Game Theory for Information Sciences

Exercises

Mark Kroon
mkroon@science.uva.nl
0306185

Nicholas Piël
nicholas@nichol.as
0104612

Stephan Schroevers
sschroev@science.uva.nl
0306053

October 17, 2008

Computer Science Excercises

Exercise A1: The Game of Matches

- Game:

> There is a pile of $n \in \mathbb{Z}^+$ matches on the table. Players in turn pick 1, 2 or 3 matches until no matches remain. The player who at this time has an odd number of matches wins the game.

(i) Who wins the game for $n = 15$?

When the game starts, both players possess an even number of matches, namely 0. Important to note is that each successive turn in the game depends on wheter the active player has an odd or even number of matches. In fact, it turns out that this property, together with the number of remaining matches, completely determines the next step by the active player (that is, how many matches to pick); if a winning strategy for this player exists at all.

This allows us to create a table that correlates these two game state properties with the number of matches that should be picked next. See Table 1. This table was created assuming that the total number of matches in the game is odd. The value for l indicates the number of matches left on the table. "Odd" and "Even" indicate what the active player should do if he possesses an odd or even number of matches respectively. For $l \in \{1, 2, 3\}$ we reasoned about the sensible number of matches to pick; dots indicate that we verified that no winning strategy exists in such state. For $l \geq 4$, the number of matches to pick was determined by referencing the table columns for $l - 1$, $l - 2$ and $l - 3$.

Arrows A, B and C indicate how this is done. For example, arrow A assumes the active player possesses an even number of matches, while $l = 7$ matches

Table 1: Correlation between game state and next step according to the winning strategy for games with an odd number of matches.

l	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
Even	1	1	3	3	•	2	2	•	1	1	3	3	•	2	2
Odd	•	2	2	•	1	1	3	3	•	2	2	•	1	1	3

remain. Note that this means that the opponent must also possess an even number of matches. Table 1 indicates that the active player should pick two matches. To see why this is, first consider the situations in which the the active player were to pick either one or three matches. He would then end up with an odd number of matches while six or four matches remain. The opponent would then pick two or three matches respectively (as his winning strategy prescribes, since he must currently also possess an even number of matches), forcing the currently active player in a position for which no winning strategy exists.

On the other hand, if the active player picks two matches, he keeps an even number of matches and $l' = 5$ matches remain on the table. In this state it does not matter how many matches the opponent picks: the currently active player has a winning strategy in all cases (that is, as a player with an even number of matches for $l'' \in \{l'-1, l'-2, l'-3\} = \{4, 3, 2\}$). This is also reflected by the fact that Table 1 indicates that no winning strategy exists for a player with an even number of matches while five matches remain—which is exactly the position in which the opponent is forced.

Now let us return to the initial question, *Who wins the game for $n = 15$?* Again note that at the start of the game both players possess an even number of matches. For $n = 15$, Table 1 indicates that that a player with an even number of matches has a winning strategy. Hence, the first person to play wins.

(ii) Who wins the game for general $n \in \mathbb{Z}^+$ when the winning condition for even n is adapted in a reasonable way?

The following winning condition turns out to be very satisfactory:

> *For even $n \in \mathbb{Z}^+$, the winner is the player that either takes away the last match and ends up with an even match count, or does not take away the last match but ends with and odd number of matches.*

To see why this condition is natural, consider Table 2, which shows the winning strategy for a game with 14 matches. Compare this table with Table 1, and notice that all values are shifted one place to the right.

Now, the winner for general $n \in \mathbb{Z}^+$ is the first player if:

- n is odd and for $l = n$ there is a suggested move in the *Even* row in Table 1.
- n is even and for $l = n$ there is a suggested move in the *Even* row in Table 2.

Table 2: Correlation between game state and next step according to the winning strategy for games with an even number of matches.

l	1	2	3	4	5	6	7	8	9	10	11	12	13	14
Even	•	1, 2	1	3	3	•	2	2	•	1	1	3	3	•
Odd	1	•	2, 3	2	•	1	1	3	3	•	2	2	•	1

Otherwise the second player wins. It is easy to see that the first player wins for $n \in \{1, 2, 3, 4, 7, 8, 9, 10, 11, 12, 15, 16, 17, 18, 19, 20, 23, \dots\}$. Let $W\colon \mathbb{Z}^+ \to \{1, 2\}$ be the function that determines which player wins a game with an initial number of n matches. Then $W(n)$ is defined as:

$$W(n) := \begin{cases} 1 & \text{if } (n+1) < 6 \pmod 8, \\ 2 & \text{otherwise.} \end{cases}$$

(iii) Show that the behavior of the game is periodic.

The concise definition of the function $W(n)$ using the modulo operation, as introduced in the previous answer immediately shows that the game is periodic.

(iv) What happens if the rules change, requiring that the number of matches removed by a player is a positive square number?

With these new rules the number of winning strategies for the first player increases. Figure 1 shows the probability of the first player winning a game with a random number of initial matches $i \leq n$, for $1 \leq n \leq 17000$.

For the original game (in which 1, 2 or 3 matches could be picked each turn) the probability of winning quickly stabalizes at 75%, in agreement with the definition of $W(n)$. When the new rules are applied, a different trend emerges: for larger n the probability that the first player has a winning strategy increases.

(v) Bonus: Consider modifying the game by changing either the set of numbers of matches both players can pick up on their turn, or the condition selecting the winner when the game terminates. Show that as long as the set of numbers of matches a player can pick-up is fixed and finite, and the winning condition is expressed by simple congruences, the behavior of the game is periodic in the number of matches in the initial position of the game.

Unfortunatly, we did not answer this question.

Exercise A2: A Game of Differences

- Game:

Consider the set $\{1, 2, \dots, n\}$, for $n \in \mathbb{Z}^+$. Players in turn select a number a_i in this set, for $i = 1, 2, \dots$. The first player to pick a number a_i such that some difference $a_i - a_j$ for $j < i$ already occurs in the set of differences $\{a_k - a_l \mid 1 \leq k, l < i\}$ loses the game.

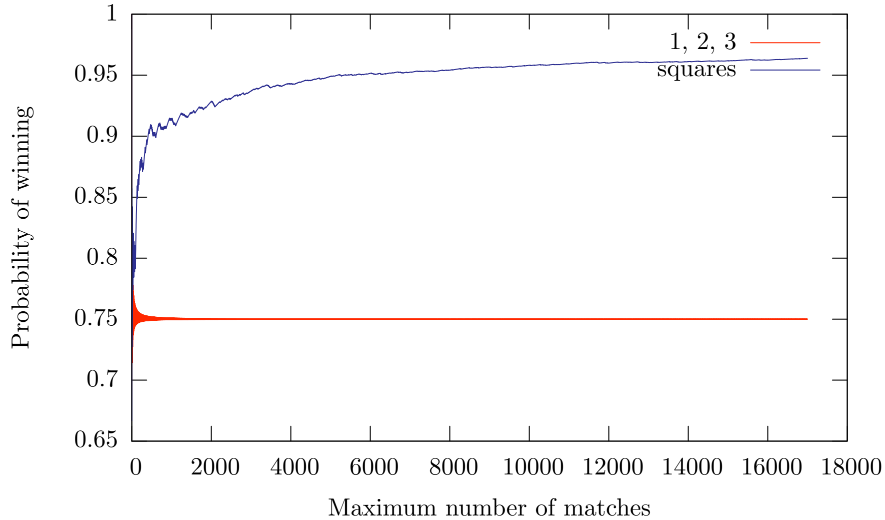

Figure 1: Probability of the first player winning a game with a random number of initial matches $i \leq n$, for $1 \leq n \leq 17000$.

Table 3: Winning player (assuming knowledge of a winning strategy) and maximum game duration for the game of exercise A2 using various initial set sizes.

Set size	1	2	3	4	5	6	7	8	9	10	11	12	13
Winner	1	2	2	1	1	1	1	1	1	1	2	2	1
Max. turns	1	2	3	3	4	4	4	4	5	5	5	5	6

Set size	14	15	16	17	18	19	20	21	22	23	24	25	26
Winner	1	1	1	1	1	1	1	1	1	1	1	1	2
Max. turns	6	6	6	6	6	7	7	7	7	7	7	7	7

(i) Analyse the game for $n = 25$.

There are 25! ways in which the elements from $L = \{1, 2, \ldots, 25\}$ may be selected (without regarding the restrictions imposed by the game rules), so in order to analyze this game, less naive methods must be devised to construct the set of valid game sequences.

For very small values of n the creation of a game tree can be done by hand (see e.g. Figure 2). For larger values this is obviously undoable. For example, for $n = 25$ there are over 6.8 million valid ways to play the game.

Table 3 shows the result of our analysis of the game for $1 \leq n \leq 26$. The table was generated using a computer program. It shows that player I has a winning strategy for $n = 25$.

The upperbound on the duration of this game for $n = 25$ is 7 moves. An example of such a game is the game that consists of the sequence of moves $\langle 1, 2, 5, 17, 23, 25, 12 \rangle$. The lower bound is 5 moves, of which the sequence $\langle 1, 2, 4, 8, 21 \rangle$ is an example.

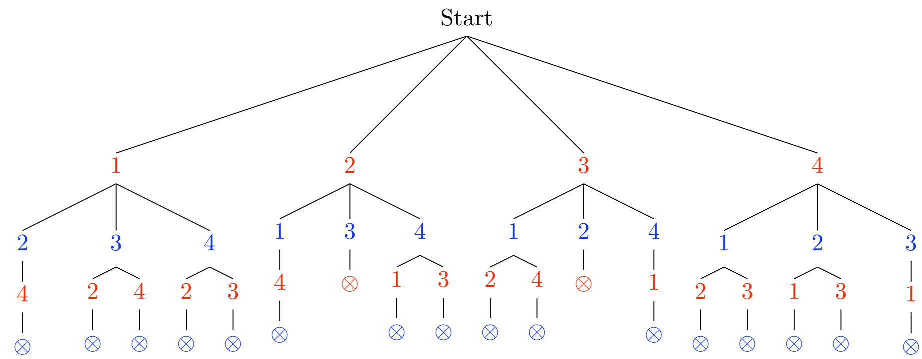

Figure 2: Game tree for $n = 4$ for the game of differences of exercise A2. Numbers picked by the first player are red, those selected by the second player are blue.

(ii) What is the best upper bound on the duration of this game you can prove for general n?

For a set $\{1, 2, \ldots, n\}$ the set of possible differences is $\{x - y \mid 1 \leq x, y \leq n\} = \{1-n, \ldots, n-1\}$. Note that since for every pair of selected numbers a_i and a_j both $a_i - a_j$ and $a_j - a_i$ will be added to the set of differences, we may restrict our focus to the non-negative part of the set of differences.

Hence the second half of the game description can be rephrased as follows: the first player to pick a number a_i such that some difference $|a_i - a_j|$ for $j < i$ already occurs in the set of differences $\{|a_k - a_l| \mid 1 \leq k, l < i\}$ loses the game[1]. Also note that $|\{|x-y| \mid 1 \leq x, y \leq n\}| = |\{0, 1, \ldots, n-1\}| = n$.

