Social Science and Management
Print ISSN: 3007-6854
Online ISSN: 3007-6862
DOI: https://doi.org/10.61784/ssm3105

BAYESIAN DATA-MINING PIPELINE FOR CLIMATE POLICY STRINGENCY DECISION SUPPORT IN CHINA

XuJie Zhao[1], XiaoBo Shi[2], WenJun Wang[1*]
[1]*Guangzhou Institute of Energy Conversion, Chinese Academy of Sciences, Guangzhou 510640, Guangdong, China.*
[2]*Guangdong Medical University, Dongguan 523808, Guangdong, China.*
[*]*Corresponding Author: WenJun Wang*

Abstract: Decision-support systems must convert heterogeneous, sparse, and noisy policy records into auditable monitoring outputs under uncertainty. We develop a Bayesian data-mining pipeline for climate policy stringency decision support in China. The pipeline first audits six annual input signals: carbon-market prices, market breadth, policy-text intensity, administrative force, green finance, and clean-energy policy actions. It then estimates a latent Climate Policy Stringency Index (CPSI) for 2000-2024 with a Bayesian dynamic factor model and stores MCMC posterior draws as reusable computational artifacts. A posterior diagnostic engine exports machine-readable fields for level uncertainty, annual change probabilities, large-increase alerts, common-component shares, acceleration evidence, jump-model comparison, and external process alignment. In the retained artifact, CPSI increases from -2.10 in 2000 to 6.30 in 2024, with more frequent large-increase alerts after 2020. The temporal-pattern diagnostics support a post-2019 acceleration label, but not a statistically necessary one-year jump at annual frequency. The contribution is a workflow and output-schema refinement: heterogeneous policy records are converted into versioned posterior artifacts, rule-based monitoring fields, and auditable diagnostic labels, while causal effect claims remain outside the measurement boundary.
Keywords: Decision support system; Data mining; Bayesian data-mining pipeline; Posterior diagnostic engine; Machine-readable monitoring outputs; Risk management and assessment; Climate policy stringency

1 INTRODUCTION

Decision-support systems increasingly rely on heterogeneous policy data, but many policy indicators still enter management analysis as fixed point estimates. This creates a computational problem. Sparse annual observations, unequal raw units, structural zeros, partial missingness, and measurement noise must be transformed into a stable signal that analysts can audit, reuse, and interpret under uncertainty. For climate policy stringency, this problem is especially visible because market prices, market coverage, policy-text signals, administrative actions, green finance, and clean-energy policies do not share a common scale or measurement process.

The pipeline is therefore framed as a knowledge-discovery workflow for management decision science. It links data selection, input auditing, probabilistic extraction, posterior artifact storage, diagnostic translation, and output reuse. This framing avoids a simple relabelling of a policy index as a computing contribution: the object of study is the workflow that turns irregular policy signals into auditable decision-support outputs.

Composite-indicator research provides the application background, but it does not by itself define the computing contribution. Existing environmental policy stringency indicators show how heterogeneous policy instruments can be converted into annual scores [1-4]. For a decision-support system, the central question is whether these inputs can be audited, modelled with uncertainty, and translated into reusable diagnostics.

The implementation treats China's climate policy stringency as a latent state inferred from six partially overlapping annual signals: carbon-market price, market breadth, policy-text intensity, administrative force, green finance, and clean-energy policy actions. These variables are not predetermined weights or causal treatments. They enter an input layer from which the Bayesian model extracts a common policy signal and the diagnostic engine generates reusable monitoring outputs.

The methodological design follows the pipeline rather than a single-model narrative. The input layer standardizes available observations and passes missing entries through an observation mask. The extraction layer estimates a Bayesian dynamic factor model with NUTS sampling. The artifact layer stores posterior draws and metadata. The diagnostic layer converts retained posterior draws into uncertainty intervals, annual-change probabilities, common-component support measures, acceleration diagnostics, jump-model comparisons, and external process-alignment checks.

The China application is used to validate whether the output schema remains traceable and interpretable over 2000-2024. The retained diagnostics are reported as posterior artifacts: a CPSI level field, large-increase alert probabilities, an acceleration label, jump-comparison evidence, and indicator-support fields. This framing keeps the empirical section focused on output validation rather than on reconstructing policy history or attributing the path to a single event.

Indicator-support diagnostics show how the pipeline records evidence breadth. Market price, policy-text intensity, and administrative force provide tighter common-component support than some broader policy-action indicators. Lower-support channels still add coverage and sectoral context, but the diagnostic does not rank instruments by effectiveness.

The contribution is a reusable posterior-diagnostic pipeline and a machine-readable output schema. The design contains three linked layers. The input-audit layer handles structural zeros, missing observations, and indicator standardization.

© By the Author(s) 2026, under the CC BY-NC-ND license (http://creativecommons.org/licenses/by-nc-nd/4.0).

The Bayesian extraction layer stores posterior draws as computational artifacts. The diagnostic layer translates those draws into monitoring, evidence-breadth, temporal-pattern, and external-alignment outputs. The China application demonstrates the workflow; the main contribution remains the data-mining decision-support procedure rather than a descriptive policy-history exercise.

2 RELATED WORK AND COMPUTATIONAL RESEARCH GAP

Three strands of research define the problem addressed here. Decision-support systems and data-mining studies treat analytic systems as workflows that select, preprocess, extract, interpret, and reuse information, with the output organized for human judgement under uncertainty [5-8]. In this view, a fitted index is only one part of the system output; the reusable object is an auditable bundle of diagnostics.

