WP 5 · THE WEAKEST CORRECT DECISION

> *Weakest Correct — among all correct options, the one that makes the fewest further commitments.*

Gerald Huesch — Global Leadership Consultancy L.L.C. · Global Leadership School · Global Leadership Society · Director, Center of Behavioural Leadership, St. Gallen Business School

Draft for discussion and further scholarly development · Version of 30 September 2026

ABSTRACT

All previous papers assumed that too little information is the problem. This paper treats the reverse case, which has become the more common one: deciding under cognitive overabundance.

At AGI-26, Michael Timothy Bennett shows that the usual selection criterion — the shortest, simplest explanation — depends on the encoding and is therefore unusable. His replacement is **weakness**: the number of commitments an option leaves open. Ben Goertzel builds on it a priority signal across dissimilar domains. David Orban shows what happens when no one decides this way: the **Freezing Equilibrium**. And Alfredo Ibias and colleagues show the property absent from every reporting system — abstention.

Keywords: Weakest Correct · weakness · extension · Freezing Equilibrium · reasoned abstention · LEAD21 · 5+1 Baskets of Decision-Making

Suggested citation: Huesch, G. (2026). The Weakest Correct Decision. Working Paper Series 2026, Series 1 "The Superintelligent Organization", WP 5. Global Leadership Consultancy.

PART 1 · THE FINDING

1 THE QUESTION

A mountain guide stands at a fork. The left path is shorter and leads safely to the intermediate camp. The right is twenty minutes longer and opens at the top onto a field from which three further routes remain available.

Both paths are correct. The experienced guide takes the right one — not because he is better informed, but because he knows that up there he will not yet know what the weather is doing.

The question of this paper: **how does one decide when information is not scarce but abundant?**

2 WHAT AGI-26 SHOWS

The shortest criterion is the wrong one. In *The Wrong Razor*, Michael Timothy Bennett attacks Ockham at the root. The rule "take the shortest correct description" is not wrong, it is **empty**: *"What counts as short depends on the encoding."* On a finite class, any prior can be made Ockham-like by a suitable code.

His replacement is a semantic quantity: **weakness**.

> $w(\pi) = |Ext(\pi)|$ — "Weakness counts how many further commitments remain compatible with correctness."

That is: how many further commitments remain compatible with this option? The more, the weaker — and the better.

Bennett's formal results: under exchangeable, as yet unknown requirements, the weakest correct policy is **Bayes optimal**. Regret equals exactly the weighted forgetting cost. And the invariance test that makes the difference: *"A serious measure must survive semantics-preserving recodings. Weakness does."*

The sentence his poster closes with is the decision rule of this paper:

> **"Do not ask which correct policy is shortest. Ask which correct policy makes the fewest further commitments."**

And from it comes a priority signal. In *Weakness is all you need*, Ben Goertzel builds on this a quantity comparable across dissimilar domains: each subsystem exports its weakness in its own

unit, the orchestration trades it off against cost.

> `ρ_S = Δw_S / cost(s)` — weakness gain per unit of effort.

His honest framing: *"Not literally all you need — but a coordination principle."* That is exactly what a leadership group needs: not the one metric, but a shared currency in which a sales initiative and an IT initiative can be compared.

When no one decides this way, the organization freezes. David Orban models coordination as an engineering problem and names the state that arises when decision speed exceeds verification capacity: the **Freezing Equilibrium**.

What is meant: when verifying a recommendation costs more than its expected benefit, both stop — the verifying and the acting. Inaction becomes the rational choice. His model predicts sharp transitions between self-healing and self-destabilizing states.

This is not a motivation problem. It is a computation.

Abstention is the rare property. With *Unsupervised Cognition*, Alfredo Ibias and colleagues build a method that refuses when the situation does not fit:

> **"I do not know." — "If no learned representation is similar enough, the Nodule returns no Projection rather than hallucinating a label — a rare property for an unsupervised learner."**

Every reporting system ought to be able to do this. Almost none can, because no field is provided for it.

Two processes, not one. Henry Ansah, Soheil Saneei and Shreya Banerjee separate semantic planning from execution — System 2 plans in language, System 1 controls in latent space. Their core finding: *"Separating semantic planning from control stabilizes policy learning."* Whoever mixes the two in the same body destabilizes both.

3 TWO INSTRUMENTS

The extension as a mandatory field of every resolution.

Every resolution gets one additional line: **what else does it rule out?** Not what it does — that is already written. But which future paths it closes.

RESOLUTION	WHAT IT DOES	EXTENSION — WHAT ELSE IT RULES OUT
We standardize CRM group-wide	one system	two markets with divergent regulation, for the next four years
We fill the post internally	fast appointment	access to a capability not present in-house

The line costs two minutes per resolution. It makes weakness visible without anyone having to have read Bennett.

Reasoned abstention, with three mandatory fields.

"We are not deciding this today" is a decision only if three fields are filled:

1. **What is missing?** Which information exactly, not "more clarity".
2. **When will it be available?** A date, not a condition.
3. **Who decides then, and what holds until then?** A person and an interim rule.

