

# AI AND PROBABILISTIC DISPUTE RESOLUTION

ZACHARY HENDERSON\*

In the United States, thousands of otherwise viable claims routinely drown in a sea of transaction costs. If you're a customer wronged by a company, chances are you have little practical legal recourse. Hire a lawyer? Good luck—an hour of attorney time might exceed the value of your claim. Small-claims court? Maybe—but it's rarely worth the time, effort, stress, and risk. Class actions aren't under your control, and anyway, that seven-dollar check is barely worth the trip to the post office. Then again, none of that is even on the table: Remember that forced arbitration clause on page thirty-nine of that clickwrap agreement?

Ballooning transaction costs create problems for companies contracting with one another as well, where the combination of uncertainty, time, litigation or arbitration fees, and asymmetrical power dynamics routinely leads to bad outcomes and suboptimal dispute resolution.

But new technologies bring new solutions, and major developments in artificial intelligence, deep learning, and computational statistics offer a powerful and elegant way to drastically reduce the transaction costs associated with resolving a wide array of claims.

This Article draws from advancements and fundamental principles in the fields of computer science, artificial intelligence, economics, and medicine to propose and predict a novel, contracts-based private mechanism enabling parties to voluntarily resolve a wide array of disputes. Parties would use a flexible AI platform with a known margin of error that they would accept at the outset, in exchange for a combination of reduced transaction costs, the chance to resolve a dispute that might otherwise be irremediable, and finality.

This Article further and relatedly argues that, if successfully implemented, such a platform could precipitate a rapid shift in how a range of disputes are resolved and could change the use cases for class actions and mass arbitrations for the better.

---

\* Visiting Assistant Professor and Shashi and Dipanjan “DJ” Deb Faculty Scholar, University of California College of the Law, San Francisco; Senior Research Scholar, Center for Innovation (C4i) and AI Law & Innovation Institute (AILII). For invaluable conversations and commentary, I thank everyone at C4i and AILII; the attendees of the 17th Annual AALS Alternative Dispute Resolution Works-in-Progress Conference; and Michael Mizrahi, Richard Deulofeut-Manzur, and Ian Schumann.

Introduction .....	217
I. The Problem: Transaction Costs and Illiquid Claims .....	222
A. The Small-Claims Wasteland: High Transaction Costs, Limited Options .....	224
1. Legal Representation for Small Claims Is Impracticable .....	225
2. Small-Claims Court Is a Mirage .....	225
3. Forced Arbitration Is Even Worse than Small- Claims Court .....	227
4. Neither Class Actions nor Mass Arbitrations Serve Claimants .....	228
5. The Through-Line Problem Is Claim Illiquidity .....	230
B. The Big Price Tag of Little Breaches .....	230
1. The Appeal of Arbitration in Business-to-Business Contracting .....	232
2. Interest in Arbitration Suggests Transaction Costs Are High Priority—and Big Business .....	233
II. Our Push-Button Future: Towards Probabilistic Dispute Resolution .....	233
A. A Thought Experiment: Pushing the Probabilistic Magic Button .....	234
B. Revisiting the Small-Claims Wasteland .....	236
C. Why Lawyers Should Embrace the Probabilistic Approach .....	236
1. The Preponderance of the Evidence Standard and Other Proofs that the Law Is Already Rooted in Statistics and Probability .....	237
2. Medicine Routinely Relies on Probability in Life-or- Death Decision-Making .....	240
3. Voluntary Dispute Resolution Is the Perfect Place To Experiment with Creating a Probabilistic Decision-Making Market .....	242
III. Using Neural Networks and Deep Learning To Create a Platform for Probabilistic Dispute Resolution .....	243
A. A Primer on Artificial Intelligence, Neural Networks, and Deep Learning .....	244
1. Machine Learning, Deep Learning, and Neural Networks .....	244
2. Training and Validating an AI Model for Dispute Resolution .....	247
B. Other Characteristics of a Probabilistic AI Dispute- Resolution Platform .....	249
1. Lack of Discovery, Honesty, and Other Evidentiary Matters .....	249
2. Beating Bias: The Black Box As a Feature Rather	

than a Bug ..... 251

3. Defining “Correct” Outcomes: Challenges and Opportunities..... 252

4. The Pace of Innovations and What’s Next ..... 253

IV. Who Will Build It, Who Will Come?..... 254

    A. Who Will Build It? ..... 254

    B. Who Will Use It? ..... 254

        1. Companies Will Want To Use the Platform with Their Customers ..... 255

        2. Some Companies Will Agree To Use the Platform To Resolve Discrete Business-to-Business Disputes... 258

    C. Why AI-Based Probabilistic Dispute Resolution Will Change Mass Dispute Resolution ..... 258

Conclusion..... 259

INTRODUCTION

*You can think of transaction costs broadly as pain-in-the-ass costs. All platforms reduce these costs in some way.*

