The Decision That Could Not Expire

Why organisations stop governing a decision on the day it is made, and what to write down before the work starts

Gareth A Holebrook
Published on Kettei, 2 October 2026: kettei.substack.com/p/the-decision-that-could-not-expire
© 2026 PerformO Publishing. Written by Gareth A Holebrook. Free to copy and use in your own organisation, with this credit kept on it.

A question that reached nobody's desk in time

On 30 August 2011 Boeing launched the 737 MAX. Its own announcement called the aircraft the new engine variant of the market-leading 737, a family that builds on the strengths of the Next-Generation 737. That was the founding decision of the programme: build the MAX as a derivative of the existing 737, with no new type certificate. Two commitments hung from it. The aircraft would be certified as a change to the existing type, under the rule that lets changed areas of a design be assessed while unchanged areas are left alone. And pilots already qualified on the 737 would move to the MAX with differences training on a computer, with no time in a simulator.

Within 4 months the second commitment acquired a price. Under a contract signed in December 2011 with Southwest Airlines, the US launch customer, Boeing was obliged to discount each MAX it delivered to Southwest by at least $1 million if the Federal Aviation Administration required simulator training for pilots moving across from the older 737. Southwest had 200 firm orders and options on 191 more. The US House Committee on Transportation and Infrastructure put the exposure at between $200 million and nearly $400 million. Reversing the training commitment now cost money that a contract had already specified.

The assumption that eventually failed concerned handling. Larger engines sat further forward, which changed how the aircraft pitched, and Boeing added flight-control software to compensate: the Maneuvering Characteristics Augmentation System, MCAS. The safety assessment the FAA saw described MCAS moving the horizontal stabiliser by up to 0.6 degrees.

In March 2016 the general manager of the 737 MAX programme and its chief project engineer approved a redesign that increased the system's authority at low speed. Boeing later gave airlines a limit of 2.5 degrees, a figure the FAA engineers who had worked from 0.6 had not seen.

The House Committee's final report records what happened next. Within hours of that approval, Boeing asked the FAA to remove references to MCAS from the Flight Crew Operations Manual, the document a flight crew relies on to understand the aircraft, and the FAA agreed. The committee could not establish whether the 2 acts were coordinated inside Boeing. What it did establish is that the FAA officials who approved the removal did not learn of the redesign until after the first crash. The regulator that signed off one decision did not know about the other, and both rested on the same founding assumption.

On the evening of 28 October 2018 a Lion Air crew flying a 737 MAX 8 from Denpasar to Jakarta met repeated automatic nose-down trim. The crew cut out the electric trim and finished the flight trimming by hand. Maintenance received a report that did not fully describe what had happened, and released the aircraft for service. The next morning it took off from Jakarta as Lion Air Flight 610 with 189 people aboard and crashed into the Java Sea minutes after take-off. On 10 March 2019 Ethiopian Airlines Flight 302 crashed after take-off from Addis Ababa. Together the 2 crashes killed 346 people. The FAA grounded the type on 13 March 2019 and cleared it to return in November 2020. In January 2021 Boeing entered a deferred prosecution agreement with the US Department of Justice worth more than $2.5 billion.

I am not an aviation engineer and I never worked at Boeing. I read the inquiry reports with the decision as the unit of analysis, where most accounts take the outcome. Aviation has rules for reconsidering a derivative. A change can be so extensive that a new type certificate is required, and the FAA's own guidance says that evolutionary changes during a certification programme may require the certification basis to be re-evaluated. The Joint Authorities Technical Review, the panel of regulators the FAA convened after the crashes, found that the changed-product process was followed and worked for discrete changes, and that it did not adequately address cumulative effects or how changes interacted across the whole aircraft. It found that design assumptions went without adequate review, update or validation as MCAS evolved, and that the certification documents were left unrevised. The House Committee found a chief project engineer who approved key design decisions while supervising none of the engineers who made those decisions, a structure it said contributed to an overall lack of accountability.

