MAKING GOOD CHOICES: AN INTRODUCTION TO PRACTICAL REASONING

CHAPTER 10: PRACTICAL REASONING IN COMPETITIVE INTERDEPENDENT DECISIONS

This chapter continues the topic begun in Chapter 9: competitive (zero sum) game decisions and strategic practical reasoning. In Chapter 9 we covered decision by the methods of dominance and maximin reasoning. However, not every competitive decision problem can be solved by these two methods of practical reasoning. We now consider mixed strategy reasoning and an interesting form of common ignorance.

10.1 Competitive decisions: decision by mixed strategy

What to do if interdependent competitive decisions can't be solved either by dominance or by maximin reasoning? Let's start with this example.

Suppose a speeder wants to travel fast, say between home and work, on a daily basis. A patrol car wants to catch speeders, especially this one. The speeder can take the highway and do 30mph over the speed limit, or go the back roads at 15mph over the speed limit. Of course, the speeder would like to avoid a speeding ticket. The patrol car can cruise the highway or can cruise the back roads looking out for speeders. To be caught speeding at 30mph over the limit is the worse outcome for the speeder, but it is the best outcome for the patrol car. Conversely, to travel on the highway at 30mph over the limit and not get caught is best for the speeder, but the worse outcome for the patrol car. We can analyze this interdependent decision problem into these decision diagrams.

