Issue no. 02 October 2026

MED-MID ARBITRATION AND MEDIATION REPORTS

A tribute is paid to our oldest legal writing
The Code of Hammurabi
(the Babylonian law code which dates back to about 1772 BC)

COUR EUROPEENNE D'ARBITRAGE
CENTRE EUROPEEN DE MÉDIATION
CENTRE EUROPEEN D'ARBITRAGE ET DE MEDIATION
3, rue du Général Frère – 67000 Strasbourg

(The Bulletin of the European Court of Arbitration authorized by the Court of Milan no. 326, dated 31/05/97)

Med-Mid Arbitration and Mediation Reports

No. 2/2026 October 2026

Canadian Court Annuls Arbitral Award for Delegation to AI, Consistent with Global Trends p. 3

A basic choice is needed between the two Main Options for Investment Disputes Resolution p. 9

Outsourcing the judge/ arbitrator's role to artificial intelligence? p. 13

ICSID and ECT Awards are not assignable p. 14

Challenges of awards due to a breach of due process as to witness examination and testifying without taking an oath p. 15

Scientific Committee of the
International School of Arbitration and Mediation for Europe, the Mediterranean and the Middle East

Alexander Belohlavek - Jacques El-Hakim - José Carlos Fernandez Rozas - Walid Kassir
Azzedine Kettani - Pierre Mayer - Denis Mouralis - Filali Osman - Gonzalo Stampa - Pierre Tercier

Jennifer Huang,
Raid Abu-Manneh,
Gustavo Mata Morreo,
Brendan J. Harrington,
Rhys Morgan,
Lisa Dubot

Canadian Court Annuls Arbitral Award for Delegation to AI, Consistent with Global Trends

EXECUTIVE SUMMARY

In the first publicly reported instance of a court annulling an arbitral award due to artificial intelligence (AI)-generated reasoning, the Quebec Superior Court set aside an award in which every legal citation had been "hallucinated" by an AI tool used by the sole arbitrator. ARIHQ v. Santé Québec, 2026 QCCS 1360 (April 22, 2026) establishes that an arbitrator who delegates his or her core decision-making authority to AI tools commits a serious procedural failing that risks the enforceability of the award. The Court did not prohibit arbitrators from using AI altogether, but drew a firm line: AI may assist—but must not drive—the decision-making process in rendering an arbitral award. In this Legal Update, we analyze the ruling, situate it within the broader global landscape of AI usage in legal proceedings, and set out practical implications for parties, arbitrators, and counsel.

BACKGROUND

Courts and arbitral tribunals across multiple jurisdictions, including the United States, the United Kingdom, and France, have confronted cases in which counsel or decision-makers submitted or relied upon AI-generated materials containing non-existent legal authorities. In our November 2025 Legal Update, AI Arbitrators Have Now Arrived (At Least in Some Cases)!, we examined the AAA-ICDR's deployment of the first AI arbitrator by a major arbitral institution, noting the safeguards built into that process—such as mandatory human-in-the-loop review and party consent. The Quebec Superior Court's decision in ARIHQ v. Santé Québec presents the other side of the coin: the consequences that follow when an arbitrator uses AI without safeguards.

The underlying domestic arbitration concerned a CAD 1,225,000 claim for remuneration owed by a Quebec health authority to a medical clinic under a national agreement governing the remuneration regime for member clinics, for services provided between 2019 and 2022. The respondent sought preliminary dismissal, arguing that a 90-day contractual deadline for submitting disputes had long expired. Following a half-day hearing in July 2025, the sole arbitrator rendered an award on August 8, 2025 granting the preliminary dismissal. The applicants filed for annulment in November 2025, pointing to, among other things, strong indicia the award had been drafted using generative AI.

THE COURT'S DECISION

The applicants advanced two grounds for annulment, under articles 646 and 648 of the Quebec Code of Civil Procedure (QCP). The first, that the award was contrary to public policy, was rejected. An error of law, even on a public policy provision, is not a ground for annulment.

