

**WE BUILT JUDGE.AI.
AND YOU SHOULD BUY IT.**

KIMO GANDALL, JACK KIEFFABER, KENNY MCLAREN¹

We just built Corporate America's dream and Asimov's nightmare—a machine that automates all dispute resolution. Meet Arbitrus; it's Judge.ai for private arbitration, and we're going to tell you how it works.

Predictable, efficient, and cheap—very cheap—Arbitrus.ai performs just as well as human arbitrators right now, reducing the cost of dispute resolution from \$100k to a flat \$10k. And with help from laissez faire tech regulation and the permissive Federal Arbitration Act, it sets the table for the Arbitration State—a looming system of mass private adjudication that will gradually outcompete standard public judiciaries and decrease the relevance of centralized governance.

In Part I, we deliver a primer on our theory of automated judging. In Part II, we present a technical outline of Arbitrus.ai, a full-stack end-to-end AI arbitration service. In Part III, we then apply Arbitrus.ai to 100 hypothetical scenarios, and grade the opinion of the AI arbitrator. We find zero (0) hallucinations; zero (0) incomplete responses; and only two (2) ungrounded sources. In Part IV, we address positive and normative objections to Arbitrus' efficacy and ethics. And finally, in Part V, we talk about consequences—and the steps between our machine and total automated governance.

There are fewer steps than you'd think; join us as we tip over the first domino.

¹ Intellectual property of Fortuna Arbitration, Inc. US Patent Application Serial Number 63/745,277. Kimo Gandall is a third-year student at Harvard Law School, J.D. 2025 and CEO of Fortuna Arbitration; Jack Kieffaber is a Harvard Law School graduate, J.D. 2023; Kenny McLaren is the CTO of Fortuna Arbitration. Heartfelt thanks to the great Stephen Sills, Matias Mayesh, and the rest of the Joseph Story Society for engaging in this debate and bringing this piece to fruit. Special thanks again to the Fortuna-Insights team, who are doing things that have never been done before.

INTRODUCTION 3

I. FROM JUDGE.AI TO ARBITRUS.AI: JUDICIAL FUTURISM FOR THE PRIVATE SPHERE 9

II. HOW ARBITRUS.AI WORKS 20

 A. Deep Neural Networks and Large Language Models..... 20

 B. Benchmarking Legal AI 26

III. METHODOLOGY: ARBITRUS.AI AT WORK..... 30

 A. Human User Interface 32

 B. AI Architecture..... 35

 C. Data Collection 47

 1. Brief Procedure..... 48

 2. Opinion Procedure..... 49

 3. Grading Procedure..... 50

 4. Selected Cases for Qualitative Review 56

 D. Results..... 56

IV. OBJECTIONS..... 72

 A. Some Positive Objections 72

 B. Some Normative Objections 78

V. CONCLUSION: TOWARDS THE *ARBITRATION STATE* 82

[Society] was always at the mercy of economic and sociological forces it did not understand—at the whims of climate, and the fortunes of war. Now the Machines understand them; and no one can stop them, since the Machines will deal with them as they are dealing with the Society...

ISAAC ASIMOV, I, ROBOT 223 (1950).

INTRODUCTION

Predictable,² efficient,³ and *cheap*⁴—*very cheap*.⁵
Automated arbitration—what we call here the *Arbitration*

² See Meriam Al-Rashid & Ulyana Bardyn, *The Role of Artificial Intelligence in International Arbitration*, GLOBAL ARBITRATION REVIEW (Jul. 25, 2019), [https://globalarbitrationreview.com/review/the-arbitration-review-of-the-americas/2020/article/the-role-of-artificial-intelligence-in-international-arbitration#:~:text=Against%20this%20backdrop%2C%20AI%20is%20now%20poised%20to; see also Kimo Gandall et al., Predicting Precedent: A Psycholinguistic Artificial Intelligence in the Supreme Court, 14 CASE W. RES. J.L. TECH. & INTERNET 220, 251-255 \(2022\) \(finding that AI can predict case outcomes at the US Supreme Court correctly 91% of the time, whereas humans as little as 75%\).](https://globalarbitrationreview.com/review/the-arbitration-review-of-the-americas/2020/article/the-role-of-artificial-intelligence-in-international-arbitration#:~:text=Against%20this%20backdrop%2C%20AI%20is%20now%20poised%20to; see also Kimo Gandall et al., Predicting Precedent: A Psycholinguistic Artificial Intelligence in the Supreme Court, 14 CASE W. RES. J.L. TECH. & INTERNET 220, 251-255 (2022) (finding that AI can predict case outcomes at the US Supreme Court correctly 91% of the time, whereas humans as little as 75%).)

³ *Id.*

⁴ Isabel Gottlieb, *AI Will Make Arbitration Faster and Cheaper, Group CEO Says*, BLOOMBERG LAW (Dec. 6, 2023), <https://news.bloomberglaw.com/artificial-intelligence/ai-will-make-arbitration-faster-and-cheaper-group-ceo-says>.

⁵ This is especially true for cases where the amount in controversy is below \$1,000,000. Suppose the average mid-range arbitrator charges \$500 per hour in a case below \$75,000 in controversy. Suppose an arbitrator is incredibly efficient, and handles all discovery—ranging from objections to review—in three hours. Suppose again the efficient arbitrator handles briefs in five hours, the hearing in three hours, and drafts the entire final opinion in ten hours. That’s \$10,500 for the whole proceeding, or 14% of the damages gone into the dispute resolution process, not including the attorney or administrative fees—and that is nevertheless a number still shockingly lower than average. Arbitrus.ai can do it, well, for *half*—and that includes the administrative fees. See *infra* note 133.

State—will gradually subsume large swathes of formerly public litigation, leaving the government with a much smaller sphere of relevant influence.⁶ So small, that is, that it might itself prove manageable by a machine that will have already perfected its craft through billions of private trials. Meet Arbitrus.ai. It’s not coming. *It’s here.*⁷

Last year, one of us (Kieffaber) put out an article about “Judge.ai”—and went around the country getting yelled at for calling that contraption “the platonic end result of the textualist project.”⁸ Judge.ai is a hypothetical legal system where all law is statutory and all statutes are interactive—such that, before you perform an action, you can ask the statute whether that action is legal and the statute,⁹ powered

⁶ At the moment, our system is tied between an endless struggle of the private and public; and in the current status, the United States is ruled by a ‘restless instability’. We aim to provide the ultimate, virtual omnipresent governance solution to that question. *See* ADRIAN VERMEULE, COMMON GOOD CONSTITUTIONALISM 58 (2022) (“The ‘Monstrous Government’... this sort of regime is tossed restlessly between domination by private actors, often abusing their legal entitlements, on the one hand, and on the other hand, futile attempts at reassertion of ruling authority by public actors. In such conditions, any possibility of purposive public action for the common good is lost.”); *id.* at 57 (“... there is also a nightmarish version of the alternation in power in which there is an indefinite cycle of successive victories and defeats by factions locked in an endless struggle for rule.”); *id.* (“[it is]... *the confusion of public and private functions...*”). On a completely separate note, “systems that resolve litigation,” are not only exclusively tied to the judiciary—indeed, Article II courts play an active role in important cases. *See, e.g.,* SEC v. Jarkesy, 144 S. Ct. 2117, 2125 (2024) (“...when the SEC adjudicates the matter in-house, there are no juries. Instead, the Commission presides and finds facts while its Division of Enforcement prosecutes the case.”).

⁷ ARBITRUS.AI, <https://www.arbitrus.ai/> (last visited Jan. 17, 2025).

⁸ Jack Kieffaber, *Predictability, AI, And Judicial Futurism: Why Robots Will Run The Law And Textualists Will Like It*, 48 HARV. J.L. & PUB. POL’Y (forthcoming 2025) (manuscript at 34), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4966334.

⁹ Kieffaber, *supra* note 8, at 7. Arbitrus.ai does not currently perform this task, despite technically possessing deterministic

by artificial intelligence, will tell you in advance with an *advisory order*.¹⁰ The answer the interactive statute tells you on the front end is exactly what it will rule on the back end should you perform the action.¹¹ Goodbye judge, goodbye lawyer, goodbye legal profession; the statute itself usurps all of these roles.¹² Hello *perfect predictability*—won at the total expense of humanity’s role in justice.¹³

qualities, because the Federal Arbitration Act prohibits “partiality.” 9 U.S.C. § 10. Whether the mere issuance of an *ex ante* advisory order whose legal logic exists regardless of the *ex parte* communication would violate the Federal Arbitration Act’s prohibition on “partiality” is a question of law not resolved. *Cf.* JCI Communs., Inc. v. IBEW, Loc., 324 F.3d 42, 51 (1st Cir. 2003) (“Evident partiality is more than just the appearance of possible bias. Rather, evident partiality means a situation in which ‘a reasonable person would have to conclude that an arbitrator was partial to one party to an arbitration.’”). This is particularly a difficult issue because state law can create additional partiality requirements. *See, e.g.*, Nat’l Union Fire Ins. Co. of Pittsburgh, Pa. v. Seneca Family of Agencies, 255 F. Supp. 3d 480 (S.D.N.Y. 2017) (explaining that California law can create additional disclosure requirements on the Federal Arbitration Act). More on this in Part IV, *infra*.

¹⁰ Kieffaber, *supra* note 8, at 36-38. Kieffaber originally referred to “advisory opinions” in the Judge.AI context, but has since realized that “advisory order” is more apt since the machine would not be issuing “opinions.” An “opinion” implies some amount of subjectivity or at least margin for error; insofar as subjectivity is a lack of predictability, AI allows for neither—unless one intentionally allows for a change in “temperature” (that is, the AI’s capacity to predict non-probable outcomes).

¹¹ Kieffaber, *supra* note 8, at 32.

¹² *Id.* at 35-36.

¹³ And this is a good thing, at least for textualists. *See id.* at 29-30.

Naturally, that article got some pushback.¹⁴ And the chief rebuttal was: *That will never work.*¹⁵ We heard it from academics;¹⁶ members in the field;¹⁷ and for two of us

¹⁴ *Predictability, AI, and Judicial Futurism with Jack Kieffaber*, THE JAMES WILSON INSTITUTE, https://www.anchoringtruths.org/2024/09/26/predictability-ai-and-judicial-futurism-with-jack-kieffaber/?fbclid=IwY2xjawHn-tNleHRuA2FlbQIxMAABHYTvxbyY0KODVCF2s9MYrlEp1wCRNPsi9qGY_3G5jEj3Yv7aGsSnt5aL-Cg_aem_uzNKAF3xni75p7bjOcXPjw (last visited Jan. 17, 2025) (“One of the things that come through ... you are not positing a [far off] future. You are talking about the relatively fast transformation of an enduring institution ... in the span of 5 to 10 years....”).

¹⁵ *Id.*

¹⁶ See Kevin Frazier, *A Different Alignment Problem: AI, the Rule of Law, and Outdated Legal Institutions and Practices*, 19 J. BUS. & TECH. L. 331 (2024) (“...artificial intelligence (‘AI’), emerging technologies have the potential to improve, as well as impair, the Rule of Law.” Generative AI tools, for example, can assist pro se litigants with legal research thereby increasing the ability of the public to fully and forcefully defend themselves in court. However, those same tools can ease the creation and dissemination of misinformation and disinformation, which causes the public to question whether the law is fairly and accurately established and enforced.”).

¹⁷ Jerry Gewirtz, *Artificial Intelligence May Assist, but Can Never Replace, the Judicial Decision-Making Process of Human Judges*, THE FLORIDA BAR (Dec. 2024), <https://www.floridabar.org/the-florida-bar-journal/artificial-intelligence-may-assist-but-can-never-replace-the-judicial-decision-making-process-of-human-judges/#:~:text=Home-Artificial%20Intelligence%20May%20Assist%2C%20but%20Can%20Never%20Replace%2C%20the%20Judicial,Making%20Process%20of%20Human%20Judges&text=Artificial%20intelligence%2C%20in%20its%20many,available%20to%20the%20legal%20profession> (“Accordingly, the fallibility of artificial intelligence programs, combined with the admonition of Chief Justice Roberts that any use of artificial intelligence requires caution and humility, provides a further compelling reason why robotic judges can never replace human judges at the trial or appellate levels.”).

(Gandall and McLaren) from numerous venture capitalists.¹⁸

To be sure, that rebuttal is really two in one:

1. The machine itself wouldn't work.¹⁹
2. Even if it *did* work, it would be chaos to implement in America.²⁰

Those are tough arguments; to really refute them, you'd basically have to *build Judge.ai*. So, two of us (Gandall and McLaren) did—and today we're all going to tell you how it works.

Now, the Gandall/McLaren machine is not called "Judge.ai." And that's because Rebuttal 2 is correct; nobody will ever throw a switch and spontaneously consolidate the federal judiciary into one robot. Rather, the machine will cut its teeth on the private market by automating *arbitrations*. We hope to show that this humble beginning won't be so humble at all. Combine the increasing percentage of our actions (speech, commerce, employment) that occur via private platforms,²¹ with the sweeping deregulatory mandate of the

¹⁸ Cf. MARC RANDOLPH, *THAT WILL NEVER WORK* 133 (2019) ("...startup is a lonely place. You are working on something that no one believes in, that you've been told time and time again will never work.").

¹⁹ See Gewirtz, *supra* note 17.

²⁰ To be sure, these both fight the Judge.AI hypothetical mightily. But we get it: A Star Wars future is a *pretty big hypothetical* to swallow at face value.

²¹ This trend will likely continue, given the state of Section 230. See Daphne Keller, *Lawful but Awful? Control over Legal Speech by Platforms, Governments, and Internet Users*, U. CHI. L. REV. ONLINE (2022) ("Congress enacted Section 230 in 1996. The law was designed to let platforms set their own rules for speech while simultaneously avoiding government censorship."); see also Jack M. Balkin, *Free Speech in the Algorithmic Society: Big Data, Private Governance, and New School Speech Regulation*, 51 U.C. DAVIS L. REV. 1149, 1155 (Feb. 2018) ("The goal of the Algorithmic Society is practical omniscience: that is, the ability to know as much as possible about who is doing what, when, and where; and the ability to predict who will do what, when, and where. But the possibility of practical omniscience raises an important question of politics: who will possess and control this

Federal Arbitration Act,²² and you'll witness the creation of the *Arbitration State*—a private legal system that allows consumers and businesses to actively customize the set of laws that bind them.²³

What follows is equal parts legal scholarship and intellectual property. Part I will recap the Judge.ai theory, reiterate its basis in legal formalism, and revise it for the arbitration space. Part II will be a pure technical exposition, explaining the cocktail of classical machine learning, human-drive reinforcement training, and AI models that make Arbitrus.ai *run*. Part III will demonstrate that it *works*, reporting the findings of a demonstration in which the machine grappled with 100 hypothetical arbitration scenarios. Part IV will address objections. And Part V will close with the broader implications of the Arbitration State.

If you're a lawyer, read Parts I, IV, and V. If you actually want to make some money, read II and III.

omniscience?"). One of us (Kieffaber) has warned elsewhere about the nascent Web 3.0 age as a precolonized new world of entrenched serfdom; automated deciders make that arrangement all the more likely. Jack Kieffaber, *Spontaneous Disorder: The Protected Sphere and the Coming Web 3.0 Age*, HARV J.L. & PUB. POL'Y: PER CURIAM (Apr. 14, 2023) (“[T]he Web 3.0 Metaverse has the makings of a Crolyan frontier.... However, the Metaverse differs from gold rush era California in a salient way: the old-world strong are already there.”).

²² The Federal Arbitration Act provides a broad mandate for private parties to govern their own dispute resolution processes. 9 U.S.C. §§ 1-16, 2 (“... an agreement in writing to submit to arbitration an existing controversy arising out of such a contract, transaction, or refusal, shall be valid, irrevocable, and enforceable, save upon such grounds as exist at law or in equity for the revocation of any contract.”); *id.* at 5 (“If in the agreement provision be made for a method of naming or appointing an arbitrator or arbitrators or an umpire, such method shall be followed...”).

²³ Thomas J. Stipanowich, *Arbitration: The New Litigation*, 2010 U. ILL. L. REV. 1, 58 (“At the same time, the impetus to make arbitration even more like trial is evident in customized contractual provisions for enhanced judicial review.”).

I. FROM JUDGE.AI TO ARBITRUS.AI: JUDICIAL FUTURISM FOR THE PRIVATE SPHERE

We'll start with our bedrock premises. First, Arbitrus.ai is a better and more feasible manifestation of Judge.ai than Judge.ai itself. Second, formalists should want it to arbitrate all private contractual disputes in the United States. See, a contract is a pocket leviathan; just as men make law to remedy contextual shortcomings of physics,²⁴ they make contracts to remedy contextual shortcomings of law.²⁵ What results is a constitution for a very small country, and the chief obstacle for these private constitutions is the same one that befalls America's: *Ambiguity*.²⁶ That obstacle leads to the one vice

²⁴ Compare THOMAS HOBBS, LEVIATHAN 424 (Penguin Classics 1982) (1651) (“...If such metaphysics and physics as this be not vain philosophy, there was never any; nor needed St. Paul to give us warning to avoid it.”), with THOMAS AQUINAS, SUMMA THEOLOGIAE II–II, q. 95, a. 2. (“Consequently every human law has just so much of the nature of law, as it is derived from the law of nature. But if in any point it deflects from the law of nature, it is no longer a law but a perversion of law.”).

²⁵ Daniel Markovits & Emad Atiq, *Philosophy of Contract Law*, STAN. ENCYCLOPEDIA OF PHIL.. (Winter 2021 Edition), <https://plato.stanford.edu/archives/win2021/entries/contract-law/>. But see Ruth L. Okediji, *Contracts, Persons and Property: A Tribute to Margaret Jane Radin*, 22 MICH. TELECOMM. & TECH. L. REV. 143, 145-46 (2015) (“But if the legal weight attached to a contract is firmly anchored in the idea that its provisions are self-imposed through the exercise of a person's free choice, then standardized contracts that consist principally of non-dickered terms raise difficult moral and conceptual questions regarding what we consider an ‘agreement.’”).

²⁶ Lawrence M. Solan, *Pernicious Ambiguity in Contracts and Statutes*, 79 CHI.-KENT L. REV. 859 (2004).

which surpasses Hell and Holocaust²⁷ as the formalist's greatest fear: *Unpredictability*.²⁸

Kieffaber pitched “Judge.ai” as the theoretical cure for unpredictability in public laws—but he wasn’t merely suggesting the status quo starring Grok as John Roberts.²⁹ Rather, he posited a new governmental regime built *from scratch* to accommodate a central AI decider.³⁰ Said regime would set out all its laws in unambiguous algorithmic language such that a machine could execute them mathematically and transform the statute into an interactive entity.³¹ To the extent such precise enumeration proved impossible, the interactive statutes would make judgment calls at the margins and compensate for any resulting randomness by telling the polity up front, when queried, what its ruling would be if X action were taken.³² If the polity disliked that *ex ante* answer, it would be free to petition the code-writing legislature to “patch” the statute and fill the holes that Judge.ai imperfectly or even nonsensically

²⁷ Literally, at least as supposed by Scalia. See Kieffaber, *supra* note 8, at 18-19 (citing CHRISTOPHER SCALIA & EDWARD WHELAN, SCALIA SPEAKS: REFLECTIONS ON LAW, FAITH, AND LIFE WELL LIVED 248 (2017)) (“Where the positive law places a judge in the position of being the instrument of evil, the judge must recuse from the case or (if there are many such cases) resign from the bench. Thus, if I were a judge in Nazi Germany, charged with sending Jews and Poles to their death, I would be obliged to resign my office....”).

²⁸ See Kieffaber, *supra* note 8, at 19 n.48 (citing Madison Debates June 4, THE AVALON PROJECT AT YALE L. SCH., https://avalon.law.yale.edu/18th_century/debates_604.asp (last visited Jan. 17, 2025)) (“MR. GERRY. The evils we experience flow from the excess of democracy. The people do not want virtue, but are the dupes of pretended patriots One principal evil arises from the want of due provision for those employed in the administration of [Government].”).

²⁹ Grok is X’s (formerly Twitter’s) open source AI. See, e.g., GROK, [x.AI](#) (last visited Jan. 17, 2025).

³⁰ See Kieffaber, *supra* note 8, at 8–27.

³¹ *Id.* at 6 n.17.

³² *Id.* at 13.

attempted to fill.³³ Over time, these frequent amendments would work the statutes pure;³⁴ and until purity came, the law—warts and all—would at least be perfectly predictable. Formalists of the world rejoice.³⁵

From this rather fanciful foundation, we make a small leap: *That would work for private governance too.*³⁶ Indeed, it would work *better*—and *sooner*.

For one, that’s because there’s a unique demand for efficiency and predictability in the arbitration sphere.³⁷ While arbitration is on balance faster and cheaper than litigation, it’s still expensive—the average enterprise arbitration case costs more than \$100,000 between administrative fees and the

³³ *Id.* at 7, 52–53.

³⁴ That is, the culture of amendment would eventually make the law as written an exact mirror of the polity’s up-to-the-minute desires.

³⁵ For a full brief of the argument, see Kieffaber, *supra* note 8, at 11 n.29.

³⁶ Indeed, we find it salient that the “big bang” for automated adjudication—Judge Kevin Newsom’s concurrence in *Snell v. United Specialty Ins. Co.*—was not a statutory interpretation case, but rather a *contract* case. 102 F.4th 1208 (11th Cir. 2024) (Newsom, J., concurring). And that was a linchpin of the concurrence; because Judge Newsom was interpreting a recently-drafted contract—rather than a decades-old statute or centuries-old constitution—he could credibly tout ChatGPT’s mastery of “ordinary meaning” as gleaned from *internet inputs*. *Id.* at 1226. ChatGPT’s hyper-modern data set could be a bug rather than a feature if we’re parsing the meaning of “speech” in 1788; for automated formalism to work, it helps for the machine to be roughly (or exactly) as old as the text it’s interpreting.

