

Cyber Phenomenon Series

The Decision Life Cycle

**Why Every Decision Eventually Expires –
and Why Leaders Must Decide Anyway**

Scott Foote

Last Updated: 9 April 2026

Phenomenati Consulting
www.phenomenati.com

6 Liberty Square, #2736
 Boston, MA 02109
 (508) 709-7990 (office)

CONFIDENTIALITY NOTICE: The contents of this document, including any attachments, are intended solely for stakeholders of Phenomenati Consulting, may contain confidential and/or privileged information, and are legally protected from disclosure.

<this page is intentionally blank>

Contents

Abstract.....	iv
1 The Central Reality: Every Decision has a Half-Life	1-1
2 Why Indecision is Usually the More Expensive Error	2-1
3 The Decision Life Cycle	3-1
4 How the Life Cycle Appears across Executive Roles	4-1
5 Operating Disciplines for Leaders Who Want to Embrace the Life Cycle	5-1
6 Conclusion	6-1
7 References	7-1

A synopsis of reflections and lessons learned from the perspective of a founder, Chief Product Officer, Chief Technology Officer, Chief Information Security Officer, Data Protection Officer, and Chief AI Officer with more than forty years of professional experience.

Core Proposition

A decision should be treated as a living commitment, not a static verdict. It begins with a judgment, continues through execution and monitoring, and ends only when it is reaffirmed, revised, or retired. In a changing environment, the failure to revisit a decision is usually more dangerous than the fact that the original decision was imperfect.

Abstract

At some point, nearly every consequential decision will appear *wrong* from a changed *future* vantage point. That is not a counsel of despair; it is simply a fact, and a challenge for leadership. Over more than forty years across founder, product, technology, cybersecurity, privacy, and AI roles, I have learned that the relevant question is not whether a decision will age. It will. The question is whether we are disciplined enough to decide in time, instrument the decision, detect the moment its assumptions expire, and revise without ego.

This paper argues for a **Decision Life Cycle** view: every material decision should be treated as a time-bounded, reviewable commitment rather than as a permanent verdict. Good framing begins by asking, "*What does success look like?*", translating the answer into explicit decision **criteria**, and documenting a *bounded decision space* that explicitly states **assumptions**, and makes **tradeoffs** visible. Drawing on bounded rationality, decision hygiene, complexity theory, dynamic capabilities, and recent cyber, privacy, and AI governance frameworks, I argue that *indecision* is often more damaging than reversible error. Mature leadership is therefore not the art of being permanently right; it is the practice of making, monitoring, revising, and learning from decisions under ever changing conditions.

1 The Central Reality: Every Decision has a Half-Life

One conclusion has followed me through every leadership role I have held: no serious decision remains *permanently* right. Some decisions are flawed at inception because human beings do not decide with perfect information, unlimited computation, or perfectly stable judgment. Herbert Simon's work on **bounded rationality** remains foundational here, and later work on bias and noise observes that even disciplined professionals vary more in judgment than they usually realize (Simon, 1955; Kahneman, Sibony, & Sunstein, 2021).

What matters even more in practice is that many decisions are not wrong *when made*; they become wrong *later*. A pricing model can fit the market that existed six quarters ago. A security control can be adequate for last year's threat landscape. A privacy position can be proportionate for one data flow and unfit after the system is expanded. An AI governance choice can be sensible for one model generation and inadequate for the next. The one constant – that time will change the conditions and inform the assumptions around every decision.

Let's acknowledge the "spectrum of *wrongness*":

- **Wrong then:** the logic, framing, or evidence was poor at the time of decision.
- **Wrong now:** the decision was reasonable when made, but its assumptions have expired.
- **Wrong in execution:** the decision may have been sound, but implementation undermined it.

Where/when "*wrongness*" is assessed matters because it keeps review honest. Without it, organizations punish adaptation, confuse hindsight with foresight, and learn the wrong lesson. Mature review asks whether the *premise* failed, the *environment* changed, or the *execution* was deficient.

