Optimal Play of the Farkle Dice Game

Matthew Busche
Todd W. Neller, Gettysburg College

Outline

- Farkle Introduction
- 2-Player Optimality Equations and Solution
- Optimal Play and Fair Komi
- Human Score Maximization Play
- Improvements For Excellent Human Approximation of Optimal Play
- Future Work

Farkle

- Farkle (a.k.a. Dix Mille, Ten Thousand, etc.) is a *jeopardy dice game* for 2+ players and six 6-sided dice.
  - Jeopardy dice game – Dice game where the dominant type of decision is whether or not to jeopardize previous gains by rolling for potential greater gains
- Relatively modern with many, many rule variants.
- We here focus on the simplest core ruleset according to (Knizia, 1999).

Farkle Rules

- Goal: be the first player to reach a banked score of 10,000 or more points.
- Player turn: Roll all 6 dice.
  1. Cannot set aside dice *combination(s)*? → Bust, i.e. end turn with no banked score gain.
  2. Set aside combination(s), and either
     - Reroll remaining dice (go to 1), jeopardizing combination scores, or
     - Hold and bank combination score(s).
- If all 6 dice are set aside, the player may reroll all 6 dice, accumulating further combinations.

Farkle Combinations

- Example: roll of 252325 → three possible combinations: two single 5s and three 2s for 50+50+200 = 300 points.
- Can choose not to set aside all combinations, e.g. can set aside only three 2s in above example.
- Must set aside at least one combination if possible.
- Set-aside combinations must be disjoint; cannot count a die for more than one combination.
- We refer to the total of the set-aside point values as the *turn total*.

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

Farkle Decisions

- An optimal player makes two types of decisions in order to maximize their probability of winning:
  - Which combinations to set aside, and
  - Whether or not to hold or reroll.
- There are thus four variables relevant to Farkle decision-making:
  - b = **b**anked points of the current player
  - d = banke**d** points of the opponent
  - n = **n**umber of dice not set aside in combinations
  - t = **t**urn total

2-Player Optimality Equations

Probability of winning from a **banking decision state**:

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

where R_n is the set of all rolls of n dice.

2-Player Optimality Equations (cont.)

Probability of winning from a **scoring decision state**:

- Define a *combination* $c = (c_N, c_P)$ where c_N is the number of dice set aside, and c_P is their point value.
- Define a *scoring* $s = (s_N, s_P)$ as a sum of a set of combinations (with possible repetition).
- Define S_r as the set of all possible scorings for roll r.
- Define *hot-dice function* $h(n)$ for resetting the number of dice to 6 when all have been set aside:

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

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

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

Optimal Play

- Optimal play is computed using value iteration as with jeopardy dice game Pig (Neller, Presser 2005).
- Initial win probabilities for 1st player with optimal play is ~0.536953 → win probability advantage of ~0.073906.
- Fairest komi (i.e. compensation points) would be for player 2 to initially have 200 points → reduces 1st player win probability to ~0.504002, and advantage to ~0.008004

Optimal Roll States (n=1)

