Game Theory Principles III

Andreas Blume
Department of Economics
University of Pittsburgh
Pittsburgh, PA 15260

Week 3: GAME THEORY IN LITERATURE: Poe, Conan Doyle, Faulkner, Philip K. Dick, (see Brams in GEB 6), games with pure conflict, the minmax theorem. Concepts: best responses, mixed strategies. Tools: basic probability, expected value, expected utility.

In *The Adventure of the Final Problem* (1893, easy to find online) Arthur Conan Doyle describes the pursuit of Sherlock Holmes by his archenemy Moriarty.

A crucial scene from the story:

1. Holmes has just boarded a train from Victoria station to Dover.
2. As the train begins moving, he spots Moriarty on the platform.
3. Holmes expects Moriarty to hire a "special" that will arrive in Dover before Holmes's train.
4. Holmes can expect to be killed by Moriarty upon his arrival in Dover.
5. Holmes decides to get off the train at Canterbury, an intermediate stop.

Both players are described as equally rational. Having seen Moriarty on the platform, Watson expresses confidence at having shaken off Moriarty. Holmes replies:

"My dear Mr Watson, you evidently did not realize my meaning when I said that this man may be taken as being quite on the same intellectual plane as myself. You do not imagine that if I were the pursuer I should allow myself to be baffled by so slight an obstacle."

"What will he do?"

"What I should do."

"What would you do then?"

"Engage a special."

"But it must be late."

"By no means. This train stops at Canterbury; and there is always at least a quarter of an hour's delay at the boat. He will catch us there."

Note that Holmes expects Moriarty to be aware of the fact that Holmes's train makes a stop at Canterbury.

Earlier in the story Holmes refers to Moriarty as having a *"phenomenal mathematical faculty."*

He describes him as *"... a genius, a philosopher, an abstract thinker. He has a brain of the first order."*

While giving instructions to Watson for meeting at Victoria station Holmes urges Watson to *"... obey them to the letter, for you are now playing a double-handed game with me against the cleverest rogue and the most powerful syndicate of criminals in Europe."*

According to the story,

1. Holmes expects Moriarty to be just as smart as Holmes himself.
2. Holmes "puts himself in Moriarty's shoes," thinking how he would act if he faced Moriarty's problem.
3. Concluding that he would hire a special, he expects Moriarty to do the same.
4. Then he reasons one step further and decides to get off at Canterbury.

At this point a natural question arises:

1. If Moriarty is indeed just as smart at Holmes, why can't he reason just as far as Holmes does?
2. Accordingly, Moriarty should expect Holmes to get off at Canterbury.
3. Therefore Holmes should expect Moriarty to stop the special at Canterbury.
4. Therefore Moriarty should expect Holmes to go on to Dover.
5. ... and so on ...

This "endless chain of reciprocally conjectural reactions and counter-reactions" was pointed out by Oskar Morgenstern (one of the founders of game theory) in his 1928-book *Wirtschaftsprognose, Eine Untersuchung ihrer Voraussetzungen und Möglichkeiten (Economic Prediction: An Examination of its Conditions and Possibilities)*.

In his 1935-paper *Vollkommene Voraussicht und Wirtschaftliches Gleichgewicht* (translated as *Perfect Foresight and Economic Equilibrium* in Selected Writings of Oskar Morgenstern (Andrew Schotter, Ed.), New York: New York University Press) Morgenstern summarizes the conundrum as follows:

Sherlock Holmes, pursued by his opponent, Moriarty, leaves London for Dover. The train stops at a station on the way, and he alights there rather than traveling on to Dover. He has seen Moriarty at the railway station, recognizes that he is very clever and expects that Moriarty will take a faster special train in order to catch him in Dover. Holmes anticipation turns out to be correct. But what if Moriarty had been still more clever and had foreseen his actions accordingly? Then, obviously, he would have traveled to the intermediate station. Holmes, again, would have had to calculate that and he himself would have decided to go on to Dover. Whereupon, Moriarty would again have reacted differently.

Because of so much thinking they might not have been able to act at all or the intellectually weaker of the two would have surrendered in the Victoria Station, since the whole flight would have become unnecessary.

In essence, Morgenstern presents the decision problem of Holmes as not having a solution.