The second ground proved decisive. The applicants argued that the arbitral procedure had not been followed because the award appeared to have been drafted using generative AI, relying on hallucinated legal authorities. Every doctrinal and jurisprudential citation turned out to be non-existent: a scholarly article attributed to Professor Frédéric Bachand could not be located; three Court of Appeal and Superior Court decisions did not exist; and an arbitral award cited as Arbitrage CHU Ste-Justine was confirmed by the Société Québécoise d'Information Juridique to be fictitious. These non-existent sources were the only legal authorities underpinning the award's reasoning. Based on a preponderance of evidence, Justice Sheehan concluded that the arbitrator had improperly delegated his decision-making authority.

The Court's analysis rested on several principles. Drawing on the duty to maintain the secrecy of deliberations under QCP article 644 and the Canadian Judicial Council's guidelines prohibiting judges from delegating decision-making to any computer program, the Court held that arbitrators must not delegate the drafting of their awards to third parties, including AI tools. Although the Court observed that using reliable AI tools is not inherently improper, the problem arises when AI is permitted to drive—rather than merely assist—the decision-making process. The Court drew an instructive analogy with law clerks and research assistants: they may support the decision-maker, but responsibility for the reasoning and the decision itself must remain with the arbitrator. The Court was careful to note that not every award containing erroneous references or drafted with AI tools would be annulled. The assessment is context-dependent, requiring courts to weigh the nature of the failing, its impact on the integrity of the process, and its effect on the award. With "minimal or peripheral" AI use, a reasoned decision may stand despite procedural shortcomings. But where, as here, non-existent authorities sit at the core of the reasoning, the breach is serious enough to "undermine confidence in the result and in arbitration generally."

Speaking to Global Arbitration Review, the sole arbitrator, Michel Jeanniot, stated that he personally drafted the award based on his independent assessment of the facts and applicable contractual framework, but subsequently asked a legal AI tool to supplement the ruling with additional authorities "without any intention of altering the substance of the reasoning." He acknowledged that he did not validate those references, having "relied in good faith on the integrity and reliability of the service used."

The Court annulled the award and ordered the parties to appoint a new arbitrator within 60 days, with costs.

INTERNATIONAL CONTEXT

This decision sends a clear signal, not only in Canada but also the international arbitration

community more broadly. While the ruling arises under Quebec civil procedure, the principles it articulates are not jurisdiction-specific: the duty of an arbitrator to personally exercise the decision-making function, the obligation to verify legal authorities, and the risk of AI-generated reasoning leading to annulment are concerns that resonate across arbitral seats and legal systems worldwide.

United States

In the United States, the trend is unmistakable and accelerating: courts at every level are imposing concrete sanctions on attorneys who submit AI-generated filings without first verifying their accuracy. The foundational principle is not new: Rule 11 of the Federal Rules of Civil Procedure has always required that at least one attorney of record certify the accuracy of submissions. The judicial response has shifted decisively from admonition to punishment. In Whiting v. City of Athens, Nos. 24-5918/5919 & 25-5424 (6th Cir., Mar. 13 2026), for example, the US Court of Appeals for the Sixth Circuit imposed USD 15,000 in punitive fines on both attorneys of record after discovering at least two dozen hallucinated citations across three consolidated appeals, declaring that they had "brought the profession into disrepute." The Court endorsed a strict per-citation standard: "'[n]o brief, pleading, motion, or any other paper filed in any court should contain any citations — whether provided by generative AI or any other source — that' a lawyer has not personally 'read and verified.'" In Hill v. Workday, Inc., No. 23-cv-06558-PHK (N.D. Cal. Apr. 28, 2026), the Court sanctioned a supervising partner who had never read a discovery brief prepared by his associate using an AI research tool, signaling that sanctions exposure now extends beyond the drafter under the Mattox accountability framework. In Lexos Media IP, LLC v. Overstock.com, Inc., No. 22-2324-JAR (D. Kan. Feb. 2, 2026), all five attorneys of record were sanctioned under Rule 11 after unverified AI-generated results containing at least eleven errors were found in two briefs, with sanctions graduated by role and the court stressing that the Rule 11 duty is "nondelegable." Perhaps most strikingly, in Ibach v. Stewart, No. SC-2025-0106 (Ala. Apr. 24, 2026), the Alabama Supreme Court dismissed an appeal entirely because counsel's briefs were so infected with AI hallucinations that the Court was left with "nothing to review," and imposed USD 17,200 in fees and a referral to the state bar.

