Chapter 7: Zero-Sum Games

Elements of Decision: Lecture Notes of Intermediate Microeconomics

Charles Z. Zheng

Tepper School of Business, Carnegie Mellon University

Last update: April 7, 2020

To apply the decision-theoretic techniques introduced in previous chapters to game theory, we start with zero sum games, which we can solve as if it were a consumer optimization problem.

1 Strategies and payoffs

A *zero-sum game* involves two players such that one player's payoff is equal to the negative of the other player's payoff (hence the appellation "zero sum"). For example, consider this matrix:

	L	R
X	2	0
Y	-1	3

This matrix describes a zero-sum game between two players, one choosing a row (X or Y), the other choosing a column (L or R). The number in the matrix is the payoff for the row player, and the negative of the number the payoff for the column player, if they choose the row and column whose intersection is the number. For instance, if the row player chooses Y, and the column player L, then the row player's payoff is equal to -1, and the column player's payoff, 1. If the row player switches to X and the column player sticks to L, then the former gets payoff 2 and the latter -2. The rows, X and Y in this game, are called *pure strategies*, or *strategies* for short, of the row player. Likewise, L and R are the (pure) strategies of the column player.

Note, from a player's perspective, choosing a strategy amounts to choosing a vector of payoffs. For example, the strategy Y to the row player is equivalent to the vector $(-1, 3)$, meaning that the player gets payoff -1 if his opponent plays L, and gets 3 if the opponent plays R. Namely, strategy Y corresponds to a vector of *contingent payoffs* (cf. §4, Chapter 4) for the row player, contingent on which strategy the column player is to choose.

For any two vectors (v_1, v_2) and (w_1, w_2), a *convex combination* between them means

$$\lambda(v_1, v_2) + (1 - \lambda)(w_1, w_2) = (\lambda v_1 + (1 - \lambda)w_1, \lambda v_2 + (1 - \lambda)w_2)$$

for some $\lambda \in [0, 1]$. Geometrically speaking, a convex combination between two points is a point that belongs to the straight segment between the two points. The set of convex combination is exactly the straight segment. In general, a convex combination among n vectors $X_1, X_2, \ldots, X_n$ is in the form of

$$\lambda_1 X_1 + \lambda_2 X_2 + \cdots + \lambda_n X_n$$

for some nonnegative real numbers $\lambda_1, \lambda_2, \ldots, \lambda_n$ such that $\lambda_1 + \lambda_2 + \ldots + \lambda_n = 1$.

The set of convex combination among any finitely many vectors $X_1, X_2, \ldots, X_n$ is called *convex hull* of $X_1, X_2, \ldots, X_n$. In the case where these vectors are all two-vectors (elements of $\mathbb{R}^2$), which is the case that concerns us in this chapter, it is easy to draw the convex hull: first, draw the vectors $X_1, X_2, \ldots, X_n$ as points on the plane; second, for any pair of these vectors, draw the straight segment between them; third, you see that some of such segments form a polygon that contains all the segments drawn in the previous step. This polygon, boundary and interior included, is the convex hull of vectors $X_1, X_2, \dots, X_n$. See Figure 1 for example.

