Skepthical review: *Latent Class Trajectories of AI-Induced Job Security: Identifying Organizational Catalysts for Professional Stability*

Summary

The paper investigates heterogeneity in employees' perceptions of how AI affects their job security, arguing that aggregate (variable-centered) analyses mask meaningful subgroups. Using survey data from $2{,}603$ employees in large global enterprises, the authors apply Latent Class Analysis (LCA) to two indicators—perceived current impact of AI on job security and expected future impact (time horizon inconsistently reported as 3 vs 4 years)—to identify three interpretable latent classes: a majority "Resiliently Optimistic" group with stable positive security perceptions, a "Stagnant Neutral" group centered on no perceived impact, and an "Anxiously Declining" minority anticipating deteriorating security (Sec. 1; Sec. 2.2; Sec. 3.1). They then reduce 16 binary emotion items into Positive and Negative Affect factors via tetrachoric EFA (Sec. 2.3.1; Sec. 3.2) and relate organizational "enablers" and "deterrents" to class membership using Elastic Net feature selection followed by multinomial logistic regression (Sec. 2.3.2; Sec. 3.3–3.4). Finally, a Monte Carlo 'policy simulation' contrasts scenarios emphasizing training versus direct involvement in AI development, concluding that participatory practices are more strongly associated with optimistic profiles than training alone (Sec. 2.3.3; Sec. 3.5). The paper is timely and offers an actionable person-centered framing, but key elements of sampling/measurement transparency, LCA specification/diagnostics (including an internal inconsistency in reported class proportions), handling of classification uncertainty in downstream regression, and the causal interpretation of scenario simulations require substantial clarification to support the paper's strongest substantive and prescriptive claims (Sec. 2–4).

Strengths

- Timely and practically important research question on AI-related job security perceptions, with a clear motivation for examining heterogeneity beyond averages (Sec. 1).
- Interpretable three-class typology ("Resiliently Optimistic", "Stagnant Neutral", "Anxiously Declining") that communicates substantive differences more clearly than aggregate sentiment summaries (Sec. 3.1).
- Methodologically coherent end-to-end workflow (LCA $\rightarrow$ affect factor reduction $\rightarrow$ predictive modeling $\rightarrow$ scenario-based predictions) aimed at translating findings into organizational levers (Sec. 2.2–2.3; Sec. 3.1–3.5).
- Action-oriented focus on structural/non-monetary practices (e.g., involvement in AI development, autonomy over AI tools, recognition, certifications, career incentives) and salient concerns (job loss, privacy, distrust, accountability), which yields concrete managerial implications (Sec. 3.3–3.5).

- Inclusion of some robustness-oriented elements (entropy/BIC reporting for class selection; VIF checks; alternative handling of "Not sure" responses) signals attention to model stability, even though key details are currently underspecified (Sec. 3.6).
- Results narration is generally aligned with reported probabilities and odds-ratio arithmetic, and visualizations are directionally helpful for communicating the story (Sec. 3.1–3.5; Figures 1–5).

Major issues

1. **The manuscript repeatedly labels the latent classes as "psychological trajectories" (Sec. 1; Sec. 3.1; Sec. 4), but the design is cross-sectional and uses two contemporaneous self-reports (current impact and expected future impact), not longitudinal observations of change. This risks overstating what the data can support (e.g., implied transitions, movement between states, or realized trajectories).**

   Recommendation: Revise framing throughout (Sec. 1; Sec. 3.1; Sec. 4) to describe the LCA output as "latent perception profiles" or "current–expected perception patterns," unless true panel data exist. In Sec. 2.2 and Sec. 4, add an explicit limitation that expected future perceptions are anticipatory beliefs and do not constitute observed longitudinal change; avoid language implying transitions/effects over time.

2. **Sampling and data-collection details are insufficient for evaluating external validity and bias. Sec. 2.1 does not adequately describe recruitment (panel vs internal survey), countries/regions, industries, participating firms, response rate, inclusion/exclusion criteria, or key demographics/roles. Given the central claims about "large global enterprises," representativeness and scope conditions are unclear (Sec. 2.1; Sec. 4).**

   Recommendation: Expand Sec. 2.1 to include: sampling frame and recruitment channel(s); enterprise/industry and geographic composition; response rate and missingness handling; respondent demographics and job characteristics (e.g., function, seniority, tenure); and any weighting/post-stratification. If the sample is convenience-based or restricted to specific firms, state this plainly and scope the claims in Sec. 4 accordingly.