Next, in order to find the best upper bound on the duration of te game, we investigate how the number of differences added to each move can be minimized. When a number a_i is selected, it can have a distance x to *at most* two numbers a_j, a_k that were selected earlier. a_1 is the first number to be selected and yields only the difference $a_1 - a_1 = 0$. Assuming an optimal placement, the minimal number of differences added by selecting number a_i for $i > 1$ is $\lceil \frac{i-1}{2} \rceil$.

A sequence of optimal placements (which is not possible to achieve in practice), then allows the size of the difference set to grow as,

$$1, 2, 3, 5, 7, 10, 13, 17, 21, 26, 31, \ldots$$

Instead of finding a formula for this sequence ourselves, we find one in the online encyclopedia of integer sequences[2]. Hence, after drawing number a_i from the set, the size $s(i)$ of the set of differences can be calculated using

$$s(n) = \lfloor \frac{n^2}{4} \rfloor + 1.$$

Now, the best upper bound on the game duration that we can prove, for a game with an initial set size of $n > 0$, is $max_{i \in \mathbb{Z}^+}(s(i) \leq n)$.

[1] Here $|x|$ denotes the absolute value of x, instead of the set size of x.
[2] http://www.research.att.com/ njas/sequences/A033638.

Table 4: Napoleon Patience game position which can be won, but which can also become lost by one single valid move.

	A	B	C	D	E	F	G	H	I	J	K	L	M	N
1	A♣	2♣	3♣	4♣	5♣	6♣	7♣	8♣	9♣	10♣	J♣	Q♣	K♣	
2	A♢	2♢	3♢	4♢	5♢	6♢	7♢	8♢	9♢	10♢	J♢	Q♢	K♢	
3	2♠	A♠	3♠	4♠	5♠	6♠	7♠	8♠	9♠	10♠	J♠	Q♠	K♠	A♡
4			2♡	3♡	4♡	5♡	6♡	7♡	8♡	9♡	10♡	J♡	Q♡	K♡

Exercise A3: Napoleon Patience

- Game:

  Take a deck of cards and randomly place all 52 of them open on a table in four rows of 13 cards. The actual length of the rows however is 14, since there is an empty place at the end of every row.
  Legal moves are to place in an empty position behind a card its successor of the same suit. An empty place in the first column may be filled by an ace of any suit, but this ace should not be placed in the first column already.
  Purpose of the game is to rearrange the initial arrangement into a completely well ordered deck: each row contains an entire suit in its natural order.

(i) Give a game position which you can win, but which becomes lost by performing the wrong (legal) move.

Figure 4 shows the requested game position. The four rows are marked by a number, while the 14 columns are denoted by the letters A through N. The game can be won by first performing the moves (3B → 4A), (3A → 4B), (4A → 3A), (4B → 3B) and (3N → 4A), after which all other cards in the bottom row can be shifted to the left by a single column.

On the other hard, if the move (3N → 4A) were to be performed while in the game position as depicted by Figure 4, then the game cannot be won anymore. This is because the resulting game position leaves no possibility to remove a card from first column, hence removing all hopes of swapping 2♠ and A♠.

(ii) How can you generalize the game so that it can be played with a pair of decks?

The game can be generalized in several ways. One thing that cannot be avoided, is that in the new situation each empty place can be filled with two instead of one card. An interesting exception are the kings. In the original game no card can be placed after a king. The introduction of an additional deck challenges this rule. Two general possibilities come to mind when playing with a pair of decks:

- The cards are positioned in *eight* rows of 14 columns.
- The cards are positioned in four rows of *27 or 28* columns (depending on whether the number of empty places is doubled or not.

In the latter case the question arises how two complete suits are positioned next to each other. Need both be of the same suit? Should the ace of the second suit immediately follow the king of the first suit? Or should there be an empty place in between? In the latter case, column 15 should receive the same status as the first column, in that aces may be placed there unless such ace was already placed in column 1 or 15.

The first option (extending the game to eight rows) seems the most general solution, since no other addional rules and restrictions need to be imposed.

(iii) How does this generalization affect the analysis of the game?

In the new game, still no card may be placed in an empty place behind a king, but for the remaining empty places (that are not in the first column), there are now *two* cards that may fill the gap. Additionally, the probability that all empty places are preceded by a king is—at least in the initial game position—much smaller than in the original single-deck game.

Still, the game can be won as well as lost.

(iv) BONUS: Prove that the game graph is acyclic and that the game thus terminates.

Note that the following proof regards the original game with a single deck of cards. It performs induction on the rank of the card that is moved during a turn.

The cards on the table are ordered in four rows of 14 columns. Let (m, n) denote the card in the nth column of the mth row, such that $(1, 1)$ is the card in the first column of the first row.

Assume an arbitrary game position P after t moves in which a card X at place (x, y) can be moved to (x', y'). That is, (x', y') is an empty place, and more importantly, $(x, y) \neq (x', y')$. Assume the player performs the move just described, resulting in the game position P' after $t + 1$ moves. Then, one of two things is the case:

- X has the lowerest rank (rank 1). Then X is an ace. The rules of the game prohibit X from being moved again, and thus there is no way that game position P can reoccur later in the game. More relevant for the remainder of this proof is the following: X will not return to its original position!
- X has rank $n + 1$ with $n > 0$. Then X is a card with a rank greater than ace. Furthermore, the predecessor W of X is at place $(x', y' - 1)$. Now assume there is a game position P'' after $t + k$ moves, for $k > 1$, such that $P = P''$. Then X must be at place (x, y) in P''. For that to happen, there must be some state Q at time $t + i$, with $1 < i < k$, where W is at position $(x, y - 1)$, *and* a state Q' at time $t + j$, with $i < j \leq k$, where W is back at position $(x', y' - 1)$. But W has rank n, and by induction it follows that W cannot return to a previously held position once moved. So after time t, state P will not reoccur.

This shows that for an arbitrary game position P, there is no sequence of moves of length greater than zero that results in a game position equal to P. Hence there is *no* game position that occurs twice in any game, leading to the conclusion that the game graph is acyclic.

By the fact that there are only finitely many game positions, the game must terminate.

Exercise A4: Finite Sequences

- Game:

A position at turn t is given by a finite sequence $S_t = \langle a_1, a_2, \ldots, a_n \rangle$ of nonnegative integers. A move consists of selecting a number a_i from the list and replacing this number by a (possibly empty) finite collection of smaller nonnegative integers. In particular if 0 is selected the move amounts to removing this occurrence of 0. The player faced with the empty list loses the game.

(i) Prove by induction that this game terminates.

Assume a game state at turn t in which the players are left with a sequence $S_t = \langle a_1, a_2, \ldots, a_n \rangle$ of nonnegative integers. One of the players selects an integer a_i from S_t. Then one of the following two things is the case:

- $a_i = 0$. Then no elements can be added to the list, so a_i is simply removed. Hence $|S_{t+1}| = |S_t| - 1$.
- $a_i = x + 1$, $x \geq 0$. Without loss of generality, assume that a_i is replaced by m integers $\langle b_1, b_2, \ldots, b_m \rangle$. Naturally, $b_j \leq x$ for all $1 \leq j \leq m$. Hence $|S_{t+1}| = |S_t| + m - 1$.

It follows that any sequence S can be extended only a finite number of times, each time by a finite amount. If $|S_0| = n$ and the size of each replacement sequence is bounded by m, then the number of turns in the game is bounded by $n + n \cdot \sum_{i=1}^{k} m^i$, where k is the largest integer in S_0. This bound is only reached if S_0 consists of n occurrences of k, and if each integer a_i removed from S_t is replaced by exactly m integers $a_i - 1$ in S_{t+1}.

(ii) Why is it impossible to perform traditional backward induction on a given starting position in order to determine which player wins the game?

Though the game always terminates, there is an infinite number of ways in which to play the game. This is because there are infinitely many collections of finite size that can be used to replace a number at each step. (Because there are infinitely many sets of finite size.)

Hence the game tree has an infinite number of branches, and therefore an infinite number of leave nodes. Traditional backwards induction starts at the leave nodes and works its way back to the root of the game tree. In this situation, such approach is not possible.

(iii) Prove that if some starting position is won (lost) for the first player, adding two copies of the same number to the list will yield another position won (lost) for the first player. Is the same true for positions won by the first player?

Unfortunatly, we did not answer this question.

(iv) Give a complete analysis of this game.

Unfortunatly, we did not answer this question.

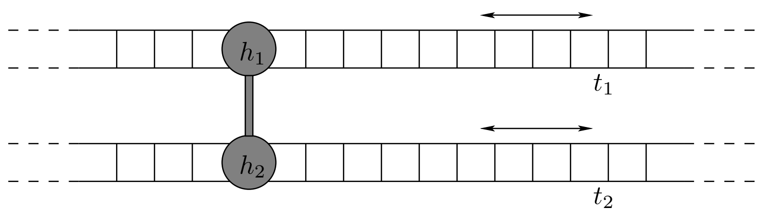

Figure 3: Two-tape Turing Machine with fixed heads and moving tapes.

Exercise B1: The Tiling Game and its use in reduction theory

- Game:

> The Tiling Game requires an infinite supply of colored tiles to be aligned in an (in)finite grid such that adjacent tiles have matching colors.
> Additionally, colors given along the edge of the grid or some initially provided tile must also be adhered.
> Turing Machine computations can be reduced to specific instances of the Tiling Game, as shown in class.

(i) Design a variant of the reduction of Turing Machine computations to the Tiling Game where a legal tiling simulates the computation of a two-tape Turing Machine.

In the original reduction of Turing Machine computations to the Tiling Game each horizontal tile boundary represents the state of the tape at some point in time, while time is running along the vertical axis. The tape does not move, but the head does (by way of the current state, which is transferred from one tile to another).

Following the notation of [vEB96], a two-tape Turing Machine is a tuple $M = \langle K, \Sigma, q_0, \Box, P\rangle$. Here K is the set of internal states, Σ is a finite set of tape symbols, $q_0 \in K$ is the start state, $\Box \in \Sigma$ is the special blank symbol and P is the program. Since a two-tape Turing Machine has two read/write heads, we have that $P \subseteq ((K \times \Sigma^2) \times ((K \cup \{\bot\}) \times (\Sigma \times \{L, 0, R\})^2))$. $\bot$ represents a termination instruction. Note that each step of the computation both heads may write on the tape.

Suppose that the tapes are fixed, positioned in parallel. Then the heads must move independently, causing the distance between them to become arbitrarily large during a computation. Still, each step in the computation the heads must act according to one and the same instruction. It is rather hard (impossible?) to enforce such restriction using tiles.

As an alternative, consider the situation in which the tapes are positioned in parallel, but are not fixed: they can move to the left and to the right. Now, the heads may be fixed, right next to eachother. See Figure 3.

