Doubly Outlier-Robust Online Infinite Hidden Markov Model

Horace Yiu[1,2] Leandro Sánchez-Betancourt[1,2] Álvaro Cartea[1,2] Gerardo Duran-Martin[2]
[1]Mathematical Institute, University of Oxford [2]Oxford-Man Institute of Quantitative Finance

ICML 2026 · Seoul, South Korea · Proceedings of the 43rd ICML (PMLR 306)

Motivation

Online iHMMs are ideal for non-stationary data: they maintain a **reusable regime library** and grow the state space on demand. But standard online iHMMs are **fragile under outliers**:

- Large residuals **corrupt observation posteriors** of the active regime
- The same extreme value **triggers spurious state creation**, fragmenting the learned structure

Illustrative Example

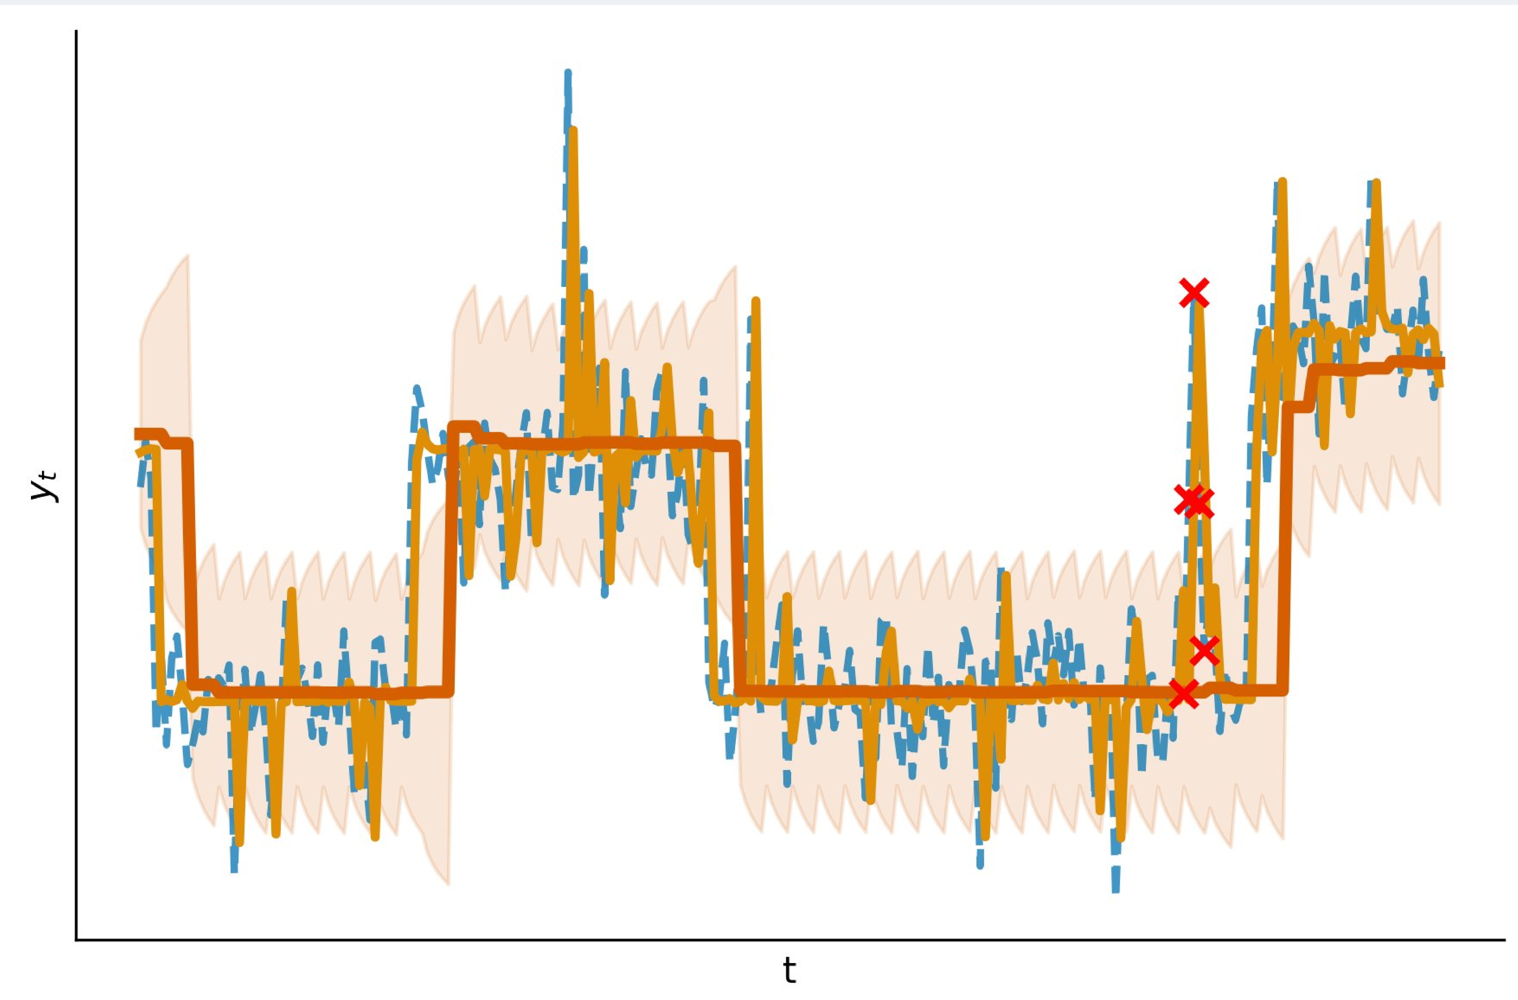

Data (blue) from a 3-state model with Student-t noise; outliers marked ×. **BR-iHMM** (red) tracks the true level and is unaffected. The standard **iHMM** (orange) is deflected and creates spurious regimes.

Formalising Robustness: The PIF

The **Posterior Influence Function (PIF)** is the KL divergence between posteriors computed with and without a contaminated observation y_t^c:

