Challenge or Pass? A Dynamic Policy for MLB's Automated Ball-Strike System

Zach Sissman[1], Lila Dodson[2], Jonathan Pipping-Gamón[3], and Abraham J. Wyner[3]
[1]Community School of Naples
[2]San Francisco University High School
[3]Department of Statistics & Data Science, University of Pennsylvania

Abstract

Major League Baseball's automated ball-strike challenge system requires teams to weigh correcting a call against preserving limited challenges for later. We develop a policy that combines the probability of an overturn, its run value, and the future cost of a failed review. A latent-signal model uses the observed call and count together with noisy impressions of pitch location to estimate overturn probabilities. Historical challenge decisions calibrate a combined sensory and action error scale for each role. Solving the Bellman equations then values the shared challenge inventory across offense and defense. We fit the model on earlier games and evaluate it on recorded opportunities in 497 later games from the 2026 season. Evaluation holds subsequent play fixed, measuring correction value rather than changes in final scores. Across both teams, the policy recovers 0.101 additional modeled correction runs per game (95% bootstrap interval $[0.075, 0.129]$), a 22% improvement over observed decisions. Both offense and defense contribute to this gain, despite a lower overturn rate (44.2% versus 54.2%).

1 Introduction

1.1 What makes a challenge valuable?

Major League Baseball's automated ball-strike (ABS) challenge system lets the batter, catcher, or pitcher immediately request a review of an umpire's call using pitch tracking. Each team begins with two challenges shared across offense and defense. A successful review preserves a challenge, whereas an unsuccessful review costs one, reducing the team's ability to correct later mistakes (Major League Baseball, 2026a). Choosing which calls to contest therefore becomes part of the team's strategy.

Consider two first-inning calls, each made with two outs and two challenges available. Kansas City's Starling Marte took a called third strike against Los Angeles with a full count and runners on first and second. The inning ended without a challenge, although public tracking places the pitch outside our reconstructed ABS zone (Figure 1). Reversing the call would have awarded ball four, loaded the bases, and kept the rally alive. Cincinnati's Nathaniel Lowe, by comparison, successfully challenged a first-pitch strike against Cleveland with the

bases empty. His review changed the count from 0–1 to 1–0 without changing the outs or runners.

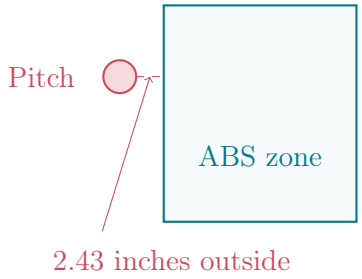

Figure 1: Marte's unchallenged third strike. Public tracking places the nearest ball edge 2.43 inches outside the reconstructed ABS zone. The pitch and zone are drawn to the same physical scale. No official review occurred.

The two corrections offer very different potential gains. Our run-expectancy estimates, based on 2023–2025 scoring data, compare the expected runs remaining in the half-inning under the corrected and original rulings. This comparison values the reversal of Marte's call at 0.782 expected runs and Lowe's at 0.032. Overturn rate gives each successful review equal weight and omits calls left unchallenged. Evaluating challenge strategy therefore requires the value of each possible correction, including opportunities a team chooses to pass.

1.2 Correct this call or preserve a future opportunity?

Those gains are conditional on an overturn, so their size alone cannot settle the decision. We must also estimate the probability of an overturn and account for the cost of an unsuccessful review. Holding that cost fixed, a reversal that preserves a rally can justify a challenge at a lower probability of success than one that only improves an early count.

The cost of failure is the future correction value forgone by losing a challenge. A batter's failed review early in a game can leave a catcher unable to contest a more valuable call later. A retained challenge can also support several successful reviews, so its value depends on the sequence of opportunities still ahead. We therefore measure this cost by comparing the expected run value of future corrections with the current inventory and with one fewer challenge. Challenging is worthwhile when the probability of an overturn multiplied by its run value exceeds the probability of failure multiplied by this future loss.

