Decoding Employee Turnover: From Prediction to Causal Intervention

Applying State-of-the-Art Machine Learning and Structural Causal Models to European HR Analytics

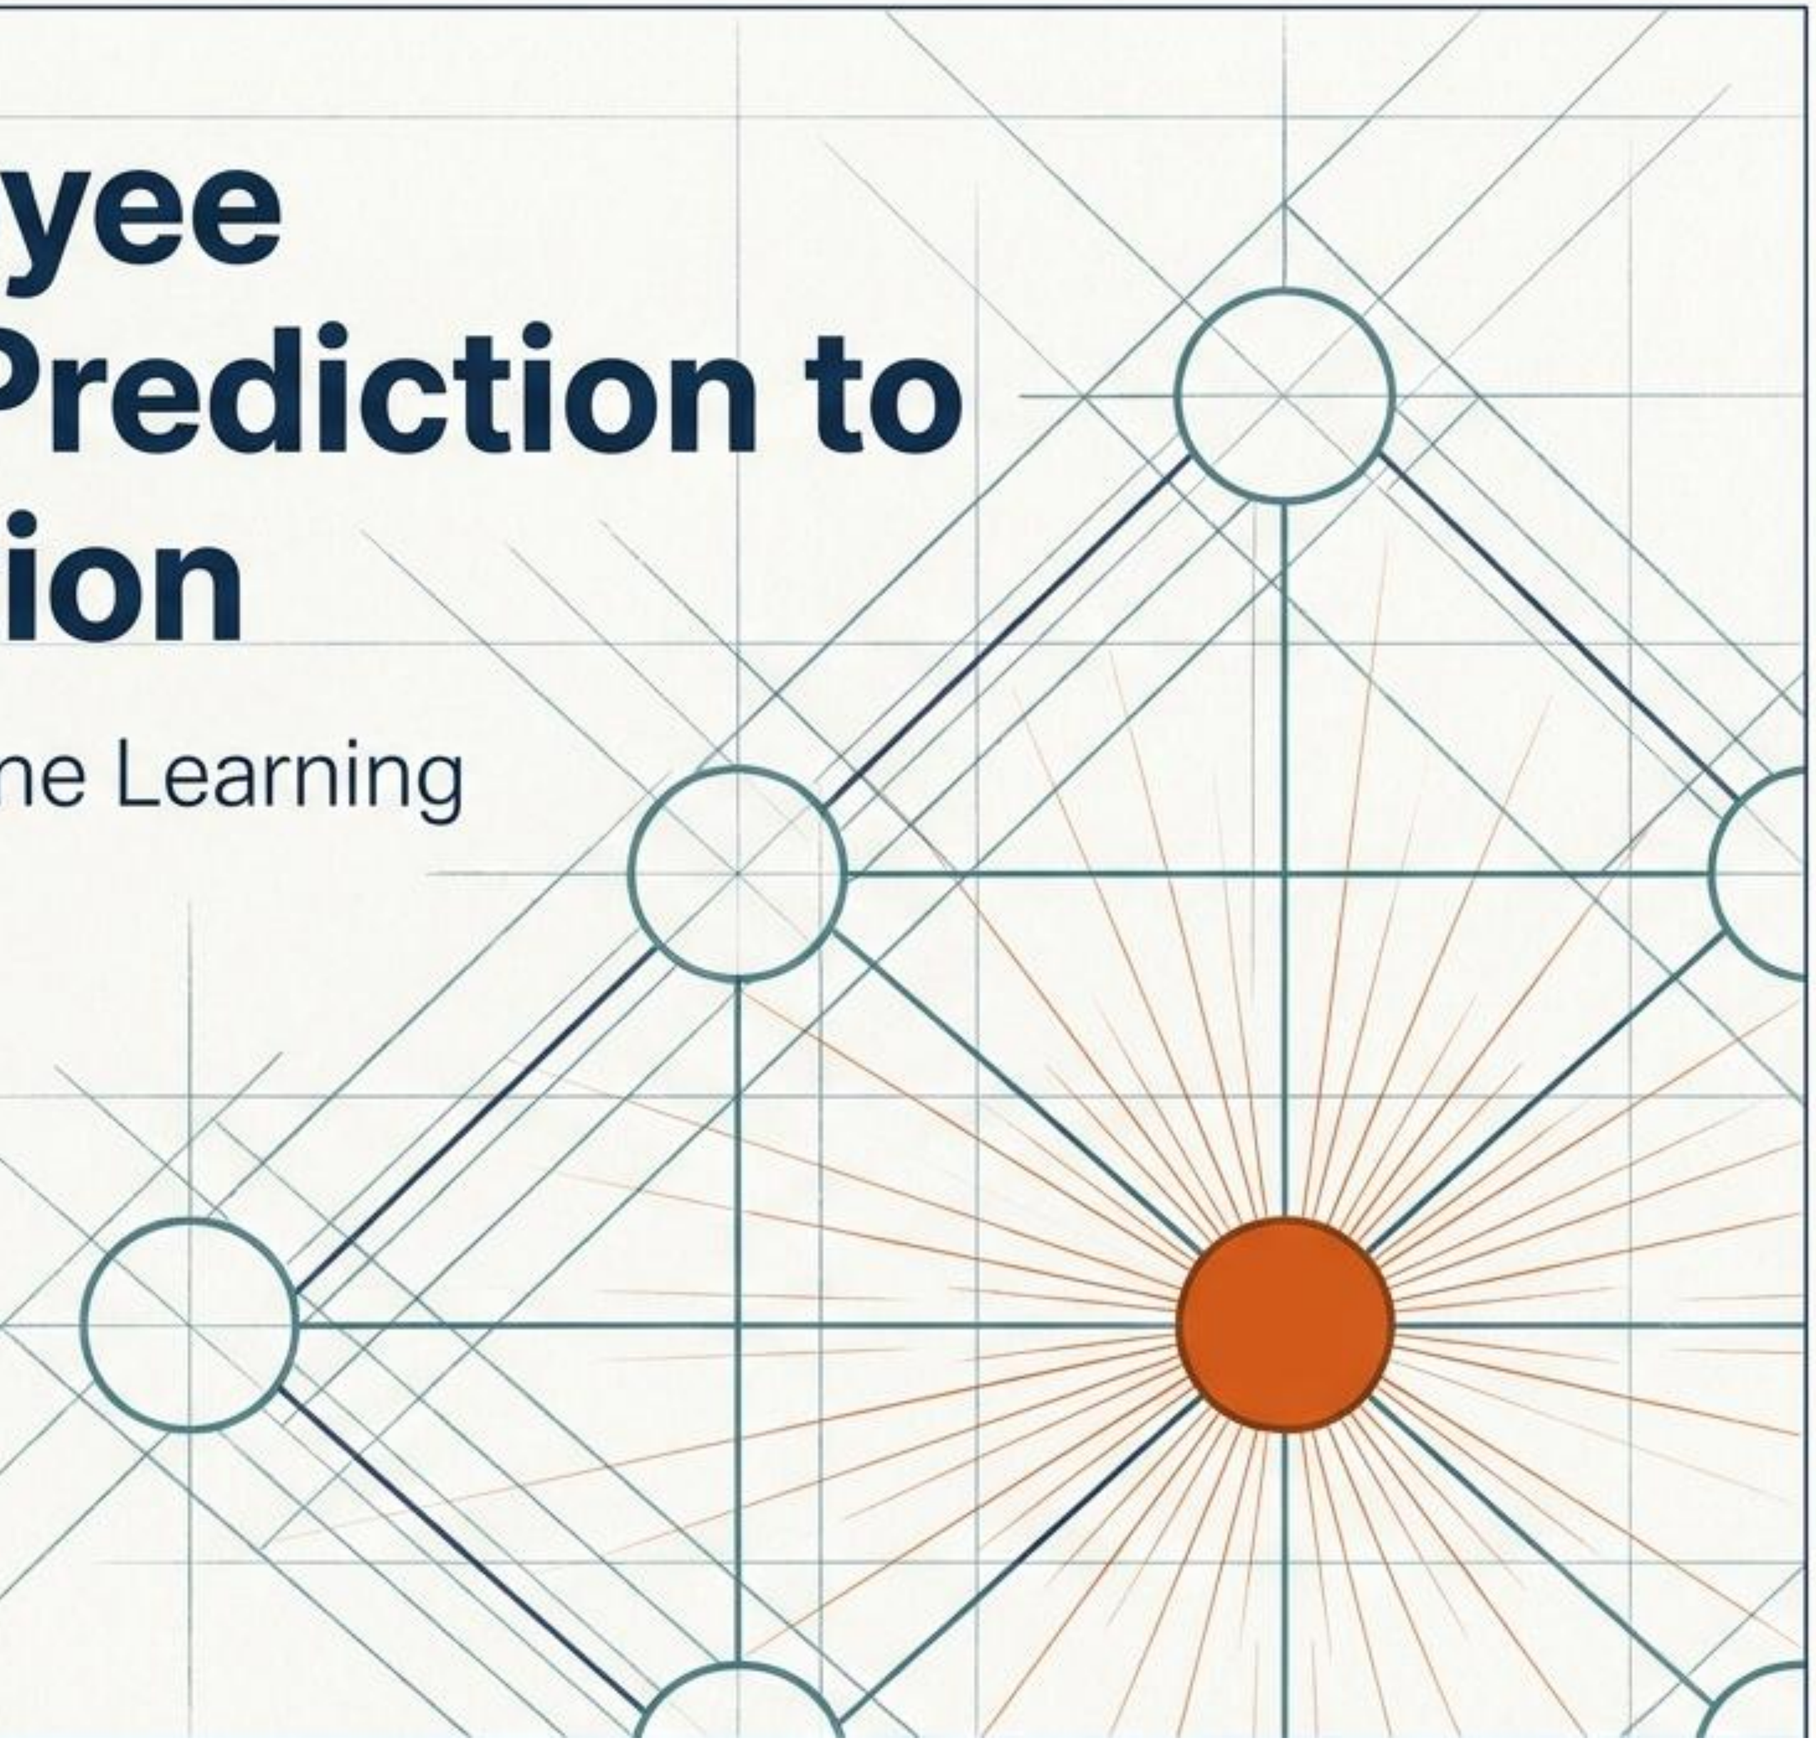

Executive Summary

The Dataset

- 9,296 localized records analyzed from 18,322 surveyed employees.
- Spanning 30 European countries, capturing 23 HR themes across 112 specific items.

The Models

- Evaluated 7 state-of-the-art Machine Learning classifiers.
- LightGBM and Logistic Regression emerged as undisputed champions for predicting departure risk.

The Insight

- Correlation is a trap. Just because an HR metric predicts turnover does not mean changing it will prevent it.
- Causal analysis proves that ‘Motivation’ and ‘Employability’ are true retention levers, while ‘Adaptability’ yields zero causal benefit.

Tracking actual turnover is an autopsy.
Predicting intention is a diagnostic warning system.

Employee turnover intention—a reported willinness to leave within a defined period (e.g., three three months)—is the single most accurate predictor of actual voluntary dysfunctional turnover. Accurately modeling this node opens a brief, highly actionable window for targeted policy interventions to retain top-performing talent.

The Data Footprint: The Global Employee Engagement Index (GEEI)