This situation can be modelled using tiles with relative ease. As before, time runs along the vertical axis. In this model however, *two* tapes run along the horizonal axis, encoded using a single color. Also, there is only one column that models the heads and the state transitions. Figure 4

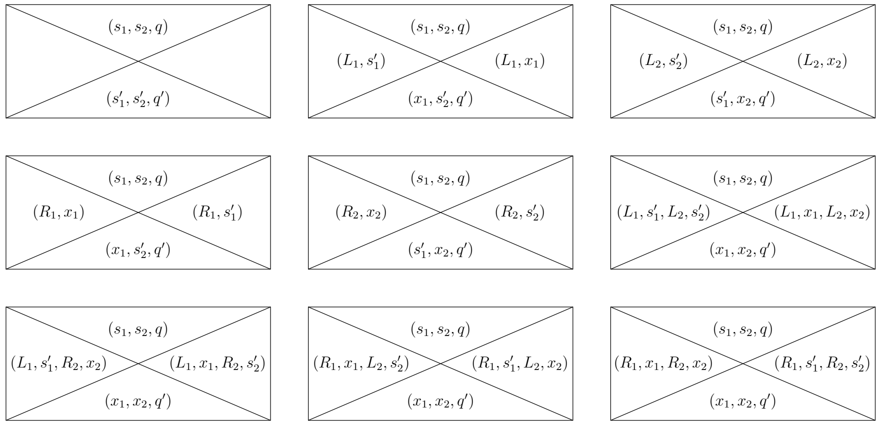

Figure 4: Tiles for a variant of the Tiling Game for which valid tilings simulate the computation of a two-tape Turing Machine.

shows nine types of tiles (which are rectangular instead of square only for formatting purposes) that can be placed in this specific column (there is one and only one such column, obviously). These tiles model the steps that a two-tape TM can take if the first head reads symbol s_1 and the second head reads symbol s_2, while in state q. Symbols s'_1 and s'_2 are written to the first and second tape, respectively, after which the TM transitions to state q'.

Either or both of the heads may be moved in either direction after writing a symbol. Since the heads are fixed, such action is modelled by moving the tapes in the opposite direction. That is, if the head over the first tape should be moved to the left, then in fact the first tape will shift one place to the right. The result is the same, in the sense that the head is then over the intended cell. Moving the head over the first (second) tape to the right is encoded by using the color modifier L_1 (L_2) for the left and right rectangle of the tile. Likewise R_1 and R_2 signify a move of one of the heads to the left. Note that for each instruction $(q, s_1, s_2, q', s'_1, \_s'_2, \_)$ the color-modifiers x_1 and x_2 range over Σ: there is a tile for each combination.

When one removes the references to states q and q' in Figure 4, then what remains are the tiles that are placed in the non-state columns (all but the one special state-column, that is).

(ii) Can you generalize this construction for the k-tape model?

Yes, one can accomodate for symbols $s_3, s_4, \ldots, s_k$ on tapes $3, 4, \ldots, k$ by extending the tiles in Figure 4: all this requires is adding more permutations. The number of different tiles increases, and along with it the number of different colors that need to be encoded. Other than that the generalization to the k-tape model is trivial.

Exercise C1: Less Restricted ATM Reduction

- Description:

  > Chlebus' reduction of Alternating Turing Machine computations to the 2-person tiling game uses the restrictions on the ATM model involved that the states alternate between universal and existential and that each instruction moves the head of the ATM.

(i) Design a variant of this reduction which doesn't use these restrictions on the ATM.

Unfortunatly, we did not answer this question.

Exercise D1: The Pebble Game

- Game:

  > The goal of the Pebble Game is to pebble the output nodes of a directed graph using as few pebbles and moves as possible and by obeying the following rules:
  >
  > – You can always pebble an input.
  >
  > – You can always remove a pebble.
  >
  > – You may pebble a node provided its ancestors are pebbled.
  >
  > – You must pebble all outputs at least once.

(i) Prove that a complete binary tree of depth $k > 1$ with 2^k leafs requires $k + 2$ pebbles in the standard version of the pebble game (no shifts are allowed).

In the trivial case with a binary tree of depth $k = 1$, there are $k + 2 = 3$ pebbles needed to pebble both the inputs and subsequently the output node after which the game ends.

To pebble a binary tree of depth $k > 1$ requires the pebbling of the two subtrees of size $k-1$ simultaneously. One pebble is reserved for the output node of the first subtree. And so, to pebble a binary tree of depth k one needs the amount of pebbles needed for a binary tree of depth $k - 1$, plus one.

Given the fact that a binary tree of depth 1 requires 3 pebbles, it follows that a binary tree of depth k requires $k + 2$ pebbles.

(ii) Is there a good reason to prefer the Cook Pyramid over this binary tree as a component in lowerbound constructions in pebbling theory?

Gilbert, Langauer and Tarjan [GLT79] show us some properties of the Cook Pyramid which a binary tree lacks. One such property is that when one of the k inputs is pebbled, the output can be pebbled using $k - 1$ additional pebbles. A Cook Pyramid can also be used to fix a pebble on a node for a certain time interval by making this node the top (or apex) of a pyramid large enough to make repebbling of the node computationally infavourable.

These properties make the Cook Pyramid useful to let a pebbling problem represent the quantified boolean formula problem for which lowerbound solutions are known.

Game Theory Excercises

Exercises 1.10.{2,7,10}: The Domino Tiling Game

- Game:

  > Dominoes can be placed on an $m \times n$ board so as to cover two squares exactly. Two players alternate in doing this. The first to be unable to place a domino loses.

(i) Draw the game tree for the case $m = 2$ and $n = 3$.

See Figure 5. Note that some board positions are omitted, since they are symmetrical to some other board position that is shown. The $\mathcal{W}$ and $\mathcal{L}$ symbols at the terminal nodes signify whether the first player wins ($\mathcal{W}$) or loses ($\mathcal{L}$).

(ii) Apply Zermelo's algorithm to the 2×4 version of the game.

Figure 6 shows the game tree for the 2×4 version of the game, along with a $\mathcal{W}$ or $\mathcal{L}$ at every node, signifying which player has a winning strategy at the corresponding game position. Again, redundant symmetrical board positions have been omitted. As can be read from the graph, the second player has a winning strategy.

(iii) Determine a winning strategy for this game for one of the players.

If a single square becomes enclosed between the edge of the board and two dominoes, then at most three, not four dominoes can be placed on the board, thus resulting in a loss for the second player. The second player, who has a winning strategy, must thus place his first domino (the second domino on the board) in such a way that no single square is locked, and that the resulting free space does not allow for a square to be locked in by the first player, either. This may sound cryptic, but all that the second player need to ensure is that the resulting free spaces are of the form 1×2 or 2×2. See the last question regarding the Domino Tiling Game below for more on these specific shapes.

(iv) Find out which player has a winning strategy if m and n are even.

The maximum number of dominoes that can be placed on a board of size $m \times n$, with both m and n even, is $\frac{mn}{2}$, which is an even number. Hence a tiling that completely fills the board causes the second player to win. Player I will want to prevent this by ensuring that the tiling contains gaps. Player II will want to neutralize the effect of this behaviour by creating *additional* gaps, such that the eventual number of dominoes on the board remains even.

It turns out that player II has a winning strategy that is easy to comprehend (though hard to come up with…). Since both m and n are even, the board can be divided in four identical rectangles by splitting it in the middle along the horizontal and vertical axis, as the green lines in Figure 7(a) show.

This allows player II to place dominoes in positions that are mirrored with respect to the domino placement of player I. In other words, all dominoes

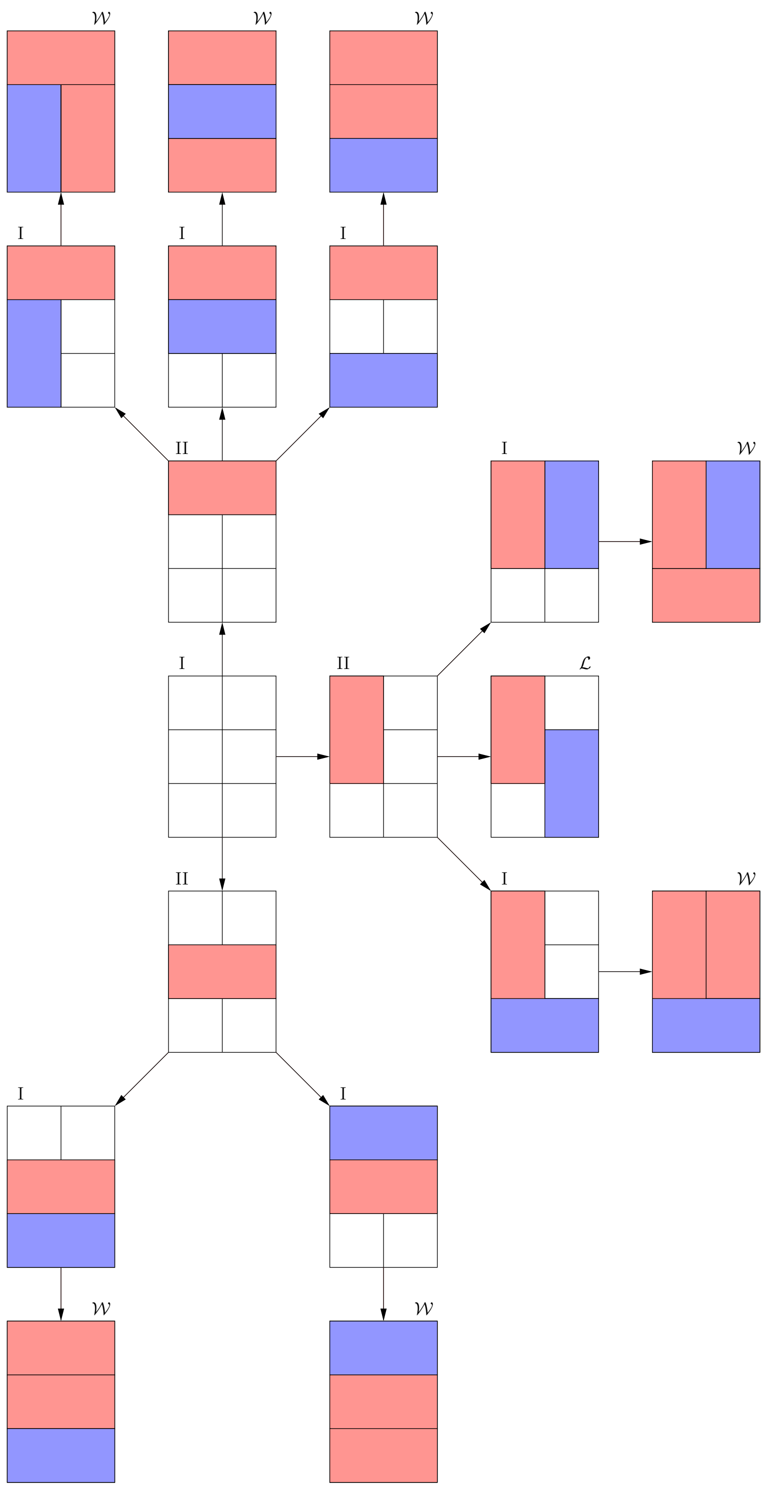

Figure 5: Game tree for the Domino Tiling Game of exercises 1.10.{2,7,10} assuming a 2×3 playing field.

