

The Right to Contest

Why Challenging an AI Decision Requires Human Authority, Not Human Presence

Dr M Maruf Hossain, PhD, GAICD

Chief AI Strategist · 42 Consulting.AI

enquire@42consulting.ai · www.42consulting.ai



CONTENTS

Table of Contents

From the Author.....5

Executive Summary.....7

A Verdict, or a Proposal?9

Authority, Not Presence.11

What a Real Review Requires.13

Built to Be Contested.15

Not a Right You Can Sign Away.17

Kept in Human Hands.19

Five Questions for Your Board.....21

42 Consulting.AI.....23

References and Sources.....25

FOREWORD

From the Author.

When an AI system decides something about a person — approves or refuses, flags or clears, ranks or excludes — the person on the receiving end has one question that matters more than how the decision was reached: can I challenge it? A decision you cannot challenge is not a decision made about you; it is a verdict handed down to you. The difference between the two is the right to contest.

People are owed an explanation of the AI decisions that affect them. But an explanation you cannot act on is only information. The right that gives it force is the right to contest — to say “I disagree, look again”, and to have a person with real authority actually look again. Transparency without contestation tells you why the door was closed; it does not let you knock.

This paper is about what that right actually requires, and about the way it most often fails — not by being denied outright, but by being met in form while hollowed out in substance. A human is placed “in the loop”; a review process exists on paper, and yet the machine’s decision stands every time, because the human was present but never had the authority, the information, or the standing to overturn it. Presence is not authority, and a review that cannot say no is not a contestation.

It is written for the board and the executive who have been told their AI decisions are “subject to human review” and want to know whether that review is real. Can the person who reviews a contested decision actually overturn it? Have they seen enough, and been given enough time and independence, to reach their own conclusion? If the honest answer is that the human almost always agrees with the machine, the organisation has built the appearance of contestation and not the substance.

M Maruf Hossain

SECTION ONE

Executive Summary.

The right to contest an AI decision is what turns accountability from a claim into a fact. An organisation can be transparent, fair, and well-governed and still leave the people it decides about with no way to challenge what was done to them. This whitepaper establishes three propositions.

First, contestation is what makes the other rights real. An explanation a person cannot challenge is information, not accountability; a fair process a person cannot question is fairness they must take on trust. The right to contest is the mechanism that lets an individual convert “here is why” into “then look again” — and without it, every other assurance about an AI decision is something the organisation grants itself, not something the affected person can enforce.

Second, contestation requires human authority, not human presence. Putting a human “in the loop” is not the same as giving someone the power to overturn the machine. A real contestation is a substantive review by a person who examines the decision afresh, owes no deference to the AI’s output, and has the standing and the time to reach an independent conclusion — including the conclusion that the system was wrong. A review that cannot, in practice, overturn the decision is a rubber stamp wearing the costume of accountability.

Third, the right cannot be waived, gated, or bolted on. It cannot be signed away in terms and conditions because it exists by virtue of the organisation’s constitutional obligations, not its contracts. It cannot be conditioned on the individual first proving harm or understanding the decision. And it cannot be added after the fact — a system that cannot support a genuine review of its decisions should not have been deployed to make them. Contestation is a precondition of deployment, not a feature request.

This paper sets out what the right to contest reframes, distinguishes human authority from human presence, specifies what a real review requires, shows why the capability has to be built in before deployment, and explains why the right cannot be signed away. It ends with five questions a board should be able to answer about whether the people it decides about can actually push back — not next quarter, but today.

SECTION TWO

A Verdict, or a Proposal?

Every AI decision that affects a person arrives as one of two things, and which one depends entirely on whether it can be contested. If there is no genuine way to challenge it, the decision is a verdict: it stands because no one can undo it, and the person it was made about has to live with it, whether or not it was right. If there is a real right to contest, the same decision is a proposal: an outcome the organisation is prepared to act on, but one that a person can require to be reconsidered by someone with the authority to reach a different conclusion.

