

When Governance Replaces Judgement

How optimisation, imperfect tools and cost pressure reduce Project Management to a patch function — and eliminate the only capability that survives complexity

The thought experiment: perfect tools as a revealer

Let us assume an intentionally extreme scenario. The tools are perfect. Artificial intelligence plans, coordinates, tracks, optimises and anticipates flawlessly. Dependencies are managed without friction. Risks are identified early. Reporting is instantaneous and accurate. In such a world, everything that traditionally constituted the execution of Project Management becomes a commodity.

The question is therefore no longer what a Project Management team does, but what still differentiates one team from another. If the reasoning is pushed to its logical conclusion, only one answer remains: judgement. Beyond a certain threshold, governance does not reduce risk; it reallocates it. Once execution and compliance reach baseline adequacy, additional governance has negative expected value because it systematically displaces judgement rather than improving decision quality. What makes a Project Management team the best team is not superior execution, but its aggregate judgement capacity. Everything else converges toward the same baseline.

The uncomfortable implication is that, in such a world, most Project Management roles as currently designed would no longer be justified. If execution converges and tools remove coordination friction, the majority of governance-heavy PM functions should disappear. What remains is not more structure, but fewer roles centred on judgement rather than control.

What is meant by judgement

Before going further, it is important to be precise about what is meant by judgement. Judgement is not expertise, seniority, intuition, or opinion. It is not decisiveness for its own sake, nor rhetorical confidence. Judgement is the capacity to take responsibility for a decision when the framework no longer determines the answer. It is exercised when multiple options are all formally acceptable but not equivalent in their consequences, when trade-offs cannot be optimised away, and when information is incomplete or asymmetrically distributed. Judgement implies ownership of outcomes rather than compliance with procedure.

Judgement is inherently situated. It depends on an intimate understanding of the system, the people involved, and the downstream effects of a decision. It develops through exposure to real situations and real consequences, not through instruction alone. For this reason, judgement cannot be standardised, certified, or reliably transferred at scale. Training can improve competence within a framework, but it cannot guarantee the ability to decide when the framework no longer applies.

This does not mean judgement cannot be cultivated at all. While it cannot be industrialised, organisations can influence the probability of its development by shaping exposure, feedback loops, and responsibility boundaries. Judgement grows through repeated contact with real decisions, rapid feedback on their consequences, and visible accountability. It does not emerge from curricula, simulations, or formal abstractions of decision-making.

Judgement is not only scarce; it is structurally illegible to managerial systems. Modern organisations are optimised for legibility, comparability, and auditability. Value that cannot be made visible, benchmarked, or justified *ex ante* is treated as unsafe. Governance is therefore not merely a reflex or a cognitive error, but a defensive response to the presence of opaque value creation that resists formalisation.

In this context, governance can be defined very simply. Governance is what an organisation relies on when it does not want to rely on individual judgement. It replaces individual decision-making with shared rules, predefined approvals and escalation paths. Instead of asking someone to decide under uncertainty, governance asks them to follow a process. Instead of ownership, it distributes responsibility across structures. Governance does not eliminate decisions; it delays them, fragments them, or pushes them upward until they become formally defensible rather than contextually appropriate.

Judgement as a scarce organisational capacity

An important implication follows from this definition. Judgement is a scarce organisational capacity. It is not evenly distributed, and it is not commoditisable in the way technical skills are. While experience and feedback can refine judgement over time, it cannot be produced on demand or rolled out uniformly. Training can improve competence within a framework; it cannot guarantee the ability to decide when the framework no longer applies.

Organisations that treat judgement as universally developable therefore tend to overestimate their collective decision-making capacity, while simultaneously constraining the few individuals who actually possess it.

Where judgement becomes visible

Judgement appears most clearly where even perfect tools reach their limits. It emerges when several options are all compliant with standards. When rules conflict. When strict application of the framework worsens the situation. When decisions are irreversible and uncertainty is irreducible. When escalation is procedurally correct but strategically harmful. Judgement is what allows a decision to be taken when optimisation no longer helps.

Judgement also includes a temporal dimension. It is not only the ability to choose well, but to choose in time. In complex systems, the cost of delayed decision-making often exceeds the cost of an imperfect decision. A theoretically superior judgement that arrives too late is an execution failure.

Project Management as a Patch Layer

In reality, of course, tools are not perfect. Most organisations operate with fragmented systems, partial integrations and misaligned processes. In such environments, Project Management progressively assumes a patch function. Rather than managing complexity, it compensates for it, bridging gaps between systems that were never designed to operate together under current clinical, regulatory and operational conditions.

This patch function is not accidental. As legacy platforms coexist with newer tools, partial automations and evolving data flows, misalignment accumulates across generations of systems. The resulting friction does not eliminate sophistication. Organisations often operate many advanced tools at the same time. The problem is that these tools were designed at different moments, for different operating models, and no longer fit together cleanly. What this produces is not robustness, but fragmentation: complexity spread across disconnected systems that require continuous human intervention to hold together. Project Management becomes the connective tissue between tools that cannot fully integrate, absorbing inconsistency through manual coordination, exception handling and continuous intervention.

