

Economics 51: Game Theory

Liran Einav

April 21, 2003

So far we considered only decision problems where the decision maker took the environment in which the decision is being taken as exogenously given: a consumer who decides on his optimal consumption bundle takes the prices as exogenously given. In perfectly competitive markets the assumption that my own actions do not influence the behavior of other agents or do not affect the market price is very reasonable. However, often this is not a good assumption. Firms which decide how to set prices, certainly take into account that their competitors might set lower prices for similar products. Furthermore in the free rider problem which can occur when there are public goods we saw that one agent's decision certainly takes into account that it might change the other agents' behavior and therefore the total supply of public goods. The tools we have studied so far are not adequate to deal with this problem.

The purpose of Game Theory is to analyze optimal decision making in the presence of strategic interaction among the players.

1 Definition of a Game

We start with abstractly defining what we mean by a game.

A game consists of

- a set of players:

In these notes we limit ourselves to the case of 2 players – everything generalizes to N players.

- a set of possible strategies for each player:

We denote a possible strategy for player $i = 1, 2$ by s_i and the set of all possible strategies by S_i .

- a payoff function that tells us the payoff that each player gets as a function of the strategies chosen by all players:

We write payoff functions directly as functions of the strategies, $v^i(s_1, s_2)$.

If each player chooses a strategy, then the result is typically some physical outcome. However, the only thing that is relevant about each physical outcome are the payoffs that it generates for all players. Therefore, we ignore the physical outcomes and only look at payoffs. Payoffs should be interpreted as von Neumann-Morgenstern utilities, not as monetary outcomes. This is important if there is uncertainty in the game.

Sometimes we write $v^i(s_i, s_{-i})$ to show that payoff for individual i depends on his own strategy s_i and on his opponent's strategy $s_{-i} \in S_{-i}$.

We will assume throughout that all players know the structure of the game including the payoff function of the opponent (one can analyze games without this assumption, but this is slightly too complicated for this class, the econ department offers a class in game theory where stuff like this is discussed; also, we will come back to this when we will talk about auctions).

We will distinguish between *normal-form games* and *extensive-form games*.

In normal form games (the reason why they have this name will become clearer later on) the players have to decide simultaneously which strategy to choose. Therefore, time is not important in these games.

In some games it might be useful to explicitly model time. I.e. when two firms change their prices over time to get a larger fraction of the market, they will certainly take into account past actions. Things like this will be modeled in extensive form games. We will come back to this in Section 3 below.

1.1 Some examples of normal-form games

It is useful to consider three different examples of games.

1.1.1 Prisoners' dilemma

Rob and Tom commit a crime, get arrested and the police interrogates them in separate rooms. If one of them pleads guilty and the other one does not, the first can cut a deal with the police and go off free. The other one is sent to jail for 20 years. If they both plead guilty, they both go to jail for 10 years. However, if they both maintain their innocence each one goes to jail only for a year. The strategies in this game are “confess” and “not confess,” for both Rob and Tom. If we assume that utility is negative of ‘time in jail’, we can summarize the payoffs in a matrix

	Confess	Don't confess
Confess	(-10,-10)	(-20,0)
Don't confess	(0,-20)	(-1,-1)

In this matrix, the horizontal player is Rob, the vertical player is Tom – each entry of the matrix gives Rob’s payoffs, then Tom’s payoffs (the convention is to write the horizontal guy’s payoff first).

1.1.2 A coordination game

A different example of a game is about how Rob and Tom might have to coordinate on what they want to do in the evening (assuming they are not in jail anymore). Rob and Tom want to go out. Tom likes hockey (he gets 5 utils from going to a hockey game), but not baseball (he gets 0 utils from that). Rob likes baseball (gets 5 utils) but not hockey (gets 0 utils). Mainly, however, they want to hang out together, so each one gets 6 utils from attending the same sport event as his friend. The payoff matrix for this game looks as follows (Rob is the horizontal guy):

	Hockey	Baseball
Hockey	(6,11)	(0,0)
Baseball	(5,5)	(11,6)

1.1.3 Serve and return in tennis

As a third and last example, suppose Ron and Tom play tennis. If Rob serves to a side where Tom stands he loses the point, if he serves to a side where Tom does not stand, he wins the point The payoff matrix to this game (Rob being the horizontal guy) is as follows:

	Left	Right
Left	(-1,1)	(1,-1)
Right	(1,-1)	(-1,1)

2 Solution concepts for normal form games

In this section we want to examine what will happen in each one of the examples, if we assume that both players are rational and choose their strategies to maximize their utility.