1.3 Our approach and contribution

Dynamic models of tennis challenges weigh overturning a call against losing later review opportunities (Abramitzky et al., 2012). In baseball, models of pitch framing account for location, count, and participants (Deshpande and Wyner, 2017), while models of plate discipline evaluate decisions by their expected run consequences (Yee and Deshpande, 2024). We bring these ideas together for challenges drawn from a shared team inventory, including opportunities that players pass.

We model the chance that a call is wrong given the umpire's ruling, the count, and a player's noisy impression of pitch location. Tracked locations among comparable adverse calls supply a prior distribution, which we update with the player's impression. Historical decisions reveal how sharply challenge behavior changes with location, allowing us to calibrate uncertainty.

Those decisions cannot separately identify errors in perception and errors in acting on it, so we keep their allocation explicit when constructing the policy.

For a specified allocation, we combine the resulting overturn probabilities with the run value of a reversal and estimate the future opportunities a team may face. Solving the Bellman equations then gives the expected correction value available from each game state and challenge inventory. The resulting Bellman-optimal policy weighs the expected gain from correcting the current call against the expected cost of failure, accounting for subsequent decisions on both offense and defense.

We fit the policy using earlier games and evaluate it on the recorded opportunities in 497 later games. In the primary scenario, we assign the fitted error scale to sensory uncertainty and assume that the recommended decisions are followed. Across both teams, the policy recovers 0.561 runs of modeled correction value per game, compared with 0.459 under observed decisions, an increase of approximately 22%. Its modeled overturn rate is lower, 44.2% versus 54.2%, but it recovers more value on both offensive and defensive opportunities.

2 Data and construction of challenge opportunities

2.1 Constructing the decision record

We join MLB's live game feed with Statcast records from Baseball Savant by game, plate appearance, and physical pitch number (Major League Baseball, 2026b,c). The feed supplies the event sequence, umpire calls, challengers, and review outcomes, while Statcast supplies pitch location, batter-specific zone bounds, and the count, outs, and runners before the pitch. Because a reviewed pitch may be recorded under its corrected ruling, we use the review outcome to reconstruct the original call. We then process game events and reviews in order to reconstruct each team's shared challenge inventory, including replenishment when a team enters an extra inning with none remaining.

We include tracked calls that meet the review rules, even when the team has no challenges remaining. A called strike creates an offensive opportunity, and a called ball creates a defensive one. Retaining unchallenged calls lets us evaluate alternative strategies, including those that could have preserved a challenge for calls made after the actual inventory was exhausted. We record actual availability separately and estimate historical challenge behavior only when a challenge was available.

2.2 Observed location and player information

These records allow us to assess whether a call was correctable even when no review occurred. We compare the tracked pitch location with the batter-specific ABS zone, accounting for the ball's radius and the zone's corners. Official ABS outcomes are available only for reviewed pitches, so we infer correctness for unreviewed calls from this geometry, subject to tracking error.

To express location on a common scale for the two roles, we define a signed margin M, measured in inches and oriented toward the potential challenger. Positive margin means outside the zone for a called strike and inside the zone for a called ball. Thus $M > 0$ indicates a correctable call for either role. At $M = 0$, the pitch counts as a strike.

We observe the reconstructed margin retrospectively, while the player knows the umpire's ruling and game situation but acts on an imperfect impression of location. Our probability model combines this unobserved impression with the call and count to represent the player's uncertainty about whether a review would reverse the ruling.

2.3 Development and evaluation samples

We estimate run expectancy from 2023–2025 Statcast data, then use 2026 games through July 19 to estimate the probability inputs and future opportunity process for the Bellman calculation. The original development schedule contains 1,490 games, including 1,488 with ABS. These ABS-enabled games supply 218,896 tracked called pitches and 6,263 observed challenges.

We evaluate the policy on 497 ABS-enabled games from July 20 through August 25, containing 72,633 tracked called pitches and 2,275 observed challenges, of which 1,233 resulted in overturns. Reported samples exclude two games without ABS infrastructure from each period.

