Dinosaur War: A Strategic Game of Utter Chance

Dave Fetterman

Obviously Unemployed

3/6/23

Abstract

We present a modified version of the simple game *War* with two identical decks, no replacement, and player card choice, invented in large part and played by my preschool-age children. Unlike *War*, there are choices that must be made by the players. But like *War*, the outcome of the game, when played by rational agents, remains 100 percent chance. This new game, *Dinosaur War*, resembles something more akin to *Rock-Scissors-Paper*[1]; knowing an opponent's guess can guarantee a win, but like *Rock-Scissors-Paper*, we show a Nash Equilibrium occurs if both players randomize their guesses uniformly across their remaining cards, whatever they may be. This result is intuitive but non-obvious.

Therefore:
- You can play optimally against your child by paying no attention at all.
- Expect a Pokemon-branded version to hit the shelves soon.

1 The Game

Children's games need to be simple. The game *Memory* has seen innumerable rebranded recreations, because the mechanic is approachable (and nominally educational) and it can be sold repeatedly, with cartoon characters, animals, or whatever to engage a short attention span. A set of *Memory* comes with matched pairs of cards with identical backs. Once the main mechanic is exhausted, the enterprising child will find some other game to create with this set. Here is that game, *Dinosaur War*, created with the cards like those in Figure 1.

1.1 Rules of Dinosaur War

- Players establish a ranking of cards, preferably under the direction of an opinionated child. Those might be "Baryonyx beats Mosasaurus beats T-Rex... beats Ap-

[1] Fine, you call it "Rock Paper Scissors". Save the pedantry for the math part.

Figure 1: Dinosaur Cards

atosaurus" in Figure 1, or the commonly-accepted Ace-to-2 if using a deck of playing cards.

- Two players each get an identical deck of these cards. Cards were unique in this set but need not be. Players conceal their hand (though the content of the hands is well-known to those tracking it).
- At each turn:
  - Each player simultaneously plays a card face up.
  - The player whose card outranks the other's gets one point. If there is a tie, no points are awarded. The two played cards are set aside.
  - Play continues until the cards are exhausted.
- The player with the most points at the end wins.

The maximum score individual score is 9 (since your opponent's 10 cannot be beaten, only tied). Game ties are relatively common.

1.2 "Strategy" in Dinosaur War

Intuitively, your hand has a certain amount of "power" that you deploy to beat an opponent; spending the minimum amount of "power" to win preserves better cards for later.

Imagine on the first turn of a 10 card deck game with hands $A = B = \{1, 2, ...10\}$, players (P_A, P_B) play respective cards $(9, 10)$. This means:

- P_B takes a one-point lead.
- The 10 card is preserved for P_A. They will necessarily win one hand in the future.

- The powerful 9 card is lost for P_A.

Alternatively, imagine the first move is $(1, 10)$. This means:

- P_B takes a one-point lead.
- 10 is preserved for P_A. They will necessarily win one hand in the future.
- 1 is lost for P_A, the worst card in the hand.
- Each card of P_A's hand beats at least one card in P_B's hand.

The second scenario *seems* better for P_A.[2] But how much better? And how can one strategically strive to lose bad cards and win "by just enough" to take tricks? This is the focus of the paper.

1.3 A reduced example

Throughout, we'll use the following conventions:

- P_A's available options are listed in bold down the left column of the payoff matrix (Fig 2a, 2b).
- P_B's available options are listed in bold across the top row.
- A trick has a payoff of 1 if P_A wins, and -1 if P_B wins. P_A is trying to get the total score as high above zero as possible, P_B below.
- For a payoff matrix M, the cell at row i, column j is the value of that trick, plus the expected value of the remaining game.

This is easy to see in Fig 2a, where the hands are identical. There are only four games of (P_A, P_B) move pairs:

- $(1, 1)$ means the first trick payoff is zero, and the rest of the game (necessarily $(3, 3)$) is determined, also of payoff zero.
- $(3, 3)$ follows similarly.
- $(3, 1)$ means the first trick payoff is 1, and the rest of the game (necessarily $(1, 3)$) pays off -1, for a total of zero.
- $(1, 3)$ follows in reverse, with another time game.

It's clear that *any strategy* is equivalent in this very boring small game. P_A could even announce his moves before P_B selects a card, and the result of the game is still determined. The expected (and only possible)value of this game is zero.

