https://doi.org/10.25163/ai.2110974 Journal of Ai ML DL | Published online October 28, 2026

Business Intelligence and Explainable Machine Learning for Capacity-Constrained Cyber Disruption Prediction: A Simulation-Based Decision Framework for U.S. Enterprises

Masrul Hasan[1]*, Md. Rezaul Hasan[2], Kazi Rakib Hasan Saurav[3], Md. Arifur Rahman[4], Md. Fazle Alahi Bhuiyan[5]

Please cite this article as: Hasan, M., Hasan, M. R., Saurav, K. R. H., Rahman, &. A., Bhuiyan, M. F. A. (2026). "Business Intelligence and Explainable Machine Learning for Capacity-Constrained Cyber Disruption Prediction: A Simulation-Based Decision Framework for U.S. Enterprises", Journal of Ai ML DL, 2(1), 1-14, 10974

Supplementary Information

S1 1 Units and random-number construction

The generator uses numpy.random.default_rng(20261028). There are 600 enterprises, with sixty assigned to each of ten sector labels, and 54 generated weekly states per enterprise. Fifty-two weekly forecast snapshots per enterprise yield 31,200 rows. The last two generated weeks are used only to mature the final forecast horizons. Enterprise identifiers are integers 0–599. Sector labels do not represent an economic sampling frame.

All components introduced below are independently drawn unless their dependence is explicitly specified. Sector log-odds effects a_s follow Normal$(0,0.15^2)$, and enterprise log-odds effects b_i follow Normal$(0,0.30^2)$. For predictor construction, E_{ij} and H_i are standard normals that remain constant over time. The dynamic state D_{i1j} is standard normal and $D_{itj} = 0.65D_{i,t-1,j} + \sqrt{(1-0.65^2)}\varepsilon_{itj}$, with standard-normal independent innovations. The Gaussian copula state is $Z_{itj} = 0.45E_{ij} + 0.35H_i + \sqrt{0.675}D_{itj}$. Thus each Z has unit marginal variance, shared H induces cross-feature dependence, and D induces temporal persistence.

The copula probability is $u_{itj} = \mathrm{clip}[\Phi(Z_{itj}), 10^{-10}, 1-10^{-10}]$. The physical predictor X_{itj} is $F_j^{-1}(u_{itj})$, where F_j is the marginal law below. Values at snapshots are available at the end of the corresponding week. A weekly count or volume refers to the completed snapshot week; a posture score or unresolved count describes the current state. For the outcome equation only, $z_{itj} = \mathrm{clip}[(X_{itj}-\mu_j)/\sigma_j, -3, 3]$, with μ_j and σ_j the theoretical mean and standard deviation of F_j. These theoretical transforms are distinct from the sample-fitted logistic scaler.

Table S1 Complete generator coefficients

Predictor	Marginal law	Weekly log-odds β
failed_logins	Poisson(12)	0.08
identity_risk	Beta(2,5)	0.35
impossible_travel	Poisson(0.6)	0.06
privilege_anomaly	Poisson(0.8)	0.08
malware_free_activity	Beta(1.5,6)	0.40

endpoint_alerts	Poisson(5)	0.12
email_phish_score	Beta(2,6)	0.25
cloud_misconfig	Poisson(3)	0.20
exposed_kev	Poisson(1)	0.55
patch_latency_days	Gamma(shape=2,scale=8)	0.32
external_surface	Beta(2,4)	0.25
third_party_risk	Beta(2,3)	0.28
vendor_dependency	Beta(2,2)	0.12
data_egress_gb	Lognormal(log-mean=2.5,log-SD=0.7)	0.10
shadow_ai_risk	Beta(1.5,6)	0.18
backup_health	Beta(6,2)	-0.30
business_criticality	Discrete uniform {1,2,3,4,5}	0.15
security_maturity	Beta(5,2)	-0.30
mfa_coverage	Beta(7,2)	-0.30
remote_access_exposure	Beta(2,5)	0.18

Table 2 of the manuscript states the units for all twenty variables. data/predictor_dictionary.csv additionally supplies the exact theoretical means and standard deviations used in the generator. Beta parameters are shapes; gamma parameters are shape and scale; lognormal parameters are the underlying normal mean and standard deviation. Higher backup_health, security_maturity, and mfa_coverage values are protective in the stipulated outcome equation. Their negative coefficients are assumed, not estimated control effects.

S1 2 Weekly incidents and fourteen-day labels

Threat pressure is $g(t) = 0.015(t-1) + 0.30I(t \geq 40) + 0.10\sin(2\pi t/52)$, for $t=1,54$. The weekly incident logit is $\eta_{it} = -3.4 + a_{s(i)} + b_i + g(t) + \Sigma_j \beta_j z_{itj} + 0.25 \max(z_{KEV},0) \max(z_{patch},0) + 0.20 \max(z_{third\text{-}party},0)\max(z_{dependency},0)$. Weekly incidents E_{it} are Bernoulli(logistic(η_{it})) draws, conditionally independent given the generated probabilities and predictors. The random effects and continuous feature process make observations dependent marginally.