This work is expensive, but its cost is systematically misattributed. It is recorded as the cost of Project Management rather than as the cost of tool misalignment, technical debt and architectural inertia. As a result, Project Management is perceived as a cost centre, even when it is merely compensating for structural deficiencies elsewhere in the system.

The organisational response is predictable. Cost pressure triggers juniorisation and increased governance. Decision authority is narrowed, controls are multiplied, and discretion is constrained. Ironically, this prevents Project Management from exercising the very capability that could differentiate it: judgement. As judgement is displaced by process, Project Management delivers less strategic value and is perceived as even more expendable.

This creates a self-reinforcing loop. The more Project Management is used as a patch layer, the more its cost is scrutinised. The more its cost is scrutinised, the more governance is imposed. And the more governance is imposed, the less judgement can be exercised. The function becomes trapped in operational firefighting, consumed by compensating for broken systems rather than making decisions under uncertainty.

Every hour of judgement spent patching misaligned tools is an hour stolen from the only activity that does not commoditise. Patching may preserve short-term operability, but it is value-destroying at the system level. It conceals technical debt rather than pricing it, creates the illusion of resilience while increasing fragility, and diverts scarce judgement away from where it is most needed.

This does not reduce the importance of judgement; it increases it. In imperfect systems, judgement is required not only to decide under uncertainty, but also to prevent structural deficiencies from dictating outcomes. The risk is that judgement becomes exhausted by compensation rather than exercised as decision-making. When this happens, organisations mistake workaround effort for value creation, and judgement is quietly depleted.

The governance reflex and the category error

This is why the thought experiment of perfect tools is revealing. It strips away the noise and exposes what Project Management would still be valued for if systems no longer required human patching. It clarifies the distinction between judgement as a compensatory mechanism and judgement as a strategic capability.

Faced with this reality, organisations almost systematically make a cardinal mistake. They attempt to reintroduce judgement through governance. When a lack of decision-making is observed, the reflex response is to add structure: new roles, new committees, formalised escalation paths, refined arbitration frameworks. Standardisation and control are presented as ways to raise decision quality and reduce risk.

This reaction is understandable, but it rests on a fundamental category error. Standardisation and control do not enhance judgement; they constrain it. They do not increase an organisation's capacity to decide under uncertainty; they reduce the space in which judgement can be exercised. What is framed as an effort to improve judgement is, in practice, an effort to make judgement unnecessary.

More importantly, governance mechanisms consume the very resource they claim to protect. Each additional layer of control introduces new tasks: documentation, justification, alignment, validation, pre-approval. These tasks do not improve judgement. They absorb the time, attention and cognitive bandwidth of those who possess it. The organisation does not create more judgement; it spreads less of it more thinly.

This reflex is also rational at the individual level. Governance changes who carries risk. Decision-making risk is pushed upward, while responsibility for following process is pushed downward. Senior layers are shielded from blame because decisions appear collectively validated and procedurally correct. Execution layers, in turn, are protected as long as they comply, even if outcomes deteriorate. The result is career-safe ambiguity at the top and procedural overload at the bottom. What damages the organisation often makes sense for individuals inside it.

Why judgement reduces escalation and improves client experience

The paradox becomes visible in day-to-day operations. When judgement is present and exercised early, many situations never become escalations. A client request that sits slightly outside standard process is assessed locally, in context. Trade-offs are made explicitly. A decision is taken, owned, and executed. From the client's perspective, the organisation appears responsive, coherent and aligned with intent rather than constrained by form.

When judgement is constrained by governance, the same situation follows a different path. Because no one feels authorised to decide, the issue is reframed as a compliance question. It is escalated "to be safe". Time is spent preparing justifications rather than choosing a direction. The client experiences delay, fragmentation and contradictory signals. What began as a manageable deviation becomes a visible organisational problem.

More judgement therefore results in fewer escalations not because risk disappears, but because it is absorbed earlier, at the right level, by people capable of weighing consequences.

The equity and upskilling illusion

Organisations often misread the opposite pattern. As judgement becomes slower and rarer under governance load, additional controls are introduced to compensate. This creates a self-reinforcing loop in which the mechanisms designed to improve decision quality systematically erode the organisation's aggregate judgement capacity.

The mistake many organisations make is to ignore this asymmetry, either in the name of equity or through the belief that judgement can be uniformly acquired through training and exposure. In both cases, the organisation designs as if judgement were abundant and scalable, when in reality it is scarce and unevenly distributed.

What this means for Clinical Project Management

While the Central Lab provides a particularly concentrated illustration of these dynamics, the underlying mechanisms apply broadly across clinical development. In this domain, they are not edge cases but structural conditions. No protocol unfolds exactly as designed. Site behaviour varies, enrolment deviates, samples arrive late or incomplete, regulatory interpretations shift, logistics fail under real-world constraints, and client priorities evolve mid-study. What is formally labelled an “exception” is, in practice, the operating environment.