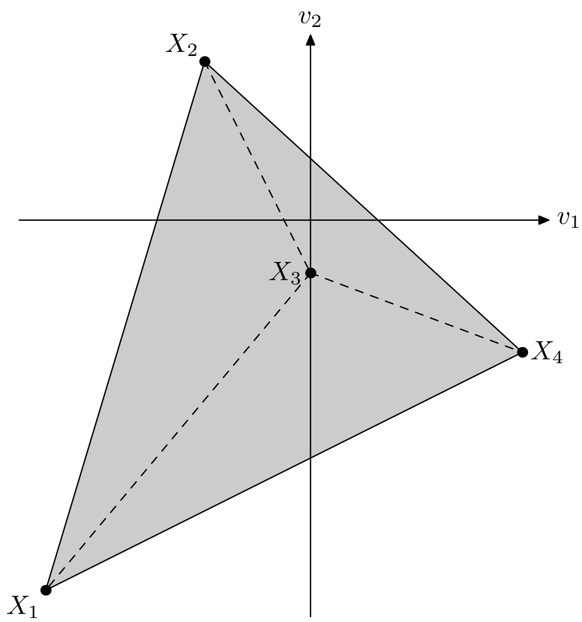

Figure 1: The shaded area is the convex hull of X_1, X_2, X_3, X_4

A convex combination among strategies is called *mixed strategy*. In the above payoff matrix, for example, $\frac{1}{4}X + \frac{3}{4}Y$ is a mixed strategy, meaning that the row player chooses X with probability $1/4$, and Y with probability $3/4$. In terms of the contingent payoff vectors that the pure strategies correspond to, this mixed strategy is

$$\frac{1}{4}X + \frac{3}{4}Y = \frac{1}{4}(2,0) + \frac{3}{4}(-1,3) = \left(\frac{1}{4} \times 2 + \frac{3}{4}(-1), \frac{1}{4} \times 0 + \frac{3}{4} \times 3\right) = \left(-\frac{1}{4}, \frac{9}{4}\right).$$

In other words, the mixed strategy $\frac{1}{4}X + \frac{3}{4}Y$, from the row player's standpoint, amounts to getting an *expected payoff* equal to $-1/4$ in the event of the opponent playing L, and an expected payoff $9/4$ in the event of R. By expected payoff we mean the expected value of a player's payoff from a lottery (cf. §4, Chapter 4). In the above example, $-1/4$ is the expected payoff from the lottery of getting payoff 2 with probability $1/4$, and getting payoff -1 with probability $3/4$; and this is the lottery that the row player gets when he plays the mixed strategy $\frac{1}{4}X + \frac{3}{4}Y$ while his opponent plays L.

Note that a pure strategy is also a mixed strategy, with the probability for the pure strategy equal to one, and the probabilities for other pure strategies equal to zero. (That is, in Figure 1, each of the four vectors X_1, X_2, X_3, X_4 is an element of the shaded area.)

2 A zero-sum game as consumer optimization

As we have seen previously, in choosing a strategy a player is essentially choosing a vector of contingent payoffs, say (v_L, v_R). That is analogous to the consumer in our earlier chapters who is choosing a vector of consumptions, (x_1, x_2). There, a component say x_1 of the vector means the quantity of good 1 the consumer is to get. Here, a component say v_L of the vector means the expected payoff the (row) player is to get in the event where his opponent plays L. Thus, we can figure out a player's decision as if he were a consumer choosing among consumption bundles, with consumption bundles corresponding to his contingent payoff vectors, and *budget set corresponding to the convex hull generated by this player's pure strategies*.

A quick way to grasp this technique is to walk through the following steps with respect to the game in Section 1:

1. Set up a coordinate system with the horizontal axis labeled u_L and vertical axis u_R.

2. Plot the point $(2, 0)$ in the coordinate system. This point represents the contingent payoff vector
$$(u(X, L), u(X, R))$$
for the row player when she plays the pure strategy X. Label the point $(2, 0)$ by X.

3. Analogously, plot the point $(u(Y, L), u(Y, R))$ (i.e., $(-1, 3)$) in the diagram and label it by Y.

4. Draw the convex hull, or the set of all the convex combinations, generated by the points X and Y in the diagram. This set is the row player's *budget set*: As a consumer's budget set consists of all the available consumption bundles, here the set of convex combinations between X and Y consists of all the contingent payoff vectors available to the row player: he has only two pure strategies, X and Y, and he can mix them in whatever way he sees fit.

