A Paradox of Blown Leads: Rethinking Win Probability in Team Sports

Jonathan Pipping and Abraham J. Wyner

The Wharton School, University of Pennsylvania

NESSIS 2025

Case Study: Super Bowl LVII

- At the start of the 2nd half, the Chiefs trailed the Eagles 14–24 and faced a 3rd and 1 at their own 34
- At that point, the Eagles' projected win probability was 78.4%
- On the next play, Jerick McKinnon converted with a 14-yard run, and the Chiefs would go on to score, eventually winning 38–35.

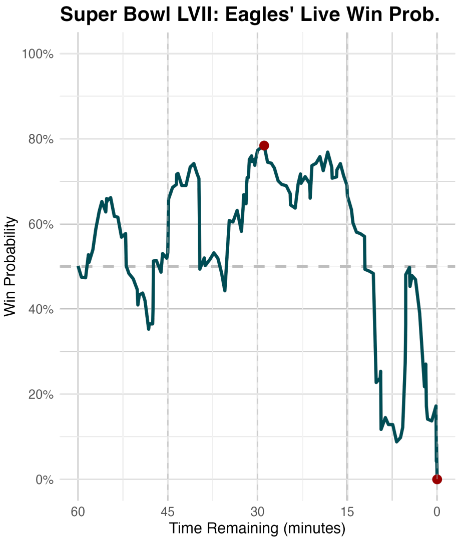

Data via nflfastR

How Often Do Blown Leads Occur?

- We'd like to understand how rare it is to blow a lead of this magnitude!

How Often Do Blown Leads Occur?

- We'd like to understand how rare it is to blow a lead of this magnitude!
- Intuitively, most people would assume it happens very rarely (about 21.6% of the time), but maybe this is wrong.

How Often Do Blown Leads Occur?

- We'd like to understand how rare it is to blow a lead of this magnitude!
- Intuitively, most people would assume it happens very rarely (about 21.6% of the time), but maybe this is wrong.
- To investigate this, we formalize the question mathematically.

Mathematical Framework

- The win probability of a team i at time t is a function of team strength S_i and game state $(\mathcal{G}, t)$. Symbolically,

$$W_i(t) = f(S_i, \mathcal{G}, t), \quad t \in [0, T]$$

Mathematical Framework

- The win probability of a team i at time t is a function of team strength S_i and game state $(\mathcal{G}, t)$. Symbolically,

$$W_i(t) = f(S_i, \mathcal{G}, t), \quad t \in [0, T]$$

- In a two-team game between teams A and B, at any time t we have

$$W_A(t) + W_B(t) = 1$$

Mathematical Framework

- The win probability of a team i at time t is a function of team strength S_i and game state $(\mathcal{G}, t)$. Symbolically,

$$W_i(t) = f(S_i, \mathcal{G}, t), \quad t \in [0, T]$$

- In a two-team game between teams A and B, at any time t we have

$$W_A(t) + W_B(t) = 1$$

- For the **eventual loser** ℓ, we know that $W_\ell(T) = 0$. So the quantity of interest (the maximum win prob. of the losing team) is given by the random variable

$$M_\ell = \max_t W_\ell(t)$$

Mathematical Framework

- The win probability of a team i at time t is a function of team strength S_i and game state $(\mathcal{G}, t)$. Symbolically,

$$W_i(t) = f(S_i, \mathcal{G}, t), \quad t \in [0, T]$$

- In a two-team game between teams A and B, at any time t we have

$$W_A(t) + W_B(t) = 1$$

- For the **eventual loser** ℓ, we know that $W_\ell(T) = 0$. So the quantity of interest (the maximum win prob. of the losing team) is given by the random variable

$$M_\ell = \max_t W_\ell(t)$$

- We will investigate the distribution of M_ℓ as a function of the team strengths S_A and S_B.

Simulation Setup

- Estimating the win probability function f is difficult in practice.

Simulation Setup

- Estimating the win probability function f is difficult in practice.
    - Game state $(\mathcal{G}, t)$ is multi-dimensional highly non-linear.

Simulation Setup

- Estimating the win probability function f is difficult in practice.
  - Game state $(\mathcal{G}, t)$ is multi-dimensional highly non-linear.
  - Team strengths S_A and S_B are non-trivial to estimate.

Simulation Setup

- Estimating the win probability function f is difficult in practice.
  - Game state $(\mathcal{G}, t)$ is multi-dimensional highly non-linear.
  - Team strengths S_A and S_B are non-trivial to estimate.
- To simplify the problem, we specify a simple model to allow exact calculation of in-game win probabilities.

Simulation Setup

- Estimating the win probability function f is difficult in practice.
  - Game state $(\mathcal{G}, t)$ is multi-dimensional highly non-linear.
  - Team strengths S_A and S_B are non-trivial to estimate.
- To simplify the problem, we specify a simple model to allow exact calculation of in-game win probabilities.
  - Each team has an equal number of possessions N.

Simulation Setup

- Estimating the win probability function f is difficult in practice.
  - Game state $(\mathcal{G}, t)$ is multi-dimensional highly non-linear.
  - Team strengths S_A and S_B are non-trivial to estimate.
- To simplify the problem, we specify a simple model to allow exact calculation of in-game win probabilities.
  - Each team has an equal number of possessions N.
  - Team strengths p_A and p_B are fixed and correspond to each team's probability of scoring on a given possession. Ties are broken by a weighted coin flip (overtime).

Simulation Setup

- Estimating the win probability function f is difficult in practice.
  - Game state $(\mathcal{G}, t)$ is multi-dimensional highly non-linear.
  - Team strengths S_A and S_B are non-trivial to estimate.
- To simplify the problem, we specify a simple model to allow exact calculation of in-game win probabilities.
  - Each team has an equal number of possessions N.
  - Team strengths p_A and p_B are fixed and correspond to each team's probability of scoring on a given possession. Ties are broken by a weighted coin flip (overtime).
- So the win prob. for each possible game state $(\mathcal{G}, t)$ is given by

$$\begin{aligned}\mathrm{WP}_A(t) = \mathbb{P}\left[\mathrm{Binom}(n_t, p_A) + \mathrm{score}_A(t) > \mathrm{Binom}(n_t, p_B) + \mathrm{score}_B(t)\right] \\ + \frac{p_A}{p_A + p_B}\mathbb{P}(\text{A and B tie})\end{aligned}$$

where $n_t = N - t$ is the number of possessions remaining.

Simulation Setup

- For each set of parameters N, p_A, p_B, we do the following:

Simulation Setup

- For each set of parameters N, p_A, p_B, we do the following:
  1. Calculate the win prob. for each possible game state $(\mathcal{G}, t)$

Simulation Setup

- For each set of parameters N, p_A, p_B, we do the following:
  1. Calculate the win prob. for each possible game state $(\mathcal{G}, t)$
  2. Simulate 10,000 games between two teams.

Simulation Setup

- For each set of parameters N, p_A, p_B, we do the following:
  1. Calculate the win prob. for each possible game state $(\mathcal{G}, t)$
  2. Simulate 10,000 games between two teams.
  3. Extract the maximum win prob. attained by the eventual loser (M_ℓ).

Simulation Setup

- For each set of parameters N, p_A, p_B, we do the following:
  1. Calculate the win prob. for each possible game state $(\mathcal{G}, t)$
  2. Simulate 10,000 games between two teams.
  3. Extract the maximum win prob. attained by the eventual loser (M_ℓ).
- We consider the distribution of M_ℓ as N, p_A, and p_B vary.

