Game Theory
Individual Optimization

Outline:

- Individual Optimization
- Security Strategies

1 Introduction

Up to this point in the course, we have investigated several types of group decision-making, and uncovered several issues which can make group decision-making challenging. However, one thing that has been absent in our discussion so far is the role that *individual* decisions may play in the collective outcomes reached by the group. In other words, *individual strategy* has so far been absent.

For example, in discussing the Social Choice problem, we always considered the individual opinions as given – thus deliberately avoiding the additional challenges of tactical voting. In a real election, a voter is not required to vote for the candidate which they actually want to win – rather, they may assess which of the candidates are actually *likely* to win and vote for their favorite among those.

As another example, the Matching lectures discussed the matching problem in detail and showed how to perform bipartite matching using the Gale-Shapley algorithm in a way that accounts for individual preferences, but we never asked whether the structure imposed by the Gale-Shapley algorithm actually incentivizes the individuals to report their preferences truthfully.

These kinds of issues prompt the next section of this course, in which we will focus on the role that individual strategy plays in collective outcomes associated with group behavior. We will ask questions such as

- If all individuals behave in a self-interested way, how does this affect the collective behavior of the group?
- What are the right models and conceptual frameworks for analyzing self-interested behavior?

2 Self-Interested Behavior

If we are to understand the effects of individual strategy on collective behavior, we first need a mathematical framework that captures individual strategy. Throughout the remainder of this course, we will adopt the modeling convention that each individual i has a *utility function,* which is a function U_i that takes in a "state of the world" as input and returns a real number that encodes how much the individual "likes" that state of the world. As a very simple example, suppose there are two states of the world: "rainy" (denoted R) and "sunny" (denoted S). Then a person's utility function would assign a number to each of these states; a bigger number means the person likes it more. For instance, if individual i has a utility function of

$$U_i(x) = \begin{cases} 5 & \text{if } x = R \\ 3 & \text{if } x = S, \end{cases}$$

this would be one way of representing the idea that person i likes rainy weather better than sunny weather, since $U_i(R) > U_i(S)$. We would say that person i receives a "payoff" of 5 when it rains, and a "payoff" of 3 when it is sunny. In essence, this says that if person i had the opportunity to choose between rain or sun, she would choose rain.

2.1 Umbrellas: Individual Choice

In reality, a person does not get to choose between rain or sun – the weather is out of an individual's control. However, this does not mean an individual has no choices to make. Importantly, in many circumstances, a person can make choices that affect her preferences over states of the world. Put differently, the choices made by a person are themselves part of the state of the world.

To see how this works in the rain/sun scenario, suppose that our person (call her Alice) owns an umbrella, and each day she has to make a choice: bring the umbrella or leave the umbrella at home? By including her choice in our description of the state of the world, we grow the number of states to 4: it either rains or is sunny, and Alice either brings her umbrella or she doesn't. Now, to fully express Alice's preferences, her utility function needs to be defined for each of these 4 possibilities. A convenient way to represent this is with a *payoff matrix*:

		The Weather	
		Rain	Sun
Alice's Choice:	Umbrella	5	2
	No Umbrella	0	3

Now that we see these expanded payoffs, we see that Alice's enjoyment of rain versus sun depends on whether she brought her umbrella. If it's sunny, she'd prefer to leave her umbrella at home. If it's rainy, she'd prefer to bring her umbrella with her. To understand Alice's preferences as encoded by this payoff matrix, it is useful to differentiate between the things under Alice's control and the things that are out of her control: all Alice can do is decide whether or not to bring her umbrella; she cannot decide whether it is rainy or sunny. Alice only has control over the *row* of the payoff matrix – she simply must accept the column as given.

Let us ask a new question: "Should Alice bring her umbrella, or leave it at home?" By reading the payoff matrix carefully, we see that the answer depends on the state of the weather! If it is raining, Alice prefers to bring her umbrella since $U_i(U, R) > U_i(\neg U, R)$ (specifically, $5 > 0$). If it is sunny, Alice prefers *not* to bring her umbrella since $U_i(U, S) < U_i(\neg U, S)$ (specifically, $2 < 3$).