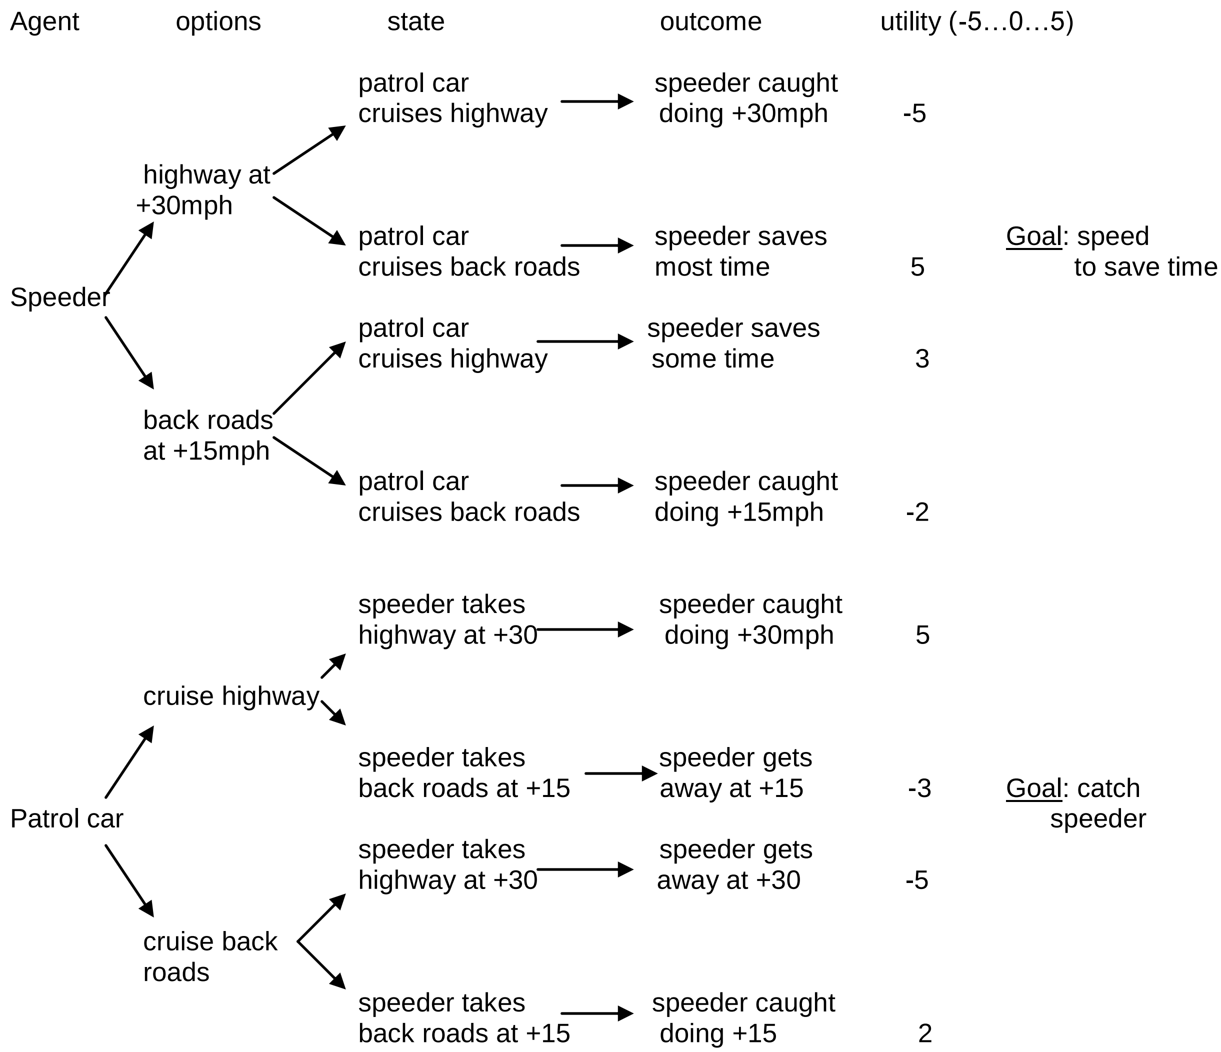

Transforming this analysis into a 2x2 outcome matrix gives us the following:

Col: patrol car

Row: speeder	C1: cruise highway	C2: cruise back roads
R1: highway at +30	-5, 5	5, -5
R2: backroads at +15	3, -3	-2, 2

This is a zero sum game, but neither Row nor Col has a dominant option. If we try to arrive at the rational choice for Row and for Col by maximin reasoning, we find that there is no saddle point cell. The minimum for R1 is -5, for R2 it is -2, and the maximin is -2. For C1 the minimum is -3, and for C2 it is -5, yielding a maximin of -3. No cell contains the outcome pair (-2, -3) (and if there were one the agents wouldn't be in a zero sum game!). There is no equilibrium cell, and this game is **unstable** in the sense that the agents, given their goals, are always better off switching and so end up "chasing" each other around and around the matrix. If Row chooses his hope limit option R1 (hoping for a 5 outcome), Col should prefer C1, for that option would give Col his hope limit (outcome 5) and hurt Row with Row's security limit (outcome -5). But if Col chooses C1, Row should switch to R2 (go from a -5 to a 3 outcome). But if Row switches from R1 to R2, Col should switch from C1 to C2 (go from -3 to 2). But if Col chooses C2, Row should switch back to R1,…, and so forth. Strategic practical reasoning does not move the agents toward one option and one outcome cell; instead the agents systematically alternate back-and-forth between options, the better outcome always resulting from switching to the other option.

In such unstable interdependent decisions, what form of strategic practical reasoning will lead to the rational choice? There are two parts to this practical reasoning. (1) One important principle is **strategic ignorance**. An agent in such a decision problem should withhold information from one's opponent. This was not necessary in competitive decisions that are stable; recall that agents can announce their rational choices to each other in stable games and it would not matter. However, withholding a decision from one's opponent is central to goal achievement in games that are unstable. What is the best way for an agent in such a decision situation to keep her intentions from his opponent, that is, to keep his opponent guessing as to how she will choose, and yet gain as much of the goal as possible? For present purposes let's think of an unstable competitive decision as one that the agents are playing repeatedly. Adopting a strategy of mixing options, switching – not systematically back-and-forth between options – but *unpredictably*, is certainly the best method to keep the opponent guessing about which option an agent will be choosing next.

(2) A second important principle is *proportion*. If an agent switches between options in a certain proportion, it will gain the agent more of the goal than any other proportion, assuming the other agent is likewise rationally doing the same thing using a mixed strategy. These two principles give us the

method of solution by optimal mixed strategies: unpredictably switching among options keeping to a certain ratio or proportion.

The well-known child's game rock-paper-scissors provides a handy example to illustrate these two practical reasoning requirements for rational choice in unstable games: (1) unpredictability of choice at (2) a suitable proportion. The rules of this game are: scissors cuts paper (win for scissors, loss for paper); paper covers rock (win for paper, loss for rock); rock breaks scissors (win for rock, loss for scissors). Each player must decide on and simultaneously indicate to each other one of these three possibilities. The matrix looks like this

		Col: scissors	Col: paper	Col: rock
	scissors	tie	w, l	l, w
Row:	paper	l, w	tie	w, l
	rock	w, l	l, w	tie

Suppose two children play this game for a penny a play and that the penny represents outcome utility: receive a penny for each win, pay a penny for each loss, neither gain nor pay a penny for a tie. The payoff matrix (with minimum outcomes indicated) looks like this:

		Col: C1 scissors	Col: C2 paper	Col: C3 rock	Row's minima:
	R1:scissors	0, 0	1, -1	-1, 1	-1
Row:	R2: paper	-1, 1	0, 0	1, -1	-1
	R3: rock	1, -1	-1, 1	0, 0	-1
	Col's minima:	-1	-1	-1	

