Tutorial 3: Dynamic Games (Solutions)

Exercise 1: Formulating an extensive form game

> *The chef, waiter, and manager of a restaurant have to decide on a menu. The chef chooses between fish and meat, the waiter between red wine and white wine. The chef prefers fish to meat, the waiter prefers red wine to white wine; they all accept that white wine and meat, and red wine and fish are not acceptable options. The manager prefers meat. The chef and waiter make their choices consecutively—the role of the manager is to decide who chooses first.*

Model this as an extensive form game, and apply Zermelo's algorithm to determine a subgame perfect equilibrium.

Solution:

The preferences over the outcomes could be given by (interpretation straightforward):

$$\begin{aligned} \text{Manager:} &\quad \text{mr} \sim \text{mw} \succ \text{fr} \sim \text{fw} \\ \text{Chef:} &\quad \text{fw} \succ \text{mr} \succ \text{mw} \sim \text{fr} \\ \text{Waiter:} &\quad \text{mr} \succ \text{fw} \succ \text{mw} \sim \text{fr} \end{aligned}$$

or, allowing for an ambiguity in the wording of the exercise, as

$$\begin{aligned} \text{Manager:} &\quad \text{mr} \succ \text{fw} \succ \text{mw} \sim \text{fr} \\ \text{Chef:} &\quad \text{fw} \succ \text{mr} \succ \text{mw} \sim \text{fr} \\ \text{Waiter:} &\quad \text{mr} \succ \text{fw} \succ \text{mw} \sim \text{fr} \end{aligned}$$

For the remainder of the analysis, we adopt the former interpretation, but the latter is equally defensible. The game tree of the extensive form game is then given in Figure 1.

Let the (pure) strategies be given by (interpretation straightforward):

$$\begin{aligned} \Sigma_{\text{Manager}} &= \{\text{chef}, \text{waiter}\} \\ \Sigma_{\text{Chef}} &= \{\text{fff}, \text{ffm}, \text{fmf}, \text{fmm}, \text{mff}, \text{mfm}, \text{mmf}, \text{mmm}\} \\ \Sigma_{\text{Waiter}} &= \{\text{www}, \text{wwr}, \text{wrw}, \text{wrr}, \text{rww}, \text{rwr}, \text{rrw}, \text{rrr}\} \end{aligned}$$

Remark: The corresponding game in *strategic* from is given in Figure 2 and its Nash equilibria in Figure 3.

Figure 4 then gives the backwards induction solution: $(\text{ffm}, \text{wrr}, \text{waiter})$. It is important that not only the *path* but the decisions at *every* decision node are specified.

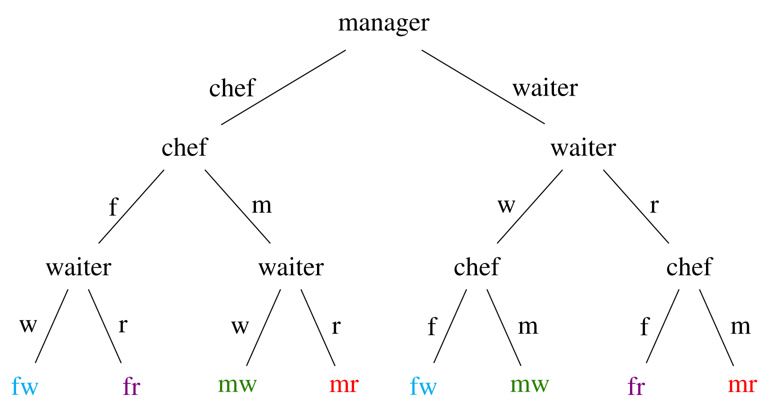

Figure 1: The extensive form of the game described in Exercise 1

Chef first

	w^3	wwr	wrw	wrr	rww	rwr	rrw	rrr
fff	fw	fw	fw	fw	fr	fr	fr	fr
ffm	fw	fw	fw	fw	fr	fr	fr	fr
fmf	fw	fw	fw	fw	fr	fr	fr	fr
fmm	fw	fw	fw	fw	fr	fr	fr	fr
mff	mw	mw	mr	mr	mw	mw	mr	mr
mfm	mw	mw	mr	mr	mw	mw	mr	mr
mmf	mw	mw	mr	mr	mw	mw	mr	mr
mmm	mw	mw	mr	mr	mw	mw	mr	mr