[2] In this paper, we measure goodness by expected tricks taken by the hand.

$$\begin{bmatrix} & \mathbf{1} & \mathbf{3} \\ \mathbf{1} & 0 & 0 \\ \mathbf{3} & 0 & 0 \end{bmatrix}$$

(a) Even 2x2 game

$$\begin{bmatrix} & \mathbf{3} & \mathbf{4} \\ \mathbf{1} & 0 & 0 \\ \mathbf{5} & 0 & 0 \end{bmatrix}$$

(b) Another even 2x2 game

$$\begin{bmatrix} & \mathbf{1} & \mathbf{3} \\ \mathbf{2} & 2 & 0 \\ \mathbf{4} & 0 & 2 \end{bmatrix}$$

(c) Uneven 2x2 game

Figure 2: Simple 2x2 games

Observe in figure 2b that starting sets $A = \{1, 5\}, B = \{3, 4\}$, while not identical, also yield this result; the choices don't matter in the end.

But for some uneven sets of cards, like in Fig 2c, things are different.

- If P_A is able to play their 2 against a 1 (on either first or second trick), they win both tricks for a score of 2.
- If P_A plays their 2 against a 3, this trick score is -1, but guaranteed to balance by the imminent (or recently played) $(4, 1)$ trick, for a total of 0.

This is more like *Rock-Scissors-Paper*: knowing your opponent's choice wins you the game. And, like RSP, the mere existence of better choices does not mean that there exists a perfect-information strategy with a nonzero expected value.

How can we quantify the goodness of one hand versus another? We introduce a metric for this particular game called the *Dominance Score*[3] and use this to compute the expected value of more complicated (larger) games.

2 Dominance Score

The dominance score of two equal-sized sets (hands) $A = \{a_1, a_2, ...a_n\}, B = \{a_1, a_2, ...a_n\}$ is defined as $\boxed{\text{D(A,B)} = \sum_{i=1}^{n} \sum_{j=1}^{n} T(a_i, b_j)}$, where

[3]This should not be conflated with a *dominant* Nash strategy.

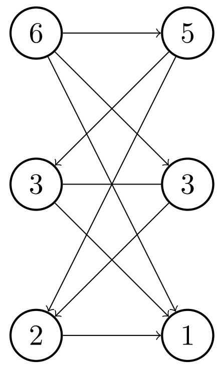

Figure 3: Dominance graph of $\{2, 3, 6\}$ vs. $\{1, 3, 5\}$

$$T(a,b) = \begin{cases} 1 & a > b \\ 0 & a = b \\ -1 & a < b \end{cases}$$

This is just adding up all possible winning tricks for A and subtracting all possible winning tricks for B, ignoring ties. For example, if $A = \{2, 3, 6\}, B = \{1, 3, 5\}$, then $D(A, B) = [T(2, 1) + T(3, 1) + T(6, 1) + T(6, 3) + T(6, 5)] + [T(3, 3)] + [T(2, 3) + T(2, 5) + T(3, 5)] = 5 + 0 - 3 = 2$.

- The identical hands in figure 2a necessarily have a dominance score of zero.
- Figure 2b's pair of a hand with the 2nd- and 3rd-highest cards versus one with the lowest and highest rank, also has a dominance score of zero.
- Figure 2c has a dominance score of two, so it's not surprising that the expected value of the game is in player A's favor (+2).

Another way to visualize hand A against hand B is a bipartite graph like Figure 3, counting "right-pointing" edges as +1, and "left-pointing" edges as -1. This shows that D($\{2, 3, 6\}$, $\{1, 3, 5\}$) = 2.

3 Main Theorem and Proof Layout

The main theorem we wish to prove states:

> Given rational players, an optimal strategy in Dinosaur War is playing all options with uniform randomness.

The steps to proving this are.

- Lemma 1: Across every row and column of the payoff matrix M determined by hands A and B, the entries sum to the graph's domination score $D(A, B)$.
- Lemma 2: A payoff matrix of such a form of width n has a Nash equilibrium[**2**] of $p(a_i) = \frac{1}{n}, p(b_j) = \frac{1}{n}$ for all $i, j \in [1, n]$.
- Lemma 3: The expected value of the game A, B is then $\frac{D(A,B)}{n}$.

We will do this *inductively* and *simultaneously*, in that, if Lemma 1, 2, and 3 are true for boards of size $n - 1$ and below, then we can prove them for a board of size n.