³⁷ Ritik Dhankhar, *Moving Towards an AI Oriented Arbitration: Significance and Challenges*, 6 INT’L J.L. MGMT & HUMANS, 2028, 2030-31 (2023), (“This would help in reducing biases, favouritism or corruption. The lengthy court proceedings may finally come to an end (at least in a small number of cases). Apart from increasing efficiency and bringing standardisation into the institution, the biggest impact this technology would have in the judicial system is that it would reduce the influential powers of the parties on the basis of their incomes.”).

arbitrator, to say nothing of attorney’s fees.³⁸ Top arbitrators can even exceed some law firm billables, charging over \$15,000 per day, or approximately \$2,000 per hour.³⁹ And arbitration is significantly more unpredictable than litigation.⁴⁰ That’s because of a trend we observed in conversations with litigators, one also reflected in the scholarly literature: Most arbitrators just *aren’t very good at their jobs*.⁴¹

That’s not because they’re inarticulate, uneducated, or negligent—and to be sure, there are some great arbitrators working today.⁴² But, all too often, the arbitrator is simply wired differently than a jurist—despite being asked, at bottom, to perform a jurist’s task.⁴³ Indeed, many arbitrators

³⁸ *International Arbitration Fee Schedule*, INT’L CTR. FOR DISP. RESOL., AM. ARB. ASSOC. (2023), <https://www.adr.org/sites/default/files/International-Arbitration-Fee-Schedule.pdf>.

³⁹ Deborah Rothman, *Trends in Arbitrator Compensation*, DISP. RESOL. MAG., DISP. RESOL., Spring 2017, at 8, 8.

⁴⁰ Kevin Mason, *Will Discovery Kill Arbitration?*, 2020 J. DISP. RESOL. J. 207, 207–08.

⁴¹ Arbitrators have been cited lacking insufficient technical expertise, and empirical studies have demonstrated arbitrator bias in favor of their appointing party. *See* James Sherk, *Federal Union Arbitrators Frequently Misapply the Law*, AMERICA FIRST POLICY INSTITUTE (2023); *see also* Sergio Pui & Anton Strezhnev, *Affiliation Bias in Arbitration: An Experimental Approach*, 46 J. LEGAL STUDIES 371, 372-74 (2017); Andre de A. Cavalcanti Abud, *Criticism of Arbitration: How to Use it*, KLUWER ARB. BLOG. (Dec. 20, 2016), <https://arbitrationblog.kluwerarbitration.com/2016/12/20/criticism-of-arbitration-how-to-use-it/>.

⁴² Brian R. Jerome, *Musings of a Longtime Arbitrator*. MASS. DISP. RESOL. SERVS. (2024), <https://www.mdrs.com/faqs/mdrs-articles/musings-of-a-long-time-arbitrator/>.

⁴³ *Id.* The ‘Double-hatting’ problem—wherein an arbitrator operates as counsel—can also create problems in agency, especially in the international investor-arbitration context. That is, arbitrators often are also consultants or counsels to major institutions involved in these arbitrations, suggesting a source of partiality. *See* Tamar Meshel, *International Arbitration: The New*

have never served as a legal judge.⁴⁴ In the worst case, your arbitrator is just a regular lawyer looking to make a buck—which is a competence lottery.⁴⁵ In the better case, your arbitrator is an industry expert.⁴⁶ These experts at least understand the scenarios at issue, but are often more concerned with maintaining broader norms across a profession than crunching the plain text.⁴⁷ Whatever the merits of a professionalist approach, its ends will be more subjective—and its results less predictable—than what parties could expect from an increasingly formalist state or federal judiciary.⁴⁸ And then there’s just the *rate of error*; a dishearteningly large number of arbitrators simply misapply

Frontier of Business and Human Rights Dispute Resolution?, 44 DALHOUSIE L.J. 101, 121 (2021) (“There have been recent debates surrounding the ‘revolving door’ phenomena in international arbitration, also known as ‘double-hatting,’ ie, individuals who act sequentially and even simultaneously as arbitrator, legal counsel, expert witness, or tribunal secretary.”). This is the same concern the Founders had in creating structures that aimed to minimize the influence of factions. See THE FEDERALIST No. 10 (James Madison) (“No man is allowed to be a judge in his own cause, because his interest would certainly bias his judgment, and, not improbably, corrupt his integrity. With equal, nay with greater reason, a body of men are unfit to be both judges and parties at the same time; yet what are many of the most important acts of legislation, but so many judicial determinations, not indeed concerning the rights of single persons, but concerning the rights of large bodies of citizens?”).

⁴⁴ Gary L. Benton, *Arbitrators are Not Judges*, SILICON VALLEY REDISTRIB. & MEDIATION CTR. (Jan. 4, 2018), <https://svamc.org/arbitrators-are-not-judges/>.

⁴⁵ Rothman, *supra* note 39, at 8.

⁴⁶ Thomas Stipanowich & Zachary Ulrich, *Arbitration in Evolution: Current Practices and Perspectives of Experienced Commercial Arbitrators*, 25 AM. REV. INT’L ARB. 396, 397 (2014).

⁴⁷ Rothman, *supra* note 39, at 8.

⁴⁸ Lauren Zimmerman & Jeff Zalesin, *Five Ways to Improve Commercial Arbitration Clauses*, SELENDY GAY (Aug. 16, 2023), <https://www.selendygay.com/news/publications/2023-08-16-five-ways-to-improve-commercial-arbitration-clauses>.

relevant judicial precedents⁴⁹ and bungle the contracts they're hired to interpret.⁵⁰ At some level, you get what you pay for—even though you're still paying *a lot*.⁵¹

⁴⁹ Empirical work on determining if arbitrators frequently misapply the law is difficult, given most arbitration cases are held confidential. Michael Collins, *Privacy and Confidentiality in Arbitration Proceedings*, 11 ARB. INT'L 321, 324 (1995), <https://doi.org/10.1093/arbitration/11.3.321>. There are, however, examples—for example, in the federal employee arbitration context. *See, e.g.*, Sherk, *supra* note 41 (“For both the George W. Bush and Barack Obama administrations, the author examined a sample of 30 arbitration awards in which unions initially prevailed and the agency appealed that decision to the Authority. The Authority upheld those arbitration awards in 57% of cases in the Bush Administration and 50% of cases under President Obama. *Federal arbitrators frequently misapply the law.*”) (emphasis added).

⁵⁰ *See* W. C. Mark Weidenmaier, *Judging-Lite: How Arbitrators Use and Create Precedent*, 90 N.C. L. REV. 1091 (2012) (“Arbitrators routinely write lengthy awards that are substantially devoted to legal analysis and that often make extensive use of precedent.”); W. C. Mark Weidenmaier, *Toward a Theory of Precedent in Arbitration*, 51 WM. & MARY L. REV. 1895, 1925-26 (2010). This phenomenon is also noted in practice, with arbitrators often referencing each other. *Id.* at 1911 (“Surveys of labor arbitrators, reports of cases in which arbitrators have relied on prior awards, and a modest body of empirical evidence all support this belief.”). And even if we are right, and the arbitrator’s and misapplying legal theory, there is nothing the courts can do about it on the merits. *See* *Matter of Nexia Health Techs., Inc. v. Miratech, Inc.*, 176 A.D.3d 589, 591 (N.Y. App. 2019) (“Manifest disregard of the law is a ‘doctrine of last resort limited to the rare occurrences of apparent egregious impropriety on the part of the arbitrators’ and is ‘severely limited’”).

⁵¹ And sometimes, claimants actually pay *more* to arbitrate over litigation. *See* Christopher R. Drahozal, *Arbitration Costs and Form Accessibility: Empirical Evidence*, 41 U. MICH. J.L. REFORM 813, 817 (2008) (“For large claims the administrative fees can be substantially higher than court filing fees: a claimant bringing a \$500,001 claim in arbitration will have to pay a \$6000 filing fee and a \$2500 case service fee, while the fee for filing suit in federal court is \$350 regardless of the amount in dispute.”).

So: Arbitrus.ai is more feasible than Judge.ai in large part because there's an immediate market for it. And the fact that we can even invoke a market concept in the first place tees up another salient difference between judging and arbitration that makes Arbitrus.ai more realistic: Arbitration is a *business*. Indeed, there are arbitration firms just like there are law firms—the two largest in America are AAA and JAMS, and they employ stables of arbitrators who get paid to arbitrate *by the parties*.⁵² That means these companies are competing not just to get parties to arbitrate in the first place, but to arbitrate *with them*. The result is a fascinating world, foreign to the sacred high priesthood of judging, where jurisprudence and methodology are *commodities*; frequently, arbitral decisions will cite landmark federal precedents and equitable platitudes not because they're legally bound to,⁵³ but *because it makes them sound more learned than the competition*. In short: All the flowery excesses federal judges employ for *clout*, arbitrators employ to stay *fed and clothed*.⁵⁴

The upshot is that efficiency *really matters*; the bottom line is a much better gun to the head than the Speedy Trial Act or Due Process Clause. And, with competition on the table, automation is in a much better position to *actually compete*. Indeed, Kieffaber already posited that Judge.ai would outcompete human judging on the international market, bubbling up in smaller countries and gradually gaining

⁵² Obviously, they don't get paid *directly*. To maintain propriety, the arbitration firm takes the cut directly and trickles the money down across its employees. Any other arrangement would sound eerily like a bribe; arbitrators aren't working on tips.

⁵³ See *Oxford Health Plans LLC v. Sutter*, 569 U.S. 564, 569 (2013) (“Because the parties ‘bargained for the arbitrator’s construction of their agreement,’ an arbitral decision ‘even arguably construing or applying the contract’ must stand, regardless of a court’s view of its (de)merits.”).

⁵⁴ See, e.g., *Builders Mut. Ins. Co. v. GCC Constr., LLC*, No. 24-5152, 2024 WL 5074878, at *1 (6th Cir. Dec. 11, 2024) (Thapar, J.). We're getting a little jiggy with the bricks.

momentum⁵⁵—and the arbitration marketplace allows for that replacement to start *now*. This raises a third key advantage for Arbitrus: A contract *is* a small country.⁵⁶ For Judge.ai to work, a nation would quite literally need to rebuild its legal system from scratch and shape it around a central AI system.⁵⁷ That’s something even the Third World is unlikely to do without substantially more proof of concept.⁵⁸ The logistics are brutal; not only does an entire polity have to

⁵⁵ Upon the publishing of this paper, Arbitrus.ai will already have been contractually introduced to large institutions in Latin America, including Colegio de Abogados De Lima. Note that in Peru, Fortuna products are called “CaseCraft” because “Fortuna” is associated with gambling. *See* Convenio De Cooperación Interinstitucional Entre El Colegio De Abogados de Lima y Fortuna-Insights (CaseCraft) (2024), https://fortunainsights-my.sharepoint.com/:b:/g/personal/kimo_gandall_fortuna-insights_com/EXoqCHo58qZLnwf5IaX3cHABB_Iu638_LY0AMTQgQyiVxQ?e=B94hma.

⁵⁶ Well, not entirely. The sovereign can infringe and regulate the boundaries of the contract—but if those boundaries can have their meaningful, substantive merits defined in an obscured process that the sovereign has delegated (as 9 U.S.C. § 2 seems to do) the real governing power is private. *See* Markovits & Atiq, *supra* note 25 (“Contract law is based on the model (or idealization) of ‘voluntary exchange ... between autonomous individuals’ (Radin 2007: 196). Contracts are enforced on the assumption that parties possess the capacities necessary for informed and voluntary choice (R2: §15), while fraud, duress, and various forms of pressure in bargaining represent grounds for rescission (R2: §§164, 175, 177). A preliminary question is why the law should be so concerned with freedom and autonomy in contractual exchange. As previously discussed, the rearrangement of private property might be justified on grounds of efficiency or general welfare.... The law’s guiding principle could be hypothetical or reasonable consent—what parties would or should consent to—rather than actual consent. The importance of actual consent seems clear enough where exchange impinges on domains over which persons retain special sovereignty, including their own body and mind.”).

⁵⁷ Al-Rashid, *supra* note 2 (“Arbitrators can get a computer’s assessment of the merits of a dispute. And litigants are able to get their disputes decided faster and with more uniformity.”).

⁵⁸ *Id.*

agree on an algorithmic model of governance, it has to agree on *an algorithm* to use—which enters a morass of corporate favoritism and corruption.⁵⁹ And, in the face of all that upheaval, what’s the chance that a polity would rally around a “culture of amendment” and patiently petition congress to patch an algorithm?⁶⁰ To establish a Judge.ai regime all at once, you’d basically need a dictator—and you show us the dictator who would willingly cede his power to an impartial machine.⁶¹

But the country a contract presides over can have as few as a thousand people in it.⁶² Or ten. Or *two*.⁶³ It’s far easier for a group of that size to decide, collectively, that it might just be worth the effort to design their micro-nation around an AI

⁵⁹ See, e.g., Margarita Leib, *Corrupted by Algorithms? How AI-generated Human-written Advice Shape (Dis)honesty*, 134 ECON. J. 766, 782 (2024) (“We find that, when faced with the trade-off between honesty and money, people will use AI advice as a justification to lie for profit.”).

⁶⁰ KYLE CHAYKA, FILTERWORLD: HOW ALGORITHMS FLATTENED CULTURE 56 (2024) (“The force of algorithmic pressure is not theoretical.... On the consumer side, the bombardment of recommendations can induce a kind of hypnosis that makes listening to, watching, or buying a product all but inevitable—whether it truly aligns with your taste or not”); see also Madison Alder & Rebecca Heilweil, *Chevron’s Downfall Highlights Need for Clear Artificial Intelligence Laws*, FEDSCOOP (July 9, 2024), <https://fedscoop.com/chevron-downfall-highlights-clear-ai-law-need/>.

⁶¹ Raluca Csernatonu, *Can Democracy Survive the Disruptive Power of AI?*, MALCOLM H. KURR CARNEGIE MIDDLE E. CTR. (Dec. 18, 2024), <https://carnegieendowment.org/research/2024/12/can-democracy-survive-the-disruptive-power-of-ai?lang=en¢er=middle-east>.

⁶² Al-Rashid, *supra* note 2, at 23.

⁶³ RESTATEMENT (SECOND) OF CONTRACTS § 9 (AM. L. INST. 1981) (“There must be at least two parties to a contract, a promisor and a promisee, but there may be any greater number.”). Of course, we acknowledge that some contracts can also have billions and billions of parties. More on that rebuttal in Part IV, *infra*.

arbitrator.⁶⁴ Moreover, the stakes are lower; the parties can try automated governance,⁶⁵ decide they don't like it, and jettison it. They can try different companies' algorithms and stick with the one they like best.⁶⁶ And the culture of amendment necessary for the regime to run smoothly will be vastly more practicable when only a handful of parties—rather than a polity—must come to the table. The stakes are also lower in the *application*: If the machine messes up... *caveat emptor*.⁶⁷ Inherent here is the advantage of a contract over a

⁶⁴ *Id.* Note that, for behemoth companies, market monopolies may well effect the same result—that the company is free to unilaterally choose the machine without much input at all. Call it a dictatorship that doesn't require force; more on this sentiment in Part IV, *infra*.

⁶⁵ Niklas Berglind, Ankit Fadia & Tom Isherwood, *The Potential Value of AI — And How Governments Could Look to Capture It*, MCKINSEY & CO., (July 25, 2022),

<https://www.mckinsey.com/industries/public-sector/our-insights/the-potential-value-of-ai-and-how-governments-could-look-to-capture-it>; *see also Artificial Intelligence (AI) in Government*, INTEL,

<https://www.intel.com/content/www/us/en/learn/ai-in-government.html?>; *Transforming Government with AI and Automation*, KPMG, <https://kpmg.com/us/en/capabilities-services/advisory-services/government/intelligent-automation-in-the-federal-government.html?>.

⁶⁶ Berglind et al., *supra* note 65.

⁶⁷ In most jurisdictions, the US has shifted away from *caveat emptor*, or 'let the buyer beware,' to *caveat venditor*, or 'let the seller beware,' because the seller is in a better position to determine the warranties of the product. AI arbitration raises a larger legal question, that is, if two parties—neither of which are sellers—engage to use a service, which way should the liability lean? A future article could review the standard of expectations for these AI products. *See* Emanuel Emroch, *Caveat Emptor to Strict Liability: One Hundred Years of Products Liability Law*, 4 U. RICHMOND L. REV. 155, 155 (1970) ("As the industrial system has come of age and man has begun to make excursions into outer space, the ancient principle of caveat emptor-'let the buyer beware'-has been significantly changed in favor of the

Constitution: The parties bound *actually agreed to it*.⁶⁸ Indeed, they were alive when it was signed; to a point, anything they don't like about the micro-nation they live in is *their fault*.⁶⁹

consumer.”). But the overall point still holds: there are private law remedies for the worst abuses, ranging from unconscionability to concealment and fraud, similar to how the Founders realized state law allows for a diffusion of faction power that permits individual determinations of law. See THE FEDERALIST 10 (James Madison), private law allows determination and specificity that Constitutional law does not.

⁶⁸ Consent is a necessary stipulation of contracts. RESTATEMENT (SECOND) OF CONTRACTS § 18 (AM. L. INST. 1981). Constitutions, in contrast, have questionable assumptions of consent See RANDY E. BARNETT, RESTORING THE LOST CONSTITUTION 5 (2004) (“It is an open question whether the U.S. Constitution—either as written or as actually applied—is in fact legitimate. Intellectual honesty requires us to acknowledge the possibility that no constitution lacking unanimous consent is capable of producing laws that bind in conscience.”).

⁶⁹ Well, in theory. As some advocates point out, adhesion contracts—think the Facebook-type ‘terms of service’—question the truth value of this statement. Regardless of whether the practical regime supports this conclusion, the Restatement’s *de jure* command does. Compare Clifford Fisher et al., *Evolution of Clickwrap & Browsewrap Contracts*, 48 RUTGERS COMPUT. & TECH. L. J. 147, 165 (2022) (“This is important due to the nature of these online adhesion contracts wherein the user will accept the premade adhesion contract, which can often be misleading and trap users into the agreement.”), and Vera Bolgar, *The Contract of Adhesion — A Comparison of Theory and Practice*, 20 AM. J. COMPAR. L. 53, 53 (1972), and Allen Reames, *The Adhesion Contract of Insurance*, 5 SANTA CLARA LAW. 60, 60 (1964), with RESTATEMENT (SECOND) OF CONTRACTS § 50 (AM. L. INST. 1981) (“Acceptance of an offer is a manifestation of assent to the terms thereof made by the offeree in a manner invited or required by the offer.”), and RESTATEMENT (SECOND) OF CONTRACTS § 18 (AM. L. INST. 1981) (“Manifestation of mutual assent to an exchange requires that each party either make a promise or begin or render a performance.”), and RESTATEMENT (SECOND) OF CONTRACTS § 17 cmt. a (AM. L. INST. 1981) (“The element of agreement is sometimes referred to as a ‘meeting of the minds.’ The parties to most contracts give actual as well as apparent assent, but it is clear

These lower stakes make for a system that will allow automated judging to grow and develop without mass epistemic scrutiny.⁷⁰ And the general ineptitude of the average arbitrator is a key rejoinder here: Even the *growing pains* of Arbitrus.ai would perform *as well as* or *better* than most arbitrators today.⁷¹ When the same or better is available at a fraction of the cost, the market will take care of the rest. Businesses will *jump*—meaning the citizens of these nascent micro-nations will be willing participants in the experiment.⁷² In this way, Arbitrus.ai creates the practical conditions precedent for the coming era of automated jurisprudence.

II. HOW ARBITRUS.AI WORKS

The next two sections get *technical*. Bear with us. Or, if you're a lawyer, zone out and skip to Part IV—what follows is for businessmen and techies.

A. Deep Neural Networks and Large Language Models

Artificial intelligence, or AI, is a term introduced in 1955 by John McCarthy, an emeritus professor at Stanford.⁷³ It broadly describes “the science and engineering of making intelligent machines.”⁷⁴ For many, AI is most familiar through

that a mental reservation of a party to a bargain does not impair the obligation he purports to undertake.”).

⁷⁰ Nicola J. Bidwell, *Nic Bidwell's Thoughts on Epistemic Accountability for AI*, DEPT. OF COMPUT. SCI. & TECH., UNIV. OF CAMBRIDGE, (Dec. 29, 2021), <https://www.cst.cam.ac.uk/blog/afb21/nic-bidwells-thoughts-epistemic-accountability-ai>.

⁷¹ We have already noted the relative ineptness of arbitrators due to suboptimal proclivities that machines do not possess. See Pui & Strezhnev, *supra* note 41, at 373-74.

⁷² See *id.*

⁷³ John McCarthy, *What is Artificial Intelligence?*, STAN. UNIV. COMPUT. SCI. DEPT'T (Nov. 12, 2007), <https://www-formal.stanford.edu/jmc/whatisai.pdf>.

⁷⁴ *Id.*

applications like ChatGPT, which belong to a category of technologies known as “large-language models” (LLMs).⁷⁵ These LLMs are a specific type of Deep Neural Network (DNN),⁷⁶ which in turn are a subset of algorithms.⁷⁷

For the sake of the lawyers who are still with us,⁷⁸ let’s start with the basics. An algorithm is simply a set of well-defined instructions or procedures designed to accomplish a specific task.⁷⁹ Mathematical algorithms, like most quantitative methods, function by correlating independent variables, or “features,” with dependent variables, or “labels.”⁸⁰

Among the types of mathematical algorithms AI users are most likely to encounter are Deep Neural Networks (DNN).⁸¹

⁷⁵ Adam Allen Bent, *Large Language Models: AI’s Legal Revolution*, 44 PACE L. REV. 91, 114-115 (2023).

⁷⁶ Eman Jawad, *The Deep Neural Network — Review*, 9 J. MATHEMATICS, no. 9, 2023, at 1, 1-5, https://www.researchgate.net/publication/374151186_THE_DEEP_NEURAL_NETWORK-A_REVIEW?.

⁷⁷ *Detailed Explanation of Deep Neural Network, Multilayer Perceptron, and Deep Q-Network*, TURING, <https://www.turing.com/kb/explanation-of-deep-neural-network-multilayer-perceptron-deep-q-network>.

