

Straight, Draw, and Fade

Geometry, Dispersion, and
Conditional Scoring

Driver · Seven Iron · Pitching Wedge

Research Question

Does deliberately curving a shot produce a tighter or more valuable outcome when nominal targets are matched and impact delivery varies?

Principal Finding

There is no universal winning shot shape in these simulations. Dispersion, carry, and conditional scoring depend on the club, delivery variability, and geometric assumptions. A matched nominal target does not force equal carry on every shot or isolate curvature as a cause.

Illustrated Research Edition—Revised October 11, 2026

10,000 Attempted Shots per Pattern per Executed Condition

Geometric and Numerical Analysis; No Measured Player Dataset

Presentation revision: conventional lateral-x/downrange-y figures and expanded methods; numerical simulation results are unchanged. The prior frozen edition is retained separately.

Abstract

This study compares straight, draw, and fade patterns for three assumed club geometries. The matched-delivery experiment solves nominal face, path, and speed for a common raw landing target, then applies paired variations in shaft rotation, path, speed, strike, and covariance. It covers shaft-rotation standard deviations of 1° and 2° , ordinary and doubled curve strength, and three forward-lean settings. The full original and supplemental inventories contain 10.44 million attempted shot rows. Supported populations, nominal-search failures, and invalid outcomes remain explicit. We examine endpoint density, lateral and longitudinal spread, illustrative flight paths, and conditional strokes gained. The selected scoring estimates do not identify one pattern that is best across these modeled clubs and delivery distributions. Face closure through shaft rotation lowers delivered loft locally, but its effect on carry is club dependent. The estimates are model sensitivities, not measured golfer performance; Monte Carlo intervals exclude physical-model and player-distribution uncertainty.

Contents

Abstract	1
1 Executive Findings	2
2 Visual Results	2
3 Model, Methods, and Numerical Results	17
A Historical Phase-One Status and Scope	33
B Inputs, Frames, and Pairing	33
C Corrected Central-Contact Impact Equations	34
D Geometric Face–Loft Coupling and Shaft Lean	34
E Flight Integration and Landing	35
F Aiming and Outcome Statistics	35
G Distance Control	36
H Club-Appropriate Scoring Estimates	36
I Corrected Monte Carlo Results	37
J Interpretation and Observability Limits	37
K Evidence and Reproduction	38
L External Context	38

1 Executive Findings

Curving the ball is not a universal accuracy advantage. The selected zero-lean scoring comparisons favor different shapes for different clubs, with some intervals spanning zero. Neither a small scoring difference nor a narrow sampling interval establishes a golfer’s preferred pattern.

The shots do not all travel the same distance. The solver matches the zero-error nominal endpoint. Noisy shots retain both lateral and downrange variation; changing variability can also change mean carry. The fixed-target controls reduce nominal distance mismatch but rematch face, path, and speed, preventing a causal attribution to curvature alone.

Face closure and loft are coupled under the shaft-rotation assumption. Closing locally lowers loft. The selected shaft-only iron and wedge cases exhibit long-left/short-right association. The selected shaft-only low-loft driver cases show the opposite association; carry also depends on launch, spin, and ball speed, so the geometric derivative alone does not determine its direction. Shaft lean changes the delivered geometry and its face-loft coupling.

Scoring effects are conditional and small in the selected examples. Across eight selected zero-lean conditions per club, Draw minus Straight SG ranges from about +0.00003 to +0.0078 for driver, −0.00136 to −0.00029 for seven iron, and +0.00010 to +0.0111 for wedge. These are ranges of per-opportunity estimates, not pooled intervals or per-round improvements. Complete paired intervals and Fade comparisons appear in the numerical results.

1.1 How to Read the Figures

The report graphics use the conventional plotting layout: reported lateral distance is on the horizontal axis and downrange distance is on the vertical axis; reported lateral is positive right. The native right-handed model frame is (X_n, Y_n, Z_n) = forward/target, left, up. Reported endpoint coordinates are $(X, Y) = (X_n, -Y_n)$, so the plotted lateral sign differs from native Y_n . Contact and club-frame vector equations use the native frame; endpoint and scoring tables use the reported frame. Density plots show endpoint populations, with 50%, 80%, and 95% binned empirical highest-density regions; these are not confidence intervals for a mean. Dots and flight paths are labelled samples. The baseline club plates use zero lean, curve strength $1\times$, shaft-twist SD 1° , and the combined2 profile: path SD 0.5° , speed SD 2 mph, and strike SD 0.25 in in each face-plane axis. Shaft-twist SD is distinct from delivered-face SD. The initial fixed-path face-SD experiment is retained separately in the historical appendix.


2 Visual Results

The following plates embed the accepted graphics directly. The paired variability and curve-strength plots use the same supported shot indices across all six populations. Counts appear in the graphics; exact source protocols and file hashes are retained in the [Presentation Figure Mapping](#). Comparisons are conditional on support; unshown invalid outcomes are still retained in the experiment inventory.

No Universal Winning Shot Shape

Selected Relative Scoring Ranges

Minimum-Maximum Across Eight Selected Zero-Lean Conditions



Distance Varies With Shape and Club. Closure Lowers Loft Here; carry direction depends on the full impact and flight response.

Relative SG effects are small and conditional on the scoring model. Straight is each cell's zero reference.


These ranges are not pooled intervals or per-round gains.

Source: Selected relative-SG table; carry only; assumed lies and historical benchmark. See numerical results for paired intervals.

Figure 1: Graphical Summary of Conditional Findings

Driver: Endpoint Dispersion

Matched Target 178.9 m · Supported 10,000–10,000/10,000 per Pattern · Shared Axes and Density Scale



Red + = Matched Target · Solid/Dashed/Dotted = 50/80/95% Binned Empirical HDR · Equal Metre Scaling
 Blue = Probability Density (m^{-2}), Shared Scale · Carry Only · Model-Conditional Simulation; Not Mean Confidence Intervals or Calibrated Player Precision.

Figure 2: Driver: Endpoint Dispersion

Driver Straight, Draw, and Fade Endpoints at a Matched 178.9 m Target. Density contours condition on supported airborne shots (10,000/10,000 supported per pattern for this configuration); axes are in metres. The zero-error nominal endpoints were matched before input variation.

Driver: Illustrative Flight Paths

Matched Target 178.9 m · 100 Replayed Shots per Pattern · Zero-Lean Combined-Noise Condition





Figure 3: Driver: Illustrative Flight Paths

Driver Straight, Draw, and Fade Endpoints at a Matched 178.9 m Target. One hundred verified replay paths per pattern illustrate the supported airborne populations (10,000/10,000 supported per pattern for this configuration); axes are in metres. The zero-error nominal endpoints were matched before input variation. Flight curves use common supported shot IDs, not the full simulated population. Read the lateral magnification stated in the plot.

Seven Iron: Endpoint Dispersion

Matched Target 127.9 m · Supported 9,996–9,996/10,000 per Pattern · Shared Axes and Density Scale



Red + = Matched Target · Solid/Dashed/Dotted = 50/80/95% Binned Empirical HDR · Equal Metre Scaling
 Blue = Probability Density (m⁻²), Shared Scale · Carry Only · Model-Conditional Simulation; Not Mean Confidence Intervals or Calibrated Player Precision.

Figure 4: Seven Iron: Endpoint Dispersion

7-Iron Straight, Draw, and Fade Endpoints at a Matched 127.9 m Target. Density contours condition on supported airborne shots (9,996/10,000 supported per pattern for this configuration); axes are in metres. The zero-error nominal endpoints were matched before input variation.

Seven Iron: Illustrative Flight Paths

Matched Target 127.9 m · 100 Replayed Shots per Pattern · Zero-Lean Combined-Noise Condition





Figure 5: Seven Iron: Illustrative Flight Paths

7-Iron Straight, Draw, and Fade Endpoints at a Matched 127.9 m Target. One hundred verified replay paths per pattern illustrate the supported airborne populations (9,996/10,000 supported per pattern for this configuration); axes are in metres. The zero-error nominal endpoints were matched before input variation. Flight curves use common supported shot IDs, not the full simulated population. Read the lateral magnification stated in the plot.

Pitching Wedge: Endpoint Dispersion

Matched Target 77.4 m · Supported 9,999–9,999/10,000 per Pattern · Shared Axes and Density Scale



Red + = Matched Target · Solid/Dashed/Dotted = 50/80/95% Binned Empirical HDR · Equal Metre Scaling
 Blue = Probability Density (m^{-2}), Shared Scale · Carry Only · Model-Conditional Simulation; Not Mean Confidence Intervals or Calibrated Player Precision.

Figure 6: Pitching Wedge: Endpoint Dispersion

Pitching Wedge Straight, Draw, and Fade Endpoints at a Matched 77.4 m Target. Density contours condition on supported airborne shots (9,999/10,000 supported per pattern for this configuration); axes are in metres. The zero-error nominal endpoints were matched before input variation.

Pitching Wedge: Illustrative Flight Paths

Matched Target 77.4 m · 100 Replayed Shots per Pattern · Zero-Lean Combined-Noise Condition




Figure 7: Pitching Wedge: Illustrative Flight Paths

Pitching Wedge Straight, Draw, and Fade Endpoints at a Matched 77.4 m Target. One hundred verified replay paths per pattern illustrate the supported airborne populations (9,999/10,000 supported per pattern for this configuration); axes are in metres. The zero-error nominal endpoints were matched before input variation. Flight curves use common supported shot IDs, not the full simulated population. Read the lateral magnification stated in the plot.

Driver: Increasing Shaft-Rotation Variability

Two Ordered Conditions · 10,000/10,000 Common Supported Indices Across All Six Populations · Shared Axes and Density Scale




Figure 8: Driver: Increasing Shaft-Rotation Variability

Driver Endpoint Response to Shaft-Rotation Variability: shaft-rotation noise SD 1° versus 2° at 0° nominal lean and curve $\times 1$. Straight, Draw, and Fade panels compare the same physical pattern across the two ordered populations at a matched 178.9 m target, conditional on 10,000 shot indices supported in all six populations (of 10,000 attempts per population); plotted distances are in metres.

Driver: Doubling Curve Strength

Two Ordered Conditions · 10,000/10,000 Common Supported Indices Across All Six Populations · Shared Axes and Density Scale




Figure 9: Driver: Doubling Curve Strength

Driver Endpoint Response to Curve Strength: curve scale $\times 1$ versus $\times 2$ at 0° nominal lean and shaft-rotation noise SD 1° . Straight, Draw, and Fade panels compare the same physical pattern across the two ordered populations at a matched 178.9 m target, conditional on 10,000 shot indices supported in all six populations (of 10,000 attempts per population); plotted distances are in metres.

Seven Iron: Increasing Shaft-Rotation Variability

Two Ordered Conditions · 9,996/10,000 Common Supported Indices Across All Six Populations · Shared Axes and Density Scale




Figure 10: Seven Iron: Increasing Shaft-Rotation Variability

7-Iron Endpoint Response to Shaft-Rotation Variability: shaft-rotation noise SD 1° versus 2° at 0° nominal lean and curve ×1. Straight, Draw, and Fade panels compare the same physical pattern across the two ordered populations at a matched 127.9 m target, conditional on 9,996 shot indices supported in all six populations (of 10,000 attempts per population); plotted distances are in metres.