In their 1944-book *Theory of Games and Economic Behavior*, John von Neumann and Oskar Morgenstern revisit the Holmes and Moriarty problem and formulate it as a game as follows

		Holmes	
		C'bury	Dover
Moriarty	Canterbury	(100,-100)	(-50,50)
	Dover	(0,0)	(100,-100)

Game III-1

The infinite regress that Morgenstern describes in his 1928 book, his "endless chain of reciprocally conjectural reactions and counter-reactions," can be reconstructed in Game III-1 as follows:

1. If Holmes expects “Dover” from Moriarty, his own best response is “Canterbury,” for a payoff of 0 instead of -100.
2. If Moriarty expects “Canterbury” from Holmes, his own best response is “Canterbury,” for a payoff of 100 instead of 0.
3. If Holmes expects “Canterbury” from Moriarty, his own best response is “Dover,” for a payoff of 50 instead of -100.
4. If Moriarty expects “Dover” from Holmes, his own best response is “Dover,” for a payoff of 100 instead of -50.

Thus, best responses cycle. None of the four strategy pairs (Dover, Canterbury), (Dover, Dover), (Canterbury, Dover) and (Canterbury, Canterbury) are mutual best responses.

Two Observations on the Moriarty-Holmes Game

1. This is a zero-sum game, i.e. one player's gain is another player's loss.
2. There is no Nash equilibrium among the strategy pairs (Dover, Canterbury), (Dover, Dover), (Canterbury, Dover) and (Canterbury, Canterbury).

Before we deal with the second observation, we examine zero-sum games more closely.

Zero-Sum Games

1. zero-sum games are games of pure conflict
2. adding a constant to all payoffs of a player does not change the incentive structure of the game: therefore constant-sum games are equivalent to zero-sum games

Iran and the West

Iran seeks to go forward with work on uranium enrichment.

The US and Europe are worried that having a uranium enrichment facility will eventually enable Iran to build a nuclear bomb.

A coarse description of their respective strategic options might be as follows:

1. The West can **ignore** the problem, engage in **diplomacy**, or **act** militarily.
2. Iran can **give up** their nuclear ambition, **slow**ly develop its enrichment capabilities, or opt for **rapid** development.

Some thoughts about payoffs for the West-Iran Game:

1. We will choose to model the conflict as a zero-sum game. A note of caution: This deliberately ignores that there are numerous dimensions of common interest between Iran and the West.
2. It suffices to keep track of the payoff of only one of players (we'll choose the West); the other player's payoffs are the same, except with the opposite sign.
3. If Iran decides to give up its nuclear ambition, then ignoring what they do is a valid option for the West, which will be indicated by a positive payoff, 5, for the West. Recall that Iran's payoff in this case must be -5.

4. Ignoring Iran's ambitions is less attractive if Iran chooses slow, indicated by a negative payoff, -2, for the West.

5. The worst outcome for the West is to ignore the problem when Iran proceeds rapidly, indicated by a payoff of -6.

6. The best outcome for the West with diplomacy would be if Iran gave up its ambition. The worst outcome in this case would be if Iran kept slowly developing its enrichment capabilities. In the case that Iran proceeded rapidly, it would then be possible to convince others, e.g. Russia and China, that diplomacy has been exhausted.

7. Military action, in the game we will set up, is the preferred option only in the event that Iran proceeds rapidly.

Payoff Matrix for West-Iran Game

Iran

West		give up	slow	rapid
	ignore	5	-2	-6
West	diplomacy	4	-1	1
	act	-5	-4	2

Game III-2

Consider the West's decision problem first:

Unlike before, assume that the West is afraid that Iran will always discover its strategy choice and that, knowing the West's strategy, Iran will always try to hold the West's payoff as low as possible.

One may ask: What is the best strategy against an omniscient opponent who is trying to minimize one's payoff?

1. Against an omniscient adversary with diametrically opposed interests, the West's payoff from ignore is -6.
2. Against an omniscient adversary with diametrically opposed interests, the West's payoff from diplomacy is -1.
3. Against an omniscient adversary with diametrically opposed interests, the West's payoff from act is -5.
4. In this scenario there is an **unambiguously best strategy for the West: Diplomacy**.

Observe that the payoff of -1 is the highest payoff that the West can guarantee for itself in this game.

The West secures at least the payoff -1 by adopting the strategy diplomacy.

Assume that Iran is equally afraid that the West will have advance knowledge of Iran's strategy and will limit Iran's payoff as much as possible.

We can think of Iran as trying to minimize the West's payoff. Thus, in our payoff matrix, Iran aims for low values.

One may ask: What is the best strategy for Iran against an omniscient West who is trying to limit Iran's payoff as much as possible?

In this world, Iran will try to **minimize** the **maximum** payoff that the West can achieve against any of Iran's strategies.

The West's highest payoff against Iran's strategy give up is 5.

Iran

		give up	slow	rapid
West	ignore	5	-2	-6
	diplomacy	4	-1	1
	act	-5	-4	2

Game III-2

The West's highest payoff against Iran's strategy slow is -1.