3. **Key measures and coding are under-documented, undermining construct validity and reproducibility. The paper frequently references internal item codes (e.g., QF_3, QKB_2_11, QKC_4) without presenting item wording, scale anchors, or recoding decisions for: (i) LCA indicators (current/expected impact on job security), (ii) the 16 affect items, and (iii) enablers/deterrents (training, involvement, autonomy, recognition, privacy concerns, distrust, accountability fear, etc.) (Sec. 2.1; Sec. 2.3.1; Sec. 3.2–3.4).**

Recommendation: Add an Appendix (or a main-text measures table) mapping every item code to: question wording, response options, coding direction (what higher values mean), and any dichotomization rules. In Sec. 2.1–2.3, explicitly define the main constructs and how binary variables were constructed. Update Sec. 3.2–3.4 and figure captions to use human-readable labels alongside (or instead of) item codes.

4. **LCA specification/diagnostics are incomplete, and there is an internal inconsistency in reported class proportions. The text reports class shares of 60.9% / 26.5% / 12.6% (Sec. 3.1; Sec. 4), while Table 1 appears to show different shares (e.g., 44.90% / 24.51% / 12.50%), and related labeling issues appear (e.g., misspelling). Beyond BIC/entropy, the paper does not report estimation settings (starts, convergence, local maxima checks), indicator category handling (including "Not sure"), or local-independence assessment (Sec. 2.2; Sec. 3.1; Table 1; Fig. 1).**

   Recommendation: First, reconcile and correct all class-size reporting across Sec. 3.1, Table 1, figures, and Sec. 4 (including any OCR/typos). Then, expand Sec. 2.2/Sec. 3.1 to report: indicator response categories and treatment of "Not sure"; estimator/software; number of random starts and best log-likelihood replication; convergence criteria; additional fit indices (AIC, SABIC; and LRT-style tests if used/available); and class-specific average posterior probabilities (not just entropy). Briefly discuss plausibility of local independence given the conceptual linkage between "current" and "expected" items.

5. **Downstream regression likely treats class membership as observed (modal assignment), which can bias coefficients and standard errors due to classification error. Entropy alone is not sufficient evidence that this is negligible (Sec. 2.3.2; Sec. 3.3–3.4).**

   Recommendation: Clarify in Sec. 2.3.2 whether the multinomial logit uses most-likely class labels or a method that propagates classification uncertainty. If using modal assignment, consider a three-step approach (e.g., Vermunt/BCH-style correction) or a one-step latent class model with covariates. At minimum, report class-specific average posterior probabilities and discuss the likely direction of bias if uncertainty is ignored.

6. **Elastic Net feature selection followed by an unpenalized multinomial logit is presented in an inferential style (odds ratios, significance-style language), but post-selection inference can be anti-conservative, and the selection procedure is under-specified (candidate predictors/interactions, α/λ tuning, CV design, inclusion rules) (Sec. 2.3.2; Sec. 3.3–3.4).**

   Recommendation: In Sec. 2.3.2, explicitly list the full candidate set (including interactions) and describe α/λ tuning (k-fold CV details; metric; stratification by class if applicable). Reframe claims where appropriate as 'selected/not selected under the chosen regularization' rather than 'no effect.' If inferential claims are central, add stability checks (bootstrap/stability selection) and/or fit a penalized multinomial model directly with cross-validated performance reporting; provide a full results table (ORs with 95% CIs and p-values, plus fit statistics) in Sec. 3.4 or an Appendix.