Bayesian measurement and probabilistic machine learning define the uncertainty layer. Dynamic factor models extract common movement from noisy indicators when the state is not directly observed [9], and Bayesian versions retain posterior distributions rather than collapsing inference into a single path [10]. Probabilistic machine learning similarly treats uncertainty as part of the model output [11]. In the proposed workflow, posterior draws feed annual change probabilities, common-component diagnostics, acceleration-versus-jump checks, and external alignment tests.

Climate-policy indicator research supplies the application data environment. OECD Environmental Policy Stringency measures and climate-policy databases show how heterogeneous market-based and non-market policy records can be converted into annual scores or structured policy-action data [1-2, 12]. Recent work on China's climate policy stringency also uses policy texts and multi-source information to quantify long-run policy evolution [4]. These outputs make policy effort observable, but they are often reused downstream as fixed scores.

The remaining gap is operational. Climate-policy indicators rarely expose the full decision-support chain: input audit, missingness handling, probabilistic extraction, posterior quality control, diagnostic translation, and output reuse. CPSI is therefore treated as a latent decision signal generated by a reusable posterior-diagnostic pipeline, not as an objective policy fact or a causal treatment variable.

3 INPUT DATA AUDIT AND PRE-PROCESSING LAYER

The first layer treats input construction as an audit and preprocessing task. Appendix A reports source categories and raw-scale descriptive statistics; the main text specifies how six annual signals enter the computational workflow and what information the downstream model is allowed to use. The retained inputs are market price, market breadth, policy-text intensity, administrative force, green finance, and clean-energy policy actions. They are handled as noisy observations of a latent policy state, not as policy weights, performance rankings, or treatment variables.

The audit layer aligns all inputs to calendar years for 2000-2024. Structural zeros are retained when they encode absence of an instrument or market signal. Other unavailable observations are passed to the Bayesian measurement model through an observation mask rather than filled by narrative interpolation. Available observations are standardized by indicator, preventing raw units such as RMB per tonne, coverage shares, text-intensity scores, administrative-event scores, and loan balances from mechanically dominating the latent CPSI.

Table 1 Input Audit and Pre-Processing Schema for the Bayesian Decision-Support Pipeline

Input code	Observed signal	Raw unit or source basis	Pre-processing and model handling
X1	Carbon-market price	RMB/t trading-price records	Structural zeros retained before market operation; standardized before estimation.
X2	Market breadth	ETS coverage share (%)	Structural zeros retained where coverage is absent; missing values are passed through the observation mask.
X3	Policy-text intensity	Climate and low-carbon document coding	Text signal is standardized by available observations and anchors the posterior sign convention.
X4	Administrative force	Event-scored enforcement and administrative pressure	Event scores enter as noisy annual observations rather than causal treatment indicators.
X5	Green finance	Outstanding green-loan balance / documented reconstruction	Reconstructed only where documented; robustness is checked by excluding X5.
X6	Clean-energy policy actions	Annual clean-energy action coding	Included as a sectoral policy-action signal with potentially larger idiosyncratic variation.

4 BAYESIAN DATA-MINING DECISION-SUPPORT PIPELINE

This section specifies the computational design that gives the study its computer-engineering contribution. The workflow has six stages: input audit, Bayesian extraction, posterior-artifact construction, posterior quality control, diagnostic-output mapping, and decision-support labelling. The design keeps uncertainty inside the workflow. Posterior information from the Bayesian model is carried forward into monitoring flags, indicator-support diagnostics, temporal-pattern labels, and boundary checks rather than reduced to a single deterministic score.

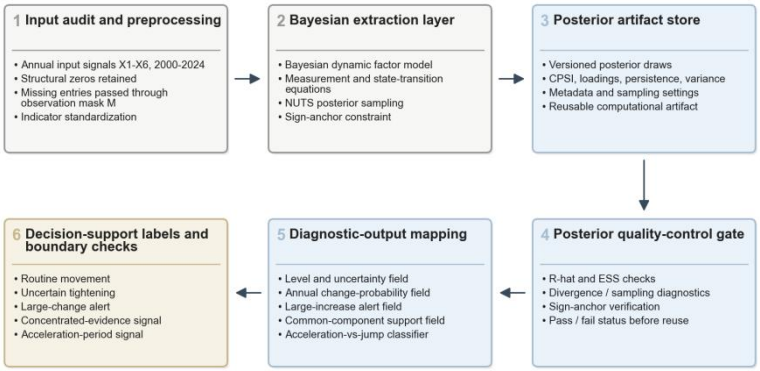


Figure 1 Six-Stage Posterior-Artifact and Diagnostic-Output Workflow for Uncertainty-Aware CPSI Decision Support

Note: The workflow begins with heterogeneous input signals and a Bayesian dynamic factor model. It stores posterior draws as reusable artifacts, applies a posterior quality-control gate, maps validated draws into machine-readable diagnostic fields, and exports decision-support labels with explicit measurement boundaries.

Algorithm 1 summarizes the computational pipeline implemented in the retained code package.

Algorithm 1. Bayesian CPSI Diagnostic Pipeline

Input: annual indicators X1-X6 and calendar-year observation mask M.

1. Audit retained indicator columns and align observations to 2000-2024.
2. Standardize observed cells and retain the observation mask for missing entries.
3. Estimate the Bayesian extraction layer by NUTS posterior sampling.
4. Sign-normalize posterior draws by anchoring the loading on X3 to be positive.
5. Save the versioned posterior artifact with CPSI, loading, persistence, and variance draws.
6. Apply the posterior quality-control gate before downstream reuse.
7. Map the validated artifact into diagnostic-output fields and decision-support labels.

