AI in Project Management

From Risk Registers to Risk Intelligence [1]

Prasad S. Kodukula and Guz Vinueza [2]

A risk register by itself does not protect project value. Its value comes from the decisions it informs and the actions those decisions trigger.

This is the fourth article in our series on AI in project management. The first article introduced the evolving landscape of Generative AI (Gen AI), AI agents, and agentic AI; the second examined governance and introduced the AI Barbell Model; and the third explored applications in project planning. This article turns to project risk management and asks a broader question:

How can these capabilities help move organizations from maintaining risk registers toward developing risk intelligence?

Three AI capability levels are central to that evolution. Gen AI responds to prompts by generating and interpreting content, challenging assumptions, and exploring alternatives. AI agents perform defined tasks by retrieving information, monitoring conditions, invoking tools, and taking permitted actions within established boundaries. Agentic AI goes further by coordinating multiple agents, tools, systems, and information sources toward a broader objective.

These capabilities are particularly relevant to project risk management because the real objective is not simply to identify risks. It is to make better decisions under uncertainty. We define risk intelligence as:

The ability to transform risk-related data, signals, experience, and analysis into timely insights that improve decisions and actions under uncertainty.

AI-enabled risk intelligence strengthens that capability by helping teams sense uncertainty earlier, analyze its potential effects, evaluate alternatives, monitor changing

[1] How to cite this work: Kodukula, P. S. and Vinueza, G. (2026). From Risk Registers to Risk Intelligence. AI in Project Management, series article, *PM World Journal*, Vol. XV, Issue XI, November.

[2] Dr. Prasad Kodukula and Guz Vinueza are the authors of the best-selling book titled *The Project Management AI Handbook: Leveraging Generative Tools in Waterfall and Agile Environments,* published in 2025 by J. Ross Publishing. See their author profiles at the end of this article.

© 2026 Prasad S. Kodukula and Guz Vinueza

conditions, connect information across functional boundaries, and deliver timely decision support.

As Figure 1 illustrates, the evolution is from a largely static risk register, through quantitative analysis where appropriate, toward continuously updated risk intelligence. The final box—better decisions and actions—is not another level of risk intelligence. It represents what good risk intelligence is intended to enable.

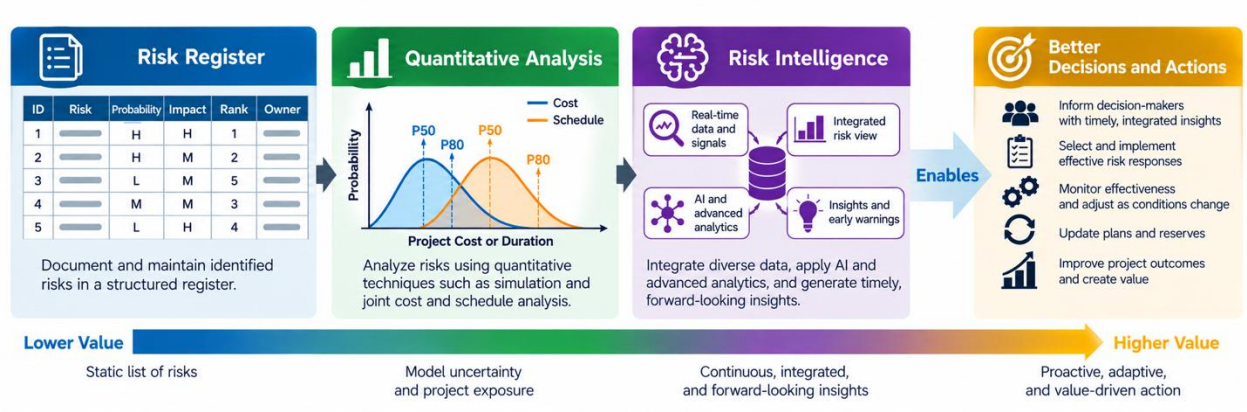


Figure 1. From Risk Registers to Risk Intelligence: Enabling Better Decisions and Actions

The objective is therefore not automated risk management. It is better information, better insight, and ultimately better human decisions.

Strategic Project Risk Management and AI-Enabled Risk Intelligence

AI can improve both the **efficiency** and **effectiveness** of project risk management.

Efficiency means performing established activities faster and with less effort: reviewing project information, identifying potential risks, analyzing and ranking them, monitoring indicators, developing responses, preparing reports, and more.

Effectiveness asks a more important question: does that work help the project recognize uncertainty earlier, make better-informed decisions, reduce cost and schedule exposure, capture opportunities, manage reserves appropriately, and respond more resiliently when uncertainty becomes reality?

That broader objective requires a strategic perspective.

© 2026 Prasad S. Kodukula and Guz Vinueza

Strategic Project Risk Management links uncertainty to project objectives, value, decision rights, management attention, and resilience. It asks not only which risks belong in a register, but which uncertainties matter most, how they may interact, what decisions they require, what responses are warranted, how much contingency is justified, when escalation is necessary, and how the project should recover when uncertainty becomes reality.

Traditional risk registers frequently emphasize **event risks**—discrete threats or opportunities that may or may not occur. Event risks are important, but they represent only one form of project uncertainty. A more complete perspective recognizes four types.

Four Types of Project Risk

Variability risk exists when variation from an estimate is virtually certain, but its magnitude and direction are uncertain. Activity durations, productivity rates, quantities, costs, equipment lead times, and other estimates rarely unfold exactly as planned.

Event risk exists when a discrete threat or opportunity may or may not occur and, if it does, may affect one or more project objectives.

Ambiguity risk arises from insufficient knowledge, unclear requirements, immature technology, uncertain assumptions, or multiple plausible interpretations.

Emergent risk consists of the **"unknown unknowns"**—risks that cannot reasonably be identified or foreseen at the outset but emerge as the project or its environment evolves.