Waiter first

	w^3	wwr	wrw	wrr	rww	rwr	rrw	rrr
fff	fw	fr	fw	fr	fw	fr	fw	fr
ffm	fw	mr	fw	mr	fw	mr	fw	mr
fmf	mw	fr	mw	fr	mw	fr	mw	fr
fmm	mw	mr	mw	mr	mw	mr	mw	mr
mff	fw	fr	fw	fr	fw	fr	fw	fr
mfm	fw	mr	fw	mr	fw	mr	fw	mr
mmf	mw	fr	mw	fr	mw	fr	mw	fr
mmm	mw	mr	mw	mr	mw	mr	mw	mr

Figure 2: Strategic game corresponding to the game in extensive form of Exercise 1

Chef first

	w^3	wwr	wrw	wrr	rww	rwr	rrw	rrr
fff	fw	fw	fw	fw	fr	fr	fr	fr
ffm	fw	fw	fw	fw	fr	fr	fr	fr
fmf	fw	fw	fw	fw	fr	fr	fr	fr
fmm	fw	fw	fw	fw	fr	fr	fr	fr
mff	mw	mw	mr	mr	mw	mw	mr	mr
mfm	mw	mw	mr	mr	mw	mw	mr	mr
mmf	mw	mw	mr	mr	mw	mw	mr	mr
mmm	mw	mw	mr	mr	mw	mw	mr	mr

Waiter first

	w^3	wwr	wrw	wrr	rww	rwr	rrw	rrr
fff	fw	fr	fw	fr	fw	fr	fw	fr
ffm	fw	mr	fw	mr	fw	mr	fw	mr
fmf	mw	fr	mw	fr	mw	fr	mw	fr
fmm	mw	mr	mw	mr	mw	mr	mw	mr
mff	fw	fr	fw	fr	fw	fr	fw	fr
mfm	fw	mr	fw	mr	fw	mr	fw	mr
mmf	mw	fr	mw	fr	mw	fr	mw	fr
mmm	mw	mr	mw	mr	mw	mr	mw	mr

Figure 3: The pure Nash equilibria (in colour) and the subgame perfect Nash equilibrium (encircled) of the game of Exercise 1.

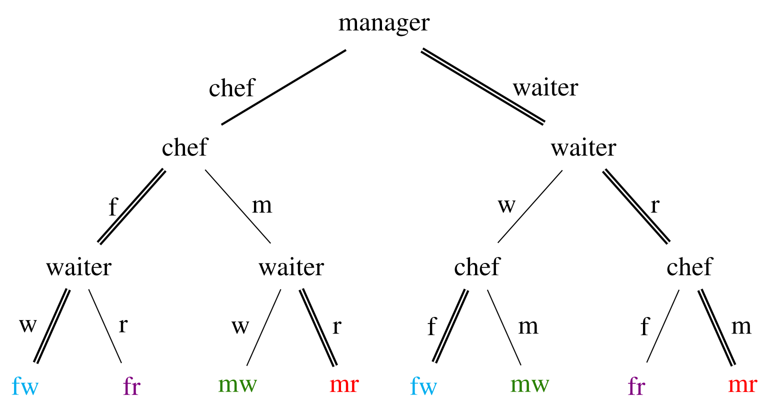

Figure 4: The backwards induction solution of the game in Exercise 1

Exercise 2: Formulating an extensive form game Consider a 2 player game in which player 1 can choose A or B. The game ends if she chooses A, while it continues to player 2 if he chooses B. Player 2 can then choose C or D with the game ending if C is chosen, and continuing again to player 1 if D is chosen. Player 1 can then choose E or F , with the game ending either choice.

(a) Model this as an extensive form game.

(b) How many pure strategies does each player have?

(c) What are the subgames of this game?

(d) Suppose that choice A gives utilities (2, 0), choice C gives (3, 1), choice E gives (1, 1), and F gives (1, 2). Then what are the Nash equilibria of the game? What SPNE outcome(s) does Zermelo's algorithm yield?

Solution:

(a) See the game tree in Figure 5.

(b) The players' pure strategies are given by:

$$\Sigma_1 = \{AE, AF, BE, BF\} \qquad\qquad \Sigma_2 = \{C, D\}$$

Hence, player 1 has four pure strategies and player 2 two.

(c) See figure.

