

Optimal minimax strategy in a dice game

F. Croce* E. Mordecki†

December 30, 2009

Abstract

Each of two players, by turns, rolls a dice several times accumulating the successive scores until he decides to stop, or he rolls an ace. When stopping, the accumulated turn score is added to the player account and the dice is given to his opponent. If he rolls an ace, the dice is given to the opponent without adding any point. In this paper we formulate this game in the framework of competitive Markov decision processes (also known as *stochastic games*), show that the game has a value, provide an algorithm to compute the optimal minimax strategy, and present results of this algorithm in three different variants of the game.

Keywords: Competitive Markov processes, Stochastic games, dice games, minimax strategy.

AMS MSC: 60J10, 60G40, 91A15

1 Introduction

Consider a two-players dice game in which players accumulate points by turns with the following rules. The player who reaches a certain fixed number of points is the winner of the game. In his turn each player rolls the dice several times until deciding to stop or rolling an ace. If he decides to stop the accumulated successive scores are added to his account; while if he rolls an ace no additional points are obtained. As a first approach to find optimal strategies for this game Roters [5] studied the optimal stopping problem corresponding to the maximisation of the expected score in one turn. The optimal solution is a good way of minimising the number of turns required to reach the objective.

Later, Roters & Haigh [3] found the strategy that minimises the expected number of turns required to reach the target. This second strategy is better than the one obtained in [5] but none of them take into account the consideration of the number of points of the opponent, that is clearly relevant in order to win the game.

*Facultad de Ciencias, Centro de Matemática. Iguá 4225, CP 11400 Montevideo. e-mail: fabian@cmat.edu.uy

†Facultad de Ciencias, Centro de Matemática. Iguá 4225, CP 11400 Montevideo. e-mail: mordecki@cmat.edu.uy

In this paper we formulate this game in the framework of competitive Markov decision processes (also known as *stochastic games*), show that the game has a value, provide an algorithm to compute the optimal minimax strategy, and present results of this algorithm in three different variants of the game.

The concept of “stochastic games” was introduced in 1953 by Shapley in [6]. In the recent book by Filar and Vrieze [2], the authors provide a general and modern comprehensive approach to this theory departing from the theory of *controlled Markov processes* (that can be considered “solitaire” stochastic games) and call the type of games we are interested in *competitive Markov decision processes*, a denomination that we find more accurate than the more usual denomination *stochastic game*.

During the preparation of this paper we found the related article by Neller and Presser [4] where the authors, following an heuristic approach, formulate the Bellman equation of the problem (that is a consequence of our results), and compute the optimal strategy of a variant of this game. It must be noted that the theory of Filar and Vrieze [2], that we follow, provide the solution of the problem in the set of all possible strategies, including non-stationary and randomised strategies, i.e. the set of behaviour strategies.

In section 2 we present the theory of *competitive Markov decision processes*, specially in the *transient* case and we conclude the section with the formulation of the theorem we need to solve our dice game. A proof of this theorem can essentially be found in [2]. In section 3 we determine the state space of our game, the corresponding action spaces for each player, the payoff function of the game, and the Markov transitions depending on each state of the process and action of the players. In section 4 we present two related games: first, in order to win, the player has to reach the target exactly (if the target is exceeded, he gives the dice to his opponent without changing his accumulated score); in the second variant the players aim to maximise the difference between their scores. In section 5 we present the conclusions.

2 Competitive Markov decision processes

A *competitive Markov decision process* (also known as a *stochastic game*) is the mathematical model of a sequential game, in which two players take actions considering the status of a certain Markov processes. Both actions determine an immediate payoff for each player and the probability distribution of the following state of the game. Our interest is centred in two-players, finite-state and finite-action, zero-sum games. To define them formally we need the following ingredients:

- (S) States: A finite set S of the possible states of the game.
- (A) Actions: For each state $s \in S$ we consider finite sets A^s and B^s whose elements are the possible actions for the players; at each step both players take his actions simultaneously and independently.
- (P) Payoffs: For each state $s \in S$ a function $r^s : A^s \times B^s \rightarrow \mathbb{R}$ determines the amount that player two has to pay to player one depending on the actions taken by both players.