Output: posterior artifact, quality-control status, diagnostic-output fields, monitoring flags, and policy-mix labels.

Algorithm 1 defines the upstream audit, extraction, artifact, and quality-control process. Algorithm 2 starts from the validated artifact and specifies the downstream decision engine. Its purpose is to turn posterior draws into machine-readable monitoring fields and decision-support labels, so it should be read as the operational layer that follows the Bayesian extraction algorithm.

Algorithm 2. Posterior Diagnostic Decision Engine

Input: validated posterior artifact A_{QC}, annual observation mask M, and large-change threshold τ.

1. Compute $\mathrm{CPSI}_{\mathrm{mean,t}}$, $\mathrm{CPSI}_{\mathrm{p05,t}}$, and $\mathrm{CPSI}_{\mathrm{p95,t}}$ from posterior draws.
2. Compute Δf_t, $\Pr(\Delta f_t > 0)$, and $\Pr(\Delta f_t > \tau)$ for each year.
3. Assign monitoring labels: routine movement, uncertain tightening, or large-change alert.
4. Compute $R^2_{\mathrm{common,i}}$ to identify which indicator channels support the latent signal.
5. Estimate breakpoint-mass and acceleration diagnostics from posterior CPSI paths.
6. Compare no-jump and jump-augmented innovation models using PSIS-LOO deviance differences.
7. Export a decision table containing annual flags, uncertainty intervals, support channels, and temporal-pattern labels.

Output: machine-readable monitoring fields, policy-mix support labels, and acceleration-versus-jump classification.

Table 2 Machine-Readable Diagnostic Fields Exported by the Posterior Decision Engine

Output field	Computation	Decision-support use
CPSI_mean_t	Posterior mean of f_t	Annual level monitoring
CPSI_p05_t / CPSI_p95_t	5th and 95th posterior percentiles	Uncertainty reporting
p_increase_t	$\Pr(\Delta f_t > 0)$	Directional tightening evidence
p_large_t	$\Pr(\Delta f_t > \tau)$	Large-change alert flag
common_component_share_i	Common-component share by indicator ($R^2_{\mathrm{common,i}}$)	Evidence-breadth and policy-mix diagnosis
breakpoint_mass	Posterior mass over candidate break years	Acceleration localization
jump_improve_prob	$\Pr(\Delta \text{ deviance} > 0)$	Check whether a one-year jump is statistically necessary

The formula sequence follows the same pipeline as Figure 1, Algorithm 1, Algorithm 2, and Table 2: input audit, Bayesian extraction, posterior-artifact construction, posterior quality control, diagnostic-output mapping, and decision-support labelling.

Input audit and pre-processing start with an observation mask. Standardization is applied to observed cells, while missing entries remain unobserved likelihood contributions rather than imputed values:

$$Z_{i,t} = \frac{X_{i,t} - X_i}{s_i}, \quad (i,t) \in D_{obs} = \{(i,t): M_{i,t} = 1\} \tag{1}$$

Equation (1) converts heterogeneous annual signals onto a common scale without assigning policy weights. The Bayesian extraction layer then links the standardized inputs to a latent annual CPSI state:

$$Z_{i,t} = \lambda_i f_t + \varepsilon_{i,t},\ \varepsilon_{i,t} \sim N(0, \sigma_i^2) \quad (2)$$

The dynamic layer gives the latent state a sparse autoregressive transition:

$$f_t = \varphi f_{t-1} + \eta_t,\ \eta_t \sim N(0,1) \quad (3)$$

This specification follows the logic of dynamic factor measurement, where several noisy observables discipline a common unobserved process [9-10]. The implemented Bayesian model uses weakly regularizing priors that match the retained code:

$$\varphi \sim Uniform(-0.98, 0.98),\ \lambda_i \sim N(0,1),\ \sigma_i = 0.05 + HalfNormal(0.30) \quad (4)$$

Posterior inference uses the No-U-Turn Sampler (NUTS) [13], with four chains, 5,000 tuning iterations, 5,000 posterior draws per chain, target acceptance 0.92, and seed 123. Factor sign indeterminacy is resolved after sampling by normalizing draws so that the loading on policy-text intensity is positive, $\lambda_3 > 0$. This convention fixes the reported direction of CPSI without changing the likelihood.

After sampling, the workflow stores posterior draws as a reusable artifact rather than reducing the model to a single mean index:

$$A = \{f_t^{(s)}, \lambda_i^{(s)}, \sigma_i^{(s)}, \varphi^{(s)} : s = 1, \ldots, S\} \quad (5)$$

This artifact is the computational boundary between Bayesian extraction and downstream decision support. Before any draw is reused, the posterior quality-control gate checks convergence and effective sample size:

$$A_{QC} = A\ if\ max_j \hat{R}_j \leq 1.01\ and\ min_j\ ESS_j \geq 400 \quad (6)$$

The retained sample satisfies this gate, with a maximum basic R-hat of 1.0055 and a minimum pooled autocorrelation ESS of 451.8 for the factor, loading, scale, and persistence summaries [14-15].