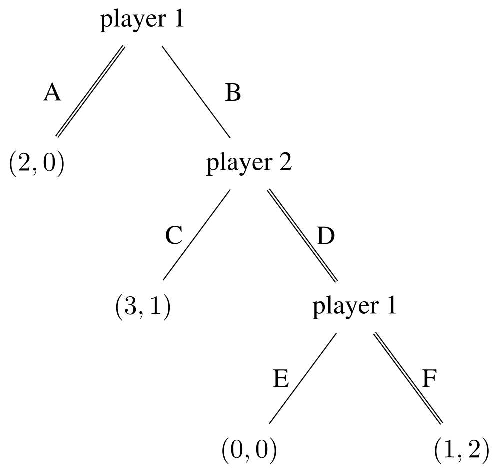

Figure 5: The extensive form of the game described in Exercise 2. Double edges indicate the choices made in the subgame perfect Nash equilibrium.

(d)

	C	D
AE	$2, 0$	$\underline{2, 0}$
AF	$2, 0$	$\underline{2, 0}$
BE	$\underline{3, 1}$	$0, 0$
BF	$3, 1$	$1, 2$

The pure Nash equilibria are: $(AE, D), (AF, D), (BE, C)$. By Zermelo's algorithm, (AF, D) can be seen to be the only the subgame perfect Nash equilibrium of this game:

Exercise 3: Formulating an extensive form game Consider the following scenario: Alice and Bob have a few drinks in The Turf, and then discover two $£1$ coins on the floor. Having been buying drinks in The Turf, they are both very short of money, and each would like to take home as much of the $£2$ as possible. They agree the following game for deciding how to share the money:

- Alice proposes a division of the money, that is, a deal of the form (x, y), where x and y are natural numbers such that $x + y = 2$, with the semantics that Alice takes $£x$ and Bob $£y$ pounds.
- Bob can either accept Alice's proposal or reject it:

- If Bob accepts the proposal, then the deal is implemented: Alice takes $\pounds x$ and Bob takes $\pounds y$. The game then ends.
- If Bob rejects the proposal, then he must make a counter proposal of the form (x', y'), with the same semantics as above— Alice must then accept or reject the proposal.
    - If Alice accepts the proposal, then it is implemented, as above, and the game ends.
    - If Alice rejects the proposal, then they give the $\pounds 2$ to the impoverished professor standing by the bar, and each goes home with nothing; the game then ends.

Note that they don't have any change (because they spent all their money on drinks), and so there are only three candidates for a proposal (x, y), that is, $(0, 2), (1, 1), (2, 0)$.

(a) Formulate this scenario as an extensive form game (assuming utility is monetary reward).

(b) Identify a Nash equilibrium of the game that maximises Alice's payoff and a Nash equilibrium that maximises Bob's payoff.

(c) What general conclusions can you draw about this game as a mechanism for dividing resources?

Solution:

(a) See Figure 7.

(b) Alice: propose $(2, 0)$ in the first round and threaten to reject every proposal by Bob in the third round unless he proposes $(2, 0)$. Then, Bob gets nothing no matter what he proposes and won't want to deviate from accepting Alice's proposal (or proposing) in the second round.

Bob: Always reject Alice's proposal and propose $(0, 2)$ in every case. Alice is then guaranteed a payoff of zero no matter which strategy she adopts. Thus, she has no incentive to deviate from always accepting Bob's proposal in the third round.

Remark: It is important to note that the reasoning is about regular Nash equilibria and not about *subgame-perfect* Nash equilibria.

(c) If Alice and Bob can make binding commitments, either player can try to "extort" the other. But such an extortion will only be effective if the extorted player is courteous and chooses in favour of the extortionist.

Looking at the SPNE of the game, there is one that yields $(1, 1)$, provided that Alice will give (out of spite) the two pounds to the professor when being offered $(0, 2)$ in the second round. If she is more charitable towards Bob, there is also an SPNE which gives Bob the two pounds. Seen thus, the mechanism appears to be biased towards Bob.

In general the mechanism would work best if Alice would be slightly spiteful (but this is not modelled in the game and cannot be banked upon.)

Another observation is that the analysis of the game much depends on whether you consider Nash equilibrium or SPNE as the correct solution concept to apply.