The distinction is not rhetorical. A verdict and a proposal can be identical in content and produced by the same system; what separates them is what happens when the affected person disagrees. Under a verdict, disagreement changes nothing — there is no route from “I think this is wrong” to a decision being reconsidered. Under the right to contest, disagreement triggers an obligation: a valid contestation requires a mandatory human review, and that reviewer has the authority to uphold, modify, or overturn the machine’s output.




Figure 1 — A verdict, or a proposal. Without a genuine right to contest, an AI decision is imposed and final; with it, the same decision becomes a proposal a person can require to be reviewed — and overturned — by someone with the authority to do so.

This is why the right to contest is the point at which control of a decision returns to a human. Between the moment a system decides and the moment its decision becomes final, the right to contest inserts a person who can intervene — not to supervise every decision, which is impossible at scale, but to be available to any individual who challenges the one that affected them. The system decides by default; a human decides on a challenge. That is what keeps the machine’s authority provisional rather than absolute.

It also reframes what the organisation is doing when it deploys the system. An organisation that treats its AI outputs as verdicts has, in effect, delegated final authority over people's outcomes to a model. An organisation that treats them as proposals has kept final authority in human hands and used the model to do the work of getting there. The right to contest is how that second posture is made real and testable, one challenged decision at a time — rather than asserted in a policy no affected person can invoke.

The rest of this paper is about what it takes for that second posture to be genuine. Because the failure mode is not usually a flat refusal to allow contestation — few organisations would defend that. The failure mode is a contestation process that exists but does not work: a human review is present in the process map but absent in practice. The next section is about that gap and the distinction that closes it.

SECTION THREE

Authority, *Not Presence.*

The most common way the right to contest fails is not denial but dilution: the organisation puts a human “in the loop” and treats that as the end of the matter. But a human in the loop is a position on an org chart, not a guarantee of anything. What the right to contest requires is not that a human is present when a decision is reviewed, but that the human has the authority to change it — and those are very different things (Hossain, 2026).

Human presence without authority takes a familiar shape. The reviewer sees only what the system chose to surface, and so reasons from the machine’s framing. They face a queue and a clock, and the fastest path through both is to agree. They sense that overturning the system is treated as friction, and that a high override rate will be questioned before a high agreement rate is. So they defer — not because they are told to, but because everything around the review is arranged to make deference the path of least resistance. The human is present. The authority is not.

Human authority is the opposite arrangement. The reviewer examines the decision afresh, with access to the underlying facts and not merely the system’s summary of them. They owe the machine no deference — its output is an input to their judgment, not a default they must justify departing from. They have the time, the competence, and the standing to reach an independent conclusion, and overturning the system when it is wrong is treated as the review working, not failing. The test of whether authority is real is simple: can this person, in practice, say no to the machine — and does it sometimes happen?

This is the distinction the framework is built to enforce, because it is the one that quietly collapses under operational pressure. It is easy to require that decisions be “subject to human review” and hard to ensure that the review is substantive. So the framework does not stop at requiring a human in the loop. It requires that the human’s review be independent of the original decision, informed by the actual decision record, and backed by genuine authority to overturn — and it treats a review that lacks any of these as a failure of the right to contest, not a discharge of it.

SECTION FOUR

What a Real *Review Requires.*

If a contestation is only real when the review is substantive, then “we have a review process” is not enough — the process must meet specific conditions, and a review that satisfies fewer than all of them yields procedural compliance without substantive accountability. The framework specifies what a genuine contestation review requires. Four conditions are essential.

Independence from the original decision. The review must be conducted by someone other than the person or process that made or approved the original decision, and it must not simply re-run the system that produced it. A review that consults the same model, or the same official, is not a fresh look — it is the first decision defending itself. Independence is what makes the second look capable of reaching a different answer.

