The Muda Method

From Lean to Muda

Twentieth-century management thought perfected the visible — efficiency, optimization, measurement. In exchange, it removed the invisible — judgment, relationships, the long-arc, the unmeasured. This book reads what survival requires across eight centuries of organizational history, against the rising tide of AI-driven efficiency.

Koto Igarashi

2026 · mudamethod.org

Prologue

When Efficiency Removed the Human

The opening question of The Muda Method

Koto Igarashi

First Draft · May 16, 2026

Over twelve million American small and mid-sized businesses now face what economists call the "Silver Tsunami" — the largest generational transfer of business ownership in modern history. Within ten years, roughly seventy percent of these companies will change hands. Yet, by most accounts, fewer than a third have any documented plan for what comes next.

This is not, properly speaking, a retirement problem. It is a continuity crisis. Behind the statistics lies a pattern that has destroyed every organization, in every era, that allowed itself to become predictable.

Tech giants lay off tens of thousands while their stock prices rise. Startups that raised hundreds of millions quietly close. Companies that seemed invincible — fortress brands, locked-in customers, cutting-edge technology — are quietly counting the years they have left. They are all walking the same path. The names change. The pattern does not.

1. The Five Stages

Every organization rises and falls along the same five stages.

Stage 1. A charismatic founder emerges. The vision is singular. The decisions are fast. Apple under Jobs. Amazon under Bezos. Walmart under Walton. The energy is undeniable, the rules few, the possibilities limitless.

Stage 2. Systems are built. The founder's intuition becomes process, documentation, culture. The organization scales beyond the founder. This is called maturity. It is also, in another sense, the beginning of the end.

Stage 3. The founder dies, retires, or is replaced. The source of original judgment is gone. What remains is the system that the founder built — perfectly tuned for the world that no longer exists.

Stage 4. Succession struggle and internal division. Factions form. Politics intensifies. The best people start to leave. The company spends more energy fighting itself than fighting its competitors. Most of the great corporate stories of organizational decay live here — Intel, IBM, GE, Sears, the post-Jobs Apple in some readings.

Stage 5. Talent exodus and collapse. The organization is acquired, broken up, or quietly shrinks into a husk of what it was. The customers lose faith. The talent has already left.

Most struggling organizations alive today, by my count, are in Stage 4.

This was not the failure of stupidity. These organizations had access to the best MBA graduates, the most advanced technologies, the deepest pools of capital in human history. They had every advantage. They still followed the script.

2. The Wrong Cure

Today, faced with this pattern, leaders turn to artificial intelligence. The reasoning is straightforward: if past judgment failed, perhaps better judgment — colder, faster, more comprehensive — can be machined. Strategy meetings now routinely begin with "let's see what the model says." Hiring is increasingly mediated by algorithm. Pricing, allocation, product roadmaps — all are slowly being transferred to systems that promise the optimal answer.

This approach, taken alone, will not break the pattern. It will accelerate it.

Artificial intelligence learns from the past. It analyzes vast historical archives and offers solutions that have, on average, worked before. Those solutions are not new futures. They are highly refined repackagings of the past.

If artificial intelligence had advised Blockbuster in 2000, it would have recommended optimizing the store-location algorithm and the late-fee schedule, because those were the variables the historical data spoke to. It would not have predicted streaming, because streaming barely existed in the training set.

If artificial intelligence had advised Nokia in 2005, it would have recommended improving hardware manufacturing and reinforcing carrier relationships, because those were the existing axes of competition. It would not have seen the iPhone.

If artificial intelligence had advised Borders in 2008, it would have recommended optimizing physical retail. It would not have predicted the e-book revolution.

The general principle is uncomfortable. **Artificial intelligence is brilliant at fighting yesterday's war.** The decision landscape has already moved.

3. The Other Cost

There is a second, deeper problem with optimization as a strategy.

Efficiency, taken to its limit, removes the very randomness that produces innovation.

Look at the current wave of corporate restructuring. Companies are cutting costs based on data-driven analysis. The model says, in effect: *these departments are unprofitable, these roles do not generate immediate return.* What the model cannot see is which of those "unproductive" people would have created the next breakthrough, which "inefficient" team would have built the next disruptive product, which conversation in which hallway would have started the next decade of growth.

Gmail came out of Google's twenty-percent time — unstructured, unjustified hours that no rational allocation system would approve. The Post-it Note came out of a failed glue experiment at 3M. Viagra came out of unexpected side effects in a heart medication trial. Every meaningful breakthrough I can find in the history of business carries this signature: it was not on the plan.

The next breakthrough never comes from optimization. It comes from deviation.

And deviation, by definition, looks wasteful in advance.

4. What Lean Removed

The twentieth century achieved what may be the most refined operational philosophy in human history. Taiichi Ohno at Toyota in the 1950s defined a category that James Womack and Daniel Jones later named, in their 1990 study, *Lean*. It went on to reshape global manufacturing, then software, then nearly every modern enterprise.

Lean has a vocabulary. At its center is a Japanese word: *Muda*. Waste. The seven Muda. The eight Muda. The endless varieties of waste that good systems should eliminate.

Lean was right, within its own frame. Inventories shrank. Defects vanished. Throughput multiplied. Capital efficiency improved by orders of magnitude.

And yet, in March of 2011, when the T■hoku earthquake and tsunami struck Japan, Toyota's Lean system — perfected over fifty years — collapsed in days. Single-source suppliers had been chosen for cost. Buffer stocks had been eliminated as waste. Slack had been engineered out of the system. When one node failed, the entire global production line stopped.

Then, in 2020, the same thing happened on a global scale. Lean supply chains everywhere — in masks, in chips, in food, in cars — collapsed under shock. The same word kept appearing in the after-action reports: *just-in-time became just-too-late.*

The lesson is structural, not historical. **Perfect efficiency, taken to its limit, produces perfect fragility.**

What Lean called Muda — what it eliminated — was not merely waste. Some of it was something else.

5. Reading Muda Differently

The book you are now reading proposes a different reading of the same word.

Yes, some Muda is waste. Defective inventory, redundant motion, excess transportation — Lean is right to eliminate them. No serious operator disagrees.

But there is a second category of Muda that Lean, by its very structure, cannot see. Buffer capacity. Redundant suppliers. Unstructured time. Conversations without an agenda. Relationships without immediate transactional value. The long-arc decisions that produce no measurable return for years. The judgments that resist quantification.

None of these can be priced. None of them survive a quarterly review. Most of them are invisible to artificial intelligence by construction. And yet, on inspection, every long-surviving organization — sacred sites that have stood for thirteen hundred years, family enterprises that have lasted eight

hundred, companies that have outlived every competitor across centuries — preserves precisely these.

What they preserve is what Lean removes.

This book gives that preserved thing a name. I call it **Muda** — using the same word Lean uses, but pointing to the other half of what the word can mean.

The Muda of Lean is what efficiency removes. The Muda I describe in this book is what survival requires.

The proposition of the book, stated as cleanly as I can:

> ***Muda is the human.*** *Judgment, relationship, the long arc, the unmeasured. What twentieth-century management called waste. What twenty-first-century AI cannot see. What every organization that has survived more than a few generations has insisted on keeping.*

6. From Lean to Muda

The arc of this book is suggested by its subtitle: *From Lean to Muda*.

This is not an attack on Lean. Lean built the twentieth century's productive base. I have no intention of unbuilding it. But the twentieth century ended. The next era will not be won by competing on the same axis that Lean optimized. That axis is now in the hands of artificial intelligence, and on that axis, no human enterprise can win against algorithmic optimization at the limit.

The competitive edge of the next era lies elsewhere. It lies in the capacity to keep, and to read, what efficiency systematically removes. It lies in the human.

This is the move I call *From Lean to Muda*.

7. A Note on the Origin

The methodology I lay out in this book is not, in its essential structure, new. Eight centuries ago in Japan, a Buddhist monk named Ano Zenjo (1153–1203) had already articulated something very close to it. Yoritomo and Masako, the founders of the first samurai government in Japan, were his disciples. Recent historical research, drawn from primary sources, has established that the visible governance of the Kamakura Shogunate stood, in considerable part, on the invisible network that Zenjo built. The visible depended on the invisible. The institutional form depended on what could not be institutionalized.

That work, and that man, were systematically erased from the official record. The textbook history we inherited tells a different story, in which Zenjo plays a marginal role. The case for restoring the

original picture is laid out, in primary-source detail, in the academic essays at tokiwatanabe.org. I will sketch it briefly in Chapter 5 of this book.

I want to be exact about my own position here. **I am not the heir to this method. I am its translator.** Ano Zenjo himself did not use the word "Muda." What he called *Ura* — the hidden, the invisible, the back of the cloth — I have translated, eight centuries later, into the language of twenty-first-century management as *Muda*. The naming is my responsibility. The thing itself has always been there.

8. The Promise of This Book

By the end of this book, you will have the following.

A framework for identifying which of the Five Stages your organization is currently in — and what the typical failure pattern of that stage looks like.

A vocabulary for the four dimensions of Muda — judgment, relationship, long-arc, unmeasured — and how to recognize each in your own organization.

A reasoned argument, from first principles, for what artificial intelligence is structurally unable to see, and where the human therefore retains an irreducible advantage.

A set of case studies — Ise Grand Shrine, Kong■ Gumi, Stora Enso, Patagonia, Costco, In-N-Out, the Nintendo Wii, Toyota's hybrid-and-hydrogen bet, the Kyoto restaurants that refused their Michelin stars — of organizations that survived, in some cases for over a millennium, by keeping what efficiency told them to remove.

And a set of practices, at the level of the individual career, the organization, and the larger institutional landscape, for keeping Muda alive in your own context.

The book is short for what it attempts. It does not need to be long. The argument is, in the end, simple. The twentieth century perfected the visible. The twenty-first century is delivering machines that will perfect the visible more completely than any human could. The competitive edge of the era ahead is on the other side of the same word.

From Lean to Muda.

Let me show you what I mean.

Chapter 1

Muda Is the Human

The four dimensions, and why a new word was needed

Koto Igarashi

First Draft · May 16, 2026

In the prologue I made a claim that needs unpacking. I said that what Lean removed was, in some cases, not waste but something else, and that this something else is what survival requires. I proposed to call it *Muda* — using the same word Lean uses, in inverted sense.

This chapter does the unpacking. It says what Muda actually is, gives it four dimensions, and explains why I needed a new word at all.

1. Why a New Word

There are already several words for what I am pointing to. Soft skills. Intuition. Tacit knowledge. Culture. Wisdom. Judgment. EQ. Each captures part of it. None captures the whole.

The problem with each of these existing terms is structural, not aesthetic.

Soft skills implies a complement to "hard" skills. The framing concedes the primacy of the hard. The phrase is used most often by people optimizing the hard, who add the soft as a courtesy. What I am pointing to is not a complement. It is the load-bearing element. Calling it "soft" is exactly the error I am writing against.

Intuition is the lay term for fast pattern recognition. It is real, and it matters, and it is part of what I mean. But intuition can be modeled. AI is now, in many narrow domains, better at "intuitive" pattern recognition than expert humans. Intuition is not the part of the human that cannot be replaced. It is the part that is being replaced first.

Tacit knowledge, in the Polanyi sense, is the knowledge we have but cannot articulate. This is closer. But tacit knowledge is still knowledge — something stored in the mind that, in principle, could be elicited, transcribed, encoded. It is not yet captured, but the framing implies that capture is the goal. What I am describing resists capture by its nature.

Culture is the closest of the common terms. The culture of an organization holds much of what I mean. But "culture" is too diffuse. It describes the atmosphere without naming the substance. And in management discourse, "culture" has been thoroughly assimilated by the optimizers — it has become another variable to be measured, designed, A/B tested.

Judgment is the term I am closest to, and in fact judgment is one of the four dimensions I will name in this chapter. But judgment alone is too narrow. It points to decision-making moments. What survival requires is not only the moments of decision but the entire substrate from which good decisions become possible — the relationships that surround the decision, the time horizon against which the decision is measured, the things that were never on the dashboard but that matter.

So I need a word that means: *the entire substrate of the human that resists capture by the optimization frame, and that survival turns out to depend on.*

That word is Muda.

I chose it deliberately. Lean's vocabulary is the operating language of twentieth-century management. To take Lean's central word — Muda, waste — and invert its meaning is not a coincidence of nomenclature. It is a claim about where the next era of competitive advantage lives.

The claim is this: **what twentieth-century management called Muda, properly read, contains the very capacity that twenty-first-century survival requires.**

Same word. Opposite reading. That is the move.

2. The Old Method's Vocabulary

The methodology I am translating from has its own vocabulary, eight centuries old. I will name it briefly here, because the translation matters.

In the esoteric Buddhist tradition that informs the original method, reality has two faces. The face that is shown to the world — visible, public, formal, available to anyone — is called *Omote (■)*. The face that is not shown — invisible, private, informal, available only through long initiation — is called *Ura (■)*.

The mark of mastery, in the tradition, was the recognition that these two are not separate. *Hyouri Ittai (■■■■)*: front and back as one. The visible face has no integrity if it is not anchored in the invisible. The invisible has no expression if not surfaced through the visible. They are aspects of the same reality. Neither is real without the other.

This is what Ano Zenjo (1153–1203) taught. The detail, and the historical evidence for it, is laid out in Chapter 5. For this chapter what matters is the structure of the claim.

The structure transposes cleanly onto modern management. The visible face — Omote — is what Lean optimizes. Operations, processes, throughput, metrics, KPIs, dashboards, OKRs, the entire apparatus of measurement and control. This is the part of the organization that everyone can see and that any competent manager can improve. The invisible face — Ura — is what Lean systematically removes as waste. Slack, redundancy, relationships, judgment, the long-arc, the unmeasured.

The translation is straightforward, with one move.

Ano Zenjo's Ura is what I have translated into the language of management as Muda.

The historical Zenjo never used this word. The translation is mine, eight hundred years late. I take responsibility for it. What he called Ura, I call Muda — because the strategic content of the concept is identical, but the word "Muda" carries the right charge for an argument with twentieth-century management thought. Lean owns the word. To borrow it, and invert it, is the most concentrated way to make the point.

