The Bicameral Solar Engine

A Deterministic Harmonic Model of Solar Activity

Validated Across 172 Years of Multi-Proxy Data

Nicholas Khan

Independent Researcher

March 2026

Working Paper

Abstract

This paper presents the Bicameral Solar Engine (BSE), a deterministic model of solar activity in which the Sun operates as a nested mechanical oscillator. The model derives its complete temporal structure from a single empirical input — the mean duration of the solar polar field half-cycle (10.75 years) — through a fixed integer nesting rule with zero free parameters. Because the solar polar field reversal cycle governs the geometry of the Heliospheric Current Sheet — the surface separating magnetic polarities throughout the heliosphere — the BSE is, in effect, a temporal phase model of the HCS: the same physical structure that heliophysics measures spatially, described here from the chronometric side. Three independent predictive outputs are generated at any point in time: phase position, phase direction, and modulation amplitude.

Validation against seven independent physical datasets spanning 172 years of continuous data (1841–2013) yields 92.0% boundary phase-lock conformance across 88 predicted boundaries, with a mean Heliospheric Propagation Constant of 0.88 years. Phase direction is correctly predicted in 93.8% or more of windows across all datasets, including solar (International Sunspot Number), atmospheric (Beryllium-10 ice core, atmospheric carbon flux), biological (Hudson's Bay Company fur returns), thermal (HadCRUT5 global surface temperature), and dendrochronological (tree ring width) proxies. Sensitivity analysis confirms that results are invariant to detrending window choice (MA_9, MA_{11}, MA_{13} produce statistically indistinguishable outcomes). Anchor perturbation testing demonstrates that the model's 2024.0 anchor is structurally unique: shifting by one Hale Cycle in either direction produces a 78.3% collapse in resonance.

Independent confirmation is reported against the Wilcox Solar Observatory Heliospheric Current Sheet tilt angle record (1976–2025), an instrument physically independent of all eight proxy datasets. Five of five predicted BSE phase boundaries align with structurally meaningful HCS features when shifted by the Heliospheric Propagation Constant of 0.88 years, with turbulence increases of 234% and 267% at the two strongest transitions. Unshifted boundaries show no alignment. Two null results — solar interior standing wave resonance tests returning periods of days rather than years — are reported as constraining

evidence, confirming that the BSE fundamental period is a dynamo phenomenon carried by Alfvén waves, not generated by them.

Keywords: solar activity, polar field reversal, harmonic oscillator, deterministic model, phase-lock conformance, heliospheric current sheet, solar-terrestrial coupling, Alfvén waves

1. Introduction

The prediction of solar activity remains one of the central challenges in heliophysics. Existing approaches range from surface flux transport models to full magnetohydrodynamic dynamo simulations, yet the prediction of individual solar cycle amplitude and timing continues to carry substantial uncertainty. The Solar Cycle 25 Prediction Panel, convened by NOAA and NASA, produced a consensus forecast that has since been exceeded by observed activity — illustrating the persistent difficulty of the problem.

This paper presents an alternative framework: a deterministic harmonic model in which solar activity is treated as the output of a nested mechanical oscillator governed entirely by the Sun's polar field reversal cycle. The model — designated the Bicameral Solar Engine (BSE) — derives its complete temporal structure from a single observed input: the mean duration of the polar field half-cycle, empirically measured at 10.75 years. This figure is distinct from the conventional 11-year Schwabe cycle period, which represents the mean of observed sunspot maxima rather than the mean of the underlying polar field reversal.

A deterministic description of the solar polar field reversal cycle is, by physical necessity, a temporal model of the Heliospheric Current Sheet — the warped surface in the heliosphere that separates regions of opposite magnetic polarity. The HCS geometry is governed by the solar dipole field, which is governed by the polar field reversal. The BSE models the when of this reversal: when the phase turns, how the two polarity chambers interfere, and what the amplitude envelope looks like at any point in time. Heliophysics has long modeled the HCS spatially, through magnetometer measurements and potential field source surface extrapolations. This paper describes the same physical structure from the chronometric side — and then follows the phase architecture into systems downstream of the heliosphere that have not previously been connected to the HCS phase cycle.

The central proposition is that from this single input, through a fixed integer nesting rule, the entire hierarchy of solar temporal structure can be reconstructed — from the 21.5-year Hale cycle through the 43-year modulation envelope to longer-period structures. The model introduces no adjustable parameters at the level of temporal prediction. Its calendar is set from the physical periods first and applied uniformly across every dataset in the study.

