Chapter 8: Dominant Strategies

Elements of Decision: Lecture Notes of Intermediate Microeconomics

Charles Z. Zheng
Tepper School of Business, Carnegie Mellon University

Last update: April 28, 2020

The theme of this chapter, going beyond zero-sum games, is to consider those non-zero sum games where our decision-theoretic technique remains applicable quite directly. In such games, a player can select his optimal strategy regardless of which strategy the other players may choose.

1 Prisoners' dilemma

Each cell in the payoff matrix of a nonzero-sum game is represented by a pair of payoffs, one for each player. The player choosing rows is usually called player 1, and the one choosing columns, player 2. In the following payoff matrix, for example, if player 1 chooses D_1 and player 2 C_2, then player 1 gets payoff 2 and player 2, -3. Obviously, this is not zero sum.

	C_2	D_2
C_1	1, 1	-3, 2
D_1	2, -3	0, 0

This game is the famous Prisoners' Dilemma, with C_i interpreted as player i's cooperating with, and D_i his defecting against, the other player. The game provides a profound explanation for the tragic end of humanity (as well as complicated instructions to possibly dodge the doom). For now, however, we simply use it to introduce the concept of *strictly dominated strategies*.

A strategy s_i for player i is *strictly dominated* by another strategy s'_i iff s'_i gives the player strictly larger expected payoff than s_i does no matter which strategy the other player is to choose. In the prisoner's dilemma, for example, C_1 is strictly dominated by D_1: If player 2 chooses C_2, C_1 yields payoff 1 while D_1 2; if player 2 chooses D_2, C_1 yields -3 while D_1 zero. Thus, player 1 would choose D_1. Likewise, C_2 is strictly dominated by D_2, hence player 2 would chooses D_2. Consequently, despite the possibility of a "win-win" outcome of $(1, 1)$ had they played (C_1, C_2), the two players end with playing (D_1, D_2) thereby getting $(0, 0)$. Hence we obtain (D_1, D_2) as the *dominant strategy equilibrium* of the game.

2 Weakly dominated strategies

Let us look at another example:

	L	R
X	1, 1	0, 2
Y	2, 3	0, 0

Here, for player 1, strategy X is worse than strategy Y when player 2 chooses L, and is as well as Y when player 2 chooses R. Thus X is *weakly* dominated by Y in the sense that X is never better than Y and is sometimes worse than Y. Thus, if he is rational, player 1 would not choose X as he would not like to take the chance that the other player never chooses the strategy (L) that renders the weakly dominated strategy (X) worse. Consequently, once strategy X is eliminated, player 2's decision is straightforward: L it is, as L gives her 3 while R only zero. Thus we have arrived at a reasonable prediction of the game: (Y, L), with payoffs being 2 to player 1 and 3 to player 2. This is another example of dominant strategy equilibrium, albeit based on a weaker notion of dominance than that in Prisoners' Dilemma.

3 Sequential-move games and backward induction

The above technique of elimination by dominance can sometimes be applied to games with sequential moves. To illustrate, consider the following graph: In this tree-like graph, the point at the