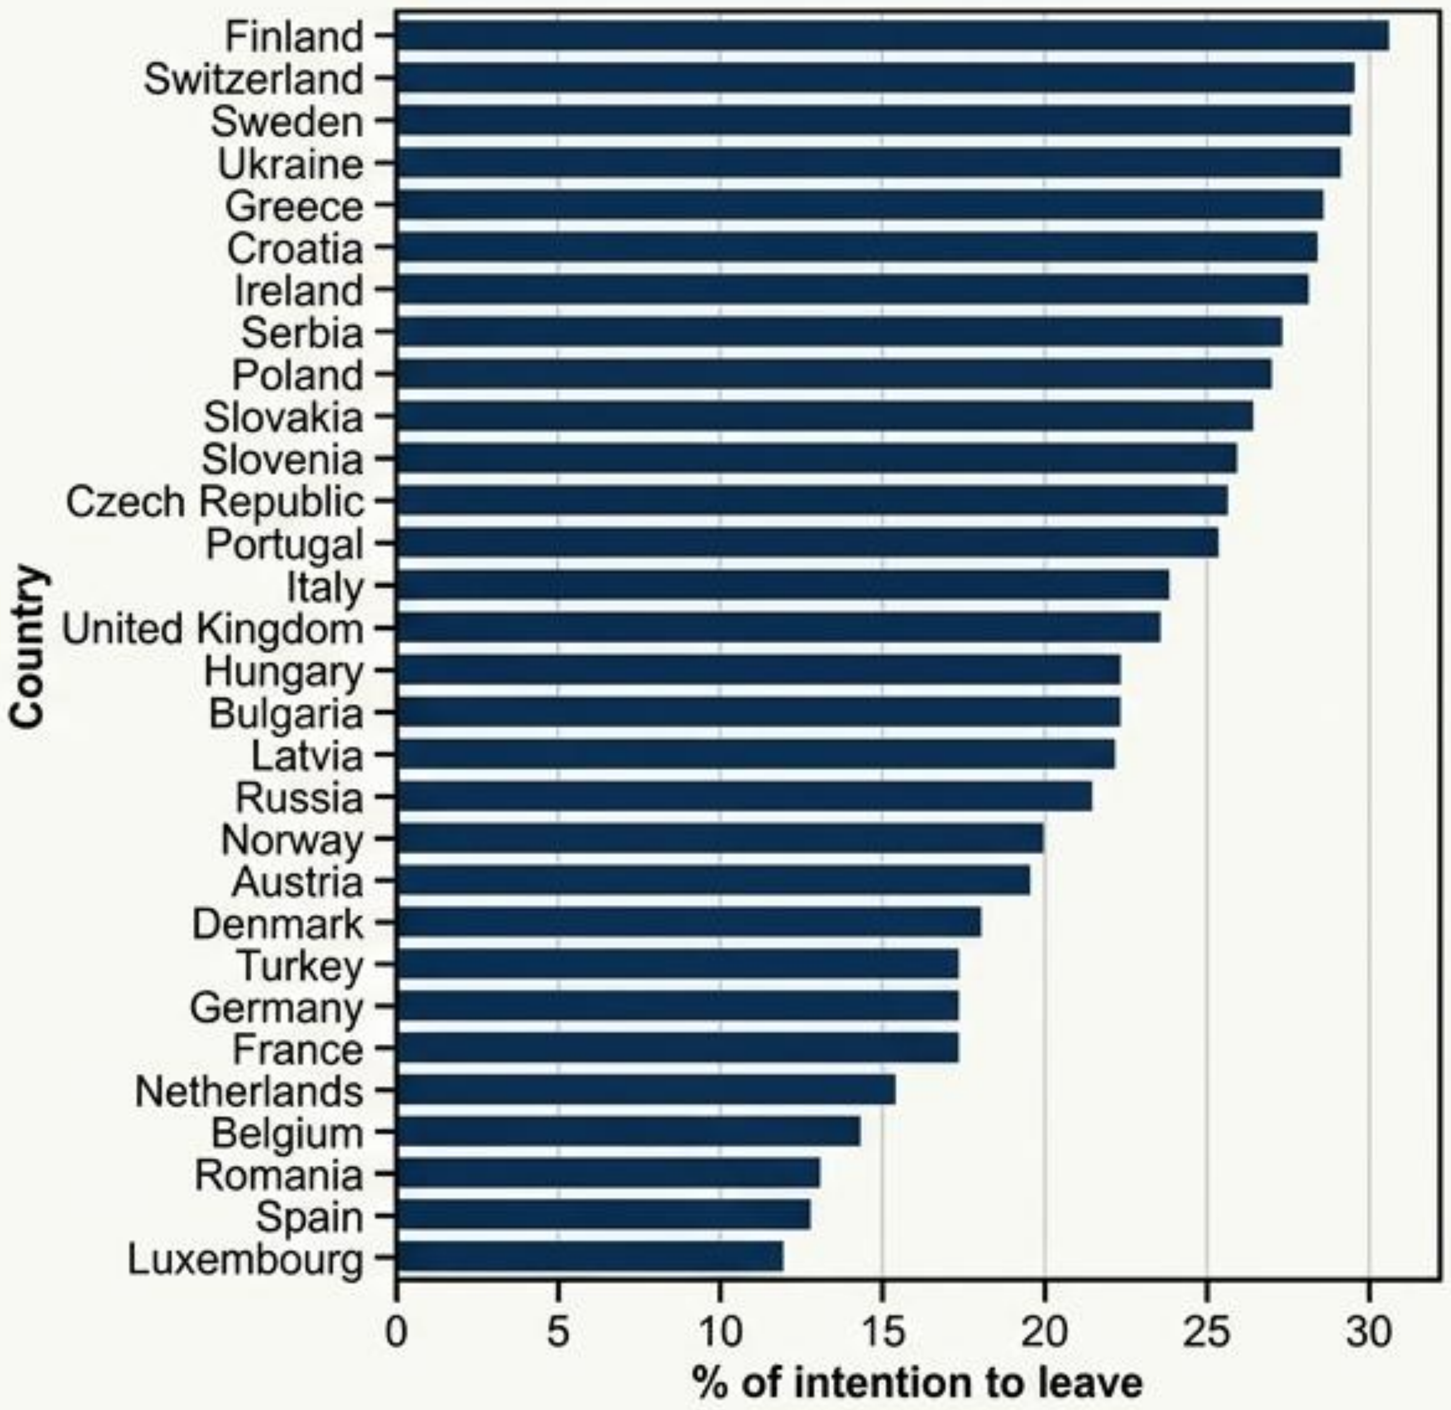

The Diagnostic Blueprint: A Three-Step Process for Employee Retention

Step 1: Predict (The Algorithm Shootout)

Objective: Find the most accurate ML model to detect employees at risk of leaving based on complex survey matrices.

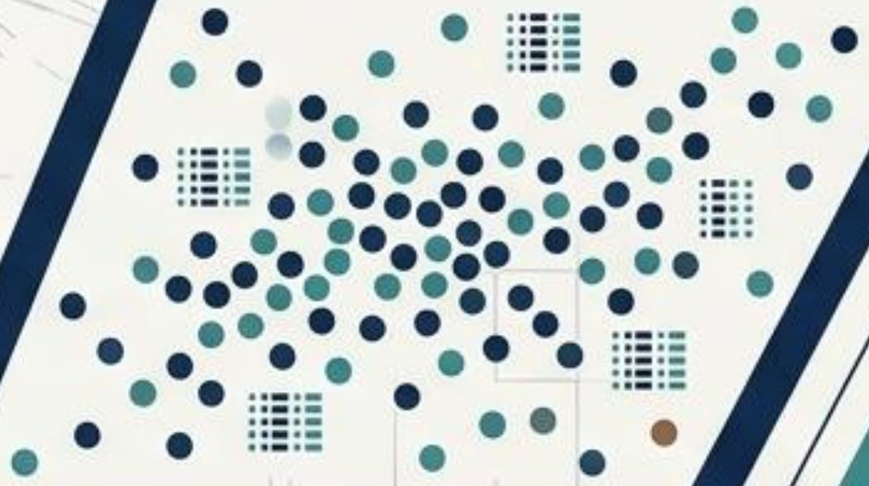

Step 2: Explain (Feature Ranking)

Objective: Utilize Explainable AI (XAI) to extract the most heavily weighted driving factors behind the algorithm's predictions.

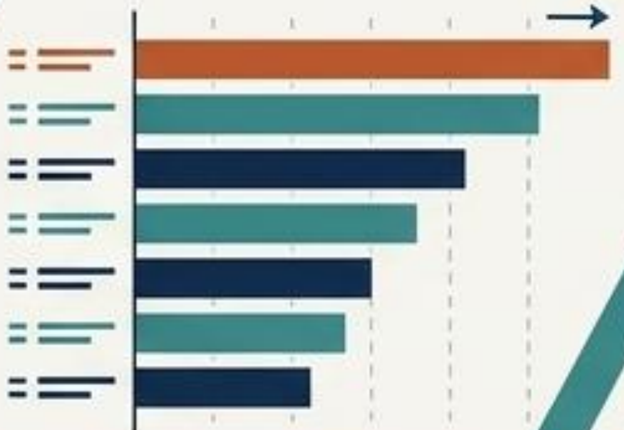

Step 3: Intervene (Proving Causality)

Objective: Move beyond correlation traps. Apply Structural Causal Models (SCM) to mathematically prove which HR policies actually yield retention ROI.

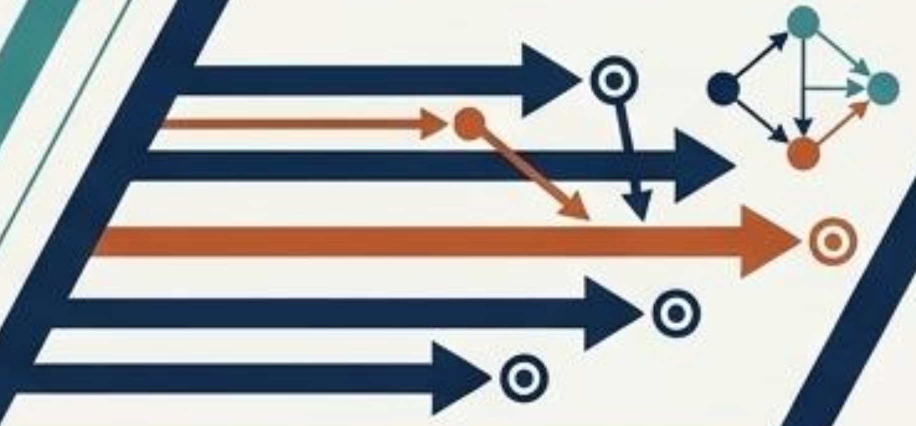

Model Performance Evaluation: The Algorithm Shootout

Model	AUC-PR Score (Performance)	Speed (Computation Time)	Interpretability (Type)