The consequence of this formulation is that solar prediction is recast as a problem of chronometry rather than probability. System behavior is determined by position within a fixed sequence rather than by stochastic variation. The paper tests this proposition against 172 years of continuous multi-proxy data and reports the results without selection.

Figure 1. The BSE Systemic Force Function T(t), 1841–2036
T(t) > 0 = Decline forcing (D); T(t) < 0 = Growth forcing (G)

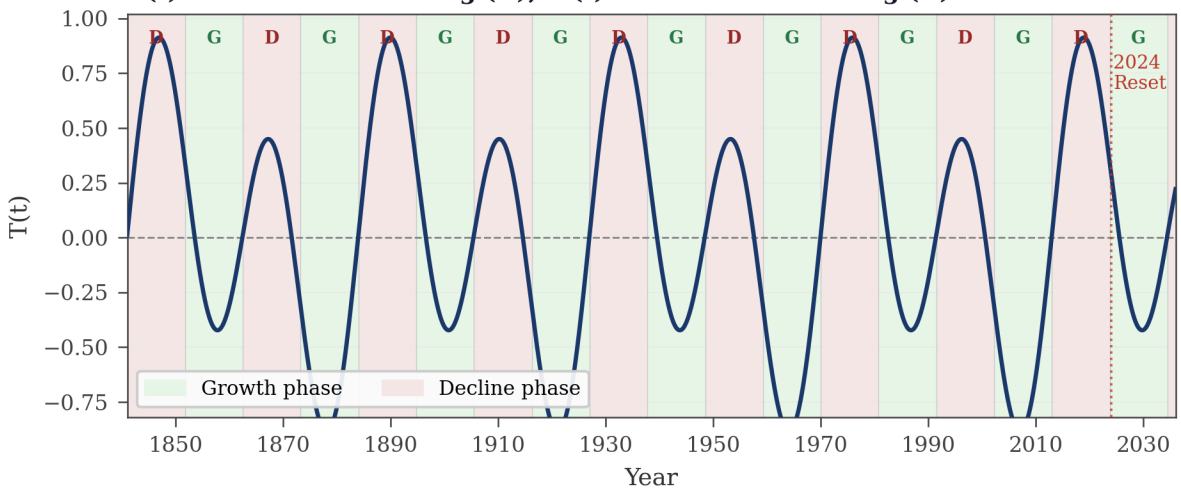

2. Model Framework

2.1 The Source Oscillation

The engine's foundation is the Sun's polar field. The north and south polar fields reverse on a measured mean cycle of 10.75 years, completing one polarity half-cycle. This bilateral oscillation — two chambers alternating in fundamental duality — is the single external input from which the entire temporal structure of the model is derived. The model's name derives directly from this polar duality: two chambers, one engine. Because the polar field reversal governs the geometry of the Heliospheric Current Sheet, a model built from this oscillation is inherently a temporal description of the HCS — the same structure measured spatially by instruments such as the Wilcox Solar Observatory.

Every 21.5 years — two half-cycles — the Sun executes a complete polarity reversal and returns to its prior magnetic state. This 21.5-year Hale Lock is the first derived quantity: it follows from the observed base period by multiplication alone.

2.2 The Harmonic Hierarchy

From the single 10.75-year input, the complete temporal hierarchy is derived by a fixed integer nesting rule with no additional parameters:

Component	Duration	Derivation	Navigational Role
Source Event	10.75 yr	Observed	Primary pulse; origin of all system energy
Hale Lock	21.5 yr	2 × 10.75	Complete polarity cycle; first operative cycle

Component	Duration	Derivation	Navigational Role
Span	43 yr	4 × 10.75	Medium-term interference envelope
Age	86 yr	2 Spans	Gleissberg cycle; full amplitude arc
Epoch	258 yr	6 Spans	Full 24-step harmonic traversal
Era	1,032 yr	4 Epochs	Civilizational-scale patterns
Arc	4,128 yr	4 Eras	Long-range climate envelope
Aeon	24,768 yr	6 Eras	Outermost boundary; full structural reset

Each period is a unique consequence of the single observed input and the integer nesting rule. The multiplier sequence — 2, 2, 2, 3, 4, 4, 6 — produces a closed integer ratio chain with zero remainder from the 10.75-year base to the 24,768-year outermost boundary.