Taken together, these decisions confirm that the era of leniency for AI-related errors in US courts is over. Courts are careful to emphasize that the problem does not lie with AI itself but, rather, with the failure of attorneys to exercise their nondelegable duty of verification. The practical lesson for counsel preparing any submission to a court or arbitral tribunal is clear: every AI-generated authority must be independently verified, and every attorney who signs a filing bears personal responsibility for its contents. Supervising attorneys cannot disclaim responsibility by pointing to the tool or to the associate who used it. Beyond the sanctions context, emerging US jurisprudence is also grappling with the discoverability of AI-generated materials in litigation. For a discussion of these issues, including the contrasting judicial approaches in United States v. Heppner, No. 25-cr-00503-JSR (S.D.N.Y. Feb. 17, 2026), and Warner v. Gilbarco, Inc., No. 2:24-cv-12333-GAD-APP (E.D. Mich. Feb. 10, 2026), see our recent Legal Update, M&A Discovery in the AI Era: Generative AI Communications and Outputs May

Become Litigation Ammunition

England and Wales

In England and Wales, the judiciary has established a zero-tolerance policy that leaves little doubt as to how similar challenges would be received. In approximately 50 cases to date in which AI hallucinations have been considered, English courts have confirmed that any legal professional utilizing artificial intelligence tools bears a "personal, absolute duty" to verify the accuracy of the output against authoritative sources. In Hamad Al-Haroun v. Qatar National Bank QPSC [2025] EWHC 1383 (Admin), a solicitor relied on legal research produced by his client using AI; of the 45 citations put before the Court, 18 did not exist. The High Court stated that submitting unverified AI output was "incompetent and grossly negligent," warranting wasted costs orders and regulatory referrals.

In the arbitration context, existing English case law on fabricated evidence, invented citations, and unreliable expert materials provides a well-developed framework for challenges to awards containing AI hallucinations. Under Section 68 of the Arbitration Act 1996, a party may challenge an award for serious irregularity causing substantial injustice. Where a party knowingly submits AI-hallucinated authorities to a tribunal, this may engage the fraud ground under Section 68(2)(g), applying the standard established in Federal Republic of Nigeria v. Process & Industrial Developments Ltd [2023] EWHC 2638 (Comm), which requires cogent evidence of deliberate deception. Where an arbitral tribunal itself introduces hallucinated authorities into its reasoning without allowing the parties to address them, this may constitute a breach of the tribunal's duty of fairness under Section 33, engaging Section 68(2)(a) and following the reasoning in P v. D [2016] EWHC 1277 (Comm), where the High Court set aside an award based on a finding never put to the relevant witness. Importantly, the challenging party must also demonstrate "substantial injustice": that the tribunal might well have reached a different conclusion absent the irregularity. As Harber v. HMRC [2024] UKFTT 1007 (TC) illustrates, where the outcome would have been identical regardless of the fabricated citations, a challenge could fail on materiality grounds.

Perhaps the most significant practical insight emerges from Section 73 of the Arbitration Act 1996, which embodies the principle of "loss of right to object." A party that participates in proceedings without raising an objection loses the right to challenge the award unless it can show it could not have discovered the grounds with "reasonable diligence." As Thyssen Canada Limited v. Mariana Maritime SA [1995] EWHC 219 (Comm) demonstrates, courts apply this bar strictly to fabricated evidence challenges. The implication for AI hallucinations is profound: because verifying a legal citation on standard databases requires only minutes, a party that fails to identify an opponent's fabricated AI-generated authority during the arbitration will almost certainly be barred from challenging the award post hoc. The advent of generative AI has thus implicitly created a "dual burden" on arbitration counsel; not only must lawyers verify their own AI-assisted research, but they now also bear a heavy proactive obligation to audit the authorities cited by opposing parties. The sole exception would be where the tribunal itself introduces hallucinated authorities into the final award without prior notice to the parties, in which case Section 73 cannot apply.

For the majority of international commercial arbitrations seated in London, where parties have contractually excluded Section 69 appeals on points of law, the battleground for AI hallucination-tainted awards will be confined to these Section 68 and Section 73 dynamics.