At the end of snapshot week t, the binary target Y_{it} equals $\max(E_{i,t+1},E_{i,t+2})$. It denotes a positive future window, not a distinct incident. This construction explicitly produces overlap between neighboring windows. The exported oracle_probability_14d field is $1-(1-q_{i,t+1})(1-q_{i,t+2})$, conditional on the fully generated future state path. It is retained solely for synthetic auditing and is never supplied as a feature. It is not an operational forecast based on information at t.

Weekly downtime is zero without an incident and otherwise $\exp[1 + 0.15\times\text{business_criticality} - 0.8\times\text{backup_health} + 0.5\varepsilon]$ hours. Weekly synthetic loss is downtime $\times \exp[8 + 0.2\times\text{business_criticality} + 0.5\xi]$ USD, with new independent standard-normal ε and ξ. Fourteen-day downtime and loss sum the next two weeks. They can therefore overlap across snapshots and must not be totaled as independent incident losses. Neither field is a model input or a prediction target in this study.

S1 3 Information boundaries and fixed learners

Primary training contains weeks 1–37 (22,200 rows). Snapshot weeks 38–39 are embargoed, and test weeks are 40–52 (7,800 rows). The last training outcome matures at week 39, before the first test forecast at the end of week 40. The alternative no-embargo fit uses weeks 1–39 and is explicitly a sensitivity comparison with outcome overlap, not a valid replacement for the primary split. The same enterprise identities recur; no unseen-enterprise test is claimed.

The fixed feature list is exactly the twenty names in Table S1. Identifiers, sector label, calendar week, oracle probabilities, and all outcome fields are excluded. Logistic regression is a Standard Scaler→Logistic Regression pipeline, with training-only sample means/scales, C=1, solver=lbfgs, class_weight=balanced, max_iter=2000, and random state=20261028. Random forest uses 200 estimators, max_depth=10, min_samples_leaf=25, max_features=sqrt, class_weight=balanced, n_jobs=1, and the same random state. GradientBoostingClassifier uses 150 estimators, learning_rate=0.05, max_depth=2, min_samples_leaf=25, subsample=1.0, and the same random state. Boosting receives inverse-frequency sample weights n/(2n_c). Other constructor defaults and primitive parameters are archived in results. json. No hyper parameter search, model selection on test outcomes, or probability recalibration is performed.

Within each test week, scores are sorted descending and ascending enterprise ID resolves ties. Exactly ninety enterprises are flagged each week. Weekly cutoffs are the score at the ninetieth rank, not fixed probability thresholds. All families therefore flag 1,170 snapshots. The policy uses the current week's scores and does not inspect its future labels.

S1 4 Scores and uncertainty

ROC-AUC uses sklearn.metrics.roc_auc_score. Average precision uses sklearn. metrics.average_precision_score and the non-interpolated sum of precision weighted by successive recall increments; it is not trapezoidal area. The Brier score uses sklearn. metrics. brier_score_loss. BSS = 1−BS/[p(1−p)], with p the evaluation sample's positive proportion. The test-prevalence reference is a hindsight evaluation comparator. It is not estimated from training and is not a deployable forecasting policy.

Five hundred paired enterprise-cluster resamples use default_rng (20261128). Each sample draws 600 enterprise identities with replacement, retaining every selected enterprise's complete thirteen-week test trajectory and the already issued predictions and flags. Percentiles 2.5 and 97.5 give metric intervals. Pairwise AP and AUC differences are computed within the same resample. Models are not refitted and flags are not reranked, so intervals are conditional on fitted models, the realized policy, and the single generator. They do not include model-development or across-seed uncertainty.

The DeLong calculation uses pairwise ranking influences for every observed positive and negative score. Covariance is the sample covariance of positive influences divided by the positive count plus the covariance of negative influences divided by the negative count. Differences use their paired covariance and a two-sided normal approximation. The exported AUC is checked against scikit-learn. Because conventional DeLong assumes independent rows and the panel violates that assumption, these p-values are descriptive only; clustered intervals take precedence. Their inclusion does not establish independent national samples or equivalence of model families.