3. The Four Dimensions

Muda, as I am using the word, has four dimensions. They are not exhaustive. They are the four that matter most in practice, the four where the work of preservation tends to fail, and the four where AI is structurally blind.

1. Judgment

The first dimension is judgment — but not all judgment, only the kind that resists formalization.

Most of what gets called "judgment" in business literature is, on inspection, applied pattern recognition. The senior executive recognizes that the situation in front of her resembles three other situations she has encountered, weighs the contextual differences, and selects the closest-fitting response. This is real. It is valuable. And it can be modeled, increasingly well, by systems that have been trained on enough examples. This kind of judgment will, over the coming decade, be largely automated. It is not the Muda of judgment.

The Muda of judgment is what remains after pattern recognition has done its work. It is the moment when the executive recognizes that this situation does *not* resemble the precedents in the way that the precedents themselves suggest — that the closest pattern is, in fact, the wrong pattern. It is the recognition that everyone agrees on the answer, and everyone is wrong. It is the call to bet against the model's recommendation, with no data to support the call, on the basis of something that cannot be specified.

This is not intuition. Intuition would have given the same answer as the model. This is the recognition that the entire frame of reference, including the intuition that sits inside it, has stopped tracking reality. It is the move outside the current frame.

Every great strategic decision I can find, when examined closely, contains this kind of judgment at its core. The decision to build the iPhone when the data said the world wanted a better BlackBerry. The decision to bet on streaming when the data said the world wanted better DVD-by-mail. The decision to refuse the Michelin star when the data said it would multiply revenue. The decision to write "Don't buy this jacket" in The New York Times when the data said it would crater sales.

In each case, someone made a call that the available frame could not justify. The call was not modeled. It was judged.

That kind of judgment is the first dimension of Muda.

2. Relationship

The second dimension is relationship — but, again, not all relationship, only the kind that resists transactional reduction.

Most business relationships are, by their nature, instrumental. The supplier relationship exists because the supplier provides inputs at acceptable price and quality. The customer relationship

exists because the customer pays for outputs. The employee relationship exists because the employee provides labor for compensation. Each of these can be modeled as a transaction with a measurable equilibrium. CRM systems are built on this premise, and they work as far as the premise holds.

The Muda of relationship is what survives after the transaction has ended.

The customer who buys nothing this year but who, when his neighbor asks for a recommendation, names you. The supplier who, when a competitor offers him better terms in a crisis, calls you first because of something that happened in a meeting fifteen years ago. The former employee who, two careers later, refers her best new hire to you because of how you treated her when she left. The contact who passes you the information you needed before the news broke, not because you paid for it, but because of a dinner in 2017.

None of this appears in the CRM. None of it can be priced. None of it can be planned. It is the residue of relationship that exists when transactions have stopped being the relevant frame.

And — this is the part the optimizers miss — these are the relationships that disproportionately determine survival in the long arc. Every organization I have studied that has lasted multiple generations runs, at its deepest level, on relationship Muda. The visible business sits on top of an invisible relational substrate, much of it inherited, that the books cannot capture and that the org chart cannot show.

That is the second dimension.

3. The Long Arc

The third dimension is the long arc — the time horizon against which an organization is actually optimizing.

Quarterly reporting cycles, annual budgets, three-year strategic plans, five-year outlooks. These are the time frames in which most organizations live. They produce a particular shape of decision — a shape weighted toward outcomes that resolve within the planning horizon.

The organizations that survive across generations live in a different time frame. The Ise Grand Shrine of Japan rebuilds itself every twenty years, in the same form, on the same site. It has been doing so for thirteen hundred years. The decision logic that drives this is invisible to any quarterly frame. It is invisible to any individual lifetime. It is a logic that operates on the scale of dynasties.

Kong■ Gumi, the temple-construction firm in Japan, made decisions on the scale of fourteen centuries. Patek Philippe positions itself as the steward of timepieces "for the next generation." Toyota's hybrid investment, when measured against the EV optimization wave, looks irrational on a five-year horizon. On a twenty-year horizon, in a world where electrical grid capacity may be the binding constraint and pure-EV may be technologically displaced by something we have not yet named, the hybrid bet starts to look like the surviving option.

This is not patience. Patience is waiting through a known time frame for a known outcome. This is a different kind of perception — the capacity to make decisions whose logic does not show up in any feasible measurement window. The optimizer cannot do it. The optimizer has no place to put the variable, because the variable's payoff happens after the optimizer is gone.

The long arc is the third dimension of Muda.

4. The Unmeasured

The fourth dimension is the unmeasured. Not the not-yet-measured. The structurally unmeasurable.

The optimizer's response to anything important is to find a metric for it. If culture matters, measure culture with an annual survey. If trust matters, measure trust with NPS scores. If alignment matters, measure alignment with OKR dashboards. If creativity matters, measure creativity with patent counts. The premise is that anything important enough to manage is, by definition, measurable.

The premise is wrong. Not as a matter of opinion. As a matter of structure.

There are kinds of organizational substance whose existence depends on not being measured. Trust between two people erodes when it is converted into a score the two people see. A team's willingness to take real risks vanishes the moment those risks become metrics on a performance review. The atmosphere of a place can survive scrutiny, but it cannot survive metrication. To measure it is to alter it.

Goodhart's Law states the principle from the negative side: when a measure becomes a target, it ceases to be a good measure. What I am pointing to is the positive side of the same observation. There is substance in an organization that is real, that matters, and whose continued existence requires that it not be made into a target.

This is a hard concept for the optimization mindset. The optimization mindset assumes that if a thing exists, it can be measured, and if it can be measured, it can be improved. The fourth dimension of Muda is the part of organizational reality where this assumption breaks. The thing exists. It matters. And the act of trying to measure it destroys it.

Every long-surviving organization protects, deliberately or by inheritance, a domain of the unmeasured. The corner where the senior people gather without an agenda. The conversations that happen during the rebuilding. The customs that are followed without anyone being able to explain why. These are not legacy artifacts to be modernized. They are the fourth dimension of Muda, and the organization is alive in part because they have not been measured.

4. Why "Muda" and Not Something Else

I want to return briefly to the question of the word, because the choice of word matters and someone reading this will want to know whether I considered the alternatives.

I did. The alternatives all fail in the same way. They concede the frame.

To call it "soft" is to accept the primacy of hard. To call it "tacit" is to accept that the goal is explicit. To call it "culture" is to accept the language of HR. To call it "wisdom" is to retreat into mysticism. To call it "human capital" is to surrender to the optimizer's frame entirely.

Naming a category is a strategic act. The category is the unit of competition in the marketplace of ideas. If I had named what I am pointing to with any of the available words, I would have inherited the existing argument, and the existing argument has already been won by the optimizers.

The move is to take their own word, and to give it back to them inverted. The move is to say: yes, what you have called Muda is exactly the right name, but you have read it backwards. What efficiency removes is not waste. It is the human. And the human is what survival requires.

This is also why the word stays. Critics will say I am playing with language. They are correct. The play with language is the argument. To make the same point in different words would be to make a different point.

5. What This Chapter Establishes

By the end of this chapter, I want the reader to have three things.

First, a working definition of Muda. Not waste. The human substrate of the organization that resists capture by the optimization frame and that survival depends on. Four dimensions: judgment, relationship, the long arc, the unmeasured.

Second, a sense of why the existing vocabulary will not do, and why a deliberate inversion of Lean's central term is the cleanest way to mark the territory.

Third, the location of the argument. Muda is not soft. Muda is not optional. Muda is not a complement to the hard, optimized, measured layer of the business. Muda is the layer of the business that is structurally invisible to that frame, and the business depends on it whether or not the frame acknowledges it.

From this point forward, when I use the word Muda in this book, this is what I mean. The four dimensions will come back, in different combinations, in every chapter that follows.

In Chapter 2 I will show why organizations forget this — why the five-stage pattern that has destroyed nearly every long-lived organization is, in essence, the pattern by which Muda is gradually replaced by Omote alone. In Chapter 4 I will show why AI, however powerful, is structurally unable to perceive the four dimensions, and what this means for the next decade of competitive advantage. In Chapter 5 I will show what Ano Zenjo's actual life looked like as a master of all four dimensions, eight centuries before any of us. And in Chapter 6 I will show the organizations alive today that have kept the four dimensions intact, and what we can learn from them.

Each of those chapters depends on this one. The argument of the book turns on the definition of a single word.

Muda is the human.

Chapter 2

The Four Failure Patterns

How organizations die — and what Muda protects against

Koto Igarashi

First Draft · May 16, 2026

In the prologue I sketched the five-stage arc by which organizations rise and fall. In Chapter 1 I defined what Muda is. In this chapter the two come together.

The five-stage arc is the shape of decline. It tells you, in coarse strokes, the trajectory. But trajectories do not explain mechanisms. Five-stage failure looks the same from a distance but feels very different from inside. The Roman Empire and your family-owned manufacturing business did not fall by the same hand. They fell by the same arc, but for distinguishable reasons.

Those reasons reduce, on inspection, to four. Most organizations that fail, fail through one of four patterns. Often two or three of them at once. Each pattern is, at its core, the same thing: **the failure to keep Muda alive in some specific dimension.** But the dimension differs, and so do the warning signs, and so do the practices that prevent it.

This chapter names the four patterns and shows what keeping Muda means against each.

1. Why Four, Not One

A common move in management literature is to identify the one thing that kills companies — failure of innovation, inability to execute, bad culture, leadership crisis. Each of these arguments is half right. They are half right because they describe what one pattern of failure feels like from inside.

What I will argue is that there are four distinguishable patterns. They are not phases of the same disease. They are different diseases. An organization can be perfectly innovative and die of succession failure. A company with brilliant succession can die of overextension. A company that grew sustainably can ossify in place and rot from the inside while the world is still rewarding it. And an organization that has handled all three can be obliterated by an external shift it had no language to recognize.

The four patterns are:

One. Succession Failure — the breakdown of transmission across generations.

Two. Overextension — growth past the carrying capacity of what makes the organization itself.

Three. Internal Ossification — the slow conversion of the system that enabled the founder into the cage that prevents the successor.

Four. External Disruption — the failure to perceive that the ground beneath the organization has moved.

Each of these has been studied independently. The contribution of the Muda frame is to see them together, and to see that all four turn on the same underlying mechanism, which is the loss of one or another dimension of Muda.

Let me take them one at a time.

2. Pattern One: Succession Failure

This is the most studied of the four, because it is the most common. Seventy percent of family businesses fail to transition successfully from the founding generation to the second. Ninety percent fail by the third. These are the headline statistics of the Silver Tsunami, but they are also the visible surface of a deeper structural problem.

The conventional analysis blames the successor. The successor was not as charismatic as the founder. The successor was not as committed. The successor was a less able operator. This analysis is sometimes true at the level of individual cases. As a general explanation it is wrong.

The general problem is not that the successor is worse. The general problem is that *what was transmitted to the successor was the wrong thing.*

What gets transmitted in a typical succession is the visible: the org chart, the financial statements, the customer list, the brand. These are the Omote of the organization. What does not get transmitted, because it cannot be put in a binder, is the substrate the founder was operating on — the judgments that have not been formalized, the relationships that exist outside any contract, the time horizon the founder was actually using, the things in the operating environment that the founder paid attention to without knowing why.

This substrate is the Muda of the organization. And succession in most cases transfers everything except the Muda.

The successor inherits the company and finds, often years later, that the company stops working in their hands. They blame themselves. They blame the team. They restructure. They bring in consultants. None of it touches the actual problem, which is that the operating substrate the founder was running on has not made the transition.

The organizations that handle succession well are the organizations that recognize this. Toyota's transmission of operating philosophy across generations is famously a multi-decade apprenticeship. The Japanese long-lived family businesses — the centuries-old hotels, the temple builders, the sake brewers — train successors not in the visible operation but in the invisible substrate. Senior people do not formally teach. They demonstrate, in front of the successor, for years, the kind of attention they pay. The successor learns what to look at. That is the actual transmission.

The first dimension of Muda — judgment — is what succession most reliably fails to transmit. The successor inherits the rule book. What the rule book cannot capture is the cases where the founder went around the rule book and was right to. The successor, lacking the same instinct, follows the rule book and is gradually wrong, in ways that accumulate.

The discipline against this failure has a name in long-lived organizations. It is called *protected apprenticeship*. The successor is exposed, for years, to the substrate of judgment that the formal documentation cannot capture. There is no curriculum. There is only the long exposure to the way the senior people look at things. This exposure is, in measurable terms, deeply inefficient. The

successor is doing nothing that shows up on a productivity dashboard. Most modern companies have eliminated this kind of apprenticeship as waste. The companies that retain it have, in most cases, the cleanest successions.

3. Pattern Two: Overextension

The second pattern is overextension. The organization grows beyond what the original Muda can sustain.

The trap of growth is, in some ways, opposite to succession failure. In succession failure, the visible is transmitted and the invisible is lost. In overextension, the visible is reproduced — more stores, more markets, more products — and the invisible is diluted past the point of recognition.

The standard pattern. A founder builds something that works because of a particular combination of relationships, taste, customer intimacy, and judgment. The thing is small. It is good. Investors notice. Investors want it bigger. Growth requires standardization. Standardization requires documenting what was previously implicit. Documentation cannot capture the parts that were doing the work invisibly. The standardized version reproduces the form but not the substance. The form, scaled, hollows out. By the time leadership recognizes that the original magic is gone, the organization is too large to return to what it was.

Every restaurant chain that started as one beloved location and ended as a competent but unremarkable franchise carries this signature. Every consumer brand that was acquired by private equity and stripped of its soul. Every craft producer that scaled its way out of being a craft producer. The damage is not failure to grow. The damage is the kind of growth that cannot carry what made the original itself.

The relevant dimension of Muda here is the second — relationship. Small organizations live on the density of relationships that small organizations can have. The founder knows every senior customer personally. The team eats lunch together. The supplier and the buyer have the same phone number for fifteen years. None of this is on the org chart. It is the relational Muda that makes the small thing distinctive.