What is the rational choice for each player in this game? Clearly, there is no one option for Row or one option for Col that would be each player's rational choice; no option dominates, and maximin reasoning can't apply. Notice that as soon as one player could predict a pattern of choice on the part of the other player, that player could use this information to form a degree of confidence about what choice to expect and adjust his decision accordingly to gain pennies. Even if there were no outright pattern one agent perceived in the repeated choices of the other player, if the other player favored an option (say, liked to choose paper with more frequency than the other options) this could be used to increase that one agent's frequency of wins. So, in order for each agent to end up with the same amount of pennies, each must:

1) avoid any pattern of choosing among options that might be detected; that is, keep the opponent guessing about the next choice by the strategy of mixing options in an unpredictable way (there should be *common ignorance*, not just *ignorance in common*!), and

2) distribute or proportion the choices so that (in the scissors-paper-stone game) each option receives 1/3rd play; that is, do not favor any option or neglect any option that would allow the opponent more gain than the absolute minimum he would receive given your best play.

This is a good example of a symmetrical game. Because the two players are equals (that is, each player has the same menu of option-outcome strategies to choose from), if these two rules of practical reasoning are followed in the game of rock-paper-scissors, each child at the end of the day will come out equal, one will not come out ahead in pennies. Each will have held the other to minimum winnings. Each will have made the rational choice – not of one option – but of the strategy of choosing the right proportion or mix of options in an unpredictable way.

10.2 Mixed strategy concepts

Before returning to the speeder/patrol car decision problem, let's define some concepts that are central to unstable competitive decision problems.

Optimal mixed strategy – this is the proportion or distribution of choices among options for an agent to make in a mixed strategy game that will guarantee the agent a minimum loss (or a minimum gain), given that the other agent chooses as rationally as possible (that is, given that the other agent plays his optimal mixed strategy). Playing the optimal mixed strategy (OMS), then, means that an agent may do better (if the other agent departs from his OMS), but cannot do worse than a certain minimum payoff or goal achievement. Equivalently, the OMS is guaranteed to holds the opponent to a minimum goal achievement that can be increased only by an agent departing from his OMS. In the rock-paper-scissors game, the OMS for each agent is 1/3rd for each option yielding a minimum payoff of equal wins and losses.

Fair – any zero sum mixed strategy game in which, if each agent plays their OMS, the result for each agent is zero. In a fair competitive decision situation, the rational choice yields each agent zero gain and zero loss of the goal. The rock-paper-scissors game is fair. Assuming that the goal is to win pennies, if each child started with 10 pennies and each played the OMS, at the game's end each child will come away with 10 pennies (in theory, that is; in reality they may have to stop playing at a moment when one child is ahead by one or more pennies, but this does not make this game unfair or biased). Note that "fair" is closely related to the concept of a *symmetrical* game; it should not matter to the agents who takes the row and who takes the col position, for they are equal in goal achievement (if each is making rational choices). But "fairness" in an OMS game might not mean the game is morally just. Justice might require that one agent gains more of the goal than other agents if, say, one was more deserving or, perhaps, more needy. "Fair" in this context means "equal," and "equal" might or might not be just when viewed by moral standards.

Biased – any zero sum mixed strategy game in which, if each agent plays their OMS, the result is a gain in goal achievement greater than zero for one agent and a goal loss for the other agent in an equal amount. A row biased game favors Row's OMS (yielding the row agent more goal achievement than the column agent), while a column biased game favors Col's OMS (yielding the column agent on average more of the goal than the row agent). In a biased competitive decision problem, the agents can have zero gain/loss only if at least one chooses irrationally. (Again, this is in theory; in reality, the agents might have to quit the decision situation at a moment when neither is ahead.) Note that "biased" is closely related to the idea of an *asymmetrical* game; it should matter very much to the agents who occupies the row and who occupies the column position, for these represent unequal goal achievement if each is playing rationally. However, “biased” does not necessarily mean that the game is morally unjust. “Biased” means “unfair” in the sense of “unequal”; by moral standards it could be perfectly just that one agent gains more of the goal at the other agent’s expense, for example if one agent were more deserving or, perhaps, more needy of the goal.

Let’s now return to speeder/patrol car problem and apply the practical reasoning principles for unstable zero sum decisions. We will not be explaining or calculating the exact OMS for each agent. There is a relatively simple formula for doing this, but presenting it would take us away from the general principles of practical reasoning and rational choice that we are focusing on. But the following will serve to reinforce the general ideas presented above in the scissors-paper-stone example that an unpredictable proportion of switching among options is the rational choice.