Simulation Setup

- For each set of parameters N, p_A, p_B, we do the following:
  1. Calculate the win prob. for each possible game state $(\mathcal{G}, t)$
  2. Simulate 10,000 games between two teams.
  3. Extract the maximum win prob. attained by the eventual loser (M_ℓ).
- We consider the distribution of M_ℓ as N, p_A, and p_B vary.
  - Case 1: $p_A = p_B$, f is symmetric across teams

Simulation Setup

- For each set of parameters N, p_A, p_B, we do the following:
  1. Calculate the win prob. for each possible game state $(\mathcal{G}, t)$
  2. Simulate 10,000 games between two teams.
  3. Extract the maximum win prob. attained by the eventual loser (M_ℓ).
- We consider the distribution of M_ℓ as N, p_A, and p_B vary.
  - Case 1: $p_A = p_B$, f is symmetric across teams
  - Case 2: $p_A \neq p_B$, f is asymmetric across teams

Symmetric Case: $p_A = p_B$

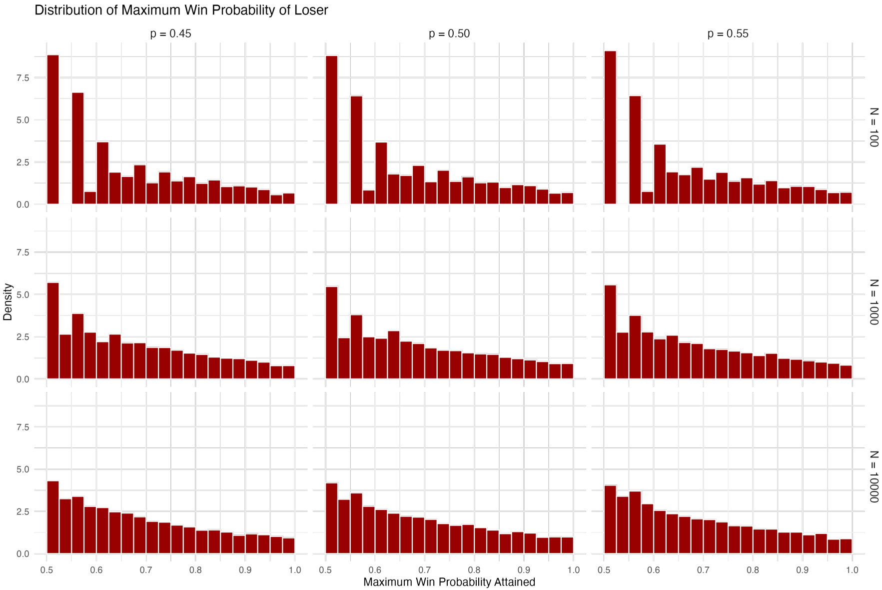

Symmetric Case: Threshold Plot

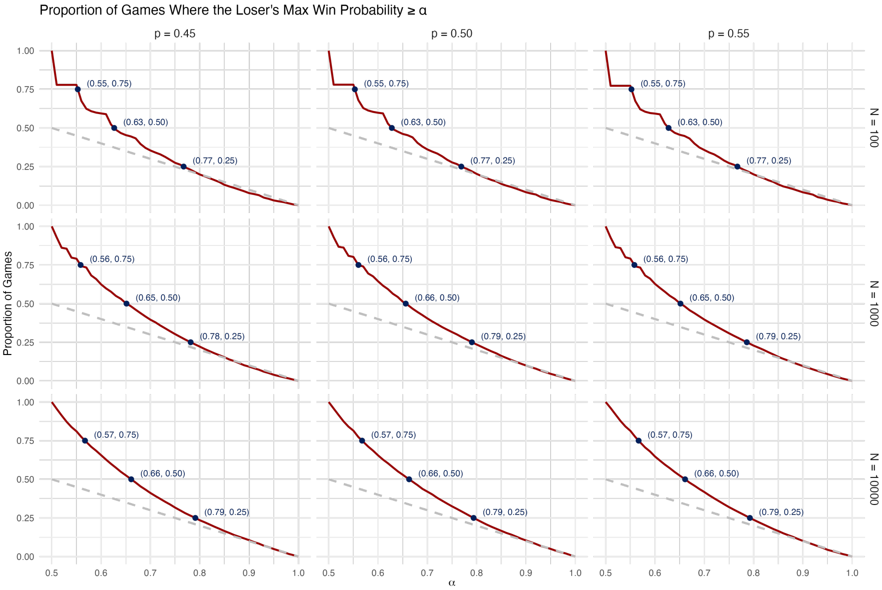

Symmetric Case: Takeaways

- When $p_A = p_B$, the support of M_ℓ is $[0.5, 1)$.

Symmetric Case: Takeaways

- When $p_A = p_B$, the support of M_ℓ is $[0.5, 1)$.
- Holding N constant, the distribution of M_ℓ is identical for all values of $p_A = p_B$.

Symmetric Case: Takeaways

- When $p_A = p_B$, the support of M_ℓ is $[0.5, 1)$.
- Holding N constant, the distribution of M_ℓ is identical for all values of $p_A = p_B$.
  - Implies that the distribution of M_ℓ may depend only on some measure of the difference in probabilities (e.g., $|p_A - p_B|$, $\frac{p_A}{p_B}$, $\ldots$).

Symmetric Case: Takeaways

- When $p_A = p_B$, the support of M_ℓ is $[0.5, 1)$.
- Holding N constant, the distribution of M_ℓ is identical for all values of $p_A = p_B$.
  - Implies that the distribution of M_ℓ may depend only on some measure of the difference in probabilities (e.g., $|p_A - p_B|$, $\frac{p_A}{p_B}$, $\ldots$).
- Holding $p_A = p_B$ constant, the distribution of M_ℓ becomes less discrete and approaches a continuous limit as N increases.

Symmetric Case: Takeaways

- When $p_A = p_B$, the support of M_ℓ is $[0.5, 1)$.
- Holding N constant, the distribution of M_ℓ is identical for all values of $p_A = p_B$.
  - Implies that the distribution of M_ℓ may depend only on some measure of the difference in probabilities (e.g., $|p_A - p_B|$, $\frac{p_A}{p_B}$, ...).
- Holding $p_A = p_B$ constant, the distribution of M_ℓ becomes less discrete and approaches a continuous limit as N increases.
- In **about half** of all games, the losing team attains a win probability of at least **66% or more**.

Asymmetric Case: $p_A \neq p_B$

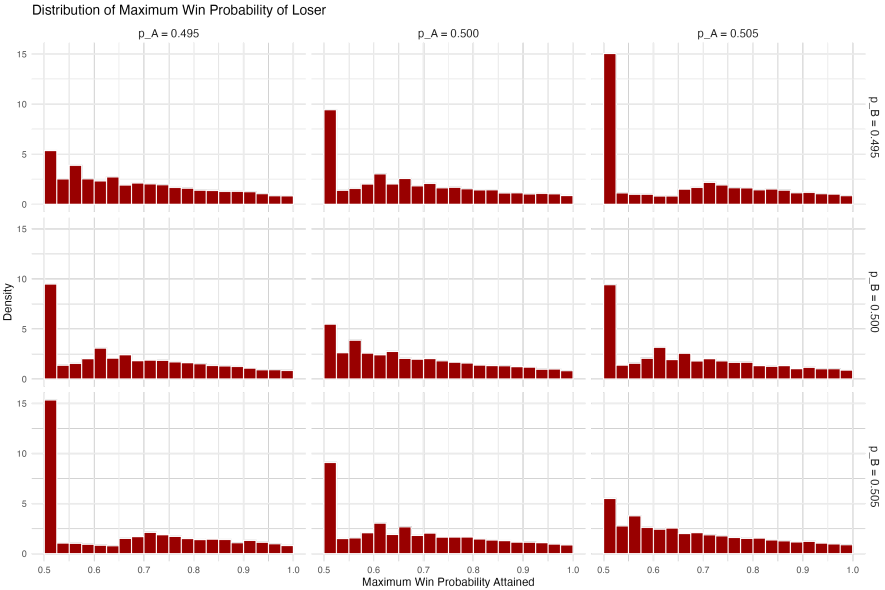

Asymmetric Case: Threshold Plot

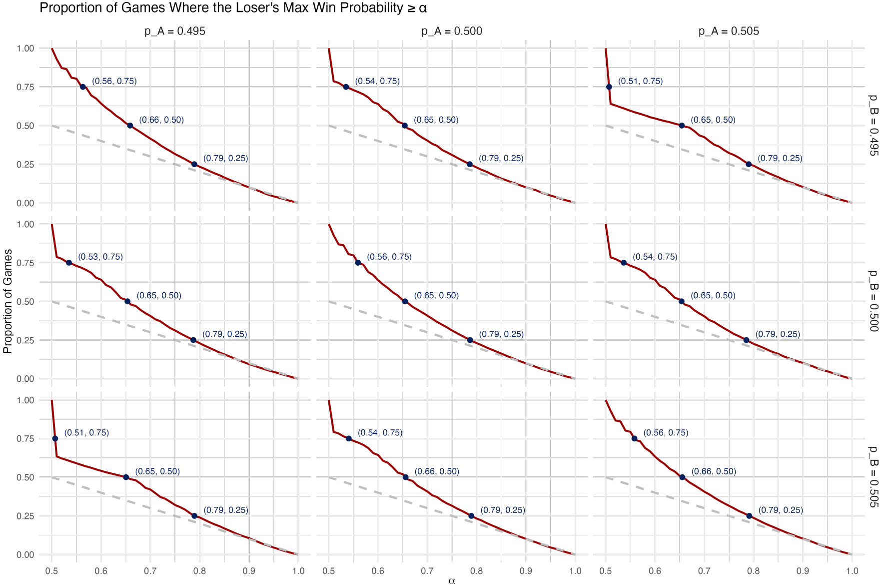

Asymmetric Case: Takeaways

- When $p_A \neq p_B$, the support of M_ℓ is $[\min\{p_A, p_B\}, 1)$.