In **product** and **market** work, the tension between exploiting what is *known* and exploring what is *emerging* never goes away. March described organizations as living inside that tradeoff between *old certainties* and *new possibilities* (March, 1991). That is why a roadmap, channel strategy, or pricing choice that once looked prudent can later become a liability.

In **technology**, every *architecture* is a *tradeoff* frozen in time. The design that maximizes speed for a start-up may later impair resilience, observability, or compliance. Dynamic capabilities scholarship is useful here because it treats sustained performance not as *static* fit but as the ability to integrate, build, reconfigure, and *adapt* competencies as the environment shifts (Teece, Pisano, & Shuen, 1997; Teece, 2007).

In **cybersecurity**, *sufficiency* decays. Threat actors adapt, the business changes its attack surface, vendors introduce new dependencies, and exceptions harden into invisible exposures. NIST's Cybersecurity Framework 2.0 explicitly frames cybersecurity through governance and a life-cycle view of risk management rather than a one-time checklist (NIST, 2024b).

In **privacy**, *accountability* is also ongoing rather than ceremonial. The GDPR requires controllers not only to comply but to demonstrate compliance, and practical guidance stresses technical and organizational measures that are reviewed and updated as necessary (European Union, 2016; Information Commissioner's Office, 2026).

In **AI**, the half-life of a decision is shorter still. Model capabilities *change*, misuse patterns *evolve*, data distributions *drift*, and acceptable-use expectations *move* under regulatory and social pressure. NIST's AI Risk Management Framework was designed to adapt as the AI landscape develops, and the 2024 Generative AI Profile extends that logic to the distinctive risks of generative systems (NIST, 2023; NIST, 2024a). A revision is likely to be released in 2026 that addresses the more significant risks of agentic platforms and rushed and reckless deployments.

2 Why Indecision is Usually the More Expensive Error

Fear of being wrong is understandable. What I have seen far more often, however, is leaders confusing *caution* with *quality*. Insecurity frequently disguises itself as more analysis, more socialization, more committees, or more process. The language stays professional, but the underlying driver is often *fear*: fear of ownership, fear of criticism, fear of reversals, or fear of being the person whose judgment becomes visible.

That fear creates what I call **decision debt**: the hidden cost of allowing circumstances, committees, or systems to decide *implicitly* through their inaction. If a product leader delays a roadmap choice, engineering capacity still gets consumed - just less coherently. If a CTO delays an architecture decision, teams still produce architecture - only accidentally. If a CISO delays a control decision, exposure continues to accumulate. If a DPO delays a retention or DPIA judgment, the organization still collects and reuses data. If a Chief AI Officer delays governance, pilots still proliferate. In each case, the absence of a formal decision is itself a decision.

Research in high-velocity environments is instructive here. Eisenhardt found that fast strategic decision-makers were not necessarily using less information or thinking less carefully; they used current information, explored alternatives, and resolved conflict in ways that preserved decision velocity (Eisenhardt, 1989). In moving environments, speed is not the *enemy* of quality. Very often, speed is one *dimension* of quality.

The practical corollary is that **decision confidence** should be calibrated to reversibility and cost of delay. Jeff Bezos popularized the distinction between *one-way-door* and *two-way-door* decisions: *reversible* decisions (two-way-doors) deserve lighter-weight processes, while more *irreversible* decisions (one-way-doors) deserve more deliberation. He also argued that waiting for *near-complete information* is often too slow and that organizations need to become good at quickly recognizing and correcting bad calls (Bezos, 2017). This is not an argument for impulsiveness. It is an argument for matching process to decision type.

It's worth an aside here to discuss the critical difference between being "**decisive**" and being "**impulsive**", given that contemporary press regularly confuse the two, celebrating the latter with praise and celebrity without acknowledging the critical downstream impacts.