Col: patrol car

Row: speeder	C1: cruise highway	C2: cruise back roads	
R1: highway at +30	-5, 5	5, -5	5
R2: backroads at +15	3, -3	-2, 2	10
	7	8	

Let’s say that Row’s OMS is to distribute choice between options R1 and R2 in a 1/3 to 2/3 proportion. (We get this by dropping the negative (disutility) sign, adding the two outcome utilities for R1 = 10, adding the two outcome utilities for R2 = 5; we <u>switch</u> these and get a 5/15 or 1/3, (.33) for R1, and a 10/15 or 2/3 (.67) proportion for R2.) Likewise, suppose that Col’s OMS is to distribute choice between options C1 and C2 in a 7/15 (.47) to 8/15 (.53) proportion. (We get this by dropping Col’s negative (disutility) signs, adding the outcome utilities for C1 = 8, adding the outcome utilities for C2 = 7; <u>switching</u> these gives Col a 7/15 frequency of choice for C1, leaving an 8/15 frequency of choice for C2.) The general idea you should see here is that the result would be that Row achieves a certain <u>average outcome utility</u>, Row’s portion of his goal (rounded off it will be .66 in this game, the sum of (-5

x .47) + (5 x .53) and (3 x .47) + (-2 x .53)). Similarly, Col would achieve an average outcome utility, Col's portion of his goal (rounded off, -.66 in this game, the sum of (5 x .33) + (-3 x .67) and (-5 x .33) + (2 x .67). Because this is a zero sum game, the sum of Row's and Col's OMS values will be zero. As you can now see, this is a Row biased game when each agent makes rational choices (OMS):
(R = .66, C = -.66).

Given these two OMS's each agent must now make sure that choices are distributed in these proportions in an *unpredictable* way. How might this be done? Speeder has a 1/3 R1 to 2/3 R2 distribution, so she can perhaps put 5 red and 10 white pieces of paper in a container, shake them up and draw one; red its highway, white its back roads. Notice that this mixed strategy works whether speeder has to make a rational choice just once or has to make a rational choice repeatedly, say on a daily basis going between home and job. Also notice the interesting element of ignorance; the agent herself does not know beforehand and can not predict what option she will be "assigned" to choose by her selection device. Even a "truth serum" could not get it out of her! This is intentional ignorance as part of the method of practical reasoning within games: **strategic ignorance**.

How about the patrol car? How will the distribution of choosing options C1 at a 7/15 rate and C2 at a 8/15 rate be made unpredictable? Again, a random selection mechanism is a handy way to achieve this. The patrol car driver might put 70 white and 80 red slips of paper in a box, shake them up and draw one each day. Clearly, neither the patrol car driver nor the speeder can predict whether the patrol car will cruise the highway or the back roads on any given day. Each agent, in assuming that the other is rational, assumes that they have common knowledge about this mutual inability to predict (a form of "common ignorance"!) how each will choose.

10.2.1 Pause for questions

You might protest at this point and think: if the patrol car will be cruising the highway roughly 7/15 of the time, a smaller fraction than 8/15, isn't it rational for speeder to stay away from the back roads and stick to speeding on the highway at 30mph over the limit? Yes, but remember that the patrol car is equally rational (the assumption of common knowledge) and will quickly discover this pattern of choice. The patrol car will realize that speeder has eliminated R2 as an option, thus creating a new decision problem very easy for the patrol car to solve in its favor and to the speeder's loss. Also, by choosing R1, speeder gives up the outcome utility 3 which would have resulted at least some of the time from choosing R2.

OK; then why wouldn't speeder drive back roads 7/15 of the time and only speed on the highway 8/15 of the time, the exact reverse of the patrol car's mixed strategy? Well, if the patrol car is patrolling the back roads 8/15 of the time, the more the speeder uses the back roads at a 7/15 rate, the more the patrol car will catch him speeding, and the same for the highways rates of choice. This same poor choice would happen if the patrol car tried to adapt to the speeder's 1/3 to 2/3 mix of choices. You can see the general problem: trying to "out-smart" the opponent is trying for goal achievement in a way that relies on the other agent's irrationality. An agent might get lucky now and then, but it is clearly not good strategic practical reasoning to try to make one's choice "rational" by making it depend on the other agent's bad decisions. Much better to respect one's opponent as a fully rational agent (that is: assume common knowledge) who will play his OMS in an unstable competition with you, against which you will only lose more of your goal by departing from your OMS than you will achieve by sticking to it.