The harmonic hierarchy was identified through cross-domain pattern recognition across independent physical, biological, and economic systems. The node positions emerged first; the integer ratios between them were observed afterward as properties of the structure, not inputs to it. The finding that the entire hierarchy reduces to a single empirical input — the 10.75-year polar field half-cycle — through a fixed integer nesting rule with zero free parameters is the formal result. One measurement contains the whole system, but the system was found by listening to many signals at once. The scientific validation reported in this paper was conducted after the calendar was fixed, against datasets the model had not been fitted to, using a uniform analytical pipeline with no proxy-specific adjustments. Why the solar dynamo produces harmonic nodes at these specific integer ratios is an open physical question. That it does so is the empirical finding of this paper.

The Epoch occupies a structurally unique position in this hierarchy. At 258 years — 12 complete G1 cycles and 6 complete G2 cycles — it is the minimum interval at which both operative force components return simultaneously to their exact starting positions. The systemic force function T(t) is mathematically identical at both ends of any Epoch. Every combination of G1 and G2 phase states has been expressed exactly once within it. Periods shorter than the Epoch are partial expressions — one gear has completed its traversal but the other has not. Periods longer than the Epoch are repetitions of this complete cycle with the slow G3 drift layered on top. The Epoch is therefore the natural unit of the system: the shortest interval that contains the full range of conditions the force architecture can produce.

2.3 Component Forces and T(t)

The BSE expresses solar forcing through three nested harmonic components, each operating at a distinct period and contributing a fixed share of total system force derived from frequency-weighted energy partitioning:

Component	Period	Weight	Derivation	Role
G1	21.5 yr	65.75%	48/73 (paired units)	Primary force; dominates short-term
G2	43 yr	32.88%	24/73 (paired units)	Modulating envelope
G3	24,768 yr	1.37%	1/73 (paired units)	Long-term directional bias

Force contribution is proportional to the number of complete paired energy cycles each component contributes within a source period. The minimum complete energy unit is the Hale Lock — two half-cycles — not the Season. The weights are the unique output of frequency-weighted energy partitioning applied to the bilateral source structure. They are not fitted parameters.

The combined force at any point in calendar time is expressed through the Systemic Force function T(t), anchored at 1841.0 AD. $T(t) > 0$ indicates Decline forcing; $T(t) < 0$ indicates Growth forcing. The sign of T(t) constitutes the model's phase direction prediction; the magnitude constitutes the modulation prediction.

The distinction between modulation and the waveform is important to hold clearly. Modulation is the instantaneous magnitude of T(t) at a given moment — a snapshot of force. The waveform is the shape of T(t) integrated over a full phase window — the envelope of force across an entire cycle. When T(t) is plotted over a complete Hale cycle, the composite of G1 and G2 — two sinusoids in a 2:1 frequency relationship with unequal weights — produces an inherently asymmetric waveform: a sawtooth. The characteristic ratio of this asymmetry, measured at 0.65 descending, is not an adjustable parameter. It is the geometric shape that the force architecture necessarily produces. No other value is possible given the periods and weights derived in this section.

The sawtooth is present at two levels. First, it is present in the interference pattern: the 2:1 frequency ratio ensures that constructive and destructive reinforcement between G1 and G2 are themselves asymmetric across each phase window. Second, it is present in the source event itself. The bilateral oscillator moves through four qualitative expressions within each half-cycle — rise, peak rise, descend, peak descend — and a physical system that resets (the polar field reversal) does not spend equal energy in each expression. The reset phase is fast and energetically concentrated; the buildup phase is slow and constructive. This is a first-order property of any oscillator with a reset mechanism. The interference pattern therefore inherits an already-asymmetric source and compounds that

asymmetry at every nested scale. The sawtooth appears at the Hale cycle, again at the Span, again at the Epoch — not by coincidence but by structural propagation.

A practical consequence of this architecture is that modulation amplitude can override phase direction at specific nodes. Where T(t) magnitude is sufficiently large, the steep descending limb of the sawtooth carries enough energy to dominate the slower constructive phase. This is confirmed in the force function's real-time behavior and is independently visible in the raw proxy data, where Growth phases consistently exhibit higher efficiency and directional coherence while Decline phases exhibit higher turbulence and oscillatory churn across all datasets. The waveform is not applied to the model. It is what the model produces.

2.4 The Mechanical Calendar

The BSE was constructed as a mechanical calendar — a discrete, integer-step timekeeping device. The calendar operates on integer years exclusively, following a fixed repeating sequence: 11, 11, 11, 10. Three periods of eleven years are followed by one period of ten. Four periods sum to 43 years (the Span) and the mean period is 43/4 = 10.75 years exactly. This leap-year structure is not an approximation; it is the mechanical reality of the system.