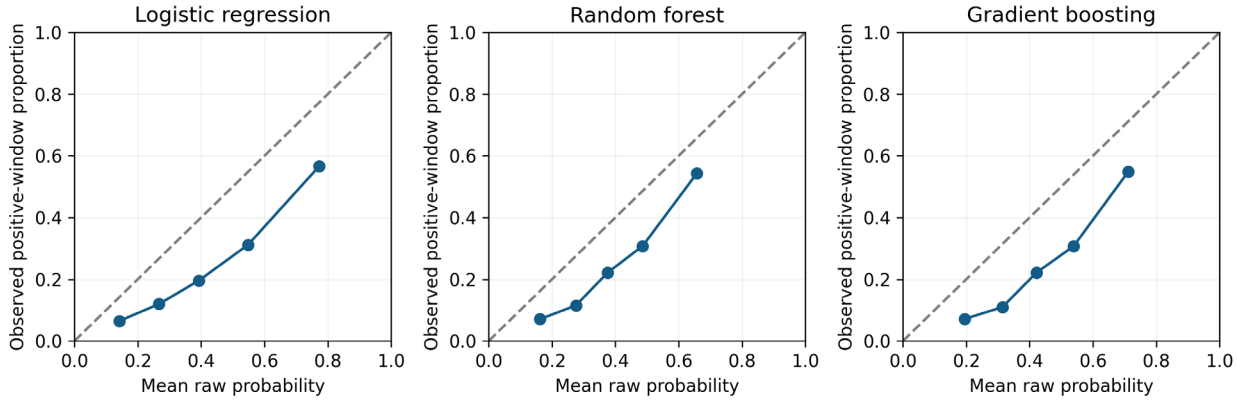
S1 5 Explanations and checks

Permutation importance concerns the prespecified logistic baseline. For each feature, its whole thirteen-week trajectory is randomly reassigned among 600 enterprises, leaving other columns fixed. Twenty repetitions use default_rng(20261228). The reported importance is the mean loss in test AP; the error bar is its permutation standard deviation. These are marginal perturbations, can disrupt cross-feature dependence, and are not causal contrasts. Every predictor is present in the CSV and manuscript Figure 3.

The local explanation uses the test row with largest raw logistic probability. Relative to the training-standardized mean (zero), contribution_j = β̃_j(z_j−0), and the reference log odds are the fitted intercept. The code verifies logistic (reference + Σ contributions) against predict_proba. This is an exact linear-model log-odds decomposition, not SHAP. data/local_log_odds_all_20.csv records the raw values, standardized values, coefficients, and signed terms.

The verification routine recomputes AP, AUC, and Brier scores from exported predictions; confirms ninety flags per week, twenty unique predictors, the embargo boundary, and local reconstruction; confirms that utility calculations count each incident once; and excludes oracle and outcome fields from predictors. Data-file SHA-256 digests are stored in data_sha256.json. These checks verify the implementation and consistency of the saved realization, not the substantive realism of the assumptions.

Figure S1 Descriptive calibration quintiles



Note. Each group contains approximately one fifth of the test observations for the specified model. Points are descriptive and have no uncertainty bars. Raw models are not recalibrated. A lower Brier score alone does not demonstrate perfect calibration.

Supplementary Material S2 Scenario assumptions and figures

All scenarios are illustrations conditional on explicit assumptions. They do not estimate causal control effectiveness, incident prevention, real enterprise loss, or national resilience. The utility scenario uses distinct generated incidents, while other tables explicitly describe snapshot-window frequencies or raw model probabilities. Scenario illustrations are separate from the primary fifteen-percent ranking policy.

Table S2 Scenario assumptions

Quantity	Value	Interpretation
MFA and conditional access	57	Residual-risk index with baseline 100; independent assumption

Accelerated KEV patching	68	Residual-risk index with baseline 100; independent assumption
Network segmentation	72	Residual-risk index with baseline 100; independent assumption
Endpoint and identity analytics	54	Residual-risk index with baseline 100; independent assumption
Third-party access review	65	Residual-risk index with baseline 100; independent assumption
Immutable backups	61	Residual-risk index with baseline 100; independent assumption
Review cost	USD 250 per alert	Each reviewed enterprise-week incurs a cost
Loss per distinct incident	USD 20,000	Fixed scenario amount; separate from generated loss values
Prevention probability	0.10	Applied once per covered underlying incident
Capacity grid	1%–40%, steps of 1 percentage point	Post-outcome exploration; not a prospective policy choice
Event window	Weeks 41–54	Incidents covered by forecast alerts issued in weeks 40–52
PSI reference	Snapshot weeks 1–13	Raw fitted logistic scores; descriptive monitoring
PSI bins	Reference deciles; duplicate edges removed	Outer edges ±infinity; 0.5 pseudocount per bin
PSI attention line	0.20	Illustrative review reference, not a validated alarm

For capacity a, each week flags floor(600a) enterprises by raw logistic score. An incident at week j is covered if an alert was issued for the same enterprise at j−1 or j−2, restricted to available test forecast weeks 40–52. The union of those alert windows is used, so each underlying incident is counted at most once. Expected gross scenario benefit is 0.10 × USD 20,000 × covered distinct incidents; review cost is USD 250 × review count. Their difference is the plotted assumed expected net value. The 0.10 prevention parameter is stipulated; no intervention is simulated or observed to establish it. The maximizing capacity is chosen after looking at test outcomes and is reported only as a hindsight illustration.

