



## The Journal of Robotics, Artificial Intelligence & Law

Editor's Note: Principles for Principals  
Victoria Prussen Spears

My AI Wrote a Check That I Must Cash: Principles for Principals  
James A. Sherer, Caleb Mabe, Emily Fedeles Czebiniak, Noam Kleinman, and  
Ayyah Saleh

Board Considerations for Public Companies Engaging with Digital Assets  
Peter I. Altman, James Joseph Benjamin Jr., John Patrick Clayton, and John C. Murphy

**The Silicon Arbiter: AI-Generated Arbitration Awards and the Federal Arbitration Act—  
Part II**  
David L. Evans

Use of Artificial Intelligence in Arbitral Institutions  
Jeremy Andrews, James Wagner, Ishan Wad, and Patton Lu

The AI M&A Playbook: Contracting for the Unknown  
C. Craig Lilly and Bryan C. Sykes

How the Adaptation of Artificial Intelligence Tools Is Impacting the Practice of Law  
Seth M. Pavsner

Securities and Exchange Commission Staff Unveils a Playbook for Tokenized Securities  
Andrew P. Blake, Sonia Gupta Barros, Teresa Wilton Harmon, Kate L. Lashley,  
Peter Y. Malyshev, Andrew J. Sioson, Charles A. Sommers, and Lilya Tessler

The EU's Product Liability Directive: What It Means for Aviation and Aerospace  
Companies  
Jamie L. Lanphear, Patrick E. Bradley, Daniel Kadar, and Gregory Speier

- 319 Editor's Note: Principles for Principals**  
Victoria Prussen Spears
- 323 My AI Wrote a Check That I Must Cash: Principles for Principals**  
James A. Sherer, Caleb Mabe, Emily Fedeles Czebiniak,  
Noam Kleinman, and Ayyah Saleh
- 343 Board Considerations for Public Companies Engaging with Digital Assets**  
Peter I. Altman, James Joseph Benjamin Jr., John Patrick Clayton,  
and John C. Murphy
- 349 The Silicon Arbiter: AI-Generated Arbitration Awards and the Federal Arbitration Act—Part II**  
David L. Evans
- 365 Use of Artificial Intelligence in Arbitral Institutions**  
Jeremy Andrews, James Wagner, Ishan Wad, and Patton Lu
- 371 The AI M&A Playbook: Contracting for the Unknown**  
C. Craig Lilly and Bryan C. Sykes
- 377 How the Adaptation of Artificial Intelligence Tools Is Impacting the Practice of Law**  
Seth M. Pavsner
- 381 Securities and Exchange Commission Staff Unveils a Playbook for Tokenized Securities**  
Andrew P. Blake, Sonia Gupta Barros, Teresa Wilton Harmon,  
Kate L. Lashley, Peter Y. Malyshev, Andrew J. Sioson,  
Charles A. Sommers, and Lilya Tessler
- 387 The EU's Product Liability Directive: What It Means for Aviation and Aerospace Companies**  
Jamie L. Lanphear, Patrick E. Bradley, Daniel Kadar, and  
Gregory Speier

**EDITOR-IN-CHIEF**

**Steven A. Meyerowitz**

*President, Meyerowitz Communications Inc.*

**EDITOR**

**Victoria Prussen Spears**

*Senior Vice President, Meyerowitz Communications Inc.*

**BOARD OF EDITORS**

**Jennifer A. Johnson**

*Partner, Covington & Burling LLP*

**Paul B. Keller**

*Partner, Allen & Overy LLP*

**Garry G. Mathiason**

*Shareholder, Littler Mendelson P.C.*

**James A. Sherer**

*Partner, Baker & Hostetler LLP*

**Elaine D. Solomon**

*Partner, Blank Rome LLP*

**Edward J. Walters**

*Chief Strategy Officer, vLex*

**John Frank Weaver**

*Director, McLane Middleton, Professional Association*

**START-UP COLUMNIST**

**Jim Ryan**

*Partner, Morrison & Foerster LLP*

THE JOURNAL OF ROBOTICS, ARTIFICIAL INTELLIGENCE & LAW (ISSN 2575-5633 (print) /ISSN 2575-5617 (online) at \$495.00 annually is published six times per year by Full Court Press, a Fastcase, Inc., imprint. Copyright 2026 Fastcase, Inc. No part of this journal may be reproduced in any form—by microfilm, xerography, or otherwise—or incorporated into any information retrieval system without the written permission of the copyright owner. For customer support, please contact Fastcase, Inc., 729 15th Street, NW, Suite 500, Washington, D.C. 20005, 202.999.4777 (phone), or email customer service at [support@fastcase.com](mailto:support@fastcase.com).