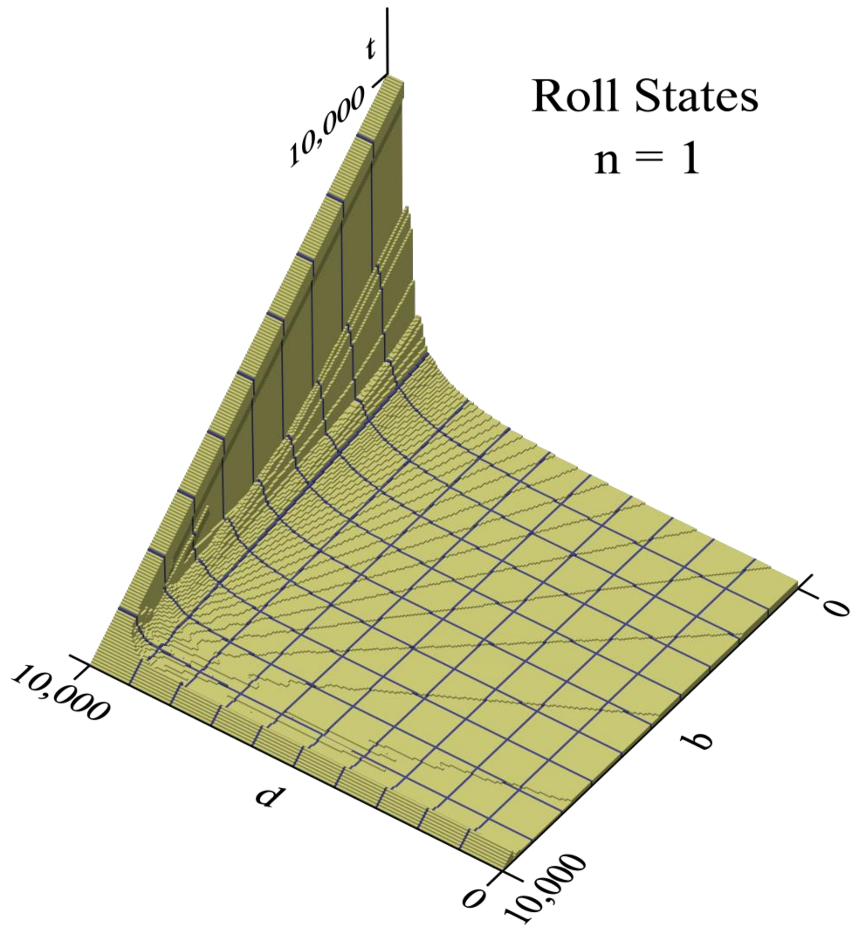

Optimal Roll States (n=2)

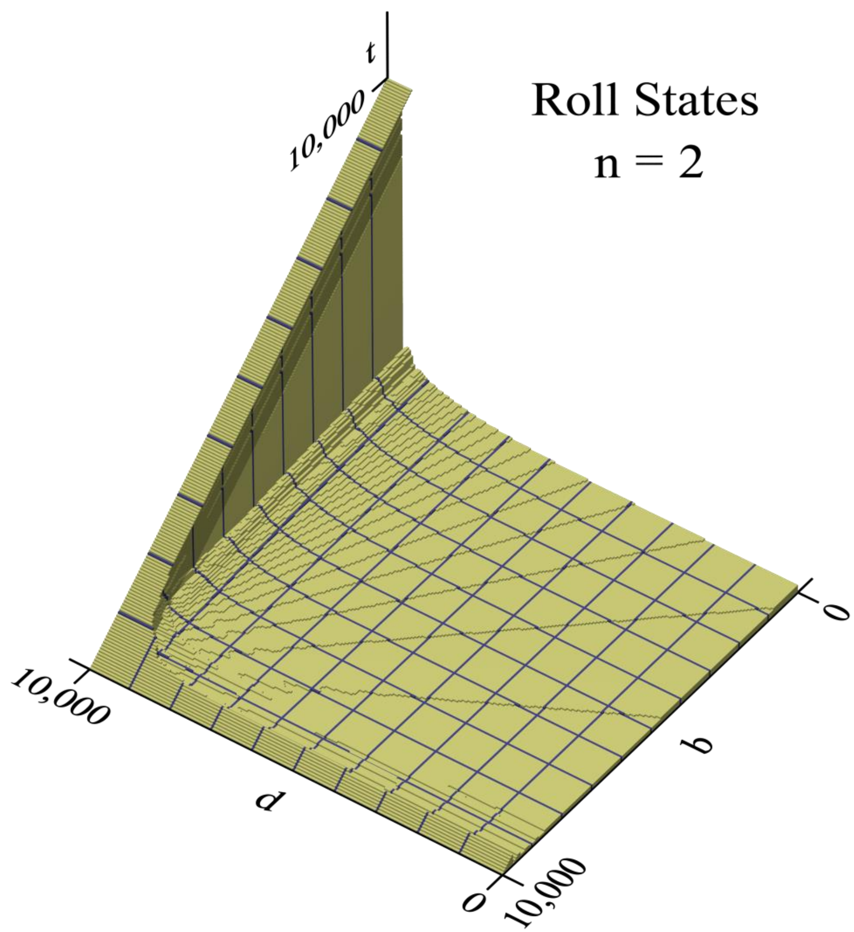

Optimal Roll States (n=3)

