AI Series #2

Why Human Judgment Matters in AI Decision Environments

lineslip

Why Human Judgment Matters in AI Decision Environments

AI-assisted systems are changing how risk professionals work. They surface patterns faster, process larger volumes of operational data, and accelerate analytical workflows that previously required significant manual effort.

What they do not change, however, is where accountability sits. Coverage adequacy decisions, renewal strategy, carrier negotiations, and board reporting still depend on human judgment. Only a person can truly determine what the data means in each context, whether it reflects operational reality, and what actions should follow.

As AI adoption scales, that judgment becomes more consequential, not less. The analytical work accelerates, but the governance responsibility attached to it does not disappear. It intensifies.

“As AI adoption scales, that judgment becomes more consequential, not less.”

What AI-Assisted Systems Cannot Determine on Their Own

Consider a practical scenario. A general-purpose AI platform flags an anomaly: a coverage limit that has not been updated in three years, while peer program benchmarks have moved significantly. The system has done exactly what it was built to do.

What it cannot determine is what the pattern means:

- Is the gap a coverage inadequacy?
- A deliberate risk retention decision made with full awareness of the benchmark spread?
- A data quality issue where the system is comparing programs that are not structurally equivalent?
- Or a negotiating position already in progress with the carrier?

Each interpretation leads to a different response. The analytical output is the same in every case. The governance judgment required to act on it correctly is entirely human, and it depends on knowledge of the program's history, the organization's risk appetite, the carrier relationship, and the broader strategic context. No amount of analytical sophistication provides that context. It has to be held by a person.

> **"No amount of analytical sophistication provides that context. It has to be held by a person."**

What Happens When Human Oversight Is Removed from a Dependency Chain

In September 2020, a ransomware attack crippled systems at Duesseldorf University Hospital. With critical systems offline and few fallback systems available, patients requiring urgent care experienced delays in treatment.

The governance lesson translates directly to risk programs: when a single-point dependency operates without human oversight capable of intervening before a failure cascades, the consequences compound quickly. In risk programs, that dynamic plays out whenever governance depends entirely on automated systems with no human checkpoints built into the dependency chain.

The gap between automated output and defensible governance decisions is exactly where human judgment needs to sit. Not to slow the process down. To make the output trustworthy.

Key Takeaways

- AI accelerates analytical workflows, but governance accountability remains human.
- As AI scales, the decisions requiring human judgment become more consequential.
- Technically correct output can still be incomplete in practice.
- Operational context cannot be automated. It has to be held by a person.
- Human judgment is what makes automated output trustworthy.
- Risk professionals become more valuable as AI adoption scales.

Human Judgment Becomes More Consequential as AI Scales

As AI-assisted workflows handle more analytical volume, the expertise and oversight required for effective governance become more critical, not less. The decisions left for human judgment increasingly become the ones that matter most: edge cases, anomalies, and situations where technically correct output is incomplete in practice.

Subject matter experts in risk management, insurance program structure, carrier relationships, and coverage adequacy become more valuable, not less, because they hold the operational context that determines what the data means. AI may enable strategy. It does not create strategy.

> **“AI may enable strategy. It does not create strategy.”**

The judgment required to evaluate AI-generated output against operational reality remains irreducibly human. Organizations that recognize that reality are the ones building governance programs that remain defensible as AI adoption scales.

The analytical work gets faster. The judgment work gets more consequential. The risk managers who understand that distinction are the ones building governance programs that stay defensible as AI adoption scales.

These are the conversations LineSlip continues to explore alongside risk and insurance leaders across the industry. More perspectives from the AI Series ahead.