2.1 Dominant strategies

In the prisoners’ dilemma, it is easy to see what’s going to happen: whatever Rob does, Tom is better off confessing. Whatever Tom does, Rob is better off confessing. So they will both plead guilty. Confessing is a dominant strategy.

Formally, a *dominant strategy* for some player i is a strategy $s_i \in S_i$ such that for all $s_{-i} \in S_{-i}$ and all $\tilde{s}_i \in S_i$

$$v^i(s_i, s_{-i}) \geq v^i(\tilde{s}_i, s_{-i})$$

When **both** players in the game have a dominant strategy we know what will happen: each player plays his dominant strategy and we can describe the outcome. We call this an *equilibrium in dominant strategies*.

There is little doubt that in games where such an equilibrium exists this often will be the actual outcome. Note, that the outcome might not be Pareto-efficient: in the prisoners' dilemma game, the *cooperative* outcome would be for both guys to keep their mouths shut.

2.2 Nash equilibrium

In the coordination game neither Rob nor Tom have a dominant strategy. If Tom goes to hockey, Rob is better off going to hockey, but if Tom goes to baseball, Rob is better off going to baseball.

In order to solve this game (i.e. say what will happen) we need an equilibrium concept which incorporates the idea that Tom's optimal action will depend on what he thinks Rob will do and that Rob's optimal action will depend on Tom's.

A strategy profile (i.e. a strategy of player 1 together with a strategy of player 2) is a Nash-equilibrium if player 1's strategy is a 'best response' to what player 2 does (i.e. given what player 2 does, player one's strategy gives him the highest payoff) and vice versa.

A little bit more formally $(s_1, s_2) \in S^1 \times S^2$ is a Nash-equilibrium if

$$v^1(s_1, s_2) \geq v^1(\tilde{s}, s_2) \text{ for all } \tilde{s} \in S^1$$

$$v^2(s_1, s_2) \geq v^2(s_1, \tilde{s}) \text{ for all } \tilde{s} \in S^2$$

Note that this definition does require that player 1 knows what player 2 is going to do and player 2 knows what player 1 is up to.

2.2.1 Pure strategy Nash equilibria

Let's go back to the coordination game. Even though there is no equilibrium in dominant strategies (because none of the players has a dominant strategy) it turns out that there are at least two Nash-equilibria: (H, H) and (B, B) .

It is very easy to verify that these are both Nash-equilibria:

- (HH) is a Nash equilibrium because

$$v^R(H, H) = 6 > v^R(B, H) = 5 \text{ and } v^T(H, H) = 11 > v^T(H, B) = 0$$

- (BB) is a Nash equilibrium because

$$v^R(B, B) = 11 > v^R(H, B) = 0 \text{ and } v^T(B, B) = 6 > v^T(B, H) = 5$$

However, in the Tennis example, we cannot find a Nash equilibrium so easily: whatever Rob does, Tom will be better off doing something else – on the other hand, whatever Tom does, Rob will be better off doing the same. So we cannot find a pair of strategies which is a Nash equilibrium. In real life, Rob and Tom would be better off to randomize over their possible strategies.

2.2.2 Mixed strategies

In some games, it is useful to allow players to randomize over their possible strategies. We say they play a *mixed strategy*. A mixed strategy is simply a probability distribution over the player's pure strategies. Sometimes we will denote the set of all mixed strategies for some player i by Σ^i and a given mixed strategy by $\sigma^i \in \Sigma^i$. If there are only two pure strategies, a mixed strategy is just the probability to play the first pure strategy - it is just a number between zero and one. When there are more than 2 pure strategies, things get a little more complicated, but you should not worry about details there.