Without these three fields a postponement is not a resolution but a Freezing Equilibrium with better manners.

4 ONE EXAMPLE, WORKED THROUGH

Illustrative example, not a case report.

A retail company must replace a merchandise management system. Three options are on the table, all three technically sound.

Option A is the cheapest and fits today's processes exactly. **Option B** costs 20 per cent more and leaves open two process variants nobody needs today. **Option C** is the most expensive, with the greatest flexibility, and would tie up eighteen months of implementation.

The usual evaluation picks A: shortest description, best price, exact fit.

Evaluation by weakness asks differently: which future commitments remain open? A rules out the two process variants — and exactly one of them will be needed as soon as the company acquires a national subsidiary, which is in the strategy. C keeps everything open but ties up eighteen months of attention needed elsewhere.

B is the weakest **correct** option: correct for today, and it leaves the most open without blocking the calendar.

Without the extension line this comparison would never have been made — because A wins in every classical evaluation.

5 HOLD ON

Objection 1: This is just indecision with a better justification.

No, and the difference is sharp. The weakest **correct** option is a decision: it is taken, it is correct, it is implemented. Indecision takes none. Bennett explicitly requires correctness as a precondition — the choice is made among correct options, not between deciding and not deciding.

The advantage: whoever decides this way decides **earlier**, because they need not wait for the information that would justify the strongest commitment.

Objection 2: Sometimes you have to commit. Keeping options open is expensive.

Correct, and that is exactly why it does not mean "keep everything open". It means: pay only for the commitments you need. Every additional one is a purchase on credit — it costs nothing today and everything on the day the situation turns.

The advantage: the extension line makes that credit visible before it is taken out.

6 THE CONDITION

> ### Condition 5 · Decision level
> Among all correct options, choose the one that forces the fewest further commitments. And every decision needs a level and a person — otherwise inaction becomes the rational choice and the organization freezes despite better information.

PART 2 · WHAT LEAD21 ANSWERS

REQUIREMENT	INSTRUMENT	MECHANISM	COVERAGE
Level and person per decision	5+1 Baskets of Decision-Making (07)	Every decision is assigned a basket and thereby a level and a person	covered
Decision path without gaps	9 Chambers of Decision (01)	Situation, goal, options, evaluation, resolution, adjustment	covered
Making what is critical visible	Laser Meeting (04)	bottom-up, fixed cadence	covered
Extension per resolution	no dedicated instrument	—	**open**
Reasoned abstention	no dedicated instrument	—	**open**
Priority across dissimilar domains	5+1 Baskets (07)	Baskets order by kind, not by gain per effort	**partial**

CRITICAL APPRAISAL

The 5+1 Baskets solve exactly the problem Orban describes formally: they ensure every decision has a level and a person. That blocks the most common route into a Freezing Equilibrium — the contradiction that belongs to nobody.

What is missing is the selection criterion itself.

The extension appears in no instrument. LEAD21 governs who decides and how the path runs. It does not govern by what standard one chooses among several correct options.

Reasoned abstention is likewise absent. The baskets give an assignment, but no admissible answer "not today, and here is how".

The priority signal is ordinal, not computational. The baskets order by kind of decision. Goertzel's quantity orders by weakness gain per effort — a different question, and the more usable one for portfolio decisions.

The mapping comes from the developer of the system and is not independently verified.

WHAT THE SERIES GIVES BACK TO LEAD21

The extension line becomes a mandatory field in Chamber 7. One line per resolution: what else does it rule out, and for how long?

Reasoned abstention becomes the sixth basket. With the three mandatory fields — what is missing, when will it be available, who decides then and what holds until then. Without it, every postponement is a silent one.

A weakness gain per effort for portfolio meetings. Not as an exact figure, but as a shared question: which initiative keeps the most open per euro spent?

FALSIFIABLE PREDICTION

> **Prediction 5.** Leadership groups with a mandatory extension and reasoned abstention revise fewer resolutions within 24 months than comparable groups without both — and postpone without a date less often.

PART 3 · ABOUT GERALD HUESCH

The question of how to decide under pressure arises most sharply in growth. Gerald Huesch accompanied a company over eighteen months from €300 million to €2.3 billion in revenue — a situation in which information is not scarce but arrives every week in greater volume than a leadership group can verify. That is exactly where what Orban describes formally as a Freezing Equilibrium arises: decisions are left lying, not because anyone hesitates, but because verifying becomes more expensive than acting.

In the accompanying hyper-growth programme, 27 leaders from finance worked over seven weeks; afterwards they reported +40 per cent efficiency and +40 per cent performance. That is an expectation survey with self-report, not a measurement — the labelling belongs with it, because this very series charges the leadership literature with leaving it out.

The same subject is treated in his piece *5 Steps for Superior Decision Making* on Medium; the podcast *Wachstum und Mitarbeiterbindung: Erfolgreiche M&A-Strategien* shows the other side of the same situation — what happens to people while the decisions pile up.

> **Systems shape behaviour.** Through the selection of the right instruments, mechanics and leadership codes — derived from values and leadership guidelines — the behaviour of everyone involved is synchronized and aligned. Not through motivation, but through a system that produces the right behaviour when applied.
>
> Sales has Salesforce. Finance has SAP. Leadership has LEAD21.