Publishing Staff

Publisher: David Nayer

Production Editor: Sharon D. Ray

Cover Art Design: Juan Bustamante

Cite this publication as:

The Journal of Robotics, Artificial Intelligence & Law (Fastcase)

This publication is sold with the understanding that the publisher is not engaged in rendering legal, accounting, or other professional services. If legal advice or other expert assistance is required, the services of a competent professional should be sought.

Copyright © 2026 Full Court Press, an imprint of Fastcase, Inc.

All Rights Reserved.

A Full Court Press, Fastcase, Inc., Publication

Editorial Office

729 15th Street, NW, Suite 500, Washington, D.C. 20005

<https://www.fastcase.com/>

POSTMASTER: Send address changes to THE JOURNAL OF ROBOTICS, ARTIFICIAL INTELLIGENCE & LAW, 729 15th Street, NW, Suite 500, Washington, D.C. 20005.

## **Articles and Submissions**

Direct editorial inquiries and send material for publication to:

Steven A. Meyerowitz, Editor-in-Chief, Meyerowitz Communications Inc.,  
26910 Grand Central Parkway, #18R, Floral Park, NY 11005, smeyerowitz@  
meyerowitzcommunications.com, 631.291.5541.

Material for publication is welcomed—articles, decisions, or other items of interest to attorneys and law firms, in-house counsel, corporate compliance officers, government agencies and their counsel, senior business executives, scientists, engineers, and anyone interested in the law governing artificial intelligence and robotics. This publication is designed to be accurate and authoritative, but neither the publisher nor the authors are rendering legal, accounting, or other professional services in this publication. If legal or other expert advice is desired, retain the services of an appropriate professional. The articles and columns reflect only the present considerations and views of the authors and do not necessarily reflect those of the firms or organizations with which they are affiliated, any of the former or present clients of the authors or their firms or organizations, or the editors or publisher.

### **QUESTIONS ABOUT THIS PUBLICATION?**

For questions about the Editorial Content appearing in these volumes or reprint permission, please contact:

David Nayer, Publisher, Full Court Press at david.nayer@clio.com or at  
202.999.4777

For questions or Sales and Customer Service:

#### **Customer Service**

Available 8 a.m.–8 p.m. Eastern Time  
866.773.2782 (phone)  
support@fastcase.com (email)

#### **Sales**

202.999.4777 (phone)  
sales@fastcase.com (email)

ISSN 2575-5633 (print)  
ISSN 2575-5617 (online)

# The Silicon Arbiter: AI-Generated Arbitration Awards and the Federal Arbitration Act—Part II

David L. Evans\*

*The emergence of artificial intelligence (AI) systems has made urgent a question that would have seemed absurd to the drafters of the Federal Arbitration Act (FAA): whether algorithms might generate binding arbitral awards without human involvement. This two-part article concludes that, under the FAA as currently written, such awards cannot be enforced. Yet this article does not end with a negative prognosis for AI's place in arbitration. Instead, it makes suggestions for legislative reform addressing informed consent, explainability, bias mitigation, accountability, and meaningful judicial review.*

*The first part of this article, published in the July-August 2026 issue of The Journal of Robotics, Artificial Intelligence & Law, framed the enforceability question precisely, distinguishing pure AI arbitration from AI-assisted human arbitration and explaining why this article focuses on judicial enforcement rather than initial contractual validity; established the textualist interpretive framework, explaining why the Supreme Court's recent decision in Loper Bright Enterprises v. Raimondo controls and what that methodology requires; applied the major questions doctrine from West Virginia v. EPA, demonstrating that recognizing AI-generated awards would constitute precisely the kind of major question requiring explicit congressional authorization; and applied textualist analysis to the FAA, demonstrating through dictionary definitions, statutory structure, and Section 10's vacatur grounds that the statute presupposes human arbitrators.*

*The conclusion of this article, published here, addresses constitutional concerns under both the Due Process Clause and Article III; responds to counterarguments based on party autonomy, engaging particularly with Professor Szalai's recent purposivist arguments; distinguishes AI-assisted arbitration from AI-generated awards and provides guidance on where the line falls; proposes a legislative framework for those who would chart a different course; and offers a conclusion.*

---

## Constitutional Considerations

Even if the Federal Arbitration Act (FAA) could be interpreted to permit artificial intelligence (AI) arbitration, constitutional

concerns would arise under both the Due Process Clause and Article III. While these constitutional questions lie beyond this article's primary focus on statutory interpretation, they warrant examination because they reinforce the statutory analysis and highlight additional obstacles to AI arbitration.