Asymmetric Case: Takeaways

- When $p_A \neq p_B$, the support of M_ℓ is $[\min\{p_A, p_B\}, 1)$.
- Holding N constant, the distribution of M_ℓ is constant whenever $|p_A - p_B|$ is constant.

Asymmetric Case: Takeaways

- When $p_A \neq p_B$, the support of M_ℓ is $[\min\{p_A, p_B\}, 1)$.
- Holding N constant, the distribution of M_ℓ is constant whenever $|p_A - p_B|$ is constant.
  - Implies that the distribution of M_ℓ depends only on the absolute difference in probabilities $|p_A - p_B|$.

Asymmetric Case: Takeaways

- When $p_A \neq p_B$, the support of M_ℓ is $[\min\{p_A, p_B\}, 1)$.
- Holding N constant, the distribution of M_ℓ is constant whenever $|p_A - p_B|$ is constant.
  - Implies that the distribution of M_ℓ depends only on the absolute difference in probabilities $|p_A - p_B|$.
- Holding N constant, the distribution of M_ℓ becomes increasingly right-skewed over the support as $|p_A - p_B|$ increases.

Asymmetric Case: Takeaways

- When $p_A \neq p_B$, the support of M_ℓ is $[\min\{p_A, p_B\}, 1)$.
- Holding N constant, the distribution of M_ℓ is constant whenever $|p_A - p_B|$ is constant.
  - Implies that the distribution of M_ℓ depends only on the absolute difference in probabilities $|p_A - p_B|$.
- Holding N constant, the distribution of M_ℓ becomes increasingly right-skewed over the support as $|p_A - p_B|$ increases.
- Larger strength differentials decrease the proportion of games where the losing team attains a high win probability.

Towards a Closed-Form Solution

- We hope to derive a closed-form solution for the distribution of M_ℓ in our simplified model. To do this, we define the following:

Towards a Closed-Form Solution

- We hope to derive a closed-form solution for the distribution of M_ℓ in our simplified model. To do this, we define the following:
  - $X_k = \text{score}_A(k) - \text{score}_B(k)$ is the score differential at time k.

Towards a Closed-Form Solution

- We hope to derive a closed-form solution for the distribution of M_ℓ in our simplified model. To do this, we define the following:
  - $X_k = \text{score}_A(k) - \text{score}_B(k)$ is the score differential at time k.
  - $\mathcal{F}_k = \sigma(X_0, \ldots, X_k)$ is the set of information available at time k.

Towards a Closed-Form Solution

- We hope to derive a closed-form solution for the distribution of M_ℓ in our simplified model. To do this, we define the following:
  - $X_k = \text{score}_A(k) - \text{score}_B(k)$ is the score differential at time k.
  - $\mathcal{F}_k = \sigma(X_0, \ldots, X_k)$ is the set of information available at time k.
  - $Y = \mathbf{1}_{\{X_N \geq 0\}}$ is the event that team A wins the game.

Towards a Closed-Form Solution

- We hope to derive a closed-form solution for the distribution of M_ℓ in our simplified model. To do this, we define the following:
  - $X_k = \text{score}_A(k) - \text{score}_B(k)$ is the score differential at time k.
  - $\mathcal{F}_k = \sigma(X_0, \dots, X_k)$ is the set of information available at time k.
  - $Y = \mathbf{1}_{\{X_N \geq 0\}}$ is the event that team A wins the game.
- Then we can define Team A's win probability at each time step k as a Doob martingale:

$$p_k = \mathbb{E}[Y \mid \mathcal{F}_k], \quad k = 0, \dots, N$$

Towards a Closed-Form Solution

- We hope to derive a closed-form solution for the distribution of M_ℓ in our simplified model. To do this, we define the following:
  - $X_k = \text{score}_A(k) - \text{score}_B(k)$ is the score differential at time k.
  - $\mathcal{F}_k = \sigma(X_0, \ldots, X_k)$ is the set of information available at time k.
  - $Y = \mathbf{1}_{\{X_N \geq 0\}}$ is the event that team A wins the game.
- Then we can define Team A's win probability at each time step k as a Doob martingale:

$$p_k = \mathbb{E}[Y \mid \mathcal{F}_k], \quad k = 0, \ldots, N$$

- From here, define the following:

Towards a Closed-Form Solution

- We hope to derive a closed-form solution for the distribution of M_ℓ in our simplified model. To do this, we define the following:
  - $X_k = \text{score}_A(k) - \text{score}_B(k)$ is the score differential at time k.
  - $\mathcal{F}_k = \sigma(X_0, \ldots, X_k)$ is the set of information available at time k.
  - $Y = \mathbf{1}_{\{X_N \geq 0\}}$ is the event that team A wins the game.
- Then we can define Team A's win probability at each time step k as a Doob martingale:

$$p_k = \mathbb{E}[Y \mid \mathcal{F}_k], \quad k = 0, \ldots, N$$

- From here, define the following:
  - $M_A = \max_{0 \leq k \leq N} p_k$ is the maximum win probability team A attains.

Towards a Closed-Form Solution

- We hope to derive a closed-form solution for the distribution of M_ℓ in our simplified model. To do this, we define the following:
  - $X_k = \text{score}_A(k) - \text{score}_B(k)$ is the score differential at time k.
  - $\mathcal{F}_k = \sigma(X_0, \ldots, X_k)$ is the set of information available at time k.
  - $Y = \mathbf{1}_{\{X_N \geq 0\}}$ is the event that team A wins the game.
- Then we can define Team A's win probability at each time step k as a Doob martingale:

$$p_k = \mathbb{E}[Y \mid \mathcal{F}_k], \quad k = 0, \ldots, N$$

- From here, define the following:
  - $M_A = \max_{0 \leq k \leq N} p_k$ is the maximum win probability team A attains.
  - $\tau_x = \min\{k \leq N : p_k \geq x\}$ is the first time p_k exceeds x.

Towards a Closed-Form Solution

- We hope to derive a closed-form solution for the distribution of M_ℓ in our simplified model. To do this, we define the following:
  - $X_k = \text{score}_A(k) - \text{score}_B(k)$ is the score differential at time k.
  - $\mathcal{F}_k = \sigma(X_0, \ldots, X_k)$ is the set of information available at time k.
  - $Y = \mathbf{1}_{\{X_N \geq 0\}}$ is the event that team A wins the game.
- Then we can define Team A's win probability at each time step k as a Doob martingale:

$$p_k = \mathbb{E}[Y \mid \mathcal{F}_k], \quad k = 0, \ldots, N$$

- From here, define the following:
  - $M_A = \max_{0 \leq k \leq N} p_k$ is the maximum win probability team A attains.
  - $\tau_x = \min\{k \leq N : p_k \geq x\}$ is the first time p_k exceeds x.
- Since τ_x is a stopping time, we invoke the optional stopping theorem to derive the distribution of M_A.