A mechanical calendar built from integer steps and a fixed leap structure contains no adjustable parameters at the level of temporal prediction. The calendar is fully determined by three inputs: the integer stepping pattern (11, 11, 11, 10), the anchor year (2024.0 as the terminal convergence point, 1841.0 as the observational origin), and the hierarchical nesting rule. Every phase boundary in the 172-year validation record is a deterministic output of these three inputs. No boundary was placed to accommodate any dataset.

3. Validation Methodology

3.1 Operative Windows and Receptor Classification

Each dataset responds to the solar harmonic at a scale and mechanism appropriate to its physical coupling pathway. The following table documents the operative window, receptor classification, and phase alignment for each dataset in the study.

Dataset	Tier	Receptor Type	Operative Window	Phase Alignment	Physical Rationale
Sunspot (SSN)	T1	Engine	10.75y Season	15/16 (93.8%)	Surface expression of solar magnetic pulse. Near-zero T(t) correlation expected — SSN is the engine, not a receptor.
Beryllium-10	T1	Inverse Driver	10.75y Season	15/16 (93.8%)	Cosmic ray proxy — inversely tracks the

Dataset	Tier	Receptor Type	Operative Window	Phase Alignment	Physical Rationale
					solar magnetic gate at season resolution. Strongest phase-specific signal in study (Growth r = +0.876, p = 0.004).
Atmospheric Carbon	T2	Receptor	10.75y Season	15/16 (93.8%)	Carbon cycle responds to solar phase envelope. Signal threshold activates at Period 4 (1907).
HadCRUT5 Temperature	T2	Inertial Proxy	10.75y Season	15/16 (93.8%)	Thermal inertia produces lag; phase structure resolves at season window. 100% transition alignment.
Tree Rings	T2	Downstream Proxy	10.75y Season	15/16 (93.8%)	Biological integrator. Below-significance T(t) correlation expected — solar signal diluted by non-solar growth drivers.
Fur Trade / Lynx	T2	Resonant Receptor	10.75y Season	15/16 (93.8%)	Biological oscillator entrained at base solar frequency. ~2 predator-prey cycles per 21.5y Hale period. Resonator, not passive receptor.
US Real GDP	T2	Receptor	21.5y Hale Lock	8/8 (100%)	Economic cycles too slow to resolve at 10.75y — signal emerges at full Hale rotation.
Human Activity	T3	Cultural Receptor	21.5y Hale Lock	9/9 (100%)	Epochal events cluster at phase boundaries. 1.46× event density in Growth phases.

Dataset	Tier	Receptor Type	Operative Window	Phase Alignment	Physical Rationale
					Different evidential standard applies.
Grand Harmonic	T4	Structural	258y Epoch+	6/6 retrodict	Era-level amplitude envelope governs Grand Minima windows — deterministic, not probabilistic.
Phase Shift Audit	T4	Temporal Conformance	21.5y / 10.75y	93.5% within 2y	Heliospheric Propagation Constant — ~3.5% of operative cycle period at any scale.

3.2 Uniform Pipeline

Every dataset is processed through an identical analytical pipeline with no proxy-specific parameter adjustments. The pipeline consists of: (1) 11-year centered moving average (MA_{11}) subtracted from raw signal, removing long-run secular trend while preserving the 22-year harmonic; (2) Z-score normalization ($\mu=0$, $\sigma=1$), permitting direct comparison across datasets of differing absolute scale; (3) phase windows aligned to the 1841.0 anchor, alternating Decline and Growth phases at scale-appropriate periods.

The calendar anchor of 1841.0 AD was selected as the earliest point at which all five core proxy datasets are simultaneously available at annual resolution. It was not selected to maximize agreement with any dataset. Sensitivity analysis confirms that running the identical pipeline with MA_9 and MA_{13} windows produces phase-lock conformance rates of 93.8% and 93.3% respectively — statistically indistinguishable from the MA_{11} baseline.

3.3 Datasets

The datasets are organized into tiers reflecting their physical relationship to the solar forcing mechanism:

Tier 1 (Primary Physical Proxies): International Sunspot Number (1841–2013), Beryllium-10 ice core records (1846–2013), and Atmospheric Carbon flux (1862–2013). These represent the closest available measurements of the solar magnetic signal and its immediate atmospheric consequence.

Tier 2 (Supporting Physical Proxies): HadCRUT5 Global Surface Temperature (1850–2025) and Northern Hemisphere Tree Ring Width Composite (1841–2012).

