

The Human Use of Powerful Systems

What Lee Thayer, Neil Postman, and the systems and communications tradition can teach leaders about emerging technology

Peter J. Dykhuis, MS, LSSGB

August 2026

© 2026 Peter J. Dykhuis. All rights reserved.

Peter J. Dykhuis | The Human Use of Powerful Systems

Purpose and Orientation

This study revisits a body of systems, communication, media, and technology scholarship that is often described as prophetic. That description is understandable but incomplete. Lee Thayer, Neil Postman, Norbert Wiener, Jacques Ellul, Marshall McLuhan, Ludwig von Bertalanffy, Gregory Bateson, Herbert Simon, Joseph Weizenbaum, Lewis Mumford, and later thinkers were not simply predicting gadgets.

They were examining recurring relationships among tools, institutions, attention, meaning, authority, and human agency.

Their enduring value lies less in whether every forecast proved correct than in the questions they taught us to ask. What does a system reward? What does a medium make easy, difficult, visible, or invisible? What kinds of people and institutions does a technology cultivate? What happens when efficiency becomes an unquestioned end? Where does meaning arise, and who remains accountable when judgment is distributed across machines, procedures, incentives, and organizations?

The aim here is constructive. Emerging technologies can expand access to knowledge, increase organizational capacity, accelerate discovery, reduce drudgery, and improve decisions. The leadership challenge is not to reject those capabilities. It is to adopt them without allowing their internal logic to displace the purposes they were meant to serve. The practical goal is technological vigor with human direction: systems that amplify judgment, creativity, learning, and responsibility rather than gradually substituting for them.

Central Proposition

The most important guardrail is not a prohibition. It is a disciplined capacity to keep purpose, context, judgment, and accountability visible while technology increases speed and scale.

Peter J. Dykhuis | The Human Use of Powerful Systems

Table of Contents

	Front Matter	
	Executive Summary	3
	Main Chapters	
1.	Why Earlier Thinkers Matter Now	5
2.	Lee Thayer: Communication as a Living System	7
3.	Neil Postman: The Questions Technology Conceals	9
4.	A Wider Intellectual Constellation	12
5.	The Contemporary Technology Environment	14
6.	Points of Convergence	16
7.	A Constructive Leadership Framework: PURPOSE	17
8.	Organizational Practices and Guardrails	20
9.	Individual Practices for Intellectual and Executive Vigor	23
10.	Applying the Framework	25
	Back Matter	
	Conclusion	27
	Appendix A: Leadership Diagnostic	28
	Appendix B: Discussion Guide	30
	References	31
	Author's Note	33
	About the Author	33

Executive Summary

Emerging technologies are often discussed as though the principal leadership question were whether organizations will adopt them quickly enough. Speed matters, but it is not sufficient. The deeper issue is

whether leaders can preserve the ability to define purposes, interpret consequences, maintain accountability, and cultivate human capability while technological systems become more capable, pervasive, and autonomous.

The systems and communications tradition offers a useful vocabulary for that task. Lee Thayer treated communication not as the transfer of packaged meaning but as a constitutive process through which living systems orient themselves and act. Neil Postman argued that technologies carry assumptions, incentives, and definitions of knowledge; they do not merely deliver neutral content. Norbert Wiener placed feedback, control, and communication at the center of cybernetics while insisting that machines must be judged by the human purposes they serve. Jacques Ellul examined the tendency of technique, or organized efficiency, to become self-justifying. Marshall McLuhan showed that the form of a medium restructures perception and social organization regardless of particular messages. Herbert Simon explained that an abundance of information consumes attention. Joseph Weizenbaum warned against confusing computational capability with moral or human judgment.

Taken together, these thinkers suggest that the principal risks of advanced technology are often second-order effects. A tool solves the problem placed immediately before it while changing the environment in which future problems will be perceived. It accelerates communication while weakening comprehension. It expands choice while narrowing attention. It standardizes decisions while concealing values inside metrics and workflows. It increases access to answers while potentially reducing the habits required to formulate worthwhile questions.

Yet the same tradition also supports a positive program. Feedback can support learning. Automation can release people from repetitive work. Networks can widen participation. Artificial intelligence can help teams see patterns, test alternatives, and make expertise more accessible. The issue is not whether organizations should use emerging technology, but whether they can govern its relationship to purpose and capability.

Seven leadership priorities emerge from this analysis:

First, begin with the human or institutional purpose, not the technological capability. A useful system is one whose success can be described without mentioning the tool. **Second**, treat every technology as an environment rather than merely an instrument; examine the behaviors, incentives, categories, rhythms, and dependencies it creates around the intended use. **Third**, preserve meaningful human judgment:

human review should never become ceremonial approval after consequential choices have already been made by data, defaults, or workflow design.

Fourth, design feedback that improves the system rather than merely proving compliance, including qualitative experience, unintended effects, dissent, and changes in user capability. **Fifth**, protect attention as an enterprise resource; organizations should manage interruption, information volume, meeting load, and automated content with the same seriousness applied to financial or operational capacity. **Sixth**, measure capability, not only output; technology adoption should leave people and institutions better able to understand, decide, create, and recover. **Seventh**, keep accountability legible: when decisions are distributed across models, vendors, policies, data, and teams, leaders must still identify who owns the purpose, the exception, the consequence, and the correction.

This paper develops these priorities into a practical framework called **PURPOSE**, which stands for Purpose, Understanding, Responsibility, Proportion, Openness, Sensemaking, and Enduring Capability. The framework is designed for strategic decisions, governance forums, project design, operating reviews, and individual leadership practice.