The four types require different analytical and management approaches, as summarized in Table 1.

There is also a dimension that cuts across all four types.

Risks are frequently managed within functional silos—procurement, engineering, construction, regulatory affairs, project controls, operations, and others. Yet several individually manageable uncertainties may interact and produce a much larger project-level exposure.

We use the term **Convergent Risk** for this combined effect.

The four risk types describe the nature of individual uncertainties. Kodukula has examined each of these four risk types individually in a series of PM+ blogs and brought them together in a subsequent synthesis, *Four Types of Project Risks You Can't Afford to Ignore*, with examples and approaches for managing each type.

© 2026 Prasad S. Kodukula and Guz Vinueza

Table 1. Four Types of Project Risk and Representative Approaches

Risk Type	Nature of Uncertainty	Representative Approaches	Potential AI Contributions
Variability	Outcome will vary; magnitude and direction are uncertain	Three-point estimating, ranges, Monte Carlo simulation, sensitivity analysis	Identify or challenge ranges, analyze historical performance, identify changing patterns, interpret simulation results
Event	Threat or opportunity may or may not occur	Probability-impact analysis, Expected Monetary Value (EMV), simulation, quantitative risk analysis	Identify events and triggers, challenge probabilities and impacts, monitor indicators, develop responses
Ambiguity	Insufficient knowledge or multiple plausible interpretations	Scenarios, decision trees, testing, prototyping, rolling-wave planning	Surface assumptions and contradictions, identify knowledge gaps, develop scenarios and decision maps
Emergent	"Unknown unknowns" become visible as conditions evolve	Vigilance, stress testing, scenario planning, resilience, rapid reassessment	Monitor weak signals, detect anomalies, assemble evidence, support rapid integrated analysis

Convergent Risk describes what can happen when those uncertainties interact. It is not a fifth risk type, but a cross-cutting project-level effect that may arise from any combination of variability, event, ambiguity, and emergent risks.

A supplier delay might be an event risk. Lower-than-planned construction productivity is a variability risk. Uncertainty regarding an immature technology may be an ambiguity risk. An unforeseen regulatory development may be an emergent risk. Any combination could converge on the same milestone, resource constraint, budget threshold, or project objective. Figure 2 illustrates this relationship.

This broader view matters because one of AI's greatest potential contributions is its ability to connect information that is otherwise fragmented across the project.

Using SAFER to Examine AI's Contribution

AI does not replace established project risk-management practices. It can strengthen them. For convenience, we use **SAFER** as a mnemonic for the familiar risk-management workflow:

- **S — Sense uncertainty**
- **A — Analyze exposure**
- **F — Formulate decisions and responses**
- **E — Execute**

© 2026 Prasad S. Kodukula and Guz Vinueza

- **R — Review, Recover, and Refine**

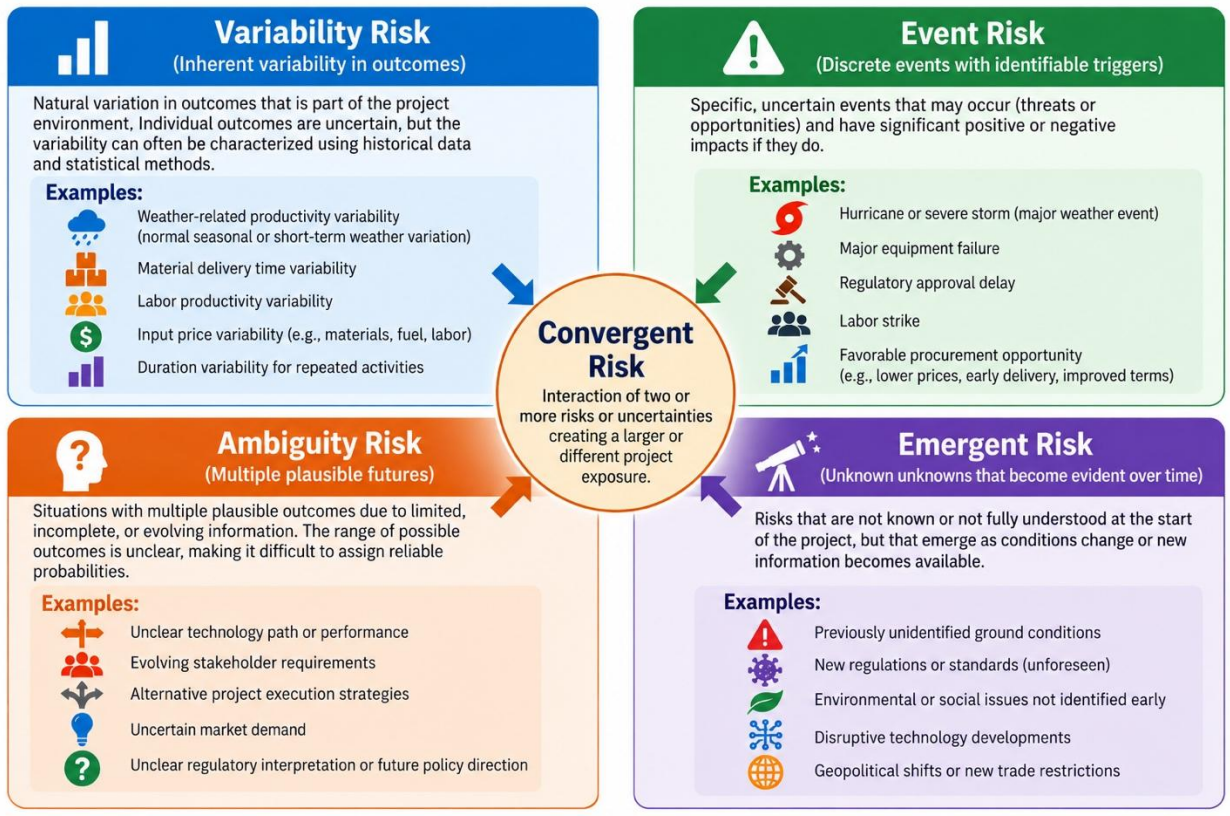


Figure 2. Four Types of Project Risk and Convergent Risk

SAFER is **not a new risk-management methodology**. It simply provides a structure for examining where Gen AI, AI agents, and agentic AI can add value.

