MAKING GOOD CHOICES: AN INTRODUCTION TO PRACTICAL REASONING

CHAPTER 9: PRACTICAL REASONING IN COMPETITIVE INTERDEPENDENT DECISIONS

In this and subsequent chapters we leave the world of an agent coming to a decision without having to worry about other agents interfering, and enter the world of interdependent decision making. Here is an example to start us off.

1) Who watches TV?

Imagine two college students who have been studying hard for a test and now want to go to the dorm lounge to watch TV for an hour break. For one student a TV break means watching a sports event; nothing else on TV will provides her a study break, and she especially dislikes political talk shows. She would rather go for a walk, not a very attractive study break, than watch anything other than sports. For the other student a political talk show is the best way to relax from studying, watching anything else, especially a sports event, is no break at all; a walk would be better, even though it is a disappointing way to spend a study break. They can't watch both sports and a talk show, for there is only one TV. They have just an hour to spare and neither can do any more studying, they each really need a break now before resuming studying. The time the TV is showing a sports event gains one student her goal, but loses the other student his goal, and vice versa for a talk show, no matter how they divvy up the time. These two students have a conflict of interests. What should each decide: watch TV or go for a walk?

In this scenario, we see agents who are confronted with a decision problem they must try to solve. But this is a decision problem unlike any we have dealt with in previous chapters. There are two agents who have completely opposing goals and each has a menu of options from which they must choose how best to reach their goals. Their decisions are interrelated: how one agent decides depends on how the other agent decides. In the theory of rational choice, such interrelated – interdependent – decisions are called games, and the kind of practical reasoning required in a game decision is called strategic; that is, it is practical reasoning that must take into account the practical reasoning of other agents. Games will be our focus in this and subsequent chapters.

9.1 Basic concepts and framework for interdependent decisions: game theory

No doubt you are familiar with the general idea of a game and can provide plenty of examples. We will narrow the concept of a game to the meaning it has in the theory of rational choice by first considering games in general. We can think of games as dividing into three kinds. Some will be games of chance. There is little or no place for skill or intelligence in games of chance; doing well is a matter of pure luck. Think of gambling opportunities like the game of roulette or a crap game or playing a slot machine, as found at a typical casino. Or take the simple game of flipping a coin to see how many times you can call it correctly. Once a person decides to play a game of chance, assuming that the game is not rigged, doing well or poorly depends entirely on events that are outside the influence of any skill or intelligence the person has.

Other games will be games of skill. Doing well in these games is very much, as the name implies, a matter of training, practice, and skill. For example, think of the many events in the winter and summer Olympic games like swimming, diving, skiing, ski jumping, the high jump, and figure skating. Great skill rather than pure luck is what, for the most part, determines the outcome of these Olympic events, and acquiring skill very much depends on things that are within a person's sphere of influence, such as the time and effort given to practice. Of course, a high degree of talent is often required, and we might think of natural talent as the result of a lucky genetic combination. But assuming that the participants are more-or-less equally talented, without skill raw talent does not go very far, as every trainer and coach well knows.

Finally, some games will be games of strategy. In games of strategy, reasoning is often the determining factor in success, although having a level of skill is typically necessary. Games of strategy include many board games like checkers, chess, and go. Many sports competitions, especially where team play is more important than the special skill of any individual team member, are games of strategy. The “war games” that are played out by armies in preparation for real war, and the “business games” that are played out by companies as practice for real business competitions count as games of strategy. But real wars as well as real business competitions are also thought of as games of strategy, so we should not think of “game” as a decision situation without serious consequences, as some of the above examples might lead you to believe.

This three-fold division of games into chance, skill, and strategy is not meant to be mutually exclusive; some games combine all three aspects. Take for example many card games. The cards you are dealt is a matter of chance; how you play them is a matter of skill and strategy. Likewise, many board games combine in various proportions chance, skill, and reasoning. Many sports require both skill and strategy, with chance being kept to a minimum.

The theory of rational choice studies interdependent decisions in the context of games of strategy, not games of chance or games of skill. For example, the scenario that we started with, two students competing for TV time, is interpreted as a games of strategy. For the purposes of practical reasoning, then, we shall define a **game** as:

> Any individual decision situation involving at least two agents in which the outcomes of each agent’s options depends on the decision of the other agent(s).

This definition is narrower than the general idea of a game because it limits games to decisions within games of strategy. But it is still a broad definition that will have to be made more precise by distinguishing various kinds of game decisions.

The study of these games takes place within several disciplines, prominently mathematics, economics, social science, philosophy, and biology – each discipline finding special areas of interest to research. The study of games is collectively known as Game Theory and is an important part of the theory of rational choice. Game Theory has introduced a somewhat specialized vocabulary that we will sometimes make use of to discuss practical reasoning and rational choice in the context of game decisions. However, in order not to overturn completely the terms we have been using in previous chapters to analyze and evaluate decisions, we will continue to use these (by now) familiar terms interchangeably with the following new vocabulary.

1) **Player**: any individual agent whose decision situation is a game. A player might be, for example, an individual person, a company, a team, a nation, an army, or a family. “Player” and “agent” are interchangeable terms in game decisions.

2) **Payoff**: the outcomes of the options in a game decision. A payoff might be something positive such as a company winning a government contract (represented by a utility number), or something negative such as a student losing the chance to watch TV during a study break (represented by a disutility or negative utility number). “Payoff” and “outcome” are interchangeable terms in game decisions. In cases where the outcome achieves or equals the whole goal, “payoff” and “goal” are interchangeable terms.

3) **Strategy**: the option-outcome alternatives in a game. Instead of saying that an agent decides on an option from the agent’s menu, as we have been speaking up to now, in game decisions it is often said that a player chooses (or “plays”) a strategy. A strategy, then, is a possible course-of-action along with its state and its outcome (payoff) that a player could decide on. “Strategy” and “option” are interchangeable terms in game decisions.

4) **Strategic rationality**: the methods of practical reasoning that should be used in games. A game decision is a decision problem that requires different methods and principles of practical reasoning than are use to analyze and solve non-game decisions (or a different approach to any methods and principles that can be carried over into games). One central feature to strategic practical reasoning is the assumption players make that they have **common knowledge** of the decision problem. Common knowledge is not just two or more agents that have *knowledge in*

common. If you and I both happen to know the same thing, for example, let's say that we both know that the Empire State Building is in midtown New York City, then we have this knowledge in common. But I might not know (or might not have known before you read this paragraph) that you know this fact about NYC, and you might not know (or might not have know before you read this paragraph) that I know this fact about NYC. It just so happens that we know the same fact about the Empire State Building: that it is located in midtown NYC. Common knowledge, however, requires something much more; it requires that I know that you know this, and that you know that I know this. And further, that I know that you know that I know this, and that you know that I know that you know this; and further, that … . Common knowledge, at least in principle, keeps compounding up all possible levels. This is what happened to the knowledge we happened to have in common about the Empire State Building after you read this paragraph: it became common knowledge between us.

