ACCESS TO JUSTICE

AI and Access to Justice in Arbitration

Stephen Rayment · Founder, CEO & Group Managing Director, Systech MDP

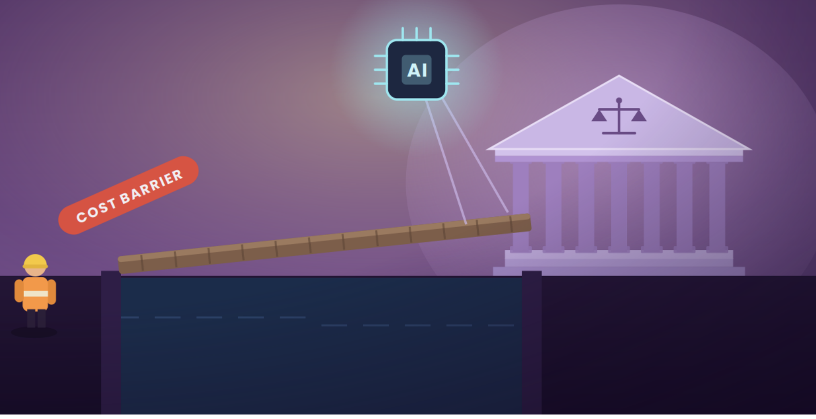

AI's cost savings can bring arbitration within reach of specialist subcontractors. Unreasonable objections to AI defend the cost barrier that has always favoured the respondent.

Arbitration was meant to be the accessible alternative: faster, more flexible, more commercial than the courts. For specialist subcontractors, that promise has always come with an asterisk. The process is expensive, the costs are front-loaded, and the claimant is usually the party least able to bear them, a subcontractor already squeezed by the very non-payment it is trying to recover. Cost, not merit, is the gatekeeper. It decides which claims are ever brought.

That is exactly why the arrival of AI in arbitration matters so much, and why the way the other side reacts to it deserves close attention. Used properly, AI is one of the few developments in a generation that genuinely **lowers the drawbridge** for smaller claimants. And a respondent who fights hard to keep that drawbridge up, through blanket or tactical objections to the claimant's use of AI, is, whether it admits it or not, defending the cost barrier that has always worked in its favour.

AI changes the economics of bringing a claim

The expensive parts of an arbitration are the labour-intensive ones: legal research, document review, bundle preparation, translation, transcription, and the analytical grind behind quantum and delay expert work. These are precisely the tasks that price a modest claim out of the process, and precisely the tasks AI is good at compressing.

The numbers are no longer speculative. The AAA-ICDR, which launched an AI-assisted arbitrator for documents-only construction cases in late 2025, reports that in those cases the tool can **cut total costs by around 35% or more** and reduce resolution times by roughly 20 to 25%, with its President projecting cost reductions of 30 to 50% as the technology matures. Independent commentary on access to justice has long identified the cost of arbitration as one of the main inhibitors to bringing a claim at all, and ADR combined with technology is increasingly seen as a way to narrow the justice gap rather than widen it.

For a specialist subcontractor, a one-third reduction in the cost of pursuing a legitimate claim is not a marginal efficiency. It is frequently **the difference between arbitrating and writing the debt off**.

Why this is an equality-of-arms point, not just a cost point

In most construction disputes the resource asymmetry is stark. The main contractor is larger, better capitalised and able to absorb legal spend that the subcontractor cannot. That imbalance has real consequences: it lets the stronger party run up the cost of a dispute knowing the weaker party may be forced to settle cheap or walk away.

AI partially levels that field. It lets a smaller claimant produce first-class research, organise a large documentary record and support its expert analysis at a fraction of the traditional cost. In other words, it does something the arbitration system has always struggled to deliver on its own: it brings genuine **equality of arms** a little closer. Anything that narrows that gap is, by definition, an access-to-justice gain.

THE GATEKEEPER

Cost, not merit, is the gatekeeper. It decides which claims are ever brought.

The new front in an old tactic

Readers who have followed the security-for-costs debate will recognise the pattern. A security-for-costs application can be a legitimate protection for a respondent, or it can be a lever to choke off a cash-constrained subcontractor's claim by demanding money the claimant cannot easily find. The tactic works precisely because the subcontractor is already stretched.