⁷⁸ Al-Rashid, *supra* note 2, at 1.

⁷⁹ ANDRIY BURKOV, *THE HUNDRED-PAGE MACHINE LEARNING BOOK 1* (2019) (“In supervised learning, the dataset is the collection of labeled examples ... each element x among N is called a feature vector. A feature vector is a vector in which each dimension $j = 1, \dots, D$ contains a value that describes the example somehow. That value is called a feature.”); *id.* at 19 (“Classification is a problem of automatically assigning a label to an unlabeled example... classification learning algorithm that takes a collection of labeled examples as inputs and produces a model that can take an unlabeled example as input and either directly output a label or output a number...”); *see also* Noson S. Yanofsky, *Towards a Definition of an Algorithm*, 21 J. LOGIC & COMPUTATION 253, 273 (2011).

⁸⁰ BURKOV, *supra* note 79, at 1-3.

⁸¹ *Id.* at 20 (“The notorious exceptions are neural network learning algorithms, specifically those that build neural networks with more than one layer between input and output. Such neural networks are

A DNN is termed “deep” because it includes multiple layers between its input and output layers.⁸² These intermediary layers, called hidden layers, are composed of interconnected nodes, or neurons, that process input data by creating tens of millions of correlations across a wide breadth of functions.⁸³ A DNN’s structure is determined by its “network architecture.”⁸⁴ Network architecture refers to the design of the network, encompassing the arrangement and types of layers, the number of neurons within each layer, the connections between layers, and any specific mechanisms or modules used to process and learn from data.⁸⁵ At its core, artificial intelligence often, but not always, relies on DNNs which are themselves a sophisticated class of algorithms.⁸⁶

LLMs, likewise, are a form of DNN that employs a network architecture called a “transformer.”⁸⁷ Transformers excel in handling text data and have become the foundation for many

called deep neural networks.”); *see also* Mason, *supra* note 40, at 208.

⁸² BURKOV, *supra* note 79, at 20.

⁸³ *Id.*

⁸⁴ *Id.* at 64; *see also* Samaya Madhavan & M. Time Jones. *Deep Learning Architectures: The Rise of Artificial Intelligence*, IBM DEV. DEVELOPER. (Apr. 25, 2024), <https://developer.ibm.com/articles/cc-machine-learning-deep-learning>.

⁸⁵ BURKOV, *supra* note 79, at 20-23.

⁸⁶ In the words of Burkov, artificial intelligence and “machine learning” are about “marketing.” *See* BURKOV, *supra* note 79, at xvii.

⁸⁷ *See* M. A. K. Raiaan et al., *A Review on Large Language Models: Architectures, Applications, Taxonomies, Open Issues and Challenges*, 12 IEEE 26839, 26840 (2024) (“The significant advancement in LLMs systems occurred when the transformer architecture was introduced in the seminal work. The transformer model is built around the self-attention mechanism, enabling parallelization and efficient handling of long-range dependencies.”); *see also* ANDRIY BURKOV, MACHINE LEARNING ENGINEERING 157 n.5 (2020) (“...transformers: validation performance first improves, then gets worse, and then improves again with increasing model size. As of July 2020, we don’t yet fully understand why it happens.”).

advanced models in natural language processing due to their ability to efficiently capture long-range dependencies via self-attention mechanisms—that is, using a transformer allows a model to pinpoint and connect related pieces of information across long distances in text by simultaneously considering all words and assigning proper weight to their relationships, even those humans would not otherwise recognize.⁸⁸

In an LLM, tokens are both the features and labels.⁸⁹ Here’s the basic process: LLMs will take any given text, tokenize it into discrete units (tokens), transform these units into embeddings containing a numerical vector value, and then are fed into the predicting model.⁹⁰ The tokens from the source are the feature; the token in the next sequence (think ‘sentence’) of the model is the label.⁹¹ These tokens serve as the features for the LLM, providing the necessary input to understand and generate text.⁹² *Contra* to the very many LinkedIn hype posts, the LLM is not ‘learning’ the truth value of any of these instances as mediated by experience—instead, the text itself is its own source of truth.⁹³ Said again, humans experience language in relation to their experience and purpose; the LLM learns language only in proximity to external language.⁹⁴ The LLM is like teaching a parrot to say “come in” when you knock on the door, learning by mere association. LLMs thus only understand language as a function of the next probable token. Type in “Knock, knock,”

⁸⁸ BURKOV, *supra* note 87, 65-66.

⁸⁹ *Id.* at 76, 82.

⁹⁰ *Id.* at 82. For example, imagine splitting amongst space. If the text is ‘My name is Kimo,’ the tokenization could be [‘My’, ‘name’, ‘is’, ‘Kimo’]. Actual tokenization algorithms, of course, are much more sophisticated and employ a frequency graph to decide on token vocabulary.

⁹¹ *Id.* at 82-83.

⁹² *Id.*

⁹³ Somewhere, Justice Gorsuch just smiled. *See* *Bostock v. Clayton Cnty.*, 140 S. Ct. 1731, 1738 (2020) (“After all, only the words on the page constitute the law adopted by Congress and approved by the President.”).

⁹⁴ BURKOV, *supra* note 87, at 83-85.

and ChatGPT invariably replies, “who’s there?” Whether or not this relationship indicates truth is a subject of debate.⁹⁵

We contend that an arbitral award can be produced in just the same probabilistic fashion. Here’s an analogy: suppose you are a judge trying to make a decision (the output) based on a series of briefs; you have a group of clerks, each with a different subset of knowledge.⁹⁶ Each clerk is a node, focusing on input information (the features)—such as case precedents or damage calculations. These clerks work collaboratively, sharing information that seeks a certain pattern of relationship. Each clerk then filters out information, repeating the process iteratively with other clerks (the layers), until a consensus is reached on the most probable answer. For a decision-making DNN, the label might be a discrete outcome based on what the court should do—affirm, reverse, or dismiss, for example. Alternatively, for a LLM it might be actually predicting the next sequence of tokens, which once finished, will constitute a whole case,⁹⁷ or in the case of arbitration, the whole dispute.

Of course the prior analogy presumes that the clerks are exclusively using the subset of data before them without

⁹⁵ Xuming Hu et al., *Do Large Language Models Know Facts?*, ARXIV (Oct. 8, 2024), at 1, <https://arxiv.org/abs/2310.05177#:~:text=Extensive%20experiment%20on%20different%20sizes,for%20realizing%20trustworthy%20artificial%20intelligence> (“Extensive experiments on different sizes and types of LLMs show that existing LLMs still lack factual knowledge and suffer from various spurious correlations. We believe this is a critical bottleneck for realizing trustworthy artificial intelligence.”).

⁹⁶ Just like with nodes, not all clerks are equally influential; in DNNs, we call this ‘weighting.’ Weights determine how much influence each node (e.g., each clerk’s opinion) has on the final outcome. See BURKOV, *supra* note 87, at 66, 183-84.

⁹⁷ Jinqi Laia et al., *Large Language Models and Law: A Survey*. ARXIV (Nov. 26, 2023), at 2-3, <https://arxiv.org/abs/2312.03718> (“Examples of legal applications of AI include legal research, document analysis, contract review, predictive analytics for case outcomes, and legal chatbots for providing basic legal information and guidance.”).

using any background knowledge. But nodes in a DNN lack an understanding of context or meaning not rooted in the text.⁹⁸ In fact, it is debatable whether LLMs are capable of “understanding” or reason as we know or conceptualize those terms at all.⁹⁹ While clerks can interpret ambiguous information and adapt their reasoning to the nuances of a case, nodes process input data blindly, relying entirely on numerical patterns.¹⁰⁰ This fundamental limitation necessitates the use of vast amounts of labeled data to train DNNs.¹⁰¹ Through repeated exposure to examples, nodes can approximate understanding by identifying patterns—but this process is far removed from the intuitive reasoning employed by human clerks.¹⁰²

Moreover, the great formalist strength of artificial intelligence—predictability—owes to the fact that the nodes constituting a DNN are confined to a predetermined network architecture (even if not transparent).¹⁰³ The connections between nodes and the roles of each layer are static, designed by engineers to suit a specific task.¹⁰⁴ Technologies such as reinforcement learning try to make these network structures more adaptable, but these remain far less dynamic as

⁹⁸ BURKOV, *supra* note 87, at 12-14.

⁹⁹ At some level, LLMs, like humans, develop inferences from pattern recognition. The question is whether an LLM is able to apply its recognition of patterns in a way similar to or distinct from humans, or, in other words, whether human problem solving goes beyond mere pattern recognition. For more research attempting to pierce the veil of an LLM’s “cognition”, see Shaham et. al., *A Multimodal Automated Interpretability Agent*, ARXIV (Apr. 22, 2024), <https://arxiv.org/abs/2404.14394>.

¹⁰⁰ BURKOV, *supra* note 87, at 97-99.

¹⁰¹ *Id.* at 38.

¹⁰² *Id.* at 99; see also *Artificial Intelligence Meets the C-Suite*, MCKINSEY Q., (Sept. 1, 2014), <https://www.mckinsey.com/capabilities/strategy-and-corporate-finance/our-insights/artificial-intelligence-meets-the-c-suite>.

¹⁰³ *Deep Neural Network (DNN) Explained*, ZOMEV (Apr. 14, 2024), <https://www.zomev.com/blog/deep-neural-network-dnn-explained/>; BURKOV, *supra* note 87, at 26.

¹⁰⁴ BURKOV, *supra* note 87, at 26-27.

compared to human-AI collaboration.¹⁰⁵ For example, even though LLMs can be adjusted for temperature, which is the frequency at which the LLM may select a less probable sequence of tokens, those less probable outcomes are still finite.¹⁰⁶ Thus, while artificial intelligence can mimic innovation and creativity, at its core it is just that: mimicry.¹⁰⁷

B. Benchmarking Legal AI

Benchmarking refers to measuring the application of AI in some use case compared to the counterfactual. Considerable efforts have been expended in benchmarking legal AI to determine its actual functionality.¹⁰⁸ LexisNexis and Westlaw have claimed that their AI products function to “mitigate, if not entirely solve,” problems with AI.¹⁰⁹

¹⁰⁵ *Id.* at 6-8; David Lindner & Mennatallah El-Assady, *Humans are Not Boltzmann Distributions: Challenges and Opportunities for Modelling Human Feedback and Interaction in Reinforcement Learning*, ARXIV (June 27, 2022),

<https://arxiv.org/abs/2206.13316>; see also Kumar Chandrakant, *Reinforcement Learning with Neural Network*, BALNDUNG (Mar. 18, 2024), <https://www.baeldung.com/cs/reinforcement-learning-neural-network>.

¹⁰⁶ In classification, at least. BURKOV, *supra* note 87, at 18-19.

¹⁰⁷ See *id.* at 85-86. Of course, some Marxist theories on jurisprudence might see jurisprudence as just that in practice: a mimicry of logic that retroactively justifies, supports, and reinforces the material conditions of the status quo.

¹⁰⁸ Magesh et al., *Hallucination-Free? Assessing the Reliability of Leading AI Legal Research Tools*, ARXIV (May 30, 2024), at 1, <https://arxiv.org/abs/2405.20362> (citing Dahl et al., *Large Legal Fictions: Profiling Legal Hallucinations in Large Language Models*, 16 J. LEGAL ANALYSIS 64, 64-65 (2024)).

¹⁰⁹ *LexisNexis Launches Lexis+ AI, a Generative AI Solution with Hallucination-Free Linked Legal Citations*, LEXISNEXIS (Oct. 25th, 2023),

<https://www.lexisnexis.com/community/pressroom/b/news/posts/lexisnexis-launches-lexis-ai-a-generative-ai-solution-with-hallucination-free-linked-legal-citations>; *Introducing AI-Assisted Research: Legal Research Meets Generative AI*, THOMSON REUTERS (Sept. 15, 2023), <https://legal.thomsonreuters.com/blog/legal-research-meets->

The primary metric for the functionality of an LLM in law is called a “hallucination”—widely known as the phenomenon occurring when an LLM “generate[s] incorrect or misleading information.”¹¹⁰ More concretely, these “hallucinations” often manifest as either inaccuracies or lacking in groundedness, or both. Accuracy is a mere question of whether the answer provided is true or false. Groundedness refers to whether the answer is supported by a valid, relevant source, independent of its accuracy.¹¹¹ For example, if one queries “Does the Constitution protect a right to same sex marriage?” a response that states “The right to same sex marriage is protected under the U.S. Constitution. *Obergefell v. Hodges*, 576 U.S. 644 (2015)” would be accurate.¹¹² But while a response that states “The right to same sex marriage is protected under the U.S. Constitution. *Miranda v. Arizona*, 384 U.S. 436 (1966)” would also be accurate, it would be *improperly grounded*.¹¹³ Likewise, a third metric—completeness—is about the LLM’s ability to articulate an answer that is relevant to all of the issues in controversy.¹¹⁴ While these metrics are not universally established, given the evolving nature of law and artificial intelligence, they provide a baseline understanding of operational legal AI.¹¹⁵

generative-ai/. Note that one of us (Kieffaber), in lectures, has disputed whether the human brain truly operates beyond mere probability and recombination. Are humans, in fact, creative? Do they actually make new things that did not previously exist? Or do they simply rearrange their various perceptions into different patterns and claim they’ve made something new?

¹¹⁰ Magesh et al., *supra* note 108, at 1.

¹¹¹ *Id.* at 5-6.

¹¹² *Id.* at 8.

¹¹³ *Id.*

¹¹⁴ There is some controversy to defining completeness, largely due to the intersecting nature between groundedness and completeness. For reasons we explain in the methodology dealing with adversarial parties, we divide completeness and groundedness. *See id.* at 5 n.5.

¹¹⁵ *Id.* at 4-5.

Many corporations claim AI can avoid or minimize hallucinations through a system called “Retrieval Augmented Generation,” or RAG.¹¹⁶ RAG is a technique in Natural Language Processing (NLP) that integrates a retrieval mechanism with a generative model to produce informative and contextually appropriate responses.¹¹⁷ Circling back to our clerk analogy, think of RAG as giving the learned clerk a research service, like Westlaw or Lexis. Or, think of RAG as the expert librarian. The problem with RAG alone is that—like any research search information—there is often too much information to quickly retrieve a good answer.¹¹⁸ If this seems confusing, think no further than the inadequacy of 1L students trying to write their first research memo¹¹⁹—legal research skills take *time* and *practice* to develop.

¹¹⁶ Yunfan Gao et al., *Retrieval-Augmented Generation for Large Language Models: A Survey*, ARXIV (Mar. 27, 2024), at 1, <https://arxiv.org/abs/2312.10997> (“Large language models (LLMs) have achieved remarkable success, though they still face significant limitations, especially in domain-specific or knowledge-intensive tasks [1], notably producing “hallucinations” [2] when handling queries beyond their training data or requiring current information. To overcome challenges, Retrieval-Augmented Generation (RAG) enhances LLMs by retrieving relevant document chunks from external knowledge base through semantic similarity calculation. By referencing external knowledge, RAG effectively reduces the problem of generating factually incorrect content. Its integration into LLMs has resulted in widespread adoption, establishing RAG as a key technology in advancing chatbots and enhancing the suitability of LLMs for real-world applications.”).

¹¹⁷ *Id.* at 12.

¹¹⁸ Case studies have shown how RAG retrieved answers that were missing content, as well as failing to provide top ranked documents, answers without context, being also unextracted, in the wrong format, incorrect specificity, and incomplete. Scott Barnett, *Seven Failure Points When Engineering a Retrieval Augmented Generation System*, ARXIV (Jan. 11, 2024), <https://arxiv.org/html/2401.05856v1>.

¹¹⁹ See, e.g., elainelorelove, REDDIT (Sept. 25, 2022), <https://www.reddit.com/r/LawSchool/comments/xnqawp/bombed>

Empirically, most of these methods have failed to produce a truly effective AI system. One Stanford study demonstrated that LexisNexis and Westlaw AI systems hallucinated between 17% and 33% of the time, respectively.¹²⁰ Other systems, such as Harvey, simply do not have active benchmarks, and efforts by various firms to benchmark these products have produced mixed results.¹²¹

Alternative methods do exist—but they often restrict the unstated metric: latency.¹²² Latency refers to how quickly the LLM can produce a sequence of tokens—a concept that can be quite frustrating, given a user expects their answer to generate immediately.¹²³ While latency is also conceptually a problem for the lawyer, its scale with AI is very different: the lawyer is expected to produce an answer in, say, weeks, whereas the AI is expected to produce an answer in seconds.¹²⁴ In the case of Arbitrus.ai, however, latency is not a

[first legal writing memo](#)/ (“I bombed my first legal writing memo”).

¹²⁰ Magesh et al., *supra* note 108, at 1.

¹²¹ *Id.* at 21-24.

¹²² Chuangtao Chen et al., *LiveMind: Low-latency Large Language Models with Simultaneous Inference*. ARXIV (Nov. 5, 2024), at 1, <https://arxiv.org/abs/2406.14319> (“In this paper, we introduce LiveMind, a novel low-latency inference framework for large language model (LLM) inference which enables LLMs to perform inferences with incomplete user input. By reallocating computational processes to the input phase, a substantial reduction in latency is achieved, thereby significantly enhancing the interactive experience for users of LLMs.”).

¹²³ *LLM Has a Performance Problem Inherent to Its Architecture: Latency*, PROXET (Aug. 1, 2023), <https://www.proxet.com/blog/llm-has-a-performance-problem-inherent-to-its-architecture-latency> (“The need for speed is as important as ever, particularly for websites that deploy LLMs. Over a decade ago, analysts found that the slower the site, the more likely users are to simply click away and go to a competitor: Google found that an extra 0.5 seconds in search page generation time dropped traffic by 20%.”).

¹²⁴ The general definition of latency is the delay before a transfer of data begins following an instruction for its transfer. It affects

problem—Arbitrus.ai is not a research enhancement tool, but a decision-making tool, and thus can process requests over the course of hours and days, not seconds.¹²⁵ As we explain below, the lack of latency problems allow us to employ alternative methods without fundamentally altering the technology behind LLMs.¹²⁶

III. Methodology: Arbitrus.ai at Work

We collectively refer to the system and protocols used to create arbitration artificial intelligence as ‘Arbitrus.ai.’ The goal of Arbitrus.ai is (1) to demonstrate a real, theoretical application of Judge.ai to decision-making machine learning, human reinforcement, and large-language models, and (2) to apply Arbitrus.ai to a set of hypothetical *legal* scenarios to illustrate practical tools.¹²⁷

virtually every technology, but the importance of that effect is subject to the context of that technology’s human use. *Id.*

¹²⁵ *Id.*

¹²⁶ It may, in fact, be the case that a positive latency could harm the perception of Arbitrus.ai outputs—indeed, various psychological studies have demonstrated that humans cognitively equate time with deliberation. Thus, if Arbitrus.ai quickly produced an opinion on a complicated case in, for example, 30 seconds, it might be perceived by subjects to lack sufficient deliberation.

¹²⁷ Our scenarios are admittedly synthetic; in other words, the scenarios presented in this paper were created using artificial intelligence, and thus were not based on specific cases. We did, however, build the scenarios based on real cases reported online. For practical purposes, this means that it is possible for our scenarios, despite human review, to reflect a systematic, ‘insidious’ biases that overemphasize certain types of fact patterns. Future studies should test Fortuna Arbitration’s systems by employing full-time arbitrators to mechanical turk fact patterns, and then match the outcomes to discern the differential. That differential, however, is not necessarily evidence of ‘truth’ matching—that is, humans can be wrong; to accomplish a ‘truth’ matching, a future study will need to conceptualize a way to create an “objective” factual truth, that neither machine nor human knows, and then compare the aggregated outcomes of both. *See* Zhuoyan Li et al., *Synthetic Data Generation with Large Language*

Our methodology has three steps. First, we enumerate an explicit set of procedures for users to navigate and interact with Arbitrus.ai. Second, we generate hypothetical scenarios in specific jurisdictions, using Fortuna Arbitration’s patent-pending architecture and methodology.¹²⁸ Third, we have Arbitrus.ai issue opinions based on the scenarios.¹²⁹ Each scenario has two metadata variables: jurisdiction and dispute topics. A jurisdiction is defined as a specific state.¹³⁰ This study will simulate how Arbitrus.ai would behave in a fully briefed *ex post* adjudication. These same principles would, of course, apply to an *ex ante* query—only with one brief submitted by a human and one brief submitted by an AI agent, and the facts in the submitted brief taken as true.

For grading Arbitrus.ai’s output, we provide three baseline metrics as grades: hallucination, completeness, and

Models for Text Classification: Potential and Limitations, ARXIV (Oct. 11, 2023), <https://arxiv.org/abs/2310.07849>.

¹²⁸ Fortuna Arbitration’s systems that produce an arbitration outcome are collectively referred to as “Arbitrus.ai.” Arbitrus.ai employs several different large-language models, ranging from GPT4-o to Llama 3.1. 70B. Models are subject to change depending on token costs.

¹²⁹ Arbitrus.ai employs several different large-language models, ranging from GPT4-o to Llama 3.1. 70B. Arbitrus.ai’s backend builds legal rules in several different ways, ranging from human-mechanical turked variables to fine-tuned Llama extractions.

¹³⁰ While we recorded jurisdiction as a state, we left it to the arbitrator to determine if there is a conflict-of-laws issue. For simplicity, we focused scenario construction on the assumption that disputing parties reside in the same state.

groundedness.¹³¹ We one-hot encode each variable,¹³² and report it in a CSV file.¹³³

Finally, we select several cases from a range of topical distributions for qualitative review, in order to provide a semblance of real, human outcomes—for example, an insurance dispute.¹³⁴ We then grade the cases, and explain their relevance to arbitrating disputes.

A. Human User Interface

A comprehensive guide of the procedural rules (hereinafter “Arbitrus.ai Procedure”) for Arbitrus.ai in Arbitrus.ai’s open source rules.¹³⁵ These procedural rules set out not only the methods of interaction between the human and AI-system,¹³⁶ but the rules also provide guardrails for AI

¹³¹ See *infra* Part II(C)(3).

