

Week 6: Dominant Strategies

Dr Daniel Sgroi

Reading: 1. Osborne sections 2.9 and 4.4;
2. Snyder & Nicholson, pp. 247–252.

With thanks to Peter J. Hammond.

Prisoner's Dilemma

In PD (prisoner's dilemma), each player has a best response ("confess" or "fink") that is best regardless of what their opponent chose.

		Player 2	
		<i>M</i>	<i>F</i>
Player 1	<i>M</i>	-1 -1	-9 0
	<i>F</i>	0 -9	-6 -6

This is an instance of a **dominant strategy**.

[Time allowing: mixed strategy equilibrium motivation]

Notation

Generally, let $u_i(s)$ denote player i 's utility from a strategy profile $s = (s_1, s_2, \dots, s_{i-1}, s_i, s_{i+1}, \dots, s_n)$.

Let $s_{-i} := (s_1, s_2, \dots, s_{i-1}, s_{i+1}, \dots, s_n)$ denote the strategy profile of all players other than i .

Evidently $s_{-i} \in S_1 \times S_2 \times \dots \times S_{i-1} \times S_{i+1} \times \dots \times S_n$ or equivalently $s_{-i} \in S_{-i} := \prod_{j \in N \setminus \{i\}} S_j$, the Cartesian product of the strategy spaces of all players other than i (player i 's "opponents").

Now we can rewrite the strategy profile as $s = (s_i, s_{-i})$ and player i 's utility $u_i(s)$ as $u_i(s_i, s_{-i})$.

Strictly Dominated Strategies

Definition

Given strategies $s_i, s'_i \in S_i$, say that s'_i is **strictly dominated** by s_i , or that s_i **strictly dominates** s'_i (written $s_i \succ_i^D s'_i$) if $u_i(s_i, s_{-i}) > u_i(s'_i, s_{-i})$ for any $s_{-i} \in S_{-i}$.

Inferior Responses

Proposition

A rational player will never play a strictly dominated strategy.

Proof.

No dominated strategy can ever be optimal because, by definition of strict dominance, there is another dominating strategy yielding a higher payoff regardless of the other players' strategies.

In fact, the dominating strategy yields a higher expected payoff regardless of the rational player's beliefs regarding other players' strategies.



A very simple, yet powerful lesson: when faced with any decision, look first for your dominated strategies, and then avoid them!

Strict Dominance in PD

		P_2	
		M	F
P_1	M	-2 -2	-5 -1
	F	-1 -5	-4 -4

Table: Prisoner's Dilemma

As argued earlier, playing M is worse than playing F for each player regardless of what the player's opponent does.

Thus, M is strictly dominated by F for both players.

For the PD, rationality alone predicts a unique outcome:

(F, F) is the only pair of strategies that are not strictly dominated.

Many games will not be as simple

in that rationality alone will not offer a unique prediction.

Cournot Duopoly

Recall the Cournot duopoly game.

Each firm i 's variable cost of producing quantity $q_i \geq 0$ was given by the cost function $c_i(q_i) = q_i^2$ and the demand was given by $p(q) = 100 - q$ (or more exactly, $p(q) = \max\{0, 100 - q\}$).

Any $q_i > 100$ will yield negative profits to player i . The reason being that $q_i > 100$ guarantees that total quantity $q = q_1 + q_2$ will exceed 100, so the resulting price will be zero.

By choosing the quantity $q_i = 0$, player i is guaranteed zero profits.

Hence any quantity $q_i > 100$ is strictly dominated by $q_i = 0$.

Strictly Dominant Strategies

In some games like prisoner's dilemma, avoiding strictly dominated strategies leaves a unique strategy that is always best, regardless of what other players do. Formally:

Definition

A strategy $s_i \in S_i$ is **strictly dominant** for i if every alternative strategy $s'_i \in S_i$ is strictly dominated by s_i — that is, if $u_i(s_i, s_{-i}) > u_i(s'_i, s_{-i})$ for all $s'_i \in S_i$ with $s'_i \neq s_i$ and for all $s_{-i} \in S_{-i}$.