7. **Potential confounding and hierarchical structure are not adequately addressed. Organizational practices (training, involvement, incentives) are plausibly correlated with unobserved firm characteristics (AI maturity, layoffs, managerial quality, HR policies), and employees may be nested within firms/regions. Affective states may also be endogenous to perceived insecurity, complicating interpretation of associations (Sec. 3.3–3.5; Sec. 4).**

   Recommendation: Add richer controls where available (role/function, seniority, tenure, AI exposure, firm AI-adoption stage). If firm identifiers exist, use cluster-robust SEs or multilevel modeling; at minimum, discuss clustering as a limitation. In Sec. 4, explicitly acknowledge omitted-variable and reverse-causality risks and scope conclusions as associational unless stronger identification is introduced.

8. **The EFA on 16 binary affect items is under-reported, and some reported validity checks are directionally confusing (e.g., negative correlations between 'Negative Affect' and deterrents), raising concerns about coding/sign conventions and construct interpretation (Sec. 2.3.1; Sec. 3.2).**

   Recommendation: In Sec. 2.3.1/Sec. 3.2, report extraction method, factor retention criteria (e.g., parallel analysis), loading thresholds, handling of cross-loadings, and factor-score construction. Provide a full loading table (Appendix acceptable) and clarify coding direction for items and deterrents; state whether factor signs were flipped for interpretability. Consider an oblique rotation robustness check (since positive/negative affect may correlate) and report factor correlations.

9. **The Monte Carlo 'policy simulation' is presented in a way that can be read causally ('effects' of training vs involvement), but it appears to be model-based scenario prediction from an observational cross-sectional model. Simulation design details are insufficient (draws, covariate setting, coefficient uncertainty propagation), and predicted probabilities are presented without uncertainty intervals (Sec. 2.3.3; Sec. 3.5; Sec. 4).**

   Recommendation: Reframe Sec. 3.5 and Sec. 4 as scenario-based predictions rather than policy effects unless causal identification is established. In Sec. 2.3.3/Sec. 3.5, specify: number of simulation draws; how covariates are set (means vs empirical distribution); scenario encodings; and whether parameter uncertainty is propagated (sampling coefficients from their covariance). Report uncertainty intervals (e.g., 95% CI) for scenario probabilities and explicitly state the assumptions required for causal interpretation.

Minor issues

1. Time-horizon inconsistency for the expected-impact indicator: Sec. 2.2 refers to "expected impact in three years," while Table 1 labels "Expected Impact … (4 years)" (Sec. 3.1).

   Recommendation: Verify the questionnaire wording and standardize the horizon across Sec. 2.2, Table 1, figures, and the narrative in Sec. 3.1 and Sec. 4.

2. Handling of "Not sure" responses is not clearly documented for the primary analyses (LCA indicators vs predictors), even though robustness checks mention alternative treatments (Sec. 3.6).

   Recommendation: In Sec. 2.1–2.3, specify for each variable family (indicators, affect items, enablers/deterrents) whether "Not sure" is modeled as its own category, recoded, or treated as missing; describe any imputation/exclusion rules and the resulting analytic N for each model.

3. Interaction effects in the multinomial model (Sec. 3.4) are potentially vulnerable to multiple-testing/overfitting concerns, especially if many candidate interactions were explored via Elastic Net.

   Recommendation: State how many interactions were considered and the rationale for those reported. Prefer a small set of theory-motivated interactions, and present marginal-effects plots on the probability scale to improve interpretability; consider multiplicity-aware reporting (e.g., stability selection frequencies).

4. The operationalization of 'training' is not sufficiently precise given the strong conclusion that training alone is less impactful than involvement (Sec. 3.3; Sec. 3.5; Sec. 4).

   Recommendation: Provide item wording and coding for training (availability vs participation vs intensity/quality). In Sec. 4, qualify conclusions as pertaining to this specific measure and note that different training designs may yield different associations.

5. Related-work positioning could be strengthened. The Introduction critiques variable-centered approaches but does not clearly situate LCA relative to alternative heterogeneity models (mixture regression, interaction-rich models) or prior person-centered work on technology/job insecurity (Sec. 1–2).

   Recommendation: Add a brief related-work subsection in Sec. 1–2 that situates LCA among other approaches and explicitly states what is novel here (e.g., organizational levers + scenario predictions in a large enterprise sample).

