
Eliminating Dominated Strategies

Um so schlimmer für die Tatsache
(So much the worse for the facts)
Georg Wilhelm Friedrich Hegel

4.1 Dominated Strategies

Suppose S_i is a finite set of pure strategies for players $i = 1, \dots, n$ in normal form game \mathcal{G} , so that $S = S_1 \times \dots \times S_n$ is the set of pure-strategy profiles for \mathcal{G} , and $\pi_i(s)$ is the payoff to player i when strategy profile $s \in S$ is chosen by the players. We denote the set of mixed strategies with support in S as $\Delta^* S = \Delta S_1 \times \dots \times \Delta S_n$, where ΔS_i is the set of mixed strategies for player i with support in S_i , or equivalently, the set of *convex combinations* of members of S_i . We denote the set of vectors of mixed strategies of all players but i by $\Delta^* S_{-i}$. We say $s'_i \in S_i$ is *strongly dominated* by $s_i \in S_i$ if, for every $\sigma_{-i} \in \Delta^* S_{-i}$, $\pi_i(s_i, \sigma_{-i}) > \pi_i(s'_i, \sigma_{-i})$. We say s'_i is *weakly dominated* by s_i if for every $\sigma_{-i} \in \Delta^* S_{-i}$, $\pi_i(s_i, \sigma_{-i}) \geq \pi_i(s'_i, \sigma_{-i})$, and for at least one choice of σ_{-i} the inequality is strict. Note that a strategy may fail to be strongly dominated by any pure strategy but may nevertheless be strongly dominated by a mixed strategy.

Suppose s_i is a pure strategy for player i such that every $\sigma'_i \neq s_i$ for player i is weakly (respectively strongly) dominated by s_i . We call s_i a *weakly* (respectively *strongly*) *dominant strategy* for i . If there is a Nash equilibrium in which all players use a dominant strategy, we call this a *dominant-strategy equilibrium*.

Once we have eliminated dominated strategies for each player, it often turns out that a pure strategy that was not dominated at the outset is now dominated. Thus, we can undertake a second round of eliminating dominated strategies. Indeed, this can be repeated until pure strategies are no longer eliminated in this manner. In a finite game, this will occur after a finite number of rounds and will always leave at least one pure strategy remaining for each player. If strongly (resp. weakly) dominated strategies are eliminated, we call this the *iterated elimination of strongly (resp. weakly)*

dominated strategies. We will refer to strategies that are eliminated through the iterated elimination of strongly (resp. weakly) dominated strategies as *recursively strongly* (resp. *weakly*) *dominated*.

In two-player games, the pure strategies that remain after the elimination of strongly dominated strategies correspond to the *rationalizable* strategies. The reader is invited to show that no pure strategy that is recursively strongly dominated can be part of a Nash equilibrium. Weakly dominated strategies can be part of a Nash equilibrium, so the iterated elimination of weakly dominated strategies may discard one or more Nash equilibria of the game. The reader is invited to show that at least one Nash equilibrium survives the iterated elimination of weakly dominated strategies.

There are games, such as the prisoner's dilemma (§3.11), that are "solved" by eliminating recursively strongly dominated strategies in the sense that after the elimination these strategies, a single pure strategy for each player remains and, hence, is the unique Nash equilibrium of the game.

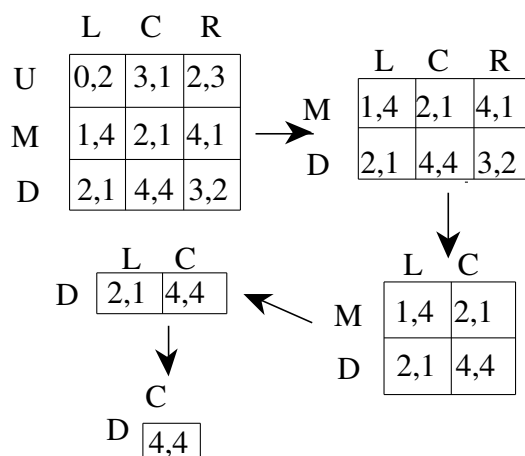


Figure 4.1. The iterated elimination of strongly dominated strategies

Figure 4.1 illustrates the iterated elimination of strongly dominated strategies. First, U is strongly dominated by D for player 1. Second, R is strongly dominated by $0.5L+0.5C$ for player 2 (note that a pure strategy in this case is not dominated by any other pure strategy but is strongly dominated by a mixed strategy). Third, M is strongly dominated by D. Finally, L is strongly dominated by C. Note that $\{D,C\}$ is indeed the unique Nash equilibrium of the game.