5. Suppose the row player assumes that, whichever strategy he plays, his payoff will be the worst possible one given his strategy, e.g., if he plays X, he assumes that the column player plays R so that his payoff equals zero. In other words, suppose that his preference is represented by the utility function
$$V(u_L, u_R) = \min\{u_L, u_R\} \tag{1}$$
for all $(u_L, u_R) \in \mathbb{R}^2$. Then draw in the diagram the row player's indifference map, which we know from earlier chapters consists of L-shape curves whose corners all belong to the 45-degree line passing through the origin.

6. Combining the indifference map and the budget set obtained above, and using the basic principle of consumer optimization in previous chapters, we easily find all the optima for the "consumer." Here there is a unique optimum, which is the intersection between the 45-degree line and the segment between X and Y. Label in the diagram the optimum for the row player.

7. Calculate the coordinates of the optimum as follows:

    a. Note that it is a convex combination between X and Y. Hence suppose it is $\lambda X + (1-\lambda)Y$ for some λ between zero and one. Correspondingly, the coordinate of this optimum is
    $$(\lambda u(X, L) + (1-\lambda)u(Y, L), \lambda u(X, R) + (1-\lambda)u(Y, R)),$$
    i.e., $(2\lambda - (1-\lambda), 3(1-\lambda))$

b. Note that the optimum lies on the 45-degree line, hence

$$2\lambda - (1-\lambda) = 3(1-\lambda),$$

i.e., $\lambda = 2/3$.

c. Thus the row player attains his optimum through mixing X and Y in the portion of 2/3 for X and 1/3 for Y, or $\frac{2}{3}X + \frac{1}{3}Y$ for short. In other words, given the utility function (1), the row player's optimal strategy is to play X with probability 2/3 and Y with probability 1/3.

8. Then in the diagram the row player's optimum is $(2(2/3) - 1/3, 3(1/3))$, i.e., $(1,1)$. In other words, through playing the mixed strategy $\frac{2}{3}X + \frac{1}{3}Y$, the row player can guarantee that his expected payoff is no less than one no matter what the column player chooses.

   In general, if (u_L^*, u_R^*) is the row player's optimum, then at this optimum the row player's utility according to (1),
   $$V(u_L^*, u_R^*) = \min\{u_L^*, u_R^*\},$$
   is called the row player's *security level*. Any mixed strategy that attains the security level is his *security strategy*.

   In the game that we have just gone through, it turns out that $u_L^* = u_R^* = 1$, so his security level $\min\{1,1\}$ is equal to one; and his security strategy is $\frac{2}{3}X + \frac{1}{3}Y$.

Note that the row player's utility function (1) assesses any contingent payoff vector according to the worst-case scenario of the vector. Thus, the optimum obtained through the above procedure corresponds to the best for the player among worst-case scenarios. In other words, *in playing a security strategy, the player secures an expected payoff no lower than his security level*.

3 What about the column player?

We can find the column player's security strategies and security level by the following steps analogous to those in the previous section:

1. Set up another diagram with the horizontal axis labeled $-u_X$ and vertical axis $-u_Y$.

2. Plot on the diagram the point $(-u(X,L), -u(Y,L))$, i.e., $(-2,1)$, and label it with L. This point represents the contingent payoff vector for the *column* player when she plays the pure strategy L. (Recall that the game is zero-sum. Thus be sure to switch the sign of the payoffs in the matrix to register payoffs for the column player.)

3. Analogously, plot the point $(-u(X,R), -u(Y,R))$, namely, $(0,-3)$—again do not forget to switch the sign from the payoff matrix—and label it by R.

4. Plot the set of all the convex combinations between the points L and R in the diagram. This set is the column player's budget set.

5. Suppose the column player assumes that, whichever strategy she plays, her payoff will be the worst possible one given her strategy, e.g., if she plays L, she assumes that the row player plays X so that her payoff equals -2. Then draw in the diagram the column player's indifference map.

6. Plot in the diagram the optimum for the column player. This is her contingent payoff vector generated by her security strategy.