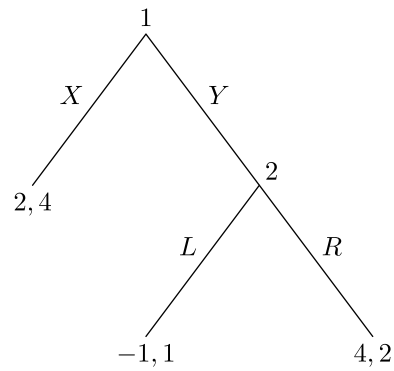

Figure 1: A sequential-move game

top is called *initial node*. It is labeled 1 here, meaning that player 1 gets to move first. The two branches from the initial node, labeled X and Y here, represent the two actions available to player 1 at the initial node. If player 1 chooses X then the game ends, with payoff vector $(2, 4)$, i.e., player 1 getting 2 and player 2 (without moving at all) getting 4. If player 1 chooses Y, then the branch leads to the next node, labeled 2, meaning that it is player 2 to move. The two branches from that node mean that player 2 chooses between L and R, and the outcomes are respectively specified by the payoff vectors $(-1, 1)$ and $(4, 2)$, with the first coordinate specifying the payoff for player 1, and the second coordinate the payoff for player 2. Any node at the bottom of the tree, labeled by a payoff vector (e.g., $(-1, 1)$ at the end of the path Y-L), is called *terminal node*, signifying an outcome of the game. The sub-tree starting from the node "2" is called *endgame*, as any branch from the node "2" leads directly to a terminal node.

We can solve such games by first turning them into the matrix format, called *normal form*:

	L	R
X	2, 4	2, 4
Y	−1, 1	4, 2

You can think of the normal form of a sequential move game as a scenario where each player plots out his entire game plan before the sequential move game starts and then each programs a proxy to carry out the plan during the operation of the game. Then a sequential-move game can be viewed as a simultaneous-move game that we have been considering in previous sections.

Given the above payoff matrix, one readily sees that, for player 2, strategy L is weakly dominated by R (verify it yourself). With L eliminated, player 1's decision between X and Y boils down to comparing between payoff 2 versus payoff 4, hence player 1 would choose Y. Thus the dominant strategy equilibrium is (Y, R), giving 4 to player 1 and 2 to player 2.

Translated back to our sequential-move game, the above solution means that the first mover, player 1, goes down the path Y thereby letting player 2 choose, and that player 2 chooses R thereby leading to the outcome $(4, 2)$.

This technique of elimination by dominance turns out to be essentially the same as an intuitive procedure called *backward induction*: Start with the endgame and locate the best action from the viewpoint of the player who is to move there, then replace that node by the terminal node resulting from that action, and then go further up to the next end game until we reach the initial node. In the above example, we start with the endgame:

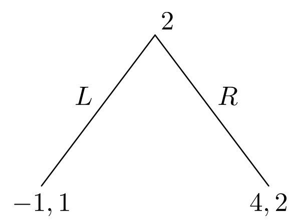

Figure 2: The endgame (player 2's decision)

Obviously player 2 would choose R in this endgame, hence this endgame is essentially the same as the outcome $(4, 2)$. Thus replace the entire endgame by the outcome $(4, 2)$ thereby shorten the original game in Figure 1 into:

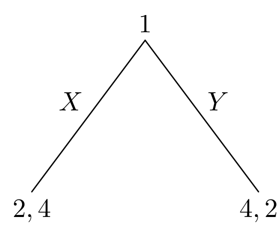

Figure 3: Player 1's decision after taking into account player 2's optimal choice

Clearly player 1 would choose Y. Thus we reach the solution (Y, R), exactly as in the previous technique of elimination by dominance.

4 Two subtle points on sequential-move games

4.1 Incredible threats

Note that, in the game of Figurefig-dominant-tree, player 2 would like the outcome to be $(2, 4)$. Should he have a chance to communicate with player 1 before the game starts, player 2 might try talking player 1 into playing X thereby getting his desired outcome $(2, 4)$: "Do play X, or I will play L to retaliate, giving you a payoff only -1 rather than 2." Should player 1 heed his threat? Our previous reasoning says No. The threat of retaliating with L is not credible: once player 1 has chosen Y, player 2's decision is the endgame in Figure 2, in which player 2's optimal choice is uniquely R, not L. What is crucial here is the sequential-move nature of the game: once the first mover has made the move, the second mover (player 2) has to take that as given to make his optimal decision. Bygone is bygone.

From the perspective of the normal form, saying that player 2's threat of choosing L is not credible is the same as saying that the strategy L is weakly dominated. Imagine that before the sequential game starts, player 2 says to player 1: "I am now programming this robot to carry out my contingency plan in the game, and the contingency plan will choose L should you choose Y in the game." To that player 1 can counter: "I don't think you would carry out such a contingency plan, which is inferior to its contrary as long as there is even a slight, ϵ probability that I might play Y: Your contingency plan would give you an expected payoff equal to $\epsilon \cdot 1 + (1 - \epsilon)4$, whereas playing R would give a higher expected payoff equal to $\epsilon \cdot 2 + (1 - \epsilon)4$."