If players play mixed strategies they evaluate their utility according to the von-Neumann Morgenstern criterion. If player one has n_1 pure strategies and player 2 has n_2 pure strategies there are generally $n_1 \times n_2$ possible outcomes - i.e. possible states of the world. The probabilities of these states are determined by the mixed strategies (this will hopefully become clear in the examples below). We can write a player h 's payoff (utility function) as a function $u^h(\sigma^1, \sigma^2)$

A Nash equilibrium in mixed strategies is then simply a profile of mixed strategies (σ^1, σ^2) (in the cases below these will just be two probabilities) such that

$$\begin{aligned} u^1(\sigma^1, \sigma^2) &\geq u^1(\tilde{\sigma}, \sigma^2) \text{ for all } \tilde{\sigma} \in \Sigma^1 \\ u^2(\sigma^1, \sigma^2) &\geq u^2(\sigma^1, \tilde{\sigma}) \text{ for all } \tilde{\sigma} \in \Sigma^2 \end{aligned}$$

Equilibrium in Tennis

Suppose Rob can serve right with probability π_R and serve left with probability $(1 - \pi_R)$. Suppose Tom can stand right with probability π_T and stand left with probability $(1 - \pi_T)$. A mixed strategy can then be represented by a number:

$$\sigma^R = \pi_R \in [0, 1] \text{ and } \sigma^T = \pi_T \in [0, 1]$$

Utilities are von-Neumann-Morgenstern, for $h = R, T$:

$$u^h(\sigma^R, \sigma^T) = \pi_R \pi_T v^h(L, L) + \pi_R(1 - \pi_T)v^h(L, R) + \pi_T(1 - \pi_R)v^h(R, L) + (1 - \pi_T)(1 - \pi_R)v^h(R, R)$$

In tennis, $\pi_R = \pi_T = 1/2$ is a Nash-equilibrium. Evidently in this equilibrium payoffs are given by

$$u^R(0.5, 0.5) = \frac{1}{4}(v^R(L, L) + v^R(L, R) + v^R(R, L) + v^R(R, R)) = 0$$

$$u^T(0.5, 0.5) = \frac{1}{4}(v^T(L, L) + v^T(L, R) + v^T(R, L) + v^T(R, R)) = 0$$

To show that it is an equilibrium we must show that neither Rob nor Tom do better than getting a payoff of 0.

$$u^R(\pi, 0.5) = \frac{1}{2}(\pi v^R(L, L) + \pi v^R(L, R) + (1 - \pi)v^R(R, L) + (1 - \pi)v^R(R, R)) = 0$$

$$u^T(0.5, \pi) = \frac{1}{2}(\pi v^T(L, L) + (1 - \pi)v^T(L, R) + \pi v^T(R, L) + (1 - \pi)v^T(R, R)) = 0$$

Given that the other guy randomizes 1/2-1/2, your decision does not affect your payoff and neither Rob nor Tom can get more than 0.

This suggests a general strategy for finding mixed strategy Nash equilibria: find probabilities for Tom such that Rob is indifferent between his two pure strategies. Find probabilities for Rob such that Tom is indifferent between his two pure strategies. These two will be a mixed strategy equilibrium: if Rob is indifferent between his two pure strategies he might as well play the mixed strategy, the same for Tom.

Mixed strategies in the coordination game

In the coordination game, we found two pure strategy Nash equilibria. As it turns out there is an additional Nash equilibrium in mixed strategies.

Suppose Tom randomizes between hockey and baseball with some probability π_T . At what point is Rob indifferent?

$$u^R(H, \pi_T) = \pi_T \cdot 6 + (1 - \pi_T) \cdot 0 = 6\pi_T$$

$$u^R(B, \pi_T) = \pi_T \cdot 5 + (1 - \pi_T) \cdot 11 = 11 - 6\pi_T$$

$$11 - 6\pi_T = 6\pi_T \Rightarrow \pi_T = \frac{11}{12}$$

Rob randomizes between hockey and baseball with π_R . When is Tom indifferent?

$$u^T(\pi_R, H) = \pi_R \cdot 11 + (1 - \pi_R) \cdot 5 = 5 + 6\pi_R$$

$$u^T(\pi_R, B) = \pi_R \cdot 0 + (1 - \pi_R) \cdot 6 = 6 - 6\pi_R$$

$$5 + 6\pi_R = 6 - 6\pi_R \Rightarrow \pi_R = \frac{1}{12}$$

Therefore a Nash equilibrium in mixed strategies is $\pi_R = 1/12$, $\pi_T = 11/12$.

3 Extensive form games

So far we considered only games in which both parties had to decide simultaneously which strategy to choose. Now we will consider situations in which one player moves first, the other player observes what the first player did and then decides on which action to take.