Being decisive is not the same thing as being fast. **Decisiveness** is disciplined commitment under *uncertainty*: the ability to frame the problem, clarify what success looks like, weigh the relevant decision criteria, examine a finite trade space, and then choose with enough conviction to move the organization forward. A decisive leader understands that *perfect information never arrives*, but still acts from a position of considered judgment. **Impulsiveness**, by contrast, is action *without* sufficient framing or reflection. It often feels energetic, bold, and confident in the moment, but it is usually driven more by emotion, ego, urgency, or discomfort with ambiguity than by an explicit understanding of alternatives, consequences, and assumptions.

The practical difference shows up in what happens next. A decisive person can explain *why* a choice was made, what *assumptions* supported it, what *risks* were accepted, and what *signals* would trigger reevaluation. In other words, decisiveness is *accountable* and *reviewable*. Impulsiveness usually leaves behind little more than motion and justification after the fact. One

creates organizational trust because people can see the reasoning, even when the outcome later proves wrong. The other *erodes trust* because it ***confuses speed with judgment***. In leadership, especially in high-velocity environments, the goal is not to avoid delay by abandoning thought; it is to make timely decisions with enough structure that action is informed, reversible where possible, and open to learning as circumstances change.

Indecision, the worst form of paralysis, is organizational exhaustion masquerading as consensus. Meetings multiply. Language gets abstract. Stakeholders wait one another out. No explicit owner decides, so the decision is ultimately made by *inertia*, loudest voice, political stamina, or calendar pressure. That is not prudence. It is unmanaged risk.

3 The Decision Life Cycle

Because decisions expire, they should be managed like living commitments. I use the term Decision Life Cycle to describe the full path from framing through retirement, and centered on continuous learning. The broader idea - that decisions should be recorded, revisited, and changed across their full life cycle, not merely at the point of initial choice - is well grounded in *decision support* literature (MacKenzie et al., 2006).

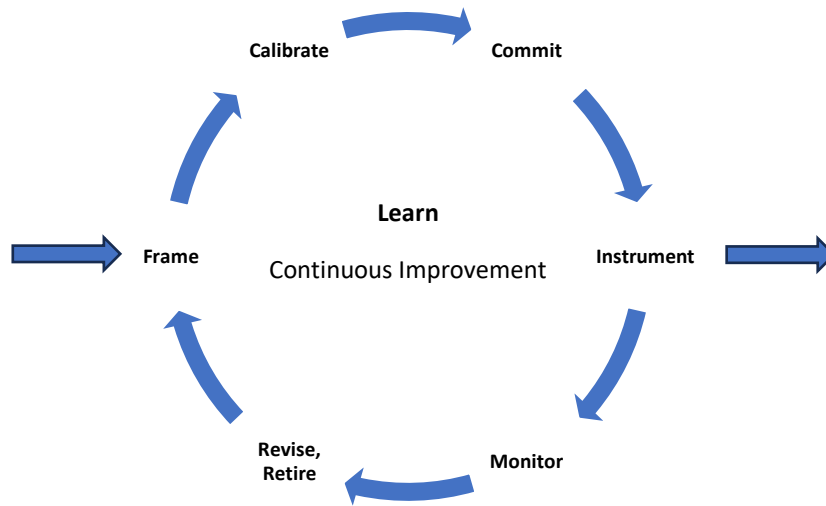


Figure 1: The Decision Life Cycle

Stage	Core Question	Minimum Artifacts
Frame	What are we deciding, for whom, under what conditions, and <i>what does success look like?</i>	Problem statement, decision owner, <i>success</i> definition, decision <i>criteria</i> , initial <i>trade space</i> , time <i>horizon</i>
Calibrate	How reversible, urgent, and regulated is this choice?	Decision class, escalation path, evidence threshold
Commit	What will we do now, and what had to be true?	Selected option, rationale against criteria, key tradeoffs, key assumptions, noted dissent
Instrument	How will we know whether to stay, adapt, or reverse?	Metrics, tripwires, review date, rollback conditions
Monitor	What is changing, or has changed, in outcomes or context?	Customer, technical, risk, legal, and operational signals
Revise or Retire	Should we reinforce, adapt, reverse, or stop?	Change record, superseding decision, communication plan
Learn	What rule, playbook, or principle should change?	Postmortem, updated standards, archived decision log

