



International Journal of Research Publication and Reviews

Journal Homepage: www.ijrpr.com ISSN 2582-7421

The Delegated Mandate: Generative AI in Arbitral Decision-Making and the Enforceability of Awards

Sneha Patidar, Padmapriya Pradhan, Abhishek Rana

| National Law Institute University, Bhopal

ABSTRACT

Generative artificial intelligence (AI) has moved from the margins of legal practice into the work of adjudicators. Existing scholarship on AI in arbitration has focused largely on efficiency gains in case management, document review and predictive analytics. This paper addresses a narrower and more consequential question: when an arbitrator relies on generative AI to analyse evidence, research the law or draft reasoning, has the arbitrator delegated a function that the parties entrusted to the arbitrator personally, and can such reliance affect the validity or enforceability of the resulting award? Adopting a doctrinal and comparative method, the paper examines the personal nature of the arbitral mandate, the grounds for setting aside and refusing enforcement under the New York Convention and the UNCITRAL Model Law, and emerging soft-law guidance. It argues that AI assistance is not unlawful in itself, but that undisclosed reliance which displaces the arbitrator's own judgment exposes awards to challenge, and it proposes a disclosure-based framework in preference to a blanket prohibition.

Keywords: *international arbitration; generative AI; arbitral mandate; delegation; New York Convention; enforcement; due process.*

I. INTRODUCTION

A. Background and Context

International arbitration derives its legitimacy from consent, and consent in turn rests on trust in the individuals whom the parties select, or whom an appointing authority selects on their behalf, to resolve the dispute. Parties choose an arbitrator for particular qualities: legal acumen, sectoral expertise and a reputation for fairness. For this reason the arbitrator's mandate is generally understood to be personal in character, a contract for the services of a particular individual which cannot freely be passed to another.¹ The arbitrator may be assisted, but the act of judging remains the arbitrator's own.

That settled understanding is now being tested by generative AI. Large language models can summarise voluminous records, propose analyses of contested law and produce fluent draft reasoning in seconds. The attraction for arbitrators, who face compressed timetables and ever-larger document sets, is evident, and the technology has already entered adjudication more broadly. In 2023 a judge of the Punjab and Haryana High Court referred to a chatbot's response when considering the jurisprudence on bail,² and the judiciary of England and Wales has issued guidance acknowledging that judicial office holders may use such tools in limited ways while warning of their risks.³

Regulators and arbitral institutions have begun to respond. The European Union's Artificial Intelligence Act treats AI systems intended to assist judicial authorities in researching and interpreting facts and the law as high-risk, and extends that treatment to systems used in a similar way in alternative dispute resolution.⁴ Arbitration-specific soft law has followed, including guidelines issued by the Silicon Valley Arbitration & Mediation Center (2024) and the Chartered Institute of Arbitrators (2025), and a report of the ICC Commission on Arbitration and ADR (2022).⁵ These instruments broadly agree that an arbitrator must not delegate decision-making responsibility to AI. What they do not do, because they cannot, is say what follows in law when an arbitrator does.

B. The Research Gap

¹Gary B Born, *International Commercial Arbitration* (3rd edn, Kluwer Law International 2021).

²Jaswinder Singh @ Jassi v State of Punjab CRM-M-22744-2022 (Punjab and Haryana HC, 27 March 2023).

³Courts and Tribunals Judiciary, *Artificial Intelligence (AI) Judicial Guidance* (first issued December 2023, subsequently updated).

⁴Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act) [2024] OJ L2024/1689, recital 61 and Annex III, point 8(a).

⁵Silicon Valley Arbitration & Mediation Center, *Guidelines on the Use of Artificial Intelligence in Arbitration* (SVAMC 2024); Chartered Institute of Arbitrators, *Guideline on the Use of AI in Arbitration* (CIArb 2025); ICC Commission on Arbitration and ADR, *Information Technology in International Arbitration* (ICC 2022).

Academic attention to AI in arbitration has so far concentrated on three themes: efficiency in case management and document review, the predictive capacity of analytics, and the more speculative prospect of an ‘AI arbitrator’. Scherer’s work on AI and legal decision-making is a leading example of the third.⁶ Much less attention has been paid to a question that is more immediate and, for practitioners, more pressing: the consequences, at the stage of annulment and enforcement, of an arbitrator’s reliance on generative AI in an otherwise conventional, human-led proceeding.

The gap matters because awards are tested not when they are issued but in national courts, which apply their own grounds and their own thresholds. A tribunal may comply with every applicable soft-law guideline in spirit and still produce an award that a losing party seeks to resist on the ground that the tribunal did not decide the case itself. The nearest analogue is the debate over tribunal secretaries, which has generated argument, though little settled doctrine, about where permissible assistance ends and impermissible delegation begins.⁷ Generative AI sharpens that debate because, unlike a human secretary, it is opaque, may be inaccurate in ways that are difficult to detect, and may process confidential material on infrastructure the parties have not approved.