$$\mathrm{PIF}_{s_t}(y_t^c) = \mathrm{KL}\big(P(s_{1:t}\,|\,\mathcal{D}_{1:t})\,\|\,P(s_{1:t}\,|\,\mathcal{D}^c_{1:t})\big)$$
$$\mathrm{PIF}_{\theta_t}(y_t^c) = \mathrm{KL}\big(P(\theta_t\,|\,s_{1:t},\mathcal{D}_{1:t})\,\|\,P(\theta_t\,|\,s_{1:t},\mathcal{D}^c_{1:t})\big)$$

Additive decomposition (chain rule of KL):

$$\mathrm{PIF}_{\theta_t,s_t} = \underbrace{\mathrm{PIF}_{s_t}}_{\text{state space}} + \sum_{s_{1:t}} P(s_{1:t}\,|\,\mathcal{D}_{1:t})\,\underbrace{\mathrm{PIF}_{\theta_{s_t}}}_{\text{obs. space}}$$

A posterior is **outlier-robust** iff its PIF is uniformly bounded.

Theorem 1 — Obs.-Space Robustness Alone is Insufficient

Even with WoLF updates (bounded PIF_{θ_t}), the state-space PIF is **unbounded**: as $\|y_t^c\|_2 \to \infty$, the path posterior concentrates on paths that instantiate a **new state**.

Proof sketch: For large y_t^c, the new state's prior variance $\Sigma(0) \succ \Sigma(n)$ for all $n>0$, making new-state paths infinitely more likely under the posterior — this holds even for Student-t observation models.