—Alex Moazed & Nicholas L. Johnson,  
authors of *Modern Monopolies*<sup>1</sup>

Transaction costs—loosely defined as the costs of exchange when transacting in a given market<sup>2</sup>—are everywhere. You’ll pay them with your wallet a dozen times this week (a shipping fee here, a tip to a barista there), and you’re constantly paying them with your time—for example, the drive to the cafe or post office, or even right now, as you trade thirty minutes of your day in the hope that this Article proves worth reading. Of course, nobody wants to pay transaction costs, and in virtually every market, there are examples of successful companies whose core business models boil down to reducing or eliminating them.<sup>3</sup>

---

1. ALEX MOAZED & NICHOLAS L. JOHNSON, MODERN MONOPOLIES: WHAT IT TAKES TO DOMINATE THE 21ST-CENTURY ECONOMY 37 (2016).

2. Transaction costs have been defined in many different ways. For example, Kenneth Arrow defines them as “the costs of running the economic system”; Yoram Barzel defines them as “the costs associated with the transfer, capture, and protection of rights”; and Thrainn Eggertsson suggests: “In general terms, transaction costs are the costs that arise when individuals exchange ownership rights to economic assets and enforce their exclusive rights. A clear-cut definition of transaction costs does not exist, but neither are the costs of production in the neoclassical model well defined.” See Alexandra Benham & Lee Benham, *The Costs of Exchange* 2 (Ronald Coase Inst., Working Paper No. 1, 2001), <http://www.coase.org/workingpapers/wp-1.pdf> [<https://perma.cc/3V4Z-FWDR>].

3. For example, Visa, PayPal, and other payment-facilitation companies reduce payment transaction costs; Uber, Lyft, and Waymo reduce transportation transaction costs; and eBay and Amazon reduce exchange-of-goods transaction costs.

But some kinds of transaction costs tend to remain high or even increase over time, and can be difficult or impossible to reduce, let alone eliminate. For example, to enjoy a live symphony orchestra performance, there's no shortcutting it: Anywhere from forty to a hundred musicians *must* collectively expend hundreds of hours to produce it.<sup>4</sup> By contrast, other kinds of transaction costs see drastic decreases when new technologies emerge. Take public transportation: In the early 1800s, getting from New York to Paris could take up to a month of your time,<sup>5</sup> plus substantial room and board on a ship—not to mention a not-insignificant risk of death at sea.<sup>6</sup> But today, thanks to advanced transportation technology, the same trip requires little more than a few hundred dollars, a mildly harried jaunt through an airport, and a long nap in a comfortable armchair.<sup>7</sup>

But few sectors of the economy have transaction costs as intractable as the law—*especially* when it comes to litigation and dispute resolution. Nor have things gotten better over time: The costs associated with resolving legal disputes increase every year<sup>8</sup> and vastly outstrip inflation and wage growth.<sup>9</sup> This shouldn't be surprising: Most of the technological developments of the past hundred years that touch on the practice of law have served to increase, not decrease, the logistical complexity of resolving legal disputes. Computers and word processing are obvious examples: Any efficiencies these technologies created for the act of lawyering have been dwarfed by the sheer volume of relevant materials and evidence they generate—all of which add to the complexity and cost of a given dispute and the amount of time it takes to resolve it. In fact, so far, it has been almost axiomatic that any technological

4. Notably, the countless hours of practice in the years-long lead up to a given performance are *themselves* transaction costs.

5. *How Long Did It Take To Sail Across the Atlantic?*, ROYAL MUSEUMS GREENWICH (Nov. 8, 2018), <https://www.rmg.co.uk/stories/blog/library-archive/18th-century-sailing-times-between-english-channel-coast-america-how> [<https://perma.cc/6TBD-G7MF>] (“[T]ypical passage times from New York to the English Channel for a well-found sailing vessel of about 2000 tons was around 25 to 30 days, with ships logging 100-150 miles per day on average.”).

6. Risk of failure, too, factors into transaction costs—more on this later.

7. Just before writing this sentence, I searched for “cheap New York to Paris flight” in Google and found one-way tickets for as little as \$139. Given that price, “comfortable” might be overstating things a bit.