Here is an example to think about. Suppose that you and your friend are each given a bowl of chocolate ice cream, but you are not sitting close together and so can't see what flavor each other has. You both, however, have this *knowledge in common*: at least one of us has a bowl chocolate ice cream (you each know this about yourself). I ask you: what flavor ice cream does your friend have? You answer (naturally): I don't know. I then ask your friend the same thing about you, and she answers (naturally): I don't know. Now both you and your friend hear me tell both of you something that you both already know: that at least one of you has a bowl of chocolate ice cream. There is no new information here. However, this piece of information now becomes *common knowledge* between you and your friend, not just knowledge in common. And now I ask your friend: what flavor ice cream does your friend have? She sees that she has chocolate and knows that at least one of you has chocolate, but answers that she doesn't know what flavor you have (because she could be the only one with chocolate). And now I again ask you: what flavor ice cream does your friend have? And now you know: it must be chocolate! If your friend didn't have chocolate, using the common knowledge that at least one of you has chocolate would have allowed her to know that you must be the one with chocolate. But because she answered that she didn't know what flavored you had, using your common knowledge allowed you to conclude that she has to have a bowl of chocolate ice cream. You could never have known this (unless you got up and looked, or asked her!) if what you already knew (that at least one of you has a bowl of chocolate ice cream) had remained knowledge in common. What allowed you to know that she has chocolate was that your knowledge in common was transformed into common knowledge.

In game decision problems, good strategic reasoning requires that players assume common knowledge of the decision situation. You will soon see first hand why it would be foolish for a player not to make this assumption (or to assume the opposite). But for now consider this: there are certain game decisions in which it is very valuable to a player that all the players (including that player!) *not know* what option he will choose, and that they all know that no one knows what option he will choose (that is: "common ignorance" will be valuable in certain games, but ignorance in common will not be valuable).

5) **2-person games**: games are divided into two categories according to the number of players. In a 2-person game there are just 2 players; for example, 2 companies, 2 nations, 2 families, or 2 teams. The 2 players need not be of the same kind; a company and a family might be in a game decision situation, or a person and a nation. In an n-person game there are more than 2 players, "n" standing for any number greater than 2. In n-person games there is the possibility that 2 or more players form alliances or coalitions and interact with 1 other player. Thus, an n-person game can become transformed into a 2-person game in which 1 player interacts with a coalition of the other players now acting as the other individual player.

6) **Competitive** verses **cooperative games**: games are divided into two categories according to the degree with which the players have a conflict of interest. In competitive games, the players have a total conflict of interest; the degree to which one player achieves the goal is the degree to which the other players loose the goal. In competitive games, one player's gain is another

player's loss; a positive outcome for one means a negative outcome for another in equal measure. If you add the gain of one player to the loss of the other(s), they must sum to zero. Thus, competitive games are also called **zero sum games**. Cooperative games form the other category. In cooperative games the players do not have a total conflict of interest; some possibility exists for all the players to gain (or to lose). Goal achievement can happen on all players' part without having to be taken from one or more players. Thus, potentially cooperative games are called **non-zero sum games**, and are popularly referred to as "win-win" situations (decisions).

Games, then, can be either 2-person or n-person, plus either competitive or cooperative. This yields 4 combined categories.

	2-person	n-person
competitive	2-person competitive games	n-person competitive games
cooperative	2-person cooperative games	n-person cooperative games

The simplest game decision problems to analyze and evaluate are 2-person competitive games. In this chapter we will stay within this category of 2-person zero sum games, first looking at those having a one strategy or pure solution and then turning to those requiring a mixture of choices for goal achievement. (In the next chapter we will take up the other category: non-zero sum games). To analyze and evaluate a game is, as with any other decision problems, to identify its type and to put its parts into a framework that allows practical reasoning to discover and justify the rational choice for each of the interacting players.

9.2 <u>Analyzing competitive interdependent decisions</u>

Here is a simple example of a 2-person competitive game that illustrates some of the concepts introduced above and also shows how game decisions are analyzed.

Imagine that two people, let's call them Rebecca (R) and Carol (C), have been separately invited to a party being given by a company hoping to hire one of them. Each is trying to decide what to wear to this important party. R wants to wear either a red or a grey outfit; the same for B, she is thinking of wearing a red or a grey outfit to the party. Now suppose that R and C know about each other and are aware that they are competing for the same desirable high-powered job. They have formed opposing goals about the outfit to wear. R knows that C will be at the party, and R wants to wear a different color outfit to the party than C wears, believing that this will outdo B in the eyes of the prospective employer. C on the other hand, knows that R will be at the party and would love to be dressed in the same color outfit as R, believing that by doing so she will outshine R at the party. Here is this decision problem framed in standard option-state-outcome form, and then put into a 2x2 game matrix.

R's goal: outdo C by wearing a different color outfit to the party than C wears

	option	state	outcome	utility (-10…0…10)
R	red outfit	C wears red →	R and C wear same color to the party	(-10)
		C wears grey →	R and C wear different colors to the party	(10)
	grey outfit	C wears red →	R and C wear different colors to the party	(10)
		C wears grey →	R and C wear same color to the party	(-10)

C's goal: outdo R by wearing the same color outfit to the party as R wears

	option	state	outcome	utility (-10…0…10)
C	red outfit	R wears red →	R and C wear same color to the party	(10)
		R wears grey →	R and C wear different colors to the party	(-10)
	grey outfit	R wears red →	R and C wear different colors to the party	(-10)
		R wears grey →	R and C wear same colors to the party	(10)

In this example of 2 rival party goers, the utility for achieving the goal is the maximum 10, and the disutility for losing the goal is the minimum -10. Each player (agent) has 2 strategies (options). Each player, R and C, are locked in a total conflict of interest in each one's decision about what color outfit to wear to the party. They are locked together in an intertwined decision situation because the payoff (outcome) of each player's options depends on the option the other player chooses. They have a total conflict of interest because goal achievement for one automatically means total goal loss for the other. In other words, each player has the power to promote the goal of the other player, but only by totally frustrating her own goal. Each player will achieve her own goal, but only at the expense of the other player's goal.

There is an efficient way to frame a 2-person, 2-option game decision using a 2x2 matrix.