C. Research Question and Thesis

This paper asks: *to what extent does an arbitrator’s use of generative AI in the decision-making process compromise the personal mandate conferred by the parties, and in what circumstances may such use justify the setting aside, or refusal of enforcement, of an award?*

Three sub-questions follow. First, what is the content of the personal mandate, and where lies the line between permissible assistance and impermissible delegation? Second, can the recognised grounds for challenge under the New York Convention⁸ and the UNCITRAL Model Law,⁹ in particular those concerning improper procedure, inability to present one’s case and public policy, accommodate a challenge founded on AI reliance? Third, what regime of disclosure and consent best reconciles technological innovation with the legitimacy of the award?

The thesis of the paper is that AI assistance is not unlawful in itself, and that a blanket prohibition would be both impractical and unnecessary. The legal risk lies instead in *undisclosed* reliance that displaces the arbitrator’s independent judgment. A disclosure-based framework, under which tribunals identify the nature and extent of any AI use and the parties have an opportunity to object, offers the best protection for the award while preserving the benefits of the technology.

D. Methodology and Scope

The paper adopts a doctrinal and comparative methodology. The primary sources are the New York Convention, the Model Law, national arbitration legislation and institutional rules; these are analysed alongside leading commentary and the soft law identified above. Reference is made to English legislation and to Model Law jurisdictions, including India, as comparators, so that the analysis is not tied to a single seat.¹⁰

Three limits on scope should be stated. First, the paper concerns international commercial arbitration; investment treaty arbitration is referred to only by way of illustration. Second, it addresses generative AI used to *assist* a human arbitrator, and does not consider fully automated adjudication or online dispute resolution platforms in which the parties have agreed to machine decision-making. Third, to the author’s knowledge no reported court decision has yet set aside or refused enforcement of an award on the ground of a tribunal’s use of AI. The analysis is therefore necessarily analogical and prospective, drawing on authority concerning arbitral delegation, tribunal secretaries, reasoned awards and the right to be heard. Because the field is moving quickly, the law is stated as at October 2026.

E. Structure of the Paper

Part II develops the conceptual framework, examining the personal mandate and its relationship with party autonomy. Part III analyses the grounds of challenge under the New York Convention and the Model Law. Part IV turns to institutional soft law and relevant case law, including the tribunal secretary controversy and the duty to give reasons. Part V offers a critical assessment and sets out the proposed disclosure framework. Part VI concludes.

II. THE PERSONAL MANDATE AND THE LIMITS OF DELEGATION

A. Arbitration, Consent and the Arbitrator’s Mandate

Arbitration occupies a hybrid position. Its authority is contractual in origin, since it depends on the parties’ agreement to submit their dispute to private adjudication, yet it is jurisdictional in effect, since the resulting award can be enforced through the coercive machinery of the state.¹¹ The point of

⁶Maxi Scherer, ‘Artificial Intelligence and Legal Decision-Making: The Wide Open?’ (2019) 36 *Journal of International Arbitration* 539.

⁷The controversy is most prominently associated with the *Yukos* awards: *Hulley Enterprises Ltd (Cyprus) v The Russian Federation* (PCA Case No AA 226); *Yukos Universal Ltd (Isle of Man) v The Russian Federation* (PCA Case No AA 227); and *Veteran Petroleum Ltd (Cyprus) v The Russian Federation* (PCA Case No AA 228), Final Awards (18 July 2014). The role of the tribunal secretary was raised in the later set-aside proceedings in the Netherlands, discussed in Part IV.

⁸Convention on the Recognition and Enforcement of Foreign Arbitral Awards (adopted 10 June 1958, entered into force 7 June 1959) 330 UNTS 38 (New York Convention).

⁹UNCITRAL Model Law on International Commercial Arbitration 1985 (with amendments as adopted in 2006) (UN Sales No E.08.V.4) (Model Law).

¹⁰Arbitration Act 1996 (UK); Arbitration and Conciliation Act 1996 (India).

¹¹Gary B Born, *International Commercial Arbitration* (3rd edn, Kluwer Law International 2021); Nigel Blackaby and others, *Redfern and Hunter on*

connection between these two characters is the tribunal. It is the tribunal that the parties' agreement constitutes, and it is the tribunal's decision that courts are asked to recognise. A legal order that enforces awards on the strength of consent therefore has a legitimate interest in asking whether the decision before it was truly made by the decision-maker to whom consent was given.