S — Sense Uncertainty: Seeing Risk Earlier

Risk intelligence begins with recognizing uncertainty early enough to act.

Traditional identification techniques—workshops, interviews, checklists, lessons learned, expert judgment, and assumptions reviews—remain important. AI adds scale, speed, and continuity.

Gen AI can review objectives, scope, schedules, estimates, assumptions, requirements, contracts, lessons learned, and other project information to help build or challenge a risk register. It can ask whether important threats or opportunities are missing, assumptions

© 2026 Prasad S. Kodukula and Guz Vinueza

are weak, requirements conflict, dependencies create exposure, or multiple entries describe the same underlying uncertainty.

AI agents add continuous monitoring. They can watch supplier milestones, schedule float, cost trends, regulatory correspondence, technical-test results, resource availability, weather conditions, and other indicators and alert the team when conditions change.

Agentic AI can connect those signals across systems and functions.

The signals vary by risk type. Variability may show up as widening duration ranges or declining productivity. Event risks may have identifiable triggers. Ambiguity may manifest as contradictory requirements, unsupported assumptions, or unresolved technical questions. Emergent risk becomes visible only after an unknown condition begins producing signals.

AI cannot predict a true unknown unknown before meaningful evidence exists. Its value lies in shortening the interval between early signals and the recognition that something important has changed.

It can also help expose Convergent Risk. Procurement may see a supplier delay, construction may see reduced productivity, and regulatory specialists may see a possible permit delay. Individually, each may appear manageable. Together, they may threaten the same milestone.

But AI's ability to generate possibilities introduces an important caution.

A plausible-sounding threat is not automatically a project risk. An apparent pattern is not automatically meaningful. If supplier history is unavailable, AI should identify the information gap rather than invent performance history. If regulatory timing is uncertain, it should not manufacture a project-specific probability and present it as fact.

AI should help you search for uncertainty—not manufacture it.

A — Analyze Exposure: Understanding What the Uncertainty Means

Once uncertainty has been identified, the next question is:

What could it mean for the project?

AI can help organize evidence, challenge assumptions, retrieve current information, invoke analytical tools, interpret results, and integrate analysis across cost, schedule,

© 2026 Prasad S. Kodukula and Guz Vinueza

procurement, technical performance, resources, and other project domains. The analytical method should reflect the type of uncertainty.

For **variability risk**, three-point estimates can represent optimistic, most-likely, and pessimistic values for uncertain durations, costs, quantities, productivity rates, or other parameters. Monte Carlo simulation can combine those ranges into distributions of project outcomes, while sensitivity analysis identifies the inputs that contribute most to uncertainty.

AI can help challenge ranges, compare them with historical information, identify inconsistent assumptions, and interpret the results. It should not generate unsupported numerical ranges merely because a model requires them.

For **event risks**, probability-impact analysis can support qualitative prioritization. Expected Monetary Value can support decisions when probabilities and financial consequences are defensible. Simulation can incorporate discrete events alongside underlying cost and schedule variability.

For major projects, integrated techniques such as Joint Cost and Schedule Confidence Level (JCL) analysis can provide management with a combined view of the likelihood of meeting both cost and schedule commitments.

AI can help structure these analyses and challenge assumptions: What evidence supports the probability? Is the impact conditional on the event occurring? Is the same consequence counted elsewhere? Does the proposed response affect probability, impact, or both?

A key principle is: **Precision is not evidence.**

An AI-generated probability of 37 percent is not inherently more credible than an experienced professional's qualitative judgment simply because it appears precise. The number's credibility depends on its evidence and assumptions.

For **ambiguity risk**, meaningful probabilities may not yet be available. The ambiguity may be technical or operational—for example, incomplete requirements, immature technology, or an untested assumption. It may also be strategic.

In project risk management, Kodukula has used the term **strategic ambiguity** to describe situations in which leaders face multiple plausible strategic paths but lack sufficient information to know with confidence which path will create the greatest value. Such uncertainty is better addressed through scenarios, staged decisions, and learning than

© 2026 Prasad S. Kodukula and Guz Vinueza

by forcing premature probabilities onto outcomes that are not yet well understood. *(See Kodukula, "3G Project Risk Management: Five Key Characteristics.")*

Decision trees can examine alternative decisions followed by uncertain outcomes. Scenario analysis can explore several plausible futures without pretending that the project knows exactly which one will occur. Testing, prototyping, or rolling-wave planning can provide additional information before an irreversible commitment is made.

AI can help surface assumptions, identify contradictory information, develop alternative scenarios, and create decision maps that show available choices, plausible outcomes, dependencies, and information gaps.

The questions are often:

What do we know? What do we not know? What alternatives are available? What outcomes could follow? What additional information would materially improve the decision?

For **emergent risk**, analysis begins only after a new condition becomes visible. Scenario planning, stress testing, and rapid integrated analysis can help establish what changed, how quickly the consequences may spread, and which objectives may be affected.

AI agents can quickly assemble relevant information. Agentic AI can coordinate cost, schedule, technical, resource, procurement, and external information to create an initial integrated view and support escalation when the situation exceeds project authority.

Analysis is especially important for **Convergent Risk**. A supplier delay may appear acceptable because float exists. A productivity decline may also appear manageable. A regulatory delay may remain within tolerance. Yet the supplier delay may consume part of the float, productivity loss may consume the remainder, and regulatory approval may then become controlling.