(TP) Transition probabilities: For each $(s, a, b) \in S \times A^s \times B^s$, $P_{s,a,b}$ is a distribution probability on S , which determines the following state of game.

We denote by S_t the state of the game at time $t = 0, 1, \dots$; the initial state is fixed ($S_0 = s_0$). At each step t players choose actions A_t and B_t in $A^{S_t} B^{S_t}$ respectively, which determine that player two has to pay to player one an amount of $r^{S_t}(A_t, B_t)$, and the distribution probability of S_{t+1} will be P_{S_t, A_t, B_t} .

The random variable $W = \sum_{t=0}^{\infty} r^{S_t}(A_t, B_t)$ is the total amount that player two pays to player one (could be negative). Note that W depends on the way in which players take their actions. The objective of the game for player one is to maximise the expected value V of the accumulated payoff W , while player two has the objective of minimising it. In principle V could be infinite. Often for economic applications the accumulated payoff is $W_\beta = \sum_{t=0}^{\infty} \beta^t r^{S_t}(A_t, B_t)$, called “the discounted sum”, where $0 < \beta < 1$ represent the devaluation of the money. This discount factor β ensures that W is finite with probability one and the existence of its expected value.

In the case of *transient* stochastic games, the situation considered in this paper, the sum defining W is finite a.s. due to the fact that the process always reaches a final state $s_f \in S$ in a (not necessarily bounded) finite number of steps. In the definition of transient stochastic game additional conditions are required, in order to V to be finite. Before the formal definition of transient stochastic games, the concept of behaviour strategy is introduced.

2.1 Strategies

Consider the set \mathbb{K} defined by

$$\mathbb{K} = \{(s, a, b) : s \in S, a \in A^s, b \in B^s\}$$

We define, for each $t = 0, 1, \dots$, the sets \mathbb{H}_t , of admissible histories up to time t , by

$$\mathbb{H}_t = \begin{cases} S & \text{if } t = 0 \\ \underbrace{\mathbb{K} \times \dots \times \mathbb{K}}_{t \text{ times}} \times S & \text{else} \end{cases}$$

Definition 1 (Behaviour strategy) *Given a stochastic game, a behaviour strategy for player one (two) in the game is a function π which associates to each history*

$$h = (s_0, a_0, b_0, \dots, s) \in \cup_{t=0}^{\infty} \mathbb{H}_t$$

a distribution probability $\pi(\cdot|h)$ in A^s (respectively $\varphi(\cdot|h)$ in B^s). In the context of a stochastic game, we denote by Π (Φ), the set of all behaviour strategies for player one (two).

Note that the previous definition is in agreement with the intuitive idea that a player can choose his action based on the history of the game. There are two relevant subclasses of strategies, *pure* and *stationary*, introduced below.

Definition 2 (Pure strategy) A behaviour strategy π is said to be pure, if for each history h there exists an action a_h such as $\pi(a_h|h) = 1$. We could say that a pure strategy chooses the action to be taken in a deterministic way.

Definition 3 (Stationary strategy) A behaviour strategy π is said to be stationary, if the probability distribution $\pi(\cdot|h)$ depends only on s , the last state of the history. In this case we use the notation $\pi(\cdot|s)$.

2.2 Probabilistic framework

We now construct the probability space in which the optimisation procedure takes place. Consider the product space

$$\Omega = (S \times \cup_{s \in S} A^s \times \cup_{s \in S} B^s)^{\mathbb{N}}$$

equipped with the product σ -algebra \mathfrak{F} , defined as the minimal σ -algebra containing the cylinder sets of Ω .

Given $\omega = (s_0, a_0, b_0, s_1, a_1, b_1, \dots) \in \Omega$, a sequence of states and actions in the product space, the coordinate processes

$$\{S_t\}_{t=0,1,\dots}, \{A_t\}_{t=0,1,\dots}, \{B_t\}_{t=0,1,\dots}$$

are defined by

$$S_t(\omega) = s_t, \quad A_t(\omega) = a_t, \quad B_t(\omega) = b_t.$$

In this framework, given π, φ behaviour strategies for players one and two and an initial state $s \in S$, it is possible to introduce a probability $\mathbf{P}_{s,\pi,\varphi}$, such that, for the random vector

$$H_t = (S_0, A_0, B_0, S_1, A_1, B_1, \dots, S_t),$$

and the finite sequence of states and actions $h_t = (s_0, a_0, b_0, \dots, s_t)$, the following assertions hold:

- the game starts in the state s , i.e., $\mathbf{P}_{s,\pi,\varphi}(S_0 = s) = 1$;
- with probability 1, H_t take their values in \mathbb{H}_t ;
- the probability distribution on the actions chosen by players at time t depends on H_t , according to

$$\mathbf{P}_{s,\pi,\varphi}(A_t = a_t | H_t = h_t) = \pi(a_t | h_t)$$

$$\mathbf{P}_{s,\pi,\varphi}(B_t = b_t | H_t = h_t) = \varphi(b_t | h_t)$$

$$\mathbf{P}_{s,\pi,\varphi}(A_t = a_t, B_t = b_t | H_t = h_t) = \pi(a_t | h_t) \varphi(b_t | h_t).$$

- the distribution probability of S_{t+1} depends only on S_t, A_t, B_t , being the transition probabilities (TP) of the game

$$\mathbf{P}_{s,\pi,\varphi}(S_{t+1} = s_{t+1} | H_t = h_t, A_t = a_t, B_t = b_t) = P_{s_t, a_t, b_t}(s_{t+1}).$$

We denote by $\mathbf{E}_{s,\pi,\varphi}$ the expected value in the probability space $(\Omega, \mathfrak{F}, \mathbf{P}_{s,\pi,\varphi})$.

2.3 Transient stochastic games

Definition 4 (Transient stochastic game) *A stochastic game is transient when there exists a final state $s_f \in S$ such that:*

- (1) $r^{s_f}(a, b) = 0, \forall a \in A^{s_f}, \forall b \in B^{s_f};$
- (2) $P_{s_f, a, b}(s_f) = 1, \forall a \in A^{s_f}, \forall b \in B^{s_f};$
- (3) *for all pair of strategies (π, φ) of player one and two respectively and for all initial state s*

$$\sum_{t=0}^{\infty} \mathbf{P}_{s, \pi, \varphi}(S_t \neq s_f) < \infty.$$

Conditions (1) and (2) ensure that, once the game falls into the final state s_f , it never changes the state again and the gain of both players is zero. The third condition ensures that the game finishes with probability one.

Definition 5 (Value of a pair of strategies) *Given $\pi \in \Pi, \varphi \in \Phi$, strategies for players one and two in a transient stochastic game, the value of the strategies is a function $V_{\pi, \varphi} : S \rightarrow \mathbb{R}$ defined by*

$$V_{\pi, \varphi}(s) = \sum_{t=0}^{\infty} \mathbf{E}_{s, \pi, \varphi}(r^{S_t}(A_t, B_t)).$$

Definition 6 (Optimal strategy) *A behaviour strategy π^* for player one in a transient stochastic game is said to be optimal if*

$$\inf_{\varphi \in \Phi} V_{\pi^*, \varphi}(s) = \sup_{\pi \in \Pi} \inf_{\varphi \in \Phi} V_{\pi, \varphi}(s), \quad \forall s \in S.$$

Analogously a behaviour strategy φ^ for player two, is said to be optimal if*

$$\sup_{\pi \in \Pi} V_{\pi, \varphi^*}(s) = \inf_{\varphi \in \Phi} \sup_{\pi \in \Pi} V_{\pi, \varphi}(s), \quad \forall s \in S.$$

We now formulate the result used to solve our dice game.

Theorem 2.1 (Value and optimal strategies) *Given a transient stochastic game the following identity is fulfilled*

$$\sup_{\pi} \inf_{\varphi} V_{\pi, \varphi}(s) = \inf_{\varphi} \sup_{\pi} V_{\pi, \varphi}(s), \quad \text{for all } s \in S. \quad (2.1)$$

The vector defined in (2.1), denoted by $(v(s))_{s \in S}$, is called the value of the game. This value is the unique joint solution of

$$x(s) = \left[r^s(a, b) + \sum_{s' \in S} P_{s, a, b}(s') x(s') \right]_{a \in A^s, b \in B^s}^*,$$

where $[\cdot]^*$ represents the value of the matrix game (in the minimax sense) obtained by considering rows $a \in A^s$ and columns $b \in B^s$. Moreover, the stationary strategies π and φ for players one and two, such that $\pi(\cdot|s)$ and $\varphi(\cdot|s)$ are the optimal strategies of the matrix game

$$\left[r^s(a, b) + \sum_{s' \in S} P_{s,a,b}(s')v(s') \right]_{a \in A^s, b \in B^s}$$

for all state $s \in S$, are optimal strategies in the transient stochastic game.