Party autonomy is the organising principle. The Model Law leaves the parties free to agree on the appointment of arbitrators and on the procedure to be followed, subject to the fundamental requirements of fairness.¹² The English Arbitration Act 1996 is to similar effect, and gives the tribunal wide latitude to decide whether, and how far, it should itself take the initiative in ascertaining the facts and the law.¹³ That latitude is significant for present purposes, because it shows that independent inquiry by the tribunal is not foreign to the arbitral process. The question is therefore not whether an arbitrator may inform themselves through tools beyond the parties' submissions, but how, and on what terms.

B. The Personal Character of the Mandate

Parties who appoint an arbitrator do not engage an office; they engage an individual. The selection is typically made after scrutiny of qualifications, experience, language and perceived approach, and it is the arbitrator's own profile on which the parties are taken to have relied. The law of challenge reflects this. An arbitrator may be challenged where circumstances give rise to justifiable doubts as to impartiality or independence, or where the arbitrator lacks the qualifications the parties agreed, and may be replaced if unable to perform the functions of the office or if failing to act without undue delay.¹⁴ Each ground presupposes that the person appointed is the person who performs the function. In English law, a failure properly to conduct the proceedings is a ground for removal, and the tribunal's duty to act fairly and impartially, adopting procedures suited to the case, is the standard against which serious irregularity is measured.¹⁵

It follows that an arbitrator may seek help but may not outsource judgment. The case law on arbitrators who rely on their own expertise illustrates the underlying logic. *Zermalt Holdings* is generally cited for the proposition that an arbitrator who proposes to act on private knowledge or opinion must give the parties a chance to address it.¹⁶ The concern is not that arbitrators use knowledge beyond the record, which is often why they were chosen, but that a private source of persuasion operates unseen and untested. Output from a generative AI system is, in this respect, a species of private source. The parties cannot cross-examine it, and unless told of it they cannot respond to it.

C. Assistance and Delegation: The Tribunal Secretary Analogy

Arbitral practice has long relied on assistants. Tribunal and administrative secretaries organise documents, prepare chronologies, check citations and sometimes prepare first drafts of non-decisional parts of the award. Institutional guidance accepts these functions but draws a firm line at decision-making. The ICC Note to Parties and Arbitral Tribunals, for example, contemplates secretaries performing organisational and administrative tasks while stating that decision-making functions may not be delegated to them in any circumstances.¹⁷ The soft law on AI follows the same pattern: the SVAMC Guidelines articulate a principle of non-delegation of decision-making responsibilities, and treat the arbitrator as answerable for the award whatever tools were used.¹⁸

The analogy is instructive but imperfect, and the imperfections are what make AI a distinct legal problem.

D. Why Generative AI Is Different

Four features distinguish generative AI from a human assistant. First, *opacity*. A secretary can be asked what was done and why; a large language model cannot give a reliable account of how an output was produced, and the arbitrator may be unable to identify what it rests on. Second, *fluent fallibility*. These systems can produce plausible but fabricated content, as the sanctions litigation in *Mata v Avianca* demonstrated when counsel filed submissions citing non-existent authorities generated by a chatbot.¹⁹ The risk for an arbitrator is that a fabricated or distorted authority, or a mischaracterisation of evidence, migrates unnoticed into the reasoning of an award.

Third, *automation bias*. Research on human interaction with automated aids shows a tendency to over-rely on automated suggestions and to reduce independent verification.²⁰ The danger is not merely that a task is performed by someone else, but that the arbitrator's own evaluative faculty is weakened,

International Arbitration (7th edn, OUP 2022).

¹²UNCITRAL Model Law on International Commercial Arbitration 1985 (with amendments as adopted in 2006) (UN Sales No E.08.V.4) (Model Law), arts 11 and 19.

¹³Arbitration Act 1996 (UK), s 34(2)(g).

¹⁴Model Law, arts 12 and 14.

¹⁵Arbitration Act 1996 (UK), ss 24(1)(d)(i), 33(1) and 68(2)(a).

¹⁶*Zermalt Holdings SA v Nu-Life Upholstery Repairs Ltd* [1985] 2 EGLR 14.

¹⁷International Chamber of Commerce, *Note to Parties and Arbitral Tribunals on the Conduct of the Arbitration* (1 January 2021), section on administrative secretaries.

¹⁸Silicon Valley Arbitration & Mediation Center, *Guidelines on the Use of Artificial Intelligence in Arbitration* (SVAMC 2024), guideline 6; see to similar effect Chartered Institute of Arbitrators, *Guideline on the Use of AI in Arbitration* (CI Arb 2025).