4.2 Normal form and the meaning of a strategy

In the game illustrated in Section 3, the strategies (X, Y, L, and R) in the normal form happen to be the same as the actions in the sequential-move game in Figure 1. In general, however, a strategy in the normal form is different from an action in the corresponding sequential-move game. Recall that a normal form describes a situation where, before the start of the sequential-move game, each player *simultaneously* chooses a *complete contingency plan* and, once the game has started, carries out the contingency plan throughout the game as if he were an automaton dictated by the contingency plan. Thus a *strategy*, namely a row or column in a normal form, should be a complete contingency plan that tells the player which action to take given which scenario, and the plan should cover *all* possible scenarios that could arise during the course of the game.

To illustrate, let us consider the following game:

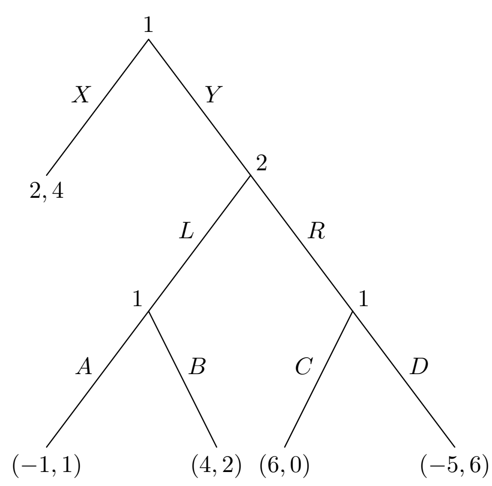

Figure 4: A game with three stages

Here, if he takes the action Y, player 1 gets to move again after player 2 has moved in response to Y (note the label "1" after the branches L and R). Thus the former's strategy, if prescribing Y, needs to specify which subsequent action to take (A or B, C or D) contingent on the action (L or R) that player 2 will take. An easy way to represent such a strategy is to write the prescribed actions in a sequence, say YBC, meaning that player 1 would take action Y to start and then play B if player 2 responds to Y with action L, and play C if player 2 responds to Y with R.

Thus, the normal form of the game in Figure 4 is:

	L	R
X	2, 4	2, 4
YAC	−1, 1	6, 0
YAD	$-1, 1$	−5, 6
YBC	4, 2	6, 0
YBD	4, 2	−5, 6

It is instructional to figure out how the payoff vector in each cell comes about. For example, consider the cell with row YAD and column R. Given these strategies, once the game starts, player 1 takes action Y; in response, player 2 takes action R; that leads to the endgame at the bottom-right in Figure 4, and the contingent plan YAD tells player 1 to take action D. Thus the outcome is $(-5, 6)$, which is the payoff vector registered in the cell at the intersection between row YAD and column R.

5 Stackelberg games

Let us use backward induction to solve a duopoly game of economic relevance, Stakelberg game: Two firms supply the same kind of outputs to a market, where the inverse demand function (cf. Chapter 1) is

$$P(q) := \begin{cases} a - bq & \text{if } 0 \leq q \leq a/b \\ 0 & \text{if } q \geq a/b \end{cases}$$

for some parameters $a, b > 0$. Firm 1 is the leader, and firm 2 the follower. First, the leader chooses an output quantity $q_1 \geq 0$. Second, knowing the q_1 that the leader has chosen, the follower chooses an output quantity $q_2 \geq 0$. Third, the market price is $P(q_1 + q_2)$, and at that price each firm i sells its output q_i and incurs the cost cq_i, where c is a parameter such that $0 \leq c < a$.

