Automate the Repeatable, Own the Judgment

A three-layer model for using AI without outsourcing responsibility

Michael Stoyanovich

Artifact Classification: Practical Guide

Version 1.1.1 | July 2026

Scope and Limitations

This paper is intended for informational and educational purposes only. The views and analyses presented - particularly those related to ethics, policy, and AI system design - reflect the author's interpretations and do not constitute legal, regulatory, or professional advice. Readers are encouraged to critically assess the content and consult appropriate experts or authorities before applying any concepts discussed herein. The author assumes no liability for any decisions or actions taken on the basis of this work.

Purpose

This 1-pager separates work that is increasingly automatable from work that remains distinctly human - and adds the missing third layer: governing the handoff from automated output to accountable decision. It also adds a practical precondition: before automating a task, determine whether the task is well formed or whether it is an unresolved coordination problem caused by unclear goals, roles, decision rights, information flows, or team norms.

How to use: Treat this as a design-and-review checklist for any workflow that blends automation and human judgment (planning, QA, governance, training).

Intended Audience

Executives, managers, governance practitioners, AI program leaders, workflow designers, auditors, and professionals responsible for introducing AI into organizational processes.

Reading Guide

This paper presents a three-layer operational model for responsible AI adoption. It begins by identifying work suitable for automation, distinguishes judgment that must remain human, and concludes with governance controls for managing the transition from automated output to accountable organizational action.

The Three Layers

1) Automate the Repeatable

Use automation (including AI systems) for tasks that are structured, repeatable, and auditable.

Common automation sweet spots include:

- Data collection & preprocessing
- Routine calculations & baselines
- Summarization & first drafts
- Basic visualizations
- Compliance checks (rule-based / checklist-driven completeness and consistency checks; not legal determinations)
- Framework cross-mapping (mechanical crosswalks, coverage matrices)

Outputs from this layer are drafts, not decisions.

In practice: Treat outputs as inputs to review, not as authoritative conclusions.

Automation can increase throughput while decreasing truth if QA gates don't scale - which is why the handoff layer matters.

Repeatable does not mean low-risk. Treat high-materiality outputs as candidates for escalation and verification.

Automation Readiness Caveat

Not every painful or repetitive task is ready for automation. Some work is repetitive because the organization has not yet clarified goals, roles, decision rights, escalation paths, or information flows. In those cases, automation may accelerate confusion rather than reduce it.

Before automating, ask: is this task well formed, or is it a coordination problem wearing the costume of work?

If the underlying problem is unclear ownership, misaligned goals, poor information flow, or inconsistent tool use, redesign the work before automating it.

2) Own the Judgment

Humans remain responsible for decisions that require accountability, prioritization, tradeoffs, and legitimacy.

Human differentiators that rise in value include:

- Critical thinking & synthesis (what matters, what's missing, what's noise)
- Strategic framing & decision context (the decision, constraints, "good" criteria)
- Risk judgment & prioritization (materiality, sequencing, accepted residual risk)

- Communication, collaboration, influence (alignment across incentives and disagreement)
- Ethical reasoning & governance (boundaries, accountability, escalation)
- Tool orchestration (designing reliable workflows with fallible tools)

Owning judgment means you can explain the decision, defend it, and revise it when new evidence arrives.

3) Govern the Handoff

Most failures happen here: automated outputs silently become "truth" (e.g., a draft risk score quietly becoming a binding decision).

Minimum handoff controls

- Provenance: What system produced this? With what inputs? When? Under what settings?
- Data boundaries: What data classes were allowed/prohibited? Any sensitive data exposure?
- Verification: What checks were performed (spot checks, reconciliations, second-source validation)?
- Uncertainty signaling: Where might it be wrong? What confidence limits were recorded?
- Stop-the-line triggers: What conditions require escalation or human-only handling?
- Accountability: Who signs off? Who is responsible for downstream impact?
- Record keeping: What is logged for auditability, traceability, and learning?

Change control: What changed since last time - model version, prompt, policy, data source, workflow?