## Due Process Concerns

The Due Process Clause requires fair procedures before government deprives persons of property. While private arbitration generally falls outside the Constitution's scope, courts enforce arbitral awards through judicial confirmation proceedings that implicate due process.<sup>42</sup> AI arbitration raises several due process concerns. First, the opacity of many AI systems may deny parties adequate notice of the basis for decisions affecting their rights. Due process may not require merely a decision, but a reasoned decision that parties can understand and, if appropriate, challenge.

The Wisconsin Supreme Court's decision in *State v. Loomis* provides instructive parallels. The court upheld the use of algorithmic risk assessments in criminal sentencing but imposed strict conditions: the algorithm's role must be limited, its operation must be disclosed, and the sentencing judge must independently exercise judgment.<sup>43</sup> These conditions recognize that due process requires human judgment in matters affecting fundamental rights, even when algorithmic tools may assist. Similar logic applies with greater force to AI arbitration, where the AI would be the decision-maker rather than merely an advisory tool.<sup>44</sup>

The *Loomis* court's conditional approval is instructive for what it required, not merely what it permitted. The algorithm survived constitutional scrutiny only because a human judge retained ultimate discretion over sentencing. The court mandated specific written warnings about the algorithm's limitations, including its proprietary nature and reliance on group-level data rather than individualized assessment. Most significantly, the court required that presentence reports include a caution that algorithmic risk scores are not intended to determine the severity of the sentence or whether an offender is incarcerated. This limitation establishes a constitutional floor: AI may inform human judgment but cannot supplant it on matters affecting legal rights. Pure AI arbitration removes precisely the human safety valve that saved the algorithm from constitutional invalidation. Where *Loomis* required a human

decision-maker to exercise independent judgment after receiving algorithmic input, pure AI arbitration eliminates the human entirely, making the algorithm's output the final word on the parties' contractual and legal rights.

The U.S. Supreme Court recognized in *Goldberg v. Kelly* that due process requires "an opportunity to be heard at a meaningful time and in a meaningful manner."<sup>45</sup> A hearing before an AI system may not be deemed meaningful in this sense because there is no hearer present. The parties present their cases, but no mind considers their arguments, no judgment weighs their evidence, no conscience struggles with the decision's difficulty.

### Article III Issues

Article III of the U.S. Constitution vests the judicial power in courts whose judges enjoy life tenure and salary protection. The Supreme Court has repeatedly emphasized that this structure serves to ensure independent judgment in the resolution of disputes.<sup>46</sup> While private arbitration lacks the state action necessary to trigger constitutional protections, the distinct role of the court in confirmation proceedings transforms a private contract outcome into a public judgment backed by the state's coercive power. Courts cannot constitutionally lend their imprimatur to a judgment generated by a process that fundamentally denies the right to be heard by a competent adjudicator.<sup>47</sup>

The Court has permitted non-Article III adjudication in various contexts, including arbitration, where parties consent and adequate judicial review remains available.<sup>48</sup> But this tolerance rests on certain premises. In *CFTC v. Schor*, the Court held that parties may consent to non-Article III adjudication, but such consent "does not eliminate all constitutional constraints."<sup>49</sup> The constitutional acceptability of arbitration depends partly on the availability of meaningful judicial review under Section 10's vacatur provisions.

AI arbitration threatens this constitutional balance. If Section 10's vacatur grounds cannot meaningfully apply to AI systems because concepts like "partiality" and "corruption" presuppose a human arbitrator, then judicial review becomes hollow. Courts would confirm AI-generated awards without any substantive basis for evaluation. The FAA's enforcement mechanism would convert algorithmic outputs into binding judgments without the meaningful judicial oversight that makes arbitration constitutionally

tolerable. This raises serious questions about whether courts can constitutionally confirm awards generated by systems to which the statutory review standards do not coherently apply.

## The Accountability Gap

Both due process and Article III concerns converge on the accountability gap in AI arbitration. When a human arbitrator renders an award, that person takes responsibility for the decision. The arbitrator's name appears on the award, their reasoning can be examined, and they can be called to account for misconduct. This accountability reflects the fundamental principle that those who exercise power over others' rights must answer for how they exercise it.

AI-generated awards sever this connection between power and accountability. The AI system cannot be held responsible for its decisions. The developers who created the system disclaim responsibility for specific outcomes. The parties who designed the arbitration agreement may not understand how the AI operates. No one bears the kind of personal responsibility that characterizes human adjudication. This accountability vacuum raises serious questions about whether AI arbitration comports with constitutional requirements for adjudication affecting property rights.