4.2 Backward Induction

We can eliminate weakly dominated strategies in extensive form games with *perfect information* (that is, where each information set is a single node) as follows. Choose any terminal node $t \in T$ and find the parent node of this terminal node, say node a . Suppose player i chooses at a , and suppose i 's highest payoff at a is attained at node $t' \in T$. Erase all the branches from a so a becomes a terminal node, and attach the payoffs from t' to the new terminal node a . Also record i 's move at a , so you can specify i 's equilibrium strategy when you have finished the analysis. Note that if the player involved moves at one or more other nodes, you have eliminated a *weakly* dominated strategy. Repeat this procedure for all terminal nodes of the original game. When you are done, you will have an extensive form game that is one level less deep than the original game. Now repeat the process as many times as is possible. If the resulting game tree has just one possible move at each node, then when you reassemble the moves you have recorded for each player, you will have a Nash equilibrium.

We call this *backward induction*, because we start at the end of the game and move backward. Note that backward induction can eliminate Nash equilibria that use weakly dominated strategies. Moreover, backward induction is *not* an implication of player Bayesian rationality, as we shall show below.

For an example of backward induction, consider figure 4.2, the Big John (BJ) and Little John (LJ) game where BJ goes first. We start with the terminal node labeled (0,0) and follow it back to the LJ node on the left. At this node, w is dominated by c because $1 > 0$, so we erase the branch where LJ plays w and its associated payoff. We locate the next terminal node in the original game tree, (4,4) and follow back to the LJ node on the right. At this node, c is dominated by w , so we erase the dominated node and its payoff. Now we apply backward induction to this smaller game tree; this time, of course, it is trivial. We find the first terminal node, (9,1), which leads back to BJ. Here c is dominated, so we erase that branch and its payoff. We now have our solution: BJ chooses w , LJ chooses cw , and the payoffs are (9,1).

You also can see from this example that by using backward induction and hence eliminating weakly dominated strategies, we have eliminated the Nash equilibrium c, ww (see figure 3.2). This is because when we assume LJ plays c in response to BJ's w , we have eliminated the weakly dominated

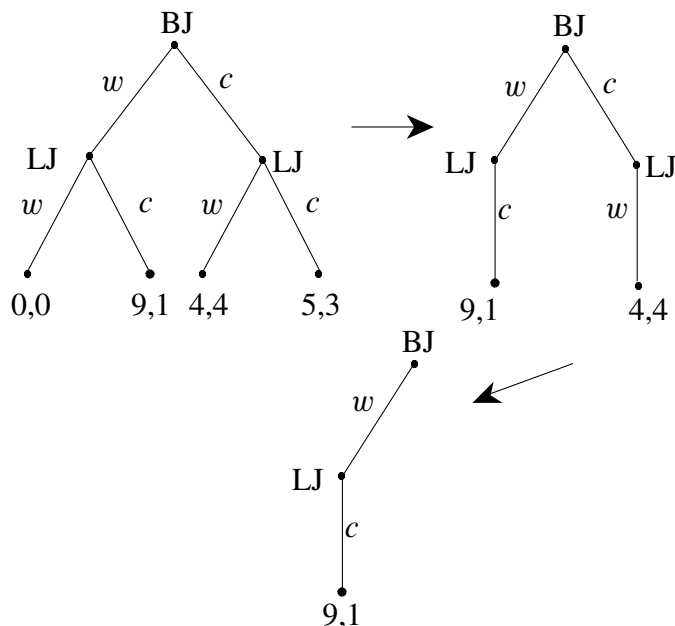


Figure 4.2. An example of backward induction

strategies ww and wc for LJ. We have called c, ww an incredible threat. Backward induction eliminates incredible threats.

4.3 Exercises in Eliminating Dominated Strategies

Apply the iterated elimination of dominated strategies to the following normal form games. Note that in some cases there may remain more than one strategy for each player. Say *exactly* in what order you eliminated rows and columns. Verify that the resulting solution is a Nash equilibrium of the game. Remember: You must eliminate a *whole row* or a *whole column*, not a single cell entry. When eliminating rows, compare the *first* of the two entries, and when eliminating columns, compare the *second* of the two entries.