3 A run-value challenge policy

3.1 The value of correcting a call

Following the run-expectancy approach to baseball strategy (Lindsey, 1963), we value a possible reversal by comparing expected runs under the ball and strike rulings. Let $\widehat{\mathrm{RE}}(x)$ be the estimated runs remaining in the half-inning from state x, defined by the count, outs, and occupied bases. For pitch i and ruling a, let $R_i(a)$ be the runs scored immediately and $x_i(a)$ the resulting state. The correction gain is

$$G_i = R_i(\text{ball}) + \widehat{\mathrm{RE}}\{x_i(\text{ball})\} - R_i(\text{strike}) - \widehat{\mathrm{RE}}\{x_i(\text{strike})\}. \tag{1}$$

For offense, this difference measures the increase in expected runs from reversing a called strike; for defense, it measures the reduction in expected runs allowed from reversing a called ball.

Both branches preserve recorded concurrent runner movements, outs, and runs, then apply the consequences of the ruling. Ball four advances the batter and any forced runners, while strike three adds an out. After the third out, remaining run expectancy is zero. We retain any negative differences produced by the empirical estimates, but the policy passes when $G_i \leq 0$. The gain G_i is conditional on a reversal and can be valued before knowing whether the call was wrong. Deciding whether to pursue that gain also requires the probability of success given the player's information.

3.2 From perceived location to overturn probability

The umpire's ruling narrows which pitch locations are plausible. For each role $r \in \{o, d\}$ and pre-pitch count c, we estimate a distribution of margins from the development opportunity sample. We place margins in 0.01-inch bins and shrink count-specific frequencies toward the distribution for that role. Let $p_{rc}(m)$ denote the resulting prior mass at bin center m. Conditioning on role incorporates the observed call, since offensive opportunities are called strikes and defensive opportunities are called balls.

The player's impression updates this prior. For an opportunity in role r, we model the perceived margin as

$$S_i = M_i + \varepsilon_{s,i}, \qquad \varepsilon_{s,i} \mid M_i, r, c \sim N(0, \sigma_{s,r}^2). \tag{2}$$

The sensory standard deviation $\sigma_{s,r}$ measures uncertainty about location in inches. For $\sigma_{s,r} > 0$, Bayes' rule gives the posterior probability $\mathbb{P}(M > 0 \mid S = s, r, c)$ as

$$q_{rc}(s; \sigma_{s,r}) = \frac{\sum_{m>0} p_{rc}(m)\phi_{\sigma_{s,r}}(s-m)}{\sum_{m} p_{rc}(m)\phi_{\sigma_{s,r}}(s-m)}, \tag{3}$$

where ϕ_σ is the normal density with mean zero and standard deviation σ. The likelihood weights each possible margin by its compatibility with the player's impression, and the numerator collects the weight on correctable calls. The probability of a successful reversal on pitch i is therefore $q_i = q_{r_i c_i}(S_i; \sigma_{s,r_i})$. Figure 2 illustrates this update for Marte's full-count call. To specify the sensory scale, we next examine what historical challenge decisions reveal about perception and action.