## The Limits of Party Autonomy

---

### The Party Autonomy Argument

The strongest argument for enforcing AI-generated awards rests on party autonomy, a cornerstone of arbitration law. Professor Imre Szalai has developed this argument, contending that the FAA's "flexible, open approach to the meaning of arbitration promoting party autonomy should be broad enough to encompass the use of AI arbitration to resolve disputes."<sup>50</sup> On this view, if parties knowingly consent to algorithmic decision-making, the FAA's pro-arbitration policy supports enforcement.

Szalai's approach emphasizes the commercial goals animating the FAA. As he has documented in his comprehensive historical discussion, the FAA's drafters were merchant organizations seeking efficient dispute resolution.<sup>51</sup> The 1924 joint hearings before Congress featured testimony from business leaders who sought to

replicate the efficient commercial arbitration available in European trading centers.<sup>52</sup> The FAA's purpose is to facilitate party-chosen dispute resolution mechanisms, and AI arbitration serves that purpose by providing faster, cheaper resolution.

The argument has intuitive appeal. The Supreme Court has repeatedly emphasized that arbitration is "a matter of consent, not coercion."<sup>53</sup> In *First Options of Chicago, Inc. v. Kaplan*, the Court observed that "arbitration is simply a matter of contract between the parties."<sup>54</sup> If parties prefer faster, cheaper algorithmic resolution, why should courts prevent them from obtaining it?<sup>55</sup>

### Why Party Autonomy Cannot Overcome the Statutory Text

The argument, however compelling as policy, conflates two distinct questions. The first is whether parties can validly agree to resolve disputes through AI processes. The answer is yes, as a matter of contract law. The second is whether the output of such processes constitutes an "arbitral award" entitled to the FAA's enforcement regime. Here the answer must be no, because the statute's terms contemplate human decision-makers. Party autonomy operates within statutory bounds, not above them.

The Supreme Court's decision in *Hall Street Associates, L.L.C. v. Mattel, Inc.* establishes this principle decisively. The Court held that parties cannot contractually expand judicial review grounds beyond those in Section 10. "[T]he text compels a reading of the §§ 10 and 11 categories as exclusive," Justice David Souter wrote.<sup>56</sup> If parties cannot expand the grounds for vacatur, they cannot expand the definition of "arbitrator." *Hall Street* establishes that the FAA defines the framework within which party autonomy operates.

Nor can the FAA's failure to define "arbitrator" be construed as implicit authorization for AI adjudication. A recurring argument for AI arbitration treats statutory silence as flexibility: because Congress did not specify that arbitrators must be human, the argument runs, parties retain freedom to designate any decision-making mechanism they choose. This reasoning inverts the proper interpretive hierarchy. Under the textualist methodology that *Loper Bright* restored, statutory silence is not delegation. When Congress leaves a term undefined, courts must determine its ordinary meaning at the time of enactment; silence directs courts to the "best reading" of the text as and when written, not an invitation for judicial creativity.<sup>57</sup> As the Court emphasized in *Bostock v. Clayton County*, when

“the express terms of a statute give us one answer and extratextual considerations suggest another, it’s no contest. Only the written word is the law.”<sup>58</sup> The FAA’s silence regarding the definition of “arbitrator” thus compels courts to apply the 1925 understanding of that term, not to treat the omission as a blank check for technological adaptation.

The Court’s decision in *New Prime Inc. v. Oliveira* reinforces this textualist approach. Justice Neil Gorsuch’s majority opinion emphasized that statutory terms must be given their ordinary meaning at the time of enactment.<sup>59</sup> “[I]t’s a ‘fundamental canon of statutory construction’ that words generally should be ‘interpreted as taking their ordinary . . . meaning . . . at the time Congress enacted the statute.’” The word “arbitrator” in 1925 meant a human decision-maker, and parties cannot contract around that meaning.

Consider an analogy. Parties might agree to resolve their dispute by consulting an oracle, flipping a coin, or following an astrologer’s instructions. They would be bound by their contractual agreement. But the resulting “decision” would not be an arbitral award entitled to confirmation under the FAA, because none of these procedures constitutes arbitration as the statute understands it. Consent creates contractual obligations but does not transform non-arbitration into arbitration. AI-generated decision-making falls into the same category.

Szalai’s purposivist approach, for all its historical richness, cannot overcome the textual obstacles. The FAA’s drafters may have sought efficient dispute resolution, but they sought it through a particular mechanism: arbitration by human arbitrators. Their purposes do not authorize courts to redefine the mechanism they chose. Statutes embody purposes at particular levels of generality, and courts cannot invoke abstract purposes to override specific textual choices. Under textualism as demanded by *Loper Bright*, the FAA’s human-centric language controls regardless of any efficiency gains.