In this context, judgement is not a luxury layered on top of a stable system. It is a condition of survival for the study. Without judgement, deviations accumulate, escalation becomes constant, and the organisation oscillates between rigid compliance and reactive firefighting.

Central Lab Project Management therefore operates permanently in the space between protocol intent and operational reality. Tools, SOPs and governance define boundaries, but they do not resolve the countless micro-arbitrages required to keep a study viable.

The more a client operates outside standard processes, the more the value of judgement increases. Complex portfolios, accelerated timelines, evolving scientific hypotheses or unconventional operational constraints simultaneously increase exposure on three critical dimensions: patient safety, data integrity and delivery timelines.

In such contexts, stricter procedural compliance does not reduce risk proportionally. What reduces risk is the ability to judge which deviations threaten patient safety, which compromise data integrity, and which delays will cascade irreversibly downstream. As client complexity increases, compliance delivers diminishing returns, while judgement delivers increasing ones.

Clinical development is not unique in this regard. It is simply more honest about uncertainty. The same dynamics apply wherever decisions are irreversible, information is incomplete, and delay

carries its own form of risk: investment committees, crisis response, military command, product leadership, and regulatory strategy.

Peer recognition as the hidden core of the system

At the centre of this model lies a mechanism that deliberately escapes traditional HR grids: peer recognition among individuals capable of judgement. More simply, it reflects a structural truth: it takes judgement to recognise judgement in others.

This recognition is based not on outcomes, but on reasoning quality. Those with judgement recognise whether the right tensions were identified, whether weak signals were prioritised, whether second-order effects were anticipated. Like chess players analysing a game, mastery is recognised through depth of anticipation, not by the final result.

A second signal is ownership. Individuals capable of judgement recognise those who do not hide behind rules but actively assume responsibility under uncertainty. This is perceptible in posture, language and intent.

This mechanism is powerful, but unstable. Without formal safeguards, it can drift into confirmation bias or homogeneity of thought. There is no clean way to regulate this without destroying it. The only stabilising force is plurality: overlapping circles of judgement, exposure to disagreement, and the absence of a single authority defining what good judgement is.

The final risk is invisibility. Because this system is deliberately unnamed, it is fragile. If those who possess judgement leave, they take with them the ability to recognise the next generation. What remains are structures optimised for compliance rather than discernment. The alternative to peer recognition is not fairness, but mediocrity enforced by process.

The Cost and Limits of Judgement

Judgement cannot be evaluated safely by outcomes alone, nor can it be audited without being neutralised. An organisation that relies on judgement must accept that some good decisions will produce bad outcomes. The distinction between a reasoned error and poor judgement cannot be determined in advance by rules, nor assessed after the fact using checklists. It can only be assessed retrospectively, by peers who themselves possess judgement, through scrutiny of how the decision space was framed, which trade-offs were made explicit, and whether responsibility was assumed before the outcome was known.

Judgement also requires exposure to real consequences. When cost pressure leads to systematic juniorisation, organisations create a judgement desert: individuals are trained in compliance under low stakes and later expected to lead under high stakes without having observed judgement in action. The solution is not equal exposure to decision authority, but asymmetric proximity to judgement: observation, shadowing, and honest post-hoc analysis of real decisions, rather than premature delegation disguised as empowerment.

Finally, judgement carries an irreducible political cost. There is no version of judgement that is simultaneously fair, scalable and governable. An organisation that wants judgement must accept structural inequality in decision authority and evaluation, tolerate ambiguity in accountability, and resist the instinct to normalise what is, by nature, unevenly distributed. Any attempt to make judgement fully legible, auditable or equitable ultimately replaces it with governance. There is no judgement framework, no maturity model, no KPI set, and no training programme that can substitute for exposure to real decisions under consequence. Any attempt to do so merely replaces judgement with its simulation.

Why this feels wrong, and why it still holds

This approach may feel uncomfortable because it is not equitable in a formal sense, because it does not appear scalable, and because it resists the natural reflex to treat governance as the solution. That reflex is understandable, but in this case it is precisely what makes the problem worse rather than better.

What is often missed is that Project Management operates across two distinct economic logics. One is the cost of compensating for misaligned, imperfect systems — a cost rooted in tools, technical debt and organisational fragmentation, but routinely attributed to Project Management. The other is the value of judgement — the only contribution that does not commoditise once execution itself converges. When these two are collapsed into a single cost category, organisations end up optimising away judgement in an attempt to reduce costs that originate elsewhere.

This is not a call for inaction. It is not an argument for abandoning structure, tools or responsibility, but for recognising where structure stops helping and starts displacing judgement. It is a call for restraint and for clearer attribution. Organisations cannot deliberately produce judgement, but they can stop pricing it as overhead. They can protect the teams that exercise it, while actively removing the conditions that suffocate it: excessive ex-ante control, punitive treatment of reasoned error, dilution of responsibility, and artificial symmetry in decision authority — while addressing patch costs where they actually belong.

Judgement is not a function.

It is not a process.

It is a flow.

And it survives only as long as organisations resist the temptation to govern it — or to account for it as if it were merely another cost centre.