8. For example, from 2016 to 2020, tort costs grew at an average rate of six percent per year. *See* U.S. CHAMBER OF COM. INST. FOR LEGAL REFORM, TORT COSTS IN AMERICA: AN EMPIRICAL ANALYSIS OF COSTS AND COMPENSATION OF THE U.S. TORT SYSTEM 3 (2022), <https://instituteforlegalreform.com/wp-content/uploads/2022/11/Tort-Costs-in-America-An-Empirical-Assessment-of-Costs-and-Compensation-of-the-U.S.-Tort-System.pdf> [<https://perma.cc/RLE4-98J3>].

9. According to the Economic Policy Institute, actual year-over-year growth for private employees is 3.8%. *See Nominal Wage Tracker*, ECON. POL'Y INST., <https://www.epi.org/nominal-wage-tracker/> [<https://perma.cc/5G6T-WTQ2>].

advance that has served to reduce *some* transaction costs to legal dispute resolution has increased *all of the other* transaction costs even more.<sup>10</sup>

But with a little nudge from economics and a leaf borrowed from the modern medicine playbook, artificial intelligence—or, more specifically, the artificial intelligence subfields of machine learning<sup>11</sup> and deep neural networks<sup>12</sup>—offers a present-day pathway to drastically decreasing transaction costs for a wide array of consumer-to-business and business-to-business disputes in a way that benefits consumers and businesses alike.

This Article has a lot of ground to cover. Part I describes some of the core transaction-costs related problems that arise in two categories of disputes: (1) small claims that an individual consumer might have against a business and (2) breach-of-contract related disputes between two businesses. With respect to the first category, this Part briefly discusses the challenges of lawyered representation for small claims against companies; talks through the mostly lawyerless small-claims court process; discusses forced arbitration; and describes the key inadequacies, from the consumer perspective, of class action lawsuits and mass arbitration. The goal here is not to suggest that any of these dispute-resolution methods are *per se* failures, but rather to highlight a category of claims that none of these processes adequately serve, for the simple reason that the transaction costs in even the most cost-efficient versions of these systems are too high to justify seeking dispute resolution through them. These claims stagnate in a kind of small-claims wasteland, wherein a core tenet of transaction-cost economics<sup>13</sup> governs: “Where transaction costs are very high, many kinds of transactions may not take place at all.”<sup>14</sup>

---

10. For example, e-discovery, for all its benefits, has resulted in cases taking much longer to resolve. See *Case Closure Rates Get Longer as E-Discovery Increases*, U.S. CTS. (Mar. 21, 2024), <https://www.uscourts.gov/news/2024/03/21/case-closure-rates-get-longer-e-discovery-increases> [https://perma.cc/GDK5-FDYW] (“Across the country, federal defender offices are struggling with increasingly complex litigation marked by escalating electronic discovery demands. . . . These factors have contributed to criminal cases taking a median of 10.4 months to close from filing to disposition in 2023, a 60 percent increase from just five years ago, when the median case closure time was 6.5 months in 2018 . . .”).

11. See Sara Brown, *Machine Learning, Explained*, MIT SLOAN (Apr. 21, 2021), <https://mitsloan.mit.edu/ideas-made-to-matter/machine-learning-explained> [https://perma.cc/AF5S-SVW7].

12. See IBM Data & AI Team, *AI vs. Machine Learning vs. Deep Learning vs. Neural Networks*, IBM (July 6, 2023), <https://www.ibm.com/think/topics/ai-vs-machine-learning-vs-deep-learning-vs-neural-networks> [https://perma.cc/KQV8-JX38].

13. Mikko Ketokivi & Joseph T. Mahoney, *Transaction Cost Economics As a Theory of Supply Chain Efficiency*, 29 PROD. & OPERATIONS MGMT. 1011, 1014 (2020) (“The name *transaction cost economics* derives from the fundamental insight that economic exchange has an *ex ante* preparation cost and an *ex post* execution cost.”).

14. Benham & Benham, *supra* note 2, at 3.

In the second category, this Part discusses the general process two companies undergo when contracting, with emphasis on the premise that, in most cases, neither party to an agreement anticipates breaching, and each is usually more concerned with the *other* party's potential for breach than the likelihood or consequences of breaching themselves. This Part goes on to discuss whether, why, and in what contexts companies tend to be amenable to agreements to arbitrate. (Primarily: "correct" outcomes; cost; time; and confidentiality.) And it finishes with an overview of the transaction costs associated with business-to-business contract dispute resolution in breach scenarios.