The posterior diagnostic engine then maps the validated artifact into a machine-readable output schema:

$$O = g(A_{QC}) = \{f_t, q_{0.05,t}, q_{0.95,t}, p_{+,t}, p_{\tau,t}, R_{common,i}^2, b_\kappa, p_{jump}\} \quad (7)$$

The annual-change fields are computed from posterior draw differences, so alert labels depend on posterior probability rather than on a deterministic threshold:

$$\Delta f_t^{(s)} = f_t^{(s)} - f_{t-1}^{(s)};\ p_{+,t} = Pr(\Delta f_t > 0),\ p_{\tau,t} = Pr(\Delta f_t > \tau) \quad (8)$$

Indicator-level diagnosis uses common-component shares. A high value means that, within this model and annual sample, the indicator tracks the common latent signal more tightly:

$$R_{common,i}^2 = \frac{Var_t(\lambda_i f_t)}{Var_t(\lambda_i f_t) + \sigma_i^2} \quad (9)$$

The final diagnostic branch asks whether the latent path is better described as gradual movement, acceleration, or a discrete annual jump. Acceleration is evaluated with a segmented trend diagnostic applied to posterior CPSI draws, using breakpoint-estimation logic [16]:

$$f_t = \alpha + \beta_1 t + \beta_2 (t - \kappa)_+ + u_t \quad (10)$$

Here, kappa is the candidate breakpoint and the positive-part term equals max(x, 0). Jump necessity is evaluated as a separate model-comparison check. The diagnostic compares the no-jump specification with a jump-augmented innovation model using posterior-draw deviance differences and PSIS-LOO logic [17]. This separation keeps the interpretation explicit: Eq. (10) tests whether the slope changes, whereas the jump comparison tests whether a one-year level shift is required.

The final module provides external process-alignment checks. Horizon-specific local projection regressions relate selected process variables to a one-standard-deviation movement in CPSI [18]. Inference uses heteroskedasticity- and autocorrelation-consistent standard errors [19]. These checks remain outside the measurement equation. They assess whether the inferred policy signal is accompanied by plausible movements in green innovation, green investment, and carbon-efficiency adjustment. They are descriptive association checks and are not interpreted as causal effects of policy.

The retained code package fixes the implementation boundary. The estimation script selects the six X1-X6 indicators, runs the Bayesian DFM, applies sign normalization, writes the posterior artifact, applies the quality-control gate, and exports the diagnostic-output schema. Downstream scripts use only the retained measurement, figure, and diagnostic outputs listed in the result-boundary file. Older dashboard outputs, Markov-switching variants, and timestamped development files are excluded. This prevents development-stage results from entering the conference paper.

5 COMPUTATIONAL OUTPUTS AND DIAGNOSTIC RESULTS

This section validates the diagnostic-output mapping defined in Eq. (7) as a set of computational decision-support artifacts. For the conference-scope contribution, validation focuses on artifact traceability, uncertainty propagation, output-schema consistency, and diagnostic stability. The reported outputs are the posterior CPSI path, 5th-95th percentile intervals, annual-change probabilities, large-increase alerts, common-component shares, acceleration diagnostics, jump-model comparison values, and external alignment checks. Each result is tied to the posterior artifact or exported CSV summaries. The section therefore evaluates how the pipeline produces reusable diagnostic outputs rather than reconstructing policy events or estimating causal effects.

Figure 2 reports the level-and-uncertainty field exported from the posterior artifact. The posterior CPSI mean increases from -2.10 in 2000, with a 5th-95th percentile interval of [-3.05, -1.35], to 6.30 in 2024, with an interval of [4.41, 8.38]. The output is a posterior summary of common variation across the six audited inputs, not a deterministic score. The wider

later-period interval is exported because uncertainty is part of the decision-support object rather than a supplementary caveat.

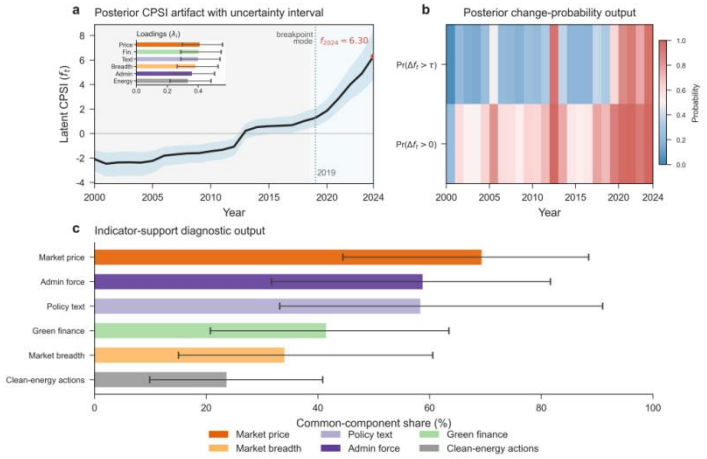


Figure 2 Diagnostic Output Bundle Generated from the Posterior CPSI Artifact

Note: Panel a reports the posterior CPSI path and 5th-95th percentile interval as the level-and-uncertainty field. Panel b reports posterior change-probability fields, including directional-increase probabilities and large-increase alert probabilities. Panel c reports common-component support fields that indicate how tightly each input channel tracks the latent policy signal within the retained Bayesian measurement model.

The annual-change module converts the posterior path into monitoring probabilities. A positive annual movement is not sufficient evidence for a substantive change, so the code reports both directional-increase probabilities and probabilities that the annual increase exceeds a posterior large-increase threshold. From 2021 onward, the large-increase probabilities are 0.86, 0.92, 0.77, and 0.96. The 2023 value remains directionally positive but does not clear the 0.80 alert threshold, which prevents the diagnostic from overstating a uniformly decisive sequence of annual shocks.