A Practical Diagnostic Tool

Use this diagnostic to classify any work product, workflow, or recurring responsibility and decide what to automate, what to keep human, what to redesign, and when to govern the handoff:

1. Repeatability: Is this task stable and pattern-based, or context-sensitive and novel?
2. Norm Clarity: Are the goals, roles, decision rights, escalation paths, information flows, and tool-use expectations clear?
3. Materiality: If this is wrong, what's the downside? Financial, legal, safety, reputational, fiduciary, operational, or employment-related?
4. Reversibility: Can we easily undo a bad decision, or would the consequences persist?
5. Legibility: Can the reasoning be explained to intended stakeholders and withstand scrutiny?

6. Ownership: Who is accountable for the decision and its consequences?

Rule of thumb:

- If it's repeatable and well formed → automate or partially automate.
- If it's painful because goals, roles, decision rights, or information flows are unclear → redesign the work before automating.
- If it's material, irreversible, or legitimacy-bound → keep human judgment in control.
- If it moves from output to action → govern the handoff.

Ethics and Governance Posture

This model assumes:

- AI outputs can be useful and still be wrong.
- Accountability cannot be delegated to a tool.
- Human oversight is not a vibe; it's a control system.
- "Faster" is not a justification for bypassing review, evidence, or stakeholder obligations.
- Data minimization and purpose limitation matter: use only the data needed, for the stated purpose.
- Automation should not be used to hide unclear ownership.
- AI should not be used to compensate for missing decision rights.
- Meeting overload, escalation patterns, and communication breakdowns may be symptoms of poor work design rather than opportunities for simple automation.
- Tools do not create alignment by themselves; teams need norms for how the tools are used, who decides, who reviews, and when issues escalate.

When stakes are high (health, benefits, safety, legal rights, employment, fiduciary duty), default to:

- tighter controls,
- clearer escalation paths, and
- documented sign-offs.

Ethics, Disclosure, and Acknowledgements

Ethical Considerations

This paper does not draw on private, sensitive, or personally identifiable data. Examples are hypothetical or anonymized. No human-subjects research was conducted. The ethical purpose is practical: to reduce responsibility leakage by keeping outputs distinct from decisions, and by specifying governance controls for high-stakes use.

Use of AI Tools

AI language models (including ChatGPT) were used as drafting and thinking aids (brainstorming, structure, and clarity checks). They can be wrong. Responsibility for all claims, framing, and conclusions remains with the human author.

Acknowledgements

Thanks to informal readers for feedback that improved clarity and structure. All errors and omissions are the author's alone.

Disclosure Statement

This work was conducted independently. The views expressed are the author's alone and do not represent any current or former employer. No conflicts of interest are declared.

License & Attribution

Trademark Notice

How to Cite This Guide

Stoyanovich, M. (2026, July). *Automate the Repeatable, Own the Judgment: A three-layer model for using AI without outsourcing responsibility* (Version 1.1.1). https://www.mstoyanovich.com

Related Works

Leading Through Alignment: A Framework for the Chief AI Officer — https://www.mstoyanovich.com

Version History and Document Status

This is a living document. As generative AI systems and their use evolve, this paper will be periodically updated to incorporate new empirical findings, theoretical insights, and policy developments. Major revisions are recorded here to preserve transparency and scholarly traceability.

Version	Date	Description
1.1.1	July 2026	Publication matter refinement and corpus normalization. Added standardized version metadata; normalized front-matter ordering and related-works presentation to align with the MJS Publication Matter Standard. No substantive changes to the framework, guidance, or recommendations.
1.1.0	June 2026	Added norm-clarity / automation-readiness precondition; revised diagnostic and rule of thumb; added work-design caveats to ethics and governance posture.
1.0.2	February 2026	Incorporated editorial refinements; added ‘How to use’ guidance; clarified handoff controls.
1.0.1	February 2026	Tightened ethics/disclosure language; added automation caveats; clarified compliance-check boundary; added change-control handoff control; added related companion links.
1.0.0	February 2026	Baseline version.