## Responding to Other Counterarguments

### *The Technological Inevitability Argument*

Some will argue that this article’s conclusion is anachronistic, that technological progress will render human arbitration obsolete, and that law must adapt. AI capabilities are advancing rapidly.

Systems that seemed impossible a decade ago now perform with remarkable sophistication. Why should arbitration remain frozen in 1925?

The response is that statutory interpretation is not prophecy. Courts interpret statutes as written, not as they might be rewritten to accommodate future developments. Congress chose particular words with particular meanings. Those meanings control until Congress changes them. Moreover, the technological inevitability argument proves too much. If it justified redefining “arbitrator” to include AI, it would equally justify redefining “judge” or “jury.” Constitutional provisions requiring jury trials would become obstacles to efficiency gains. The appropriate response to technological change is legislative action—whether amending existing law or deliberately declining to do so—not judicial redefinition of foundational legal concepts.

### *The Efficiency Argument*

A related objection emphasizes efficiency. AI arbitration could be faster, cheaper, and more accessible than human arbitration. High-volume, low-value disputes that currently go unresolved might find resolution through algorithmic processes. These efficiency gains are real and significant. But efficiency is not the only value the legal system serves. The FAA’s confirmation and vacatur provisions protect parties from fundamentally unfair proceedings. Courts cannot override that judgment by invoking efficiency concerns Congress did not authorize them to consider. Furthermore, efficiency gains remain speculative. The 2025 Queen Mary Survey found significant skepticism about AI’s reliability in dispute resolution contexts.<sup>60</sup> If AI awards require extensive human oversight to be trustworthy, the efficiency case weakens considerably.

### *The Functional Equivalence Argument*

A more sophisticated objection proposes functional equivalence: perhaps AI systems can satisfy the FAA’s requirements not literally but functionally. Algorithmic bias might serve as the functional equivalent of human partiality. In other words, processing failures might constitute functional misconduct.

The functional equivalence argument rests on a seductive analogy: if a machine can perform the tasks associated with a role, it should qualify for that role regardless of its nature. A dictionary

from 1920 would define “chef” as a human, the argument runs, yet we would accept a robot that chops, heats, and plates food as performing the chef’s function. By extension, if an AI can process evidence and render a decision, it should qualify as an “arbitrator.”

This analogy collapses under scrutiny. A robot chef performs physical tasks: heating, chopping, mixing, plating. These are objective, replicable actions measurable against external standards. Arbitration is categorically different. When arbitrators determine whether conduct was “reasonable,” “in good faith,” or “unconscionable,” they engage in normative evaluation that requires understanding context, assessing credibility, weighing equities, and applying community standards of fairness. These are not tasks reducible to pattern matching or probabilistic token prediction. Current AI systems simulate decisional outputs through statistical correlation, but simulation is not performance. The outputs may resemble human judgments, but resemblance does not satisfy a statute whose enforcement mechanisms presuppose the capacity for moral agency and reasoning.

Moreover, statutory interpretation does not track colloquial language evolution. Society might informally call a Roomba a “cleaner” or a food-preparation robot a “chef,” but statutes are not living organisms that mutate with slang. Under the *Loper Bright* framework, an “arbitrator” in the FAA remains what it was in 1925: a human decision-maker capable of judgment, discretion, and moral accountability. The functional equivalence argument attempts to decouple the arbitrator’s role from the human condition, reframing arbitration as a service rather than a status conferred on persons capable of bearing the responsibilities the statute imposes. This move finds no support in the statutory text.

## AI-Assisted Versus AI-Generated: A Crucial Distinction

---

### The Proper Role of AI in Arbitration

Nothing in this analysis suggests that AI has no place in arbitration. AI tools may appropriately assist human arbitrators: organizing documents, identifying precedents, analyzing quantitative evidence, calculating damages, even drafting portions of awards for human review and approval. The 2025 Queen Mary-White &

Case International Arbitration Survey found that 91 percent of respondents expected to use AI for research and data analytics, and 77 percent considered it appropriate for arbitrators to use AI to calculate damages.<sup>61</sup> What AI cannot do under the FAA is replace the human arbitrator as the final decision-maker who issues the award.

The American Arbitration Association's (AAA's) approach to its AI Arbitrator illustrates this distinction. Despite the "AI Arbitrator" branding, the system operates within a "human-in-the-loop framework" where human arbitrators "review, refine and issue the final decision."<sup>62</sup>

The 2025 AAA Guidance on Arbitrators' Use of AI Tools emphasizes that arbitrators must "retain independent decision-making" and use AI "as a support tool, not as a substitute." This guidance correctly captures the line between permissible AI assistance and impermissible AI replacement.