Logistic Regression (LR) ★	0.635	13.6s	White-box
K-Nearest Neighbors (KNN)	0.513	55.0s	White-box
Decision Trees (DT)	0.538	23.3s	White-box
Random Forests (RF)	0.613	64.9s	Black-box
XGBoost (XGB)	0.614	49.6s	Black-box
LightGBM (LGBM) ★	0.641	35.1s	Black-box
TabNet	0.561	7489s	Black-box

METHODOLOGY NOTE: Models were tested using repeated stratified 10-fold cross-validation to guarantee generalization. Deep learning tabular models (TabNet) proved far too slow (7489s) without offering performance gains, validating the selection of tree-based and regression models.

White-Box vs. Black Box Convergence

LightGBM (Black-box)

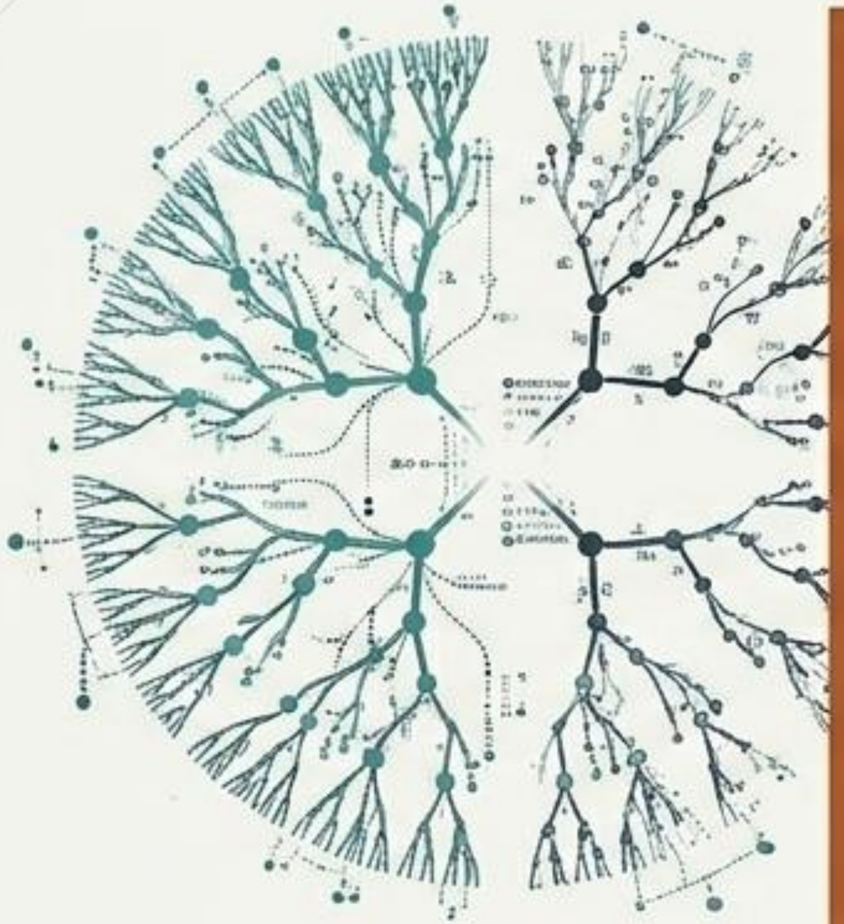

Because a highly complex, non-linear model (LGBM) and a straightforward, linear model (LR) converged on nearly identical high-performance predictions, we can confirm the turnover signals in the data are extraordinarily robust, not algorithmic artifacts.

