Pigtail: A Pig Addendum

Todd W. Neller
Clifton G.M. Presser
Department of Computer Science
300 N. Washington St.
Campus Box 402
Gettysburg College
Gettysburg, PA 17325–1486
{tneller,cpresser}@gettysburg.edu

Introduction to Pig

The object of the jeopardy dice game Pig is to be the first player to reach 100 points. Each turn, a player repeatedly rolls a die until either a 1 is rolled or the player holds and scores the sum of the rolls (i.e., the *turn total*). At any time during a player's turn, the player is faced with two choices: *roll* or *hold*. If the player rolls a 1, the player scores nothing and it becomes the opponent's turn. If the player rolls a number other than 1, the number is added to the player's turn total and the player's turn continues. If the player instead chooses to hold, the turn total is added to the player's score and it becomes the opponent's turn.

In our original article [Neller and Presser 2004], we described a means to compute optimal play for Pig. Since that time, we have also solved a number of Pig variants. In this addendum, we review the optimality equations for Pig, show how these equations change for several Pig variants, and show how the resulting optimal policies change accordingly.

Maximizing the Probability of Winning

Let $P_{i,j,k}$ be the player's probability of winning if the player's score is i, the opponent's score is j, and the player's turn total is k. In the case where $i + k \geq 100$, $P_{i,j,k} = 1$ because the player can simply hold and win. In the general case where $0 \leq i, j < 100$ and $k < 100 - i$, the probability of an optimal player winning is

$$P_{i,j,k} = \max\left(P_{i,j,k,\mathrm{roll}}, P_{i,j,k,\mathrm{hold}}\right),$$

The UMAP Journal 26 (4) (2005) 443–458.

where $P_{i,j,k,\text{roll}}$ and $P_{i,j,k,\text{hold}}$ are the probabilities of winning if one rolls and holds, respectively. These probabilities are given by:

$$P_{i,j,k,\text{roll}} = \frac{1}{6}\left[(1 - P_{j,i,0}) + P_{i,j,k+2} + P_{i,j,k+3} + P_{i,j,k+4} + P_{i,j,k+5} + P_{i,j,k+6}\right],$$
$$P_{i,j,k,\text{hold}} = 1 - P_{j,i+k,0}.$$

The probability of winning after rolling a 1 or holding is the probability that the other player will not win beginning with the next turn. All other outcomes are positive and dependent on the probabilities of winning with higher turn totals.

These equations, and those of the variants that follow, can all be solved using value iteration as described by Neller and Presser [2004]. The solution to Pig is visualized in **Figure 1**. The axes are i (player 1 score), j (player 2 score), and k (the turn total). The surface shown is the boundary between states where player 1 should roll (inside the surface) and states where player 1 should hold (outside the surface).

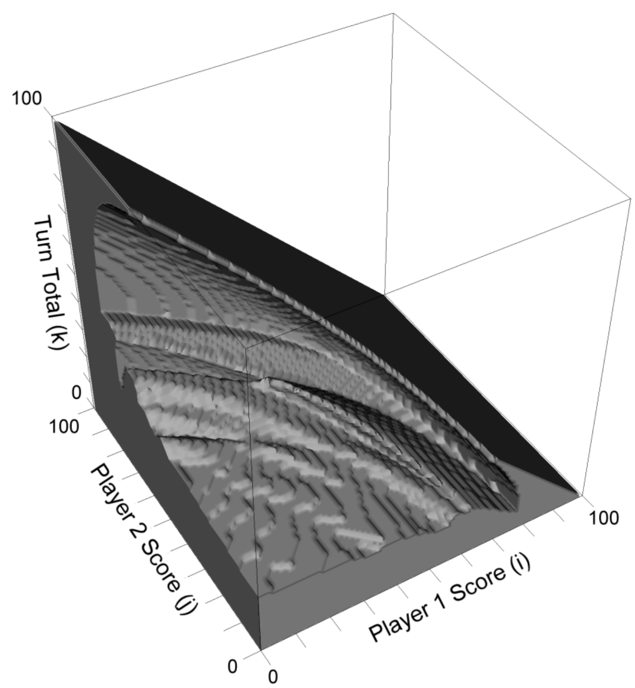

Figure 1. Roll / hold boundary for optimal Pig policy.

In each of the following subsections (except Hog), the state space is the same, and the same notation will be used to underscore similarities.

Parker Brothers Pig Dice

Pig Dice®[1] is a commercial 2-dice variant of Pig that has had surprisingly little influence on the rules of modern variants. This is in part due to the fact that the game requires specialized dice. One die is a standard die with a pig head replacing the 1. The other is a standard die with a pig tail replacing the 6. Such rolls are called Heads and Tails, respectively. The goal score is 100, yet after a player has met or exceeded 100, all other players have one more turn to achieve a higher score. We ignore this restriction for symmetry in analysis.

As before, players may hold or roll, risking accumulated turn totals. However, there is no undesirable single die value in this game. Rather, rolling dice that *total 7* ends the turn without scoring. Rolling a Head and a Tail doubles the current turn total. Rolling just a Head causes the value of the other die to be doubled. Rolling just a Tail causes the value of the other die to be negated. However, the turn total can never be negative. If a negated die would cause a negative turn total, the turn total is set to 0. All other non-7 roll totals are added to the turn total.

