Chapter 5

Regret Minimization and Security Strategies

Until now we implicitly adopted a view that a Nash equilibrium is a desirable outcome of a strategic game. In this chapter we consider two alternative views that help us to understand reasoning of players who either want to avoid costly 'mistakes' or 'fear' a bad outcome. Both concepts can be rigorously formalized.

5.1   Regret minimization

Consider the following game:

	L	R
T	$100, 100$	$0, \quad 0$
B	$0, \quad 0$	$1, \quad 1$

This is an example of a coordination problem, in which there are two satisfactory outcomes (read Nash equilibria), (T, L) and (B, R), of which one is obviously better for both players. In this game no strategy strictly or weakly dominates the other and each strategy is a best response to some other strategy. So using the concepts we introduced so far we cannot explain how come that rational players would end up choosing the Nash equilibrium (T, L). In this section we explain how this choice can be justified using the concept of **regret minimization**.

With each finite strategic game $G := (S_1, \ldots, S_n, p_1, \ldots, p_n)$ we first associate a **regret-recording game** $G := (S_1, \ldots, S_n, r_1, \ldots, r_n)$ in which each payoff function r_i is defined by

$$r_i(s_i, s_{-i}) := p_i(s_i^*, s_{-i}) - p_i(s_i, s_{-i}),$$

where s_i^* is player's i best response to s_{-i}. We call then $r_i(s_i, s_{-i})$ player's i **regret of choosing** s_i **against** s_{-i}. Note that by definition for all s we have $r_i(s) \geq 0$.

For example, for the above game the corresponding regret-recording game is

	L	R
T	0, 0	1, 100
B	100, 1	0, 0

Indeed, $r_1(B, L) := p_1(T, L) - p_1(B, L) = 100$, and similarly for the other seven entries.

Let now

$$regret_i(s_i) := \max_{s_{-i} \in S_{-i}} r_i(s_i, s_{-i}).$$

So $regret_i(s_i)$ is the maximal regret player i can have from choosing s_i. We call then any strategy s_i^* for which the function $regret_i$ attains the minimum, i.e., one such that $regret_i(s_i^*) = \min_{s_i \in S_i} regret_i(s_i)$, a **regret minimization strategy** for player i.

In other words, s_i^* is a regret minimization strategy for player i if

$$\max_{s_{-i} \in S_{-i}} r_i(s_i^*, s_{-i}) = \min_{s_i \in S_i} \max_{s_{-i} \in S_{-i}} r_i(s_i, s_{-i}).$$

The following intuition is helpful here. Suppose the opponents of player i are able to perfectly anticipate which strategy player i is about to play (for example by being informed through a third party what strategy player i has just selected and is about to play). Suppose further that they aim at inflicting at player i the maximum damage in the form of maximal regret and that player i is aware of these circumstances. Then to miminize his regret player i should select a regret minimization strategy. We could say that a regret minimization strategy will be chosen by a player who wants to avoid making a costly 'mistake', where by a mistake we mean a choice of a strategy that is not a best response to the joint strategy of the opponents.

To clarify this notion let us return to our example of the coordination game. To visualize the outcomes of the functions $regret_1$ and $regret_2$ we put the results in an additional row and column:

	L	R	$regret_1$
T	0, 0	1, 100	1
B	100, 1	0, 0	100
$regret_2$		1	100	

So T is the minimum of $regret_1$ and L is the minimum of $regret_2$. Hence (T, L) is the unique pair of regret minimization strategies. This shows that using the concept of regret minimization we succeeded to single out the preferred Nash equilibrium in the considered coordination game.

It is important to note that the concept of regret minimization does not allow us to solve all coordination problems. For example, it does not help us in selecting a Nash equilibrium in symmetric situations, for instance in the game

	L	R
T	1, 1	0, 0
B	0, 0	1, 1

Indeed, in this case the regret of each strategy is 1, so regret minimization does not allow us to distinguish between the strategies. Analogous considerations hold for the Battle of Sexes game from Chapter 1.

