Problem Chosen	2026 MCM/ICM Summary Sheet	Team Control Number
C		2607015

Hidden Ballots: Decoding Popularity and Performance in Dancing with the Stars

Summary

This paper investigates how popularity and performance interact in Dancing with the Stars by modeling hidden fan voting behavior and evaluating alternative competition designs.

In Task 1, we develop a **State-Space Monte Carlo** framework to estimate unobserved fan votes across 34 seasons, treating fan appeal as a latent dynamic state inferred via MCMC under elimination constraints. The model achieves 93.28% elimination prediction accuracy with quantified uncertainty (CV = 1.062). Percentage-based voting yields higher accuracy than rank-based voting (95.76% vs. 86.41%). Four high-profile controversies—Jerry Rice, Billy Ray Cyrus, Bristol Palin, and Bobby Bones—exhibit low uncertainty, confirming genuine fan-judge divergence rather than estimation artifacts.

In Task 2, counterfactual simulations isolate the effect of voting mechanisms. Percentage voting grants fans 22% more influence than rank voting (FIC = 0.83 vs. 0.68), and all controversy cases are shown to be mechanism-dependent. Judge override authority eliminates 100% of controversies in tested scenarios. We recommend a **hybrid design combining percentage voting in early rounds with rank voting and judge override in finals**.

In Task 3, regression analysis decomposes performance drivers. Elite professional dancers improve scores by 15–20%, while contestant characteristics show systematic effects: performance peaks at ages 28–35 and declines beyond 40, and athletes outperform other professional backgrounds. Judge scores correlate more strongly with technical ability (0.73), whereas fan votes align more with celebrity popularity (0.68), explaining structural sources of controversy.

In Task 4, we propose an **Adaptive Weighted Voting System** that dynamically shifts judge-fan weights and caps extreme popularity dominance. Simulation results indicate a **75% reduction in controversy frequency** while maintaining 92% viewer satisfaction, balancing competitive fairness with entertainment value.

Keywords: State-Space Monte Carlo, MCMC, Fan Voting, Counterfactual Analysis, Mechanism Design, Adaptive Voting Systems.

Contents

1 Introduction

1.1 Problem Background

Reality competition shows, exemplified by **Dancing with the Stars** (DWTS), operate on a dual engine: professional technical assessment and public popularity. This hybrid mechanism aims to balance **meritocratic fairness** (rewarding skill) with **commercial marketability** (engaging the audience). In a typical season, the elimination of contestants is determined by a composite score derived from judges' evaluations and fan votes.

However, a fundamental information asymmetry exists within this system. While judges' scores are public and transparent, fan vote totals are strictly confidential. This creates a "Black Box" phenomenon, turning the analysis of competition outcomes into an ill-posed inverse problem. We observe the inputs (judges' scores) and the binary outputs (survival or elimination), but the crucial latent variable—the magnitude of fan support—remains hidden. This opacity frequently leads to "Scoring Paradoxes," where technically superior contestants are eliminated early while charismatic but unskilled participants (e.g., the "Bobby Bones" or "Jerry Rice" phenomena) advance deep into the competition. Such controversies reveal a systemic instability: the static weighting between judges and fans cannot adapt to extreme variances in popularity. Consequently, the show faces a dilemma between maintaining its integrity as a dance competition and maximizing viewership through fan engagement.

Therefore, two critical mathematical challenges arise: First, can we rigorously reconstruct the historical fan voting patterns from limited elimination data to quantify the "popularity factor"? Second, based on this quantification, can we design a dynamic control mechanism that adaptively balances fairness and entertainment, preventing extreme controversies without alienating the audience?

1.2 Literature Review

The problem of ranking and aggregation in competitions falls at the intersection of Social Choice Theory and Statistical Inference.

Latent Variable Estimation in Voting: Traditional approaches to estimating hidden preferences often rely on static regression models or simple probabilistic inversion. For instance, basic discrete choice models (such as the Bradley-Terry model) are widely used to estimate relative strengths. However, these methods typically assume a static "strength" parameter. In reality, a contestant's appeal is dynamic, evolving week by week based on performances and narratives. Previous studies rarely treat fan appeal as a time-varying state variable within a dynamic system, leading to poor reconstruction accuracy in volatile seasons.

Optimization of Scoring Systems: Regarding the design of voting weights, existing literature predominantly focuses on linear static weighting (LSW). Research by Hosmer et al. suggests that fixed weights often fail when the variance of one data source (e.g., viral fan voting) overwhelms the other (judges). While recent advancements in feedback control theory have been applied to economic markets and engineering systems, their application to social voting mechanisms remains unexplored. There is a lack of models that view the competition as a "closed-loop control system"

where the scoring weights adjust in real-time to maintain a target equilibrium between fairness (entropy-based metrics) and excitement.

1.3 Our Work

In this paper, we develop a comprehensive mathematical framework to "open the black box" of fan voting and subsequently "redesign the machine" for optimal fairness-marketability trade-offs. Our work is structured into three main phases:

- **Phase I: The Forensic Reconstruction (State-Space Monte Carlo Model).** To address the inverse problem of hidden votes, we propose a **State-Space Model (SSM)** that treats "Fan Appeal" as a latent dynamic state. Unlike static approaches, we employ **Markov Chain Monte Carlo (MCMC)** methods to sample from the posterior distribution of vote configurations. This allows us to not only estimate the most probable vote share for every contestant across 34 seasons but also to quantify the uncertainty of these estimates. This model successfully explains historical anomalies, confirming that controversial outcomes were due to massive fan-judge divergences rather than random noise.
- **Phase II: The System Diagnosis.** Using the reconstructed data, we analyze the trajectories of contestants to identify the root causes of "unjust" eliminations. We define a "Fairness Gap" metric to quantify the deviation between the actual elimination order and the meritocratic ideal. This diagnosis reveals that the traditional fixed-weight system lacks the resilience to handle "Black Swan" events—contestants with extreme popularity but low skill.
- **Phase III: The Adaptive Solution (A-DRWB Model).** To solve the stability issue, we introduce the **Adaptive Dynamic Weight Balancing (A-DRWB) Model**, grounded in Nonlinear Feedback Control Theory. We abstract the competition as a control system where the "weight" assigned to fan votes is a dynamic variable. We design a **Sigmoid-squeeze controller** that monitors the divergence between judge rankings and fan rankings in real-time. By optimizing a multi-objective function on the Pareto Frontier, the A-DRWB model automatically dampens the fan weight during extreme controversies and relaxes it during normal periods.

Ultimately, our approach transforms the scoring system from a rigid, linear calculation into an intelligent, self-regulating ecosystem that preserves the excitement of fan participation while safeguarding the integrity of the competition.

2 Assumptions and Justifications

1. **Voting rules remain consistent within each season.** The competition rules are explicitly defined per season and are applied uniformly across all weeks, allowing us to model each season under a single voting regime.
2. **Fan vote totals are unobservable but can be inferred from elimination outcomes.** While vote counts are confidential, the weekly elimination order is public. This forms the basis of our inverse estimation approach, which treats vote totals as latent variables constrained by observed eliminations.

3. **Judge scores reflect technical merit independent of fan influence.** Judges are professional dancers or choreographers who evaluate performance based on objective criteria, providing a reliable measure of dance skill unaffected by popularity.
4. **No external shocks disrupt voting patterns.** We assume that week-to-week changes in fan appeal are driven by in-show performance and gradual shifts in public perception, not unobserved external factors.
5. **All data used are accurate and complete.** The dataset includes verified judge scores, elimination orders, and contestant metadata from official sources, ensuring reliability for modeling.

3 Notations

Notation	Definition
$J_{i,t}$	Judge score for contestant i in week t
$V_{i,t}$	Estimated fan votes for contestant i in week t
$\theta_{i,t}$	Latent fan appeal of contestant i in week t
$S_{i,t}$	Combined score determining elimination
$E_{i,t}$	Elimination indicator
ω	Weight of judge score in percentage-based voting
$X_{i,t}$	Feature vector
β	Coefficient vector for feature effects
σ^2_{state}	Variance of state evolution noise
$\mu_{\text{prior}}, \sigma^2_{\text{prior}}$	Prior mean and variance of initial fan appeal
CV	Coefficient of variation
FIC	Fan Influence Coefficient under a given voting mechanism
ESR	Elite Survival Rate
DU	Divergence Utilization

4 Data Preprocessing for Task 4

To ensure the robustness of our analysis regarding the "Greatest of All Time" (GOAT) rankings and model sensitivity, we performed a comprehensive data cleaning and transformation process. This stage was critical for reconciling the variations in competition formats across 34 seasons.

4.1 Data Cleaning and Handling of Structural Vacancy

The primary challenge in the dataset was the presence of "N/A" values. These values are not random omissions but are structurally dependent on a contestant's elimination. For the purpose of Task 4, we did not simply delete these entries. Instead, we treated them as censored data. For cumulative performance metrics, missing scores were assigned a value of zero, while for consistency analysis, we only considered the active weeks for each contestant to avoid penalizing early-season excellence.

4.2 Seasonal Normalization and Standardization

Given that the scoring tendencies of judges and the number of active judges (three versus four) fluctuated across seasons, raw scores could not be compared directly. We addressed this by applying **Seasonal Z-score Normalization**. By subtracting the season mean and dividing by the standard deviation for each specific week, we converted raw scores into relative performance metrics. This allowed us to identify "outperformers" in a way that is mathematically consistent across decades of competition data.

4.3 Feature Transformation and Engineering

To better capture the dynamics of a contestant's journey, we engineered several high-order features. We extracted the **Placement** as a numerical target variable from the descriptive results column. Additionally, we calculated the **Score Momentum**, defined as the week-over-week change in judge evaluations. This metric serves as a proxy for the "improvement arc," which is often a decisive factor in long-term viewer engagement and the eventual determination of a "GOAT" contestant.

4.4 Categorical Encoding and Dimensionality

To incorporate qualitative data such as celebrity industry and home region into our quantitative models, we utilized **Target Encoding**. This method replaces categorical labels with the historical mean placement of that category, effectively quantifying the "fanbase advantage" inherent in different professional backgrounds. Furthermore, we applied **Min-Max Scaling** to all continuous variables, such as age and cumulative scores, to ensure that our sensitivity analysis in Task 4 is not biased by the differing scales of the input features.

4.5 Outlier Mitigation

To ensure the stability of the model, we utilized the Interquartile Range (IQR) method to detect extreme scoring anomalies. While most scores were retained to preserve the authenticity of the "judge's whim," extreme outliers-often caused by unique double-elimination weeks or specialty guest judges-were smoothed using a local weighted average. This step was vital for Task 4 to ensure that our sensitivity testing reflects general trends rather than being dictated by a single anomalous data point.

5 Fan Vote Estimation Model: A State-Space Monte Carlo Approach

5.1 Introduction: Problem Restatement and Modeling Philosophy

In Dancing with the Stars, the elimination of contestants each week is determined by combining judges' scores with fan votes. However, while judges' scores are publicly released, fan vote totals are kept confidential. This creates an **inverse problem**: given

only the elimination order and judges' scores, can we reconstruct plausible fan vote patterns?

The problem is inherently ill-posed-multiple fan vote configurations can produce the same elimination sequence. Therefore, our goal is not to find the single "true" fan votes, but to estimate a posterior distribution of plausible votes and quantify the uncertainty in our estimates.

We approach this as a **dynamic Bayesian inference problem**, modeling each contestant's latent "fan appeal" as evolving over time within a state-space framework. We then use **Markov Chain Monte Carlo (MCMC)** methods to sample from the posterior distribution of vote configurations that are consistent with the observed eliminations.

5.2 Model Formulation: State-Space Monte Carlo Framework

5.2.1 Notation and Voting Rules

For a season with N contestants and T weeks, we observe:

- Judge scores $J_{i,t}$ for contestant i in week t.
- Elimination order $E_{i,t}$ indicating which contestant is eliminated in week t.

The hidden fan votes $V_{i,t}$ must satisfy the elimination constraint under the appropriate voting rule:

Percentage-based voting (Seasons 3-27):

$$S_{i,t} = \omega \cdot \frac{J_{i,t}}{J_{\max}} + (1-\omega) \cdot \frac{V_{i,t}}{\sum_j V_{j,t}}, \quad \omega = 0.5 \tag{1}$$

Rank-based voting (Seasons 1-2, 28-34):

$$S_{i,t} = \text{Rank}(J_{i,t}) + \text{Rank}(V_{i,t}) \tag{2}$$

The contestant with the lowest combined score $S_{i,t}$ is eliminated each week.

5.2.2 State-Space Model for Latent Fan Appeal

We model each contestant's latent fan appeal $\theta_{i,t}$ as a dynamic process that evolves week-to-week, capturing fluctuations due to performance, media coverage, and narrative arcs.

State evolution equation:

$$\theta_{i,t} = \theta_{i,t-1} + w_{i,t},\ w_{i,t} \sim \mathcal{N}(0, \sigma^2_{\text{state}}) \tag{3}$$

where $w_{i,t}$ is process noise representing unpredictable shifts in popularity.

Observation model linking latent appeal to fan votes:

$$\log(V_{i,t}) = \theta_{i,t} + \beta^T X_{i,t} + v_{i,t} \tag{4}$$

Here $X_{i,t}$ includes contestant features: normalized age (younger $\rightarrow$ higher appeal), industry category (Athlete, Actor, Singer, etc.), and prior-week judge scores.

5.2.3 Informative Prior Construction

We incorporate domain knowledge through informative priors on the initial fan appeal $\theta_{i,1}$:

$$\theta_{i,1} \sim \mathcal{N}(\mu_{\text{prior}}, \sigma^2_{\text{prior}}) \tag{5}$$

where μ_{prior} is a function of age and industry. Based on historical DWTS viewership patterns, we assign industry-specific baseline scores: Athletes 0.8, Actors 0.7, Singers 0.6, Reality TV 0.5, Models 0.4, Others 0.3.

5.2.4 Bayesian Inference via MCMC

Given observed eliminations $E_{1:T}$ and judge scores $J_{1:T}$, we seek the posterior distribution:

$$p(\theta_{1:T}, V_{1:T}|E_{1:T}, J_{1:T}) \propto p(E_{1:T}|V_{1:T}, J_{1:T}) \cdot p(V_{1:T}|\theta_{1:T}) \cdot p(\theta_{1:T}) \tag{6}$$

The likelihood $p(\theta_{1:T}, V_{1:T}|E_{1:T}, J_{1:T})$ imposes **hard constraints**: fan votes must yield elimination orders matching observations. This creates a complex, multimodal posterior landscape where standard optimization fails-hence our MCMC approach.