Part II introduces a novel thought experiment that lays the groundwork for the proposals and arguments that appear later on. Suppose any two parties had a magic button they could push that would (1) resolve their dispute; (2) instantly; (3) cost hundreds of times less than the next cheapest dispute-resolution mechanism; (4) bind both parties to that result; and (5) require the losing party to pay appropriate compensation—with just one catch: (6) *the magic button gets it right only 90% of the time, and gets it wrong 10% of the time*. The thought experiment begs the question: *Who would push the button and under what circumstances?* How about if the right/wrong probability were 99% to 1%? What if it were even better than that?

In response to this thought experiment, this Part concludes by making several related arguments. First, it proposes that many claimants, and especially (but not exclusively) those who live in the consumer-to-business small-claims wasteland described in Part I, would gladly push the button. Second, it argues that this kind of probabilistic and statistical resolution is not antithetical to, nor even especially different from, legal dispute resolution as currently conducted. In furtherance of this argument, this Part points to our current standards of evidence—including our civil standard of a "preponderance of the evidence"—as proof that we *already* have a probabilistic system, even if we don't often explicitly acknowledge it. And to counter the rejoinder that probabilistic and statistical determinations would be somehow unethical or socially unacceptable, this Part directs the reader's attention to modern medicine—a profession that is arguably more respected than the law, and yet openly and routinely relies on statistics and probability in making life or death decisions, in contexts as critical as diagnosis, recommending treatment options, and describing potential surgical outcomes. Finally, this Part argues that, for historical, economic, and freedom-of-contract reasons, *voluntary* commercial dispute resolution is the perfect context in which to experiment with probabilistic dispute resolution in the form this Part contemplates.

In the first half of Part III, this Article explains why AI systems based on deep learning are precisely the tool for the job of making the "magic button" described in Part II a reality, and why it is likely that the

technology needed to achieve a greater-than-90% “proper outcome” rate already exists today. Alas, this Part does get a bit technical—a necessity to have any chance of convincing the reader that modern AI systems are indeed capable of offering statistically reliable dispute resolution. But I do make an effort to keep the computer science and computational statistics at the 101 level, and where I can’t, I rely heavily on analogy.

The second half of this Part describes some of the unique characteristics of an AI platform like the one proposed and discusses some of the potential concerns with it. This includes questions about evidence (on what factual and legal basis would the platform make decisions?); decision transparency (how would we know that the AI system is fair?); bias (would this system *create* or *mitigate* bias?); the difficulty of defining “correct” outcomes (a problem relevant to *all* dispute resolution methods); and how we should think about even more advanced platforms that can reach the right result even more often (and perhaps far more often than human arbiters).

After all the talk about algorithms and neural networks is wrapped up, Part IV describes some of the ways that an AI platform like the one proposed could be practically deployed, as well as some of the levers at our disposal to shift the incentives around choosing to use the AI platform to resolve disputes. Many of these levers include changing the costs of use. For example, in the consumer-to-business context, a typical approach might be to charge a claimant a small fee to use the platform. In other cases, it might make sense for a company to incentivize using the platform by shrinking the fee, or possibly, even by making it free. And in extreme cases—for example, when a company has reason to believe it has become susceptible to a class action that it hopes to head off—it may make sense for a company to actually *pay* its past and current customers to use the dispute-resolution platform, as a way of shrinking the potential class action pool, and/or to give itself a stronger basis for defeating class certification under the Federal Rules of Civil Procedure.<sup>15</sup>

Finally, this Article concludes by emphasizing the flexibility of a platform like the one proposed, the importance of its voluntariness, and a reminder that while a platform like this one can (and in my view, should) be deployed as soon as possible, the technology is only going to improve, and the bull case for statistical and probabilistic rulings will only get stronger. The Article ends by reprising the future-looking version of the “magic button” question: How many people would push it if it reached the “right” result 99% of the time? How about 99.9%? And, as a closing teaser, at what point would the *platform* begin to define “right” results, and at what point would it be unconscionable *not* to deploy systems like these, not only in the voluntary, contractual context

---

15. See FED. R. CIV. P. 23(a).

like this Article proposes, but in the more formal judicial contexts that other articles contemplate, and that, indeed, some jurisdictions are already deploying?