Optimality equations for Pig Dice are the same as those for Pig with the following difference:

$$P_{i,j,k,\text{roll}} = \frac{1}{36} \sum_{d_1,d_2 \in [1,6]} \begin{cases} P_{i,j,2k}, & \text{if } d_1 = 1 \text{ and } d_2 = 6; \\ P_{i,j,k+2d_2}, & \text{if } d_1 = 1 \text{ and } d_2 \neq 6; \\ P_{i,j,\max(0,k-d_1)}, & \text{if } d_1 \neq 1 \text{ and } d_2 = 6; \\ 1 - P_{j,i,0}, & \text{otherwise if } d_1 + d_2 = 7; \\ P_{i,j,k+d_1+d_2}, & \text{otherwise.} \end{cases}$$

Examining the optimal policy (**Figure 2**), we see another reason this variant is not popular. Optimal play essentially consists of waiting for two turns that score approximately half of the points each. If both players have no points, it is optimal to hold at 46. If the opponent has scored between 46 and 51 points, the optimal player holds in the range of 61–65 points. Having once scored, a player seeks to score the remaining number of points to win.

2-Dice Pig

According to game analyst Bill Butler, one of the simplest and most common variants, which we will call *2-dice Pig*, was produced commercially under the name "Pig" in the 1950s. The rules are the same as our 1-die Pig except:

- Two standard dice are rolled. If neither shows a 1, their sum is added to the turn total.
- If a single 1 is rolled, the player's turn ends with the loss of the turn total.

[1] ©1942 Parker Brothers

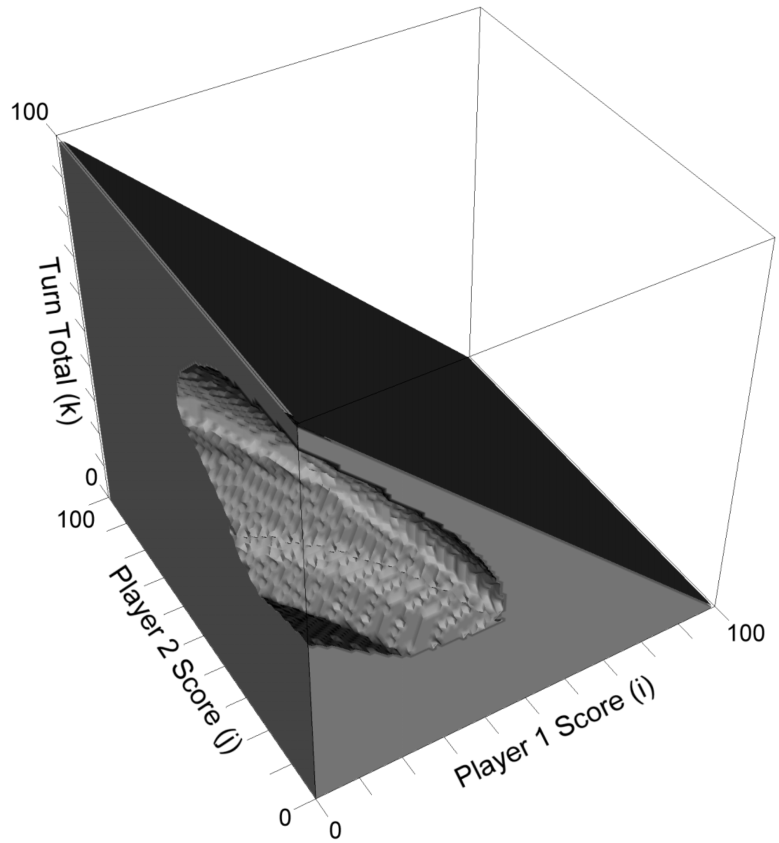

Figure 2. Roll / hold boundary for optimal Pig Dice policy.

- If two 1s are rolled, the player's turn ends with the loss of the turn total *and score*.

Optimality equations for Two-Dice Pig are the same as those for Pig with the following difference:

$$P_{i,j,k,\text{roll}} = \frac{1}{36} \sum_{d_1, d_2 \in [1,6]} \begin{cases} 1 - P_{j,0,0}, & \text{if } d_1 = d_2 = 1; \\ 1 - P_{j,i,0}, & \text{if either } d_1 = 1 \text{ or } d_2 = 1; \\ P_{i,j,k+d_1+d_2}, & \text{otherwise.} \end{cases}$$

Examining the optimal policy (**Figure 3**), we see that it is very similar to the optimal policy for Pig, with the greatest differences in the endgame, where "troughs" in the roll / hold boundary are shifted, and one seeks to roll all remaining points at higher scores. Even so, optimal play is very similar to Pig, despite the addition of a second die.