6. Several figures and captions remain hard to interpret without additional context: cryptic item codes, missing estimator/model details, unclear error bars/uncertainty, and incomplete class-enumeration reporting (Figures 1–5; Sec. 3.1–3.5).

Recommendation: Enhance figure captions with: decoded item labels; sample sizes per model; key modeling choices (e.g., rotation type for Fig. 2; reference class and coding for Fig. 4); and explicit uncertainty definitions for error bars in Fig. 5. For Fig. 1, consider adding additional indices (AIC/SABIC; LRT-based tests if available) and visually marking the chosen solution.

7. Ethics and sensitive-topic handling (job insecurity, privacy concerns) are not explicitly discussed (Sec. 2.1; Sec. 4).

   Recommendation: Include an ethics statement (Sec. 2.1 or a dedicated subsection) covering consent, anonymity/confidentiality, ethics approval (if applicable), and safeguards. In Sec. 4, discuss ethical considerations for recommended organizational interventions (e.g., avoiding coercion in 'participation').

8. Model reporting is mostly verbal; core equations/notation are not presented, limiting reproducibility and making it harder to verify reference-category and interaction parameterization (Sec. 2.2–2.3).

   Recommendation: Add concise model equations for: (i) the LCA likelihood and class-conditional response probabilities, (ii) multinomial logit link and reference class, and (iii) the elastic net objective with α/λ; define $\mathrm{OR} = \exp(\beta)$ and interaction construction explicitly.

Very minor issues

1. Typographical/label inconsistencies occur (e.g., "Stegment Neutral" vs "Stagnant Neutral"), alongside inconsistent capitalization/quoting of class names and occasional line-break artifacts (Sec. 1; Sec. 3.1; Table 1; Sec. 4).

   Recommendation: Proofread to standardize class naming, capitalization, and formatting across text/tables/figures; correct misspellings; and remove line-break artifacts.

2. Numerical/statistical formatting varies (decimal places, OR/CI notation), and at least one percentage statement appears sensitive to rounding (e.g., 92.4% vs 92.3% from displayed components) (Sec. 3.1; Sec. 3.4).

   Recommendation: Standardize reporting (e.g., consistent decimals for percentages; consistent OR/CI formatting). Where rounding matters, either compute text from unrounded internal values and state so, or align text to the rounded table values.

3. Figure accessibility could be improved (font sizes, crowded labels, colorblind-safe palettes), and some axes/scales may be misread without clearer annotation (Figures 1–5).

   Recommendation: Increase font/marker sizes, adopt colorblind-safe palettes, and clarify axes (avoid dual/truncated scales unless clearly annotated). Provide vector/high-resolution outputs.

Mathematical consistency audit

This section audits **symbolic/analytic** mathematical consistency (algebra, derivations, dimensional/unit checks, definition consistency).

Maths relevance: light

The paper primarily provides methodological descriptions and reports model outputs (class probabilities, odds ratios, correlations) but contains essentially no explicit equations or step-by-step derivations. As a result, only limited internal algebraic sanity checks (probability normalization and basic transformations like $\mathrm{OR} \leftrightarrow$ percent change) are directly auditable from the PDF text/tables/figures.

Checked items

1. ✔ **Normalization of LCA conditional probabilities (Class 1)** (Table 1, Sec. 3.1, p.6)
   - **Claim:** Within each class, the conditional response probabilities across response categories sum to 1 for each indicator (current and expected).
   - **Checks:** normalization/constraints
   - **Verdict:** PASS; confidence: high; impact: moderate
   - **Assumptions/inputs:** Displayed probabilities are rounded decimals., Categories shown exhaust the response options for each indicator.
   - **Notes:** Class 1 current: $0.000 + 0.030 + 0.124 + 0.419 + 0.423 + 0.004 = 1.000$. Class 1 expected: $0.000 + 0.046 + 0.120 + 0.377 + 0.449 + 0.008 = 1.000$.