Clearly this is a sequential-move game. Writing it into its normal form would be cumbersome because each firm has infinitely many strategies. (Consider the leader, for example, any nonnegative real number corresponds to a pure strategy.) Backward induction, however, is a lot simpler: Let us start with the endgame, where firm 2, given q_1, solves the problem

$$\max_{q_2 \in \mathbb{R}_+} \left(P(q_1 + q_2) - c\right) q_2.$$

This is similar to the monopoly problem we have considered in Chapter 1, with q_1 here treated as a parameter to the decision maker (firm 2). (Here is the time to recollect the first- and second-order conditions for maximization in Chapter 1.)

Given any $q_1 \geq 0$, there are only two possibilities:

Case 1: $a - bq_1 - c \leq 0$ Then

$$a - b(q_1 + q_2) - c \leq a - bq_1 - c \leq 0$$

for any $q_2 \geq 0$. Namely, the market price $P(q_1 + q_2)$ is always below the marginal cost c. Then firm 2's per-unit profit is always nonpositive, and negative if firm 2 chooses any $q_2 > 0$. Thus, firm 2's best action is $q_2^* = 0$.

Case 2: $a - bq_1 - c > 0$ Then $a - b(q_1 + q_2) - c > 0$ for any sufficiently small (but positive) q_2. That is, firm 2's per-unit profit can be positive for some $q_2 > 0$. Thus, the firm does not choose $q_2 = 0$, which yields only zero profit, while any positive q_2 that is sufficiently small yields a positive profit. It follows that firm 2's decision becomes

$$\max_{q_2 \in (0,\infty)} \left(a - b(q_1 + q_2) - c\right) q_2. \tag{1}$$

Note that, based on the above reasoning, we have replaced the domain of firm 2's decision by the open set $(0, \infty)$. Now that firm 2's best action q_2^* is necessarily an interior solution, it satisfies the first-order condition

$$a - bq_1 - bq_2^* - c - bq_2^* = 0,$$

i.e.,

$$q_2^* = \frac{a - c - bq_1}{2b}. \tag{2}$$

(This equation suffices q_2^*'s being firm 2's profit-maximum because the second-order condition is always satisfied: taking the derivative of the objective in (1) and then taking the derivative of that derivative, we see that

$$\frac{\partial^2}{\partial q_2^2} \left(\left(a - b(q_1 + q_2) - c\right) q_2\right) = \frac{\partial}{\partial q_2} \left(a - b(q_1 + q_2) - c - bq_2\right) = -2b < 0.$$

Hence the derivative of the objective is decreasing in q_2.)

Now reason backward to firm 1's decision on q_1, taking into account firm 2's best response to q_1. First, the above Case 1 is trivial: If firm 1 is to choose some q_1 such that $a - bq_1 - c \leq 0$, then $q_1 = 0$, otherwise firm 1 suffers a self-inflicted loss. Hence the only rational outcome within this case is that the firm gets zero profit.

Second, consider any q_1 belonging to Case 2. Then firm 2's best response is given by Eq. (2). Thus, firm 1's decision becomes

$$\max_{q_1 \in \mathbb{R}_+} \left(P\left(q_1 + \frac{a - c - bq_1}{2b}\right) - c\right) q_1.$$

Note that firm 2's quantity is indirectly determined by firm 1's output via (2), just like in the game of Figure 1, where player 2's action (R) is indirectly determined by player 1's choice of Y. The above maximization problem is again similar to the monopoly problem in Chapter 1. Plug in the inverse demand function to rewrite the problem as

$$\max_{q_1 \in \mathbb{R}_+} \left(\frac{a - c}{2} - \frac{bq_1}{2}\right) q_1. \tag{3}$$

One can check that firm 1's best action q_1^* is determined by its fist-order condition

$$\frac{a - c}{2} - \frac{bq_1^*}{2} - \frac{bq_1^*}{2} = 0.$$

Thus

$$q_1^* = \frac{a-c}{2b}.$$

Since $a > c$ by assumption, one readily sees that the maximand of (3) is positive. Thus the interior solution $q_1^* = \frac{a-c}{2b}$ is better than the zero quantity and hence is the global maximum for firm 1.