Table 1: Phases of the Decision Life Cycle

The **first** stage is **framing**. The real question is not simply 'What do we prefer?' but 'What *problem* are we truly solving, for *whom*, under what *conditions*?' In my own career, framing has long been founded on a single starting question: "**What does success look like?**" Unless leaders can answer that clearly, they are usually debating preferences, personalities, or technologies rather than making a decision. Bad framing creates *false certainty*. In complex contexts, *cause* and *effect* are often clear only in retrospect, which means *experimentation* and *iterative* sensing are more reliable than prediction dressed up as certainty (Snowden & Boone, 2007).

The answer should not remain a slogan. It should be translated into a reasonable set of **decision criteria** relevant at the time of the choice: dates, cost, expected return, risk, quality, strategic alignment, regulatory exposure, operational impact, reversibility, and any other factor that materially shapes the judgment. Those criteria do two jobs. They discipline the original discussion by making tradeoffs explicit, and they improve later review by showing which assumptions were embedded in the decision at the time. As circumstances change, some of those *criteria will change* with them (Kahneman, Sibony, & Sunstein, 2021; Teece, 2007).

Options should then be documented inside a finite **Decision Space**, or *trade space*: a bounded set of plausible alternatives stated clearly enough that their benefits, costs, dependencies, and risks can be compared. Not an endless brainstorm, and not a false binary, but a manageable set of real choices. A documented trade space improves judgment in the moment and strengthens learning later, because it preserves not only what was *chosen* but what was consciously *rejected* and why (Simon, 1955; Kahneman, Lovallo, & Sibony, 2011).

The **second** stage is **calibration**. Leaders should classify the decision before they glorify it. Is it reversible or irreversible? Local or enterprise-wide? Operational or safety-critical? Lightly coupled or tightly coupled to privacy, cyber, legal, and financial risk? Calibration determines the appropriate level of rigor, documentation, and escalation.

The **third** stage is **commitment**. A decision without commitment is commentary. At this point the leader selects a course from the documented decision space, records why it best satisfied the relevant criteria *at that moment*, and identifies the assumptions that had to hold for the decision to make sense. Decision quality improves when alternatives are explicitly considered and cognitive distortions are challenged, not when commitment is indefinitely postponed (Kahneman, Lovallo, & Sibony, 2011).

The **fourth** stage is **instrumentation**. This is the *most neglected* stage in most organizations, and its presence/absence is a strong indicator of the experience and competency of a leadership team. Every material decision should have an owner, success measures derived from the agreed criteria, leading indicators, tripwires, a review date, and reversal conditions. Without instrumentation, the organization remembers the *conclusion* and forgets the *assumptions*.

The **fifth** stage is **monitoring**. This is where reality answers back. Customers behave differently than expected. Systems strain in new ways. Exceptions accumulate. Threats evolve. Regulators clarify expectations. Models drift. Dependencies move. Monitoring means watching for *assumption drift* before failure becomes public.

The *sixth* stage is *revision* or *retirement*. Mature leaders do not cling to yesterday's rationality once conditions change. They strengthen a decision if the evidence supports it, adapt it if the environment has shifted, reverse it if the premise has failed, or retire it if the business context has moved on. Refusing to revisit a decayed decision is usually vanity disguised as conviction.

At the center of this life cycle is *learning* that drives continuous *adaptation* and *improvement*. The obligation of all decision making (professional or even personal) is not simply to remember what happened, but to capture the lesson at the correct level. Some lessons are local. Others belong in architecture standards, risk thresholds, product principles, data governance rules, or AI deployment policy. Over time, this is how organizations balance *exploration* with *exploitation* and convert experience into dynamic capability (March, 1991; Teece, 2007).