		Iran		
		give up	slow	rapid
	ignore	5	-2	-6
West	diplomacy	4	-1	1
	act	-5	-4	2

GAME III-2

The West's highest payoff against Iran's strategy rapid is 2.

Iran

West		give up	slow	rapid
	ignore	5	-2	-6
West	diplomacy	4	-1	1
	act	-5	-4	2

GAME III-2

Recall our assumption that Iran's goal is to minimize the maximum payoff the West can achieve against Iran's strategy.

This goal is achieved by adopting the strategy **slow**.

Iran's strategy slow guarantees that the West can achieve no higher payoff than -1.

Observe that both the West and Iran can guarantee the payoff -1 in this game.

The West can guarantee that its payoff will be at least -1.

Iran can guarantee that the West's payoff is no higher than -1.

In a two-person zero-sum game, a payoff with this property is referred to as the **Minimax Value** of the game

A higher level of abstraction

If player 1 adopts action a_1 and player 2 adopts action a_2, then we can write player 1's payoff as a function of these actions:

$$\pi_1(a_1, a_2)$$

In a zero sum game, player 2's payoff $\pi_2(a_1, a_2)$ satisfies the condition:

$$\pi_2(a_1, a_2) = -\pi_1(a_1, a_2)$$

The highest payoff player 1 can guarantee for herself is

$$\max_{a_1} \min_{a_2} \pi_1(a_1, a_2)$$

The highest payoff player 2 can guarantee for himself is

$$\max_{a_2} \min_{a_1} \pi_2(a_1, a_2)$$

The highest payoff player 2 can guarantee for himself also equals

$$\max_{a_2} \min_{a_1} \{-\pi_1(a_1, a_2)\},$$

which itself equals

$$-\min_{a_2}\max_{a_1}\ \pi_1(a_1,a_2).$$

A strategy α_1 that guarantees player 1 at least the payoff

$$\max_{a_1} \min_{a_2} \pi_1(a_1, a_2)$$

is called her **maxmin strategy**

A strategy α_2 that guarantees player 2 at least the payoff

$$-\min_{a_2} \max_{a_1} \ \pi_1(a_1, a_2)$$

is called his **minmax strategy**

In the West-Iran game, the West's maxmin strategy is diplomacy.

In the West-Iran game, Iran's minmax strategy is slow.

Observe that in the West-Iran game

$$\max_{a_1} \min_{a_2} \pi_1(a_1, a_2) = \min_{a_2} \max_{a_1} \pi_1(a_1, a_2).$$

Whenever the condition

$$\max_{a_1} \min_{a_2} \pi_1(a_1, a_2) = \min_{a_2} \max_{a_1} \pi_1(a_1, a_2)$$

holds in a two-person zero-sum game, we call the quantity

$$v \equiv \max_{a_1} \min_{a_2} \pi_1(a_1, a_2) = \min_{a_2} \max_{a_1} \pi_1(a_1, a_2)$$

the (minimax) value of the game.

We saw that the West-Iran game has a minimax value, which equals -1.

A player in a two-person game will want to adopt her maxmin strategy (his minmax strategy) if she (he) is paranoid, or has reason to believe that he (she) faces an omniscient adversarial opponent.

		Holmes: C'bury	Holmes: Dover
Moriarty	Canterbury	(100,-100)	(-50,50)
	Dover	(0,0)	(100,-100)

GAME III-1

What should Sherlock Holmes do if he believes Moriarty to be omniscient?

It appears that against an omniscient Moriarty, Holmes will always end up with a payoff of -100.

A suggested way out of Holmes's dilemma: If Holmes himself does not know which action he will choose, Moriarty cannot take advantage of knowing Holmes's planned action.

Ignorance of one's own action is achieved via

randomization,

e.g. by having one's action be determined by a coin flip.

Randomization works because it produces unpredictability: It makes it impossible for an opponent to predict one's action with perfect accuracy.

Suppose Holmes determines his action via flipping a fair coin.

Then

1. If Moriarty chooses Canterbury, Holmes will receive a payoff of 50 with probability one-half.
2. If Moriarty chooses Dover, Holmes will receive a payoff of 0 with probability one-half.
3. Hence, no matter what Moriarty chooses, by using a fair coin, Holmes has a one-half chance of receiving a payoff greater than -100.

Hence, against an omniscient Moriarty, Holmes is better off using a fair coin than by committing to an action.

We will see later that Holmes can make further improvements against an omniscient Moriarty by fine-tuning his randomization, e.g. by using a biased coin.

Arthur Conan Doyle

1. neither endows Moriarty with the power of omniscience,
2. nor does he see the benefit of randomization.