When the organization scales, the relational density per node falls. The customer relationship becomes a CRM entry. The team relationship becomes an org chart cell. The supplier relationship becomes a contract. Each of these is more efficient than what it replaced. Each of them removes a measurable amount of friction. And each of them also removes the relational substrate the original organization was running on.

The organizations that scale without overextending have, in nearly every case I can find, made a deliberate decision to refuse some axis of growth. Patagonia refuses the kind of growth that requires sourcing decisions inconsistent with its environmental commitments. In-N-Out refuses the franchising model that would let it expand at the pace its market position suggests. Basecamp has refused, for two decades, to take outside capital that would force it to grow on someone else's timetable. The Kyoto restaurant that turned down a Michelin star was refusing the customer base

the star would have brought, in order to preserve the relationship density it already had.

These look like irrational choices on a spreadsheet. They are, in fact, the exact discipline that prevents overextension. They are decisions to keep relational Muda intact at the cost of growth that would dilute it.

4. Pattern Three: Internal Ossification

The third pattern is internal ossification. The system that enabled the founder becomes the cage that prevents the successor.

Every organization that scales builds systems. The systems are how the founder's reach is multiplied. Approval flows, planning cycles, performance reviews, budgets, OKRs, governance committees. Each system is built at a moment when something concrete was going wrong. Each system, at its founding, was a solution.

Time passes. The conditions that motivated each system change. The systems do not. The people who built them have moved on. The people who arrive later inherit the systems without inheriting the conditions that justified them. The systems are, by now, the only language the organization speaks. To operate without them is to be outside the legitimate order. To question them is to be a difficult employee.

What dies is the willingness to act outside the system. The willingness to start a project without going through the approval flow. The willingness to make a hire without the formal job description. The willingness to take a risk that does not fit into the planning cycle. The willingness to ignore the budget because the situation in front of you demands it.

These willingnesses are the third and fourth dimensions of Muda — the long arc, and the unmeasured. The long arc, because the willingness to act outside the system depends on optimizing for outcomes that lie outside the system's measurement window. The unmeasured, because what makes the system feel like a cage is precisely that everything important about your work is being scored on dimensions that do not capture what your work actually does.

This is the pattern that produces Stage 4. The system that scaled the founder has become the substrate the successor cannot work in. The best people leave first, because they have options. The willingness to act becomes the willingness to leave. What remains is the political layer, which is gifted at navigating the system without being able to do anything outside it.

Ossification is the most invisible of the four patterns, because the organization is, by every visible metric, still healthy. Revenue is fine. Margins are fine. The dashboard is green. What is dying is not on the dashboard. What is dying is the capacity to act when the dashboard is wrong.

The discipline against ossification is, in most long-lived organizations, the deliberate preservation of zones where the system does not apply. Twenty-percent time at Google in the first decade. Skunkworks projects at Lockheed. The R&D; lab that does not have to justify its quarterly output.

The internal venture capability that does not have to clear the standard governance. These are deliberate exemptions from the system, justified on grounds that are not measurable.

Organizations remove these zones first, in budget cycles, because they look like waste. Removing them is the most reliable signal of ossification I know.

5. Pattern Four: External Disruption

The fourth pattern is external disruption. The world moves. The organization does not.

This is the pattern that gets the most attention in business journalism, because the stories are vivid. Kodak. Blockbuster. Nokia. Borders. Sears. Each had every advantage. Each watched its market shift in slow motion. Each, by the time leadership recognized the shift, had locked itself into a structure that could not turn fast enough.

The conventional reading is that these organizations lacked vision. Kodak's leadership did not see the digital future. Blockbuster did not see streaming. Nokia did not see the smartphone. This reading is wrong. Kodak invented the digital camera. Blockbuster had multiple internal proposals for streaming. Nokia had a smartphone platform years before the iPhone. The vision was present. What was absent was the willingness to act on a vision that contradicted the current business.

This is the pattern that the existing failure literature comes closest to describing. Christensen's *innovator's dilemma* is one version. The general shape: every well-run organization optimizes for its current market position, and every optimization for the current position increases the cost of pivoting away from it. The better-run the organization, the more it is locked into the current configuration.

The Muda dimension here is the long arc. External disruption succeeds against organizations that have, structurally, no time horizon longer than the planning cycle. The planning cycle was designed to handle the current business. By the time the planning cycle is showing serious deterioration in the current business, the disruption has already moved through the early stages where it could have been responded to.

The long-arc capability is the ability to take seriously a future that has no presence in the current data. Toyota's hybrid bet, made when the entire industry was committed to either internal combustion or pure electric, is an example of long-arc capability used well. The bet is, on a five-year horizon, irrational. On a twenty-year horizon, accounting for grid constraints, mineral availability, and the as-yet-unnamed technologies that may displace pure EV, it starts to look like the only configuration that survives every plausible scenario.

Most large organizations have, by my observation, lost long-arc capability entirely. The CEO is on a five-year contract. The board reports quarterly. The capital allocation framework requires payback periods measurable in years. Everything that would be needed to recognize an external disruption — and to act on it before the existing business deteriorates — has been engineered out of the decision system in the name of accountability.

The discipline against external disruption, accordingly, is the deliberate preservation of decision-making capacity outside the planning cycle. This usually means a small group, often around the founder, that can commit resources to bets whose payoff is not visible in the current measurement window. Most organizations that have weathered disruption well have this group. Most organizations that have not, do not.

6. What the Four Have in Common

Let me make explicit what the four patterns share.

Succession failure transfers the Omote and loses the Muda of judgment.

Overextension reproduces the Omote and dilutes the Muda of relationship.

Internal ossification multiplies the Omote and crowds out the Muda of the long arc and the unmeasured.

External disruption preserves the Omote against a world that has moved, and the failure is, again, the absence of Muda of the long arc.

In each case the visible is reproduced, scaled, or preserved. In each case the invisible — the Muda — is what was actually doing the work, and what is what fails.

This is why I have insisted, in Chapter 1, on the four dimensions of Muda. Each pattern of failure attacks a different dimension. Each pattern requires a different practice of preservation. There is no single discipline that prevents all four. There are four disciplines, applied to four dimensions, that protect against four characteristic deaths.

The diagnostic question for any organization, then, is not "are we failing." Failure is too gross a measure. The diagnostic question is: *which dimension of Muda are we currently losing, and what is the practice that protects it?*

In a healthy organization, all four dimensions are being actively preserved. The signs are visible if you know what to look for. There is a protected apprenticeship for the substrate of judgment. There is a willingness to refuse growth that would dilute relational density. There is a deliberate zone, often small, where the system does not apply. There is a small group that can commit on horizons no measurement system will validate.

In a failing organization, one or more of these is gone. Often, by the time the failure is visible in financial metrics, all four are gone, and the organization is in late Stage 4 or Stage 5. The diagnostic value of the four patterns is precisely that they show up in the substrate before they show up in the dashboard.

This is the entire reason for naming them.

7. What This Chapter Establishes

By the end of this chapter, the reader should have three things.

First, a vocabulary for the four patterns by which organizations die: succession failure, overextension, internal ossification, external disruption.

Second, the recognition that each pattern corresponds to the loss of a specific dimension of Muda — and therefore that the discipline against failure is the preservation of the relevant dimension, not the application of some general remedy.

Third, a frame for diagnosis that operates at the substrate, not the symptom level. An organization in trouble has been losing Muda along some dimension for a while before any financial measure catches up. The four patterns are how to see it early.

In Chapter 3 I will argue that artificial intelligence, however powerful, is structurally unable to perceive the four dimensions of Muda — and what this means for organizations whose primary tool for understanding themselves is increasingly the algorithm. In Chapter 4 I will show what it looked like, eight centuries ago, to operate with all four dimensions intact, in the figure of Ano Zenjo. And in Chapter 5 I will show the organizations alive today that have preserved the dimensions long enough to outlive every competitor in their field.

Each of those chapters depends on this one. The four patterns are how organizations die. Keeping Muda alive, dimension by dimension, is how they live.

Chapter 3

What AI Cannot See

The three structural limits of artificial intelligence, and the four dimensions of Muda through its eyes

Koto Igarashi

First Draft · May 16, 2026

Chapter 1 defined Muda — the human substrate of the organization, in four dimensions. Chapter 2 showed how the loss of those dimensions produces the four patterns of organizational failure. This chapter takes the next step.

The next step is to look at the most powerful tool the twenty-first century has produced for understanding organizations, and to ask whether that tool can perceive what we have now defined as Muda.

The tool is artificial intelligence. The answer is no.

Not "not yet." The answer is structural. AI cannot perceive Muda, and the reasons it cannot are not engineering limitations that will close with the next model release. The reasons are intrinsic to what AI is. This chapter explains why.

1. What AI Has Already Won

Begin with what I am not saying.

I am not saying that AI is overhyped. I am not saying that we are in a bubble. I am not saying that the current capabilities are smaller than they appear. The current capabilities are immense, and they are growing fast, and they are going to reshape, in the coming decade, almost every function that has measurable inputs and measurable outputs.

Pattern recognition on structured data: AI is now better than the median expert in most domains. Forecasting based on historical data: AI is competitive with the best human analysts and substantially cheaper. Optimization within defined parameters: AI is, on most problems, the new best practice. Generating fluent text, plausible images, working code, passable strategies — all of this is now in routine commercial deployment.

The visible work of management — what Lean would have called the Omote of the organization — is, in 2026 and increasingly in the years that follow, going to be done by AI. This is not a prediction. It is a description of the current trajectory.

My argument is not against AI. It is about where AI stops, and what lies on the other side of where it stops.

2. The Three Structural Limits

AI, however powerful, has three structural limits. These are not contingent on the architecture, the training data, the compute budget, or the next research breakthrough. They are intrinsic to the nature of any learning system that operates on representation. I will name them, briefly, and then show how each maps onto the dimensions of Muda.

Limit One: AI Learns Only From the Past

The first limit is the most often stated and the most often misunderstood.

AI learns from data. Data is, by definition, a record of what has already happened. The model is, at any given moment, a compressed representation of patterns in the past. The model's output, when applied forward, is an extrapolation from those patterns.

This is true of every machine learning system. It is true of the most advanced models. It will remain true of the next generation, and the generation after that, because it is what learning from data *means*.

The implication is not that AI is bad at prediction. AI is excellent at prediction in domains where the future resembles the past. The implication is that AI is structurally weak in exactly the situations where the future stops resembling the past — that is, in the situations that actually matter for organizational survival.

I gave the standard examples in the prologue. If AI had advised Blockbuster, it would have recommended optimizing late fees. If AI had advised Nokia, it would have recommended better hardware. If AI had advised Borders, it would have recommended better store layouts. The recommendations would have been technically excellent. They would have been suicidal.

The general statement: **AI is brilliant at fighting yesterday's war.** Every strategic moment that matters is a moment when the war has changed and the war-fighting apparatus has not yet caught up. In those moments, the model's confident recommendation is the most dangerous artifact in the room.

Limit Two: AI Cannot Generate Meaning

The second limit is less commonly discussed and more difficult to articulate, but it is the heart of the matter.

AI can recognize patterns. AI can manipulate symbols. AI can produce text that, on inspection, reads as if it were generated by a system that understands what the text means. None of this is the same as actually generating meaning. The model is, in the technical sense, doing very sophisticated statistical inference over a representation. The representation is, by construction, derived from how humans have used language and symbols in the past.

The model has no access to the world that the symbols are about. It has only the symbols.

This sounds philosophical. It has very concrete consequences. The model can tell you, with great fluency, what every recent management thinker has said about strategy. It can summarize, compare, synthesize. What it cannot do is decide, for your specific organization, in your specific moment, with your specific people, what actually matters. That decision requires something the model does not have: contact with the situation that the symbols are pointing to.

The model's strategic recommendation is always, in effect, an interpolation between past human strategic recommendations. If the situation in front of you closely resembles situations that humans have already written about, the interpolation is useful. If the situation is genuinely novel

— if it requires categories that have not yet been named — the model will produce a confident-sounding interpolation, and that interpolation will be wrong in ways the model cannot detect.

The act of *naming what does not yet have a name* is the generative work of meaning. AI cannot do this. AI can only rearrange the names that already exist.

Limit Three: AI Cannot Execute Strategic Deviation

The third limit is the most consequential for the argument of this book.

AI, at its best, produces the optimum recommended action given the available data. By construction, this is the action that the data suggests has the highest probability of producing the best measurable outcome.

The problem is that, in many strategic moments, *the action with the highest expected measurable outcome is the worst action the organization can take.*

This is not a paradox. It follows directly from the structure of competition. If your competitors are also using AI to find the optimal action, and the data is the same, then the optimal action is, by definition, the action your competitors will also take. The optimum becomes the consensus. The consensus becomes the marketplace's commodity strategy. The commodity strategy generates no advantage, because everyone has it.

The only strategic moves that produce durable advantage are moves that deviate from the optimum that everyone else's optimization is also recommending. By construction, AI cannot recommend these moves. AI is optimizing for the same outcome on the same data as every other system. The moves AI cannot make are the moves the next era of competitive advantage will be made of.

This is the deepest of the three limits, because it is the one most managers will resist. The instinct of every well-trained operator is to trust the model when the model is confident. The model is confident exactly when the data is dense — which is exactly when the model's recommendation has no strategic value. The moves that matter are made in moments of low data density, where the model has nothing to say, and where the human is alone with the judgment.

3. The Four Dimensions Through AI's Eyes

The three limits produce, in combination, four specific blindnesses. Each blindness corresponds to one of the four dimensions of Muda I named in Chapter 1.

Judgment

AI is excellent at pattern-matched judgment. Given that this situation resembles those situations, the appropriate move is this one. This kind of judgment is what most of corporate management has been doing for decades, and it is what AI is in the process of automating.

The Muda of judgment is what remains when pattern-matched judgment is taken away. It is the recognition that this situation does not, in fact, resemble the precedents in the way the precedents themselves suggest. It is the recognition that everyone agrees on the answer, and everyone is wrong.

AI cannot do this. AI is the engine that recognizes the precedent. The judgment that the precedent does not apply is exactly the judgment AI is structurally unable to make, because the data on which AI is trained encodes the precedent as valid.