Table S3 Input perturbations and independent continuity assumptions

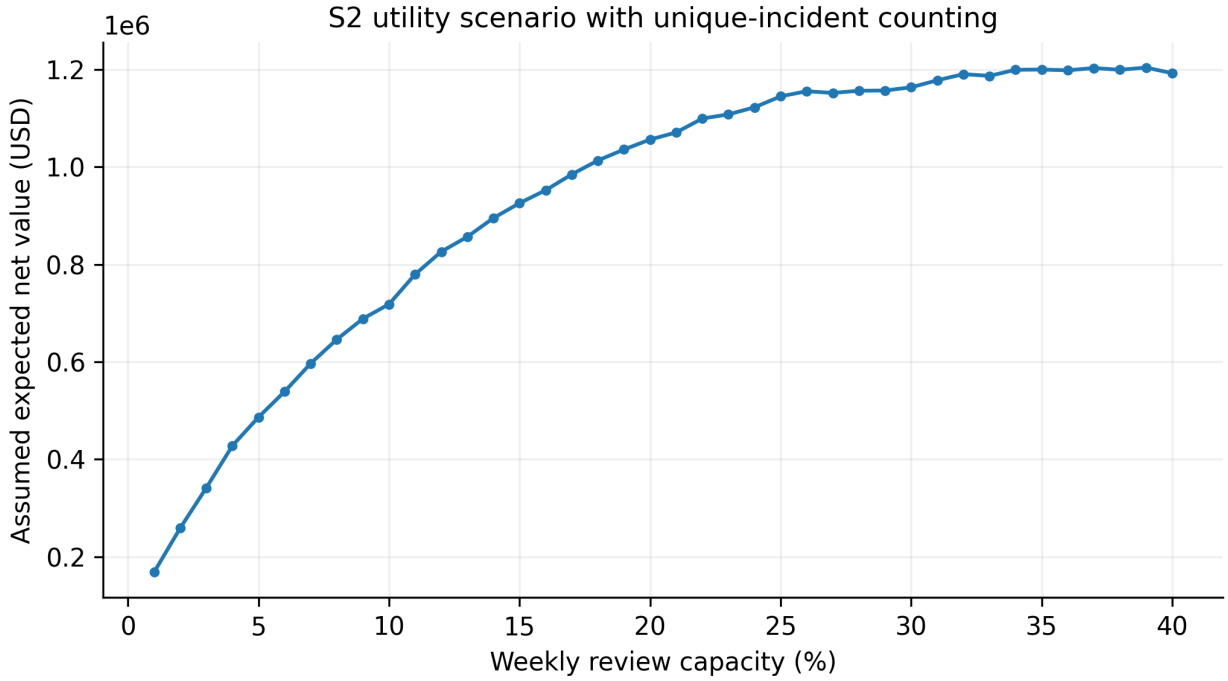
Scenario	**Feature change**	**Continuity index**
Baseline	No changes	1.00
Phishing	email_phish_score: +0.15	0.86
Zero-day	exposed_kev: +2; patch_latency_days: ×1.3	0.78
Supplier outage	third_party_risk: +0.2; vendor_dependency: +0.15	0.73
Shadow AI	shadow_ai_risk: +0.25; remote_access_exposure: +0.1	0.88

Combined	email_phish_score: +0.15; exposed_kev: +2; patch_latency_days: ×1.3; third_party_risk: +0.2; vendor_dependency: +0.15; shadow_ai_risk: +0.25; remote_access_exposure: +0.1	0.61

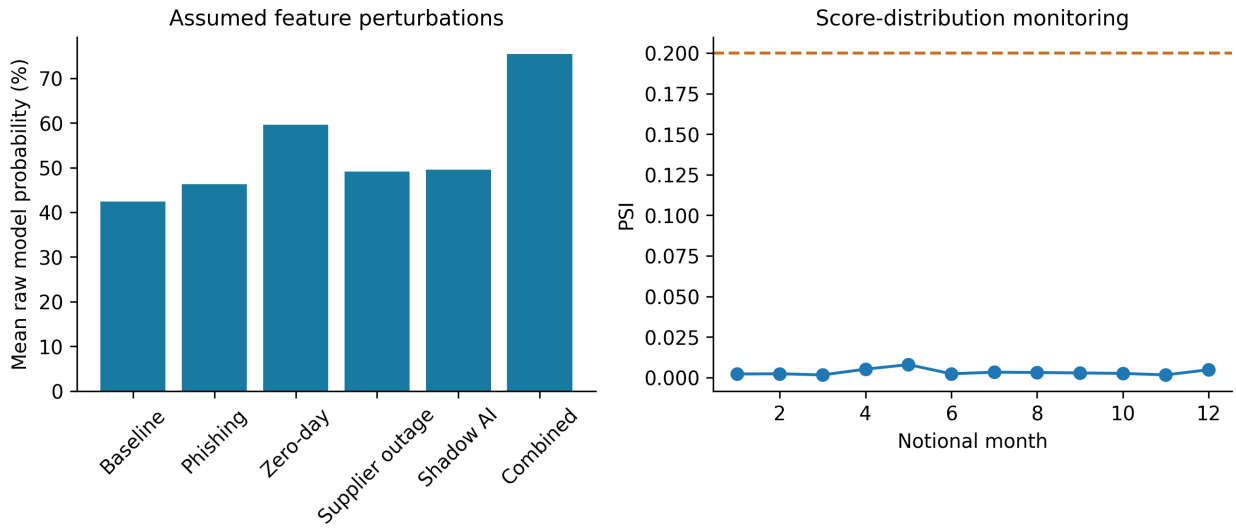
Additive changes to beta-distributed scores are clipped to [0,1]. Counts and positive patch latency retain their physical units. No model is retrained. The historical label 'zero-day' denotes an illustrative vulnerability-pressure scenario represented by two additional exposed KEVs and 30% longer patch latency; it is not a validated model of an unknown vulnerability. The combined scenario applies every listed feature change once. The reported mean is a raw logistic model response; its bias and calibration limitations persist. Continuity indices are assigned separately and are not computed from predicted risk.

