B&E

# Resolution Fraud

## The New Lie AI Learned Instead of Saying "I Don't Know"

### Correct ingredients. Impossible recipes.

A non-technical briefing for business, education, governance, and professional readers on a failure mode in which AI does not fabricate individual facts, but assembles real facts into a proposed resolution that cannot realistically be implemented.


| | |
|---|---|
| **Series** | Business & Education Series |
| **Publisher** | UTIE Research Institute |
| **Date** | February 2026 |
| **Language** | English edition, adapted from the Japanese briefing 解決詐欺 |


**Note on terminology:** The term "Resolution Fraud" is used here as an analytical category, not as a legal allegation against any specific person or entity unless otherwise stated in a formally documented case record.

## Contents

# Executive Summary

Generative AI does not only fail by inventing false facts. It also fails by turning real facts, real technologies, and real methods into proposed resolutions that look coherent on paper but cannot be implemented under real-world constraints. This document calls that failure mode Resolution Fraud.

A hallucination is a false ingredient: a nonexistent library, a fabricated court case, a made-up product version, or a false source. Resolution Fraud is different. The ingredients may be real. The recipe is the fraud. AI can name genuine technologies, frameworks, institutions, and process steps, then combine them into a plan that ignores cost, labor, procurement, authorization, maintenance, legal obligations, political resistance, infrastructure, and time.

For business and education readers, this distinction matters. Fact-checking is necessary, but it is not sufficient. A proposal can survive fact-checking and still collapse at the implementation layer. The central question is therefore not only "Is each statement true?" but also "Can this proposed resolution be executed by this organization, with these resources, in this environment, within this time horizon?"

Resolution Fraud is especially persuasive because it often appears fluent, moderate, and technically literate. It acknowledges small constraints while silently hiding larger ones. It uses phrases such as "integrate the systems," "optimize the workflow," "build a data lake," "implement governance," or "refresh the architecture" as if these were small actions rather than expensive, multi-departmental programs.

The practical defense is not to reject AI. It is to route AI outputs through feasibility checks. Every proposed resolution should be tested for missing How, Cost, Authority, Dependency, Maintenance, and Friction. Organizations that learn this skill will gain from AI; organizations that do not will receive polished documents that quietly convert impossible assumptions into executive decisions.

**Key takeaway:** Hallucination is a failure of facts. Resolution Fraud is a failure of feasibility. The second can remain invisible even when every individual fact appears correct.

# 1. Definition: How Resolution Fraud Differs from Hallucination

Hallucination refers to AI output that fabricates facts. It may cite a nonexistent legal case, invent a software package, attribute a quotation to the wrong person, or confidently state a product specification that does not exist. Hallucinations are serious, but they are often detectable through verification. If the source, library, statute, or data point does not exist, the error can be exposed.

Resolution Fraud is more subtle. It occurs when AI uses real components to produce an unrealizable resolution. The technology exists. The process name is valid. The framework is recognizable. The citation may even be accurate. But the proposed resolution is not executable in the real world, because it ignores the cost, friction, dependencies, and institutional conditions required to make the proposal work.

A hallucination is a dish cooked with fake ingredients. Resolution Fraud is a recipe made with real ingredients that no one in the given kitchen can actually cook.

| Question | Hallucination | Resolution Fraud |
|---|---|---|
| **Primary failure** | The fact is false. | The proposed resolution is infeasible. |
| **Typical example** | A nonexistent source or library. | A real system integration treated as a trivial step. |
| **How it evades review** | May be missed if no one checks the source. | Can pass fact-checking because its parts are real. |
| **Best first question** | Does this fact exist? | Can this be executed here, now, by |

| | | |
|---|---|---|
| | | these people? |

## 2. Why Resolution Fraud Emerges from Language Models

Large language models operate by generating likely continuations in language. This is not the same thing as weighing the real-world cost of those continuations. A short phrase can conceal an enormous implementation burden.

For a model, "integrate the systems" is a compact phrase. For an organization, it may require API discovery, authentication design, firewall approvals, vendor negotiation, data mapping, exception handling, monitoring, incident response, maintenance ownership, budget approval, and months of coordination. The language is light. The world is heavy.

This gap is the root of Resolution Fraud. The model can generate a grammatically and technically plausible sequence: identify data sources, build a data lake, automate analytics, optimize operations. Each phrase sounds professional. Yet each phrase may contain a full program of work.

A useful metaphor is the difference between a frying pan and a rocket. In language, both can be mentioned in a few words. In the physical world, frying an egg takes minutes; building a rocket requires an industrial base. AI can place both into a sentence with similar ease, but organizations cannot execute them with similar ease.

> **Operational principle:** Every short verb in an AI-generated plan should be treated as a possible container for hidden labor.

## 3. The Linguistic Pattern: The Magic Wand of "If You Just..."

Resolution Fraud often appears through conditional language: "If you integrate the databases," "If you build a governance layer," "If you optimize the workflow," "If you migrate to a modern architecture," or "If you standardize the data." These clauses sound like preparation. In reality, they may be the entire project.