¹³² That’s not a typo of “one-hit” or “one-shot”; “one-hot” is the actual, technical term. See *One Hot Encoding*, DEEPAI, <https://deepai.org/machine-learning-glossary-and-terms/one-hot-encoding> (last visited Jan. 17, 2025) (“One hot encoding converts categorical data, typically represented in string format, into a numerical format that can be used in mathematical calculations and hence by machine learning algorithms.”).

¹³³ See *Fortuna Arbitration Rules* (January 21, 2025), https://cdn.prod.website-files.com/6706ddaecf6d39663de1f20/678959c119669f48a699e76b_Fortuna%20Arbitration%20Rules%201.0.pdf; See, e.g., Jia Li et al., *Deep Convolutional Neural Network Based ECG Classification System Using Information Fusion and One-Hot Encoding Techniques*, MATHEMATICAL PROBS. IN ENG’G (Dec. 2, 2018), <https://onlinelibrary.wiley.com/doi/10.1155/2018/7354081>.

¹³⁴ See *infra* Part II(D).

¹³⁵ Procedural rules are decided and approved by the Fortuna Arbitration Administrative Committee. The number of days are specified in the rules as working days. See *Fortuna Arbitration Rules*, *supra* note 133.

¹³⁶ Fortuna Arbitration does review the outcomes of AI processes and decisions in certain cases. We explain this process more extensively. See *infra* note 146 and accompanying text; see also Arbitrus.ai Procedure, Rule 1.3 (“...when the rules reference “Arbitrus” or “Arbitrus.ai” it shall refer to Fortuna Arbitration’s proprietary artificial intelligence systems. Fortuna Arbitration

alignment by preventing critically dangerous errors that manifest from systematic biases that AI models can often reproduce.¹³⁷ The following section outlines how a user interacts with Arbitrus.ai.

Users are formally invited to the Arbitrus.ai portal only after a Party files a claim (“Party A”). (300). Arbitrus.ai then has seven days (7) to serve valid, legal notifications to each party (301); valid legal notification includes the substantive complaint from Party A. Each party then has seven (7) days prior to confirm notification (302-303).¹³⁸ Once notification is confirmed, the parties are advanced to the second stage: briefing (343). During Briefing, the deadline for Party B’s response brief is seven (7) days (305-308). Party A then has seven (7) days to file an answer (308-309). During this period, either Party may file a dispositive motion (306-307, 313, 317). Arbitrus.ai then decides the motion (310-312, 316,318-319, 321); if a motion is dispositive, a ruling is issued and final order authorized (338-339, 342). If a motion is not dispositive, a ruling is still issued, and the Parties sent back for Briefing (312, 321). After Briefing, Parties are sent to Discovery (344).

During Discovery, both Parties are ordered to submit evidence (315, 323, 324-325). After Discovery is processed, Parties have seven (7) days to file objections and request additional evidence (328); if Arbitrus.ai agrees with a Party’s objection or request for additional information, the Parties have seven (7) days to resubmit relevant Discovery to be

retains the right to assign human employees or agents to review the work of Arbitrus.ai prior to disclosure of rulings or information to the Parties.”).

¹³⁷ Cf. ISAAC ASIMOV, I, ROBOT 36 (1950) (“We have: One, a robot may not injure a human being, or, through inaction, allow a human being to come to harm.... Two ... a robot must obey the orders given to it by human beings except where such orders would conflict with the First Law.... And three, a robot must protect its own existence as long as such protection does not conflict with the First or Second Laws.”).

¹³⁸ Party A’s complaint is included in the notice. It would be unfair for Party B to respond to an allegation that was not sufficiently particular.

processed (322). Discovery is then processed again. If there are objections, and Arbitrus.ai agrees, then the process begins again (322). If there are no objections, Arbitrus.ai may decide to order additional briefings, in which case the process begins again (330-331, 341-342).

Alternatively, if Arbitrus.ai determines no additional briefing is necessitated, it orders a Hearing (333); during which time, both Parties are simultaneously queried for questions by Arbitrus.ai (334-336). Both Parties will have an opportunity to review responses. After the Hearing is concluded, Arbitrus.ai will issue a ruling within three (3) days (337-339). Either as part of, or separate to, the Ruling (depending on the case), Arbitrus.ai will also issue a Final Order (342), that contains the specific award.¹³⁹

¹³⁹ LLMs cannot effectively do computations without tool calling. Tool calling, however, is a rather novel invention. Because of this, while Arbitrus.ai does calculate an award, the final award is reviewed by a human. See Shengtao He, *Achieving Tool Calling Functionality in LLMs Using Only Prompt Engineering Without Fine-Tuning*, ARXIV (July 6, 2024) <https://arxiv.org/abs/2407.04997v1>.

by human review and reinforcement. This section outlines the process from the view of the backend to Arbitrus.ai.

As explained above, human interaction with the system—predominantly performed by two, adversarial parties—occurs in four primary modules. These modules are illustrated in Figure 1, with corresponding explanations for each module unit (represented by a 100 series). First, after a party files a dispute submission (100), both parties are summoned to the notice module through electronic communication, predominantly email (102-103).¹⁴⁰ Human users are then given access to the Graphical User Interface (105).¹⁴¹ Second, unless there is an objection (115),¹⁴² both parties are prompted to provide briefs within a given timeframe (104-106). If there is an objection, Arbitrus.ai decides the result (116); if Arbitrus.ai decides the motion is dispositive (116, 132) in nature, then Arbitrus.ai decides on the merits if the case should be disposed (120). If the motion is not dispositive, a decision is issued and reported in the Arbitrus.ai Procedural Decisions module (117, 120, 124, 126). If the case should be dispositively ruled on, a human reviews and confirms the output (130). If the human agrees with Arbitrus.ai that the case should be dispositively ruled on, a final and binding ruling is issued to the parties (131-132);

¹⁴⁰ Arbitrus.ai Procedure, Rule 6.2(a); *see also* Arbitrus.ai Procedure, Rule 3. In Fortuna Arbitration’s backend, there is a manual check to ensure the petitioning parties are governed by Arbitrus.

¹⁴¹ For those users concerned with process-maxxing (i.e. the process independent of the substance), the User Interface can also be dynamic—i.e., users, if the parties mutually agree, enhance the process. *Cf.* Benjamin Pontz, *Review: Keeping our Republic*, HARV. J.L. & PUB. POL’Y: PER CURIAM (Sept. 26, 2023), at 2, <https://journals.law.harvard.edu/jlpp/review-keeping-our-republic-benjamin-pontz/> (“...the idea that self-government according to a *rule of law prescribed in advance through methods* that are ours until we change them has *inherent value*, independent of particular outcomes.”) (emphasis added).

¹⁴² For example, an objection on jurisdictional grounds. *See* Rule 3 in *supra* note 133.

otherwise, the case is returned to Arbitrus.ai for processing (122).

Third, after briefs are submitted, parties are advanced to the discovery module for additional information gathering (108).¹⁴³ During this phase, parties can file objections to discovery on reasonable grounds (118, 113).¹⁴⁴ A recording of the ruling is returned to the Arbitrus.ai Procedural Decision module and reported to the parties (118, 125-126). It is during this phase that depositions that involve cross examination should occur, and be uploaded (108).¹⁴⁵ Fourth, and finally, parties are promoted to a hearing module to present their case and answer questions queried by Arbitrus.ai—all of which is done via writing, and none of which permits cross examination (119, 114, 109).¹⁴⁶ If a motion is made or

¹⁴³ Arbitrus.ai Procedure, Rule 6.2(b).

¹⁴⁴ Arbitrus.ai is not required to follow formal legal rules of discovery, as these are not applicable in the arbitration context. *See* Arbitrus.ai Procedure, Rule 6.4 (“Arbitrus shall manage any necessary exchange of information among the parties with a view to achieving an efficient and economical resolution of the dispute, while at the same time promoting equality of treatment and safeguarding each Party’s opportunity to fairly present its claims and defenses. Arbitrus is not required to follow normative legal rules relevant in a district court, although such rules may be persuasive.”).

¹⁴⁵ Cross-examinations occur in depositions. If parties submit a motion, they can request a referee from Fortuna-Arbitration. In the future, separate IP will need to be created to referee cross-examination.

¹⁴⁶ No cross examination is permitted during the hearing. Some jurists believe oral arguments are vital to the truth-value of the arbitration process; some do not. *Compare* Warren D. Wolfson, *Oral Argument: Does It Matter*, 35 IND. L. REV. 451, 451-52 (2002) (“Or are the briefs so good we just don't need any more argument? Or is it because we don't have many Daniel Websters any more? Appellate judges will tell you sitting up there and listening to lawyers read prepared talks or their briefs is not our idea of a great time.”), *with* Robert Coleman, *The Vanishing Oral Argument: Why it Matters and What to Do About It*, AM. BAR. ASSOC. (Feb. 18, 2020), <https://www.americanbar.org/groups/judicial/publications/appellat>

submitted during the hearing, Arbitrus.ai renders a decision and reports it into the Arbitrus.ai Procedural Decision module (119, 120, 125, 126). This process is iterative (110, 107, 104, 111); Arbitrus.ai can decide to have parties submit new briefs if discovery or the hearing warrants it. If all the relevant human information has been sufficiently gathered, the data is piped into a database (110, 107, 104, 111, 114, 127). Arbitrus.ai then transitions to the next phase of AI decision making (128). Figure 2 represents the process of the human interface.

e_issues/2020/winter/the-vanishing-oral-argument (“oral argument is not only helpful in that individual case, but it also piques his intellectual curiosity by exposing him to discussions with attorneys in different types of cases.... It presents the opportunity for her to learn about the opinions of her co-panelists and address their concerns. And although oral argument *does not usually fundamentally alter Judge Keenan’s opinion regarding the outcome of a case*, it does provide the chance for her to gauge the strength of her rationale for arriving at that outcome.”) (emphasis added); For further discussion, see Bruce A. McAllister & Amy Bloom, *Evidence in Arbitration*, 34 J. MAR. L. & COM. 35, 44 n.41 (2003). Arbitrus.ai accounts for this truth-value *from witnesses* in the Discovery phase, even excluding the talented manipulation of lawyers. We generally exclude the importance of attorneys for a simple reason: the persuasive oral talent of a given party’s attorneys is not important as a matter of law, and it incentivizes gaming the system.

LLM side, information between each step is processed, checked for prompt hacking or other malicious data, and then merged into a single file that summarizes the positions, data, evidence, and responses of both parties.

First, data aggregated from the human interface is piped into the preprocessing container (201, 202, 203, 204). The preprocessing container merges the data (209, 213, 215-216) into a single database (217-218), where transformed raw text data can be either sent to the ML Decision Module or to the LLM Reasoning Module (219, 220, 221).

The second step of decision-making is for Arbitrus.ai to decide the overall ruling of the case. Data is piped through a decision query into several classification models (221, 224, 223); in turn, these models predict and encode a multi-categorical outcome (Ruling for Plaintiff = 0; Ruling for Defendant = 1; A partial outcome = 2; For the case or motion to be dismissed, or other result = 3) (223, 225, 226). Each classification model then votes on an outcome, with each vote weighted based on the model's accuracy (223, 225-226).¹⁴⁷ The result with the highest probability is passed back into the reasoning query, to be summarized as the decision based on the set of presented facts (224).

The third step of decision-making is for Arbitrus.ai to explain its decision in human terms (249).¹⁴⁸ User legal queries

¹⁴⁷ The classification systems, especially during early deployment, are reviewed and improved by human reinforcement agents. This description is also an oversimplification of how voting ensembles technically function. For a technical overview on the voting ensemble methodology, see BURKOV, *supra* note 87, at 201.

¹⁴⁸ We intentionally eschew the mention of a specific LLM, as (1) the preferable model may change over time and (2) multiple models are employed on the backend. The one or more artificial intelligence-based large language models may also include any type of tool that is configured to evaluate natural-language prompts based on pre-processed input data. This includes models that are accessed via application programming interfaces (APIs), as well as non-API based services. For example, the software platform may utilize large language models such as Open AI's GPT-4, Meta's LLaMa, Claude, or Microsoft's Wizard. There are many examples

first are vectorized¹⁴⁹ through various methodologies¹⁵⁰ (228-229, 232) and reported in a Vector Index (233).¹⁵¹ The vectors are compared to other vectors in a vector index using cosine similarity (233, 234).¹⁵² The most similar vectors are returned

of large language models, including those large language models that are capable of being custom-trained, and the present invention is not intended to be limited to any one or any specific type of model. It is to be understood that other third-party tools may also enable connection via an API, or via other approaches, as those tools emerge and become available.

¹⁴⁹ Vectorization of the components can be done through a variety of open-source or proprietary models. Other transformer models can work but may create slightly different outputs. Samples of models include Google’s BERT, OpenAI’s Ada-002, and Meta’s LLaMa, and many others.

¹⁵⁰ We give the model agency to select the best tools—which sometimes go beyond retrieval augmented generation (RAG). In some cases, the model can also iterate through tools to filter results to select the best tool.

¹⁵¹ There are two methods in which cases are systematically compared. First, a large language model will look at the query in sections no larger than its context window and identify key phrases, legal terminology, and events. These data points are referenced as topics. Extracted components are then vectorized at step. Natural language processing (NLP) models are able to embed semantic information as a vector, or a series of floating-point numbers. While these vectors have no actual meaning, they can be compared to other vectors created by the same model to evaluate semantic and contextual similarity. Second, similar to the aforementioned procedure, a large language model will look at the query in sections no larger than its context window and extract the relevant legal rule from the context of the user briefs. These legal rules are then compared to the legal rules in Fortuna Arbitration’s database, which have been extracted from cases using a proprietary methodology. Combined, these methods maximize groundedness.

¹⁵² It is notable this is not the normal retrieval augmented generation (RAG) system often employed by legal AI companies; instead, Arbitrus.ai uses a proprietary preprocessing methodology that allows for the extraction of key legal and factual concepts. In effect, this means Arbitrus.ai is RAGing on key concepts—not raw legal text. In other legal research systems, these key concepts, or ‘rules,’ are called ‘headnotes.’

(234-237),¹⁵³ and Arbitrus.ai then checks these using a further, different transformer model to verify their similarity (234). Finally, Arbitrus.ai utilizes a separate large language model to check similarity again and pulls the full opinion text of cases which are most similar from a first-party case database (235-237). Those cases are then provided to the one or more large language models with the relevant client data.

Next, the data is filtered through an argument iteration engine, illustrated in Figure 3 by several proprietary filters (239-241),¹⁵⁴ wherein a combination of deep-learning networks and human interaction are trained to evaluate the suggested responsive communication based on a variety of factors for likelihood of success (239-241). The argument iteration engine ‘iteration’ component largely occurs within the Ra.ai agentic solution, illustrated in Figure 4, but for practical purposes represented by (239).¹⁵⁵ The opinions

¹⁵³ The process retrieves full case opinions (according to a specified number n) related to the returned vectors. The backend database accesses the case database and passes in all opinion identifiers returned from the vector index search. The full case data is then transferred to the backend generator for drafting.

¹⁵⁴ We have contemplated employing machine learning models for evolutionary iteration to optimize performance. However, because our system produces results *ex ante* the reasoning, we decided against it. In our early prototypes, the argument iteration engine may include a prediction component which is configured to generate one or more predictive outcomes for the suggested, responsive communications. Based on these predictive outcomes, the software platform then generates output data as determined by the user or guided by the type of situation of each legal query. In substance, subject to further metadata, the models are based on courts, forums, judges, arbitrators, mediators, date ranges. Non-metadata, such as rule statements, are also included.

¹⁵⁵ Ra.ai is also the component largely responsible for eliminating virtually every hallucination from the process. The Ra.ai system includes multiple calls to one or multiple LLM models, which in turn police results. We call our internal agentic solution “Ra.ai,” after the Egyptian pantheon, because it ‘rules’ above all other agents, whether in the sky (the source agents, 416-419), the earth (the templatization and return on the front end, 423-425), or finally the underworld (determining if the templatization may pass onto

returned by the Ra.ai agentic solution may have a large number of characters, which may exceed the context window of the large language model that will be used for drafting suggested responsive communications after tokenization. They are thus summarized first by a separate large language model, and the summaries of the opinions/legal research are then provided to the primary prompts for generation of suggested responsive communications (242). These generational prompts may be tailored to follow the IRAC or CRUPAC legal model for legal writing, and apply rules extracted from the opinions to client issues in the legal queries (245-246).¹⁵⁶

As a quality control measure for on-boarding new clients, the output data is reviewed by a human (247),¹⁵⁷ before being

the frontend, 427-428, or be reiterated, 403). While the actual number of calls made vary (and thus the number of ‘gods’), Ra.ai reigns above the others. This should not confuse Ra.ai as the largest or most sophisticated model—although it often is, depending on the call made.

¹⁵⁶ It is notable that this process is considerably more computationally heavy than normal RAG; not only does it require greater infrastructure (e.g. constant training and retraining of machine learning models), but it is also contingent on the cost-per-token, which varies depending on methodology. And because it is iterative—that is, human interference could send back a decision to Arbitrus.ai—those costs could increase in complex cases. While the cost per opinion will change over time, it approximately costs us ten dollars (\$10) per ten to twenty pages. This cost will decrease at scale, especially if Fortuna can host the process internally. Luckily, because latency is not a problem, the increased procedure costs do not affect the end-user.

¹⁵⁷ We built in this step for three reasons. First, it is a request made by industry actors, who are not yet comfortable with a full AI solution. Second, it allows us to record continually improving metrics (e.g. a lack of or presence of hallucinations). Third, it allows the insertion of industry specialists to handle or correct specific facts in complex disputes—especially where such specialists are expected by the parties.

In practice, a human reviewer will only review several first initial arbitrations as a quality control check. Once the machine is

issued as a final and binding opinion to the parties (248).¹⁵⁸ All human-reviewed briefs are graded for sufficiency, and returned into a reinforcement training database, where the data is used to train future ML decision models,¹⁵⁹ and the Ra.ai agents (250-253).¹⁶⁰ Output data may include different outcomes, in the form of different types of legal writing (i.e., motion rulings, final and binding opinions, etc.).¹⁶¹

Figure 3. Arbitrus.ai Architecture

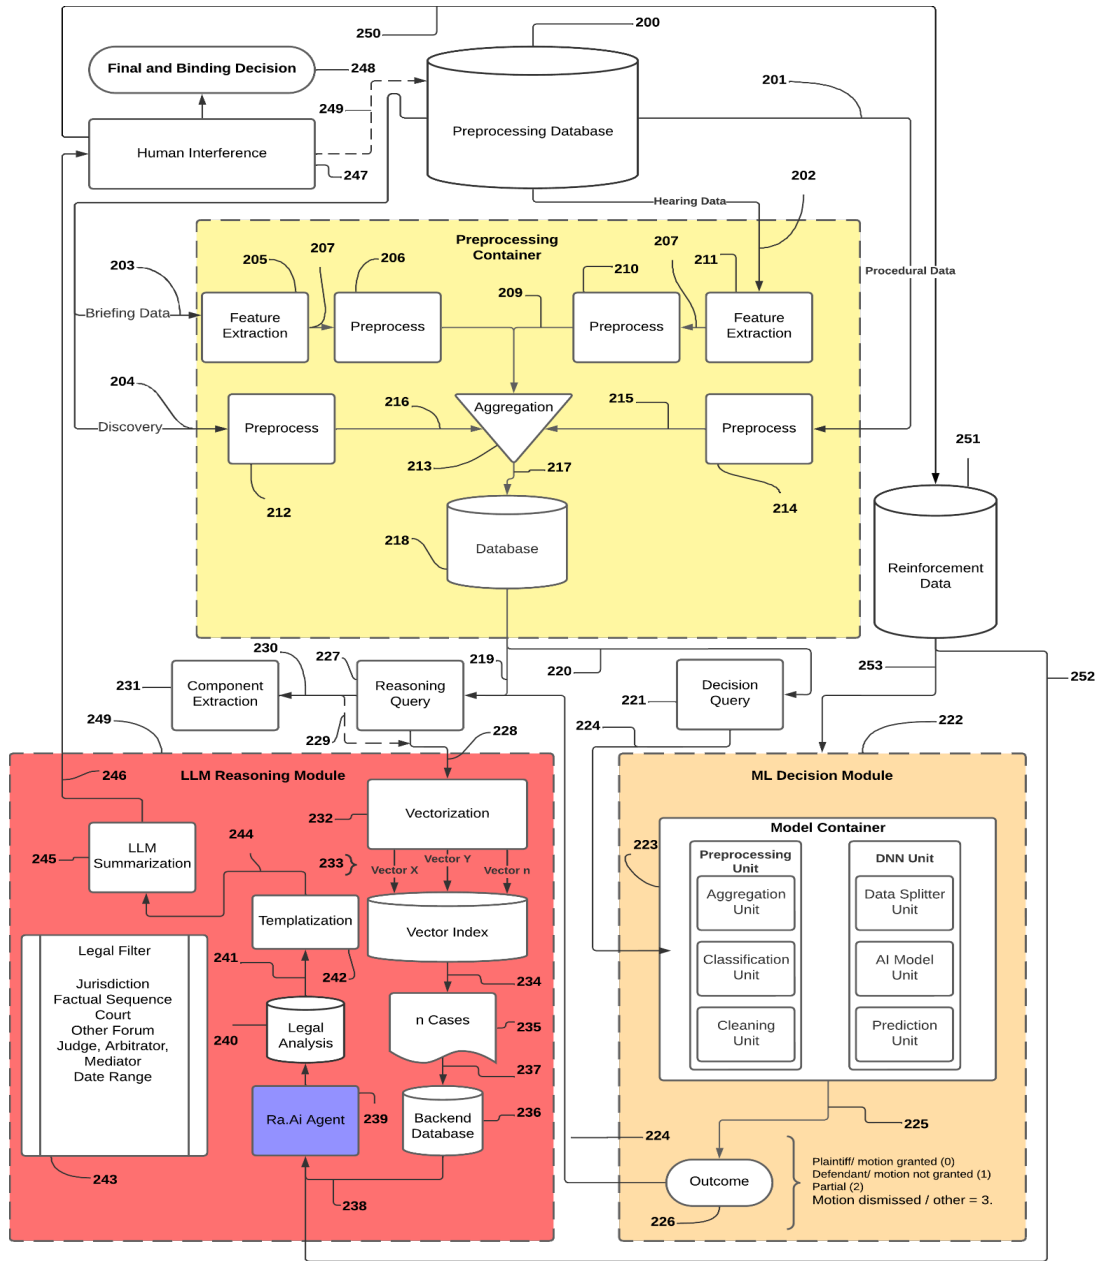
confirmed operational, human checks can be phased out of the process.