Seven Iron: Doubling Curve Strength

Two Ordered Conditions · 9,996/10,000 Common Supported Indices Across All Six Populations · Shared Axes and Density Scale




Figure 11: Seven Iron: Doubling Curve Strength
 7-Iron Endpoint Response to Curve Strength: curve scale $\times 1$ versus $\times 2$ at 0° nominal lean and shaft-rotation noise SD 1° . Straight, Draw, and Fade panels compare the same physical pattern across the two ordered populations at a matched 127.9 m target, conditional on 9,996 shot indices supported in all six populations (of 10,000 attempts per population); plotted distances are in metres.

Pitching Wedge: Increasing Shaft-Rotation Variability

Two Ordered Conditions · 9,999/10,000 Common Supported Indices Across All Six Populations · Shared Axes and Density Scale




Figure 12: Pitching Wedge: Increasing Shaft-Rotation Variability

Pitching Wedge Endpoint Response to Shaft-Rotation Variability: shaft-rotation noise SD 1° versus 2° at 0° nominal lean and curve ×1. Straight, Draw, and Fade panels compare the same physical pattern across the two ordered populations at a matched 77.4 m target, conditional on 9,999 shot indices supported in all six populations (of 10,000 attempts per population); plotted distances are in metres.

Pitching Wedge: Doubling Curve Strength

Two Ordered Conditions · 9,999/10,000 Common Supported Indices Across All Six Populations · Shared Axes and Density Scale




Figure 13: Pitching Wedge: Doubling Curve Strength

Pitching Wedge Endpoint Response to Curve Strength: curve scale $\times 1$ versus $\times 2$ at 0° nominal lean and shaft-rotation noise SD 1° . Straight, Draw, and Fade panels compare the same physical pattern across the two ordered populations at a matched 77.4 m target, conditional on 9,999 shot indices supported in all six populations (of 10,000 attempts per population); plotted distances are in metres. The second population comes from the independently executed wedge primary face-bounds protocol.

Driver: Forward-Lean Comparison

Lean 0° Versus 5° · Shaft-Rotation SD 2° · Curve $\times 1$ · 100 Replayed Shots per Pattern and Condition




Figure 14: Driver: Forward-Lean Comparison

Driver Landing-Location Shift Across Nominal Shaft Lean. Six hundred verified replay flight paths (100 per pattern at each lean), plus dashed nominal references, illustrate 0° with 5° lean while curve scale ($\times 1$) and shaft-rotation noise SD (2°) stay fixed. Straight, Draw, and Fade are matched at a 178.9 m nominal target; the separate paired endpoint and scoring analysis uses 9,940 indices supported in all six populations (of 10,000 attempts per population). The plotted carry and lateral coordinates are in metres; this model perturbation is descriptive rather than an independently validated causal mechanism. Flight curves use common supported shot IDs, not the full simulated population. Read the lateral magnification stated in the plot.

3 Model, Methods, and Numerical Results

The current phase-two study supersedes the aimed, fixed-delivery comparisons recorded below. Its calculation-level design is maintained in [the Matched-Delivery Design](#); its accepted execution, independent audits, limitations, and remaining publication work are tracked in [the Study Turnover](#). Reproduction instructions are in [the Saved-Data Reproduction Guide](#).

For each feasible straight, draw, and fade condition, nominal face, path, and clubhead speed are solved so that the unadjusted zero-error landing reaches the same physical target within 0.05 m. No post-flight aim rotation supplies that match. The reference static lofts are 10° , 34° , and 46° for driver, 7-iron, and pitching wedge, respectively. Each is evaluated at 0° , 5° , and 10° forward shaft lean, with Gaussian shaft rotation standard deviations of 1° and 2° , and two curve scales. These are geometric modeling assumptions, not measured player distributions. The earlier delivered-loft presets and lean sensitivity values below describe phase one only.

Nine delivery-noise profiles include shaft rotation alone, path SD 0.5° , speed SD 1 or 2 mph, horizontal and vertical strike SD 0.25 in each, combined perturbations, and signed shaft-rotation/path covariance. The interactive tool supports covariance across all seven modeled delivery variables. Each feasible condition uses 10,000 paired shot indices per pattern. Unsupported impacts and infeasible nominal matches remain explicit; they are not discarded from the experiment inventory.

Driver scoring uses a hypothetical tee-shot fairway/rough context, while 7-iron and wedge scoring use hypothetical approaches. First ground contact is treated as final lie. Individual-supported dispersion, common-three-pattern scoring, and common-six-population paired contrasts have different denominators and must not be conflated. The fixed-target control separates comparisons at a common distance from within-lean matching; re-matching delivery still prevents a causal claim that curvature or shaft lean alone improves accuracy.

3.1 From Delivery Assumptions to Scored Endpoints

The calculation has five linked stages. A nominal shot is first matched by forward simulation; a paired error vector is then added to the club delivery; rigid-body impact computes ball and club post-impact motion; the resulting ball state is integrated to first ground contact; and a stated course scenario maps the endpoint to a conditional score. This order matters: endpoint location and scoring emerge from delivery, contact, and flight together. There is no direct rule that assigns extra carry to a closed face or extra curvature to a draw.

Matched Nominal Delivery $\rightarrow$ **Paired Delivery Errors** $\rightarrow$ **Rigid Club–Ball Impact**
 $\rightarrow$ **RK4 Ball Flight to First Contact** $\rightarrow$ **Conditional Dispersion and Course Score**

Nominal Matching and Paired Sampling. For each club, lean, and curve scale, the study forward-simulates candidate deliveries and solves nominal face and clubhead speed. The prescribed face-to-path difference sets path from face, while the target remains $(X_T, 0)$ in the native forward/left/up frame. Acceptance is based on the raw two-coordinate endpoint residual $\sqrt{(X - X_T)^2 + Y^2} \leq 0.05$ m. No post-flight rotation or per-shot aim correction is used. Each pattern therefore begins at the same nominal point, but its matched face, path, and speed can differ. A failed bounded search is recorded as unavailable; the optimizer’s exit alone neither proves reachability nor physical impossibility. For prescribed face-to-path difference δ , path is $p = f - \delta$. The bounded least-squares problem is

$$(f_*, v_*) = \arg \min_{(f,v) \in B} \frac{1}{2} \|(X_n(f, v; \delta) - X_T, Y_n(f, v; \delta))\|^2.$$

The original search box uses reference face $f_{\text{ref}} \pm 3^\circ$ and speed $[0.75, 1.25]v_{\text{ref}}$, initialized at the reference face and speed clipped to the bounds. The reference face is 0° for Straight and $\pm 1.5^\circ$ times curve scale for Draw/Fade; their face-to-path differences are respectively 0° , -1.5° and $+1.5^\circ$ times curve scale. SciPy's bounded least-squares solver uses $\text{xtol} = \text{ftol} = \text{gtol} = 10^{-11}$ and $\text{max_nfev} = 80$. Both a successful solver termination and the independently evaluated endpoint tolerance are required. This local search provides neither a uniqueness guarantee nor a global feasibility proof. Within-lean targets use the zero-error straight reference at that lean; fixed-target controls use its zero-lean target. The explicit expanded face and speed bounds of the supplement are retained in the across-lean controls below.

For each executed condition a single seeded multivariate Gaussian sample is drawn and reused by shot index for Straight, Draw, and Fade. Its ordered components are shaft rotation, path, speed, toe/heel strike, high/low strike, forward lean, and attack angle. The modeled units are degrees, degrees, metres per second, metres, metres, degrees, and degrees. In selected protocols, shaft-rotation SD is 1° or 2° , path SD is 0.5° , speed SD is 1 or 2 mph, and strike SD is 0.25 inch on each face axis. Paired indices make contrasts less sensitive to Monte Carlo noise; they do not represent repeated measurements of the same golfer. The Gaussian distributions and any correlations are user-specified sensitivity inputs: no empirical golfer dataset was available to estimate them. Shaft rotation is converted nonlinearly into delivered face and loft, so the imposed rotation/path covariance is not itself the realized face/path covariance.

The executed matrices use seed 20261009. With D the diagonal matrix of primitive SDs and C their correlation matrix, the sampler implements

$$\epsilon_i = DU\Lambda^{1/2}\mathbf{z}_i, \quad C = U\Lambda U^T, \quad \mathbf{z}_i \sim \mathcal{N}(\mathbf{0}, I_7), \quad \Sigma = DCD.$$

NumPy's default generator supplies the independent standard-normal rows; correlation eigenvalues are clipped only for validated numerical roundoff. The same seed and field ordering preserve pairing across the comparisons. Forward-lean and attack-angle SDs are zero in all nine executed profiles, although the interface supports them. Each row below also includes the chosen shaft-twist SD of 1° or 2° .

Profile	Path SD ($^\circ$)	Speed SD (mph)	Strike SD (in/Axis)	$\rho_{\theta,p}$
shaft_only	0	0	0	0
shaft_path	0.5	0	0	0
shaft_speed1	0	1	0	0
shaft_speed2	0	2	0	0
shaft_strike	0	0	0.25	0
combined1	0.5	1	0.25	0
combined2	0.5	2	0.25	0
cov_facepath_positive	0.5	1	0.25	-0.6
cov_facepath_negative	0.5	1	0.25	+0.6

Here 1 mph = 0.44704 m/s and 0.25 in = 0.00635 m. Zero entries in the last column indicate no imposed correlation; correlation with a constant dimension is not statistically defined. The first seven profiles use $C = I_7$. For the two covariance profiles the active five-dimensional block is $C_5 = AA^T$, with $q = -0.6$ or $+0.6$, respectively:

$$A = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 \\ q & 0.8 & 0 & 0 & 0 \\ 0.2 & 0 & \sqrt{0.96} & 0 & 0 \\ 0 & 0.3 & 0 & \sqrt{0.91} & 0 \\ 0 & 0 & 0.25 & 0 & \sqrt{0.9375} \end{pmatrix}.$$

Thus these profiles contain additional speed/strike correlations, not only rotation/path correlation. Correlations involving the two zero-SD dimensions do not change the modeled covariance. Positive shaft

twist closes the face, explaining the opposite signs in the legacy profile names. Requested primitive covariance, realized nonlinear face/path covariance, and covariance calculated only among supported shots are distinct quantities; the reported delivered face/path diagnostic is conditioned on each pattern's supported shots.

Club and Contact Assumptions. Static loft and lie are taken from the cited model-year specifications; all remaining club properties are explicit hybrid research assumptions, not a reconstruction of those products. The finite club head has mass, a center of gravity behind the face, and a diagonal inertia tensor stated in the local face-normal/heel/face-up axes:

Club	Mass (kg)	CG (mm)	I_n (10^{-4} kg m ²)	I_h	I_u	$h_{\max}$ (mm)	$v_{\max}$ (mm)
Driver	0.200	35	2.5	4.0	4.5	50	27
7-Iron	0.272	10	1.2	2.0	2.5	35	23
Pitching Wedge	0.300	8	1.0	2.0	2.2	35	25

Here CG denotes depth behind the face; $h_{\max}$ and $v_{\max}$ are the face half-width and half-height, respectively. The tensor is rotated into world coordinates for each delivery. A toe/heel or high/low offset changes both the contact lever from head CG and, for the assumed driver bulge and roll radii of 300 mm, the local face normal. The iron and wedge faces are modeled as flat. Thus an off-center strike can change translation and rotation through the same contact impulse. The model sets pre-impact club angular velocity and ball velocity/spin to zero; shaft twist specifies club orientation, not a closure angular rate. Club angular velocity generated by impact is retained in the contact solution.