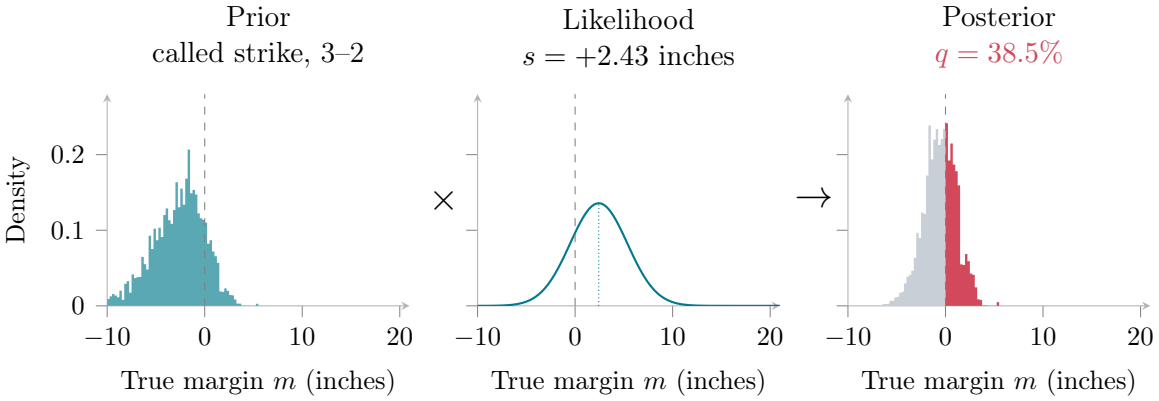

Figure 2: For an illustrative impression at Marte's tracked margin, $s = +2.43$ inches, the full-count prior and Gaussian likelihood yield a posterior with 38.5% of its mass on correctable calls. The primary allocation uses a sensory standard deviation of 2.94 inches (Section 3.3).

3.3 What historical decisions tell us about uncertainty

To connect perceived location to observed choices, let $t_r(X_i)$ be a threshold in inches, fixed given context X_i, and let $\varepsilon_{a,i}$ be an action error with standard deviation $\sigma_{a,r}$. The threshold captures systematic reasons to challenge or conserve inventory, while action error captures variation around it. We assume that sensory and action errors are independent centered

Gaussians and independent of margin given context. For opportunities with a challenge available, the historical challenge indicator A_i^{hist} satisfies

$$\begin{aligned} A_i^{\text{hist}} &= \mathbf{1}\{M_i + \varepsilon_{s,i} + \varepsilon_{a,i} > t_r(X_i)\}, \\ \mathbb{P}(A_i^{\text{hist}} = 1 \mid M_i, X_i, r) &= \Phi\left(\frac{M_i - t_r(X_i)}{w_r}\right), \qquad w_r^2 = \sigma_{s,r}^2 + \sigma_{a,r}^2, \end{aligned} \tag{4}$$

where Φ is the standard normal distribution function. The fitted probit coefficient on margin, b_r, implies a combined error scale of $w_r = 1/b_r$ inches.

We fit this coefficient separately by role within three inches of the boundary, adjusting for gain, count, inning, score, outs, remaining challenges, pitch family, handedness matchup, and participant and team differences. The fitted widths are $w_o = 2.940$ inches and $w_d = 1.817$ inches.

These fitted widths combine sensory and action error without identifying the two components separately. Our primary analysis assumes $\sigma_{s,r} = w_r$ ($\kappa_r = 1$) and that players follow the recommended policy given their impressions. Section 4.2 examines other allocations while retaining action error. With the overturn probabilities specified, the recommended decision depends on the current gain and the value of preserving a challenge for later opportunities.

3.4 Future value and the Bellman-optimal policy

Keeping a challenge preserves the chance to correct later calls. Let $C_k(z)$ be the expected run value of those later corrections with k challenges remaining and optimal subsequent decisions. Here z records the inning, half-inning, outs, and role, while k is shared across offense and defense. The cost of losing one challenge is

$$L_k(z) = C_k(z) - C_{k-1}(z), \qquad k = 1, 2. \tag{5}$$

A successful review earns G and keeps the challenge; only failure reduces future availability. Thus qG is the expected current gain and $(1-q)L_k(z)$ is the expected future loss. This gives the same challenge tradeoff as Abramitzky et al. (2012), expressed here in correction runs. When a challenge is available, the rule is

$$\text{challenge if } qG > (1-q)L_k(z). \tag{6}$$

A valuable correction can therefore justify challenging with less confidence, while valuable future opportunities favor conserving inventory. Because posterior overturn probability increases with perceived margin, this rule can also be written as $S > t$, with the cutoff t determined by the opportunity's role, count, gain, and inventory cost.

We estimate those future opportunities from development games, following each team's adverse calls in order. These sequences show where the next call occurs and when opportunities end. States with few observations borrow information across innings with the same half-inning, outs, and role, treating extra innings separately. The same games supply the joint distribution of counts and gains, which we combine with the margin priors and signal model to describe

possible future (q, G) pairs. When comparing a challenge with a pass, we hold the distribution of later opportunities fixed.