C: wear same colors to the party

		wear red	wear grey	
R: wear different colors to the party	wear red	-10, 10	10, -10	← row
	wear grey	10, -10	-10, 10	← row
		↑ column	↑ column	

This matrix contains 4 cells. If R and C had 3 options each, we would have a 3x3 game matrix containing 9 cells. The <u>left-side number</u> in each cell is the utility/disutility of R's outcome (payoff). Looking at R's options you'll see two rows, one on top and the other on the bottom. From this point of view, C's options are the alternative states-of-the-world of R's options, yielding different outcomes. So, if R wears red and the state is that C wears red, the outcome for R in the upper-left cell is -10 (that is, R loses her goal). But if R wears red and the state is that C wears grey, the outcome for R in the upper-right cell is 10 (R gains her goal). Similarly, the bottom row is R's option of wearing grey. There are 2 states-of-the-world depending on what C decides to wear, yielding 2 outcomes for R, the bottom-left left one is that R achieves her goal (10), and the bottom-right left one is that she loses her goal (-10).

Now let's look at the 2 columns, one on the left and one on the right. These are C's options and from this point of view R's options are now the states-of-the-world that will yield different outcomes for C. The right side number in each cell is the utility/disutility of C's outcome. Going down on the left column, if C wears red and the state is that R wears red, C's goal will be achieved, the outcome for C in the upper-left cell is 10. But if B wears red and the state is that R wears grey, C's outcome in the lower-left cell is -10 having lost her goal. Similarly for the right column; C's option has 2 possible states depending on what R decides to wear, yielding 2 outcomes; the right top cell showing that C's goal is achieved, and the lower right cell showing that R has frustrated C's goal achievement.

This example of a 2-person competitive decision is a zero sum game. If you add up the outcome utility/disutility in each cell, the sum for each cell is zero.

In analyzing or framing game decisions in order to represent them in matrix form, here are the general rules to follow:

1) form the decision problem into a branching diagram with options, states and outcomes for each agent.
2) assign utility values to the outcomes from an appropriately wide interval scale, using the goal of each agent as the single criteria of evaluation.
3) to form the game matrix, make a row for each of the row agent's options: two horizontal rows if the row agent has two options from which to choose, three rows if there are three options, …, etc.
4) make a column for each of the column agent's options: two vertical columns if the column agent has two options, three vertical columns if the column agent has three options, …, etc.
5) in each cell, the utility/disutility of the outcome for the row options always goes on the left and the utility/disutility of the outcome for the column options always goes on the right.

Using the above two example – the two party goers trying to out-dress each other and the opening example of two students taking a TV break from studying – let's examine certain key features of competitive games.

1) By definition, these games are strictly and completely competitive decision situations. Once a person decides to become a player, it is rational to maximize gain and minimize loss of the payoff. A player is trying for the best outcome – as much goal achievement as the other player allows. The other player, being equally rational, tries to do the same. Caution: even though competitive games are decisions that involve more than one agent, they are individual decisions, not group decisions (social choices). Also, for each player the other player is not a stakeholder.

2) Competitive games have a fixed or constant total goal value, the total outcome utility does not increase or diminish; instead it becomes redistributed between the players as the result of their decisions. One player's gain must come from somewhere and the only source is the other player who now has a loss equal to that gain. One player in a zero sum game gains only at the expense of the other player. In addition to the two above examples, here are some other familiar examples. Suppose you and your friend apply for the same job. Only one person can be hired, so if you achieve your goal (get the job) then by the necessity of the situation your friend has lost her goal (she is not hired); your gain is your friend's loss. Likewise, if your friend gets the job, this means that you have lost that goal; her gain equals your loss, for it is the same job you were both competing for. The friend who was not hired may be happy that her friend, and not a stranger, got the job, but this is not the point; the idea is that goal achievement of one agent is necessarily frustrated by goal achievement of the other agent; one agent has to be disappointed about this, even if otherwise happy for the success of the other agent. Take two teams competing to win a sports event: if one team is ahead by, say, 10 points, the other team must be down by exactly 10 points. If the team that is down by 10 now gains 5 points, it has to be taken from the other team

that has now lost half of a 10-point lead. If one team scores a run, the other team has been made to give up a run. If the final score is 10 to 5, then the winning team has won by 5 points and the losing team has lost by 5 points. Take the case of business competition: if one store gains a new customer, the competition has lost that new customer (assuming the customer can't do business with both stores). If one salesperson closes a deal, a competing salesperson has now lost that deal. Finally, look at many forms of human interaction: if one person "gets her way", the other person has given in and so has lost "getting her way." In all such examples, by trying to maximize gain each player <u>necessarily</u> tries to inflict loss on the other player; and the more one player can make the other player lose the more that player gains. Thus, in a zero sum game players are necessarily opponents and do not cooperate with each other. If one or both players decide to stop competing and start cooperating (say, one tries to help the other by intentionally making an irrational choice), then by definition they are no longer in a competitive game.

3) We should distinguish two different ways that individuals compete in zero sum games. One way is illustrated by a criminal trial in which the goals of each legal team are opposed. One team (the prosecutor) has the goal: verdict of guilt:, the other team (the defense) has the goal: verdict not guilty. These two goals are *not* logically independent; they are not just different goals, they are logically incompatible goals. They are mutually exclusive in the sense that both cannot exist at the same time. If both agents had goals that were not logically opposed, goals that both could achieve – let's say, in the outfit example, R desired to wear a red outfits to the party and C wanted to wear a grey outfit – then by definition they are no longer in a zero sum game for there is the possibility that they can cooperate in a way that each gains their (new) goals. So, one way that agents in a zero sum game compete is for the agents to have opposite goals; one agent achieving his goal to any degree deprives the other agent from achieving his different (incompatible) goal to the same degree.

But what if the agents in a zero sum game have the same goal? What if they both desire the same thing? In the example two paragraphs above in which you and your friend apply for the same job, you both desire the same goal; you do not have conflicting goals. If two people want the same thing, doesn't this make them allies and not opponents? Well, not if the goal in question can't be divided or shared. In this case the agents are competitors not because they have opposite goals, only one of which can be achieved, but because more than one agent at a time can't achieve the same goal. If two agents could share the same goal – say, both you and your friend could be hired for the same job – then by definition it would not be a zero sum game and the agents could cooperate in a way that allowed both to have goal achievement.

4) Because players in a zero sum game try to make each other lose the payoff, it does not mean that they dislike each other or want to hurt each other out of hostility. They might be good friends sitting down to a friendly game of cards, checkers or chess. They might be teams that have the highest respect and admiration for each other and yet fiercely compete on the playing field. A criminal trial can be thought of as a decision problem: a zero sum game in which the players are the opposing defense and prosecution lawyers, the payoff is the jury verdict, and the judge makes sure that the "game" is played strictly by the rules of a fair trial. One side's loss is the other side's gain, but the opposing lawyers might be good friends who have the highest professional regard for each other and even admire each other's legal strategies during the trial.