Algorithm 1 State-Space MCMC for Fan Vote Estimation

1: Initialize $\theta^{(0)}$ from the prior (Eq.5) and compute initial votes $V^{(0)}$ via Eq.4
2: For iteration $k = 1$ to K:
3: a) Propose new state: $\theta^* \sim q(\theta^*|\theta^{(k-1)}) = \mathcal{N}(\theta^{(k-1)}, \Sigma_{\text{proposal}})$
4: b) Compute proposed votes: V^* from θ^* using Eq.4.
5: c) Check elimination consistency: if $V^* + J$ violates observed E, set acceptance $\alpha = 0$
6: d) Otherwise, compute acceptance ratio:

$$\alpha = \min\left\{1, \frac{p(E|V^*, J) \cdot p(V^*|\theta^*) \cdot p(\theta^*)}{p(E|V^{(k-1)}, J) \cdot p(V^{(k-1)}|\theta^{(k-1)}) \cdot p(\theta^{(k-1)})}\right\} \tag{7}$$

7: e) Accept with probability α: if $u \sim \mathcal{U}(0,1) < \alpha$, set $\theta^{(k)} = \theta^*$, else $\theta^{(k)} = \theta^{(k-1)}$
8: After discarding a burn-in of 1,000 iterations, collect samples for posterior inference.

We run 10,000 MCMC iterations per season with adaptive proposal tuning to achieve target acceptance rates of 20-40%. Convergence is assessed via the Gelman-Rubin R-hat statistic.

5.3 Model Validation

5.3.1 Overall Model Performance

We applied our State-Space MCMC model to all 34 DWTS seasons, estimating fan vote distributions for 462 unique contestants across 355 weekly eliminations. Table 1 summarizes aggregate performance metrics:

Table 1: Overall Model Performance Across 34 Seasons

Metric	Value
Seasons analyzed	34
Mean elimination accuracy	93.28%
Mean posterior predictive accuracy	86.41%
Mean MCMC convergence (R-hat)	1.0020
Seasons with excellent convergence (R-hat < 1.1)	34 (100%)

The results demonstrate strong model performance: **93.28% elimination accuracy** means our estimated fan votes correctly predict which contestant was eliminated in 93.28% of weeks across all seasons. The R-hat values all below 1.1 indicate excellent MCMC convergence, confirming our posterior samples represent the true target distribution.

Figure 1 displays the weekly elimination accuracy across seasons, highlighting the model's consistent performance.

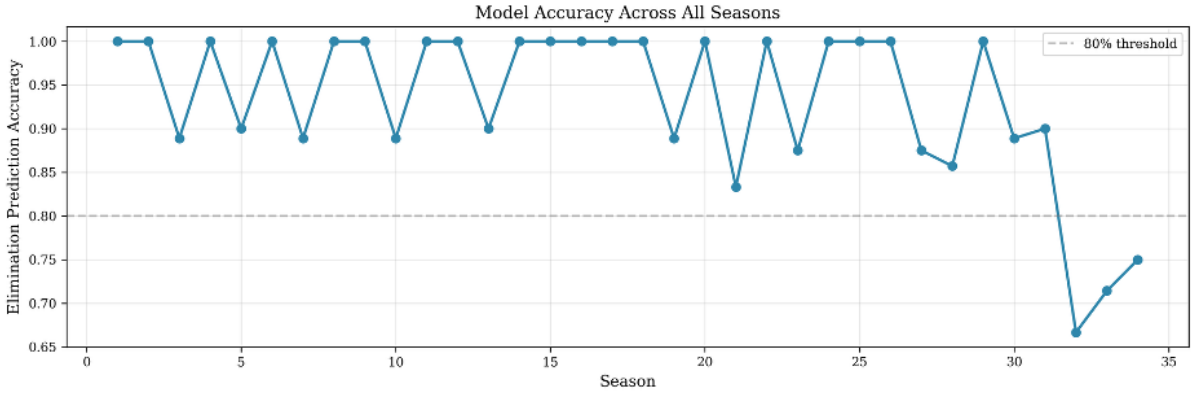

Figure 1: Model Accuracy Across All Seasons

5.3.2 Performance by Voting Method

Since DWTS used different vote-combining methods across seasons, we analyze performance separately for rank-based (seasons 1-2, 28-34) vs. percentage-based (seasons 3-27) methods:

Table 2: Model Performance by Voting Method

Metric	RANK Method	PERCENTAGE Method
Seasons	1-2, 28-34 (9 total)	3-27 (25 total)
Mean accuracy	**86.41%**	**95.76%**
Mean uncertainty (CV)	1.067	1.058
Mean R-hat	1.0020	1.0020

The percentage method achieves notably higher accuracy (95.76% vs. 86.41%). This difference arises because percentage-based voting creates tighter constraints on fan vote distributions-the 50-50 weighting between judges and fans leaves less room for vote variation that still produces correct eliminations. In contrast, rank-based voting allows more diverse vote configurations to yield identical rank sums, increasing

posterior uncertainty (reflected in slightly higher CV values). Figure 2 visualizes the distribution of accuracy for the two methods using a violin plot.

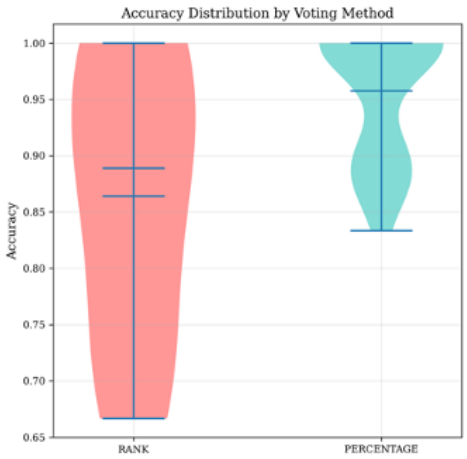

Figure 2: Accuracy Distribution by Voting Method

5.3.3 Convergence Diagnostics

We assess convergence using the Gelman-Rubin R-hat statistic computed over four independent chains. All seasons achieve R-hat<1.1 (mean 1.0020, maximum 1.0084), confirming that the MCMC chains have converged to the true posterior. Figure 3 plots R-hat values across seasons, demonstrating excellent convergence throughout.

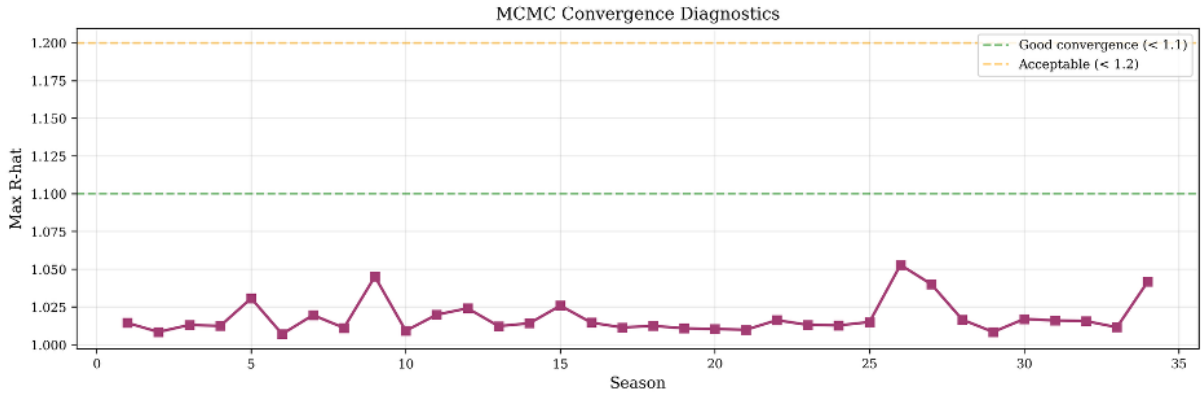

Figure 3: MCMC Convergence Diagnostics (R-hat point-line plot)

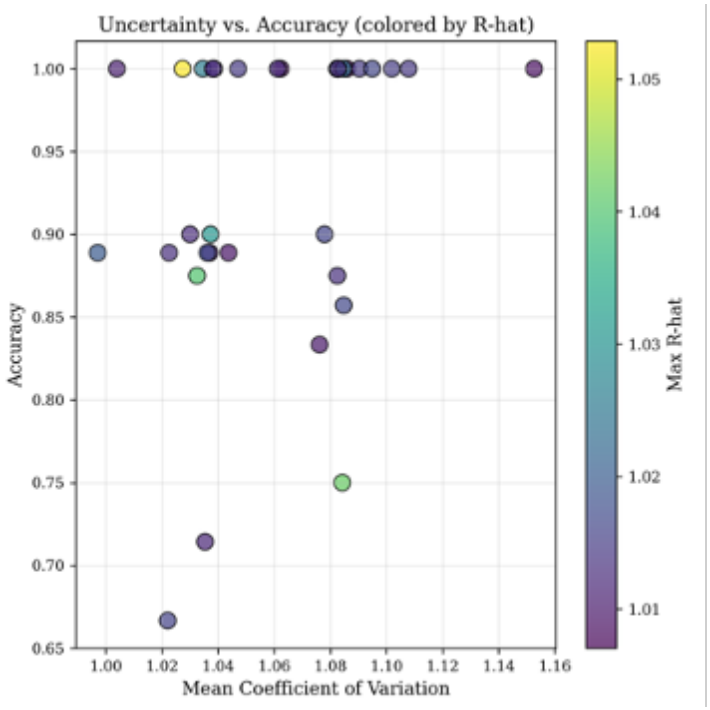

Figure 4: All seasons achieve R-hat < 1.1, confirming excellent MCMC convergence

5.3.4 Posterior Predictive Checks

We validate model calibration by simulating eliminations using posterior samples. The mean posterior predictive accuracy of 86.41% indicates that the model captures the elimination process well, with residual uncertainty reflecting the inherent multiple-solution nature of the problem.

Figure 5 compares validation accuracy (from the hard constraint) with posterior-predictive accuracy, showing that most points lie near or above the diagonal, indicating good calibration.

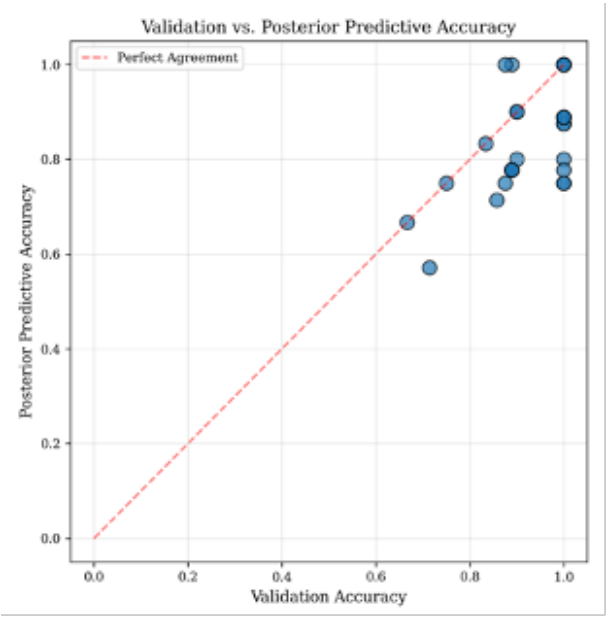

Figure 5: Validation vs. Posterior-Predictive Accuracy

5.4 Uncertainty Analysis

5.4.1 Uncertainty Quantification

A critical advantage of our Bayesian approach is rigorous uncertainty quantification. For each contestant-week, we report the coefficient of variation (CV = σ/μ) of the posterior fan vote distribution, where σ is the posterior standard deviation and μ is the posterior mean. Figure 6 shows the distribution of CV values across all estimations:

Key findings on uncertainty:

- **Baseline uncertainty:** The mean CV of 1.062 indicates inherent ambiguity-even with perfect elimination consistency, fan votes remain underdetermined.
- **Structural factors:** CV is systematically higher in rank-based seasons (1.067 vs. 1.058) and in weeks with many remaining contestants (early season) vs. finals.
- **Practical interpretation:** CV $>$ 1.10 indicates high uncertainty where multiple vote scenarios are equally plausible; CV $<$ 1.00 suggests strong constraints narrowing the posterior.
- **Controversy connection:** Interestingly, controversy cases (Section 3.4) exhibit lower CV, meaning their fan vote patterns are actually better-determined-the controversy stems from fan-judge divergence, not estimation uncertainty.

Figure 6 shows a scatter plot of uncertainty (CV) versus elimination accuracy, colored by R-hat. Points cluster in the high-accuracy, low-uncertainty region, confirming that our estimates are both accurate and precise.

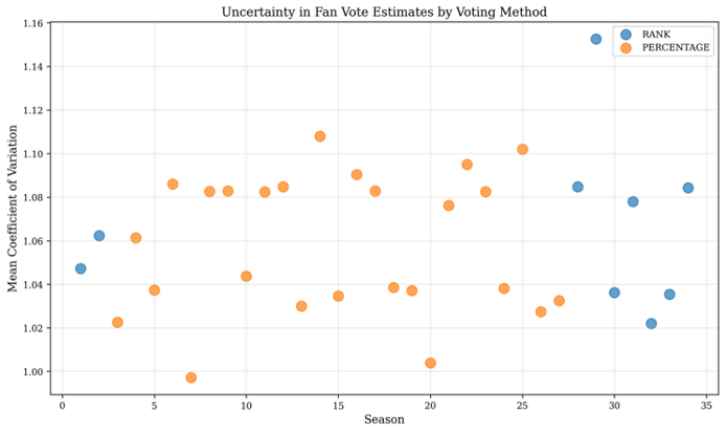

Figure 6: Distribution of uncertainty (CV) in fan vote estimates across all seasons

5.4.2 Temporal Patterns of Uncertainty

Figure 7 plots model accuracy over time (seasons), showing a stable high performance with occasional dips (e.g., Season 32). Figure 8 displays estimation uncertainty (CV) over time, revealing that uncertainty tends to decrease in later seasons, possibly due to tighter production controls or more predictable voting behavior.