Regret minimization is based on different intuitions than strict and weak dominance. As a result these notions are incomparable. In general, only the following limited observation holds. Recall that the notion of a dominant strategy was introduced in Exercise 9 on page 34.

Note 14 (Regret Minimization) *Consider a finite game. Every dominant strategy is a regret minimization strategy.*

Proof. Fix a finite game $(S_1, \ldots, S_n, p_1, \ldots, p_n)$. Note that each dominant strategy of player i is a best response to each $s_{-i} \in S_{-i}$. So by the definition of the regret-recording game for all $s_{-i} \in S_{-i}$ we have $r_i(s_i, s_{-i}) = 0$. This shows that s_i is a regret minimization strategy for player i, since for all joint strategies s we have $r_i(s) \geq 0$. $\qquad\square$

The process of removing strategies that do not achieve regret minimization can be iterated. We call this process the ***iterated regret minimization***. The example of the coordination game we analyzed shows that the process of regret minimization may yield to a loss of some Nash equilibria. In fact, as we shall see in a moment, during this process all Nash equilibria can be lost. On the other hand, as recently suggested by J. Halpern and R. Pass, in some games the iterated regret minimization yields a more intuitive outcome. As an example let us return to the Traveller's Dilemma game considered in Example 1.

Example 12 (Traveller's dilemma revisited)
Let us first determine in this game the regret minimization strategies for each player. Take a joint strategy s.

Case 1. $s_{-i} = 2$.
Then player's i regret of choosing s_i against s_{-i} is 0 if $s_i = s_{-i}$ and 2 if $s_i > s_{-i}$, so it is at most 2.

Case 2. $s_{-i} > 2$.
If $s_{-i} < s_i$, then $p_i(s) = s_{-i} - 2$, while the best response to s_{-i}, namely $s_{-i} - 1$, yields the payoff $s_{-i} + 1$. So player's i regret of choosing s_i against s_{-i} is in this case 3.

If $s_{-i} = s_i$, then $p_i(s) = s_{-i}$, while the best response to s_{-i}, namely $s_{-i} - 1$, yields the payoff $s_{-i} + 1$. So player's i regret of choosing s_i against s_{-i} is in this case 1.

Finally, if $s_{-i} > s_i$, then $p_i(s) = s_i + 2$, while the best response to s_{-i}, namely $s_{-i} - 1$, yields the payoff $s_{-i} + 1$. So player's i regret of choosing s_i against s_{-i} is in this case $s_{-i} - s_i - 1$.

To summarize, we have

$$regret_i(s_i) = \max(3, \max_{s_{-i} \in S_{-i}} s_{-i} - s_i - 1) = \max(3, 99 - s_i).$$

So the minimal regret is achieved when $99 - s_i \leq 3$, i.e., when the strategy s_i is in the interval $[96, 100]$. Hence removing all strategies that do not achieve regret minimization yields a game in which each player has the strategies in the interval $[96, 100]$. In particular, we 'lost' in this way the unique Nash equilibrium of this game, (2,2).

We now repeat this elimination procedure. To compute the outcome we consider again two, though now different, cases.

Case 1. $s_i = 97$.

The following table then summarizes player's i regret of choosing s_i against a strategy s_{-i} of player i:

strategy of player $-i$	best response of player i	regret of player i
96	96	2
97	96	1
98	97	0
99	98	1
100	99	2

Case 2. $s_i \neq 97$.

The following table then summarizes player's i regret of choosing s_i, where for each strategy of player i we list a strategy of player $-i$ for which player's i regret is maximal:

strategy of player i	relevant strategy of player $-i$	regret of player i
96	100	3
98	97	3
99	98	3
100	99	3

So each strategy of player i different from 97 has regret 3, while 97 has regret 2. This means that the second round of elimination of the strategies that do not achieve regret minimization yields a game in which each player has just one strategy, namely 97. □

Recall again that the unique Nash equilibrium in the Traveller's Dilemma game is (2,2). So the iterated regret minimization yields here a radically different outcome than the analysis based on Nash equilibria. Interestingly, this outcome, (97,97), has been confirmed by empirical studies.