Additionally: There are no equilibria with a higher expected value in the game, since according to Von Neumann's Minimax theorem[**3**], all Nash equilibria of a zero-sum game have the same value. So though there may be other equally good strategies (duplicate cards in a hand allows the strategy to play *one* versus the *other* with some variation; also see the fait accompli games of Figure 2a, 2b), there are none more optimal than a uniform strategy.

Finally, we compute our payoff matrix for a larger game of card sets $A = B = [1, 10]$.

4 Lemmas: Base case

4.0.1 Base cases

Base case, n=1: We see in Figure 4a that D is equal to function H at $n = 1$: a 1 if player A's single card outranks B's, a 0 for a tie, and a -1 if B's outranks. With a single element and therefore single row and column, the Lemma 1 is clearly true. There is only one total strategy in the game, so Lemma 2 is true. And Lemma 3 is equivalent to Lemma 1 when $n = 1$.

(Extra) base case, n=2: To see how this extends to a 2x2 matrix, consider Fig. 4b displaying the payoff matrix built from hands $A = (a_1, a_2) = \{3, 6\}; B = (b_1, b_2) = \{3, 5\}$.

- The upper-left element is $T(a_1, b_1) + D(\{a_2\}, \{b_2\}) = T(a_1, b_1) + T(a_2, b_2)$
- The upper-right element is $T(a_1, b_2) + D(\{a_2\}, \{b_1\}) = T(a_1, b_2) + T(a_2, b_1)$
- The lower-left element is $T(a_2, b_1) + D(\{a_1\}, \{b_2\}) = T(a_2, b_1) + T(a_1, b_2)$
- The lower-right element is $T(a_2, b_2) + D(\{a_1\}, \{b_1\}) = T(a_2, b_2) + T(a_1, b_1)$

It is clear that summing the top row, bottom row, left column, or right column yields $T(a_1, b_1) + T(a_1, b_2) + T(a_2, b_1) + T(a_2, b_2) = D(A, B)$.

At the 2x2 size, it should be clear that:

$$\begin{bmatrix} & \mathbf{5} \\ \mathbf{6} & 1 \end{bmatrix}$$

(a) 1x1 base case

$$\begin{bmatrix} & \mathbf{3} & \mathbf{5} \\ \mathbf{3} & 1 & 0 \\ \mathbf{6} & 0 & 1 \end{bmatrix}$$

(b) 2x2 base case

Figure 4: Base cases

- The upper left and lower right values are necessarily equal, as are the upper right and lower left. There are four possible game sequences (completely determined by the first move), and those starting with, say, $(3, 3)$ in figure 5b need to end with $(6, 5)$, equivalent to $[(6, 5), (3, 3)]$.
- Because of this identity, the rows and columns all sum to the same value. (proving Lemma 1 at $n = 2$)
- An equilibrium strategy of this zero sum game is playing each option with probability $\frac{1}{2}$, as the players are essentially playing a game of Matching Pennies[**1**] (showing Lemma 2).
- Therefore, the expected payoff is the average of a(ny) row or column (showing Lemma 3).

5 Inductive case

For the *inductive case*, consider Figure 5a and Figure 5b. In each, the vertical axis is the set of moves A for player P_A, and the horizontal the move set B for player P_B.

In Figure 5a, the element (i, j) represents the game continuation (or subgame) should the next play be (a_i, b_j) (so, excluding cards a_i and b_j). In Figure 5b, the element (i, j) is of the form ("immediate payoff from move (a_i, b_j)" and "expected payoff of remaining subgame"). So, the upper right element of 5b is ($\mathit{1} + .5 = 1.5$); the italicized 1 represents that card 2 beats card 1. .5 points is the expected payoff of the subgame $\{3, 6\}$ vs. $\{3, 5\}$, totaling 1.5.