Captures deep non-linear behavioral relationships.

Logistic Regression (White-box)

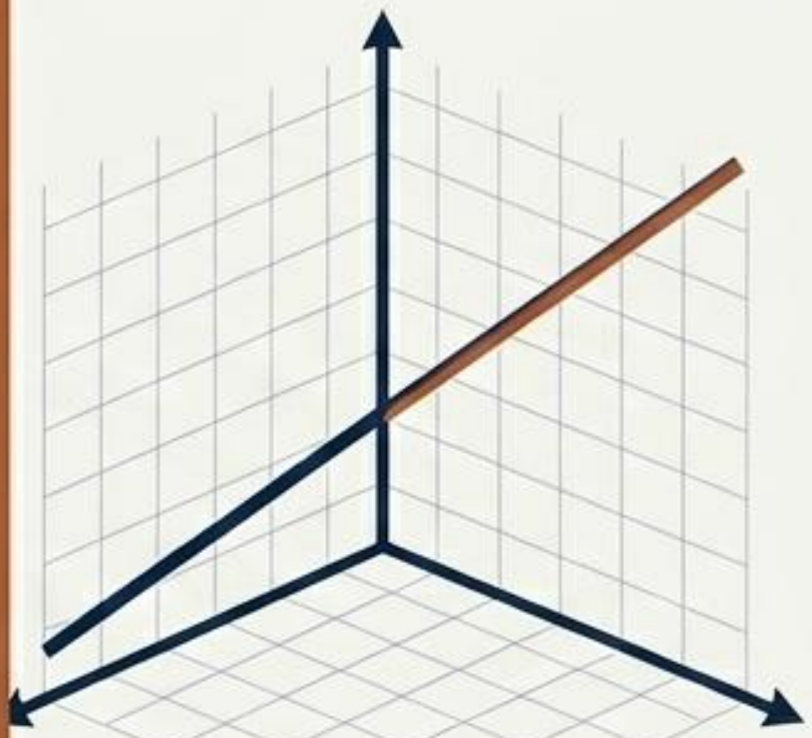

Delivers speed, interpretability, and baseline stability.