Closed-Form Distributions: Discrete Case

- **Theorem 1:** The distribution of M_A satisfies:

$$F_{M_A}(x) \geq 1 - \frac{p_0}{x}, \quad x \in [p_0, 1)$$

with equality when $\mathbb{P}(\tau_x = N) = 0$ (continuous limits).[1]

[1] Note that M_A has a point mass of weight p_0 at 1.

Closed-Form Distributions: Discrete Case

- **Theorem 1:** The distribution of M_A satisfies:

$$F_{M_A}(x) \geq 1 - \frac{p_0}{x}, \quad x \in [p_0, 1)$$

  with equality when $\mathbb{P}(\tau_x = N) = 0$ (continuous limits).[1]

- **Theorem 2:** The conditional distribution of M_A given that team A loses satisfies:

$$F_{M_A|Y=0}(x) \geq 1 - \left(\frac{p_0}{1-p_0}\right) \cdot \left(\frac{1-x}{x}\right), \quad x \in [p_0, 1)$$

  with equality when $\mathbb{P}(\tau_x = N) = 0$ (continuous limits)

[1] Note that M_A has a point mass of weight p_0 at 1.

Closed-Form Distributions: Discrete Case

- What about the distribution for team B?

Closed-Form Distributions: Discrete Case

- What about the distribution for team B?
  - Recall that team B's path is just a reflection of team A's!

Closed-Form Distributions: Discrete Case

- What about the distribution for team B?
  - Recall that team B's path is just a reflection of team A's!
  - So M_B has a similar form to M_A, but with p_0 replaced by $1 - p_0$.

Closed-Form Distributions: Discrete Case

- What about the distribution for team B?
  - Recall that team B's path is just a reflection of team A's!
  - So M_B has a similar form to M_A, but with p_0 replaced by $1 - p_0$.
- So then what about the distribution of the **eventual** loser M_ℓ?

Closed-Form Distributions: Discrete Case

- What about the distribution for team B?
  - Recall that team B's path is just a reflection of team A's!
  - So M_B has a similar form to M_A, but with p_0 replaced by $1 - p_0$.
- So then what about the distribution of the **eventual** loser M_ℓ?
  - This is a **mixture** of the distributions $(M_A | A \text{ loses})$ and $(M_B | B \text{ loses})$ with weights $(1 - p_0)$ and p_0, respectively.

Closed-Form Distributions: Discrete Case

- What about the distribution for team B?
  - Recall that team B's path is just a reflection of team A's!
  - So M_B has a similar form to M_A, but with p_0 replaced by $1 - p_0$.
- So then what about the distribution of the **eventual** loser M_ℓ?
  - This is a **mixture** of the distributions $(M_A | A \text{ loses})$ and $(M_B | B \text{ loses})$ with weights $(1 - p_0)$ and p_0, respectively.
  - But wait! In the region $[\min(p_0, 1 - p_0), \max(p_0, 1 - p_0))$, the distribution of M_ℓ comes entirely from the underdog! So we must define the distribution of M_ℓ piecewise.

Closed-Form Distributions: Discrete Case

- What about the distribution for team B?
  - Recall that team B's path is just a reflection of team A's!
  - So M_B has a similar form to M_A, but with p_0 replaced by $1 - p_0$.
- So then what about the distribution of the **eventual** loser M_ℓ?
  - This is a **mixture** of the distributions $(M_A | A \text{ loses})$ and $(M_B | B \text{ loses})$ with weights $(1 - p_0)$ and p_0, respectively.
  - But wait! In the region $[\min(p_0, 1 - p_0), \max(p_0, 1 - p_0))$, the distribution of M_ℓ comes entirely from the underdog! So we must define the distribution of M_ℓ piecewise.
- **Theorem 3:** Since team labels are arbitrary, let team A be the favorite ($p_0 \geq 0.5$). Then the distribution of M_ℓ satisfies:

$$F_{M_\ell}(x) \geq \begin{cases} 1 - \dfrac{1 - p_0}{x} & \text{if } x \in [1 - p_0, p_0) \\ 2 - \dfrac{1}{x} & \text{if } x \in [p_0, 1) \end{cases}$$

with equality when $\mathbb{P}(\tau_x = N) = 0$ (continuous limits)

From Inequalities to Equality: The Brownian Limit

- The inequalities in Theorems 1-3 are tight, but what if we want something stronger, like a density equality? What's the limitation?

From Inequalities to Equality: The Brownian Limit

- The inequalities in Theorems 1-3 are tight, but what if we want something stronger, like a density equality? What's the limitation?
- In discrete time, there are **finitely-many** levels for WP_A and WP_B, so the process can "jump over" level x at the final step.

From Inequalities to Equality: The Brownian Limit

- The inequalities in Theorems 1-3 are tight, but what if we want something stronger, like a density equality? What's the limitation?
- In discrete time, there are **finitely-many** levels for WP_A and WP_B, so the process can "jump over" level x at the final step.
  - For example, consider the simplest case: $N = 1$ and $p_0 = 0.5$.

From Inequalities to Equality: The Brownian Limit

- The inequalities in Theorems 1-3 are tight, but what if we want something stronger, like a density equality? What's the limitation?
- In discrete time, there are **finitely-many** levels for WP_A and WP_B, so the process can "jump over" level x at the final step.
  - For example, consider the simplest case: $N = 1$ and $p_0 = 0.5$.
  - Then $WP_A(0) = 0.5$ and $WP_A(1) = \{0, 1\}$, jumping the rest!

From Inequalities to Equality: The Brownian Limit

- The inequalities in Theorems 1-3 are tight, but what if we want something stronger, like a density equality? What's the limitation?
- In discrete time, there are **finitely-many** levels for WP_A and WP_B, so the process can "jump over" level x at the final step.
  - For example, consider the simplest case: $N = 1$ and $p_0 = 0.5$.
  - Then $WP_A(0) = 0.5$ and $WP_A(1) = \{0, 1\}$, jumping the rest!
  - What we need to prevent this behavior is **continuous sample paths**, which prevents the process from jumping over levels.

From Inequalities to Equality: The Brownian Limit

- The inequalities in Theorems 1-3 are tight, but what if we want something stronger, like a density equality? What's the limitation?
- In discrete time, there are **finitely-many** levels for WP_A and WP_B, so the process can "jump over" level x at the final step.
  - For example, consider the simplest case: $N = 1$ and $p_0 = 0.5$.
  - Then $WP_A(0) = 0.5$ and $WP_A(1) = \{0, 1\}$, jumping the rest!
  - What we need to prevent this behavior is **continuous sample paths**, which prevents the process from jumping over levels.
- How do we get continuous sample paths? By letting $N \to \infty$.

From Inequalities to Equality: The Brownian Limit

- The inequalities in Theorems 1-3 are tight, but what if we want something stronger, like a density equality? What's the limitation?
- In discrete time, there are **finitely-many** levels for WP_A and WP_B, so the process can "jump over" level x at the final step.
  - For example, consider the simplest case: $N = 1$ and $p_0 = 0.5$.
  - Then $WP_A(0) = 0.5$ and $WP_A(1) = \{0, 1\}$, jumping the rest!
  - What we need to prevent this behavior is **continuous sample paths**, which prevents the process from jumping over levels.
- How do we get continuous sample paths? By letting $N \to \infty$.
  - In this limit, the process becomes continuous and we can derive exact closed-form expressions for the distributions of M_A, M_B, and M_ℓ.

From Inequalities to Equality: The Brownian Limit

- The inequalities in Theorems 1-3 are tight, but what if we want something stronger, like a density equality? What's the limitation?
- In discrete time, there are **finitely-many** levels for WP_A and WP_B, so the process can "jump over" level x at the final step.
  - For example, consider the simplest case: $N = 1$ and $p_0 = 0.5$.
  - Then $WP_A(0) = 0.5$ and $WP_A(1) = \{0, 1\}$, jumping the rest!
  - What we need to prevent this behavior is **continuous sample paths**, which prevents the process from jumping over levels.
- How do we get continuous sample paths? By letting $N \to \infty$.
  - In this limit, the process becomes continuous and we can derive exact closed-form expressions for the distributions of M_A, M_B, and M_ℓ.
  - In addition, **Donsker's Invariance Principle** allows us to approximate the scoring process X_k with a Brownian motion!