2. ✔ **Normalization of LCA conditional probabilities (Class 2)** (Table 1, Sec. 3.1, p.6)
   - **Claim:** Within Class 2, current and expected conditional probabilities sum to 1.
   - **Checks:** normalization/constraints
   - **Verdict:** PASS; confidence: high; impact: moderate
   - **Assumptions/inputs:** Rounding to three decimals can cause small deviations from exact 1.000.
   - **Notes:** Class 2 current sums to 1.001 $(0.051 + 0.861 + 0.052 + 0.037)$, consistent with rounding. Class 2 expected sums exactly to 1.000 $(0.152 + 0.628 + 0.173 + 0.047)$.

3. ✔ **Normalization of LCA conditional probabilities (Class 3)** (Table 1, Sec. 3.1, p.6)
   - **Claim:** Within Class 3, current and expected conditional probabilities sum to 1.
   - **Checks:** normalization/constraints

- **Verdict:** PASS; confidence: high; impact: moderate
- **Assumptions/inputs:** Rounding to three decimals can cause small deviations from exact 1.000.
- **Notes:** Class 3 current: $0.222 + 0.701 + 0.041 + 0.036 + 0 + 0 = 1.000$. Class 3 expected sums to 1.001 $(0.296 + 0.510 + 0.130 + 0.064 + 0.001)$, consistent with rounding.

4. ✔ **Class size proportions sum to 100%** (Sec. 3.1 bullets, p.5 (also Table 1, p.6))

   - **Claim:** Reported class proportions constitute a full partition of the sample.
   - **Checks:** constraint/sanity check
   - **Verdict:** PASS; confidence: high; impact: minor
   - **Assumptions/inputs:** Percentages reported are rounded to one decimal place in text and two decimals in Table 1.
   - **Notes:** $60.9\% + 26.5\% + 12.6\% = 100.0\%$ (text). Table 1 values $60.90\% + 26.51\% + 12.59\% = 100.00\%$.

5. ✔ **Derived cumulative probability: positive current impact for Class 1** (Sec. 3.1 paragraph after Table 1, p.6)

   - **Claim:** Class 1 has cumulative probability 84.2% of reporting a positive current impact.
   - **Checks:** algebra/arithmetic from table
   - **Verdict:** PASS; confidence: high; impact: minor
   - **Assumptions/inputs:** Positive current impact = (Slightly Positive + Significantly Positive).
   - **Notes:** From Table 1 Class 1 current: 0.419 (Sl. Pos.) $+0.423$ (Sig. Pos.) = $0.842 = 84.2\%$.

6. ✔ **Derived cumulative probability: negative current impact for Class 3** (Sec. 3.1 paragraph after Table 1, p.6)

   - **Claim:** Class 3 has 92.4% probability of reporting a negative current impact.
   - **Checks:** algebra/arithmetic from table
   - **Verdict:** PASS; confidence: high; impact: minor
   - **Assumptions/inputs:** Negative current impact = (Significantly Negative + Slightly Negative).
   - **Notes:** From Table 1 Class 3 current: 0.222 (Sig. Neg.) $+0.701$ (Sl. Neg.) $= 0.923 = 92.3\%$, which rounds to 92.4% depending on rounding conventions.

7. ✔ **OR-to-percent increase conversion (examples)** (Sec. 3.4, p.8)

   - **Claim:** $\mathrm{OR} = 1.156$ corresponds to a 15.6% increase in odds; $\mathrm{OR} = 1.111$ corresponds to an 11.1% increase in odds.

- **Checks:** algebra/arithmetic
- **Verdict:** PASS; confidence: high; impact: minor
- **Assumptions/inputs:** Percent change in odds computed as $(\mathrm{OR} - 1) \times 100\%$.
- **Notes:** $(1.156 - 1) \times 100\% = 15.6\%$. $(1.111 - 1) \times 100\% = 11.1\%$.

8. ✔ **OR-to-percent decrease conversion (job loss fear example)** (Sec. 3.4, p.8)
   - **Claim:** $\mathrm{OR} = 0.174$ implies odds reduced by over 82%.
   - **Checks:** algebra/arithmetic
   - **Verdict:** PASS; confidence: high; impact: minor
   - **Assumptions/inputs:** Percent decrease in odds computed as $(1 - \mathrm{OR}) \times 100\%$.
   - **Notes:** $(1 - 0.174) \times 100\% = 82.6\%$, consistent with "over 82%."

