The Big Match

Joshua Ruiter
Virginia Wesleyan University

March 25, 2025

Contents

1 The matching game

The matching game is a classic game between two players (A and B). Each player chooses a 1 or 2, then reveal choices simultaneously. If the sides match, A gets $+1$ point and B gets -1 point. If the sides don't match, A gets -1 point and B gets $+1$ point. Game theorists summarize this with a matrix.

	1	2
1	$(1, -1)$	$(-1, 1)$
2	$(-1, 1)$	$(1, -1)$

To read this matrix: A chooses a row, B chooses a column (simultaneously, not after knowing A's choice). A tuple (x, y) means A receives x and B receives y. These numbers are called **payoffs**. This is called a **zero-sum game** because the sum of payoffs to all players is zero, i.e. one player's gain is always another player's loss. Because it's zero-sum, we can drop the y components, and just track payoff to A.

	1	2
1	1	-1
2	-1	1

1.1 Play the game - 1st time

Activity time! Find a neighbor, decide on roles of A and B, then play 10 rounds. Record the results on a piece of paper.

- Who in the room got the most number of points?
- What strategies did people employ? What was your process for choosing 1 or 2?
- What is the best way to play this game, if you are trying to maximize the number of points you have after 10 rounds?

2 Terminology from game theory

Definition 1. A **pure strategy** in a game is a deterministic strategy which involves always making the same deterministic decision in a given situation.

Example 2. In the matching game, the pure strategies are "Always play 1" and "Always play 2."

Definition 3. A **mixed strategy** is a strategy which involves randomizing among some set of pure strategies, with some kind of weighted probability.

Did anyone play a pure strategy in the first 10 rounds?

Example 4. In the matching game, mixed strategies include 50/50, 80/20, 100/0, etc. A pure strategy is also a mixed strategy, just choosing one strategy 100% of the time. In the matching game, a mixed strategy is uniquely specified by a number $p_1 \in [0, 1]$, giving the probability of choosing 1.

Definition 5. If players play some mixed stragies, the **average expected payoff** to a player is the weighted average of their possible payoffs, where payoffs are weighted by the probability of reaching that payoff. If A plays strategy p_1 and B plays strategy q_1, the average expected payoff to A is denoted

$$v_A(p_1, q_1)$$

Then the average expected payoff to B is $v_B(p_1, q_1) = -v_1(p_1, q_1)$, so we just focus on v_1.

Example 6. Suppose A plays the strategy $p_1 = 0.4$ (playing 1 40% of the time) and B plays the strategy $q_1 = 0.3$. Then the outcome (1, 1) happens 28% of the time. Below is a version of the game matrix with payoffs replaced by the probability of that outcome.

		0.3 1	0.7 2
0.4	1	0.12	0.28
0.6	2	0.18	0.42

With these probabilities, the average expected payoff to A is

$$(0.12)(1) + (0.28)(-1) + (0.18)(-1) + (0.42)(1) = 0.08$$

In other words, with these strategies, A will *on average* earn 0.08 points per round (and B will *on average* lose 0.08 points). That is, if they play 100 rounds, A will probably be around +8 and B will be around −8 by the end.

2.1 Play the game - 2nd time

Activity time! Now that you've had a bit of time to digest and think about the game, try again. Return to your neighbor and play 10 more rounds, making sure you still record the outcomes on paper.

- Who in the room got the most number of points in these 10 rounds?
- Did you approach the game differently this time?

Definition 7. A **Nash equilibrium** of a game is a set of (mixed or pure) strategies for each player which is "regret-free" for all players. In other words, neither player can increase their value function by unilaterally changing strategy (even knowing the opponent's strategy). More formally, (p, q) is an equilibrium if

$$\begin{aligned} v_A(p, q) &\geq v_A(p', q) \qquad \text{for any alternative strategy } p' \text{ for A} \\ v_B(p, q) &\geq v_B(p, q') \qquad \text{for any alternative strategy } q' \text{ for B} \end{aligned}$$

In other words, in equilibrium, both players are maximizing their value function, holding strategies of other players fixed.

