

Bayesian Belief Layers Under Non-Stationarity

Core conclusion

A Bayesian belief layer should not treat non-stationarity as “strong evidence for the same prior.” It should treat it as **model uncertainty about whether the data-generating process has changed**. Formally, the decision is not “update versus reset” inside one model; it is **continuation model versus changepoint/regime model**, with action chosen by posterior expected loss. Under equal misclassification costs, a break is declared when the posterior probability of a changepoint exceeds 0.5; under unequal costs, the threshold shifts with the relative costs of false alarms and missed breaks. In Bayes-factor terms, the relevant quantity is posterior odds, which are prior odds multiplied by the Bayes factor. ¹

For an operational belief system, the practical implication is: do **not** literally erase history whenever a break is suspected. Instead, maintain a posterior over structural states, instantiate a new regime posterior when break probability is high, preserve prior regime posteriors for possible recurrence, and produce predictions from a weighted mixture while regime uncertainty remains unresolved. That is the common structure across BOCPD, online multiple-changepoint filtering, and Markov-switching models. ²

Formal criterion for discarding a prior

The formal criterion is a **Bayesian model comparison problem**. Let M_0 denote “the current regime continues” and M_1 denote “a changepoint occurred and a new regime started.” For new evidence x_t , compare

$$BF_{10,t} = \frac{p(x_t \mid x_{1:t-1}, M_1)}{p(x_t \mid x_{1:t-1}, M_0)}.$$

Then choose the action that minimizes posterior expected loss. With symmetric 0–1 loss, this reduces to declaring a regime break when $P(M_1 \mid x_{1:t}) > 0.5$. With asymmetric costs, the threshold should be higher or lower depending on the operational cost of false reset versus missed break. The key point is that **“discard the prior” is never justified by a large residual alone**; it is justified when the data are better explained by a structural-break model than by continued updating under the old regime. ¹

In BOCPD language, the central sufficient statistic is the posterior over **run length** r_t , the time since the last changepoint. A switch is indicated when posterior mass moves toward $r_t = 0$, meaning that the newest datum is more compatible with a fresh segment than with growth of the current one. Adams and MacKay formulate this as exact online inference over the most recent changepoint, using a hazard distribution over changepoint intervals and message passing over run lengths. ³

For a belief layer, the continuation model should use the current regime posterior, while the break model should use a **regime-entry prior** or “fresh-segment prior.” The criterion for “discarding” is therefore best implemented as: if the posterior probability of a break exceeds a cost-calibrated threshold, stop borrowing information from the active regime for local prediction and reinitialize local sufficient statistics from the

regime-entry prior. Preserve the old regime posterior in an archive or regime library rather than deleting it. That preserves information for recurrence, audit, and retrospective smoothing. ⁴

A useful refinement is to separate **predictive surprise** from **structural surprise**. The literature on Bayes-factor surprise in volatile environments shows that adaptation should depend on a probability ratio that compares the new observation under the prior belief versus the current belief, not on raw improbability alone. That is the correct direction for belief systems: a datum that is improbable under the current regime but also improbable under the entry prior is an outlier candidate, not necessarily a regime break. ⁵

What BOCPD and BCP imply for online regime detection

BOCPD is explicitly an **online** detector. Adams and MacKay derive an exact online algorithm for the most recent changepoint and compute the run-length posterior with a message-passing recursion. Fearnhead and Liu generalize online inference to multiple changepoints, showing exact filtering and posterior simulation for changepoint positions, with particle-filter approximations when linear-time updating is needed. ⁶

Barry and Hartigan's product-partition approach, by contrast, is primarily **retrospective/offline segmentation**. The standard `bcp` implementation explicitly states that it uses the Barry and Hartigan product-partition model with MCMC, a prior probability of change p_0 at each location, and returns posterior probabilities of change over the sequence. That makes BCP valuable for re-segmentation, posterior calibration, and retrospective relabeling of history, but it is not the natural core of a real-time detector. ⁷

The correct response to a detected changepoint in an online belief system is therefore layered:

1. **Local reset**: initialize a new active-segment posterior from the regime-entry prior.
2. **Historical freeze**: stop pooling new observations into the prior active regime.
3. **Uncertainty retention**: keep posterior mass over alternative run lengths or regime labels for several steps rather than committing instantly.
4. **Retrospective smoothing**: optionally run an offline BCP-style pass over a sliding window to merge, split, or relabel breaks after more evidence arrives. ⁸

The robust BOCD literature is also operationally important. Altamirano, Briol, and Knoblauch show that standard online Bayesian changepoint methods can be brittle under outliers or model misspecification, and propose a generalized-Bayes variant with improved robustness. For a belief layer, this means changepoint triggers should be based on **persistent posterior evidence** or robust predictive scores, not on a single anomalous edge update. ⁹

What Markov-switching VAR models imply for prior architecture

Markov-switching models handle non-stationarity by introducing a **latent regime variable** s_t with Markov transitions. Kim's extension of Hamilton's framework embeds switching into a general state-space model and adds filtering and smoothing over latent states. Krolzig's MS-VAR framework generalizes this to multivariate dynamics, explicitly allowing regime-specific dynamics in vector autoregressions. ¹⁰