To capture the sequential structure of the game, we will depict sequential games by using **game trees**.

It is important to clarify what a strategy for a player is in extensive form games. A strategy for a player who moves second will be a contingent plan: for all possible things the first player could have done, the second player needs to specify his optimal action.

3.1 Examples

We will consider two examples and show what game theory predicts about the outcomes.

R and D Problem

The first example is a simple research and development problem. Suppose a small company (player A) has to decide how much to invest in research and development - if it invests a lot it will invent a new method of production. However, a big firm moves (player B) second and has the choice to imitate the small firm. If the small firm gets to use its invention it gets a payoff of 9, the big firm gets 1. If the big firm imitates the small firm, the small firm only gets 1 and the big firm gets 9.

If the small firm does not spend any money on research and development both firms get 3. Figure 1 shows the game tree for this game.

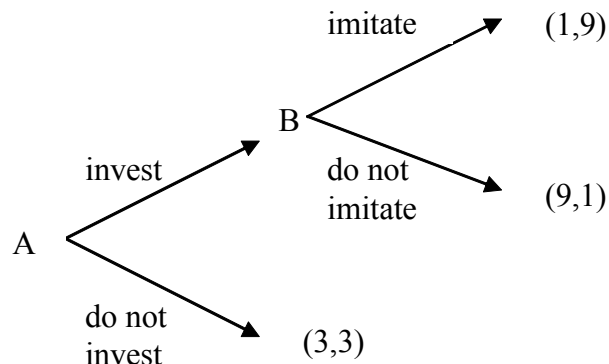



Figure 1:

In order to figure out how Nash-equilibria look like, we want to ask, what are the possible strategies in this game. This part of the notes might seem mind-boggling at first (in fact it still confuses me even now when I teach it) but if you think a little about it, it makes sense. Obviously player A's strategies are $S^A = \{invest, dont\ invest\}$. Naively one would think (and for inexplicable reasons some undergraduate textbooks (who know better) make it sound like this) that Player B's strategies are $S^B = \{imitate, dont\ imitate\}$. However, this is false. Player B knows what player one has done when it is his turn to move. So his actual strategy has to specify what he does in each possible situation - his strategies can differ depending on player A's action.

We will see below why it is important to treat this issue carefully and why this formulation gives us some problems with the concept of Nash equilibrium

Centipede game

The second example is a famous game which shows that there can be a real problem in sequential games. It's called *centipede game* after the looks of its game tree. The basic idea here is that Rob and Tom go to a party together. Rob is there first; if he leaves without waiting for Tom he gets a payoff of one util (say the food which was served at the party). If he waits for Tom, it'll only be fun if Tom stays for a while. If Tom leaves right away, Tom gets a payoff of one (from the food) while Rob gets zero (the waiting got on his nerves and his overall payoff is zero). However, if Tom decides to stay, Rob has to make one final decision - he can stay the rest of the evening in which case he bores Tom to death, he gets 3 and poor Tom gets 0. Or he can leave after he and Tom talked for a while in which case they both get 2. Figure 2 shows the game tree for this game.



3.2 Nash-equilibrium and backward induction

We now want to figure out how to solve these two games.

It makes sense to solve a sequential game by **backward induction**. We start at the end and ask what is the players best move, given that he ever gets to this node of the tree. Then we ask what will the player who moves first do, given that he knows that the other player will react optimally.

Solving R and D by induction

- Player B:
 - If A chooses to invest, player B should choose to imitate.
 - If A chooses not to invest, player B has no choice.
 - Player A:
 - If A chooses to invest, then B will choose to imitate and A will get a payoff of 1.
 - If A chooses not to invest, then A will get 3.
- ⇒ **Nash equilibrium:** A will choose not to invest.

Nash equilibrium in R and D

Note, however, that this game would have had a second Nash equilibrium if player B (the big firm) could threaten player A to fight it if it does not invest. Suppose, for example, that player B can cut prices and make payoffs for both firms go to zero (instead of 3). If we take the strategies as described above, player B could play (fight if no investment, imitate

if investment). Given that it plays this, player A's best response will be to invest. The outcome will be (1, 9) and player B is very happy.

Recall the definition of a Nash equilibrium. Check that these strategies satisfy the definition of a Nash equilibrium.

Intuitively, player B **threatens** to fight even if A chose not to invest. A thinks that the threat is credible and therefore invests.