Proposition 8. *The only Nash equilibrium in matching pennies is* $(p_H, q_H) = \left(\frac{1}{2}, \frac{1}{2}\right)$. *The average expected payoff to each player at this equilibrium is zero.*

Proof. The second sentence requires no proof, and in fact a more general statement is true: if either player plays 1 randomly half the time, the average expected payoff for both players is zero. For example, if A plays $\frac{1}{2}$ and B plays strategy q,

$$v_A\left(\frac{1}{2}, q\right) = \frac{1}{2}q(1) - \frac{1}{2}(1-q)(-1) + \frac{1}{2}q(-1) + \frac{1}{2}(1-q)(1) = 0$$

If B plays 1 more than half the time, then A can exploit this by playing 1 all the time, in which case $v_A > 0$ and $v_B < 0$. This cannot be an equilibrium, because B has regret – they prefer to revert to playing 1 half the time.

Similarly, if B plays 2 more than half the time, A can exploit this by playing 2 all of the time, getting $v_A > 0$ and $v_B < 0$. This also cannot be an equilibrium, since B is not regret-free.

By the same reasoning, A cannot player 1 or 2 more than half of the time in equilibrium, because B exploits either by playing the opposite move 100% of the time. □

In summary, in the matching game, the best strategy is to be as unpredictable as possible and evenly split your moves between the two possibilities. Any predictability will cost you.

Example 9. Geometrically, a Nash equilibrium is just a saddle point of a surface. We can graph the value function

$$v_A(p, q) = pq - p(1-q) - q(1-p) + (1-p)(1-q) = 4pq - 2p - 2q + 1$$

on the square $[0, 1] \times [0, 1]$. An equilibrium is a point where this is maximized in the p-direction, and minimized in the q-direction, i.e. a saddle point. Below are some views of this surface, focusing on the saddle at $\left(\frac{1}{2}, \frac{1}{2}, 0\right)$. The height of the surface at point (x, y) is the average expected payoff to A when players play strategies (x, y).

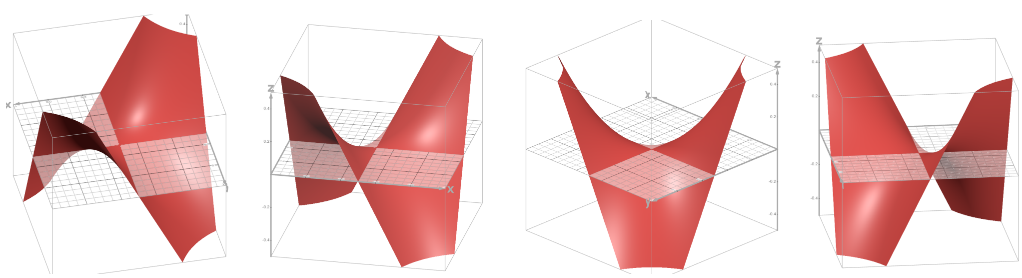

3 The big match

Now we will throw in a wrinkle: If A ever plays 2, it is a "big round" and both players will be locked into playing their move from that round in all future rounds. In the game matrix, we denote this by adding an asterisk to those entries.

	1	2
1	1	-1
2	-1^*	1^*

This version of the game is called "the big match," so named for the critical round in which A plays 2. The true big match game is played with no limit on the number of rounds. Because of this unboundedness, the sum of accumulated payoffs can diverge to $\pm\infty$, so instead we say that the payoff to a player is their average payoff per round.

There are two ways the game can go: A can play 2 at some point, or A can play 1 forever. If A plays 2 at some point, then their payoff is 1 or -1, depending on the results of the "big" round. The earlier rounds don't really matter, since we're averaging the values of an eventually constant sequence.

Round	1	2	3	$\cdots$	Last 1	First 2	2 forever	$\cdots$	
A	1	1	1	$\cdots$	1	**2**	*2*	*2*	$\cdots$
B	?	?	?	$\cdots$	?	**X**	X	X	$\cdots$

In the table above, if $X = 2$ the the payoffs are $(1, -1)$ and if $X = 1$ the payoffs are $(-1, 1)$.

On the other hand, A might never play 2, in which case their payoff is the average over the infinite sequence of 1's and 2's chosen by B.

Round	1	2	3	$\cdots$		
A	1	1	1	$\cdots$	1	$\cdots$
B	?	?	?	$\cdots$	?	$\cdots$

For example, B might choose 1 30% of the time and 2 70% of the time, in which case the payoffs would be $(-0.4, 0.4)$.

