

Optimal Play of the Farkle Dice Game

Matthew Busche¹ and Todd W. Neller²

¹ mtbusche@gmail.com

² Gettysburg College

Department of Computer Science

tneller@gettysburg.edu

Abstract. We present and solve optimality equations for the 2-player, jeopardy dice game Farkle (a.k.a. Dix Mille, Ten Thousand). For fairest play, we recommend 200 compensation points at the beginning of the game for the second player. We then compute the strategy that maximizes expected score, demonstrate a means for replicating such play with mental mathematics, and augment this method so as to enable human Farkle play against which complex optimal play maintains only a small win advantage of $\sim 1.7754\%$.

1 Introduction

Farkle (a.k.a. Dix Mille, Ten Thousand, etc.) is a jeopardy dice game [6, Ch. 6] played with two or more players and six 6-sided dice. As a game with folk origins, Farkle exists in many variations. In this paper, we examine what could be described as the simplest form of Farkle from a scoring perspective, yet with maximum freedom in scoring decisions. We restrict our analysis to the two-player game.

The goal of Farkle is to be the first player to reach a banked score of 10,000 or more points. A player's turn begins by rolling all six dice, and thereafter consists of alternating decisions concerning (1) how to best set aside combinations of dice (if possible) to increase the turn total, and (2) whether or not to continue the turn by re-rolling the dice not yet set aside. If a player cannot set aside a combination of dice, the player is said to "farkle" (or "bust"), scores no points for their turn, but retains prior banked points. However, a player that "holds", i.e. voluntarily ends their turn before farkling, adds their turn total to their banked score.

A combination is a subset of rolled dice that can contribute points to a player's turn total. The combinations are listed in Table 1. Multiple combinations from a single roll may be taken to accumulate points, but only dice from the most recent roll may be used to form a combination, and each die may only be used in a single combination.

If the roll contains one or more combinations, the player must set aside at least one combination and may set aside more, adding the combinations' associated points to their turn total. The player may then either re-roll the remaining dice in an attempt to increase their turn total, or may ("hold" or "bank"), adding their turn total to their banked score and ending their turn. If

all six dice are set aside in combinations, the player may continue their turn re-rolling all six dice. This is called *hot-dice* (or sometimes *turning-the-corner* [4] or a *free-roll* [2]).

Combination	Point Value
one 1	100
one 5	50
three 1s	1000
three 2s	200
three 3s	300
three 4s	400
three 5s	500
three 6s	600

Table 1. Dice combinations and their point value.

In Section 2, we present optimality equations for two-player Farkle and describe the method used to solve them. We next interpret optimal play (Section 3) and compare it to play that maximizes expected score gain (Section 4). A human-playable modification to maximum scoring play is shown to provide a good approximation to optimal play in Section 5, and we conclude with a summarization and possibilities for future work.

2 Optimality Equations and Solution Method

Farkle’s game tree consists of chance nodes where 1-6 dice are rolled and maximizing choice nodes of two kinds: banking decision states concerning whether to roll or hold (i.e. bank), and scoring decision states concerning how to set aside dice combinations for points. Let banking decision states be denoted as the 4-tuple $\langle b, d, n, t \rangle$ where b is the banked score of the current player, d is the banked score of the opponent, n is the number of dice not set aside in combinations, and t is the turn total. Let scoring decision states be denoted as the 5-tuple $\langle b, d, n, t, r \rangle$, where in addition to previous variables we add a roll sequence $r \in \{1, \dots, 6\}^n$. Let the set of all roll sequences of n dice be denoted R_n . Let the probability of winning from a banking decision state be denoted $W(b, d, n, t)$ and the probability of winning from a scoring decision state be denoted $W(b, d, n, t, r)$.

The probability of winning from a banking decision state is

$$W(b, d, n, t) = \begin{cases} 1 & \text{if } b + t \geq 10,000 \\ \sum_{r \in R_n} \frac{1}{6^n} W(b, d, n, t, r) & \text{if } t = 0, \text{ and} \\ \max \left(\begin{array}{l} 1 - W(d, b + t, 6, 0), \\ \sum_{r \in R_n} \frac{1}{6^n} W(b, d, n, t, r) \end{array} \right) & \text{otherwise.} \end{cases} \quad (1)$$

In the first case, a rational player will bank a turn total t that is sufficient to win the game, so we treat the state as a terminal victory state with win probability 1.