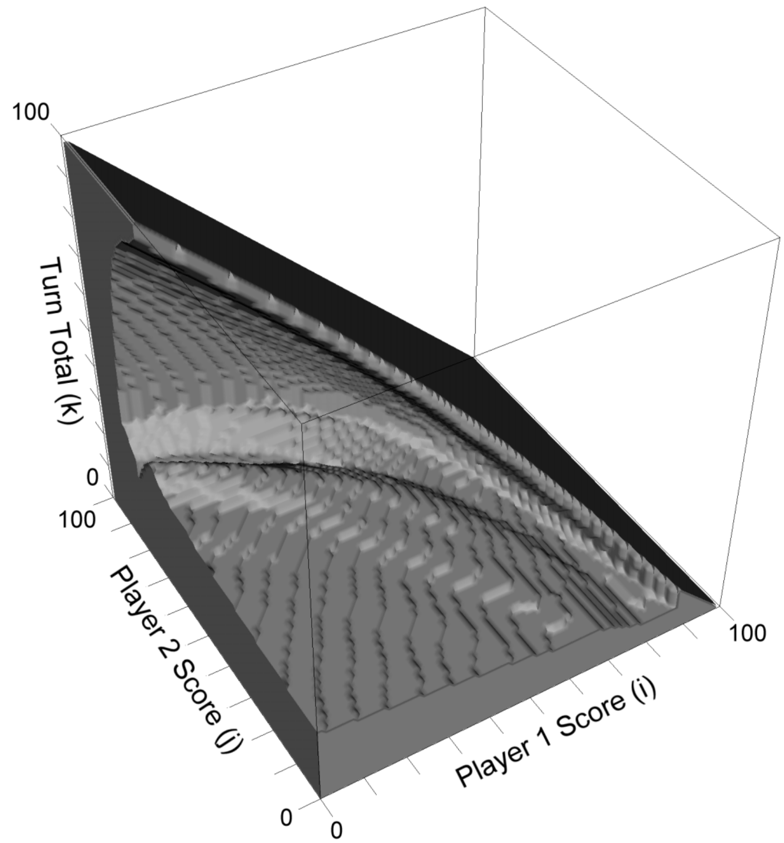

Figure 3. Roll / hold boundary for optimal Two-Dice Pig policy.

Frey's Pig

Skip Frey describes a 2-dice Pig variation [1975] that differs only in how doubles are treated and how the game ends.

- If two 1s are rolled, the player adds 25 to the turn total and it becomes the opponent's turn.
- If other doubles are rolled, the player adds twice the value of the dice to the turn total, and the player's turn continues.
- Players are permitted the same number of turns. If the first player scores 100 or more points, the second player must be allowed the opportunity to exceed the first player's score and win.

For symmetry in analysis we ignore the "same number of turns" requirement, awarding victory to the first player to reach 100 points.

Optimality equations for Two-Dice Pig are the same as those for Pig with the following difference:

$$P_{i,j,k,\text{roll}} = \frac{1}{36} \sum_{d_1,d_2 \in [1,6]} \begin{cases} 1 - P_{j,i+k+25,0}, & \text{if } d_1 = d_2 = 1; \\ P_{i,j,k+4d_1}, & \text{if } d_1 = d_2 \neq 1; \\ 1 - P_{j,i,0}, & \text{if either } d_1 = 1 \text{ or } d_2 = 1; \\ P_{i,j,k+d_1+d_2}, & \text{otherwise.} \end{cases}$$

Examining the optimal policy (**Figure 4**), we see that the roll / hold boundary has the same shape, yet play is significantly different from higher expected turn totals. At the beginning of the game, one holds at 28 rather than 21. An optimal player seeks to reach the goal score in 1-3 scoring turns depending on the opponent's score.

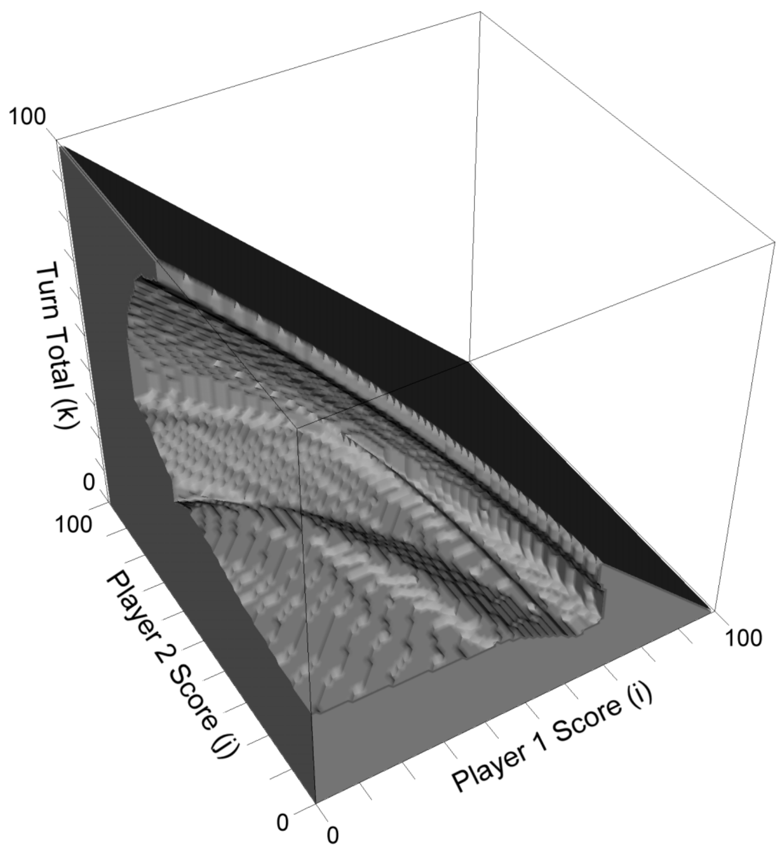

Figure 4. Roll / hold boundary for optimal Frey's Pig policy.