At this stage, one of the most crucial things to understand is that when Alice is trying to decide what to do, she only has control over the rows of the payoff matrix. Simply put, she only ever reasonably compares numbers across rows (that is, up-and-down) – there is no point in comparing numbers across columns. To see this, suppose she has her umbrella and it is sunny (so her payoff is 2 – she is lugging the umbrella around for no reason). She can reasonably say "I wish I hadn't brought my umbrella, because then my payoff would be 3." She could have left her umbrella at home, and she is solely responsible for the regret she feels for bringing it. However, if she says "I wish it were raining, because then I would be getting a payoff of 5" she is just whining. She has no control over the weather (that is, no control over the column of the matrix), so the fact that she prefers rain to sun is just wishful thinking. In a practical sense, it doesn't matter that there is a hypothetical way for her to be getting a payoff of 5, since she has no ability to get there herself.

2.2 Security Strategies

Now we know how to think about Alice's choices in *retrospect*: if she decides to bring the umbrella and the weather turns out to be rainy, she's happy. If she decided to leave the umbrella at home and it turns out to be sunny, she's happy. But how should Alice make these decisions in the first place? In other words, if Alice wakes up in the morning and has no idea whether it will rain today, should she bring her umbrella or not?

It turns out the answers to this question are varied and nuanced. However, as a starting point, let's think about how Alice would make her decision if she is very pessimistic. In other words, suppose Alice's main goal is to avoid low payoffs. To do this, she can look at her choices, and for each one ask "if I choose this option, what is the worst that could happen to me?" Then, having thought about worst-case scenarios for each choice, she can just pick the one with the *best* worst-case scenario. One way to think about this in terms of payoff matrices is that she writes down the smallest number in each row of the payoff matrix, and then treats that number as the payoff for that particular action. Here, we write the original payoff matrix on the left, and on the right we write the worst-case payoff (the smallest number in each payoff row) for each of her potential choices:

		The Weather	
		Rain	Sun
Alice's Choice:	Umbrella	5	2
	No Umbrella	0	3

	Worst-case payoff
Umbrella	2
No Umbrella	0

Using this technique, Alice can hedge her bets and always choose to bring her umbrella, knowing that her payoff will never be worse than 2 no matter what the weather is.

Is this the best Alice can do? It depends. If Alice is committed to choosing just one option always, then yes – she can't ever *guarantee* that she'll ever get more than 2. However, notice that if Alice has committed to always bring her umbrella, an odd thing happens with this example: on days when it is sunny and she actually gets a payoff of 2, she experiences some regret – she wishes she'd left her umbrella at home. Alice would like to exploit this fact somehow and improve her worst-case payoffs.

To do so, she needs to think bigger: she needs to think in terms of *average* or *expected* payoffs. The idea here is that she knows that some days are sunny, so sometimes she'll leave her umbrella at home. Perhaps not too often, because leaving her umbrella at home puts her at risk of getting rained on. But if she can properly balance the risk of being rained on with the benefits of getting to walk around in the sun without an umbrella weighing her down, then on average she could potentially be better off than if she simply always brought the umbrella.

To analyze this, we need a notion of average payoffs, which requires a notion of randomized choices. Suppose Alice brings her umbrella a p fraction of the time, and leaves it at home a $(1-p)$ fraction of the time. That is, if $p = 1/3$, this means that on average, she has her umbrella with her a third of the days. Then Alice can continue to apply her worst-case analysis: just as with the simple deterministic case, she will ask for each possible value of p, what is the worst thing that could happen? In other words, what is her expected payoff as a function of p, for each of the possible states of weather? Figure 1 captures this.