Technical note: In general, the limit of the average payoff over the first n rounds as $n \to \infty$ might fail to converge if B does something more complicated, so in general the payoff is defined as the lim sup of the average over the first n rounds.

$$v_A(\text{overall}) = \limsup_{n\to\infty} \frac{1}{n} \sum_{k=1}^{n} v_A(\text{round k})$$

Example 10. Here is an example game. In the first two rounds, A plays 1, so they keep choosing the next round. In round 3, A plays 2, making round 3 the "big" round. Since B played 2 in that critical round, they'll have to play 2 forever after. A will also play 2 in future rounds.

Round	1	2	3	4	5	$\cdots$
A	1	1	**2**	*2*	*2*	$\cdots$
B	2	1	**2**	*2*	*2*	$\cdots$

The resulting payoff from 10 rounds would then be $+0.9$ for A. If the game was infinite, it would be $+1$. If in round 3, B had played 1 instead, then the game would have gone

Round	1	2	3	4	5	$\cdots$
A	1	1	**2**	*2*	*2*	$\cdots$
B	2	1	**1**	*1*	*1*	$\cdots$

If the game was 10 rounds, the resulting payoff would be -0.9 for A, or -1 in an infinite game.

3.1 Play the big match - 2 times

Activity time! Play 20 rounds of big match with your partner, recording the results. Then switch roles, and play a new game of 20 rounds.

- Who in the room won the most number in a set of 20 rounds?
- Did it feel different than the regular matching game?
- How many games had a big round?
  Of those, how many were (2,1) and how many were (2,2)?
- Did A play 1 20 times in a row in any games?
- When playing as A, what were you thinking about?
- When playing as B, what were you thinking about?

3.2 Stationary and history-based strategies

Finally, let's analyze the big match formally. We need one last bit of terminology.

Definition 11. In a repeated game, a **stationary strategy** is a strategy which doesn't depend on time. In other words, a player playing a stationary strategy will implement the same strategy (possibly mixed) in each round, ignoring the number of rounds played or history of moves. In contrast, a **history-based strategy** takes the history into account.

So far, we have only talked about stationary strategies. "Play 1 half the time" is a stationary strategy. We need history-based strategies to properly analyze the big match. Some examples of history-based strategies include:

- Play the opposite of my opponent's move in the previous round.
- Play 1 40% of the time, then 30% of the time, then 60% of the time, then repeat this.
- If my opponent has played 1 x times in all n rounds played, I will play 1 with probability x/n.

In general, non-stationary strategies can be extremely convoluted.

3.3 Analysis from B's perspective

To analyze the big match from the perspective of B, we just need a stationary strategy.

Proposition 12 (Blackwell and Ferguson, 1968)**.** *In the big match, if B plays the stationary strategy $q_1 = \frac{1}{2}$, the average expected payoff is 0 regardless of A's strategy.*

Proof. If A ever plays 2 (a "big round"), then in that round we have $\frac{1}{2}$ chance of 1 and $\frac{1}{2}$ chance of 2 from B, so the payoff will be ± 1 depending on the big round. Basically, players have agreed to bet everything on a single round of the original matching game.

If A only plays 1, then on average half the rounds will be a match and half the rounds a mismatch, resulting in an average payoff of zero. Under the hood, this assertion relies on the strong law of large numbers. □

3.4 Analysis from A's perspective

Proposition 13 (Lester Dubins)**.** *There is no strategy for A which guarantees exactly zero.*

Proof. Suppose A has a strategy S that guarantees an average expected payoff zero against any sequence of moves from B. This strategy for A cannot be 1 in every round, since B can exploit this by playing 2 in every round, resulting in a payoff of -1.

Let m be the first round in which A assigns probability $\varepsilon > 0$ to playing 2. That is, in round m, A plays $p_2 = \varepsilon$.

Round	1	$\cdots$	$m-1$	m	$\cdots$
Strategy S	1	$\cdots$	1	$p_2 = \varepsilon > 0$	$\cdots$

I claim this strategy doesn't guarantee zero for A against every strategy for B, because B has a strategy which gives A a negative average expected payoff. That strategy is: B plays 2 for rounds 1 through $m-1$, 1 in round m, then $q_1 = \frac{1}{2}$ in future rounds.

