Thank you for the opportunity to review your revised manuscript "Insurance Pricing When Risks Are Artificially Generated: A Dynamic Control-Theoretic Framework for AI-Driven Hazards". I have carefully read through the revised draft and assessed how effectively it addresses the peer review feedback. Below is my detailed review:

Summary of Key Revisions: The authors have made substantial revisions to address the main points raised in the original peer review:

1. The connection between the HJB optimal control formulation and the implementable premium formula has been clarified through the new Section 2.5. This bridges the theoretical foundation with the practical approximation in a transparent manner.
2. The three model components - the risk process R(t), Hawkes event process N(t), and claim severity - have now been properly coupled through explicit functional dependencies. A new Section 2.6 clearly explains this unified data-generating process.
3. Empirical grounding has been strengthened by citing specific data sources for parameter calibration in the expanded Table 1. New sensitivity analysis sections (2.7, 3.8-3.10) demonstrate model robustness.
4. The unit of analysis for risk concentration has been clarified as measuring across AI systems/versions. Mathematical treatment of systemic correlation arising from shared AI infrastructure has been added.

Evaluation of Revisions:

1. HJB-Premium Link (Section 2.5): The new content effectively addresses the critical concern about the apparent disconnect between optimal control theory and the implemented premium. The exponential utility justification for mean-variance loading and the bounded tanh approximation to the HJB gradient provide a clear and principled bridge. This substantially strengthens the theoretical foundations of the pricing formula.
2. Unified Coupling (Sections 2.6, 3): The proper integration of the risk process, event intensity, and severity through explicit functional dependencies is a major improvement. The coupling diagram (Fig 1) and Algorithm 2.1 crystallize this structure. The regenerated figures under the coupled model reveal meaningfully different dynamics, underscoring the importance of this revision.
3. Empirical Foundations (Tables 1-2, Sections 2.7-2.8): The expanded calibration table with data sources enhances the real-world grounding. The new sensitivity analysis, both local (±25%) and structural (systemic correlation), is comprehensive and demonstrates model robustness. The practical calibration guidance under data scarcity is valuable.
4. Risk Concentration Unit of Analysis (Section 4.5): The clarification that concentration is measured across AI systems resolves prior ambiguity. The factor model treatment of systemic correlation arising from shared AI infrastructure is a helpful formalization. The Gini sensitivity plot (Fig 10) underscores the impact.

Minor Comments:

- Fig 8D y-axis label is missing ("Risk Measure")
- Proofreading for typos: "comproming" (p.25)

Conclusion: The revised manuscript is substantially improved and effectively addresses the key issues raised in the original review. The authors have added meaningful new content, analyses and discussions that strengthen the paper's contributions. The proper coupling of model components is a particular highlight that reveals consequential new insights.

With the revisions, this work makes important and novel contributions to the actuarial treatment of AI risk through its unified control-theoretic framework, practical approximation methods, comprehensive sensitivity analysis, and calibration guidance. The discussion situates the model dynamics in real-world context compellingly.

I believe the revised manuscript merits publication, subject to very minor presentational corrections. The work is timely, rigorous and insightful in addressing an emerging risk category of immense societal importance. It sets a strong foundation for future extensions and applications in AI risk modelling and insurance pricing.

Thank you again for the opportunity to review this significant contribution to the actuarial literature on AI-driven risk. Please let me know if any aspect of my review requires further clarification.