That is, $s_i \in S_i$ is strictly dominant provided that $s_i \succ_i^D s'_i$ for all $s'_i \in S_i$ with $s'_i \neq s_i$.

In the PD example, $F \succ_i^D M$ for each of the two players.

So each player has a unique strictly dominant strategy: fink.

Strict Dominant Strategy Equilibrium

Suppose that, as in the PD,
each player has one fixed strategy that is uniquely best
regardless of what other players do.
Then each player has a unique dominant strategy
that we should expect to be chosen.

Definition

The strategy profile $s^D \in S$
is a **strict dominant strategy equilibrium**
if $s_i^D \in S_i$ is a strict dominant strategy for all $i \in N$.

Proposition

*Any strictly dominant strategy equilibrium s^D
in a game $\Gamma = \langle N, (S_i)_{i=1}^n, (u_i)_{i=1}^n \rangle$ is unique.*

Dominant Strategy Equilibria May Not Exist

In BoS neither strategy dominates the other, for either player.

		P_2	
		B	S
P_1	B	2 0	1 0
	S	0 0	1 2

Table: Bach or Stravinsky (or Battle of the Sexes)

So there is no dominant strategy equilibrium.

Pareto Efficiency

Definition

One strategy profile $s \in S$ **Pareto dominates** any other strategy profile $s' \in S$ if $u_i(s) \geq u_i(s')$ for all $i \in N$, with $u_i(s) > u_i(s')$ for at least one $i \in N$ (in which case we also say that s' is **Pareto dominated** by s .)

We might like players to find ways of coordinating to reach Pareto optimal outcomes, and to avoid those that are Pareto dominated.

But in many games like PD, the Nash or dominant strategy equilibrium profile like (F, F) is Pareto dominated by an alternative (M, M) .

Notice in PD that all the three other profiles (M, M) , (M, F) and (F, M) are Pareto efficient, since no other profile Pareto dominates any of them.

Institutional Inefficiency

This inefficient outcome of PD is *not* a failure of the solution concept.

Pareto inefficiency does imply, however, that the players might benefit from modifying the **institutional environment** reflected in our model of the game.

The players of PD might want to create a norm of conduct designed to rule out talking to the authorities (“honour among thieves” perhaps). Organized crime would be one such institution.

PD with no-talking norm

Let v denote the cost of violating *omertà*, measured in the extra years of jail time.

		P_2	
		M	F
P_1	M	-2 -2	-5 -1 - v
	F	-1 - v -5	-4 - v -4 - v

Table: “No-talking norm”-modified Prisoner’s Dilemma

Any $v > 1$ alters the predicted equilibrium outcome by making M the new strictly dominant strategy for each player. It also ensures that there is a strictly dominant strategy profile $s^* \in S$ satisfying $u_i(s^*) > u_i(s)$ for all $i \in N$ and all $s \in S$ satisfying $s \neq s^*$. So playing strictly dominant strategies is Pareto efficient in the “no-talking norm”-modified PD.

Institutional Design

Notice that if the players find it credible that the criminal community will punish any who violate the “no-talking norm”, then there is actually no need to punish anybody — the players choose not to fink, so there is no punishment.

However, to be sure that such an institution really is self-enforcing, we need to model the behaviour of potential outside enforcers, and whether they themselves have the selfish incentives to carry out the enforcement activities.

Of course then we move into the world of repeated games which we studied in an earlier lecture. We will also return to multi-period games later but in a slightly different guise (dynamic games).

A currently important topic in economics is the design of institutions like markets, auctions, stock exchanges, banks.

Weakly Dominated Strategies

Definition

Given strategies $s_i, s'_i \in S_i$,

say that s'_i is **weakly dominated** by s_i , and write $s_i \succ_i^W s'_i$,

if $u_i(s_i, s_{-i}) \geq u_i(s'_i, s_{-i})$ for any $s_{-i} \in S_{-i}$,

with $u_i(s_i, s_{-i}) > u_i(s'_i, s_{-i})$ for some $s_{-i} \in S_{-i}$.