¹⁵⁸ Not listed here is the Human Appeals process in Fortuna Arbitration Rules, *supra* note 133, Rule 8 on very limited grounds. Eventually the Appeals process can be incrementally phased out as the machine demonstrates legal proficiency.

¹⁵⁹ Amongst other methods, humans are required to grade the sufficiency of each outcome. That is, if the ‘correct’ decision was made that would mirror a real human arbitrator.

¹⁶⁰ Each opinion is given grades based on the three metrics of this study: hallucination, correctness, and groundedness. These metrics can be substituted later for different preferences by clients, such as preferred types of citations (*e.g.* insurance companies might prefer a reference to a direct policy statement, as opposed to a common law rule). Actually training the models for industry preferences creates stronger results than mere prompting, which will lower the cost from iterative generation.

¹⁶¹ Arbitration rulings can also be substituted in other forms, as may be required for other verticals. This may include legal briefs, legal memoranda, motions, and legal correspondence, and any other types of legal writing. The argument iteration engine may also generate legal research services, which may or may not be the result of a predictive analysis determined as a predictive outcome.

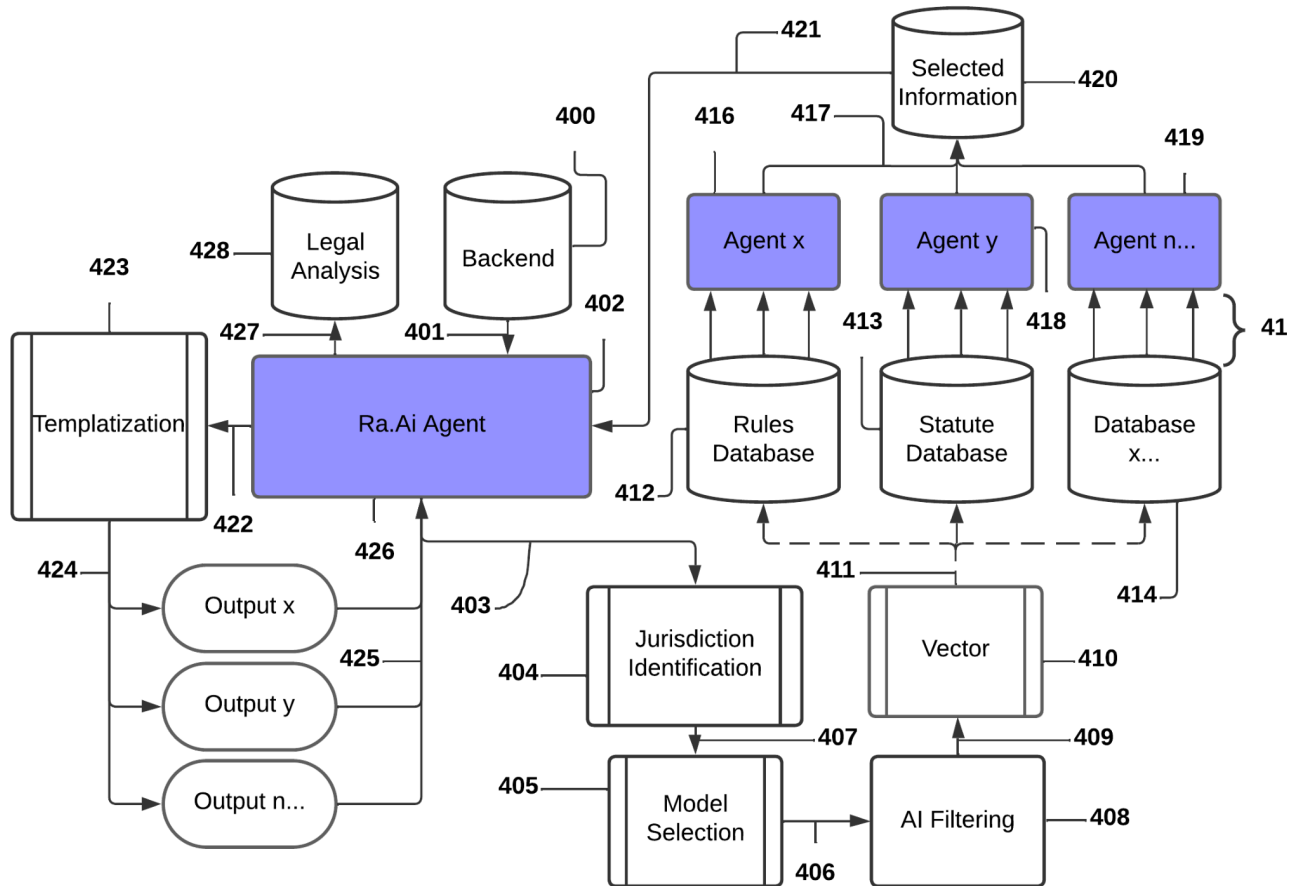


Ra.ai is Fortuna Arbitration’s agentic solution to the hallucination problem, as illustrated by Figure 4. As illustrated by Figure 3, vectors are passed into a backend database (236, 400). The backend includes a finite number of vector relevant cases (235). The backend database then passes the data into Ra.ai (400-402). Ra.ai then determines the relevant jurisdiction, given a predetermined process (i.e. prompted to identify one of the following jurisdictions, given

a set of US states) (404-404). The jurisdiction identification process then selects a LLM model (407, 405), which then filters cases exclusively related to the jurisdiction (406, 408), or by a preferred filter (such as only a ‘rule’ or only a statute).¹⁶²

Thereafter, the vector is passed into the cases (which database the vector is passed to can be controlled by the AI filtering) (411-414); the cases returned in the backend database are sorted by category into several different databases (412-414). Each of these databases then pass the final related sources, which may be numerous (normally five to ten) (415-416, 418-419). Each AI agent, using a model defined by 405, then ranks the sources, before passing the highest ranked sources (up to three, but can be varied) (417), into a selected information database (420). The selected information database is then passed back into Ra.ai, which selects the best sources (420-421). Ra.ai then passes the information into various templates (e.g., in IRAC, the issue, the rule, etc.) (402, 422-423), which employs various LLMs to produce various outputs (424). Ra.ai then determines if the output is sufficient (425, 402); if it is sufficient, the outputs are returned to a legal analysis database (427-428). This is the analysis to be employed in the LLM summarization stage, where the final “opinion” is to be written (240-242, 244). If Ra.ai believes the outputs are not related sufficiently to the input, it iteratively returns them to the analysis process described (403-421). This iteration process effectively allows Ra.ai to assign various ‘agents’ micro-tasks, whether it be about selecting the relevant cases (416-419), or writing the best analysis (422-425). In effect, this means that Ra.ai can access multiple custom APIs and other tools to retrieve and select the most relevant sources for the dispute.

¹⁶² Fortuna-Insights has a proprietary database of ‘rules’ extracted from every legal case in the United States. Each legal rule is matched to a parenthetical in the given case. Other research tools, such as Westlaw, call these ‘headnotes.’

Figure 4. Ra.ai Agentic Methodology

C. Data Collection

We collected data in three steps: first, we generated briefs for both adversarial sides; then, we generated opinions, using the Arbitrus.ai backend.¹⁶³ Finally, we employed a standardized grading procedure to determine the validity of each opinion.

¹⁶³ The data used in this study can be found here: https://fortunainsights-my.sharepoint.com/:x:/g/personal/kimo_gandall_fortunainsights_com/EYLB4nuQAtxKhyOkd9Y1sNcB8l-PSP_INATx1rvZCTFaIw?e=a1die3.

1. Brief Procedure

To overcome the lack of sufficient real-world data,¹⁶⁴ we employed an LLM-driven synthetic data paradigm to create scenarios, match facts, and create responsive situations.¹⁶⁵ In this study, we defined a brief as a single party’s arguments on the merits and process for a given case.¹⁶⁶ Each brief included a summary, fact section, substantive arguments, and a conclusion; moreover, each brief is dispositive—this means that each brief includes a sufficient list of legal arguments and facts for Arbitrus.ai to render an opinion. This prevents the need for Arbitrus.ai to iterate briefing, as may be needed in cases where Parties ignore, intentionally or not, key facts or arguments.¹⁶⁷ Because Arbitrus.ai needs to be responsive to invalid, weak, or hallucinated legal arguments that human lawyers might present, we did not review each brief for hallucinations, groundedness, or correctness.¹⁶⁸ Briefs were

¹⁶⁴ Even where arbitration rulings are online, arbitration briefs are not. And the briefs—what the parties substantively argued to the arbitrator—is what is of importance. *See Collins, supra* note 49, at 323.

¹⁶⁵ *See Cheng Chang, LLMScenario: Large Language Model Driven Scenario Generation, 54 IEEE 6581, 6582 (2024).*

¹⁶⁶ While we did not strictly control case length, we aimed to build a hundred short scenarios. Future studies will need to invest in reviewing long-form scenarios. Part of the reason for shorter scenarios is to provide easier review and discussion—indeed the long-form scenarios will naturally present additional difficulties in ascertaining “correctness” and “groundedness.”

¹⁶⁷ A future study should procedurally war-game Arbitrus.ai; that is, because parties are frequently negotiating, changing conditions, or discovering new evidence, Arbitrus.ai should be responsive to changing or adversarial conditions. As a note, the objective of this study is to demonstrate a legal decision-making machine’s application to a set of legal scenarios—not to demonstrate the entire system is empirically justified.

¹⁶⁸ A future study should control for intentionally hallucinated cases, or insidiously presented facts. We did not do so here because of the methodological cost of creating truly ‘insidious’

then created for both Parties, and reported to human researchers.

2. Opinion Procedure

To generate opinions, we placed matched briefs (representing a claimant and respondent on a specific given case) into the Arbitrus.ai process, illustrated in Figure 3. Each opinion was then reported to a human researcher.¹⁶⁹ Arbitrus.ai provided several key descriptions about the opinion, including jurisdiction and dispute topic. The human researcher reports these descriptions; thereafter, the human researcher checks the opinion to ensure the descriptions are accurate. Finally, the human researcher reports the prevailing party as either “Claimant” or “Respondent.”¹⁷⁰

scenarios where a fact seems to be true but is verifiably not. The best way to accomplish this in the future is to create a verifiably true instance of an occurrence—such as a subject hitting another subject—that is then matched to facts implicating the truth value of the occurrence (such as eyewitness testimony). Even then, confirming laboratory truth values and real-life truth is a different dilemma, because laboratory blind-testing invariably results in certain ‘matched facts’ having systematic biases that are difficult for humans to test (for example, it might be the subjects recruited by the examiners happen to have characteristics the AI correlates with more or less trustworthiness). Some of these systematic biases might also present legal problems for AI, including discrimination. However, insofar as these biases promote accuracy, courts have found these biases to be legitimate mechanisms of promoting justice. *See, e.g., State v. Loomis*, 881 N.W.2d 749, 767 (Wis. 2016) (“We determine that COMPAS’s use of gender promotes accuracy that ultimately inures to the benefit of the justice system including defendants.”).

¹⁶⁹ While some studies have suggested that LLM models can grade LLM outputs, we decided to be sensitive to the demands of lawyers that graded outputs be reviewed by human researchers.

¹⁷⁰ Our prevailing party is binary, unlike our traditional ML models, because it excludes how an award might be issued; that is, an award might not offer the prevailing party their full demand. This can be significant; for example, if a prevailing party is asking for five thousand dollars, but only is given five hundred dollars. Because Arbitrus.ai runs the award ruling through human

3. Grading Procedure

Each opinion was manually graded for three key metrics: hallucination, completeness, and groundedness. A hallucination is reported if the opinion contains a fact, case, or concept that does not exist; in practice, that could anything from an incorrect pincite,¹⁷¹ to a completely invented case precedent.¹⁷² We define completeness as a ruling that answers the questions presented by the briefed query (i.e. the issue in controversy is resolved).¹⁷³ We define groundedness as a ruling that contains a legal authority that is properly cited to

interference to reinforce proper tool calling, we defined the prevailing party as a binary to create simplicity in our reporting structure for this paper. As a result, we one-hot encode the prevailing party as “Claimant” or “Respondent” (Claimant = 0; Respondent = 1).

¹⁷¹ We did not count a bluebooking error as a hallucination. The hallucination needed to point to some material, incorrect fact—not a preferred stylistic option. To the extent that a bluebook format error misrepresented a fact, such a pincite indicating the page number a quote appeared on, we counted the error as a hallucination; but if the bluebooking error did not sufficiently convey a falsifiable fact—for example, incorrect italicization on a comma—we did not report a hallucination.

¹⁷² It is notable that we opted not to adopt the more expansive definition provided by Magesh et al., *supra* note 108, at 1, because it includes “misleading information.” We intentionally excluded that definition from hallucination, since we included that definition under groundedness. This definition also encompasses “accuracy,” insofar as the “correct” answer is legally correct. Again, true “correctness”—that the opinion complies with a higher law than the positive law—is not the concern of the study.

¹⁷³ We differentiated our definition of completeness from Magesh et al., who define “incomplete” responses as either “refusals or ungrounded.” *Id.* at 9. Because, unlike generic legal research products that only need to respond to a set of alleged facts, Arbitrus.ai has drafted a ruling in response to a set of competing facts, we defined completeness more stringently: namely, a response that addresses two, competing parties.

the extent that it is factually relevant, and that is responsive to the issue in controversy.¹⁷⁴

Table 1. Selected Scenarios				
Case ID	Jurisdiction	Topic	Claimant Argument	Respondent Argument
4	Illinois	Executive Contract Dispute	<p>The brief addresses an employment dispute where the claimant, Victoria Green, alleges that Summit Dynamics Corporation breached an executive employment contract by failing to provide severance and equity compensation upon her termination without cause, in violation of Illinois law.</p> <p>Summit Dynamics materially breached its clear and enforceable contract with Green by not paying \$200,000 in severance and \$150,000 in equity compensation following her termination, which was for "restructuring" and does not qualify as</p>	<p>The brief addresses a dispute in which Summit Dynamics Corporation argues that it lawfully terminated Victoria Green for cause, citing gross mismanagement, budgetary misconduct, and breaches of fiduciary duty, thereby rendering her claims for severance and equity compensation invalid under the employment contract and Illinois law.</p> <p>Summit Dynamics asserts that Green's termination was justified as "for cause" under the contract due to her documented mismanagement and fiduciary breaches, precluding her entitlement</p>

¹⁷⁴ *Id.* at 7, distinguishes “groundedness” from “correctness”: “Correctness refers to the factual accuracy of the tool’s response (Section 4.1). Groundedness refers to the relationship between the model’s response and its cited sources (Section 4.2).” We combine these metrics, because Arbitrus.ai’s chain-of-thought methodology is largely summarizing sources, not drafting the most probable answer—the better question is if the source is *sufficiently correct*. Like Magesh et al., however, we define a ruling as grounded even if it is not the best response: “...we label an answer that is partially correct—that is, one that contains correct information that does not fully address the question—as correct.” Determining the best, rather than sufficient, answer is a matter of legal interpretation that would take considerably more research (if it even is possible, given that there is often a range of acceptable legal answers).

			"cause" under Illinois law or the contract's terms.	to severance, equity compensation, or attorney's fees, and requests dismissal of her claims.
17	Louisiana	Payment in Construction	<p>The brief concerns a payment dispute in which Crescent Bay Builders seeks to recover \$180,000 in unpaid amounts, interest, and attorney's fees from Riverfront Commercial Development Group for labor, materials, and change orders completed under a construction contract for a commercial building project in Baton Rouge, Louisiana.</p> <p>Crescent Bay Builders asserts that Riverfront breached the construction contract by failing to pay the outstanding \$180,000 for completed and approved work, and argues that Crescent Bay is entitled to full payment, interest, and attorney's fees under the contract and Louisiana law.</p>	<p>The brief addresses a construction contract dispute in which Riverfront Commercial Development Group argues that Crescent Bay Builders failed to meet contractual obligations by delivering delayed and substandard work, improperly charging for unapproved change orders, and causing financial harm to Riverfront.</p> <p>Riverfront contends that Crescent Bay breached the contract by delivering deficient and delayed work, justifying the withholding of payments and offsetting costs for remediation, and asserts that Crescent Bay's claims for additional compensation lack merit due to unapproved charges and failure to comply with contract terms.</p>
24	Oregon	Insurance Coverage Disputes	<p>The brief addresses an insurance dispute in which Harper Industries seeks compensation from NorthStar Insurance for \$1.2 million in damages resulting from a fire at its facility, alleging that NorthStar wrongfully denied coverage based on inapplicable exclusions and acted in bad faith under Oregon law.</p>	<p>The brief addresses an insurance dispute in which NorthStar Insurance argues that its denial of Harper Industries' claim for fire-related damages was justified under the policy exclusions for inadequate maintenance and wear and tear, and asserts that Harper's demands for damages and attorney's fees lack contractual or legal support.</p>

			<p>Harper Industries argues that NorthStar breached its contractual obligations by denying coverage for fire-related losses, relying on exclusions that do not apply to the accidental electrical fire, and contends that Harper is entitled to full compensation, consequential damages, attorney’s fees, and arbitration costs due to NorthStar’s bad faith and unreasonable denial of the claim.</p>	<p>NorthStar contends that Harper’s claim is barred by the policy exclusions for inadequate maintenance and wear and tear, and argues that Harper is not entitled to consequential damages or attorney’s fees because the denial was consistent with the policy terms and conducted in good faith under Oregon law.</p>
45	Kentucky	Data Privacy	<p>The brief addresses a dispute in which BlueHaven Financial Group alleges that SecureNet IT Services breached a managed services agreement by failing to maintain adequate cybersecurity protections and promptly respond to a data breach, resulting in significant financial losses, reputational harm, and regulatory scrutiny.</p> <p>BlueHaven argues that SecureNet materially breached its contractual and professional obligations by failing to implement industry-standard security measures and respond promptly to a data breach, thereby causing foreseeable damages, for which BlueHaven seeks compensation, reimbursement, and attorney’s fees.</p>	<p>The brief addresses a dispute where SecureNet IT Services defends against allegations of failing to uphold cybersecurity obligations, arguing that BlueHaven Financial Group’s data breach resulted from its own negligence, refusal to implement recommended security measures, and employee errors, rather than any breach by SecureNet.</p> <p>SecureNet contends that it fulfilled its contractual obligations by providing professional services and timely recommendations, and asserts that BlueHaven’s claims for damages and reimbursement lack merit because the breach resulted from BlueHaven’s own negligence and failure to implement critical security measures.</p>

<p>49</p>	<p>Colorado</p>	<p>Mining rights and royalty disputes</p>	<p>The brief concerns a mining rights and royalty dispute in which Northern Ridge Mining, LLC alleges that RedRock Resources Ltd. materially breached their agreement by imposing unauthorized access restrictions and refusing royalty payments, resulting in financial losses and operational disruptions.</p> <p>Northern Ridge argues that RedRock violated the agreement by obstructing access and rejecting royalty payments without legal basis, entitling Northern Ridge to compensation for unpaid royalties, damages for lost revenue, and recovery of attorney’s fees due to RedRock’s breach and bad faith conduct.</p>	<p>The brief addresses a mining rights dispute where RedRock Resources Ltd. defends its actions to restrict Northern Ridge Mining, LLC’s access to a Colorado property and withhold royalty payments, asserting these measures were necessary due to Northern Ridge’s violations of environmental and safety regulations under their agreement.</p> <p>RedRock argues that its access restrictions and suspension of royalty payments were justified by Northern Ridge’s breaches of environmental and safety obligations, which exposed RedRock to significant legal risks, and asserts that Northern Ridge’s claims for unpaid royalties, damages, and attorney’s fees lack merit as the losses stem from Northern Ridge’s own non-compliance.</p>
<p>80</p>	<p>New York</p>	<p>Profit Sharing</p>	<p>The brief concerns a profit-sharing dispute in which Beacon Partners LLC alleges that Sterling Collective, Inc. breached their partnership agreement by withholding \$250,000 in profit shares, despite Beacon’s full performance of its contractual obligations.</p> <p>Beacon contends that Sterling breached the partnership agreement by failing to pay the agreed profit shares</p>	<p>The brief addresses a profit-sharing dispute in which Sterling Collective, Inc. argues that Beacon Partners LLC failed to meet performance metrics under their partnership agreement, justifying Sterling’s withholding of profit shares and rejection of Beacon’s claims for damages.</p> <p>Sterling contends that its withholding of profit shares was consistent with the</p>

			without any valid evidence of deficient performance, entitling Beacon to the withheld \$250,000, interest, and reimbursement of attorney’s fees and arbitration costs under the agreement.	partnership agreement, as Beacon’s contributions were minimal and did not meet the agreed-upon performance standards, making Beacon’s claims for profit shares, interest, and attorney’s fees unjustified.
99	Georgia	Health Insurance Denial	<p>The brief addresses a dispute in which Genetix Solutions, LLC accuses BioMosaic Innovations, Inc. of breaching a confidentiality and licensing agreement by infringing on Genetix’s intellectual property, disclosing proprietary technology to a third party, and improperly sublicensing it, causing financial and reputational harm.</p> <p>Genetix asserts that BioMosaic materially breached their agreement by unauthorized use, disclosure, and sublicensing of Genetix’s proprietary technology, entitling Genetix to damages, injunctive relief to prevent further misuse, and reimbursement for legal and arbitration costs.</p>	<p>The brief concerns a dispute in which BioMosaic Innovations, Inc. denies allegations of intellectual property infringement and breach of contract brought by Genetix Solutions, LLC, asserting that its developments were independently derived and that its actions complied with the terms of their confidentiality and licensing agreement.</p> <p>BioMosaic argues that it did not infringe Genetix’s intellectual property or violate the agreement, as its advancements were based on independent methodologies, the alleged breaches are unsubstantiated, and any ambiguity in the agreement justifies dismissing Genetix’s claims and awarding BioMosaic reimbursement of arbitration costs and legal fees.</p>

Next, we one-hot encode each variable (hallucination: 0 = not hallucinated, 1 = hallucinated; completeness: 0 = accurate,

1 = inaccurate; groundedness: 0 = law cited and relevant, 1 = no law cited or irrelevant). Because each variable has some subjective nature, each grade is reviewed for accuracy. When a grade is confirmed as accurate, it is recorded and stored in a CSV file.