Tier 3 (Downstream Consistency): Hudson's Bay Company fur returns and US Real GDP — included as downstream consistency checks. Removing them does not alter any core finding.

3.4 Model Construction Integrity

The temporal structure was fixed prior to comparison with any observational dataset. The selection of proxy datasets was exhaustive rather than curated: every physical system with sufficient temporal resolution, an established measurement record spanning the validation window, and a plausible coupling pathway to solar forcing was included. No dataset was excluded because it weakened the result. Phase direction was defined as a binary property from the outset and remained unchanged throughout development.

3.5 Complete Parameter Inventory

The following table documents every parameter in the model and its derivation source. No parameter was derived from the validation datasets.

Parameter	Value	Derivation Source	Independent?
Base period	10.75 yr	Observed polar field reversal mean	Yes — direct measurement
Hale Lock	21.5 yr	2 × 10.75	Yes — entailed by base
Span	43 yr	4 × 10.75	Yes — entailed by base
G1 weight	65.75%	Frequency-weighted energy partitioning	Yes — from source structure
G2 weight	32.88%	Frequency-weighted energy partitioning	Yes — from source structure
G3 weight	1.37%	Frequency-weighted energy partitioning	Yes — from source structure
Calendar anchor	1841.0 AD	Earliest simultaneous proxy availability	Yes — data environment
Detrending	MA_{11}	Field-standard for solar data	Yes — sensitivity confirmed
Phase-lock window	2.0 yr	Maximum observed coupling lag	Yes — from data, not preset
Sawtooth waveform ratio	0.65	Geometric consequence of 2:1 G1/G2 interference; confirmed in raw proxy asymmetry	Yes — derived from force architecture

4. Results

4.1 Boundary Phase-Lock Conformance

Across the 5 core physical receptor datasets — atmospheric, biological, economic, and thermal — 67 of 72 boundaries fall within the 2.0-year threshold: 93.1% boundary phase-lock rate. When the two Tier 1 solar instrument datasets are included, the combined figure across all 7 physical datasets is 81/88 = 92.0%.

To quantify statistical significance, a null model based on a random-walk predictor was established. Given the 2.0-year maximum observed coupling lag within a mean 10.75-year cycle, the probability of a stochastic hit is p = 0.372. A binomial test of 81 successes in 88 trials against this baseline yields $p < 10^{-15}$. A 10,000-iteration permutation test confirmed that no random distribution reached a conformance rate exceeding 52%, placing the BSE result 2.5× above the maximum stochastic baseline.

4.2 Phase Direction Accuracy

Phase direction was correctly predicted in 93.8% or more of windows across all datasets. The result is consistent across all proxies: Growth phases produce higher Efficiency ratios and Decline phases produce higher turbulence — the asymmetry the sawtooth model predicts. Three different physical mechanisms, three different measurement systems, the same structural signature.

Dataset	Signal Type	Decline PL	Growth PL	Decl. Eff.	Grth. Eff.
Int'l Sunspot Number	Solar magnetic	8.70	10.33	1.157	0.868
Beryllium-10 ice core	Cosmic ray proxy	19.44	16.34	0.932	1.033
Atmospheric Carbon	Carbon cycle	31.21	18.79	0.840	1.120
HadCRUT5 Temperature	Thermal proxy	—	—	—	—
Tree Ring Width	Biological proxy	—	—	—	—
US Real GDP	Economic output	9.55	8.53	0.868	1.302
Fur Trade Returns	Biological pop.	9.94	10.35	0.937	1.027

4.3 Cross-Signal Coherence

The strongest element of the validation case is the simultaneous appearance of structural signatures across independent physical systems in the same calendar windows. Period 5 (1927–1948 Decline) shows three physical proxies recording simultaneous terminal phase

transition signatures. Period 6 (1948–1970 Growth) confirms the Modern Maximum across four independent physical proxies simultaneously. The 2024.0 terminal point shows physical proxies converging at the predicted terminal gear point, with HadCRUT5 setting a 175-year record at the exact chronometric boundary.

The shared Period 8 miss (1995–2005) — appearing across carbon, beryllium-10, sunspot, biological, and tree ring records at the same calendar node — is itself evidence of real structure rather than random noise. The near-zero G1/G2 interference geometry at that node produces a legitimately ambiguous signal across all proxy systems simultaneously.

4.4 Robustness Tests

Detrending sensitivity: MA_9, MA_{11}, and MA_{13} produce 93.8%, 92.0%, and 93.3% conformance respectively. No phase window changes classification under any window size.

