The page numbers on which the references are cited are given in brackets after each item.

Nagel, Rosemarie (1995), "Unraveling in guessing games: an experimental study", *American Economic Review* **85**, 1313–1326. [4]

Ochs, Jack (1995), "Coordination problems", in *Handbook of experimental economics* (John H. Kagel and Alvin E. Roth, eds.), 195–251. Princeton: Princeton University Press. [4]

Van Huyck, John B., Raymond C. Battalio, and Richard O. Beil (1990), "Tacit coordination games, strategic uncertainty, and coordination failure", *American Economic Review* **80**, 234–248. [4]