Big Pig

Knizia [1999] describes Frey's Pig as being a variant of "Big Pig", which differs only in that the player's turn continues after rolling a 1-1 ("big pig").

Thus, the difference in equations is:

$$P_{i,j,k,\mathrm{roll}} = \frac{1}{36} \sum_{d_1,d_2 \in [1,6]} \begin{cases} P_{i,j,k+25}, & \text{if } d_1 = d_2 = 1; \\ P_{i,j,k+4d_1}, & \text{if } d_1 = d_2 \neq 1; \\ 1 - P_{j,i,0}, & \text{if either } d_1 = 1 \text{ or } d_2 = 1; \\ P_{i,j,k+d_1+d_2}, & \text{otherwise.} \end{cases}$$

Examining the optimal policy (**Figure 5**), we note that it is similar to that of Frey's Pig, with minor variations in the roll / hold boundary.

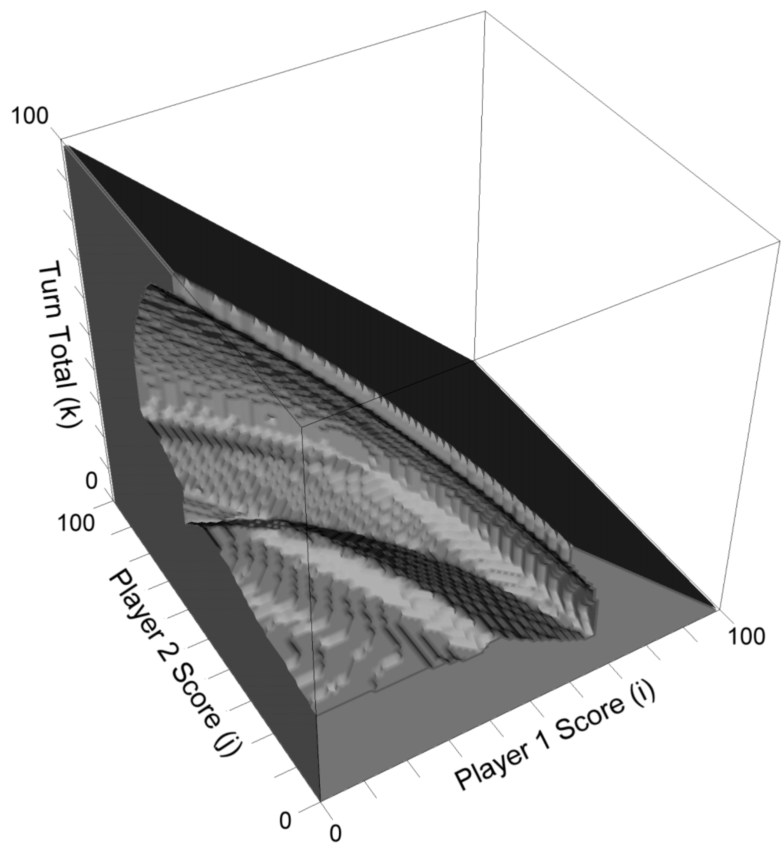

Figure 5. Roll / hold boundary for optimal Big Pig policy.

Piggy

Ivars Peterson [2000] describes Piggy, which varies from 2-dice Pig in that there is no single bad dice value. Rather, *doubles* have the same consequences as a single 1 in 2-dice Pig. Thus, the difference in equations is:

$$P_{i,j,k,\mathrm{roll}} = \frac{1}{36} \sum_{d_1,d_2 \in [1,6]} \begin{cases} 1 - P_{j,i,0}, & \text{if } d_1 = d_2; \\ P_{i,j,k+d_1+d_2}, & \text{otherwise.} \end{cases}$$

Examining the optimal policy (**Figure 6**), we see that optimal play of Piggy is very similar to that of Parker Brothers Pig Dice. That is, one seeks to reach the goal score in two scoring turns, holding at roughly half the goal score.