In contrast in The Solar Lottery, Philip K. Dick

1. explicitly endows the players on one side of the game with omniscience by making them telepaths (teeps)
2. and has the players on other side of the game adopt a clever randomization scheme designed to defeat the telepathic abilities of their adversaries.

The solar lottery - plot summary - and comments

- Ted Benteley loses his job at one of the hills (the industrial basis of the society).
- He plans to leave the hill system and to join the Directorate (the political structure).
- At the top of the Directorate is the Quizmaster, currently Reese Verrick.

- Benteley cannot be sure to be signed on by the Directorate because randomness is an essential part of the entire social structure: *"Quizmaster Verrick's hiring was integrated on the basic principle of Minimax: positional oaths were apparently passed out on a random basis. In six days Benteley hadn't been able to plot a pattern."* (p.3)
- Comments: This is not as crazy as it may sound. An example of randomness in today's society can be found in tax-auditing.
- An integral part of social organization is "the bottle" whose random motion can move one up in the classification system.

- Benteley does end up being signed up by the Quizmaster.
- Benteley, however, takes an oath directly to Reese Verrick rather than to the position as Quizmaster.
- Soon after, he learns that a random twitch of the bottle has removed Reese Verrick from his position of Quizmaster.
- Comment: So, even the selection of who is in the top political position has an explicit random element. There are historical examples of random assignments of administrative positions. Under Athenian democracy, for example, a sizable number of administrative posts were regularly determined by a lottery.

- A new Quizmaster, Leon Cartwright, arrives.
- A key feature of the political structure is that the incumbent Quizmaster can be removed by an assassin who is publicly determined at a challenge convention.
- Part of the system is that the Quizmaster is defended by a group of telepaths (teeps). All teeps are available to the Quizmaster and only the Quizmaster has teeps.
- The challenge convention is manipulated by the former Quizmaster, Reese Verrick.

- One of Verrick's men, Herb Moore, has designed a randomized scheme to defeat the teeps:
  - Moore has constructed an android, Keith Pellig, who has no volition of his own
  - Pellig is hooked up to a system that permits him to be endowed with the mind of a remote operator.
  - The remote operators are switched randomly.
  - Thus, the teeps can never be entirely sure of Pellig's intentions.

- Upon the advice of the teep corps, Leon Cartwright tries to hide on the moon.
- When the Pellig body does not find Cartwright, it heads to the moon, under the guidance of Herb Moore.
- On the moon, Ted Benteley is switched into the Pellig body.

- Benteley discovers that the Pellig body contains a bomb that would destroy him along with Cartwright.
- Benteley switches sides, and destroys the bomb in the Pellig body, thus violating his oath to Verrick.
- Verrick appears on the moon to kill Benteley and to offer a deal to Cartwright.
- Through a clever scheme Cartwright, seemingly accepting Verricks offer, gets out of the Quizmaster postion, kills Verrick, and puts Benteley in as the new Quizmaster.

Comments on the appearance of game theory in The Solar Lottery

- The society portrayed in the Solar Lottery is organized around the M-game:

  "In the early twentieth century the problem of production had been solved ... The Quizzes had helped, a trifle. If people couldn't afford to buy the expensive manufactured goods, they could still hope to win them. ... Gradually, over the years, prizes in the quizzes grew from material commodities to more realistic items: power and prestige. ... Statistical prediction became popular ... People lost faith in the belief that they could control the environment. p.16-17

"The theory of Minimax – the M-game – was a kind of stoic withdrawal, a nonparticipation in the aimless swirl in which people struggled. ... Minimax, the method of surviving the great game of life, was invented by two twentieth century mathematicians, von Neumann and Morgenstern. It had been used in the Second World War, the Korean War and in the Final War. Military strategists and then financiers had played with the theory. In the middle of the century von Neumann was appointed to the U.S. Atomic Energy Commission: recognition of the burgeoning significance of his theory. And in two centuries and a half it became the basis of government." p.17

- Assassin's can be thought of as providing "check's and balances" in the M-game. Quizmasters have to prove their competence against each assassin.
- *"The challenge ... protects us from incompetents, from fools and madmen."* p.37
- Except, as Leon Cartwright puts it: *"Assassins have killed every unk the bottle ever twitched. ... The checks and balances of this system work to check us and balance them. ... As far as they're concerned, I broke the rules by just wanting to play."* p.21