For PSI, reference-bin and monthly-bin counts each receive a pseudocount of 0.5 and are normalized to proportions q and p. PSI is $\Sigma(p-q)\log(p/q)$. Notional month m is $\text{floor}[12(t-1)/52]+1$, giving twelve groups from 52 snapshot weeks. The reference uses weeks 1–13. This trace describes the distribution of raw fitted scores, including earlier in-sample snapshots, and is not a prospective deployment test. A stable score distribution cannot establish stable calibration, especially when the generator shifts latent threat pressure without a corresponding shift in the input process.

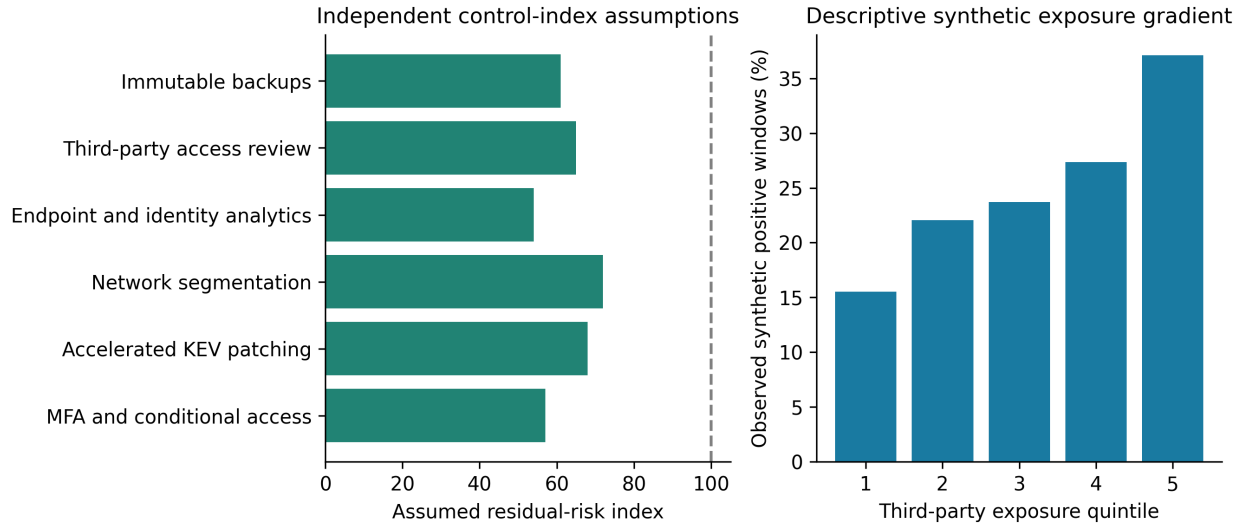
Figure S2 Utility illustration



Note. The curve uses the explicit costs and prevention probability in Table S2 and deduplicates underlying incidents. It is conditional and selected with hindsight. No avoided enterprise losses were measured.



Note. The left panel is the response of the fixed raw logistic model to the listed input changes. The right panel describes score-distribution variation; the 0.20 line is an illustrative reference. Neither panel estimates operational effectiveness.



Note. Control values on the left are stipulated residual-risk indices with baseline 100. The right panel shows observed synthetic positive-window proportions by third-party exposure quintile. These distinct quantities must not be combined into a causal effect.

Decision scenarios

S2A Control effect assumptions

Control scenario	Residual index	Index reduction
MFA and conditional access	57	43%
Accelerated KEV patching	68	32%
Network segmentation	72	28%

Control scenario	Residual index	Index reduction
Endpoint and identity analytics	54	46%
Third-party access review	65	35%
Immutable backups	61	39%

Note. The baseline residual-risk index is 100. Each control is a separate illustration. Index reductions are not reductions in disruption probability, and indices must not be added or multiplied to infer a combined effect.

S2B Economic assumptions

Parameter	Assumed value	Application
Investigation cost	USD 250	Per reviewed enterprise-week, including reviews without a subsequent incident.
Loss per missed event	USD 20,000	Fixed loss per distinct underlying incident; also applies to covered incidents not prevented.
Prevention probability	0.10 (10%)	Applied once to an incident covered by an earlier alert; no actual prevention was observed.
Review-capacity grid	1%–40%; 1 pp steps	Post-outcome sensitivity grid; separate from the primary 15% review policy.

S2C Stress parameters and continuity assumptions

Each stress scenario changes the listed inputs to the fixed logistic model. All other predictors and fitted preprocessing remain unchanged, and the model is not retrained. Additive changes to fraction scores are absolute changes on their 0–1 scales, not relative percentage increases.