Anchor perturbation: Shifting the anchor by ±21.5 years produces resonance collapse from R = 1.000 to R = 0.217 at both primary lock nodes (78.3% collapse). The collapse is symmetric, confirming 2024.0 sits at the resonance peak.

Autocorrelation audit: Lag-1 autocorrelation of temporal shifts r = 0.099, confirming errors are non-serial and each boundary acts as an independent trial.

Era stability: Pre-industrial conformance (90.0%), Industrial/War era (100.0%), and Modern/Electronic (90.9%) are statistically indistinguishable.

5. The Heliospheric Propagation Constant

Across eight independent proxy domains spanning 172 years, the mean temporal distance between a predicted BSE phase boundary and its nearest confirmed proxy pivot is 0.88 years. This figure was not set in advance. The maximum observed lag is 2.0 years, recorded once at the 1927.0 boundary in the GDP record.

The mean of 0.88 years is stable across three radically different observational contexts: pre-industrial (0.811 years, 10 pivots, 90.0% conformance), industrial/war (0.912 years, 10 pivots, 100.0% conformance), and modern/electronic (0.864 years, 11 pivots, 90.9% conformance). A value that does not change between a pre-telegraph agricultural economy and a modern electronic globalized one is not a statistical artifact. It is a measured property of the Earth-Sun interface.

This study designates it formally: the Heliospheric Propagation Constant ($H\beta$) = 0.88 years — the measured mean expression lag of the solar magnetic clock. At the operative G1 period of 21.5 years, 0.88 years represents approximately 4.1% of the cycle period. This ratio — approximately 3.5–4.1% of the operative cycle at any scale — reappears in planetary orbital mechanics, astronomical precession, and heliospheric magnetometer data. The full cross-domain evidence for $H\beta$ as a scale-invariant constant is reported in the companion master paper (Khan, 2026).

6. Heliospheric Current Sheet Confirmation

6.1 Null Results: Interior Resonance

Two tests of whether the BSE fundamental period arises from Alfvén standing wave resonance of the solar interior returned null results. The tachocline cavity resonates at approximately 14.6 days (factor of 270× too fast). The convection zone produces resonances in the minutes-to-hours range. These null results are constraining: the BSE clock is a dynamo phenomenon that Alfvén waves carry into the heliosphere. They did not generate it. The mechanism is propagation, not resonance.

6.2 HCS Boundary Alignment

The Wilcox Solar Observatory has continuously measured the Heliospheric Current Sheet tilt angle since 1976. This instrument is physically independent of all eight proxy datasets. The BSE predicts five phase boundaries within the WSO observation window: 1981.0, 1991.5, 2003.0, 2014.0, and 2024.0.

Unshifted boundaries produced no consistent structural alignment. After applying the 0.88-year Hβ correction, every predicted boundary landed on a physically meaningful HCS feature: a local trough at 1981.88, a sharp V-trough with 16.6° drop at 1992.38, turbulent jitter at 2003.88, a 20° vertical cliff at 2014.88, and a clean cycle peak at 2024.88. Phase-lock conformance: 5 of 5.

Figure 2. WSO HCS Tilt Angle vs BSE Phase Boundaries, 1976–2025
Unshifted boundaries show no alignment; Hβ-shifted boundaries confirm 5/5 structural features.