EXERCISE: Analyze (frame) this decision problem according to the 5 steps given above.

It's Sunday afternoon and Roger has finished all his final exams; he wants to celebrate with friends in his dorm room tonight. His options: he can have 8 friends over for a wild time or just 3 friends for a lower-key party, depending on what his roommate Carl is going to do. If Roger has 8 friends over, he'll have the best time; and if he has 3 friend over it will still be fun but not as much. Carl, meanwhile, has 2 options; he still has 2 final exams Monday, one of which he'll need to prepare for tonight: either history for which he must still do some reading and will need a reasonably quiet dorm room, or advanced violin for which he must practice on his instrument. If Carl decides to study history with 8 of Roger's friends partying, it will be a lost evening; and he'll be able to get only a little of his reading done if Roger celebrates with just 3 friends. If, however, Carl practices violin, it will be a terrible evening for Roger and his 8 friends; they will leave early, giving Carl a full evening of good practice. If Roger has his 3 friends over, they will stay the whole evening, but their partying will have to be seriously toned down because of Carl's violin practice; Carl, however, will have only minor distractions from the 3 friends while he practices. What should Roger and Carl decide to do Sunday evening, given their options and these outcomes? (Suggestion: to make this a clear zero sum decision, use utility 10 and –10 for best and worse outcomes, and utility 5 and –5 for 2nd best and 2nd worse outcomes.)

9.2.1 Finding the rational choice in competitive interdependent decisions

Before attempting to solve competitive decision problems, we will first see why using expected utility won't work. Return to the above example of the two party goers who have opposing goals: R wants to wear a different color outfit to the party than C wears, and C wants to wear the same color outfit to the party that R wears, each trying to outdo the other in how they dress. Suppose R recollects past parties they both attended, and estimates that C wore a red outfit to 4 out of 6 and wore a gray outfit to the other two. Of course, C also knows what she wore to the last 6 parties she went to! Now R tries to make a rational choice by expected utility.

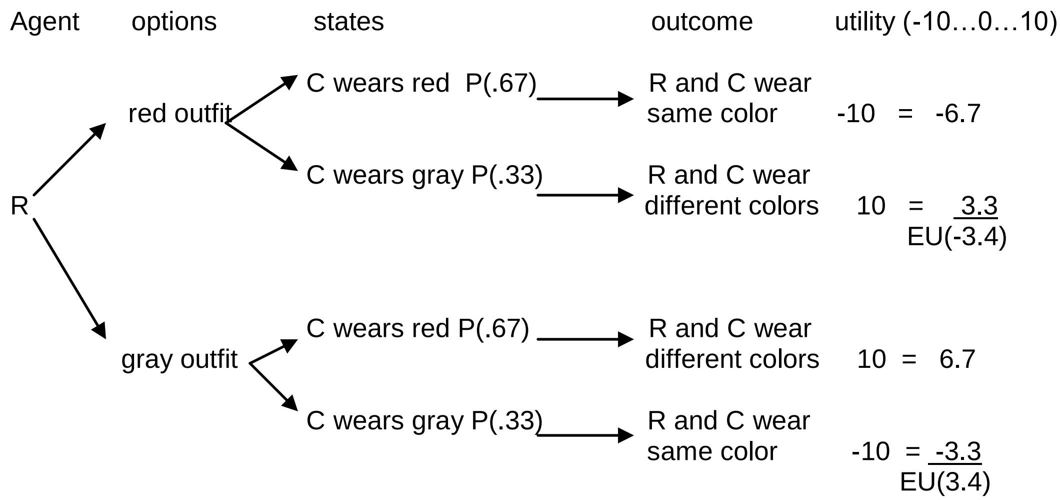

So, by the expected utility rule, the rational choice for R appears to be to wear a gray outfit to the party. But now add to this the fact that C knows these same probabilities and, being just as good at practical reasoning as R, will quickly realize her rational choice is to wear gray, for she expects R to wear gray and C's goal is to wear the same color outfit as R wears to the party.

But now R will expect C to wear gray (for C expects R to wear gray based on EU) and so the rational choice for R is to switch to a red outfit, given R's goal.

And C will expect R to make such a switch, and will likewise switch to a red outfit, which in turn will be expected by R who will now switch back to a gray outfit, which will be expected by C …. .

You clearly see the problem: solution by expected utility is useless for finding the rational choice in competitive decisions, for it makes the rational choice for each agent depend on an irrational choice on the part of the other agent. For example, sticking to a decision by expected utility on R's part in the competitive game we just discussed is irrational because, in effect, it hands C her goal, and in doing so R loses her goal. And if R shouldn't stick to a decision by expected utility, why use it to reach a decision?

The failure of solution by expected utility in zero sum games serves to highlight a very important principle of practical reasoning in these kinds of interdependent decisions: practical reasoning should be *strategic*. Each agent should assume that the other agent is equally rational, equally good at practical reasoning, and that they have *common knowledge*. This means that each agent assumes that each will use the same information as the other uses, if it will help the other to achieve the goal, and that each is aware that they making this assumption about each other. It would be irrational – as well as asking for trouble – to assume that your opponent in a competitive decision is irrational, or doesn't know what you know, and as a result will decide in such a way that hands you your goal. Yet this is what would happen in a solution by expected utility: one player's choice is "rational" only because the other player's choice is irrational. For each player in a game decision, it would be much better to have standards of rational choice and methods of practical reasoning that don't depend on the other player being irrational. Let's look at how each agent in a zero sum game can discover the rational choice independent of the other player's possible irrationality. To be sure that we are not relying on anyone's possible irrationality as our standard of a rational choice, we will accept the above principle of strategic rationality and assume that both players are equally rational, that each will do all the practical reasoning required to discover the rational choice, and that they have common knowledge of the relevant information concerning the decision that has to be made.

9.3 Competitive decisions: decision by dominance

Imagine that it's income tax time and the tax cheater (TC) realizes that this year he can cheat on his tax return in many little ways for a total of $2000 or cheat in a few ways for $225, but he can't do both. The first way of cheating on the return is time-consuming and challenging, while the second way is quick and easy to do. TC's goal, of course, is to pay as little taxes as possible; anything kept by cheating is a positive outcome for TC. The Government Tax Collecting Agency (IRS), on the other hand, wants to minimize tax cheating and collect as much owed taxes as possible; any taxes owed that the IRS loses to tax cheating is a negative outcome. But the IRS can't audit every return, for auditing is expensive and time-consuming. Let's suppose that this year the IRS can audit tax returns of people in TC's income bracket at a rate of 3 in 10 or audit at a rate of 7 in 10. Here we have 2 players with opposing goals whose decisions are interrelated. If he'll be audited TC will cheat the easy way for $225, but if he won't be audited, it makes sense, given his goal, to make the extra effort and cheat for $2000. The IRS, on the other hand, would find it worthwhile to audit TC at a 7 in 10 rate to catch $2000, but for $225 would rather audit at a 3 in 10 rate and put resources into other tax collection efforts. Each agent's outcome depends on the other agent's decision, and any money that goes to one agent is taken from the other agent. What's the rational choice for each?