Figure 3 validates the temporal-pattern classifier. The breakpoint field places the posterior mode at 2019, with posterior mass of 0.592 on 2019 and 0.83 on 2019-2020 combined. The estimated post-break slope increment is 0.81, with a 5th-95th percentile interval of [0.27, 1.26]. The piecewise acceleration field also dominates the linear benchmark, with a posterior summary of 28.8 and interval [11.3, 46.8]. These outputs trigger a rate-shift label for the retained artifact around the end of the 2010s.

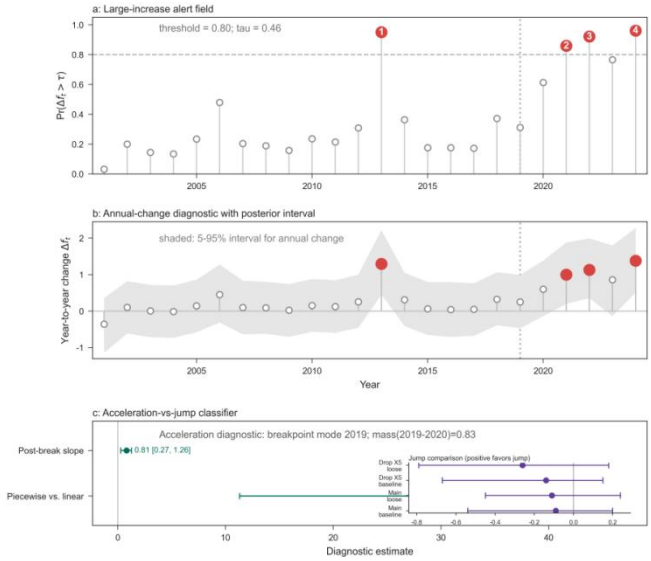


Figure 3 Temporal-Pattern Diagnostic Outputs Exported by the Posterior Decision Engine

Note: Panel a reports the large-increase alert field based on posterior probabilities that annual CPSI changes exceed the monitoring threshold. Panel b reports the annual-change diagnostic field with 5th-95th percentile intervals. Panel c reports the acceleration-versus-jump classifier, comparing segmented acceleration evidence with jump-model evidence; positive jump-comparison values would favor a one-year jump model.

The jump comparison gives the complementary result. Across retained specifications, deviance-improvement summaries are below zero. They are -0.09 for the main baseline specification, -0.11 under the looser jump prior, -0.14 when the green-finance indicator is excluded, and -0.26 in the corresponding loose-prior sensitivity run. Because positive values would favor the jump model, these estimates do not support the claim that a one-year discontinuity is statistically necessary at annual frequency. The posterior probabilities that the jump model improves fit are also modest, ranging from

0.151 to 0.304 across these checks. This result should not be stated as evidence that abrupt policy actions never occurred. It is a measurement statement: after uncertainty is propagated through the annual CPSI, the observed non-smoothness is better represented as acceleration than as a required discrete jump.

The indicator diagnostics show why the common-factor interpretation is plausible but not uniform across indicators. The common-component share is highest for the market price signal, at 69.3 percent. Administrative force and policy-text intensity follow closely, at 58.8 and 58.3 percent. Green finance contributes a moderate common share of 41.5 percent, while market breadth and clean-energy actions are lower, at 34.0 and 23.6 percent. These values should not be treated as normative weights or as estimates of policy importance. They indicate how tightly each indicator tracks the latent policy signal within this annual sample. In practical terms, the CPSI is anchored mainly by price, administrative, and policy-text information, while the remaining indicators add coverage and sectoral context with larger idiosyncratic components.

For decision support, the three empirical layers should be read jointly. The CPSI level describes the accumulated strength of the latent policy signal. The annual-change probabilities identify years in which the posterior evidence for additional tightening warrants attention. The indicator diagnostics show whether the estimated signal is broadly supported or dominated by a narrow subset of inputs. This joint reading is more operational than a single ranked index because it tells analysts when the signal changes, how confident that change is, and which data channels carry it.

Taken together, the exported results validate three properties of the workflow. First, the posterior artifact provides a reproducible uncertainty-bounded CPSI measurement for 2000-2024. Second, the temporal classifier labels the dominant non-smooth feature as later-period acceleration rather than a statistically necessary annual jump. Third, the output schema reports the probability of large annual increases and the indicator support behind the latent signal. These properties make the artifact usable for monitoring and decision support while preserving the measurement boundary: the section validates diagnostic outputs, not causal policy effects.

6 DECISION-SUPPORT IMPLICATIONS

The decision-support contribution begins from the labelled output schema generated by the posterior diagnostic engine. CPSI is not used as a deterministic scorecard for ranking climate policy performance. It is used as a diagnostic object that combines five machine-readable outputs: the posterior CPSI path, the 5th-95th percentile interval, annual large-increase probabilities, common-component shares, and the acceleration-versus-jump classification. This output bundle is consistent with decision-support and business-intelligence work in which analytic systems structure judgement without replacing it [7-8].

The first diagnostic use is monitoring. For each year, the analyst should read three quantities together. The posterior mean gives the direction and level of the latent policy signal; the posterior interval reports measurement precision; and the large-increase probability determines whether the year should receive a monitoring flag. A positive movement in the mean therefore does not automatically become an alert. It becomes operationally salient only when the posterior probability of a large annual increase is high enough to justify closer review. Compared with deterministic composite-index reporting [3], this rule keeps posterior uncertainty inside the monitoring signal.