Two architectural consequences follow. First, the belief system should maintain **regime-specific parameter posteriors**, because the relationships among variables differ by regime. Second, it should produce **mixture predictions** during transitions, because the active regime is uncertain. In Krolzig's forecasting analysis, predicted densities are mixtures of normals weighted by regime probabilities; the forecast need not come from a single hard-assigned regime. ¹¹

For practice, the most defensible design is **both separate priors and a mixture model**:

- separate posteriors $p(\theta_k \mid D_k)$ for each inferred regime k ;
- a transition model $P(s_t \mid s_{t-1})$;
- a predictive mixture $\sum_k p(y_{t+1} \mid \theta_k, s_{t+1} = k) P(s_{t+1} = k \mid D_{1:t})$.

This matches the logic of MS-VARs: regime identity is latent, transitions are probabilistic, and predictions integrate over that uncertainty. Hard switching without mixture is appropriate only after posterior regime probability is near-degenerate. ¹²

If regimes recur, separate priors are especially useful because they let the system **re-enter** a previously learned regime instead of relearning from scratch. If drift is gradual rather than abrupt, time-varying transition probabilities or dynamic model averaging are the appropriate generalization. The time-varying-transition literature shows that constant transition matrices are often too restrictive, and dynamic model averaging explicitly changes model weights over time via forgetting. ¹³

What supervisory stress testing implies for adversarial belief evaluation

Supervisory stress testing does not treat scenarios as forecasts; it treats them as **structured, severe, internally coherent counterfactuals** designed to reveal vulnerabilities that ordinary historical calibration can miss. Guidance from the Federal Reserve ¹⁴ states that supervisory scenarios are hypothetical and that the severely adverse scenario should include salient risks that would be underrepresented if the design relied only on past recessions or ordinary historical relationships. The Office of the Comptroller of the Currency ¹⁵ similarly frames company-run stress tests as forward-looking assessments of resilience under economic and financial stress rather than forecasts. ¹⁶

The decisive idea for belief systems is **joint extremity**. Stress-testing guidance from the Bank for International Settlements ¹⁷, the Federal Reserve, the OCC, the Federal Deposit Insurance Corporation ¹⁸, and the International Monetary Fund ¹⁹ emphasizes that vulnerabilities appear when multiple channels move together: market prices, liquidity, counterparty quality, collateral values, funding conditions, and institution-level interdependence. The IMF notes that modest but correlated shocks can generate extreme outcomes, especially because correlations often rise in distress. The OCC explicitly warns that stress testing should not be limited to higher volatility or historical confidence intervals, but should model interactions among market moves, liquidity, counterparty credit quality, and collateral. Federal Reserve scenarios add a global market shock and, for certain firms, default of the largest counterparty. ²⁰

Translated to a knowledge graph, adversarial belief evaluation should therefore test **correlated evidence failures**, not only single-edge perturbations. The meaningful adversarial scenarios are:

- simultaneous insertion of mutually reinforcing false triples across several relations;
- coordinated deletion of the most influential supporting triples for a target fact;
- synchronized corruption of provenance-heavy or high-centrality nodes;
- time-clustered contradictory evidence from multiple sources;
- “largest counterparty default” analogs, where the system loses its most influential source, relation type, or subgraph at once. ²¹

The graph-robustness literature supports this mapping. Survey work on robust learning on graphs notes that perturbations propagate through relational structure and that robustness should be evaluated with explicit robustness metrics. Knowledge-graph attack papers show that adversarial additions and deletions of triples can substantially degrade link-prediction performance, and recent KG-based robustness frameworks generate adversarial prompts or perturbations from graph triplets specifically to evaluate worst-case behavior. ²²

For a belief layer, the natural stress-test outputs are not only accuracy metrics but also **belief-risk metrics**: threshold crossings, posterior flips, posterior concentration under attack, recovery time after attack removal, cascade depth through neighbors, and blast radius across the graph. Those are the graph analogue of capital depletion, loss amplification, and contagion under CCAR/DFAST-style supervisory scenarios. ²³

Non-stationarity and the belief storm problem

These are structurally different phenomena. **Non-stationarity** means the data-generating mechanism changed: likelihoods, transition structure, dependence, or measurement relation are no longer the same. A **belief storm**, as described by “high evidence churn,” is an inference-layer instability pattern: many local beliefs are revised rapidly because evidence arrives in bursts, conflicts, or propagates through highly connected structure. The latter may occur with or without a true regime change. ²⁴

The distinction matters because robust changepoint methods are designed precisely to avoid conflating outliers or transient disturbances with genuine structural breaks. The robust BOCD literature shows that standard BOCD can create false changepoints under anomalous points, and the broader volatile-environment literature emphasizes that a learner must detect true changes while refraining from overreacting to noise. A belief storm is therefore closer to **high propagation sensitivity, threshold fragility**, or **outlier-induced churn** than to non-stationarity itself. ²⁵