### **Where the Line Falls: Guidance for Practitioners**

Because hybrid models occupy the middle of the spectrum, practitioners need guidance on where permissible assistance ends and impermissible replacement begins. The key question is whether a human arbitrator exercises genuine independent judgment over the award's substance.

Several factors indicate that human judgment remains controlling. First, the human arbitrator must review the AI's output rather than accepting it automatically. Second, the arbitrator must possess the expertise and information necessary to evaluate the AI's analysis. If the arbitrator lacks access to the underlying evidence or cannot assess the AI's reasoning, the review is pro forma. Third, the arbitrator must be empowered to reject or modify the AI's proposal when warranted. Fourth, the arbitrator must take personal responsibility for the award's contents.

Conversely, several factors suggest impermissible replacement. If the AI generates the award and the human's role is limited to checking for obvious errors, the human is not functioning as an arbitrator. If the arbitrator lacks time, expertise, or information to meaningfully evaluate the AI's output, the review is cosmetic. If institutional pressures discourage deviation from AI recommendations, the human's judgment is not genuinely independent.

The AAA's current model falls on the permissible side of this line. Human arbitrators review draft awards, possess relevant expertise, and retain authority to modify or reject AI outputs. The ultimate award is the decision of the arbitrator.

## Parallels in Other Professional Domains

This distinction parallels developments in other professional domains. Physicians may use AI to assist in diagnosis, but medical judgment remains the physician's. Lawyers may use AI for research, but legal advice must come from a licensed professional. Judges may use AI for research or drafting, but judicial decisions must be rendered by human judges who take responsibility for them. Similarly, arbitrators may use AI as a tool, but the arbitral judgment must be their own.

The human-in-the-loop is not a procedural formality but an essential requirement. The human arbitrator brings capacities that AI lacks: the ability to assess credibility, to exercise equitable discretion, to understand context and nuance, to bear moral responsibility for the decision rendered. These capacities are constitutive of what it means to arbitrate, and imperative for an enforceable award under the FAA.

## A Framework for Legislative Reform

---

### The Case for Congressional Action

If society determines that AI-generated arbitral awards should be enforceable, the appropriate path is legislative action rather than judicial reinterpretation. Congress could amend the FAA to authorize algorithmic arbitration, defining conditions under which such procedures would be valid and specifying standards for judicial review. Such processes would allow democratic deliberation about the questions AI adjudication raises.

These questions are not merely technical. They implicate fundamental values about justice, judgment, and the relationship between humans and machines. Should parties' fates be decided by processes they cannot fully understand? Should AI developers bear responsibility for the awards those systems generate? How should courts review decisions whose reasoning they cannot trace? These

questions are for legislatures, not for courts engaged in statutory interpretation.

## Elements of a Legislative Framework

Should Congress choose to authorize AI arbitration, legislation should consider the following elements:

First, informed consent. Parties must specifically agree to AI adjudication, with clear disclosure of what that entails. Standard-form arbitration clauses should not suffice. Parties should understand that a computer program rather than a human being will analyze their submissions and render a decision. The Executive Order on AI and the NIST (National Institute of Standards and Technology) AI Risk Management Framework provide useful models for consent disclosure requirements.<sup>63</sup>

Second, explainability. AI systems used in arbitration should be capable of generating reasoned explanations for their decisions. Black box systems whose reasoning cannot be traced should be impermissible. The NIST framework's emphasis on transparency and explainability provides useful guidance.<sup>64</sup> Parties should be able to understand why they won or lost, and courts should have a meaningful record for judicial review.<sup>65</sup>

Third, bias mitigation. AI systems should be tested for discriminatory bias and monitored for disparate impacts. Training data should be audited for historical biases that might produce systematically unfair outcomes. Ongoing monitoring should identify and correct bias that emerges through operation.

Fourth, accountability. Legislation should specify who bears ultimate responsibility for the AI systems. Developers, institutions administering AI arbitration, and parties invoking such systems should have defined roles and liabilities. The current accountability gap is incompatible with adjudicatory legitimacy.

Fifth, judicial review. Legislation should define standards for reviewing AI-generated awards. The current Section 10 grounds do not translate naturally to algorithmic decision-makers. New grounds addressing algorithmic malfunction, systematic bias, and processing failures should be specified. Courts should have authority to examine the AI system's functioning, not merely the award's facial validity.

Sixth, categorical limitations. Legislation might exclude certain categories of disputes from AI arbitration. Employment

discrimination claims, civil rights matters, and disputes involving fundamental rights might be reserved for human decision-makers capable of exercising moral judgment. The Court's recognition in *Mitsubishi Motors*<sup>66</sup> and *Gilmer*<sup>67</sup> that arbitration of statutory claims is permissible rested partly on the assurance that human arbitrators would apply the law. AI systems present different concerns.