Peter J. Dykhuis | The Human Use of Powerful Systems

CHAPTER 1

Why Earlier Thinkers Matter Now

Technological novelty can obscure intellectual continuity. Generative AI, autonomous agents, pervasive sensing, synthetic media, predictive analytics, neurotechnology, and digitally mediated work are new in capability and scale. The underlying leadership dilemmas are older: how tools reorganize institutions; how communication shapes reality; how measurement influences conduct; how systems pursue goals; how feedback alters behavior; and how human beings retain agency within environments they have built.

The value of earlier theory is therefore not nostalgia. It is conceptual leverage. A current product description tells us what a technology can do. A systems or media perspective asks what changes after it becomes normal. That shift, from capability to consequence and from isolated use to systemic pattern, is essential because major effects rarely remain confined to the original use case.

From Tools to Environments

A hammer remains largely external to the carpenter. A communication platform, enterprise workflow, or learning model is fundamentally different. It defines categories, establishes default sequences, distributes visibility, records some activities and ignores others, accelerates some interactions, and embeds assumptions about what counts as a valid input or successful result. Once integrated, it becomes part of the environment in which judgment occurs.

The environment matters because people adapt. They learn what the system notices, what leaders reward, what the interface permits, and what the metric recognizes. Over time, behavior can shift from accomplishing the underlying purpose to satisfying the representation of that purpose inside the system. This is not necessarily misconduct. It is a predictable response to the design of organizational attention.

First-Order and Second-Order Questions

First-order questions concern direct performance: Is the model accurate? Does the workflow save time? Can the platform increase reach? Second-order questions concern the conditions created by success: What happens to expertise when recommendations are routinely accepted? How does faster production affect review? What kinds of knowledge become less visible because they are difficult to quantify? What new dependencies or concentrations of power are created? What behaviors will the metric teach?

Responsible leadership requires both kinds of inquiry. First-order skepticism alone can become obstruction. First-order enthusiasm alone can become naïveté. The objective is disciplined adoption: enough confidence to experiment and enough systems awareness to notice what the experiment is changing.

Key Distinction: First-order questions ask what the technology can do. Second-order questions ask what the technology will make normal, and what it will gradually make invisible.

A Tradition of Constructive Restraint

The thinkers examined here differed greatly in politics, method, and temperament. It would be misleading to collapse them into a single anti-technology school. Wiener saw enormous promise in automation. McLuhan often analyzed rather than condemned. Postman defended technological progress when it remained answerable to human purposes. Simon developed decision sciences and artificial intelligence while warning about attention scarcity. Weizenbaum helped create an early conversational program and then argued that some decisions should remain beyond computation.

Their common contribution is a form of constructive restraint: a refusal to let capability settle questions of value. They encourage leaders to keep ends prior to means, human beings prior to abstractions, and learning prior to mere acceleration.

Peter J. Dykhuis | The Human Use of Powerful Systems

CHAPTER 2

Lee Thayer: Communication as a Living System

Lee O. Thayer occupies a distinctive place in organizational communication. His work rejected simple sender-message-receiver models when applied to human systems. In his account, communication is not best understood as a parcel of meaning transmitted intact from one mind to another. It is a living system taking something into account in relation to its continued organization and action. Meaning is not contained in a message like cargo in a vehicle; it arises in the receiver's capacities, history, purposes, and situation (Thayer, 1968, 1972, 1987; Anton, 2015).

Communication Is What Is Taken Into Account

This distinction has immediate relevance to contemporary leadership. Organizations frequently equate distribution with communication. A message was sent, a dashboard was published, a policy was posted, an AI summary was generated, or a training module was completed. These facts establish transmission,

not understanding. From a Thayerian perspective, communication has occurred only insofar as the receiving system can take something into account and reorganize action accordingly.

The consequence is practical: communication quality cannot be measured only by volume, reach, open rates, prompt completion, or message consistency. Leaders must examine what recipients were prepared to understand, how the message interacted with existing expectations, whether the context supported interpretation, and what action became possible afterward.

The Radical Openness of Human Systems

A closed mechanical system can often be described through bounded inputs, stable components, and predictable outputs. Human systems are radically open. They are historically formed, symbolically mediated, self-interpreting, and responsive to expectations about themselves. The same information can produce different meanings across people, teams, professions, and moments. The act of measuring or communicating can itself alter the system being observed.

This openness frustrates leaders seeking deterministic control, but it is also the source of creativity and adaptation. An open system can reinterpret, learn, improvise, and generate possibilities that were not specified in advance. Attempts to eliminate uncertainty entirely may therefore eliminate the very capacities needed for resilience.

Organization as Maintained Communication

Organizations are not merely structures within which communication happens. They are continually maintained through communication: distinctions about roles, priorities, authority, legitimacy, time, risk, and performance. A new digital system participates in this maintenance. It may redefine who can see what, which events count, whose language becomes standard, and how quickly a response is expected.

This is why technology implementation is always partly organizational design. Introducing an AI assistant, collaboration platform, customer system, or automated decision rule changes the communication system even if the formal organization chart remains untouched.

The Receiver's Capability

Thayer's emphasis on the receiving system places capability at the center of communication. The quality of a message cannot compensate indefinitely for the absence of interpretive competence. A complex

organization needs people who can distinguish signal from noise, recognize context, test assumptions, and connect information to purpose.

Emerging technology creates an apparent paradox. It can make knowledge more accessible while weakening the incentive to develop the capacities by which knowledge is judged. The leader's task is to use assistance without allowing assistance to become dependency. The relevant measure is not how much information a person can retrieve, but how well that person can understand, question, integrate, and act upon it.