In sum, we obtain the solution where firm 1 chooses the quantity $q_1^* = \frac{a-c}{2b}$ and firm 2 chooses

$$q_2^* = \frac{a-c-bq_1^*}{2b} = \frac{a-c-b\cdot\frac{a-c}{2b}}{2b} = \frac{a-c}{4b}.$$

6 Exercises

1. Solve the following game by iteratively eliminating strictly dominated strategies:

	L	M	R
U	3, 0	0, −3	0, −4
C	1, −1	3, 1	−2, 4
D	2, 4	4, 5	−1, 8

Hint: You might find it useful to follow these steps:

a. Consider the mixed strategy $\frac{1}{2}L + \frac{1}{2}R$ for player 2 (column player). Does it strictly dominate M for him? If yes then eliminate the column for M.

b. In the matrix resulting from eliminating M, eliminate all pure strategies that are strictly dominated for player 1 (row player).

c. In the matrix resulting from the elimination(s) in the previous step, which pure strategy is the best for player 2?

2. In a *second-price auction*, a single object is being auctioned off; each bidder writes down secretly his bid (in dollar values) and submits it in a sealed envelope; the auctioneer sells the object to the bidder who bids the highest (with ties broken by a fair coin toss) at the price equal to the highest bid among all the other bidders. Suppose that the object being auctioned off is the dollar bill.

a. Why is submitting \$1.40 weakly dominated by submitting \$1.00?

b. Why is submitting \$0.99 weakly dominated?

c. Find a dominant-strategy equilibrium of the game.

Hint: As far as a bidder concerns, no matter how complicated the strategies played by the other bidders are, it suffices to consider the highest bid among one's rivals.

3. Use iterative elimination of strictly dominated strategies to solve the following Guess the Average game:

> First, each player secretly writes on a piece of paper a number in the set $\{0, 1, \ldots, 100\}$ and his name. Then the papers are collected, and a prize of \$1 is split equally among all the players whose number is closest to 1/3 of the average number.

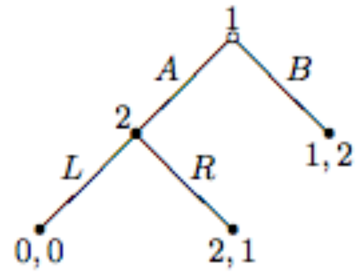

Figure 5: Exercise 4

4. Use backward induction to solve the following sequential-move game (Figure 5).

5. For the sequential game depicted in Figure 4:

   a. Solve it by backward induction.

   b. Solve its normal form by iterative elimination of weakly dominated strategies.

   c. Do the two methods yield the same solution(s)?

6. Use backward induction to solve the following "Centipede Game" (Figure 6).

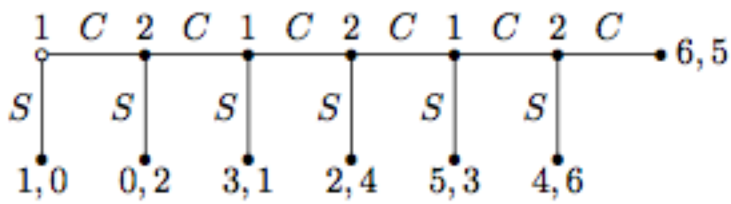

Figure 6: The Centipede Game

7. Suppose $a = 10$, $b = 2$ and $c = 0$ in the Stackelberg games. Work out the equilibrium through backward induction. (The general formula at the end of the chapter of course applies, but it is more instructional to work out the backward induction by yourself rather than simply plugging the numbers into the general formula.)

8. Given the general parameters of the Stackelberg games considered in this chapter, suppose, before the game starts, a second-price auction is held between the two firms to decide who gets to be the leader: the higher bidder becomes the leader after paying a price (to the auctioneer) equal to the other firm's bid; and this other firm becomes the follower; ties are broken by fair coin toss. Solve this auction game. (That is, how much would each firm bid for the first-mover advantage?)