7. Calculate the security strategy:

   a. Suppose it is $\alpha L + (1-\alpha)R$ for some α between zero and one. Correspondingly, the coordinate of this optimum is

   $$\alpha L + (1-\alpha)R = \alpha \begin{bmatrix} -2 \\ 1 \end{bmatrix} + (1-\alpha) \begin{bmatrix} 0 \\ -3 \end{bmatrix} = \begin{bmatrix} -2\alpha \\ \alpha - 3 + 3\alpha \end{bmatrix} = \begin{bmatrix} -2\alpha \\ 4\alpha - 3 \end{bmatrix}$$

   b. Note that the optimum lies on the 45-degree line, hence $-2\alpha = 4\alpha - 3$, i.e., $\alpha = 1/2$.

   c. Then you obtain the column player's security strategy: $\frac{1}{2}L + \frac{1}{2}R$, namely, playing L and R randomly with equal probabilities. The value of α obtained above also gives the coordinate of the column player's optimum in the diagram: $(-2(1/2), 4(1/2) - 3)$, i.e., $(-1, -1)$. From the coordinate you can tell that the column player's security level, i.e., his lowest possible expected payoff from playing his security strategy, is equal to -1.

8. Recall from the previous section that the security level for the row player in this game is equal to one. Thus, in this game, each player's security level is equal to the negative of the other's.

4 The Minimax Theorem

The observation in Step 8 of the previous section has a profound implication. Recall from the previous sections that we obtain a player's security strategy based on the assumption that he assesses his choices according to their worst-case scenarios, as in (1). Thus one would ask Would he still find it optimal to play his security strategy if he is less pessimistic or paranoid? The answer is Yes if his security level is equal to the negative of his opponent's (as in the game analyzed above) and if he expects his opponent to play her security strategy. Let us illustrate that with the numbers in the previous game. Expecting the column player to play her security strategy, which guarantees her an expected payoff no less than her security level -1, the row player knows that his expected payoff cannot go higher than his security level 1 (because the game is zero-sum). Thus, whatever strategy that gives him an expected payoff no less than one is the best he can do. And his security strategy does exactly that. Hence the player's *best response* to his opponent's security strategy is to play his own security strategy, whether he is paranoid or not.

Now switch the table to the column player. The same reasoning in the previous paragraph applies to her as well. Thus the column player would play her security strategy even if she does not assess her options according to their worst-case scenarios, as long as she expects her opponent to play his security strategy. Furthermore, according to the reasoning in the previous paragraph, the column player can expect her opponent to play his security strategy without having to believe that her opponent assess his options according to their worst-case scenarios: she needs only to believe that her opponent believes that she—the column player—will play her security strategy! In other words, *playing security strategies against each other is a self-fulfilling prophesy* in a zero-sum game.

Note that the above reasoning uses a condition that the two players' security levels are the negative of each other. While this is true in the game analyzed above (cf. Step 8, previous section), is this condition true in general? The answer is Yes for all zero-sum games. That is the celebrated Minimax Theorem due to von Neumann: *The security level of the column player is equal to the negative of the security level of the row player*. This theorem gives us a trick to obtain column player's security strategy much quicker than the steps in the previous section:

1. By Step 8 in Section 2, we already know that the security level for the row player is equal to one, hence the security level for the column player is equal to -1 by the minimax theorem.

2. Let the security strategy for the column player be $\alpha L + (1-\alpha)R$ for some α between zero and one. Then the contingent payoff vector for the column player is equal to

$$\alpha\,(-2,1) + (1-\alpha)\,(0,-3) = (-2\alpha, \alpha - 3(1-\alpha))\,.$$

3. Since column player's security level is equal to -1, it follows that

$$\min\{-2\alpha, \alpha - 3(1-\alpha\} = -1. \tag{2}$$

4. Solve the above equation and we get $\alpha = 1/2$.[1] Hence the column player's security strategy is $\frac{1}{2}L + \frac{1}{2}R$.