The second use is breadth-of-evidence assessment. The common-component shares indicate which observed channels anchor the latent signal in a given measurement system. In this application, market price information, administrative force, and policy-text intensity provide the tightest support for the common factor, while market breadth, green finance, and clean-energy actions contribute additional coverage with larger idiosyncratic components. These shares are not normative weights. They tell the analyst whether a CPSI movement is broadly supported across instrument families or mainly carried by a narrower subset of data channels.

The third use is input-channel interpretation. When the CPSI rises and common-component support is concentrated in price, administrative, and text channels, the management question is not which instrument is most effective. The question is whether the exported signal is being carried mainly by market, enforcement, or agenda-setting information, and whether missing or weakly aligned channels require additional data checks. This use converts the index from a single headline number into an input-support diagnosis.

The fourth use is temporal-pattern classification. The acceleration diagnostic localizes a faster-growth window around 2019, whereas the jump comparison asks whether a discrete one-year level shift is statistically necessary. At annual frequency, the retained evidence supports acceleration rather than a stable jump year. This distinction is useful for management interpretation because it prevents a visible intensification from being mechanically reclassified as an institutional regime break.

The practical reporting format is therefore a compact diagnostic panel rather than a policy ranking. Each annual update should report the CPSI mean, interval width, large-increase probability, leading common-component channels, and temporal-pattern label. A monitoring note can then classify the observation as routine movement, uncertain movement, large-increase alert, concentrated-evidence signal, or acceleration-period signal. Such labels help analysts decide where to focus review effort while preserving probabilistic and adaptive decision-making under uncertainty [11, 20].

Table 3 Decision-Rule Interpretation for Annual CPSI Monitoring

Monitoring label	Posterior rule	Interpretation
Routine movement	$\Pr(\Delta f_t > \tau) < 0.50$	No large-change alert; continue regular monitoring.

Monitoring label	Posterior rule	Interpretation
Uncertain tightening	$\Pr(\Delta f_t > 0) > 0.60 \quad \Pr(\Delta f_t > \tau) < 0.80$	Direction is positive but not strong enough for an alert.
Large-change alert	$\Pr(\Delta f_t > \tau) \geq 0.80$	Substantive increase is sufficiently probable for focused review.
Concentrated-evidence signal	High CPSI movement with support concentrated in few channels	Check whether the movement is driven by a narrow subset of indicators.
Acceleration-period signal	Breakpoint mass concentrated around a candidate window	Treat the change as rate-shift evidence, not as a one-year regime break.

These implications also define the boundary of the framework. CPSI helps organize heterogeneous evidence about climate policy stringency; it does not estimate or identify causal effects of climate policy on emissions, innovation, investment, or macroeconomic outcomes. The contribution for intelligent management decision science is a measurement-and-diagnosis workflow. It assembles heterogeneous policy signals, propagates uncertainty, translates posterior draws into operational flags, and states when a diagnostic output should not be read as causal policy evaluation.

7 COMPUTATIONAL VALIDATION AND BOUNDARY CONDITIONS

This section is reframed as computational validation rather than a conventional robustness appendix. The checks evaluate whether the decision-support outputs remain stable under three perturbations: measurement perturbation, structural diagnostic perturbation, and external-alignment inference perturbation. This distinction responds to the engineering requirement that a decision-support pipeline should be reproducible, quality-controlled, and interpretable under alternative computational settings.

Table 4 is organized into two panels. Panel A reports measurement robustness and temporal-pattern diagnostics, while Panel B reports external process-alignment checks. The columns separate point estimates from sensitivity values or p-values, so the table does not mix deterministic fit gains, posterior diagnostics, and association-test evidence under a single statistic.

Panel A first shows that the acceleration result is not driven by the green-finance indicator. When X5 is excluded, the breakpoint mode remains 2019, and the posterior mass on 2019 increases from 0.592 in the main specification to 0.733. The deterministic fit gain for the piecewise acceleration diagnostic remains positive. The magnitude is smaller under the perturbation, but the qualitative conclusion is unchanged: the later-period CPSI path is better read as acceleration than as a smooth linear drift.

The complementary structural check concerns jump stability. Across the main and drop-X5 measurements, and under both baseline and more permissive jump priors, no year crosses the 0.3 or 0.5 posterior-probability thresholds for a stable jump year. This does not prove that no abrupt policy action existed. It means that, at annual frequency and under the retained measurement model, a discrete jump component is not statistically necessary for explaining the non-smoothness in CPSI. This conservative claim is consistent with Sections 5 and 6.

Panel B defines the external process-alignment boundary condition. These checks are construct-validity diagnostics rather than outcome-effect estimates. A one-standard-deviation CPSI movement is associated with subsequent increases in green innovation and green investment, and with a lagged improvement in carbon efficiency. The carbon-efficiency association is strongest at horizon 3, where the coefficient is about 0.52 and remains significant under the HAC sensitivity check. At the same time, placebo-lead tests are mixed: green innovation and investment lead coefficients are imprecise, while carbon efficiency has a negative h = -2 lead. This pattern supports the external alignment of CPSI as a measurement signal while keeping the causal boundary visible.

The boundary conditions are explicit. First, annual data have limited power to detect short-lived discontinuities; the no-necessary-jump conclusion should not be extrapolated to monthly or event-level data. Second, a national latent index compresses regional and sectoral heterogeneity. Third, process-alignment checks are association checks, not treatment-effect estimates. The proposed framework therefore gives a reproducible uncertainty-aware measurement of climate policy stringency, an acceleration diagnostic, and a bounded account of external alignment evidence.