Key Turnover Drivers: Macro Themes vs. Micro Items

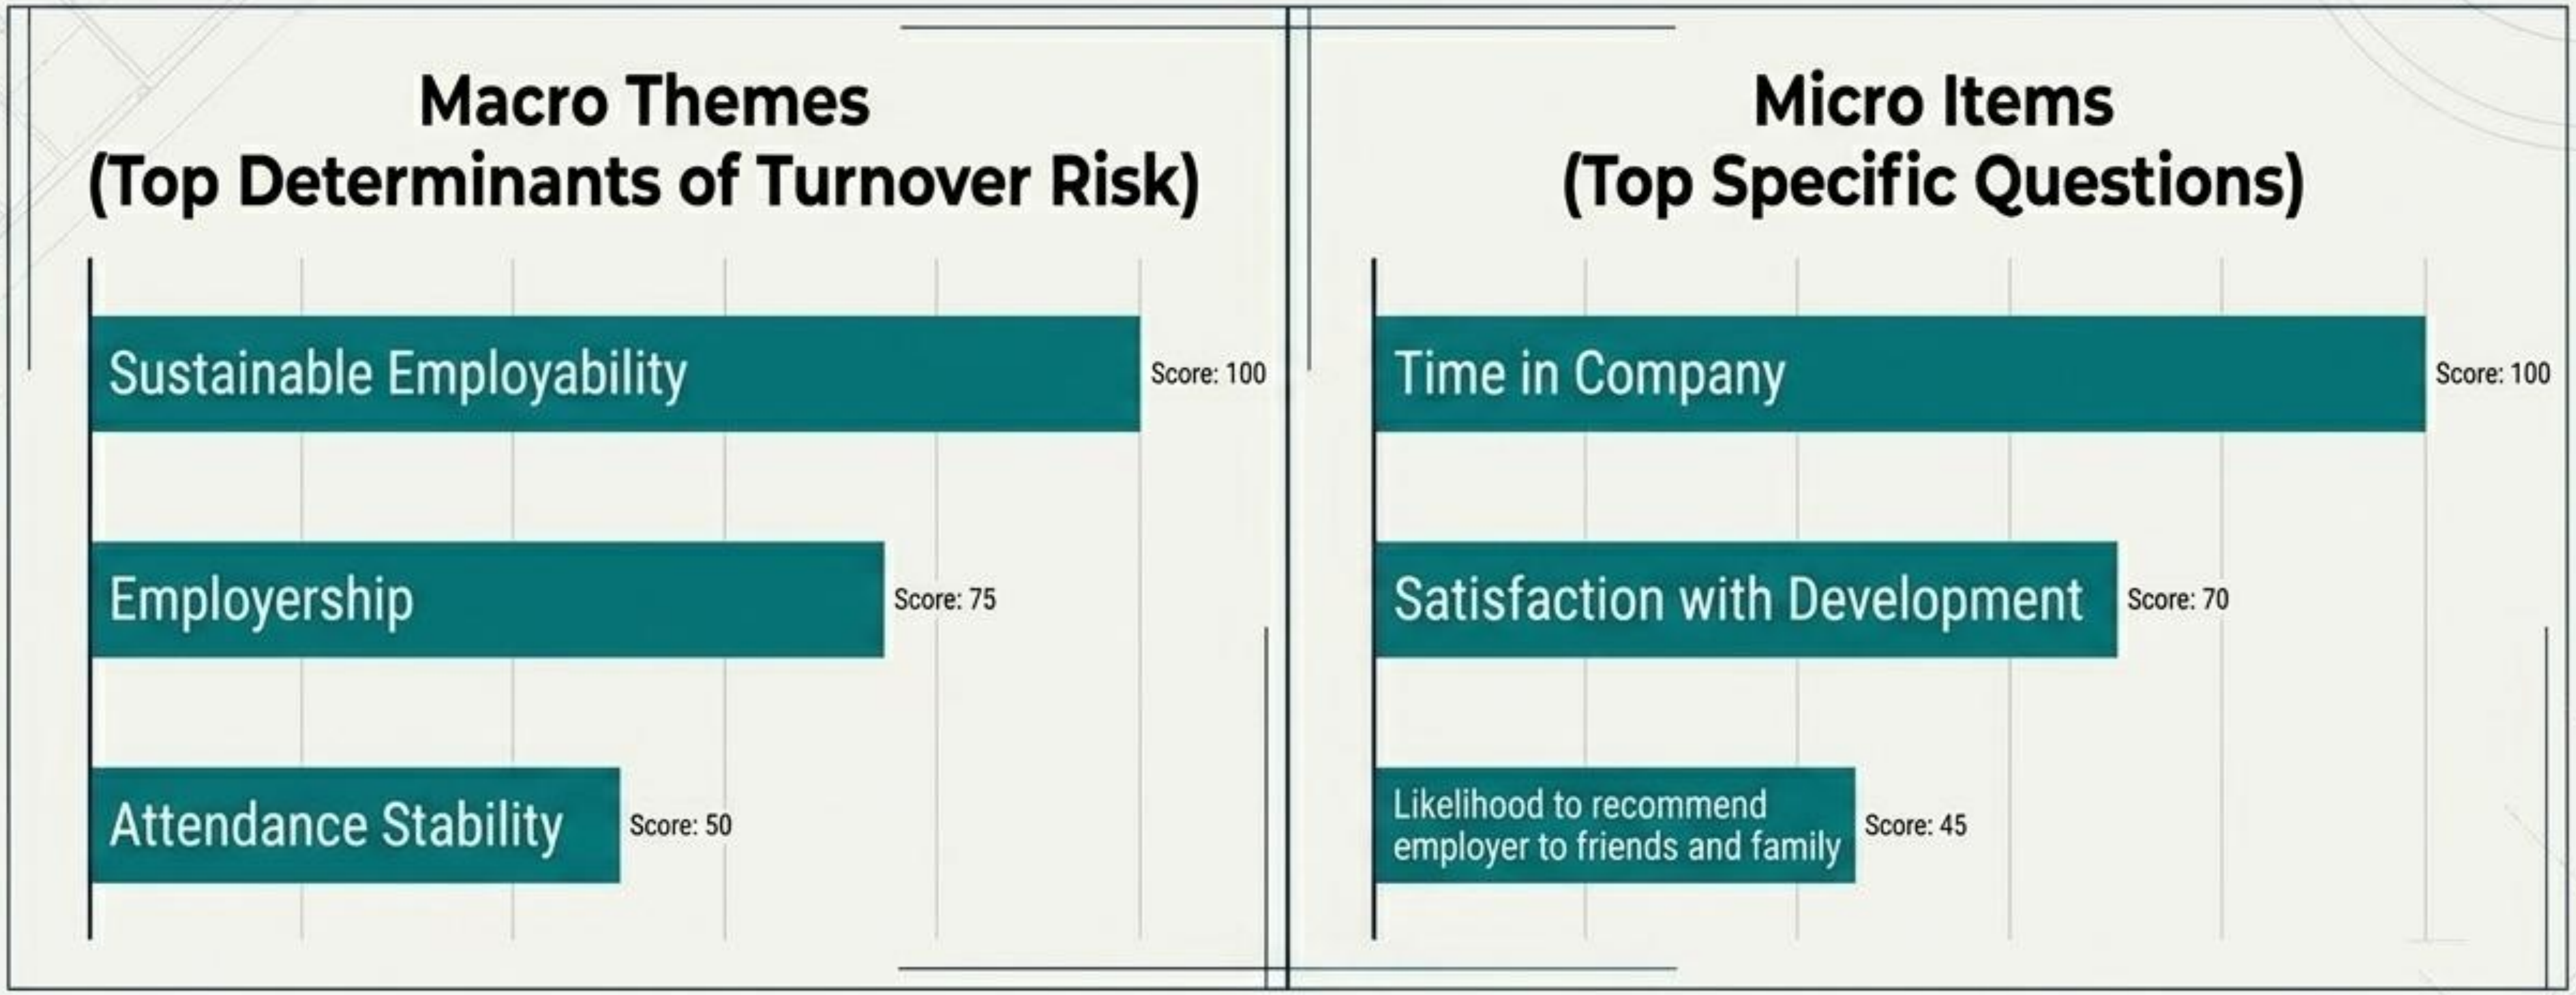

Feature importance was aggregated via a novel procedure ranking determinants across 100 experimental folds, ensuring massive resilience against data anomalies.

The Missing Variables: What the Algorithm Ignored

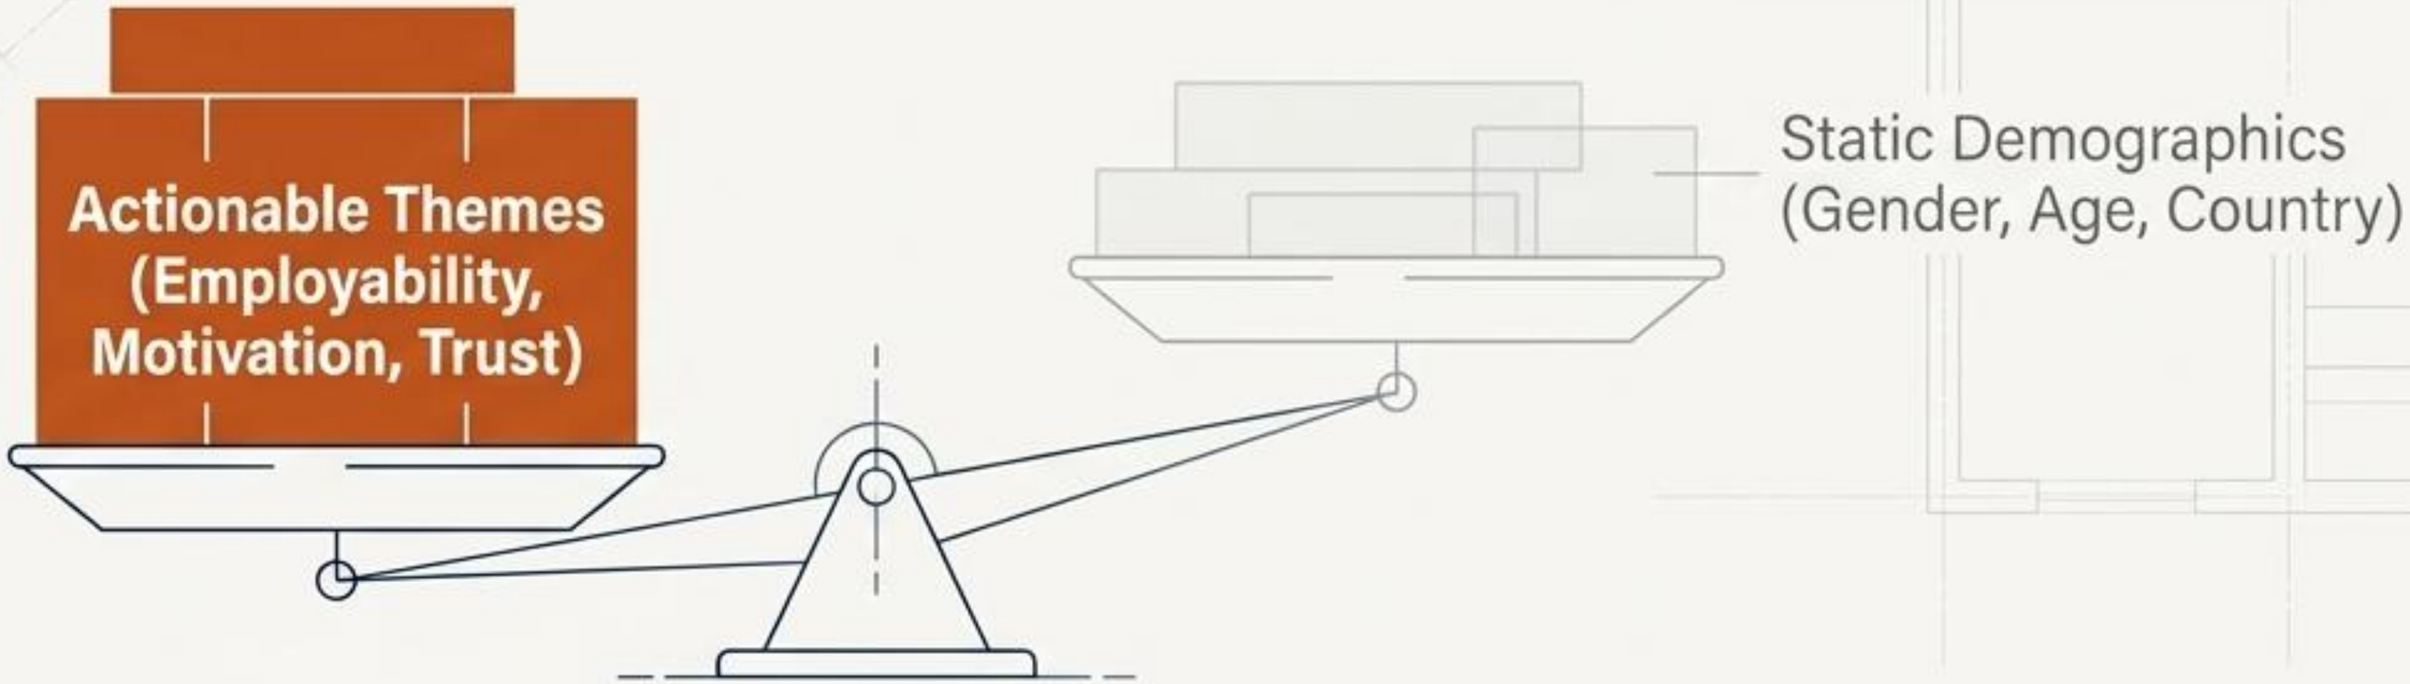

The Expectation	The Reality
Traditional HR literature assumes demographic markers (e.g., Gender, Age bands) or specific geographic location (Country) are the primary drivers of turnover intention.	When controlled properly within top-tier ML models, Gender failed to rank among the top features entirely. No isolated Country-specific effect emerged as a primary predictive driver.