So the rules and the mechanisms existed. What neither produced, when MCAS quadrupled in authority, was a reassessment of the founding assumption against the whole aircraft, by a person accountable for that assumption, with the regulator informed. Regulators asked whether the aircraft was still a derivative in the sense its certification depended on in 2019, after 2 crashes.

Everything has an owner except the decision

Here is the claim in its shortest form. Organisations make consequential decisions with care, usually with a business case and a signature, and then govern the work that follows far more closely than the decision itself. The programme has an owner. The budget has an owner. The risk register has an owner. The continuing validity of the decision that authorised all 3 belongs to nobody.

Nobody writes down what would have to remain true for the decision to stay valid, so no assumption can expire. Nobody records what depends on it, so when it changes nothing downstream is told. A reversible commitment gets defended with the same ferocity as an irreversible one, because nobody classified which it was. The conditions under which it ends were never agreed. The decision cannot expire, cannot be challenged by events and cannot be surfaced for review when reality moves. It can only persist, and it does, until stopping looks more expensive than going on, whatever the arithmetic says.

The research on why organisations keep going explains the rest. In Barry Staw's 1976 experiment, people who felt personally responsible for a failing decision committed more resources to it after bad news. Diane Vaughan's study of the Challenger launch decision described how the standard for acceptable risk moved a little with each flight that returned safely, until the deviation had become normal. NASA had a written expectation that the O-rings would seal without erosion, and that expectation was itself reinterpreted flight by flight, so writing a condition down is no guarantee on its own. What a condition fixed in advance gives a review is a stable reference, and it holds best in hands other than the people defending the current rationale. Simonson and Staw showed the effect in the laboratory in 1992: participants who wrote down minimum targets before receiving performance feedback put markedly less money into a losing option than the baseline group, while thorough deliberation on its own changed nothing.

Escalation of commitment is structural. A culture programme can make people more willing to raise a doubt, and the incentives in the room decide whether raising it is safe. What a culture programme does not do is name the person who owns a decision's validity, write down the condition that would end it, or oblige anyone to look on a date. Without those, the doubt has nowhere to land.

The same failure in a shop in Christchurch

I have run this failure myself, at the scale of one studio, and I have written about the personal cost before. Here I want the decisions.

In January 2004 I raced my first Half Ironman in Auckland, where an altitude training franchise called BodyO_2 was the naming sponsor. At its tent after the race I asked whether a studio was planned for Christchurch, where I lived. The franchisor said the company was looking for franchisees. I bought a coffee from a burger van, thought about it for 10 minutes and went back for a business card. On 19 April I opened BodyO_2 Christchurch.

The founding decision was that the franchise would be the revenue-generating entry point to a technology startup. Its earnings would fund the development of my own performance technology. That decision stood on 2 assumptions I never wrote down: that the studio would earn enough to fund development, and that the franchise behind it would stand. By September my cash flow was gone. I borrowed from the bank and finance companies and brought in an investor to buy time. By the end of the year, revenue for the entire business was lower than a single salary. The first assumption had failed inside 9 months, and nothing obliged anyone, least of all me, to say so.

The second failed from upstream. The franchise holding company collapsed, its assets were sold to another franchisee and legal conflict followed. Nothing in my business brought the founding decision back for review, because I had never written down that it depended on the franchisor at all. I exited the franchise and continued under my own brand, PerformO, from 2006 to 2010, using portable altitude simulators I had helped develop. The simulators were prototypes and not robust enough for the field, and failed my clients. By the time I closed, most of my revenue came from retailing nutrition and clothing. The technology startup that justified the whole enterprise was stillborn. The business had become a retailer, and at no point had I decided it would.

The ownership question is the uncomfortable one. I owned that decision completely. Authority, capital and conviction were all mine. What I did not have was any obligation attached to the ownership: no condition to watch, no date on which I was bound to look, no person entitled to ask. Ownership without an obligation to watch validity is authority without governance, and the owner is usually the person least able to see it.