This means that for some $s_{-i} \in S_{-i}$

the weak inequality can hold with equality,

while for at least one $s_{-i} \in S_{-i}$ it must hold strictly.

Definition

Define a strategy $s_i \in S_i$ to be **weakly dominant**

if $s_i \succ_i^W s'_i$ for all $s'_i \in S_i$ with $s'_i \neq s_i$.

Weak dominance is also useful since player i never does worse by replacing strategy s'_i with the weakly dominating s_i , and may do better.

Uniqueness of Weakly Dominant Strategies

Like a strictly dominant strategy,
a player can have at most one weakly dominant strategy;
if $s_i \succ_i^W s'_i$, then $s'_i \succ_i^W s_i$ is impossible;
so if $s_i \succ_i^W s'_i$ for all $s'_i \in S_i$ with $s'_i \neq s_i$,
then s_i must be the only weakly dominant strategy.

Weakly Dominant Equilibrium

Call the strategy profile $s^W \in S$

a **weakly dominant strategy equilibrium**

if $s_i^W \in S_i$ is a weakly dominant strategy for all $i \in N$.

That is, if $s_i^W \succsim_i^W s'_i$ for all $s'_i \in S_i$, for all $i \in N$.

Obviously BoS is an example of a game with no weakly dominant strategy equilibrium.

Uniqueness?

There is an alternative weaker relation (e.g. used in Steve Tadelis's notes): \succsim_i^W defined so that $s_i \succsim_i^W s'_i$ if $u_i(s_i, s_{-i}) \geq u_i(s'_i, s_{-i})$ for any $s_{-i} \in S_{-i}$.

Notice that one can have both $s_i \succsim_i^W s'_i$ and $s'_i \succsim_i^W s_i$; indeed, both hold

iff [i.e., if and only if] $u_i(s_i, s_{-i}) = u_i(s'_i, s_{-i})$ for all $s_{-i} \in S_{-i}$.

Therefore, an important difference between \succ_i^W and \succsim_i^W is that if a \succsim_i^W -dominant equilibrium exists, it need not be unique (whereas a \succ_i^W -dominant equilibrium must be unique).

Iterated Elimination of Strictly Dominated Strategies (or IESDS)

The individual rationality hypothesis has two important implications:

1. a rational player will never play a dominated strategy;
2. a rational player with a strictly dominant strategy will play it.

Question: Doesn't the first conclusion imply the second?

Answer: Yes, but only in the relatively few games where the rational player does have a strictly dominant strategy.

Conclusion 2 leads to the strict dominance solution concept that is very appealing, but often fails to exist.

A more general theory of strategic choice relies on an alternative approach that offers predictions for a broader class of games.

Common Knowledge of Rationality . . .

Conclusion 1 claims that a rational player never chooses a dominated strategy.

This implies that some strategies will not be played **under any circumstances**.

We assume that both the structure of the game and players' rationality are **common knowledge** among the players — i.e., all players know these facts, know that all know them, know that all know they know them, and so on *ad infinitum*.

Common knowledge of rationality implies that **all players know** that each player will never play a strictly dominated strategy.

So each player can “eliminate” all strictly dominated strategies for **every player**, since everybody knows these will never be played.

... and Iterated Elimination

Suppose that in the original game there are some players with some strictly dominated strategies. Then all the players know they effectively face a “smaller” game with those dominated strategies eliminated. Also, since it is common knowledge that all players are rational, everyone knows that everyone knows that the game is effectively smaller.

In this smaller game, again everyone knows that players will not play strictly dominated strategies. So a player's strategy that was originally undominated may become dominated after eliminating some of the other players' strategies in the first stage above. Since it is common knowledge that all the players will perform this kind of reasoning again and again, this process can go on until it can eliminate no more strategies.