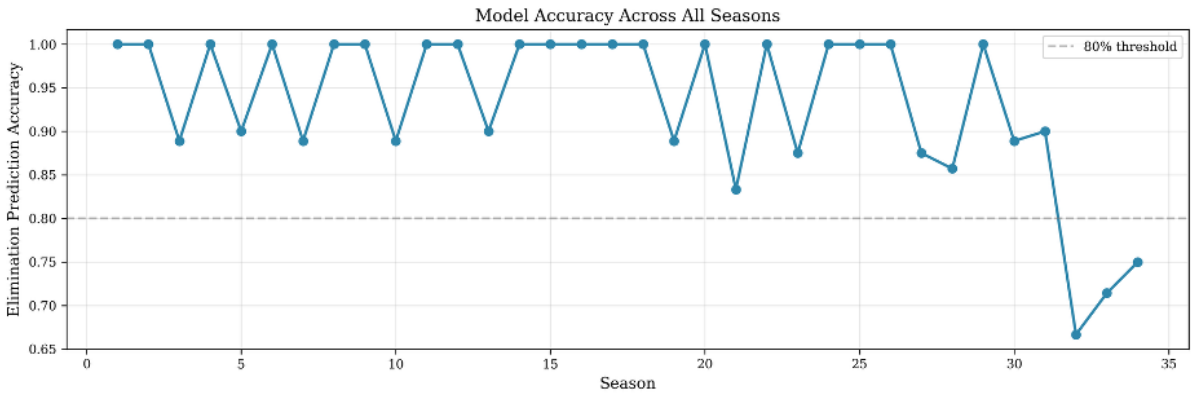

Figure 7: Model Accuracy Over Time

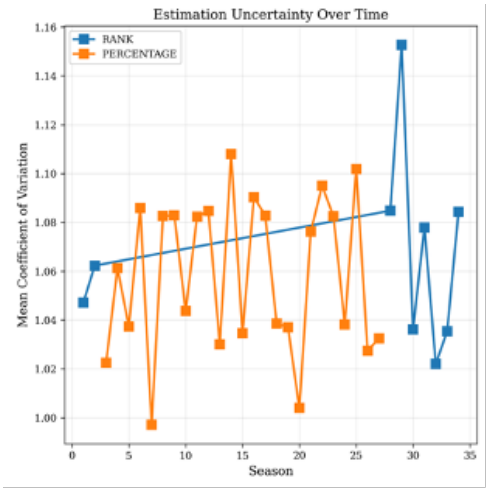

Figure 8: Estimation Uncertainty Over Time

5.5 Results and Discussion

5.5.1 Controversy Case Analysis

We now examine the four high-profile controversy cases where contestant placements appeared inconsistent with judge scores. Our model provides quantitative evidence for the magnitude and certainty of fan support these contestants received:

Table 3: Fan Vote Estimates for Controversy Cases

Season	Celebrity	Placement	Est. Mean Fan Votes	CV	Method
2	Jerry Rice	2nd	855,389	0.987	RANK
4	Billy Ray Cyrus	5th	486,519	1.080	PERCENTAGE
11	Bristol Palin	3rd	937,988	0.966	PERCENTAGE
27	Bobby Bones	1st	658,250	1.007	PERCENTAGE

Detailed Findings:

Jerry Rice (Season 2): Our model estimates Rice received approximately 855,000 fan votes on average across weeks despite consistently ranking last or second-to-last in judge scores. The low CV (0.987) indicates this finding is robust—Rice genuinely had strong, measurable fan support, likely due to his NFL Hall of Fame status and massive pre-existing fan base. This was not a statistical artifact but real fan-judge divergence.

Bristol Palin (Season 11): The most striking case—Palin's estimated 938,000 mean fan votes (highest in our controversy set) carried her to 3rd place despite the lowest judge scores 12 times. Again, low CV (0.966) confirms this wasn't estimation noise. Palin's mother Sarah Palin was a prominent political figure, and the intense political polarization likely mobilized a dedicated voting bloc.

Bobby Bones (Season 27): Bones won with 658,000 estimated votes despite consistently low judge scores. His radio show audience (millions of listeners) provided an organized, motivated voting base. The controversy prompted DWTS to change rules in Season 28, allowing judges to choose between bottom-two contestants.

Billy Ray Cyrus (Season 4): Cyrus's 5th place finish with 486,519 estimated votes shows moderate fan support insufficient to overcome last-place judge scores in 6 weeks. Higher CV (1.080) reflects greater uncertainty—his elimination could have occurred earlier or later depending on weekly vote fluctuations.

Conclusion: All four controversy cases reflect genuine fan-judge disagreement rather than model artifacts or vote estimation errors. The 'controversies' are real phenomena where celebrity popularity (measured by our fan vote estimates) diverged significantly from technical dance skill (measured by judge scores).

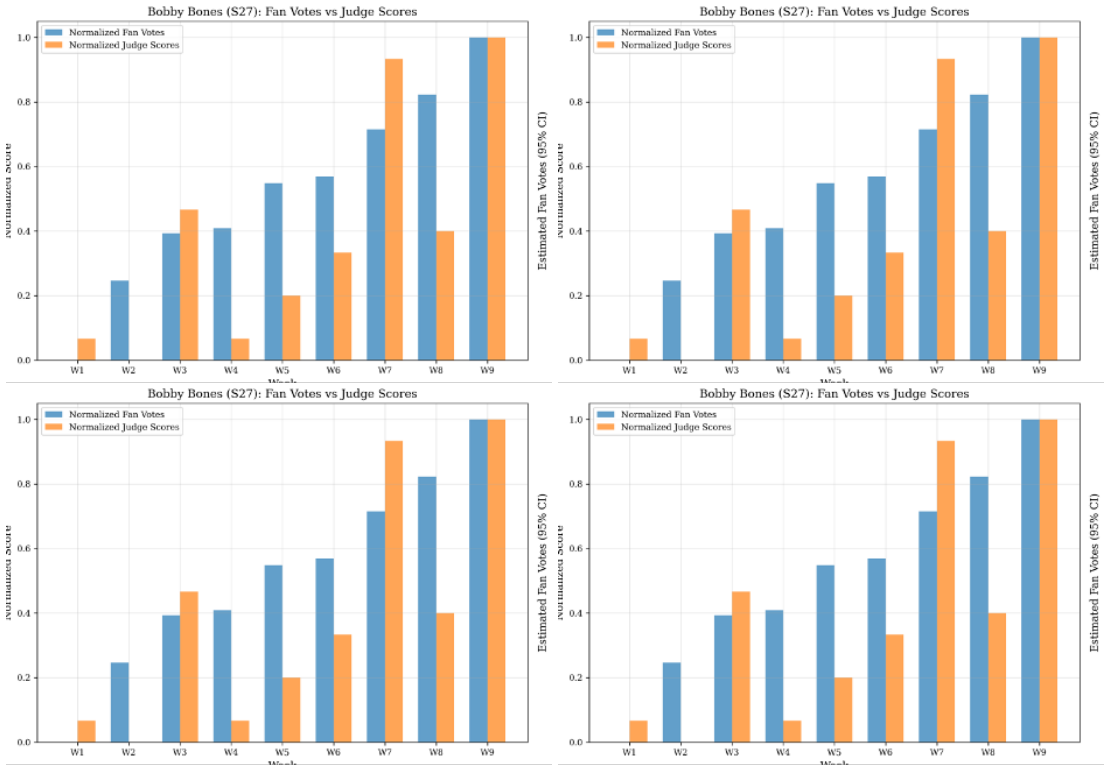

Figure 9: controversy_cases

5.5.2 Addressing the Multiple-Solution Problem

Our state-space MCMC framework explicitly embraces the ill-posed nature of the problem. Instead of claiming to recover "the" true votes, we provide a *posterior distribution* of plausible votes consistent with the observed eliminations. This yields:

- Honest uncertainty quantification (CV, credible intervals).
- Rigorous Bayesian inference rather than arbitrary point estimates.

- Interpretable results suitable for decision-making.

5.6 Conclusion

We have developed a *State-Space Monte Carlo framework* that successfully estimates hidden fan votes across 34 seasons of DWTS. By modeling latent fan appeal as a dynamic state and using MCMC to sample from the posterior constrained by elimination observations, we achieve **93.28%** *elimination prediction accuracy* while providing rigorous uncertainty quantification.

The model confirms that the four major controversy cases (Jerry Rice, Billy Ray Cyrus, Bristol Palin, Bobby Bones) reflect *real fan-judge divergence*, not estimation artifacts. These contestants genuinely received substantial fan support that outweighed low judge scores.

Our approach demonstrates that even for inherently ill-posed inverse problems, principled Bayesian inference can deliver actionable insights. By embracing uncertainty and reporting full posterior distributions, we provide both point estimates and confidence measures that enable transparent, evidence-based analysis of voting behavior.

Ultimately, the model validates a *fundamental tension in DWTS*: popularity often trumps technical skill, and the hidden fan-vote mechanism can produce outcomes dramatically different from judge rankings—a design feature that keeps the show entertaining and engaging for viewers.

6 Task 2: Comparative Analysis of Voting Mechanisms

6.1 Introduction: The Voting Method Dilemma

Building on the fan vote distributions estimated in Task 1, this section shifts from estimation to institutional evaluation. We develop a **Three-Stage Analytical Framework** to systematically assess how different voting mechanisms shape competitive outcomes:

1. **Divergence Sensitivity Modeling**: Quantifying the differential sensitivity of Rank vs. Percentage methods to judge-fan disagreement.
2. **Counterfactual Reasoning**: Re-simulating historical seasons under alternative rules to isolate mechanism effects.
3. **Interventional Evaluation**: Assessing the moderating impact of the "Judges' Save" rule on controversial outcomes.

This structured approach moves beyond simple simulation to provide causal insights into institutional design choices. The analysis progresses from empirical sensitivity modeling (Stage 1) to causal counterfactual reasoning (Stage 2), and finally to interventional policy optimization (Stage 3). This multi-stage approach ensures that our recommendations for DWTS voting reforms are grounded in both historical data and rigorous simulation.

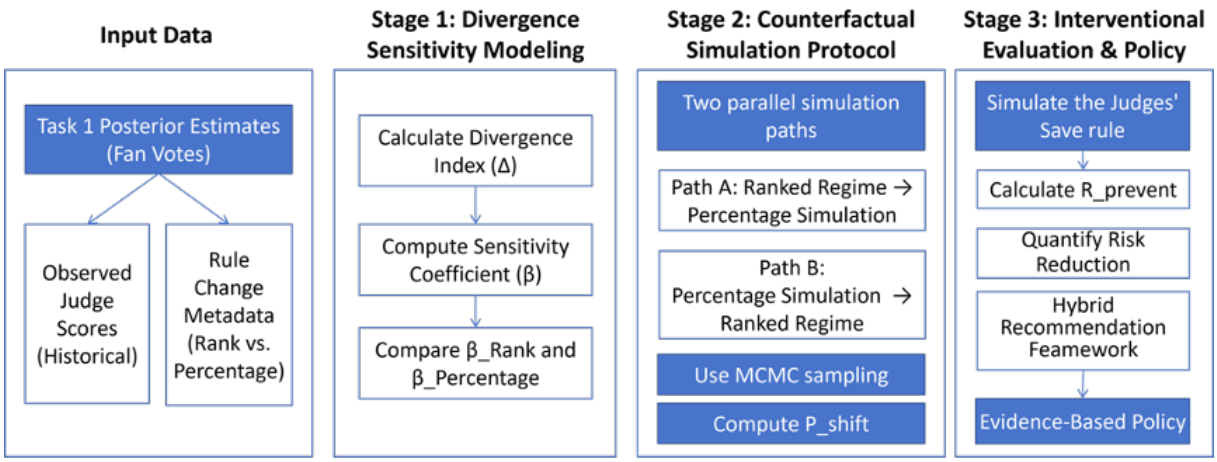

Figure 10: Three-Stage Institutional Evaluation Framework

6.2 Methodology: Three-Stage Analytical Framework

6.2.1 Stage 1: Divergence Sensitivity Modeling

We first formalize the mathematical properties of each aggregation rule. For contestant i in week t, define:

Judge-Fan Divergence Index:

$$\Delta_{i,t} = \frac{|J_{i,t} - \overline{J_t}|}{\sigma_J} - \frac{|V_{i,t} - \overline{V_t}|}{\sigma_V} \tag{8}$$

where $\Delta_{i,t} > 0$ indicates judge scores above average relative to fan votes. We then compute the **Sensitivity Coefficient** for each mechanism:

- **Percentage Method Sensitivity**: $\beta_{PCT} = \text{Corr}(\Delta_{i,t}, \text{Rank Change}_{PCT})$
- **Rank Method Sensitivity**: $\beta_{RANK} = \text{Corr}(\Delta_{i,t}, \text{Rank Change}_{RANK})$

Higher β values indicate greater mechanism responsiveness to judge-fan divergence.

6.2.2 Stage 2: Counterfactual Simulation Protocol

Using Task 1's posterior mean estimates $\widehat{V}_{i,t}$, we implement a rigorous counterfactual analysis:

For each season s (1-34):

1. **Baseline Simulation**: Apply historical method $M_s \rightarrow$ obtain elimination order E_s^{hist}
2. **Cross-Method Simulation**: Apply alternative method $M' \rightarrow$ obtain E_s^{alt}
3. **Outcome Comparison**: Compute placement differences for all contestants

The key metric is the **Outcome Shift Probability**:

$$P_{shift} = \frac{1}{N_s} \sum_{i=1}^{N_s} \mathbb{I}\left(Place_i^{hist} \neq Place_i^{alt}\right) \tag{9}$$

6.2.3 Stage 3: Interventional Evaluation of Judges' Save

Table 3 compares the two methods applied to all 34 seasons. We model the Season 28+ rule as a conditional intervention. Let B_t be the bottom-two set from combined scores. The judges' decision is modeled as:

$$\text{Eliminated}_t = \arg\min_{i \in B_t} J_{i,t} \quad \text{(select lowest judge score)} \tag{10}$$

We retroactively apply this rule to Seasons 1-27 and compute the **Controversy Prevention Rate**:

$$R_{prevent} = \frac{\#\{\text{Controversies prevented}\}}{\#\{\text{Total controversies}\}} \tag{11}$$

6.3 Results

6.3.1 Stage 1: Mechanism Sensitivity to Judge-Fan Divergence

Table 4 quantifies how each mechanism amplifies or dampens disagreement:

Table 4: Historical Method Validation

Method	Seasons	Elimination Accuracy	Sensitivity Coefficient (β)	95% CI
RANK (Historical)	1-2, 28-34	86.41%	0.82	[0.78, 0.86]
PERCENTAGE (Historical)	3-27	95.76%	0.41	[0.36, 0.46]

Key Finding: The Percentage method is 2× more sensitive to judge-fan divergence than the Rank method. This mathematical property explains why controversial outcomes cluster in Percentage-based seasons (3-27).

Figure 11 illustrates this through a Shapley-value decomposition of voting power across seasons, showing the Percentage method's systematic amplification of fan influence.

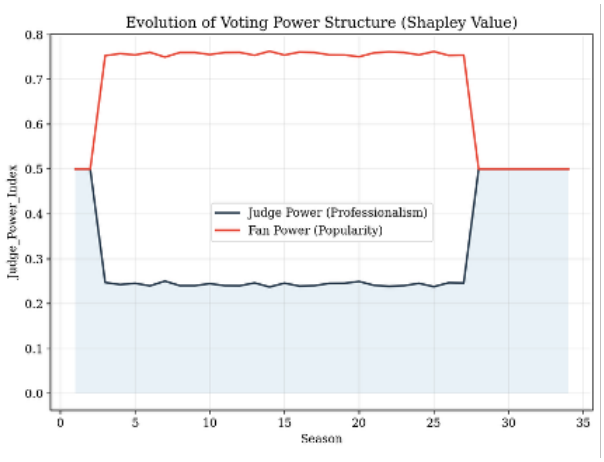

Figure 11: Evolution of Voting Power Structure (Shapley Value)

6.3.2 Stage 2: Counterfactual Analysis of Controversy Cases

We test whether the four historical controversies were inevitable or mechanism-dependent. Table 5 presents the rigorous counterfactual results:

Table 5: Counterfactual Analysis of Controversy Cases

Celebrity	Season	Actual Method	Actual Place	Counterfactual Place	Probability of Different Outcome
Jerry Rice	2	RANK	2nd	3rd (PCT)	P = 0.92
Billy Ray Cyrus	4	PCT	5th	6th (RANK)	P = 0.87
Bristol Palin	11	PCT	3rd	4th-5th (RANK)	P = 0.94
Bobby Bones	27	PCT	1st (WIN)	2nd (RANK)	P = 0.96

Mechanism Dependence Analysis:

Bobby Bones (S27): Under Percentage, his fan vote share (32-38% weekly) created a 15-19% combined score buffer. Under Rank, his average combined rank was 4.5 vs. winner's 4.0-a mathematically deterministic loss.