The signals were there, and none had a route into a decision. In February 2005 the bank withdrew my credit line. One of my staff reduced her hours and accepted shares in lieu of pay. She could see the reality I was still denying, and nothing in how the business ran gave her view any standing. It cost me the best part of a million New Zealand dollars.

A single franchise studio in Christchurch bears no comparison with Boeing in scale, and I am not equating the 2. Boeing wrote its assumptions into a certification basis and had rules for revisiting that basis. I wrote nothing down and had no rules. What the cases share is that a founding assumption failed and nothing brought the founding decision back to a person obliged to reassess it in time.

The decision as an object

Kettei (決定) is an ordinary Japanese word for a decision or a determination. I use it for the decision in its binding sense, the commitment as an established fact, because Japanese business practice has separated the stages of deciding for decades. Nemawashi builds agreement informally, ringi circulates the formal proposal for approval, and kettei is the moment the decision becomes binding. What I am proposing keeps the kettei alive after the signature.

In my first book I gave individuals a handwritten card for this. A card works when the intent remains inside one skull. In an organisation the decision is spread across many owners who depend on each other, and personal honesty does not propagate from one desk to the next. Structure does.

The scope is narrow on purpose. This governs the steering decisions inside an endeavour, the handful of choices that set its direction, constraints and commitments and that waste months or millions when wrong. Board strategy, regulatory certification and specialist safety assessment sit outside it, as constraints it operates inside. Boeing's founding decision is a steering decision. The certification and the MCAS safety assessment are constraints, and both depended on it.

The principle is Decision as a Governable Object, which I shorten to DAAGO. A consequential decision is treated as a first-class object with 7 mandatory properties, defined before it moves from proposed to active. If any property cannot be defined, the decision stays Proposed.

PROPERTY	WHAT IT ANSWERS
Owner	One person accountable for the decision's continuing validity. Never a committee, never shared and never vacant.
Reversibility classification	How hard the decision is to unmake. The harder it is, the tighter the governance.
Lifecycle state	Proposed, Active, Suspended, Superseded or Revoked. Every change of state is recorded with a date, a name and the evidence.
Coupling dependencies	What this decision depends on and what depends on it, so a change upstream reaches everything downstream.
Scope	What the decision covers, what it leaves out and what authority it delegates.
Abort criteria	The conditions under which the decision is revoked outright, set before execution begins.
Validity boundary	The assumptions the decision stands on, a date by which it must be reviewed regardless, and the triggers that force an earlier review.

The owner is one person accountable for the decision's continuing validity. Never a committee, never shared, never vacant. The owner watches the boundary and answers when a review is called. Reversibility classification records how hard the decision is to unmake, and it sets the

governance weight: the more irreversible, the tighter the boundary and the stricter the conditions for stopping. It is fixed at creation. If accumulated commitment later makes the decision harder to reverse, that entrenchment is itself a trigger for review.

The lifecycle state is one of 5: Proposed, Active, Suspended, Superseded or Revoked. Every transition is recorded, with the evidence. Coupling dependencies declare what the decision depends on and what depends on it. When a decision is superseded or revoked, every decision that declared a dependency on it is reviewed by its own owner. When a decision is suspended, its dependents go on watch, with no new irreversible commitments until the review resolves. Scope and exclusion say what the decision covers, what it leaves out and what authority it delegates.

The last 2 are the ones almost nobody writes, and the discipline keeps the 2 apart on purpose. Abort criteria are the conditions under which the decision is revoked. Revoked means stopped, with no revalidation, because those conditions were agreed while stopping was cheap and stay closed to renegotiation now that it is expensive. The validity boundary names the assumptions the decision stands on, a horizon after which review is mandatory whatever the news, and the triggers that force an earlier review. A trigger compels a conversation. An abort criterion commits the organisation, in advance, to an ending. Each trigger has a signal, a threshold, an owner and a response window. A calendar review is the thinnest possible validity boundary. The difference is between asking in June whether the strategy is still right and being told in February that the assumption it rested on has stopped holding.