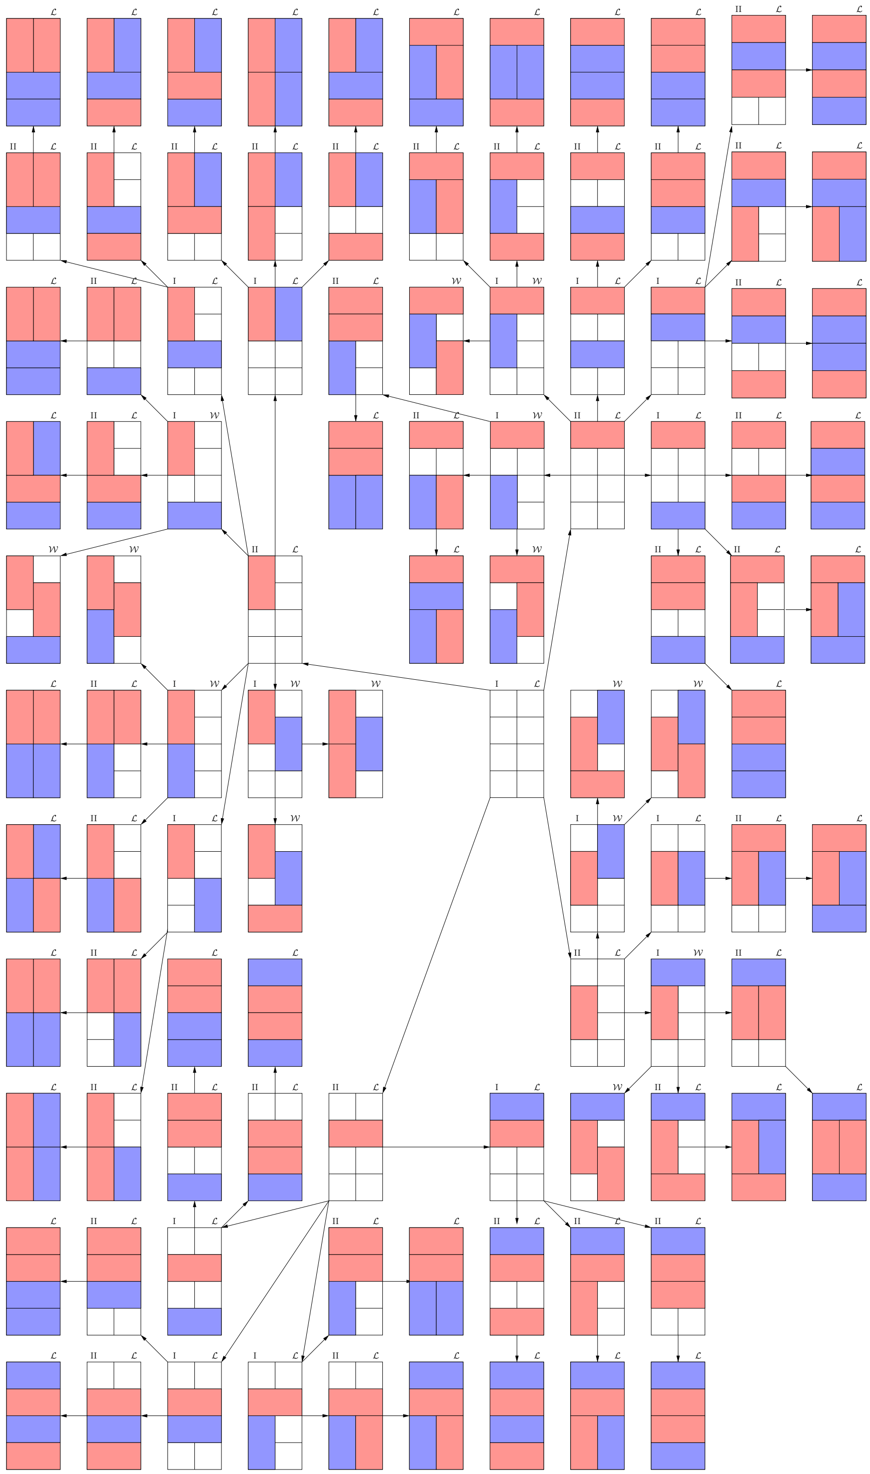

Figure 6: Game tree for the Domino Tiling Game of exercises 1.10.{2,7,10} assuming a 2×4 playing field.

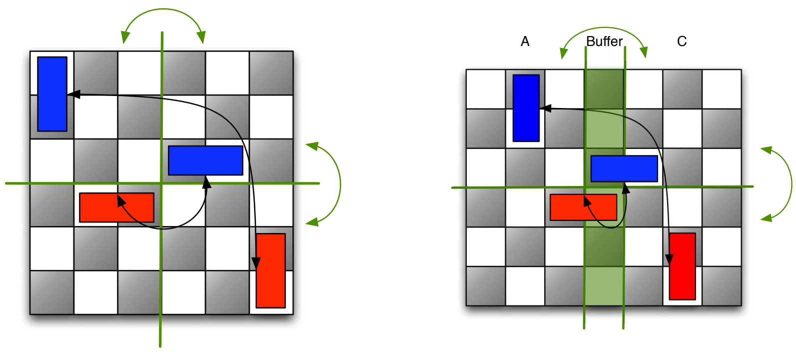

(a) 6×6 board.

(b) 6×7 board with buffer.

Figure 7: Example domino tiling games in which a symmetrical strategy is applied.

are mirrored along both axes, thus ensuring that the board has a two-fold rotational symmetry. This method ensures that whenever player I can place a domino in some quadrant, then so can player II in the opposite quadrant. Hence the second player wins. Again, see Figure 7(a) for an example.

(v) Who wins if m is even and n is odd?

When m is even and n is odd, the board can be split along the center ($\frac{n+1}{2}$th) column. The result consists of three smaller boards. For example, a board of size 6×7 can be split into boards A (size 6×3), B (size 6×1, in the middle, call this the *buffer*) and C (which is identical to A). See Figure 7(b).

Ignoring the buffer for a moment, any domino placement by player I on either board A or C can be copied by player II on the alternate board, as described above, as long as player II employs this strategy consistently.

However, player I has a trick up its sleeve. Contrary to a board where m and n are even, the boards currently under consideration *do* to allow a domino to be placed exactly at the center of the board, namely in the middle of the buffer (board B). By placing the first domino exactly at the center of the board, player I causes the roles to be reversed. Obviously player II cannot mirror this move, since there is only one board center. Hence player II is forced to place his domino somewhere else. From then on, player I can mirror the behaviour of player II, thus ensuring a win. Hence player I has a winning strategy.

(vi) For $m = n = 3$, who wins?

The second player will win. First, the first player places a domino anywhere on the board. Then, the second player places a domino right next to the first domino, thus creating a 2×2 square in one of the corners of the board. The squares along the borders opposite to this corner are then still free. Regardless of how the first player then places his second domino, the second player will be able to place one final domino, thus winning the game.

(vii) Give examples of boards such that the winner of the game does not depend at all on the strategies used by the players.

There are exactly five rectangular game boards that meet this criterium: 1×1, 1×2, 1×3, 1×5 and 2×2. It is easy to see that the number of dominoes that can be placed on these boards does not depend on the actual placement of the dominoes: the boards will allways accomodate 0, 1, 1, 2 and 2 dominoes respectively. The winner is thus determined from the start, regardless of the strategy employed.

Next, one can prove that there are in fact no other boards that meet the criterium. The following proof crafts specific board positions for all combinations of m and n not mentioned above. Without loss of generality, it is assumed that $m \leq n$:

- $m = 1$:
  - $n = 4$: If the first player places a domino in the middle of the board, the second player loses. On the other hand, if the first player aligns his domino to one side of the board, then the second player is left with enough room to place a second domino, thus winning th game. See Figure 8(a) for a graphical representation of the former situation.
  - $n \geq 6$: Figure 8(b) shows a possible placement of dominoes after two turns for boards of size $1 \times n$, $n \geq 6$. Obviously the left domino can be shifted a single position to the left or right to accomodate exactly one additional domino, thus yielding another winner.
- $m = 2$:
  - $n = 3$: Depending on the domino placement, a 2×3 field can either contain two or three dominoes. In the former case (see also Figure 8(c)) the second player wins, in the latter the first player wins.
  - $n \geq 4$: By strategically placing a domino on a board of size $2 \times n$, $n \geq 4$, the same situation as in the previous case can be created. Thus, again the first as well as the second player can win, depending on the strategy. See Figure 8(d).
- $m \geq 3$:
  - $n = 3$: Again, both players can win. Figure 8(e) shows a board position, that, if rearranged, could accomodate another domino, thus allowing the other player to win.
  - $n \geq 4$: Identical argument, see Figure 8(f).

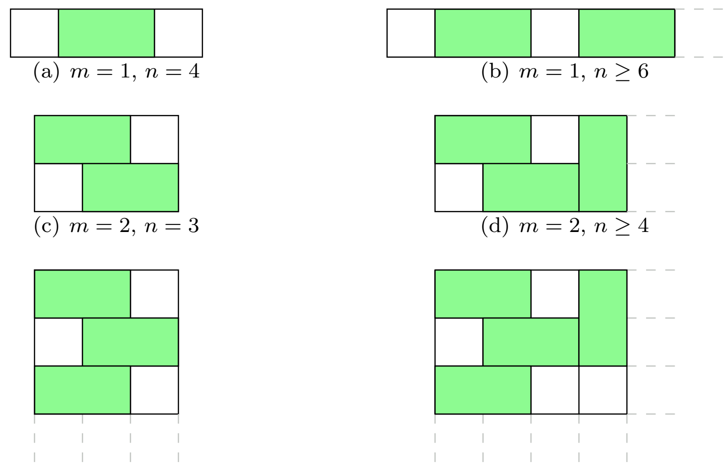

(a) $m = 1$, $n = 4$ (b) $m = 1$, $n \geq 6$

(c) $m = 2$, $n = 3$ (d) $m = 2$, $n \geq 4$

(e) $m \geq 3$, $n = 3$ (f) $m \geq 3$, $n \geq 4$

Figure 8: Specific non-optimal board fillings that prevent (at least) one additional domino from being placed.

Exercise 1.10.3: Blackball

- Situation:

> A committee of three club members (I, II and III) has to select one from a list of four candidates (A, B, C and D) as a new member of the club. Each committee member is allowed to blackball (veto) one candidate. This right is exercised in rotation, beginning with player I and ending with player III.

(i) Create the game tree for this game, labelling the edges with blackballed candidates, and labelling terminal nodes with the winning candidate.

See Figure 9.

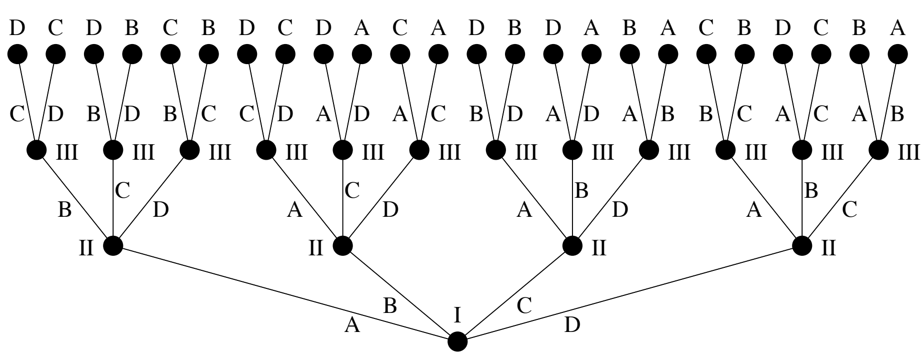

Figure 9: The game of Blackball of exercise 1.10.3.