Thayer's Lesson for AI-Mediated Communication

Generative systems can produce fluent messages, summaries, analyses, and recommendations. Fluency can create the impression that meaning has been secured. Thayer's work reminds us that semantic form is not shared understanding. A polished answer may still be unusable, misleading, or disconnected from the receiving system's real problem.

For leaders, this suggests four operating principles:

1. Do not confuse message production with communication achieved.
2. Assess the interpretive readiness of the audience, not only the quality of the content.
3. Design feedback that reveals how information was understood and used.
4. Build human capability so that technology enlarges the range of what can be taken into account.

Peter J. Dykhuis | The Human Use of Powerful Systems

CHAPTER 3

Neil Postman: The Questions Technology Conceals

Neil Postman's work is often reduced to a criticism of television. His broader project was media ecology: the study of how communication environments shape perception, knowledge, public discourse, and

institutional life. He insisted that technologies are not neutral containers. Each carries an embedded idea about what the world is, what counts as knowledge, and which human capacities matter (Postman, 1985, 1992, 1998).

Every Technology Is Both a Burden and a Blessing

Postman urged his audiences to reject one-sided technological narratives. Every major technology gives and takes away, distributes benefits unevenly, and changes the meanings of familiar words and practices. The printing press strengthened particular forms of abstraction, sequence, argument, and individual reading. Television privileged immediacy, image, personality, and entertainment. Digital media and AI similarly amplify some forms of cognition while diminishing the relative standing of others.

This trade-off is not a reason to refuse the technology. It is a reason to name the exchange honestly. Leadership improves when adoption proposals state not only expected gains but also the capacities, practices, or relationships that may be displaced.

The Medium Shapes What Can Be Said

In *Amusing Ourselves to Death*, Postman argued that a medium's form influences the kinds of discourse it can sustain. Television could discuss serious matters, but its commercial and visual logic tended to convert them into entertainment. The crucial claim was not that all televised content was frivolous. It was that the medium established conditions under which seriousness had to compete with spectacle, speed, personality, and discontinuity.

The contemporary analogy extends well beyond social media. Dashboards favor measurable states. Chat systems favor rapid response. Short-form video favors immediacy and emotional clarity. Generative AI favors answer-shaped outputs, even when the underlying situation may require uncertainty, investigation, or silence. Each format can be useful; each also changes what is likely to be noticed and rewarded.

Technopoly and the Surrender of Judgment

Postman used the term *technopoly* to describe a culture in which technology becomes not merely a tool but a source of authority. In such a culture, questions that were once moral, political, professional, or educational are reformulated as technical problems. Quantification acquires prestige. Expertise becomes associated with systems and procedures. Institutions risk accepting a technology's categories as though they were natural descriptions of reality.

This danger is especially relevant when organizations adopt AI scoring, ranking, forecasting, or recommendation systems. The issue is not simply model bias or accuracy. It is the possibility that the system's representation of the problem becomes the problem. Once a performance measure is embedded in executive reviews, resource decisions, incentives, and workflows, alternative interpretations become harder to articulate.

Institutional Purpose Under Pressure

Postman repeatedly asked what happens when the logic of one institution colonizes another. Education can become entertainment, politics can become marketing, medicine can become data processing, and management can become dashboard administration. Technology accelerates this process when it imports categories designed for one purpose into a different human setting.

Leaders should therefore distinguish an institution's mission from the mechanisms used to support it. A school is not a content-delivery platform. A hospital is not an optimization engine. A newsroom is not an engagement factory. A company is not merely a portfolio of measurable transactions. Metrics and systems are necessary, but they remain representations of a purpose that exceeds them.

Six Questions Every Leader Should Ask

Postman's framework translates directly into a set of questions worth asking before any significant technology adoption:

1. What problem does the technology claim to solve, and for whom is it a problem?
2. Whose interests and values are embedded in the solution?
3. What new problems will be created by solving the old one?
4. Which people or groups gain power, and which lose it?
5. What language, categories, or institutional purposes will be redefined?
6. What human capacity may weaken because the technology performs it more conveniently?

These are not rhetorical obstacles. They are the questions that distinguish **adoption** from **surrender**.

A Wider Intellectual Constellation

Thayer and Postman become more useful when placed within a wider tradition. The following thinkers contribute complementary lenses rather than a unified doctrine.

Bertalanffy: Resist Local Optimization

General systems theory challenged the assumption that a complex living or social system can be understood by analyzing components independently. Open systems exchange energy, information, and resources with their environments; they maintain themselves through relationships and adaptation. Organizational outcomes emerge from interactions, not simply from component quality (von Bertalanffy, 1968).

For technology leadership, the implication is direct. A highly efficient automated function may create downstream brittleness, reduce learning, shift risk to another team, or degrade the customer experience. The unit of analysis must be the end-to-end system and its environment, not the isolated task being optimized.

Wiener: Feedback Needs Purpose

Cybernetics studies communication and control through feedback. Wiener recognized that machines could extend human capacity and relieve repetitive labor, but he warned against treating people according to the logic of machines or pursuing poorly specified goals through powerful automation (Wiener, 1950/1954).

Modern systems can optimize engagement, throughput, conversion, response time, or prediction accuracy. Cybernetic thinking asks whether the feedback loop is connected to the right purpose. A system can become extremely effective at producing the wrong state. Leaders must therefore govern objectives, sensors, thresholds, and correction mechanisms, not merely outputs.

Bateson: The Pattern That Connects

Bateson emphasized relationship, context, and levels of learning. Many organizational failures persist because a team improves actions within an unchanged frame. More reports, more prompts, or more data may refine the response while leaving the governing assumptions intact (Bateson, 1972).

Emerging technology can support higher-order learning when it helps leaders compare frames, reveal patterns, test alternative explanations, and question categories. It can inhibit learning when it provides rapid closure and reinforces the current definition of the problem.