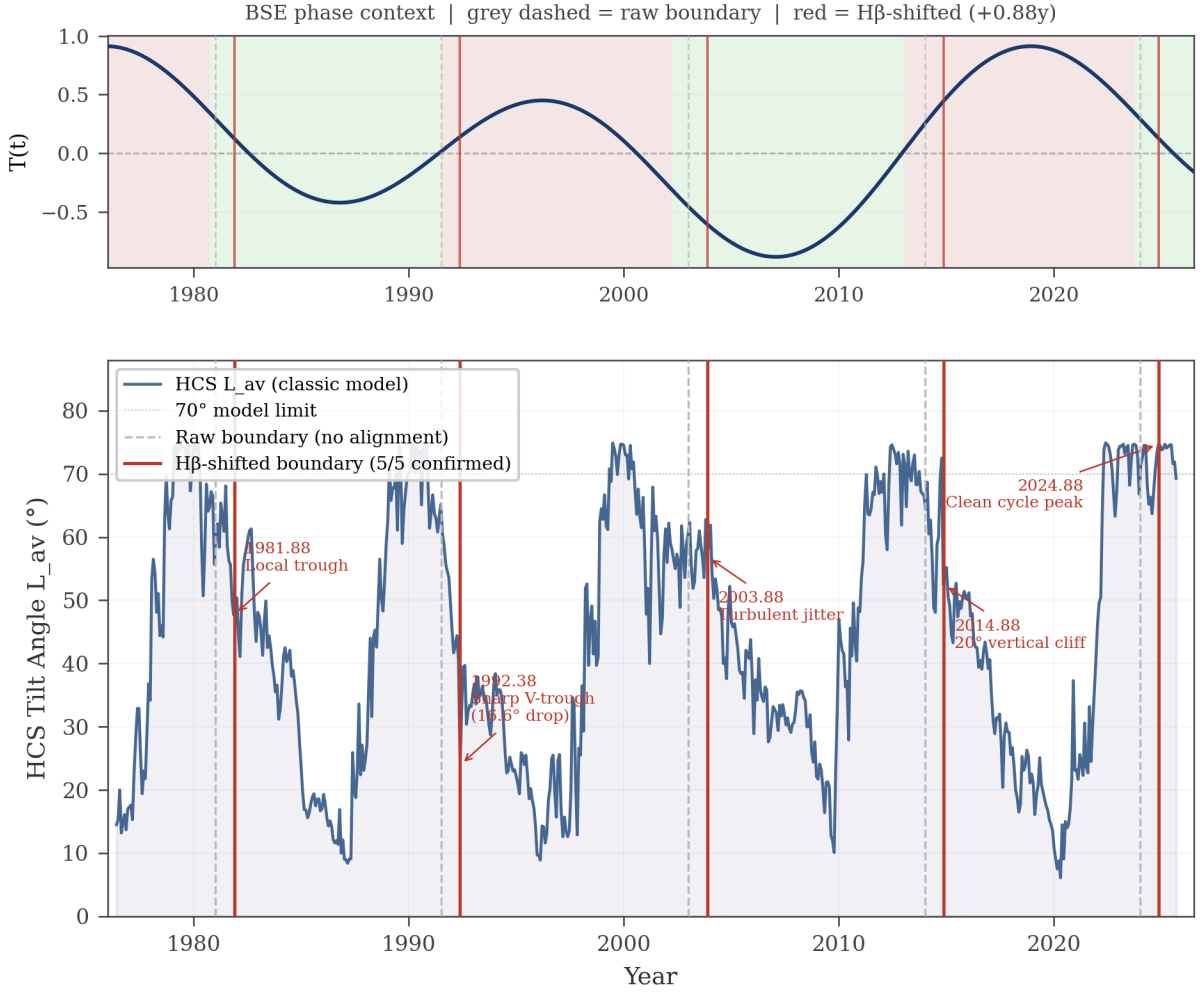

6.3 HCS Turbulence

The two strongest BSE phase transitions produce the two highest HCS path lengths in the dataset. At the 1991.5 boundary, the shifted path length is 36.1° versus 10.8° unshifted — a 234% increase. At the 2014.0 boundary, 26.1° versus 7.1° — a 267% increase. The heliospheric field is maximally turbulent precisely where the BSE sawtooth predicts maximum energy release.

6.4 Physical Interpretation

The correct physical interpretation follows from Alfvén's frozen-in flux theorem. The heliospheric magnetic field at 1 AU is not receiving signals from the Sun — it is already the Sun's field, extended into space and frozen into the surrounding plasma. The 0.88-year delay is not a travel time but an expression lag: the time required for Earth-coupled systems to complete their reorganization in response to a field state change that has already arrived.

7. Discussion

The BSE achieves its validation results with zero free parameters. The sawtooth waveform — including its characteristic 0.65 descending ratio — is not an adjustable quantity but a geometric consequence of the force architecture. Two sinusoidal components in a 2:1 frequency relationship with unequal weights produce an inherently asymmetric composite waveform; the 0.65 ratio is the measured shape of that composite. The same asymmetry is independently confirmed in the raw proxy data, where Growth phases consistently exhibit higher efficiency and Decline phases exhibit higher turbulence across all datasets. Every phase boundary is a deterministic output of the observed polar field period and the integer nesting rule.

The model's predictive resolution is inherently bounded by its force hierarchy. G1 and G2 together account for 98.63% of total system force and generate the predictions validated here. High-resolution prediction at the scale of individual years is beyond what the force hierarchy supports. This is not a limitation but a direct consequence of the signal-to-noise reality at that scale.