Access to the actual decision record. The reviewer must work from the specific facts of the contested decision — the principal factors that drove it, the inputs it relied on, any human oversight that applied — not from a general account of how the system behaves. A review of the system in the abstract, rather than of the decision in particular, cannot find the error in this case. Access to the record is what makes the review substantive rather than notional.

Authority to overturn, and the standing to use it. The reviewer must be able to uphold, modify, or overturn the decision, and that authority must be real: resourced with enough time, protected from penalty for overturning, and exercised often enough to show it is not theoretical. A reviewer who can technically overturn but never does, or is quietly discouraged from doing so, has authority on paper and presence in fact.

These conditions are cumulative, not a menu. A review that is independent but works from a summary cannot find the error; one that has the record but no authority cannot fix it; one that has authority but no time will not use it. The right to contest is discharged only when all of them hold together — and the timeframe within which the review completes matters too, because a decision overturned long after it has already caused harm is a remedy owed, not a contestation honoured. Specifying the conditions is what makes “we reviewed it” answerable to what a real review actually is.

The four conditions are set out below. They are cumulative rather than a menu: a review satisfying three of them produces procedural compliance without substantive accountability.

Condition	What it requires	What its absence produces
Independence	A reviewer other than the person or process that made or approved the decision, not re-running the same system	The first decision defending itself, presented as a second look
Access to the record	The principal factors, the inputs relied on, and any human oversight that applied to this decision	A review of the system in the abstract, which cannot find the error in this case
Authority to overturn	Standing to uphold, modify or overturn, resourced with time and protected from penalty for overturning	Authority on paper, exercised rarely enough to show it is theoretical
A timeframe that matters	Completion while the outcome can still be changed for the person affected	A decision overturned long after the harm has already been taken

Table 1 — The four cumulative conditions of a genuine contestation review.

Reading the third column is the fastest way to audit an existing process. Each failure mode is one most organisations would still describe as having a review in place.

SECTION FIVE

Built to *Be Contested.*

A genuine capability for contestation cannot be improvised after a challenge arises. If the decision record was never captured, there is nothing for a reviewer to examine; if no independent reviewer exists, there is no one to conduct the review; if the individual has no way to reach the process, the right is theoretical. So the ability to be contested is not a customer-service function bolted on after launch. It is a property the system and the process around it must be built to have — which makes it a decision taken before deployment, not after the first complaint.

The system must preserve what a review will need. The specific decision record — the factors, inputs, and any human oversight — must be captured and retained at the moment of decision, because a review cannot reconstruct what the system never stored. A system that does not preserve its own decisions in a reviewable form cannot support contestation, however willing the organisation is to conduct a contestation.

The process must be reachable, and the reviewer must exist. There must be open channels through which any affected individual can lodge a challenge — and the organisation may not narrow them to a single hard-to-find form; a complaint made through any reasonable route is a valid contestation. And an independent reviewer with authority to overturn must actually be resourced and available, not a role that exists only in the process diagram.




Figure 2 — The contestation gate. A system that decides about people is fit for deployment only if it can genuinely be contested; one that cannot support a real review of its decisions is not fit to make them.

Contestability is verified before deployment, not assumed. No system reaches production for decisions affecting individuals without demonstrating that it can be contested — that the record is retrievable, the channels are open, the reviewer is available, and the authority to overturn is

real. A system that cannot pass this gate is not ready to make decisions about people, because a decision that cannot be challenged should not be made about a person.

The capability is maintained, not left to decay. Contestability degrades if it is not protected: a logging change can break record retrieval, a reorganisation can remove the independent reviewer, a process “improvement” can quietly close a channel. So the capability is monitored in production, and anything that breaks it is treated as a fault that suspends the system’s authority to decide, not an inconvenience to be tolerated.

SECTION SIX

Not a Right *You Can Sign Away.*

Everything so far depends on one structural commitment: that the right to contest belongs to the individual and cannot be removed by the organisation that would find it convenient to remove. This is where contestation regimes most often quietly fail — not by refusing challenges, but by narrowing the right until it is easy to satisfy and hard to use.