Agentic AI can help reveal such interactions by combining information that individual risk owners may never see together.

More sophisticated analysis, however, does not automatically mean better analysis. Data sources, assumptions, dependencies, correlations, and model logic must remain visible and subject to professional judgment.

Ultimately, analysis exists to answer management questions:

- Is the exposure acceptable?

© 2026 Prasad S. Kodukula and Guz Vinueza

- Do we need to act?
- Do we need more information?
- How much contingency is justified?
- Should the plan change?
- Should the issue be escalated?
- Which alternative offers the best balance of cost, schedule, opportunity, and residual exposure?

The goal is not more analysis, but better intelligence for deciding what to do.

F — Formulate Decisions and Responses

Once exposure is understood, the project must determine an appropriate response.

For event risks, threats may be avoided, mitigated, transferred, accepted, or escalated; opportunities may be exploited, enhanced, shared, accepted, or escalated. The real work lies in translating a strategy into specific actions and comparing alternatives by cost, feasibility, implementation time, resource requirements, effects on probability or impact, and possible secondary risks.

A project with numerous significant risks can generate dozens of possible actions. Gen AI can expand the set of alternatives, compare advantages and disadvantages, challenge classifications, and rank options against criteria established by the project team.

AI can also evaluate responses across several risks. A **risk-response matrix**, for example, can show whether a candidate action reduces, increases, or has little effect on multiple risks. An action that addresses several important exposures may create greater project value than one focused narrowly on a single risk.

Agents can check whether proposed responses are feasible using current information on resources, funding, schedules, suppliers, contracts, and regulatory constraints. Agentic AI can compare combinations of responses across functions.

The appropriate response also depends on the type of uncertainty. Variability may be addressed through better estimates, flexibility, contingency, sequencing, or resource adjustments. Ambiguity may call for more information before commitment through testing, prototyping, clarification, or phased decisions. Emergent risk may require rapid scenario development and escalation.

Responses also need to be evaluated at the project level.

© 2026 Prasad S. Kodukula and Guz Vinueza

Accelerating procurement may reduce schedule exposure but increase cost. Adding construction resources may reduce duration but create congestion and lower productivity. Resequencing may protect one milestone while creating another technical or regulatory dependency.

The question is not only, **"What is the best response to this risk?"** It is also, **"What does this response do to the rest of the project?"**

AI can expand the alternatives and expose trade-offs. It does not own the decision.

The appropriate risk owner, project manager, sponsor, or governance body remains responsible for selecting the response and accepting any residual exposure.

E — Execute: Turning Responses into Effective Controls

Selecting a response does not reduce risk. The response must be implemented and shown to work. A useful progression is:

Risk → Response → Action → Control → Effectiveness → Residual Exposure

A response defines what the project intends to do. An action implements it. A control is the operating measure intended to prevent, detect, or limit the cause or consequence. Effectiveness asks whether that control is actually changing the exposure.

Gen AI can translate responses into specific actions, owners, completion criteria, monitoring requirements, triggers, and evidence of effectiveness.

AI agents are particularly useful here because execution requires continuing attention. They can track whether actions have owners, due dates, funding, schedule links, and approvals, and monitor whether the intended results are occurring.

Suppose a supplier-delay risk is addressed through increased expediting and more frequent fabrication reviews. Holding the meetings does not prove the risk has been reduced. Evidence might include improved fabrication progress, recovery of schedule float, achievement of interim milestones, or increased confidence in the delivery date.

AI can therefore help move risk management from **tracking completed actions to monitoring whether those actions actually reduce exposure**.

© 2026 Prasad S. Kodukula and Guz Vinueza

Agentic AI adds the ability to evaluate response effects across the project. A procurement action may affect cost. A schedule response may affect resources. A design change may create regulatory or construction exposure elsewhere.

When risks converge, responses may converge as well. Individually reasonable actions can compete for the same resources, consume the same contingency, or create new constraints. The central question is simple: **Is the response actually reducing the exposure?**

R — Review, Recover, and Refine

Risk management is not a one-time exercise. It is a **continuous process that begins during project initiation and continues through planning, execution, monitoring and control, and closeout**.

Conditions change. Risks disappear, increase, or occur. New information changes assumptions. Responses succeed or fail. New uncertainty emerges.

Review

Review asks whether the project's current understanding of uncertainty still reflects reality.

Agents can continuously monitor schedule, cost, supplier, technical, regulatory, resource, and risk information. Agentic AI can integrate changes across functions and identify when a local development alters project-level exposure.

Review is also the point at which **contingency reserves** should be reassessed.

Contingency is established for identified uncertainty. As risks change, occur, are retired, or are mitigated, the amount and timing of contingency required may also change.

AI can link current risk assessments to cost and schedule forecasts, track which risks remain active, evaluate residual exposure, and identify where contingency may be insufficient or no longer needed.

AI does not decide whether reserves should be consumed or released. It provides evidence for the authorized decision-maker.

© 2026 Prasad S. Kodukula and Guz Vinueza

Emergent risk requires different treatment.

Because emergent risks are unknown unknowns, they cannot be specifically included in identified-risk contingency at the outset. When such a risk becomes visible and exceeds the resources or authority available within the project, **management reserve** may need to be considered.

AI can help assess the new condition, estimate potential cost and schedule effects, develop response scenarios, and prepare the information required for sponsor or senior-management decisions.

The authority distinction remains important: project managers may manage contingency within established limits, whereas use of management reserve generally requires higher-level authorization.

Recover

When uncertainty materializes into an issue, the questions change. Before a risk event occurs, the team may ask: *What is the probability and potential impact?*

After occurrence, the event is no longer uncertain. The relevant questions become:

- What are the actual consequences?
- What uncertainty remains?
- What should we do now?

AI can rapidly assemble current information, update forecasts, distinguish confirmed facts from remaining uncertainty, and help develop recovery alternatives.