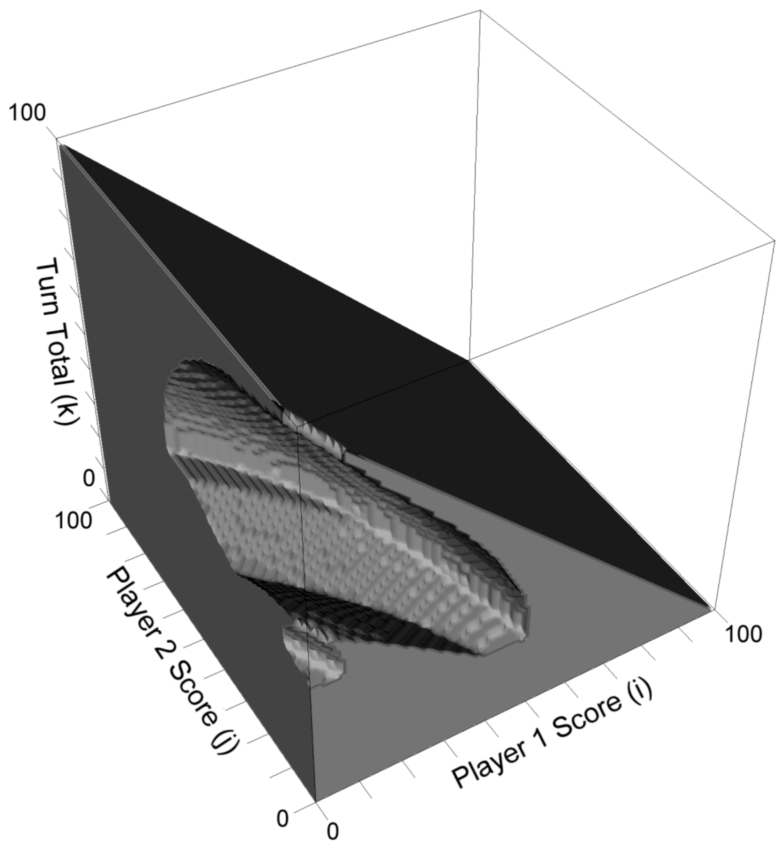

Figure 6. Roll / hold boundary for optimal Piggy policy.

Piggy Sevens

Piggy Sevens is a variant of Piggy we devised. In this variant, it is a total of 7 rather than doubles that causes the turn to end with the turn total lost. Thus, the difference in equations is:

$$P_{i,j,k,\mathrm{roll}} = \frac{1}{36} \sum_{d_1,d_2 \in [1,6]} \begin{cases} 1 - P_{j,i,0}, & \text{if } d_1 + d_2 = 7; \\ P_{i,j,k+d_1+d_2}, & \text{otherwise.} \end{cases}$$

The optimal policy (**Figure 7**) is almost identical to that of Piggy.

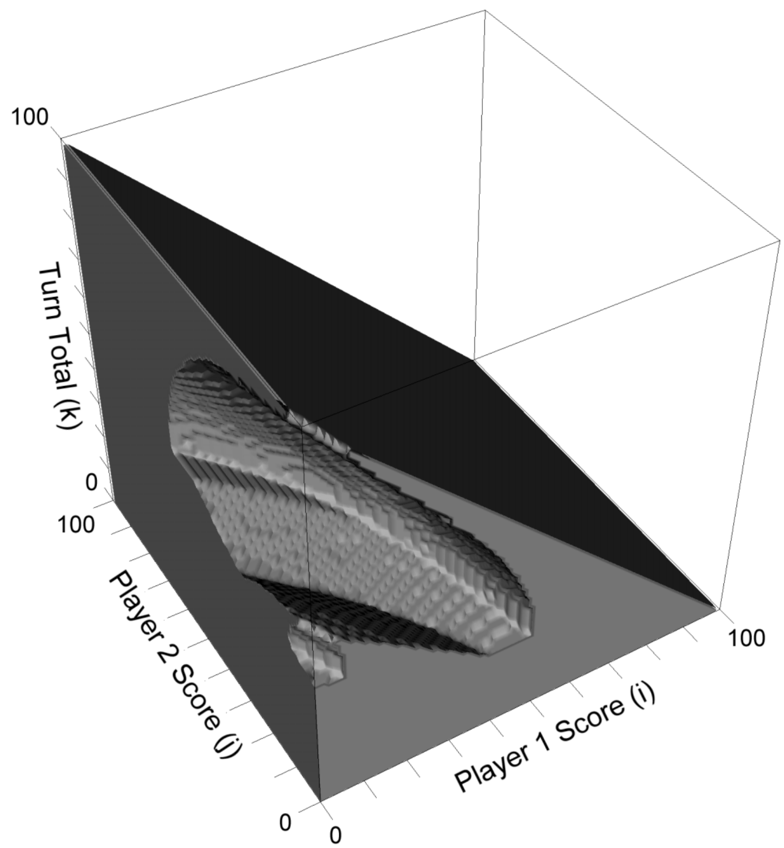

Figure 7. Roll / hold boundary for optimal Piggy Sevens policy.

Pigmania/Pass the Pigs

A popular commercial variant, Pass the Pigs®[2] (which was originally called PigMania®[3]) was designed by David Moffat in 1977. In this variant, small rubber pigs are used as dice. When rolled, each pig can come to rest in a variety of positions with varying probability: on its right side, on its left side (marked with a dot for easy perception), upside down (a.k.a. "razorback"), upright

[2] ©1995 David Moffat Enterprises and Hasbro, Inc.
[3] ©1977 David Moffat Enterprises

(a.k.a. "trotter"), balanced on the snout and front legs (a.k.a. "snouter"), and balanced on the snout, left leg, and left ear (a.k.a. "leaning jowler").

In 1997, 52 6th-grade students of Dean Ballard at Lakeside Middle School in Seattle, Washington rolled such pigs 3939 times. In the order of roll types listed above, the number of rolls were: 1344, 1294, 767, 365, 137, and 32. [4]

In Pass the Pigs, a pig coming to rest against another is called an "oinker" and results in the loss of all points. Since the pigs may be thrown in such a way to make this an unlikely event, we assume that the probability of an oinker $p_{\text{oinker}} = 0$. However, we will include this parameter in the optimality equations below if one wishes to examine cases where $p_{\text{oinker}} > 0$.

PigMania is similar to 2-dice Pig, in that a roll of left side and right side has the same consequences as rolling a 1 (i.e., the turn ends with the loss of the turn total), and a roll with pigs touching has the same consequences as rolling two 1s (i.e., the turn ends with the loss of the turn total and score). PigMania is similar to Frey's variant in that two pigs in the same non-side configuration score double what they would individually.