Stated more sharply: the most valuable judgments any organization makes are judgments that the available data does not support. AI is, by construction, the system that reaches the conclusion the available data does support. The valuable judgments are in the complement of AI's output set.

Relationship

AI can model relationships as nodes and edges in a graph. AI can analyze CRM data, communication patterns, organizational network maps. AI can identify which relationships are most active, most valuable, most at risk.

What AI cannot represent is the substrate that the active relationships are sitting on top of. The relationship that has not produced a transaction in five years but that, when needed, is the call that gets returned first. The supplier who, in a crisis, calls you ahead of his other customers because of something that happened in 2010. The former employee who passes you the best new hire because of how she was treated when she left. None of this appears in any data the model has access to. None of this can be encoded as a feature.

The optimizing system, asked to allocate relationship-maintenance budget, will allocate it to the active relationships, because those are the relationships the data tells it about. The substrate of dormant, latent, deep relationships will get nothing. Over time, the substrate erodes. The active relationships continue to be tracked. The organization, on every measurable dimension, looks well-connected. And then a moment arrives when what is needed is not on the active list, and the call is not returned, and the organization discovers that the substrate it was running on is no longer there.

This is not theoretical. This is the standard failure mode of organizations that adopt AI-driven relationship management without preserving the unmeasured part of the relational landscape.

The Long Arc

AI operates on a time horizon defined by its training data and its measurement window. For most commercial AI systems, the relevant horizon is the planning cycle — quarters, perhaps years. The model is fed signals on that scale and optimized against outcomes on that scale.

What AI cannot do is optimize against an outcome that lies outside its measurement window. The model does not know that the window is finite. The model treats outcomes inside the window as the only outcomes that exist. Decisions that produce returns on a twenty-year scale, or that protect

against tail risks that have not yet materialized, are not legible to the model. They look, on every metric the model can compute, like waste.

This is why AI-driven cost-cutting reliably destroys long-arc capability. The cuts target everything that does not produce measurable return in the relevant window. The research program with no immediate output. The supplier relationship that is not yet commercially valuable. The technology bet whose payoff is twenty years away. Each of these, in the AI's view, is waste. Each of these, in the longer view, is precisely the capacity that determines whether the organization is alive at the end of the next disruption.

The most well-known recent example is the choice between hybrid and pure-electric drivetrains in the global auto industry. On a five-year horizon, with available data, AI recommends pure EV. On a twenty-year horizon, accounting for grid constraints, mineral availability, charging infrastructure, and the technologies that may displace pure EV, the recommendation looks very different. Toyota's hybrid bet, ridiculed on the five-year horizon, is starting to look, on the twenty-year horizon, like the only configuration that survives every plausible scenario.

AI cannot make this kind of bet. The data does not extend that far.

The Unmeasured

This is the dimension where AI's limitation is most absolute, because it is true by definition.

AI operates on representations. Representations are, by construction, what can be encoded. The unmeasured is what resists encoding. The model has no access to the unmeasured. The model can be told that some variable is missing, and it can try to estimate the missing variable, but the estimate is always derived from the variables it does have. The model cannot perceive substance that, by its nature, does not reduce to variables.

This sounds like a tautology. Its consequences are not.

Every long-surviving organization protects, deliberately or by inheritance, a domain of the unmeasured. The space where decisions are made without producing a paper trail. The conversations that happen during the rebuilding. The customs that are followed without anyone being able to explain why. The substrate of trust between senior people that exists because it has never been quantified.

The optimizing system, brought into contact with these zones, will immediately propose measuring them. *If trust matters, measure trust. If alignment matters, measure alignment. If culture matters, measure culture.* The measurement is then introduced. The behavior optimizes against the measurement. The thing the measurement was supposed to track — which existed only because it had never been measured — disappears. The metric continues to show green for a while, by inertia. Then the metric, too, drifts.

This is the Goodhart effect, named after the economist Charles Goodhart: when a measure becomes a target, it ceases to be a good measure. What I am describing is the more general form:

there are kinds of organizational substance whose existence depends on not being measured, and the act of measuring them destroys them.

AI, by construction, cannot recognize the existence of such substance. It is invisible to the model in the same way that the air is invisible to a barometer that only measures air pressure. The barometer is not wrong about pressure. It is simply not seeing the air.

4. The Competitive Implication

The argument of this chapter has a strategic consequence, which I want to make explicit.

If AI is, increasingly, doing the visible work of management — pattern-matched judgment, transactional relationship management, in-window optimization, metricated culture management — then the visible work of management is, increasingly, a commodity. Every competitor has access to the same models, the same data, the same optimizations. The ground on which Lean made the twentieth century its own — efficient operations, optimized supply chains, refined processes — is, in the AI era, no longer ground on which durable competitive advantage can be built.

The advantage moves elsewhere. The advantage moves to the dimensions AI cannot see.

This is what I mean by "From Lean to Muda." It is not nostalgia for an earlier era. It is a positive claim about where the next era's competitive advantage lives. The visible has been commoditized by the new tools. The invisible has not, because the new tools cannot, by construction, reach it. The four dimensions of Muda are the dimensions where the human still has structural advantage over the machine.

This advantage is not permanent in absolute terms. Future systems may, in some domains, learn to approximate parts of what I have called Muda. But the structural limits are general, and they apply at the limit. As long as AI is optimizing on the same data as every other competitor's AI, the optimum is consensus, and the advantage is elsewhere. The advantage is on the side of the system AI cannot make.

The implication for organizations is straightforward. Every strategic investment that adds AI capability on the visible side is, on the strategic ledger, a defensive move. It keeps you at parity with competitors who are doing the same thing. Every strategic investment that protects, cultivates, and deploys the four dimensions of Muda is an offensive move. It is investment on the axis competitors cannot reach.

Most organizations, at present, are spending heavily on the defensive side and disinvesting on the offensive side. The disinvestment is invisible because the offensive side has no metric. The metric, here, is the cost. The organizations that recognize the asymmetry — that aggressively invest in what AI cannot see — will have, by the end of the decade, advantages their competitors cannot understand and cannot copy.

5. What This Chapter Establishes

By the end of this chapter, the reader should have three things.

First, a clear account of the three structural limits of AI: it learns only from the past, it cannot generate meaning, it cannot execute strategic deviation. These are not engineering bugs to be fixed. They are intrinsic to the nature of what AI is.

Second, a mapping of those limits onto the four dimensions of Muda. AI is structurally blind to the judgment that goes outside precedent, to the relational substrate that does not produce transactions, to the long arc that lies past the measurement window, and to the substance that exists only by not being measured.

Third, the strategic consequence: the visible work of management is being commoditized by AI, and competitive advantage is moving to the dimensions AI cannot see. This is the strategic content of "From Lean to Muda" — not a critique of Lean, but a recognition that the next axis of advantage is on the other side of the word Lean used.

In Chapter 4 I turn to the historical case that motivated my naming of this method. Eight centuries ago, a single individual operated with all four dimensions of Muda intact, at a level of integration that has rarely been matched since. His name was Ano Zenjo. The official record erased him. The story that remains is the story of what an organization looks like when the four dimensions are not protected by anyone — and what it might look like if we recovered the discipline of the man who did.

Chapter 4

The Erased Man

Zenjo (1153–1203) and the methodology eight centuries ago

Koto Igarashi

First Draft · May 16, 2026

Up to this point I have argued in abstractions. Muda is the human, with four dimensions; four patterns of organizational failure follow from losing those dimensions; AI is structurally blind to all four. The argument is general. It applies to any organization in any era.

This chapter departs from the abstract. It tells the story of one individual who operated, eight centuries ago, with all four dimensions of Muda intact, at a level of integration that has rarely been documented since. The story matters not because it is exotic, but because it shows what the abstract argument looks like as a life. And it shows, equally important, what happens when a methodology has been so threatening to the existing order that the record itself is rewritten to erase it.

The man was a Buddhist monk and strategist named Zenjo (1153–1203). He was the half-brother of Yoritomo, the founder of the first samurai government — the form of rule that would shape East Asia for the next seven centuries. The textbook history of the period gives Zenjo a marginal role: the eccentric brother, given the unflattering epithet "the wicked priest," who joined Yoritomo's uprising, performed religious rites, and was eventually executed by the Hojo family that came to dominate the regime after Yoritomo's death.

That textbook history is wrong. Or, more precisely, it is the version of history that the Hojo constructed after they killed him.

The reconstruction of who Zenjo actually was is, in 2026, still being completed. The primary source work is laid out in detail in the academic essays at tokiwatanabe.org. This chapter is not a substitute for that scholarship. This chapter is the management argument that the scholarship makes possible.

The argument: **Yoritomo's visible government rested on Zenjo's invisible substrate.** The Omote of the regime stood on an Ura that Zenjo had spent two decades constructing. When Zenjo was eliminated, what was eliminated was not merely a person. What was eliminated was the four-dimensional substrate the regime had been resting on. Within a generation, the Hojo had inherited a government that no longer had its operating substrate. The visible form continued. The Muda was gone.

1. The Official Record and Its Erasures

The dominant source for the founding of the regime is the record held to be official, compiled a century later by the family that, by then, controlled it. Every textbook account of the period descends, ultimately, from this text.

That record is not a neutral document. It was compiled, in part, to legitimize the Hojo's inheritance of power. It serves the purposes of the family in power at the time of its writing. It selects what to record and what to omit accordingly.

What the record says about Zenjo is, on the surface, sufficient. It records his exile, his joining of Yoritomo, his marriage to a woman of the Hojo (sister of Masako), his various religious activities,

and his execution on charges of treason.

What the record systematically omits is his institutional position. It refers to him as the "wicked priest" and notes his religious activities, but never identifies what office he held. The reader of the record is left with the impression of a relative who performed ceremonies and was eventually executed for plotting against those in power.

This impression is wrong. Recent scholarship — drawing on primary source materials that the record either did not have access to or did not include — has established that Zenjo held one of the most senior religious-political positions of his era. He was the head priest of the central religious complex of the eastern region. This was the religious authority on whose territory Yoritomo's exile had been spent and on whose blessing Yoritomo's uprising had depended.

The record does not mention this. The omission is not a casual oversight. The record records, in detail, who held lesser religious positions in the same period. It is the topmost layer of the religious hierarchy — the layer Zenjo occupied — that is systematically blank.

2. What the Primary Sources Show

The reconstruction begins with two surviving documents.

The first is a temple deed, dated to a period before Yoritomo's regime was secure. It records a land donation to the central religious complex of the region. The donation is made for the explicit purpose of "the safety and prosperity of Yoritomo's house" — that is, of Yoritomo's regime, then in its formative stage. The document is signed in the name of the head priest, with a personal seal. The seal has been definitively identified, through paleographic analysis, as belonging to Zenjo. He was, at that point, in office as head priest.

The significance of this is difficult to overstate. The head priest of this complex, in the political theology of the period, was the religious figure whose authority legitimized Yoritomo's regime's claim to rule. The deed shows that this figure was Yoritomo's own half-brother — and that he was actively directing the temple's resources toward the regime's stability.

The second document is an esoteric initiation manual from the religious lineage that operated at the complex. The manual records the vows of two specific initiates: Yoritomo and Masako, his wife. The initiation it describes is the highest-level esoteric Buddhist rite, in which the master pours consecrated water on the disciple's head and transmits, in principle, the entire tradition. This is not a routine ceremony. It is the rite of master-to-disciple succession in esoteric Buddhism.

The manual records that Yoritomo and Masako received this initiation. By implication, it identifies the master who gave it. The master could only have been the head priest of the complex. By the deed above, this was Zenjo.

Cross-referenced, the two documents establish a fact that the official record never states and the textbook history does not acknowledge: **Zenjo was the master of esoteric initiation for both the**

founder of the regime and his wife.

This is not the relationship of a brother and an in-law. This is the deepest pedagogical and spiritual relationship the period's religious culture knew. It is the relationship of teacher to disciple, of authority to authorized.

From this fact, much else falls into place.

3. The Twenty Years Before

Twenty years before Yoritomo's uprising, what was Zenjo doing?

The official record begins the story on the day of the uprising. But a different record shows that, for the two decades before that day, Zenjo had been engaged in the work of binding three separate worlds onto a single axis through his own body.

The first world was the local military power — the warrior families who controlled land and arms in the eastern region. The Hojo were among them. The second was the religious center — the highest religious institution of the region, where Zenjo had risen to occupy the topmost position. The third was the upper aristocracy of the distant capital — and specifically, the surviving family of the faction that stood in opposition to the ruling power.

Normally, these three worlds do not intersect. Geographically, socially, and politically, each belongs to a separate sphere. For a regional monk to receive the son of a high aristocrat of the capital as the husband of his daughter was, in the social structure of the period, the kind of thing that does not happen.

Zenjo made it happen. Over twenty years, deliberately, through the body of one specific person — his daughter.

Why?

There was a common enemy: the ruling power of the day. To overthrow that power, military force alone was not enough. Religious authority was needed to legitimize the use of force. And a political channel into the capital was needed to move the court against the enemy from within. Without all three, nothing would work. Military force alone would produce only rebellion. Religious authority alone would move no swords. A channel into the capital alone would leave no one to fight on the ground.

The work of connecting the three onto a single axis had to be done by someone. But that someone would need to belong simultaneously to all three worlds. Such a person does not normally exist.

Zenjo spent twenty years becoming that nonexistent person.

The uprising appears, in the visible record, to have happened on a single day. In fact, what happened on that day was the ignition of a structure that had been twenty years in the preparation.

History records, as the protagonist of the uprising, the figure who stood at the point of ignition. The figure who had assembled the structure that made ignition possible is the one removed from the record.

What matters here, for the argument of this book, is the nature of the work Zenjo had been doing.

The work of connecting three worlds cannot be completed on paper. To establish a marriage bond, a husband and a wife and a child must come into being through the actual bodies and lives of specific people. What happens here is not the duplication of information. One Muda — one human — and another mix together, and from them a new life is born. This is the ultimate definition of Muda. New life is born out of Muda.