¹⁹*Mata v Avianca, Inc* 678 F Supp 3d 443 (SDNY 2023).

²⁰Raja Parasuraman and Dietrich H Manzey, 'Compacency and Bias in Human Use of Automation: An Attentional Integration' (2010) 52 Human Factors 381.

so that adoption of the output becomes a substitute for judgment rather than a product of it. Fourth, *confidentiality and data governance*. Entering case materials into a third-party system may breach confidentiality obligations arising from the law of the seat, the institutional rules or the parties' agreement. Because confidentiality is an incident of the arbitrator's role, a breach may be pleaded alongside delegation as part of a challenge to the arbitrator's conduct.

E. A Functional Spectrum of AI Use

Because the line between assistance and delegation turns on function rather than technology, this paper proposes a four-tier functional spectrum of AI use, set out in Table 1. The tiers are the author's own analytical framework, developed to structure the legal analysis in the Parts that follow.

Table 1: Functional spectrum of AI use in arbitral decision-making

Tier	Typical use	Principal risk	Provisional legal position
1. Administrative	Transcription, scheduling, formatting, file management	Breach of confidentiality; data security	Generally unproblematic if confidentiality obligations are respected; disclosure not ordinarily required.
2. Research and retrieval	Locating legal sources; summarising the record	Inaccuracy; omission; fabricated authority	Permissible if outputs are independently verified against primary materials; disclosure advisable where reliance is material.
3. Drafting	Procedural orders; procedural history; summaries of submissions	Error; anchoring of the arbitrator's thinking	Permissible with full review and adoption by the arbitrator; disclosure advisable.
4. Evaluative	Weighing evidence; assessing credibility; proposing outcomes; drafting reasoning on contested issues	Substitution of the arbitrator's judgment	High risk. Impermissible where it displaces the arbitrator's own judgment; if used at all, disclosure and party consent are required (see Part V).

Two tests give the spectrum operational content. The *substitution test* asks whether the AI output informed the arbitrator's own evaluative judgment on a contested issue of fact, law or quantum, or replaced it. The *verification test* asks whether the arbitrator independently checked the output against the primary materials before relying on it. The tiers are not watertight. Summarising a record (Tier 2) can shape which evidence the arbitrator reads closely, and so slide towards evaluation. The tiers should accordingly be read as setting presumptions, not categorical rules, with the substitution test deciding hard cases.

Party autonomy cuts both ways. Parties may agree to the use of AI tools, and the Model Law leaves room for such agreement on procedure.²¹ Informed consent is the most reliable means of reducing the legal risk of AI-assisted decision-making at Tiers 2 to 4, because a party that has agreed to a procedure will find it difficult to impugn the award on the basis of that procedure. Whether parties can validly waive human judgment altogether is outside the scope of this paper, as explained in Part I. Part III now turns to the grounds on which an award tainted by impermissible delegation may be set aside or refused enforcement.

III. GROUNDS OF CHALLENGE AND ENFORCEMENT

A. The Legal Architecture

An award may be attacked at two points: by an application to set it aside before the courts of the seat, and by resistance to recognition or enforcement before the courts of the state in which enforcement is sought. The grounds at each point are closely aligned and are exhaustive.²² Three are relevant here: inability of a party to present its case;²³ a tribunal composition or arbitral procedure that departed from the parties' agreement or, failing agreement, from

²¹UNCITRAL Model Law on International Commercial Arbitration 1985 (with amendments as adopted in 2006) (UN Sales No E.08.V.4), art 19(1).

²²Convention on the Recognition and Enforcement of Foreign Arbitral Awards (adopted 10 June 1958, entered into force 7 June 1959) 330 UNTS 38 (New York Convention), art V; UNCITRAL Model Law on International Commercial Arbitration 1985 (with amendments as adopted in 2006) (UN Sales No E.08.V.4) (Model Law), arts 34(2) and 36(1).

²³New York Convention, art V(1)(b); Model Law, art 34(2)(a)(ii).

the law of the seat;²⁴ and conflict with public policy.²⁵ English law reaches similar territory through the concept of serious irregularity, and the Indian statute follows the Model Law grounds.²⁶

Three features of this architecture shape the analysis. First, courts do not review the merits, so an error of reasoning, however regrettable, is not in itself a ground of challenge. Second, the grounds are applied narrowly, in keeping with the pro-enforcement policy of the New York Convention.²⁷ Third, relief is discretionary or conditional: enforcement may be refused rather than must be, and the grounds are subject to waiver and to a requirement that the irregularity mattered to the outcome. A challenge founded on AI use must therefore fit within a recognised ground. There is no free-standing ground of unlawful use of technology.