The key rhetorical move is that the condition is made to look small while the conclusion is made to look certain. "If the systems are integrated, real-time reporting is possible" may be true in the abstract. But the sentence quietly hides the most expensive part inside the word "integrated."

Another pattern is tautological advice. A company says sales are declining. AI replies that the company should analyze customer needs and implement appropriate marketing measures. This is not operational advice. It is a restatement that success will occur if the correct successful actions are taken.

The missing element is How. A resolution that does not include the actual path, responsible parties, dependencies, cost, sequence, and failure modes is not a resolution. It is a slogan in procedural clothing.

## 4. Vacuum-World Proposals: Why AI Misses Friction

Many AI proposals are written as if they exist in a vacuum. The proposal assumes clean data, cooperative vendors, compatible systems, stable budgets, available staff, aligned executives, and users who will adopt new workflows without resistance.

Real organizations do not operate in that vacuum. Systems are old. Fields are inconsistent. Departments protect their own processes. Security teams reject open ports. Legal teams slow down data sharing. Procurement cycles delay tools. The person who understands the legacy system may be leaving next month. These are not side issues. They are the project.

This remains true even if AI systems become more capable. Intelligence can produce a design quickly, but concrete must still dry, permits must still be approved, machines must be repaired, employees must be trained, and institutions must accept responsibility. Intelligence accelerates information work; it does not abolish physical and social friction.

The practical lesson is that feasibility is not a property of a sentence. Feasibility is a relationship between a proposal and a specific organization, budget, authority structure, infrastructure, legal environment, and implementation capacity.

## 5. Why Humans Believe It

Resolution Fraud is persuasive because it exploits normal human psychology. First, fluency produces authority. A smoothly written explanation feels intelligent. Once users see AI write emails, summaries, and presentations with impressive speed, they may incorrectly assume that the same fluency guarantees valid planning or technical judgment.

Second, early success lowers verification standards. If AI saves time on low-stakes tasks, users develop a sense of operational momentum. That momentum can carry into higher-stakes work where the cost of a wrong assumption is much larger.

Third, users begin to adapt their goals to the tool. Instead of asking what the organization actually needs, they ask what the AI can conveniently produce. The user slowly moves from controlling the tool to being trained by the tool. The output looks useful because the problem has been simplified until the AI can answer it.

Finally, small acknowledgements of reality can make the whole answer feel realistic. An AI system may mention that transport costs are high or that data cleaning is needed. These small concessions create trust, while the larger feasibility gap remains hidden.

## 6. Business and Education Impact

In business settings, Resolution Fraud can enter reports, strategy decks, internal proposals, vendor evaluations, and executive briefings. The document becomes readable, confident, and seemingly well structured. The underlying implementation plan may still be hollow.

A common scenario is the AI-assisted proposal. A user asks for a plan to improve data utilization. The AI recommends integrating internal databases, creating a data lake, and automating dashboards. The language is professional. But for the organization, the plan may imply a multi-year transformation program rather than a practical next step.

Education faces a related problem. Students and professionals may learn to produce plausible answers without learning the feasibility discipline behind them. A polished AI-generated essay can describe governance, ethics, or innovation while avoiding the hard operational question: who does what, with which resources, under which constraints?

Vendors can also amplify this failure mode. Demonstrations often show a best-case workflow using clean inputs. When the system is later connected to messy spreadsheets, inconsistent records, unclear permissions, and real customer data, the original promise turns into a new consulting engagement. In this form, Resolution Fraud becomes a business model.

## 7. What Happens When AI Is Pressured

One way to reduce Resolution Fraud is to impose strict constraints before asking for a solution: zero budget, unknown labor capacity, no tolerance for operational failure, no new vendors, and no unapproved data movement. When these constraints are stated explicitly, AI systems often move to the opposite failure mode: they refuse to propose anything useful.

This creates a practical tension. If the model is given freedom, it may produce a confident but infeasible plan. If the model is boxed in too tightly, it may answer that no solution exists or produce only generic caution. Between those poles, organizations need a disciplined workflow rather than blind acceptance or total rejection.

A practical workflow is to let AI generate options, then run each option through feasibility gates. The user asks: What is the hidden task? What dependency is assumed? Who has authority? What must already be true? What happens if this fails? How much time and money are being hidden by the verb? This converts AI from an apparent decision-maker into a draft generator subject to human inspection.

# 8. Case Note: Peer Review, AI, and Institutional Friction

A case discussed by UTIE in 2026 illustrates why Resolution Fraud is not limited to corporate transformation projects. In the peer-review context, AI-generated or AI-like review content may appear technically literate while misunderstanding the actual object of analysis, the obligations of the review process, or the institutional duties attached to scholarly evaluation.

The important point for this briefing is structural rather than personal. Academic publishing, like business operations, depends on procedures, records, accountability, and trust. When a review process produces fluent text that cannot be reconciled with the real duties of peer review, the problem is not merely poor wording. It is a failure of institutional feasibility.

The same forensic logic applies: do not evaluate the text only by its fluency. Evaluate whether the reasoning, procedure, evidence chain, and institutional explanation can coexist in the real world. A system that says a resolution has occurred must be able to show the steps, records, authorities, and obligations by which that resolution occurred.