9. ⚠ **Sign consistency of correlations with 'Negative Affect'** (Sec. 3.2, p.6)
   - **Claim:** Negative Affect is associated with deterrents (distrust, fear of job loss, privacy concerns) as validation evidence.
   - **Checks:** symbol/definition consistency, sign sanity check
   - **Verdict:** UNCERTAIN; confidence: medium; impact: moderate
   - **Assumptions/inputs:** Typical interpretation: higher negative affect should align with higher deterrence/concern measures, unless coding is reversed or factor sign is flipped.
   - **Notes:** Reported correlations are negative ($r = -0.246,\ -0.173,\ -0.127$). This can be consistent if deterrent variables are coded such that higher values mean less deterrence (or if the factor score sign is reversed), but the paper does not state coding/sign conventions, so the direction cannot be verified.

10. ⚠ **Interaction OR interpretation (buffering effect)** (Sec. 3.4, p.8 (also Fig. 4 caption p.9 mentions QKB_2_9_x_QF_7))
    - **Claim:** An interaction $\mathrm{OR} = 1.281$ indicates that freedom to choose AI tools buffers the negative effect of privacy concerns on optimism.
    - **Checks:** model-structure consistency, interpretation consistency
    - **Verdict:** UNCERTAIN; confidence: low; impact: moderate
    - **Assumptions/inputs:** Underlying model is logit with linear predictor including main effects and an interaction term., Privacy-concern main effect is negative on the optimistic-vs-reference log-odds, and tool-choice moderates it via a positive interaction.
    - **Notes:** Without the explicit multinomial logit equation and the signs/magnitudes of the main effects, the claimed 'buffering' direction cannot be checked. OR (interaction)> 1 is compatible with buffering only under specific coding and coefficient sign conditions that are not shown.

11. ⚠ **Elastic net 'zero coefficient' implies no effect** (Sec. 3.3, p.7)

    - **Claim:** A coefficient of zero for training implies training is insufficient to foster optimism.
    - **Checks:** logical implication check
    - **Verdict:** UNCERTAIN; confidence: medium; impact: minor
    - **Assumptions/inputs:** Elastic net coefficients reflect penalized estimation with feature selection.
    - **Notes:** In a penalized model, an exact zero indicates non-selection under the chosen penalty/tuning, not necessarily a true null relationship in the data-generating process. The stated substantive conclusion is stronger than the mathematical implication of coefficient shrinkage alone.

Limitations

- The provided PDF text contains essentially no explicit equations (no likelihoods, link functions, penalty terms, or probability mappings), so most derivation/algebra checks are not possible.
- Several key quantities (BIC, normalized entropy, McDonald's Omega, tetrachoric-based EFA, multinomial logit interactions) are referenced only by name; without formal definitions in the document, internal verification is necessarily limited to basic sanity checks on reported summaries.
- Figures (e.g., forest plot) are referenced, but the underlying coefficient tables and exact variable codings are not fully specified in the text provided, limiting sign/interpretation audits.

Numerical results audit

This section audits **numerical/empirical** consistency: reported metrics, experimental design, baseline comparisons, statistical evidence, leakage risks, and reproducibility.

All 19 candidate numeric checks passed. The audit verified internal arithmetic consistency for class proportions (summing to 100%), conditional-probability rows (summing to ~ 1 within rounding tolerance), derived cumulative percentages from component probabilities, OR-to-percent-change conversions (including an 'over 82%' inequality), repeated sample-size consistency ($N = 2,603$), a sample-size difference identity ($2603 - 2538 = 65$), small OR range stability across sensitivity specifications, and example VIF values being below 5.0. One item shows a minor text-vs-table rounding difference (92.4% reported vs 92.3% from displayed components) but remains within the stated tolerance.

Checked items

1. ✔ **C1__class__percentages__sum__100** (Page 5, Section 3.1 (bulleted class descriptions))

- **Claim:** The three latent class proportions are 60.9%, 26.5%, and 12.6% of the sample.
- **Checks:** percentage_sum_to_100
- **Verdict:** PASS
- **Notes:** $60.9 + 26.5 + 12.6 = 100.0$.