	(a)	a	b	c		(b)	a	b	c
A		2,12	1,10	1,11		A	1,1	-2,0	4,-1
B		0,12	0,10	0,11		B	0,3	3,1	5,4
C		0,12	0,10	0,13		C	1,5	4,2	6,2

(c)

	N_2	C_2	J_2
N_1	73,25	57,42	66,32
C_1	80,26	35,12	32,54
J_1	28,27	63,31	54,29

(d)

	a	b	c	d	e
A	63, -1	28, -1	-2, 0	-2, 45	-3, 19
B	32, 1	2, 2	2, 5	33, 0	2, 3
C	54, 1	95, -1	0, 2	4, -1	0, 4
D	1, -33	-3, 43	-1, 39	1, -12	-1, 17
E	-22, 0	1, -13	-1, 88	-2, -57	-3, 72

(e)

	a	b	c	d	e
A	0, 1	0, 1	0, 1	0, 1	0, 1
B	0.81, 0, 19	0.20, 0.80	0.20, 0.80	0.20, 0.80	0.20, 0.80
C	0.81, 0.19	0.49, 0.51	0.40, 0.60	0.40, 0.60	0.40, 0.60
D	0.81, 0.19	0.49, 0.51	0.25, 0.75	0.60, 0.40	0.60, 0.40
E	0.81, 0.19	0.49, 0.51	0.25, 0.75	0.09, 0.91	0.80, 0.20
F	0.81, 0.19	0.49, 0.51	0.25, 0.75	0.09, 0.91	0.01, 0.99

(f)

	a	b	c	d	e	f	g	h
A	-1, 1	-1, 1	-1, 1	-1, 1	1, -1	1, -1	1, -1	1, -1
B	1, -1	0, 0	1, -1	0, 0	1, -1	0, 0	1, -1	0, 0
C	-1, 1	-1, 1	0, 0	0, 0	-1, 1	-1, 1	0, 0	0, 0

(g)

	0	1	2	3	4	5
0	4, 5	4, 14	4, 13	4, 12	4, 11	4, 10
1	13, 5	3, 4	3, 13	3, 12	3, 11	3, 10
2	12, 5	12, 4	2, 3	2, 12	2, 11	2, 10
3	11, 5	11, 4	11, 3	1, 2	1, 11	1, 10
4	10, 5	10, 4	10, 3	10, 2	0, 1	0, 10

4.4 Subgame Perfection

Let h be a node of an extensive form game \mathcal{G} that is also an information set (that is, it is a *singleton* information set). Let \mathcal{H} be the smallest collection of nodes including h such that if h' is in \mathcal{H} , then all of the successor nodes of h' are in \mathcal{H} and all nodes in the same information set as h' are in \mathcal{H} . We endow \mathcal{H} with the information set structure, branches, and payoffs inherited from \mathcal{G} , the players in \mathcal{H} being the subset of players of \mathcal{G} who move at some information set of \mathcal{H} . The reader is invited to show that \mathcal{H} is an extensive form game. We call \mathcal{H} a *subgame* of \mathcal{G} .

If \mathcal{H} is a subgame of \mathcal{G} with root node h , then every pure strategy profile s_G of \mathcal{G} that reaches h has a counterpart s_H in \mathcal{H} , specifying that players in \mathcal{H} make the same choices with s_H at nodes in \mathcal{H} as they do with s_G at the same node in \mathcal{G} . We call s_H the *restriction* of s_G to the subgame \mathcal{H} . Suppose $\sigma_G = \alpha_1 s_1 + \dots + \alpha_k s_k$ ($\sum_i \alpha_i = 1$) is a mixed strategy of \mathcal{G} that reaches the root node h of \mathcal{H} , and let $I \subseteq \{1, \dots, k\}$ be the set of indices such that $i \in I$ iff s_i reaches h . Let $\alpha = \sum_{i \in I} \alpha_i$. Then, $\sigma_H = \sum_{i \in I} (\alpha_i / \alpha) s_i$ is a mixed strategy of \mathcal{H} , called the *restriction* of σ_G to \mathcal{H} . We are assured that $\alpha > 0$ by the assumption that σ_G reaches h , and the coefficient α_i / α represents the probability of playing s_i , conditional on reaching h .

We say a Nash equilibrium of an extensive form game is *subgame perfect* if its restriction to every subgame is a Nash equilibrium of the subgame. It is clear that if s_G is a Nash equilibrium for a game \mathcal{G} , and if \mathcal{H} is a subgame of \mathcal{G} whose root node is reached with positive probability using s_G , then the restriction s_H of s_G to \mathcal{H} must be a Nash equilibrium in \mathcal{H} . But if \mathcal{H} is *not* reached using s_G , then it does not matter what players do in \mathcal{H} ; their choices in \mathcal{H} can be completely arbitrary, so long as these choices do not alter the way people play in parts of the game that *are* reached with positive probability. But how, you might ask, could this make a difference in the larger game? To see how, consider the following Microsoft-Google game.