#### I. THE PROBLEM: TRANSACTION COSTS AND ILLIQUID CLAIMS

*“[T]o the buyer all costs are transaction costs . . . .”*

—Michael Munger, Professor of Political Science,  
Economics, and Public Policy, Duke University<sup>16</sup>

Resolving legal disputes is time-consuming and expensive<sup>17</sup>—and any dispute can turn into a legal one so long as one party is upset, determined, or misguided enough.<sup>18</sup> Many disputes involve damages that are significant enough to justify the time, expense, and risk of legal action.<sup>19</sup> But as legal and arbitration costs grow, an increasingly large number of claims are becoming totally illiquid—which is to say, the transaction costs necessary to resolve them are higher than the value of even the best-case outcome,<sup>20</sup> such that even if a claimant were certain to win, it would be more economical to accept the harm and move on than

16. Michael Munger, *To Consumers, All Costs Are Transaction Costs*, DAILY ECON. (June 13, 2023), <https://thedailyeconomy.org/article/to-consumers-all-costs-are-transaction-costs/> [<https://perma.cc/RT5Q-Z82Z>].

17. According to Clio’s 2024 Legal Trends Report, the average hourly rate for lawyers in California is \$391 per hour. CLIO, LEGAL TRENDS REPORT 131 (2024) [hereinafter 2024 LEGAL TRENDS REPORT], <https://www.clio.com/wp-content/uploads/2024/10/NA-2024-Legal-Trends-Report-Full-Publication.pdf> [<https://perma.cc/L3FW-SF4X>]. As for what companies spend on litigation per lawsuit, for companies with revenue less than \$100 million, each litigation matter costs, on average, \$50,000. For companies with revenues in excess of \$1 billion, that cost balloons to \$200,000 per matter on average. See Robert Freedman, *Big Discrepancy in Average Litigation Costs Between Large, Small Companies*, LEGAL DIVE (Oct. 24, 2022), <https://www.legaldive.com/news/big-discrepancy-average-litigation-costs-between-large-small-companies-everlaw-acc-lawfirm-legalfees/634811/> [<https://perma.cc/KCU4-JVBM>].

18. For example, Richard Overton sued Anheuser-Busch for deceptive public advertisements that led him to mistakenly believe drinking their beer would result in fantasies coming to life, including encounters with beautiful women “engaged in unrestricted merriment.” *Overton v. Anheuser-Busch*, 517 N.W.2d 308, 308 (Mich. Ct. App. 1994).

19. *Product Liability Verdict Statistics*, MILLER & ZOIS ATT’YS L., <https://www.millerandzois.com/maryland-injury-victim-help-center/product-liability-statistics/> [<https://perma.cc/J7TL-9TGY>] (showing median ranges in product liability cases from roughly \$280,000 to \$1.5 million).

20. See, e.g., Mark Fotohabadi, *How Much Does Arbitration Cost?*, ADR TIMES (Oct. 2, 2023), <https://adrimes.com/how-much-does-arbitration-cost/> [<https://perma.cc/XA2G-GQTY>] (listing the potential costs of arbitration as filing fees, hearing fees, administrative and administration fees, hearing room rental, arbitrator fees, discovery costs, and attorneys’ fees).



to pursue the claim.<sup>21</sup> Would-be disputes like these fall into a kind of small-claims wasteland: Taking *any* step to vindicate the claim, however righteous, would serve only to worsen the claimant's economic situation.

Another class of claims includes those that are economically and *functionally* illiquid. This is because, even if the *best-case* outcome would yield a net benefit, the *expected value* of the claim—that is, the probability-weighted value of a claim, which takes into account all possible outcomes and their probabilities<sup>22</sup>—is so low in comparison to the transaction costs that entering into a dispute-resolution process would result in an expected economic loss.<sup>23</sup>

This Part discusses two specific claim scenarios, each of which falls into one of the two buckets described above: Totally illiquid claims a human customer might have against a business, and economically illiquid claims that a business might have against another business (for example, as a result of a contractual breach).

---

21. See, e.g., J. Maria Glover, *Mass Arbitration*, 74 STAN. L. REV. 1283, 1329 (2022) (pointing out that in the context of arbitration, “just the filing fee for the arbitration demand can exceed the value of any individual claim”).

22. See CFI Team, *Expected Value*, CORP. FIN. INST., <https://corporatefinanceinstitute.com/resources/data-science/expected-value/> [https://perma.cc/XP6G-G64W] (“Expected value (also known as EV, expectation, average, or mean value) is a long-run average value of random variables. It also indicates the probability-weighted average of all possible values.”).

23. That is, the probability of the loss, multiplied by the probability of the event occurring, would be less than the expected value of the claim. See, e.g., *Understanding Our Methodology*, LLOYD'S, <https://www.lloyds.com/news-and-insights/futureset/futureset-insights/systemic-risk/research-methodology> [https://perma.cc/NPV6-EAF2].