Proof 1 This theorem is essentially a particular case of theorem 4.2.6 in [2]. A detailed proof can be found at [1].

Remark 2.1 The previous theorem ensures the existence of optimal strategies for both players. Particularly they are in the subclass of stationary strategies. In the proof of this theorem a map U such that

$$Uv(s) = \left[r^s(a, b) + \sum_{s' \in S} P_{s,a,b}(s')v(s') \right]_{a \in A^s, b \in B^s}^*, \quad (2.2)$$

which is a n -step contraction, is considered. Afterwards, we used the map U to implement a numerical method to find a unique fixed point, that is the value of the game.

3 The dice game

In this section we describe the states, actions, payoffs, and transition probabilities (defined in section 2) corresponding to our dice game, and present the numerical results, showing the optimal strategy for a player. This strategy, optimal in the class of behaviour strategies, ensures a player to win with probability of at least $1/2$ independently of the opponent strategy. The optimal strategy is pure and stationary, and consists in a simple rule indicating whether *to roll* or *to stop*, depending on the scores of the player and its opponent.

3.1 Modelling the dice game

To solve the dice game (compute optimal strategies), we model it as a transient stochastic game. We have to specify the set of states, possible actions, payoffs and transition probabilities.

- (S) States: During the dice game, there are four aspect varying: the player j who has the dice ($j = 1, 2$), the accumulated score α of player one, the accumulated score β of player two and the turn score τ of the player j . So, we consider states (j, α, β, τ) . We also need to consider two special states: an initial state s_0 , and a final state s_f .

Table 1: Possible actions for each player depending on the state of the game.

	player one	player two
s_i, s_f	<i>to wait</i>	<i>to wait</i>
$(1, \alpha, \beta, 0)$	<i>to roll</i>	<i>to wait</i>
$(1, \alpha, \beta, \tau)_{0 < \tau < 200 - \alpha}$	<i>to roll, to stop</i>	<i>to wait</i>
$(1, \alpha, \beta, \tau)_{\tau \geq 200 - \alpha}$	<i>to stop</i>	<i>to wait</i>
$(2, \alpha, \beta, 0)$	<i>to wait</i>	<i>to roll</i>
$(2, \alpha, \beta, \tau)_{0 < \tau < 200 - \beta}$	<i>to wait</i>	<i>to roll, to stop</i>
$(2, \alpha, \beta, \tau)_{\tau \geq 200 - \beta}$	<i>to wait</i>	<i>to stop</i>

If the score of either of the players is greater or equal than 200 the game is over, then, it is in the state s_f . Because of that, states (j, α, β, τ) , only make sense if $\alpha < 200$, $\beta < 200$. The same happen if τ is big enough to reach 200 stopping. So the finite set S of possible states is

$$S = \{s_0, s_f\} \cup S_1 \cup S_2$$

where S_1 is the set of states of the player one:

$$S_1 = \{(1, \alpha, \beta, \tau) : 0 \leq \alpha \leq 199, 0 \leq \beta \leq 199, 0 \leq \tau \leq 205 - \alpha\}$$

and S_2 is the set of states of the player two:

$$S_2 = \{(2, \alpha, \beta, \tau) : 0 \leq \alpha \leq 199, 0 \leq \beta \leq 199, 0 \leq \tau \leq 205 - \beta\}.$$