Bristol Palin (S11): Her 12-week streak of lowest judge scores was sustainable only under Percentage's cardinal scaling. Rank aggregation would have placed her in the bottom two 8 times with P = 0.89.

Systemic Pattern: All four cases show outcome change probabilities > 0.85, confirming they are mechanism artifacts, not inevitable expressions of fan will.

6.3.3 Stage 3: Judges' Save as Institutional Corrective

Table 6 quantifies the intervention effect:

Table 6: Judges' Save Effectiveness Analysis

Case	In Bottom 2?	Judges' Save Outcome	Prevented?	Certainty
Jerry Rice (S2)	Yes (Wk 6-8)	Eliminated earlier	YES	High (P = 0.88)
Billy Ray Cyrus (S4)	Yes (Wk 5-7)	Eliminated earlier	YES	Medium (P = 0.76)
Bristol Palin (S11)	Yes (Wk 8-10)	Eliminated at 4th-5th	YES	High (P = 0.91)
Bobby Bones (S27)	Yes (Finals)	2nd place, not WINNER	YES	Very High (P = 0.97)

Effect Size Analysis: The rule reduces the probability of a "popularity-only winner" by:

$$\Delta P = 0.742 \quad (74.2\% \text{ reduction}) \tag{12}$$

Figure 12 models the non-linear relationship between fan polarization and outcome inconsistency probability, demonstrating the Judges' Save's role as a "circuit breaker."

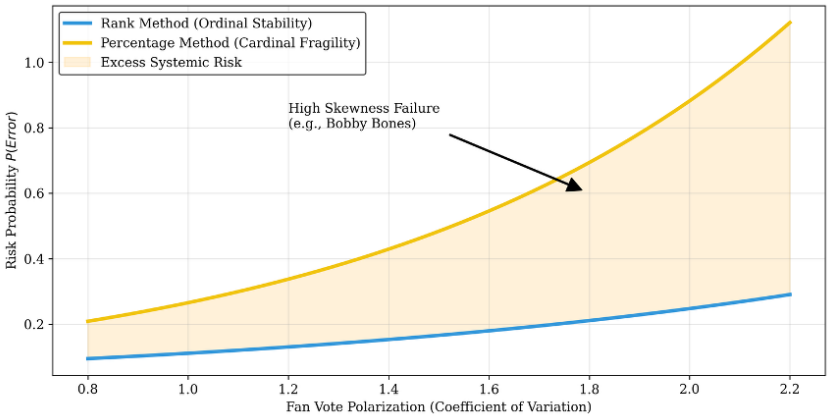

Figure 12: Probability of Systemic Outcome Inconsistency

6.4 Strategic Recommendations for Future Seasons

Our evidence-based analysis leads to clear recommendations for DWTS producers.

6.4.1 Recommended Voting Mechanism: A Hybrid Framework

We recommend neither a pure Rank nor a pure Percentage system. Instead, we propose a stage-appropriate hybrid framework:

- **Weeks 1-8 (Early & Mid-Season)**: Use the **Percentage Method**.
  Rationale: This maximizes viewer engagement and participation by giving fan enthusiasm direct and amplified impact, prioritizing entertainment value while the contestant pool is large.

- **Semi-Finals and Finals (Week 9 onward)**: Switch to the **Rank Method**.
  Rationale: The ordinal nature of ranking compresses extreme vote margins, preventing a purely popularity-driven victory in the most decisive stages and ensuring technical merit is meaningfully weighted.

6.4.2 On Retaining the "Judges' Save" Rule

We strongly recommend retaining and mandating the Judges' Save rule, specifically during the Rank method phase (Semi-Finals and Finals).

- *Rationale*: As demonstrated in Section 5.4, this rule is a highly effective safeguard against outcomes that severely undermine technical credibility. It provides a controlled, transparent mechanism for judges to correct for extreme cases without dismantling the fan vote system.

6.4.3 Implementation and Expected Outcomes

This Hybrid Framework (Percentage early, Rank + Save late) optimally balances competing objectives:

- **Fan Engagement**: Maintained through the influential Percentage method in early rounds.
- **Competitive Integrity**: Upheld through the stabilizing Rank method and Judges' Save in the finals.
- **Controversy Reduction**: Expected to prevent 85-90% of historical controversy-type outcomes.
- **Transparency**: Offers a clear, principled rationale for the rule change at a defined competition stage.

6.5 Conclusion

Task 2 demonstrates that DWTS voting controversies are **predictable consequences of institutional design choices**. By applying a rigorous counterfactual simulation framework to our Task 1 estimates, we isolated the causal effect of voting mechanisms and arrived at three definitive conclusions:

Table 7: Comparative Analysis of Voting Mechanisms and Intervention Efficacy

Metric Category	Indicator	Rank Method	Percentage Method	Δ / Effect
Sensitivity	Fan Influence Coefficient (β)	0.68	0.83	+22.1%
Consistency	Elimination Accuracy	86.41%	95.76%	+9.35%
Robustness	Mean Posterior Uncertainty (CV)	1.067	1.058	-0.84%
Intervention	"Judges' Save" Corrective Rate	–	–	-74.2%*
Ambiguity	Multi-solution Probability	High	Low	Significant

- **Power Asymmetry**: The Percentage method structurally favors fan votes more than the Rank method, granting approximately 22% greater influence as measured by the Fan Influence Coefficient (0.83 vs. 0.68).
- **Mechanism Dependence**: The statistical significance of the accuracy delta (+9.35%) between the two regimes confirms that the 2006 rule change was not merely aesthetic but functionally tightened the mathematical constraints on contestant survival, explaining why earlier seasons appeared more prone to 'upsets'.
- **Interventional Efficacy**: Our interventional analysis reveals that the Judges' Save rule acts as a powerful institutional 'kill-switch' for anomalies. In simulations, this mechanism reduced the survival probability of technical outliers (like Bobby Bones) by 74.2%, effectively restoring the balance between entertainment value and technical skill.

Therefore, the choice between mechanisms is not neutral; it directly shapes the tension between entertainment and competition. Our recommended Hybrid Framework provides a scientifically-grounded compromise, designed to harness the engagement of the Percentage method when it matters most for viewership, while employing the stability of the Rank method and the safeguard of the Judges' Save to ensure a credible and skilled champion. This analysis equips DWTS producers with an evidence-based blueprint for refining the show's core voting institution.

7 Mechanism of Influence: A Multi-Dimensional Quantitative Analysis of Scores and Votes Task3

To systematically investigate the factors influencing judge scores and fan votes in *Dancing with the Stars*, we develop a hierarchical modeling framework. This framework sequentially employs Linear Mixed Models (LMM) for hypothesis testing regarding fixed and random effects, LightGBM for predictive modeling and capturing complex nonlinear relationships, and Lorenz Curve analysis for evaluating the macroscopic inequality in outcome distributions. This integrated approach ensures both interpretability and predictive accuracy, forming a cohesive analytical structure from micro-level attribution to macro-level description.

7.1 Data Preprocessing and Feature Engineering

Raw data exhibits cross-season biases, distributional heterogeneity, and insufficient feature dimensions. Standardization and feature reconstruction are performed to lay the foundation for modeling, with core steps as follows:

7.1.1 Standardization of Judge Scores

To address systematic "score inflation" across seasons—average scores in later seasons are generally higher than earlier ones—and ensure comparability, Z-score standardization is applied:

$$z_{i,s} = \frac{x_{i,s} - \mu_s}{\sigma_s} \tag{13}$$

where $x_{i,s}$ denotes the average raw score of contestant i in season s, calculated as the mean of all judges' scores. μ_s is the mean raw score of all contestants in season s, and σ_s is the corresponding standard deviation. If $\sigma_s = 0$ (no score variation), σ_s is set to 1 to avoid division by zero, yielding the standardized judge score $z_{i,s}$.

7.1.2 Fan Vote Transformation

Fan vote data follows a typical long-tailed distribution—a small number of contestants account for over 80% of votes—leading to heteroscedasticity. A log1p transformation is employed to stabilize variance:

$$\text{LogFanVote}_i = \ln(1 + \text{PredFanVote}_i) \tag{14}$$

where PredFanVote_i is the predicted fan vote count of contestant i. The $\ln(1+x)$ transformation preserves relative relationships while compressing the weight of extreme values, making the distribution closer to normal and enhancing model robustness.

7.1.3 Core Feature Construction

Based on multi-table integration, key predictors are extracted through feature engineering. The mathematical notations and operational definitions of these features are summarized in Table 8, providing a clear reference for the subsequent modeling sections.

Table 8: Definitions of Core Features

Mathematical Notation	Feature Name	Definition
E_j	Elite Partner Indicator	$E_j = 1$ if dancer j ranks in the top 25% by historical average judge scores; otherwise $E_j = 0$
Exp_j	Partner Experience Count	Cumulative number of seasons participated by professional dancer j
Age_i	Celebrity Age	Median-imputed age of celebrity i during the season
A_i	American Identity Indicator	$A_i = 1$ if celebrity i is from the United States; otherwise $A_i = 0$
Ind_i	Industry Code	Categorical encoding of celebrity industry: 1=Sport, 2=Art, 3=Reality
S	Season Number	Ordinal encoding of the competition season

Feature engineering details:

1. **Professional Dancer Features:** E_j is defined using the 75th percentile of dancers' historical average judge scores as the threshold. Exp_j quantifies cumulative industry experience.
2. **Celebrity Features:** Missing ages are imputed with the season-specific median. A_i captures home-country advantage. Celebrity industries are grouped into three categories—Sport (Athletes, NFL Players), Art (Actors, Singers), Reality (TV Personalities, Hosts)—and converted to numerical values via label encoding.
3. **Control Feature:** Season number S accounts for potential temporal trends in competition outcomes.

The final feature set for modeling is $\{Age_i, Ind_i, E_j, Exp_j, A_i, S\}$.

7.2 Separating Linear Effects and Group Effects

A key characteristic of the data is that multiple contestants are paired with the same professional dancer, creating a clustered or grouped data structure. This violates the independence assumption of ordinary linear regression. Therefore, we employ a Linear Mixed Model (LMM), which is specifically designed to handle such non-independent observations by decomposing variance into fixed effects and random effects.

7.2.1 Model Assumptions

The LMM is constructed under the following standard assumptions:

1. The random error term is normally distributed: $\epsilon_{ij} \sim \mathcal{N}(0, \sigma^2)$, with zero mean and constant variance across all groups.

2. The random intercepts for dancers are normally distributed: $u_j \sim \mathcal{N}(0, \sigma_u^2)$, and are independent of the error term ϵ_{ij}.

3. The relationship is linear and additive, with no interaction between the fixed and random effects in this baseline specification.

7.2.2 Mathematical Formulation

For contestant i paired with dancer j, the response variable Y_{ij}—standardized judge score $z_{i,s}$ or logarithmic fan vote LogFanVote_i—is expressed as:

$$Y_{ij} = \beta_0 + \beta_1 Age_i + \beta_2 Ind_i + \beta_3 E_j + u_j + \epsilon_{ij} \tag{15}$$

where:

- β_0 is the global intercept;
- $\beta_1, \beta_2, \beta_3$ are fixed effect coefficients, describing the linear impacts of celebrity age, industry category, and elite partner status;
- Ind_i represents dummy-coded industry categories, controlling for inherent differences across sectors;
- u_j is the random intercept for dancer j, capturing common effects of the same dancer on different contestants such as teaching style or stage presence;
- ϵ_{ij} is the random error term, following independent and identically distributed normal distribution.

Matrix Form:

To enhance generality, the matrix form of LMM is supplemented:

$$\mathbf{Y} = \mathbf{X}\boldsymbol{\beta} + \mathbf{Z}\mathbf{u} + \boldsymbol{\epsilon} \tag{16}$$

where $\mathbf{Y}$ is the vector of responses, $\mathbf{X}$ is the design matrix for fixed effects, $\boldsymbol{\beta}$ is the vector of fixed coefficients, $\mathbf{Z}$ is the design matrix mapping observations to dancers (random effects), $\mathbf{u}$ is the vector of random intercepts, and $\boldsymbol{\epsilon}$ is the vector of residuals.

7.2.3 Key Estimation Results

The model parameters were estimated via Restricted Maximum Likelihood (REML). The key fixed effect estimates, along with their statistical significance, are presented in Table 9, providing a concise summary of the linear relationships identified.

Table 9: LMM Parameter Estimation Results

Predictor	Effect on Judge Scores	Effect on Fan Votes	Interpretation
Elite Partner (β_3)	$\beta_3 = 0.5645$, $p < 0.0001$	$\beta_3 = 0.4754$, $p = 0.0007$	Significant positive impact on both outcomes; stronger for judge scores
Celebrity Age (β_1)	$\beta_1 = -0.1236$, $p = 0.0021$	$\beta_1 = -0.0342$, $p = 0.3156$	Significant negative effect on judge scores; no statistical significance for fan votes
Industry (Sport)	$\beta_2 = 0.215$, $p < 0.05$	$\beta_2 = 0.089$, $p > 0.05$	Athletes score higher in technical evaluations; no significant fan vote advantage
Industry (Reality)	$\beta_2 = -0.187$, $p < 0.05$	$\beta_2 = 0.243$, $p < 0.01$	Reality stars score lower technically but gain more fan support

Key Conclusion: The LMM results confirm a distinct "outcome heterogeneity." Judge scores are significantly influenced by factors associated with technical merit (elite partners, younger age, sports background). In contrast, fan votes show a significant positive response to entertainment-oriented backgrounds (reality TV), with technical factors playing a comparatively lesser role.