The Correlation vs. Causation Trap

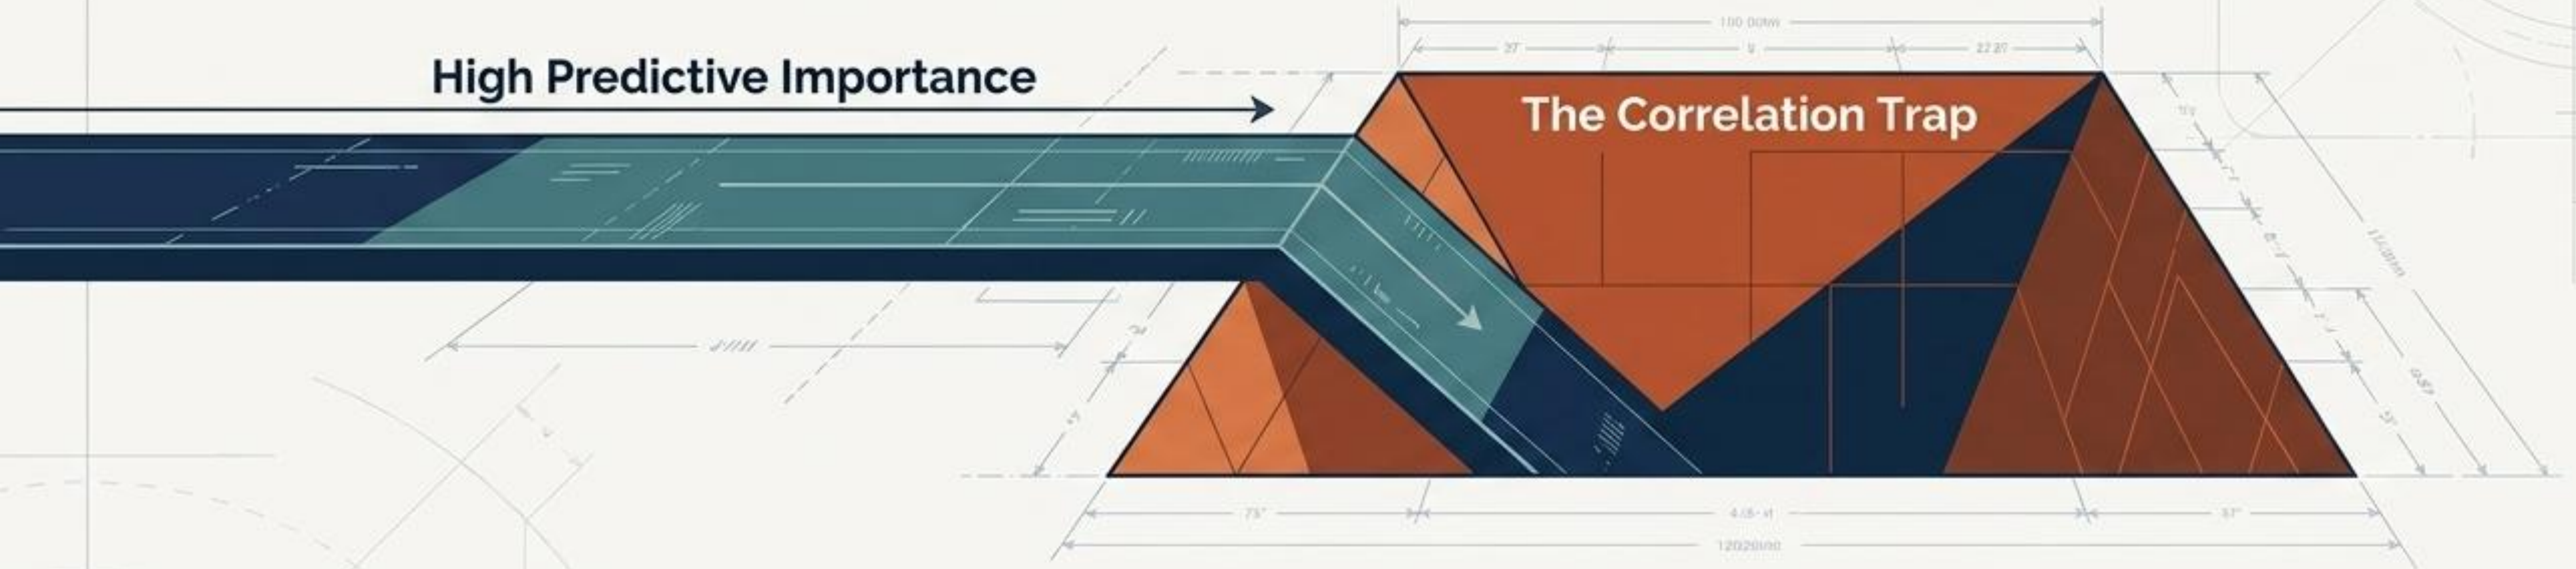

Just because an HR theme is critical to the algorithm's prediction doesn't mean intervening on that theme will cause a change in employee retention.

Analogy Callout

Ice cream sales perfectly predict sunburns (high correlation). But banning ice cream will not prevent sunburns (zero causation).

Resolution

To allocate HR budgets effectively, we must mathematically isolate the pure impact of an HR intervention on turnover intention. We require

Structural Causal Models (SCM).