Recovery does not necessarily mean restoring the original plan at any cost. It means identifying the best path forward under the new conditions.

Actions may include resequencing work, reallocating resources, accelerating selected activities, changing procurement strategies, revising scope, using contingency or management reserve, renegotiating commitments, or modifying milestones.

AI can help generate and compare alternatives; agents can test their feasibility; agentic AI can evaluate their cross-functional effects.

The recovery decision itself remains human.

© 2026 Prasad S. Kodukula and Guz Vinueza

Refine

Every significant risk event, response, control failure, successful intervention, reserve decision, and recovery action creates information that should improve future risk management.

Actual performance can refine variability ranges. Expected and actual outcomes can improve event-risk assessments. Ambiguity analysis can reveal which information reduced uncertainty most effectively.

For emergent risk, an especially useful question is: **When did the first meaningful signals appear, and could we have recognized them sooner?**

Convergent Risk raises similar questions. If individually manageable risks combined into a much larger exposure, why was the interaction not visible earlier? Were data fragmented? Were assumptions inconsistent? Were risks managed too narrowly within silos? Did one response create exposure elsewhere?

AI can help capture and reuse these lessons systematically while the project is still underway. What the project learns today should improve what it senses tomorrow.

Challenges and Guardrails for AI-Enabled Risk Management

AI can strengthen risk management while introducing its own risks. As AI moves from generating content to monitoring systems, invoking analytical tools, coordinating agents, and preparing actions, the consequences of error increase. The principal challenges and corresponding guardrails are summarized in Table 2.

The Orchestration Challenge

AI orchestration deserves particular attention because increasingly capable risk systems may involve multiple models, agents, analytical tools, enterprise systems, data sources, and human decision-makers.

One agent may monitor suppliers. Another may analyze schedule effects. Another may estimate cost consequences. Another may track regulatory information. Another may monitor resource availability.

The challenge is not merely whether each component performs correctly. It is whether the entire workflow is coordinated correctly:

© 2026 Prasad S. Kodukula and Guz Vinueza

- Which agent performs which task?
- In what sequence?
- Which data source is authoritative?
- What information passes between agents?
- How are conflicting outputs reconciled?
- Which tools and systems may an agent access?
- What permissions does it have?
- When should the workflow stop and escalate?
- Which actions may AI prepare but not execute?
- How is the entire chain kept traceable?

Table 2. Key Challenges and Guardrails for AI-Enabled Risk Management

Challenge	Why It Matters in Risk Management	Key Guardrails
False precision and unsupported inference	AI may generate precise-looking probabilities, impacts, contingency estimates, or causal relationships without adequate evidence.	Validate assumptions and sources; distinguish judgment from fact; require qualified review of material risk analyses.
Data quality, currency, and context	Risk assessments can be misleading when AI uses outdated schedules, stale risk registers, incomplete cost data, obsolete assumptions, or fragmented project information.	Use current, authoritative project data; reconcile conflicting sources; flag missing information and changed assumptions.
Automation bias	Teams may begin accepting AI-generated risk rankings, response recommendations, or reserve estimates without sufficient challenge.	Preserve expert judgment; require human review of material risks, responses, reserve decisions, and risk acceptance.
AI orchestration and cascading errors	An error in one agent can propagate into schedule impacts, cost forecasts, response recommendations, and reserve decisions.	Define agent roles, sequencing, handoffs, permissions, escalation rules, and human intervention points.
Traceability, security, and accountability	Significant risk decisions should be traceable from evidence and assumptions through analysis, response, reserve use, escalation, and approval.	Maintain audit trails, access controls, decision rights, approval thresholds, risk ownership, and accountability.

Poor orchestration can produce cascading errors.

Suppose one agent incorrectly interprets a regulator's request for additional information as a permit rejection. A schedule agent might then correctly calculate the resulting delay.

© 2026 Prasad S. Kodukula and Guz Vinueza

A cost agent might correctly estimate prolongation costs. Another agent might correctly prepare a recovery scenario.

Every downstream calculation could be correct while the entire analysis rests on a false premise.

As AI becomes more agentic, the reliability of **handoffs, permissions, routing, escalation rules, and human intervention points** becomes as important as individual model performance.

Applying the AI Barbell Model to Risk Intelligence

The challenges above reinforce the principle behind the **AI Barbell Model**, introduced in Article 2 of this series. AI creates the greatest value when strong inputs on one side are matched by appropriate human controls on the other.

Rather than repeat the original Barbell graphic, **Figure 3 applies that logic specifically to risk intelligence**. It connects project information, expertise, and signals to AI capabilities and then to human decision controls, with governance spanning the entire system.