Note, however, that the threat of B to fight if A chooses not to invest is not credible. Once A has chosen not to invest, B will understand that he hurts himself by choosing to fight and that he would do better by choosing not to fight. Hence, this Nash equilibrium is not convincing.

In order to rule out these types of unconvincing Nash equilibria we require that in a sequential game an equilibrium has to be "subgame perfect".

Definition 1 (Subgame perfect equilibrium) *A Nash equilibrium is **subgame perfect**, if the strategies of all players form a Nash equilibrium not only in the game as a whole, but also in every subgame of the game. That is, after every possible history of the game the strategies of the players have to be mutually best responses.*

The Nash equilibrium [invest, (fight if no investment, imitate if investment)] is not subgame perfect. It turns out (we will not prove this here, because the issue of subgame perfection is a little too advanced for this class) that every equilibrium one obtains by backward induction is subgame perfect. Therefore it makes sense to simply solve sequential games by backward induction without spending too much time thinking about the possible Nash equilibria of the game.

Centipede and induction

We now want to solve the centipede game by backward induction: in the last stage, Rob will certainly stay, since this gives him a payoff of 3, leaving would give him only a payoff of 2. Knowing that Rob will stay at the end, Tom knows that we will get 0 if he stays. If he leaves he will get 1 - therefore, when it is Tom's turn to decide, his optimal decision is to leave. Knowing this, what is Rob going to do? If Rob stays at the very beginning he knows that Tom will leave and he will get 0, therefore he leaves right away, gets 1 and the game is over.

The unique subgame-perfect Nash equilibrium of this game is therefore: Rob leaves right away.

The problem with this outcome is that it is somewhat counterintuitive. True, at the end Rob might stay and therefore Tom does not trust him, but in many experiments people did it was actually observed that people played the cooperative outcome: stay, stay, leave

which gives a payoff of 2 for both players. Also, a more careful examination reveals that there is a slight problem with backward induction here: Suppose Rob stays, when it is Tom's turn, how is he going to reason?

One possibility is that he'll say: If I stay, Rob will stay (because he is a selfish player and does not care about my payoffs) and I will get 0, so I better leave now and I will get 1.

However, if Rob really is a rational (selfish) player, why did he stay in the first place? This was not optimal, given that he knows Tom is rational. So may-be, Rob stayed because he is not the rational selfish guy Tom thought he is and may-be he will not leave if Tom stays.

3.3 Chess

As a last example for extensive form games, suppose Rob and Tom play chess. The final payoffs are $(1, 0)$ if Rob wins, $(0, 1)$ if Tom wins and $(0.5, 0.5)$ if it is a draw. Drawing the game tree for this game is possible in principle but since there are billions of possible sequences of moves I'll pass on this.

A strategy for a player is a full contingent plan, given any possible sequence of previous moves it must specify what the player should do in this situation.

Since the game will end after a finite number of moves for sure (recall that after three times of the same thing, the game ends in a draw), we could in principle solve it by backward induction. Note that there must be a strategy for one of the players which will ensure that he either always wins or always draws - we just don't know it, but it must be true that in principle either black or white always wins or chess always ends in a draw.

4 Application: Duopoly

Suppose Rob and Tom both produce bananas at zero costs (you can do the whole thing with positive costs – this makes everything slightly more complicated, but in my view it does not add any additional insights) and face an aggregate demand curve for bananas $D(p) = 10 - p$. Rob and Tom are the only producers of bananas, so they realize that in equilibrium they can control the price of bananas.

There are 4 different things that can happen:

- *Collusion*: Rob and Tom become buddies and try to maximize their joint profit.
- *Cournot*: Rob and Tom independently decide how much bananas to supply
- *Stackelberg*: First Rob decides how many bananas to supply, then Tom who lives closer to the market sees this and makes his decision

- *Bertrand*: Rob and Toms set different prices for their bananas.

You should remember from Econ 50 how to solve for the collusion outcome. If they maximize joint profits, they will behave like a monopoly, i.e. they set a quantity q to maximize profits $\pi(q) = q(10 - q)$. Therefore, we get $10 - 2q = 0$ or $q = 5$. They jointly produce 5 bananas, get 25 dollars of profit and split them 50-50, so they each end up with 12.50.

I now want to discuss the other three situations (which are ‘non-cooperative’) and show how to model them as games and how to solve for the Nash equilibria.