Every Material Decision Deserves a Record

Decision statement and accountable owner; definition of success and decision criteria in force at the time; date, time horizon, and reversibility classification; documented decision space or trade space, including options considered and the option chosen; key assumptions and material dependencies; metrics, tripwires, and review date; retirement or rollback conditions.

4 How the Life Cycle Appears across Executive Roles

The vocabulary changes by function, but the *pattern* does not. In each role, a decision begins as a *hypothesis*, generates operating *consequences*, and must eventually be revisited.

Role	Common Trap	Life-Cycle Discipline
Founder	Waiting for certainty on market, people, or capital allocation	Use staged bets, explicit revisit points, and runway-aware learning
Chief Product Officer	Over-researching roadmap choices or confusing feedback with unanimity	Frame hypotheses, ship in increments, and review leading indicators
Chief Technology Officer	Treating a provisional architecture as permanent or endlessly debating platforms	Use architecture decision records with expiry dates and review triggers
Chief Information Security Officer	Treating controls or risk exceptions as static	Use threat-informed reviews, exception expiries, and rapid escalation paths
Data Protection Officer	Treating compliance artifacts as one-time paperwork	Refresh records, DPIAs, retention schedules, and vendor reviews at defined intervals
Chief AI Officer	Either shipping without governance or waiting for perfect policy	Use phased deployment, red-teaming, monitoring, human oversight, and rollback conditions

What changes from role to role is not the life cycle itself but the *feedback latency*. Product decisions may show feedback in days or weeks. Architecture decisions may reveal their cost over years. Privacy and AI governance decisions often have long quiet periods followed by sharp regulatory, operational, or reputational discontinuities. The life cycle still applies; only the telemetry changes.

5 Operating Disciplines for Leaders Who Want to Embrace the Life Cycle

Separate identity from decisions. If leaders treat reversal as humiliation, teams suppress evidence and reinterpret reality to defend old commitments. The organization becomes *performative* instead of *adaptive*.

Put expiry conditions on assumptions. Assumptions should not live forever. Major decisions need either *calendar-based* review points, *event-based* triggers, or both. Risk exceptions should expire. PIAs and DPIAs should be refreshed. Model approvals should be revisited. Architecture choices should have review criteria.

Use both leading and lagging indicators. By the time a lagging indicator confirms failure, value has usually already been destroyed. Good instrumentation watches for early signs of context drift, not only for final outcomes.

Reward early correction. A leader who reverses a deteriorating decision early has demonstrated competence, not weakness. The organization should learn to applaud disciplined correction before external reality turns it into an embarrassing crisis.

Build decision memory. *Decision logs* prevent endless re-litigation, make audit and accountability easier, and preserve the link between the chosen option, the rejected alternatives, the definition of success, the *criteria* in force, and the *context* of the judgment (MacKenzie et al., 2006; European Union, 2016; NIST, 2023).

Govern portfolios, not only single decisions. The most damaging failures are often cross-functional. Product wants faster data collection, privacy wants minimization, security wants control, technology wants scale, and AI wants automation. Senior leaders must *govern* how decisions interact, not merely how each looks in isolation.

Useful Questions for a Decision Review

- What *did* success look like, and what had to be true, when we made this decision?
- Which decision *criteria* mattered most at the time?
- Which *assumptions* or *criteria* are now weaker, false, or no longer relevant?
- Is the problem in the premise, the environment, or the execution?
- Which rejected option now deserves reconsideration?
- What is the smallest viable correction?
- What rule, standard, or playbook should change because of this lesson?

6 Conclusion

After more than four decades across roles including founder, product, technology, security, privacy, and AI leadership, I no longer believe that the mark of an excellent leader is the ability to be permanently right. That standard is impossible, and pretending otherwise produces brittle institutions. The better standard is to decide with appropriate speed, define success explicitly, document decision criteria and tradeoffs, instrument consequences, revisit without ego, and reverse without shame.