We conclude this section by showing that iterated regret minimization is not order independent. To this end consider the following game:

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

The corresponding regret-recording game, together with the recording of the outcomes of the functions $regret_1$ and $regret_2$ is as follows:

	L	R	$regret_1$
T	0, 2	1, 0	1
B	2, 0	0, 1	2
$regret_2$	2	1	

This shows that (T, R) is the unique pair of regret minimization strategies in the original game. So by removing from the original game the strategies B and L that do not achieve regret minimization we reduce it to

	R
T	0, 3

On the other hand, if we initially only remove strategy L, then we obtain the game

	R
T	0, 3
B	1, 1

Now the only strategy that does not achieve regret minimization is T. By removing it we obtain the game

	R
B	1, 1

5.2   Security strategies

Consider the following game:

	L	R
T	0, 0	101, 1
B	1, 101	100, 100

This is an extreme form of a **Chicken game**, sometimes also called a **Hawk-Dove game** or a **Snowdrift game**.

The game of Chicken models two drivers driving at each other on a narrow road. If neither driver swerves ('chickens'), the result is a crash. The best option for each driver is to stay straight while the other swerves. This yields a situation where each driver, in attempting to realize his the best outcome, risks a crash.

The description of this game as a snowdrift game stresses advantages of a cooperation. The game involves two drivers who are trapped on opposite sides of a snowdrift. Each has the option of staying in the car or shoveling snow to clear a path. Letting the other driver do all the work is the best option, but being exploited by shoveling while the other driver sits in the car is still better than doing nothing.

Note that this game has two Nash equilibria, (T, R) and (B, L). However, there seems to be no reason in selecting any Nash equilibrium as each Nash equilibrium is grossly unfair to the player who will receive only 1.

In contrast, (B, R), which is not a Nash equilibrium, looks like a most reasonable outcome. Each player receives in it a payoff close to the one he receives in the Nash equilibrium of his preference. Also, why should a player risk the payoff 0 in his attempt to secure the payoff 101 that is only a fraction bigger than his payoff 100 in (B, R)?

Note that in this game no strategy strictly or weakly dominates the other and each strategy is a best response to some other strategy. So these concepts are useless in analyzing this game. Moreover, the regret minimization for each strategy is 1. So this concept is of no use here either.

We now introduce the concept of a **security strategy** that allows us to single out the joint strategy (B, R) as the most reasonable outcome for both players.

Fix a, not necessarily finite, strategic game $G := (S_1, \ldots, S_n, p_1, \ldots, p_n)$. Player i, when considering which strategy s_i to select, has to take into account which strategies his opponents will choose. A 'worst case scenario' for player i is that, given his choice of s_i, his opponents choose a joint strategy for which player's i payoff is the lowest[1]. For each strategy s_i of player i once this lowest payoff can be identified a strategy can be selected that leads to a 'minimum damage'.

---

[1]We assume here that such s_i exists.

To formalize this concept for each $i \in \{1, \ldots, n\}$ we consider the function[2]

$$f_i : S_i \to \mathcal{R}$$

defined by

$$f_i(s_i) := \min_{s_{-i} \in S_{-i}} p_i(s_i, s_{-i}).$$

We call any strategy s_i^* for which the function f_i attains the maximum, i.e., one such that $f_i(s_i^*) = \max_{s_i \in S_i} f_i(s_i)$, a **security strategy** or a **maxminimizer** for player i. We denote this maximum, so

$$\max_{s_i \in S_i} \min_{s_{-i} \in S_{-i}} p_i(s_i, s_{-i}),$$

by $maxmin_i$ and call it the **security payoff** of player i.