B. Three Routes of Challenge

Composition and agreed procedure. The most direct argument is that a tribunal which has handed its decision-making to an AI system is not the tribunal the parties agreed to, and that the procedure followed was not the agreed procedure. The argument is strongest at Tier 4 of Table 1, where the system's output effectively determines contested issues. It is also strongest where the parties have agreed an AI protocol, or where institutional rules or a procedural order require disclosure, because departure from an agreed procedure is the textual core of the ground. Absent such agreement, courts are likely to confine composition to questions of appointment and independence, and to treat the use of research aids as part of the tribunal's procedural discretion. The argument is therefore realistic only where delegation is substantial and demonstrable.

The right to be heard. The more probable route is the right to present one's case. If an AI system supplies a legal theory, a characterisation of the evidence or an authority that the tribunal adopts without putting it to the parties, the result resembles the surprise decisions criticised in the case law on arbitrators' private knowledge, discussed in Part II. English law treats such failings as potential serious irregularity, and courts in Model Law jurisdictions apply comparable principles of natural justice. These courts nonetheless require the breach to be connected to the making of the award and to have caused prejudice or substantial injustice.²⁸ A fabricated authority that did not affect the outcome will not suffice.

Public policy. Public policy is the residual ground, and the hardest to satisfy. Courts have repeatedly confined it to violations of the most basic notions of justice and morality, and some jurisdictions have narrowed it further by statute and case law.²⁹ An award substantially written by an undisclosed machine might be said to offend fundamental ideas of due process, but a court is unlikely to reach that conclusion unless the delegation is shown to have been wholesale. A breach of confidentiality, which may follow from uploading case materials to external platforms, is more likely to sound in damages or injunction than in annulment, unless it infects the fairness of the procedure.

C. Proof, Prejudice and Waiver

Three practical obstacles limit all three routes. The first is proof. Arbitral deliberations are confidential and courts are reluctant to inquire into how a tribunal reached its decision. A challenger will usually have only textual indications: fabricated citations, stylistic anomalies, or reasoning that fails to engage with the parties' arguments. This produces an uncomfortable asymmetry. Careless use of AI is detectable and therefore challengeable, whereas careful but undisclosed use is invisible. A regime that penalises only the careless is a poor guarantee of legitimacy, which is a central reason why disclosure, rather than after-the-fact review, should carry the weight (Part V).

The second obstacle is prejudice, as noted above. The third is waiver. A party that knows of a departure from the agreed procedure and continues without objection is taken to have waived the right to object,³⁰ so a party who learns of AI use during the proceedings must object promptly. Applied to Table 1, Tier 1 use raises no realistic ground of challenge; Tier 2 and Tier 3 use becomes legally relevant only where it yields an error or surprise that is material; and Tier 4 use alone engages the composition ground directly.

D. A Comparative Note: India

The Indian statute illustrates how the analysis shifts with local doctrine. Section 34 of the Arbitration and Conciliation Act 1996 contains grounds corresponding to inability to present one's case and to non-compliance with the agreed composition or procedure, and the public policy ground has been narrowed by statute and judicial interpretation.³¹ For domestic awards, section 34(2A) adds patent illegality appearing on the face of the award. It is

²⁴New York Convention, art V(1)(d); Model Law, art 34(2)(a)(iv).

²⁵New York Convention, art V(2)(b); Model Law, art 34(2)(b)(ii).

²⁶Arbitration Act 1996 (UK), s 68; Arbitration and Conciliation Act 1996 (India), s 34.

²⁷Albert Jan van den Berg, *The New York Arbitration Convention of 1958: Towards a Uniform Judicial Interpretation* (Kluwer Law and Taxation Publishers 1981); UNCITRAL Secretariat, *Guide on the Convention on the Recognition and Enforcement of Foreign Arbitral Awards (New York, 1958)* (2016 edn).

²⁸*Soh Beng Tee & Co Pte Ltd v Fairmount Development Pte Ltd* [2007] 3 SLR(R) 86 (CA); Arbitration Act 1996 (UK), s 68(2).

²⁹*Parsons & Whittemore Overseas Co Inc v Société Générale de l'Industrie du Papier (RAKTA)* 508 F 2d 969 (2d Cir 1974); International Law Association, *Recommendations on the Application of Public Policy as a Ground for Refusing Recognition or Enforcement of International Arbitral Awards* (ILA Resolution 2/2002); *Ssangyong Engineering & Construction Co Ltd v National Highways Authority of India* (2019) 15 SCC 131.

³⁰Model Law, art 4; Arbitration Act 1996 (UK), s 73.