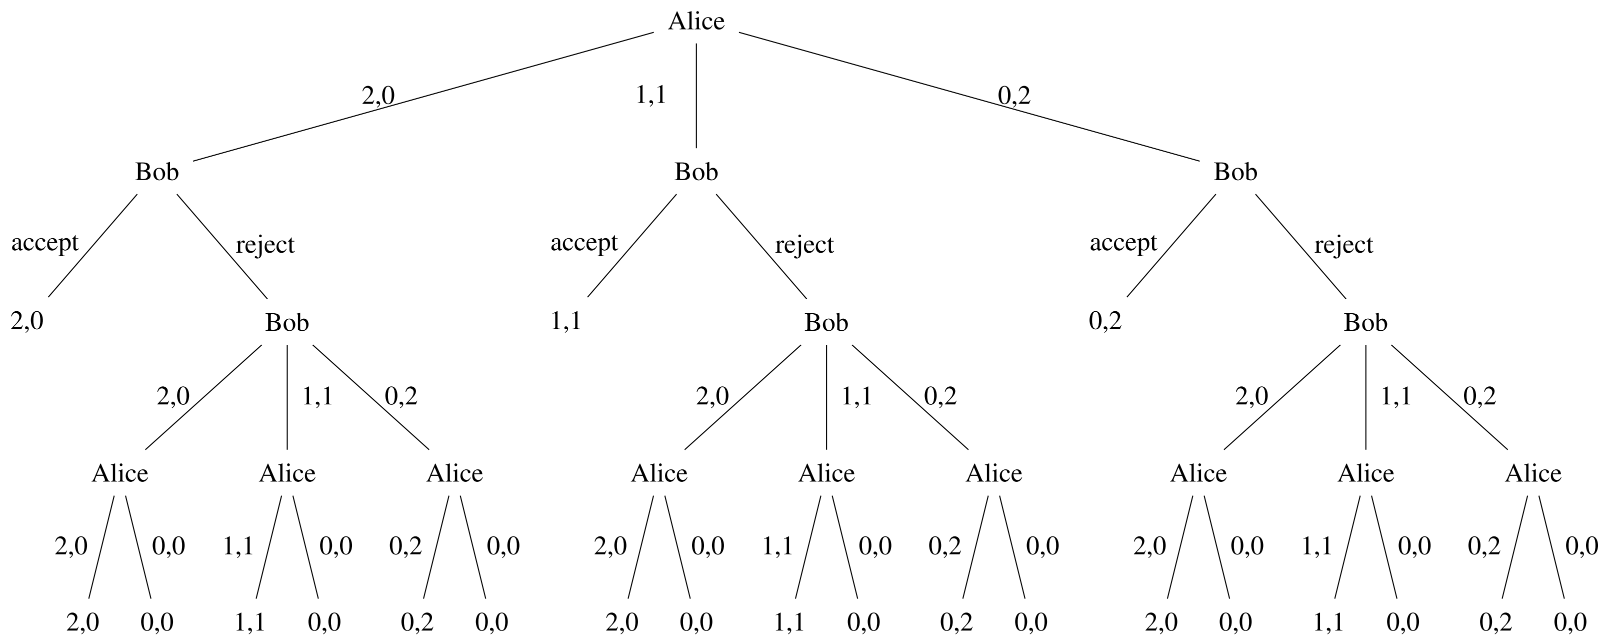

Figure 6: The extensive form of the game described in Exercise 3

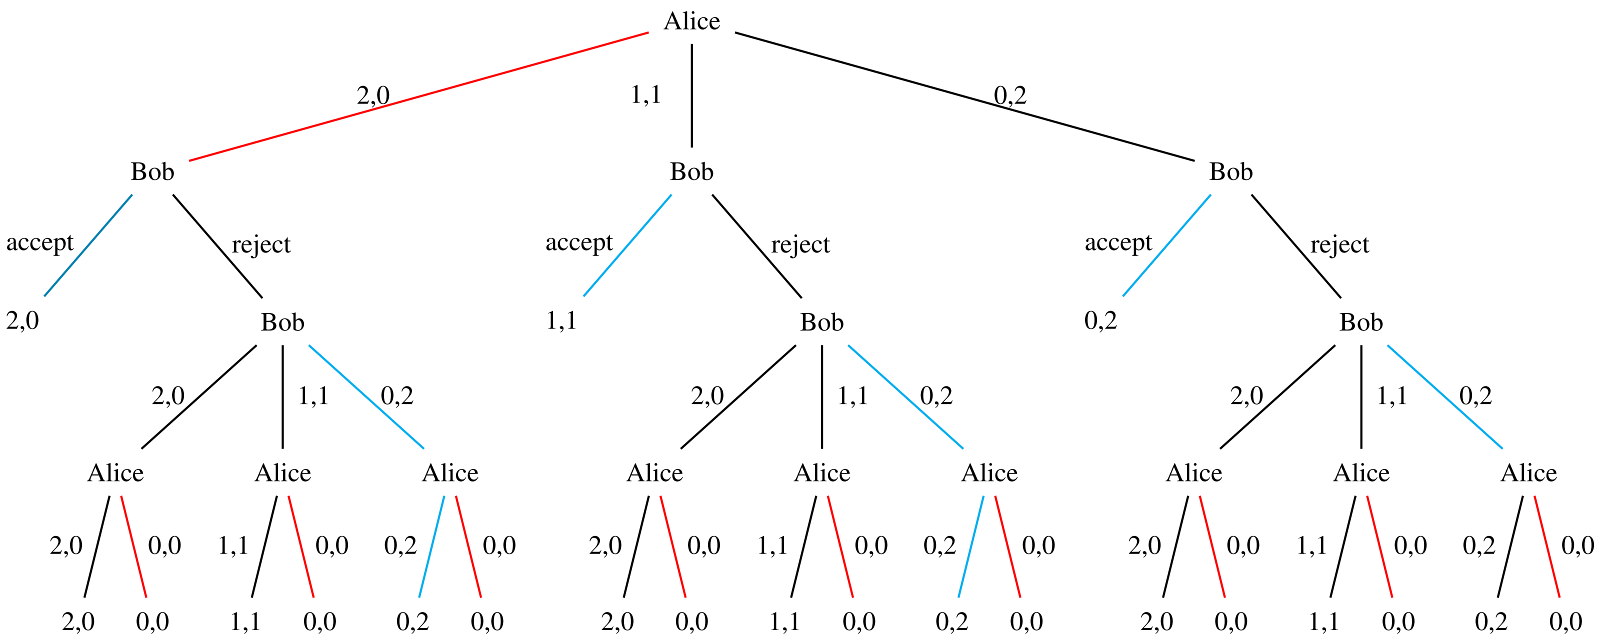

Figure 7: A Nash equilibrium maximising