Now put Boeing's founding decision beside those properties. Aviation already had the reversibility question, in the rule on when a change is extensive enough to need a new type certificate, and a standing review of discrete changes. What this discipline adds is a named owner of the derivative assumption itself, a written trigger on that assumption, such as any increase in the authority of the software compensating for the new engines, and a declared coupling between the founding decision, the MCAS design and the training commitment, so that a change to any one surfaces the other 2. The March 2016 redesign would have crossed that trigger. A trigger creates an obligation: a review, by a named person, inside a response window, with the regulator told. Whether the review would have been honest, and whether anyone would have acted on it, depends on the people in the room, and the limits below say how often the answer is no. The obligation is the part that can be written down.

The square nobody has filled

I compared 7 established systems for deciding and learning: OODA, Plan-Do-Study-Act, double-loop learning, The Flow System, Toyota Kata, Amplio and DecisionProtocol. Every one has a Decide. None has a lifecycle for the decision after it is made. Architecture decision records, which Michael Nygard described in 2011, gave software decisions an identity, a status and supersession by reference. US bank supervisory guidance on quantitative risk tools, SR 11-7 from the same year, required ongoing monitoring of whether changed conditions meant a tool had to be adjusted or replaced, and a review of each one at least once a year. Both do for a narrow class of object what I am describing for general organisational commitments. For those commitments, in everything I reviewed through 2026, the square is empty. The best coverage I found in any single method was 4 of the 7 properties fully present and a fifth in part, and the 2 missing from every method were the validity boundary and the governed authority to stop.

Each individual property has prior art, and I name it. Reversibility has Jeff Bezos's doors, the ones you can walk back through and the ones you cannot, and Pankaj Ghemawat's work on commitment before him. Kill criteria belong to Annie Duke. Tripwires belong to Chip and Dan Heath. Assumption-based planning with signposts came out of RAND in 1993. What nobody had done, in the sources I reviewed, was make the properties mandatory, interdependent and enforceable at the level of the single decision. Others have the features. Nobody I have read has the discipline. If you know a method that holds all 7, I want to read it.

The contrast is useful. On 18 September 2011 Reed Hastings published a blog post that opened with "I messed up" and announced that Netflix would split its DVD service into a separate business called Qwikster. Customers revolted. On 10 October, 3 weeks later, Netflix killed Qwikster and kept the price increase it had announced in July. One owner, a fast reversal and a clear separation between the part of the decision that had failed and the part that still held. In that same September post Hastings wrote that he had slid into arrogance based upon past success.

The public record shows one owner and a fast reversal. It shows no condition written in advance that would bring the split back for review, and I have found none. The trigger was customers leaving, discovered after the fact. A reversal that fast is rare, and structure is what would make it ordinary in organisations whose leaders cannot count on personal standing to carry one.

What pre-registered stopping can and cannot do

I owe you the limits. Pre-registered stopping conditions create an obligation: a record, a named owner and a conversation due on a clock. Whether the conversation happens depends on somebody watching for the breach and honouring the obligation, and even then the conditions

do not reliably produce stopping. US law has required the Pentagon to report a major weapons programme to Congress when its unit cost breaches set thresholds since the early 1980s, and since 2009 a critical breach requires termination unless the Secretary of Defense certifies the programme to continue. The Government Accountability Office counted 74 breaches across 47 programmes between 1997 and 2009 and could name 7 programmes terminated. Most of that record predates the 2009 rule and does not test it. In the GAO's own verdict, by the time the mechanism triggers the programme has usually already had significant problems. A breach creates the moment. Human beings still have to act in it.