- Leon Carwright is described as a "chaotic jumble of passion." It is perhaps enlightening to relate this to **Nixon's madman strategy**. Compare the following quote from the Boston Globe (June 14, 2005): *"I call it the madman theory, Bob," Richard Nixon said to Robert Haldeman. ... "I want the North Vietnamese to believe," he went on, "that I've reached the point that I might do anything to stop the war. We'll just slip the word to them that for God's sake, you know Nixon is obsessed about communism. We can't restrain him when he's angry, and he has his hand on the nuclear button, and Ho Chi Minh himself will be in Paris in two days begging for peace."*

- Philip K. Dick deliberately uses telepaths to justify minimax reasoning. At one point he makes Verrick ask: *“How can you plan with teeps around? Teeps perfectly fulfill the pessimistic expectations of Minimax: they find out every strategy.”* p.59

- Philip K. Dick fully understands and lets Herb Moore explain the rationale for using a mixed strategy in a zero-sum game: *"If you act randomly your opponent can find out nothing about you because even* you *don't know what you're going to do."* p.59

- Philip K. Dick connects randomization in games with other random phenomena, e.g. Heisenberg's "Uncertainty Principle" (The more precisely the position is determined, the less precisely the momentum is known in this instant, and vice versa. –Heisenberg, uncertainty paper, 1927). Philip K. Dick lets Herb Moore say: *"You see, Pellig is Heisenberg's random particle. The teeps can trace his path: directly to Cartwright. But not his velocity. Where Keith Pellig will be along that path at a given moment nobody knows."* p.79

- Philip K. Dick realizes that in strategic situations the value of information may be negative. Eleanor Steven's explains: *"A lot of non-teeps won't live with a teep ... ."* p.81

- Toward the end of the novel it becomes clear that a number of the players have engaged in deception and bent the rules. This includes Herb Moore, who did not reveal certain features of the Pellig body to Verrick; Leon Cartwright, who manipulated the bottle; and Philip K. Dick himself, who for most of the novel keeps all this information from the reader.

Payoffs in games with mixed strategies

1. When we permit players in a game to randomize, we effectively enlarge their set of available strategies.
2. This prompts us to assign payoffs to combinations of mixed strategies.
3. The outcome of a combination of mixed strategies is randomly determined.
4. Therefore we have to be able to assign payoffs to "lotteries".

		Holmes	
		C'bury	Dover
Moriarty	Canterbury	(100,-100)	(-50,50)
	Dover	(0,0)	(100,-100)

GAME III-1

Payoff from lotteries

1. Example: Recall the Moriarty-Holmes game. We will assume that a lottery that results in a payoff of 100 with probability one-half and a payoff of -50 with probability one-half yields a payoff of

$$\frac{1}{2}100 + \frac{1}{2}(-50) = 25.$$

2. This is Moriarty's expected payoff from going to Canterbury if Holmes randomizes and puts equal probability on going to Canterbury and going to Dover.

Moriarty's payoffs against equal-probability mixing by Holmes

Expected payoff from Canterbury:

$$\frac{1}{2}100 + \frac{1}{2}(-50) = 25$$

Expected payoff from Dover:

$$\frac{1}{2}0 + \frac{1}{2}100 = 50.$$

Two important observations

1. An omniscient Moriarty would choose "Dover" against equal-probability mixing by Holmes.
2. Observe that Holmes can lower Moriarty's payoff (and thereby increase his own payoff) by increasing the probability on Canterbury.

The indifference principle

For Holmes's mixed strategy to be optimal against an omniscient Moriarty, it must make Moriarty indifferent between both of his strategies.

Applying the indifference principle

1. If Holmes goes to Canterbury with probability p, then Moriarty is indifferent between Canterbury and Dover if

$$p100 + (1-p)(-50) = p0 + (1-p)100,$$

or equivalently if

$$250p = 150,$$

or

$$p = \frac{3}{5}.$$

Holmes's minmax strategy

In the mixed extension of the Holmes-Moriarty game, Holmes's minmax strategy is to go to Canterbury with probability $p = \frac{3}{5}$.

Applying the indifference principle once more

1. If Moriarty goes to Canterbury with probability q, then Holmes is indifferent between Canterbury and Dover if

$$q(-100) + (1 - q)0 = q50 + (1 - q)(-100),$$

or equivalently if

$$250q = 100,$$

or

$$q = \frac{2}{5}.$$

Moriarty's maxmin strategy

In the mixed extension of the Holmes-Moriarty game, Moriarty's maxmin strategy is to go to Canterbury with probability $q = \frac{2}{5}$.

The value of the Holmes-Moriarty game is $\frac{2}{5} \times 100 = 40$.

Some important facts about minimax values

If a game has a minimax value, then it has a Nash equilibrium.

A Nash equilibrium payoff in a two-person zero-sum game is the game's minimax value.

The mixed extension of every finite two-person zero-sum game has a value.