State-Space PIF: Batching Bounds It

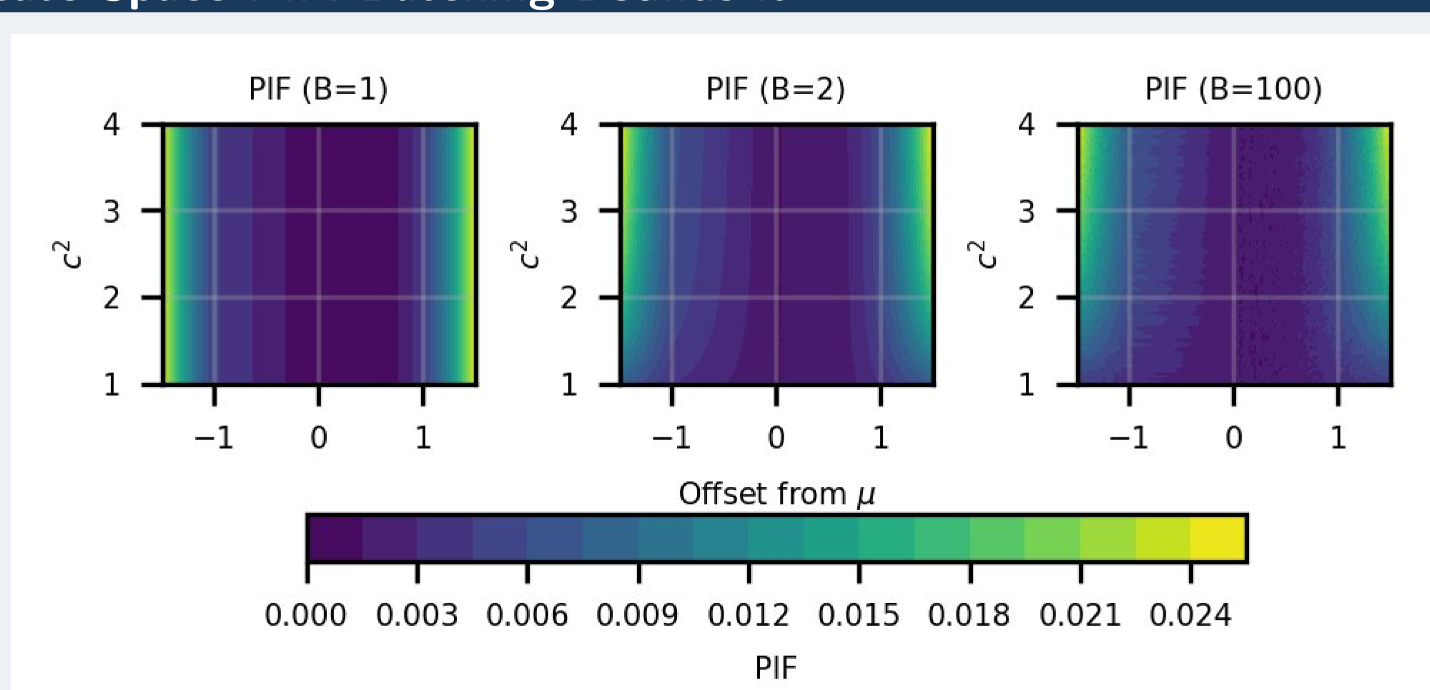

Non-batched PIF ($B=1$) grows unbounded with observation offset. Batched PIFs ($B\geq 2$) are uniformly bounded for all IMQ scales c^2.

BR-iHMM achieves *doubly*-robust online iHMM inference by combining two complementary mechanisms:

WoLF updates — IMQ-weighted generalised Bayes in obs. space → **Batched inference** — Degenerate sticky HDP prior in state space → **BR-iHMM** — **Bounded Joint PIF**

Component 1 — Observation-Space Robustness (WoLF)

Replace the standard likelihood with a **generalised Bayes update**:

$$P(\theta_{1:t}\,|\,s_{1:t},\mathcal{D}_{1:t}) \propto P(y_t\,|\,\theta_{s_t},x_t)^{W(y_t,\hat{y}_t)^2}\prod_k P(\theta_k\,|\,\Psi_{k,t-1})$$

using the **Inverse Multi-Quadric (IMQ)** weight: $W(y,\hat{y})^2 = \left(1 + c^{-2}\|y-\hat{y}\|^2_{R_t}\right)^{-1}$

- Large residuals ⇒ weight ≈ 0; update is damped
- **Preserves conjugacy** ⇒ closed-form Kalman-style updates
- Mahalanobis norm makes the weight **scale-aware** w.r.t. R_t

Component 2 — State-Space Robustness (Batched Inference)

Key idea: pool evidence over $B>1$ observations before permitting any regime switch. Batched log-posterior score for candidate path $s_{t+1:t+B}$:

$$\log \nu(s_{1:t+B}) = \sum_{b=1}^{B} w^2_{s_{t+b},t+b|t}\,\log P(y_{t+b}\,|\,\cdot) + \log P(s_{t+1}\,|\,s_t,\Phi_t)$$

with constraint $s_{t+1} = \cdots = s_{t+B}$ (intra-batch persistence), enforced via the **Degenerate Sticky HDP Prior**:

$$\kappa_t = \begin{cases} 0 & t \equiv 1 \pmod B \text{ (switching allowed)} \\ \infty & \text{otherwise (forced self-transition)} \end{cases}$$