Alice's payoff as in Exercise 3

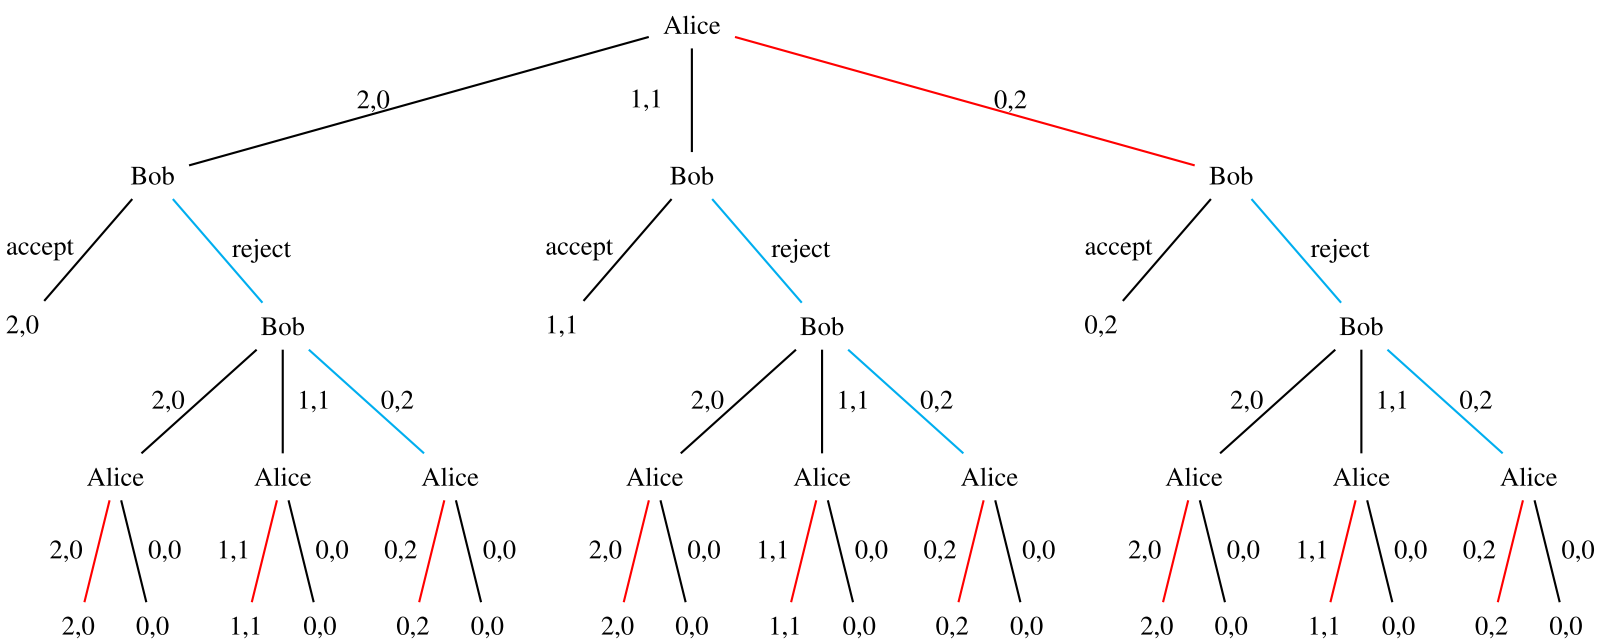

Figure 8: A Nash equilibrium maximising Bob's payoff as in Exercise 3

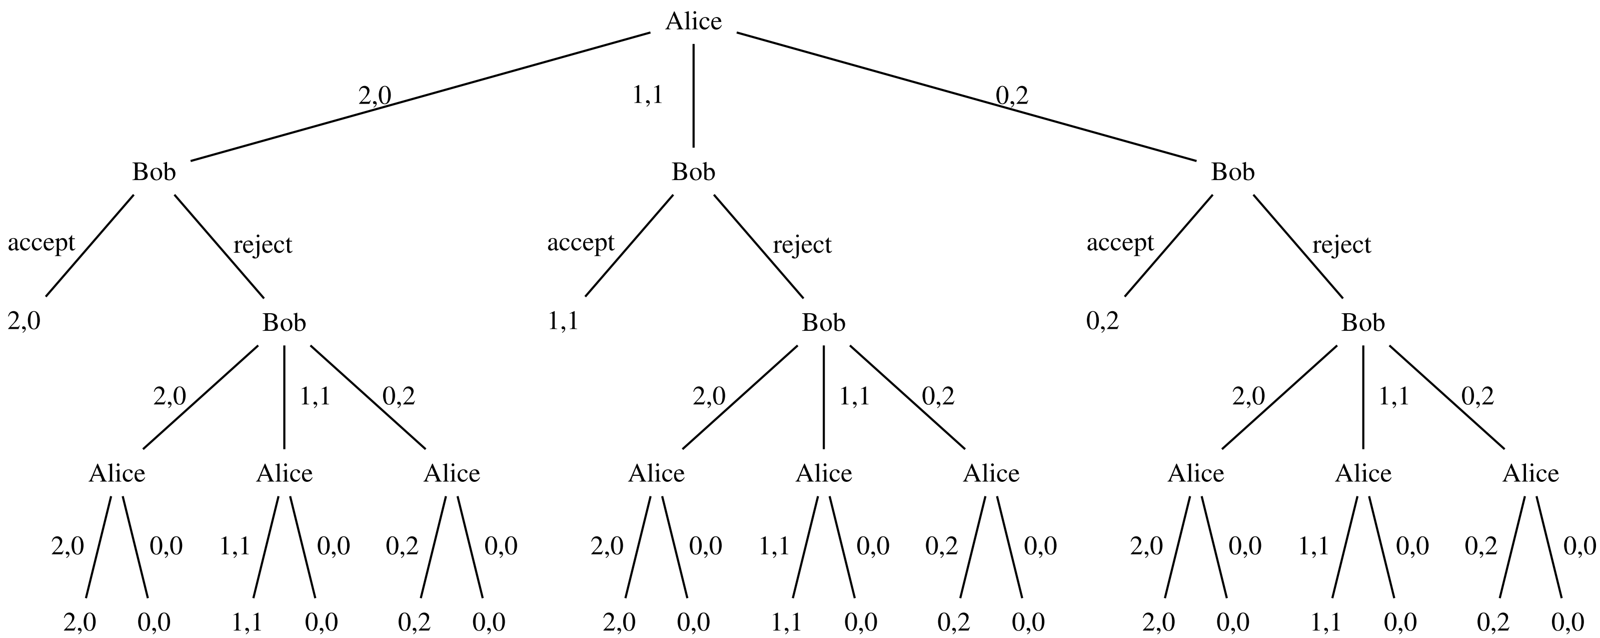

Figure 9: Subgame perfect Nash equilibrium for the game in Exercise 3

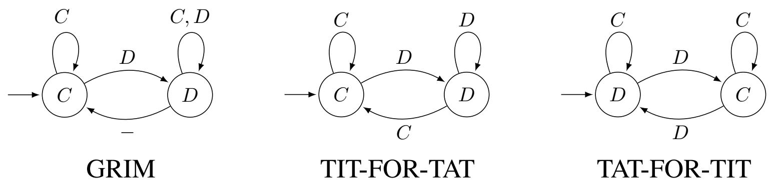

Figure 10: Three finite state machines that play the iterated prisoner's dilemma.