We must first work out the utility of the outcomes, given the risks, and then structure this decision problem into a 2x2 matrix.

Goals: TC – keep maximum taxes
IRS – loose minimum taxes to cheating

Agent	options	states		expected monetary outcomes	utility (-10…0…10)
TC	cheat for $2000	IRS audits at 3/10 rate (.3 x $2000 means $600, on average, not gained)	=	$1400	10
		audit at 7/10 rate (.7 x $2000)	=	$ 600	6
	cheat for $225	IRS audits at 3/10 rate (.3 x $225)	=	$157.50	3
		audit at 7/10 rate (.7 x $225)	=	$67.50	2
IRS	audits 3/10	TC cheats for $2000 (.3 x -$2000)	=	- $1400	-10
		TC cheats for $225 (.3 x -$225)	=	-$157.50	-3
	audit 7/10	TC cheats for $2000 (.7 x -$2000)	=	-$600	-6
		TC cheats for $225 (.7 x -$225)	=	-$67.50	-2

In this example, whatever amount TC manages to keep in taxes by cheating is a loss for the IRS, and so the utility assignments must show this fact using both positive and negative values.

Now let's transform this analysis into a 2x2 matrix, using only utility values to represent outcomes. It does not matter which agent we make Row and which we make Column.

		Col: IRS	
		C1: audits at 3/10 rate	C2: audit at 7/10 rate
Row: TC	R1: cheat for $2000	10, -10	6, -6
	R2: cheat for $225	3, -3	2, -2

This is clearly a zero sum game, for the sum of Row's and Col's outcomes in each cell equals zero. What is the rational choice for Row (TC) if we assume that Col (IRS) will choose rationally (rather than irrationally)? Likewise, what is the rational choice for Col., independent of the fact that Row might make a mistake and choose irrationally? Let's start with Row; the first thing to look for is a *dominant strategy*. Recall that one option dominates another if its worst outcome is equal to or better than the best outcome of the other, and at least one outcome is better. Clearly, Row has a dominant strategy. The outcomes of cheating for $2000 are better than the outcomes of cheating for $225 no matter what Col does (10 is better than 3 if Col audits at a 3/10 rate, and 6 is better than 2 if Col does not audit at a 7/10 rate). The option of cheating for $225 is dominated and drops out. Thus, Row has discovered the rational choice without having to depend on Col choosing irrationally. For Row: choose (R1) over (R2).

How about Col? Does the IRS have a dominant strategy? Clearly yes; auditing taxpayers in TC's income bracket at a 7/10 rate is better (less loss) no matter what Row does (-6 is a better outcome than -10 if TC decides to cheat for $2000, and -2 is better than -3 if TC cheats for $225). The option of auditing taxpayers at a 3/10 rate is dominated and Col drops it as an option. The rational choice for Col has been discovered independently of Row's decision. For Col: choose (C2) over (C1).

There are three interesting points to note in this example.

1) Look at what happens if one player plays his dominant strategy and the other player chooses irrationally (plays the dominated strategy). The rational agent's outcome is improved and the other player's outcome suffers in equal measure as a result of the irrational choice. In a zero sum interactive decisions, the only way to help the other agent is by being irrational.

2) Both Row and Col achieve part of their goals, but Row clearly does better in this interactive decision than Col does. Col gets his 3rd best outcome while Row gets his 2nd best outcome. Such a zero sum game is **asymmetrical**: when agent's make rational choices, one does better in goal achievement than the other. (If rational choices resulted in agents ending up equal in goal achievement/loss the game is called **symmetrical**.) But by choosing rationally Col clearly keeps Row from achieving his best outcome. Likewise, by choosing rationally Row keeps Col from achieving his best outcome. As we just noted above, each agent can achieve his goal maximally only with the "help" of the other agent in the form of an irrational choice. The flip-side of this is that it is in the very nature of a zero sum game that being rational inflicts frustration – inflicts some degree of goal failure – on the other agent.

3) In a game, two options (or two outcomes) are in **equilibrium** if each is the best possible decision in response to the decision of the other player. Given the other player's choice, each player cannot do better by switching to another option. In the above example, options R1 and C2 (outcomes 6, -6) are in equilibrium. A game is said to be **stable** if it has equilibrium outcomes, stable in the sense that practical reasoning moves the players toward the equilibrium options from which there is then no reason to switch. (Note: the principle of equilibrium will grow in importance as we advance in game decision problems in this and subsequent chapters; be sure to get a good grasp of this concept.)

Suppose one agent has no dominant strategy. Suppose, to vary the above story to suit our purpose, that TC enjoys seeing the IRS waste time and effort auditing a small amount and so rates an audit of $100 cheating a utility of 4. Likewise, the IRS does not like wasting time and effort for an audit of $100 cheating, and rates it -4. In this case, the matrix would look like this.

		Col: IRS	
		C1: audits 3/10	C2: audit 7/10
Row: TC	R1: cheat for $2000	10, -10	6, -6
	R2: cheat for $100	3, -3	4, -4

In this case, Col has no dominant option. If Row chooses R1, Col should choose C2 for –6 is better than –10. But if Row chooses R2, Col does better by switching to C1 for –3 is better than – 4. What is the rational choice for Col? In game decisions, an important rule of practical reasoning is <u>not</u> to underestimate the rationality of one's opponent. To repeat the point that was stressed above: each player should respect the other by assuming that the other player is at least as rational as that player is. If we apply this strategic reasoning principle here, it is clear to Col that Row has a dominant option, so Col should not expect Row to be irrational and choose R2; Col should put itself in Row's shoes and ask: what would the rational choice be if I were in Row's position? For Row, R2 is dropped as a dominated option. So, for both Row and Col the decision is actually like this.

		Col: IRS	
		C1: audit 3/10	C2: audit 7/10
Row: TC	R1: cheat for $2000	10, -10	6, -6

It is now clear to Col what the rational choice is: C2, for –6 is better than –10 given Col's goal. The solution to this game, then, is (R1,C2).

Suppose that we have analyzed and framed a zero sum interactive decision, and the matrix looks like this (we are skipping a story and the steps of analysis that would lead to this matrix in order to focus on finding the rational choice).