5.1 Inductive case for Lemma 1

In 5a, the subgame at (i, j) is the full game $(A - \{a_i\}, B - \{b_j\})$. This is "the rest of the game" with cards a_i, b_j already used up. By inductive hypothesis, the payoff of the

$$
\begin{bmatrix}
 & \mathbf{1} & \mathbf{3} & \mathbf{5} \\
\mathbf{2} & \begin{bmatrix} & \mathbf{3} & \mathbf{5} \\ \mathbf{3} & 1 & 0 \\ \mathbf{6} & 0 & 1 \end{bmatrix} & \begin{bmatrix} & \mathbf{1} & \mathbf{5} \\ \mathbf{3} & 2 & 0 \\ \mathbf{6} & 0 & 2 \end{bmatrix} & \begin{bmatrix} & \mathbf{1} & \mathbf{3} \\ \mathbf{3} & 2 & 1 \\ \mathbf{6} & 1 & 2 \end{bmatrix} \\
\mathbf{3} & \begin{bmatrix} & \mathbf{3} & \mathbf{5} \\ \mathbf{2} & 0 & 0 \\ \mathbf{6} & 0 & 0 \end{bmatrix} & \begin{bmatrix} & \mathbf{1} & \mathbf{5} \\ \mathbf{2} & 2 & 0 \\ \mathbf{6} & 0 & 2 \end{bmatrix} & \begin{bmatrix} & \mathbf{1} & \mathbf{3} \\ \mathbf{2} & 2 & 0 \\ \mathbf{6} & 0 & 2 \end{bmatrix} \\
\mathbf{6} & \begin{bmatrix} & \mathbf{3} & \mathbf{5} \\ \mathbf{2} & -2 & -1 \\ \mathbf{3} & -1 & -2 \end{bmatrix} & \begin{bmatrix} & \mathbf{1} & \mathbf{5} \\ \mathbf{2} & 0 & 0 \\ \mathbf{3} & 0 & 0 \end{bmatrix} & \begin{bmatrix} & \mathbf{1} & \mathbf{3} \\ \mathbf{2} & 1 & 0 \\ \mathbf{3} & 0 & 1 \end{bmatrix}
\end{bmatrix}
$$

(a) Recursive game matrix

$$
\begin{bmatrix}
 & \mathbf{1} & \mathbf{3} & \mathbf{5} \\
\mathbf{2} & (\mathit{1} + .5 = 1.5) & (-\mathit{1} + 1 = 0) & (-\mathit{1} + 1.5 = .5) \\
\mathbf{3} & (\mathit{1} + 0 = 1) & (\mathit{0} + 1 = 1) & (-\mathit{1} + 1 = 0) \\
\mathbf{6} & (\mathit{1} + -1.5 = -.5) & (\mathit{1} + 0 = 1) & (\mathit{1} + .5 = 1.5)
\end{bmatrix}
$$

(b) Payoff matrix

Figure 5: $\{2, 3, 6\}$ vs. $\{1, 3, 5\}$

game at matrix entry $(1,1)$ is $D(A-\{a_i\}, B-\{b_j\})$, which is the sum of $T(a_m, b_n)$ for all combinations of m, n except all those where $m = i$ or $n = j$.

Without loss of generality, consider row a_1 (P_A plays card 2) in Figure 5a (the subgames) and in figure 5b. Note that for $i = 1$, summing across all expected subgame payoffs $\frac{D(A-\{a_1\},B-\{b_j\})}{n-1}$ counts each pair $T(a_i, b_j), i \neq 1$ exactly $n-1$ times. $a_i, i \neq 1$ is always in each subgame, and b_j is in every subgame except those in column j, for a total of $n-1$. So, with each D term fully expanded, the sum includes $n-1$ terms $\frac{T(a_i,b_j)}{n-1}, i \neq 1$. These correspond to the numbers on the right side of the addition sign in figure 5b.

The immediate payoffs of playing a_1 (card 2 in this case) are just $T(a_1, b_1), T(a_1, b_2)...T(a_1, b_n)$, corresponding to the (italicized) number on the left side of the addition sign in 5b. Of course, summing the immediate payoffs ($\sum_{i=2}^{n} \sum_{j=1}^{n} T(a_i, b_j)$) and the subgame payoffs ($\sum_{j=1}^{n} T(a_1 b_j)$) is the definition of $D(A, B)$. This proves Lemma 1.

This logic holds equivalently for any other row or any column.

5.2 Lemma 2

Consider a payoff matrix where each row and each column sum to the same value (in our case, this is always $D(A, B)$).

A Nash equilibrium occurs for players P_A, P_B when any deviation of P_A from the equilibrium benefits P_B, assuming P_A announces his strategy (and same for P_B, P_A).