In the second nonterminal case where the turn total t is zero, we are necessarily at the start of a turn where rolling is mandatory, so we average over the win probabilities associated with the result of each possible roll r performed in the current state.

In the third nonterminal case where the turn total t is nonzero, the player has a choice between holding (i.e. banking) and rolling, and rationally chooses the action that maximizes the player's win probability. If the player holds, the player banks the turn total, accumulating it to the player's banked score and it becomes the opponent's turn. Since there are no draws in Farkle, the win probability after holding is one minus the probability that the opponent wins at the beginning of their turn with the player having new banked score $b + t$. The opponent begins their turn with 6 dice and a turn total of 0, so this holding win probability is expressed as $1 - W(d, b + t, 6, 0)$. However, if the player rolls, then the win probability is the same as expressed in case 2 where the roll action is required.

We next turn our attention to the problem of writing an equation for the probability of winning from a scoring decision state. Define a *combination* c as a length two vector $c = (c_N, c_P)$ where c_N is the number of dice used in c and c_P is its point value. The set of available combinations C are defined by the dice scoring rules of the game shown in Table 1. Define a *scoring* $s = (s_N, s_P)$ as a sum of an arbitrary set of combinations (with repetition). Define S_r as the set of all possible scorings that can be formed from combinations from roll r . As an example, for roll $r = (4, 5, 3, 4, 4, 5)$

$$S_r = \{(1, 50), (2, 100), (3, 400), (4, 450), (5, 500)\}.$$

Define a *hot-dice* function $h(n)$ for resetting the number of available dice back to 6 when a scoring successfully uses all dice from a roll

$$h(n) = \begin{cases} 6, & \text{for } n = 0. \\ n, & \text{otherwise.} \end{cases}$$

Then, the probability of winning from a scoring decision state is

$$W(b, d, n, t, r) = \begin{cases} 1 - W(d, b, 6, 0) & \text{if } S_r = \emptyset, \text{ and} \\ \max_{s \in S_r} (W(b, d, h(n - s_N), t + s_P)) & \text{otherwise.} \end{cases} \quad (2)$$

In the first case, there is no possible scoring, so the result is the same as holding in the equation before, yet without a change to the player’s banked score. In the second case, the player chooses the scoring s that leads to a resulting number of dice and turn total that maximizes the player’s win probability. Note that this does not necessarily imply picking a scoring with a maximum s_p . Both the number of dice and the turn total matter for optimal play.

Substituting (2) into (1) eliminates $W(b, d, n, t, r)$, reducing the unknowns to a number that can easily be modeled by computer and solved by a generalization of value iteration as in [8]. We began with arbitrary probability guesses for all unknown win probabilities. Iteratively, we reevaluated each equation right-hand-side given our current estimates and let the result be our new estimate for the win probability of the equation left-hand-side. When estimates converged such that the maximum estimate change magnitude was less than 1×10^{-14} , we terminated value iteration.

3 Overview of Optimal Strategy

The initial win probabilities for the first and second players in an optimal game is ~ 0.536953 and ~ 0.463047 , respectively, so the first-player has a win probability advantage of ~ 0.073906 . First-player advantages are common in jeopardy race games [9], so we were interested to see what the fairest komi (i.e. compensation points) would be for player 2 to receive at the beginning of the game. If player 1 and 2 begin the game with 0 and 200 points, the fairest komi, then the player 1 win probability is reduced to ~ 0.504002 , and the advantage is thus reduced to ~ 0.008004 in player 1’s favor, or about 8 more won games than player 2 per 1000 games.

The optimal roll states are visualized in Figure 1. Each graph depicts the set of states in which an optimal player should roll and seek to improve their turn total. Thus, the current player with b points facing an opponent with d points begins their turn with turn total $t = 0$ and proceeds to roll and set aside combinations, until either (1) the player farkles, (2) t has increased to the point where the state has passed upward beyond the optimal roll states and the player thus holds, or (3) t has increased to the point where $b + t$ has reached the goal score so the player holds and wins.