More precisely, let the central delivered face frame be $B = [\mathbf{n}_0, \mathbf{h}, \mathbf{u}]$, with $\mathbf{h}$ heelward and $\mathbf{u}$ face-up. For toe-positive offset a , high-positive offset b , CG depth d , and curvature radii R_B, R_R , the adapter uses a local *parabolic* surface:

$$\mathbf{r}_c = \left[d - \frac{1}{2}(a^2/R_B + b^2/R_R) \right] \mathbf{n}_0 - a\mathbf{h} + b\mathbf{u},$$

$$\mathbf{n} = \frac{\mathbf{n}_0 - (a/R_B)\mathbf{h} + (b/R_R)\mathbf{u}}{\|\mathbf{n}_0 - (a/R_B)\mathbf{h} + (b/R_R)\mathbf{u}\|}, \quad I_c = B \operatorname{diag}(I_n, I_h, I_u) B^T.$$

A zero radius denotes a flat direction: its reciprocal terms are omitted. The normal used in contact is the local $\mathbf{n}$, while inertia is oriented by the central frame B ; the radii do not define an exact spherical face. The rectangular supported face requires $|a| \leq h_{\max}$ and $|b| \leq v_{\max}$. For right-positive path p , attack angle γ , and speed v_c , the pre-impact translational velocity is

$$\mathbf{v}_c^- = v_c (\cos \gamma \cos p, -\cos \gamma \sin p, \sin \gamma)^T.$$

All trigonometric evaluations use radians after converting the degree-valued public inputs.

The primary impact law is a two-body rigid-body impulse calculation with Coulomb friction coefficient $\mu = 0.4$ and *energetic restitution* $e_E = 0.83$. In compact form, let $\mathbf{c}$ be the club-minus-ball relative velocity at contact, $\mathbf{J}$ the impulse on the ball, m_c, m_b the club and ball masses, $\mathbf{r}_c$ and $\mathbf{r}_b$ the CG-to-contact levers, and I_c, I_b the world-frame inertia tensors. The ball is modeled as a solid sphere with mass $m_b = 0.04593$ kg, radius $R = 0.021335$ m, and $I_b = (2/5)m_b R^2$ (the scalar multiplying the identity tensor). The contact-relative velocity and its impulse update are

$$\mathbf{c} = \mathbf{v}_c + \boldsymbol{\omega}_c \times \mathbf{r}_c - \mathbf{v}_b - \boldsymbol{\omega}_b \times \mathbf{r}_b.$$

The relative contact velocity changes as

$$\mathbf{c}^+ = \mathbf{c}^- - K\mathbf{J}, \quad K = \left(m_c^{-1} + m_b^{-1} \right) I - [\mathbf{r}_c]_{\times} I_c^{-1} [\mathbf{r}_c]_{\times} - [\mathbf{r}_b]_{\times} I_b^{-1} [\mathbf{r}_b]_{\times}.$$

The impulse updates both bodies' linear velocities and creates angular velocity through its moment arm:

$$\begin{aligned}\Delta \mathbf{v}_b &= \mathbf{J}/m_b, & \Delta \mathbf{v}_c &= -\mathbf{J}/m_c. \\ \Delta \boldsymbol{\omega}_b &= I_b^{-1}(\mathbf{r}_b \times \mathbf{J}), & \Delta \boldsymbol{\omega}_c &= -I_c^{-1}(\mathbf{r}_c \times \mathbf{J}).\end{aligned}$$

For the ball, $\mathbf{r}_b = -R\mathbf{n}$, so tangential impulse contributes spin about the ball center; for the club, an off-center impulse changes head rotation through the finite inertia tensor. Contact is integrated in normal-impulse increments. At each increment the tangential impulse is selected by a maximum-dissipation Coulomb-disk solve, with $\|d\mathbf{J}_t\| \leq \mu dJ_n$; normal compression stops when approach velocity reaches zero. During rebound the accumulated returned normal work is set to e_E^2 times compression work. This is why the primary coefficient is an energy restitution and not a prescribed ratio of normal velocities when normal and tangential motion are coupled. For an orthonormal tangent basis T , current relative velocity $\mathbf{c}$, and positive normal increment dP , the implemented disk problem is

$$\mathbf{q}_* = \arg \min_{\|\mathbf{q}\| \leq \mu dP} \left\{ \frac{1}{2} \mathbf{q}^T (T^T K T) \mathbf{q} - [T^T (\mathbf{c} - dP K \mathbf{n})]^T \mathbf{q} \right\}, \quad d\mathbf{J} = dP \mathbf{n} + T \mathbf{q}_*.$$

The positive compression and rebound work are

$$E_c = \int_{\text{compression}} c_n dP, \quad E_r = - \int_{\text{rebound}} c_n dP = e_E^2 E_c, \quad c_n = \mathbf{n}^T \mathbf{c}.$$

The nominal normal-impulse resolution is $N = 128$, defining $\Delta P_0 = c_n^-/[N(\mathbf{n}^T K \mathbf{n})]$; it is not a guarantee of exactly 128 compression increments. Work uses trapezoidal quadrature, and 48 bisection iterations locate the final compression and rebound increments. An eligible frictionless or block-diagonal sticking case uses the analytical fast path instead. The disk solve likewise uses a 48-iteration multiplier bisection when its unconstrained solution exceeds the Coulomb radius. Refinement evidence at $N = 256, 512$ is linked by the design reference.

A separate maximum-dissipation/Newton comparator enforces an instantaneous normal velocity restitution condition with the same numeric value 0.83 and a Coulomb impulse cone. That value has a different law definition there. The comparator is a sensitivity contrast, not the primary contact history and not a second empirical calibration. Conservation, passivity, and refinement checks support numerical consistency only. The rigid point-contact model does not represent calibrated face compliance, shaft flexibility, grooves, local friction variation, or measured ball/club launch behavior; external empirical validation remains open.

Flight and Endpoint. The computed post-impact ball velocity and spin initialize the same frozen Rust flight kernel documented for the historical study, with the phase-two contact model supplying the launch state. Its state is $(\mathbf{r}, \mathbf{v}, \omega)$, with constant spin-axis direction $\hat{\boldsymbol{\omega}}$ and angular-speed magnitude ω :

$$\dot{\mathbf{r}} = \mathbf{v}, \quad m_b \dot{\mathbf{v}} = m_b \mathbf{g} + \mathbf{F}_D + \mathbf{F}_M, \quad \dot{\omega} = -\lambda \omega, \quad \lambda = 0.05 \text{ s}^{-1}.$$

The launch adapter converts impact spin magnitude from rad/s to RPM at the Python launch boundary and the native wrapper converts it back to rad/s. The ODE has exponential spin decay; its seventh state is integrated by RK4, not an exact exponential update. The actual Python-supplied ball radius is $R = 0.021335 \text{ m}$, with the impact mass above and area $A_b = \pi R^2$. Still-air parameters are $\rho = 1.225 \text{ kg/m}^3$, $\mu_a = 1.81 \times 10^{-5} \text{ Pa s}$, and $\mathbf{g} = (0, 0, -9.80665) \text{ m/s}^2$. For relative air velocity $\mathbf{v}_a = \mathbf{v} - \mathbf{w}$ ($\mathbf{w} = 0$ here), $V = \|\mathbf{v}_a\|$, and $\boldsymbol{\omega} = \omega \hat{\boldsymbol{\omega}}$, the active forces are

$$\mathbf{F}_D = -\frac{1}{2} \rho C_D A_b V \mathbf{v}_a, \quad S = R \|\boldsymbol{\omega}\| / V, \quad C_L = \min(0.26, 0.70 S^{0.645}),$$

$$\mathbf{F}_M = \frac{1}{2} \rho C_L A_b V^2 \frac{\boldsymbol{\omega} \times \mathbf{v}_a}{\|\boldsymbol{\omega} \times \mathbf{v}_a\|}.$$

Only this single spin-induced force is added. Define

$$\text{Re}_c = \text{clip}(\rho V(2R)/\mu_a, 10^3, 10^7), \quad H(z; z_0, w) = \frac{1}{2}[1 + \tanh((z - z_0)/w)].$$

For the supplied base drag coefficient 0.25, the exact active drag curve is

$$C_D = \text{clip}(0.50 - 0.25H(\log_{10} \text{Re}_c; \log_{10}(70000), 0.16) \\ + 0.05H(\log_{10} \text{Re}_c; \log_{10}(700000), 0.45), 0.15, 0.55).$$

These coefficients are the frozen implementation’s empirical assumptions; they are not a new calibration to the modeled ball. The spin force uses total spin magnitude rather than its component perpendicular to air velocity. Its magnitude therefore does not decrease smoothly with axis alignment except at the near-zero cross-product cutoff. The source returns zero spin force when speed, spin magnitude, or the spin–velocity cross-product norm is below 10^{-6} in their respective SI numerical units; drag is zero below the same speed threshold. Arbitrary off-center spin predictions inherit this model limitation. Internal contact conservation does not validate the aerodynamic law.

All seven states use fourth-order Runge–Kutta at $\Delta t = 0.02$ s with a nominal 12 s horizon. The wrapper allows $\lfloor 12/\Delta t \rfloor + 1$ steps, giving a nominal upper bound of 12.02 s at the production step. The native ground-stop test is enabled only after $t > 0.05$ s to exclude the initial zero-height state. Failure to return a landing is a numerical failure, not a substituted airborne endpoint. Launch position is the origin at zero height. Between the last above-ground state and first nonpositive-height state, landing is

$$\mathbf{r}_{\text{land}} = \mathbf{r}_j + \frac{z_j}{z_j - z_{j+1}}(\mathbf{r}_{j+1} - \mathbf{r}_j).$$

The startup guard is an implementation qualification for extremely short flights: it can pass an earlier ground crossing before termination is enabled. This linear event interpolation does not inherit a fourth-order landing-time accuracy guarantee merely from RK4 state integration. First ground contact is treated as the hypothetical final lie for scoring. Bounce, roll, turf response, obstacles, wind variation, and elevation are not simulated. Unsupported off-face, nonapproaching, or nonairborne samples are kept with their failure status, never clipped or resampled. Endpoint distributions are therefore conditional on supported airborne shots when failures occur.

Conditional Scoring and Interpretation. The driver case assumes a flat 400 m hole with a 30 m fairway corridor; its carry endpoint is assigned a fairway or rough lie. Iron and wedge cases assume a fairway approach to a centered 15 m-radius green, with rough outside. A source-backed historical PGA expected-strokes benchmark is interpolated at those assumed lie and distance states. Fairway and rough values come from the historical published table; short-putt values use an anchor-reconciled approximate putting fit. The result is an approximate 2003–2010 population benchmark, not a current player-specific model. SG is consequently conditional on supported endpoints and these geometry/baseline choices; it is not measured performance or an expected per-round improvement. Within-condition paired SG contrasts use shot indices supported in all three patterns; between-condition contrasts use the common supported indices across all six populations. Bootstrap intervals quantify sampling variation for those simulated populations only. They do not include uncertainty in club properties, contact law, flight physics, human delivery distributions, course geometry, or the scoring benchmark.