- (A) Actions: We have to specify the set of actions per state for each player. Possible actions in this game are *to roll* and *to stop*. We add an extra action *to wait*, which represents that is not the turn of the player. There are some constraints to ensure the transient condition of the stochastic game: in the states $(1, \alpha, \beta, 0)_{\alpha < 200}$ does not make sense for player one to take the action *to stop* because there's nothing to loose. If in our model we permit taking the action *to stop* with 0 points in the turn ($\tau = 0$), is easy to see that there exist a pair of strategies that make the game infinite. The same happen if the action *to roll* is possible when stopping is enough to win. The table 1 shows the set of possible actions per state.
- (P) Payoffs: Because we want to maximise the probability of winning, we define the payoff function in such a way that maximising the probability of winning is equivalent to maximising the expected value V of the payoffs accumulated along the game. The model of transient stochastic game allow us to define a payoff for each pair $(state, action)$ but in this case is enough to define the payoff depending only on the state as follows:

$$r^s = \begin{cases} 1 & \text{if } s = (1, \alpha, \beta, \tau) \text{ with } \alpha + \tau \geq 200 \\ 0 & \text{else} \end{cases}$$

(TP) Transition probabilities: To represent graphically the transition probabilities we use the following representation for the states:

$$(1, \alpha, \beta, \tau) = \textcircled{\alpha_\beta^\tau} \quad (2, \alpha, \beta, \tau) = \boxed{\beta_\alpha^\tau}.$$

The dynamic of the game and the semantic of the states determine transition probabilities between states. In the figure 1, the probability transitions from a state $(1, \alpha, \beta, \tau)$ with $\alpha + \tau < 200$, depending on the player decision are presented.

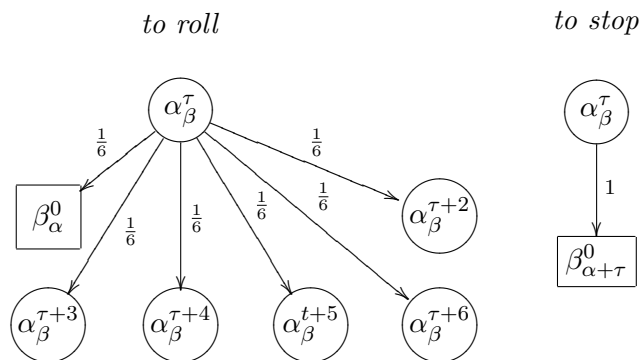


Figure 1: Transition probabilities from a state $(1, \alpha, \beta, \tau)$ with $\alpha + \tau < 200$, depending on the player decision.

Figure 1 shows that, when the decision is *to roll*, the distribution probability on the states is associated with the results of rolling a dice; particularly the probability of losing the turn is $1/6$. In the winning states of player one (i.e. $(1, \alpha, \beta, \tau)$ with $\alpha + \tau \geq 200$) the transition is, with probability one, to the final state (s_f). Transitions for player two are completely symmetric. As is showed in figure 2 in the special states s_i and s_f transitions do not depend on the actions taken by players, indeed they do not have options.




Figure 2: Transition probabilities from the initial and final states.

Now we verify that the stochastic game defined above is transient. We then prove that s_f satisfies conditions (1), (2) and (3) in definition 4. Conditions (1) and (2) are trivially fulfilled, it remains to verify that

$$\sum_{t=0}^{\infty} \mathbf{P}_{s, \pi, \varphi}(S_t \neq s_f) < \infty,$$

for all initial state s , and every pair of strategies π, φ .

Due to the fact that at the beginning of a turn the only option is *to roll*, and that in a state in which the accumulated score is enough to win the player has *to stop*, the game can not continue indefinitely.

For example, if a 6 is rolled consecutively 70 times it is impossible to avoid reaching the final state. Denoting by γ the probability of rolling 70 times a 6, i.e. $\gamma = 1/6^{70} > 0$, it is easy to see that γ is a lower bound for the probability of $S_t = s_f$ for $t \geq 70$. By a similar argument, it follows that, for $n = 0, 1, \dots$

$$\mathbf{P}(S_t \neq s_f) < (1 - \gamma)^n \quad \text{if} \quad 70n \leq t < 70(n + 1).$$

Then

$$\sum_{t=0}^{\infty} \mathbf{P}_{s,\pi,\varphi}(S_t \neq s_f) < 70 \sum_{n=0}^{\infty} (1 - \gamma)^n < \infty$$

and our model is transient.