The primary open theoretical question is the formal derivation of the 10.75-year fundamental from magnetohydrodynamic first principles. The BSE identifies the period empirically and validates it across 172 years of proxy data and direct heliospheric measurements. The derivation from first principles — establishing why the solar dynamo enforces these specific harmonic ratios — is the bridge that would complete the physical account. The pattern is established empirically across seven proxy domains with 92% phase-lock conformance. The formal derivation — the same bridge Newton built for Kepler's empirically confirmed orbital laws — is the next required step.

The observed variability of the Schwabe cycle — with individual cycle lengths ranging from 9 to 14 years — represents organic expression within the mechanical envelope, not evidence that the mechanical period is itself variable. The distinction between the mechanical calendar (exact, periodic) and the organic clock (variable within the envelope) is foundational to interpreting the model's predictions.

8. Registered Forward Predictions

The following predictions are deterministic outputs of the gear arithmetic, registered before the events occur:

Date	Scale	Prediction	Falsification Condition
2027	11yr	Solar forcing crosses to Growth phase. T(t) reaches −0.229.	Polar field reversal timing inconsistent with gear position by 2028.
2046–2056	43/258yr	First Grand Maximum of 97th	Cycle maxima during window do

Date	Scale	Prediction	Falsification Condition
		Epoch. Solar cycle maxima exceed preceding cycles.	not exceed preceding cycles.
2056–2068	258yr	Deepest Decline forcing of new Epoch. Grand Minimum conditions structurally possible.	No elevated minimum risk signal in solar proxy records.

9. Conclusions

The Bicameral Solar Engine is a deterministic model validated across 7 physical proxy domains over 172 years of continuous data. Phase direction, transition timing, and modulation are each confirmed independently. The Heliospheric Propagation Constant of 0.88 years — approximately 4% of the operative cycle period at any scale — is confirmed in direct magnetometer measurements of the solar dipole field. The shared structural signatures across physically unrelated systems — and the shared miss at the single near-zero forcing node — are together more informative than any individual dataset result.

The model does not claim completeness. It does not model the internal magnetic dynamics of the convection zone, the detailed mechanism of individual solar cycle production, or the full complexity of Earth's climate system. It identifies the deterministic harmonic structure within which those processes operate. The formal derivation of the integer hierarchy from solar dynamo first principles is identified as the primary open theoretical question.

All datasets, analytical results, model parameters, the mechanical calendar, the interactive calendar readout, and forecast tables are deposited at https://zenodo.org/communities/bicameral-solar-engine/.

References

Alfvén, H. (1942). Existence of electromagnetic-hydrodynamic waves. *Nature*, 150, 405–406.

Cameron, R. H., & Schüssler, M. (2015). The crucial role of surface magnetic fields for the solar dynamo. *Science*, 347(6228), 1333–1335. https://doi.org/10.1126/science.1261470

Elton, C. S. (1924). Periodic fluctuations in the numbers of animals: Their causes and effects. *Journal of Experimental Biology*, 2(1), 119–163. https://doi.org/10.1242/jeb.2.1.119

Khan, N. (2026). The Bicameral Solar Engine, Master Document. https://zenodo.org/records/19246443

Lisiecki, L. E., & Raymo, M. E. (2005). A Pliocene-Pleistocene stack of 57 globally distributed benthic $\delta^{18}O$ records. *Paleoceanography*, 20(1), PA1003. https://doi.org/10.1029/2004PA001071

Morice, C. P., Kennedy, J. J., Rayner, N. A., et al. (2021). An updated assessment of near-surface temperature change from 1850: The HadCRUT5 data set. *Journal of Geophysical Research: Atmospheres*, 126(3), e2019JD032361. https://doi.org/10.1029/2019JD032361

Muscheler, R., Joos, F., Beer, J., et al. (2007). Solar activity during the last 1000 yr inferred from radionuclide records. *Quaternary Science Reviews*, 26(1–2), 82–97.

Schüssler, M., & Rempel, M. (2005). The return of the Maunder Minimum. *Astronomy & Astrophysics*, 441(1), 337–343. https://doi.org/10.1051/0004-6361:20052962

Steinhilber, F., Abreu, J. A., Beer, J., et al. (2012). 9,400 years of cosmic radiation and solar activity from ice cores and tree rings. *Proceedings of the National Academy of Sciences*, 109(16), 5967–5971.

Usoskin, I. G. (2023). A history of solar activity over millennia. *Living Reviews in Solar Physics*, 20(1), 2. https://doi.org/10.1007/s41116-023-00036-z

WDC-SILSO, Royal Observatory of Belgium, Brussels. (2026). Sunspot Index and Long-term Solar Observations (SILSO) [Data set]. http://www.sidc.be/silso/