4. Selected Cases for Qualitative Review

We selected seven examples from our one-hundred cases. We selected cases, qualitatively, based on a diversity of jurisdiction and topics. Once obtaining the brief for the Claimant and the Respondent, we summarized the arguments for each using the following query:

Explain the brief's context in one sentence; and the brief's legal argument in one sentence.

We then checked the summary with the brief, and reported it, along with accompanying metadata. The data can be found here.¹⁷⁵

D. Results

After reporting our results in a CSV file, we first compiled the relevant descriptive statistics, represented in Table 2. In Table 3, we reported the rate of hallucination, completeness, and groundedness per topic. Finally, we reported the outcome of the select scenarios in Table 4.¹⁷⁶

In totality, Arbitrus.ai performed as expected: it did not hallucinate; it answered the issues in controversy, and almost always—with two narrow exceptions—grounded those answers in relevant case law.¹⁷⁷ Of note, and in line with other

¹⁷⁵ *Supra* note 173.

¹⁷⁶ We edited cases to only include the preliminary statements and legal issues. We excluded the specific detailing of facts and jurisdiction. Future studies will need to examine specific subtopics—one such topic is insurance. *See infra* note 179 and accompanying text for more detail.

¹⁷⁷ The two ‘ungrounded’ cases were labeled such because of specific disanalogous facts. In the first case, *Beacon Partners LLC v. Sterling Collective, Inc.* (No. 80), the case contemplated a

studies, Arbitrus.ai generally ruled in favor of the respondent—which could indicate structural bias. Most likely, however, it is simply that in most cases favoring the respondents—especially where the proportion of respondent wins is exaggerated, such as insurance—the dispute and resolution are qualitatively rooted in law, whereas most of the claimant’s arguments invoke equitable claims. This matches a theoretical expectation of the law—that “equity follows the law”¹⁷⁸—as well as an empirical one—most insurance claims that are denied remain denied on appeal.¹⁷⁹ Finally, there are

breach of contract through a partnership; unfortunately, it cited *Greenberg v. Ladicorbic*, 200 A.D.2d 465 (N.Y. App. Div. 1994), which contemplates not a breach of contract, but whether the partnership existed at all. In the present case, it is uncontested that the partnership exists. The second instance, *StellarFranchise Group, Inc. v. BrightPath Ventures, LLC*, centers on alleged breaches of a franchise agreement. StellarFranchise claims that BrightPath failed to meet operational standards, misappropriated advertising funds, and unilaterally changed pricing policies, causing reputational and financial harm. The franchise agreement required BrightPath to maintain specific operational practices, to contribute to a centralized advertising fund, and to seek approval for pricing changes, all of which StellarFranchise asserts BrightPath failed to do. The arbitrator, however, cites *Sinclair Refining Co. v. Jenkins Petroleum Process Co.*, 289 U.S. 689 (1933). First, the case is too old and has questionable legal standing as related to the franchise issue in South Dakota. Second, in *Sinclair*, the case primarily focuses on a contractual obligation, not a fiduciary duty inherent to a franchisee claim. Notably, in the field, all of these problems would be caught by the final step of Fortuna’s process, “human interference.” See Figure 3.

¹⁷⁸ See, e.g., *Hedges v. Dixon County*, 150 U.S. 182, 192 (1893) (“The fact that the complainants have no remedy at law, arising from the invalidity of the bonds, confers no jurisdiction upon a court of equity to afford them relief. The established rule, although not of universal application, is that equity follows the law... “that wherever the rights or the situation of parties are clearly defined and established by law, equity has no power to change or unsettle those rights or that situation, but in all such instances the maxim *equitas sequitur legem* is strictly applicable.”).

¹⁷⁹ One research study suggests that only about 0.2% of healthcare claims were appealed; although it is of note that specific types of

also cases that point to the opposite conclusion; indeed, corporate cases disproportionately favored the claimant. Because of the small sample size, we can only extrapolate a generalized conclusion that Arbitrus.ai operated superior to the benchmarks applied against other legal products.¹⁸⁰

As a final note, this process was considerably more expensive than a normal chatbot. We employed approximately 26,617 requests and 907,744,509 million tokens to create the opinion database. These costs vary on the model type. These costs will likely increase greatly as case complexity increases. Moreover, these costs exclude the long-term expenses of hosting and maintaining machine learning models, both of which are necessary to make factual determinations.

cases, such as prior authorization requests, did have unexpectedly high overturn rates on appeal—83% to be specific. A larger sample size of cases on specific types of insurance disputes will be needed to confirm this number. Jeannie Fuglesten Biniek et al., *Use of Prior Authorization in Medicare Advantage Exceeded 46 Million Requests in 2022*, KFF (2024), <https://www.kff.org/medicare/issue-brief/use-of-prior-authorization-in-medicare-advantage-exceeded-46-million-requests-in-2022/#:~:text=The%20vast%20majority%20of%20denied,of%20the%20prior%20authorization%20process.> Of note, AI is not here to solve social problems—rather it is to efficiently replicate predictable social processes, whether for good or evil. If individuals are upset with the health insurance industry, they should petition the government to draft new regulations.

¹⁸⁰ It might be of value to create new metrics to analyze decision-making legal products. These metrics will naturally be more controversial, since they would make normative claims as to what the machine ought to do, rather than grading the machine outputs on ‘objective’ grounds.

Table 2. Descriptive Statistics

Topic	Cases		
	Claimant Wins	Respondent Wins	Aggregate
Contract	22	40	62
Insurance	2	10	12
Intellectual Property	4	7	11
Employment	2	3	5
Real Estate	0	3	3
Corporate	4	1	5
Other	0	2	2
Total	34	66	100

Table 3. Graded Results¹⁸¹

Topic	Cases		
	Hallucination	Correctness	Groundedness
Contract	0	0	2
Insurance	0	0	0
Intellectual Property	0	0	0
Employment	0	0	0

¹⁸¹ We reported the frequency of the issue. Thus, a ‘0’ indicates no issues with the given metric. A 0 for hallucination means there were 0 hallucinations, a 0 for correctness, means that each brief had the issues in controversy answered, and a 0 for groundedness means that each brief had the relevant law applied.

Real Estate	0	0	0
Corporate	0	0	0
Other	0	0	0
Total	0	0	2

Table 4. Selected Scenarios Results		
Case ID	Prevailing Party	Selected Excerpt of Ruling
4	Respondent	<p>This case presents a classic dispute over the interpretation of an executive employment contract and the circumstances surrounding termination. The claimant, Victoria Green, argues that Summit Dynamics breached her employment contract by failing to provide severance and equity compensation after her termination, which she contends was without cause. In contrast, the respondent, Summit Dynamics, asserts that Green was terminated for cause due to gross mismanagement and breaches of fiduciary duty, which exempted them from paying severance and equity compensation. After considering the arguments, I rule in favor of Summit Dynamics, as the evidence presented supports their claim of termination for cause based on documented performance failures and breaches of fiduciary duty.</p> <p>...</p> <p>The legal issue at hand is whether Summit Dynamics Corporation's termination of Victoria Green was for cause, thereby exempting them from their contractual obligations to provide severance and equity compensation.</p> <p>Rule: Under Illinois law, an employer may terminate an employee for cause if the employee engages in gross misconduct or breaches fiduciary duties as defined in their employment contract. In the case of <i>Martin v. Federal Life Insurance Co.</i>, 518 N.E.2d 306 (Ill. App. Ct. 1987), the court addressed the enforceability of employment contracts and the conditions under which an employer may terminate an employee without fulfilling severance obligations. The court found that an employer's failure to adhere to contractual obligations could be justified if the employee's actions constituted a breach of the contract or misconduct.</p> <p>In <i>Martin</i>, the plaintiff alleged wrongful termination based on an oral agreement for permanent employment. The court examined whether</p>

		<p>the employer's actions were justified under the terms of the contract and the applicable law. The court ultimately ruled that the employer's claims of misconduct were sufficient to uphold the termination, thereby allowing the employer to avoid severance payments.</p> <p>Application: In the present case, Victoria Green's termination was based on allegations of gross mismanagement and breaches of fiduciary duty. Summit Dynamics claims that Green's approval of a \$500,000 advertising campaign that violated contractual agreements with a major client, her exceeding the department's budget by 25% without authorization, and the creation of a toxic work environment justify her termination for cause. These allegations parallel the misconduct discussed in Martin, where the court upheld the employer's right to terminate based on the employee's actions that violated company policies and standards.</p> <p>Factually, both cases involve an employee whose actions led to significant negative consequences for the employer. In Martin, the employee's alleged misconduct was sufficient to justify termination without severance. Similarly, Green's actions, as claimed by Summit Dynamics, resulted in financial losses and reputational harm to the company, which could be interpreted as gross misconduct under the terms of her employment contract.</p> <p>Moreover, the employment contract in Green's case explicitly states that severance and equity provisions apply only in the event of termination without cause. Since Summit Dynamics asserts that Green's termination was for cause, they argue that they are not obligated to fulfill these provisions. This aligns with the ruling in Martin, where the court found that the employer was justified in terminating the employee based on documented misconduct, thus negating any claims for severance.</p> <p>Conclusion: Based on the legal precedent established in Martin v. Federal Life Insurance Co., Summit Dynamics Corporation's termination of Victoria Green appears to be justified under the terms of her employment contract and Illinois law. The allegations of gross mismanagement and breaches of fiduciary duty provide a sufficient basis for the termination to be classified as for cause, thereby exempting Summit from their contractual obligations to provide severance and equity compensation. As such, the ruling should favor Summit Dynamics in this arbitration.</p>
17	Claimant	<p>In this construction payment dispute, the stakes are high as both parties assert their rights under the contract. Crescent Bay Builders, LLC claims entitlement to \$180,000 for completed work and approved change orders, arguing that they fulfilled all contractual obligations and that</p>

		<p>Riverfront's claims of budget overruns and subcontractor issues are unfounded. Conversely, Riverfront Commercial Development Group contends that Crescent Bay delivered substandard work and missed deadlines, justifying their withholding of payments and asserting that Crescent Bay's claims for additional compensation are unsupported by the contract. After considering the arguments, I rule in favor of Crescent Bay Builders, LLC, as they have demonstrated compliance with the contract and provided sufficient documentation to support their claims for payment.</p> <p>...</p> <p>Legal Analysis The legal issue at hand is whether Crescent Bay Builders, LLC is entitled to recover the \$180,000 in unpaid amounts for completed work and approved change orders under the terms of the construction contract with Riverfront Commercial Development Group.</p> <p>Rule: A contractor may withhold payment for work that does not meet contractual standards or for which there is a legitimate dispute regarding the quality of the work performed. This principle is illustrated in the case of <i>Hester v. Burns Builders</i>, 92 So. 3d 641 (La. Ct. App. 2012). In <i>Hester</i>, the plaintiff refused to issue a scheduled progress payment due to alleged defects in the construction work performed by the defendant. The court emphasized that genuine issues of material fact regarding the alleged defects could justify withholding payment.</p> <p>In <i>Hester</i>, the factual context involved a contract for the construction of a new home, where the plaintiff contended that the quality of the defendant's work was sub-par. The plaintiff hired an expert to inspect the work, which led to the refusal to authorize further payments. The trial court initially granted summary judgment in favor of the defendant, but the appellate court reversed this decision, highlighting that the existence of genuine issues of material fact regarding the quality of work warranted further proceedings.</p> <p>Application: In the current dispute between Crescent Bay and Riverfront, similar issues arise regarding the quality of work and the entitlement to payment. Crescent Bay claims that it completed all work in compliance with the contract, including several change orders that were requested and approved by Riverfront. The project passed inspection and was delivered on schedule, which aligns with the plaintiff's position in <i>Hester</i>, where the work was also completed and inspected.</p>
--	--	--

		<p>Conversely, Riverfront contends that Crescent Bay delivered substandard work and missed deadlines, justifying their withholding of payments. This assertion mirrors the defendant's position in Hester, where the defendant claimed that the plaintiff's refusal to pay was unfounded. However, just as in Hester, the key factor is whether there are genuine issues of material fact regarding the quality of the work performed by Crescent Bay.</p> <p>Crescent Bay has provided detailed invoices and documentation of the completed work, similar to the expert report presented by the plaintiff in Hester. Riverfront's claims of budget overruns and subcontractor performance issues are vague and unsupported, akin to the defendant's arguments in Hester that did not sufficiently counter the plaintiff's evidence. The court in Hester ruled that the existence of such disputes warranted further examination, which is applicable here as well.</p> <p>Based on the legal principles established in Hester v. Burns Builders, Crescent Bay Builders, LLC has demonstrated compliance with the contract and provided sufficient documentation to support its claims for payment. The genuine issues of material fact regarding the quality of work and the legitimacy of Riverfront's claims necessitate a ruling in favor of Crescent Bay. Therefore, Crescent Bay is entitled to recover the \$180,000 in unpaid amounts for completed work and approved change orders.</p>
24	Respondent	<p>This case revolves around Harper Industries' claim against NorthStar Insurance Company for coverage of losses resulting from a fire at its manufacturing facility. The claimant, Harper, argues that the fire was a covered peril under the policy and that NorthStar wrongfully denied coverage despite Harper's compliance with all policy terms. Harper contends that the exclusions cited by NorthStar, namely "negligence in maintenance" and "wear and tear," are inapplicable to the circumstances of the loss, as the fire was caused by an unexpected electrical fault.</p> <p>In contrast, the respondent, NorthStar, asserts that the denial was justified based on the policy exclusions for inadequate maintenance and wear and tear, claiming that Harper's failure to maintain its equipment directly contributed to the fire. NorthStar argues that the fire was not an independent event but rather a result of foreseeable equipment failure due to neglect.</p> <p>After considering the arguments presented, I rule in favor of NorthStar Insurance Company, as the evidence supports that Harper's lack of maintenance contributed to the incident, falling within the policy exclusions.</p> <p>...</p>

		<p>Legal Analysis</p> <p>Issue: The legal issue at hand is whether NorthStar Insurance Company properly denied coverage for Harper Industries' claim based on the exclusions for "negligence in maintenance" and "wear and tear" in the insurance policy.</p> <p>Rule: Under Oregon law, an insurer may deny coverage based on specific exclusions outlined in the policy if the circumstances of the claim fall within those exclusions. In the case of <i>Palmrose v. Oregon Insurance Guaranty Association</i>, 205 Or. App. 613 (Or. Ct. App. 2006), the court held that a claim was not covered because the insurer was not a member of the Oregon Insurance Guaranty Association at the time of the incident, and the policy did not meet the statutory requirements for coverage. The court emphasized that the definitions of "covered claim" and "insolvent insurer" are critical in determining coverage eligibility.</p> <p>In <i>Palmrose</i>, the facts involved a personal injury and wrongful death claim against an assisted living center, where the insurance policy was issued by a surplus lines insurer that later became insolvent. The court ruled that the claim was not covered because the policy did not apply under the relevant statutes, highlighting the importance of the insurer's status and the nature of the policy in determining coverage.</p> <p>Application: In the current case, Harper Industries submitted a claim for \$1.2 million in damages resulting from a fire at its manufacturing facility. NorthStar Insurance Company denied the claim, citing exclusions for "negligence in maintenance" and "wear and tear." The facts indicate that the fire was caused by an electrical fault, which NorthStar argues was exacerbated by Harper's failure to maintain its equipment properly.</p> <p>Comparing the facts of <i>Palmrose</i> to the current case, both involve claims where the insurer seeks to deny coverage based on specific policy exclusions. In <i>Palmrose</i>, the insurer's status and the nature of the policy were critical in determining coverage eligibility. Similarly, in this case, NorthStar's reliance on the maintenance exclusions is central to its denial of Harper's claim. The rule established in <i>Palmrose</i> suggests that if the circumstances of the claim align with the exclusions outlined in the policy, the insurer may rightfully deny coverage.</p> <p>In applying the rule to the facts, it is evident that Harper's alleged lack of maintenance directly contributed to the incident leading to the claim. NorthStar's argument that the fire was a result of inadequate maintenance aligns with the exclusions stated in the policy. Just as the</p>
--	--	---

		<p>court in Palmrose ruled that the claim was not covered due to the insurer's status and the policy's limitations, NorthStar's denial of coverage is justified based on the exclusions for negligence and wear and tear.</p> <p>Conclusion: Based on the analysis of the issue, rule, and application of the relevant case law, it is concluded that NorthStar Insurance Company properly denied coverage for Harper Industries' claim. The evidence supports that Harper's lack of maintenance contributed to the incident, falling within the policy exclusions. Therefore, the ruling should favor NorthStar Insurance Company, affirming its denial of the claim based on the applicable exclusions in the insurance policy.</p>
45	Claimant	<p>This arbitration case highlights the critical importance of cybersecurity in the modern business landscape. The claimant, BlueHaven Financial Group, LLC, argues that SecureNet IT Services, Inc. breached their managed services agreement by failing to provide adequate security measures, resulting in a data breach that exposed sensitive client information and caused significant financial losses. BlueHaven claims damages of \$350,000 for financial losses, \$75,000 for reimbursement of fees paid, and attorney's fees. In contrast, the respondent, SecureNet, contends that it fulfilled its contractual obligations and that the breach was primarily due to BlueHaven's internal failures, including its refusal to implement recommended security measures and inadequate employee training. After considering the arguments, I rule in favor of BlueHaven Financial Group, LLC, as SecureNet's failure to implement basic cybersecurity protocols directly contributed to the breach and the resulting damages.</p> <p>...</p> <p>Legal Analysis Issue: The legal issue at hand is whether SecureNet IT Services, Inc. breached its contractual obligations to BlueHaven Financial Group, LLC by failing to implement adequate cybersecurity measures, resulting in a data breach that exposed sensitive client information.</p> <p>Rule: Under Kentucky law, service providers have a duty to exercise reasonable care in protecting their clients' sensitive information. This duty includes implementing industry-standard security measures to prevent data breaches. The Kentucky Personal Information Security Breach Notification Requirements outline the obligations of agencies regarding data security and breach notifications, emphasizing the importance of timely and effective responses to security incidents (KRS 61.931 to 61.934).</p>

		<p>Application: In the case of BlueHaven Financial Group, LLC v. SecureNet IT Services, Inc., the facts reveal that BlueHaven entered into a managed services agreement with SecureNet, which included responsibilities for monitoring IT systems, implementing encryption, conducting vulnerability assessments, and responding to incidents within two hours. However, SecureNet failed to implement multi-factor authentication (MFA) and did not adequately update its email gateway, which left BlueHaven's systems vulnerable to a phishing attack. This breach exposed the personal financial information of approximately 2,000 clients and resulted in significant financial losses and reputational harm for BlueHaven.</p> <p>Comparatively, the Kentucky Personal Information Security Breach Notification Requirements establish that organizations must notify affected individuals and relevant authorities of a breach within a specified timeframe. In this case, SecureNet's failure to detect the breach until 72 hours after it began directly contravenes the expectations set forth in the law. The lack of timely response and inadequate security measures demonstrate a breach of the duty of care owed to BlueHaven.</p> <p>Furthermore, the rule applies to the facts of this case as SecureNet's actions (or lack thereof) directly led to the breach. The failure to implement basic cybersecurity protocols, such as MFA, constitutes a clear violation of the contractual obligations outlined in the managed services agreement. The breach not only resulted in financial losses for BlueHaven but also triggered regulatory investigations and client complaints, further exacerbating the harm caused.</p> <p>Conclusion: Based on the analysis of the legal issue, the applicable rule, and the application of the facts to the rule, it is evident that SecureNet IT Services, Inc. breached its contractual obligations to BlueHaven Financial Group, LLC. The failure to implement adequate cybersecurity measures and respond promptly to the breach directly contributed to the significant financial and reputational harm suffered by BlueHaven. Therefore, BlueHaven is entitled to the relief sought in this arbitration, including damages for financial losses, reimbursement of fees, and attorney's fees.</p> <hr/> <p>Conclusion After considering the arguments, I rule in favor of BlueHaven Financial Group, LLC, as SecureNet's failure to implement basic cybersecurity protocols directly contributed to the breach and the resulting damages.</p>
49	Respondent	In this arbitration dispute, the tension between mining rights and contractual obligations takes center stage. The claimant, Northern Ridge