*A. The Small-Claims Wasteland: High Transaction Costs,  
Limited Options*

Let's set the stage. Suppose you're a previous customer of ServiceCo, a company that provides a subscription-based online service, at the cost of \$200 charged annually. You haven't used the service for seven months and meant to cancel your subscription but never got around to it. Alas, you just discovered that ServiceCo renewed your subscription three months ago and raised the price by \$10 (meaning you're now \$210 in the hole for a service you haven't used and don't want). You reach out to ServiceCo's customer support to try to get them to refund you, but all you get is an automated response telling you that you aren't eligible for a refund, because you waited too long to cancel. (They don't respond to your follow-up emails.) If you had used your credit card, you might have been able to get the charges reversed (though no guarantees there—ServiceCo would get an opportunity to disagree with you). But you paid via an automatic deduction from your checking account, and unlike your credit card, it has no charge-reversal safety net. What to do?

Situations like these are far from far-fetched.<sup>24</sup> In fact, for many subscription-based companies, a huge proportion of their revenue comes from exploiting customer inattention, whether intentionally or unintentionally.<sup>25</sup> And in our plausible hypothetical case, there is a genuine dispute that turns on a combination of the specific facts, whatever agreements exist between you and the company, and any applicable consumer protection laws and regulations. Perhaps the company *is* legally entitled to refuse you a refund; perhaps not. But whatever the proper outcome might be in principle, in practice you'll never reach it, because for each of the mechanisms for dispute resolution available to you, the transaction costs outstrip the value of your best-case outcome. Welcome to the small-claims wasteland.

---

24. One interviewee described an almost-identical experience with a company's automatic subscription:

Since its customer service office is only open on weekdays, I had to wait until Monday morning to call. The person I reached told me that I had renewed for one year. Any cancellation would only be effective 12 months from now, when that new term would end. He said that I had to notify them before the renewal date and stated that since I had not done so, the \$249 would not be refunded.

See Abe Wischnia, *Are Automatic Subscription Renewals a Scam? Here's What Consumers Need To Know*, CONSUMER RESCUE (Dec. 23, 2022), <https://consumerrescue.org/news-and-alerts/how-protect-against-automatic-subscription-renewals/> [https://perma.cc/J5E5-LN8Y].

25. Liran Einav, Ben Klopck & Neale Mahoney, *Selling Subscriptions 2* (Stanford Inst. for Econ. Pol'y Rsch., Working Paper No. 23-23, 2023) (“[I]nattention raises seller revenues by between 14% and more than 200%.”).

### 1. Legal Representation for Small Claims Is Impracticable

Let's get the obvious one out of the way first: Hiring a lawyer is probably not even an option. Many of the plaintiffs' attorneys that represent consumers operate on contingency, meaning they front the costs of litigation and take a cut if you settle or win.<sup>26</sup> But no lawyer working on contingency would take your case: Even at the extreme high end of acceptable contingency fee percentages—say, 50%<sup>27</sup>—a lawyer would stand to make only \$105 out of your \$210 claim.

Nor would it make any sense to pay a lawyer's hourly rate to assist you. Even if you were to find an attorney that charged one-third the average hourly rate of attorneys in your state,<sup>28</sup> within an hour and a half you would have exceeded the total value of your claim. (And that's before any court and administrative costs.)

So hiring a lawyer is a no-go—but that isn't news to anyone. After all, small-claims courts have been around for decades, and ostensibly exist to address exactly this problem.<sup>29</sup> Let's turn to that option next.

### 2. Small-Claims Court Is a Mirage

Small-claims courts<sup>30</sup> exist in every state in the country<sup>31</sup> and have at least the *intention* to be effective avenues for resolving a wide range of low-dollar disputes.<sup>32</sup> But for very small claims (such as the small consumer claim in our hypothetical example), the value that small-claims courts offer is often illusory at best. On the one hand and on a positive note, the filing fees are often much lower than the face value even of very small claims, depending on the state and municipality. For example,

---

26. See *Contingency Fees*, CONSUMER ATT'YS CAL., <https://www.caoc.org/index.cfm?pg=contingencyfees> [<https://perma.cc/2S4J-QBXW>] (“In a typical contingency fee agreement, the plaintiff is only responsible for paying their attorney if they win the case, with the payment coming as a percentage of the winnings.”).

27. Lester Brickman, *Contingent Fees Without Contingencies: Hamlet Without the Prince of Denmark?*, 37 UCLA L. REV. 29, 30 (1989) (“According to conventional wisdom virtually all contingent fee percentages exceeding fifty percent are illegal and excessive, but most lower percentages are valid.”).