3.2 Numerical results

In this section the results of the theorem 2.1 are applied to the particular case of the transient stochastic game defined above. Rewriting the definition of application U , defined in equation (2.2), we obtain:

$$U v(s) = \begin{cases} \frac{1}{2}v(1, 0, 0, 0) + \frac{1}{2}v(2, 0, 0, 0) & \text{if } s = s_i \\ v(1, \alpha, \beta, 0)_{roll} & \text{if } s = (1, \alpha, \beta, 0) \\ \max \{v(1, \alpha, \beta, \tau)_{stop}, v(1, \alpha, \beta, \tau)_{roll}\} & \text{if } s = (1, \alpha, \beta, \tau) : \alpha + \tau < 200 \\ 1 & \text{if } s = (1, \alpha, \beta, \tau) : \alpha + \tau \geq 200 \\ v(2, \alpha, \beta, 0)_{roll} & \text{if } s = (2, \alpha, \beta, 0) \\ \min \{v(2, \alpha, \beta, \tau)_{stop}, v(2, \alpha, \beta, \tau)_{roll}\} & \text{if } s = (2, \alpha, \beta, \tau) : \beta + \tau < 200 \\ 0 & \text{in other case} \end{cases}$$

where

$$\begin{aligned} v(1, \alpha, \beta, \tau)_{roll} &= \frac{1}{6}v(2, \alpha, \beta, 0) + \frac{1}{6} \sum_{k=2}^6 v(1, \alpha, \beta, \tau + k) \\ v(1, \alpha, \beta, \tau)_{stop} &= v(2, \alpha + \tau, \beta, 0) \\ v(2, \alpha, \beta, \tau)_{roll} &= \frac{1}{6}v(1, \alpha, \beta, 0) + \frac{1}{6} \sum_{k=2}^6 v(2, \alpha, \beta, \tau + k) \\ v(2, \alpha, \beta, \tau)_{stop} &= v(1, \alpha, \beta + \tau, 0). \end{aligned}$$

Note that in the equations above we have replaced the value of the matrix games in equation (2.2) by a maximum, in the states in which player one has to take the decision, and by a minimum when is player two the one that has to do it. In the states in which both players have only one choice the value of the matrix game is the only entry of the matrix. Since there are no states in which both players have to decide simultaneously, the stationary strategy that emerges from the theorem turns out to be pure, i.e. each player takes an action with probability 1. To determine the complete solution is necessary to specify which action should be taken in about 4 000 000 states. In figure 3 the optimal strategy for player one for some states is shown. The complete solution can be found at




Figure 3: Part of the optimal strategy for player one (for player two is symmetric). It includes states with opponent score $\beta = 0, 150, 180$ & 185 . In the gray zone the optimal action is *to roll* and in the black zone is *to stop*.

www.cmat.edu.uy/cmat/docentes/fabian/documentos/optimalstrategy.pdf.

Some observations about the solution:

- At the beginning of the game, when both players have low scores, we see that the optimal action is *to roll* when $\tau < 20$ and *to stop* in the other case, following the strategy found by Roters [5] maximising the expected value of a turn score. The heuristic interpretation of this fact is: when far away from the target it is optimal to approach it in steps as large as possible.
- When the opponent score β becomes larger the optimal strategy becomes riskier. This can be explained because there are less turns to reach the target.
- For opponent scores greater or equal than 187 ($\beta \geq 187$), the graphic becomes absolutely gray. In other words, when the opponent is close to win, giving him the dice is a bad idea.

- To compare the optimal strategy with the one found by Haigh & Roters [3], we simulate 10 000 000 games. Our simulation showed that in 52% of the games, the winner was the player with the optimal strategy.

4 Two related games

4.1 Reaching exactly the target

It is interesting to explore how the optimal strategy changes when the game is modified. In this section we consider the same dice game, with the only difference being that the condition to win is to reach exactly 200 points. If the sum of accumulated and turn score is greater than 200 the turn finishes without capitalising any point. The formulation of the game is quite similar, the difference appears when the accumulated score plus the turn score is greater than 194, situation in which one roll of the dice can exceed the target. As an example of the mentioned difference, in figure 4 we show the transition probabilities when the accumulated score is 180 and the turn score is 16 ($180 + 16 > 194$). Note that the probability of losing the




Figure 4: Example of transition probability in the variant of the game presented in section 4.1.

turn is the probability of rolling a 1,5,6. In figure 5 part of the optimal strategy for this variant of the game is shown. The complete optimal strategy is available in www.cmat.edu.uy/cmat/docentes/fabian/documentos/optimal exactly.pdf.