- A single outlier **cannot dominate** the pooled likelihood
- Eliminates rapid-switching and spurious state proliferation
- **Faster**: $O(L_t)$ paths per batch vs $O(L_t^B)$ unconstrained

Theorem 2 — Joint Robustness (Main Result)

Let $W(\cdot,\cdot)$ satisfy the IMQ boundedness conditions. The **batched generalised posterior** with the **degenerate sticky HDP prior** yields:

$$\sup_{y^c_{t+1:t+B}\in\mathbb{R}^d} \mathrm{PIF}^B_{s_{t+B},\theta_{t+B}} < \infty$$

Both the state-space and observation-space PIFs are **simultaneously bounded** for contamination batches of any size B.

BR-iHMM Algorithm (Particle Learning)

1. **Predict** B-step-ahead $\hat{y}^{(i)}_{t+1:t+B}$; compute IMQ weights $w^{(i)}_{\ell,t+b|t}$
2. **Score** WoLF-weighted batch log-likelihood $\nu^{(i)}_{\ell,t+B}$ for each state ℓ
3. **Resample** if $\mathrm{ESS} \leq \tau_{\mathrm{ESS}}$, else normalise
4. **Sample** $\hat{s}^{(i)}_{t+1}$; propagate: $\hat{s}^{(i)}_{t+b} = \hat{s}^{(i)}_{t+1}$ for all b
5. **Update** $\Psi^{(i)}$ via WoLF; update $\Phi^{(i)}$ *once* per batch (count matrix incremented only at inter-batch step)
6. **Advance** $t \leftarrow t+B$

Hyperparameters optimised via Bayesian optimisation on a held-out partition: batch size B, IMQ scale c^2, ESS threshold τ_{ESS}.

Experiments

Four datasets:

- **Synthetic**: 2 500 obs., $d=100$ covariates, 3 regimes, 1% outliers $\sim \mathcal{U}[-600,600]$
- **Electricity**: 31 912 hourly demand obs. (2017–2020), COVID-19 structural break
- **OFI**: MSFT limit-order-book, March 2020 (8 273 obs.)
- **Well-log**: 4 050 NMR measurements; outliers retained

One-Step-Ahead RMSE (100 runs)

Model	Synthetic	Electricity	OFI
BR-iHMM	**46.1**$_{\pm.003}$	**0.47**$_{\pm.04}$	**0.616**$_{\pm.08}$
WoLF-iHMM	103.8$_{\pm.012}$	0.63$_{\pm.03}$	0.623$_{\pm.09}$
iHMM	101.7$_{\pm.026}$	0.57$_{\pm.03}$	0.620$_{\pm.08}$
offline-iHMM	2.9	0.32	0.552
BOCD	123.1$_{\pm.014}$	0.80$_{\pm.11}$	0.733

BR-iHMM reduces online forecasting error by up to **55%** over iHMM on synthetic and **41%** over BOCD on electricity.