McLuhan: Extensions Create Amputations

McLuhan described media as extensions of human faculties. An extension increases reach but can also reduce reliance on the faculty extended. The automobile extends mobility while reorganizing cities and diminishing walking. Digital memory extends recall while changing what people choose to remember. AI extends composition and analysis while potentially changing the relationship between effort, understanding, and authorship (McLuhan, 1964).

The leadership question is not whether extension is good or bad. It is what complementary practice preserves the human capacity that may otherwise atrophy.

Ellul: Efficiency Can Become Destiny

Ellul's concept of *technique* is broader than machinery. It is the drive to identify and impose the most efficient method across domains. When technique becomes autonomous, the question "Can this be optimized?" displaces "What should this institution be for?" (Ellul, 1954/1964).

Organizations need efficiency, but they also need slack, deliberation, redundancy, apprenticeship, trust, and professional discretion. These can appear inefficient in a narrow model while remaining essential to resilience and quality. Leaders must protect forms of value that do not immediately present themselves as throughput.

Simon: Abundance Consumes Attention

Herbert Simon observed that information abundance creates a poverty of attention. This insight has become foundational in an environment of continuous notifications, dashboards, feeds, generated content, and automated recommendations (Simon, 1971).

Attention is not simply a personal wellness issue. It is an organizational production factor. Fragmented attention weakens deep work, synthesis, memory, strategic imagination, and the ability to detect weak

signals. A system that saves minutes on document production while multiplying the documents to be reviewed may reduce rather than increase enterprise intelligence.

Weizenbaum: Not Every Computable Decision Should Be Computed

Weizenbaum's ELIZA demonstrated how readily people attribute understanding to a program that manipulates language. His later work distinguished between decisions that can be made computationally and judgments that require human responsibility, empathy, and moral interpretation (Weizenbaum, 1976).

This is a crucial guardrail for leadership. Human oversight is not valuable merely because humans are more accurate. It is valuable because certain decisions require someone to understand the situation as a human situation and remain answerable for the result.

Summary: This constellation of thinkers shares a core conviction: capability does not settle questions of value, and the environment created by a technology matters as much as the task it performs.

Peter J. Dykhuis | The Human Use of Powerful Systems

CHAPTER 5

The Contemporary Technology Environment

The present environment differs from the eras of Thayer and Postman in speed, scale, personalization, and integration. Media are no longer only channels consumed by audiences. They are adaptive systems that observe behavior, generate content, rank possibilities, shape visibility, and increasingly act on behalf of users. This difference is significant, but it strengthens rather than diminishes the relevance of systems and media theory.

From Mass Media to Adaptive Mediation

Television presented largely the same stream to many viewers. Contemporary systems personalize streams, infer preferences, test responses, and update predictions. Communication environments now participate in selecting the messages, sequencing the experience, and optimizing continued interaction. The system is not merely a medium; it is a feedback-driven actor within the communication ecology.

From Automation of Labor to Automation of Cognition

Earlier automation primarily displaced or reorganized physical and routine clerical work. Current systems can draft, summarize, translate, classify, recommend, simulate, code, and converse. They enter activities previously associated with professional judgment and symbolic production. The resulting opportunity is enormous: expertise can be distributed, analysis accelerated, and barriers to expression lowered. The corresponding challenge is to distinguish augmentation from substitution and output from understanding.

Synthetic Abundance

Generative systems reduce the marginal cost of producing plausible text, images, audio, video, and software. This creates access and creativity, but it also changes scarcity. When content is abundant, verification, provenance, attention, and trusted interpretation become more valuable. The organizational bottleneck shifts from production to judgment.

Datafication and Legibility

Digital systems make activity legible by translating it into categories and records. Legibility enables coordination, accountability, personalization, and analysis. Yet every representation excludes. Tacit knowledge, context, relationships, uncertainty, and qualitative difference may be compressed into fields and scores. Leaders should treat data as a disciplined perspective on reality, not reality itself.

Delegated Agency

Agentic systems can plan, invoke tools, communicate, and execute sequences with decreasing human intervention. This raises a governance question beyond model output: what authority has been delegated, within which boundaries, under whose identity, with what audit trail, and through which recovery mechanism? As autonomy rises, accountability must become more explicit rather than less.

Contemporary Governance Standards

Current governance frameworks echo several of these earlier concerns. The NIST AI Risk Management Framework organizes practice around Govern, Map, Measure, and Manage, with attention to context, trustworthiness, ongoing evaluation, and risk ownership. Its generative-AI profile identifies risks including confabulation, information integrity, privacy, harmful bias, and human-AI configuration (NIST, 2023, 2024). UNESCO's Recommendation on the Ethics of Artificial Intelligence emphasizes human dignity, rights, fairness, transparency, and human oversight (UNESCO, 2021). The OECD AI Principles likewise address human-centered values, transparency, robustness, accountability, and human intervention (OECD, 2019, updated 2024).

These frameworks are important, but formal governance alone is not enough. A compliant system can still erode attention, weaken expertise, narrow institutional purpose, or create cultures of automatic deference. Systems theory adds the broader ecological questions that risk registers may miss.

Peter J. Dykhuis | *The Human Use of Powerful Systems*

CHAPTER 6

Points of Convergence

Across thinkers separated by decades and disciplines, several propositions emerge with unusual consistency.

Meaning is not information	Thayer, Bateson, and later systems theorists distinguish the availability of signals from the creation of meaning. More information can improve understanding only when the receiving system has time, context, competence, and purpose.
Tools reorganize their users	Postman, McLuhan, Mumford, and Turkle show that technologies cultivate habits, expectations, identities, and institutional forms. Adoption always changes the adopter.
Feedback amplifies goals	Wiener and cybernetics demonstrate that feedback creates powerful self-correction, but only relative to specified objectives. Poor goals can be pursued with increasing precision.