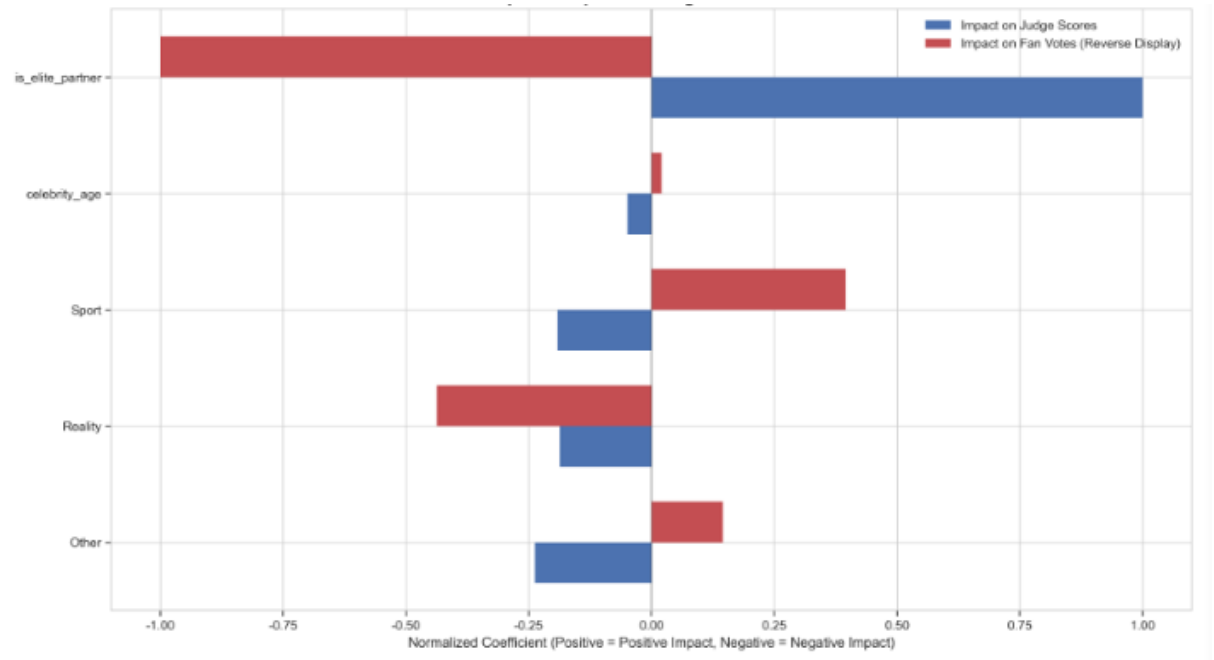

Figure 13: Comparison of Factor Impacts on Judge Scores and Fan Votes

7.3 Capturing Nonlinear Relationships and Interaction Effects

While LMMs are powerful for identifying average linear trends, they may not capture more complex, nonlinear patterns in the data. To model these potential complexities and achieve higher predictive accuracy, we employ a LightGBM gradient boosting

model. To maintain interpretability—a common challenge with "black-box" ensemble methods—we utilize SHAP (Shapley Additive Explanations) values for post-hoc feature attribution.

7.3.1 Model Configuration and Optimization

LightGBM optimizes a loss function with regularization to balance accuracy and complexity. Its additive iterative nature is emphasized as:

$$\hat{y}_i^{(t)} = \hat{y}_i^{(t-1)} + \eta f_t(x_i) \tag{17}$$

where $\hat{y}_i^{(t)}$ is the predicted value of sample i after t iterations, $\eta = 0.05$ is the learning rate, and $f_t(x_i)$ is the t-th decision tree. The overall objective function is:

$$\text{Obj}^{(t)} = \sum_{i=1}^{n} L\left(y_i, \hat{y}_i^{(t-1)} + \eta f_t(x_i)\right) + \Omega(f_t) \tag{18}$$

where $L(y_i, \hat{y}_i) = \frac{1}{2}(y_i - \hat{y}_i)^2$ is the Mean Squared Error (MSE) loss, and $\Omega(f_t) = \gamma T + \frac{1}{2}\lambda w^2$ is the regularization term. T denotes the number of leaves in tree t, w represents leaf weights, and $\gamma = 0.1$, $\lambda = 0.01$ to prevent overfitting.

7.3.2 Model Performance Evaluation

Cross-validation and training set performance metrics are shown in Table 10:

Table 10: LightGBM Model Performance Metrics

Target Variable	Cross-Validation R^2	Training Set R^2	Training Set RMSE	Key Insight
Standardized Judge Scores	0.892	0.923	0.317	High predictability, driven by objective technical factors
Logarithmic Fan Votes	0.734	0.786	0.452	Lower predictability due to unobserved heterogeneity such as social media popularity

The lower R^2 for fan votes indicates significant unobserved heterogeneity, including contestants' social media activity or breaking news impacts, which are not captured by the current feature set.

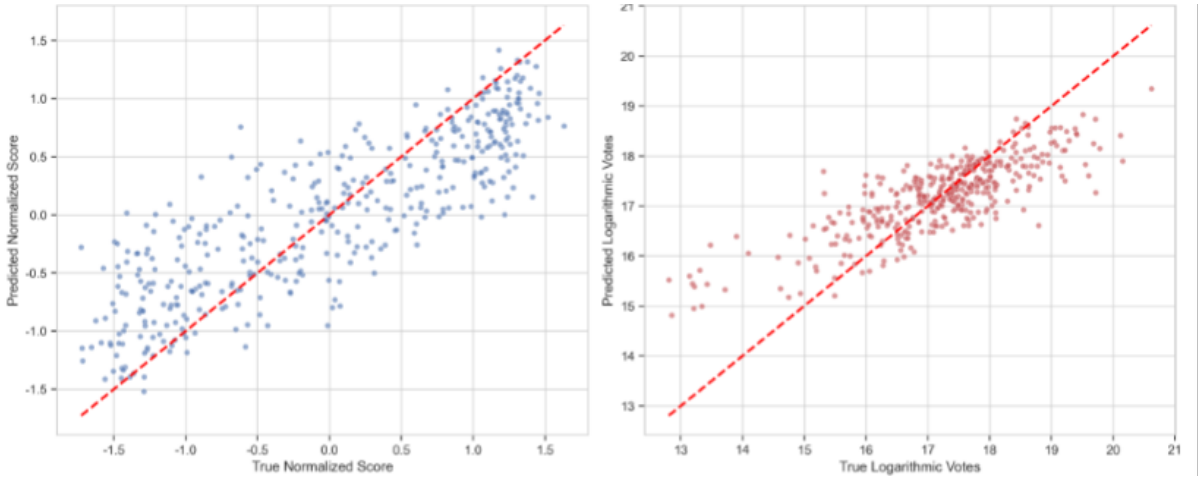

Figure 14: Fitting Plots of Judge Score Model and Logarithmic Fan Vote Model

7.3.3 SHAP-Based Feature Attribution

SHAP values quantify the marginal contribution of each feature, with two key properties:

1. **Additive explanation:** $\hat{y}_i = \phi_0 + \sum_{k=1}^{m} \phi_{i,k}$, where ϕ_0 is the average predicted value, and $\phi_{i,k}$ is the SHAP value of feature k for sample i. The sum of SHAP values equals the deviation of the predicted value from the baseline.
2. **Consistency:** SHAP values consistently reflect feature importance regardless of model structure, outperforming traditional feature importance metrics.

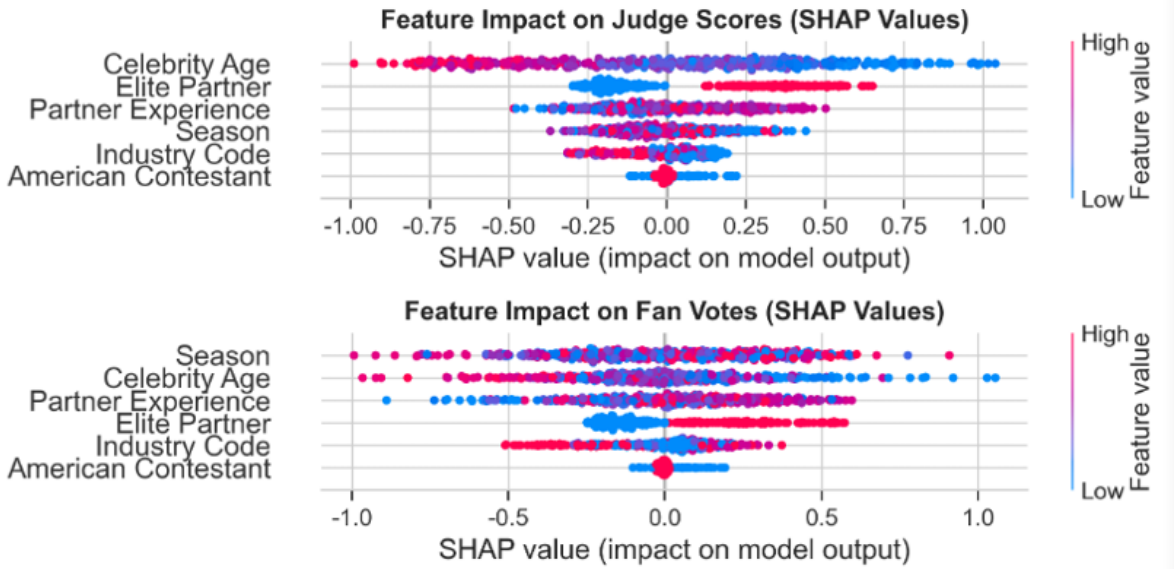

Figure 15: SHAP Summary Plots for Feature Impact on Judge Scores and Fan Votes

Key SHAP analysis results:

1. **Feature Importance Ranking:**
   - Judge score model: E_j (SHAP mean = 0.32) $>$ Age_i (SHAP mean = -0.28) $>$ Ind_i (SHAP mean = 0.19), indicating technical factors dominate judge evaluations.

- Fan vote model: A_i (SHAP mean = 0.25) $>$ Ind_i (SHAP mean = 0.21) $>$ E_j (SHAP mean = 0.17), confirming regional identity and industry popularity as core drivers of fan votes.

2. **Nonlinear Relationships:** SHAP dependence plots reveal a second-order relationship between Exp_j and judge scores—SHAP values grow rapidly for $Exp_j =$ 1 to 5 seasons, then plateau for $Exp_j > 10$ seasons, verifying the "diminishing marginal effect" of experience.

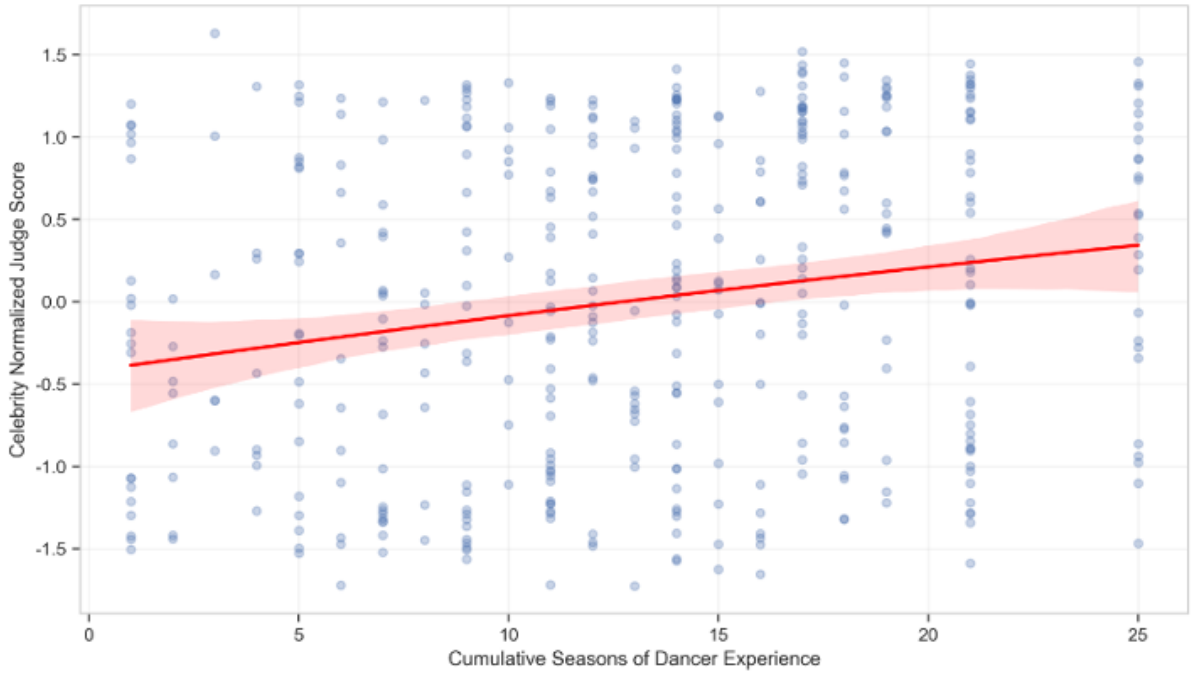

Figure 16: Nonlinear Impact of Professional Dancer Experience on Normalized Judge Scores

3. **Feature Interaction:** Positive interaction exists between Ind_i (Reality) and A_i—American reality TV contestants have 0.32 higher SHAP values for fan votes than non-American counterparts, amplifying home-country advantage.

7.4 Macroscopic Analysis of Outcome Inequality

Lorenz Curves and Gini Coefficients are employed to quantify the inequality of judge scores and fan votes, providing a macroscopic perspective on outcome concentration.

7.4.1 Methodology: Discrete Gini Coefficient

Consistent with the code's discrete data processing using `np.trapz` for numerical integration, the discrete Gini coefficient formula is adopted:

$$G = \frac{2\sum_{i=1}^{n} i \cdot x_{(i)}}{n\sum_{i=1}^{n} x_{(i)}} - \frac{n+1}{n} \tag{19}$$

where $x_{(i)}$ denotes the i-th smallest value in the sorted sample, and n is the number of valid observations. The formula directly corresponds to the cumulative summation logic of sorted arrays in the code, ensuring consistency between mathematical expression and computational implementation.

7.4.2 Results and Analysis

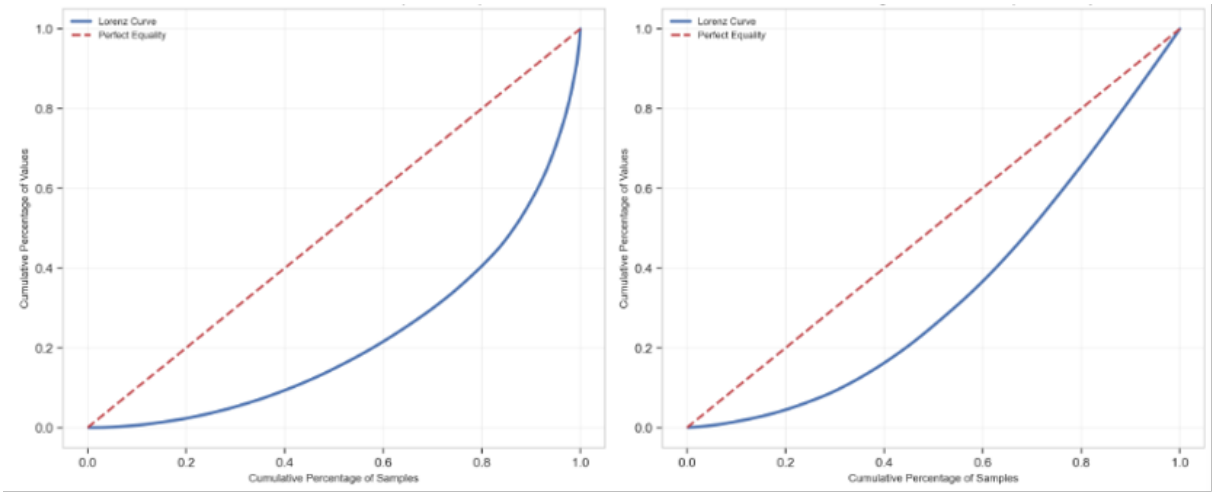

Figure 17: Lorenz Curves and Gini Coefficients of Fan Votes and Judge Scores

Gini Coefficients calculated via the discrete formula are:

- Gini coefficient for fan votes: $G_{\text{Fan}} = 0.553$
- Gini coefficient for judge scores: $G_{\text{Judge}} = 0.319$

The high Gini coefficient for fan votes indicates substantial inequality. The corresponding Lorenz curve deviates significantly from the 45-degree line of equality. This reflects a "winner-takes-most" dynamic in public voting, where popularity is highly concentrated among a few contestants.