Towards a Continuous Closed-Form

- We hope to derive a continuous closed-form for the distribution of M_ℓ in our simplified model. Like before, we have:

Towards a Continuous Closed-Form

- We hope to derive a continuous closed-form for the distribution of M_ℓ in our simplified model. Like before, we have:
  - $X_k = \text{score}_A(k) - \text{score}_B(k)$ is the score differential at time k.

Towards a Continuous Closed-Form

- We hope to derive a continuous closed-form for the distribution of M_ℓ in our simplified model. Like before, we have:
  - $X_k = \text{score}_A(k) - \text{score}_B(k)$ is the score differential at time k.
  - $\mathcal{F}_k = \sigma(X_0, \ldots, X_k)$ is the set of information available at time k.

Towards a Continuous Closed-Form

- We hope to derive a continuous closed-form for the distribution of M_ℓ in our simplified model. Like before, we have:
  - $X_k = \text{score}_A(k) - \text{score}_B(k)$ is the score differential at time k.
  - $\mathcal{F}_k = \sigma(X_0, \ldots, X_k)$ is the set of information available at time k.
- Then define the following quantities:

Towards a Continuous Closed-Form

- We hope to derive a continuous closed-form for the distribution of M_ℓ in our simplified model. Like before, we have:
  - $X_k = \text{score}_A(k) - \text{score}_B(k)$ is the score differential at time k.
  - $\mathcal{F}_k = \sigma(X_0, \ldots, X_k)$ is the set of information available at time k.
- Then define the following quantities:
  - $\mu = p_A - p_B$ is the difference in scoring probabilities.

Towards a Continuous Closed-Form

- We hope to derive a continuous closed-form for the distribution of M_ℓ in our simplified model. Like before, we have:
  - $X_k = \text{score}_A(k) - \text{score}_B(k)$ is the score differential at time k.
  - $\mathcal{F}_k = \sigma(X_0, \ldots, X_k)$ is the set of information available at time k.
- Then define the following quantities:
  - $\mu = p_A - p_B$ is the difference in scoring probabilities.
  - $\sigma^2 = p_A(1 - p_A) + p_B(1 - p_B)$ is the variance of each step in the scoring process.

Towards a Continuous Closed-Form

- We hope to derive a continuous closed-form for the distribution of M_ℓ in our simplified model. Like before, we have:
  - $X_k = \text{score}_A(k) - \text{score}_B(k)$ is the score differential at time k.
  - $\mathcal{F}_k = \sigma(X_0, \ldots, X_k)$ is the set of information available at time k.
- Then define the following quantities:
  - $\mu = p_A - p_B$ is the difference in scoring probabilities.
  - $\sigma^2 = p_A(1 - p_A) + p_B(1 - p_B)$ is the variance of each step in the scoring process.
- Then by Donsker's Invariance Principle, we have that

$$\frac{X_{\lfloor Nt \rfloor}}{\sigma\sqrt{N}} \xrightarrow{d} B_t + \mu^* t, \quad \mu^* = \frac{\sqrt{N}\,\mu}{\sigma}, \ t \in [0, 1]$$

where B_t is standard Brownian motion.

Towards a Continuous Closed-Form

- From here, note the following:

Towards a Continuous Closed-Form

- From here, note the following:
  - The discrete terminal event $\{X_N \geq 0\}$ converges to $\{B_1 + \mu^* \geq 0\}$.

Towards a Continuous Closed-Form

- From here, note the following:
  - The discrete terminal event $\{X_N \geq 0\}$ converges to $\{B_1 + \mu^* \geq 0\}$.
  - So $Y = \mathbf{1}_{\{B_1+\mu^* \geq 0\}}$ is the event that team A wins the game.

Towards a Continuous Closed-Form

- From here, note the following:
  - The discrete terminal event $\{X_N \geq 0\}$ converges to $\{B_1 + \mu^* \geq 0\}$.
  - So $Y = \mathbf{1}_{\{B_1 + \mu^* \geq 0\}}$ is the event that team A wins the game.
  - $p_0 = \mathbb{P}(Y = 1) = \Phi(\mu^*)$.

Towards a Continuous Closed-Form

- From here, note the following:
  - The discrete terminal event $\{X_N \geq 0\}$ converges to $\{B_1 + \mu^* \geq 0\}$.
  - So $Y = \mathbf{1}_{\{B_1 + \mu^* \geq 0\}}$ is the event that team A wins the game.
  - $p_0 = \mathbb{P}(Y = 1) = \Phi(\mu^*)$.
- Now we can define Team A's win probability as a Doob martingale:

$$p_t = \mathbb{E}[Y \mid \mathcal{F}_t], \quad t \in [0, 1]$$

Towards a Continuous Closed-Form

- From here, note the following:
  - The discrete terminal event $\{X_N \geq 0\}$ converges to $\{B_1 + \mu^* \geq 0\}$.
  - So $Y = \mathbf{1}_{\{B_1+\mu^* \geq 0\}}$ is the event that team A wins the game.
  - $p_0 = \mathbb{P}(Y = 1) = \Phi(\mu^*)$.
- Now we can define Team A's win probability as a Doob martingale:

$$p_t = \mathbb{E}[Y \mid \mathcal{F}_t], \quad t \in [0, 1]$$

- Then we can define the following quantities:

Towards a Continuous Closed-Form

- From here, note the following:
  - The discrete terminal event $\{X_N \geq 0\}$ converges to $\{B_1 + \mu^* \geq 0\}$.
  - So $Y = \mathbf{1}_{\{B_1 + \mu^* \geq 0\}}$ is the event that team A wins the game.
  - $p_0 = \mathbb{P}(Y = 1) = \Phi(\mu^*)$.
- Now we can define Team A's win probability as a Doob martingale:

$$p_t = \mathbb{E}[Y \mid \mathcal{F}_t], \quad t \in [0, 1]$$

- Then we can define the following quantities:
  - $M = \sup_{0 \leq t \leq 1} p_t$ is the maximum win probability team A attains.

Towards a Continuous Closed-Form

- From here, note the following:
  - The discrete terminal event $\{X_N \geq 0\}$ converges to $\{B_1 + \mu^* \geq 0\}$.
  - So $Y = \mathbf{1}_{\{B_1+\mu^* \geq 0\}}$ is the event that team A wins the game.
  - $p_0 = \mathbb{P}(Y = 1) = \Phi(\mu^*)$.
- Now we can define Team A's win probability as a Doob martingale:

$$p_t = \mathbb{E}[Y \mid \mathcal{F}_t], \quad t \in [0, 1]$$

- Then we can define the following quantities:
  - $M = \sup_{0 \leq t \leq 1} p_t$ is the maximum win probability team A attains.
  - $\tau_x = \inf\{t \in [0, 1] : p_t = x\}$ is the first time p_t exceeds x.

Towards a Continuous Closed-Form

- From here, note the following:
  - The discrete terminal event $\{X_N \geq 0\}$ converges to $\{B_1 + \mu^* \geq 0\}$.
  - So $Y = \mathbf{1}_{\{B_1 + \mu^* \geq 0\}}$ is the event that team A wins the game.
  - $p_0 = \mathbb{P}(Y = 1) = \Phi(\mu^*)$.
- Now we can define Team A's win probability as a Doob martingale:

$$p_t = \mathbb{E}[Y \mid \mathcal{F}_t], \quad t \in [0, 1]$$

- Then we can define the following quantities:
  - $M = \sup_{0 \leq t \leq 1} p_t$ is the maximum win probability team A attains.
  - $\tau_x = \inf\{t \in [0, 1] : p_t = x\}$ is the first time p_t exceeds x.
- Then we invoke the optional stopping theorem and use properties of Brownian motion to derive the distribution of M_A, M_B, and M_ℓ.