In these graphs, we can appreciate how significantly roll/hold behavior varies according to the number of dice available for re-rolling. Whereas a player will be very conservative and hold at low turn totals with few dice, a farkle event has much lesser probability when many dice are rolled, giving the player a much greater willingness to continue their turn with much higher turn totals.

One might wonder how such high turn total states would be reachable for non-extreme scores where holding is optimal for much lower turn totals as the number of dice decrease. The answer is that a player can only reach such states through a succession of hot-dice rolls where the player is never without sufficient dice to merit continuation with the current turn total.

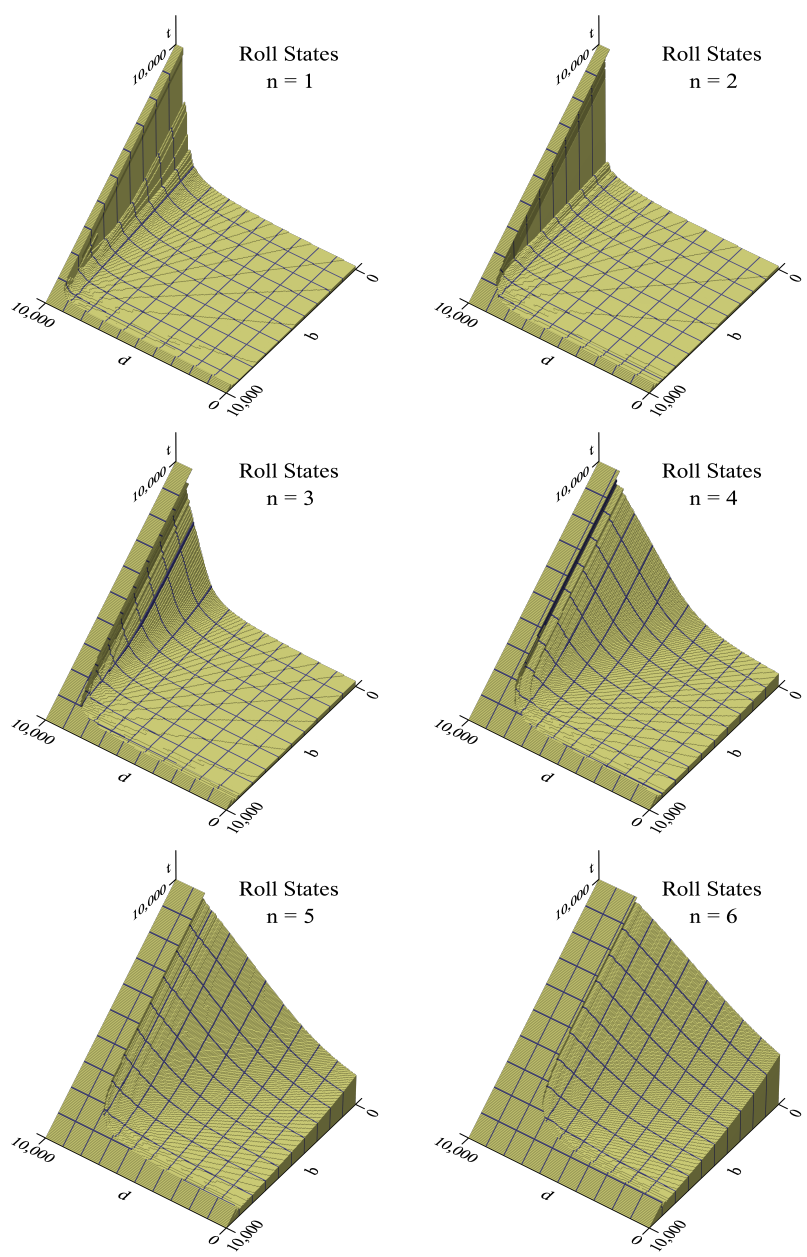


Fig. 1. Game states from which you should roll.