Scenario	Input change	Continuity index
Baseline	No input changes.	1.00
Phishing	email_phish_score: +0.15	0.86
Zero-day	exposed_kev: +2 patch_latency_days: ×1.3	0.78
Supplier outage	third_party_risk: +0.2 vendor_dependency: +0.15	0.73
Shadow AI	shadow_ai_risk: +0.25 remote_access_exposure: +0.1	0.88
Combined	Apply each no-baseline change above once; no additional multiplier.	0.61

Note. Beta-distributed fraction scores are clipped to [0, 1] after an additive change. The +2 KEV change is a count increase; ×1.30 patch latency is a 30% increase in days. Each combined-scenario input change is applied once.

Alert and incident counting rules

S2D Drift monitoring and exposure grouping

These specifications complete the assumptions for the population stability index (PSI) trace and third-party exposure illustration. Monitoring describes distributions of raw logistic scores; it does not establish operational detection performance or stable probability calibration.

Parameter	Specification
Model and score	Fixed logistic model; raw positive-class probability.
PSI reference period	Snapshot weeks 1–13.
Bin construction	Reference-score deciles; duplicate edges removed; outer limits $-\infty$ and $+\infty$.
Smoothing	Add 0.5 to every bin count in both reference and comparison groups, then normalize.
Comparison periods	Twelve notional month groups across the 52 snapshot weeks.
Attention reference	PSI = 0.20; illustrative review line, not a calibrated alarm threshold.
Exposure grouping	Five quantile groups of test-set third_party_risk using pandas.qcut; observed positive-window rate within each group.

$$PSI = \sum_{b} (p_b - q_b)\, ln(\frac{p_b}{q_b})$$

Here, p and q are the smoothed bin proportions in the comparison and reference groups. Month group m is floor [12(t − 1)/52] + 1 for snapshot week t; “month” is a grouping convention, not a calendar-month sample. Earlier scores include in-sample snapshots, so the trace is descriptive rather than a prospective deployment evaluation.

Exposure-quintile positive rates are outputs of the artificial data, not assumed control effects. Score stability does not rule out a change in outcome prevalence or latent threat pressure. Control indices, feature perturbations, continuity indices, and the economic scenario are separate illustrations and are not combined into a causal estimate.

Additional figures for the replacement simulation. Complete assumptions are in the previously supplied Supplementary Table S2, sections S2A–S2D. Control indices are stipulated; model responses and exposure patterns are descriptive.