28. See 2024 LEGAL TRENDS REPORT, *supra* note 17, at 131.

29. Alan M. Roodhouse, *Small Claims Court—What Should It Provide and How Well Does It Do So*, 51 CAL. ST. BAR J. 126, 127 (1976) (“The role that a small claims court is supposed to play . . . is clear: a quick, efficient and low-cost means whereby rich and poor, entity and individual, can settle minor disputes under the aegis of justice.”).

30. For a succinct history of small-claims courts, see Jill I. Gross, *AT&T Mobility and the Future of Small Claims Arbitration*, 42 SW. L. REV. 47, 58–59 (2012).

31. Bruce Zucker & Monica Her, *The People's Court Examined: A Legal and Empirical Analysis of the Small Claims Court System*, 37 U.S.F. L. REV. 315, 317 (2003).

32. See Roodhouse, *supra* note 29, at 128.

in a state like California with especially affordable small-claims systems, the filing fee could be as low as \$30.<sup>33</sup> But on the other hand, the nonmonetary transaction costs are much harder to reduce or eliminate. Even in the most streamlined systems, a claimant will still need to: spend time learning about the small-claims process;<sup>34</sup> try to learn the law, and grapple with any governing contracts or agreements; create a plan for what to say at a hearing, and practice it; and then attend a hearing, which may require transportation costs, taking off work, and/or childcare costs. On top of all this, there are other potential costs that the litigant may be subject to—for example, fee shifting, countersuits, and the possibility of the other party seeking to appeal.<sup>35</sup> And finally, there are the emotional and mental costs of the process: For many people, even small-claims court can be an incredibly stressful experience.<sup>36</sup>

These costs are impossible to precisely quantify, but it suffices to say that the actual total cost of resolving many very small claims is significantly higher than the filing fee. Unless one has a strong emotional basis for pursuing a \$200–\$300 claim, the all-in costs of resolving it in small-claims court probably exceed the claim’s actual value by a considerable margin. (And that’s *without* factoring in the claimant’s likelihood of success, which will inevitably be less than 100%—though it is generally true that most claimants who file with small-claims court *do* end up winning.<sup>37</sup>) In short, small-claims court may be great for “larger” small claims, but they don’t help us get out of the small-claims wasteland.

---

33. DOREATHEA JOHNSON, ANITA SCURI & ALBERT Y. BALINGIT, CAL. DEP’T CONSUMER AFFS., *THE SMALL CLAIMS COURT: A GUIDE TO ITS PRACTICAL USE*, [https://www.dca.ca.gov/publications/small\\_claims/small\\_claims.pdf](https://www.dca.ca.gov/publications/small_claims/small_claims.pdf) [<https://perma.cc/T5K5-TG9D>] (stating that, in California, “[t]he fee for filing in small claims court depends on the amount of the claim: \$30 if the claim is for \$1,500 or less, \$50 if the claim is for more than \$1,500 but less than or equal to \$5,000, or \$75 if the claim is for more than \$5,000”).

34. California’s small-claims guidebook is both dense and over fifty pages long. *See id.*

35. William J. Woodward, Jr., *Legal Uncertainty and Aberrant Contracts: The Choice of Law Clause*, 89 CHI.-KENT L. REV. 197, 202 (2014) (noting that fee-shifting contract clauses are “widely found in contracts promulgated by the finance industry”).

36. Peter Bowal & Jacqueline Bowal, *Small Claims Court: A Venue Made for Self-Represented Litigants*, LAWNOW, July/Aug. 2015, at 29, 31 (“There are . . . non-economic costs of suing such as adverse publicity, stress and mental anguish, uncertainty, time to prepare and prosecute the case, and destruction of relationships. We think you should not sue unless it is for a relatively large amount of money and you have a good chance of getting paid if successful.”).

37. *See* Julie Bradley, Barbara Sherman & W. Keith Bryant, *Winning in Small Claims Courts: An Empirical Analysis*, 16 J. CONSUMER AFFS. 112, 124 (1982) (“[T]he probability is very high that small claims court cases initiated by consumers are won by consumers.”).