Operationally, the diagnostic split is:

- if posterior instability is explained by a high $P(r_t = 0 \mid x_{1:t})$ or high posterior mass on a new regime across several consecutive observations, the issue is non-stationarity;
- if posterior instability is large but changepoint probability stays low and sensitivity to local perturbations is high, the issue is a belief storm;
- if both occur together, the system is undergoing a real regime break **and** has fragile propagation dynamics. ²⁶

Recommended belief-layer design

A defensible design for a Bayesian comparator layer in a knowledge graph is a **three-level structure**:

- **Global structural layer**: posterior over continuation, changepoint, and regime identity.
- **Regime layer**: archived regime-specific priors/posteriors and transition probabilities.
- **Local comparator layer**: subject-decision-type sufficient statistics for the currently active regime.

27

The update policy should be:

$$p(\text{predict at } t + 1) = \sum_k p(\text{predict} \mid \theta_k, s_{t+1} = k) P(s_{t+1} = k \mid D_{1:t}),$$

with a changepoint action triggered only when posterior expected loss favors the break model. When triggered, reset only the **active local segment** to a regime-entry prior; retain old regimes in memory; and use retrospective smoothing on a replay window to clean up false positives and merge short-lived pseudo-regimes. For gradual drift, add discounting or time-varying transition probabilities rather than repeated hard resets. For adversarial evaluation, stress the graph with coherent multi-signal perturbations and measure threshold crossings, posterior reversals, and recovery dynamics, not only average predictive loss.

28

Open questions and limitations

The cited changepoint, regime-switching, and supervisory-stress literature formalizes structural breaks, regime uncertainty, salience, and robustness, but it does **not** appear to use “belief storm” as a standard technical term. In this report, that term is therefore treated as a systems-level concept: high-churn belief propagation under volatile or adversarial evidence, which is related to but distinct from non-stationarity. The mapping from bank stress testing to knowledge-graph belief evaluation is also a translation across domains rather than a direct result stated in the supervisory documents. 29

1 14 <https://sites.stat.washington.edu/raftery/Research/PDF/bollen1993.pdf>
<https://sites.stat.washington.edu/raftery/Research/PDF/bollen1993.pdf>

2 3 4 8 26 27 <https://lips.cs.princeton.edu/pdfs/adams2007changepoint.pdf>
<https://lips.cs.princeton.edu/pdfs/adams2007changepoint.pdf>

5 19 <https://arxiv.org/abs/1907.02936>
<https://arxiv.org/abs/1907.02936>

6 29 <https://arxiv.org/abs/0710.3742>
<https://arxiv.org/abs/0710.3742>

7 <https://www.rdocumentation.org/packages/bcp/versions/1.6/topics/bcp>
<https://www.rdocumentation.org/packages/bcp/versions/1.6/topics/bcp>

9 25 <https://proceedings.mlr.press/v202/altamirano23a/altamirano23a.pdf>
<https://proceedings.mlr.press/v202/altamirano23a/altamirano23a.pdf>

- ¹⁰ <https://www.sciencedirect.com/science/article/pii/S0304407694900361>
<https://www.sciencedirect.com/science/article/pii/S0304407694900361>
- ¹¹ ¹² <https://www.nuff.ox.ac.uk/economics/papers/2000/w31/msvarfor.pdf>
<https://www.nuff.ox.ac.uk/economics/papers/2000/w31/msvarfor.pdf>
- ¹³ https://research.vu.nl/ws/portalfiles/portal/267687090/Time_Varying_Transition_Probabilities_for_Markov_Regime_Switching_Models.pdf
https://research.vu.nl/ws/portalfiles/portal/267687090/Time_Varying_Transition_Probabilities_for_Markov_Regime_Switching_Models.pdf
- ¹⁵ ¹⁸ ²³ <https://www.federalreserve.gov/supervisionreg/srletters/sr1207.htm>
<https://www.federalreserve.gov/supervisionreg/srletters/sr1207.htm>
- ¹⁶ <https://www.federalreserve.gov/frs/regulations/appendix-a-policy-statement-on-the-scenario-design-framework-for-stress-testing.htm>
<https://www.federalreserve.gov/frs/regulations/appendix-a-policy-statement-on-the-scenario-design-framework-for-stress-testing.htm>
- ¹⁷ ²¹ <https://www.federalreserve.gov/publications/2026-stress-test-scenarios.htm>
<https://www.federalreserve.gov/publications/2026-stress-test-scenarios.htm>
- ²⁰ <https://www.imf.org/external/np/pp/eng/2012/082212.pdf>
<https://www.imf.org/external/np/pp/eng/2012/082212.pdf>
- ²² https://galina0217.github.io/works/robust_survey.pdf
https://galina0217.github.io/works/robust_survey.pdf
- ²⁴ <https://www.sciencedirect.com/science/article/pii/S0022249622000517>
<https://www.sciencedirect.com/science/article/pii/S0022249622000517>
- ²⁸ <https://hbiostat.org/bayes/bet/background>
<https://hbiostat.org/bayes/bet/background>