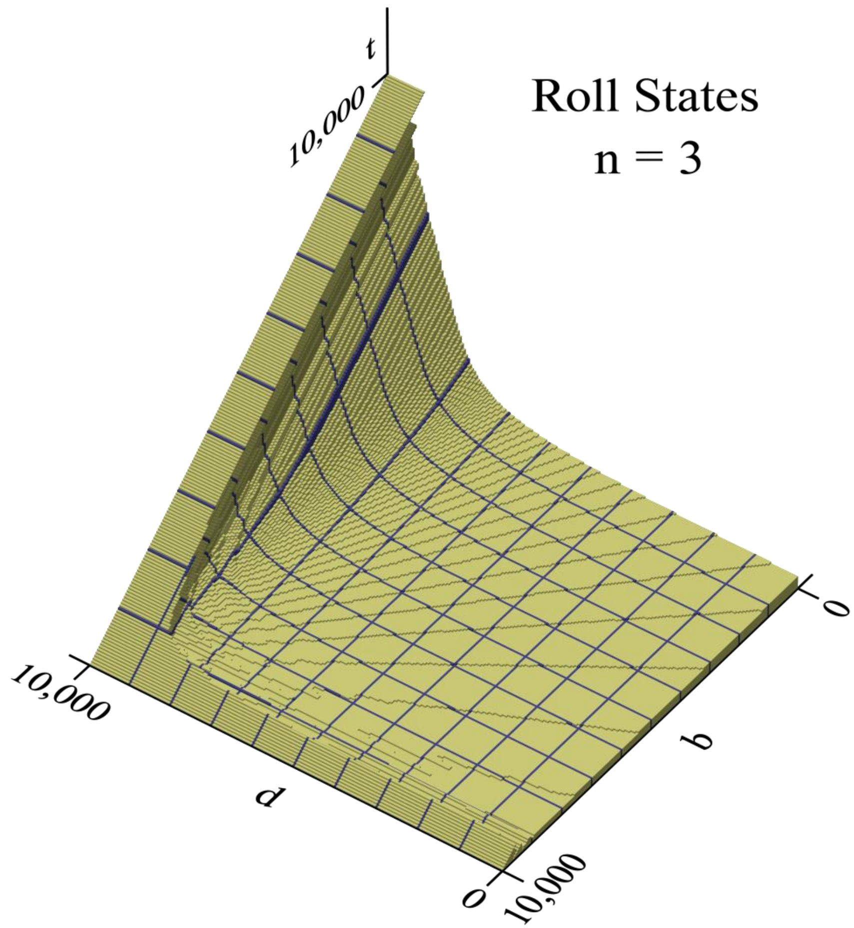

Optimal Roll States (n=4)

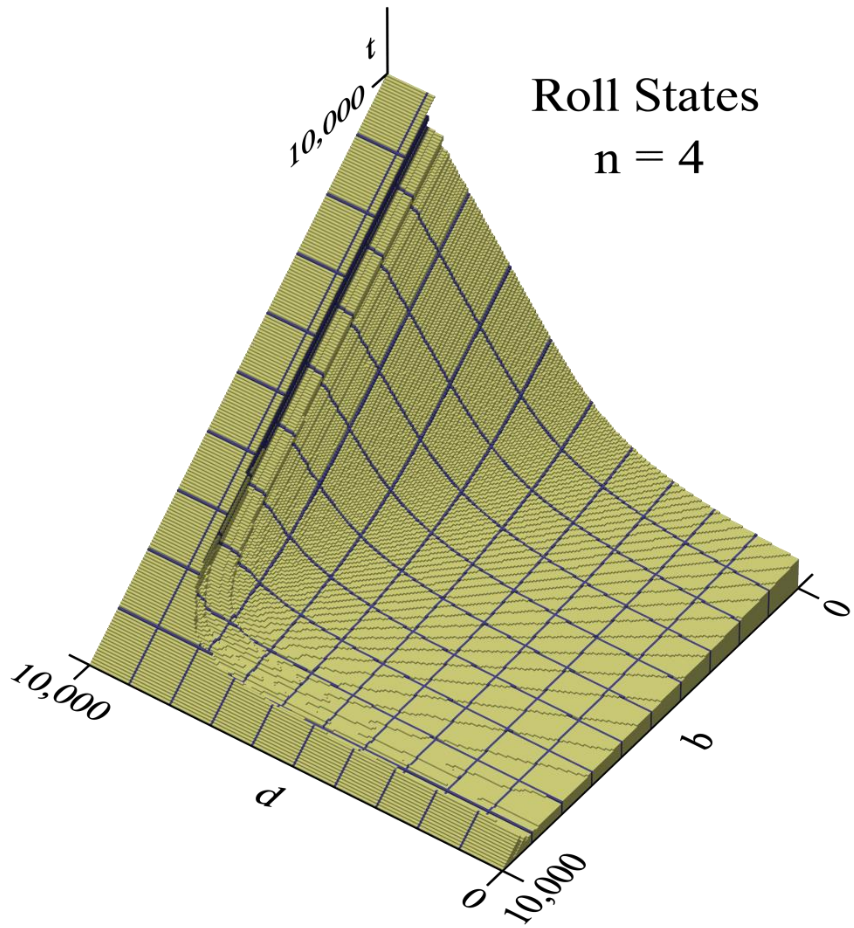

Optimal Roll States (n=5)