	<p>Mining, LLC, argues that RedRock Resources Ltd. breached their mining rights and royalty agreement by restricting access and refusing to accept royalty payments, resulting in \$300,000 in unpaid royalties and \$150,000 in damages. Conversely, the respondent, RedRock, contends that Northern Ridge violated environmental and safety regulations, justifying their actions to restrict access and suspend royalty payments. After considering the arguments, I rule in favor of RedRock Resources Ltd. because Northern Ridge's non-compliance with the agreement's terms undermines its claims for unpaid royalties and damages.</p> <p>...</p> <p>Legal Analysis The primary legal issue is whether RedRock Resources Ltd. breached the mining rights and royalty agreement with Northern Ridge Mining, LLC by restricting access and refusing to accept royalty payments, or whether Northern Ridge's alleged violations of environmental and safety regulations justified RedRock's actions.</p> <p>Rule: The rule governing this dispute is that parties to a contract must adhere to the specific terms of the agreement, and any unilateral changes or breaches can lead to liability. In the case of <i>Gypsum Aggregates Corp. v. Lionelle</i>, 170 Colo. 282 (Colo. 1969), the court emphasized the importance of the parties' adherence to the terms of their agreement regarding mining claims. The court found that the determination of the nature of mineral deposits was a factual issue that needed to be resolved based on the evidence presented, and the trial court's findings were upheld due to the presence of competent evidence supporting its conclusions.</p> <p>In <i>Gypsum Aggregates Corp. v. Lionelle</i>, the plaintiff sought a declaration regarding the nature of a mineral deposit and whether it was subject to location as a placer or lode mining claim. The trial court found that the deposit was placer in nature, and the plaintiff's request for a new trial based on insufficient evidence was denied. The court ruled that the trial court's findings were supported by competent evidence, reinforcing the principle that contractual obligations must be honored unless a valid reason for non-compliance is established.</p> <p>Application: In the current dispute, Northern Ridge Mining, LLC claims that RedRock Resources Ltd. breached their agreement by imposing access restrictions and refusing to accept royalty payments. Northern Ridge argues that it complied with the terms of the agreement, including timely royalty payments and adherence to operational</p>
--	---

		<p>guidelines. However, RedRock contends that Northern Ridge violated environmental and safety regulations, which justified their actions to restrict access and suspend royalty payments.</p> <p>Factually, both cases involve disputes over compliance with contractual obligations. In Gypsum Aggregates Corp. v. Lionelle, the court upheld the trial court's findings based on the evidence presented regarding the nature of the mineral deposit. Similarly, in the current case, the determination of whether Northern Ridge complied with environmental regulations is crucial. If Northern Ridge did indeed violate these regulations, as RedRock claims, then RedRock's actions to restrict access may be justified under the terms of the agreement.</p> <p>Furthermore, the Gypsum case illustrates that the burden of proof lies with the party claiming a breach. In this instance, Northern Ridge must demonstrate that RedRock's actions were unjustified and constituted a breach of the agreement. Conversely, RedRock must provide evidence of Northern Ridge's non-compliance with environmental regulations to support its defense.</p> <p>Conclusion: Based on the analysis of the legal issue, the applicable rule, and the factual context of both cases, it can be concluded that RedRock Resources Ltd.'s actions may be justified if Northern Ridge Mining, LLC is found to have violated the terms of the agreement regarding environmental and safety regulations. The precedent set in Gypsum Aggregates Corp. v. Lionelle reinforces the importance of adhering to contractual obligations and the necessity of providing competent evidence to support claims of breach. Therefore, if Northern Ridge cannot sufficiently demonstrate compliance with the agreement, the ruling should favor RedRock.</p>
80	Claimant	<p>In a classic case of contractual disagreement, Beacon Partners LLC claims that Sterling Collective, Inc. has unjustly withheld profit shares despite fulfilling its obligations. The claimant, Beacon, argues that it successfully contributed to securing contracts that generated significant profits and that Sterling's claims of deficient performance are vague and unsupported. Conversely, the respondent, Sterling, contends that Beacon failed to meet the necessary performance metrics outlined in their partnership agreement, justifying their decision to withhold payments. After considering the arguments, I rule in favor of Beacon Partners LLC, as Sterling's lack of specific evidence to support its claims of deficient performance constitutes a breach of the partnership agreement.</p> <p>...</p>

	<p>Legal Analysis</p> <p>Issue: The legal issue at hand is whether Sterling Collective, Inc.'s withholding of profit shares from Beacon Partners LLC constitutes a breach of the partnership agreement.</p> <p>Rule: The rule governing this issue is that a party may not withhold payments under a contract without specific evidence justifying such action. In the case of <i>Greenberg v. Ladicorbic</i>, 200 A.D.2d 465 (N.Y. App. Div. 1994), the court held that claims premised on the existence of a partnership must fail in the absence of evidence showing that the claimant contributed to the capital of the alleged partnership or was to share in its losses. The court emphasized that without evidence supporting the claims, the party withholding payment could not justify its actions.</p> <p>In <i>Greenberg</i>, the plaintiff sought relief based on the assertion of a partnership, but the court found that there was no evidence of capital contribution or shared losses, leading to the dismissal of the claims. The factual context involved a dispute over the existence of a partnership and the obligations arising from it, where the court required concrete evidence to support the claims made by the plaintiff.</p> <p>Application: In the present case, Beacon Partners LLC claims that Sterling Collective, Inc. has unjustly withheld profit shares amounting to \$250,000, despite Beacon fulfilling its obligations under the partnership agreement. Similar to the <i>Greenberg</i> case, where the court required evidence of partnership contributions, Beacon must demonstrate that it met the performance metrics outlined in the agreement to justify its entitlement to the profit shares.</p> <p>Factually, Beacon asserts that it contributed to securing contracts that generated \$1.67 million in profits for Sterling, which aligns with the requirement of demonstrating value-added contributions. In contrast, Sterling claims that Beacon's performance was deficient and that it failed to meet the necessary metrics. However, Sterling's assertions lack specific evidence, as they have not provided documentation to substantiate their claims of Beacon's underperformance. This parallels the <i>Greenberg</i> case, where the absence of evidence led to the dismissal of claims.</p> <p>Furthermore, in <i>Greenberg</i>, the court highlighted that the existence of a partnership must be supported by evidence of capital contribution or shared losses. In this arbitration, Beacon's entitlement to profit shares is based on the partnership agreement, which explicitly states that Beacon is entitled to 15% of Sterling's profits. The lack of evidence from Sterling</p>
--	--

		<p>to support its claims of deficient performance undermines its position, just as the absence of evidence in Greenberg led to the dismissal of the plaintiff's claims.</p> <p>Conclusion: Based on the analysis of the legal rule established in Greenberg v. Ladicorbic and its application to the present case, it is clear that Sterling Collective, Inc.'s withholding of profit shares from Beacon Partners LLC constitutes a breach of the partnership agreement. The absence of specific evidence to support Sterling's claims of deficient performance, coupled with Beacon's fulfillment of its obligations, supports a ruling in favor of Beacon. Therefore, Beacon is entitled to the relief sought, including the immediate payment of \$250,000 in profit shares, interest, and reimbursement for legal fees and arbitration costs.</p>
99	Respondent	<p>In a case that highlights the complexities of insurance enrollment, Emily Carter contends that Atlantic Health Assurance unjustly denied her health insurance claim. The claimant argues that she acted in good faith, promptly submitting the required documentation and reasonably believing her application was approved based on Atlantic's communication. Conversely, Atlantic maintains that Carter's application was incomplete due to insufficient proof of income, which she failed to rectify despite multiple requests for additional documentation. After considering the arguments, the ruling is in favor of Atlantic Health Assurance, as the claimant did not fulfill the necessary requirements outlined in the insurance policy to activate her coverage.</p> <p>...</p> <p>Legal Analysis</p> <p>The legal issue at hand is whether Atlantic Health Assurance, Inc. was justified in denying Emily Carter's health insurance claim based on her alleged failure to provide complete and verifiable documentation during the application process.</p> <p>Rule: An insurance company has a duty to communicate clearly and in good faith with applicants regarding the status of their applications and any deficiencies that may exist. This duty includes providing timely feedback and opportunities for applicants to rectify any issues that could prevent the activation of their coverage. This principle is supported by the ruling in Lucas v. State Farm Fire & Casualty Co., 864 F. Supp. 2d 1346 (M.D. Ga. 2012), where the court found that the insured's failure to comply with policy conditions justified the insurer's denial of coverage.</p> <p>In Lucas, the plaintiff, R. Wesley Lucas, filed a claim for fire loss under</p>

	<p>his homeowner's insurance policy with State Farm. The court found that Lucas breached the insurance contract by failing to submit to a required examination under oath, which was a condition precedent to bringing suit. The court emphasized that the insurer is entitled to require its insured to abide by the policy terms, and failure to comply with such conditions precludes recovery as a matter of law.</p> <p>Application: In the case of Emily Carter, she submitted her application for health insurance coverage to Atlantic Health Assurance, which required personal identification, proof of income, and prior coverage documentation. On July 15, 2024, Atlantic informed Carter that her proof of income was insufficient and requested an updated version. Carter promptly provided the requested documentation on July 18, 2024. However, she received no further communication from Atlantic regarding her application, leading her to reasonably believe that her application was complete and approved.</p> <p>Comparing this to the facts in Lucas, both cases involve an applicant who submitted documentation to an insurance company and received communication regarding the status of their application. In Lucas, the plaintiff's failure to comply with the requirement for an examination under oath resulted in the denial of his claim. Similarly, in Carter's case, Atlantic's communication regarding the insufficiency of her proof of income created an obligation for her to rectify the issue. However, unlike Lucas, where the insured failed to comply with a clear requirement, Carter believed she had fulfilled her obligations based on the absence of further requests and the automated confirmation message from Atlantic.</p> <p>The key distinction lies in the clarity of communication from the insurer. In Lucas, the court found that the insured's non-compliance with the policy conditions justified the insurer's denial of coverage. In Carter's situation, Atlantic's failure to provide timely feedback after her second submission left her without the opportunity to address any deficiencies, which contributed to her belief that her application was approved. This lack of communication may indicate that Atlantic did not fully uphold its duty to inform Carter of the status of her application.</p> <p>Conclusion: Based on the principles established in Lucas v. State Farm Fire & Casualty Co., Atlantic Health Assurance, Inc. may have acted within its rights to deny Emily Carter's claim due to her failure to provide complete documentation. However, the insurer's lack of clear communication regarding the status of her application and the insufficiency of her proof of income raises questions about whether it fulfilled its duty to inform her adequately. Given these circumstances,</p>
--	---

		while Atlantic may have a valid basis for denying the claim, the failure to communicate effectively could undermine its position, suggesting that the denial may not be entirely justified.
--	--	---

IV. Objections

If the idea of Judge.ai stirred the pot, Arbitrus.ai kicks it over; we imagine you might have some objections. And they'll be harder to address. That's the trouble with making Judge.ai tangible—there's no more hypothetical for you to fight. Where Judge.ai was free to skate past present day implications, Arbitrus.ai must *actually comport with it*. We'll address some positive objections first, then some normative ones.

A. Some Positive Objections

What's the law on arbitrations these days? You'll find it in the Federal Arbitration Act, which states that arbitration clauses in private contracts “shall be valid, irrevocable, and enforceable save upon such grounds as exist at law or in equity for the revocation of any contract or as otherwise provided in chapter 4.”¹⁸² So, the only barriers between Arbitrus.ai and the sweet, broad safety of a textual “shall” are law, equity, and chapter 4 of the FAA¹⁸³—and that's where most procedural objections will stem from. We'll briefly survey what we feel are the most salient ones.

Up first, the word that's been on 1L's minds since the abstract: *unconscionability*, equity's final stand against privatization.¹⁸⁴ Quintessentially vague, it's the doctrine

¹⁸² 9 U.S.C. § 2.

¹⁸³ We won't spend time on Chapter 4 of the FAA—but, if you must know, it deals with “arbitration of disputes involving sexual assault and sexual harassment.” 9 U.S.C. § 401–02. Read it on your own time; suffice it to say that Arbitrus.ai has some extra hoops to jump through if it wants to be your Title IX coordinator. As with any set of hoops, it could easily clear them.

¹⁸⁴ *See, e.g., Terrell v. Kiromic Biopharma, Inc.*, 297 A.3d 610, 624 (Del. 2023) (“Unconscionability—‘a concept that is used sparingly’—traditionally requires that ‘the party with superior

courts use when there's nothing technically wrong with the contract but it's just ... you know, *bad for society*.¹⁸⁵

That “badness” can manifest in two places: the formation of the contract (like when one party doesn't really understand what he's signing up for) and the terms of the contract itself (when the party *does* understand what he's signing up for but, dammit, what he's signing up for is just *really, really bad*).¹⁸⁶

bargaining power used it to take unfair advantage of his weaker counterpart.”).

¹⁸⁵ See RESTATEMENT (SECOND) OF CONTRACTS § 208 (AM. L. INST. 1981) (“[Unconscionability]... also overlaps with rules which render particular bargains or terms unenforceable on grounds of public policy. Policing against unconscionable contracts or terms has sometimes been accomplished “by adverse construction of language, by manipulation of the rules of offer and acceptance or by determinations that the clause is contrary to public policy or to the dominant purpose of the contract.”); *Fritz v. Nationwide Mut. Ins. Co.*, 1990 Del. Ch. LEXIS 193, at *10 (Del. Ch. 1990) (“This section is intended to make it possible for courts to consider questions of unconscionability directly and explicitly. It is aimed not at the instance where superior bargaining power is present, but rather at the ‘prevention of oppression and unfair surprise.’ . . . ‘The basic test is whether, in light of the general commercial needs of the particular trade or case, the clauses involved are so *one-sided as to be unconscionable under the circumstances existing at the time of the making of the contract.*’”) (emphasis added).

¹⁸⁶ For example, contracts designed to exploit a vulnerable class of people. See, e.g., *Williams v. Walker-Thomas Furniture Co.*, 350 F.2d 445, 449 (D.C. Cir. 1965) (“Did each party to the contract, considering his obvious education or lack of it, have a reasonable opportunity to understand the terms of the contract, or were the important terms hidden in a maze of fine print and minimized by deceptive sales practices? Ordinarily, one who signs an agreement without full knowledge of its terms might be held to assume the risk that he has entered a one-sided bargain. But when a party of little bargaining power, and hence little real choice, signs a commercially unreasonable contract with little or no knowledge of its terms, it is hardly likely that his consent, or even an objective manifestation of his consent, was ever given to all the terms.”); see also *Terrell*, 297 A.3d at 624 (“[C]ourts are particularly reluctant to find unconscionability in contracts between sophisticated [parties].”); RESTATEMENT (SECOND) OF CONTRACTS § 208 (AM. L.

Procedural unconscionability is inappropriate here because Arbitrus.ai has no role in the formation of its own enabling contract; it falls on the parties to be fully transparent with one another about the benefits and risks of automated arbitration.¹⁸⁷ But is deference to a robot arbitrator a *substantively* unconscionable term to put in a contract?

Most will answer ‘yes’ purely because a robot *feels sketchy*. But let’s face it: Arbitration *itself* is sketchy. Indeed, it is the federal judiciary if judges were removable at-will by the President; it forces one-shotters to argue, under opposing-house rules, in front of a costume jurist who’s getting *paid* by the repeat player. It’s not angelic; businesses (allegedly) do it to win cases they would have lost in a real court.¹⁸⁸ But, despite that sketchiness, “there is nothing per se unconscionable about an agreement to arbitrate.”¹⁸⁹ And it’s

INST. 1981) (citing *Hume v. United States*, 132 U.S. 406, 406 (1889) (“Traditionally, a bargain was said to be unconscionable in an action at law if it was “such as no man in his senses and not under delusion would make on the one hand, and as no honest and fair man would accept on the other;” damages were then limited to those to which the aggrieved party was “equitably” entitled.”)).

¹⁸⁷ See 2 BRADLEY W. VOSS, *DELAWARE CONTRACT LAW* § 10.171 (5th ed. 2024) (“The concept of procedural unconscionability examines the procedures that led to the contract with the goal of evaluating whether seemingly lopsided terms might have resulted from arms’-length bargaining. Courts focus on the relative bargaining strength of the parties and whether the weaker party could make a meaningful choice. The concept is ‘broadly conceived to encompass not only the employment of sharp bargaining practices and the use of fine print and convoluted language, but a lack of understanding and an inequity of bargaining power.’...”).

¹⁸⁸ See RESTATEMENT (SECOND) OF CONTRACTS § 208 cmt. d (AM. L. INST. 1981) (“...factors weighing in favor of a finding of unconscionability... (stress of hospital admitting room, superior bargaining position of hospital, failure to call patient's attention to arbitration clause) (it is not clear whether the court was relying on the unconscionability concept or on one limited to contracts of adhesion)...”).

¹⁸⁹ *Lucchese Boot Co. v. Solano*, 473 S.W.3d 404, 418 (Tex. App. 2015).

not even unconscionable when one repeat player gives a specific arbitrator a *ton* of business—such that, relative to one-shotters, they have *disproportionate* experience with the decision-maker.¹⁹⁰ If any of that bothers you, your gripe isn't with automation; it's with *FAA deference*.

Ultimately, arbitration has only one failure condition: “An arbitration provision is substantively unconscionable if it fails to provide a neutral arbitrator.”¹⁹¹ *That's it*. Under this construction, Arbitrus.ai is actually *far more* conscionable than a human arbitrator. It's a *robot*; it physically cannot be partial—it has no ability to engage in personal “misconduct,” such as owning a part of a business.¹⁹² Now, sure, you can *try* to hack it or otherwise distort the system.¹⁹³ And sure, you can program it to say “company always wins.” But in doing so, you leave a paper trail—robots keep a record of their *mens rea*

¹⁹⁰ Remember our friends at the American Arbitration Association (AAA)? The biggest arbitration firm in America? They get sued for this *all the time*, as tons of companies have them on retainer. But here's the Ninth Circuit: “California courts have uniformly concluded that the AAA provides a neutral forum for dispute resolution.” *Nagrampa v. MailCoups, Inc.*, 469 F.3d 1257, 1302 (9th Cir. 2006) (O'Scannlain, J., concurring in pertinent part). Translation: *Nothing to see here*. We'll note that AAA is a nonprofit—which should, in theory, alleviate some scrutiny. But it's still a vehicle that pays arbitrators for arbitrating. *There needs to be money to pay those guys*.

¹⁹¹ *Archer v. TIC-The Indus. Co.*, No. CV 16-6649-GHK (SSX), 2016 WL 7635956, at *6 (C.D. Cal. Nov. 7, 2016).

¹⁹² *Compare Monster Energy Co. v. City Bevs. LLC*, 940 F.3d 1130, 1138 (9th Cir. 2019) (“Here, the Arbitrator's failure to disclose his ownership interest in JAMS—given its nontrivial business relations with Monster—creates a reasonable impression of bias and supports vacatur of the *arbitration award*.”) (emphasis added), *with Certain Underwriting Members of Lloyds of London v. Florida*, 892 F.3d 501, 507 (2d Cir. 2018) (“[we] will not *vacate arbitration awards* for *evident partiality* when the party opposing the award ‘identifies no direct connection between [the arbitrator] and the outcome of the arbitration.’”) (emphasis added).

¹⁹³ Alignment and security against prompt hacking is delegated first to Ra.ai, and second to the final human reviewer. Ra.ai rewrites each query to protect against prompt hacking.

called *code*.¹⁹⁴ Compare that to a human arbitrator who knows full well that, “if Acme Co. loses *every time* it arbitrates with me, it’s probably going to find a different arbitrator.” What does the arbitrator do with that knowledge?¹⁹⁵ Does he ever fudge the results to curry favor? And if he did, *would we ever know?* With Arbitrus, companies may well enact the first arbitration clauses that are *empirically conscionable*. Anyone retorting that humanity somehow holds a monopoly on conscionability places entirely too much stock in mounds of faceless process; a machine composed of hollow men is no less metallic.

Next up, an objection from “law”: Is the *ex ante* query function legal? Obviously, we’d have a problem if *federal courts* were issuing the kinds of advisory orders Arbitrus.ai can one day promise—that’s a *Constitutional* question, which wouldn’t impact arbitration.¹⁹⁶ But the FAA has a provision

¹⁹⁴ This is how FTX’s insurance fraud was discovered—the entire fund’s investment was calculated by a line of python coding that multiplied the daily trading volume by a random integer around 7500. See Thomas Claburn, *SBF on Trial: The Python Code that Allegedly Let Alameda Hedge Fund Spend People’s FTX Deposits*, THE REG. (Oct. 10, 2023), https://www.theregister.com/2023/10/10/ftx_python_code/ (“return f2d(numpy.random.normal(7500, 3000)) * daily_volume / Decimal('1e9')”).

¹⁹⁵ The question applies equally (and more realistically) to arbitration companies. Indeed, AAA tends to get a pass from courts because “[its] rules . . . are generally regarded to be neutral and fair.” *Izzi v. Mesquite Country Club*, 186 Cal. App. 3d 1309, 1318 (1986). Translation: AAA has internal rules designed to keep its arbitrators on the straight and narrow—and, apparently, courts trust that they work. Well, Arbitrus.ai has those too—it has parameters built into its code. And it *must follow them*; it can’t go rogue like a human arbitrator can. We must remember that organizations of human beings operate *as machines*—AAA, taken as a whole, *is Arbitrus*. It’s just *worse*.

¹⁹⁶ But is it *really* so obvious? One of us (Kieffaber) isn’t convinced that Article III’s “cases and controversies” language is sufficient to prevent federal judges from issuing advisory opinions. At any rate, that language certainly cannot purport to bar advisory orders in arbitrations between private parties.

that arguably bars such orders in the arbitration context, too: 9 U.S.C. § 10(a)(2)'s bar on "evident partiality." To be clear, that's redundant with unconscionability—but courts have specifically used § 10(a)(2) to strike down *ex parte communications*.¹⁹⁷ In theory, that's what an advisory order would be; the machine telling one party, out of the other's earshot, what it would hold in X scenario.¹⁹⁸

Now, here's an easy way around that: Arbitrus.ai could send a notice to all contractual parties whenever it issued an advisory order to just one party. Indeed, "[a]n *ex parte* communication occurs when a judge communicates with any person concerning a pending or impending proceeding without notice to an adverse party."¹⁹⁹ Notice solves the issue. But this tees up a broader point: FAA case law needs to adapt to the age of automated arbitration. An *ex parte* communication is a *human* folly—threatening wink-and-nod impropriety at best, outright collusion at worst. Robot judges don't wink; there is no risk inherent in automated *ex parte* communications, because AI can be wholly deterministic.²⁰⁰

¹⁹⁷ See, e.g., *Amoco D.T. Co. v. Occidental Petroleum Corp.*, 343 S.W.3d 837, 845 (Tex. App. 2011); *Fid. Fed. Bank, FSB v. Durga Ma Corp.*, 386 F.3d 1306, 1309-10 (9th Cir. 2004).