The Causal Confounder Map

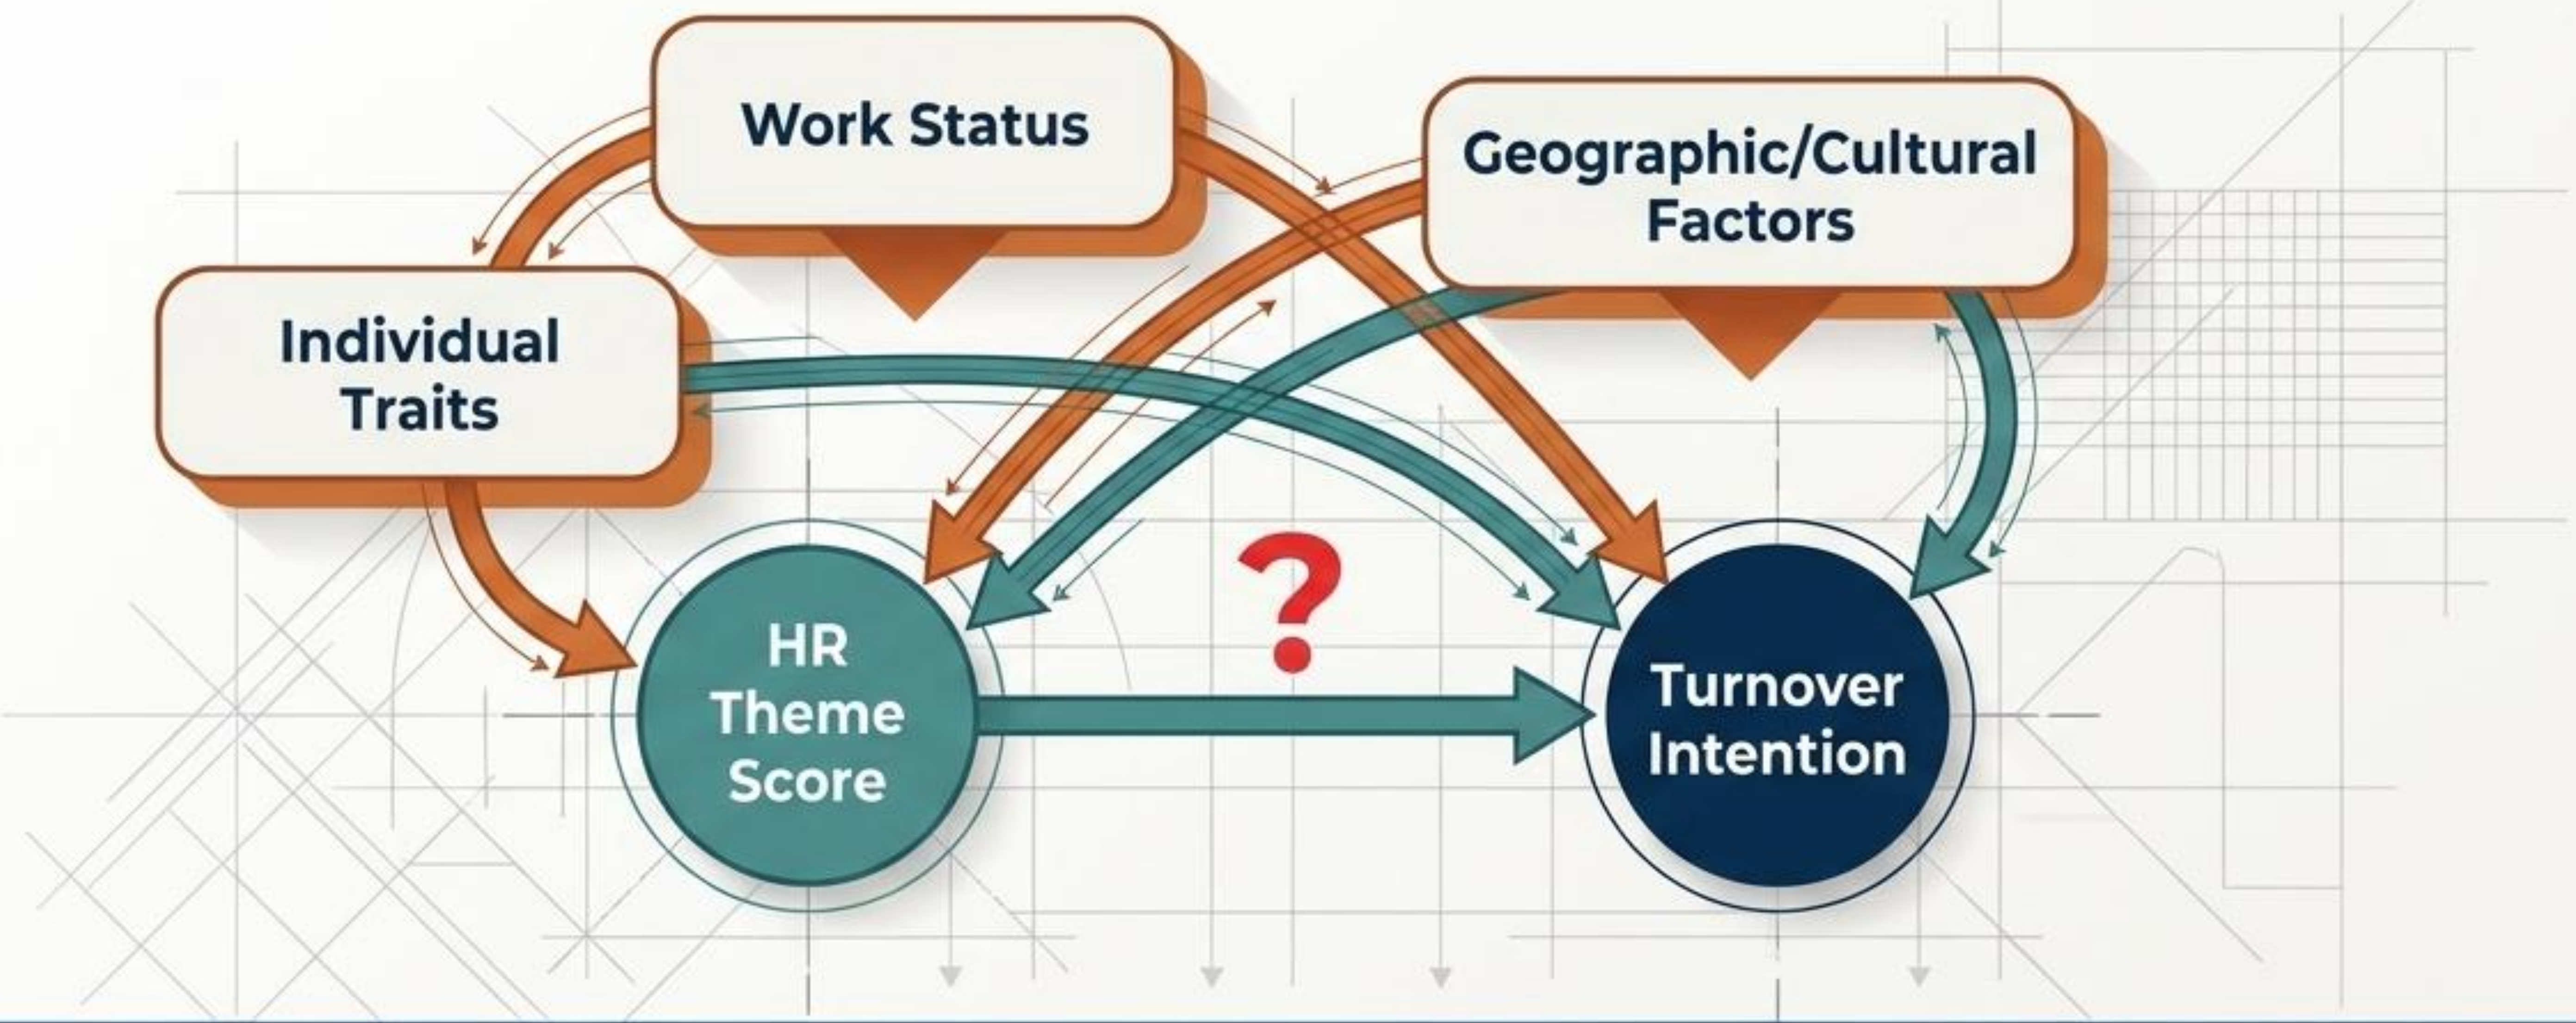

Hidden variables (Geography, Individual traits, Work status) secretly influence both how an employee scores a survey theme and their actual intention to leave. This muddies the waters, making it impossible to see the true impact of the HR Theme with standard analytics.

Intervention Impact Curves: Motivation vs. Adaptability

High Causal Impact (Motivation)