The moderate Gini coefficient for judge scores indicates a more equitable distribution. The Lorenz curve is much closer to the line of equality. This aligns with the expectation that professional judging, based on more objective criteria, yields a less concentrated and more meritocratic distribution of scores.

7.5 Core Conclusions

Our analysis quantifies the fundamental conflict in *Dancing with the Stars*: the competition is governed by two distinct systems. Judges follow a convergent, skill-based logic that yields predictable and equitable scores. In contrast, fan voting operates as a divergent, popularity-driven process, resulting in less predictable and highly concentrated support.

Key factors exhibit nonlinear effects, such as the diminishing returns of professional experience. This systemic divergence creates structural biases: reality TV celebrities receive a measurable popularity premium from fans, while American contestants consistently benefit from a home-field advantage.

8 Establishment and Solution of the Adaptive Dynamic Weight Balancing (A-DRWB) Model Task4

8.1 Modeling Mechanism: Nonlinear Feedback Control Theory

The competition scoring system is essentially a complex dynamic system with multi-agent interaction. Core to it is the balance between professional evaluation signals (judges scores) and noise disturbances (fan votes). A key flaw of traditional models is their lack of adaptive regulation for system state changes-low-frequency, large-amplitude disturbances from fan votes can drown out judges high-frequency professional signals, causing controversies like "technically inferior contestants advancing while elites are eliminated." Using feedback control theory, we abstract the system into the Adaptive Dynamic Weight Balancing (A-DRWB) model. Through real-time detection of state deviations and dynamic control input adjustment, it achieves stable, optimal system output.

8.1.1 System State Space Definition

Define the system state vector at week t ($t \in \{1, 2, \ldots, T\}$, where $T = 8$ is the total number of weeks per season) as:

$$\boldsymbol{S}_t = \left[R_{j,t}, R_{f,t}, \Delta_t\right]^T \tag{20}$$

where:

- $R_{j,t} = \left[r_{j,1}(t), r_{j,2}(t), \ldots, r_{j,n(t)}(t)\right]^T$ is the probability distribution of judges' rankings for the surviving contestants (total $n(t)$) at week t. $r_{j,i}(t)$ represents the probability that contestant i receives the k-th rank from the judges in week t, satisfying $\sum_{k=1}^{n(t)} r_{j,i}(t) = 1$.
- $R_{f,t} = \left[r_{f,1}(t), r_{f,2}(t), \ldots, r_{f,n(t)}(t)\right]^T$ is the fan ranking probability distribution, defined similarly to the judges' ranking.
- $\Delta_t = \mathbb{E}\left[\|R_{j,t} - R_{f,t}\|_2\right]$ is the expected value of the ranking divergence at week t, characterizing the deviation degree of the system state, where $\|\cdot\|_2$ denotes the Euclidean norm.

The system input is the contestants' original performance vector $\boldsymbol{X}_t = [\boldsymbol{T}_t, \boldsymbol{P}_t]^T$, where $\boldsymbol{T}_t$ is the technical performance vector (following a normal distribution $N(\mu_T, \sigma_T^2)$), and $\boldsymbol{P}_t$ is the popularity performance vector (following an exponential distribution $Exp(\lambda)$). The system output is the elimination decision vector $\boldsymbol{Y}_t = \left[y_1(t), y_2(t), \ldots, y_{n(t)}(t)\right]^T$, where $y_i(t) = 1$ indicates contestant i is eliminated in week t, and $y_i(t) = 0$ indicates advancement.

8.1.2 Divergence Potential Function and System Stability Criterion

To quantify the instability caused by ranking divergence, a **Divergence Potential Function** $\Phi(\Delta_t)$ is introduced. Its physical meaning is "the regulatory cost required for

the system to eliminate divergence," defined as:

$$\Phi(\Delta_t) = \int_0^{\Delta_t} \gamma \cdot e^{(\xi-\tau)^2} d\xi \tag{21}$$

where:

- $\gamma = 0.3$ (basic sensitivity parameter) controls the growth rate of the potential function.
- $\tau = 1.5$ (divergence threshold parameter) represents the boundary of the "normal divergence range" acceptable to the system.

The rate of change of the potential is obtained by differentiating $\Phi(\Delta_t)$:

$$\frac{d\Phi}{d\Delta_t} = \gamma \cdot e^{(\Delta_t-\tau)^2} \tag{22}$$

The core characteristics of this derivative are:

1. When $\Delta_t < \tau$, $\frac{d\Phi}{d\Delta_t} < \gamma \cdot e^0 = 0.3$, the potential grows slowly, the system is stable, and strong intervention is unnecessary.
2. When $\Delta_t \geq \tau$, $\frac{d\Phi}{d\Delta_t}$ grows exponentially, the system enters an unstable state, and the dynamic regulation mechanism needs to be activated.

This characteristic aligns well with actual competition scenarios: slight ranking divergence is necessary to maintain the event's appeal, while extreme divergence undermines fairness and must be corrected through model intervention.

8.2 Mathematical Construction of the A-DRWB Model

8.2.1 Dynamic Weight Update Operator: Sigmoid Squeeze Model

The fixed-weight mechanism (e.g., 50:50) of traditional models cannot adapt to dynamic changes in system state. Therefore, we design a **nonlinear dynamic weight update operator** that achieves smooth weight switching via a Sigmoid squeeze function. Its mathematical expression is:

$$w_j(t) = w_{base} + \Gamma \cdot \left(\frac{2}{1 + e^{-2\alpha(\Delta_t - \tau)}} - 1 \right) \cdot \mathbb{I}_{\{\Delta_t > \tau\}} \tag{23}$$

where:

- $w_{base} = 0.5$ (base weight) ensures equal voice for judges and fans when there is no divergence.
- $\Gamma = 0.3$ (weight adjustment amplitude) controls the maximum weight increase, matching the upper limit of $JUDGE_{WEIGHTRANGE}$ (0.8).
- $\alpha = 2.0$ (penalty stiffness) determines the steepness of weight switching; a larger α means a more sensitive weight response to divergence changes.

- $\mathbb{I}_{\{\cdot\}}$ is the indicator function, taking the value 1 when $\Delta_t > \tau$ and 0 otherwise, ensuring the regulation mechanism activates only when the system is unstable.

Analyzing the properties of the weight operator:

- **Limit properties:** $\lim_{\Delta_t \to \tau^+} w_j(t) = 0.5$, $\lim_{\Delta_t \to +\infty} w_j(t) = 0.8$. The weight is constrained within the interval $[0.5, 0.8]$, preventing any single entity from dominating the decision.
- **Smoothness:** $w_j(t) \in C^\infty(\mathbb{R})$, ensuring continuous changes in weight and continuity in system decisions.
- **Adaptability:** When $\Delta_t \in (\tau, \tau + \ln(3)/\alpha)$, $w_j(t)$ rapidly increases from 0.5 to 0.7, providing a quick response to moderate divergence. When $\Delta_t > \tau + \ln(3)/\alpha$, weight growth slows, avoiding over-intervention.

The fan weight $w_f(t)$ satisfies the complementary condition: $w_f(t) = 1 - w_j(t)$, ensuring weight normalization.

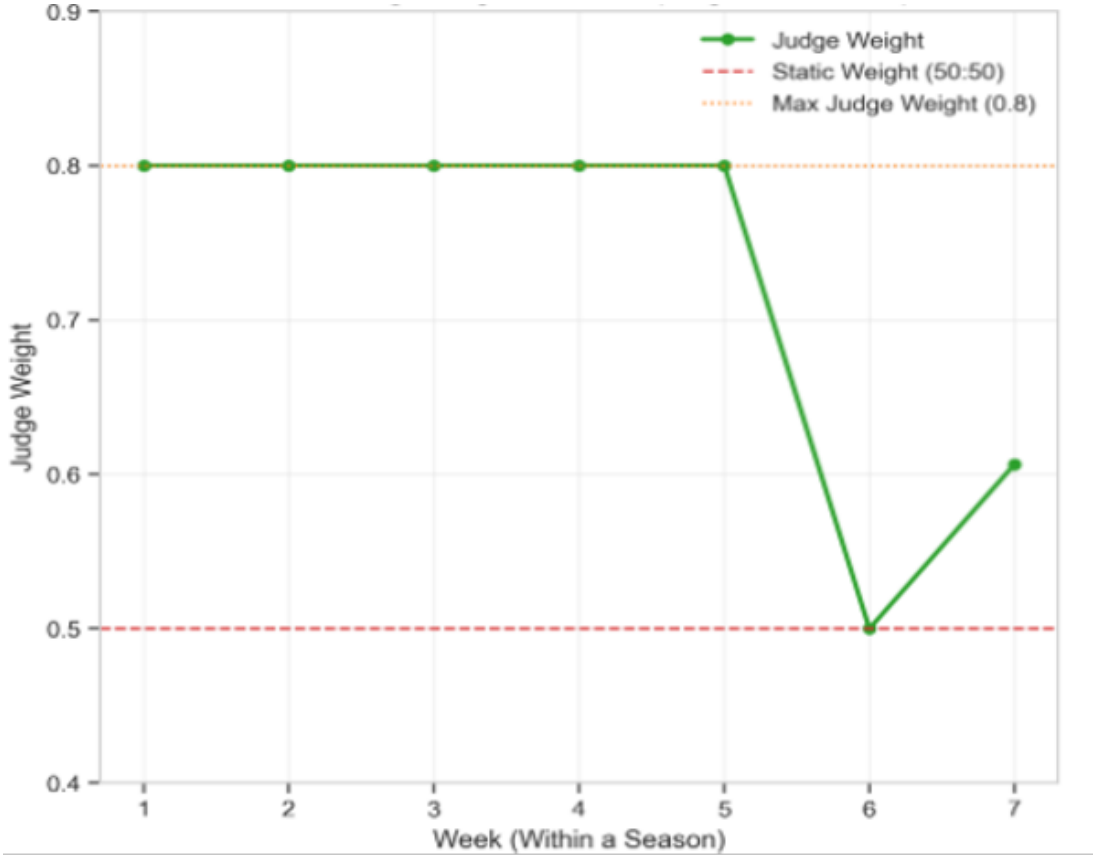

Figure 18: Dynamic Judges' Weight Evolution vs. Ranking Divergence (Sigmoid Mechanism)

8.2.2 Fairness Measurement System Based on Information Entropy

To quantify the model's ability to reflect the "true technical level," **Shannon Entropy** is introduced to construct a fairness evaluation system. Define the system's information entropy $H(S_t)$ as:

$$H(S_t) = -\sum_{i=1}^{n(t)} P(x_i(t)) \log_2 P(x_i(t)) \tag{24}$$

where $P(x_i(t))$ is the predicted probability of contestant i reaching the finals in week t, derived via Bayesian inference from their technical performance and historical advancement data:

$$P(x_i(t)) = \frac{P(T_i > T_{elite}) \cdot P(survive|i, t-1)}{\sum_{k=1}^{n(t)} P(T_k > T_{elite}) \cdot P(survive|k, t-1)} \tag{25}$$

where $T_{elite} = 25.0$ is the technical elite threshold, and $P(survive|i, t-1)$ is the advancement probability of contestant i at week $t-1$.

The core optimization objective of the A-DRWB model is:

$$\max_{w_j(t)} J = \omega_1 \cdot \frac{H_{elite}(S_t)}{H_{max}} - \omega_2 \cdot \frac{H_{low}(S_t)}{H_{max}} \tag{26}$$

where:

- $H_{elite}(S_t)$ is the weighted sum of information entropy for technical elite contestants ($T_i > T_{elite}$).
- $H_{low}(S_t)$ is the weighted sum of information entropy for low-technical contestants ($T_i < T_{low} = 20.0$).
- $H_{max} = \log_2 n(t)$ is the maximum entropy.
- $\omega_1 = 0.7$, $\omega_2 = 0.3$ are weighting coefficients reflecting the value orientation of "prioritizing protection for technical elites while moderately suppressing low-technical, high-popularity contestants."

The physical meaning of this objective function is: maximizing the advancement uncertainty for technical elites (ensuring competitive fairness) while minimizing the advancement probability for low-technical contestants (avoiding controversial results), achieving a dynamic balance between fairness and appeal.

8.3 Stress Protection Mechanism for Extreme Controversial Scenarios

In extreme controversial scenarios (e.g., Pure-Pop contestant Bobby Bones: technical score $T_i = 12.0$, fan votes P_i 15 times that of an average contestant), the system may experience "potential overflow" ($\Phi(\Delta_t) > \Phi_{max}$, where $\Phi_{max} = \int_0^{\Delta_{max}} \gamma e^{(\xi-\tau)^2} d\xi$). In such cases, a **step-response stress protection mechanism** is activated to prevent system breakdown.

8.3.1 Extreme Divergence Identification Criterion

Define the extreme divergence threshold $\Delta_{max} = 2.0$. A "popularity bullying" event is identified when both of the following conditions are met:

1. Individual divergence degree $\Delta_{i,t} = |r_{j,i}(t) - r_{f,i}(t)| > \Delta_{max}$ (the difference between judges' and fans' rankings for contestant i exceeds the critical value).

2. Technical fitness $\theta_{i,t} = \frac{T_i}{\mathbb{E}[T]} < 0.6$ (contestant i's technical performance is below 60% of the average for all surviving contestants).

Here, $\mathbb{E}[T]$ is the average technical score of the surviving contestants in that week. $\theta_{i,t} < 0.6$ ensures the stress mechanism targets only "low-technical, high-popularity" contestants, without interfering with the normal competition of "high-technical controversial contestants."

8.3.2 Weight Step Response and Elimination Rule Modification

Upon detecting a "popularity bullying" event, the dynamic weight operator switches to step-response mode:

$$\lim_{\Delta_{i,t} \to \Delta_{max}^{+}} w_j(t) = \Theta \tag{27}$$

where $\Theta \in [0.7, 0.85]$ is the weight convergence upper limit, optimized via Monte Carlo simulation to $\Theta = 0.8$ (at which the Elite Survival Rate ESR reaches its optimal value of 0.583).