Microsoft and Google are planning to introduce a new type of Web browser. They must choose between two platforms, Java and ActiveX. If they introduce different platforms, their profits are zero. If they introduce the same platform, their profits are 1, plus Microsoft gets 1 if the platform is ActiveX and Google gets 1 if the platform is Java. The game tree is as shown in figure 4.3, where we assume that Microsoft chooses first.

Because Microsoft has one information set and two choices at that information set, we can write its pure-strategy set as $\{\text{ActiveX}, \text{Java}\}$. Because

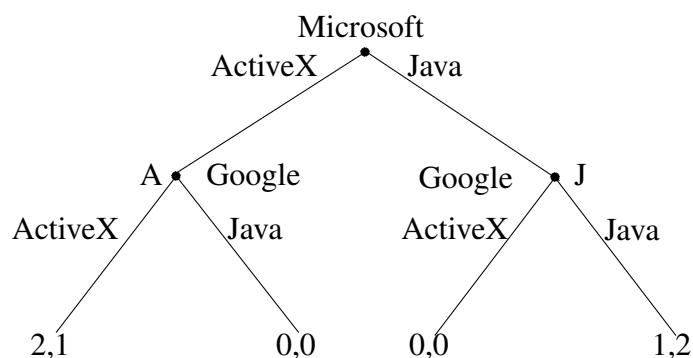


Figure 4.3. The Microsoft-Google game

Google has two information sets and two choices at each, its pure-strategy set has four elements $\{JJ, JA, AJ, AA\}$. Here, the first letter says what Google does if Microsoft plays ActiveX, and the second says what Google does if Microsoft plays Java. Thus, “AJ” means “play ActiveX if Microsoft plays ActiveX, and play Java if Microsoft plays Java.”

It is easy to see that $\{A, AA\}$, $\{A, AJ\}$, and $\{J, JJ\}$ are all Nash equilibria (prove this formally!), but only the second is subgame perfect, because the first is not a Nash equilibrium when restricted to the subgame starting at J, and the third is not a Nash equilibrium when restricted to the subgame at A. The outcomes from $\{A, AA\}$ and $\{A, AJ\}$ are the same, because ActiveX is chosen by both players in either equilibrium. But the outcome from $\{J, JJ\}$ is that Java is chosen by both players. This is obviously to the benefit of Google, but it involves an *incredible threat*: Google threatens to choose Java *no matter what Microsoft does*, but Microsoft knows that when it actually comes time to choose at A, Google will in fact choose ActiveX, not Java.

It is easy to see that a simultaneous-move game has no proper subgames (a game is always a subgame of itself; we call the whole game an *improper* subgame), because all the nodes are in the same information set for at least one player. Similarly, a game in which Nature makes the first move and the outcome is not known by at least one other player also has no proper subgames.

At the other extreme, in a game of perfect information (that is, for which all information sets are singletons), *every* nonterminal node is the root node of a subgame. This allows us to find the subgame perfect Nash equilibria of such games by backward induction, as described in section 4.2. This line of

reasoning shows that in general, backward induction consists of the iterated elimination of weakly dominated strategies.

4.5 Stackelberg Leadership

A Stackelberg leader is a player who can precommit to following a certain action, so other players effectively consider the leader as “going first,” and they predicate their actions on the preferred choices of the leader.

	t_1	t_2
s_1	0,2	3,0
s_2	2,1	1,3

Stackelberg leadership is a form of power flowing from the capacity to precommit. To see this, note that a form of behavior that would be an incredible threat or promise without the capacity to precommit becomes part of a subgame perfect Nash equilibrium when precommitment is possible. Formally, consider the normal form game in the matrix.

- Write an extensive form game for which this is the corresponding normal form, and find all Nash equilibria.
- Suppose the row player chooses first and the column player sees the row player’s choice, but the payoffs are the same. Write an extensive form game, list the strategies of the two players, write a normal form game for this new situation, and find the Nash equilibria.
- On the basis of these two games, comment on the observation “The capacity to precommit is a form of power.”

4.6 The Second-Price Auction

A single object is to be sold at auction. There are $n > 1$ bidders, each submitting a single bid b_i , simultaneously, to the seller. The value of the object to bidder i is v_i . The winner of the object is the highest bidder, but the winner pays only the next highest bid. Thus if i wins and j has made the next highest bid, then i ’s payoff is $\pi_i = v_i - b_j$, and the payoff for every other bidder is zero.