In that sense, the Decision Life Cycle is more than a governance technique. It is an *ethic* of leadership. It *replaces* bravado with humility, impulsiveness with decisiveness, paralysis with disciplined action, and blame with learning. It also insists that leaders make their premises inspectable:

- what success looked like,
- which criteria mattered, and
- which alternatives formed the trade space.

A decision is not a monument to defend. It is a commitment to manage until reality changes - and reality *always* changes.

7 References

- Bezos, J. (2017). 2016 Letter to Shareholders. Amazon.com, Inc.
 URL: <https://www.aboutamazon.com/news/company-news/2016-letter-to-shareholders>
- Eisenhardt, K. M. (1989). Making Fast Strategic Decisions in High-Velocity Environments. *Academy of Management Journal*, 32(3), 543-576.
 URL: <https://doi.org/10.5465/256434>
- European Union. (2016). Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 (General Data Protection Regulation).
 URL: <https://eur-lex.europa.eu/eli/reg/2016/679/oj/eng>
- Information Commissioner's Office. (2026). Guide to accountability and governance.
 URL: <https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/accountability-and-governance/guide-to-accountability-and-governance/>
- Kahneman, D., Lovallo, D., & Sibony, O. (2011). Before You Make That Big Decision. *Harvard Business Review*, 89(6), 50-60, 137.
 URL: <https://hbr.org/2011/06/the-big-idea-before-you-make-that-big-decision>
- Kahneman, D., Sibony, O., & Sunstein, C. R. (2021). *Noise: A Flaw in Human Judgment*. Little, Brown Spark.
 URL: <https://www.hachettebookgroup.com/titles/daniel-kahneman/noise/9780316451406/>
- MacKenzie, A., Pidd, M., Rooksby, J., Sommerville, I., Warren, I., & Westcombe, M. (2006). Wisdom, decision support and paradigms of decision making. *European Journal of Operational Research*, 170(1), 156-171.
 URL: <https://doi.org/10.1016/j.ejor.2004.07.041>
- March, J. G. (1991). Exploration and Exploitation in Organizational Learning. *Organization Science*, 2(1), 71-87.
 URL: <https://doi.org/10.1287/orsc.2.1.71>
- National Institute of Standards and Technology. (2023). Artificial Intelligence Risk Management Framework (AI RMF 1.0). NIST AI 100-1.
 URL: <https://nvlpubs.nist.gov/nistpubs/ai/nist.ai.100-1.pdf>
- National Institute of Standards and Technology. (2024a). Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile. NIST AI 600-1.
 URL: <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.600-1.pdf>
- National Institute of Standards and Technology. (2024b). The NIST Cybersecurity Framework (CSF) 2.0. NIST CSWP 29.
 URL: <https://nvlpubs.nist.gov/nistpubs/CSWP/NIST.CSWP.29.pdf>
- Simon, H. A. (1955). A Behavioral Model of Rational Choice. *Quarterly Journal of Economics*, 69(1), 99-118.
 URL: <https://doi.org/10.2307/1884852>
- Snowden, D. J., & Boone, M. E. (2007). A Leader's Framework for Decision Making. *Harvard Business Review*, 85(11), 68-76.
 URL: <https://hbr.org/2007/11/a-leaders-framework-for-decision-making>

Teece, D. J. (2007). Explicating Dynamic Capabilities: The Nature and Microfoundations of (Sustainable) Enterprise Performance. *Strategic Management Journal*, 28(13), 1319-1350.

URL: <https://doi.org/10.1002/smj.640>

Teece, D. J., Pisano, G., & Shuen, A. (1997). Dynamic Capabilities and Strategic Management. *Strategic Management Journal*, 18(7), 509-533.

URL: <https://sms.onlinelibrary.wiley.com/doi/abs/10.1002/%28SICI%291097-0266%28199708%2918%3A7%3C509%3A%3AAID-SMJ882%3E3.0.CO%3B2-Z>