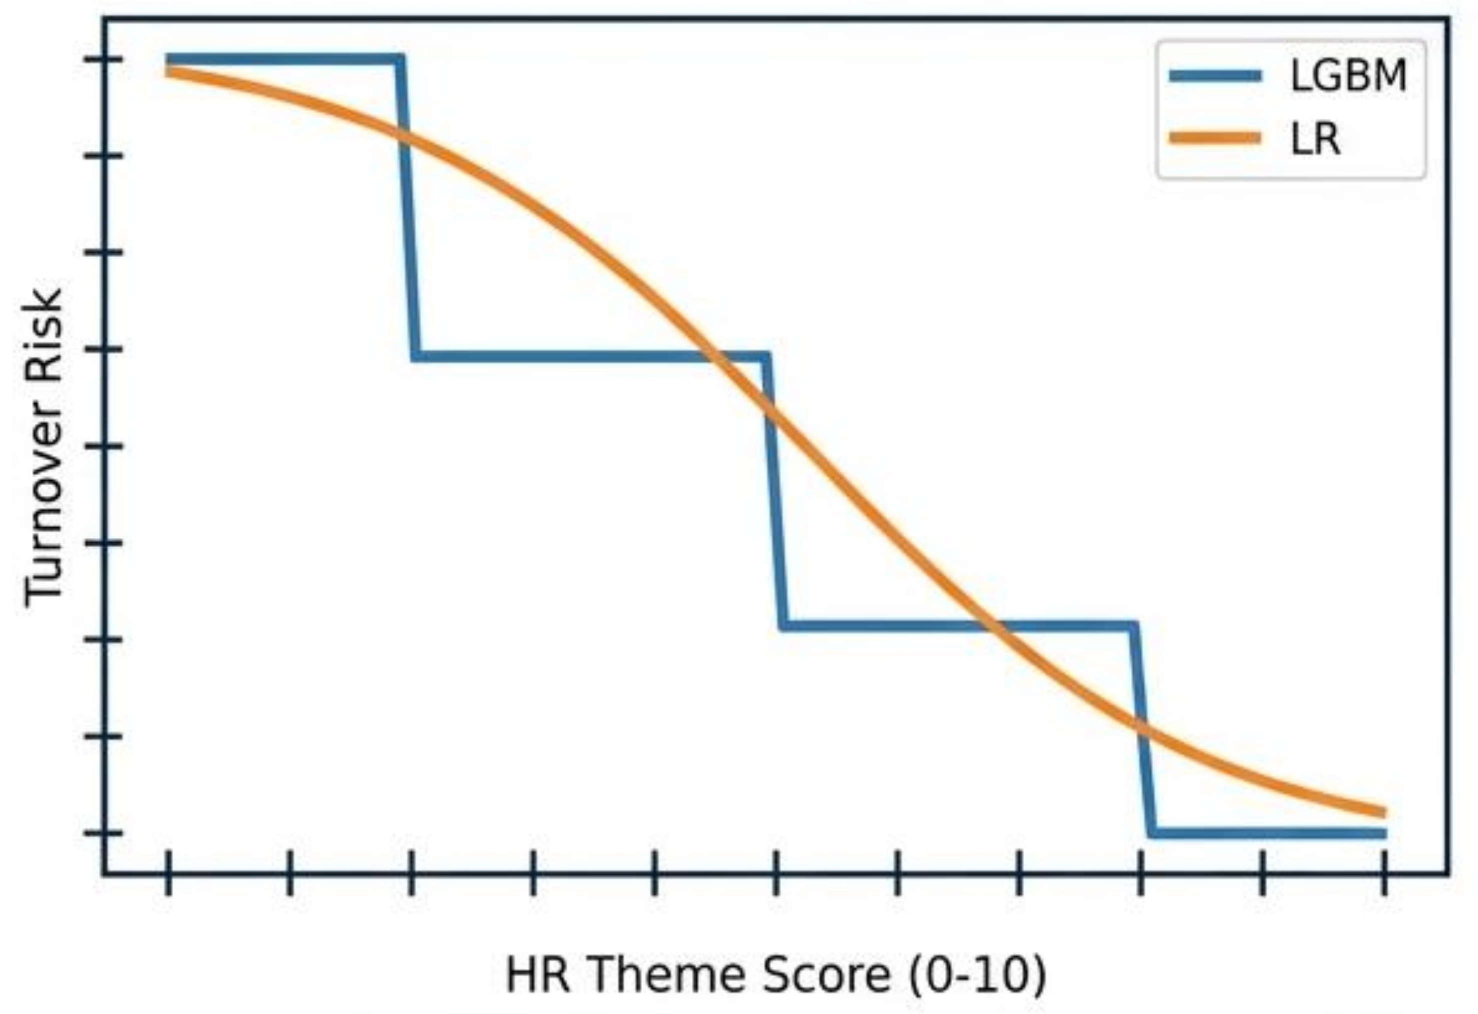

As motivation scores increase, the predicted probability of turnover demonstrably drops.

Zero Causal Impact (Adaptability)

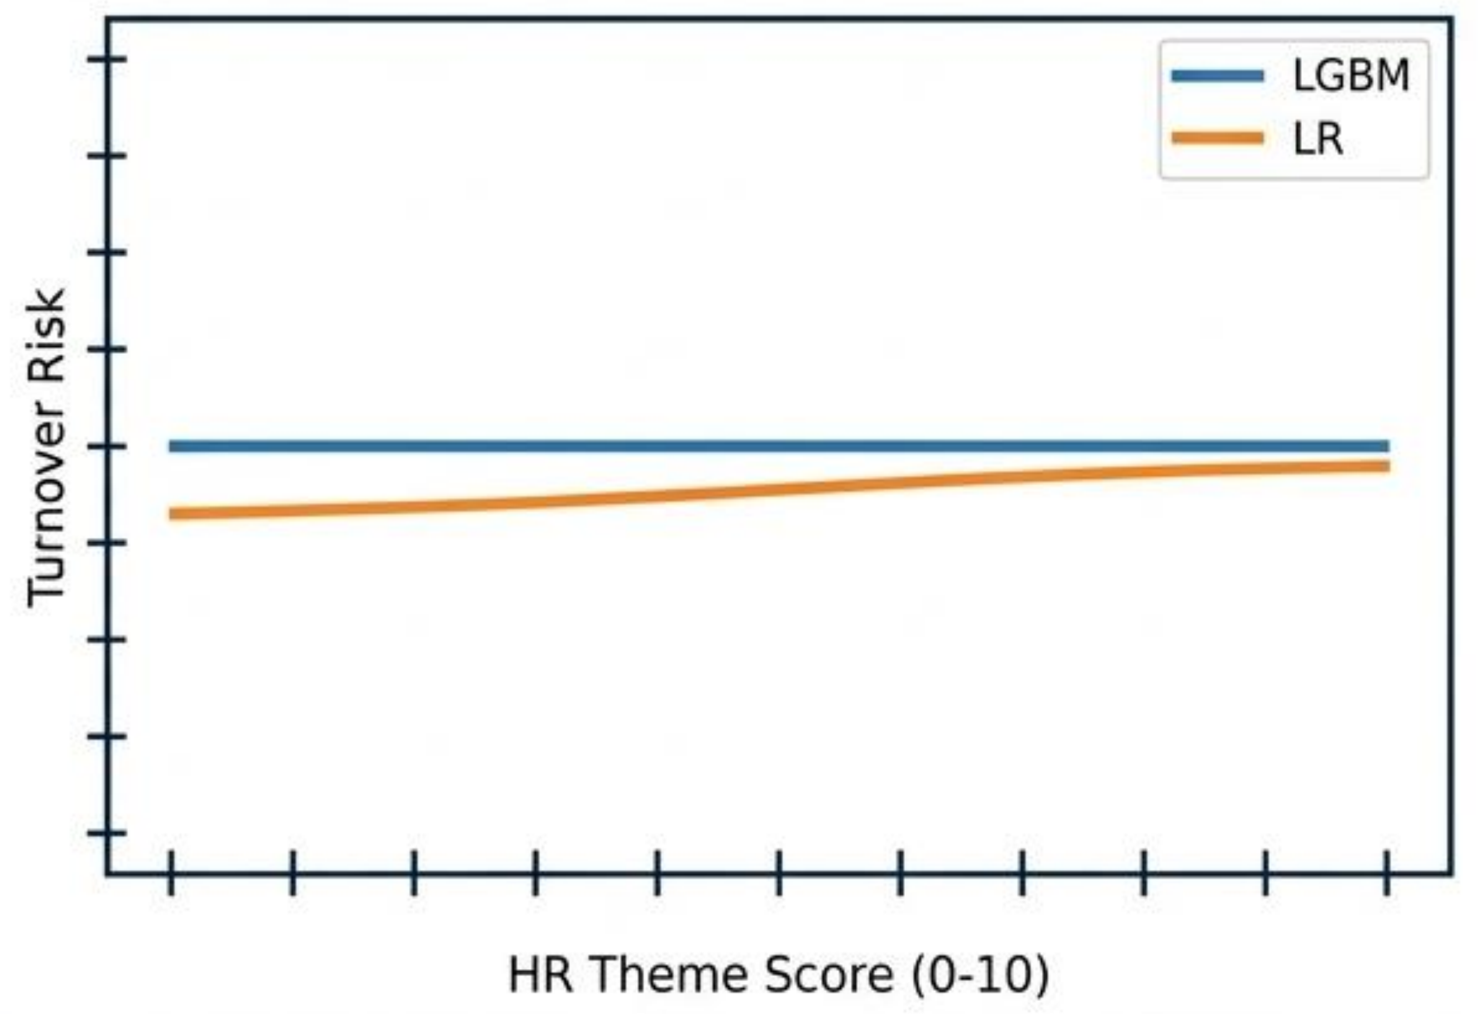

Despite correlation in raw data, forcing ‘adaptability’ initiatives yields absolutely zero retention benefit.

Visual proof that investing in Motivation directly prevents turnover, while investing in Adaptability wastes resources

Strategic Mandates for Organizational Retention

1	**Shift the Target** Stop conducting autopsies on actual turnover. Deploy analytics to accurately capture Turnover Intention, opening the critical Intervention Window before high-performers exit.
2	**Upgrade the Analytics** Abandon standard HR correlation dashboards. Utilize combined Black-Box (LightGBM) and White-Box (LR) modeling paired with Structural Causal Models to eliminate demographic noise and confounder bias.
3	**Focus the Budget** Isolate policy investments strictly to high-causality themes. Prioritize resources toward Sustainable Employability and Motivation, while defunding interventions targeting non-causal variables like Adaptability.