The narrowing takes familiar forms. A clause in the terms and conditions purports to make the AI decision “final”. The process requires the individual to first prove they were harmed or first understand the decision before they may contest it. The window to challenge is made short, the channel obscure, the burden high. Each step is defensible in isolation; together they leave a right that exists on paper and is unreachable in practice — which is indistinguishable, for the person affected, from having no right at all.

The framework forecloses this by placing the right beyond the organisation’s reach. The right to contest exists by virtue of the constitutional obligation the board has adopted, not by virtue of any contract — so a clause purporting to waive or limit it is simply unenforceable, and its presence in the organisation’s own documents is a governance concern in itself (Hossain, 2026). The right is not conditional on proving harm or on understanding the decision first; a person may contest a decision they do not yet understand. And no operational convenience — volume, cost, delay — is a ground for narrowing it.

So the right to contest is not a service level that the organisation sets and can revise. It is a constraint the organisation operates under, owed to the people it decides about and enforceable by them. An organisation that has narrowed the right until almost no one can use it has met the form of contestation and defeated its purpose — and it has done so while able to claim, truthfully but misleadingly, that its decisions are “subject to challenge”. Placing the right beyond the contract is what gives that claim meaning.

SECTION SEVEN

Kept in *Human Hands.*

A right that depends on the organisation choosing to honour it, decision by decision, is only as strong as its least convenient moment. The framework therefore does not leave the right to contest to good intentions. It wires contestation into the same machinery that governs the rest of the AI estate, so that the capability is built, the reviews are recorded, and the right cannot quietly lapse.

The line runs from constitution to evidence. The individual's right to contest an AI decision, and to a substantive human review of it, is adopted at board level as a governing principle; the framework defines what a genuine review requires; the deployment gate requires contestability to be demonstrated before a system ships; and the register that tracks each system carries its contestation capability and status. Conducting a review, and maintaining the capacity to conduct one, are recorded acts with owners — not favours granted when someone complains loudly enough.

Crucially, contestation attaches to the existing enforcement spine rather than sitting beside it as goodwill. The gate that clears a system for deployment is extended to require demonstrated contestability. The register that records a system's authority determines whether its decisions can genuinely be challenged and overturned. The people who supervise operational systems are briefed to treat a broken contestation channel as a fault that suspends the system's authority, not as a backlog to clear later. The right becomes a precondition that the system must satisfy and continue to satisfy, not a promise in a policy.

The effect is that contestation becomes an artefact rather than an assurance. When a regulator asks whether a person could challenge the decision that affected them, the answer is the review that was conducted, its independence, and its outcome — not a statement that decisions are “subject to review”. When asked whether the review could have overturned the machine, the answer is the record of reviews that did. A right that can be shown, review by review, is one that was real; one that can only be asserted was always just presence wearing the name of authority.

None of this removes the machine from the work — the system still makes the decisions at the scale only a system can. What the architecture provides is a guarantee that the machine's authority stays provisional: that behind every decision it makes about a person stands a human who can be reached, who can look again, and who can say no. That is what it means to keep final authority in human hands — not to slow the machine down, but to ensure that the last word, when a person asks for it, is still a human's.

SECTION EIGHT

Five Questions *for Your Board.*

The measure of a contestation regime is not whether a review process exists. It is about whether the people the organisation makes decisions about can actually challenge those decisions and have them overturned. Each of the following is a question a regulator, a court, or an affected individual could put to the board — and each should be answerable today, with an artefact rather than an assurance.

Can the person actually challenge the decision — and reach someone who can change it? For each consequential system, can an affected individual lodge a contestation through a reachable channel and trigger a mandatory human review — or is the right buried, gated, or purely nominal?