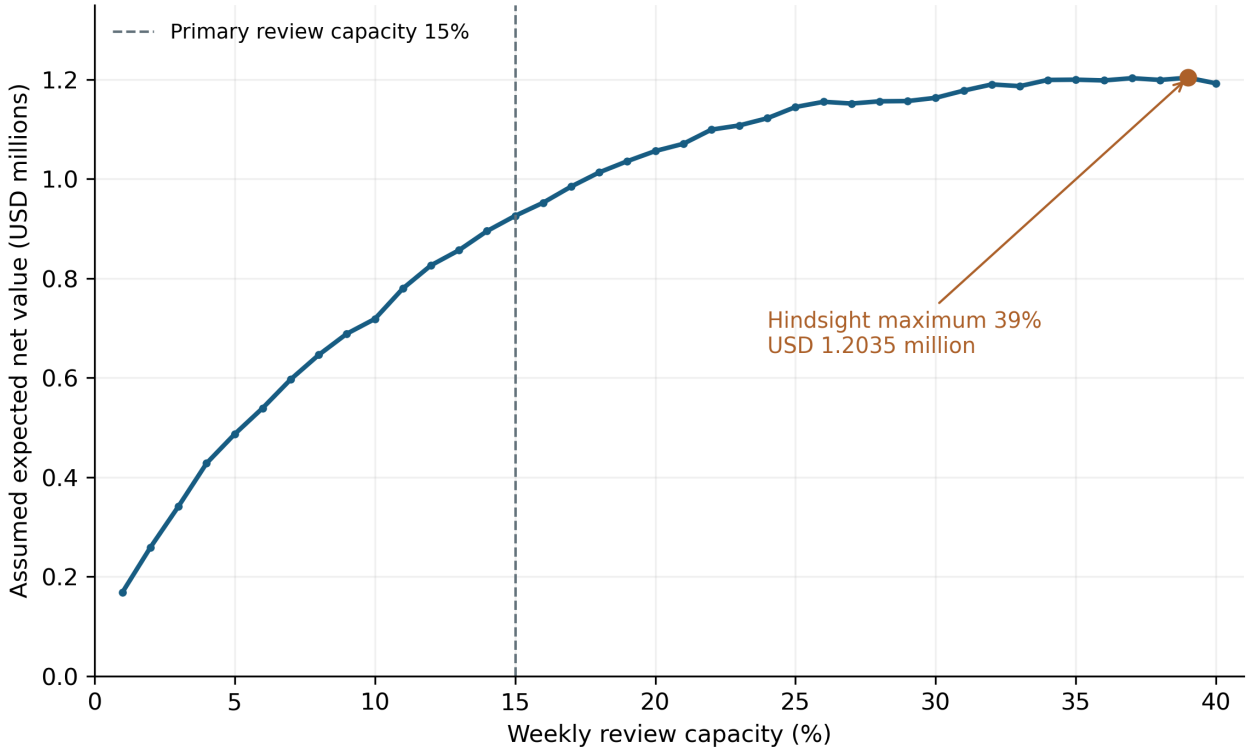


Figure S2 Expected net value across weekly review capacities
At 15% capacity, 1,170 reviews cover 609 distinct incidents, giving USD 925,500 assumed expected net value. The hindsight maximum occurs at 39% capacity (USD 1,203,500). Each incident is counted once. Economic assumptions: Table S2B.
Data: utility_curve.csv.

Feature stress and score stability

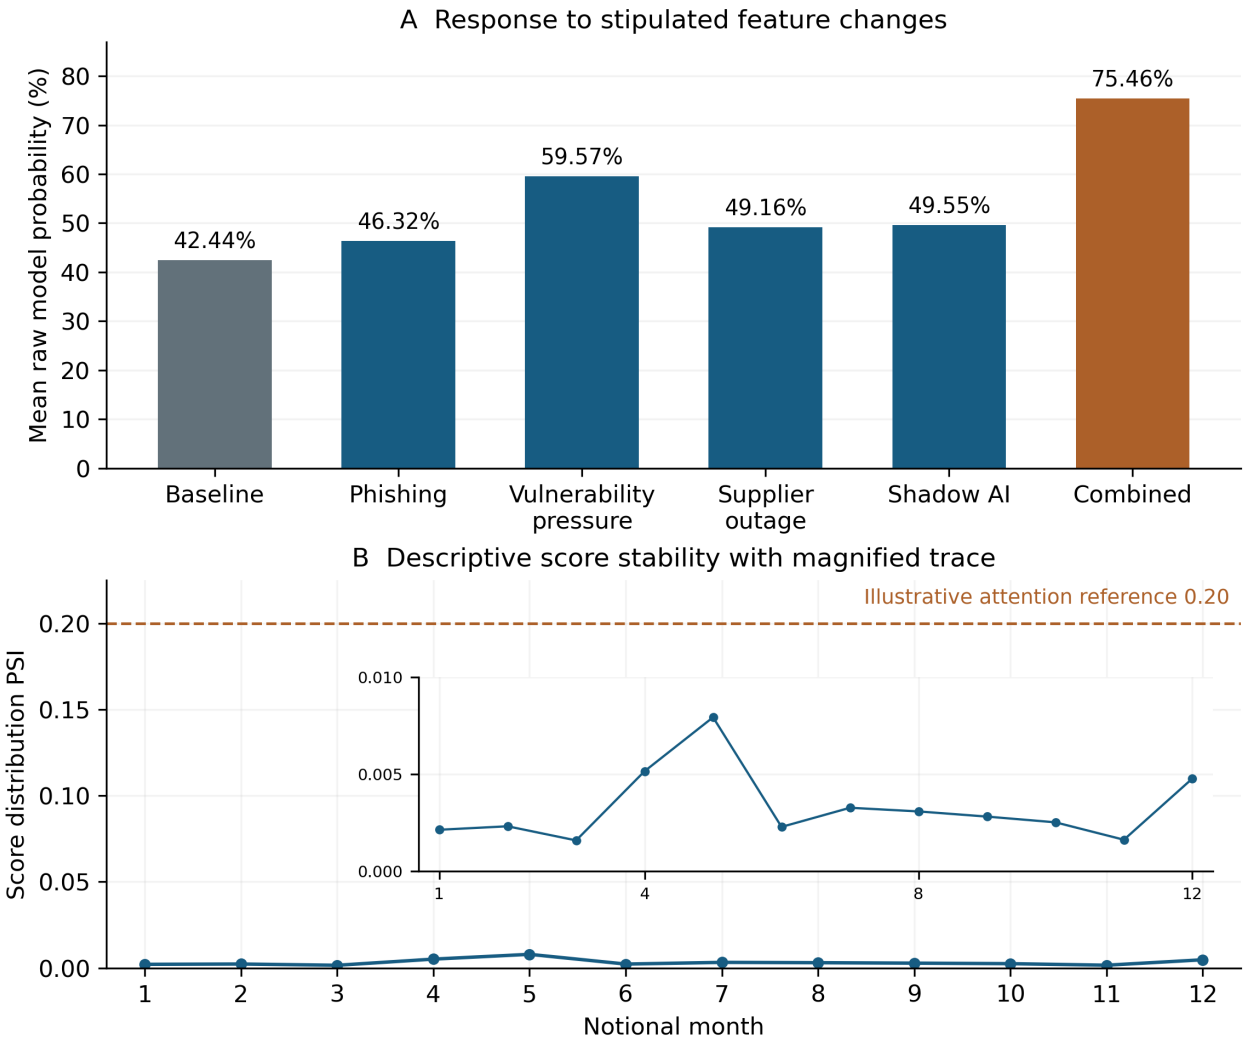


Figure S3 Model responses and descriptive score stability

A. Raw logistic model responses to the feature changes in Table S2C. B. Monthly score PSI with a magnified inset; maximum PSI = 0.007936. The 0.20 line is illustrative. Monitoring assumptions: Table S2D. Stable scores do not establish stable calibration.