At this point, let's return briefly to Chapter 9 and consider the opening example of two party-goers who are trying to outdo each other in how they dress for the party. You will see from the matrix that neither agent has a dominant option, nor can their decision problem be solved by the maximin method. They are in an unstable competitive game and so each should use an unpredictable mixed strategy. As in the scissors-paper-stone game, the proportion of choice Row assigns to R1 and R2 should be equal, and the same for Col. You can see this by noticing that each has the same menu of options and that each option, unlike the case of the speeder/patrol car, yields equal payoffs. This, then, is a fair game – each agent achieving (and loosing) her goal an equal number of times (at least in theory!) providing each plays her OMS.

10.3. Summary

Here is a summary of the strategic practical reasoning steps for analyzing and evaluating zero sum interdependent decisions (Chapters 9 and 10).

1. For each agent, frame the decision into the standard option-outcome branching diagram and, using the goal as a single criterion, assign utility/disutility values to outcomes from a wide enough interval scale.

2. Construct an outcome matrix, making sure that each cell sums to zero.

3. Eliminate any dominated options for Row and for Col. This might solve the game and the rational choice for each agent will be clear.

4. If not, find the maximin for Row and for Col. If there is a saddle point outcome cell, this will be an equilibrium outcome pair that solves the game and the rational choice for each will be clear.

5. If there is no maximin equilibrium, the game is unstable and the rational choice is for each agent to play her OMS (that is: estimate the proportion of switching among options that yields the highest average utility, keeping the pattern of switching unpredictable).

Note: every 2-person competitive game has a rational choice by the methods of either dominance, maximin reasoning, or OMS.

6. If the game is biased against an agent, the rational choice is to opt out as soon as this option becomes available; opting out will dominate any OMS biased against an agent. (Note that if an agent desires to remain in a game biased against the agent's goal achievement, say because the agent is having fun or likes competing or gets to meet new people or whatever, then it is not a zero sum game and the true goal should be made clear and its value reflected in the outcome utilities.)

Before turning to our next topic, potentially cooperative interdependent decisions, let's highlight three of the methods of practical reasoning that we have used to discover the rational choice in zero sum games. These same three methods will be at the center of the practical reasoning used in the next chapter. We saw that the important method of expected utility would not work to solve zero sum games, and this method will not be used for potentially cooperative games in the next chapter either. Also, for reasons that will become clear in the next chapter, we will not use the method of mixing options to solve cooperative decision problems. There are three methods of practical reasoning by which a rational choice is discovered that we will carry over to the next chapter. 1) Dominance -- an agent should eliminate dominated options and should choose a dominant option, if any. 2) Maximin -- an agent should eliminate options and choose by maximin reasoning. 3) Equilibrium -- agents should choose the option whose outcomes are in equilibrium (from which it would be irrational to be the only agent to switch).

EXERCISES:

For the following decision problems, explain which are stable and which unstable. For those that are stable, can it be solved by the method of dominance or maximin reasoning? For those that are not stable, estimate if the game is fair or biased.

a)

cat \ mouse	run	stay
spring	-5, 5	10, -10
stalk	0, 0	2, -2

b)

he \ she	leave	stay
drink	- 8, 8	3, -3
smoke	- 4, 4	6, -6

d)

player \ dealer	ace	king	jack
bid	7, -7	5, -5	-10, 10
hold	-1, 1	0, 0	-2, 2
quit	-7, 7	-6, 6	-5, 5

e)

student \ teacher	quiz	lesson
play	0, 0	-5, 5
sleep	3, -3	1, -1
study	8, -8	-4, 4

c)

cops \ robbers	walk	run
walk	2, -2	-2, 2
run	-2, 2	2, -2

f)

patient \ nurse	respond	ignore
ring	7, -7	-5, 5
yell	-6, 6	0, 0
silence	1, -1	-1, 1

Sources and suggested readings:

Chapters 9 and 10 draws on material from Davis (1983) Chapters 1 - 3, Luce and Raiffa (1957) Chapter 4, Mullen and Roth (2002) Chapter 8, Rapoport (1966) Chapters 3 - 7, Resnik (1987) Chapter 5, and Straffin (1993) Chapter 1. Davis has become a classic and is especially recommended for its non-technical presentation and range of vivid examples. Straffin is highly recommended to anyone with an interest in mathematics or wishing to pursue either the full method for finding OMS or the shortcut used in this chapter. Allingham (2002) Chapter 5 offers a very compact tour of zero sum games as well as several uncluttered examples. Casti (1996) Chapter 1 offers a clear popular introduction to games with a focus on the centrality of von Neumann's minimax theorem, while Resnik presents an accessible proof of this theorem (called the maximin theorem) for the 2x2 case. Poundstone (1992) Chapter 3 is an engaging popular historical approach to games presented in a lively style.