To value these opportunities, let $V_k(z)$ be expected correction value from the current call onward, before its q and G are known. After observing these inputs, the player chooses between passing, which adds zero, and challenging, which adds $qG - (1-q)L_k(z)$. Averaging the better choice over possible opportunities gives

$$V_k(z) = C_k(z) + \mathbb{E}[\max\{0, qG - (1-q)L_k(z)\} \mid z], \qquad k = 1, 2. \tag{7}$$

For calls that must be passed, the contribution inside the expectation is set to zero. With no challenges, $V_0(z) = C_0(z)$. This value can still be positive because an extra inning may restore a challenge.

We solve these equations by value iteration (Puterman, 1994). Starting from zero values, we average the next opportunity's V values to update C, then apply the Bellman equation to update V. These averages include the chance that no opportunities remain and any extra-inning replenishment. We repeat until the values stabilize.

4 Policy evaluation

4.1 Evaluating decisions through whole games

For the primary analysis, we hold the fitted probability model and Bellman cost function fixed in test games. Conceptually, we draw a sensory error for each pitch, add it to the tracked margin, and use the resulting impression to obtain a posterior overturn probability. The policy then gives a deterministic decision. Averaging these decisions over possible draws gives the challenge probability at each inventory level. We compute this average directly rather than simulating the draws.

Starting with two shared challenges per team, we follow its recorded calls and track the probabilities of having zero, one, or two challenges. These inventory probabilities weight the policy's decisions at each level. The reconstructed margin determines whether a challenge would succeed. Failed challenges reduce inventory, while passes and successes preserve it. Entry into an extra inning restores one challenge when none remain.

Expected recovered run value equals the overall challenge probability times G_i on correctable calls and zero on correct calls. We sum these gains across both teams to report value per game, keeping the recorded pitches and baseball states fixed. The result measures correction value from those opportunities rather than a change in final scores.

4.2 Benchmarks, sensitivity, and uncertainty

We compare the policy with observed decisions and a perfect-execution oracle that knows each pitch's exact location. On the same recorded opportunities, the oracle challenges every correctable call with $G_i > 0$. All its reviews succeed, so it retains its challenges and captures all positive correction value available on this record.

We allocate the fitted error variance between sensory and action error, retaining both under the new policy:

$$\sigma_{s,r} = \kappa_r w_r, \qquad \sigma_{a,r} = w_r\sqrt{1-\kappa_r^2}. \tag{8}$$

Here $\kappa_r \in [0,1]$, with sensory variance share κ_r^2 and total error variance w_r^2 throughout. At $\kappa_r = 0$, perception is exact but action-error standard deviation remains w_r. The oracle separately removes both errors.

With a challenge available and $G > 0$, a margin cutoff t gives the implemented decision $S + \varepsilon_a > t$. Averaging this decision over sensory and action errors gives the challenge probability used in replay. It has the same probit form as Equation (4), with the historical threshold replaced by t and the combined width w_r unchanged. Holding the other fitted inputs fixed, every candidate cutoff therefore gives the same expected gains and inventory transitions across allocations. Within this cutoff family, Bellman optimization gives the same optimal policy value, so the primary $\kappa_r = 1$ solution also solves the problem for the other allocations.

We form 95% percentile intervals from 5,000 bootstrap resamples of the 497 test games. Each resample keeps both teams and all decisions within a game together, using the same games to compare the policy with observed decisions. The fitted inputs and policy remain fixed, so these intervals capture variation across test games and exclude uncertainty from model estimation.

5 Results

5.1 Recovered run value

Across the 497 test games, the Bellman-optimal policy recovers 0.561 modeled correction runs per game, compared with 0.459 under observed decisions. The gain is 0.101 runs per game (95% interval $[0.075, 0.129]$), an increase of 22%. Both offense and defense contribute to this improvement, with offensive opportunities accounting for about two-thirds of the total gain (Table 1).