Let the rolls "right sider", "left sider", "razorback", "trotter", "snouter", and "leaning jowler" be numbered 0, 1, 2, 3, 4, and 5, respectively. Let r_i denote the number of times roll i occurred in the data. Let r denote the total number of rolls $\sum_{i=0}^{5} r_i$. Let the point values $p_0, \ldots, p_5$ be set to 0, 0, 5, 5, 10, and 15, respectively. Then the difference in optimality equations is:

$$P_{i,j,k,\text{roll}} = p_{\text{oinker}}(1 - P_{j,0,0}) + (1 - p_{\text{oinker}})P_{i,j,k,\text{roll,non-oinker}},$$

$$P_{i,j,k,\text{roll,non-oinker}} = \frac{1}{r^2} \sum_{d_1,d_2 \in [0,5]} r_{d_1} r_{d_2} \begin{cases} 1 - P_{j,i,0}, & \text{if } d_1 + d_2 = 1; \\ P_{i,j,k+1}, & \text{if } d_1 = d_2 \le 2; \\ P_{i,j,k+4p_{d_1}}, & \text{if } d_1 = d_2 > 2, \text{ or} \\ P_{i,j,k+p_{d_1}+p_{d_2}}, & \text{otherwise.} \end{cases}$$

The optimal policy (**Figure 8**) is very similar to the optimal policy for Pig. In fact, it is one of the closest matches for all known Pig variants. On the first turn, the player should hold at 21. Assuming two optimal players, the starting player will win with a probability of approximately 0.5310, a 6% advantage over the opponent.

Hog

Hog [Bohan and Shultz 1996; Feldman and Morgan 2003; Mathematical Sciences Education Board 1993] is a variation of Pig in which players have only one roll per turn but may roll as many dice as desired. If no 1s are rolled, the sum of the dice is scored. If any 1s are rolled, no points are scored for the turn. It is as if a Pig player must commit to the number of rolls in a turn before the turn begins.

[4] Data source: `http://members.tripod.com/~passpigs/prob.html`.

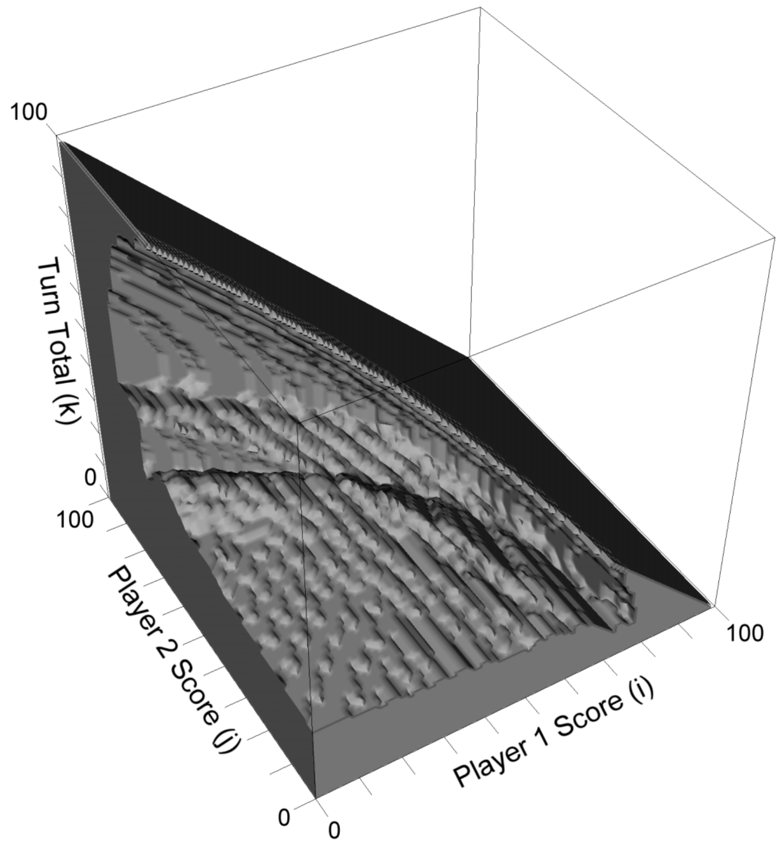

Figure 8. Roll / hold boundary for optimal Pass the Pigs policy.

This is technically not a jeopardy dice game, as one is never faced with a roll / hold decision where a turn total is in jeopardy. Indeed, the state space has no turn total, so non-terminal states are indexed by (i, j) where the player's score is $0 \leq i < 100$, and the opponent's score is $0 \leq j < 100$.