## The Importance of Democratic Deliberation

The path forward should be charted through democratic processes rather than judicial interpretation. Legislatures can hold hearings, receive expert testimony, consider competing interests, and craft balanced solutions. Courts interpreting century-old statutes cannot perform these functions. If AI arbitration is to become part of our legal landscape, its contours should be shaped by the people's representatives, not by judges stretching statutory language beyond recognition.

For those optimistic about AI's potential in dispute resolution, this conclusion is a call to action rather than a final roadblock. The analysis here identifies what legal changes would be necessary to permit AI arbitration. Proponents should advocate for those changes through appropriate democratic processes.

For those who remain concerned about AI's place in arbitration, this conclusion provides the narrowest and most defensible ground for resisting the displacement of human judgment. The argument rests not on speculation about AI's future capabilities or philosophical objections to algorithmic decision-making, but on the text of statutes that presuppose human arbitrators.

## Conclusion

As the law currently stands, fully AI-generated arbitral awards cannot be confirmed under the FAA. This conclusion rests on straightforward textual analysis: the statute presupposes human arbitrators exercising judgment, discretion, and moral reasoning that current AI systems cannot replicate. The statute employs language—"corruption," "fraud," "partiality," "misconduct"—whose settled meaning presupposes human decision-makers capable of intentional wrongdoing, bias, and ethical lapses.

This textual reality does not reflect technophobia or reflexive resistance to innovation. It reflects the simple fact that Congress in 1925 assumed arbitrators would be human. They used language whose ordinary meaning presupposed human decision-makers. Textualist interpretation requires applying statutes as written, not rewriting them to accommodate technological developments the drafters never contemplated. The Supreme Court's recent decisions in *Loper Bright* and *West Virginia v. EPA* reinforce this interpretive approach.

The practical implications are immediate and significant. The AAA's AI Arbitrator, operating under meaningful human review, likely falls within acceptable bounds because human arbitrators retain ultimate decision-making authority. By contrast, systems purporting to offer pure AI arbitration without human oversight cannot produce awards reliably enforceable under the FAA. Parties using such systems face substantial legal risk: they may achieve quick resolution but discover their "award" lacks judicial enforceability.

For parties negotiating arbitration agreements, the implications are straightforward: clauses requiring pure AI arbitration without human oversight face serious enforceability risks. Until the law changes, the safer path uses AI as a tool assisting human arbitrators rather than replacing them.

Party autonomy, while central to arbitration, cannot overcome the statutory text. *Hall Street* establishes that parties cannot contract around the FAA's requirements. Parties can agree to resolve disputes through pure AI, but the resulting outputs are not arbitral awards entitled to confirmation under the FAA. This is not a restriction on party autonomy but a recognition that statutory terms have meanings courts must enforce.

Recognizing that technological pressure will continue and that legislative change may come, the framework identified here suggests issues that AI arbitration legislation should address. Informed consent, explainability, bias mitigation, accountability, meaningful judicial review, and categorical limitations represent elements of a legislative scheme that takes seriously both the promise and the peril of algorithmic adjudication.

Today's AI systems can process, predict, and generate, but they cannot exercise the distinctively human capacity the law calls judgment. Until legislatures determine otherwise that capacity remains the essential ingredient of enforceable arbitral awards under the

FAA. The silicon arbiter may yet find its place in our system of dispute resolution, but only through democratic deliberation that grapples honestly with what we ask of those who decide our disputes, and only if the frameworks adopted preserve the fundamental fairness that legitimates adjudication in a society committed to the rule of law.

## Notes

\* The author, counsel in the Boston office of Murphy & King, may be contacted at [devans@murphyking.com](mailto:devans@murphyking.com). Endnote numbering continues from Part I.

42. *Mathews v. Eldridge*, 424 U.S. 319, 335 (1976) (establishing framework for due process analysis).

43. *State v. Loomis*, 881 N.W.2d 749, 760-66 (Wis. 2016) (upholding use of algorithmic risk assessment in sentencing subject to strict conditions including judicial oversight and limitations on algorithm's role).

44. The criminal sentencing context involves different constitutional stakes than private arbitration. Criminal defendants possess Sixth Amendment rights to confront witnesses and to a jury trial that do not apply in arbitration. Nonetheless, Eric Loomis's reasoning about algorithmic transparency and judicial oversight provides a useful framework for thinking about AI in adjudicative contexts more broadly. If due process concerns constrain algorithmic use even in advisory roles for criminal sentencing, similar concerns apply a fortiori when algorithms serve as primary decision-makers in binding arbitration.