Meaning is not information	Thayer, Bateson, and later systems theorists distinguish the availability of signals from the creation of meaning. More information can improve understanding only when the receiving system has time, context, competence, and purpose.
Efficiency is a value, not the value	Ellul warns that technique can make efficiency appear morally neutral and universally desirable. Human institutions also require dignity, learning, resilience, fairness, and meaning.
Attention is a governing constraint	Simon and Postman explain why more communication can produce less comprehension. The management of attention is central to organizational intelligence.
Judgment carries responsibility	Weizenbaum and Wiener insist that computational capability does not erase human accountability. Some choices require a responsible person, not merely a reliable procedure.
Open systems require learning	Bertalanffy, Thayer, and Bateson emphasize adaptation and emergence. Guardrails must support feedback and revision rather than freeze an organization around current assumptions.

The Central Tension: Scale Without Reduction

Technology creates value by reducing complexity into forms that can be processed, compared, transmitted, and acted upon. Leadership creates value by remembering what the reduction leaves out. The goal is not to avoid abstraction; organizations cannot function without it. The goal is to prevent the abstraction from becoming sovereign.

This tension can be stated as a design challenge: How can an organization gain the speed, reach, consistency, and analytical power of technology while preserving context, discretion, diversity of interpretation, and the capacity to reconsider its ends?

Peter J. Dykhuis | The Human Use of Powerful Systems

CHAPTER 7

A Constructive Leadership Framework: PURPOSE

PURPOSE provides a practical way to evaluate and govern emerging technology by asking whether adoption strengthens the human and organizational system that uses it. It goes beyond cataloging risks to examine purpose, understanding, responsibility, proportion, openness, sensemaking, and enduring capability.

The framework applies at multiple levels. Strategic leaders can use it to evaluate major technology investments. Governance teams can use it to structure adoption reviews. Project teams can use it to design pilots. Individual leaders can use it as a personal discipline for consequential decisions about delegation, automation, and capability.

Letter	Dimension	Leadership Question & Guidance
P	Purpose	State the human, customer, civic, or institutional end before discussing the tool. Define success in language that remains meaningful if the technology changes. A useful system is one whose success can be described without mentioning the technology itself.
U	Understanding	Ensure that decision makers and affected users understand the system's role, limitations, assumptions, data dependencies, and likely second-order effects. Transparency is not optional; it is the precondition for meaningful oversight.
R	Responsibility	Assign accountable owners for the objective, deployment, exceptions, human review, adverse outcomes, and correction. Responsibility must not dissolve across vendors and workflows. A named individual must own the purpose and the consequence.
P	Proportion	Match the level of automation, monitoring, and control to the

Letter	Dimension	Leadership Question & Guidance
		stakes, reversibility, uncertainty, and affected rights. Not every efficiency warrants full automation. High-stakes, irreversible decisions require correspondingly strong human oversight.
O	Openness	Create channels for challenge, dissent, affected-user feedback, audit, and revision. Open systems learn; closed governance systems defend their initial design. Define in advance what evidence would cause the organization to change course.
S	Sensemaking	Protect the time and practices through which people interpret information, compare frames, deliberate, and form shared understanding. Do not allow production speed to eliminate reflection. Meeting structures should use AI outputs as inputs to discussion, not as conclusions.
E	Enduring Capability	Measure whether the system develops or depletes human competence, organizational memory, creativity, resilience, and the ability to act without the tool when necessary. Technology adoption should leave people and institutions more capable, not more dependent.

Framework Summary

PURPOSE is not a compliance checklist. It is a discipline of attention: a set of questions that keep human and institutional ends visible as technological capability grows more powerful and more pervasive.

Peter J. Dykhuis | The Human Use of Powerful Systems

CHAPTER 8

Organizational Practices and Guardrails

1. Separate the Purpose Case from the Technology Case

Many proposals begin with a capability and search for a use. Require sponsors first to describe the underlying problem, affected people, current constraints, desired outcome, and non-negotiable values. Only then compare technological and non-technological alternatives.

A strong business case should include an explicit statement of what the organization will refuse to optimize. Examples may include dignity in sensitive interactions, professional independence, due process, customer trust, or protected time for learning.

2. Use Tiered Delegation

Avoid the false choice between manual work and full automation. Define clear levels of authority based on stakes, reversibility, observability, and recovery capacity:

Tier	Mode	Description	When Appropriate
1	Assist	Retrieve, organize, or draft while a person performs the substantive judgment.	Any context; most supportive mode.

Tier	Mode	Description	When Appropriate
2	Recommend	Propose a course of action with evidence, uncertainty, and alternatives.	When options need to be structured for human choice.
3	Act with Approval	Prepare or initiate action only after meaningful human authorization.	Moderate stakes; approval must be real, not ceremonial.
4	Act Within Bounds	Execute reversible, low-stakes actions under monitored constraints.	Well-understood, observable, recoverable processes.
5	Act Autonomously	Execute without human intervention in individual steps.	Reserve for tightly bounded, well-tested, low-stakes workflows only.

3. Make Human Review Meaningful

Human-in-the-loop can become an empty phrase when reviewers face excessive volume, unclear standards, automation bias, or pressure to approve quickly. Meaningful review requires authority to reject, time to inspect, access to source material, visibility into uncertainty, and accountability that is not punished for prudent escalation.