Closed-Form Distributions: Continuous Case

- **Theorem 4:** The distribution of M_A satisfies:

$$F_{M_A}(x) = 1 - \frac{p_0}{x} = 1 - \frac{\Phi(\mu^*)}{x}, \quad x \in [\Phi(\mu^*), 1)$$

Closed-Form Distributions: Continuous Case

- **Theorem 4:** The distribution of M_A satisfies:

$$F_{M_A}(x) = 1 - \frac{p_0}{x} = 1 - \frac{\Phi(\mu^*)}{x}, \quad x \in [\Phi(\mu^*), 1)$$

- **Theorem 5:** The conditional distribution of M_A given that team A loses satisfies:

$$\begin{aligned} F_{M_A|Y=0}(x) &= 1 - \left(\frac{p_0}{1-p_0}\right) \cdot \left(\frac{1-x}{x}\right), \quad x \in [p_0, 1) \\ &= 1 - \left(\frac{\Phi(\mu^*)}{1-\Phi(\mu^*)}\right) \cdot \left(\frac{1-x}{x}\right), \quad x \in [\Phi(\mu^*), 1) \end{aligned}$$

Closed-Form Distributions: Continuous Case

- What about the distribution for team B?

Closed-Form Distributions: Continuous Case

- What about the distribution for team B?
  - In the discrete case, we just replaced p_0 with $1 - p_0$.

Closed-Form Distributions: Continuous Case

- What about the distribution for team B?
  - In the discrete case, we just replaced p_0 with $1 - p_0$.
  - So in the continuous case, we just replace $\Phi(\mu^*)$ with $\Phi(-\mu^*)$.

Closed-Form Distributions: Continuous Case

- What about the distribution for team B?
  - In the discrete case, we just replaced p_0 with $1 - p_0$.
  - So in the continuous case, we just replace $\Phi(\mu^*)$ with $\Phi(-\mu^*)$.
- We saw in the discrete case that the distribution of M_ℓ is a piecewise mixture. We use a similar logic to derive the distribution of M_ℓ.

Closed-Form Distributions: Continuous Case

- What about the distribution for team B?
  - In the discrete case, we just replaced p_0 with $1 - p_0$.
  - So in the continuous case, we just replace $\Phi(\mu^*)$ with $\Phi(-\mu^*)$.
- We saw in the discrete case that the distribution of M_ℓ is a piecewise mixture. We use a similar logic to derive the distribution of M_ℓ.
- **Theorem 6:** Since team labels are arbitrary, let team A be the favorite ($\mu^* \geq 0$). Then the distribution of M_ℓ satisfies

$$F_{M_\ell}(x) = \begin{cases} 1 - \dfrac{1 - p_0}{x} & \text{if } x \in [1 - p_0, p_0) \\ 2 - \dfrac{1}{x} & \text{if } x \in [p_0, 1) \end{cases}$$

where $p_0 = \Phi(\mu^*)$.

Closed-Form Examples

- In evenly-matched games, we have the following:

Closed-Form Examples

- In evenly-matched games, we have the following:
  - Each team's starting win probability (p_0) is 0.5.

Closed-Form Examples

- In evenly-matched games, we have the following:
  - Each team's starting win probability (p_0) is 0.5.
  - $F_{M_\ell}(x) = 2 - \frac{1}{x}$ uniformly over all $x \in [0.5, 1)$.

Closed-Form Examples

- In evenly-matched games, we have the following:
  - Each team's starting win probability (p_0) is 0.5.
  - $F_{M_\ell}(x) = 2 - \frac{1}{x}$ uniformly over all $x \in [0.5, 1)$.
- In unevenly-matched games, we have the following:

Closed-Form Examples

- In evenly-matched games, we have the following:
  - Each team's starting win probability (p_0) is 0.5.
  - $F_{M_\ell}(x) = 2 - \frac{1}{x}$ uniformly over all $x \in [0.5, 1)$.
- In unevenly-matched games, we have the following:
  - Starting win probabillities differ ($p_0 \neq 0.5$).

Closed-Form Examples

- In evenly-matched games, we have the following:
  - Each team's starting win probability (p_0) is 0.5.
  - $F_{M_\ell}(x) = 2 - \frac{1}{x}$ uniformly over all $x \in [0.5, 1)$.
- In unevenly-matched games, we have the following:
  - Starting win probabillities differ ($p_0 \neq 0.5$).
  - M_ℓ is distributed **piecewise** over $x \in [p_0^{(\text{dog})}, 1)$.

Closed-Form Examples

- In evenly-matched games, we have the following:
  - Each team's starting win probability (p_0) is 0.5.
  - $F_{M_\ell}(x) = 2 - \frac{1}{x}$ uniformly over all $x \in [0.5, 1)$.
- In unevenly-matched games, we have the following:
  - Starting win probabillities differ ($p_0 \neq 0.5$).
  - M_ℓ is distributed **piecewise** over $x \in [p_0^{(\text{dog})}, 1)$.
    - From $[p_0^{(\text{dog})}, p_0^{(\text{fav})})$, only the underdog's maxima contribute.

Closed-Form Examples

- In evenly-matched games, we have the following:
  - Each team's starting win probability (p_0) is 0.5.
  - $F_{M_\ell}(x) = 2 - \frac{1}{x}$ uniformly over all $x \in [0.5, 1)$.
- In unevenly-matched games, we have the following:
  - Starting win probabillities differ ($p_0 \neq 0.5$).
  - M_ℓ is distributed **piecewise** over $x \in [p_0^{(\text{dog})}, 1)$.
    - From $[p_0^{(\text{dog})}, p_0^{(\text{fav})})$, only the underdog's maxima contribute.
    - From $[p_0^{(\text{fav})}, 1)$, both teams' maxima contribute.

But Wait. . . Does This Even Matter?

- The math is clean, but it's still a major simplification of sports!

But Wait. . . Does This Even Matter?

- The math is clean, but it's still a major simplification of sports!
  - Scoring often isn't binary!

But Wait...Does This Even Matter?

- The math is clean, but it's still a major simplification of sports!
    - Scoring often isn't binary!
    - True team scoring probabilities may vary throughout a game!

But Wait. . . Does This Even Matter?

- The math is clean, but it's still a major simplification of sports!
  - Scoring often isn't binary!
  - True team scoring probabilities may vary throughout a game!
  - Teams often have an unequal number of possessions!

But Wait. . . Does This Even Matter?

- The math is clean, but it's still a major simplification of sports!
  - Scoring often isn't binary!
  - True team scoring probabilities may vary throughout a game!
  - Teams often have an unequal number of possessions!
  - Weather, injuries, game strategy, and other external factors matter!

But Wait...Does This Even Matter?

- The math is clean, but it's still a major simplification of sports!
  - Scoring often isn't binary!
  - True team scoring probabilities may vary throughout a game!
  - Teams often have an unequal number of possessions!
  - Weather, injuries, game strategy, and other external factors matter!
- Let's take a look at the distribution of M_ℓ for real-life games...

Evenly-Matched NFL Games (2002–2024, |Spread| < 2)

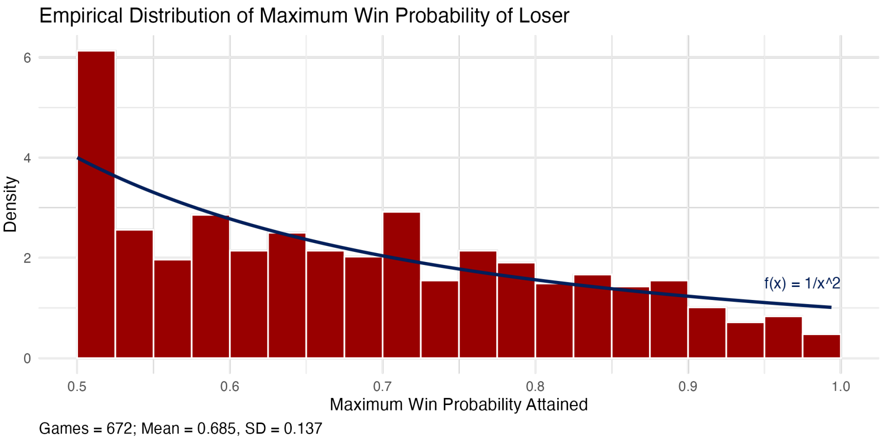

Data via nflfastR (2002–2024)

Evenly-Matched NFL Games (2002–2024, |Spread| < 2)