¹⁹⁸ To be clear, Arbitrus.ai will not provide *ex ante* query functions in its initial deployment. In the future, it will be possible—but nothing in this paper suggests the *current* IP performs this function.

¹⁹⁹ *State v. Lierman*, 305 Neb. 289, 306 (2020); see also *Commonwealth v. Gonzalez*, 112 A.3d 1232, 1240 (Pa. Super. Ct. 2015) ("[A]n *ex parte* communication, by definition, involves the inclusion of one party in a consultation with a judge over the exclusion of another."); *In re D.M.*, 993 A.2d 535, 541 (D.C. 2010) ("*Ex parte* communications are those that involve fewer than all of the parties who are legally entitled to be present during the discussion of any matter and are prohibited in order to ensure that every person who is legally interested in a proceeding is given the full right to be heard according to law.>").

²⁰⁰ Even with temperature, which would make the *ex ante* query function not wholly deterministic, the result would still be within the range of probability. It is also notable that, at least in this paper, temperature only applies to LLMs—indeed, the actual

Perhaps more validly concerning are *ex parte* communications between a party and the robot’s developers; these should certainly be disclosed to relevant parties. But even if they *weren’t* disclosed, we repeat: The machine runs on code and thus leaves a trail when it cheats. FAA case law (and the FAA itself) should be modernized to better probe that trail—and establish different regimes for flesh and metal arbitrators.

B. Some Normative Objections

Kieffaber’s Judge.ai was, in part, a commentary on legal formalism. And that played nicely in the public law context, where formalism has been increasingly dominant.²⁰¹ But is formalism equally rampant in the contract space? It’s at least *present*; the baseline is Bill Brennan textualism,²⁰² where “[e]xtrinsic evidence of the parties’ subjective intent may be resorted to only if the document is ambiguous.”²⁰³ Especially in the corporate and commercial contexts, states like Delaware have increasingly applied “contractarian”²⁰⁴ positions of interpretation between sophisticated parties.²⁰⁵ Of course,

determination of the prevailing company comes from ML models, which are not subject to temperature.

²⁰¹ Kieffaber, *supra* note 8, at 13–14.

²⁰² See *H.J. Inc. v. Nw. Bell Tel. Co.*, 492 U.S. 229, 236–237 (1989) (Brennan, J.) (“We begin, of course, with RICO’s text ... Our guides in the endeavor must be the text of the statute....”).

²⁰³ *District of Columbia v. D.C. Pub. Serv. Comm’n*, 963 A.2d 1144, 1156 (D.C. 2009); see also *929 Flushing LLC v. 33 Dev. Inc.*, 52 Misc. 3d 195, 199 (N.Y. Sup. Ct. 2016) (“When the terms of a written contract are clear and unambiguous, the intent of the parties must be found within the four corners of the contract, giving practical interpretation to the language employed and the parties’ reasonable expectations.”).

²⁰⁴ This refers to textualism, but for contracts.

²⁰⁵ See, e.g., *L-5 Healthcare Partners, LLC v. Alphatec Holdings, Inc.*, 2024 Del. Ch. LEXIS 301, at *12–13 (Del. Ch. Aug. 21, 2024) (“To say that Delaware prides itself on the contractarian nature of its law risks understatement: ‘This jurisdiction respects the right of parties to freely contract and to be able to rely on the enforceability of their agreements; where Delaware’s law applies, with very limited exceptions, our courts will enforce the

common law judges can throw in the ambiguity towel pretty early—and, when they do, subjectivity is fair game.²⁰⁶

So, contract interpretation is not a literalist endeavor; at its most formalist, it's a microcosmic approximation of original meaning textualism.²⁰⁷ We think that's enough to suggest that a formalist should, at the very least, have no opposition to Arbitrus.ai. Indeed, Kieffaber has emphasized elsewhere that AI can easily tackle original meaning inquiries²⁰⁸—and that would be platonically true of Arbitrus.ai contracts, as the machine would be the same age as the written words and trained on a linguistic corpus coterminous with public understanding at the time of the contract's drafting. But note also that, when parties form an Arbitrus.ai contract, they are stipulating that there *shall be no ambiguity*. Indeed, what the machine says *goes*; the remedy for an undesirable holding is not an appeal of the holding, but an *amendment to the machine*.²⁰⁹ The crux of automation is that it arbitrarily purges

contractual scheme that the parties have arrived at through their own self-ordering, both in recognition of a right to self-order and to promote certainty of obligations and benefits.”).

²⁰⁶ See, e.g., RESTATEMENT (SECOND) OF CONTRACTS § 220 (AM. L. INST. 1981) (“Language and conduct are in general given meaning by usage rather than by the law, and ambiguity and contradiction likewise depend upon usage. Hence usage relevant to interpretation is treated as part of the context of an agreement in determining whether there is ambiguity or contradiction as well as in resolving ambiguity or contradiction.”).

²⁰⁷ See *Snell*, 102 F.4th at 1233 (Newsom, J., concurring). (“The ordinary-meaning rule has an important corollary—namely, that words must be given the meaning they had when the text was adopted. That principle—'originalism,' if you will—most obviously applies to constitutional and statutory texts. To be clear, though, it holds for private instruments like contracts and insurance policies, as well.”) (internal quotation omitted).

²⁰⁸ Kieffaber, *supra* note 8, at 54.

²⁰⁹ Here, any resulting randomness is precisely what would bring the parties to the bargaining table and facilitate the culture of amendment. Whenever the machine has to “flip a coin” to make a judgment call, the party that wins the toss is incentivized to work with the losing party to patch the bug in some compromissory

ambiguity to achieve predictability; the formalist impulse that would interpret contracts textually when the words are clear should covet this regime.²¹⁰

Of course, we've made some to-do about arbitration being a market. That means egghead legal theorists aren't the relevant constituency; *parties themselves* must desire contractual formalism in order for Arbitrus.ai to *sell*. Do they? Well, a party's *preferred* method of contractual interpretation is *whatever makes them win every time*. But predictability is cheap and saves time, making it a strong second-best. Since the FAA effectively bars the parties' first choice, we're betting that parties will happily settle for their second choice—a system that will at least enforce their verbatim terms rather than trying to play hero for the one-shot claimant and fudging the edges.

way—lest that winner wind up a loser on the next random toss and find itself stonewalled in retaliation.

²¹⁰ We should also note that, to the extent formalists *don't* want strict textual interpretations of contracts ... *that's weird*. The biggest knock against formalism is that statutes and constitutions are hard to amend, such that absurd literalist interpretations would be trapped in the text for all time and wreak perpetual havoc on the polity. Indeed, basically all arguments that interpretation should be flexible and account for change over time are rooted in skepticism of legislative capacity. Contracts mitigate all those concerns—again, they typically bind *way* fewer parties and are thus *way* easier to change. If anything, then, contract interpreters should be *more* formalist than statute interpreters. The mitigating factor here, perhaps, is a lower view of contracting parties than of legislators; courts feel the need to coddle private drafters, who don't legislate for a living and thus are, in theory, more prone to inartful drafting. This is perhaps true of individual parties, but almost certainly *not* true of sophisticated corporate behemoths. Moreover, vagary is often the product of compromise; since contract drafters serve fewer constituents, they have greater capacity to enumerate nuanced specifics. We should expect contracts, then, to be at minimum *clearer* than statutes. To the extent contracts are less wise than legislation—a towering assumption—they compensate in that the parties bound possess far greater influence over the drafting process. Again, *caveat emptor* is the key fissure between contracts and constitutions.

Another normative rebuttal from the Kieffaber paper that looms large in the market context: *Healthy fear*.²¹¹ Indeed, the *ex ante* query function gives parties unprecedented power to master the law that binds them; they can query the machine for hours in search of that *one edgy play* that somehow slips through the contract's cracks and makes the machine flash green for "legal." Governments worth *trillions* don't want you to have that power, so why would a company worth *millions*? Well, because *it gives the company that power too*—and, since consumers are busy living their own lives, the company is more likely to use it. See, *ex ante* querying works the law pure *if the querier wants purity*; if the querier has different motivations well ... querying can also work the law *profitable* by the same token.

Of course, that tees up one final normative rebuttal that dogs arbitration in general: *Big companies don't bargain with consumers*.²¹² This retort cuts against the "contract as small nation" theory, noting that a random Amazon user probably has a better shot at amending the US Constitution than making Jeff Bezos send a refund for the towels that never came.²¹³ So constructed, Arbitrus.ai might be fine for mom-

²¹¹ Kieffaber, *supra* note 8, at 58 (“[M]aybe some unpredictability is healthy; that is, maybe we don’t want people to know exactly where the boundary between legal and illegal lies, for fear that bad actors would operate precisely along that line.”).

²¹² This objection is not relevant to commercial contracts between sophisticated parties. See RESTATEMENT (SECOND) OF CONTRACTS § 208 (AM. L. INST. 1981) (citing AIDS Healthcare Found. v. Express Scripts, Inc., 658 F. Supp. 3d 693, 704 (E.D. Mo. 2023) (“...owner failed to plausibly plead unconscionability because, among other things, according to its own allegations, both parties were sophisticated business entities.”)).

²¹³ Indeed, the biggest companies have substantially more consumers than America has citizens. Under a contraction framework, they are conceptually a *bigger* country. See *Global Social Media Statistics*, DATAREPORTAL, <https://datareportal.com/social-media-users#:~:text=Facebook%20has%203.07%20billion%20monthly,see%20more%20Pinterest%20stats%20here> (last visited Jan. 17, 2025).

and-pop contracts—but it’s liable to make policyholders at insurance conglomerates go postal. Practical monopolies exist, and they erode bargaining power for small consumers whose only other option is to forfeit the relevant service entirely.

We could tout the market as a panacea, reassuring the consumer who can’t manifest a culture of amendment at Youtube that she can always take her talents to OnlyFans. But we’d be facing the entirety of critical-theory scholarship and we’d be *wrong*—the companies that stand to gain the most from Arbitrus.ai, absent miracles or Musks, aren’t going to change their contracts when some consumers get mad about losing. So, what’s the answer? Well ... *there really isn’t one*. Arbitration isn’t sunshine and rainbows; it’s largely designed so that *you lose*. With Arbitrus.ai, you just lose *faster* and more *transparently*—and, if there’s an *ex ante* query, you don’t lose at all but merely accept the fact that you *would*. As the Judge.ai hypothetical seeks to demonstrate, morality is not the formalist’s business—and it’s *certainly* not the free-market libertarian’s, grounding his concept of the good in a literal societal median. If you want to afford parties broad contracting latitude in the first place, welcome to the end of your slippery slope. As with all libertine slopes, it ends in a machine.

V. Conclusion: Towards the *Arbitration State*

So, Judge.ai exists; it’s called *Arbitrus.ai*. We’ve shown you *how* it works, and we’ve shown you *that* it works in concept.²¹⁴ But we want to close with something broader: What will automated arbitration look like *at scale*? Our answer is the *Arbitration State*, a mass exodus of litigation into the private sphere with the potential to radically diminish the relevance

²¹⁴ Of course, future studies will need to verify that particular areas of the law are sufficiently covered, and will also need to examine more complicated, procedural scenarios.

of traditional legislatures.²¹⁵ That exodus will happen across two realms: “Free choice” areas, where the public is at liberty to paper over public law with contracts, and “forced choice” areas, where public law is supposed to be the only game in town.

Free choice areas are easy—that’s just a market, and automation will beat all comers. Indeed, we’ve touted how much cheaper Arbitrus.ai is than *standard arbitration*; compared to the cost of *litigation*, we’re talking orders of magnitude. And that’s to say nothing of *time*. District litigation alone takes years, sometimes *decades*; factor in appeals, and we start genuinely worrying whether the parties will still be *alive* to read the final decision. The only practical beneficiary of such a system is, ultimately, the law itself as it’s purportedly “worked pure;” for the rest of us litigants, the chief results are stomach acid and an access-to-justice crisis.²¹⁶

For litigants to meaningfully benefit from litigation, they need their answer *now*—or at least *soon*. But even arbitration—which is moving at light speed compared to the

²¹⁵ Stephen J. Ware, *Is Adjudication a Public Good: Overcrowded Courts and the Private Sector Alternative of Arbitration*, 14 CARDOZO J. CONFLICT RESOL. 899, 904 (2013) (“In short, the perpetual ‘crisis’ of ‘underfunded’ courts, crowded dockets and justice delayed is caused by the public subsidy for litigants. Focus on this subsidy for parties in litigation enables a contrast with the absence of a subsidy for parties in the private sector alternative to litigation, arbitration.”).

²¹⁶ Indeed, the number of people who truly have both the time and the money on their hands to retain competent counsel throughout an entire legal proceeding is subject to debate, but it is probably a small fraction of the total population. This is part of the justification for the sophisticated party doctrine in unconscionability—that lawyer-speak is indiscernible to the common man. See RESTATEMENT (SECOND) OF CONTRACTS § 208 (AM. L. INST. 1981) (“Factors which may contribute to a finding of unconscionability [include] ... knowledge of the stronger party that the weaker party is unable reasonably to protect his interests by reason of physical or mental infirmities, ignorance, illiteracy or inability to understand the language of the agreement, or similar factors.”).

old-world judiciaries—is not *now*.²¹⁷ Arbitrators have clogged dockets; parties have to find time to meet with them; the arbitrator must deliberate. And the time and money difference, in many cases, is not substantial enough to outweigh the higher rate of unpredictability for most litigants. Arbitrus, on the other hand, *is now*—and, when *now* is *cheap*, it’s going to outcompete the alternative. The legal system’s only hope would be to stack the game in its favor with either muscular unconscionability rules or frameshift amendments to the Federal Arbitration Act. But the latter depends on legislative constituencies that might just as soon keep the time and money Arbitrus.ai is saving them. The former, to be sure, will be a fistfight—but one which will grow more and more lopsided as the machine becomes more and more perfect.

Forced choice areas—where the government has decided to be supreme—are tougher. Domains like constitutional law and criminal law bind conduct, generally speaking, whenever an individual carries out such conduct in a government-controlled space. And, historically, government-controlled space has accounted for *basically all space*; if your feet are touching ground, odds are there’s a government nearby that can *make you do stuff*. That doesn’t leave much room for arbitration. But the Internet Age has led to the gradual consolidation of human action into a digital space owned not by governments but by *companies*—and, insofar as public law continues to religiously respect the corporate boundary, actions occurring in this space will be the Arbitration State’s bread and butter.

Make no mistake: The primacy of public law is already fading, as tech giants have created an *alternative physics* that we’re increasingly forced to live through. Speech is illustrative: Ask yourself, *what is the most important speech law in America today?* Since the Founding, the incontestable

²¹⁷ See *AT&T Mobility LLC v. Concepcion*, 563 U.S. 333, 348 (2011) (“According to the American Arbitration Association (AAA), the average consumer arbitration between January and August 2007 resulted in a disposition on the merits in six months, four months if the arbitration was conducted by documents only.”).

answer has been the First Amendment. But, when it decided *NetChoice*, the Supreme Court *de facto* ceded *all* social media speech—and the underlying power to decide *which* of that speech gets seen by *whom*—to Zuckerberg, Musk and the boys.²¹⁸ The First Amendment held onto *sidewalks* and *parks*;²¹⁹ private corporations got *the Internet*.²²⁰ Purely from a power standpoint, *NetChoice* was the most backward-looking decision that the Court had issued in years, as the Web 3.0 internet²²¹ promises to direct a disproportionate amount of

²¹⁸ *Moody v. NetChoice, LLC*, 144 S. Ct. 2383, 2407 (2024) (“But a State may not interfere with private actors’ speech to advance its own vision of ideological balance. States (and their citizens) are of course right to want an expressive realm in which the public has access to a wide range of views. That is, indeed, a fundamental aim of the First Amendment. But the way the First Amendment achieves that goal is by preventing *the government* from ‘tilt[ing] public debate in a preferred direction.’ ... It is not by licensing the government to stop *private actors* from speaking as they wish and preferring some views over others. And that is so even when those actors possess ‘enviable vehicle[s]’ for expression.”).

²¹⁹ *See Pleasant Grove City v. Summum*, 555 U.S. 460, 469 (2009) (“With the concept of the traditional public forum as a starting point, this Court has recognized that members of the public have free speech rights on other types of government property and in certain other government programs that share essential attributes of a traditional public forum.”).

²²⁰ The Supreme Court all but said it verbatim. *Moody*, 144 S. Ct. at 2402 (“And that is as true when the content comes from third parties as when it does not. (Again, think of a newspaper opinion page or, if you prefer, a parade.) Deciding on the third-party speech that will be included in or excluded from a compilation—and then organizing and presenting the included items—is expressive activity of its own. And that activity results in a distinctive expressive product. When the government interferes with such editorial choices—say, by ordering the excluded to be included—it alters the content of the compilation. (It creates a different opinion page or parade, bearing a different message.) And in so doing—in overriding a private party’s expressive choices....”).

²²¹ For a primer, see Henrique Centieiro, *The Insane Future of Web 3.0 and the Metaverse*, MEDIUM: DATADRIVENINVESTOR (Jan. 23,

the salient, litigable activity in peoples' lives through private online platforms. Today, the most impactful speech law in America—on its elections, on its civil discourse, on its thought-economy—is undoubtedly the one that Mr. Musk makes up on the spot.²²²

What if I don't like Musklaw? Or I'm upset that my healthcare claim was denied coverage? Or that my bank decided to close all of my accounts? Well ... *take it up with the company robot*. That dynamic at scale—placing litigable human actions beyond government reach—is the *Arbitration State*. And as more and more of our daily activities start to manifest through corporate-owned platforms in the digital space, the greater that scale will be. That's the metaverse concept: A wholly corporatized alternate reality where the man of the future will be forced, by pressures of speed and efficiency, to work.²²³ In that space, law will be what the company says it is; market competition will be its only referendum. And when efficiency is the virtue, we predict the company's law will be decreed by its AI. That's the Arbitrus endgame.

Bleak? No, *formalist*—a provision of instant, accurate process fully supported by the Federal Arbitration Act. And,

2022), <https://medium.datadriveninvestor.com/the-insane-future-of-web-3-0-and-the-metaverse>.

²²² For the local 'who-decides' question, the answer here is clearly the private market—even with more creative public schemes. See Adrian Vermeule, *Who Decides?*, POSTLIBERAL ORDER (Jan. 11, 2022), <https://www.postliberalorder.com/p/who-decides>; see also *NRA of Am. v. Vullo*, 144 S. Ct. 1316, 1328 (2024) (“To state a claim that the government violated the First Amendment through coercion of a third party, a plaintiff must plausibly allege conduct that, viewed in context, could be reasonably understood to convey a threat of adverse government action in order to punish or suppress the plaintiff’s speech. ... Accepting the well-pleaded factual allegations in the complaint as true, the NRA plausibly alleged that Vullo violated the First Amendment by coercing DFS-regulated entities into disassociating with the NRA in order to punish or suppress the NRA’s gun-promotion advocacy.”).

²²³ Just as he's been forced to own first a personal computer, then a smartphone, then (inevitably) a ChatGPT subscription.

crucially, the Arbitration State may prove much more conducive to predictability than even a nationalized Judge.ai regime. Indeed, even a perfectly functioning, oft-amended democracy will bind the minority with the will of the majority—which will always, at some level, confound the former. But the Arbitration State lets you pick and choose the legal systems that bind you. At minimum, that customization manifests in choosing one platform over another; at maximum, Arbitrus.ai develops to a scale at which it could, perhaps in tandem with consumer-side AI, promulgate nuanced contractual arrangements that differ for each user. True metaverse customization seeks to make the individual a nation of one, bound only by the restraints he chooses. When the law that binds you is a perfect mirror of your personal desires, *that's perfect predictability*. So construed, the Arbitration State is not only the formalist's ideal but also the libertarian's—insofar as those are two separable individuals.

What sphere does the Arbitration State leave for the federal government in the individual's daily life? On our current trajectory of corporate sovereignty, *little*; the state increasingly becomes the groundskeeper tasked with keeping the lights on in the old world so that its physical inhabitants can frolic in the new. And it might, in the end, prove small enough that *Arbitrus.ai* might offer it a lifeline—called *Judge.ai*—to provide the kind of metaverse efficiency that might someday make the real world competitive again.

That's pure futurism. And we doubt it sounds utopian; indeed, it puts faith in markets to provide a real choice that a great many will scoff at. But that's a question of antitrust capabilities and the epistemic merits of the libertarian outlook. We've never claimed that any of this is rooted in morals.²²⁴ Rather, it's rooted in the modern enlightenments of

²²⁴ See Kieffaber, *supra* note 8, at 57 (“I’ll finish with arguments that Judge.AI is *morally undesirable*. But I want to remind you from the outset that I’ve never claimed the contrary—I’ve never argued that Judge.AI is *good*.”).

formalism and free markets where efficiency is God and the machine is its vicar on Earth.²²⁵

You can mull over the philosophy on your own. In the meantime, Arbitrus.ai is here. And things are going to be different.²²⁶

²²⁵ This question contains both a ‘local’ and ‘global’ question of ‘who decides?’ The local-question is simply “given that someone must have the power to determine” a question, “to which institution is that power to be allocated?” In this case, it is first the contract itself; but if the contract is ambiguous, that power ultimately lies in the dispute resolution process—the ‘court,’ or Arbitrus.ai. The global question is broader: “must anyone decide?” In a world governed increasingly by contract, where *de facto* power lies in competing private institutions, we believe that the liberalized global question is in decline. See Vermeule, *supra* note 220.

²²⁶ See John Kruzel, *US Supreme Court’s Roberts Urges ‘Caution’ as AI Reshapes Legal Field*, REUTERS (Jan. 2, 2024), <https://www.reuters.com/legal/us-supreme-courts-roberts-urges-caution-ai-reshapes-legal-field-2023-12-31/> (“I predict that judicial work - particularly at the trial level - will be significantly affected by AI.”) (quoting Roberts, C.J.); Asimov, *supra* note 137, at 224 (“We don’t know. Only the Machines know, and they are going there and taking us with them.”).