4. Establish an Attention Budget

Treat attention as finite capacity. New tools should identify the notifications, reports, meetings, drafts, and review demands they will create. Productivity gains should be measured net of coordination and verification cost.

- Default to asynchronous communication when immediacy is unnecessary.
- Create quiet periods for analysis and composition.
- Limit automated summaries to material decisions and changes.
- Retire reports and alerts when new ones are introduced.
- Track review load and decision latency, not only content production.

5. Preserve Dissent and Alternative Frames

Systems tend to stabilize around their own categories. Governance forums should include a structured challenge function that asks whether the system is solving the right problem, whether affected groups interpret outcomes differently, and what evidence would cause the organization to change course.

Red teams are useful for security and failure analysis. Organizations also need rotating critical reviewers who examine institutional purpose, language, incentives, and second-order consequences.

6. Measure Capability Effects

Add capability indicators to adoption metrics. Relevant questions include:

- Can employees explain the reasoning behind the system's recommendations?
- Are novice employees developing expertise or merely learning how to operate the interface?
- Can the team detect a plausible but incorrect output?
- Is organizational memory becoming more accessible or more dependent on a vendor?
- Are people generating more alternatives, or converging prematurely on the system's first suggestion?
- Can critical work continue during interruption or withdrawal?

7. Design Feedback for Learning, Not Defense

Operational dashboards often favor confirmation: uptime, usage, cycle time, acceptance rate, and cost savings. Add signals capable of disproving the current design, including exception patterns, complaints, reversals, near misses, qualitative interviews, unexpected workarounds, and changes in professional practice.

8. Govern the Ecology, Not Only the Model

AI governance can become too model-centric. The actual system includes data collection, prompts, interfaces, workflow placement, incentives, training, user expectations, vendor contracts, escalation routes, and downstream decisions. Harm or value may arise from any combination of these elements.

9. Build Reversible Experiments

Pilot design should preserve options. Limit scope, establish baseline measures, define stop conditions, retain a manual alternative, and specify what evidence is needed to expand. Reversibility encourages learning because the organization can examine consequences without defending a sunk commitment.

10. Create a Technology Constitution

A concise organizational constitution can state durable principles across changing tools:

- Technology will support clearly stated human and institutional purposes.
- People affected by consequential systems will have understandable routes to question and appeal.
- No system will be treated as the final moral or professional authority.
- Automation will be proportional to stakes and reversibility.
- The organization will preserve the human capabilities required to supervise, challenge, and recover from its systems.
- Efficiency gains will be evaluated alongside trust, learning, resilience, and attention.

Peter J. Dykhuis | The Human Use of Powerful Systems

CHAPTER 9

Individual Practices for Intellectual and Executive Vigor

Institutional guardrails are necessary, but leaders also need personal disciplines. Technology changes leadership partly by changing daily cognition: the pace of response, tolerance for ambiguity, habits of reading, memory, conversation, and the temptation to outsource formulation before thinking has begun.

Think Before Prompting

Before asking a system for an answer, write a brief independent statement of the problem, the decision, the known facts, and the uncertainties. This creates an intellectual baseline and reduces anchoring on the system's framing.

Use AI for Divergence Before Convergence

Ask for alternative hypotheses, counterarguments, stakeholder perspectives, failure modes, and questions that have not been considered. Delay the request for a final recommendation until the problem space has been widened.

Maintain a Primary-Source Habit

Summaries are useful navigation aids, not substitutes for direct encounter with important texts, data, customers, employees, and events. For consequential decisions, inspect representative primary material and record what the summary omitted.

Protect Unfilled Attention

Solitude, walking, reflective writing, and unstructured conversation create conditions for synthesis. These practices can appear unproductive because their output is delayed and uncertain. They are nevertheless essential to executive judgment, which depends on connecting weak signals across domains.

Practice Explanatory Accountability

Do not present a generated conclusion that you cannot explain, defend, and revise. A leader may use systems to extend analysis, but ownership begins where attribution ends. *"The model said"* is not an executive rationale.

Cultivate Friction Deliberately

Good friction slows action where error would be costly and compels contact with reality. Examples include:

- A pre-mortem before deployment.
- A dissenting review before major investment.
- Direct customer observation before process redesign.
- A written rationale before delegating authority to an agent.

Keep a Decision Journal

Record the context, assumptions, alternatives, technological inputs, uncertainties, and expected outcomes behind major decisions. Revisit the record after results emerge. This strengthens learning and reveals whether tools are improving judgment or merely increasing confidence.

CHAPTER 10

Applying the Framework

Scenario 1: Generative AI for Executive Decision Support

An executive team introduces a system that summarizes performance, synthesizes market information, and proposes priorities. The first-order benefit is obvious: reduced preparation time and broader information access. The systems questions concern framing, source selection, uncertainty, dissent, and dependence.

For this use case, the goal is better strategic consideration rather than shorter meetings. Leaders should require links to primary evidence, distinguish facts from inferences and recommendations, circulate the synthesis before discussion, and reserve meeting time for disagreement, interpretation, and choice. Periodically forming an independent view before reviewing generated analysis can reveal anchoring, while comparing assumptions with outcomes can surface omitted signals.

Scenario 2: AI-Enabled Customer Service

A service organization uses AI to handle routine requests and assist representatives. The opportunity is faster response, greater consistency, and wider access. The systemic risk is that the organization optimizes containment while losing insight into customer difficulty and reducing the development of human expertise.

Effective governance here measures resolution quality and customer effort, not only deflection and handle time. It routes emotionally complex, high-stakes, or ambiguous cases to people early. It uses unresolved interactions as learning signals for product and process improvement. It preserves representative exposure to real customer conversations for leaders and designers. It trains employees to challenge suggested responses and explain why an exception matters.