About the author. Gerald Huesch is the creator of the leadership operating system category – "one of the most professional and creative minds in the domain of leadership and difficult change projects" (Boston Business School). With LEAD21 he built the first operating system for leadership – 21 evidence-based methods, the comprehensive solution and the missing link for complete transformation, extreme growth and high-volume post-merger integration: a unique leadership system for human coordination at scale, in which alignment and culture emerge as a by-product. He has trained more than 30,000 executives in over 30 countries (NASA, Rolls-Royce, Airbus, Omnicom and others) and delivers executive leadership training, leadership development and executive mentoring for executives in Germany, the Americas, Asia Pacific and the UAE. Participants from 300+ companies report +87% clarity, +58% efficiency, +50% performance,

+46% alignment and −50% negative stress after LEAD21 (participant-reported); a one-year study with 567 leaders (Prof. Beate Hanson & Prof. Mohit Bhandari, McMaster University) showed improvement in 8 of 9 climate dimensions. Dr. Charles Pellerin (NASA) on LEAD21: "Unique in the world, and a capability that nobody else has." **gerald-huesch.com/about-en**

PART 4 · OUTLOOK

STATUS OF THE SERIES

QUESTION	STATUS
Why is leadership the architectural condition of AI adoption?	answered in WP 0
How does one recognize that an organization perceives itself?	answered in WP 1
How does the organizational world model become explicit?	answered in WP 2
What governs entry into durable belief?	answered in WP 3
How is leadership climate regulated as a system?	answered in WP 4
How does one decide under cognitive overabundance?	**answered in WP 5**
How much legibility does delegation require?	open, WP 6
What does membership instead of control mean?	open, WP 7
Does LEAD21 affect AI adoption?	not substantiated — Prediction 0
Can openness be produced by a routine?	not substantiated — Prediction 1
Does a Perspective Audit produce clearer directional decisions?	not substantiated — Prediction 2
Does a recorded verdict shorten Correction Latency?	not substantiated — Prediction 3
Do target bands lower Correction Latency?	not substantiated — Prediction 4
Does a mandatory extension lower the revision rate?	**not substantiated** — Prediction 5

WHAT COMES NEXT

WP 6 · The Legible Organization. *Legibility* and *Epistemic Distance* — what it costs to understand someone well enough to evaluate them. This paper showed how decisions are made. The next shows why delegation fails on a cost quantity that is rarely named.

OPEN QUESTIONS FROM THIS PAPER

1. Can weakness be operationalized for leadership decisions without lapsing into false precision?
2. Are there situations in which the strongest commitment is the right one — and how are they recognized in advance?
3. How often may a leadership group abstain with reasons before the abstention itself becomes the pattern?

NOTES

On method. Conception, argument and responsibility rest with the author. An AI system (Claude, Anthropic) was used as a working instrument for research, structuring, the transposition of the AGI-26 contributions and the bilingual drafting, with the explicit instruction to assess the relation of LEAD21 to AGI critically and with the greatest possible objectivity, including against the author's interest. Readers are expressly invited to evaluate the contribution of LEAD21 to AGI critically, with maximum objectivity of their own, and to draw their own conclusions.

Declaration of interest. The author is the developer and owner of LEAD21 and has a commercial interest in its dissemination. Every statement about LEAD21 is assigned a degree of coverage and an evidence status.

On sources. The author evaluated the posters he documented at AGI-26 as well as titles, abstracts and the proceedings tables of contents — not the full texts. Verbatim quotations come from the posters.

REFERENCES

AGI-26

Ansah, H., Saneei, S., & Banerjee, S. (2026). Thinking fast and planning slow: Language-guided latent control with hierarchical value functions. In M. Iklé, A. Franz, A. Kemendo & H. Lowy (Eds.), *Artificial General Intelligence. AGI 2026, Part I* (LNCS 16854, pp. 1–16). Springer.

Bennett, M. T. (2026). The wrong razor. In Iklé et al. (Eds.), *AGI 2026, Part I* (LNCS 16854, pp. 48–61). Springer.

Goertzel, B. (2026). *Weakness is all you need: A generalized Occam bias for artificial general intelligence.* Conference poster, AGI-26, San Francisco State University, 27–30 July 2026. — *Poster documented first-hand; proceedings entry open.*

Ibias, A., Antona, H., Ramirez-Miranda, G., Guinovart, E., & Alarcón, E. (2026). *Unsupervised Cognition.* Conference poster, AGI-26; arXiv:2409.18624. — *Poster documented first-hand; proceedings entry open.*

Orban, D. (2026). Civilizational metamaterials: Engineering coordination under capability gradients and structural turbulence. In Iklé et al. (Eds.), *AGI 2026, Part II* (LNCS 16855, pp. 118–136). Springer. https://doi.org/10.1007/978-3-032-33195-3_10

Further works cited on the posters

Bennett, M. T. (2023). The optimal choice of hypothesis is the weakest, not the shortest. arXiv:2301.12987.