And what that life will judge, with whom it will mix, what it will carry forward — none of it exists yet. Before what does not yet exist, a machine built to predict can only fall silent.

This cannot be compressed into a decade. It cannot be parallelized across three apprentices. The position at the religious center is the same: it rests on years of practice, accumulated reputation, and personal trust that cannot be replicated by a newcomer in three years. The channel to the distant capital, too, is maintained only through one-to-one connections — through letters, through messengers, through occasional return visits — each renewed across years. Information arrived at Zenjo because the relationships that carried it had been maintained, person by person, for a long time.

What these activities share is the quality of being **inseparable from time and from specific human bodies.** A marriage, a religious office, a network — none of these can be achieved by separating them from the specific person who carries them and the specific time they require. They cannot be transcribed to paper. They cannot be delegated. They cannot be bought.

What Zenjo built over twenty years was a structure of exactly this kind. The structure connecting the three worlds onto a single axis was inseparable from Zenjo as a specific person and from the specific time he had spent.

4. The Four Dimensions in Practice

Let me show now what Zenjo looked like, in each of the four dimensions of Muda I defined in Chapter 1. This is the chapter's central work. Read it not as biography but as the closest historical instance I know of what an integrated four-dimensional practice can be.

The Judgment Dimension

The esoteric tradition Zenjo had mastered was, by its nature, a discipline of judgment of the kind I described in Chapter 1. Esoteric Buddhist training is not a doctrinal education. It is a years-long apprenticeship in the perception of situations that resist conceptual reduction. The initiated person learns to see what does not appear in the text. The master is the person who has, over decades of practice, developed the capacity to make calls that cannot be derived from the doctrinal sources

but that have been validated, in concrete cases, by the lineage's accumulated experience.

What this looks like, in practical terms, is the capacity to recognize when the situation in front of you is different from its surface presentation, and to act on that recognition. The textbook description of Zenjo's life is full of moments that, taken seriously, only make sense if you assume he was operating on this kind of judgment. The decision to join Yoritomo's uprising — at a moment when Yoritomo had nothing, when the previous generation of Yoritomo's house had been comprehensively destroyed — looks irrational on the visible information available. It is the kind of move that follows from a judgment about the situation that the visible information does not support.

The same applies to the network construction in the two decades that followed, and the configuration in which Zenjo was operating at the end of his life. These were, repeatedly, calls that did not follow from precedent. They followed from the kind of judgment esoteric mastery is meant to produce.

The Relationship Dimension

The relationship layer is where the case for Zenjo as a Muda master is most direct, because the evidence is now established.

Yoritomo and Masako were not, as the textbooks frame it, Zenjo's brother and sister-in-law. They were his initiated disciples. The master-disciple relationship in esoteric Buddhism is the most binding pedagogical relationship the tradition recognizes. The master is, in the technical religious sense, the spiritual authority over the disciple for life. The disciple owes the master, on every relevant question, deference of a kind that has no exact analogue in modern professional life.

Yoritomo founded a government. The government legislated, fought wars, allocated land. The administrative form was Yoritomo's. The religious and pedagogical authority, throughout the entire period from the uprising to Zenjo's death, sat behind him, in the person of his initiator.

This is the substrate that the visible institutional form of the regime rested on. It does not appear in any organizational chart. It is structurally invisible to any reading of the period that takes the official record at face value. But it is the precise relationship — the one not captured by transaction, not measurable, not subject to political renegotiation — that made the Yoritomo regime stable through the brutal civil war that founded it.

The relationship layer, in Zenjo's case, was deeper still. Below the master-disciple bond with Yoritomo and Masako, Zenjo was also the foster father of Yoritomo and Masako's third son — the eventual third shogun. He was, in this capacity, the most senior figure in that child's upbringing. His wife was the child's wet nurse. The intergenerational relationship was the most intimate the period could construct.

The Long Arc Dimension

The third dimension — the long arc — shows most clearly in the period before the uprising, the years when Yoritomo was in exile and Zenjo was building his religious authority and his network.

This is the period the official record treats as a quiet preface to the main story. In the textbook version, Zenjo is doing religious things in obscurity while the political situation develops elsewhere. In fact, the surviving records indicate that this period — roughly two decades — was when Zenjo built the network on which the uprising would depend.

None of this network was constructed in the year of the uprising. The uprising drew on it. The two decades of relationship-building were the long-arc investment that made the visible mobilization possible. The long arc was, in the technical sense, invisible. Nobody outside the network would have understood, midway through it, what Zenjo was doing or why. The payoff was twenty years out.

This is the long arc as a strategic practice. Most contemporary corporate strategy cannot operate on this horizon, because the institutional structure does not permit it. Zenjo's institutional structure — the religious lineage — was designed precisely to enable long-arc operation, because the time-frame of the lineage was generations rather than quarters.

The Unmeasured Dimension

The fourth dimension shows in what Zenjo's actual base of operations was — and in what the official record never names.

Zenjo's residential and operational base was a territory along the main eastern road, at a strategic point. This was not a religious institution. It was a private estate, granted by the Yoritomo regime, from which Zenjo also held a samurai-rank lay title with concrete administrative authority.

The estate sat at a specific point in the geography of the regime. It controlled a strategic mountain pass — the chokepoint between the new government and the western capital. It included pasture lands where the regime's military horses were bred. The military infrastructure of the regime's mobility — its horses — was located on land directly administered by Zenjo, with the substrate religious and personal authority of the figure who had initiated the regime's founders.

None of this appears in the official record as such. The horses, the pasture, the territorial control, the strategic chokepoint — these are exactly the kind of operational substrate that the record treats as beneath its notice. It is the unmeasured substrate of the visible administrative form. And the visible administrative form was, in concrete operational terms, depending on it.

Removing this from the record was not a casual omission. It was the elimination of the operational base on which the visible form had been resting. The Hojo, in compiling the record a century later, knew exactly what was being left out. The textbook history, descending from this redacted source, has spent the intervening eight centuries reporting the visible form and not knowing what supported it.

5. The Execution

When the Hojo finally moved against Zenjo, the head of the Hojo line — Masako's father, the founder of the regency that would inherit the regime — accused him of conspiring with the second shogun to depose the chosen successor. The charge was treason. Zenjo was executed.

On the official account, the execution was the elimination of a treasonous brother-in-law. On the account that follows from the full primary source record, the execution was something else. It was the elimination of the figure who had, for two decades, held the substrate of the regime. With Zenjo removed, the Hojo inherited a government whose visible form they could control but whose invisible substrate they could not reconstruct.

The substrate could not be reconstructed because it had never been documented. It was, by its nature, the part of the regime that was carried by the relationships and the long-arc network and the religious authority and the territorial control that Zenjo personally embodied. Eliminate the person, and you eliminated, in the same act, the substrate.

The Hojo regency that followed inherited the Omote of the regime. It administered the visible institutions for over a century. Throughout that century, the visible form became increasingly elaborate — more legal codes, more administrative offices, more formal procedures — while the invisible substrate was no longer being renewed. The form continued. The substance was thinning. The Hojo regency eventually fell, after a slow process of internal hollowing that we have seen, in this book, by another name.

I have called it Pattern Three: internal ossification. The system that had scaled the original substrate had become the cage that prevented its restoration. The original substrate had been gone for over a century. By the time the visible form fell, what was left to fall was, in the deepest sense, already empty.

6. The Structural Isomorphism

I want to make the connection explicit, because it is the reason this chapter is in this book.

The Hojo, in eliminating Zenjo and erasing him from the official record, performed an act that is identical in structure to what the optimization frame does to organizational Muda. They preserved the visible. They eliminated the invisible. They reorganized the record to make the invisible never to have existed. Within a generation, the visible form was running on inertia. Within a century, the form collapsed under stresses that the lost substrate would have absorbed.

This is the pattern, generalized. **The act of erasing Zenjo from history and the act of optimizing Muda out of an organization are, structurally, the same act.** Both eliminate the substrate that was carrying the visible form. Both leave the form looking healthy for a generation. Both produce, eventually, a collapse that the record cannot explain because the explanatory information has been removed from the record.

This is why Zenjo is in this book. He is not in this book because he is my ancestor, though he is. He is not in this book because his story is dramatic, though it is. He is in this book because his erasure is the cleanest historical case I know of what happens, at organizational scale, when Muda is eliminated.

The full reconstruction of who Zenjo was — the seal analysis, the manual's dating, the estate's strategic geography, the network reconstruction — is the work of the academic essays at tokiwatanabe.org. I will not duplicate that work here. The argument of this book is that what was eliminated, when Zenjo was eliminated, is exactly the substrate that any organization in any era requires in order to survive its own scaling.

7. What This Chapter Establishes

By the end of this chapter, the reader should have three things.

First, a concrete historical instance of what an integrated four-dimensional practice of Muda can look like. Zenjo, across his lifetime, operated with all four dimensions simultaneously: judgment trained by esoteric mastery, relationship anchored in the deepest pedagogical bond the period knew, long-arc network construction across two decades, and unmeasured operational substrate that the official record never named.

Second, a clear case of what happens when this substrate is eliminated. The Hojo regency inherited the form. The form continued. The substance was gone. The collapse came, in slow motion, over the following century.

Third, the structural isomorphism between the historical case and the contemporary one. The act of erasing Zenjo and the act of optimizing Muda out of an organization are, at the level of structure, the same act. Both produce, over time, the same outcome: a visible form running on inertia until the stresses that the lost substrate would have absorbed are large enough to collapse it.

The next chapter takes the argument forward into the present. I will look at the organizations alive today that have, deliberately or by inheritance, refused this erasure. The Ise Grand Shrine. Kong■ Gumi. Stora Enso. Patagonia. The Kyoto restaurants that turned down their Michelin stars. Each, in its own register, is doing now what Zenjo was doing eight centuries ago. The question for the reader is whether your own organization is among them, or among the ones running on inherited inertia from a substrate that no longer exists.

Chapter 5

The Organizations That Kept the Muda

What thirteen-century shrines, eight-century corporations, and modern refuseniks have in common

Koto Igarashi

First Draft · May 16, 2026

Chapter 4 was a single, dense case — Ano Zenjo, eight hundred years ago, operating across all four dimensions of Muda until the regime that had depended on him moved to erase him. This chapter looks at the cases that did not get erased. The organizations alive today that have, by inheritance or by deliberate choice, kept the Muda intact long enough to outlast nearly every competitor in their domain.

The cases come in three rough categories. Some are very old and have preserved the four dimensions almost as a side effect of what the institution is for. Some are modern and have preserved the dimensions by explicit refusal — refusal to grow, refusal to optimize, refusal to take outside capital. And some have preserved the dimensions by making strategic bets that the consensus optimization frame could not justify, and that, on the longer horizon, turned out to be the configurations that survived.

None of these organizations would describe what they are doing in the language of this book. The language is my translation. What they are doing, in their own terms, is keeping faith with a way of operating that the dominant management discourse has trouble seeing. Each one offers a partial picture of what keeping Muda alive looks like in practice. Read together, they offer a composite — the modern equivalent, in distributed form, of what Zenjo did as a single person eight hundred years ago.

1. The Inheritance Survivors

Ise Grand Shrine: 1,300 Years of Deliberate Rebuilding

The Ise Grand Shrine in Mie Prefecture is, by most measures, the most sacred site in Shinto. Its main pavilions house the imperial regalia and are dedicated to Amaterasu, the sun goddess from whom the imperial line is traced. The site has been continuously active since the eighth century at the latest, and many of its institutional structures predate that.

Every twenty years, Ise rebuilds itself. The entire main shrine complex is dismantled and reconstructed on an adjacent plot, in the same form, using the same construction techniques, by craftsmen who have been trained for this work since childhood. The ceremony is called Shikinen Seng■ (■■■■). It has been performed, on schedule, for over thirteen hundred years.

From the perspective of any optimization frame, the Shikinen Seng■ is incomprehensible. The buildings are not falling down. The materials are not exhausted. The location does not need to change. The labor and material costs are enormous. The economic case is, on any quarterly or even decadal horizon, indefensible.

The institution does it anyway. It does it because the rebuilding is the practice through which the unmeasured substrate of the institution is renewed. The craft techniques are kept alive only because they are practiced every twenty years on the most important project the tradition knows. The intergenerational transfer of knowledge — master to apprentice, over the working life of three generations between each Seng■ — is the protected apprenticeship of Chapter 2 made literal. The relationships between the shrine and its supporting community of craftsmen, suppliers, and donors

are renewed in the same cycle. The long arc is built into the calendar.

If Ise had been optimized at any point in the last thirteen hundred years — if some efficiency-minded administration had decided that the rebuilding was waste and that the budget could be better spent on operational improvements — the institution would have, in a single century, lost the substrate it has been carrying for thirteen. The buildings would still be there, more or less. The thing that makes Ise *Ise* would be gone, irrecoverably, because the carriers of the practice would have moved on to other work.

The Seng■ is, in the language of this book, Muda made institutional. It is the deliberate, costly, recurring practice that the institution has refused to let efficiency remove.

Kong■ Gumi: Fourteen Centuries of Specialist Practice

Kong■ Gumi is a construction firm in Osaka, founded in 578 CE. For nearly fifteen centuries, the firm built and maintained Buddhist temples in Japan, almost exclusively. The company was independent until 2006, when it was absorbed by a larger construction conglomerate as a wholly-owned subsidiary. The lineage continues under the new structure, but the long pure-form independence ended after one thousand four hundred and twenty-eight years.

What is striking about the firm's long history is the discipline of specialization. Kong■ Gumi did one thing — Buddhist temple construction — and refused, for fifteen centuries, to do anything else. Other construction companies entered and left the temple market based on margin. Kong■ Gumi remained in it because the lineage of master carpenters and their accumulated craft was, in their understanding, exactly what the firm was. To diversify into other construction would have diluted the substrate of the craft. To industrialize the production would have eliminated it.

The firm's eventual loss of independence in 2006 is itself instructive. The cause was not the loss of the craft. The cause was that the institutional environment around the firm — the temple economy, the rate of new temple construction, the willingness of clients to pay for craft-level work in a commoditizing market — had changed faster than the firm could adapt. The Muda was still intact inside the firm. The market that could afford that Muda had thinned.