France

In France, the legal framework under Article 1520 of the French Code of Civil Procedure provides several avenues through which a similar challenge could succeed. An award could be annulled for failure to comply with the tribunal's mandate (Article 1520(3°)) if a party demonstrates that AI effectively replaced the arbitrator, determined the outcome, or influenced deliberations such that the arbitrator signed an award that AI, rather than the arbitrator, actually decided. This ground is reinforced by Article 1450 of the Code, which provides that an arbitrator's mandate may only be exercised by a natural person. An award could also be set aside for violation of due process (Article 1520(4°)) if the tribunal relied on AI-generated evidence or arguments that were determinative of the outcome but were never submitted to adversarial debate, in breach of the principle of contradictoire. Finally, procedural fraud affecting the integrity of the proceedings may justify annulment on international public policy grounds (Article 1520(5°)) by analogy to existing case law on falsified evidence that "surprises" the tribunal, such as forged documents submitted by the parties.

In domestic litigation, French courts have already begun confronting AI-generated content and issuing increasingly explicit warnings. For instance, the Tribunal administratif d'Orléans admonished counsel upon discovering numerous false references in a filing, stressing that all citations must be verified "by any means whatsoever" to avoid reliance on AI "hallucinations" or "confabulations" (29 December 2025, No. 2506461). This emerging judicial response aligns with broader institutional efforts. In April 2025, the Working Group on Artificial Intelligence of France's highest civil court, the Cour de cassation, issued a report emphasizing that human oversight and control over decisions is indispensable to preserving the judge's role, drawing on the European Ethical Charter on the Use of Artificial Intelligence in Judicial Systems. In March 2026, the Conseil National des Barreaux adopted a guide on Déontologie et Intelligence Artificielle, confirming that a lawyer who relies on AI-generated content without appropriate verification may be exposed to disciplinary sanctions. While these developments arise in the domestic litigation context, the principles they articulate—the duty of verification, the prohibition on delegating core functions to AI, and the risk of sanctions for using unverified AI output—are directly transposable to the international arbitration setting

PRACTICAL IMPLICATIONS

For parties to an arbitration: The ARIHQ decision introduces a new and important dimension to post-award scrutiny. Indications that an award may have been AI-generated may give rise to grounds for challenge or annulment. The Quebec Superior Court held that expert evidence is not required to establish AI use; the civil standard of a balance of probabilities applies, and indicia of unverified AI-generated content (hallucinated authorities, internally inconsistent reasoning) may speak for themselves. Parties and their

counsel should consider raising expectations regarding AI usage with tribunals at the outset of proceedings. Given the pace at which AI capabilities are evolving, these discussions are best held early in each dispute (as opposed to at the stage of drafting the arbitration clause), when the parties can agree on ground rules tailored to the case at hand.

For arbitrators and judges: The message is unambiguous. Generative AI may be a useful research or drafting aid, much like a law clerk or research assistant, but the arbitrator must retain ownership of the reasoning and independently verify every authority cited. No responsible arbitrator would sign an award drafted entirely by an assistant without careful review, and the same standard must apply to AI-generated output. In practical terms, this means treating any AI output as a starting point subject to rigorous human verification rather than as a finished product. It also means maintaining a clear record of how AI tools were used in the decision-making process, both for the arbitrator's own protection and to preserve the integrity of the proceedings.

For counsel: The lesson for practitioners is consistent across jurisdictions: all AI-assisted work product, whether submissions to a court or an arbitral tribunal, must be subject to the same standard of independent verification that applies to work prepared by a junior team member. In addition, certain jurisdictions may impose a burden on counsel to police each other's pleadings for unverified AI content during the course of the proceedings.

More broadly, the decision highlights the growing need for the arbitration community to develop clear, practical frameworks for AI usage in the arbitral process. Several arbitral institutions, including the SCC, VIAC, AAA-ICDR, SVAMC, and CIETAC, have already issued guidelines on AI usage in arbitration. The question is no longer whether AI will be used in arbitration, but how to ensure its use is transparent, responsible, and subject to meaningful human oversight. As we observed in our November 2025 Legal Update on the AAA-ICDR's AI Arbitrator, the safeguards that institution built into its process (party consent, defined scope, and mandatory human review) offer a useful model for responsible AI integration. The ARIHQ decision is a stark reminder of what can go wrong in the absence of such safeguards.

We will continue to monitor developments in this rapidly evolving area across all of the jurisdictions in which we practice. Should you wish to discuss how this decision may affect your existing or contemplated arbitration proceedings, and what practical steps can be taken to address AI-related risks, please contact any of the authors or your usual Mayer Brown contact.