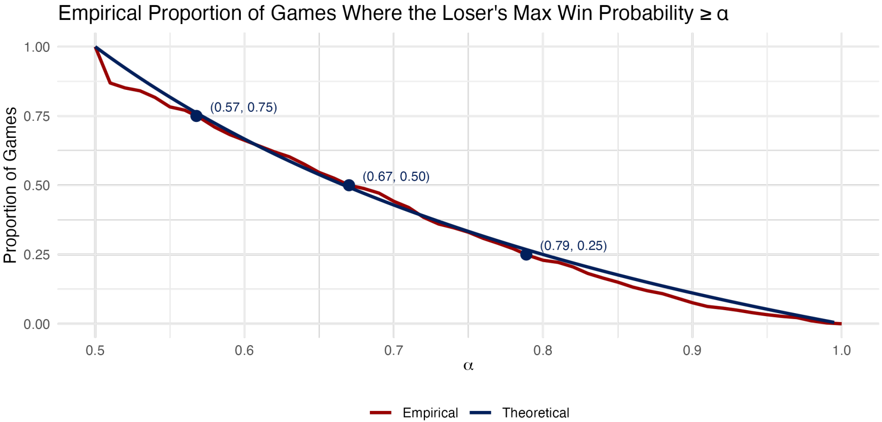

Data via nflfastR (2002–2024)

Unevenly-Matched NFL Games: 2002–2024, $|\text{Spread}| = 3$

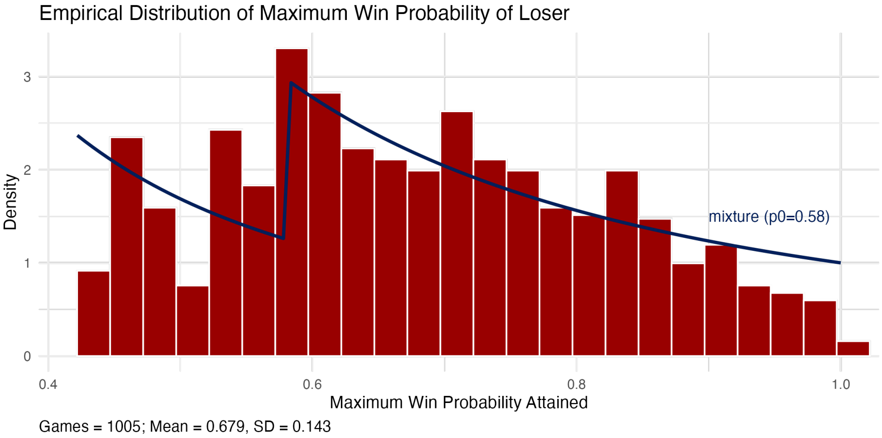

Data via nflfastR (2002–2024)

Unevenly-Matched NFL Games: 2002–2024, $|\text{Spread}| = 3$

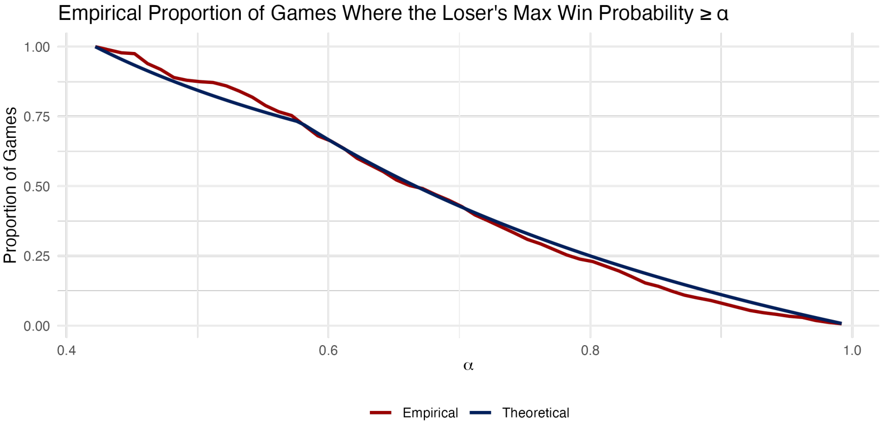

Data via nflfastR (2002–2024)

Conclusions

- Our solutions closely match the empirical distribution of the maximum win probability of the eventual loser (M_ℓ)!

Conclusions

- Our solutions closely match the empirical distribution of the maximum win probability of the eventual loser (M_ℓ)!
- Blown leads happen **all the time**, especially in even matchups!

Conclusions

- Our solutions closely match the empirical distribution of the maximum win probability of the eventual loser (M_ℓ)!
- Blown leads happen **all the time**, especially in even matchups!
- Conditioning on an eventual loss **fundamentally changes** the distribution of the maximum win probability attained.

Conclusions

- In Super Bowl LVII, the Eagles reached a maximum win probability of 78.4% before ultimately losing.

Conclusions

- In Super Bowl LVII, the Eagles reached a maximum win probability of 78.4% before ultimately losing.
- From this position (assuming a well-calibrated model) **it's true** that the Eagles only had a 21.6% chance of losing the game.

Conclusions

- In Super Bowl LVII, the Eagles reached a maximum win probability of 78.4% before ultimately losing.
- From this position (assuming a well-calibrated model) **it's true** that the Eagles only had a 21.6% chance of losing the game.
- However, the event that the eventual loser of this game reached 78.4% is **provably closer** to 30%.

Acknowledgements

- Special thanks to Professor Jiaoyang Huang, Dr. Paul Sabin, and Dr. Ryan Brill for their helpful feedback.
- All work supported by the Wharton Sports Analytics & Business Initiative (WSABI)

References

- Donsker, M. D. (1951). An invariance principle for certain probability limit theorems. *Memoirs of the American Mathematical Society*, 6.
- Doob, J. L. (1953). *Stochastic Processes*. New York: John Wiley & Sons.
- nflfastR: Carl, S. and Baldwin, B. (2025). *nflfastR: Functions to Efficiently Access NFL Play by Play Data*. R package version 5.1.0.9000. Available at `https://www.nflfastr.com/`.

Proof of Theorem 1

- **Theorem 1:** When $p_A = p_B$, the cumulative distribution of M satisfies
$$F_M(x) \geq 1 - \frac{p_0}{x}, \quad x \in [p_0, 1)$$
with equality exactly when $\mathbb{P}(\tau_x = N) = 0$.

Proof of Theorem 1 (continued)

- If $x \leq p_0$, then $\tau_x = 0$ a.s., so $M \geq x$ a.s. and $F_M(x) = 0$.
- For $x > p_0$, by the optional stopping theorem on the bounded martingale (p_k) at $\tau_x \wedge N$:

$$\mathbb{E}[p_{\tau_x \wedge N}] = p_0$$

- Decomposing this:

$$p_0 = \mathbb{E}[p_{\tau_x} \mathbf{1}_{\{\tau_x < N\}}] + \mathbb{E}[p_N \mathbf{1}_{\{\tau_x = N\}}]$$

- Since $p_{\tau_x} = x$ on $\{\tau_x < N\}$ and $p_N = 1$ on $\{\tau_x = N\}$:

$$p_0 = x\mathbb{P}(\tau_x < N) + \mathbb{P}(\tau_x = N)$$

Proof of Theorem 1 (continued)

- Since $\mathbb{P}(M \geq x) = \mathbb{P}(\tau_x < N) + \mathbb{P}(\tau_x = N)$:

$$p_0 = x\mathbb{P}(M \geq x) + (1 - x)\mathbb{P}(\tau_x = N)$$

- For $x < 1$, $\mathbb{P}(\tau_x = N) \geq 0$ and we have:

$$\mathbb{P}(M \geq x) \leq \frac{p_0}{x}, \quad x \in [p_0, 1)$$

  with equality exactly when $\mathbb{P}(\tau_x = N) = 0$.

- When $x = 1$, $\tau_x = N$ a.s., so:

$$\mathbb{P}(M \geq 1) = \mathbb{P}(M = 1) = p_0$$

Proof of Theorem 2

- **Theorem 2:** When $p_A = p_B$, the cumulative distribution of M conditional on $Y = 0$ satisfies

$$F_{M|Y=0}(x) \geq 1 - \left(\frac{p_0}{1-p_0}\right) \cdot \left(\frac{1-x}{x}\right), \quad x \in [p_0, 1)$$

with equality exactly when $\mathbb{P}(\tau_x = N) = 0$.