In our computation of optimal policy, we assume there is a finite maximum number of dice $d_{\max}$ to be rolled. For $d_{\max} \geq 26$, the optimal policy remains the same. Let $\pi(d, k)$ denote the probability that rolling $0 < d \leq d_{\max}$ dice will result in a turn score of $k \geq 0$. $\pi(d, k)$ can be computed using dynamic programming[5] with the following equations:

[5] Dynamic programming is described in Cormen et al. [2001, Ch. 15], Campbell [2002], and Neller [2004].

$$\pi(d,k) = \begin{cases} \frac{1}{6}, & d = 1 \text{ and } k \in \{0,2,3,4,5,6\}; \\ 0, & d = 1 \text{ and } k \notin \{0,2,3,4,5,6\}; \\ \pi(d-1,0) + \frac{1}{6}\sum_{k=2}^{6(d-1)} \pi(d-1,k), & d > 1 \text{ and } k = 0; \\ \frac{1}{6}\sum_{r=2}^{\min(6,k-2)} \pi(d-1,k-r), & \text{otherwise.} \end{cases}$$

The decision a player is faced with is not whether to roll or hold, but rather how many dice d to roll. The optimality equations for play thus take on a very different form. Let $P_{i,j}$ denote the player's probability of winning. In the case where $i \geq 100$, we have $P_{i,j} = 1$ because the player can simply hold and win. In the general case where $0 \leq i, j < 100$, the probability of winning is

$$P_{i,j} = \max_{0<d\leq d_{\max}} \sum_{k=0}^{6d} \pi(d,k)(1 - P_{j,i+k}).$$

Interestingly, the optimal die roll policy for Hog (**Figure 9**) has a similar shape to that of the optimal Pig roll / hold boundary. Multiplying the optimal number of dice to roll by 4 (the average good die roll) approximates the optimal roll / hold boundary for Pig. Assuming two optimal players, the starting player will win with a probability of approximately 0.5299, a 6% advantage over the opponent. On the first turn, the player should roll 6 dice.

If we compare the expected score per turn for Pig and Hog (**Figure 10**), we see that a player seeks a similar expected turn score in most situations.

Future Challenges

Endgame Variations

For symmetry and state minimization in analyzing some of these games, we have simplified some endgame conditions. In all cases, the first player to reach the given goal score wins. However, all of these games could be played such that other players have one more chance to win by achieving the highest score. Here are two common variations:

- A player reaching the goal score begins the last round of turns. Once a player reaches the goal score, all other players have one more opportunity to achieve the highest score.

- All players are allowed an equal number of turns. Once a player reaches the goal score, all other players *remaining in the turn order* have one more opportunity to achieve the highest score.

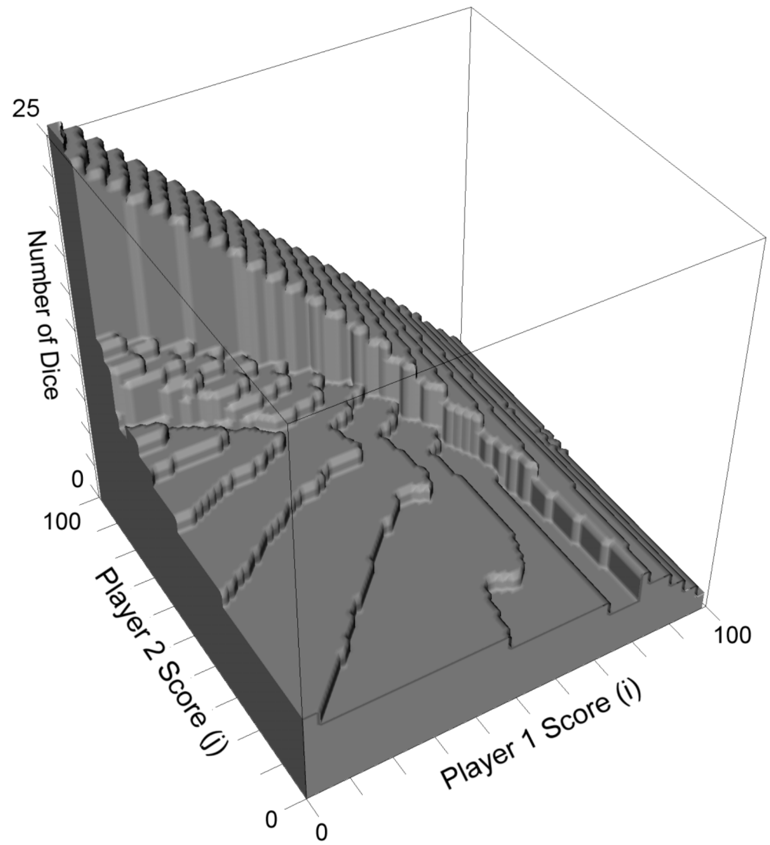

Figure 9. Dice to roll for optimal Hog policy.

If, after the last turn, the maximum score is tied, the first player to reach that score wins.

These endgame variations presumably encourage players to exceed the goal score by a margin in close games, but by how much? What fraction of the time does the first optimal player win against an optimal opponent? These are open questions left to the curious reader. Optimality equations will now have to contain additional state information to model endgame play.

Three or More Players

There is no reason the same analyses and techniques cannot be extended to Pig games with three or more players. Of course, the size of the state space grows exponentially with the number of players.