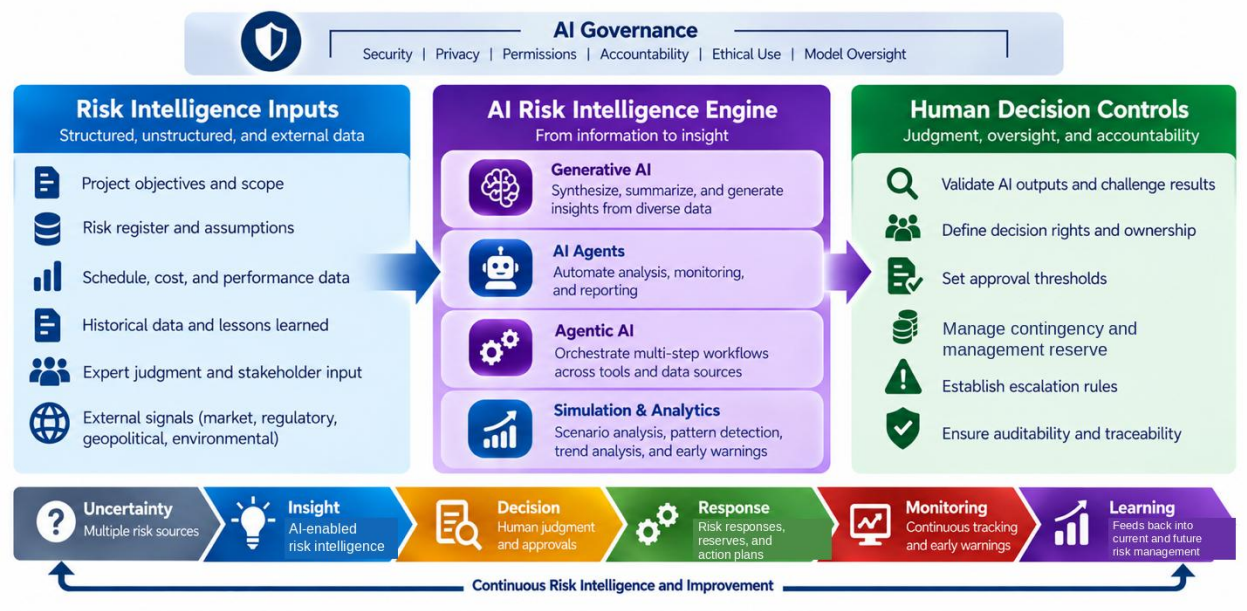


Figure 3. AI-Enabled Risk Intelligence: Inputs, AI Capabilities, and Human Controls

© 2026 Prasad S. Kodukula and Guz Vinueza

The model highlights two points particularly relevant to risk management.

First, AI can analyze only the uncertainty it can see and understand. An agentic system cannot identify interactions among procurement, schedule, cost, resource, technical, and regulatory conditions if the relevant information is inaccessible, outdated, inconsistent, or disconnected. Likewise, better analytical capability cannot compensate for unsupported duration ranges, probabilities, correlations, or response assumptions.

Second, human controls become increasingly important as AI moves from generating information to recommending and potentially preparing actions.

Gen AI may draft a risk register or explain simulation results. AI agents may monitor thresholds, invoke tools, or update draft records. Agentic AI may integrate information across domains, compare response packages, and prepare coordinated actions.

As autonomy rises, the consequences of error can rise as well. The principle introduced earlier in this series therefore applies directly:

Higher autonomy + lower reversibility = stronger—and sometimes different—controls.

An AI system that summarizes a risk workshop may need relatively light review. An agent monitoring supplier milestones and generating an internal alert may operate with greater autonomy. An agentic system preparing a baseline change, management-reserve request, supplier instruction, contractual commitment, or regulatory communication requires much stronger controls and explicit approval before execution.

AI may identify risks, quantify exposure, recommend a response, monitor a trigger, or prepare an action. It does **not** accept risk on behalf of the organization.

Risk acceptance, reserve decisions, baseline changes, contractual commitments, regulatory submissions, and other consequential actions remain with the people or governance bodies authorized to make them.

Project Pinot: From Prompt-Based Gen AI to Risk Intelligence

Project Pinot, introduced in Article 3, is a fictitious industrial wastewater-treatment project for Acme Chemical Corporation in Baton Rouge, Louisiana. The project has a budget of approximately $3 million and a planned duration of about 15 months (65 weeks). It involves the OmniTreat treatment technology and a treatability study, an effluent-

© 2026 Prasad S. Kodukula and Guz Vinueza

discharge permit required before operation, custom equipment with long procurement lead times, and exposure to hurricane-related construction disruption.

These characteristics make Pinot useful for illustrating variability, event, ambiguity, emergent, and Convergent Risk.

In our book, *The Project Management AI Handbook: Leveraging Generative Tools in Waterfall and Agile Environments*, we used Project Pinot to demonstrate how **Gen AI can support the traditional project risk-management workflow**.

Through a series of structured prompts for individual risk processes and workflows, we developed the risk register, identified threats and opportunities, assessed and ranked risks, developed response actions, and supported other risk-management activities. That work demonstrated how prompt-based Gen AI can improve the efficiency and thoroughness of established risk practices.

Here we extend the same case further. The question is no longer simply: **How can Gen AI help perform individual risk-management activities?** It becomes: **How can Gen AI, agents, and agentic AI work together to create continuing risk intelligence?**

That progression moves from prompt-based assistance to continuous sensing, monitoring, integrated analysis, cross-functional orchestration, and timely decision support.

Pinot contains all four forms of uncertainty.

Construction productivity, activity durations, quantities, costs, and equipment lead times create **variability risk**. Permit delays, late delivery of custom equipment, hurricane disruption, and procurement opportunities represent **event risks**. The limited full-scale operating experience of the OmniTreat technology creates **ambiguity**, which the treatability study is intended to reduce. Unforeseen regulatory, environmental, technical, supply-chain, or other conditions may emerge as **unknown unknowns**.

The larger value of AI-enabled risk intelligence becomes especially apparent when these uncertainties interact.

Suppose the custom-equipment supplier begins slipping fabrication milestones. Procurement concludes that the delay can be absorbed by available schedule float.

At the same time, construction productivity is below the original estimate, but the construction team expects to recover.

© 2026 Prasad S. Kodukula and Guz Vinueza

Meanwhile, the regulatory team sees indications that permit review could take longer than anticipated.

Specialist resources required for installation and commissioning may also have limited availability during the affected period.

Each condition may appear manageable independently.

Together, they may not be.

The equipment delay consumes some float. Lower construction productivity consumes more. The permit path may then become controlling just as installation and commissioning approach a period of constrained specialist availability.

The project may now face a substantially higher cost, schedule, and commissioning exposure than any individual risk owner sees within a functional silo. **This is Convergent Risk.** Figure 4 illustrates the interaction.

This is also where agentic AI can provide particular value.