		Col: C1	Col: C2
Row:	R1:	6, -6	-5, 5
	R2:	2, -2	-3, 3

In this example, Row cannot find the rational choice by dominance. If Col chooses C1, then Row should pick R1; but if Col chooses C2, then R2 is Row's rational choice. But if Row looks at this decision problem strategically from Col's perspective, Row sees that Col has a dominated option. Row must reason that Col will not choose C1, for Row should assume that Col is as good at practical reasoning as Row is. For both Row and Col, then, the decision problem becomes this.

		Col: C2
Row:	R1:	-5, 5
	R2:	-3, 3

It is now clear to Row that R2 is the rational choice, for –3 is not as bad an outcome for Row as –5. The rational choice solution for this interactive decision problem is (R2,C2).

2-person game decisions might have any number of options for Row or for Col. Suppose, varying our example above of the Tax Cheater and the IRS, that the IRS has a new computer-auditing program that it would like to test on someone not cheating, for it sometimes makes mistakes and fails to catch cheating. Suppose TC can cheat for $200 in one place on her tax return, or for $225 in another part of her return (but not both), or not cheat at all for she has heard about the new computer-auditing program. Suppose we did the decision analysis and it results in the following 3x3 matrix.

		Col: IRS		
		C1: old audit	C2: computer audit	C3: no audit
	R1: cheat for $200	1, -1	- 3, 3	5, -5
Row: TC	R2: cheat for $225	2, -2	-4, 4	7, -7
	R3: don't cheat this year	-3, 3	-5, 5	-10, 10

No matter how many options each player has, practical reasoning in strictly competitive decision situations proceeds in the same way: first try to eliminate any dominated options, each player assuming the other is equally rational in doing so. For Row, R3 is dominated by both R1 and R2; no matter what option Col chooses, Row will gain more of her goal (or loose less of it) with options R1 or R2. Thus, both Row and Col, assuming each other to be equally good practical reasoners and to have common knowledge, drop R3 as an option for Row. Thus, Col's hope limit (10) and Row's security limit (-10) are gone as possible outcomes. Col, meanwhile, sees that C3 is dominated by C1 and C1 is dominated by C2; Col would never choose C1 or C3 as long as C2 was an available option in the menu. Row sees this too and so both Row and Col drop C1 and C3 from Col's menu of options. With R3 and with C3 and C1 eliminated, it is clear that the rational choice for Col is C2. Row has followed Col's reasoning step by step and now reasons that R1 should clearly be chosen over R2. The rational choice solution, then, is: (R1,C2) giving Row an outcome of -3 and Col an outcome of 3.

The game matrix makes it easy to see that these 2 rational choices (R1,C2) are in *equilibrium*. Row would be irrational to choose another option, given Row's goal and given that Col chooses C2. Doing so would only hurt Row's goal achievement more than R1 does, and increase Col's goal achievement in equal measure. Likewise, Col would be irrational to switch to another option given Col's goal and given that Row chooses R1, for exactly the same reasons. If Row, perhaps out of fear of the new IRS computer-auditing program, were to choose R3, then Col should switch to C3, for this would move Col from outcome utility 3 to Col's hope limit outcome 10. But note that in this scenario, C3 would only be a "rational choice" because R3 would be an irrational choice. It is not that C3 would be a rational choice because it results from good practical reasoning; its "rationality" would depend directly on the "help" it gets from Row by way of the poor decision R3. Likewise, if Row expected Col to pick C3 (because Col expected Row to choose R3 out of fear of the new computer-auditing program)), then Row should switch from R3 to R1 and go from -10 to outcome utility 5. But Row would make a "wise decision" in this case only because Col was irrational in choosing C3, and not because R1 is the rational choice independently of Col's decision.

The general point here is that the ideal of a rational choice could not be a norm of practical rationality if it depended on other agent making bad decisions. Other people's worse stupidity does not turn my stupidity into wisdom.

(It should be noted that the example we have been working with, a cheating taxpayer, is a case of an agent having a goal that, by standards outside the field of practical reasoning, is illegal and probably – though not necessarily – morally wrong.)

EXERCISE:

1) Solve the following (already analyzed) decision problems by dominance. For each, explain the concept: outcomes in equilibrium.

a)

		Jack: wine	Jack: beer
Jill	wine	7, -7	0, 0
	beer	-1, 1	-2, 2

b)

		Plato: math	Plato: poetry
Aristotle	excess	-5, 5	-9, 9
	mean	2 , -2	10, -10
	deficient	-6, 6	4, -4

c)

		Corporation XYZ		
		radio ad	TV ad	newspaper ad
Corporation ABC	radio ad	0, 0	-12, 12	6, -6
	TV ad	10, -10	0, 0	2, -2
	internet ad	3, -3	2, -2	5, -5

2) Put the following decision problem in matrix form, and solve by dominance. Explain the concept of equilibrium outcomes in this problem.

You work part-time planning events for a busy local caterer. You plan for the meals, the entertainment, and even for decorations, so you have significant responsibilities. Your boss, unfortunately, is not a very pleasant person, always thinking that her workers are not doing enough, especially the part timers. It's Thursday evening and you have to work Friday, but hope to leave early for a three-day weekend trip you have been looking forward to. The earlier you can get away Friday the better: that's your goal. Any time before 3:00 that you can get away is a plus, but any time after 3:00 cuts into your trip. You know that your boss, however, would like you to work the full day, and more: that's her goal. There are three big catering jobs that have to be planned on Friday, and you will be assigned to plan for one of them: jobs A (a wedding reception), B (a 50th marriage anniversary), or C (a celebration to honor a soldier returning from military duty). Meanwhile, you can get a jump on the planning Thursday night and save yourself some time Friday by working up details for one of three event planning outlines: I, II, or III. Any prior prep work you can do Thursday night is time saved on Friday. Based on your experience, you estimate the hours saved or lost (using 3:00 as your base), depending on the job you will be assigned on Friday and the event-outline you prepare Thursday. If you work on outline I and are assigned job A, you leave 1 hour early (a plus for you), and your boss feels she looses an hour of work from you (a minus for her). If you prep I and are assigned job B, it delays you 4 hours (and gains her 4 hour of work from you). And if you prep I and are assigned C, you get to leave 3 hours early (and your boss is out 3 hour of your time). If you do outline II, here are the outcomes: for A you'll have to stay 2 hours past 3:00; for B you leave 1 hour before 3:00; for C you must work 1 hour past 3:00. Finally, If you prep event-outline III Thursday night, here are the outcomes: for job A you leave 2 hours early; for B you save 1 hour on Friday; for job C you get to leave a full 4 hours early. Using the hours gained/lost by you as utility values, what catering job is the rational choice for your boss to assign you Friday, and which event-outline is the rational choice for you the prep Thursday night?