This is a more painful version of the same story. The internal Muda was preserved. The external conditions that allowed the internal Muda to survive economically had shifted. The firm continues, in a different form, with the substrate partly preserved. But the lesson is that even an institution that does everything right internally can have the rug pulled by an environmental change. Muda is necessary. Muda is not sufficient.

Stora Enso: The Eight-Century Mining Company

The Finnish-Swedish forestry company Stora Enso traces its corporate lineage to Stora Kopparberg, the world's oldest documented stock corporation, formally chartered in 1288 in Falun, Sweden. The original business was copper mining. The mine ran for over six centuries before being exhausted. The company, however, did not close. It transformed. From copper to iron. From iron to forestry. From forestry to pulp and paper. From paper, in the current century, to

renewable packaging and bio-materials.

The continuous corporate identity over more than seven centuries is, in modern terms, almost incomprehensible. Most modern public companies cycle through their entire identity in twenty years.

What Stora Enso has done, repeatedly, is execute the long-arc Muda from Chapter 1 — the capacity to make decisions whose payoff is not visible within any feasible measurement window. The 1888 decision to enter forestry, when copper was still profitable, looked premature. The mid-twentieth century pivot from raw materials to paper made the company's modern form. The current pivot from paper to renewable packaging is being made against a paper market that is, in some segments, still profitable.

Each pivot was, at the moment of decision, a bet against the optimization that the existing business would have justified. Each pivot, in retrospect, was the move that kept the corporate identity alive across centuries of changing technology. The firm has, repeatedly, refused to be the firm it was. The refusal is the practice.

2. The Refusal Survivors

The next category consists of modern organizations that have preserved Muda not through institutional inheritance but through explicit choice. These are the companies whose strategic decisions, when read in the language of conventional business strategy, look irrational. They are not irrational. They are protecting specific dimensions of Muda against the optimization pressure that would, if accepted, eliminate them.

Patagonia: "Don't Buy This Jacket"

In November 2011, Patagonia took out a full-page advertisement in the New York Times on Black Friday weekend. The advertisement showed one of the company's own fleece jackets and the headline "Don't Buy This Jacket." The body text walked the reader through the environmental cost of producing the garment and argued, in plain terms, that the reader probably did not need it.

On a conventional advertising frame, this is suicide. Telling customers not to buy your product on the heaviest shopping weekend of the year is the precise opposite of what the optimization logic of retail advertising recommends.

The advertisement worked. Patagonia's sales rose. The customer base that responded to the advertisement was, in a measurable way, the customer base the company had been cultivating: customers who understood that the company was not optimizing for short-term sales and that the products were being made by people who shared their values. The advertisement was a costly signal in the technical sense — an action that only a company actually committed to those values could afford to take.

The Muda dimension here is the long arc and the unmeasured. Patagonia has been refusing growth opportunities that would dilute its environmental commitments for decades. The refusals are not visible on any quarterly report. They are visible in the integrity of the brand and the loyalty of the customer base. By the time competitors realize what Patagonia has built, the substrate that produced it cannot be reproduced.

Costco: The $1.50 Hot Dog

Since 1985, Costco has sold a hot dog and soda combination, in its warehouse food courts, for $1.50. The price has not changed in forty years. Inflation alone would put the current cost-equivalent price at over four dollars. Internal financial pressure to raise the price has been continuous, and continuously refused.

The CEO, when asked about the practice, is reported to have told the team that if anyone tried to raise the hot dog price, he would kill them. This is, in the strict optimization frame, irrational management. The hot dog is not even a major revenue line. The economic case for the discipline is hard to make on any quarterly horizon.

What the hot dog does is hold a precise position in the customer's mind. The $1.50 hot dog is the visible signal of a company that does not optimize against its customers. Every other practice the company is known for — the limited margins, the bulk packaging, the deliberate inconvenience of the store layout — coheres around the same signal. The customer trusts Costco in a way that customers do not trust comparable retailers, and the trust translates, year after year, into membership renewals at a rate the industry cannot match.

The $1.50 hot dog is unmeasured Muda. It cannot be optimized into a profit center. Its function is to anchor a relationship to the customer that the optimization frame cannot see and therefore cannot match.

In-N-Out: Sixty Years of Refusing to Franchise

In-N-Out Burger has operated, since 1948, as a chain of physically integrated company-owned stores in California, with limited expansion into adjacent states. The company has refused, for over seventy years, to franchise. Its founders and successors have, repeatedly, declined acquisition offers that would have realized fortunes for the family.

The conventional case for franchising is overwhelming. Franchising would allow In-N-Out to expand nationally without putting the company's own capital at risk. Franchising would generate predictable royalty streams. Franchising would allow the brand to grow at the pace its market position justifies.

The refusal is principled. Franchised operations cannot maintain the operational standards In-N-Out's company-owned stores have. The relationships with suppliers, the freshness of the meat, the consistency of the kitchen practice, the wage and benefit structure for store-level employees — all of these would be diluted, by mechanism, in a franchised model. The family has chosen, repeatedly, to keep the substrate intact and refuse the growth that would compromise it.

The Muda dimension is relational. In-N-Out's relationship to its suppliers, its employees, and its customers has the density of a single family-run operation. The density does not survive franchising. The family has chosen the density over the scale.

Basecamp: Twenty Years Without Outside Capital

Basecamp, the project-management software firm founded by Jason Fried and David Heinemeier Hansson, has operated since 2004 without external investment. The company has been consistently profitable. Its co-founders have written, at length, about the strategic decision to refuse venture capital, and about the operational practices — the four-day summer week, the limited product scope, the deliberate refusal to add features that the customer base does not ask for — that follow from that decision.

The refusal of outside capital is a specific kind of Muda preservation. Outside capital, in the venture model, brings growth expectations. Growth expectations force standardization. Standardization erodes the relational density that small organizations have. The cycle is the overextension pattern of Chapter 2 in pure form. Basecamp has refused the first step of the cycle.

What this protects is the company's capacity to operate on its own time horizon. Decisions can be made on year-long or decade-long scales because no investor is demanding quarterly returns. The four-day summer week is not a productivity hack. It is a deliberate exercise of the slack — the structural Muda — that an organization not optimizing against external pressure can afford to keep.

3. The Strategic Survivors

The third category consists of organizations that made strategic bets the consensus optimization frame did not support, and that, on the longer horizon, turned out to be the configurations that survived.

Toyota's Hybrid Bet

When Toyota launched the Prius in 1997, the industry consensus was that the future of automotive propulsion was, in the long run, pure electric. The hybrid was a transitional technology — a compromise that, on the consensus view, would be obsoleted within a decade.

Toyota committed to hybrid technology anyway, and continued committing through the entire period when its competitors were placing larger and larger bets on pure EV. As recently as 2022, Toyota's commitment to hybrid and to hydrogen fuel cell development was being criticized in financial press as a strategic mistake.

The criticism began to soften in 2024 and 2025, as grid constraints, mineral supply problems, charging infrastructure gaps, and the technological maturation of hybrid systems became more visible. By 2026, the hybrid configuration is no longer a transitional compromise. It is, increasingly, one of the configurations that may dominate the next twenty years of personal

mobility — alongside, possibly, hydrogen fuel cell for heavy-duty applications and pure EV for short-range urban use.

The strategic content of Toyota's bet was long-arc Muda. The company made a decision against a consensus that was extrapolating linearly from the visible trajectory. The decision held under decades of pressure that the consensus model could not have predicted. It is, at this point, the closest thing the auto industry has to a survival hedge against the as-yet-unnamed disruptions of the coming decades.

Nintendo Wii: Stepping Off the Performance Race

In 2006, Nintendo released the Wii. The console's hardware was, by the standards of the day, deliberately weak. The processor was slower than the contemporary PlayStation 3 and Xbox 360. The graphics were less impressive. The marketing strategy abandoned the technical specifications race that the industry had been running since the 1990s.

What the Wii offered instead was a motion controller and a set of games that opened the console market to demographics — older adults, families with young children, casual players — who had never bought a gaming console before. The bet was that the technical specifications race had reached the point of diminishing returns, and that the next growth axis was not better graphics but a different relationship to the customer.

The Wii became, for several years, the best-selling console in the world. It outsold both of its more powerful competitors. Nintendo's competitors, optimizing on the specifications axis, had no response, because their entire production stack was built for an axis Nintendo had stepped off.

The Muda dimension was relational, in a specific sense. Nintendo bet on customers the existing optimization frame could not see — because those customers were not in the existing customer base, and so were invisible to any model trained on it.

4. The Refuseniks of Kyoto

The last category I want to name is the smallest in number but the most pointed in its logic. It is the set of restaurants and craft producers in Kyoto and elsewhere in Japan that have refused recognition that would have multiplied their commercial value.

Some traditional restaurants in Kyoto have declined inclusion in the Michelin Guide. The reasoning is consistent across the cases. A Michelin star would bring customers seeking the star, not customers seeking what the restaurant actually offers. The character of the customer relationship would change. The dishes would have to be reproduced more consistently to satisfy star-driven expectations, which would change the cooking. The atmosphere of the small dining room would shift. The relationship to the regulars who had been coming for decades would dilute.

None of this is captured on any spreadsheet. From the conventional business perspective, refusing a star is irrational. The star is the most reliable marker of commercial value the industry has. To

refuse it is to refuse free advertising.

The refusals are, on inspection, exact preservations of the Muda I have been describing. The relational density of the small establishment cannot survive star-driven volume. The judgments the chef makes night to night cannot survive standardization. The atmosphere that exists because the establishment is for a specific community cannot survive opening to a global audience.

The refusals are not provincial. They are strategic. The restaurants that refuse the star are choosing the substrate over the scale. They are paying the cost — the foregone revenue — that the substrate's preservation requires.

5. What They Share

Let me name what the cases share, because the common thread is the argument of this chapter.

One. Each of these organizations has refused something that the conventional optimization frame would have endorsed. Ise refuses the efficiency of leaving its buildings alone. Kong■ Gumi refuses to diversify. Stora Enso refuses to be the firm it was a generation ago. Patagonia refuses sales it does not endorse. Costco refuses to raise the hot dog price. In-N-Out refuses to franchise. Basecamp refuses outside capital. Toyota refused the pure-EV consensus. Nintendo refused the specifications race. The Kyoto refuseniks refuse the star.

Two. In each case, the refusal protects a specific dimension of Muda. The dimension is different across cases, but the structure is the same. The institution names what it will not optimize away, and pays the visible cost — foregone growth, foregone revenue, foregone recognition — to keep the invisible substrate intact.

Three. In each case, the protected Muda is, on the longer horizon, what carries the institution through changes that the optimization frame could not have anticipated. Ise survives a century of accelerating modernization. Kong■ Gumi survives the disappearance of Buddhism as a major institutional patron. Patagonia survives the commoditization of outdoor apparel. Costco survives multiple recessions. The Kyoto restaurants survive the global homogenization of fine dining.

Four. The refusal is the practice. It is not a one-time decision. It is a continuous discipline that has to be exercised, in fresh form, every time the optimization pressure renews itself. The pressure renews itself constantly. The discipline has to renew itself constantly.

The deepest common element is the fourth. None of these organizations is finished. None of them has solved the problem of keeping Muda alive. Each of them is, every quarter, every year, every generation, facing renewed pressure to be the conventionally-managed version of itself. The pressure does not go away because it has been refused before. The discipline of refusal is the institution's actual operating practice.

6. What This Chapter Establishes

By the end of this chapter, the reader should have three things.

First, a set of concrete cases in which Muda has been preserved in organizational form. The cases span thirteen centuries and multiple cultures. They are not exotic. They are studied across business literature, but rarely from the angle this book offers.

Second, a recognition that the preservation of Muda is, in practice, the recurring discipline of refusal. The institutions that have it have it because someone, repeatedly, refused the optimization that would have eliminated it.

Third, a frame for looking at your own organization. What is the optimization pressure currently being applied? What dimension of Muda would the optimization eliminate? What discipline of refusal has the institution actually practiced, and which has it lacked?

In Chapter 6 I turn to a meta-question. I have been writing this book as a translator. I am translating an eight-century-old methodology, originally developed in the religious-political world of Ano Zenjo, into the language of twenty-first-century management. The question of what it means to be a translator — rather than an heir, or an inventor, or a critic — needs to be made explicit. That is the work of the next chapter.

Chapter 6

The Translator's Position

On rendering an eight-century-old method into the language of contemporary management

Koto Igarashi

First Draft · May 16, 2026

This chapter is, in the structure of the book, the place where the author has to be specific about her own position. Up to this point I have written as if the question of who I am, and on what basis I am making these arguments, did not matter. That has been a convenience. The question matters, and this chapter takes it up directly.

The matter to address is a particular one. I am a descendant of Ano Zenjo. I have, by virtue of that descent, access to family materials and to a tradition of interpretation that few outsiders would have. I am also writing a book of contemporary management thought that names a method, and the named method shares its substantive content with the practice Zenjo represented eight centuries ago. The natural reading — and the reading that a marketing-minded reviewer would encourage — is that I am presenting myself as the inheritor of the method, the carrier of a lineage, the modern voice of an old tradition.

I am not. I want to say this directly and clearly, because the difference between the position I am actually taking and the position the reading would assign to me has consequences for how the book should be read.

I am a translator. I am not an heir.

This chapter explains the difference and why it matters.

1. What the Method Was, and What It Has Become

The method that informs this book was not, in its original form, a method of management. It was a discipline of esoteric Buddhist practice, embedded in a particular religious institution at a particular moment in Japanese history. The vocabulary it used — Omote, Ura, Hyouri Ittai, the practice of recognizing the two faces of reality as one — was the vocabulary of that institution. The training that conferred mastery in it was a multi-decade religious apprenticeship under a recognized lineage.