Table 4 Computational Validation Diagnostics and External Process-Alignment Checks

Check	Estimate	Sensitivity or p-value	Interpretation
Panel A. Measurement robustness and temporal-pattern diagnostics			
Acceleration timing	Breakpoint mode = 2019; posterior mass on 2019 = 0.592.	Drop-X5: breakpoint mode = 2019; posterior mass = 0.733.	Later-period acceleration is not driven by X5.
Acceleration magnitude	Posterior piecewise-vs-linear gain = 28.8 [11.3, 46.8].	Deterministic BIC gain = 39.73 in the main fit.	Positive values favor acceleration over linear drift; the deterministic BIC value is not the posterior summary.
Jump stability	No stable jump year at thresholds 0.3 or 0.5.	Loose-prior and drop-X5 checks: no stable year at 0.5; maximum improvement probability <= 0.304.	A one-year jump is not statistically necessary at annual frequency.
Panel B. External process-alignment checks			

Check	Estimate	Sensitivity or p-value	Interpretation
Green innovation alignment	Coefficient = 0.276 at h = 0.	HAC p = 0.149.	Positive but imprecise external alignment evidence.
Green investment alignment	Coefficient = 0.441 at h = 0.	HAC p = 0.099.	Positive but marginal external alignment evidence.
Carbon-efficiency alignment	Coefficient = 0.517 at h = 3.	HAC p = 0.00009; placebo lead h = -2 coefficient = -0.796, p = 0.016.	Stronger downstream alignment, but not causal identification.

8 CONCLUSION

The study designed and validated a Bayesian data-mining decision-support workflow for uncertainty-aware policy-signal processing. The framework treats climate policy stringency as a latent annual state inferred from six heterogeneous policy signals. The Bayesian dynamic factor model carries posterior draws into a diagnostic engine. That engine links input audit, probabilistic extraction, posterior quality control, machine-readable output generation, model comparison, and management decision support in one reproducible pipeline.

The retained 2000-2024 posterior artifact exports a CPSI path that rises from -2.10 in 2000 to 6.30 in 2024. It also exports three diagnostic outputs: a later-period acceleration label centered around 2019, jump-comparison evidence that does not require a one-year discontinuity at annual frequency, and common-component support fields. These fields show that market price, policy-text intensity, and administrative force carry the latent signal more tightly than some broader policy-action indicators. These outputs support the interpretation of CPSI as an uncertainty-bounded monitoring signal, not as a causal effect estimate.

For management decision science, the contribution lies in the reusable computational workflow as much as in the China application. Analysts can use posterior paths, uncertainty intervals, large-increase probabilities, indicator-support diagnostics, and rule-based labels to flag years requiring closer review, audit evidence breadth, and check whether an apparent change is broad-based or indicator-specific. The framework remains extensible to additional regions, sectors, or higher-frequency policy signals when comparable data become available. Its boundary is also explicit: annual CPSI is suited to measurement and diagnosis, while causal policy evaluation, treatment-effect estimation, event-level attribution, and regional attribution require separate research designs.

COMPETING INTERESTS

The authors have no relevant financial or non-financial interests to disclose.

FUNDING

This work was supported by the Energy Foundation Beijing Representative Office [grant number G-2505-36421]; the Guangdong Basic and Applied Basic Research Foundation [grant number 2025A1515012990]; the 2025 Guangdong Provincial Higher Education Institutions' Characteristic Innovation Project [grant number 2025WTSCX035]; and the Rockefeller Brothers Fund [grant number 23-246].

AI USE DISCLOSURE

OpenAI ChatGPT (GPT-5.5 Thinking; accessed on 5 June 2026) was used only for English language polishing, manuscript-format checking, and assistance in planning the schematic layout of Figure 1. The authors independently verified all scientific content, including data sources, indicator construction, code execution, statistical analysis, numerical results, references, equations, tables, and figures. No generative AI tool was used to collect or synthesize data, alter research data, manipulate figures, conduct statistical analysis, or generate research results.

REFERENCES

[1] Kruse T, Dechezleprêtre A, Saffar R, et al. Measuring environmental policy stringency in OECD countries: An update of the OECD composite EPS indicator. OECD Economics Department Working Papers, No. 1703. Paris: OECD Publishing, 2022. DOI: 10.1787/90ab82e8-en.

[2] Botta E, Koźluk T. Measuring Environmental Policy Stringency in OECD Countries: A Composite Index Approach. OECD Economics Department Working Papers, No. 1177. Paris: OECD Publishing, 2014. DOI: 10.1787/5jxrjnc45gvg-en.

[3] Nardo M, Saisana M, Saltelli A, et al. Handbook on Constructing Composite Indicators: Methodology and User Guide. OECD Statistics Working Papers, No. 2005/03. Paris: OECD Publishing, 2005. DOI: 10.1787/533411815016.

[4] Li B, Fu E X, Yang S H, et al. Measuring China's Policy Stringency on Climate Change for 1954-2022. Scientific Data, 2025, 12: 188. DOI: 10.1038/s41597-025-04476-0.

[5] Fayyad U, Piatetsky-Shapiro G, Smyth P. From Data Mining to Knowledge Discovery in Databases. AI Magazine, 1996, 17(3): 37-54. DOI: 10.1609/aimag.v17i3.1230.