The next features of Figure 1’s graphs that we would note are the L-shaped “go for it” regions in the (b, d) -plane. These regions show where at least one player is close enough to win the game such that a player would risk any turn total in order to win the game in that turn. We note that the “go for it” margin thickness decreases with a decrease in the number of dice to roll. In contrast to the analogous “go for it” regions in the jeopardy race Pig game family [9], these are more angular and L-shaped than curved in the region where both players advance near to the goal score with similar scores.

Finally, we observe that play varies considerably according to player score differences when one has 4 or more dice, whereas much of the roll/hold boundary could be fitted with a shallow-sloped plane for 3 or fewer dice. Put simply, aggressive/conservative optimal play according to a player’s score disadvantage/advantage, respectively, is mainly seen in roll/hold decisions with many dice. With few dice, the high probability of loss with continuation gives little reward for the “push your luck” dynamic of such games.

Note that these graphs do not give us insight into the complex nuances of optimal scoring. In many cases, an optimal player does not set aside the maximum combination that a roll permits. It is not simply how much a combination adds to the turn total; the win probability of the resulting state is what matters.

4 Maximum Scoring Strategy

In this section, we discuss the maximum *scoring* strategy where the objective is to maximize the expected number of points to be added to one’s banked score, and contrast this with the optimal, or maximum *winning*, strategy. Note that the maximum scoring strategy has no dependence on either b or d . To find this strategy, we can follow the same approach as used to find the optimal strategy. Define $T(n, t)$ to be the expected number of points to be added to your banked score (at the end of your current turn) from banking decision state (n, t) ; and define $T(n, t, r)$ to be the expected number of points to be added to your banked score from scoring decision state (n, t, r) . The expression for expected change to banked score at a banking decision state is identical in form to (1):

$$T(n, t) = \begin{cases} \sum_{r \in R_n} \frac{1}{6^n} T(n, t, r) & \text{if } t = 0, \text{ and} \\ \max \left(\begin{array}{c} t, \\ \sum_{r \in R_n} \frac{1}{6^n} T(n, t, r) \end{array} \right) & \text{otherwise.} \end{cases} \quad (3)$$

Similarly, the expression for the expected change to one’s banked score at a scoring decision state is identical in form to (2):

$$T(n, t, r) = \begin{cases} 0 & \text{if } S_r = \emptyset, \text{ and} \\ \max_{s \in S_r} (T(h(n - s_N), t + s_P)) & \text{otherwise.} \end{cases} \quad (4)$$

Substituting (4) into (3) eliminates $T(n, t, r)$. For sufficiently large t (where the best bank/roll decision is always to bank), $T(n, t) = t$. Other $T(n, t)$ values

may be found by iterating down to successively smaller t values, and simply evaluating (3) as you go.

The resulting strategy yields an average of ~ 446.57144 points per turn. $\sim 20.5964\%$ of turns end with a farkle. Before making a comparison with the optimal strategy, we first modified the maximum scoring strategy by forcing selection of a scoring option that achieves a turn total sufficient to bank and win the game. Under these conditions, the maximum scoring player wins $\sim 51.3812\%$ of games as first player, $\sim 43.8470\%$ of games as second player, and $\sim 47.6141\%$ of games overall. The optimal player thus has a $\sim 4.7718\%$ win advantage overall.

We have devised a means by which a human can play the maximum scoring strategy perfectly by memorizing a small table of integers. To see how, let us first define the value of continuing the turn in state (n, t) to be $V(n, t) = T(n, t) - t$. V is the expected future turn total increase for the rest of the turn. When V is zero, there is no value in continued rolling and the maximum scoring strategy will bank. Table 2 shows the $V(n, t)$ for small t . To make scoring decisions, one considers all states (n, t) achievable with the current roll, and chooses the state that maximizes $t + V(n, t)$.

t	6	5	4	3	2	1
0	446.571	*	*	*	*	*
50	*	291.561	*	*	*	*
100	*	278.777	162.486	*	*	*
150	*	*	147.597	66.904	*	*
200	*	*	134.168	51.681	4.551	*
250	*	*	*	37.488	0.000	0.000
300	397.543	*	*	23.321	0.000	0.000
350	390.959	227.676	*	*	0.000	0.000
400	384.381	219.761	90.767	0.000	0.000	0.000
450	377.983	211.854	82.745	0.000	0.000	0.000
500	372.298	203.954	74.730	0.000	0.000	0.000

Table 2. Table of $V(n, t)$ values showing the expected future increase in turn total. States marked with an asterisk are unreachable.