5 Saddle points and the value of a game

A *saddle point* of any zero-sum game means a pair of security strategies, one for each player. For example, in the game analyzed above, the pair $\left(\frac{2}{3}X + \frac{1}{3}Y, \frac{1}{2}L + \frac{1}{2}R\right)$ is a saddle point.

By the first two paragraphs of Section 4, a saddle point of a zero sum game is an *equilibrium* in the following sense. Say (σ_1, σ_2) is a saddle point, with σ_1 denoting a mixed strategy of the row player, and σ_2 that of the column player. Then the row player cannot do better than σ_1 as long as the column player sticks to σ_2, and neither can the column player do better than σ_2 as long as the former sticks to σ_1. Note that the worst-case-scenario preference assumption we made in Step 5 of Sections 2 and 3 is not needed at all to the equilibrium property of saddle points. Rather, it is just a solution technique that is valid due to the minimax theorem.

A zero sum game may have multiple saddle points. However, the security level is unique for a player regardless of which security strategy he chooses. This unique security level for the row player is called the *value* of the game. Note, by the minimax theorem, that the negative of the value of the game is equal to the security level of the column player.

[1] Details: Eq. (2) is the same as saying that either $-2\alpha = -1 \geq \alpha - 3(1-\alpha)$, or $\alpha - 3(1-\alpha) = -1 \geq -2\alpha$. Each of the two cases corresponds to a system of simultaneous equations/inequalities. Solve each system. If a system yields no solution, the corresponding case is impossible. Else the solution for that system is a solution for (2). See Section 6 for more on solving such "minimum" equations.

To clarify the fact that a player may have multiple security strategies while his security level is always unique, look at Figure 2. There the row player has two pure strategies and their corresponding contingent payoff vectors have the same vertical coordinate. Thus his budget set is the horizontal segment XY in the figure. This coupled with his L-shape indifference curves implies that the part of XY that coincides with a leg of an L-shape curve is the set of optima for the player. That part is the thick segment in the figure. Note that there is a continuum of optima, namely, the row player has a continuum of security strategies. Nevertheless, all these security strategies yield the same security level: for any point in the thick segment in Figure 2, the minimum (recalling the definition of security level) between its coordinate is equal to the vertical coordinate of X (and that of Y).

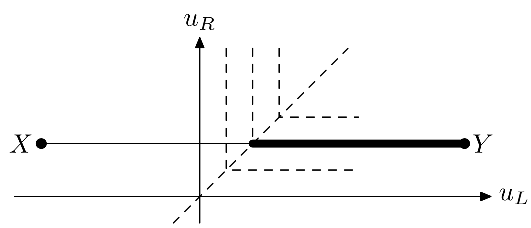

Figure 2: The solid segment XY: the budget set; the thick segment: the optima

Specifically, Figure 2 corresponds to the following game.

	L	R
X	-3	1
Y	5	1

It is instructional to walk through the following steps with this game (cf. Exercise 2).

1. Use the method illustrated above to calculate the security strategies for the row player. Note that there are multiple security strategies in this case. For each security strategy, calculate the row player's security level. Are these security levels equal to one another?

2. Calculate the column player's security strategy. Note that there is only one security strategy in this case. What is the column player's security level?

6 Generalizing to $n \times 2$ Zero-Sum Games

So far we have been considering only 2-by-2 zero sum games, meaning that each player has only two pure strategies. The method sketched above applies generally to n-by-m zero sum games, with row player having n pure strategies, column player m pure strategies, for any natural numbers n and m. However, to use the method on two-dimensional diagrams, we shall restrict attention to only n-by-2 zero sum games, where at least one of the two players has only two pure strategies.