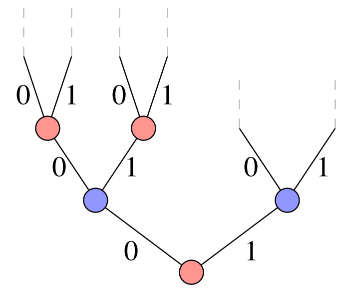

Figure 10: Partial game tree for the Real Number Game of exercise 1.10.5.

(ii) Assume that every player has a different favourite candidate and will be disappointed if that candidate is not selected. Show that in this case no player has a winning strategy.

This is immediate, from the fact that no player can prevent another player from blackballing a certain candidate. More elaborately put:

- Player I may blackball a candidate that is not of his liking, but he cannot prevent his favourite candidate from being blackballed by one the two other players that exercise their right to blackball next.
- Player II may find his favourite candidate blackballed by player I, which is enough to conclude that player II has no winning strategy.
- The same can be said about player III.

Exercise 1.10.5: The Real Number Game

- Game:

Two players alternately choose either 0 or 1 forever. A play of this infinite game can therefore be identified with a sequence of 0s and 1s. Such sequence can be interpreted as the binary expansion of a real number x satisfying $0 \leq x \leq 1$. For a given set $E \subset \mathbb{R}$, player I wins if $x \in E$ but loses if $x \notin E$.

(i) Begin to draw the game tree.

See Figure 10.

(ii) Bonus: Prove that the second player has a winning strategy in the case that E is the set of rational numbers in the unit interval $[0, 1]$.

The binary expansion of every rational number consists of a finite (possibly of length zero) initial part followed by an infinite repetition of some finite string. This repeating part may consist of just zeros, as is e.g. the case for $\frac{1}{2}$.

The second player will want to ensure that $x \notin E$, which means that x must be made irrational. An easy way to achieve this, is for player II to enumerate the binary expansion of the irrational number π: since this sequence of bits does not repeat, player I has no hope of creating an infinite repeating sequence. Note that the binary expansion of π is computably enumerable, thus providing player II with a viable winning strategy.

(iii) BONUS: How about countable sets E in general?

Since both players choose an infinite number of bits, they are essentially interleaving the binary expansion of two reals. The players therefore have an uncountable ($2^{\aleph_0}$) number of strategies. Hence there is no countable set E for which player I can prevent player II from placing x in the complement of E (whether on purpose or by accident). Clearly, player I does not have a winning strategy.

(iv) BONUS: What happens if E is an interval $[0, y)$ for some y in $(0, 1)$?

If $E = [0, y)$ for some $0 < y < 1$, then player I wants to ensure that $x < y$, while player II will aim for $x \geq y$.

Let $x = x_1 x_2 \ldots$, such that x_{2i+1} are the bits played by player I (for all $i \geq 0$). Let $x \in E$, and hence $x < y$. Now, suppose that $x_{2i+1} = 1$ for some $i \geq 0$. Let $x'_{2i+1} = 0$ and define $x' = x_1 x_2 \ldots x_{2i} x'_{2i+1} x_{2i+2} x_{2i+3} \ldots$. Then clearly $x' < x < y$ and hence $x' \in E$. This holds for all $i \geq 0$, and hence player I's strategy to play only zeros is at least as good as any other, regardless of the strategy of player II: it is optimal. Similarly the optimal strategy for player II is to play only ones.

Let player I and II employ their optimal strategies. Then

$$x = 0.0101010\ldots = \sum_{i=1}^{\infty} 2^{-2i}.$$

Hence player I has a winning strategy if and only if $y > \sum_{i=1}^{\infty} 2^{-2i}$. Otherwise player II wins.

Exercise 2.6.4: A bookie unlikely to make any profit...

- Situation:

> A bookie offers odds of $a_k : 1$ against the kth horse in a race being the winner. There are n horses in the race, and
>
> $$\frac{1}{a_1 + 1} + \frac{1}{a_2 + 1} + \cdots \frac{1}{a_n + 1} < 1.$$

(i) How should you bet if you want to be certain of winning?

For any $m > 0$, one should bet $\frac{m}{a_k+1}$ on the kth horse, for all $1 \leq k \leq n$. Let w be the winning horse. First note that

$$\frac{1}{a_1 + 1} + \frac{1}{a_2 + 1} + \cdots + \frac{1}{a_n + 1} < 1,$$

and so,

$$\frac{m}{a_1 + 1} + \frac{m}{a_2 + 1} + \cdots + \frac{m}{a_n + 1} < m.$$

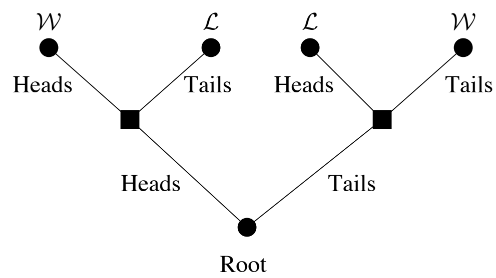

Figure 11: Game tree for the chance game of exercise 2.6.7.

Now, since w is the winning horse, we receive $r = \frac{ma_w}{a_w+1}$ from the bookie for horse k. We will need to pay

$$\begin{aligned} p &= \frac{m}{a_1+1} + \frac{m}{a_2+1} + \cdots + \frac{m}{a_{w-1}+1} + \frac{m}{a_{w+1}+1} + \cdots + \frac{m}{a_n+1} \\ &< m - \frac{m}{a_w+1} = \frac{m(a_w+1)}{a_w+1} - \frac{m}{a_w+1} = \frac{ma_w}{a_w+1} = r. \end{aligned}$$

We receive more that we have to pay, so the payoff is positive.

Exercise 2.6.7: The role of independence in the model

- Game:

  Figure 11 illustrates a game with only chance moves. Each chance move represents the independent toss of a fair coin.

(i) Express the situation as a simple lottery.

See Figure 12(a).

(ii) Express the situation as a simple lottery in the case that all chance moves refer to a single toss of the same coin.

See Figure 12(b).

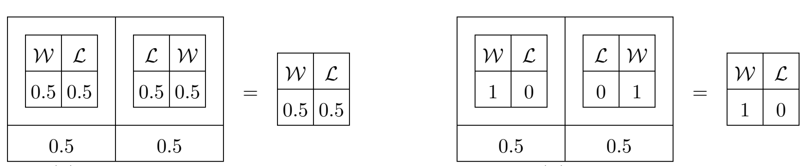

(a) Independent coin tosses.

(b) Single coin toss.

Figure 12: Lotteries for the game of exercise 2.6.7.

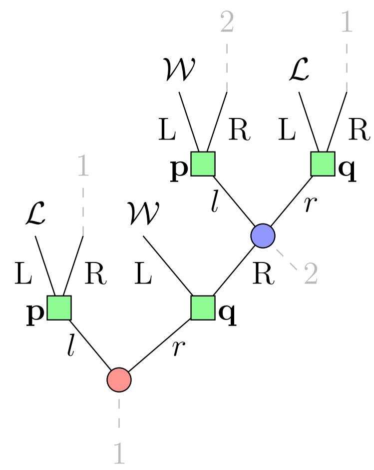

Figure 13: Concise representation of the game of exercise 2.6.11. The labeled dashed gray lines at the top of the tree indicate how the tree continues indefinately, by matching up with the numbers at the bottom of the tree.

Exercise 2.6.11: Analysis of a possibly infinite game with probabilistic moves

- Game:

  Player I can choose l or r at the first move in a game G. If he chooses l, a chance move selects L with probability p or R with probability $1 - p$. If L is chosen, the game ends with outcome $\mathcal{L}$. If R is chosen, a subgame identical in structure to G is played. If player I chooses r, then a chance move selects L with probability q or R with probability $1 - q$. If L is choosen, the game ends in the outcome $\mathcal{W}$. If R is chosen, a subgame is played that is identical to G, *except* that the outcomes $\mathcal{W}$ and $\mathcal{L}$ are interchanged together with the roles of player I and II.

(i) Begin the game tree.

The bottom red circle in Figure 13 represents the root of the game tree. By adding gray dashed lines to the tree, the figure gives a concise representation of the whole game. For a somewhat better understanding of how to interpret the diagram, read its caption.

(ii) Why is this an infinite game?

The game may potentially continue forever. For example, if $p = q = 0$, then there is no way the game finishes.

(iii) With what probability will the game continue forever if player I always chooses l?

If player I always chooses l, then after every turn there is a probability p of losing. Player II never gets to play. If $p = 0$, then obviosuly the game continues forever. If $p > 0$, then $\lim_{n\to\infty}(1 - p)^n = 0$, so if player I sticks with this strategy, then the probability that the game continues forever is 0.

(iv) If the value of G is $\mathbf{v}$, show that $v = q + (1 - q)(1 - v)$.

If player I chooses l, then he has no chance of winning. If he chooses r, then he has probability q of winning, and probability $1-q$ that the tables turn. In the latter case player II will have probability v of winning, which means that player I wins with probability $1-v$. Hence indeed $v = q+(1-q)(1-v)$.

(v) Work out the probability v that player I will win if both players use optimal strategies.

Neither player is ever served by choosing l, while r always gives the possibility of winning. So both players only choose r. Hence $v = q+(1-q)(1-v)$ as shown above. Then,

$$v = q + (1 - q)(1 - v) = q + qv - q - v + 1 = qv - v + 1.$$

So $2v = qv + 1$, which can be rewritten to $2 = q + \frac{1}{v}$. Hence $\frac{1}{v} = 2 - q$, which allows us to conlude that $v = \frac{1}{2-q}$.

(vi) What is v when $q = \frac{1}{2}$?

Following the above answer, we have $v = \frac{1}{2-q} = \frac{1}{2-\frac{1}{2}} = \frac{2}{3}$.

Exercise 2.6.25: Alternative set of Gale's Roulette wheels

- Game:

> In an alternative version of Gale's Roulette, each of the three roulette wheels is labeld with *four* equally likely numbers. The numbers on the first wheel are 2, 4, 6 ad 9; those on the second wheel are 1, 5, 6 ad 8; and those on the third wheel are 3, 4, 5 and 7. If the two wheels chosen by the players stop on the same number, the wheels are spun again and again until someone is a clear winner.

(i) If player I chooses the first wheel and player II chooses the second wheel, show that the probability p that player I will win satisfies $p = \frac{1}{2} + \frac{1}{16}p$.

There are 16 outcomes. In eight of these instances player I is the clear winner: $(2, 1)$, $(4, 1)$, $(6, 1)$, $(6, 5)$, $(9, 1)$, $(9, 5)$, $(9, 6)$, $(9, 8)$. This makes up for half of the event space. In one case there is a draw, which will require the wheels to be spun again. Then player I again has probability p of winning. Hence indeed $p = \frac{1}{2} + \frac{1}{16}p$.

(ii) What is the probability that player I will win the whole game if both players choose optimally?

Table 5 shows the probability p that player I wins the game, for each possible combination of selected disks. Knowing that player II will choose rationally, player I will select disk 1, because this ensures him the largest chance of winning. In this case it does not matter which disk the second player chooses.

Now, we have $p = \frac{1}{2} + \frac{1}{16}p$, which rewrites to $\frac{15}{16}p = \frac{1}{2}$, and so we conclude that $p = \frac{1}{2} \cdot \frac{16}{15} = \frac{8}{15}$.

Table 5: Probability p of player I winning for each combination of selected disks.