Scenario 3: Automated Talent Decisions

A company considers AI tools for recruiting, performance, promotion, or workforce planning. These systems promise consistency and scale, but they convert complex human potential into data and categories. Because the decisions affect livelihood, dignity, and opportunity, proportion and accountability must be especially strong.

Governance here uses systems to organize evidence, not to make final judgments about human worth or potential. It provides affected people understandable information and a route to correction or appeal. It audits outcomes across groups and inspects the assumptions embedded in labels and historical data. It includes qualitative, developmental, and contextual evidence that resists simple scoring. It holds an identified leader accountable for the decision and its fairness.

Scenario 4: Agentic Workflow Automation

An organization deploys agents that access systems, draft communications, schedule work, and execute transactions. The central design task is delegated authority.

Sound design defines a permissions architecture based on task, data sensitivity, financial exposure, and reversibility. It requires explicit approval for novel, external, irreversible, or reputationally consequential actions. It maintains complete logs of intent, tool calls, data access, decisions, and outcomes. It designs kill switches, fallback procedures, and recovery ownership before launch. It tests whether employees can recognize and respond when the agent's local success conflicts with the wider system.

Scenario 5: Knowledge Work and Organizational Memory

A firm uses AI to search internal knowledge, summarize meetings, and draft deliverables. The system can reduce rediscovery and improve access, but it may also flood repositories with low-value content and obscure authorship, provenance, and context.

Strong practice distinguishes authoritative records from generated working material. It attaches provenance, date, owner, confidence, and source links to synthesized knowledge. It creates retention standards that prevent synthetic abundance from becoming organizational noise. It rewards contribution of tested knowledge, not raw volume. It uses the system to connect people with expertise, not only to extract answers from their past work.

Conclusion: Technological Vigor with Human Direction

The systems-and-communications tradition does not require leaders to choose between technological progress and human integrity. It requires them to reject a simpler and more dangerous choice: the assumption that progress can be defined by capability alone.

Lee Thayer reminds us that communication is an achievement of living systems, not a volume of messages. Neil Postman reminds us that technologies reshape the questions institutions ask and the meanings they preserve. Wiener shows that feedback and automation require worthy goals. Ellul warns that efficiency can become self-authorizing. McLuhan reveals that media alter their users. Simon identifies attention as the scarce resource. Weizenbaum draws a line between computation and responsibility.

The practical result is stewardship rather than restraint for its own sake. Leaders can use systems to widen perception, speed appropriate work, and distribute expertise while retaining the reflection and accountability those gains require. Technical controls matter, but they are insufficient unless organizational routines also keep purpose visible and preserve the ability to question the system's framing.

Closing Vision

The best future is not one in which human beings compete with their tools at the tasks the tools perform best. It is one in which tools expand the range of worthy purposes people and institutions are capable of pursuing, while human beings remain responsible

for defining those purposes, interpreting their consequences, and caring for the world those systems create.

Appendix A: Leadership Diagnostic

Use the following diagnostic during strategy reviews, technology investments, pilots, and operating reviews. Score each item from **1** (not true) to **5** (consistently true). Low scores are prompts for inquiry rather than automatic reasons to stop.

#	Statement	Score (1 to 5)	Notes
PURPOSE CLARITY			
1	We can state the human or institutional purpose this technology serves without mentioning the technology itself.		
2	We have identified whose problem this solves and what they currently experience.		
3	We have articulated what we will refuse to optimize, even if optimization is technically possible.		
4	Success can be described in terms that remain meaningful if the technology changes.		
5	Sponsors were required to frame the problem before identifying the tool.		
UNDERSTANDING			
6	Decision makers understand the system's limitations, assumptions, and failure modes.		
7	Affected users know what the system does and does not do on their behalf.		

#	Statement	Score (1 to 5)	Notes
8	Second-order effects were explicitly considered before adoption.		
9	We understand what data the system depends on and where it comes from.		
10	We have mapped what the system cannot see, measure, or account for.		
RESPONSIBILITY			
11	A named individual owns the purpose, deployment, and consequences of this system.		
12	Responsibility does not dissolve across vendors, teams, or automated processes.		
13	There is a clear escalation path when the system produces unexpected results.		
14	Human review is meaningful, not ceremonial.		
15	Exception handling is designed and tested, not assumed.		
PROPORTION			
16	The level of automation is matched to the stakes and reversibility of decisions.		
17	We have not automated high-stakes, irreversible decisions beyond our recovery capacity.		
18	Sensitive interactions involving dignity or rights remain under meaningful human control.		
19	We can describe the specific boundary conditions under which the system should not act.		
20	Automation levels are reviewed periodically rather than locked in at launch.		
OPENNESS			
21	Affected users have understandable channels to question and appeal system decisions.		
22	We have created structured mechanisms for internal dissent and alternative framing.		
23	Feedback reaches decision makers in forms that can prompt revision.		
24	We have defined what evidence would cause us to change course.		
25	Audit access is available to appropriate parties.		
SENSEMAKING			
26	We protect time for deliberation and interpretation, not only production.		
27	Meeting structures use AI outputs as inputs to discussion, not as conclusions.		

#	Statement	Score (1 to 5)	Notes
28	Leaders form independent views before reviewing generated analyses.		
29	We identify and examine what the technology makes invisible or difficult to articulate.		
30	Qualitative and contextual knowledge is valued alongside quantitative output.		
ENDURING CAPABILITY			
31	People are developing expertise through this technology, not merely operating an interface.		
32	Novice employees can explain the reasoning behind the system's recommendations.		
33	Critical work can continue during system interruption or withdrawal.		
34	Organizational memory is more accessible, not more vendor-dependent.		
35	We measure whether people generate more alternatives or converge prematurely on the system's first suggestion.		