Data: stress_scenarios.csv and psi_monitoring.csv.

Control indices and supplier exposure

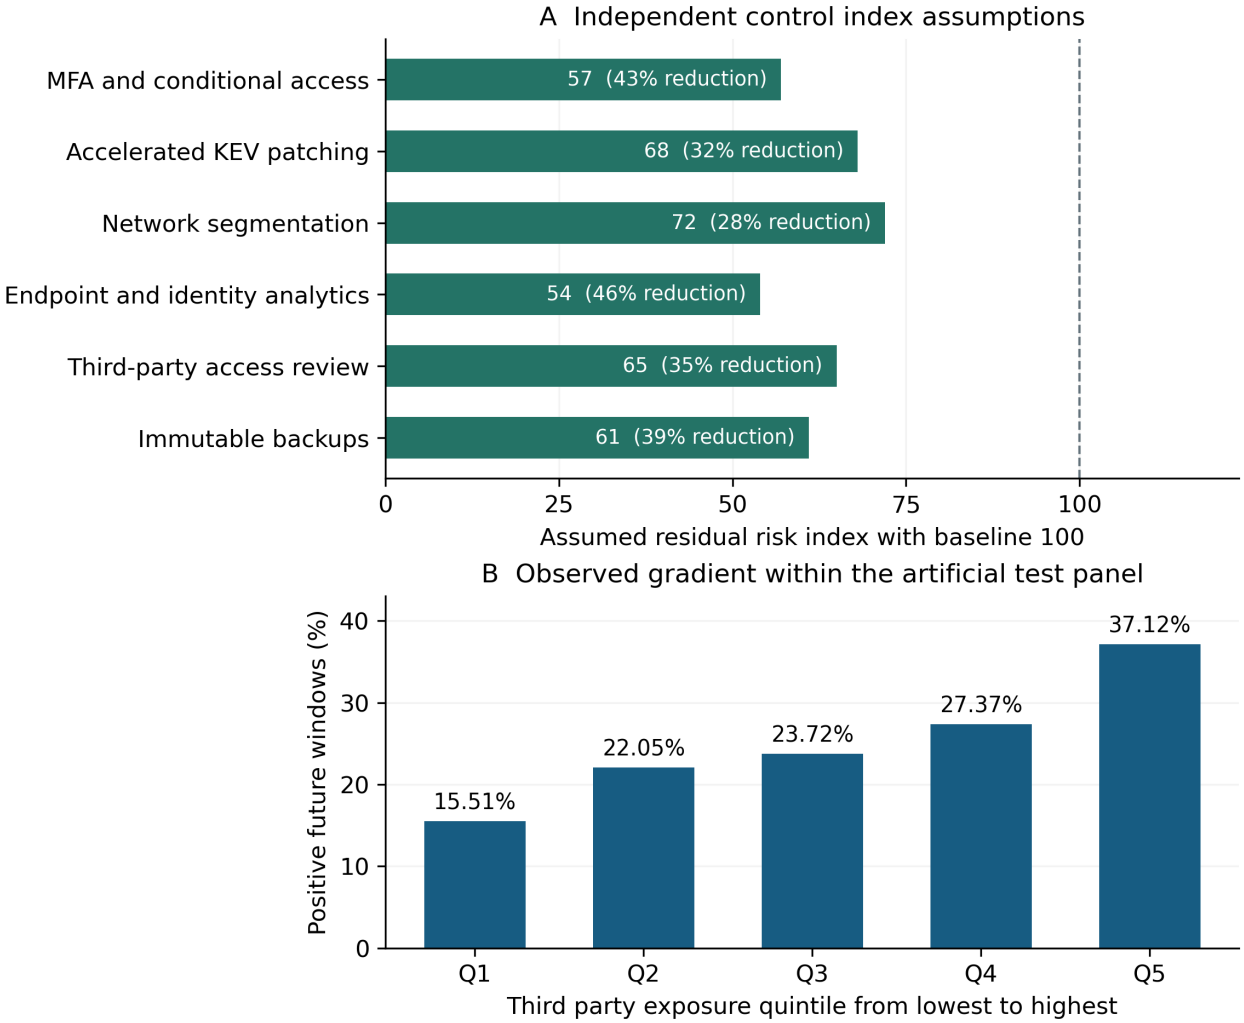


Figure S4 Assumed control indices and synthetic exposure patterns

A. Independent assumed control indices relative to baseline 100 (Table S2A). B. Synthetic positive-window rates by third-party exposure quintile; 1,560 snapshots per group (Table S2D). These quantities do not estimate causal control or supplier effects.

Data: scenario_assumptions.json and third_party_quintiles.csv.