There is also a cost at the start. Ariella Kristal and Charles Dorison published preregistered experiments with 7,759 participants in the Journal of Applied Psychology, online in December 2024. Leaders who had publicly committed in advance to change course under a stated future condition were trusted more on later winding a failing project down. The same precommitment made those leaders look less confident, and less motivating to followers, at the outset. Both halves hold for what I am proposing. The second is the objection I expect first in a steering meeting, and that is my prediction from experience, which the study does not test.

I would take that trade every time. A decision that holds its own ending conditions looks a little less heroic on the day it is announced. It is also the kind that can be stopped with trust intact, which is the half of that finding that pays.

What to write down this week

So I will ask you directly. Pick one decision your organisation made more than a year ago that is still consuming budget and people's time. Who owns its continuing validity today, by name? What would have to remain true for it to stand? What does it depend on, and what depends on it? Under what observable condition would it stop? When is anyone obliged to look?

If those questions take longer to answer than to read, find out why before the next steering meeting. Somebody may be governing the decision's validity in a way that has never been written down. More often, in my experience, the work is governed far more closely than the decision that authorised it. That is what the 737 MAX programme and my studio in Christchurch had in common: a founding assumption failed, and the decision resting on it went unreassessed until too late.

The instrument I use is a single page per endeavour, which I call a Watchcard. It holds the intent of the endeavour, the conditions that must not change, the few decisions steering it with the owner and state of each, and the triggers being watched. Write one this week for the endeavour

that worries you most. Name the owner of its founding decision. Under that, write the assumptions it stands on, the condition that would stop it and the date someone is obliged to look. Put it on the wall of the room where the steering meeting happens.

Agree the ending conditions now, while the case for the decision is still fresh and nobody has had to defend it against bad news. Every month that passes adds another reason to keep going and another person invested in not asking. The conditions only get harder to write.

These are the 16 failure modes my research identified. I have experienced every one of them, at one or several of the companies I have worked for.

PHASE	FAILURE MODE	WHAT IT LOOKS LIKE
Decide	Decision Vacancy	No single person holds the decision's continuing validity. "Who is driving this?"
Decide	Authority Capture	Authority taken in a crisis stays after the crisis, and commitments made inside execution are never declared as decisions.
Decide	Strategic-Tactical Chasm	A board decision reaches delivery as a task list, with no scope or conditions the people doing the work can govern against.
Decide	Coupling Blindness	2 coupled decisions are treated as independent. One changes and the other never hears.
Forge	Unforged Governance	A practice is adopted and adapted before anyone drilled it to competence.
Forge	Documentation-Reality Divergence	The written process and the process people actually run have separated.
Forge	Capability-Governance Mismatch	Decision authority grows faster than the capability to govern it.
Persist	Assumption Drift	The world moved and the decision stayed put. "We decided this last year."
Persist	Trigger Register Decay	Triggers sit on paper and nobody watches the register.
Persist	Normalisation of Deviance	Deviations become background noise. "That error rate is normal for us."
Persist	Signals Ignored	The trigger fires and the owner does not act.
Persist	Escalation of Commitment	Persisting because of what has been spent. "We have come too far to stop."
Persist	Ungoverned Persistence	No triggers and no abort criteria, so nothing defines when to stop.
Reflect	Governance Theatre	Every dashboard green and every forum held, with no decision ever changed by either.
Reflect	Reflection Theatre	The post-mortem is filed, and 12 months later the same failure recurs.
Reflect	Institutional Moralisation	Questioning the decision reads as questioning the institution, so it is never revalidated.

Kettei: The Discipline of Governed Intent is on Leanpub: https://leanpub.com/kettei-dogi

About the author

Gareth A Holebrook is the author of *Ishi: The Discipline of Determined Intent*, with a foreword by L. David Marquet, author of *Turn the Ship Around!*, and is writing *Kettei: The Discipline of Governed Intent*, on governing organisational decisions after the signature. He trained as a weapons engineer in the Royal Navy and spent nearly 3 decades in technology, from product development at Sony and programme management at Navman as it grew from 30 to 200 engineers, to building and leading a 30-person delivery practice at director level at Fraedom, a Visa subsidiary. In recent years he has coached chief executives and leadership teams on organisational culture, governance and leadership behaviour. He finished the Ironman World Championship in Kona in 2022, runs a smallholding in Northland, New Zealand, with his wife Thiri, and writes weekly on Ishi at ishidodi.substack.com.