Let A_{kc} be the supported shot-index set for pattern k in condition c , $A_{3c} = \bigcap_{k \in \{S,D,F\}} A_{kc}$, and $A_6 = A_{3c_1} \cap A_{3c_2}$. With endpoint scores $SG_{ikc} = J(d_{\text{start}}, \ell_{\text{start}}) - 1 - J(d_{ikc}, \ell_{ikc})$, the within-condition benefit and between-condition interaction are

$$\hat{B}_{kc} = \frac{1}{|A_{3c}|} \sum_{i \in A_{3c}} (SG_{ikc} - SG_{iSc}),$$

$$\widehat{\Delta B}_k = \frac{1}{|A_6|} \sum_{i \in A_6} [(SG_{ikc_2} - SG_{iSc_2}) - (SG_{ikc_1} - SG_{iSc_1})].$$

The order $c_2 - c_1$ follows the named contrast in each result table. Changing support can change the conditional population as well as the response; these estimands do not recover the unknown scores of unsupported outcomes. Failure probabilities remain separately reported, and unconditional expected scoring is unavailable without defensible terminal states for those failures. The saved mean-difference intervals use 500 paired bootstrap resamples of the appropriate common index set and empirical 2.5th/97.5th percentiles. Differences are formed before resampling. These pointwise intervals have finite bootstrap endpoint precision; a borderline zero crossing is not strong evidence of a stable sign. No multiplicity-adjusted discovery claim is made.

Endpoint Statistics and Binned Density. For a pattern's n supported reported endpoints $\mathbf{z}_i = (X_i, Y_i)^T$ and target $\mathbf{z}_T = (X_T, 0)^T$, let

$$\bar{\mathbf{z}} = n^{-1} \sum_i \mathbf{z}_i, \quad \hat{\Sigma}_z = \frac{1}{n-1} \sum_i (\mathbf{z}_i - \bar{\mathbf{z}})(\mathbf{z}_i - \bar{\mathbf{z}})^T, \quad \mathbf{b} = \bar{\mathbf{z}} - \mathbf{z}_T.$$

Lateral and downrange SDs are the square roots of the corresponding diagonal entries; covariance describes orientation, while target RMSE includes bias:

$$\text{RMSE}^2 = n^{-1} \sum_i \|\mathbf{z}_i - \mathbf{z}_T\|^2 = \|\mathbf{b}\|^2 + \frac{n-1}{n} \text{tr}(\hat{\Sigma}_z).$$

Radial target quantiles use $\|\mathbf{z}_i - \mathbf{z}_T\|$, and containment is the fraction within the stated radius. These are distinct from radial carry $\|\mathbf{z}_i\|$. A narrow one-dimensional endpoint locus is not evidence of a calibrated two-dimensional player dispersion ellipse.

The unsmoothed histogram uses common spatial bin edges across the compared patterns. For bin j of area a_j and count n_j , density is $\hat{f}_j = n_j/(na_j)$, so $\sum_j \hat{f}_j a_j = 1$ separately for each supported pattern population. Bins are ordered by density and retained through the threshold needed for 50%, 80%, or 95% cumulative mass; all tied bins at the cutoff are included. Achieved mass can therefore exceed the named level. Bin resolution and achieved mass are reporting quantities, not physical precision; these regions depend on the grid and are not smooth population density estimates or confidence regions for a mean. No Gaussian width or KDE is imposed on singular or nearly one-dimensional samples.

Supplemental Flight-Trace Replay. For the three baseline plates and Driver lean comparison, 100 distinct shot IDs are chosen at evenly spaced ranks of the support intersection across all five displayed cells and three patterns. IDs 0–4 are reserved for verification. Exact saved delivery inputs are replayed through the original impact and flight pathway; no statistical outcomes are resampled or replaced. Before accepting additional traces, each cell's nominal path and first five sampled paths must reproduce all 41 saved coordinates to 10^{-6} m absolute tolerance. Selected shots must reproduce saved carry, launch and delivered geometry (normally 10^{-6} in their stated units; spin rate 10^{-5} rpm). Runtime, native binary and physics-source identities must match the recorded execution. The separate replay receipt records all input hashes, common IDs, tolerances and output hash. These are illustrative replays conditional on common support, not another independent experiment; all original numerical summaries and manifests remain unchanged.

The equations above follow these preserved study adapters and shared kernels:

Stage	Implementation
Nominal Matching	<code>nominal_matching.py</code>
Delivery and Face Geometry	<code>matched_physics.py</code>
Primitive Sampling	<code>noise.py</code>
Incremental Contact	<code>incremental_contact_kernel.py</code>
Aerodynamic Forces	<code>aerodynamics.rs</code>
Flight Integration	<code>ball_flight.rs</code>

Their execution identities are recorded in the [Accepted Source Transition](#); the saved-data workflow retains those identities separately from later reporting changes. The local parabolic face, contact law, force coefficients, stochastic covariance, and hypothetical course geometry jointly define the model whose conditional outputs are tabulated below.

The saved original and supplemental simulations and independent data audits are complete. Numerical narrative and standalone graphics have scoped independent review. This illustrated article requires its own compiled-page review before publication. The retained phase-one results below must not be presented as phase-two outcomes.

3.2 Phase-Two Inventory and Statistical Populations

The original within-lean inventory contains 324 planned cells: 270 completed and 54 explicitly unavailable, with 8,100,000 attempted shot rows in completed cells. The fixed-target inventory contains 72 planned cells: 52 completed and 20 unavailable, with 1,560,000 attempted rows. The separately declared bounds supplement contains 26 distinct executed cells and 780,000 attempted rows; 779,712 are supported and 288 remain unsupported. Four control aliases reuse existing zero-lean supplemental cells and add no shots. Original unavailable records remain unchanged. Expanded bounds resolve some nominal-search failures; an original bounded-search failure is not proof of physical impossibility. See the [Primary Audit](#), [Control Audit](#), and [Supplement Completion](#).

Dispersion and directional statistics use each pattern's supported endpoints. Within-cell SG benefits use the common three-pattern supported indices. Between-condition relative benefits use the intersection of all six populations, with paired differences formed before bootstrap resampling. These denominators are different; unsupported shots are retained in accounting rather than retried. The selected examples below come from 24 zero-lean cells and 72 pattern rows, covering three clubs, two curve scales, two shaft SDs and shaft-only versus combined delivery noise. Four PW strong-curve cells use the explicitly labelled supplement. This selection is not the complete factorial inventory.

3.3 Rigid Shaft Twist, Delivered Loft, and Forward Lean

Phase two pitches a fixed physical club, unlike the historical phase-one axis-only lean control. Let static loft be L , lie be β , forward lean be ℓ , nominal right-positive face direction be f , and shaft twist be θ . In the native forward/left/up frame,

$$\mathbf{n}_0 = (\cos L, 0, \sin L)^T, \quad \mathbf{s}_0 = (0, \cos \beta, \sin \beta)^T,$$

$$Q = R_z(-f)R_y(\ell), \quad \mathbf{s} = Q\mathbf{s}_0, \quad \mathbf{n} = R_s(\theta)Q\mathbf{n}_0.$$

Here R_s is Rodrigues rotation about the delivered shaft axis. Zero twist gives delivered loft $\lambda = L - \ell$. The local sensitivities are

$$\frac{df}{d\theta} = -\frac{\sin \beta \cos L}{\cos(L - \ell)}, \quad \frac{d\lambda}{d\theta} = -\cos \beta, \quad \frac{d\lambda}{df} = \cot \beta \frac{\cos(L - \ell)}{\cos L}.$$

Thus closing the face by this assumed rotation tends to reduce loft. Fixed shaft-twist SD has the same first-order loft SD across lean, but not the same face-angle SD. Twist is an orientation error, not clubhead closure angular velocity. Finite-angle geometry and impact/flight determine the endpoints; these derivatives do not require every closed-face miss to travel farther. The [Phase-Two Design](#) and [Independent Derivative Checks](#) provide the equations and their numerical checks.

3.4 Club-Dependent Long-Left and Short-Right Behavior

For the straight, zero-lean, ordinary-curve reference with shaft-only noise, let X be downrange distance and Y the reported right-positive lateral coordinate. “Left” means $Y < 0$, “right” means $Y > 0$, and long/short refer to $X - X_T$, not total radial range. Each row below has 10,000 supported shots. The side-conditioned mean errors are descriptive sample quantities, not causal effects of selecting a side and not confidence intervals.

Club	Twist SD (°)	$\rho(Y, X)$	Left Error (m)	Right Error (m)	Mean Error (m)
Driver	1	+0.9788	-4.214	+3.190	-0.505
Driver	2	+0.9186	-9.513	+5.473	-2.005
7-Iron	1	-0.9960	+0.878	-0.984	-0.055
7-Iron	2	-0.9841	+1.637	-2.069	-0.220
Pitching Wedge	1	-0.9982	+0.896	-0.939	-0.023
Pitching Wedge	2	-0.9930	+1.740	-1.922	-0.095

Targets are 178.949 m for driver, 127.882 m for 7-iron, and 77.402 m for PW. In this conditional model the 7-iron and PW exhibit long-left/short-right association, while the driver exhibits the opposite association. This is consistent with loft reduction affecting launch, spin and ball speed together; it is not evidence that a universal carry rule follows from face closure. Increasing shaft variability also shifts realized mean downrange distance, even though the zero-error nominal is matched to the same target. The [Exact Selected Numeric Rows](#) and [Independent Field-by-Field Review](#) bind the source summaries and full-precision values. No new simulation or statistical fit supplies this table.

3.5 Distance Confounding and Scoring Interpretation

Nominal landing matching removes an initial target mismatch; it does not force equal realized carry distributions. Within-lean targets may differ across lean. Fixed-target controls hold the physical reference distance while rematching face, path and speed. Consequently, a change in spread or SG between these conditions cannot be uniquely attributed to shaft lean, curvature or distance alone. Lateral SD, downrange SD, bias and target error must be read together. Marginal statistics cannot partition their causal contributions.

For example, the zero-lean driver, ordinary-curve, one-degree shaft-only Draw has relative SG $+7.46156 \times 10^{-5}$ per shot, with paired 95% Monte Carlo interval $[-7.10405 \times 10^{-5}, 2.18961 \times 10^{-4}]$. Its sign alone does not establish an advantage; the interval includes zero. The estimate uses the hypothetical tee/fairway/rough context and the common three-pattern population of 10,000, not a measured golfer or a round score. Irons instead use the approach/green context. A smaller lateral width does not necessarily improve scoring when downrange error, target bias and lies change. Small sampling intervals do not include head-parameter uncertainty, impact discrepancy or the approximation in the historical SG benchmark.

3.6 Dependence on Assumed Head Properties

A separate deterministic sensitivity check holds face, path and shaft twist at zero, retains preset speed and attack angle, and does not rematch targets. It examines all three clubs and lean settings, with center strikes and 0.5 in toe, heel, high and low offsets. Of 300 variant evaluations, 268 are supported and 32 are unsupported. The offsets are fixed test positions, not the stochastic study's 0.25 in per-axis strike standard deviation.