9.4 Competitive decisions: decision by maximin reasoning

As we have seen above, an agent in a competitive decision situation might not have a dominant option. How would each agent find the rational choice if neither agent has a dominant option? Let's look at a simple 2x2 example in which one agent has a dominant option in order to illustrate finding the rational choice by maximin reasoning. We will then practice this method of practical reasoning on more complex examples without dominant options.

Suppose Miss Row wants to avoid Mr. Col, but Mr. Col wants to meet face-to-face with Miss Row (perhaps one is a bill collector and the other is trying to avoid receiving the bill, or one is a court officer trying to serve a summons and the other is trying to avoid the summons, or one is pressuring the other for a date and the other wishes to avoid having to say no). Row and Col can each go to a party to which each has been invited, or each can go to the movies. It would be very unpleasant for Row to be confronted by Col at the party. She knows that Col likes going to the movies much more than going to parties, and dealing with Col in the movies offers some degree of protection from Col's advances. Col, for his part, likes going to movies more than parties generally, but if it comes to meeting Row, he would prefer to do so at the party much more than at the movies. Given this background, let's suppose that we have analyzed and framed this decision problem into this 2x2 matrix.

		Col: (meet up with Row)	
		C1: go to the party	C2: go to the movies
Row: (avoid Col)	R1: go to the party	-10, 10	-5, 5
	R2: go to the movies	5, -5	-3, 3

It is easy to see that this zero sum game has a rational choice solution by dominance, but let's put this aside for now and not use dominance. How should each agent reason (other than by dominance) in order for each to discover the rational choice, assuming that they are equally excellent practical reasoners and have common knowledge of all information relevant to the decision? We might think that one way would be for each agent to find the best outcome and choose the option that contains it. The outcome utility 5 is Row's *hope limit*, so Row chooses R2. The outcome utility 10 is Col's best, so Col chooses C1. The result, however, is an irrational choice for Col ((R2, C1) yields Col an outcome of -5 disutility, Col's *security level* – the worst outcome Col could receive – not utility 10). Also, this would make Row's choice R2 rational only because of the "help" received from Col in the form of a bad decision. Clearly, "Choose the option containing the outcome having maximum utility" is a bad principle of practical reasoning in zero sum games, for it only works if an agent knows that the other agent will make an irrational choice.

Let's instead have each agent reason strategically. A nice way to envision this is to have each agent pretend to be the other agent. Each agent imaginatively puts herself in the other agent's decision situation and asks: how would I choose if I were competing against myself? What is the best I can do if I were my own competition trying to *minimize* the utility of my choice? Each agent is looking for, not the best outcome (maximum utility), but instead <u>the best of the minimum utilities</u>. In Chapter 8 in which decision under ignorance is covered, we called this the **maximin** method of finding the rational choice, the method of practical reasoning for finding the "best of the worse." Here are the steps applied to zero sum games.

1) For each option (going across, looking at left-side utilities) Row finds the outcome having minimum utility. For R1 it is -10, and for R2 it is -3. From these 2 minimum utilities, Row selects the maximum utility (or minimum negative utility/minimum disutility). For Row it is -3, Row's maximin outcome utility.

2) Likewise, for each option (going down, looking at right side utilities) Col finds the outcome having minimum utility. For C1 it is -5, for C2 it is 3. From these 2 minimum utilities, Col selects the maximum utility (or minimum negative utility/minimum disutility). For Col it is 3, Col's maximin outcome utility.

3) If both Row's and Col's maximin outcomes are in the same cell of the matrix, called the **saddle point** (because each utility in the cell is the "lowest" of all the outcomes in the options containing them, and if you were to draw it you would get a figure with a dip or valley that imaginatively depicts a saddle), the rational choice for Row and for Col is the option from each one's menu containing the saddle point. The rational choice solution to the above decision problem by maximin reasoning, then, is (R2, C2).

Here is the above matrix for the Miss Row—Mr. Col decision problem with the maximin solution added.

		Col: (meet up with Row) C1: go to the party	C2: go to the movies	
Row: (avoid Col)	R1: go to the party	-10, 10	-5, 5	-10
	R2: go to the movies	5, -5	saddle point -3, 3	-3 (Row's maximin)
		-5	3 (Col's maximin)	

As we did in the case of solution by dominance in the previous section, let's note points of interest about maximin reasoning before practicing this method of finding the rational choice on more complex decision problems. The following four points of interest are interrelated.

1) If one player plays his maximin strategy and the other player chooses irrationally (plays an option that does not contain the maximin outcome), the rational agent's outcome is improved and the other player's outcome suffers in equal measure as a result of the irrational choice. For

example, if Row chooses R2 and Col chooses C1, Col suffers a loss of goal achievement and "helps" Row improve goal achievement. This point is worth repeating: in a zero sum interdependent decision, the only way to help the other agent is by being irrational. One agent's rational choice guarantees, by necessity of the decision problem, that outcome utility will improve in the event that the other agent makes an irrational choice. In other words, there is no way to "out-smart" one's opponent; you can tell each other your decisions and the information would not help. If the decision was the rational choice, you could not possible do better by switching from your rational choice; and if your opponent told you she was going to choose an option that was not her rational choice, it would be irrational to believe her (how do you know you were not being set-up?) and switch away from your rational choice. If you did, and you were in luck with an outcome that was better than what your rational choice would have gained you, you made a "good" decision only because your opponent made a bad one, not because it was a good decision in its own right, on the basis of good practical reasoning. By the standards of rational choice theory, you would both be irrational.

2) Sometimes the best we can do by being rational in an interdependent decision situation is less than what the other agent does in goal achievement: the decision problem is asymmetrical. The reality of our situation should be recognized, even if we don't like it. Yet each agent keeps the other from achieving the hope limit. In the above decision problem, Row can only get her 2nd best outcome if Col chooses rationally, and Col can only get his 3rd best outcome, if Row makes a rational choice. To repeat an earlier point, it is in the very nature of a zero sum game that being rational inflicts frustration – inflicts some degree of failure – on the other agent. In the example we are working with, Row does this to a greater degree to Col than Col does to Row, but at least Col does this to some degree to Row.

3) In a competitive game, the two options resulting in the saddle point (or two outcomes in the saddle point cell) are in *equilibrium*; each is the best possible decision in response to the decision of the other player in the context of the decision problem. Given the other player's choice, each

player cannot do better by switching to an option that does not result in the saddle point. In the above example, options R2 and C2 (or outcomes -3, 3) are in equilibrium. A zero sum game containing a saddle point cell is stable, meaning that practical reasoning will move the agents toward the saddle point options from which there is then no practical reason for either alone to switch to a different option.

4) If a zero sum game can be solved by dominance, it can be solved by the maximin method. (Go back to the previous section where 4 zero sum games were solved by dominance and find the rational choice by the maximin method. You'll see that each has the same solution by these two methods of practical reasoning.) But the converse is not true; a zero sum game that is solved by the maximin method can't always be solved by dominance. This makes maximin a more powerful method of strategic practical reasoning than the method of dominance.