Synthetic Task — Regime Recovery

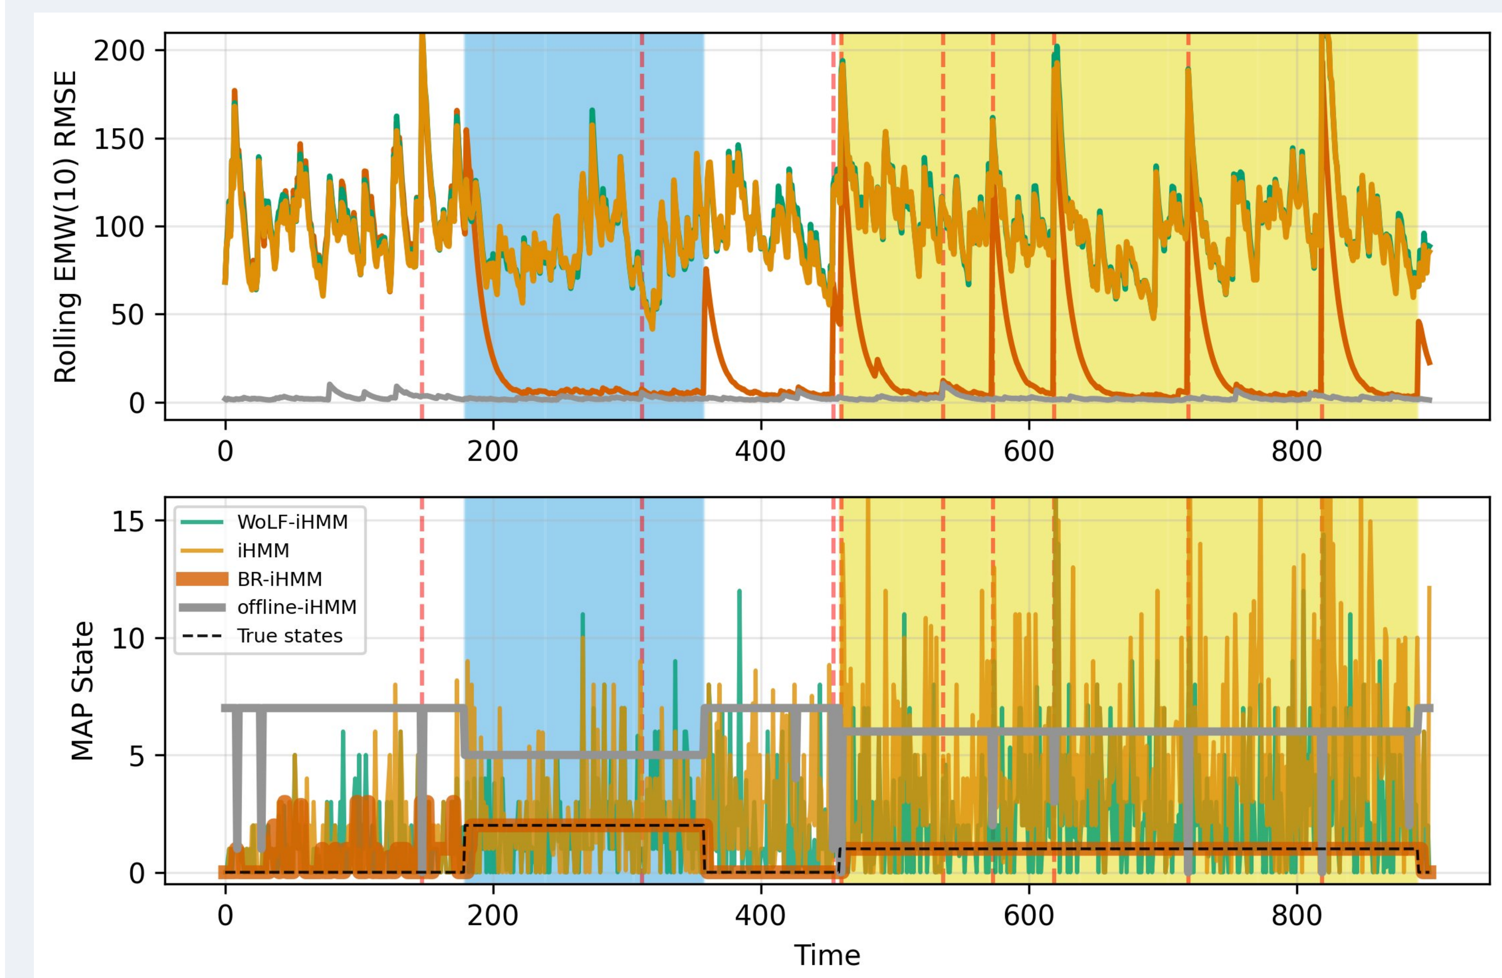

BR-iHMM recovers the 3 true regimes with near-zero RMSE (bottom: MAP state hugs true dashed line). iHMM and WoLF-iHMM create >30 spurious states driven by outliers (red dashes). Yellow/blue = true regimes 2 and 3.

Regime Detection (Synthetic, 100 runs)

Model	PPV	TPR	Delay
BR-iHMM	**0.963**$_{\pm.001}$	**0.996**$_{\pm.001}$	3.4
WoLF-iHMM	0.533$_{\pm.002}$	0.344$_{\pm.002}$	**0.0**
iHMM	0.551$_{\pm.002}$	0.363$_{\pm.008}$	0.2
offline-iHMM	0.996	0.997	0.0

Near-offline precision/recall at the cost of a small, bounded detection delay ($\leq B$).