Table 1: Total modeled runs recovered across both teams per game in the 497 test games. Gain is the policy's improvement over observed decisions, with paired bootstrap 95% intervals in brackets.

Role	Observed	Bellman-optimal	Oracle	Gain (95% interval)
Offense	0.193	0.261	0.744	0.068 [0.050, 0.087]
Defense	0.266	0.299	0.655	0.033 [0.014, 0.052]
Combined	0.459	0.561	1.399	0.101 [0.075, 0.129]

The policy captures 40.1% of the correction value available to the oracle.

5.2 How the policy uses challenges

The policy makes more challenges: 6.17 per game across both teams, compared with 4.58 observed. Its modeled overturn rate is 44.2%, compared with 54.2% under observed decisions,

but the additional attempts produce more successful reviews. Those successes also recover more value on average, 0.206 runs compared with 0.185 under observed decisions. The improvement thus combines more corrections with more valuable corrections, while accepting additional failed reviews.

The cost of those failures depends on the opportunities still ahead. Figure 3 shows how the fitted continuation costs translate into break-even overturn probabilities for a 0.10-run correction, following a home team at calls with no outs. Losing the last challenge costs more than losing one of two, and both costs fall as regulation progresses. The declining probability thresholds show how the policy protects scarce inventory early while becoming more willing to challenge as fewer opportunities remain.

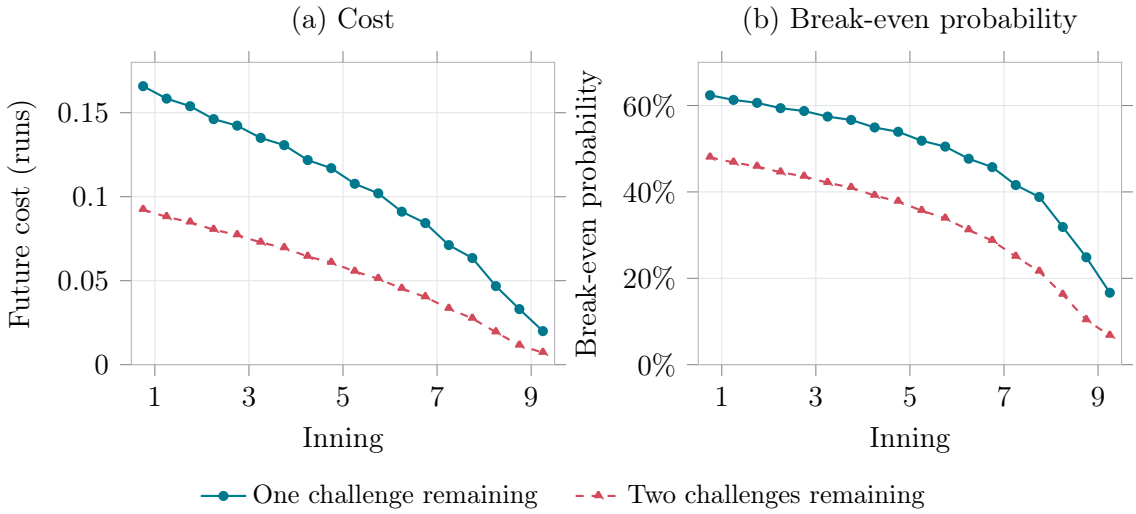

Figure 3: Future correction value lost on failure (left) and break-even overturn probabilities for a 0.10-run correction (right), with one or two challenges remaining. Points follow a home team through regulation, with defense in each top half followed by offense in the bottom half, evaluated at calls with no outs. All costs include future opportunities in both roles.

For example, a home batter facing a 0.10-run correction opportunity with no outs in the first inning needs an overturn probability above about 47% with two challenges remaining, compared with 61% when only one remains. The higher confidence requirement reflects the cost of losing future opportunities, even though the current call offers the same gain.