Does the reviewer have authority, or only presence? Is the human review independent of the original decision, informed by the actual decision record, and backed by real authority to overturn — or is it a human in the loop who, in practice, agrees with the machine almost every time?

Do reviews ever actually overturn the system? Can you show cases where a contestation review changed or reversed an AI decision — evidence that the authority to overturn is real — or is the override rate effectively zero?

Have we tried to sign right away? Do your terms, policies, or processes purport to make AI decisions final, require proof of harm before a challenge, or otherwise narrow the right to contest — clauses that are unenforceable and are themselves a governance concern?

Was contestability built in before deployment? For systems already deciding people's outcomes, was the ability to be genuinely contested demonstrated before they shipped — or could a person challenge a decision today and find there is no record to review and no one with the authority to change it? A framework that cannot answer these questions has not kept authority in human hands. It has handed it to the machine and hoped no one would ask.

ABOUT

42 Consulting.AI.

42 Consulting.AI builds constitutional AI governance for regulated enterprises. Its core work, the MANDATE Suite, is a comprehensive governance architecture comprising 64 normative instruments across five frameworks — governing the delegation of machine decision authority, the ethical and rights boundaries on its use, the management system that maintains it, and the regulatory extensions for EU and group structures. The Suite is designed for organisations in financial services, insurance, superannuation, healthcare, government, and other regulated sectors operating autonomous AI systems at scale.

The architecture is supported by an academic preprint series published on Zenodo, which establishes the theoretical basis for each governance mechanism, and by the forthcoming book *Governance-First AI*. The Suite's distinguishing conviction is that the AI governance that regulated enterprises require cannot be delivered by a compliance framework alone — it requires a constitutional architecture, adopted by the board, enforced by the infrastructure, and continuously tested against current conditions.

Dr M Maruf Hossain, PhD, GAICD · Chief AI Strategist

enquire@42consulting.ai · www.42consulting.ai

The preprint series is available at: zenodo.org/communities/mandate/records

© 2026 42 Consulting.AI · All rights reserved · This whitepaper is general commentary and is not legal, financial, or compliance advice. Organisations should obtain their own professional advice on the application of the referenced regulatory instruments.

REFERENCES

References

and Sources.

The regulatory instruments referenced in this paper were verified at the primary source. The governance mechanisms are set out in full in the MANDATE preprint series, published through Zenodo, and in the forthcoming book Governance-First AI. The full MANDATE preprint series is available at: zenodo.org/communities/mandate/records

REGULATORY AND INSTITUTIONAL SOURCES

European Union (2024). Regulation (EU) 2024/1689 (Artificial Intelligence Act), Article 14 — human oversight, including the ability to intervene in or override the operation of a high-risk AI system.

ASIC (2015). Regulatory Guide 271: Internal Dispute Resolution — RG 271 obligations for accountability, review, and human reconsideration of automated decisions affecting consumers.

Office of the Australian Information Commissioner. Privacy Act 1988 (Cth), Australian Privacy Principles 1, 12 and 13 — transparency, access to, and correction of personal information underpinning automated decisions.

Commonwealth of Australia. Corporations Act 2001 (Cth) s.912A — the obligation to provide financial services efficiently, honestly and fairly, bearing on the fair reconsideration of automated decisions.

United Nations (2011). Guiding Principles on Business and Human Rights — access to effective remedy for individuals affected by business decisions, including automated ones.

BOOK

Hossain, M Maruf and Ponnusamy, Ahilan (2026). Governance-First AI. Apress (forthcoming) — the primary public articulation of the responsible AI governance doctrine.

MANDATE PREPRINT SERIES

Hossain, M Maruf (2026). The Individual Rights Response Obligation: Why the Right to Contestation in Enterprise AI Governance Requires Human Authority, Not Human Presence. Preprint. Zenodo. <https://doi.org/10.5281/zenodo.20731491>.



Consulting.AI

Scaling Intelligence

enquire@42consulting.ai · www.42consulting.ai