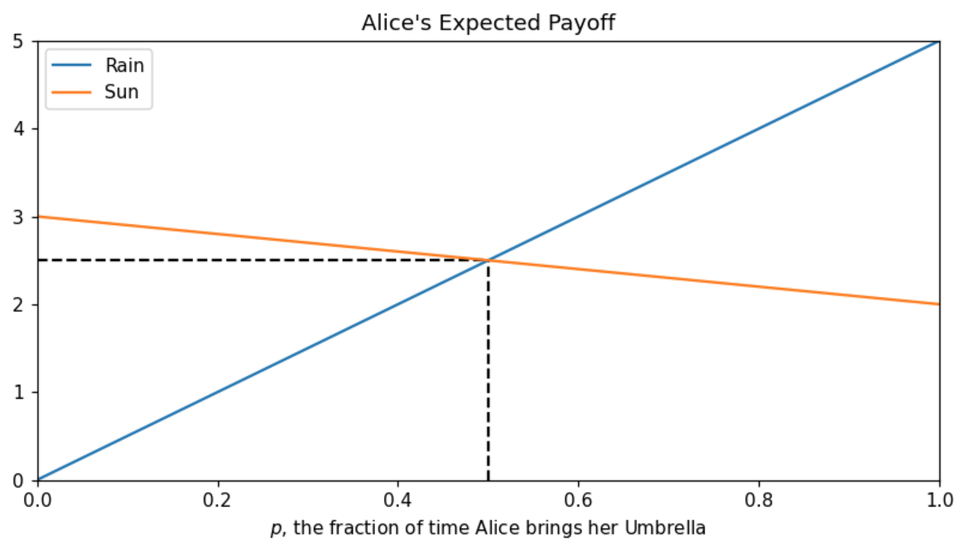

Figure 1: Plot depicting Alice's expected payoff as a function of p, the fraction of days on which she brings her umbrella. The blue line corresponds to her expected payoff when it rains, and the orange line corresponds to her expected payoff when it is sunny. Note that when she always brings her umbrella (i.e., $p = 1$), she receives a payoff of 5 when it is rainy and 2 when it is sunny. This means that her worst-case guaranteed minimum payoff is only 2. However, if she brings her umbrella only half the time (i.e., $p = 1/2$), her worst-case (expected) guaranteed minimum payoff increases to 2.5. This is because regardless of whether it rains or is sunny, her expected payoff is *always* exactly 2.5 when she brings her umbrella half the time.

For each p, our pessimistic Alice wants to know what is the worst thing that could happen. To find out what this is, Alice plugs in values of p, and then checks whether the orange (sunny) line or the blue (rainy) line results in a lower expected payoff. For instance, it is useful to consider the deterministic case (i.e., $p \in \{0, 1\}$). When $p = 0$ (Alice always leaves her umbrella at home), her expected payoff is equal to 0 when rainy and 3 when sunny, so the worst thing that could happen is rain. Likewise, when $p = 1$ (Alice always brings her umbrella), her expected payoff is equal to 5 when rainy and 2 when sunny, so the worst thing that could happen is sun.

The interesting and less-intuitive case is when Alice starts to decrease p from 1; that is, she starts to leave her umbrella at home occasionally. Why would Alice do this? Remember that Alice is focused on worst-case outcomes; when $p = 1$, the worst thing that could happen is a sunny day. What she finds is that when she occasionally leaves her umbrella at home, the worst thing that could happen is still a sunny day, but now she enjoys more sunny days. It happens that she also enjoys her rainy days a little less – but rainy days are still far from being the worst thing that could happen, so she doesn't worry about that potential loss of enjoyment.

To see this on the plot in Figure 1, for each p simply find the minimum of the Rain (blue) and Sun (orange) traces; this value is Alice's minimum expected payoff for that value of p. Since Alice's goal is to *maximize* her minimum expected payoff, she wants to select p such that the minimum of the two lines is the largest-possible; for this particular problem, this "maxi-min" payoff occurs at $p = 0.5$, and is marked in Figure 1 by the dashed black line. That is, a pessimistic Alice should bring her umbrella exactly half the time on average – and this will result in an expected payoff *no less than* 2.5.

In game-theoretic terms, we call this type of strategy is a *security strategy*, or a *maximin* strategy: Alice is playing so as to maximize her guaranteed expected payoff. Note that in computing her security strategy, Alice made no assumptions on the weather – her strategy is completely independent of what the weather actually does. It depends only on her own payoffs.

2.3 The Math

To model these things formally, we will write $\mathcal{A}_i$ to denote the set of choices (or *pure strategies*) that player i can select, and $\mathcal{A}_{-i}$ to denote the set of states of the world *other than the ones that player i controls.* With this, we write $\mathcal{A} := \mathcal{A}_i \times \mathcal{A}_{-i}$ to denote the complete state of the world. Given this, player i's utility function is defined as $U_i : \mathcal{A} \to \mathbb{R}$. We think of player i as an individual optimizer; that is, player i attempts to select an action $a_i \in \mathcal{A}_i$ that maximizes her utility function, given the state of the world. We will often speak of player i's *best response function*, $B_i : \mathcal{A}_{-i} \to 2^{\mathcal{A}_i}$ which takes a state of the world as an input and returns the set of actions which maximize the player's utility function:

$$B_i(a_{-i}) := \underset{a_i \in \mathcal{A}_i}{\arg\max}\, U_i(a_i, a_{-i}). \tag{1}$$