In other words, s_i^* is a security strategy for player i if

$$\min_{s_{-i} \in S_{-i}} p_i(s_i^*, s_{-i}) = maxmin_i.$$

Note that $f_i(s_i)$ is the minimum payoff player i is guaranteed to secure for himself when he selects strategy s_i. In turn, the security payoff $maxmin_i$ of player i is the minimum payoff he is guaranteed to secure for himself in general. To achieve at least this payoff he just needs to select any security strategy.

The following intuition is helpful here. Suppose the opponents of player i are able to perfectly anticipate which strategy player i is about to play. Suppose further that they aim at inflicting at player i the maximum damage (in the form of the lowest payoff) and that player i is aware of these circumstances. Then player i should select a strategy that causes the minimum damage for him. Such a strategy is exactly a security strategy and it guarantees him at least the $maxmin_i$ payoff. We could say that a security strategy will be chosen by a 'pessimist' player, i.e., one who fears the worst outcome for himself.

To clarify this notion let us return to our example of the chicken game. Clearly, both B and R are the only security strategies in this game. Indeed, we have $f_1(T) = f_2(L) = 0$ and $f_1(B) = f_2(R) = 1$. So we succeeded to

---

[2]In what follows we assume that all considered minima and maxima always exist. This assumption is obviously satisfied in finite games. In a later chapter we shall discuss a natural class of infinite games for which this assumption is satisfied, as well.

single out in this game the outcome (B, R) using the concept of a security strategy.

The following counterpart of the Regret Minimization Note 14 holds.

Note 15 (Security) *Consider a finite game. Every dominant strategy is a security strategy.*

Proof. Fix a game $(S_1, \ldots, S_n, p_1, \ldots, p_n)$ and suppose that s_i^* is a dominant strategy of player i. For all joint strategies s

$$p_i(s_i^*, s_{-i}) \geq p_i(s_i, s_{-i}),$$

so for all strategies s_i of player i

$$\min_{s_{-i} \in S_{-i}} p_i(s_i^*, s_{-i}) \geq \min_{s_{-i} \in S_{-i}} p_i(s_i, s_{-i}).$$

Hence

$$\min_{s_{-i} \in S_{-i}} p_i(s_i^*, s_{-i}) \geq \max_{s_i \in S_i} \min_{s_{-i} \in S_{-i}} p_i(s_i, s_{-i}).$$

This concludes the proof. $\qquad\square$

Next, we introduce a dual notion to the security payoff $maxmin_i$. It is not needed for the analysis of security strategies but it will turn out to be relevant in a later chapter.

With each $i \in \{1, \ldots, n\}$ we consider the function

$$F_i : S_{-i} \to \mathcal{R}$$

defined by

$$F_i(s_{-i}) := \max_{s_i \in S_i} p_i(s_i, s_{-i}).$$

Then we denote the value $\min_{s_{-i} \in S_{-i}} F_i(s_{-i})$, i.e.,

$$\min_{s_{-i} \in S_{-i}} \max_{s_i \in S_i} p_i(s_i, s_{-i}),$$

by $minmax_i$.

The following intuition is helpful here. Suppose that now player i is able to perfectly anticipate which strategies his opponents are about to play. Using this information player i can compute the minimum payoff he is guaranteed to achieve in such circumstances: it is $minmax_i$. This lowest payoff for player

i can be enforced by his opponents if they choose any joint strategy s^*_{-i} for which the function F_i attains the minimum, i.e., one such that $F_i(s^*_{-i}) = \min_{s_{-i} \in S_{-i}} F_i(s_{-i})$.

To clarify the notions of $maxmin_i$ and $minmax_i$ consider an example.

Example 13 Consider the following two-player game:

	L	M	R
T	3, −	4, −	5, −
B	6, −	2, −	1, −

where we omit the payoffs of the second, i.e., column, player.

To visualize the outcomes of the functions f_1 and F_1 we put the results in an additional row and column:

	L	M	R	f_1
T	3, −	4, −	5, −	3
B	6, −	2, −	1, −	1
F_1	6	4	5	

That is, in the f_1 column we list for each row its minimum and in the F_1 row we list for each column its maximum.