It turns out that our solution technique for n-by-2 zero sum games is essentially the same as that for 2-by-2 ones. Let us illustrate this with the next 3-by-2 zero-sum game.

	t_1	t_2
s_1	-2	0
s_2	3	-1
s_3	-3	1

Following the graphical method illustrated in Section 2, you will see that the row player's budget set is the shaded area in Figure 3, and his optimum corresponds to the thick dot O in the figure, which is the unique point where the budget set shares with the highest possible L-shape indifference curve (the thick dashed L-shape in the figure).[2] To find out the row player's

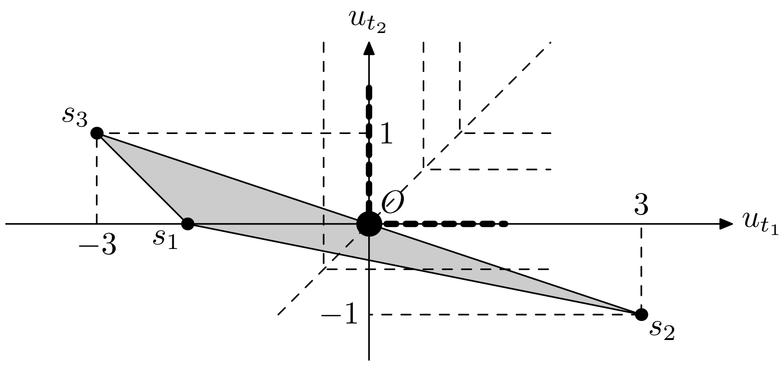

Figure 3: A 3-by-2 game

security strategy, note that the optimum O belongs to the segment between s_2 and s_3. That is, the row player's optimum corresponds to a convex combination between s_2 and s_3. Thus his security strategy is in the form of

$$\lambda s_2 + (1-\lambda)s_3 = \lambda \begin{bmatrix} 3 \\ -1 \end{bmatrix} + (1-\lambda) \begin{bmatrix} -3 \\ 1 \end{bmatrix} = \begin{bmatrix} 6\lambda - 3 \\ -2\lambda + 1 \end{bmatrix}$$

for some $\lambda \in [0,1]$. Since point O belongs to the 45-degree line, we have

$$6\lambda - 3 = -2\lambda + 1,$$

namely, $\lambda = 1/2$. Thus, the row player's security strategy is $\frac{1}{2}s_2 + \frac{1}{2}s_3$, and his security level is

$$\min\{6\lambda - 3, -2\lambda + 1\}|_{\lambda=1/2} = \min\{6(1/2) - 3, -2(1/2) + 1\} = 0.$$

Thus, it follows from the minimax theorem that the column player's security level is equal to zero. Since she has only two pure strategies, t_1 and t_2, her security strategy is in the form of

$$\alpha t_1 + (1-\alpha)t_2 \overset{!}{=} \alpha \begin{bmatrix} 2 \\ -3 \\ 3 \end{bmatrix} + (1-\alpha) \begin{bmatrix} 0 \\ 1 \\ -1 \end{bmatrix} = \begin{bmatrix} 2\alpha \\ -3\alpha + 1 - \alpha \\ 3\alpha - 1 + \alpha \end{bmatrix} = \begin{bmatrix} 2\alpha \\ -4\alpha + 1 \\ 4\alpha - 1 \end{bmatrix}$$

[2]You may have guessed from Figure 3 that the point O is the origin, $(0,0)$, of the diagram. But do not assume so yet, as drawing by hand could be imprecise. If the guess is correct it will turn out as an outcome of the calculation.

for some $\alpha \in [0, 1]$, where the "!" on the first equality is just to remind you of switching the signs for the column player's payoffs. Recall that we have deduced that the column player's security level is zero. By definition of security level we have

$$\min\{2\alpha, -4\alpha + 1, 4\alpha - 1\} = 0.$$

This equation is the same as saying that there are only three possible cases:

i. $2\alpha = 0$ and $-4\alpha + 1 \geq 0$ and $4\alpha - 1 \geq 0$: this is impossible because it implies $\alpha = 0$ and hence $1 \geq 0$ and $-1 \geq 0$.

ii. $-4\alpha + 1 = 0$ and $2\alpha \geq 0$ and $4\alpha - 1 \geq 0$. This is reduced to $\alpha = 1/4$ (consistent with the inequalities: $2(1/4) \geq 0$ and $4(1/4) - 1 \geq 0$).

iii. $4\alpha - 1 = 0$ and $2\alpha \geq 0$ and $-4\alpha + 1 \geq 0$: same as Case (ii).