Exercise 4: Imperfect information games Consider an extensive form imperfect information game in which each player i has k information sets, that is, $|\mathcal{I}| = k$ for all $1 \leq i \leq n$.

(a) If a player has an identical number m of possible actions in each information set, how many pure strategies does he have?

(b) If a player has m_j actions in the j'th information set ($1 \leq j \leq k$) how many pure strategies does he have?

Solution

(a) m^k pure strategies.

(b) $\prod_{j=1}^{k} m_j = m_1 * \cdots * m_k$ pure strategies.

Exercise 5: Iterated games Let's consider playing the infinitely repeated prisoner's dilemma, using finite state automata strategies, and measuring utility over infinite runs, as discussed in the lecture. Figure 10 shows three two-state strategies for playing the iterated prisoner's dilemma.

	C	D
C	$-1, -1$	$-3, 0$
D	$0, -3$	$-2, -2$

(a) Informally explain what TAT-FOR-TIT does.

(b) Consider each strategy playing against each other strategy (including itself). Compute the runs that would be generated, and identify the finite but in- finitely repeating sequence of outcomes. Use this repeating sequence to compute the utility obtained by each strategy in each pairing.

	GRIM	TIT-FOR-TAT	TAT-FOR-TIT
GRIM	$\underline{C}\cdots$ -1 $\underline{C}\cdots$ -1 *	$\underline{C}\cdots$ -1 $\underline{C}\cdots$ -1 *	$C\underline{DD}\cdots$ -1 $D\underline{DC}\cdots$ $-2\frac{1}{2}$
TIT-FOR-TAT	$\underline{C}\cdots$ -1 $\underline{C}\cdots$ -1 *	$\underline{C}\cdots$ -1 $\underline{C}\cdots$ -1 *	$\underline{CDD}\cdots$ $-1\frac{2}{3}$ $\underline{DDC}\cdots$ $-1\frac{2}{3}$
TAT-FOR-TIT	$D\underline{DC}\cdots$ $-2\frac{1}{2}$ $C\underline{DD}\cdots$ -1	$\underline{DDC}\cdots$ $-1\frac{2}{3}$ $\underline{CDD}\cdots$ $-1\frac{2}{3}$	$D\underline{C}\cdots$ -1 $D\underline{C}\cdots$ -1 *

Figure 11: Table with solutions for Exercise 5

(c) Which of these pairs of strategies do you think forms a Nash equilibrium? (An informal argument will suffice.)

Solution:

(a) TAT-FOR-TIT starts with defecting and then repeats its action as long as his opponent cooperates, if his opponent defects, he changes his action. Repeat action if opponent cooperate, toggle if opponent defects.

(b) See Figure 11. Periods are underlined. A plus indicates that the respective strategy profile is a Nash equilibrium, a minus that it is not.

(c) *Remark:* Playing two machines strategies at the same time defines a concurrent system. Observe that, if both machines involved have only two states, the "concurrent system" has only four states in all. If moreover at both states the output of both machines is different, the states of the concurrent system (profile) are fully determined by the two outputs.

Remark: Observe that to establish that some machine strategy profile is *not* a Nash equilibrium, you need to formulate a better response for one of the players, and it may suffice to show this for one of the named machine strategies. However, to establish that some profile of machine strategy *is* a Nash equilibria, you have to quantify over all possible strategies and show that none of them is a better response!

(1) Against GRIM it is always a best response to adopt a strategy that, against GRIM, always cooperates, otherwise you will get a guaranteed a payoff between -2 and -3.

(2) Against TIT-FOR-TAT a strategy that guarantees cooperation throughout is a best response. To see this, observe that $(C, C)^{\omega}$ guarantees an average payoff of -1. By defecting one round against TIT-FOR-TAT, one may perhaps obtain a maximal payoff of 1 in one round only to loose this advantage in the next round.

(3) By playing TAT-FOR-TIT against either GRIM or TIT-FOR-TAT, you had better deviate to playing either GRIM or TIT-FOR-TAT.

(4) Playing against TAT-FOR-TIT any strategy that guarantees a period $(C, C)^{\omega}$ is a best response. To see this, first observe that by adopting a strategy that always chooses C, results in the least attractive possible payoff of -3. So you have to

adopt a strategy that defects at least once against TAT-FOR-TIT. After that you will have to defect at some point so as to gain more than a payoff of -1. This yields a return of 0 in the next round. The round after that however this gain will be nullified as TAT-FOR-TIT will then defect as well.