Some remarks about the solution:

- As in the classical game, when the target is far, the strategy is similar to “stop if $\tau > 20$ ”.
- Unlike the classical game, the optimal strategy in this case never becomes so risky. This is easy to understand because the probability of winning in any turn is less or equal than $1/6$, even in the case of being very close to the target.




Figure 5: Part of the optimal strategy for player one (for player two is symmetric) for the game presented in subsection 4.1. It includes states with opponent score $\beta = 0, 150, 180$ & 198 . In the gray zone the optimal action is *to roll* and in the black zone is *to stop*.

- When $\alpha + \tau = 194$ there is a “roll zone” larger than usual, because 194 is the largest score in which there is no risk of losing in one roll but it is possible to win rolling a 6.

4.2 Maximising the expected difference

In the second variant of the game, the winner, when reaching the target, obtains from the loser the difference between the target and the loser’s score.

Again, the model of the game is very similar to the classical model, the only difference is the payoff function:

$$r^s = \begin{cases} 200 - \beta, & \text{if } s = (1, \alpha, \beta, \tau) \text{ with } \alpha + \tau \geq 200, \\ \alpha - 200, & \text{if } s = (2, \alpha, \beta, \tau) \text{ with } \beta + \tau \geq 200, \\ 0, & \text{else.} \end{cases}$$

Figure 6 shows the optimal strategy for some opponent scores. The complete optimal strategy can be found at www.cmat.edu.uy/cmat/docentes/fabian/documentos/optimalmaxdif.pdf. The main difference with the optimal strategy in the classical case, is that when one player is close to win (taking into account his current turn score), he takes the risk of rolling, this feature being observed for any score of the opponent.




Figure 6: Part of the optimal strategy for player one (for player two is symmetric) for the variant presented in subsection 4.2. It includes states with opponent score $\beta = 0, 150, 170$ & 180 . In the gray zone the optimal action is *to roll* and in the black zone is *to stop*.

5 Conclusion

In this paper we model a dice game in the framework of Markov competitive decision processes (also known as *stochastic games*) in order to obtain optimal strategies for a player. Our main results are the proof of the existence of a value and an optimal minimax strategy for the game, and the proposal of an algorithm to find this strategy.

We base our results on the theory of transient stochastic games exposed by Filar and Vrieze in [2].

Previous mathematical treatments of this problem include the solution of the optimal stopping problem for a player that wants to maximise the expected number of points in a single turn (see Roters [5]) and the minimisation of the expected number of turns required to reach a target (see Haigh and Roters and [3]). Another previous contribution was done by Neller and Presser [4], who found the optimal strategy in the set of stationary pure strategies, departing from a Bellman equation.

We also provide an algorithm to compute explicitly this optimal strategy (that coincides with the optimal strategy in the larger class of behaviour strategies) and show how this algorithm works in three different variants of the game.

Acknowledgements

This work was partially supported by Antel-Fundaciba agreement “Análisis de Algoritmos de Codificación y Cifrado”

References

- [1] CROCCE, F. (2009). *Juegos estocásticos transitorios y aplicaciones*. Master thesis, PEDECIBA, Montevideo, Uruguay.
- [2] FILAR, J. & VRIEZE, K. (1997). *Competitive Markov Decision Processes*. Springer-Verlang, New York.
- [3] HAIGH, J. & ROTERS, M. (2000). Optimal Strategy in a Dice Game. *Journal of Applied Probability* **37**, 1110–1116.
- [4] NELLER, T. & PRESSER, C. (2004). Optimal Play of the Dice Game Pig. *The UMAP Journal* **25.1**, 25–47.
- [5] ROTERS, M. (1998). Optimal Stopping in a Dice Game. *Journal of Applied Probability* **35**, 229–235.
- [6] SHAPLEY, L.S. (1953). Stochastic games. *Proc. of the Nat. Acad. Sciences* **39**, 1095-1100.