Objecting to a claimant's use of AI can operate in exactly the same way. Strip away the language of "integrity" and "due process" and some objections are, in substance, an attempt to **re-load cost and friction onto the weaker party**, to take away the very tools that made the claim affordable. The tell-tale forms are familiar:

- **Blanket objections.** Demands that the claimant be barred from using AI at all, untethered to any specific, demonstrated problem.
- **Disproportionate disclosure demands.** Fishing requests about every tool, prompt and process, designed less to protect the record than to generate satellite litigation and expense.
- **Weaponising a real risk.** Taking the genuine problem of AI "hallucinations" and inflating it into a general attack on a claimant who has, in fact, verified its material.

Each of these dresses up a cost tactic as a principled concern. Each, if it succeeds, restores the cost barrier that AI had just lowered.

Where the legitimate line actually is

None of this means objections to AI are illegitimate. They are often essential, and the difference between a proper objection and a tactical one is not hard to state.

A legitimate objection is **specific and evidenced**. If an opponent has put fabricated authorities before the tribunal, that must be raised: the duty to verify AI output against authoritative sources is now firmly established, and the English High Court's 2025 rulings on fake AI citations show how seriously it is taken. If an expert has fed privileged or confidential project data into an open, third-party tool, that is a real confidentiality breach. If an expert has effectively outsourced the opinion itself to a model, that goes to the independence and reliability of the evidence. These objections protect the process and should be pressed.

An unreasonable objection is the opposite: a blanket prohibition, a disclosure fishing expedition, or an objection that cannot point to any actual prejudice and exists only to raise the other side's costs. The leading guidance already draws this line. Both the **CIArb Guideline (2025)** and the **SVAMC Guidelines** treat disclosure of AI use as a proportionate, case-by-case question tied to the integrity of the process, not an automatic entitlement, and caution tribunals against over-regulating the parties' private use of efficient tools. A party demanding a blanket ban is asking the tribunal to do the very thing the guidance counsels against.

What tribunals can, and should, do

Tribunals hold the answer, and they have the tools. When an objection to AI is raised, the fair and efficient response is to make the objector do the work: identify the specific prejudice, the specific tool or output, and the specific risk to the record. An objection that survives that test deserves a remedy. An objection that cannot meet it should be seen for what it is.

Two further points keep the incentives honest. First, proportionality and equality of arms cut against allowing a better-resourced party to strip a weaker one of cost-saving tools absent real harm. Second, and this echoes the costs argument familiar from security-for-costs disputes, a respondent who runs unreasonable satellite battles over AI is generating cost that a tribunal can later place at its own door. In a costs regime where conduct matters, the party that manufactures expense risks paying for it, potentially on the indemnity basis. **The tactic can be an own goal.**

THE OWN GOAL

In a costs regime where conduct matters, the party that manufactures expense risks paying for it.

The honest counter-argument

To be fair to the other side, not every concern about AI is a tactic, and a credible case for caution exists.

The hallucination problem is real: general-purpose models still fabricate citations, and unverified reliance has already cost lawyers wasted-costs orders, regulatory referrals and sanctions on both sides of the Atlantic. Confidentiality risk is real where sensitive material meets open tools. And there is a genuine academic debate about whether AI will truly democratise dispute resolution or instead entrench a two-tier system in which the well-resourced get the best tools, and a parallel argument that AI will not automatically make legal services cheaper unless the savings are actually passed on. These are serious points, and a thoughtful tribunal should weigh them.

But the answer to those concerns is **discipline, not prohibition**. Verify every output against primary sources; keep confidential data out of open tools; disclose where use affects the integrity of the process; keep the human, lawyer or expert, accountable for the result. That is what the guidance already requires, and it neutralises the integrity and confidentiality risks without throwing away the access-to-justice gain. The risk of a two-tier system is, if anything, an argument for getting these tools into the hands of smaller claimants, not against it.

The bottom line

Cost has always been the quiet filter that decides whose claims get heard in arbitration. AI is one of the first developments capable of loosening that filter, bringing arbitration genuinely within reach of the specialist subcontractor who would otherwise have no realistic route to recover what it is owed.

Objections that point to a real, evidenced problem protect the integrity of the process and should always be available. Objections that exist to rebuild the cost barrier do the opposite. The arbitration community, and tribunals in particular, should be able to tell the two apart, and **should not let a tool that widens access to justice be quietly objected out of reach.**

Where do you draw the line between a principled objection to AI and a tactical one? I'd be interested in others' experience.

This article is written from the perspective of claimant-side specialist subcontractors in commercial and construction arbitration, principally in the UK and the US. It expresses a viewpoint and is general commentary, not legal advice; AI law and guidance in this area are developing quickly, and any specific situation turns on its facts, the seat of the arbitration, the governing rules and the applicable professional codes.