4.1 Cournot’s story

Augustin Cournot was a French economist who first studied oligopoly. In our modern formulation, his idea is as follows. Rob’s strategy is to pick some amount of bananas (up to 10) to supply, i.e. $s^R \in [0, 10]$ and Tom picks how many bananas to supply $s^T \in [0, 10]$. The equilibrium price will satisfy

$$10 - p = s^R + s^T \Rightarrow p = 10 - s^R - s^T$$

and Rob’s and Tom’s payoffs will therefore be

$$v^R(s^R, s^T) = s^R \cdot (10 - s^R - s^T)$$

$$v^T(s^R, s^T) = s^T \cdot (10 - s^R - s^T)$$

Obviously this is a static game. What is a Nash-equilibrium for this game? The optimal choice for Rob must satisfy:

$$\frac{\partial v^R(s^R, s^T)}{\partial s^R} = 10 - s^T - 2s^R = 0$$

Same for Jerry:

$$\frac{\partial v^T(s^R, s^T)}{\partial s^T} = 10 - s^R - 2s^T = 0$$

Since $s^R = s^T$ we get that $s^R = s^T = 10/3$. So in the Nash equilibrium they will each supply 10/3 - total supply will be 6.67 and total profits will be $20/30 \cdot (10 - 20/3) = 22.22$, less than under the collusion outcome.

4.2 Stackelberg’s story

Stackelberg included some time-component into Rob and Tom’s banana business. We assume that Rob chooses first how many bananas to supply; he then walks by Tom’s

farm with his bananas, Tom observes this and adjusts optimally. Since the players do not move at the same time, we need to model the situation as an extensive form game: Rob moves first and decides to supply $s^R \in [0, 10]$ bananas. Then it is Tom's turn to supply some $s^T \in [0, 10]$ (it is difficult to draw a game-tree because there are infinitely many strategies).

We solve this game by backward induction: Given that Rob has chosen some fixed \bar{s}^R Tom solves:

$$\frac{\partial v^T(\bar{s}^R, s^T)}{\partial s^T} = 10 - \bar{s}^R - 2s^T = 0$$

For any choice for Rob, Tom's optimal solution satisfies

$$s^T = 5 - 1/2\bar{s}^R$$

When it is Rob's turn to decide he knows that Tom will always choose $s^T = 5 - 1/2\bar{s}^R$ and he will take this into account when picking s^R . Therefore Rob maximizes

$$s^R(10 - (s^R + s^T)) = s^R(10 - (s^R + 5 - 1/2s^R))$$

His first order conditions imply

$$10 - s^R - 5 = 0 \Rightarrow s^R = 5$$

Therefore Rob just goes out with 5 bananas, Tom sees this and has no choice but to only supply 2.5 bananas.

So in this situation, the total supply of bananas is 7.5 (even larger than in the Cournot case) but the profits are not split equally. Rob has a *first mover advantage*.

4.3 Bertrand

Bertrand assumed that Rob and Tom both bring huge amounts of bananas to the market but that they do not necessarily have to sell them all. Instead they set some price and sell whatever amount of bananas they can at that price. So Rob's strategy now is to set some $p^R \in [0, 10]$ and Tom's strategy is to set some $p^T \in [0, 10]$. The buyers obviously only buy at the lower price, so payoffs now are

$$v^R(p^R, p^T) = \begin{cases} p^R \cdot (10 - p^R) & \text{if } p^R < p^T \\ 1/2p^R \cdot (10 - p^R) & \text{if } p^R = p^T \\ 0 & \text{otherwise} \end{cases}$$

$$v^T(p^R, p^T) = \begin{cases} p^T \cdot (10 - p^T) & \text{if } p^T < p^R \\ 1/2p^T \cdot (10 - p^T) & \text{if } p^R = p^T \\ 0 & \text{otherwise} \end{cases}$$

One can verify that the unique Nash equilibrium is $(p^R, p^T) = (0, 0)$. Whenever $p^T > 0$ Ben can increase his payoff by setting $p^R = (1 - \epsilon)p^T$ for some small $\epsilon > 0$. However, the same is true for Tom; he can increase his payoff by setting $p^T = (1 - \epsilon)p^R$ for some small $\epsilon > 0$. Only if $p^R = p^T = 0$ none of them can do better.