Thus, there is exactly one solution for the above equation: $\alpha = 1/4$. That is, the column player's security strategy is $\frac{1}{4}t_1 + \frac{3}{4}t_2$.[3] Both players' security strategies and security levels pinned down, we have hence solved the game. (The value of the game is equal to zero, because the security level for the row player is zero according to our previous calculation.)

Finally, what about 2-by-n zero-sum games such as the following one?

Table 1: A 2-by-4 zero sum game

	t_1	t_2	t_3	t_4
s_1	-2	0	4	-5
s_2	3	-1	-2.5	-7

To handle them we need only one little trick: just transpose the payoff matrix such that there are only two rows. That is, simply rewrite the game from the column player's viewpoint:

	s_1	s_2
t_1	2	-3
t_2	0	1
t_3	-4	2.5
t_4	5	7

Note that the the payoffs have each switched signs. Then use the previous technique to solve this transposed game (cf. Exercise 4).

7 Exercises

1. Consider the zero-sum game described by the following payoff matrix:

[3] Why do you think we use the minimax theorem rather than the procedures in Section 3 to find this strategy?

	L	R
X	1	-2
Y	-3	2

a. The column player's payoff is equal to ________ if she plays R and the row player plays X.

b. When the row player plays the mixed strategy $(1/3)X + (2/3)Y$, his expected payoff is equal to ____________ if the column player plays L, and his expected payoff is equal to ____________ if the column player plays R.

c. On a diagram where the horizontal axis stands for the row player's expected payoff if the column player plays L, and the vertical axis the expected payoff if the column player plays R, draw the "budget set" for the row player. Be explicit about the coordinates of the boundary points of the budget set.

d. The column player's expected payoff from playing strategy L is equal to ____________ when the row player plays $(1/3)X + (2/3)Y$.

2. Consider the zero-sum game in Section 5 and complete the two steps listed thereof. (For the row player, follow the steps in Section 2, but note, differently from the example there, the multiplicity of optima. For the column player, either follow the steps in Section 3, or the quicker method in Section 4.)

3. Consider the zero-sum game described by the following payoff matrix:

	L	R
X	3	-2
Y	0	2
Z	-3	1

a. On a diagram with u_L being the horizontal axis and u_R the vertical axis, graph the row player's choice set (shade it with dashed lines) and indifference map (assume that the player evaluates a strategy by its worst-case scenario). Label the row player's optimal contingent payoff vector by the symbol $*$.

b. Row player's security strategy is of the form

i. $\lambda X + (1-\lambda)Y$ for some λ between 0 and 1.

ii. $\lambda Y + (1-\lambda)Z$ for some λ between 0 and 1.

iii. $\lambda X + (1-\lambda)Z$ for some λ between 0 and 1.

iv. $\alpha X + \beta Y + \gamma Z$ for some α and β such that $0 < \alpha < 1$ and $0 < \beta < 1$ and $0 < 1 - \alpha - \beta < 1$.

c. Row player's security strategy is equal to ______________, and his security level equal to ______________.

d. Given the row player's security level derived above, use the minimax theorem to do:

i. the column player's security level is equal to ______________

ii. the lowest possible expected payoff for the column player, if she plays her security strategy, is equal to ______________.

iii. the column player's security strategy is equal to ________________.

4. Use the technique in Section 6 to solve the 2-by-4 zero sum game in Table 1. That is, calculate the security level and all security strategies for each player, and the value of the game.