Player I	Player II		
	Disk 1	Disk 2	Disk 3
Disk 1		$p = \frac{1}{2} + \frac{1}{16}p$	$p = \frac{1}{2} + \frac{1}{16}p$
Disk 2	$p = \frac{7}{16} + \frac{1}{16}p$		$p = \frac{9}{16} + \frac{1}{16}p$
Disk 3	$p = \frac{7}{16} + \frac{1}{16}p$	$p = \frac{3}{8} + \frac{1}{16}p$	

Exercise 3.7.{9,12}: Definition von Neumann and Morgenstern utility function

- Situation:

> The preferences of a rational person satisfy $\mathcal{L} \prec \mathcal{D}_1 \prec \mathcal{D}_2 \prec \mathcal{W}$. The person regards $\mathcal{D}_1$ and $\mathcal{D}_2$ as being equivalent to certain lotteries whose only prizes are $\mathcal{L}$ and $\mathcal{W}$ (see Figure 14).

(i) Find a Von Neumann and Morgenstern utility function u that represents these preferences.

Since the preferences of the agent satisfy $\mathcal{L} \prec \mathcal{D}_1 \prec \mathcal{D}_2 \prec \mathcal{W}$, we must have $u(\mathcal{L}) < u(\mathcal{D}_1) < u(\mathcal{D}_2) < u(\mathcal{W})$. Let $u(\mathcal{L}) = 0$ and $u(\mathcal{W}) = 1$. Now, $\mathcal{D}_1$ and $\mathcal{D}_2$ represent lotteries in which $\mathcal{W}$ occurs with probability 0.4 and 0.8, respectively. Hence we define $u(\mathcal{D}_1) = 0.4$ and $u(\mathcal{D}_2) = 0.8$. Now u is a Von Neumann and Morgenstern utility function.

(ii) What is the user's preference between the lotteries $\mathbf{L}_1$ and $\mathbf{L}_2$ that are depiced in Figure 14?

Since u defined above is a Von Neumann and Morgenstern utility function, $\mathcal{E}u$ is a utility function for the agent's preferences over lotteries. So we calculate:

$$\mathcal{E}u(\mathbf{L}_1) = 0.25 \cdot 0 + 0.25 \cdot 0.4 + 0.25 \cdot 0.8 + 0.25 \cdot 1 = 0.55$$

$$\mathcal{E}u(\mathbf{L}_2) = 0.2 \cdot 0 + 0.15 \cdot 0.4 + 0.5 \cdot 0.8 + 0.15 \cdot 1 = 0.61$$

We see that $\mathcal{E}u(\mathbf{L}_1) < \mathcal{E}u(\mathbf{L}_2)$ and hence $\mathbf{L}_1 \prec \mathbf{L}_2$. So lottery $\mathbf{L}_2$ is preferred.

$\mathcal{D}_1 \sim$

$\mathcal{L}$	$\mathcal{W}$
0.6	0.4

$\mathcal{D}_2 \sim$

$\mathcal{L}$	$\mathcal{W}$
0.2	0.8

$\mathbf{L}_1 =$

$\mathcal{L}$	$\mathcal{D}_1$	$\mathcal{D}_2$	$\mathcal{W}$
.25	.25	.25	.25

$\mathbf{L}_2 =$

$\mathcal{L}$	$\mathcal{D}_1$	$\mathcal{D}_2$	$\mathcal{W}$
.20	.15	.50	.15

Figure 14: Lotteries for exercise 3.7.{9,12}.

$\mathbf{J}$ =

5m	1m	0m
0	1	0

$\mathbf{K}$ =

5m	1m	0m
.01	.89	.10

$\mathbf{L}$ =

5m	1m	0m
.89	.11	0

$\mathbf{M}$ =

5m	1m	0m
.9	0	.1

Figure 15: Altered lotteries of Allais' Paradox, for exercise 3.7.14. Prizes are in dollars.

(iii) If $u\colon \Omega \to \mathbb{R}$ is a Von Neumann and Morgenstern utility function and $\mathbf{L}$ is a lottery, explain why it is true that $\mathcal{E}(Au(\mathbf{L}) + B) = A\mathcal{E}u(\mathbf{L}) + B$.

Please recall that a Von Neumann and Morgenstern utility function defined over any lottery $\mathbf{L}$ can be reduced to a Von Neumann and Morgenstern utility function over a simple lottery $\mathbf{p}$ in which the only prizes are drawn from the set $\Omega = \{\mathcal{W}, \mathcal{L}\}$. then, for a simple lottery, the following holds,

$$\begin{aligned} A\mathcal{E}u(\mathbf{p}) + B &= A(pu(\mathcal{W}) + (1-p)u(\mathcal{L})) + B \\ &= pAu(\mathcal{W}) + (1-p)Au(\mathcal{L}) + B \\ &= \mathcal{E}(Au(\mathbf{p}) + B). \end{aligned}$$

Exercise 3.7.14: Modifying the lotteries in Allais' Paradox

- Situation:

  The prizes \$0m and \$5m in the lotteries of Allais Paradox are reversed, as in Figure 15.

(i) Would Savage's original preferences ($\mathbf{J} \succ \mathbf{K}$ and $\mathbf{L} \prec \mathbf{M}$) still be inconsistent in this setting?

Yes, they are. Fix $u(\$0\text{m}) = 0$ and $u(\$5\text{m}) = 1$. Let $u(\$1\text{m}) = x$. Then his preferences $\mathbf{J} \succ \mathbf{K}$ and $\mathbf{L} \prec \mathbf{M}$ still impose a contradiction. We first note that:

$$\begin{aligned} \mathcal{E}u(\mathbf{J}) &= 0 \cdot u(\$0\text{m}) + 1 \cdot u(\$1\text{m}) + 0 \cdot u(\$5\text{m}) = x \\ \mathcal{E}u(\mathbf{K}) &= 0.10 \cdot u(\$0\text{m}) + 0.89 \cdot u(\$1\text{m}) + 0.01 \cdot u(\$5\text{m}) = 0.89x + 0.01 \\ \mathcal{E}u(\mathbf{L}) &= 0 \cdot u(\$0\text{m}) + 0.11 \cdot u(\$1\text{m}) + 0.89 \cdot u(\$5\text{m}) = 0.11x + 0.89 \\ \mathcal{E}u(\mathbf{M}) &= 0.1 \cdot u(\$0\text{m}) + 0 \cdot u(\$1\text{m}) + 0.9 \cdot u(\$5\text{m}) = 0.9 \end{aligned}$$

The first constraint on u is that $\mathcal{E}u(\mathbf{J}) > \mathcal{E}u(\mathbf{K})$, since $\mathbf{J} \succ \mathbf{K}$. This implies:

$$\begin{aligned} x &> 0.89x + 0.01 \\ 0.11x &> 0.01 \\ x &> \frac{1}{11} \end{aligned}$$

However, the second constraint on u is that $\mathcal{E}u(\mathbf{L}) < \mathcal{E}u(\mathbf{M})$. Hence,

$$\begin{aligned} 0.11x + 0.89 &< 0.9 \\ 0.11x &< 0.01 \\ x &< \frac{1}{11} \end{aligned}$$

Obviously, the constraints $x > \frac{1}{11}$ and $x < \frac{1}{11}$ cannot be satisfied at the same time. So in the new setting Savage's preferences would also be inconsistent.

Additional exercise for chapter 3

- Situation:

Assume that the utility of \$$x$ for an agent equals $x^{\frac{2}{3}}$.

(i) Evaluate the amount of money this agent is willing to pay for participating in the St. Petersburg Lottery **L**.

The agent's expected utility for the St. Petersburg Lottery is:

$$\begin{aligned} \mathcal{E}u(\mathbf{L}) &= \frac{1}{2}u(\$2) + \left(\frac{1}{2}\right)^2 u(\$2^2) + \left(\frac{1}{2}\right)^3 u(\$2^3) + \cdots \\ &= \frac{1}{2}2^{\frac{2}{3}} + \left(\frac{1}{2}\right)^2 (2^2)^{\frac{2}{3}} + \left(\frac{1}{2}\right)^3 (2^3)^{\frac{2}{3}} + \cdots \\ &= \frac{1}{2}2^{\frac{2}{3}} + \left(\frac{1}{2}\right)^2 \left(2^{\frac{2}{3}}\right)^2 + \left(\frac{1}{2}\right)^3 \left(2^{\frac{2}{3}}\right)^3 + \cdots \\ &= \frac{1}{2}2^{\frac{2}{3}} + \left(\frac{1}{2}2^{\frac{2}{3}}\right)^2 + \left(\frac{1}{2}2^{\frac{2}{3}}\right)^3 + \cdots \\ &= 2^{-\frac{1}{3}} + \left(2^{-\frac{1}{3}}\right)^2 + \left(2^{-\frac{1}{3}}\right)^3 + \cdots \\ &= 2^{-\frac{1}{3}} \cdot \left(1 + 2^{-\frac{1}{3}} + \left(2^{-\frac{1}{3}}\right)^2 + \left(2^{-\frac{1}{3}}\right)^3 + \cdots\right) \\ &= 2^{-\frac{1}{3}} \cdot \frac{1}{1 - 2^{-\frac{1}{3}}} = \frac{2^{-\frac{1}{3}}}{1 - 2^{-\frac{1}{3}}} = \frac{1}{2^{\frac{1}{3}} - 1} \approx 3.84732 \end{aligned}$$

The player will be indifferent between the lottery **L** and a sum of money

x iff their utilities are the same. Thus we calculate,

$$u(\$x) = \mathcal{E}u(\mathbf{L})$$
$$x^{\frac{2}{3}} = \frac{1}{2^{\frac{1}{3}} - 1}$$
$$x \approx 7.54637$$

We conclude that the agent is willing to pay at most a little over \$7.5 for participating in the lottery.

(ii) Show that the Paradox is not solved for an agent ascribing utility $\frac{x}{\log x}$ for an amount of $\$x$ (for $x > 1$).

Let us assume the base 2 logarithm. Then,

$$\mathcal{E}u(\mathbf{L}) = \sum_{n=1}^{\infty} \left(\frac{1}{2}\right)^n u(\$2^n) = \sum_{n=1}^{\infty} \frac{1}{2^n} \frac{2^n}{\log_2 2^n} = \sum_{n=1}^{\infty} \frac{1}{\log_2 2^n} = \sum_{n=1}^{\infty} \frac{1}{n}.$$

This is the harmonic series, which is divergent and hence sums to infitity (this was proved by Nicole Oresme). For any base larger than 2 (specifically e and 10), it is obvious that the series goes to infinity as well: the divisor grows more slowly, and hence the series will grow faster.

Hence $\mathcal{E}u(\mathbf{L})$ is infinite, and thus the Paradox is not solved.

Exercise 4.8.2: Cat and Mouse Game

- Game:

Jerry can hide in the bedroom, the den or the kitchen. Tom can search in one and only one of these locations. If he searches where Jerry is hiding, he catches Jerry for certain. Otherwise Jerry escapes.

(i) Assign appropriate Von Neumann and Morgenstern utilities to the possible outcomes.