³¹Arbitration and Conciliation Act 1996 (India), s 34(2)(a)(iii), (2)(a)(v), (2)(b)(ii) and (2A); *Ssangyong Engineering & Construction Co Ltd v National Highways Authority of India* (2019) 15 SCC 131.

arguable that reasoning resting on a fabricated authority, visible on the face of the award and central to the result, could meet that threshold, although courts have confined patent illegality to defects going to the root of the matter. For international commercial arbitrations seated in India the narrower public policy ground and the absence of patent illegality make the due-process routes more important. The comparison confirms that the same AI-related defect may attract different treatment depending on the seat, which strengthens the case for dealing with the issue through the agreed procedure rather than relying on the vagaries of national review.

IV. SOFT LAW AND JUDICIAL RESPONSES

A. Institutional Soft Law: Content and Status

The principal instruments converge on several points.³² Arbitrators remain personally responsible for the award and must not delegate decision-making to AI; AI outputs must be verified before they are relied on; confidentiality and data security must be protected; and the parties' right to be heard must not be compromised by reliance on material they cannot test. On disclosure the instruments are more cautious, generally favouring a case-by-case assessment over a mandatory rule.

Their legal status matters. They are not binding, and a court asked to set aside an award will not treat departure from them as a ground of challenge in itself. They nonetheless matter in three ways. They are likely to inform the standard of fairness expected of a competent tribunal when a court applies concepts such as serious irregularity. They can be incorporated into the agreed procedure by the parties or through a procedural order, at which point departure engages the agreed-procedure ground. And they shape professional expectations, so that failure to follow them will weigh in any challenge to an arbitrator. Soft law is thus a bridge to legal consequence, but one that the parties must cross deliberately by adopting it.

B. Lessons from the Tribunal Secretary Controversy

The nearest litigated analogue is the controversy over the role of the tribunal's assistant in the *Yukos* arbitrations, in which the Russian Federation argued before the Dutch courts that the assistant had performed tasks belonging to the arbitrators.³³ As reported, the courts' approach treated the question as one of fact: what the assistant in fact did, and whether the arbitrators themselves reached the decisions. Three lessons transfer to AI. First, courts do not condemn assistance as such; the line is substantive, not formal. Second, the decisive evidence concerns the actual division of labour, which a challenger can rarely obtain. Third, the controversy arose from opacity about the assistant's role, which supports the argument that early and clear disclosure is preferable to after-the-fact justification.

C. The Duty to Give Reasons

Most arbitration laws require an award to state its reasons unless the parties have agreed otherwise.³⁴ The duty serves a function directly relevant here: reasons are the main evidence that the tribunal actually thought through the dispute. Reasons need only be intelligible and adequate, and an AI-assisted draft that is reviewed, adopted and, where necessary, rewritten by the arbitrator will meet that standard. The risk lies in generic or formulaic reasoning that fails to engage with the parties' specific arguments. Such reasoning may support an inference that the arguments were not considered, and so link the reasons requirement to the right-to-be-heard ground analysed in Part III.

D. Early Judicial Attitudes

Courts have so far addressed AI chiefly in the conduct of litigants and lawyers. English courts have emphasised that professional duties of verification apply to any use of AI in legal submissions and have warned of serious consequences for placing fabricated authorities before the court.³⁵ Those decisions concern counsel and not adjudicators, but they indicate judicial intolerance of unverified output, which is likely to colour how courts regard an arbitrator who relies on AI without verification.

E. The Regulatory Overlay: The EU Artificial Intelligence Act

The Artificial Intelligence Act adds a regulatory dimension. By classifying as high-risk those systems intended to assist judicial authorities, or used in a similar way in alternative dispute resolution, it signals a legislative expectation of human oversight where AI informs adjudication.³⁶ Whether an individual arbitrator is a deployer for these purposes, and which tools fall within the classification, are open questions, and the Act does not regulate arbitral awards as such. Its relevance to challenge is indirect. A party seeking annulment or resisting enforcement in a Member State might invoke the Act's underlying principles as an expression of public policy or of the fairness expected of a tribunal. That argument is untested, but it shows that the

³²Silicon Valley Arbitration & Mediation Center, *Guidelines on the Use of Artificial Intelligence in Arbitration* (SVAMC 2024); Chartered Institute of Arbitrators, *Guideline on the Use of AI in Arbitration* (CIArb 2025); ICC Commission on Arbitration and ADR, *Information Technology in International Arbitration* (ICC 2022).

³³Gerechtshof Den Haag, 18 February 2020, *Russian Federation v Veteran Petroleum Ltd, Yukos Universal Ltd and Hulley Enterprises Ltd*; Hoge Raad, 5 November 2021. The underlying awards are cited in Part I.

³⁴Model Law, art 31(2); Arbitration Act 1996 (UK), s 52(4).