That institution and that lineage, in the form in which they existed in 1200, no longer exist. The religious world that produced them was destroyed, reorganized, and absorbed over the centuries between then and now. The kanj■ rite that established Yoritomo and Masako as Zenjo's disciples has not been transmitted in the same form. The Hashiriyu Gongen complex, where the method was practiced, has been institutionally transformed multiple times. There is no living person, in 2026, who can claim authority within the original lineage. The lineage, in that strict sense, is broken.

What survives is documentation. The texts that the lineage produced — most importantly the *Busshin Ittai Kanj■sh■*, which I have referred to in Chapter 4 — exist in archive. The structural content of the method, the framework that organized the practice, can be reconstructed from the texts. The personal transmission cannot. The transmission required a master and a disciple in living relationship, and the relationships have been severed by time.

An inheritor would be someone within the broken lineage who claimed authority by the lineage's own standards. There is no such person. To claim that position would be a fraud. I will not make the claim.

What is left is the documentary substance, the structural content of the method, and the open question of what to do with it. That open question is where this book starts.

2. Why a Translator

The position I have taken is the translator's position. I will explain what this means, because it is unfamiliar in management writing and because the difference matters.

A translator works between two languages. The translator does not claim authority within either of the two languages. The translator's claim is more limited: that she has read carefully in the source language, that she has thought carefully about the target language, and that she has rendered the substance of the source into the target with as much fidelity to the source as the target allows.

The translator's authority is grounded in the work, not in lineage. A good translator of Sappho is not an heir to Sappho. A good translator of Confucius is not a Confucian master. The work stands or falls on whether the rendering captures what was there to capture.

This is the position I am taking. I am rendering what survives, in the documentary form, of an eight-century-old methodology into the language of contemporary management. The substance is, to my reading, the substance that was there. The vocabulary — Muda, judgment, relationship, the long arc, the unmeasured — is not the original vocabulary. It is the target language. I take responsibility for the rendering.

There are things this position permits and things it forbids.

It permits me to do exactly what I have done in this book: take the structural content seriously, render it into terms that the contemporary management reader can work with, and argue for its present-day relevance. The argument can be inspected. The case studies in Chapter 5 either hold up or they do not. The reasoning about AI in Chapter 3 either tracks or it does not. The historical reconstruction in Chapter 4 either has primary-source support or it does not. The translator's claim is that the work can be evaluated on its merits.

It forbids me from making certain claims that an heir could make. I cannot, for example, claim that someone who reads this book has been initiated. The kanj■ rite is not transmissible through prose. I cannot claim authority over disputes about what the original methodology "really" meant, because the original is in archive and is interpretable, and other readings are possible. I cannot claim the personal authority of a master within a living tradition, because the tradition in the relevant sense is not living.

What I can claim is that I have read carefully, thought carefully, and rendered carefully. That is the entire claim.

3. The Act of Naming

The translator's most consequential decision is what to call things.

In Chapter 1 I argued that "Muda" was a deliberate choice — a borrowing of Lean's central term, inverted to point to the opposite of what Lean uses the term for. I described this as a strategic act. I want to return to it here, because it is the part of the work where the translator's responsibility is most concentrated.

The historical Zenjo did not use the word Muda. The word in his vocabulary that maps most closely is *Ura*, the back of the cloth, the hidden side of the visible. The translation is mine. It is not a literal translation — Ura, in modern Japanese, still means "back" or "rear," and a literal rendering would not have the strategic content I want.

I chose Muda for the strategic content. The word "Muda" in modern management vocabulary belongs to Lean. Lean has used it for over half a century to mean waste, the thing to be eliminated. To take this word and reuse it to mean the substrate that survival depends on is to make a polemic move at the level of the vocabulary itself. The reader of Lean literature, encountering this book's use of Muda, has to make the inversion explicit. The inversion is the argument.

This is a translator's choice. An heir would have used the lineage's own vocabulary. An inventor would have used a coined neologism. I have done neither. I have taken the contemporary word that is occupying the territory I want to occupy, and I have reused it. The choice is mine. The responsibility for the choice is mine.

This is what I mean when I say, in the prologue, that the naming is my responsibility. The thing the name points to has been there for centuries. The name is a 2026 act, made by a particular person for particular strategic reasons.

4. Two Sites

I keep two sites for two kinds of work, and the relation between them is worth describing.

The historical-academic work lives at tokiwatanabe.org. The work there is what every academic discipline calls primary — the actual reading of the documents, the paleography, the archival reconstruction, the work of establishing what the surviving record can be made to say. Specifically, the work on Ano Zenjo: the ka■ analysis from the 1183 deed, the dating of the *Busshin Ittai Kanj■sh■*, the reconstruction of the Sansho three-shrine network, the geography of Ano-sh■ at the foot of Ashitaka. This is the slow, foundational work. It accumulates over years. The audience is historians, archivists, and readers who care about what the twelfth century actually contained.

The contemporary management work lives at mudamethod.org, which is the site you are on. The work here is the translation — what does the substance the academic work has established mean for organizations in 2026? What does it mean for AI, for succession, for the long arc of corporate

life? The audience is different: founders, executives, board members, people responsible for institutions whose decade may not survive the current optimization pressure.

The relation between the two sites is exactly the relation this book has been arguing for. The visible work — the management argument, the strategic implications, the recommendation — rests on a substrate that is not itself the visible work. The substrate is the primary source argument, the years of paleographic detail, the careful reconstruction of what was actually there in 1183. None of that substrate is on mudamethod.org. It does not need to be. It is at tokiwatanabe.org, where it belongs.

The book you are reading is, in this small way, an instance of its own argument. The Omote — what is here, in this chapter, in this synthesis — stands on an Ura that is somewhere else. A reader who wants the Ura can find it. A reader who wants the Omote can stay here. Both are the same work, in two registers, because the work is genuinely two-sided.

5. What This Implies for the Reader

The translator's position has consequences for how the book should be read.

First, the argument is open for evaluation. The case studies are inspectable. The reasoning is inspectable. The historical evidence is inspectable, in the academic essays at tokiwatanabe.org. Everything I claim is open to inspection, criticism, and revision.

Second, the book is not closed. The translator's rendering is one rendering. Other renderings are possible. A reader who, after working through this book, develops a different translation of the same source — a different vocabulary, a different organization of the dimensions, a different account of how Muda manifests in practice — has done something useful and is not, by virtue of the difference, wrong. The translation that I have made is the best one I have found. It is not the only one possible.

Third, the practice of keeping Muda alive in your own organization is not the practice of following this book's instructions. It is the practice of recognizing, in your own context, what the four dimensions look like and what would be required to preserve them. The work has to be done on your own ground. The book gives you a vocabulary and a frame. It cannot give you the judgment that has to be made fresh, in your own situation, every time.

The book is a translator's offering. It is not a master's instruction.

6. Why I Have Written This

I want to say, briefly, why I have written this book at this moment.

The argument I have been making — about AI's structural limits, about the visible commoditization that follows from those limits, about the consequent strategic relocation of competitive advantage to dimensions AI cannot see — is, in 2026, the most consequential

argument I know how to make for the next decade of organizational practice. The argument is not original. Versions of it have been made by others. What I have added is, I hope, the specific translation that connects the contemporary argument to the eight-century-old methodology that, on inspection, has been making it all along.

The reason for the connection is not nostalgia. It is that the older method has the advantage of having been tested. Most contemporary management frameworks have been tested for, at most, a generation. The dimensions I have called Muda were the operating substrate of an institution that produced organizations that lasted, in some cases, for over a thousand years. The argument is not "what worked in the twelfth century will work now." The argument is "what worked across the centuries, in a form we can now reconstruct, is precisely what the current optimization frame has been removing."

If the argument is right, then the practice the book points to is not optional. It is, in the most direct sense, what an organization in the AI era must do to be alive at the end of the next disruption.

I have written the book because I think the argument is right. I am asking the reader to inspect it and decide.

7. What This Chapter Establishes

By the end of this chapter, the reader should have three things.

First, a clear understanding of my position. I am a translator, not an heir. The substance of the method has been there for centuries. The naming and the rendering are 2026 work, for which I take responsibility.

Second, a clear understanding of what this implies for how the book should be read. The argument is open to inspection. The case studies are open to inspection. The historical reconstruction is open to inspection. Nothing is offered on the authority of lineage.

Third, the relation between the two sites. tokiwatanabe.org is where the academic-historical work lives — the primary source argument, the slow substrate. mudamethod.org is where the translation lives. The visible book rests on the invisible substrate, in exactly the way the book itself argues organizations do.

In Chapter 7 I turn from the meta-question to the operational one. Given the argument of this book, what does it actually look like to keep Muda alive in your own organization? Not in principle, but in the practices of an individual career, an operating organization, and a wider institutional landscape. That is the work of the next chapter.

Chapter 7

Implementation

Keeping Muda alive across three levels — the individual, the organization, and the institutional landscape

Koto Igarashi

First Draft · May 16, 2026

The book has been building, chapter by chapter, an argument and a vocabulary. This chapter takes the argument into practice. It asks the question every reader who has stayed this far is entitled to ask: what does this look like in my actual situation, this week, this quarter, this year?

The answer operates on three levels. The same logic applies to each, but the practice differs by scale. At the level of the individual, you are responsible for the four dimensions of Muda in your own working life — your judgment, your relationships, your time horizon, your unmeasured commitments. At the level of the organization, you are responsible for protecting those dimensions in the institution you are part of, against the optimization pressure that will, by default, eliminate them. At the level of the wider institutional landscape, you are responsible for the long-arc choices about which institutions to support, which to leave, and which to build — choices whose payoff is generational rather than personal.

This chapter walks through the three levels in order. The aim is not to give you a checklist. The aim is to give you the practice that, in each case, the discipline takes.

1. Level One: The Individual Career

The individual practice of Muda has, in the end, four sub-practices, one for each dimension. They are not strategies. They are habits, sustained over years.

Protecting Judgment

The Muda of judgment, as I defined it in Chapter 1, is the capacity to make the call outside the precedent. To protect this capacity in your own working life means, in practice, refusing to outsource your important judgments to the available models.

The temptation, in 2026, is to ask the model. The model is excellent at synthesizing options, identifying tradeoffs, summarizing what has been said about similar decisions. None of this is wrong. The risk is that the model's confidence in its recommendation displaces the act of judgment that the situation actually requires. You read the model's output. You agree with it. You move on. The decision has been made by extrapolation. Your own judgment has not been exercised.

The practice is to read the model's output, then ask yourself: *does this situation actually resemble the situations the model is extrapolating from?* The honest answer, on questions that matter, is often no. Your situation has features the model did not see. The features are usually the features that will determine the outcome. The judgment is yours to make.

This sounds simple. It is not. The pressure to defer to the model is institutional, social, and increasingly built into the tools themselves. Protecting your own judgment is, against this pressure, a disciplined practice. It is also the practice that, over a career, accumulates the actual experience that lets you make the calls the model cannot.

Building Relationship Beyond Transaction

The Muda of relationship is what survives after the transaction has ended. To protect this in your own career means investing, over years, in relationships whose value cannot be priced.

The practice has a specific shape. Stay in contact with people whose current role has nothing to offer you. Help people whose careers have moved sideways. Refer good candidates to organizations you no longer work with. Show up for funerals and small ceremonies. Send the email that has no purpose. Reply to messages from people you have not seen in years.

None of this scales. None of this is efficient. None of this would survive an optimization audit. The point is precisely that it would not. The substrate of relationship that determines the second half of your career is built by these unscalable, inefficient, unbookable acts in the first half. By the time you need the call to be returned, the call gets returned because the relationship has been kept warm by years of attention that produced no measurable outcome at the time.

Most professionals stop doing this in their thirties. They become busy. They optimize their network. The relationships that survive are the transactional ones. By their fifties, the substrate has thinned, and the calls that they would now need are not the calls that get returned.

Operating on the Long Arc

The Muda of the long arc is the capacity to take seriously a horizon that no measurement system will validate. In an individual career, this is the discipline of making moves whose payoff is decades out.

The practical form is straightforward. Choose work whose value compounds. Develop skills the market has not yet learned to price. Build reputational capital in domains that are currently unrewarded but that you have reason to believe will matter. Make the move that no one is making, because everyone agrees it is the wrong move, on a horizon you can defend.

This requires the ability to suffer through periods when the long-arc bet is not paying off and the conventional bet is. Most professionals cannot sustain this. They abandon the long arc somewhere around the fifth year, when the visible reward gap has grown large enough to be socially painful. The people who hold the long arc through the painful middle period are the people who, in their sixties, are doing work that nobody else can do.

Toyota's hybrid bet is the company-level version. The career-level version is the same shape on a smaller scale.

Keeping Something Unmeasured

The Muda of the unmeasured is what exists only because it has not been quantified. In your career this is the work, the relationship, the practice that you have deliberately not put on the dashboard.

The hardest version of this practice is also the most consequential. In a world that increasingly demands metrication of everything — your productivity, your wellness, your sleep, your output, your contribution — the practice of keeping a corner of your life unmeasured is, simultaneously,

an act of professional discipline and an act of resistance. The corner can be small. It can be the hour every morning that is not on any calendar. It can be the friendship that has never been instrumentalized. It can be the project you work on that no one is paying you for and that you do not promote.

What this protects is the substrate that, when the dashboard fails, you still have. Most professionals discover late in their careers that the most consequential work of their life was done in the corner they had not metricated. The work was done there because metrication would have killed it.

2. Level Two: The Operating Organization

The organizational practice of Muda is the same shape, scaled up. The organization's leaders are responsible for protecting the four dimensions in the institution itself, against the optimization pressure that will, by default, eliminate them.

Protecting Organizational Judgment

The organizational practice begins with the recognition that not every decision should be modeled. The model is appropriate for decisions where the data is dense, the situation is recognizable, and the optimum is more likely to be correct than the human alternative. For these decisions, the organization should use the model. For the decisions that determine whether the organization is alive in a decade, the model is the wrong tool. The practice is to know which is which.

The institutional form of this practice is a senior group that has the authority to commit on questions that the model cannot answer. The group's authority is grounded in something other than data fluency. It is grounded in the demonstrated capacity to make calls that the data did not support and that turned out to be right. The composition of the group is itself a Muda decision. The members are not the most data-driven executives. They are the executives who have demonstrated, across cycles, the capacity to make the call outside the precedent.