Let's now gain practice with some examples that require maximin reasoning to discover the rational choice that can't be solved by dominance.

Suppose Miss Row is doing a successful job avoiding Mr. Col. She has three options for the coming weekend: R1 – go to the beach, R2 – visit her parents, and R3 – stay home (which is near the beach parking lot). All the outcomes of these options have an even or better chance of avoiding Mr. Col (outcome utility 0 for an even chance, outcome utility 10 for her best chance). Mr. Col is increasingly frustrated with his failure to meet face-to-face with Miss Row (his best chance is outcome utility 0, and his other outcomes are all disutilities). Mr. Col can (C1) wait near Miss Row's parent's house, or (C2) wait near Miss Row's house, or (C3) try waiting for Miss Row in the large beach parking lot. Suppose we analyzed this game, assigned utilities on a scale (-10…0…10) to the expected outcomes, and the following matrix resulted.

Col: (meet up with Miss Row)

		C1: parent's house	C2: Row's house	C3: beach parking lot
	R1: beach	8, -8	8, -8	7, -7
Row: (avoid Mr. Col)	R2: parents	0, 0	10, -10	4, -4
	R3: home	9, -9	0, 0	1, -1

Neither Row nor Col can solve this decision problem by dominance; neither has a dominated option. Row's *hope limit* is outcome 10, but should she choose the option with that outcome (R2)? Col's *hope limit* is 0 outcome (only an even chance of meeting up with Row); should he go for one of the options containing this outcome?

The maximin solutions makes (R1,C3) the rational choices. The minimum outcome utility for R1 is 7, for R2 it's 0, and for R3 it's 0. The maximum of these three minima is 7. For Col's options we have minima C1 = -9, C2 = -10, C3 = -7. The maximum is -7. We see that there is a saddle point cell, one outcome cell in the matrix that contains both Row's and Col's maximin outcome utilities. The rational choice is (R1, C3), and these outcomes are in equilibrium (any other choice but the rational choice will decrease an agent's goal achievement and contribute in equal measure to the other agent's goal achievement, given the other agent stays with the rational choice).

Here is a final example of a 4x4 interdependent competitive decision problem requiring maximin reasoning to find the rational choice for each agent.

A college professor has to prepare to teach just one of four different equally important topics in a 75 minute class. It is the night before, and the class next day is the last one of the semester. This professor would like to keep the class for the full 75 minutes, getting in as much material as possible. The students in the class can prepare the night before to ask just one of four sets of general questions covering the four possible topics. The students would like to get out of class as soon as possible – at this point in the semester any time spend in class has negative utility! Depending on the topic the professor chooses and the question set the students prepare, the class could be finish up very quickly or continue for the full 75 minutes. Let the utility numbers represent minutes spent in class for each topic the professor could teach, depending upon the question set the students have prepared to ask. (If we didn't do this, we would have to analyze this decision problem into a branching diagram in order to assign utility values to outcomes, 4 states to each of the 4 options of each agent; this would be a good size diagram with 16 outcomes for each agent.)

Col: students

Row: Prof.		ask question set #1	ask question set #2	ask question set #3	ask question set #4
	teach topic #1	30, -30	12, -12	10, -10	40, -40
	teach topic #2	45, -45	25, -25	75, -75	15, -15
	teach topic #3	60, -60	30, -30	55, -55	35, -35
	teach topic #4	5, -5	30, -30	20, -20	40, -40

The students, of course, would like to prepare question set #1, for that option has the payoff of finishing class in just 5 minutes, but only if the professor decides to teach topic #4. If the professor expected the students to prepare question set #1, the professor would prepare topic #3 and keep the class for 60 minutes. The professor, on the other hand, would like to teach topic #2, for that option has the payoff of keeping the class in session for the full 75 minutes, but only providing the class asks question set #3. The students, of course, would prepare question set #4 and get out of class in 15 minutes, if they expected the professor to prepare topic #2. As you can see, Row and Col can chase each other all over this matrix if each thinks that a rational choice depends on the other agent making an irrational choice. Can Row eliminate any options by domination and thereby simplify the decision problem? Can Col? What topic is the rational choice for the professor to prepare the night before the last class, and what question set is the rational choice for the students to prepare? Why?

9.4.1 <u>Rational choice rules for decision by dominance or by maximin equilibrum</u>

In zero sum games, an agent should choose the dominant option, if there is one. In the absence of a dominant option, an agent should choose the maximin equilibrium option. Not every competitive decision problem can be solved by these two methods of practical reasoning, as we will see in the next chapter. Let's combine these two rational choice rules into one that contains this condition.

In any 2-person zero sum game, for any Row outcome x and any Col outcome y:

> (i) if the outcome pair x,y form a saddle point cell (or form an equilibrium) by dominance or by maximin reasoning, then options (Rx,Cy) are the rational choices; that is, for Row (Rx) is chosen over any other R option, and for Col (Cy) is chosen over any other C option.

EXERCISE:

1) Here is a well-known 2-person zero sum decision problem called "cut the cake".

Row wants more cake than Col, in fact Row would like the whole cake. Col likewise wants more cake than Row, the whole cake if possible. Row has to divide the cake and can cut it almost in half (but not perfectly so) with utility 1 for the slightly bigger piece, or can cut it very unequally hoping for the larger portion with utility 5. Col gets to pick his piece of cake and can pick 1st or let Row go 1$^{st.}$ Explain the rational choice for each. How might this zero sum game serve as a model for a rational division of a good like a piece of land or an amount of money when each party desires all of the good?

		Col: more cake than Row	
		C1: choose first	C2: choose last
Row: more cake than Col	R1: cut unevenly	-5, 5	5, -5
	R2: cut almost in half	-1, 1	1, -1

2) Select one of the following possible conflicts of interest and create a zero sum decision problem that can be solved either by dominance or by maximin reasoning. You should create a narrative that provides enough information to: (i) identify each agent's goal, (ii) identify each agent's options (at least 2 for each!), (iii) describe outcomes, (iv) assign utility values. Feel free to play with the utility values so that the decision problem can be solved by the methods of dominance or maximin equilibrium, and yet accurately represents the details of your story.

a) labor leader vs. management (conflict of interest over salaries increases)

b) illegal parker vs. parking ticket officer (conflict of interest over space to park)

c) government troops vs. rebels forces (conflict of interest over an area of land)

d) Corporation X vs. Corporation Y (conflict of interest over potential customers)

e) teacher's union vs. school board (conflict of interest over time in class)

<u>Sources and suggested readings</u>: (See Chapter 10 for references for both Chapters 9 and 10)