³⁵R (*Ayinde*) v *London Borough of Haringey* [2025] EWHC 1383 (Admin).

³⁶Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act) [2024] OJ L2024/1689, art 14, recital 61 and Annex III, point 8(a).

soft-law principle of non-delegation is converging with binding regulation, and that arbitrators sitting in or enforcing in the European Union cannot assume the matter is left entirely to the parties.

V. CRITICAL ASSESSMENT AND A DISCLOSURE-BASED FRAMEWORK

A. Two Rejected Approaches

A blanket prohibition on AI use by arbitrators has superficial appeal but should be rejected. It would be difficult to police, given how far AI features are embedded in ordinary research, translation and drafting software. It would draw a boundary that the functional analysis in Part II shows to be unstable. And it would forgo real gains in cost and speed, which are themselves part of arbitration's appeal. The opposite approach, leaving the matter to arbitral discretion and non-binding guidance, is equally unsatisfactory. It exposes awards to challenge on grounds which the parties cannot anticipate, and it leaves the careful and the careless arbitrator in the same uncertainty.

B. The Proposed Framework

This paper proposes a framework of four elements, calibrated to the tiers in Table 1.

(1) *Early agreement.* At the first case management conference the tribunal should raise AI use and record the parties' agreement in the procedural order, identifying the permitted tools and tiers, data-security requirements and the disclosure regime. Because the agreement forms part of the agreed procedure, it converts soft law into binding procedure and gives the parties a clear standard against which departure can be measured.

(2) *Tiered disclosure.* Tier 1 use needs no disclosure beyond compliance with confidentiality obligations. Tier 2 and Tier 3 use should be disclosed in general terms where material, identifying the category of tool and the nature of the use. Tier 4 use should be impermissible unless all parties give express and informed consent.

(3) *Opportunity to object.* Parties should be given a short, defined period in which to object. Failure to object should be treated as waiver under the applicable law, which gives finality in return for transparency.

(4) *Verification and certification.* The tribunal should confirm, in the award or an accompanying statement, that the arbitrators personally deliberated and decided the issues and verified any AI-assisted output against the primary materials. Such a statement need not reveal prompts or the content of deliberations, and so preserves their confidentiality, but it fixes the arbitrators' accountability in a form a court can consider. Arbitral institutions are best placed to standardise these elements through model procedural orders and rule revisions, which would lower the cost of agreement.

C. An Illustration

Suppose a three-member tribunal is hearing a construction dispute involving some 20,000 pages of documents. The presiding arbitrator uses a generative AI tool to produce a chronology and to summarise the expert reports (Tier 2), and to draft the procedural history for the award (Tier 3). Under the proposed framework, the use would be addressed in the first procedural order, disclosed in general terms and subject to a short objection period, and the tribunal's certification would confirm that the arbitrators read the key documents and verified the summaries. If the tool had also been asked which expert was more credible, and the award adopted that assessment without independent evaluation (Tier 4), the use would fall outside the framework absent express consent, and the award would be vulnerable on the composition and right-to-be-heard grounds identified in Part III. The example shows that legal risk is concentrated where the machine's output replaces judgment, not where it merely accelerates preparation.

D. Objections and Limits

Three objections deserve a response. First, disclosure may invite tactical challenges and satellite litigation. The risk is real but is mitigated by the waiver rule and by the requirement of materiality, and a disclosed procedure that has been consented to is harder to attack than an undisclosed one. Second, certification may seem to intrude on the confidentiality of deliberations. The proposal is confined to confirming *that* the arbitrators decided the case, not *how* they did so. Third, disclosure cannot cure everything. If an arbitrator in fact adopts an AI-generated analysis wholesale, a bland certification would be inaccurate, and consent obtained on the basis of inaccurate disclosure would be vitiated. The framework therefore depends on arbitrator integrity, which is in the end the foundation of the personal mandate itself.

The framework also has limits of scope. It is designed for international commercial arbitration; its application in investment treaty arbitration, where transparency obligations and state interests differ, would need separate consideration. And it is untested, since no court has yet ruled on these questions.

VI. CONCLUSION

This paper asked to what extent an arbitrator's use of generative AI compromises the personal mandate, and when it may justify setting aside or refusing enforcement of an award. Its answer has three parts. First, the personal mandate requires that evaluative judgment on contested issues remain the arbitrator's own, but it does not forbid assistance; the line is functional, as the four-tier spectrum shows. Second, the recognised grounds of challenge, namely composition and agreed procedure, the right to be heard and public policy, can accommodate a challenge founded on AI use, but only where delegation is substantial, the breach is material and the use can be proved. The pro-enforcement bias of the Convention and the confidentiality of deliberations make successful challenges unlikely, though not impossible. Third, because careless use is detectable while careful undisclosed use is not,

the law's incentives are poorly aligned with legitimacy, and a disclosure-based framework combining early agreement, tiered disclosure, an opportunity to object and arbitrator certification offers the better remedy.