Thus, (1) and (2) guarantee that (GRIM,GRIM), (GRIM,TIT-FOR-TAT), (TIT-FOR-TAT,GRIM), and (TIT-FOR-TAT,TIT-FOR-TAT) are Nash equilibria. Furthermore, (3) explains why (TAT-FOR-TIT,GRIM), (TAT-FOR-TIT,TIT-FOR-TAT), (GRIM,TAT-FOR-TIT), and (TIT-FOR-TAT,TAT-FOR-TIT) are not Nash equilibria. Finally, (4) explains why (TAT-FOR-TIT,TAT-FOR-TIT) is another Nash equilibrium.

Exercise 6: Iterated games *This question is for personal development only, and is not strictly speaking part of the assessment!*

Consider all possible two-state finite machines for playing the iterated prisoner's dilemma.

(a) Suppose one of these machines has both states labelled with the same action (we have either C in both ovals, or D in both ovals). What behaviour does it generate? Can you simplify the automaton at all?

(b) Let us say two automata strategies σ_1 and σ_2 are distinct if there is an automaton σ^* such that the sequence of outcomes generated by playing σ^* against σ_1 is different to the sequence of states generated by playing σ^* against σ_2. Claim: there are precisely 26 distinct one and two state automata. (The one-state automata are ALLD and ALLC.) Of these 26 automata strategies, 13 will start by playing C, and the other 13 start by playing D. Draw the 13 automata that start by playing C. Given these 13 automata, you can very easily obtain the 13 automata that start by playing D. So do it.

Solution:

(a) The automata with both states labelled with C will always cooperate no matter what action the other player chooses. The opponent therefore will also always read C. The trivial one-state automaton of the with only state C will therefore induce exactly the same behaviour both by the player as well as by her opponent. A similar argument of course pertains to the two-state automata with both states labelled with D.

(b) See Figures 12 and 13.

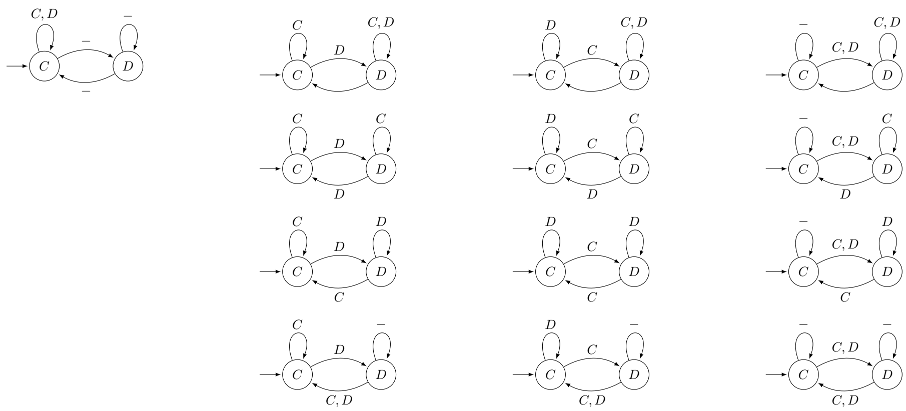

Figure 12: The 13 automata that start cooperating.

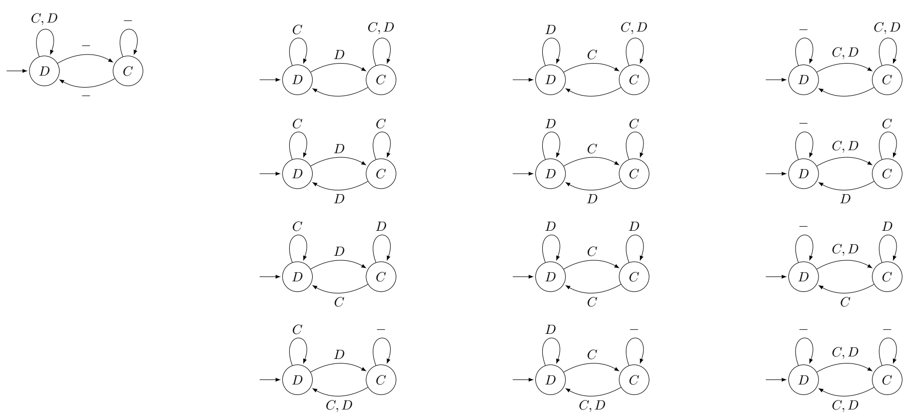

Figure 13: The 13 automata that start defecting.