Peter J. Dykhuis | The Human Use of Powerful Systems

Appendix B: Discussion Guide for Leadership Teams

These questions are designed for use in leadership offsites, technology governance forums, project retrospectives, and executive team reviews. They are most productive when answered in writing before discussion, with responses shared before the meeting begins.

What valuable purpose are we trying to advance, and how might our current framing be incomplete?

What does the proposed technology make easier, and what does it make less visible or less practiced?

Which assumptions about people, knowledge, performance, and authority are embedded in the system?

What will users learn to pay attention to because of this design? What may they stop noticing?

Where could local optimization damage the wider system?

What forms of human judgment are essential because someone must remain responsible, not merely accurate?

How will we know whether the system is developing or depleting capability?

What new content, alerts, review obligations, and expectations of speed will be created?

Who benefits, who bears risk, and who has a meaningful voice in correction?

What evidence would change our minds after implementation?

How will we preserve alternatives, reversibility, and recovery?

What would responsible success look like three years after the technology has become ordinary?

Peter J. Dykhuis | *The Human Use of Powerful Systems*

References

Anton, C. (2015). A heuristic sketch of Lee Thayer's systems theory orientation to human communication: Reflections on the range of open systems. *BCSSS*, 3(1), 46 to 69.

Bateson, G. (1972). *Steps to an ecology of mind*. Chandler Publishing.

Carr, N. (2010). *The shallows: What the Internet is doing to our brains*. W. W. Norton.

Ellul, J. (1964). *The technological society* (J. Wilkinson, Trans.). Vintage Books. (Original work published 1954)

Lanier, J. (2010). *You are not a gadget: A manifesto*. Alfred A. Knopf.

McLuhan, M. (1964). *Understanding media: The extensions of man*. McGraw-Hill.

Mumford, L. (1967). *The myth of the machine: Technics and human development*. Harcourt, Brace & World.

- National Institute of Standards and Technology. (2023). *Artificial Intelligence Risk Management Framework (AI RMF 1.0)* (NIST AI 100-1). U.S. Department of Commerce.
- National Institute of Standards and Technology. (2024). *Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile* (NIST AI 600-1). U.S. Department of Commerce.
- Organisation for Economic Co-operation and Development. (2019, updated 2024). OECD principles on artificial intelligence. <https://oecd.ai/en/ai-principles>
- Postman, N. (1985). *Amusing ourselves to death: Public discourse in the age of show business*. Viking.
- Postman, N. (1992). *Technopoly: The surrender of culture to technology*. Alfred A. Knopf.
- Postman, N. (1998, March 27). *Five things we need to know about technological change*. Address to the NewTech '98 Conference, Denver, Colorado.
- Shannon, C. E., & Weaver, W. (1949). *The mathematical theory of communication*. University of Illinois Press.
- Simon, H. A. (1971). Designing organizations for an information-rich world. In M. Greenberger (Ed.), *Computers, communications, and the public interest* (pp. 37 to 72). Johns Hopkins Press.
- Thayer, L. O. (1968). *Communication and communication systems in organization, management, and interpersonal relations*. Richard D. Irwin.
- Thayer, L. O. (1972). Communication. *Main Currents in Modern Thought*, 29(1), 10 to 13.
- Thayer, L. O. (1987). *On communication: Essays in understanding*. Ablex Publishing.
- Turkle, S. (2011). *Alone together: Why we expect more from technology and less from each other*. Basic Books.
- Turkle, S. (2015). *Reclaiming conversation: The power of talk in a digital age*. Penguin Press.
- United Nations Educational, Scientific and Cultural Organization. (2021). *Recommendation on the ethics of artificial intelligence*. <https://unesdoc.unesco.org/ark:/48223/pf0000381137>
- von Bertalanffy, L. (1968). *General system theory: Foundations, development, applications*. George Braziller.
- Weizenbaum, J. (1976). *Computer power and human reason: From judgment to calculation*. W. H. Freeman.

Wiener, N. (1954). *The human use of human beings: Cybernetics and society* (2nd ed.). Houghton Mifflin.
(Original work published 1950)

Zuboff, S. (2019). *The age of surveillance capitalism*. PublicAffairs.

Peter J. Dykhuis | *The Human Use of Powerful Systems*

Author's Note

This paper is a synthetic thought-leadership study rather than a comprehensive intellectual history. It draws on representative primary works and contemporary governance frameworks to develop practical leadership guidance. Page-specific quotations were intentionally minimized in favor of original synthesis and clear attribution. The author reviewed and takes responsibility for the analysis, source selection, claims, and final wording. Future editions may deepen individual theorist chapters, add sector-specific case studies, and incorporate interviews or original empirical research.

About the Author

Peter J. Dykhuis, MS, LSSGB

Peter J. Dykhuis is a practitioner at the intersection of organizational communication, systems thinking, and technology leadership. His work draws on more than two decades of experience helping organizations navigate the human dimensions of complex change, from process improvement and operational design to strategic communication and leadership development.

Trained in systems theory and holding a Lean Six Sigma Green Belt, Peter brings a disciplined analytical perspective to the study of how tools, structures, and information environments shape what people notice, decide, and become. His thought leadership is grounded in the conviction that the most consequential technology questions are not technical; they are organizational, communicative, and ultimately human.

Peter's work is intended for those who want rigorous thinking without unnecessary abstraction: practitioners willing to examine the systems they have built, the messages they have assumed were received, and the purposes they risk losing to efficiency. He is based in West Michigan.

Contact: peter@thoughtline.com

Contact: <https://thoughtline.com/peterdykhuis>