Now, to describe randomized strategies for player i, let s_i denote a vector of probabilities with which player i selects each of her actions. That is, s_i denotes a probability distribution over $\mathcal{A}_i$, so each component of s_i lies in the interval $[0, 1]$ and the sum of s_i's components is equal to 1. We call s_i a *mixed strategy* for player i. We write the set of possible mixed strategies as $\Delta(\mathcal{A}_i)$; likewise, we write the set of probability distributions over states of the world as $\Delta(\mathcal{A}_{-i})$. When a player selects a mixed strategy s_i and the probability distribution over states of the world is given by s_{-i}, player i's expected payoff is written

$$\mathcal{U}_i(s_i, s_{-i}) = \mathbb{E}_{(s_i, s_{-i})} U_i(a_i, a_{-i}). \tag{2}$$

The goal of a pessimistic player i (such as Alice in our example in the text) is to select a mixed strategy which *maximizes* her worst-case expected payoff; we call such a strategy a *security strategy* s_i^*, given by

$$s_i^* \in \arg\max_{s_i \in \Delta(\mathcal{A}_i)} \min_{s_{-i} \in \Delta(\mathcal{A}_{-i})} \mathcal{U}_i(s_i, s_{-i}). \tag{3}$$

When player i uses security strategy s_i^*, we say that the player obtains a *value* of

$$\underline{v} = \max_{s_i \in \Delta(\mathcal{A}_i)} \min_{s_{-i} \in \Delta(\mathcal{A}_{-i})} \mathcal{U}_i(s_i, s_{-i}). \tag{4}$$

To encode the example with Alice and the weather, we identify Alice with index i and write Alice's action set as $\mathcal{A}_i = \{U, \neg U\}$ and the other states of the world as $\mathcal{A}_{-i} = \{R, S\}$. Then the complete state of the world is a member of the set $\mathcal{A} = \{(U, R), (U, S), (\neg U, R), (\neg U, S)\}$. Alice's utility function is conveniently written as a payoff matrix:

		Weather	
		R	S
Alice	U	5	2
	$\neg U$	0	3

Now, let Alice's mixed strategy be given by $s_i = (p, 1 - p)$; that is, she brings her umbrella with probability p and leaves it at home with probability $1 - p$. Likewise, suppose the state of the weather is drawn from a probability distribution $s_{-i} = (q, 1 - q)$; that is, it rains with probability q and is sunny with probability $1 - q$. Then Alice's expected utility is given by

$$\mathcal{U}_i(s_i, s_{-i}) = 5pq + 2p(1 - q) + 3(1 - p)(1 - q). \tag{5}$$

For any fixed p, note that Alice's expected utility is a linear-affine function of q; this guarantees that her worst-case expected utility occurs either when $q = 0$ or when $q = 1$ (that is, when the weather is either always sunny or always rainy, respectively). Thus, Alice may compute her security strategy s_i^* by considering only the $q = 1$ and $q = 0$ cases of (5):

$$s_i^* \in \arg\max_{p \in [0,1]} \min\{5p,\ 2p + 3(1 - p)\}, \tag{6}$$

and her value is

$$\underline{v} = \max_{p \in [0,1]} \min\{5p,\ 2p + 3(1 - p)\}. \tag{7}$$

To see this math in action, see Figure 1, and note that the blue (rainy) line is given by the expression $5p$, and the orange (sunny) line is given by the expression $2p + 3(1 - p) = 3 - p$.

3 Conclusion

In this lecture we started to look at the problem of strategic decision-making in uncertain environments. Any decision-making process for such scenarios must incorporate an underlying model for how the environment will make its decision. This lecture focused on the most pessimistic of models, i.e., a worst-case model. One of the advantages of this model is its universality; performance guarantees for a worst-case model immediately extend to any other model of the environment. While these performance guarantees could be low in general, we demonstrated that randomization can be exploited to improve these worst-case performance guarantees.

4 Questions

1. Consider the following payoff matrix. As in the Alice/Umbrella example, the decision-maker's actions correspond to the rows of the matrix, while the other possible states of the world correspond to the columns:

$$\begin{pmatrix} 1 & 3 \\ 4 & 2 \end{pmatrix}$$

   (a) What is the decision-maker's mixed security strategy for this payoff matrix?

   (b) What is the decision-maker's value?

2. Consider the following payoff matrix. As in the Alice/Umbrella example, the decision-maker's actions correspond to the rows of the matrix, while the other possible states of the world correspond to the columns. Note that here, there are 3 other possible states of the world.

$$\begin{pmatrix} 0 & 3 & 1 \\ 4 & 1 & 2 \end{pmatrix}$$

   (a) What is the decision-maker's mixed security strategy for this payoff matrix?

   (b) What is the decision-maker's value?