45. *Goldberg v. Kelly*, 397 U.S. 254, 267 (1970).

46. *Stern v. Marshall*, 564 U.S. 462, 484 (2011) (emphasizing that Article III reserves the judicial power to courts capable of exercising independent judgment).

47. *Northern Pipeline Constr. Co. v. Marathon Pipe Line Co.*, 458 U.S. 50, 58-59 (1982) (plurality opinion) (discussing the constitutional significance of adjudication by Article III judges).

48. *Thomas v. Union Carbide Agric. Prods. Co.*, 473 U.S. 568, 589 (1985) (upholding non-Article III adjudication where parties consented and adequate judicial review remained available).

49. *CFTC v. Schor*, 478 U.S. 833, 851 (1986) (holding that parties may consent to non-Article III adjudication but such consent does not eliminate all constitutional constraints).

50. Imre Stephen Szalai, *Stranger Disputes: When Artificial Intelligence Turns Arbitration Upside Down*, 25 *Pepp. Disp. Resol. L.J.* 133, 150 (2025).

51. See Imre Stephen Szalai, *Outsourcing Justice: The Rise of Modern Arbitration Laws in America* at 45-67 (2013) (documenting FAA's drafting history).

52. H.R. Rep. No. 68-96, at 1-2 (1924).

53. *Stolt-Nielsen*, 559 U.S. at 681.

54. *First Options of Chi., Inc. v. Kaplan*, 514 U.S. 938, 943 (1995).

55. The party autonomy argument for AI arbitration has attracted scholarly attention from multiple perspectives. Professor David Horton reaches a conclusion similar to this article through textualist analysis but proposes a different remedial framework: because the FAA only applies to adjudication by a “person,” state law fills the gap, allowing individual jurisdictions to decide whether to permit robot arbitration. David Horton, *Forced Robot Arbitration*, 109 *Cornell L. Rev.* 679, 685 (2024). Professor Dimitrios Ioannidis, by contrast, argues that federal arbitral laws do not prohibit AI arbitrators and that party autonomy should be the decisive factor supporting their use. Dimitrios Ioannidis, *Will Artificial Intelligence Replace Arbitrators Under the Federal Arbitration Act?*, 28 *Rich. J.L. & Tech.* 505, 587 (2022). We engage with Professor Szalai’s argument because it represents the most developed purposivist case for AI arbitration under the FAA.

56. *Hall Street*, 552 U.S. at 584.

57. *Loper Bright Enters. v. Raimondo*, 603 U.S. 369, 400 (2024) (holding that courts must determine statutory meaning rather than “speculating about the desires of hypothetical reasonable legislators”).

58. *Bostock v. Clayton County*, 590 U.S. 644, 653 (2020).

59. *New Prime Inc. v. Oliveira*, 586 U.S. 105 (2019).

60. 2025 *Queen Mary-White & Case International Arbitration Survey* 32-35 (2025).

61. *Id.* at 28-31.

62. The Chartered Institute of Arbitrators issued its own Guideline on the Use of AI in Arbitration in March 2025, which similarly emphasizes human decision-making authority. Article 8.2 provides that “[a]rbitrators should not relinquish their decision-making powers to AI but may use AI to support more accurate and efficient processing of submitted information, always ensuring independent judgment.” Chartered Inst. of Arbitrators, *Guideline on the Use of AI in Arbitration* Art. 8.2 (2025). The convergence of institutional guidance around this principle suggests an emerging consensus that AI assistance is appropriate only when human arbitrators retain genuine decisional authority.

63. Exec. Order No. 14,110, 88 Fed. Reg. 75,191 (Oct. 30, 2023) (establishing federal AI governance framework)(subsequently revoked Jan. 2025).

64. National Inst. of Standards & Tech., *Artificial Intelligence Risk Management Framework (AI RMF 1.0)* 23-31 (Jan. 2023).

65. The requirement of explainability finds support in regulatory frameworks. The EU’s Artificial Intelligence Act classifies AI systems used in “the administration of justice and democratic processes” as high risk, requiring human oversight, transparency, and record-keeping. Regulation (EU) 2024/1689, arts. 6, 14, 2024 O.J. (L 1689). While the EU AI Act does not directly govern U.S. arbitration, its framework reflects a growing international

consensus that adjudicative AI requires heightened safeguards. Any U.S. legislative framework should consider these emerging international standards.

66. *Mitsubishi Motors Corp. v. Soler Chrysler-Plymouth, Inc.*, 473 U.S. 614, 628 (1985).

67. *Gilmer v. Interstate/Johnson Lane Corp.*, 500 U.S. 20, 30 (1991).