Assume both strategies are uniform: $p(a_1) = p(a_2) = ... = \frac{1}{n} = p(b_1) = p(b_2) = ...p(b_n)$. Then the payoff of any given row i would be, by Lemma 1, $\sum_{j=1}^{n} M_{i,j} = n \cdot \frac{1}{n} D(A, B) = D(A, B)$; the same follows for columns. The sum of all payoffs across the matrix is then $n \cdot D(A, B)$.

If P_A shifts his strategy to a different distribution $p(a_1), p(a_2), ..., p(a_n)$, note that the sum of all payoffs $n \cdot D(A, B)$ does not change, since this is simply reallocating the probabilities, which sum to 1, over different rows each of expected value $D(A, B)$.

- Case 1: If this reallocation sees each column j's expected payoff $\sum_{i=1}^{n} M_{i,j} p(a_i)$ still summing to $D(A, B)$, then the equilibrium condition (and the same expected payoff) is maintained.
- Case 2: If, however, a column j sums to a value greater than $D(A, B)$, then $p(b_j) = 1; p(b_k) = 0$ for $(k \neq j)$ increases P_B's payoff. As a zero-sum game, this decreases P_A's payoff, and P_A would not choose it. Note that if Case 1 does not hold, there *must* be a column like j, since the sum across all columns remains $n \cdot D(A, B)$. In other words, if we're not in Case 1, column sums can't all be less than or equal to $D(A, B)$, or the sum wouldn't be $n \cdot D(A, B)$, so at least one column has a payoff greater than $D(A, B)$.

	1	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**
1	0	−8/9	−2/3	−4/9	−2/9	0	2/9	4/9	2/3	8/9
2	8/9	0	−8/9	−2/3	−4/9	−2/9	0	2/9	4/9	2/3
3	2/3	8/9	0	−8/9	−2/3	−4/9	−2/9	0	2/9	4/9
4	4/9	2/3	8/9	0	−8/9	−2/3	−4/9	−2/9	0	2/9
5	2/9	4/9	2/3	8/9	0	−8/9	−2/3	−4/9	−2/9	0
6	0	2/9	4/9	2/3	8/9	0	−8/9	−2/3	−4/9	−2/9
7	−2/9	0	2/9	4/9	2/3	8/9	0	−8/9	−2/3	−4/9
8	−4/9	−2/9	0	2/9	4/9	2/3	8/9	0	−8/9	−2/3
9	−2/3	−4/9	−2/9	0	2/9	4/9	2/3	8/9	0	−8/9
10	−8/9	−2/3	−4/9	−2/9	0	2/9	4/9	2/3	8/9	0

Figure 6: $A = B = [1, 10]$ payoff matrix

With Lemma 2 proven, Lemma 3 follows quickly. Each $M_{i,j}$ occurs with probablity $\frac{1}{n^2}$, and the matrix entries sum to $n \cdot D(A, B)$, so the expected payoff is $\frac{D(A,B)}{n}$.

6 Considerations and Examples

Computing the expected value of the matrices in Figure 5a recursively gets computationally expensive, but with this formula in hand, the code [4] becomes very simple: just apply Lemmas 1-3, and the 10 v. 10 game payoffs are easily generated (Figure 6).

Note that:

- The matrix is obviously symmetric.
- Winning the first trick is neither a net positive nor negative.
- Winning by exactly $(n/2)$ rank spots is neutral; winning a trick by less than that yields a net positive value (for player P_A), and more than that, a negative value.

Though the double-uniform strategy is optimal, and no other perfect information strategies outperform it (Lemma 2), certainly, like *Rock-Scissors-Paper*, having an inkling of what the opponent will do can yield advantage.

Note again that there are cases where the choice of move truly does not matter, (like fig 2a and 2c), or, in some version of the game, where k duplicate cards could render the choice between them irrelevant (as long as the sum of the probabilities remains $\frac{k}{n}$.

This paper has only dealt with the expected value of the game, not the chance of winning,

[4] `https://github.com/fettermania/mathnotes/tree/main/dino/clj/dino/src/dino/core.clj`

though my intuition is that this would follow similarly[5].

References

[1] Wikipedia: https://en.wikipedia.org/wiki/Matching_pennies

[2] Wikipedia: https://en.wikipedia.org/wiki/Nash_equilibrium

[3] Wikipedia: https://en.wikipedia.org/wiki/Minimax_theorem

[5] Though this is not a proof!