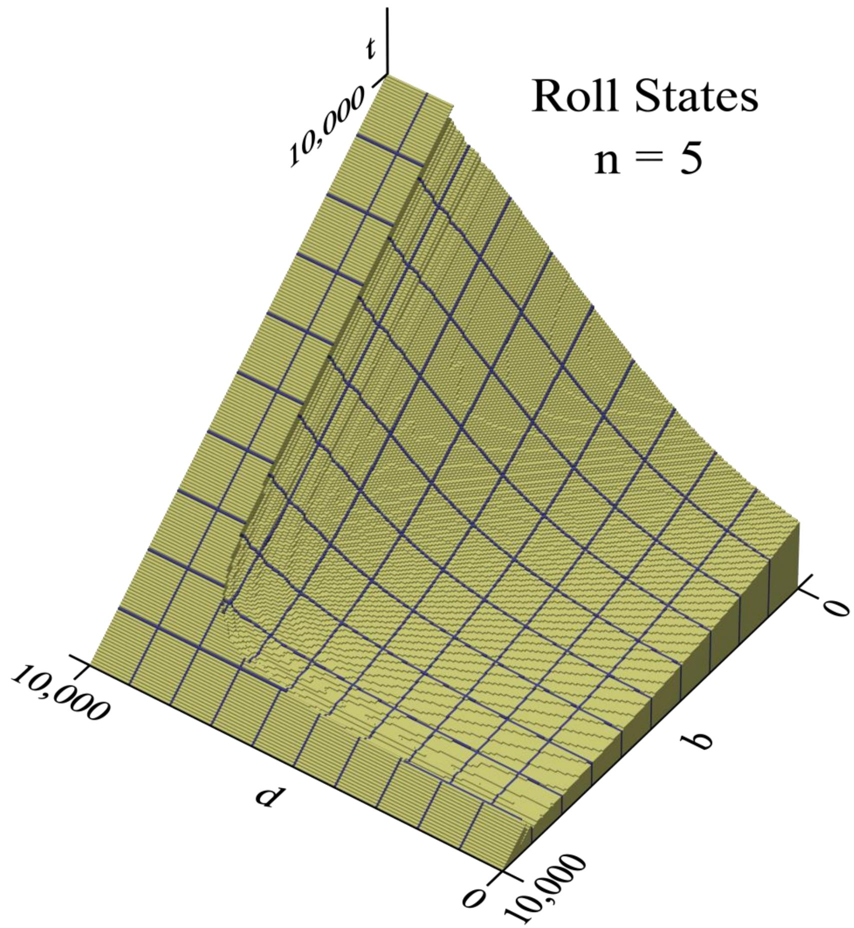

Optimal Roll States (n=6)