The thesis is therefore confirmed in qualified form: AI assistance is not unlawful in itself, and the greatest legal risk lies in undisclosed reliance that displaces the arbitrator's judgment. Two limitations should be acknowledged. The analysis is necessarily prospective in the absence of reported decisions, and it is confined to AI-assisted rather than AI-decided arbitration. Future research might examine the first judicial decisions as they emerge, empirical practice among arbitrators, and the question whether parties can validly agree to dispense with human decision-making altogether. Until courts speak, the prudent course for arbitrators, institutions and parties alike is to address the question openly at the outset of the proceedings, rather than at the enforcement stage.

BIBLIOGRAPHY

Cases

- *Jaswinder Singh @ Jassi v State of Punjab* CRM-M-22744-2022 (Punjab and Haryana HC, 27 March 2023)
- *Hulley Enterprises Ltd (Cyprus) v The Russian Federation* (PCA Case No AA 226), Final Award (18 July 2014)
- *Mata v Avianca, Inc* 678 F Supp 3d 443 (SDNY 2023)
- *Parsons & Whittemore Overseas Co Inc v Société Générale de l'Industrie du Papier (RAKTA)* 508 F 2d 969 (2d Cir 1974)
- *R (Ayinde) v London Borough of Haringey* [2025] EWHC 1383 (Admin)
- *Russian Federation v Veteran Petroleum Ltd, Yukos Universal Ltd and Hulley Enterprises Ltd* (Gerechtshof Den Haag, 18 February 2020; Hoge Raad, 5 November 2021)
- *Soh Beng Tee & Co Pte Ltd v Fairmount Development Pte Ltd* [2007] 3 SLR(R) 86 (CA)
- *Ssangyong Engineering & Construction Co Ltd v National Highways Authority of India* (2019) 15 SCC 131
- *Veteran Petroleum Ltd (Cyprus) v The Russian Federation* (PCA Case No AA 228), Final Award (18 July 2014)
- *Yukos Universal Ltd (Isle of Man) v The Russian Federation* (PCA Case No AA 227), Final Award (18 July 2014)
- *Zermalt Holdings SA v Nu-Life Upholstery Repairs Ltd* [1985] 2 EGLR 14

Legislation and International Instruments

- Arbitration Act 1996 (UK)
- Arbitration and Conciliation Act 1996 (India)
- Convention on the Recognition and Enforcement of Foreign Arbitral Awards (adopted 1. June 1958, entered into force 7 June 1959) 330 UNTS 38
- Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act) [2024] OJ L2024/1689
- UNCITRAL Model Law on International Commercial Arbitration 1985 (with amendments as adopted in 2006) (UN Sales No E.08.V.4)

Books

- Blackaby N and others, *Redfern and Hunter on International Arbitration* (7th edn, OUP 2022)
- Born GB, *International Commercial Arbitration* (3rd edn, Kluwer Law International 2021)
- van den Berg AJ, *The New York Arbitration Convention of 1958: Towards a Uniform Judicial Interpretation* (Kluwer Law and Taxation Publishers 1981)

Journal Articles

- Parasuraman R and Manzey DH, 'Complacency and Bias in Human Use of Automation: An Attentional Integration' (2010) 52 Human Factors 381
- Scherer M, 'Artificial Intelligence and Legal Decision-Making: The Wide Open?' (2019) 36 Journal of International Arbitration 539

Institutional and Other Materials

- Chartered Institute of Arbitrators, *Guideline on the Use of AI in Arbitration* (CIArb 2025)
- Courts and Tribunals Judiciary, *Artificial Intelligence (AI) Judicial Guidance* (first issued December 2023, subsequently updated)

-
- ICC Commission on Arbitration and ADR, *Information Technology in International Arbitration* (ICC 2022)
 - International Chamber of Commerce, *Note to Parties and Arbitral Tribunals on the Conduct of the Arbitration* (1 January 2021)
 - International Law Association, *Recommendations on the Application of Public Policy as a Ground for Refusing Recognition or Enforcement of International Arbitral Awards* (ILA Resolution 2/2002)
 - Silicon Valley Arbitration & Mediation Center, *Guidelines on the Use of Artificial Intelligence in Arbitration* (SVAMC 2024)
 - UNCITRAL Secretariat, *Guide on the Convention on the Recognition and Enforcement of Foreign Arbitral Awards (New York, 1958)* (2016 edn)