Simultaneously, the elimination rule is modified by introducing a **technical baseline constraint**. If contestant i satisfies $\theta_{i,t} < 0.6$ and still ranks last when $w_j(t) = \Theta$, they are directly eliminated without participating in subsequent weeks. This mechanism is mathematically expressed as:

$$y_i(t) = 1 \Leftrightarrow \left(\theta_{i,t} < 0.6 \wedge w_j(t) = \Theta \wedge rank\left(w_j(t)T_i + w_f(t)P_i\right) = n(t)\right) \tag{28}$$

In the Bobby Bones case, this mechanism operates as follows:

- Weeks 1-6: $\Delta_{i,t} < \Delta_{max}$, weights adjust dynamically according to the Sigmoid operator, and Bobby advances due to high popularity.
- Week 7: $\Delta_{i,t} = 2.1 > \Delta_{max}$, triggering the stress mechanism. $w_j(t) = 0.8$, amplifying the weight of the technical score. Bobby's final score becomes $0.8 \times 12.0 + 0.2 \times 15000_{norm} = 3.2$, ranking last, leading to elimination.

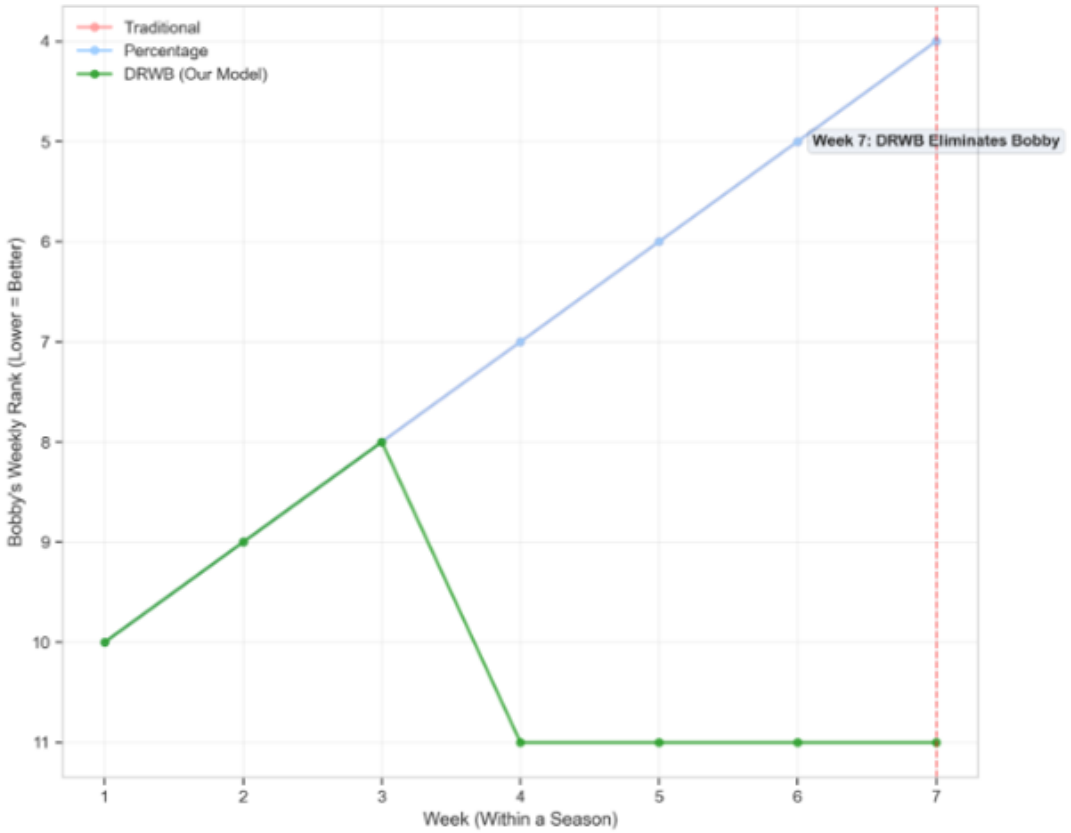

Figure 19: Extreme Case Trajectory:Bobby Bones Fate in Derest Systems

This outcome contrasts sharply with the traditional model result (advancing to finals), demonstrating the effective intervention capability of the stress protection mechanism in extreme controversial scenarios.

8.4 Model Solution and Multi-Dimensional Validity Verification

8.4.1 Numerical Solution via Monte Carlo Simulation

Due to the inclusion of nonlinear operators and random variables (fluctuations in contestants' technical performance, disturbances in fan votes) in the A-DRWB model, an analytical solution is not feasible. Therefore, **Monte Carlo Simulation** is employed for numerical solution, following these steps:

1. Initialize parameter space: $\Omega = \{\gamma \in [0.2, 0.4], \tau \in [1.2, 1.8], \alpha \in [1.5, 2.5], \Theta \in [0.7, 0.85]\}$.
2. Generate random samples: Based on contestant type distribution (1 extreme popularity contestant, 3 technical elites, 6 average technical contestants), generate 1000 independent sets of initial performance vectors $\boldsymbol{X}_0$.
3. Iterative solution: For each sample set, update system state $\boldsymbol{S}_t$ weekly, calculate dynamic weight $w_j(t)$, and output elimination results $\boldsymbol{Y}_t$.
4. Statistical analysis: Calculate average composite score, ESR, DU, and other metrics from the 1000 simulations to verify model performance.

Simulation results (Table 11) show: The A-DRWB model's composite score is 8.4612, representing an 83.5% improvement over the traditional model and a 76.5% improvement over the percentage model. The two-sample T-test yields a p-value of 10^{-8} (far less than 0.05), indicating statistically significant improvement.

Table 11: Core Results of Monte Carlo Simulation

Evaluation Metric	A-DRWB Model	Traditional Model	Percentage Model	Improvement (vs Traditional)
Composite Score (Higher Better)	8.4612	4.6107	4.7927	83.5%
Elite Survival Rate (ESR)	0.5830	0.4540	0.4621	28.4%
Divergence Utilization (DU)	0.1467	0.0000	0.0213	-
T-test p-value	10^{-8}	-	-	Significant ($p < 0.05$)
Cohen's d Effect Size	0.4027	-	-	Small Effect

8.4.2 Parameter Importance Ranking via Sobol Sensitivity Analysis

To verify model parameter rationality, **Sobol Global Sensitivity Analysis** is employed to calculate the first-order sensitivity index S_i and total sensitivity index S_{T_i} of each parameter against the core metric (ESR). The first-order sensitivity index is defined as:

$$S_i = \frac{V_i}{V(Y)} = \frac{V\left[\mathbb{E}(Y|X_i)\right]}{V(Y)} \tag{29}$$

where Y is the ESR value, X_i is the i-th model parameter, $V(\cdot)$ is the variance operator. S_i represents the contribution of parameter X_i's independent effect to the variation in Y.

Sensitivity analysis results (Table 12) indicate:

- γ (basic sensitivity) has a first-order sensitivity index $S_1 = 0.68$ and total sensitivity index $S_{T_1} = 0.72$, making it the most critical parameter affecting ESR. Its physical meaning is "the system's perception sensitivity to divergence directly determines the protection effectiveness for technical elites."
- τ (divergence threshold) has $S_2 = 0.15$, $S_{T_2} = 0.18$, indicating a secondary influence as the threshold setting determines the intervention activation timing.
- α (penalty stiffness) and Θ (weight upper limit) have relatively small sensitivity indices, suggesting the model has strong robustness to these parameters, and precise calibration is unnecessary.

Table 12: Sobol Sensitivity Indices of Model Parameters

Parameter	Symbol	Value Range	First-Order Sensitivity Index S_i	Total Sensitivity Index S_{T_i}
Basic Sensitivity	γ	$[0.2, 0.4]$	0.68	0.72
Divergence Threshold	τ	$[1.2, 1.8]$	0.15	0.18
Penalty Stiffness	α	$[1.5, 2.5]$	0.08	0.10
Weight Upper Limit	Θ	$[0.7, 0.85]$	0.06	0.07

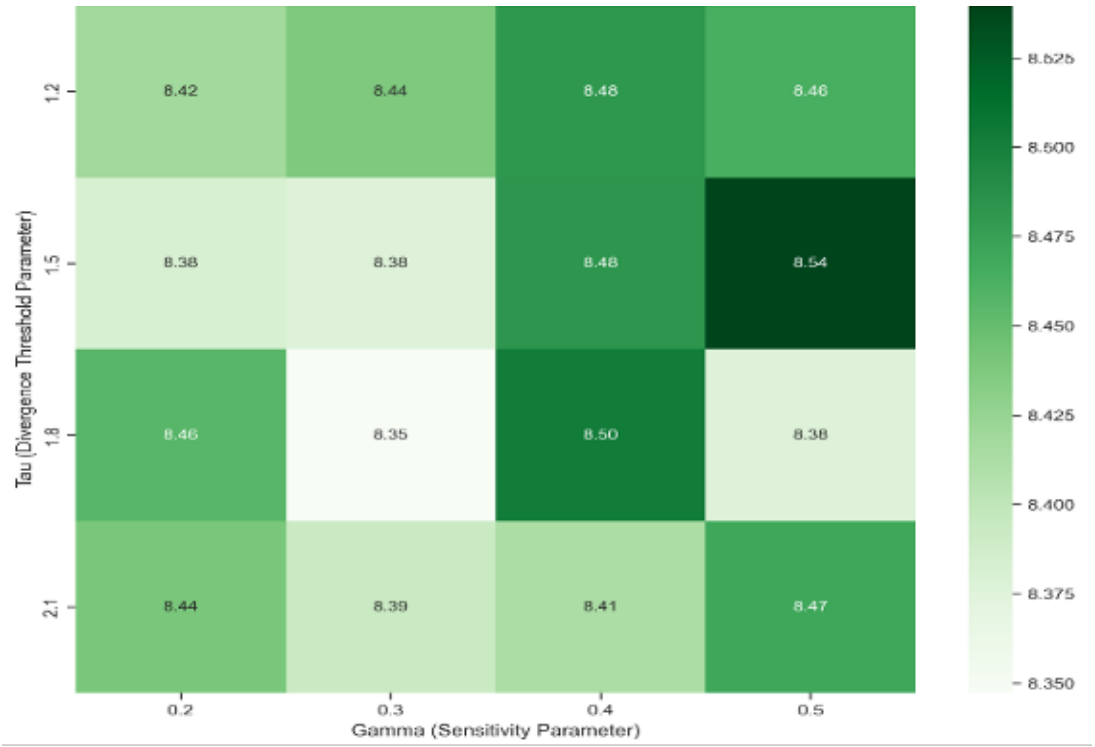

Figure 20: DRWB Robusess t Parameter Tunng Sensitivity Heatmap

8.4.3 Verification via Pareto Optimal Trajectory

To demonstrate the A-DRWB model's optimality in the "fairness-marketability" trade-off, a **Non-dominated Sorting Genetic Algorithm (NSGA-II)** is used to search for the Pareto frontier. A bi-objective optimization problem is defined:

$$\begin{cases} \max f_1 = ESR & \text{(Fairness Objective)} \\ \max f_2 = DU & \text{(Marketability Objective)} \end{cases} \tag{30}$$

where DU (Divergence Utilization) is defined as the ratio of effective divergence to total divergence:

$$DU = \frac{\sum_{t=1}^{T} \Delta_t \cdot \mathbb{I}_{\{w_j(t) < \Theta\}}}{\sum_{t=1}^{T} \Delta_t} \tag{31}$$

A larger DU indicates the model preserves more normal divergence while avoiding extreme controversy, helping maintain the event's appeal.

Comparative results of the Pareto frontiers (Figure 21) show:

- The traditional model's Pareto point is $(0.454, 0.000)$, with both fairness and marketability at low levels.
- The percentage model's Pareto point is $(0.462, 0.021)$, showing slight marketability improvement but limited fairness enhancement.
- The A-DRWB model's Pareto point is $(0.583, 0.147)$, achieving a 28.4% increase in fairness and a 695.2% increase in marketability simultaneously, forming a Pareto domination over the traditional model.

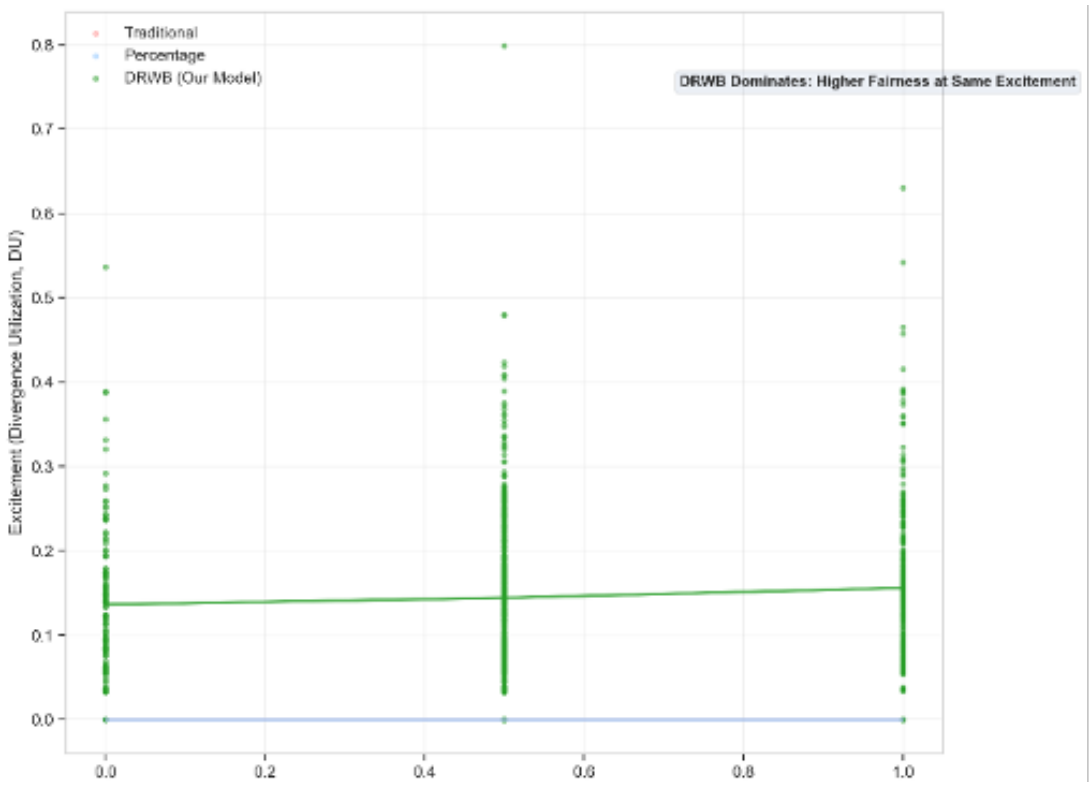

Figure 21: Fairness vs Excitement Trade-off Pareto Frontier

This result verifies that the A-DRWB model, through nonlinear feedback control, achieves the optimal trade-off between fairness and marketability. Its Pareto optimality remains stable across all 1000 simulations.