Sources

Bezos, J. (2016). Letter to shareholders, Amazon.com, for the 2015 fiscal year.

Board of Governors of the Federal Reserve System and Office of the Comptroller of the Currency (2011). SR 11-7, *Supervisory Guidance on Model Risk Management*, 4 April 2011.

Boeing (2011). Boeing introduces 737 MAX with launch of new aircraft family. Press release, 30 August 2011.

Dewar, J. A., Builder, C. H., Hix, W. M. and Levin, M. H. (1993). *Assumption-Based Planning: A Planning Tool for Very Uncertain Times*. RAND.

Duke, A. (2022). *Quit: The Power of Knowing When to Walk Away*. Portfolio.

Gates, D. (2019). Flawed analysis, failed oversight: how Boeing, FAA certified the suspect 737 MAX flight control system. *The Seattle Times*, 17 March 2019.

Ghemawat, P. (1991). *Commitment: The Dynamic of Strategy*. Free Press.

Government Accountability Office (2011). *DOD Cost Overruns: Trends in Nunn-McCurdy Breaches and Tools to Manage Weapon Systems Acquisition Costs*. GAO-11-499T, 29 March 2011.

Hastings, R. (2011). An explanation and some reflections. Netflix blog, 18 September 2011, and the Netflix announcement of 10 October 2011 ending the Qwikster plan.

Heath, C. and Heath, D. (2013). *Decisive: How to Make Better Choices in Life and Work*. Crown Business.

Holebrook, G. (2026). *Ishi: The Discipline of Determined Intent*. Leanpub.

Joint Authorities Technical Review (2019). *Boeing 737 MAX Flight Control System: Observations, Findings, and Recommendations*. Report submitted to the FAA, 11 October 2019.

Komite Nasional Keselamatan Transportasi (2019). *Aircraft Accident Investigation Report, PT. Lion Mentari Airlines, Boeing 737-8 (MAX), PK-LQP*. Final report.

Kristal, A. S. and Dorison, C. A. (2025). Precommitment can allow decision makers to maintain trust when de-escalating commitment. *Journal of Applied Psychology*, 110(10), 1395-1411. Published online December 2024.

Federal Aviation Administration (2016). *Advisory Circular 21.101-1B, Establishing the Certification Basis of Changed Aeronautical Products*, 11 March 2016.

Nygard, M. (2011). Documenting architecture decisions.

Simonson, I. and Staw, B. M. (1992). Deescalation strategies. *Journal of Applied Psychology*, 77(4), 419-426.

Staw, B. M. (1976). Knee-deep in the big muddy. *Organizational Behavior and Human Performance*, 16(1), 27-44.

US Department of Justice (2021). Boeing charged with 737 MAX fraud conspiracy and agrees to pay over $2.5 billion. Press release, 7 January 2021.

US House of Representatives, Committee on Transportation and Infrastructure (2020). *Final Committee Report: The Design, Development and Certification of the Boeing 737 MAX*.

Vaughan, D. (1996). *The Challenger Launch Decision: Risky Technology, Culture, and Deviance at NASA*. University of Chicago Press.

KETTEI

More Kettei case studies on Substack

The remaining Kettei case studies, the Kettei essays and the weekly Kettei Decision Digest are on Substack.

KETTEI ON SUBSTACK

kettei.substack.com

THE FLAGSHIP CASE: THE DECISION THAT COULD NOT EXPIRE: THE BOEING 737 MAX

kettei.substack.com/p/the-decision-that-could-not-expire

Since 8 October 2026, Kettei uses one Watchcard per steering decision and one Watchlist per endeavour; the essay is kept as published.