One might approximate optimal n-player play policy by

- abstracting all of a player's opponents as a single opponent with their maximum score, and

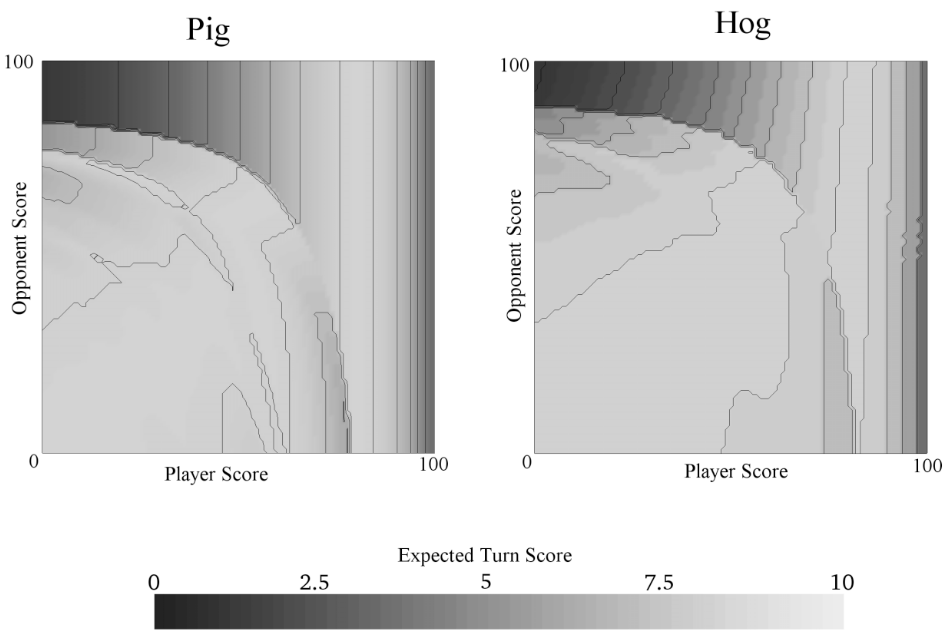

Figure 10. Expected turn scores for Pig and Hog.

- applying optimal 2-player play policy.

How good an approximation is this? Again, this is an open question left to the curious reader.

Further Investigations

Further related exercises may be found as part of [Neller 2004]. These curricular materials for introductory artificial intelligence were developed as part of NSF DUE CCLI-A&I Award Number 0409497.

Conclusion

Judging from the archival and online mathematical literature on Pig, this simple jeopardy dice game has served well as an entertaining focus problem in probability education. In this addendum, we have surveyed optimal play for several variants of Pig, observing roll / hold surfaces of optimal play.

We conclude by noting that mathematical insight can aid in game design. Pig variants where optimal play takes four or five successful turns to win are much more popular than those that take two. One can compute values as simple as the expected game duration, or as complex as the expected probability of visiting a state. These mathematical insights combined with human insights of perceived rule complexity, game tension, average player attention span, etc. suggest a parting question for the reader: Considering such objective and subjective measures, can you design a "better" jeopardy dice game?

References

Bohan, James F., and John L. Shultz. 1996. Revisiting and extending the Hog game. *Mathematics Teacher* 89 (9): 728–733.

Campbell, Paul J. 2002. Farmer Klaus and the Mouse. *The UMAP Journal* 23 (2): 121–134. Errata, 24 (2003) (4): 484; 25 (2004) (4): 446.

Cormen, Thomas H., Charles E. Leiserson, Clifford Stein, and Ronald L. Rivest. 2001. *Introduction to Algorithms*. 2nd ed. Cambridge, MA: MIT Press.

Feldman, Larry, and Fred Morgan. 2003. The pedagogy and probability of the dice game HOG. *Journal of Statistics Education* 11 (2). `http://www.amstat.org/publications/jse/v11n2/feldman.html` .

Frey, Skip. 1975. *How to Win at Dice Games*. 1997. Reprint. North Hollywood, CA: Wilshire Book Co.

Knizia, Reiner. 1999. *Dice Games Properly Explained*. Brighton Road, Lower Kingswood, Tadworth, Surrey, KT20 6TD, UK: Elliot Right-Way Books.

Mathematical Sciences Education Board. 1993. *Measuring Up: Prototypes for Mathematics Assessment*. Washington, DC: National Academy Press.

Neller, Todd W. 2004. Solving the dice game Pig: An introduction to dynamic programming and value iteration. `http://cs.gettysburg.edu/~tneller/nsf/pig/index.html` .

________, and Clifton G.M. Presser. 2004. Optimal play of the dice game Pig. *The UMAP Journal* 25 (1): 25–47.

Peterson, Ivars. 2000. Weird dice. *Muse Magazine* (May / June 2000): 18.

About the Authors

Todd W. Neller is Assistant Professor of Computer Science at Gettysburg College. A Cornell University Merrill Presidential Scholar, he received a B.S. in Computer Science with distinction in 1993. He received a Stanford University Lieberman Fellowship in 1998, where he received the George E. Forsythe Memorial Award for excellence in teaching and a Ph.D. in Computer Science in 2000. His thesis concerned extension of artificial intelligence search algorithms to hybrid dynamical systems, and the refutation of hybrid system properties through simulation and information-based optimization. Recent works have concerned the application of reinforcement learning techniques to the control of optimization and search algorithms.

Clifton G.M. Presser is also Assistant Professor of Computer Science at Gettysburg College. He received a B.S. in Mathematics and Computer Science from Pepperdine University in 1993. Clif received his Ph.D. in Computer Science at the University of South Carolina in 2000, where he received the Outstanding Graduate Student Award in the same year. Clif's dissertation research was on automated planning in uncertain environments. Currently, his research concerns computer visualization of high-dimensional geometry, algorithms and information.