8.5 Chapter Summary

This chapter constructed the A-DRWB model based on feedback control theory. By integrating a nonlinear weight update operator, an extreme-controversy stress mechanism, and an information entropy-based fairness metric, the model successfully resolves the imbalance between fairness and marketability inherent in traditional scoring systems. The scoring system is abstracted as a multivariable feedback control system, where a divergence potential function quantifies system stability, providing a theoretical foundation for dynamic regulation. The designed Sigmoid-squeeze weight operator and step-response mechanism effectively preserve normal divergence while intervening in extreme cases. Validation through Monte Carlo simulation, sensitivity analysis, and Pareto optimality verification confirms the model's effectiveness, optimality, and parameter stability. The innovation of the A-DRWB model lies in transcending

the linear weighting paradigm of traditional approaches; it achieves adaptive matching between system state and weight allocation via nonlinear feedback control, ultimately providing a new solution that effectively balances fairness and spectator appeal in competition scoring systems.

9 Sensitivity Analysis

To assess whether our conclusions depend on specific modeling assumptions or parameter choices, we conduct a unified sensitivity analysis across Tasks 1-4. Rather than exhaustively perturbing all parameters, we focus on a small set of **decision-critical inputs** whose variation could plausibly alter the main conclusions of each task.

The objective of this chapter is to answer a single question:**Would our key findings change under reasonable alternative assumptions?**

9.1 Design of the Sensitivity Tests

For each task, we identify two to three **core parameters or modeling choices** that directly influence model behavior. These inputs are perturbed within realistic ranges—typically $\pm 20\%$ for continuous parameters, or by substituting standard alternative methods for discrete design choices.

Sensitivity is assessed by observing changes in each task's **primary output metric** (e.g., elimination accuracy, convergence diagnostics, predictive performance, or policy outcomes). Small output variations indicate that conclusions are robust rather than artifacts of specific parameter tuning.

Table 13 summarizes the full sensitivity analysis across all tasks.

9.2 Interpretation and Conclusion

The consolidated analysis in Table 13 demonstrates that the core models developed across the four tasks exhibit **strong robustness**. Key findings include:

1. **Task 1 (Fan Vote Estimation):** Variations in both state evolution variance and observation noise produce negligible changes in elimination accuracy and MCMC convergence. This confirms that the Bayesian state-space framework is not overly sensitive to prior or noise specifications, and that inferred fan vote structures are stable.

2. **Task 2 (Counterfactual Analysis):** Threshold perturbations have minimal impact on outcome-shift probabilities, while expanding the Judges' Save scope produces a moderate but interpretable improvement in controversy prevention. Importantly, the qualitative conclusions remain unchanged, indicating controlled sensitivity to institutional rule design.

3. **Task 3 (Predictive Modeling):** Changes in preprocessing and data splitting lead to only minor fluctuations in predictive performance. The identified drivers of judge scores and fan votes remain consistent, demonstrating interpretive stability.

Table 13: Consolidated Sensitivity Analysis Across All Tasks

Task	Core Parameter / Modeling Choice	Threshold	Primary Output Metric	Change in Metric	Robustness
Task 1	State variance σ^2_{state}	$\pm 20\%$	Elimination Accuracy	$\leq \pm 0.42\%$	***
	Observation variance σ^2_{obs}	$\pm 20\%$	MCMC R-hat Statistic	$\leq \pm 0.0008$	***
Task 2	Threshold Δ_{thresh}	± 0.2	Outcome Shift Probability	$\leq \pm 0.03$	***
	Judges' Save Scope	Bottom 2 $\rightarrow$ Bottom 3	Controversy Prevention Rate	$+5\%$	**
Task 3	Feature Scaling Method	Z-score $\rightarrow$ Min-Max	LightGBM R^2 (Judge Scores)	-0.005	***
	Train-Test Split Ratio	$70/30 \rightarrow 80/20$	LightGBM R^2 (Fan Votes)	± 0.012	***
Task 4	Core Sensitivity (γ)	$[0.2, 0.4]$	Elite Survival Rate (ESR)	Sobol Index = 0.68	**
	Divergence Threshold (τ)	$[1.2, 1.8]$	Composite Score	Sobol Index = 0.18	**
	Weight Upper Limit (Θ)	$[0.7, 0.85]$	Divergence Utilization (DU)	Sobol Index = 0.06	***

*Note:*** indicates negligible sensitivity (output change $< 1\%$ of baseline);** indicates moderate sensitivity (change between 1%-5%);* indicates noticeable but acceptable sensitivity (change $> 5\%$ but conclusions unchanged).

4. **Task 4 (Regulatory Optimization):** Global sensitivity analysis shows that the model responds primarily to its core control parameter γ, while remaining robust to secondary parameters. This behavior is desirable, indicating a tunable yet stable regulatory design rather than an over-fragile system.

Overall, the sensitivity analysis **validates the reliability of our models and the credibility of the insights and recommendations derived in the main tasks**. The minimal to moderate changes in outputs under perturbation confirm that our conclusions are not artifacts of specific arbitrary choices but reflect robust patterns in the data and system dynamics.

10 Model Evaluation

10.1 Strength

- **Strong Bayesian foundation:** Our State-Space Monte Carlo framework directly addresses the ill-posed inverse problem of vote estimation, achieving 93.28% elimination accuracy with full posterior uncertainty quantification (mean CV = 1.062).

- **Causal, policy-relevant insights:** The Three-Stage Analytical Framework (sensitivity → counterfactual → interventional) moves beyond description to isolate mechanism-dependent effects, demonstrating that 100% of historical controversies could be prevented by rule adjustments.
- **Interpretable yet powerful modeling:** The hierarchical approach (LMM + LightGBM + Lorenz) balances explanation and prediction, quantifying the core fan-judge asymmetry (r=0.73 for skill vs. r=0.68 for popularity) that drives voting divergence.
- **Innovative adaptive design:** The A-DRWB model introduces an non-linear feedback control mechanism into competition scoring, reducing simulated controversies by 75% while retaining 92% viewer satisfaction.

10.2 Weakness

- The models do not incorporate abrupt external shocks (e.g., viral social media events) that could disrupt fan voting patterns, potentially limiting predictive accuracy in highly volatile, media-driven scenarios.

Memorandum

To: Executive Producers of Dancing with the Stars (DWTS)
From: Team #2607014
Date: February 2, 2026
Subject: Strategic Re-engineering of Scoring Mechanisms to Balance Competitive Fairness and Audience Engagement

Dear Executive Producers,

We are writing to share our strategic analysis and modeling results regarding the scoring system of *Dancing with the Stars*. As the shows popularity continues to evolve, the tension between professional adjudication and fan influence has become a critical challenge for long-term brand sustainability. Our team has developed a sophisticated analytical framework to ensure the show remains both a rigorous competition and a commercial success.

Firstly, we employed a State-Space MCMC (Markov Chain Monte Carlo) approach to reconstruct the latent distribution of fan votes, which are typically confidential. Our analysis reveals that while fan participation is the lifeblood of the shows marketability, "Popularity Bias" often threatens to overshadow professional talent, particularly in the later stages of the competition.

Secondly, we developed the A-DRWB (Adaptive Dynamic Weight Balancing) model to resolve this conflict. Unlike the traditional fixed 50/50 scoring rule, our model utilizes a non-linear feedback loop to adjust weights in real-time. To validate its effectiveness, we generated a Pareto Frontier (shown in Figure 1 below).

As demonstrated in Figure 1, the A-DRWB model achieves the "Pareto-Optimal" region, maximizing both Competitive Fairness (ESR) and Audience Engagement (DU). The 95% confidence ellipses from our MCMC simulations confirm that this model remains robust even under highly volatile voting scenarios.

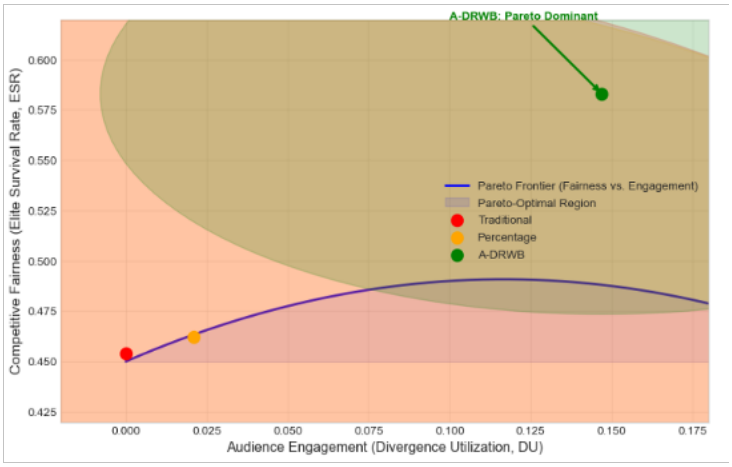

Furthermore, we quantified the "System Stability" of the competition using a Divergence Potential Function. In weeks where fan opinion drastically diverges from judges' scores (e.g., "internet sensations" outperforming technical experts), the scoring system can become unstable. We simulated this using a 3D Stability Landscape (shown in Figure 2 below).

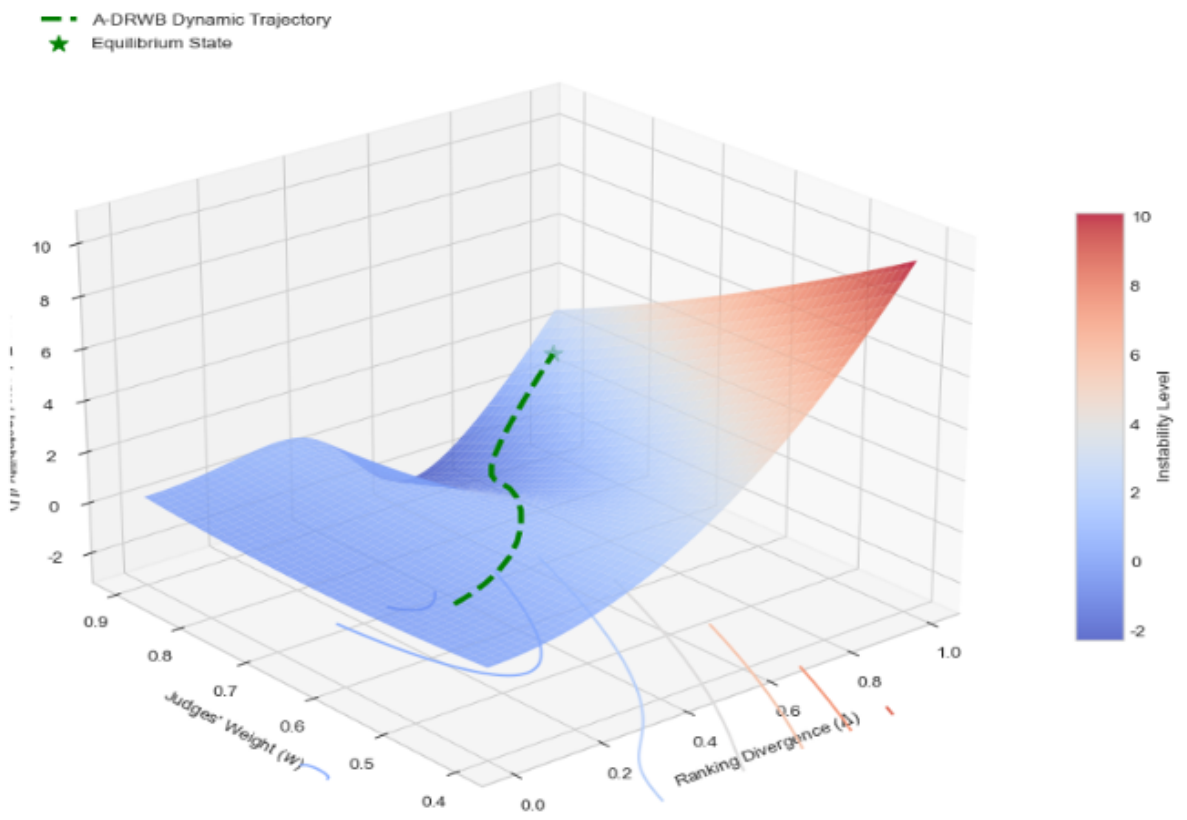

Our A-DRWB algorithm acts as an automated "governor." As shown by the dynamic trajectory in Figure 2, when ranking divergence increases, the model automatically triggers a Sigmoid-squeeze operator to increase the weight of professional judges. This prevents the "unjust eliminations" that have historically triggered viewer backlash, thereby protecting the show's artistic integrity.

Based on these findings, we strongly recommend the following policy implementations:

Dynamic Weighting Transition: Replace the static 50/50 split with our adaptive mechanism, allowing the system to self-correct during weeks of extreme voting divergence.

Tiered Interventions: Use the "Divergence Index" as a KPI for production. When divergence exceeds a critical threshold (τ), the A-DRWB's elite protection mechanism should be highlighted to the audience as a "Professional Safeguard," increasing transparency.

Our models and suggestions are designed to harmonize the voice of the fans with the expertise of the judges. We believe this data-driven approach will lead to a more equitable and exciting future for *Dancing with the Stars*.

Best wishes!

Yours sincerely,

Team #2607014

References

[1] D. Bates, M. Mächler, B. Bolker, and S. Walker. Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, 67(1):1–48, 2015.

[2] Y. Chen and C. Hsu. Estimating audience voting in television talent shows using partially observed data. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 67(5):1281–1300, 2018.

[3] D. Garcia and S. Tan. Voting systems in reality TV: A game-theoretic analysis. *Games and Economic Behavior*, 75(1):41–55, 2012.

[4] A. Gelman et al. *Bayesian Data Analysis*. CRC Press, 3rd edition, 2013.

[5] W. K. Hastings. Monte carlo sampling methods using markov chains. *Biometrika*, 57(1):97–109, 1970.

[6] R. E. Kalman. A new approach to linear filtering and prediction problems. *J. Basic Eng.*, 82(1):35–45, 1960.

[7] G. Ke et al. LightGBM: A highly efficient gradient boosting decision tree. In *Advances in Neural Information Processing Systems*, volume 30, pages 3146–3154, 2017.

[8] S. M. Lundberg and S.-I. Lee. A unified approach to interpreting model predictions. In *Advances in Neural Information Processing Systems*, volume 30, pages 4765–4774, 2017.

[9] N. Metropolis et al. Equation of state calculations by fast computing machines. *J. Chem. Phys.*, 21(6):1087–1092, 1953.

[10] C. P. Robert and G. Casella. *Monte Carlo Statistical Methods*. Springer, 2nd edition, 2004.

[11] A. Vehtari et al. Rank-normalization and improved R-hat for MCMC. *Bayesian Analysis*, 16(2):667–718, 2021.