5.3 Returning to the opening examples

The opening examples face similar costs of failure but very different potential gains. Losing one of two challenges costs 0.091 expected runs for Marte's opportunity and 0.086 for Lowe's. Their correction gains of 0.782 and 0.032 runs therefore imply sharply different confidence requirements: the policy challenges above a 10.4% posterior probability of an overturn for Marte, compared with 72.8% for Lowe. Although both calls occur early in the game, preserving Marte's rally justifies acting with substantially less certainty.

With two challenges available, averaging over possible impressions gives challenge probabilities of 94.9% for Marte's pitch and 0.8% for Lowe's. The policy therefore usually contests the

valuable call left unchallenged and passes on the small correction that Lowe successfully obtained. These probabilities reflect both the pitches' locations and the different confidence requirements. A sufficiently convincing impression could still justify Lowe's challenge, but the successful outcome alone does not show whether his confidence met that requirement.

6 Discussion

6.1 Implications for challenge strategy

The policy's gains come from changing decision thresholds while retaining the fitted error scales. Within this model, teams can improve challenge allocation while uncertainty about location persists. Overturn rate excludes opportunities a team passes, so it can favor caution even when a lower-confidence challenge has positive expected value. Valuing both current corrections and future opportunities makes that tradeoff explicit.

Because batters and catchers draw on the same inventory, their decisions also have consequences across roles. Preserving a challenge on offense can leave it available for a valuable defensive opportunity later in the game. The framework therefore provides a common basis for reviewing team decisions and developing situational guidance: the confidence required to challenge depends on the correction's run value and the opportunities still ahead.

6.2 Limitations and next steps

Reported gains measure correction value along the recorded sequence of opportunities. A different ruling could change the pitches and situations that follow, so these gains do not estimate a change in final scores. Errors in tracking or zone reconstruction can also change which calls we classify as correctable. Extending the objective to win probability would allow the score to affect how much a reversal is worth.

Players' impressions are also unobserved. Gaussian errors and a common error scale within each role approximate how players perceive pitches and act on them, while priors condition on the call and count. Individual differences and additional private information could change the recommended decisions. Both the policy's optimality and the allocation result in Section 4.2 depend on these assumptions. Historical choices do not identify sensory and action error separately. Direct measures of player confidence could help assess these assumptions and support individual calibration.

Bootstrap intervals capture variation across test games with fitted inputs held fixed. Resampling the estimation data and refitting the inputs and policy would also reflect estimation uncertainty.

7 Reproducibility

The analysis data and source code are available at https://github.com/WhartonSABI/abs. The repository includes instructions for reproducing the analysis.

References

Abramitzky, R., Einav, L., Kolkowitz, S., and Mill, R. (2012). On the optimality of line call challenges in professional tennis. *International Economic Review*, 53(3), 939–964. https://doi.org/10.1111/j.1468-2354.2012.00706.x.

Deshpande, S. K. and Wyner, A. (2017). A hierarchical Bayesian model of pitch framing. *Journal of Quantitative Analysis in Sports*, 13(3), 95–112. https://doi.org/10.1515/jqas-2017-0027.

Lindsey, G. R. (1963). An investigation of strategies in baseball. *Operations Research*, 11(4), 477–501. https://doi.org/10.1287/opre.11.4.477.

Major League Baseball (2026a). ABS Challenge Dashboard. Baseball Savant. https://baseballsavant.mlb.com/abs. Accessed October 17, 2026.

Major League Baseball (2026b). MLB Stats API Documentation. https://docs.statsapi.mlb.com/. Accessed October 17, 2026.

Major League Baseball (2026c). Statcast Search CSV Documentation. Baseball Savant. https://baseballsavant.mlb.com/csv-docs. Accessed October 17, 2026.

Puterman, M. L. (1994). *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons. https://doi.org/10.1002/9780470316887.

Yee, R. and Deshpande, S. K. (2024). Evaluating plate discipline in Major League Baseball with Bayesian Additive Regression Trees. *Journal of Quantitative Analysis in Sports*, 20(1), 5–20. https://doi.org/10.1515/jqas-2023-0048.