Round	1	$\cdots$	$m-1$	m	$m+1$	$\cdots$
A	1	$\cdots$	1	$p_2 = \varepsilon > 0$	?	$\cdots$
B	2	$\cdots$	2	1	$q_1 = \frac{1}{2}$	$\cdots$

If A ends up playing 2 in round m, then they lose the "big round" and end up with payoff of -1. If they play 1 in round m, they end up with an average expected payoff of zero from the infinite tail of rounds $m+1$ and beyond. Overall, A's average expected payoff from both players implementing these strategies is $-\varepsilon < 0$. $\square$

Next we need to understand the history-based strategy for A that Blackwell and Ferguson describe.

Definition 14. Fix a non-negative integer N (think large). As A, after n rounds of us playing 1, before playing round number $n+1$, we count the number of times B has played 1 and 2. We calculate

$$E_1 = \text{excess 1's} = (\#1\text{'s by B}) - (\#2\text{'s by B})$$

Then in round $n+1$, we play 2 with probability

$$p_2(E_1) = \frac{1}{(E_1 + N + 1)^2} = \frac{1}{(\#1 - \#2 + N + 1)^2}$$

We call this strategy $S(N)$.

The intuition of strategy $S(N)$ is that we as A are more likely to play 2 if B has played 2 a lot of times in the past, and more likely to play 1 if B has played 1 a lot in the past.

Note that the denominator is never zero, because if E_1 ever reaches $-N$, then $p_2 = 1$, guaranteeing a big round. So there will not be a decision to make in a round where $E_1 = -N - 1$.

Example 15. Suppose $N = 9$. In round 1, there are no previous rounds so $E_1 = 0$. So in round 1 as A playing $S(N)$, we play 2 1% of the time.

$$p_2(0) = \frac{1}{(0+9+1)^2} = \frac{1}{10^2} = \frac{1}{100}$$

Supose that after 49 rounds, and B has played 29 1's and 20 2's. Then $E_1 = 9$, so in round 50 we play 2 even less often.

$$p_2(9) = \frac{1}{(9+9+1)^2} = \frac{1}{19^2} = \frac{1}{361}$$

On the other hand, if B has played 20 1's and 29 2's, then $E_1 = -9$ so in round 50 we would play 2 with probability 1.

$$p_2(-9) = \frac{1}{(-9+9+1)^2} = \frac{1}{1^2} = 1$$

Theorem 16 (Blackwell and Ferguson, 1968)**.** *If A plays strategy $S(N)$, then their average expected payoff (regardless of what B does) is at least $\frac{-1}{2N+1}$. So for $\varepsilon > 0$, A can choose N sufficiently large that $v_A \geq -\varepsilon$.*

In technical terms, for any $\varepsilon > 0$, A has an ε-optimal strategy.

Proof. The proof is about a page long, but involves some intricate manipulation of inequalities. The key idea is: supposing A plays $S(N)$,

- If B plays a sequence of moves in which E_1 eventually reaches $-N$ (i.e B plays a lot more 2's than 1's), then the game is guarateed to end in a finite number of rounds, and the inequality follows from an induction argument.
- If B plays a sequence of moves in which E_1 never reaches $-N$ (i.e. B never plays more than N more 2's than 1's), then proportion of 1's and 2's is never too far from $\frac{1}{2}$, so the average expected payoff can't get too far from zero.

□

Corollary 17. *Neither player can guarantee bettter than an average expected payoff of zero in the big match.*

Proof. B can guarantee themselves zero by playing each move half the time. If B could guarantee themselves a strictly positive number $\varepsilon > 0$, this would mean A gets no better than $-\varepsilon$, but we know A can do at least this well for themselves, so B cannot do this. □

If you are interested in learning more, the broad area of mathematics is known as stochastic game theory. Stochastic games is a huge area of active research, including things like Markov Decision Processes (1 player stochastic games), stochastic games with incomplete information. The big match is something known as a repeated game with absorbing states.

4 Stochastic games

In the years following Blackwell and Ferguson's analysis of the big match, more general results were gradually proven, culminating in the major capstone result of Mertens and Neyman in 1979.