A Simple Example: First Round

		P_2		
		L	C	R
P_1	U	4 3	5 1	6 2
	M	2 1	8 4	3 6
	D	3 0	9 6	2 8

Table: Original Game

Exercise: Find every instance of strict dominance.

Answer: For player 1, no strategy strictly dominates any other. For player 2, one has $R \succ_2^D C$, but no other strict dominance pair. So neither player has a strictly dominant strategy. But player 2's strategy C is strictly dominated by R . So both players argue we can exclude C from the game.

A Simple Example: Second Round

One round of elimination results in the following reduced game:

		P_2	
		L	R
P_1	U	4 3	6 2
	M	2 1	3 6
	D	3 0	2 8

Table: After One Round of Elimination

We move on to try a second round of elimination.

In the reduced game, after eliminating C for player 2, one has both $U \succ_1^D M$ and $U \succ_1^D D$.

So both M and D are now strictly dominated for player 1, and should be eliminated on the second round.

A Simple Example: Third Round

Eliminating both M and D on the second round yields the following greatly reduced game:

		P_2	
		L	R
P_1	U	4 3	6 2

Table: After Two Rounds of Elimination

Here player 2 has a strictly dominant strategy, L . For this example IESDS yields a unique prediction: the strategy profile we expect is (U, L) , yielding the payoff pair $(4, 3)$.

Caveat

Caveat: Beware the common error of referring to the payoff pair $(4, 3)$ as the solution.

The solution should always be described as the strategy profile(s) that the players will choose.

Predictions, or equilibria, should always refer to **what players do** as the equilibrium, not to their payoffs.

The Iterated Elimination Algorithm

The elimination process is an algorithm for constructing recursively a shrinking sequence of strategy sets S_i^k ($k = 0, 1, 2, \dots$) for each player $i \in N$ as follows:

- Step 1:** Starting with $k = 0$, set S_i^0 as the game's original strategy set S_i of each player i .
- Step 2:** At stage k , for **every** player $i \in N$ ask whether any strategies $s_i \in S_i^k$ are strictly dominated by others in the same set.
- Step 3:** While the answer in Step 2 remains 'yes' for some players, add one to k , and define a new game Γ^k whose strategy sets S_i^k exclude those strategies that have just been identified as strictly dominated, then go back to Step 2.
Once the answer in Step 2 becomes 'no' for all players, stop iterating and go on to the final Step 4.
- Step 4:** The strategies remaining in each set S_i^k are reasonable predictions for behavior.

IESDS

This iterative process is called

Iterated Elimination of Strictly Dominated Strategies (or IESDS).

When the strategy sets are finite, it must terminate after finitely many iterations with non-empty sets for each player.

IESDS builds iteratively on common knowledge of rationality:

1. the first stage of dominated strategy elimination is due to **level 1** rationality;
2. the second **level 2** stage is due to each player's knowledge that other players are level 1 rational;
3. the third **level 3** stage is due to each player's knowledge that other players are level 2 rational.

And so on, through as many rationality levels as are needed for the elimination process to come to an end.

Iterated Elimination Equilibrium

Definition

Any strategy profile $s^{IE} = (s_1^{IE}, \dots, s_n^{IE})$
that survives the IESDS process
will be called an **iterated-elimination equilibrium**.

This is another solution concept.

Like a strictly dominant strategy equilibrium,
an iterated-elimination equilibrium presumes player rationality.
But not on just one level of rationality.

Instead, the added assumption of common knowledge of rationality
implies order k rationality, where k is large enough
for the iterative elimination process to be completed.

Comment

Strictly speaking, the procedure discussed here is one of iterated elimination of strategies that are **strictly dominated by other pure strategies**.

Sometimes strategies are strictly dominated by mixed strategies, even if they are not dominated by any pure strategy.

You will be asked to work through an exercise to show this.

Sometimes, with infinite strategy sets, the IESDS process does not stop after any finite number of rounds.

Instead, it may go on eliminating strategies until it reaches a well defined limit as the number of rounds tends to infinity.