Club	CG Depth	Inertia	Face Radii	Contact Law
Driver	5.14	7.10	22.63	11.33
7-Iron	0.54	1.76	—	1.22
Pitching Wedge	0.18	1.40	—	0.33

Values are maximum landing displacements in metres among comparable supported baseline/variant pairs. CG depth, principal inertias together, and driver bulge/roll radii together are varied by $\pm 20\%$. These deliberately chosen ranges are not measured confidence intervals. The last column instead changes contact law: the primary energetic restitution and comparator Newton velocity coefficient both use the number 0.83 but have different definitions. Maxima can occur at different deliveries; they are not additive or dispersion SDs. The driver's largest face-radius response occurs at 5° lean with a low strike and radius scale 0.8, illustrating sensitivity at low launch.

The [saved summary](#) and [full execution receipt](#) preserve these cases and their unsupported outcomes. This model dependence is distinct from finite-sample Monte Carlo uncertainty; small statistical intervals do not establish empirical impact-model accuracy.

3.7 Selected Zero-Lean Conditions and Exact Source Identities

The selected analysis contains 24 zero-lean conditions and 72 pattern rows: three clubs, two curve scales, two shaft-twist SDs, and two noise profiles. Each pattern attempts 10,000 shared-seed deliveries. These are selected strata of the full matrix, not independent player trials.

Protocol keys identify distinct source records: P = `original_primary/1`; OC = `original_control/1`; WP = `pw-primary-face-bounds/1`; WC = `pw-control-face-bounds/1`; DC = `driver-control-speed-bounds/1`.

The four PW curve-2 headline conditions use WP. Original bounded-search unavailable records remain preserved. Original matching searches reference face $\pm 3^\circ$; WP and WC use absolute face bounds $[-6, 6]^\circ$ with speed factors $[0.75, 1.25]$. These supplemental bounds were declared after feasibility review, before supplemental sampling; they are not the original protocol.

Keys identify the cell names `<club>_lean0_curve<C>_rotation<SD>_<profile>`. Club identifiers are `driver`, `seven_iron`, and `pitching_wedge`. S denotes `shaft_only`; C denotes `combined2`. The profile C adds path SD 0.5° , speed SD 2 mph (0.89408 m/s), and horizontal and vertical strike SD 0.25 in (0.00635 m) each, with independent primitive errors. S holds the other primitive errors at zero. Neither profile makes delivered face SD equal to shaft-twist SD.

Table 1: Selected Condition and Protocol Keys

Key	Club	Curve	Twist	SD (°)	Profile	Protocol
C01	Driver	1	1		S	P
C02	Driver	1	1		C	P
C03	Driver	1	2		S	P
C04	Driver	1	2		C	P
C05	Driver	2	1		S	P
C06	Driver	2	1		C	P
C07	Driver	2	2		S	P
C08	Driver	2	2		C	P
C09	7-Iron	1	1		S	P
C10	7-Iron	1	1		C	P
C11	7-Iron	1	2		S	P
C12	7-Iron	1	2		C	P
C13	7-Iron	2	1		S	P
C14	7-Iron	2	1		C	P
C15	7-Iron	2	2		S	P
C16	7-Iron	2	2		C	P
C17	Pitching Wedge	1	1		S	P
C18	Pitching Wedge	1	1		C	P
C19	Pitching Wedge	1	2		S	P
C20	Pitching Wedge	1	2		C	P
C21	Pitching Wedge	2	1		S	WP
C22	Pitching Wedge	2	1		C	WP
C23	Pitching Wedge	2	2		S	WP
C24	Pitching Wedge	2	2		C	WP

Targets are Driver 178.948758 m, 7-Iron 127.881942 m, and PW 77.402477 m. Face, path and speed are solved before sampling, with no post-flight aim correction. The maximum residual among the 72 selected nominal solutions is $+1.999733\text{e-}13$ m, below the predeclared 0.05 m criterion. Matching nominal endpoints does not constrain sampled mean carry or remove longitudinal error.

3.8 Dispersion, Absolute Scoring, and Within-Condition Benefits

All geometry columns below use each pattern's individual supported endpoints. n_e is that endpoint count; every row attempts 10,000. Absolute SG and relative SG use the common-three scored population n_3 shown in the following tables, which can differ from individual support. Unsupported cases remain in accounting and are not resampled. RMSE is radial target error; σ_Y and σ_X are lateral and downrange SDs, all in metres. SG units are strokes per modeled opportunity. Numerical notation such as $2.5\text{e-}5$ means 2.5×10^{-5} ; interval endpoints retain seven significant digits. Display rounding does not add model accuracy.

Table 2: Driver Endpoint Geometry and Absolute SG

Key	Pattern	n_e	σ_Y	σ_X	RMSE	Mean SG
C01	Straight	10000	6.177	4.672	7.764	-0.395737
C01	Draw	10000	5.996	5.219	7.973	-0.3956624
C01	Fade	10000	6.266	4.182	7.549	-0.3956688
C02	Straight	10000	6.686	8.561	11.181	-0.4125945
C02	Draw	10000	6.534	8.863	11.337	-0.4125676
C02	Fade	10000	6.777	8.290	11.018	-0.4125523

Table 2: Driver Endpoint Geometry and Absolute SG (Continued)

Key	Pattern	n_e	σ_Y	σ_X	RMSE	Mean SG
C03	Straight	10000	12.011	9.781	15.648	-0.4428029
C03	Draw	10000	11.660	10.817	16.090	-0.4399103
C03	Fade	10000	12.210	8.808	15.190	-0.4440281
C04	Straight	10000	12.133	12.236	17.746	-0.4577065
C04	Draw	10000	11.804	13.072	18.160	-0.4555113
C04	Fade	10000	12.328	11.461	17.313	-0.4586342
C05	Straight	10000	6.177	4.672	7.764	-0.395737
C05	Draw	10000	5.713	5.800	8.168	-0.395412
C05	Fade	10000	6.285	3.764	7.340	-0.3955027
C06	Straight	10000	6.686	8.561	11.181	-0.4125945
C06	Draw	10000	6.312	9.186	11.476	-0.4120319
C06	Fade	10000	6.824	8.048	10.854	-0.4124999
C07	Straight	10000	12.011	9.781	15.648	-0.4428029
C07	Draw	10000	11.152	11.868	16.493	-0.4350053
C07	Fade	10000	12.276	7.940	14.735	-0.4442635
C08	Straight	10000	12.133	12.236	17.746	-0.4577065
C08	Draw	10000	11.336	13.932	18.533	-0.4510242
C08	Fade	10000	12.407	10.774	16.882	-0.4585412

Table 3: Driver Paired Benefit Relative to Straight

Key	Pattern	n_3	Estimate	Lower 95%	Upper 95%
C01	Draw	10000	+7.461556e-05	-7.104047e-05	+0.0002189608
C01	Fade	10000	+6.825953e-05	-5.227718e-05	+0.0001704984
C02	Draw	10000	+2.691177e-05	-0.0001641751	+0.0002203737
C02	Fade	10000	+4.217545e-05	-0.000173919	+0.0002645741
C03	Draw	10000	+0.002892622	+0.002462301	+0.003378319
C03	Fade	10000	-0.00122512	-0.001619667	-0.0008628542
C04	Draw	10000	+0.00219516	+0.001694901	+0.00263646
C04	Fade	10000	-0.0009276973	-0.001382425	-0.0005146906
C05	Draw	10000	+0.0003250171	+0.0001041999	+0.000551286
C05	Fade	10000	+0.000234387	+5.011171e-05	+0.0004289665
C06	Draw	10000	+0.0005625981	+0.0002392218	+0.0008977671
C06	Fade	10000	+9.461603e-05	-0.0002218223	+0.0004418635
C07	Draw	10000	+0.007797655	+0.007132512	+0.00851801
C07	Fade	10000	-0.001460539	-0.001974284	-0.0009630427
C08	Draw	10000	+0.00668231	+0.005938053	+0.007437962
C08	Fade	10000	-0.0008347119	-0.001478401	-0.0002621467

Table 4: 7-Iron Endpoint Geometry and Absolute SG

Key	Pattern	n_e	σ_Y	σ_X	RMSE	Mean SG
C09	Straight	10000	2.559	1.164	2.812	+0.5365352
C09	Draw	10000	2.598	1.086	2.816	+0.5360296
C09	Fade	10000	2.513	1.242	2.804	+0.5375821
C10	Straight	9996	2.551	5.090	6.215	+0.1818424
C10	Draw	9996	2.588	5.074	6.217	+0.1810974
C10	Fade	9996	2.509	5.107	6.211	+0.1827358
C11	Straight	10000	5.104	2.330	5.617	+0.2950796

Table 4: 7-Iron Endpoint Geometry and Absolute SG (Continued)

Key	Pattern	n_e	σ_Y	σ_X	RMSE	Mean SG
C11	Draw	10000	5.178	2.175	5.622	+0.2945804
C11	Fade	10000	5.015	2.483	5.602	+0.2961466
C12	Straight	9996	4.993	5.448	7.853	+0.07221942
C12	Draw	9996	5.063	5.387	7.858	+0.07141414
C12	Fade	9996	4.909	5.513	7.844	+0.07333032
C13	Straight	10000	2.559	1.164	2.812	+0.5365352
C13	Draw	10000	2.629	1.006	2.815	+0.5361373
C13	Fade	10000	2.460	1.319	2.792	+0.5390877
C14	Straight	9996	2.551	5.090	6.215	+0.1818424
C14	Draw	9996	2.619	5.058	6.219	+0.1804867
C14	Fade	9996	2.463	5.125	6.208	+0.1838173
C15	Straight	10000	5.104	2.330	5.617	+0.2950796
C15	Draw	10000	5.236	2.018	5.618	+0.2947913
C15	Fade	10000	4.911	2.635	5.581	+0.2976612
C16	Straight	9996	4.993	5.448	7.853	+0.07221942
C16	Draw	9996	5.120	5.328	7.857	+0.07116655
C16	Fade	9996	4.812	5.579	7.831	+0.07451322

Table 5: 7-Iron Paired Benefit Relative to Straight

Key	Pattern	n_3	Estimate	Lower 95%	Upper 95%
C09	Draw	10000	-0.0005056061	-0.000513691	-0.0004968345
C09	Fade	10000	+0.001046956	+0.001033213	+0.001059063
C10	Draw	9996	-0.0007449974	-0.0008996072	-0.0005794542
C10	Fade	9996	+0.0008933618	+0.0007089666	+0.001077776
C11	Draw	10000	-0.0004991576	-0.0005776303	-0.0004544381
C11	Fade	10000	+0.001066956	+0.0009860656	+0.001182346
C12	Draw	9996	-0.0008052724	-0.001037565	-0.0005917861
C12	Fade	9996	+0.001110901	+0.0008441734	+0.001355484
C13	Draw	10000	-0.0003978168	-0.0004105794	-0.0003851089
C13	Fade	10000	+0.002552569	+0.002520034	+0.002580022
C14	Draw	9996	-0.00135572	-0.001651328	-0.001025277
C14	Fade	9996	+0.001974931	+0.001622698	+0.002286496
C15	Draw	10000	-0.0002882714	-0.0004426203	-0.0001305862
C15	Fade	10000	+0.002581634	+0.002452823	+0.002736362
C16	Draw	9996	-0.001052864	-0.001428442	-0.0006655063
C16	Fade	9996	+0.002293802	+0.001930552	+0.002645474