Clearly, memorizing $V(n, t)$ is impractical. However, one can memorize Table 3 and use it to mentally calculate an estimate of $V(n, t)$ as follows: for each scoring option, determine your prospective new state (n, t) . Beginning with the row corresponding to n , find the leftmost column of that row with a value $\leq t$, and estimate $V(n, t)$ as the top value label for that column. Any scoring option that uses all remaining dice (yielding $n = 6$) is always superior to any other option. Furthermore, $V(6, t) > 0$ for all $t < 10,000$, so one never banks if six dice can instead be rolled. Thus the $n = 6$ row is omitted from the table.

As an example, assume you have a turn total of 2000 and make the six-die roll (3, 6, 6, 1, 1, 4). You have two scoring options: take one 1 for 100 points, or two 1s for 200 points. Using the table, we find the corresponding estimates $V(5, 2100) = 100$, and $V(4, 2200) = 0$. In both cases the quantity $t + V(n, t) =$

n	0	50	100	150	200	250	300
5	2900	2250	1600	950	550	250	0
4	1000	700	350	150	0		
3	400	250	0				
2	250	0					
1	0						

Table 3. Table for approximating $V(n, t)$. Beginning with the row corresponding to n , find the leftmost column of that row with a value $\leq t$, and estimate $V(n, t)$ as the top value label for that column.

2200. Ties are broken by with the following rule: *in the case of ties, choose an option that leaves you in a banking state; otherwise if no such option is available among those that tie, choose the option that leaves you with more dice to roll.* So in this case, we take both 1s for 200 points to move to a banking state and end the turn.

A player who memorizes this table of 13 nonzero values (with 11 distinct values), and follows these rules will perfectly reproduce the strategy that maximizes expected score.

5 Approximately Optimal Play

In this section, we modify the strategy that maximizes expected scores to achieve a significantly better approximation to optimal play.

From Figure 1, it is clear that for any particular n value, there exist thresholds for both b and d above which one should never bank short of winning the game. Accordingly we modified the banking rule for the maximum score strategy such that if $b \geq B_n$ or $d \geq D_n$ then do not bank. The optimal values for B_n and D_n are shown in Table 4. Using these thresholds to establish the go-for-it region of the game, the optimal opponent's win probability advantage is reduced to $\sim 1.7754\%$.

n	B_n	D_n
6	*	*
5	*	7900
4	8950	8600
3	9350	9350
2	9550	9550
1	9600	9500

Table 4. Table showing thresholds for the player's banked score, B_n , and the opponent's banked score, D_n , at which one should not bank. Entries marked with an asterisk require no change to the banking policy.

This table presents an additional 9 integers to memorize, yet there are only 7 distinct integer values as we note that $B_n = D_n$ for $n = 2, 3$. Thus, a player memorizing the 18 distinct values of Tables 3 and 4 can play a good approximation of optimal play, winning $\sim 49.1124\%$ of games against an optimal player where the first player is chosen at random.

6 Conclusions

In this paper, we have presented the fundamental rule set for Farkle (a.k.a. Dix Mille, Ten Thousand), expressed and solved optimality equations for the 2-player game for the first time. We noted that, between optimal players, the first-player win advantage is $\sim 7.3906\%$. Having the second player begin the game with 200 points is the fairest komi (i.e. compensation points), reducing the first-player win advantage to $\sim 0.8004\%$.

We then contrasted optimal strategy with strategy that maximizes a player's expected score. Optimal strategy has a win advantage of $\sim 4.7718\%$ over maximum scoring strategy when the first player is chosen at random. A mental mathematical means of replicating maximum scoring strategy was developed and demonstrated.