Our next Cournot duopoly example illustrates this possibility.

The Cournot Duopoly Model Revisited

Consider a simpler version of the Cournot Duopoly example, where firms' costs are zero.

As before, let demand be given by $p(q) = \max\{0, 100 - q\}$, where $q = q_1 + q_2$.

Firm 1's profit (utility) function

is $u_1(q_1, q_2) = p(q)q_1 = \max\{0, (100 - q_1 - q_2)q_1\}$, which is just its revenue, of course.

We need to see which strategies are dominated

To do so, restrict attention to strategies satisfying $q_1, q_2 \in [0, 100)$, since any $q_i \geq 100$ gives a zero price and zero revenue.

(Any such strategy would be strictly dominated if there were even very small costs.

In any case, it is weakly dominated.)

Cournot Revisited: Round 1

Denote the partial derivative $\frac{\partial}{\partial q_1} u_1(q_1, q_2)$ by $u'_1(q_1, q_2)$.

Because $u_1(q_1, q_2) = \max\{0, (100 - q_1 - q_2)q_1\}$,
firm 1's marginal profit (or revenue) is $u'_1(q_1, q_2) = 100 - 2q_1 - q_2$
provided that $100 - q_1 - q_2 > 0$.

Because $q_2 \geq 0$, one has $u'_1(q_1, q_2) \leq 100 - 2q_1 < 0$ when $q_1 > 50$.
So any $q'_1 > 50$ is strictly dominated by $q_1 = 50$.

But no other strictly dominated strategies exist.

Indeed, any particular $q_1 \in [0, 50]$ is firm 1's best response
when firm 2's output happens to be $q_2 = 100 - 2q_1$.

So the quantity interval $[0, 50]$ remains
after all strictly strategies have been eliminated.

On the first round, level 1 rationality and symmetry
imply that each firm chooses $q_i \in [0, 50]$.

Cournot Revisited: Round 2

On the second round, because $q_2 \leq 50$, one has $u'_1(q_1, q_2) = 100 - 2q_1 - q_2 \geq 50 - 2q_1 > 0$ whenever $q_1 < 25$.

So any $q'_1 < 25$ is strictly dominated by $q_1 = 25$.

But no other strictly dominated strategies exist.

Indeed, any particular $q_1 \in [25, 50]$ is firm 1's best response when firm 2's output happens to be $q_2 = 100 - 2q_1 \in [0, 50]$.

On the second round, therefore, level 2 rationality and symmetry together imply that each firm chooses $q_i \in [25, 50]$.

Cournot Revisited: Round 3

On the third round, because $q_2 \geq 25$,
one has $u'_1(q_1, q_2) = 100 - 2q_1 - q_2 \leq 75 - 2q_1 < 0$
whenever $q_1 > 37.5$.

So any $q'_1 > 37.5$ is strictly dominated by $q_1 = 37.5$.

But no other strictly dominated strategies exist.

Indeed, any particular $q_1 \in [0, 37.5]$ is firm 1's best response
when firm 2's output happens to be $q_2 = 100 - 2q_1 \in [25, 50]$.

Only the quantity interval $q_i \in [25, 37.5]$ remains.

Cournot Revisited: Further Rounds

We can prove by induction that after round k , only the quantity interval $q_i \in [a_k, b_k]$ remains, where $a_0 = 0$, $b_0 = 100$, and:

1. $a_{k+1} = a_k$ and $b_{k+1} = \frac{1}{2}(a_k + b_k)$ if k is even;
2. $a_{k+1} = \frac{1}{2}(a_k + b_k)$ and $b_{k+1} = b_k$ if k is odd;
3. $b_k - a_k = 100 \cdot 2^{-k}$,

so the quantity interval keeps halving in length;

In the limit, therefore, as $k \rightarrow \infty$,

only a degenerate interval of zero length remains.

In fact IESDS leads to the unique **Cournot solution** with $q_1 = q_2 = 100/3$.