Proof of Theorem 2 (continued)

- Decompose the event $\{M \geq x\}$:

$$\mathbb{P}(M \geq x) = \mathbb{P}(M \geq x, Y = 0) + \mathbb{P}(M \geq x, Y = 1)$$

- On $\{Y = 1\}$, $p_N = 1 \geq x$ a.s. So $\{M \geq x\} \cap \{Y = 1\} = \{Y = 1\}$:

$$\begin{aligned}\mathbb{P}(M \geq x) &= \mathbb{P}(Y = 1) + \mathbb{P}(M \geq x, Y = 0)\\ &= \mathbb{P}(Y = 1) + \mathbb{P}(Y = 0)\mathbb{P}(M \geq x \mid Y = 0)\\ &= p_0 + (1 - p_0)\mathbb{P}(M \geq x \mid Y = 0)\end{aligned}$$

- From Theorem 1: $\frac{p_0}{x} \geq p_0 + (1 - p_0)\mathbb{P}(M \geq x \mid Y = 0)$
- Solving for $\mathbb{P}(M \geq x \mid Y = 0)$:

$$\mathbb{P}(M \geq x \mid Y = 0) \leq \left(\frac{p_0}{1 - p_0}\right) \cdot \left(\frac{1 - x}{x}\right)$$

with equality exactly when $\mathbb{P}(\tau_x = N) = 0$.

Proof of Theorem 3

- **Theorem 3:** When $p_A = p_B$, the distribution of M_ℓ satisfies

$$F_{M_\ell}(x) \geq 2 - \frac{1}{x}, \quad x \in [0.5, 1)$$

with equality exactly when $\mathbb{P}(\tau_x = N) = 0$.

Proof of Theorem 3 (continued)

- Since $p_A = p_B$, we have $p_0 = 0.5$ and the mixture weights are equal.
- From Theorem 2, both conditional distributions are identical:

$$F_{M_A|Y=0}(x) = F_{M_B|Y=1}(x) \geq 1 - \left(\frac{0.5}{0.5}\right) \cdot \left(\frac{1-x}{x}\right) = 1 - \frac{1-x}{x}$$

- Since both teams have identical distributions, the mixture is simply:

$$F_{M_\ell}(x) \geq 0.5 \cdot F_{M_A|Y=0}(x) + 0.5 \cdot F_{M_B|Y=1}(x) = 1 - \frac{1-x}{x} = 2 - \frac{1}{x}$$

Proof of Theorem 4

- **Theorem 4:** When $p_A \neq p_B$, the cumulative distribution of M satisfies

$$F_M(x) = 1 - \frac{p_0}{x} = 1 - \frac{\Phi(\mu^*)}{x}, \quad x \in [p_0, 1)$$

Proof of Theorem 4 (continued)

- By the optional stopping theorem on the bounded martingale (p_t) at $\tau_x \wedge 1$:

$$\mathbb{E}[p_{\tau_x \wedge 1}] = p_0$$

- Decomposing this:

$$p_0 = \mathbb{E}[p_{\tau_x} \mathbf{1}_{\{\tau_x < 1\}}] + \mathbb{E}[p_1 \mathbf{1}_{\{\tau_x = 1\}}]$$

- Since $p_{\tau_x} = x$ on $\{\tau_x < 1\}$ and $p_1 = 1$ on $\{\tau_x = 1\}$:

$$p_0 = x\mathbb{P}(\tau_x < 1) + \mathbb{P}(\tau_x = 1)$$

- Since $\mathbb{P}(M \geq x) = \mathbb{P}(\tau_x < 1) + \mathbb{P}(\tau_x = 1)$:

$$p_0 = x\mathbb{P}(M \geq x) + \mathbb{P}(\tau_x = 1)$$

Proof of Theorem 4 (continued)

- For $x < 1$, $\mathbb{P}(\tau_x = 1) = 0$ (continuous paths), so:

$$\mathbb{P}(M \geq x) = \frac{p_0}{x} = \frac{\Phi(\mu^*)}{x}, \quad x \in [p_0, 1)$$

- When $x = 1$, $\tau_x = 1$ a.s., so:

$$\mathbb{P}(M \geq 1) = \mathbb{P}(M = 1) = p_0 = \Phi(\mu^*)$$

Proof of Theorem 5

- **Theorem 5:** When $p_A \neq p_B$, the conditional distribution of M given $Y = 0$ satisfies

$$F_{M|Y=0}(x) = 1 - \left(\frac{p_0}{1 - p_0}\right) \cdot \left(\frac{1 - x}{x}\right), \quad x \in [p_0, 1)$$

where $p_0 = \Phi(\mu^*)$ and $M = 0$ for $x \in [0, p_0)$.

Proof of Theorem 5 (continued)

- Decompose the event $\{M \geq x\}$:

$$\mathbb{P}(M \geq x) = \mathbb{P}(M \geq x, Y = 0) + \mathbb{P}(M \geq x, Y = 1)$$

- On $\{Y = 1\}$, $p_1 = 1 \geq x$ a.s. So $\{M \geq x\} \cap \{Y = 1\} = \{Y = 1\}$:

$$\begin{aligned}
\mathbb{P}(M \geq x) &= \mathbb{P}(Y = 1) + \mathbb{P}(M \geq x, Y = 0) \\
&= \mathbb{P}(Y = 1) + \mathbb{P}(Y = 0)\mathbb{P}(M \geq x \mid Y = 0) \\
&= p_0 + (1 - p_0)\mathbb{P}(M \geq x \mid Y = 0)
\end{aligned}$$

- From Theorem 4: $\frac{p_0}{x} = p_0 + (1 - p_0)\mathbb{P}(M \geq x \mid Y = 0)$
- Solving for $\mathbb{P}(M \geq x \mid Y = 0)$:

$$\mathbb{P}(M \geq x \mid Y = 0) = \left(\frac{p_0}{1 - p_0}\right) \cdot \left(\frac{1 - x}{x}\right)$$

Proof of Theorem 6

- **Theorem 6:** When $p_A \neq p_B$, let team A be the favorite ($\mu^* \geq 0$). Then the distribution of M_ℓ satisfies

$$F_{M_\ell}(x) = \begin{cases} 1 - \dfrac{1-p_0}{x} & \text{if } x \in [1-p_0, p_0) \\ 2 - \dfrac{1}{x} & \text{if } x \in [p_0, 1) \end{cases}$$

where $p_0 = \Phi(\mu^*)$ and $M_\ell = 0$ for $x \in [0, 1-p_0)$.

Proof of Theorem 6 (continued)

- The distribution of M_ℓ is a mixture: $F_{M_\ell}(x) = (1 - p_0)F_{M_A|Y=0}(x) + p_0 F_{M_B|Y=1}(x)$.
- From Theorem 5, we have the conditional distributions:

$$F_{M_A|Y=0}(x) = \begin{cases} 0 & x \in [0, p_0) \\ 1 - \left(\frac{p_0}{1-p_0}\right) \cdot \left(\frac{1-x}{x}\right) & x \in [p_0, 1) \end{cases}$$

$$F_{M_B|Y=1}(x) = \begin{cases} 0 & x \in [0, 1 - p_0) \\ 1 - \left(\frac{1-p_0}{p_0}\right) \cdot \left(\frac{x}{1-x}\right) & x \in [1 - p_0, 1) \end{cases}$$

Proof of Theorem 6 (continued)

- For $x \in [1-p_0, p_0)$: $F_{M_A|Y=0}(x) = 0$ (since $x < p_0$), so

$$F_{M_\ell}(x) = (1-p_0)\cdot 0 + p_0 \cdot \left(1 - \frac{1-p_0}{p_0}\cdot\frac{x}{1-x}\right) \\ = 1 - \frac{1-p_0}{x}$$

- For $x \in [p_0, 1)$: Both terms contribute, giving

$$F_{M_\ell}(x) = (1-p_0)\left(1 - \frac{p_0}{1-p_0}\cdot\frac{1-x}{x}\right) + p_0\left(1 - \frac{1-p_0}{p_0}\cdot\frac{x}{1-x}\right) \\ = 2 - \frac{1}{x}$$