Show clearly and carefully that a player cannot gain by deviating from truth telling, no matter what the other players do.

- Does this answer depend on whether other players use this strategy?
- Does this answer still hold if players are allowed to place a higher bid if someone bids higher than their current bid?

- c. Show that if you know every other bid, you still cannot do better than bidding $b_i = v_i$, but there are weakly dominated best responses that affect the payoff to other players.
- d. Show that if the bids arrive in a certain order, and if later bidders condition their bids on previous bids, then bidding $b_i = v_i$ may not be a dominant strategy. Hint: consider the case where bidder i sets $b_i = 0$ if any previous bidder has placed a bid higher than v_i .
- e. Consider the following *increasing price auction*: There are n bidders. The value of the vase to bidder i is $v_i > 0$, $i = 1, \dots, n$. The auctioneer begins the bidding at zero, and raises the price at the rate of f\$1 per ten seconds. All bidders who are willing to buy the vase at the stated price put their hands up simultaneously. This continues until there is only one arm raised. This last and highest bidder is the winner of the auction and must buy the vase for \$1 less than the stated price. Show that this auction has the same normal form as the second-price auction. Can you think of some plausible Nash equilibria distinct from the truth-telling equilibrium?
- f. What are some reasons that real second-price auctions might not conform to the assumptions of this model? You might want to check with eBay participants in second-price auctions. Hint: think about what happens if you do not know your own v_i , and you learn about the value of the prize by tracing the movement over time of the currently high bid.

4.7 The Mystery of Kidnapping

It is a great puzzle as to why people are often willing to pay ransom to save their kidnapped loved ones. The problem is that the kidnapers are obviously not honorable people who can be relied upon to keep their promises. This implies that the kidnapper will release rather than kill the victim *independent* of whether the ransom is paid. Indeed, assuming the kidnapper is *self-regarding* (that is, he cares only about his own payoffs), whether the victim should be released or killed should logically turn on whether the increased penalty for murdering the victim is more or less than offset by the increased probability of being identified if the victim is left alive. Viewed in this light, the threat to kill the victim if and only if the ransom is not paid is an incredible threat. Of course, if the kidnapper is a known group that has a reputation to uphold, the threat becomes credible. But often this is not the case. Here is a problem that formally analyzes this situation.

Suppose a ruthless fellow kidnaps a child and demands a ransom r from the parents, equal to \$1 million, to secure the release of the child, whom the parents value the amount v . Because there is no way to enforce a contractual agreement between kidnapper and parents to that effect, the parents must pay the ransom and simply trust that the kidnapper will keep his word. Suppose the cost to the kidnapper of freeing the child is f and the cost of killing the child is k . These costs include the severity of the punishment if caught, the qualms of the kidnapper, and whatever else might be relevant.

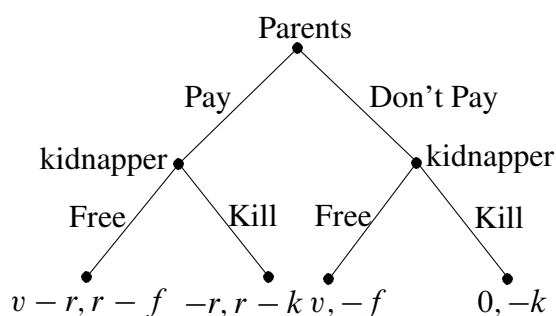


Figure 4.4. The mystery of kidnapping

Because the parents are the first movers, we have the extensive form game shown in figure 4.4. Using backward induction, we see that if $k > f$ the kidnapper will free the child, and if $k < f$ the kidnapper will kill the child. Thus, the parents' payoff is zero if they do not pay, and $-r$ if they pay, so they will not pay.

Parents	kidnapper			
	ff	fk	kf	kk
pay	$v - r, r - f$	$v - r, r - f$	$-r, r - k$	$-r, r - k$
don't pay	$v, -f$	$0, -k$	$v, -f$	$0, -k$

Figure 4.5. Normal form of the mystery of kidnapping game

The corresponding normal form of the game is shown in figure 4.5. Here, ff means “Free no matter what,” fk means “Free if pay, kill if not pay,” kf means “Kill if pay, free if not pay,” and kk means “kill no matter what.”

We may assume $v > r$, so the parents will pay the ransom if that ensures that their child will be released unharmed. We now find the Nash equilibria we found by applying backward induction to the extensive form game. Moreover, it is easy to see that if $k < f$, there is no additional Nash equilibrium where the child is released. The threat of killing the child is not