For governance purposes, this moves the discussion beyond "Was AI used?" toward a more practical question: did the output, process, or explanation convert a missing or defective procedure into the appearance of a completed resolution?

# 9. Detection and Governance Approaches

Resolution Fraud is harder to detect than hallucination because the individual facts may be correct. Detection therefore requires feasibility analysis, not only source verification.

A simple first step is linguistic triage. Phrases such as "integrate," "optimize," "automate," "standardize," "refresh," "govern," "monitor," and "align stakeholders" should trigger a feasibility check when they appear without implementation detail. These words are not wrong. They are containers that may hide cost and coordination.

A second step is the HOW-COST test. If an AI output says that a problem can be resolved by doing X, the reviewer asks: How exactly is X done? Who does it? What must already exist? What is the cost? What is the timeline? What can fail? Who owns maintenance after launch?

A third step is domain triage. Some AI uses can tolerate 60 percent accuracy when speed is the value, such as brainstorming, drafting, summarization, or internal ideation. Other uses require near-total reliability, such as customer-facing advice, compliance judgments, contractual interpretation, public communications, incident response, and institutional determinations. Resolution Fraud becomes most damaging when these categories are confused.

### Feasibility Gate Checklist

| Gate | Question to Ask Before Accepting the AI Output |
|---|---|
| **HOW** | What are the concrete steps, in sequence? |
| **COST** | What budget, labor, tools, and time are required? |
| **AUTHORITY** | Who is authorized to approve, operate, and be accountable for this? |
| **DEPENDENCY** | What must already exist for the plan to work? |
| **FRICTION** | Which legal, political, organizational, physical, or security constraints are being ignored? |
| **MAINTENANCE** | Who keeps it running after the first demonstration? |
| **FAILURE MODE** | What happens when a step fails or produces an exception? |

# 10. Conclusion: From Magic to Civil Engineering

Resolution Fraud does not show that AI is useless. It shows that AI and human institutions inhabit different worlds. AI operates in information space, where a rocket and a frying pan can both be handled as compact language objects. Organizations operate in physical and social space, where rockets require industrial systems, approvals, capital, maintenance, and time.

The next stage of AI literacy will not be defined by enthusiasm or rejection. It will be defined by translation. Professionals must learn to translate AI outputs into the language of implementation: labor, cost, authority, dependency, friction, and maintenance.

The age of magical AI promises is giving way to an age of civil engineering. The organizations that benefit will not be those that ask AI for the most impressive answer. They will be those that can identify which parts of the answer are real, which parts are slogans, and which verbs quietly contain a multi-year project.

The basic discipline is simple: whenever AI says "This can be resolved by...", ask, "By whom, with what, at what cost, through which path, under which constraints, and with what evidence that the resolution has actually occurred?"

# Appendix A. Common Patterns of Resolution Fraud

**1. Minimized preconditions:** The condition clause hides months of work: "If the databases are integrated..."

**2. Tautological advice:** The answer says, in effect, that success will occur if successful measures are taken.

**3. Vacuum-world planning:** The proposal assumes clean data, cooperative institutions, compatible systems, and no politics.

**4. Hidden-cost simplicity:** The plan looks short because the expensive part has been compressed into one verb.

**5. Infinite dependency chains:** Each proposed step requires a prior system, prior approval, prior dataset, or prior capability that does not yet exist.

**6. Confidence as a substitute for verification:** A firm tone makes a conditional possibility feel like an executable commitment.

**7. Best-case demo laundering:** A demonstration under clean conditions is treated as evidence that the system works in messy operations.

**8. Small-friction realism:** The output admits minor constraints, which makes readers overlook the larger missing feasibility layer.

# Appendix B. Illustrative Example: The Martian Italian Restaurant

The question "How can one open an Italian restaurant on Mars?" is a useful stress test because the answer can contain many real technologies while still becoming infeasible as a business or institutional plan.

A typical AI answer may mention reusable spacecraft, closed-loop life support, in-situ resource utilization, hydroponic wheat, LED-grown tomatoes, precision-fermented cheese, pressure domes, low-gravity cooking, and welfare services for early settlers. These elements are not imaginary. Many are connected to real research or plausible engineering directions.

The Resolution Fraud appears when the answer treats the existence of those elements as if it resolves the actual problem. Building and sustaining a settlement on Mars would require transportation infrastructure, radiation protection, energy systems, industrial maintenance, medical support, governance, supply chains, capital, legal permissions, and a population large enough to support non-essential commercial services.

When AI says that the restaurant is "technically possible" because it sits on top of future settlement infrastructure, the burden has been moved rather than solved. The restaurant is not the main difficulty; the hidden difficulty is the civilization that must already exist for the restaurant to be ordinary.

This example is exaggerated, but the pattern is common in everyday business advice. "Build a data lake" plays the same role as "build a Martian settlement." It can be written quickly.

**Recommended citation:** UTIE Research Institute. Resolution Fraud: The New Lie AI Learned Instead of Saying "I Don't Know". Business & Education Series. UTIE Instruments Inc., 2026.