Table 6: Pitching Wedge Endpoint Geometry and Absolute SG

Key	Pattern	n_e	σ_Y	σ_X	RMSE	Mean SG
C17	Straight	10000	1.102	1.145	1.589	+0.569051
C17	Draw	10000	1.120	1.092	1.564	+0.5734849
C17	Fade	10000	1.082	1.198	1.615	+0.5646011
C18	Straight	9999	1.074	4.001	4.512	+0.1911866
C18	Draw	9999	1.092	3.987	4.504	+0.1912817
C18	Fade	9999	1.056	4.015	4.520	+0.1912805
C19	Straight	10000	2.201	2.286	3.176	+0.3439965
C19	Draw	10000	2.237	2.180	3.126	+0.3495231

Table 6: Pitching Wedge Endpoint Geometry and Absolute SG (Continued)

Key	Pattern	n_e	σ_Y	σ_X	RMSE	Mean SG
C19	Fade	10000	2.161	2.392	3.227	+0.3385095
C20	Straight	9999	2.138	4.431	5.260	+0.1011112
C20	Draw	9999	2.172	4.380	5.232	+0.1015956
C20	Fade	9999	2.101	4.483	5.287	+0.1007775
C21	Straight	10000	1.102	1.145	1.589	+0.569051
C21	Draw	10000	1.137	1.038	1.540	+0.577929
C21	Fade	10000	1.060	1.252	1.640	+0.5601166
C22	Straight	9999	1.074	4.001	4.512	+0.1911866
C22	Draw	9999	1.109	3.971	4.495	+0.1914379
C22	Fade	9999	1.036	4.028	4.526	+0.1914785
C23	Straight	10000	2.201	2.286	3.176	+0.3439965
C23	Draw	10000	2.268	2.073	3.076	+0.3551136
C23	Fade	10000	2.118	2.498	3.278	+0.3330305
C24	Straight	9999	2.138	4.431	5.260	+0.1011112
C24	Draw	9999	2.203	4.328	5.203	+0.1021577
C24	Fade	9999	2.060	4.535	5.315	+0.100484

Table 7: Pitching Wedge Paired Benefit Relative to Straight

Key	Pattern	n_3	Estimate	Lower 95%	Upper 95%
C17	Draw	10000	+0.004433926	+0.004363272	+0.004501976
C17	Fade	10000	-0.004449852	-0.004516689	-0.004380068
C18	Draw	9999	+9.508286e-05	-6.220824e-05	+0.0002799946
C18	Fade	9999	+9.389912e-05	-3.145028e-05	+0.0002411445
C19	Draw	10000	+0.005526609	+0.005465531	+0.005576312
C19	Fade	10000	-0.005486948	-0.005537317	-0.005426622
C20	Draw	9999	+0.0004843982	+0.0003075227	+0.0006522765
C20	Fade	9999	-0.0003337114	-0.0005016453	-0.0001720514
C21	Draw	10000	+0.008878062	+0.008737846	+0.009015448
C21	Fade	10000	-0.008934324	-0.00906791	-0.008797752
C22	Draw	9999	+0.000251305	-3.604852e-05	+0.0005325912
C22	Fade	9999	+0.0002919228	+3.619769e-05	+0.0005921619
C23	Draw	10000	+0.0111171	+0.01099354	+0.01121934
C23	Fade	10000	-0.01096597	-0.01106806	-0.01084432
C24	Draw	9999	+0.001046484	+0.0007406396	+0.001332323
C24	Fade	9999	-0.0006271744	-0.000968112	-0.0002936083

The scoring context is a 400 m tee opportunity for Driver, with a 15 m fairway half-width and rough outside. The 7-Iron and PW use fairway approach opportunities at their respective nominal target distances, with a centered circular green of radius 15 m and rough outside. First ground contact is treated as the stopping lie: bounce, roll and hazards are not modeled. The baseline is the historical Broadie-derived approximation, version `table9-full-plus-anchor-reconciled-putting/2`; positive-distance unholed putting expectations are bounded below by one stroke. Absolute SG is benchmark-dependent and is not a predicted handicap or per-round gain.

3.9 Paired SD and Curve-Strength Interactions

For pattern P , each interaction is the rowwise difference $[(SG_P - SG_S)_B - (SG_P - SG_S)_A]$ on the common supported indices of all six populations. Its interval is computed directly from paired resampling, not

by subtracting marginal confidence limits. The source pairs below use OC where indicated, even when their cell names also occur in the headline P matrix. A C-key in this table identifies the cell name only; the adjoining protocol code completes the identity.

Table 8: Ordered Interaction Source Keys

Key	Contrast	A	B
E01	Twist SD 1 to 2	C01/OC	C03/OC
E02	Curve 1 to 2	C03/OC	C07/OC
E03	Twist SD 1 to 2	C02/OC	C04/OC
E04	Curve 1 to 2	C04/OC	C08/OC
E05	Twist SD 1 to 2	C09/OC	C11/OC
E06	Curve 1 to 2	C11/OC	C15/OC
E07	Twist SD 1 to 2	C10/OC	C12/OC
E08	Curve 1 to 2	C12/OC	C16/OC
E09	Twist SD 1 to 2	C17/OC	C19/OC
E10	Curve 1 to 2	C19/OC	C23/WP
E11	Twist SD 1 to 2	C18/OC	C20/OC
E12	Curve 1 to 2	C20/OC	C24/WP

Table 9: Direct Common-Six Relative-Pattern Interactions

Key	Pattern	n_6	Estimate	Lower 95%	Upper 95%
E01	Draw	10000	+0.002818007	+0.002345731	+0.003287813
E01	Fade	10000	-0.001293379	-0.001707546	-0.0008956275
E02	Draw	10000	+0.004905033	+0.004382983	+0.005452311
E02	Fade	10000	-0.0002354194	-0.0005740319	+0.0001123525
E03	Draw	10000	+0.002168249	+0.001600113	+0.002706174
E03	Fade	10000	-0.0009698728	-0.001502144	-0.0005027993
E04	Draw	10000	+0.00448715	+0.003937666	+0.005066104
E04	Fade	10000	+9.298539e-05	-0.0002735216	+0.0004641632
E05	Draw	10000	+6.448547e-06	-7.089107e-05	+4.9342e-05
E05	Fade	10000	+1.999969e-05	-5.979478e-05	+0.0001357949
E06	Draw	10000	+0.0002108861	+9.163611e-05	+0.0003577932
E06	Fade	10000	+0.001514679	+0.00143332	+0.001628772
E07	Draw	9996	-6.027502e-05	-0.0002966133	+0.0001539902
E07	Fade	9996	+0.0002175397	-7.965288e-05	+0.0005113381
E08	Draw	9996	-0.0002475919	-0.0004948098	+2.780392e-05
E08	Fade	9996	+0.001182901	+0.0009817452	+0.001364682
E09	Draw	10000	+0.001092683	+0.001026405	+0.001163565
E09	Fade	10000	-0.001037096	-0.001104793	-0.0009711341
E10	Draw	10000	+0.005590492	+0.005528753	+0.005643235
E10	Fade	10000	-0.00547902	-0.005530369	-0.005418491
E11	Draw	9999	+0.0003893153	+0.0002184871	+0.0005505622
E11	Fade	9999	-0.0004276105	-0.0005452449	-0.0003149902
E12	Draw	9999	+0.0005620857	+0.0004164435	+0.0007166807
E12	Fade	9999	-0.0002934631	-0.0004910577	-8.723418e-05

In these selected controls, increasing twist SD from 1° to 2° increases Draw relative to Straight for Driver and PW in both profiles. The 7-Iron SD interactions span zero; that is not proof of equivalence or negligible effects. Increasing curve scale from 1 to 2 at twist SD 2° increases Driver Draw, increases 7-Iron Fade, and increases PW Draw while decreasing PW Fade in both profiles. These are conditional comparisons of complete matched deliveries. They do not isolate curvature as a causal intervention.

3.10 Fixed-Target Across-Lean Controls

The two selected controls both compare lean 0° to 5° , curve scale 2, twist SD 2° , and combined2 noise. F1 is Driver: OC/driver_lean0_curve2_rotation2_combined2 to DC/driver_lean5_curve2_rotation2_combined2. F2 is PW: WP/pitching_wedge_lean0_curve2_rotation2_combined2 to WC/pitching_wedge_lean5_curve2_rotation2_combined2. F2's zero-lean data also supplies a verified fixed-target control alias; it is reused data, not another simulation.

Driver F1 holds target 178.948758 m; PW F2 holds 77.402477 m. DC uses absolute face bounds $[-6, 6]^\circ$ and speed factors $[0.75, 1.30]$, whereas its OC reference uses the original reference-centered face search and $[0.75, 1.25]$ speed factors. PW WP and WC both use absolute $[-6, 6]^\circ$ and $[0.75, 1.25]$. The Driver contrast therefore includes an explicitly qualified speed-bound exception, not merely a lean change.

Table 10: Signed Nominal Inputs for Fixed-Target Controls

Key	Pattern	Face ($^\circ$)	Path ($^\circ$)	Speed (m/s)	Residual (m)
F1A	Straight	0.00000	0.00000	45.00000	+0
F1A	Draw	4.03262	7.03262	45.25281	+3.032587e-14
F1A	Fade	-4.03262	-7.03262	45.25281	+3.032587e-14
F1B	Straight	0.00000	0.00000	55.40794	+1.421085e-13
F1B	Draw	4.89485	7.89485	56.34507	+3.69519e-13
F1B	Fade	-4.89485	-7.89485	56.34507	+3.69519e-13
F2A	Straight	0.00000	0.00000	32.00000	+0
F2A	Draw	-0.26046	2.73954	32.04591	+1.136886e-13
F2A	Fade	0.26046	-2.73954	32.04591	+1.136886e-13
F2B	Straight	0.00000	0.00000	29.11782	+1.563194e-13
F2B	Draw	-0.07671	2.92329	29.16131	+2.845955e-14
F2B	Fade	0.07671	-2.92329	29.16131	+2.845955e-14

Table 11: Fixed-Target Common-Six Relative-Pattern Interactions

Key	Pattern	n_6	Estimate	Lower 95%	Upper 95%
F1	Draw	9940	+0.005980075	+0.004831458	+0.006995762
F1	Fade	9940	-0.00139745	-0.002337964	-0.0004807456
F2	Draw	9999	-0.0003513245	-0.0005602612	-0.000142443
F2	Fade	9999	+0.0006868012	+0.0004851194	+0.0008858669

The Driver common-six intersection excludes 60 of the 10,000 shared indices; the PW intersection excludes one. These counts are union exclusions across six populations, not counts of one diagnosed failure mechanism. At zero twist the geometric loft is static loft minus lean: Driver changes from 10° to 5° , PW from 46° to 41° . The solved speeds increase for Driver and decrease for PW as shown above. This is a model result, not evidence that increased kinetic energy caused unsupported shots or that an aerodynamic optimum was identified. Nominal face and path also change. Holding the target fixed removes the between-condition nominal target-distance difference, but neither holds realized carry fixed nor identifies a pure causal lean, curvature, or distance effect.