Let b, d, k (B, D, K) signify Jerry's (Tom's) choice to go to the bedroom, den or kitchen, respectively. Then, since Jerry does not want to be caught, his utility function u_J can be defined as,

$$u_J(bD) = u_J(bK) = u_J(dB) = u_J(dK) = u_J(kB) = u_J(kD) = 1$$
$$u_J(bB) = u_J(dD) = u_j(kK) = 0.$$

Naturally, Tom's utility function is defined as $u_T(x) = 1 - u_J(x)$.

(ii) Draw the game tree for the case when Tom can see where Jerry hides before searching. Find the 3×27 bimatrix game that is the corresponding strategic form. (Jerry is player I.)

The game tree is show in Figure 16(a). The bimatrix game is depicted in Table 6. Note that the table is cut in three due to space considerations.

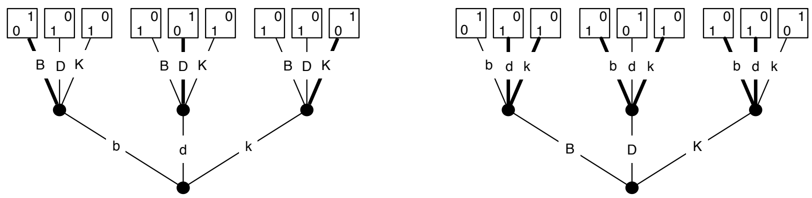

(a) Tom knows. (b) Jerry knows.

Figure 16: Game trees for exercise 4.8.2, which correspond to the cases when Tom and Jerry make their decisions knowing the other's choice.

Table 6: The 3×27 bimatrix game for exercise 4.8.2 that is the strategic form corresponding to the situation in which Tom can see where Jerry hides.

	BBB	BBD	BBK	BDB	BDD	BDK	BKB	BKD	BKK
b	0, ①	0, ①	0, ①	0, ①	0, ①	⓪, ①	0, ①	0, ①	0, ①
d	①, 0	①, 0	①, 0	0, ①	0, ①	⓪, ①	①, 0	①, 0	①, 0
k	①, 0	①, 0	0, ①	①, 0	①, 0	⓪, ①	①, 0	①, 0	0, ①

	DBB	DBD	DBK	DDB	DDD	DDK	DKB	DKD	DKK
b	①, 0	①, 0	①, 0	①, 0	①, 0	①, 0	①, 0	①, 0	①, 0
d	①, 0	①, 0	①, 0	0, ①	0, ①	0, ①	①, 0	①, 0	①, 0
k	①, 0	①, 0	0, ①	①, 0	①, 0	0, ①	①, 0	①, 0	0, ①

	KBB	KBD	KBK	KDB	KDD	KDK	KKB	KKD	KKK
b	①, 0	①, 0	①, 0	①, 0	①, 0	①, 0	①, 0	①, 0	①, 0
d	①, 0	①, 0	①, 0	0, ①	0, ①	0, ①	①, 0	①, 0	①, 0
k	①, 0	①, 0	0, ①	①, 0	①, 0	0, ①	①, 0	①, 0	0, ①

Table 7: Transposition of the 27×3 bimatrix game for exercise 4.8.2 that is the strategic form corresponding to the situation in which Jerry can see where Tom is searching.

	bbb	*bbd*	*bbk*	*bdb*	*bdd*	*bdk*	*bkb*	*bkd*	*bkk*
B	(1), 0	(1), 0	(1), 0	(1), 0	(1), 0	(1), 0	(1), 0	(1), 0	(1), 0
D	0, (1)	0, (1)	0, (1)	(1), 0	(1), 0	(1), 0	0, (1)	0, (1)	0, (1)
K	0, (1)	0, (1)	(1), 0	0, (1)	0, (1)	(1), 0	0, (1)	0, (1)	(1), 0

	dbb	*dbd*	*dbk*	*ddb*	*ddd*	*ddk*	*dkb*	*dkd*	*dkk*
B	(0), (1)	(0), (1)	0, (1)	0, (1)	0, (1)	0, (1)	(0), (1)	(0), (1)	0, (1)
D	(0), (1)	(0), (1)	0, (1)	(1), 0	(1), 0	(1), 0	(0), (1)	(0), (1)	0, (1)
K	(0), (1)	(0), (1)	(1), 0	0, (1)	0, (1)	(1), 0	(0), (1)	(0), (1)	(1), 0

	kbb	*kbd*	*kbk*	*kdb*	*kdd*	*kdk*	*kkb*	*kkd*	*kkk*
B	(0), (1)	(0), (1)	0, (1)	0, (1)	0, (1)	0, (1)	(0), (1)	(0), (1)	0, (1)
D	(0), (1)	(0), (1)	0, (1)	(1), 0	(1), 0	(1), 0	(0), (1)	(0), (1)	0, (1)
K	(0), (1)	(0), (1)	(1), 0	0, (1)	0, (1)	(1), 0	(0), (1)	(0), (1)	(1), 0

(iii) Draw the game tree for the case when Jerry can see where Tom searches before hiding. Find the 27×3 bimatrix game that is the corresponding strategic form.

The game tree is show in Figure 16(b). The bimatrix game is depicted in Table 7. Note that the table is transposed and cut in three due to space considerations.

(iv) Draw two game trees that both correspond to the case when Tom and Jerry each make their decisions in ignorance of the other's choice. Find the 3×3 bimatrix game that is the corresponding strategic form.

Figure 17(a) shows the first game tree, in which Tom is the first to make his move. Figure 17(b) shows the other game tree, in which Jerry is first to move. The bimatrix game is depicted in Table 8.

(v) In each case, find all pure strategy pairs that are Nash equilibria.

The previously mentioned Tables 6, 7 and 8 have the payoffs of each player's best responses given the other player's action encircled. Cells with two circles correspond to strategy pairs that are Nash equilibria.

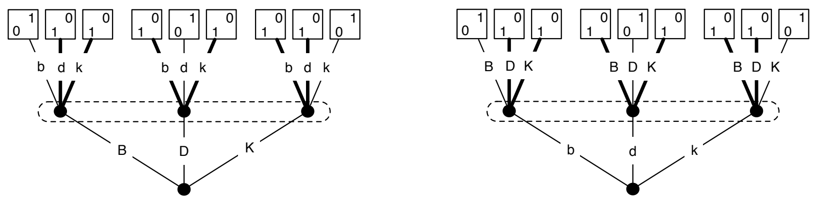

(a) Tom moves first. (b) Jerry moves first.

Figure 17: Game trees for exercise 4.8.2, which correspond to the case when Tom and Jerry each make their decisions in ignorance of the other's choice.

Exercise 4.8.16: One of the recurring Blotto-Baloney campaigns

- Game:

> Colonel Blotto can send each of his five companies to one of ten locations whose importance is valued at $1, 2, 3, \ldots 10$ respectively. No more than one company can be sent to any location. His opponent, Count Baloney, must simultaniously do the same with his four companies. A commander who attacks an undefended location, captures it. If both commanders attack the same location, the result is a stand-off at that location. A commander's payoff is the sum of the values of of the locations he captures minus the sum of the values of the locations captured by the enemy.

(i) What would Colonel Blotto do if he knew what a dominated strategy was?

First, let us see what the effect is of sending a company to a location valued at n, instead of not sending the company at all. The opponent may or may not send a company to the same location. If the opponent does send a company, then countering the attack results in a stand-off, adding 0 to the payoff. If no company is send to counter the attack, then the payoff will be diminished by n. On the other hand, if the opponent does not send a company to the location valued at n, then sending a company will result in a capture, thus increasing the payoff by n, versus a payoff increase of 0 if no company is sent.

Table 8: The 3×3 bimatrix game for exercise 4.8.2 that is the strategic form corresponding to the situation in which Tom and Jerry are in ignorance of the other's choice.

	B	D	K
b	0, (1)	(1), 0	(1), 0
d	(1), 0	0, (1)	(1), 0
k	(1), 0	(1), 0	0, (1)

Table 9: The strategic form of the game of exercise 5.9.10.

	dD	dU	uD	uU
dD	4, 4	0.5, 5	2.5, 1	2.5, 1
dU	5, 0.5	2.5, 2.5	2.5, 1	2.5, 1
uD	1, 2.5	1, 2.5	1.5, 1.5	1.5, 1.5
uU	1, 2.5	1, 2.5	1.5, 1.5	1.5, 1.5

Hence, in each instance, sending a company to a location valued at n yields a payoff that is n higher than the payoff of not sending the company at all.

Now, suppose Colonel Blotto sends a company to a location valued at $n \leq 5$. Then there will be some location valued at $6 \leq m \leq 10$ to which he does not send a company (since he has only five companies).

If Colonel Blotto would change his strategy by sending the company to the location valued at m instead of the location valued at n, then this would give him a netto increase in payoff of $m - n > 0$. Hence sending a company to the location valued at n is a dominated strategy.

If we combine this with the insight that all companies should be sent out in order to achieve the highest payoff, we conclude that Colonel Blotto has only one non-dominated strategy, namely sending his companies to the locations valued at 6, 7, 8, 9 and 10.

Exercise 5.9.10: An economics model

- Game:

Consider a game's strategic form as presented in Table 9.

(i) Explain why (dD, dD) is a Pareto-efficient pure strategy pair of the game.

Recall that an outcome x in a feasible set X is Pareto-efficient if there is no other outcome y in X that both players like at least as much and at least one player likes more than x. More formally, Pareto-efficiency requires the condition that, for all y,

$$y > x \Rightarrow y \notin X.$$

It is easy to infer from Table 9 that this condition holds for (dD, dD). Although some players may prefer the outcomes (dU, dD) or (dD, dU), not all players like these outcomes at least as much as (dD, dD).

(ii) Is this a Nash equilibrium of the game?

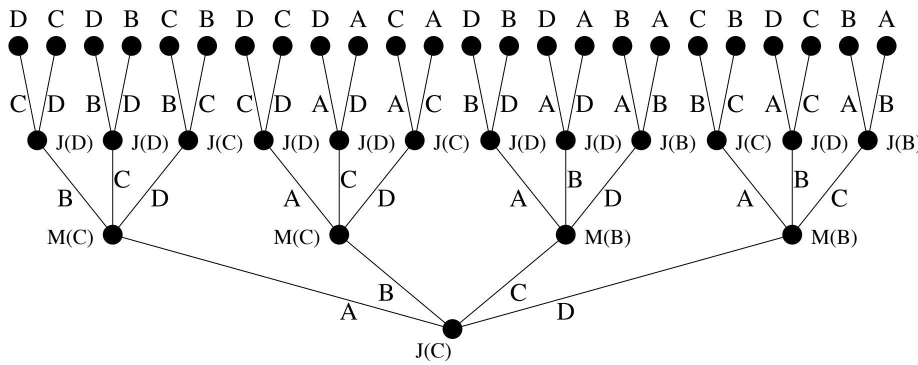

Figure 18: Game tree for exercise 5.9.27 in the case that John (J) gets to veto first. The vetos $10:90$, $20:80$, $50:50$ and $60:40$ are represented by A, B, C and D respectively.

A strategy pair (s,t) is a Nash equilibrium in this game if s is a best reply to t and t is a best reply to s. Table 9 shows that this is not the case for the stratey pair (dD, dD), since player I will prefer strategy dU over dD in reply to strategy dD of player II. The same preference also holds for player II when player I's strategy is dD. Therefore (dD, dD) is not a Nash equilibrium.