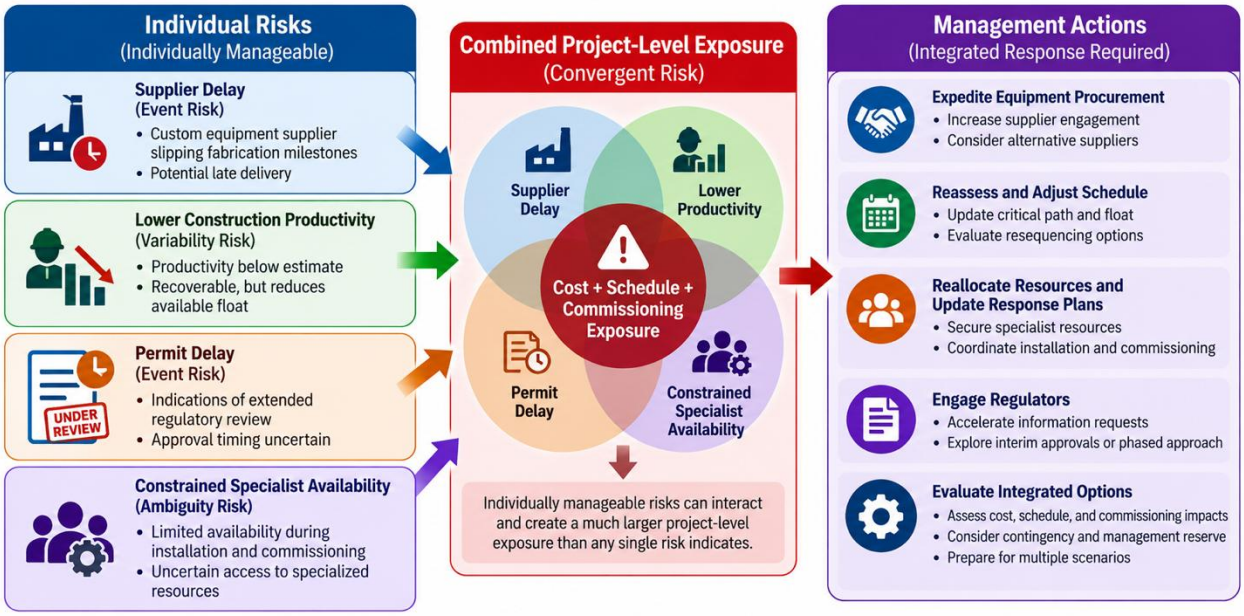


Figure 4. Project Pinot—Convergent Risk in Practice

A procurement agent can monitor supplier milestones. A schedule agent can track float and critical-path changes. A cost agent can estimate prolongation effects. A resource agent can monitor specialist availability. A regulatory agent can follow the permit process.

© 2026 Prasad S. Kodukula and Guz Vinueza

Through appropriate orchestration, those inputs can be combined to determine how conditions interact, which project objectives are threatened, what response alternatives exist, and what the combined exposure may be.

The objective is not for the system to make the decision. Its value is to make integrated risk intelligence available early enough for management action.

The same principle applies to responses. Expediting equipment may reduce schedule exposure but increase cost. Adding resources may reduce duration while creating congestion. Resequencing work may protect one milestone while creating another dependency.

AI can help evaluate those interactions before the project commits to an action.

As Pinot progresses, AI can also support reserve management. Agents can monitor whether the risks underlying cost and schedule contingency remain active and whether responses have reduced residual exposure. If an emergent condition exceeds planned contingency and project authority, AI can help assemble the evidence needed to consider management reserve.

Project Pinot therefore illustrates the progression at the heart of this article:

- Gen AI can improve individual risk-management activities.
- AI agents can make those activities more continuous.
- Agentic AI can connect and orchestrate information across the project.

Together, these capabilities can move risk management from maintaining a register toward developing **risk intelligence**—and that intelligence can enable better decisions and actions.

Conclusion: Better Risk Intelligence for Better Decisions

Project risk management is ultimately about making better decisions under uncertainty.

A risk register remains useful, but documenting risk is only the beginning. Quantitative analysis can deepen understanding of exposure. AI can extend that capability by sensing changing conditions, integrating information across functional boundaries, evaluating more alternatives, monitoring response effectiveness, and helping project leaders recognize interactions that might otherwise remain hidden.

© 2026 Prasad S. Kodukula and Guz Vinueza

The four risk types reinforce the need for this broader perspective. Variability, event, ambiguity, and emergent risk require different approaches. Convergent Risk reminds us that individually manageable uncertainties can combine into much larger project-level exposure.

Risk management must also remain continuous. Assumptions change. Responses change exposure. Contingency requirements change. New risks emerge. Some uncertainties materialize and become recovery problems rather than probability questions. The lessons from those experiences should improve what the project senses next.

AI can make that process more efficient and more effective, but greater capability also introduces new risks. False precision, unsupported inference, stale data, automation bias, orchestration failures, cascading errors, security concerns, and unclear accountability can undermine the very intelligence AI is intended to improve.

That is why human judgment does not become less important as AI capability increases. It becomes more important in different places.

Strong inputs, effective orchestration, appropriate controls, clear decision rights, traceability, and governance are essential if AI-supported analysis is to become trusted risk intelligence.

AI can expand the evidence. It can accelerate the analysis, connect information across silos, illuminate alternatives, and monitor changing conditions continuously.

But the project manager, risk owner, sponsor, or appropriate governance body still decides what action to take and remains accountable for the outcome.

The goal is not automated risk management.

It is better risk intelligence for better project decisions.

In the next article in this series, we move from managing uncertainty to managing performance—examining how Gen AI, AI agents, and agentic AI can support project execution, monitoring, and control as the project moves from plan to results.

Authors' Note

Portions of this article are adapted from our recent book, The Project Management AI Handbook: Leveraging Generative Tools in Waterfall and Agile Environments. We used ChatGPT extensively as a writing and editing tool in preparing this manuscript.