2. ✔ **C2_table1_class_percentages_sum_100** (Page 6, Table 1 (class sizes in parentheses))

   - **Claim:** Table 1 reports class sizes of (60.90%), (26.51%), (12.59%).
   - **Checks:** percentage_sum_to_100
   - **Verdict:** PASS
   - **Notes:** $60.90 + 26.51 + 12.59 = 100.00$.

3. ✔ **C3_table1_row_sums_current_class1** (Page 6, Table 1, Class 1 (Current Impact on Job Security))

   - **Claim:** Class 1 current-impact conditional probabilities should sum to 1 across response categories.
   - **Checks:** probabilities_sum_to_1
   - **Verdict:** PASS
   - **Notes:** Sum = 1.000 using displayed 3-decimal probabilities.

4. ✔ **C4_table1_row_sums_expected_class1** (Page 6, Table 1, Class 1 (Expected Impact on Job Security (3 years)))

   - **Claim:** Class 1 expected-impact conditional probabilities should sum to 1 across response categories.
   - **Checks:** probabilities_sum_to_1
   - **Verdict:** PASS
   - **Notes:** Sum = 1.000 using displayed 3-decimal probabilities.

5. ✔ **C5_table1_row_sums_current_class2** (Page 6, Table 1, Class 2 (Current Impact on Job Security))

   - **Claim:** Class 2 current-impact conditional probabilities should sum to 1 across response categories.
   - **Checks:** probabilities_sum_to_1
   - **Verdict:** PASS
   - **Notes:** Displayed values sum to 1.001; within rounding tolerance for 3-decimal reporting.

6. ✔ **C6_table1_row_sums_expected_class2** (Page 6, Table 1, Class 2 (Expected Impact on Job Security (3 years)))

   - **Claim:** Class 2 expected-impact conditional probabilities should sum to 1 across response categories.

- **Checks:** probabilities_sum_to_1
- **Verdict:** PASS
- **Notes:** Sum = 1.000 using displayed 3-decimal probabilities.

7. ✔ **C7_table1_row_sums_current_class3** (Page 6, Table 1, Class 3 (Current Impact on Job Security))

   - **Claim:** Class 3 current-impact conditional probabilities should sum to 1 across response categories.
   - **Checks:** probabilities_sum_to_1
   - **Verdict:** PASS
   - **Notes:** Sum = 1.000 using displayed 3-decimal probabilities.

8. ✔ **C8_table1_row_sums_expected_class3** (Page 6, Table 1, Class 3 (Expected Impact on Job Security (3 years)))

   - **Claim:** Class 3 expected-impact conditional probabilities should sum to 1 across response categories.
   - **Checks:** probabilities_sum_to_1
   - **Verdict:** PASS
   - **Notes:** Displayed values sum to 1.001; within rounding tolerance for 3-decimal reporting.

9. ✔ **C9_positive_current_impact_class1_cumulative** (Page 6, after Table 1 (text: 'Resiliently Optimistic ... cumulative probability of 84.2% of reporting a positive current impact'))

   - **Claim:** Individuals in the 'Resiliently Optimistic' class have a cumulative probability of 84.2% of reporting a positive current impact.
   - **Checks:** derived_percentage_from_components
   - **Verdict:** PASS
   - **Notes:** $(0.419 + 0.423) * 100 = 84.2\%$.

10. ✔ **C10_negative_current_impact_class3_cumulative** (Page 6, after Table 1 (text: 'Anxiously Declining ... 92.4% probability of reporting a negative current impact'))

    - **Claim:** Those in the 'Anxiously Declining' class have a 92.4% probability of reporting a negative current impact.
    - **Checks:** derived_percentage_from_components
    - **Verdict:** PASS
    - **Notes:** Displayed components imply 92.3% $(0.222 + 0.701)$, while text reports 92.4%; difference 0.1 percentage points, within abs_tol.

11. ✔ **C11_sample_size_consistency_2603** (Pages 1, 3, 10-11 (Abstract/Methods/Conclusions mention sample size))

- **Claim:** Sample size is stated as $2,603$ employees (also shown as $N = 2603$ in robustness).
- **Checks:** repeated_constant_consistency
- **Verdict:** PASS
- **Notes:** Both extracted sample-size statements equal 2603.