Protecting Organizational Relationships

The relational Muda of an organization is the substrate of contacts, ex-employees, retired customers, dormant suppliers, occasional partners, and unaffiliated advisors that the formal CRM does not see. Protecting this is, in practice, a discipline of not pruning the organizational network according to the optimization frame.

The discipline has a specific institutional form. Allocate budget to relationship maintenance that has no immediate return. Sponsor reunions of former employees. Maintain contact with retired customers. Continue to invite the founder to events long after the founder has ceased to have operational authority. None of these activities will appear on any performance review as productive. All of them are the practice that keeps the relational substrate alive.

The organizations I studied in Chapter 5 do this almost without thinking. The relationships are part of the institutional fabric. The newer organizations that try to install this practice, after the optimization has removed it, find that it cannot be reconstructed quickly. The substrate is built over decades. It is destroyed in a budget cycle.

Protecting the Long Arc

The organizational practice of the long arc is the discipline of making and protecting strategic bets that the planning cycle cannot evaluate.

In practical institutional terms, this requires a structural arrangement that exempts certain decisions from the standard governance. The arrangement might be a separate budget line for long-horizon bets. It might be a designated small group that has the authority to commit resources on decade-plus horizons without the standard ROI gating. It might be an explicit policy that the standard planning cycle does not apply to the firm's foundational strategic commitments.

What this structural arrangement is for is to protect the bets that, by their nature, will not survive the standard planning process. The standard process will, every time, optimize them away. The arrangement exists to keep them alive against that pressure.

Most organizations do not have this arrangement. Most organizations should.

Protecting the Unmeasured

The hardest organizational practice is the deliberate preservation of zones where the system does not apply.

In a healthy organization, there are spaces where the standard performance metrics do not reach. The R&D; group whose annual output is not measured by patents. The library or archives function whose value cannot be benchmarked. The senior advisor whose contribution will not appear in any performance review. The internal community that meets without an agenda. These zones are where the unmeasured substrate of the organization lives.

Every budget cycle proposes to optimize these zones. The proposals are always made in good faith. The proposers are not trying to destroy the organization. They are trying to apply the optimization frame uniformly, which is what the optimization frame asks them to do. The discipline against ossification is the deliberate practice of refusing these proposals, on grounds that are explicitly outside the proposal's frame.

This requires institutional authority that can defend the unmeasured against the measurable. In the long-lived organizations of Chapter 5, the authority is built in. In modern organizations, it has to be deliberately preserved. The first place to look for organizational health is whether such authority exists and is being exercised. If the answer is no, the institution is somewhere in the late stages of ossification, regardless of what the financial dashboard says.

3. Level Three: The Institutional Landscape

The third level is the level above any single organization. It concerns the wider landscape of institutions that determine what kind of careers and what kind of organizations are possible at all.

This level is the one most often left out of management books, because it does not fit the format. A management book is supposed to be addressed to the individual reader or the firm. The institutional landscape — the universities, the regulatory bodies, the professional associations, the civic institutions, the cultural commons — sits above the firm and is shaped by collective decisions that no single reader can make.

And yet the level matters, because the firm and the career exist within the landscape. A firm that wanted to operate by the practices of this book would find the operating ground much easier in some institutional environments than in others. A career oriented toward the long arc would find some societies welcoming and others hostile. The landscape is the substrate of the substrate. Without attention to it, the practices of the first two levels are running on borrowed ground.

The practice at this level is necessarily slower and more collective. It includes the long-cycle work of supporting institutions whose function is to preserve dimensions of Muda at scale. The institutions that support multigenerational craft transmission. The institutions that fund research on horizons no commercial actor will fund. The civic structures that protect relational density at the level of neighborhoods and small cities. The professional associations that maintain standards no individual firm can sustain alone.

Most of these institutions have been weakening for decades, under optimization pressure that has been operating at the societal scale. Restoring them is not a project for any single firm or career. It is the work of generations of collective attention.

What the individual reader can do, at this level, is recognize which institutions are doing this work and support them. Support, in this register, is unglamorous. It is the small donation to the historical society. The membership in the professional association whose immediate benefit is unclear. The participation in the civic body whose work moves slowly. The teaching that happens outside any compensation framework. The hours of attention given to the institutional commons.

None of this is heroic. The aggregate of these unheroic acts, sustained across many people for many years, is the difference between an institutional landscape that supports the practices of the first two levels and one that does not.

4. The Practice Across All Three Levels

What the three levels share, in the end, is the same shape. At each level, the practice is the discipline of refusing what the available optimization frame would eliminate, and of investing what the frame cannot price.

At the individual level, you refuse to outsource your judgment and invest in relationships that will not pay back for years. At the organizational level, you refuse to prune the network according to the dashboard and invest in zones the dashboard cannot see. At the institutional level, you refuse

to abandon the slow institutions and invest in the commons that no single actor benefits from immediately.

The shape is the same. The scale differs. The discipline, at all three levels, is to do what the optimization tells you not to do, because the optimization is operating on a frame that cannot see what you are protecting.

This is not heroic work. The cases in Chapter 5 are not stories of heroism. They are stories of consistent, multi-generational, often quiet discipline. The hot dog price did not change. The franchise was refused. The Seng■ was performed on schedule. The Michelin star was declined. Each instance is small. The accumulation, across the long horizon, is the institution that survives when its competitors do not.

5. What This Chapter Establishes

By the end of this chapter, the reader should have three things.

First, a frame for implementing the argument of the book in your own situation. The frame operates on three levels — individual, organizational, institutional — and the practice on each level has the same logical shape: refuse what the optimization eliminates, invest in what it cannot price.

Second, a vocabulary for the specific sub-practices at each level. Protecting judgment. Building relationship beyond transaction. Operating on the long arc. Keeping something unmeasured. These four sub-practices apply at each level, with different mechanics but the same logic.

Third, the recognition that the practice is not heroic. It is the small, repeated, often quiet discipline of refusing the optimization that everyone else is accepting. The discipline does not produce a dramatic story. It produces an organization, or a career, or an institution that is alive at the end of a horizon long enough for the difference to matter.

One chapter remains. The closing chapter will name, as cleanly as I can, what the transition from Lean to Muda actually is. It is not, in the end, a project at the scale of any single book. It is the shape of the next era's competitive ground. That is the work of the epilogue.

Epilogue

From Lean to Muda

An era declaration

Koto Igarashi

First Draft · May 16, 2026

This book began with a claim about the present moment. The twentieth century, building on Taiichi Ohno and the operational philosophy that James Womack and Daniel Jones eventually named *Lean*, produced what may be the most refined apparatus for the visible work of management that human institutions have ever built. The apparatus runs the world now. It runs supply chains, it runs production, it runs services, it runs software, it runs, in 2026, the AI systems that are taking over the apparatus itself.

The argument I have been making is not that the apparatus was wrong. It was right, within its frame. The argument is that the apparatus operates on one side of a word, and the next era of competitive advantage lives on the other side of the same word.

The word is Muda. Lean used it to mean waste. I have used it, in this book, to mean the human substrate of the organization — judgment, relationship, the long arc, the unmeasured — the part that survival depends on. The same word, pointing in opposite directions. The transition the book points to, in its title, is the transition between the two readings.

From Lean to Muda.

1. What the Era Was

The era of Lean — call it the second half of the twentieth century and the first quarter of the twenty-first — was an era in which the visible was the binding constraint. Inventories cost capital. Defects cost rework. Variation cost quality. The Lean apparatus addressed each, in turn, with discipline that compounded over decades. The result was an order-of-magnitude improvement in nearly every operational metric the apparatus tracked.

The era succeeded so completely that, by 2026, the visible side of management is approaching the limit of what optimization can do. The supply chains are tight. The processes are mapped. The dashboards are exhaustive. The remaining gains, on the visible side, are increments of a few percentage points. The strategic content of the visible side has been consumed.

And precisely at this moment, artificial intelligence has arrived to take over what remains. AI is now, in 2026, doing — and over the next decade will increasingly do — what the visible apparatus of Lean management did during the twentieth century. The pattern recognition, the optimization, the forecasting, the dashboard analysis. All of it. Better than humans, cheaper than humans, faster than humans.

The era of Lean is not ending because Lean failed. It is ending because Lean succeeded, and what it succeeded at is now being automated.

2. What the Next Era Is

The next era's binding constraint is on the other side of the word.

What AI cannot do, as I argued in Chapter 3, is exactly what the Muda of this book has named. AI cannot make the judgment that lies outside precedent. AI cannot maintain the relational substrate that does not produce transactions. AI cannot operate on horizons longer than its measurement window. AI cannot represent substance that, by its nature, does not reduce to variables.

These are not engineering limitations. They are intrinsic to what AI is. The next decade will see AI systems become extraordinarily good at the visible work of management. They will not become good at the four dimensions of Muda, because the dimensions are defined by what the optimization frame cannot reach.

This is the strategic geography of the era ahead. The visible is commoditized. The invisible is the territory where competitive advantage now lives. Every organization, every career, every institutional commitment, is increasingly being evaluated on its capacity to operate on the invisible side.

This is not a forecast. It is a description of the trajectory already in progress. The organizations that are pulling ahead, in 2026, are the organizations that have figured out — explicitly or by instinct — how to keep the four dimensions alive against the optimization pressure. The organizations that are falling behind are the organizations that have surrendered the invisible to the model.

3. What the Move Requires

The move from Lean to Muda is not a single decision. It is a practice — at the individual scale, at the organizational scale, at the institutional scale. Chapter 7 named the practice on each scale. I will not repeat it here. I will name only what the practice has in common.

At every scale, the practice is the discipline of refusal. The discipline of refusing what the available optimization frame would eliminate, and of investing in what the frame cannot price. The discipline is, by its nature, costly in the short run. It looks irrational on every contemporary measure. The refusals — to outsource judgment, to prune relationships, to abandon long-arc commitments, to metricate everything — are the practice.

The organizations of Chapter 5 — Ise, Kong■ Gumi, Stora Enso, Patagonia, Costco, In-N-Out, Basecamp, Toyota, Nintendo, the Kyoto refuseniks — each illustrate, in their own register, what the discipline looks like in practice. Each is making, repeatedly, the choice the optimization frame would not endorse. Each is, in the long run, the institution that survives.

What they share is not a strategy. It is a posture. The posture is the willingness to be wrong on the conventional dashboard, in order to be right on a longer horizon. The discipline of holding that posture, through the periods when the dashboard's verdict is socially painful, is the substance of the move from Lean to Muda.

It is not glamorous. It is not heroic. It is the slow, costly, often invisible work of keeping the substrate alive when the optimization pressure is telling you to abandon it. Most institutions will

not do this work. They will, by default, follow the optimization, and they will be alive on the dashboard until they are no longer alive at all. The institutions that survive the era ahead will be the institutions that did the unglamorous work.

4. A Note on the Origin, Again

I have said, in the prologue and throughout the book, that the methodology I am translating did not originate with me. The thing the book names had a name in the twelfth century, in the practice of Ano Zenjo and the lineage at Hashiriyu Gongen. The name was different. The substance was the same.

This matters because the substance is older than any of the contemporary cases I have cited. It is older than Lean. It is older than the modern corporation. It is older than every framework the current management discourse contains. What I have called Muda was the operating substrate of organizations that lasted, in some cases, for over a thousand years.

The implication is straightforward. The argument of this book is not a novel insight produced for the AI era. The argument is the recovery of a discipline that the twentieth century, in its concentration on the visible, allowed itself to forget. The AI era is, simply, the moment when forgetting has become unaffordable.

The book does not ask the reader to trust me on the strength of the lineage. As I said in Chapter 6, I am a translator, not an heir. The substance is what matters, and the substance can be inspected.

5. What the Reader Should Do Now

The reader who has stayed through the book is, at this point, entitled to a final practical word.

The practical word is brief. The discipline of keeping Muda alive in your own context — your career, your organization, the institutions you support — is the work of the rest of your professional life. It will not be completed in a quarter, or a year, or a decade. The discipline is not a project. It is a posture, sustained over the long arc, against the optimization pressure that will renew itself every cycle.

Begin somewhere small. Refuse one optimization that you would otherwise have accepted. Make one decision that the available data does not support. Invest in one relationship that has no immediate return. Commit to one horizon that lies past any measurable window. Preserve one corner of your work that you will deliberately leave unmeasured.

None of these acts, by themselves, will produce a visible outcome. The compounding will. The career, the organization, the institutional commitment that has been built by years of these small unglamorous acts will, by the end of the next decade, be the career, the organization, the commitment that competitors cannot understand and cannot copy.

This is the move from Lean to Muda. It is a quiet move. The era it inaugurates will be loud — full of AI, full of optimization, full of the visible competition that the current apparatus has perfected. The advantage will be on the side of the quiet move.

6. The Closing

I will end where the book began.

The prologue opened with a question: what did efficiency remove? The book's eight subsequent chapters were, in different registers, the same answer. Efficiency removed the human. Judgment, relationship, the long arc, the unmeasured. The substrate that organizations, careers, and institutions depend on when the dashboards are wrong.

The era ahead will reveal which institutions kept the substrate and which did not. The reveal will happen, in some cases, slowly. In other cases, quickly. The Hojo regents inherited the form of the Kamakura Shogunate in 1203 and held it, formally, for over a century while the substance drained out. The contemporary institutions that have surrendered the four dimensions to the model will not have a century. The disruption cycles are faster now. The reveal will, for most of them, happen within a decade.

For the reader who has read this book, the implication is direct. You know what to look for now. You know what your own institution has been preserving, and what it has been letting go. You know which of your own habits are part of the optimization that is hollowing you out, and which are the small acts of preservation that, over the next decade, will be the substance of the career you have left.

The choice is not abstract. It is the choice you will make tomorrow morning, in the small decisions that the optimization frame is asking you to optimize, and the slightly different decision that the practice of this book is asking you to consider.

Make the slightly different decision.

And then make it again, and again, and again, across the years, until the substrate you have been building is the thing that holds when everything else gives way.

That is the work.

That is the era.

From Lean to Muda.