Finally, we presented an improvement to this mental mathematical strategy that recognized "go for it" regions of the state space where one player is close enough to winning that a player should be willing to risk any turn total in an attempt to win that turn. This reduced the optimal strategy win advantage to $\sim 1.7754\%$ when the first player is chosen at random. That the memorization of 18 unique integer constants plus simple mental mathematics allows such good performance against the subtleties of the optimal strategy visualized in Figure 1 is a pleasing result.

7 Future Work

Although Farkle appears to be a relatively new folk dice game, rule variations are already abundant. We have here chosen the simplest rule-set found in the earliest descriptions found in our research. However, more recent rule-sets have simplified the game in an undesirable respect by requiring a player to set aside *all* possible dice combinations. While we are striving for a simplest, fundamental form of the game, we reject such trivialization of game decision. That said, there are many different variations of rules in the Farkle/Ten Thousand dice game family, including:

- Variation in the goal score (e.g. 5,000 points [7,6,5]),
- A minimum turn total required for a player's first bank (e.g. 500 points [4], 1000 points [7]),
- A minimum turn total required to bank each turn (e.g. 300 points [2], 350 points [4,6]),

- A minimum turn total that is greater than the opponent’s scored turn total of the previous turn, or some minimum (e.g. 350 points [6]) otherwise,
- Additional points awarded for four-, five-, or six-of-a-kind [4,1],
- Points awarded for a six-die roll of 1-2-3-4-5-6 (e.g. 1000 [7], 1500 points [4,5,1], 3000 points [3,6]),
- Points awarded for a six-die roll of three pairs (e.g. 750 points [4], 1500 points [3,4,6,1], 2000 points [4]),
- Points awarded for two-triplets (e.g. 2500 points [1], 3000 points [4]),
- Points awarded for a four-of-a-kind and a pair (e.g. 1500 [1]),
- Cumulative bonus points awarded for each single-die hot-dice roll within a single turn (e.g. 500 for the first, 1000 for the second, etc. [3]),
- Variation in endgame rules either by ensuring each player gets an equal number of turns [3,4], or by giving all other players a final turn once a player reaches the goal [1],
- A penalty to banked score for multiple consecutive farkle turns (e.g. 3 farkles → -500 points [2], 3 farkles → -1000 points [7]),
- The possibility of an opponent stealing a player’s turn total after a farkle roll [6],
- A disaster roll consisting of four or more 2s resulting in the loss of all banked points [3,6],
- An immediate win for rolling six-of-a-kind [4],
- Points awarded for rolls that would normally be treated as a six-die farkle (e.g. 500 points [2]),
- Scoring of the turn total when farkling [6], and
- The requirement to set aside all scoring combinations [7].

Generally speaking, good game design consists not in the elaboration of rules, but in the quality of strategic considerations for a given rule-set. One would like to find the simplest rules for which decisions are both interesting and comprehensible. We are therefore curious whether or not any of the variations substantively add to the interest of Farkle game play.

We thus see potential future interesting work in the analysis of such variations and the comparison of optimal play policies to observe the impacts of such game design decisions.

References

1. Patch Products Farkle Rules. http://www.boardgamecapital.com/game_rules/farkle.pdf (2007), [Online; accessed 11-March-2017]
2. Zilch. <http://zilch.playr.co.uk/> (2008), [Online; accessed 11-March-2017]
3. Arnold, P. (ed.): The Book of Games. Exeter Books, New York (1985)
4. Jacobs, G.: The World’s Best Dice Games. John N. Hanson Co., Inc. (1993)
5. Keller, M.: Ten Thousand. <https://web.archive.org/web/20020320064638/http://members.aol.com/dicetalk/rules/10000.txt>, [Online; accessed 11-March-2017]
6. Knizia, R.: Dice Games Properly Explained. Elliot Right-Way Books, Brighton Road, Lower Kingswood, Tadworth, Surrey, KT20 6TD U.K. (1999)

7. Mohr, M.S.: *The New Games Treasury*. Houghton Mifflin, New York (1997)
8. Neller, T.W., Presser, C.G.: Optimal play of the dice game pig. *The UMAP Journal* 25(1), 25–47 (2004)
9. Neller, T.W., Presser, C.G.: Pigtail: A pig addendum. *The UMAP Journal* 26(4), 443–458 (2005)