12. ✔ **C12_excluded_not_sure_count_difference** (Page 10, Section 3.6 (sensitivity analysis sample variants))

    - **Claim:** Full sample ($N = 2603$) vs excluding 'Not sure' respondents ($N = 2538$).
    - **Checks:** difference_check
    - **Verdict:** PASS
    - **Notes:** $2603 - 2538 = 65$ (implied excluded count).

13. ✔ **C13_or_to_percent_increase_1156** (Page 8, Section 3.4 (text: 'increased the odds by 15.6% (OR $= 1.156$)'))

    - **Claim:** Providing career progression incentives increased the odds by 15.6% (OR $= 1.156$).
    - **Checks:** or_to_percent_change
    - **Verdict:** PASS
    - **Notes:** $(1.156 - 1) * 100 = 15.6\%$ (to rounding).

14. ✔ **C14_or_to_percent_increase_1111** (Page 8, Section 3.4 (text: 'increased them by 11.1% (OR $= 1.111$)'))

    - **Claim:** Direct employee involvement in AI development increased odds by 11.1% (OR $= 1.111$).
    - **Checks:** or_to_percent_change
    - **Verdict:** PASS
    - **Notes:** $(1.111 - 1) * 100 = 11.1\%$ (to rounding).

15. ✔ **C15_or_to_percent_decrease_0174** (Page 8, Section 3.4 (text: 'reduced by over 82% (OR $= 0.174$)'))

    - **Claim:** Fear of job loss reduces odds by over 82% (OR $= 0.174$).
    - **Checks:** or_to_percent_change_inequality
    - **Verdict:** PASS
    - **Notes:** $(1 - 0.174) * 100 = 82.6\%$, which is strictly greater than 82%.

16. ✔ **C16_correlation_sign_consistency_negative_affect** (Page 6, Section 3.2 (correlations reported with Negative Affect))

    - **Claim:** Negative Affect is said to be significantly associated with deterrents, with correlations $r = -0.246$, -0.173, -0.127 (all $p < 0.001$).
    - **Checks:** basic_numeric_sanity_and_ordering

- **Verdict:** PASS
- **Notes:** All r values are within $[-1, 1]$, all are negative, and the magnitude ordering as written holds.

17. ✔ **C17_sensitivity_or_stability_qkb211** (Page 10, Section 3.6 (sensitivity analysis ORs for QKB_2_11))

    - **Claim:** Employee involvement (QKB_2_11) ORs are 1.111 (full), 1.119 (reduced), 1.120 (imputed).
    - **Checks:** relative_difference_small
    - **Verdict:** PASS
    - **Notes:** Range $= 1.120 - 1.111 = 0.009 \leq 0.02$.

18. ✔ **C18_sensitivity_or_stability_qf3** (Page 10, Section 3.6 (sensitivity analysis ORs for QF_3))

    - **Claim:** Job loss worry (QF_3) ORs are 0.174, 0.149, 0.175 across specifications.
    - **Checks:** relative_difference_small
    - **Verdict:** PASS
    - **Notes:** Range $= 0.175 - 0.149 = 0.026 \leq 0.03$.

19. ✔ **C19_vif_below_threshold_examples** (Page 10, Section 3.6 (VIF statement with examples))

    - **Claim:** All primary predictors had VIF values well below 5.0; examples: VIF $= 1.75$ and VIF $= 2.27$.
    - **Checks:** threshold_check
    - **Verdict:** PASS
    - **Notes:** Both example VIF values $(1.75, 2.27)$ are < 5.0.

Limitations

- Audit is restricted to the provided parsed text; numeric values shown only in figures (plots/bars/forest plots) are treated as non-extractable unless explicitly written in the text or tables.
- No raw survey data, model output tables, or coefficient matrices are included, preventing recomputation of fit statistics (BIC/entropy), McDonald's omega, correlations' p-values, regression OR confidence intervals, or Monte Carlo simulation results.
- Checks focus on cheap arithmetic consistency (sums, conversions between OR and percent change, and probability normalization) and cannot validate substantive causal interpretations.