[6] Provost F, Fawcett T. Data Science and its Relationship to Big Data and Data-Driven Decision Making. Big Data, 2013, 1(1): 51-59. DOI: 10.1089/big.2013.1508.
[7] Shim J P, Warkentin M, Courtney J F, et al. Past, present, and future of decision support technology. Decision Support Systems, 2002, 33(2): 111-126. DOI: 10.1016/S0167-9236(01)00139-7.
[8] Chen H C, Chiang R H L, Storey V C. Business Intelligence and Analytics: From Big Data to Big Impact. MIS Quarterly, 2012, 36(4): 1165-1188. DOI: 10.2307/41703503.
[9] Stock J H, Watson M W. Forecasting using principal components from a large number of predictors. Journal of the American Statistical Association, 2002, 97(460): 1167-1179. DOI: 10.1198/016214502388618960.
[10] Aguilar O, West M. Bayesian Dynamic Factor Models and Portfolio Allocation. Journal of Business & Economic Statistics, 2000, 18(3): 338-357. DOI: 10.1080/07350015.2000.10524875.
[11] Ghahramani Z. Probabilistic machine learning and artificial intelligence. Nature, 2015, 521(7553): 452-459. DOI: 10.1038/nature14541.
[12] Nachtigall D, Lutz L, Cárdenas Rodríguez M, et al. The Climate Actions and Policies Measurement Framework: A Structured and Harmonised Climate Policy Database to Monitor Countries' Mitigation Action. OECD Environment Working Papers, No. 203. Paris: OECD Publishing, 2022. DOI: 10.1787/2caa60ce-en.
[13] Hoffman M D, Gelman A. The No-U-Turn Sampler: Adaptively Setting Path Lengths in Hamiltonian Monte Carlo. Journal of Machine Learning Research, 2014, 15: 1593-1623. https://jmlr.org/papers/v15/hoffman14a.html.
[14] Gelman A, Rubin D B. Inference from iterative simulation using multiple sequences. Statistical Science, 1992, 7(4): 457-472. DOI: 10.1214/ss/1177011136.
[15] Vehtari A, Gelman A, Simpson D, et al. Rank-normalization, folding, and localization: An improved R-hat for assessing convergence of MCMC. Bayesian Analysis, 2021, 16(2): 667-718. DOI: 10.1214/20-BA1221.
[16] Muggeo V M R. Estimating regression models with unknown break-points. Statistics in Medicine, 2003, 22(19): 3055-3071. DOI: 10.1002/sim.1545.
[17] Vehtari A, Gelman A, Gabry J. Practical Bayesian model evaluation using leave-one-out cross-validation and WAIC. Statistics and Computing, 2017, 27(5): 1413-1432. DOI: 10.1007/s11222-016-9696-4.
[18] Jordà Ò. Estimation and inference of impulse responses by local projections. American Economic Review, 2005, 95(1): 161-182. DOI: 10.1257/0002828053828518.
[19] Newey W K, West K D. A Simple, Positive Semi-definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix. Econometrica, 1987, 55(3): 703-708. DOI: 10.2307/1913610.
[20] Haasnoot M, Kwakkel J H, Walker W E, et al. Dynamic adaptive policy pathways: A method for crafting robust decisions for a deeply uncertain world. Global Environmental Change, 2013, 23(2): 485-498. DOI: 10.1016/j.gloenvcha.2012.12.006.
[21] Wind Information Co., Ltd. China carbon allowance trading price records, 2000-2024. Wind Economic Database, 2024.
[22] International Carbon Action Partnership (ICAP). China National ETS. ICAP ETS Detailed Information, 2025.
[23] The State Council of the People's Republic of China. National Policy Document Library. China Government Website, n.d. Retrieved January 2, 2026 from https://sousuo.www.gov.cn/zcwjk/policyDocumentLibrary.
[24] Ministry of Ecology and Environment of the People's Republic of China. Environmental Inspection Management. Ministry of Ecology and Environment website, n.d. Retrieved January 2, 2026 from https://www.mee.gov.cn/ywgz/zysthjbhdc/dcjl/.
[25] People's Bank of China (PBoC). Statistical Reports on Loan Directions of Financial Institutions (2000-2024). PBoC public statistical releases, various years. Retrieved January 2, 2026 from http://www.pbc.gov.cn/goutongjiaoliu/113456/index.html.

APPENDIX A. DATA-SOURCE CONVENTIONS AND DESCRIPTIVE STATISTICS

Table A1 Source Conventions and Raw-Scale Descriptive Statistics for CPSI Input Indicators

Code	Indicator (unit)	Mean	SD	Min	Max	Source or construction basis
X1	Carbon price level (RMB/t)	20.98	25.72	0.00	90.00	Carbon allowance trading price records [21]
X2	ETS coverage share (%)	9.72	17.19	0.00	55.00	Emissions trading coverage information [22]
X3	Policy-text intensity (index)	123.22	87.02	5.12	280.68	Climate and low-carbon policy-text coding [23]
X4	Administrative force (index)	50.15	34.60	0.00	95.79	Event-scored enforcement and administrative-pressure records [24]
X5	Outstanding green loans (trillion RMB)	7.93	9.65	0.77	39.08	Green-loan balance, reconstructed where needed [25]
X6	Clean-energy policy actions (index)	0.790	0.455	0.101	1.730	Annual clean-energy policy-action coding [4]

Notes: Statistics are computed from the raw, unstandardized annual series for 2000-2024. Model estimation uses standardized series as described in Section 3. Structural zeros are retained where they indicate the absence of a policy instrument or market signal; other unavailable observations are not filled by narrative interpolation. The six inputs are used as noisy signals of latent climate policy stringency, not as predetermined weights or causal treatments.