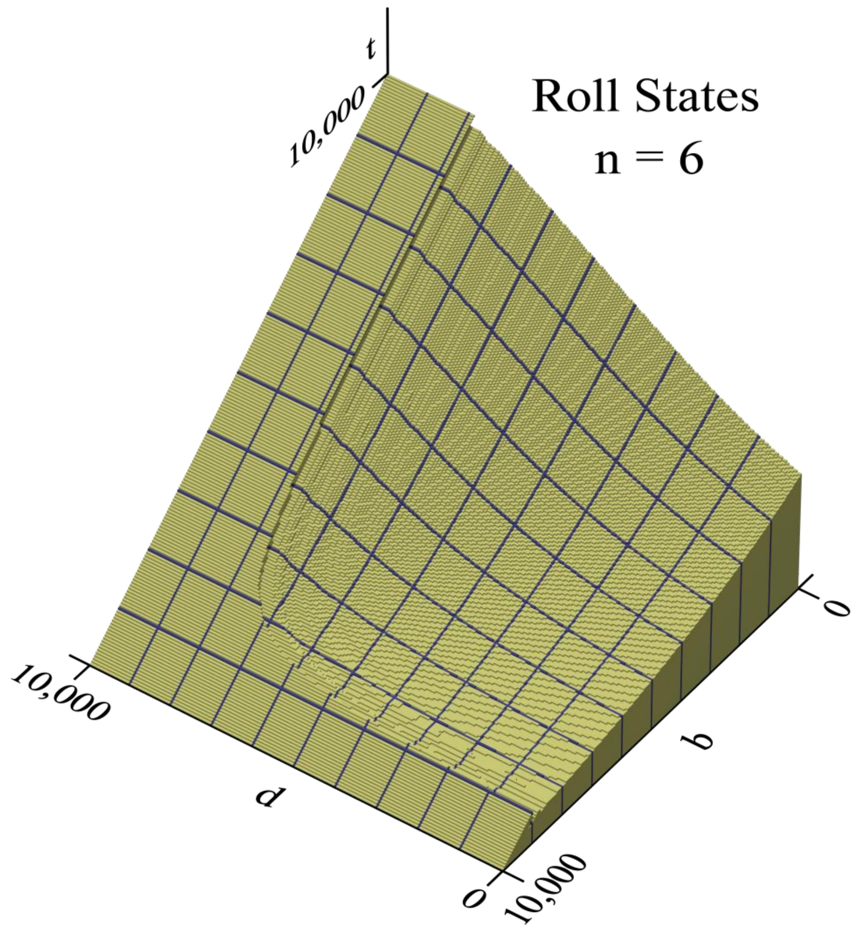

Optimal Scoring Decisions

- The previous graphs do nothing to help us understand the nuances of scoring decisions in optimal play.
- There are many cases where an optimal player doesn't set aside the maximal roll scoring.
- What matters is the win probability of the state to which scoring leads.
- That said, we wondered how well a policy that maximizes expected turn score would perform…

Score Maximization Equations

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

where

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

Score Maximization Performance

- Averages ~446.57144 points per turn
- ~20.5964% of turns end with a farkle (bust).
- A player that maximizes expected score per turn (holding when the turn total is sufficient to win) wins ~47.6141% of games against an optimal player.
- The optimal player thus has a ~4.7718% win probability advantage over score maximization.

Score Maximization Play

- Let V(n, t) = T(n, t) – t be the expected future turn total increase for the rest of the turn.
- Consider each scoring (n, t) achievable with the current roll, and choose the scoring that maximizes t + V(n, 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

Human Score Maximization Play

- While a human can't memorize the V(*n*, *t*) table, we present a smaller table below that allows human score maximization play.
  - For row *n*, find the leftmost column entry ≤ *t*.
  - The top column value approximates V(*n*, *t*) well enough to yield equivalent score maximization play.
  - Breaking ties: If possible, prefer an option that leaves you in a banking state; otherwise, prefer the option with greater dice to roll *n*.

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						

Human Score Maximization Play Example

- Example:
  - Initial roll of 5 5 2 2 2 6
    - Score one 5: t = 50, n = 5 → t + V(n, t) ~= 50 + 300 = **350**
    - Score two 5s: t = 100, n = 4 → t + V(n, t) ~= 100 + 200 = 300
    - Score three 2s: t = 200, n = 3 → t + V(n, t) ~= 200 + 100 = 300
    - Score one 5, three 2s: t = 250, n = 2 → t + V(n, t) ~= 250 + 0 = 250
    - Score two 5s, three 2s: t = 300, n = 1 → t + V(n, t) ~= 300 + 0 = 300
  - Best to score one 5 and reroll remaining 5 dice for greater exp. score.
  - Note: Prefer "two 5s, three 2s" (bank) to "two 5s" (>n) to "three 2s".

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						

Play Improvements

- Human score maximization play can be improved by approximating “go for it” regions of the optimal roll states.
- Augment human score maximization play s.t. if $b \geq B_n$ or $d \geq D_n$ then do not bank:
- Reduces optimal win probability advantage to ~1.7754%
- Memorization of 18 distinct values

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

Future Work

- We have analyzed the core form of Farkle with greatest rule consensus, but many variants exist, e.g.:
  - Different goal scores (e.g. 5000),
  - Minimum turn totals to bank first/every turn,
  - Additional combinations (e.g. 4/5/6-of-a-kind, 1-2-3-4-5-6, 3 pairs, 2 triplets),
  - Requirement to set aside all combinations,
  - etc.
- More complex rules do not imply a deeper game, but do any of the variants create more interesting decisions that merit increased rule complexity? Interesting game design issues abound!

Conclusion

- Optimal play of 2-player Farkle
  - Fair komi of 200 points for the 2nd player
  - Has ~4.7718% win probability advantage over score maximization play
- Our best human-playable policy
  - Memorization of two tables with 18 distinct values permits score maximization play augmented with "go-for-it" states
  - Reduces optimal play win probability advantage to ~1.7754%