Segmentation — Well-Log Benchmark

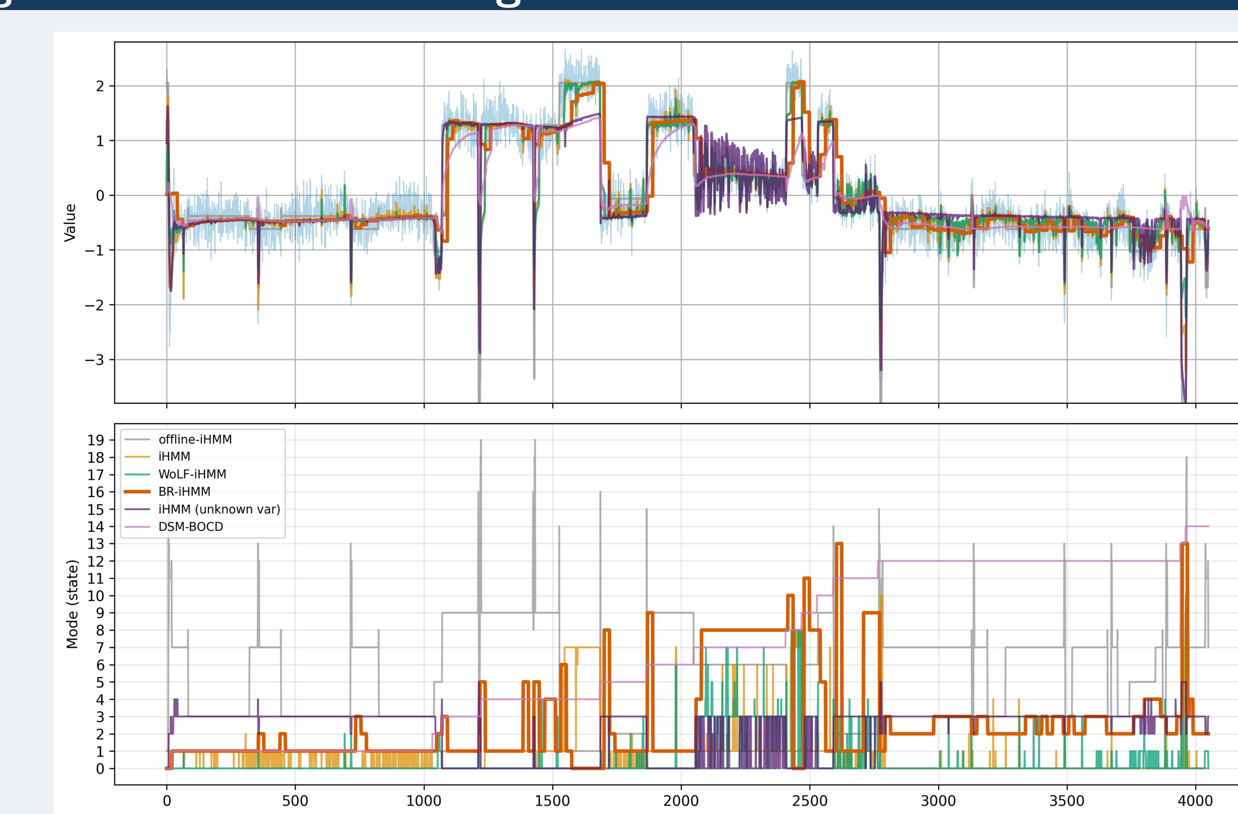

BR-iHMM resists burst-window outliers (obs. 1000–1500) and exploits reusable states. Even iHMM (unknown-var) with Student-t predictive creates spurious states — consistent with Theorem 1.

Key Contributions

1. **First** provably doubly-robust online iHMM — simultaneous robustness in obs. and state spaces
2. Obs.-space robustness alone is **provably insufficient** (Theorem 1 covers Gaussian *and* Student-t models)
3. Batched degenerate sticky HDP bounds the **joint PIF** for arbitrary contamination batches (Theorem 2)
4. **Computationally cheaper** than standard iHMMs; degrades gracefully to WoLF-iHMM when $B=1$ is optimal
5. Up to **67% RMSE reduction** vs competing online Bayesian methods across finance, energy, and geophysics