Exercise 5.9.27: Another Divide the Dollar Game

- Game:

Suppose John and Mary play the following bargaining game about the division of a dollar donated by a philanthropist. The philanthropist specifies that only the splits $10:90$, $20:80$, $50:50$ and $60:40$ are to be permitted. Moreover, John and Mary are to alternate in vetoing splits that they regard as unacceptable.

(i) What split will result if subgame-perfect strategies are used and John has the first opportunity to veto?

Figure 18 shows the game tree of the bargaining game in the case that John (J) gets to veto first. Every edge represents an action and is labeled with the split that is vetoed. For clearity the actions are labeled with a character (A means veto $10:90$, B means veto $20:80$, C means veto $50:50$ and D means veto $60:40$). Every output node is labeled with the split that results from that branch of the game tree (the labels refer to the same splits as with the edges). Every internal node is labeled with the player and the value of the node fot this player determined with backward induction.

The game tree shows that John will favor vetoing either $10:90$ or $20:80$, these actions have the highest values for John. Mary (M) will respond by vetoing the split that is least attractive for her: $60:40$. John will finish with vetoing the unfavourable split he tolerated during the first move. The resulting split is $50:50$.

(ii) What split will result if Mary begins?

When Mary gets to move first, we can determine the outcome in a similar fashion. In this case, Mary is able to veto two of her most unfavourable splits while John can only veto one. And so Mary will veto 60 : 40 and 50 : 50 while John will veto 10 : 90, which results in a split of 20 : 80.

Exercise 6.10.4: Determination of minimax and maximin values

- Situation:

Assume the following matrices:

$$A = \begin{pmatrix} 1 & 2 \\ 3 & 4 \end{pmatrix} \qquad C = \begin{pmatrix} 2 & 4 & 6 & 3 \\ 6 & 2 & 4 & 3 \\ 4 & 6 & 2 & 3 \end{pmatrix}$$

$$B = \begin{pmatrix} 1 & 3 \\ 4 & 2 \end{pmatrix} \qquad D = \begin{pmatrix} 3 & 2 & 2 & 1 \\ 2 & 3 & 2 & 1 \\ 2 & 2 & 3 & 1 \end{pmatrix}$$

(i) Find the maximin and minimax values of these matrices.

We determine the minimax $\overline{m}$ and maximin $\underline{m}$ values of four zero-sum games. The games are represented by the matrices above and we denote the entry in row $s \in S_M$ and column $t \in T_M$ of each matrix M by $\pi_M(s,t)$.

$$\begin{aligned}
\overline{m}_A &= \min_{t \in T_A}\{\max_{s \in S_A} \pi_A(s,t)\} = \min\{3,4\} = 3 \\
\underline{m}_A &= \max_{t \in T_A}\{\min_{s \in S_A} \pi_A(s,t)\} = \max\{1,3\} = 3 \\
\overline{m}_B &= \min_{t \in T_B}\{\max_{s \in S_B} \pi_B(s,t)\} = \min\{4,3\} = 3 \\
\underline{m}_B &= \max_{t \in T_B}\{\min_{s \in S_B} \pi_B(s,t)\} = \max\{1,2\} = 2 \\
\overline{m}_C &= \min_{t \in T_C}\{\max_{s \in S_C} \pi_C(s,t)\} = \min\{6,6,6,3\} = 3 \\
\underline{m}_C &= \max_{t \in T_C}\{\min_{s \in S_C} \pi_C(s,t)\} = \max\{2,2,2\} = 2 \\
\overline{m}_D &= \min_{t \in T_D}\{\max_{s \in S_D} \pi_D(s,t)\} = \min\{3,3,3,1\} = 1 \\
\underline{m}_D &= \max_{t \in T_D}\{\min_{s \in S_D} \pi_D(s,t)\} = \max\{1,1,1\} = 1
\end{aligned}$$

(ii) For which matrices is it true that $\underline{m} < \overline{m}$?

This is true for matrices B and C.

(iii) For which matrices is it true that $\underline{m} = \overline{m}$?

This holds for matrices A and D.

Exercise 6.10.14: Dominance by Mixed Strategies

- Situation:

> If one of the matrices in Exercise 6.10.4 is taken to be player II's payoff matrix in a game, then player II has a pure strategy that is strongly dominated by a mixed strategy but not by any pure strategy.

(i) Which of A, B, C and D is the matrix with this property? What is the dominated pure strategy? What is the dominating mixed strategy?

The game matrix C of exercise 6.10.4 is given by:

	t_1	t_2	t_3	t_4
s_1	2	4	6	3
s_2	6	2	4	3
s_3	4	6	2	3

Player II's strategy t_4 is not strongly dominated by any other pure strategy, but is strongly dominated by the mixed strategy $q = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3}, 0)^\top$. For any pure strategy of player I, the expected outcome of player II with strategy t_4 is 3 while his expected outcome with strategy q is $2 \cdot \frac{1}{3} + 4 \cdot \frac{1}{3} + 6 \cdot \frac{1}{3} = 4$.

Exercise 6.10.15: Computation of a mixed security strategy

- Situation:

> Player I's payoff matrix in a game is
>
> $$\begin{pmatrix} 1 & 2 & 3 & 4 & 5 \\ 9 & 7 & 5 & 3 & 1 \end{pmatrix}$$
>
> The matrix has no saddle point, and hence player I's security strategies are mixed.

(i) Find player I's security level in the game and a mixed security strategy for player I.

Since there is no saddle point we will have to work out the expected payoff $x = E_k(r)$ that player I will get when player II will play her pure strategy and he uses the mixed strategy $(1-r, r)$. We then have the following:

$$\begin{aligned} E_1(r) &= 1(1-r) + 9r = 1 + 8r \\ E_2(r) &= 2(1-r) + 7r = 2 + 5r \\ E_3(r) &= 3(1-r) + 5r = 3 + 2r \\ E_4(r) &= 4(1-r) + 3r = 4 - r \\ E_5(r) &= 5(1-r) + r = 5 - 4r \end{aligned}$$

When we graph the above equations (see Figure 19), we see that they all cross at exactly the same point. Therefore to obtain the crossing point we can take any of the above formulae and compare them with each other. For example solving $E_1(r) = E_5(r)$ equals $1 + 8r = 5 - 4r$ which tells us that the lines cross for $r = \frac{1}{3}$.

Now Player 1 can ensure an expected payoff of at least $E_1(\frac{1}{3}) = 3\frac{2}{3}$ (his security level) by playing the mixed strategy of $(\frac{2}{3}, \frac{1}{3})$

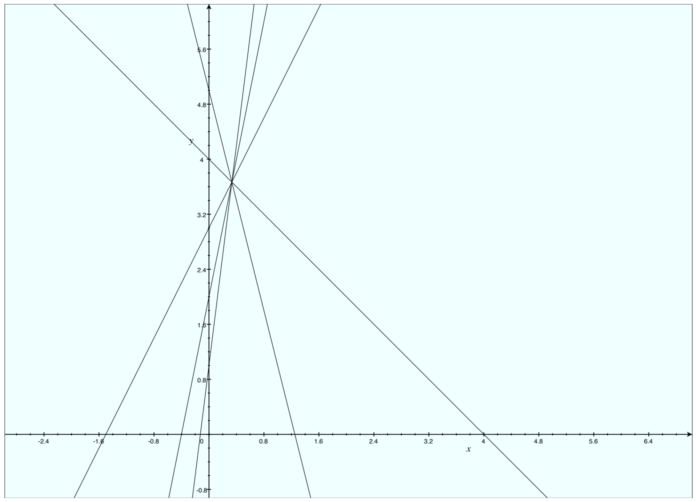

Figure 19: Graph of the equations of exercise 6.10.15.

Exercise 6.10.34: A game on distributed warfare...

- Situation:

> Colonel Blotto has four companies that he can distribute among two locations in three different ways: $(3,1)$, $(2,2)$, $(1,3)$. His opponent, Count Baloney, has three companies that he can distribute among the same two locations in two different ways: $(2,1)$ or $(1,2)$. Suppose that Blotto sends m_1 companies to location 1 and Baloney sends n_1 companies to location 1. If $m_1 = n_1$, then the result is a standoff, and each commander gets a payoff of zero for location 1. If $m_1 \neq n_1$, the larger force overwhelms the smaller force without loss to itself. If $m_1 > n_1$, Blotto gets a payoff n_1 and Baloney gets a payoff of $-n_1$ for location 1. If $m_1 < n_1$, Blotto gets a payoff $-m_1$, and Baloney gets a payoff of m_1 for location 1. Each player's total payoff is the sum of his payoffs at both locations.

(i) Find the strategic form of this simultaneous-move game and show that it has no saddle point.

From the situation described above we can extract the strategic form that is shown in Table 10.

Table 10: Strategic form for the situation described by exercise 6.10.34.

	$(2,1)$	$(1,2)$
$(3,1)$	2, -2	0, 0
$(2,2)$	1, -1	-1, 1
$(1,3)$	0, 0	2, -2

Because this is a zero-sum game, the optimal strategy for a player is a security strategy. Moreover, we know that a pair of security strategies is guaranteed to be a Nash equilibrium.

(ii) Determine a mixed-strategy Nash equilibrium.

Since count Baloney has only two actions to choose from, we can find his security strategy by maximizing the expected payoff with respect to the assignment of probability r of choosing action $(1, 2)$ and probability $1 - r$ of choosing action $(2, 1)$. This leads to the following equations:

$$
\begin{aligned}
E_1(r) &= 2(1-r) + 0r = 2 - 2r \\
E_2(r) &= 1(1-r) + r = 1 \\
E_3(r) &= 0(1-r) + 2r = 2r
\end{aligned}
$$

It is easy to see that the strategy with $r = \frac{1}{2}$ with $\tilde{q} = (\frac{1}{2}, \frac{1}{2})^T$ maximizes Baloney's payoff. Colonel Blotto on the other hand, has three actions to choose from. We can use the technique of seperating hyperplanes to find the optimal strategy for him.

Because we already know the security strategy of Baloney, we can find the value of the game $v = 1$ by looking at the intersection of the seperating line and the line $x_1 = x_2$. We find the security strategy of Blotto by solving the following equation:

$$
\begin{aligned}
\begin{pmatrix} 2 & 1 & 0 \\ 0 & 1 & 2 \end{pmatrix} \begin{pmatrix} \tilde{p}_1 \\ \tilde{p}_2 \\ \tilde{p}_3 \end{pmatrix} &= \begin{pmatrix} 1 \\ 1 \end{pmatrix} \\
\begin{pmatrix} \tilde{p}_1 \\ \tilde{p}_2 \\ \tilde{p}_3 \end{pmatrix} &= \begin{pmatrix} \frac{1}{3} \\ \frac{1}{3} \\ \frac{1}{3} \end{pmatrix}
\end{aligned}
$$

And so we have found a Nash equilibrium with two mixed strategies: $\tilde{p} = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})^T$ for Blotto, and $\tilde{q} = (\frac{1}{2}, \frac{1}{2})^T$ for Baloney.

References

[GLT79] J.R. Gilbert, T. Lengauer, and R.E. Tarjan, *The pebbling problem is complete in polynomial space*, Proceedings of the eleventh annual ACM symposium on Theory of computing (1979), 237–248.

[vEB96] Peter van Emde Boas, *The convenience of tilings*, Tech. report, University of Amsterdam, july 1996.