3.11 Conditional Synthesis and Limits of Inference

Table 12: Selected Relative-SG Ranges and Pointwise Interval Directions

Club	Pattern	Minimum	Maximum	Pos/Neg/Spans 0
Driver	Draw	+2.691177e-05	+0.007797655	6/0/2
Driver	Fade	-0.001460539	+0.000234387	1/4/3
7-Iron	Draw	-0.00135572	-0.0002882714	0/8/0
7-Iron	Fade	+0.0008933618	+0.002581634	8/0/0
Pitching Wedge	Draw	+9.508286e-05	+0.0111171	6/0/2
Pitching Wedge	Fade	-0.01096597	+0.0002919228	1/6/1

For the eight selected 7-Iron conditions, Fade intervals are positive and Draw intervals negative. Driver and PW have different signs and zero-spanning cases as tabulated. Consequently these selected modeled comparisons do not support a general recommendation to prefer one shape across clubs and delivery distributions. Counts of pointwise interval directions are descriptive, not multiplicity-adjusted discoveries. No interval here includes uncertainty in the physical model, course model, clubhead assumptions or human delivery distribution.

The geometric derivative above establishes that face closure and delivered loft are coupled for this shaft-axis convention; it does not by itself establish carry direction. The earlier source-bound Straight examples show opposite lateral-downrange correlation signs for Driver versus the two irons. Their observed long-left/short-right or short-left/long-right patterns are local results of the complete impact and flight model. Marginal dispersion or carry changes cannot quantitatively decompose the SG difference into a distance effect and a shape effect.

3.12 Numerical Traceability and Publication Status

The selected 24 conditions and 72 rows are bound by [Headline Traceability](#) to 52 summary/manifest hashes. The direct interactions and controls use the actual selectors in `final_report_inputs.py` and the [384-Comparison Numerical Ledger](#): 164 SD, 164 curve-strength and 56 across-lean comparisons. All seven numerical artifacts were byte-reconciled against [Expected Numerical Hashes](#). The numerical qualification and the separate saved-graphic visual reviews are complete. This revised edition replots 13 plates from those saved inputs with conventional axis placement and adds a graphical summary. The [Revised Article Review](#) binds the assembled PDF to its source and figure hashes; the frozen original edition and its receipts remain separate.

The full matrix inventories remain 324 primary cells (270 completed, 54 unavailable; 8.1 million attempted rows), 72 original fixed-target cells (52 completed, 20 unavailable; 1.56 million rows), and 26 supplemental cells (780,000 rows, 779,712 supported and 288 unsupported). Four verified aliases add no shots. Source audits are [Primary](#), [Control](#), and [Supplemental Completion](#). Unavailable bounded-search records are not relabeled as physical impossibility.

All results remain conditional on assumed head geometry/inertia, the incremental Coulomb/energetic-restitution contact law, the numerical flight model, primitive Gaussian errors and scoring approximations described in the design reference. Gaussian tails are not certified by finite deterministic refinement grids. Shared seeded draws correlate cells; the pointwise bootstrap intervals express Monte Carlo sampling uncertainty, not independent empirical replications. The documented head-parameter and contact-law sensitivity is not a measured uncertainty distribution. Empirical qualification remains deferred under issue #11984. This editable article compiles successfully with local Tectonic 0.15.0. The editor preview remains unavailable because its dependency-bundle download fails before source parsing (issue #12120). Repository release gates and empirical qualification remain separate from document compilation.

A Historical Phase-One Status and Scope

All sections from this point onward describe phase one. Their delivery geometry, impact model, aiming procedure, seed and results are historical; they are not the phase-two matched-delivery method or results.

The original four driver result bundles are historical controls from a superseded impact approximation: tangential impulse changed spin but not linear velocity. The completed corrected three-club experiments replace their conclusions. Fixed loft also excluded face-loft covariance.

This is an editable standalone experiment reference, separate from the governed engineering design manual. The canonical manual remains the QMD source under `manuals/upstreamdrift`. This experiment does not qualify the underlying impact or flight model against measured shot populations. It estimates conditional outcome variance from face-angle error, not a golfer’s ability to produce a chosen pattern. All results are first ground contact (carry); bounce and roll are excluded.

B Inputs, Frames, and Pairing

User-facing face and path angles are measured from the target, positive right for a right-handed golfer. Straight uses $(f, p) = (0, 0)$ degrees, draw uses $(1.5, 3)$, and fade uses $(-1.5, -3)$. The doubled-curve trials use $(3, 6)$ and $(-3, -6)$ while retaining the straight baseline. The draw face is therefore 1.5 degrees left of its path. The repository solver uses x forward, y left, and z up; the adapter changes the sign of horizontal input and output angles once. The CSV diagnostic `spin_axis_tilt_deg` instead retains the native axis tilt convention, positive toward draw/left curvature. It is not Trackman’s right-positive spin-axis convention. Every pattern receives the same 10,000 independent errors $\epsilon_i \sim \mathcal{N}(0, \sigma_F^2)$, with separate trials at $\sigma_F = 1$ and 2 degrees. In each trial, $f_i = \bar{f} + \epsilon_i$ and $p_i = \bar{p}$. Shared errors provide paired comparisons; finite draw and fade samples need not be exact mirrors. The random seed is 20261008, using NumPy’s default generator.

The corrected study uses three illustrative hybrid delivery presets:

Club	Speed (m/s)	Loft (°)	Attack (°)	Lie (°)	Mass (kg)
Driver	45	12.8	-0.9	58.5	0.200
7-Iron	36	24.0	-4.0	63.0	0.272
Pitching Wedge	32	36.7	-5.0	64.0	0.300

These are assumptions for comparative modeling, not measured player inputs or reproduced specifications of one complete club. Trackman’s dynamic-loft page reports driver 12.8° as a Tour average; its PW 36.7° is an optimizer example at 72 mph, not a measured Tour PW mean. The 7-iron 24° delivered loft is an explicit assumption. Manufacturer static lie angles supply assumed shaft elevations. Speeds, attack angles, and masses are illustrative inputs unless the source receipt separately identifies a measurement. Nominal lean is zero; axis sensitivity explores 0, 10, and 20 degrees where the geometry is valid.

The experiment matrix is three clubs, two face standard deviations, two curve magnitudes, and two delivery modes (fixed loft and assumed shaft rotation): 24 bundles, each with 10,000 shots per shape. The ball begins stationary without spin; strike is central. Restitution, friction, ball mass and radius, and atmospheric parameters are recorded in receipts. There is no wind, speed error, strike error, or path error. Driver scoring uses a hypothetical tee/fairway/rough context; 7-iron and PW scoring use approach/green contexts. The earlier 45 m/s, 10.9° loft, zero-attack driver case is retained only in superseded historical bundles.

C Corrected Central-Contact Impact Equations

The experiment calls the public `ImpactSolverAPI` rigid-body model. For unit face normal $\mathbf{n}$, relative centre velocity $\mathbf{v}_r = \mathbf{v}_c^- - \mathbf{v}_b^-$, ball mass m_b , club mass m_c , restitution e , ball radius R , and inertia $I_b = 2m_b R^2/5$:

$$m_* = (m_b^{-1} + m_c^{-1})^{-1}, \quad J_n = (1 + e)m_* \mathbf{v}_r \cdot \mathbf{n}.$$

The ball contact point is $-R\mathbf{n}$ from its centre. Tangential surface slip, including initial ball spin, and its effective inverse mass are

$$\mathbf{u}_t = \mathbf{v}_r - (\mathbf{v}_r \cdot \mathbf{n})\mathbf{n} + R(\boldsymbol{\omega}_b^- \times \mathbf{n}), \quad D = m_b^{-1} + m_c^{-1} + R^2/I_b.$$

For positive normal approach, Coulomb friction supplies

$$\mathbf{J}_t = \min\{\mu J_n, \|\mathbf{u}_t\|/D\} \frac{\mathbf{u}_t}{\|\mathbf{u}_t\|}, \quad \mathbf{J} = J_n \mathbf{n} + \mathbf{J}_t.$$

The tangential impulse is zero at zero slip or zero friction. Both linear velocities and ball spin receive their corresponding impulse:

$$\mathbf{v}_b^+ = \mathbf{v}_b^- + \mathbf{J}/m_b, \quad \mathbf{v}_c^+ = \mathbf{v}_c^- - \mathbf{J}/m_c, \quad \boldsymbol{\omega}_b^+ = \boldsymbol{\omega}_b^- + (R/I_b)(\mathbf{J}_t \times \mathbf{n}).$$

Earlier result bundles omitted $\mathbf{J}_t$ from both translational updates and used the infinite-club sticking cap $(2/7)m_b\|\mathbf{u}_t\|$. Those bundles are superseded historical controls. Regression tests first failed against those defects, then passed with the correction. Independent arbitrary-3D tests check linear and angular momentum, normal restitution, nonincreasing kinetic energy, sticking contact, initial spin, and rotation covariance. These establish internal consistency of central contact with a translating club; off-centre club rotation, shaft response, and measured launch accuracy remain unqualified.

D Geometric Face–Loft Coupling and Shaft Lean

This is an explicitly assumed rigid geometry, not observed player data. The native right-handed frame is $+X$ forward, $+Y$ left, and $+Z$ up. User face angle is positive right. At zero twist, the delivered face normal is $\mathbf{n}_0 = (\cos L, 0, \sin L)$ for nominal delivered loft L . Given assumed shaft elevation ℓ and forward lean a :

$$s_x = \sin \ell \tan a, \quad \mathbf{s} = (s_x, \sqrt{\cos^2 \ell - s_x^2}, \sin \ell).$$

This construction is valid only when its square-root argument is nonnegative. Rodrigues rotation about the unit shaft axis gives

$$\mathbf{n}(\theta) = \mathbf{n}_0 \cos \theta + (\mathbf{s} \times \mathbf{n}_0) \sin \theta + \mathbf{s}(\mathbf{s} \cdot \mathbf{n}_0)(1 - \cos \theta).$$

Delivered face and loft are extracted as

$$F = \text{atan2}(-n_y, n_x), \quad L_d = \text{atan2}(n_z, \sqrt{n_x^2 + n_y^2}).$$

Near zero twist,

$$\frac{dF}{d\theta} = -(s_z - s_x \tan L), \quad \frac{dL_d}{d\theta} = -s_y, \quad \frac{dL_d}{dF} = \frac{s_y}{s_z - s_x \tan L}.$$

With zero lean, $dL_d/dF = \cot \ell$: closing tends to reduce loft in this specific rotation mode. Rotation about the vertical axis instead preserves loft. Face angle alone therefore does not identify delivered loft.

The main comparison applies this geometry to the face error $F - F_0$ around each pattern's nominal face F_0 , then rotates the axis and normal together by the nominal target yaw. All shapes share the same zero-error delivered loft, speed, attack angle, and strike. This isolates sensitivity to the assumed face-loft coupling without giving one shape a different baseline loft. Fixed-loft trials provide the corresponding control. Separate shaft-lean sensitivity holds nominal loft fixed and changes only the axis; it does not represent pitching a fixed physical club forward. That latter operation changes nominal loft as well and must be analyzed separately.

Less delivered loft does not necessarily increase carry. Launch, spin, speed, and aerodynamic trajectory jointly determine carry, so long-left/short-right behavior is measured from the modeled endpoint covariance rather than imposed. Static manufacturer lie is used only as an assumed delivered shaft elevation; dynamic lie, shaft bending, and golfer covariance are unmeasured.