Since $f_1(T) = 3$ and $f_1(B) = 1$ we conclude that $maxmin_1 = 3$. So the security payoff of the row player is 3 and T is a unique security strategy of the row player. In other words, the row player can secure for himself at least the payment 3 and achieves this by choosing strategy T.

Next, since $F_1(L) = 6, F_1(M) = 4$ and $F_1(R) = 5$ we get $minmax_1 = 4$. In other words, if the row player knows which strategy the column player is to play, he can secure for himself at least the payment 4.

Indeed,

- if the row player knows that the column player is to play L, then he should play B (and secure the payoff 6),

- if the row player knows that the column player is to play M, then he should play T (and secure the payoff 4),

- if the row player knows that the column player is to play R, then he should play T (and secure the payoff 5). $\qquad\square$

In the above example $maxmin_1 < minmax_1$. In general the following observation holds. From now on, to simplify the notation we assume that s_i and s_{-i} range over, respectively, S_i and S_{-i}.

Lemma 16 (Lower Bound)

(i) For all $i \in \{1, \ldots, n\}$ we have $maxmin_i \leq minmax_i$.

(ii) If s is a Nash equilibrium of G, then for all $i \in \{1, \ldots, n\}$ we have $minmax_i \leq p_i(s)$.

Item (i) formalizes the intuition that one can take a better decision when more information is available (in this case about which strategies the opponents are about to play). Item (ii) provides a lower bound on the payoff in each Nash equilibrium, which explains the name of the lemma.

Proof.
(i) Fix i. Let s_i^* be such that $\min_{s_{-i}} p_i(s_i^*, s_{-i}) = maxmin_i$ and s_{-i}^* such that $\max_{s_i} p_i(s_i, s_{-i}^*) = minmax_i$. We have then the following string of equalities and inequalities:

$$maxmin_i = min_{s_{-i}} p_i(s_i^*, s_{-i}) \leq p_i(s_i^*, s_{-i}^*) \leq max_{s_i} p_i(s_i, s_{-i}^*) = minmax_i.$$

(ii) Fix i. For each Nash equilibrium (s_i^*, s_{-i}^*) of G we have

$$min_{s_{-i}} max_{s_i} p_i(s_i, s_{-i}) \leq max_{s_i} p_i(s_i, s_{-i}^*) = p_i(s_i^*, s_{-i}^*).$$

$\square$

To clarify the difference between the regret minimization and security strategies consider the following variant of a coordination game:

	L	R
T	$100, 100$	$0, \quad 0$
B	$1, \quad 1$	$2, \quad 2$

It is easy to check that players who select the regret minimization strategies will choose the strategies T and L which yields the payoff 100 to each of them. In contrast, players who select the security strategies will choose B and L and will receive only 1 each.

Next, consider the following game:

	L	M	R
T	5, 5	0, 0	97, 1
B	1, 0	1, 0	100, 100

Here the security strategies are B and R and their choice by the players yields the payoff 100 to each of them. In contrast, the regret minimization strategies are T (with the regret 3) and R (with the regret 4) and their choice by the players yields them the respective payoffs 97 and 1.

So the outcomes of selecting regret minimization strategies and of security strategies are incomparable.

Finally, note that in general there is no relation between the equalities of the $maxmin_i = minmax_i$ and an existence of a Nash equilibrium. To see this let us fill in the game considered in Example 13 the payoffs for the column player as follows:

	L	M	R
T	3, 1	4, 0	5, 1
B	6, 1	2, 0	1, 1

We already noted that $maxmin_1 < minmax_1$ holds here. However, this game has two Nash equilibria, (T, R) and (B, L).

Further, the following game

	L	M	R
T	3, 1	3, 0	5, 0
B	6, 0	2, 1	1, 0

has no Nash equilibrium and yet for $i = 1, 2$ we have $maxmin_i = minmax_i$.

In a later chapter we shall discuss a class of two-player games for which there is a close relation between the existence of a Nash equilibrium and the equalities $maxmin_i = minmax_i$.