© 2026 Prasad S. Kodukula and Guz Vinueza

However, the ideas, interpretations, analyses, recommendations, and conclusions presented here are solely our own.

References

Kodukula, P. S., & Vinueza, G. (2026a). The Landscape, the Opportunity, and the Limits. AI in Project Management, Series Article 1. PM World Journal, Vol. XV, Issue VIII, August. Available online at https://pmworldjournal.com/wp-content/uploads/2026/08/pmwj167Aug2026-Kodukula-Vinueza-AI-in-PM-article-1.pdf

Kodukula, P. S., & Vinueza, G. (2026b). The AI Barbell Model: Balancing Value Creation and Human Control in Project Management. AI in Project Management, Series Article 2. PM World Journal, Vol. XV, Issue IX, September. Available online at https://pmworldjournal.com/wp-content/uploads/2026/09/pmwj168-Sep2026-KodukulaVinueza-Barbell-Model-AI-in-PM-article-2.pdf

Kodukula, P. S., & Vinueza, G. (2026c). AI for Project Planning: From Plausible Artifacts to a Coherent Plan. AI in Project Management, Series Article 3. *PM World Journal*, Vol. XV, Issue X, October. Available online at https://pmworldjournal.com/wp-content/uploads/2026/10/pmwj169-Oct2026-Kodukula-Vinueza-AI-for-Project-Planning.pdf

Kodukula, P. S., & Vinueza, G. (2025). *The Project Management AI Handbook: Leveraging Generative Tools in Waterfall and Agile Environments*. J. Ross Publishing.

Kodukula, P. S. Blogs (2021):
https://kodukula.com/four-types-of-project-risks-you-cant-afford-to-ignore/
https://kodukula.com/managing-variability-risk-a-virtually-guaranteed-risk/
https://kodukula.com/ambiguity-risk-what-is-it-how-to-manage-it/
https://kodukula.com/six-key-strategies-to-manage-known-unknowns/
https://kodukula.com/building-a-resilient-organization-to-survive-even-thrive-unknown-unknowns-2/
https://kodukula.com/3g-project-risk-management-five-key-characteristics/

© 2026 Prasad S. Kodukula and Guz Vinueza

About the Authors



Dr. Prasad S. Kodukula

Illinois, USA



Dr. Prasad S. Kodukula, PMP, PgMP, PMI-ACP, DASM, DASSM, BCES, is a USA Today and Amazon #1 best-selling author, PMI Fellow, thought leader, and entrepreneur with over 35 years of professional experience. A global ambassador for project management, Dr. Kodukula has lectured in nearly 50 countries and worked with more than 40 Fortune 100 companies (e.g., Allstate, Abbott Labs, BP, Caterpillar, Dow, Chrysler, IBM, JP Morgan Chase, Motorola, Stryker, United Technologies, Whirlpool) across all 11 S&P industrial sectors. He serves as an Adjunct Industry Professor at Illinois Tech. He teaches how to recover troubled projects at NASA. He has delivered project management training in professional education programs at Stanford University, Duke University, and the University of Chicago. As co-founder and CEO of Kodukula & Associates, Inc. and NeoChloris, Inc., he leads the firms in project management and renewable energy, respectively.

His work with AI began in 1988, when he created award-winning expert systems in environmental engineering. In the late 1990s, he co-invented an AI-powered smart monitoring and control system for water and wastewater treatment plants, combining early neural network, sensor, and automation technologies. This invention received a U.S. patent in 2005 and was later licensed to prominent control and automation firms.

Recognized three times by the Project Management Institute as “Best of the Best in Project Management,” he has received multiple accolades, including the Illinois Tech Alumni Association Professional Achievement Award and honors from the U.S. Environmental Protection Agency and the states of Arizona, Kansas, and Illinois for his outstanding leadership in education and training, environmental improvement, and innovation. Most recently, he was recognized by Project Controls Academy with a Project Controls Excellence Award. An accomplished author, Dr. Kodukula has co-authored or contributed to 12 books and over 40 articles and holds four U.S. patents. He can be contacted at: https://www.linkedin.com/in/prasadkodukula/.

© 2026 Prasad S. Kodukula and Guz Vinueza



Guz Vinueza

Ecuador

Guz Vinueza, M.S., MBA, is a recognized expert in data analytics, machine learning, and quantitative risk management, with over 25 years of experience helping organizations utilize advanced technologies to create business value and manage uncertainty. He currently works as Consulting Director at The Ferryfield Group, a boutique consulting firm focused on quantitative risk software, where he manages global projects across various industries and regions. Previously, he was Data Director at Betterfly, Latin America's first social-impact unicorn, overseeing enterprise data architecture, analytics, and machine learning initiatives to expand the company's digital ecosystem. Earlier in his career, as Director of Consulting and Analytics Solutions at Palisade, he transformed the consulting division into an agile software and analytics team, tripling revenue and delivering solutions across Latin America, the U.S., and the Middle East.

An accomplished consultant, educator, and keynote speaker, Guz has trained thousands of professionals from leading organizations such as the U.S. Army Corps of Engineers, Ontario Power Generation, Amway, DEWA, Mitsubishi Pharmaceuticals, and Network Rail. He has also served as an Adjunct Professor in Analytics and Business Intelligence at institutions in Argentina, Ecuador, and Spain, teaching advanced analytics, agile methodologies, and project management. Guz holds an MBA from Universidad Torcuato Di Tella, an M.S. in Financial Direction from Universidad Adolfo Ibáñez, and a Postgraduate Certificate in AI and Machine Learning from the University of Texas at Austin. His work continues to shape the integration of AI, quantitative risk, and data-driven decision-making across global enterprises. He can be contacted at: https://www.linkedin.com/in/gustavovinueza/

© 2026 Prasad S. Kodukula and Guz Vinueza