E Flight Integration and Landing

Each shot calls the existing `BallFlightSimulator`, backed by the `upstream-physics` Rust kernel. Launch velocity uses metres per second, launch and azimuth angles radians, and spin magnitude RPM at the Python boundary, converted to radians per second inside the kernel. The spin axis is normalized. The solver integrates position, velocity and decaying spin with RK4, drag and spin-induced aerodynamic force:

$$\dot{\mathbf{r}} = \mathbf{v}, \quad m_b \dot{\mathbf{v}} = m_b \mathbf{g} + \mathbf{F}_D + \mathbf{F}_M, \quad \dot{\omega} = -\lambda \omega.$$

The production step is 0.02s with a 12s horizon. Landing is linearly interpolated between the last above-ground and first nonpositive-height samples. Failure to land raises an error; it never silently treats the final airborne sample as a landing. The Rust coefficient path is distinct from the constant-spin Waterloo/Penner model in `flight_models.py`; their calibrations must not be mixed. The kernel applies a single spin force (not separate additive backspin lift and sidespin Magnus forces). With $A = \pi R^2$ and relative airspeed V :

$$\mathbf{F}_D = -\frac{1}{2} \rho C_D A V \mathbf{v}, \quad S = R|\omega|/V, \quad C_L = \min(0.26, 0.70S^{0.645}),$$

$$\mathbf{F}_M = \frac{1}{2} \rho C_L A V^2 \frac{\omega \times \mathbf{v}}{|\omega \times \mathbf{v}|}.$$

The zero-cross-product case returns zero spin force. The shared drag-crisis curve evaluates Reynolds number $\rho V(2R)/\mu$ with $\mu = 1.81 \times 10^{-5}$ Pa s and base drag coefficient 0.25, using the exact source-bound constants in `aerodynamics.rs`. Spin decay uses $\lambda = 0.05 \text{ s}^{-1}$.

F Aiming and Outcome Statistics

Raw outcomes keep the requested target-relative angles. A second comparison rotates each entire pattern by a single angle computed from its nominal, zero-error landing, placing that nominal landing on the target line. This is an aiming adjustment, not per-shot correction or a sample-mean fit. The fixed target distance is the straight nominal carry; differences in curved carry remain in the error metric.

$$\alpha = \text{atan2}(y_0, x_0), \quad \begin{pmatrix} x' \\ y' \end{pmatrix} = \begin{pmatrix} \cos \alpha & \sin \alpha \\ -\sin \alpha & \cos \alpha \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix}.$$

Lateral variance is the unbiased sample variance with denominator $n - 1$; standard deviation is its square root. Central 90% width is the empirical 95th minus 5th percentile. Target RMSE is

$$\text{RMSE} = \sqrt{n^{-1} \sum_i [(x_i - x_T)^2 + y_i^2]}.$$

Target-hit probability counts landings within a 15 m radius circle, not a fairway strip. Paired percentile-bootstrap variance-ratio intervals resample common shot indices 500 times. These intervals describe Monte Carlo error only; they exclude model discrepancy and parameter uncertainty.

G Distance Control

Carry is not fixed: face error changes impact and spin, and curved patterns have shorter mean carry. The supplementary common-range diagnostic maps each nominal-aimed endpoint to

$$(x'', y'') = \frac{x_T}{\sqrt{x'^2 + y'^2}}(x', y').$$

It preserves landing bearing and sets every radial range to the same straight nominal distance. It is geometric normalization, not a new flight simulation, and cannot isolate a causal physical mechanism. The superseded historical controls suggested a small residual effect after normalization. Their values do not establish the corrected-model result; all final shape comparisons require the corrected experiment matrix.

H Club-Appropriate Scoring Estimates

The driver scenario is a tee shot on a hypothetical flat 400 m hole with a fairway corridor of total width 30 m. Endpoints inside that corridor use the fairway benchmark; outside use rough. Remaining distance is $d_i = \sqrt{(400 - x_i)^2 + y_i^2}$. Hole distances 350 and 450 m and corridor widths 20 and 40 m provide layout sensitivities. The tee starting benchmark must come from the published tee table, not a fairway approach substitute.

For the 7-iron and pitching wedge, the scenario is an approach from fairway at the straight nominal carry, with the hole at that target and a circular green of radius 15 m. Inside the green the benchmark uses putting distance; outside it uses unobstructed rough distance. Green radii 10 and 20 m provide layout sensitivities. There are no bunkers, water, out-of-bounds, slopes, bounce, or roll. Carry endpoints are assumed final resting locations in both contexts; actual driver total distance and wedge stopping behavior are not modeled. The original approximately 195 m driver-to-green experiment is only a superseded historical control.

For historical-tour expected strokes $J(d, \ell)$ at distance d and lie ℓ , the existing source-backed Tools analysis computes

$$SG_i = J(d_{\text{start}}, \ell_{\text{start}}) - 1 - J(d_i, \ell_i).$$

All patterns share the starting state, so their relative benefit depends only on the difference in expected strokes remaining. Reported benefits are per shot, not per round. Multiplying by an assumed number of identical approaches is an illustration, not a full-course simulation. The green benchmark represents an *unholed* putting state. With valid one-, two-, and three-putt probabilities,

$$J_{\text{putt}} = p_1 + 2p_2 + 3p_3 = 2 - p_1 + p_3 \geq 1, \quad \lim_{d \downarrow 0} J_{\text{putt}}(d) = 1.$$

A separately explicit already-holed terminal state has zero strokes remaining. No endpoint in this carry-only study is assumed holed; hole capture is not modeled. The original baseline incorrectly placed terminal zero in the same putting interpolation table. Between zero and its first positive knot at 5.3975 cm, it awarded partial hole-outs. At 1 cm this produced approximately 0.815 artificial extra SG. In the fixed-loft one-degree controls it inflated straight means by approximately 0.0053 strokes for 7-iron and 0.0103

for PW. Direct API/cache parity reproduced this error and therefore did not establish physical validity. The corrected baseline uses the unholed tap-in anchor of one, a new version/digest, and rescoring of all saved endpoints. Pre-correction scoring artifacts are retained as failed experiments.

The historical fairway and rough values come from Broadie’s published benchmark table. The putting approximation and its source reconciliation are recorded in the scoring provenance. In the 2011 preprint, literal evaluation of the displayed three-putt formula with the listed coefficient order gives negative probabilities and conflicts with the stated anchors; that literal experiment was rejected. A reconciled approximation must not be represented as an independently verified exact reproduction of the author’s fit. The implemented approximate three-putt expression uses putt distance d_f in feet:

$$\tilde{p}_3 = \frac{1}{1 + \exp(5.49 - 0.106d_f + 0.000563d_f^2)} - 0.00398, \quad p_3 = \min\{\tilde{p}_3, 1 - p_1\}.$$

The upper bound prevents a negative two-putt probability for tap-ins; the published one-putt expression uses yards. The baseline supports only its explicit distance range and rejects unsupported queries rather than extrapolating. Sampling intervals exclude putting-model uncertainty, green-layout error, and impact-model discrepancy.

For the corrected matrix, the public Tools API evaluates every explicit benchmark distance knot for each fixed starting state and finishing lie. Endpoint SG is reconstructed by bounded piecewise-linear interpolation of those returned SG values. This reuses the public API’s outputs without copying its expected-strokes formula. Knots, midpoints, and random supported queries are checked against direct API evaluation to 10^{-12} strokes. Receipts distinguish shots scored from API-evaluated states. Unsupported states are reported unavailable; neither starts nor endpoints are extrapolated.

I Corrected Monte Carlo Results

All 24 cells and 720,000 flights completed. Baseline version 2 re-scores identical saved landings with the unholed tap-in correction. The adjacent README provides all 24 rows and the overview retains paired 95% sampling intervals. Representative shaft-rotation, doubled-curve, two-degree-SD results follow; columns are straight/draw/fade lateral SD in metres and draw/fade relative SG per shot.

Club	Lateral SD S/D/F	Draw ΔSG	Fade ΔSG
Driver	16.29/15.63/15.98	-0.0046	-0.0053
PW	3.94/4.02/3.80	+0.0037	-0.0061
7-Iron	8.38/8.51/7.99	-0.0084	+0.0096

Carry is not constant. A smaller lateral SD does not necessarily improve tee scoring when carry is lost. Equal-range normalization preserves bearing and diagnoses distance scaling, not a causal mechanism. These conditional predictions do not establish a universal curvature benefit or an empirically measured player effect.

J Interpretation and Observability Limits

The relevant question is whether curved patterns reduce spread or target error after equal aiming treatment, not whether their nominal balls curve back toward the line. Fixed-loft centered still-air delivery is mirror symmetric. Keeping the same right-handed shaft axis in coupled mode does not preserve that

reflection: $L_d(-\delta) \neq L_d(\delta)$ in general. Thus conditional draw/fade differences can arise structurally from assumed face-loft covariance and aiming as well as finite sampling; they do not establish a player-specific preference. A nonzero face-to-path mean also does not eliminate opposite curvature: a 1.5-standard-deviation offset leaves about 6.68% probability of crossing zero face-to-path under the assumed normal error distribution. At two degrees SD, the theoretical crossing probability rises to about 22.66%. Real golfers may have correlated face/path errors, different variability by intention, strike-dependent gear effect, wind, varying speed or loft, and asymmetric hazards. Those require measured inputs and additional experiments. With only one uncertain input, the landing set is a narrow one-dimensional curve in the horizontal plane, not a measured two-dimensional dispersion ellipse. No ellipse-area advantage is inferred from this experiment. For intuition, a local linear landing model $Y \simeq aF + b(F - P)$ with fixed path gives $\text{Var}(Y) \simeq (a + b)^2 \sigma_F^2$, independent of the mean face and path if the sensitivities are constant. Any simulated spread difference therefore arises from changing nonlinear sensitivities, not the mere existence of a nonzero curvature mean. This linearization explains the comparison; it is not used to generate any trajectory or landing.

K Evidence and Reproduction

The adjacent README and checked-in result bundle record numerical results, validation commands, dependency/source hashes and graphic provenance. The graphics show a limited set of representative flights while the dispersion and statistics use all 10,000 shots per pattern. Full fresh reproduction uses a new output directory so preserved historical controls are never overwritten. Existing corrected bundles may instead be validated and summarized by the overview command.

```
python3 -m src.tools.shot_pattern_analysis.matrix \
    shot_pattern_reproduction/corrected_results
python3 -m src.tools.shot_pattern_analysis.overview \
    shot_pattern_reproduction/corrected_results \
    --output-dir shot_pattern_reproduction/overview
python3 -m src.tools.shot_pattern_analysis.sensitivity \
    --output shot_pattern_reproduction/sensitivity.json
python3 -m src.tools.shot_pattern_analysis.numerics \
    --output shot_pattern_reproduction/numerics.json
```

The native Rust dependency must be installed from this checkout. Compilation of this reference is a document check, not scientific acceptance.

L External Context

Trackman defines face-to-path as the horizontal difference between face and path and relates it to curvature; see <https://www.trackman.com/blog/face-to-path> and <https://www.trackman.com/blog/spin-axis>. Historical strokes-gained methodology and benchmarks are described in Broadie, *Assessing Golfer Performance on the PGA TOUR*, https://business.columbia.edu/sites/default/files-efs/pubfiles/4996/assessing_golfer_performance.full.pdf. These definitions support the input interpretation, not validation of this repository's numerical predictions.