### 3. Forced Arbitration Is Even Worse than Small-Claims Court

From a claimant’s perspective, small-claims court is still a more affordable option than most of the alternatives—and it is *much* more affordable than the kind of arbitration many consumers are forced to agree to in their customer contracts.<sup>38</sup> To begin with, even under some of the most consumer-friendly arbitration rules, the consumer filing fees are high. For example, under the American Arbitration Association’s Consumer Arbitration Rules, the nonrefundable filing fee for a single consumer is set at \$225.<sup>39</sup> So the moment our exemplar claimant files, she’s underwater—never mind the nonmonetary transaction costs we discussed earlier,<sup>40</sup> most of which are still present in the arbitration context. And it gets much worse from there. First, the consumer win rates in consumer arbitration are *extremely* low:

Rank	Sector	Consumer Win Rate
1	Retail	0.6%
2	Transport	0.7%
3	Healthcare	1.4%
4	Restaurants	1.7%
5	Financial Services	1.8%
6	Cable	2.1%
7	Other	2.8%
8	Car Sales/Lease	7.6%
9	Home Construction/Renovation	16.4%
10	Tech	16.5%

Table 1. AAA Financial Services Companies Win Rate in Forced Arbitration, 2017–21<sup>41</sup>

38. Glover, *supra* note 21, at 1308, 1329 (“Many claims eliminated by forced arbitration and class-action waivers, then, are not so much low merit as they are low value. . . . In many cases, just the filing fee for the arbitration demand can exceed the value of any individual claim. Indeed, the initial filing fee is the reason that most individual consumer and employment demands, at least if unconnected to a mass arbitration, are never brought.”).

39. *Consumer Arbitration Rules: Costs of Arbitration*, AM. ARB. ASS’N (Jan. 15, 2024), <https://go.adr.org/consumerfeeschedule.html> [<https://perma.cc/XSY8-VSCF>].

40. *See supra* Section I.A.2.

41. *Forced Arbitration and Big Banks: When Consumers Pay To Be Ripped Off*, AM. ASS’N FOR JUST. (Sept. 2022) [hereinafter *Forced Arbitration*], <https://www.justice.org/resources/research/forced-arbitration-big-banks> [<https://perma.cc/LZK7-4W7Y>].

Second, the consequences of losing in forced arbitration are often significant, thanks to fee-shifting provisions.<sup>42</sup> These provisions are a major concern: In some industries such as banking, the amount of money consumers end up *paying* after losing in arbitration can be catastrophic. For example, according to one study, “[i]n more than 100 cases [filed between 2017 and 2021], Americans brought a forced arbitration case against a bank, only to be ordered to pay the bank . . . an average of approximately \$24,000 each.”<sup>43</sup> There isn’t much more to say here, other than to restate the obvious: Arbitration, forced or not, is not a way out of the small-claims wasteland.

#### 4. Neither Class Actions nor Mass Arbitrations Serve Claimants

Whatever one thinks about commercial class actions and mass arbitration, one thing is (mostly) indisputable: Even if these tools are effective at changing company behavior prospectively, they do a middling job, at best, of making individual claimants who have suffered actual harm completely whole.<sup>44</sup> This is true for several reasons.

First, both processes tend to be lawyer-driven and not consumer-driven.<sup>45</sup> This is because, by design, both vehicles involve a very large number of claimants, and so both class actions and mass arbitrations

---

42. See Christopher R. Leslie, *Conspiracy to Arbitrate*, 96 N.C. L. REV. 381, 397–98 (2017) (noting that arbitration potentially offers companies a path to circumvent pro-plaintiff fee-shifting laws).

43. *Forced Arbitration*, *supra* note 41.

44. See William P. Barnette, *There Is No Conservative Case for Class Actions*, 22 FEDERALIST SOC’Y REV. 192, 195 (2021) (“It is well known . . . that the claim rates in consumer class settlements are abysmal; typically, well less than 10% of class members even bother to make claims.”). See also Andrew Pincus, Archis Parasharami & Daniel Jones, *Mass Arbitration Abuse Boosts Attorneys’ Fees, Not Fair Results*, BLOOMBERG L. (Feb. 8, 2024, 3:30 AM), <https://news.bloomberglaw.com/us-law-week/mass-arbitration-abuse-boosts-attorneys-fees-not-fair-results> (“Some commentators are celebrating the rise of mass arbitrations, but there’s little cause for celebration. The sole beneficiaries of most mass arbitrations are the plaintiffs’ lawyers who reap huge fees from settlements. Often the settlement amounts are wholly unrelated to the merits of the underlying claims. Who’s hurt? Parties with legitimate claims and the consumers and workers who foot the bill for these unjustified settlements and the accompanying attorneys’ fees.”); Scott Medintz, *How Consumers Are Using Mass Arbitration To Fight Amazon, Intuit, and Other Corporate Giants*, CONSUMER REPS., <https://www.consumerreports.org/money/contracts-arbitration/consumers-using-mass-arbitration-to-fight-corporate-giants-a8232980827/> [<https://perma.cc/QBN4-JHAE>] (Aug. 13, 2021) (“The huge cost of launching a mass arbitration limits the number of lawyers able to take on these cases, and pushes those that can to focus on the cases with the largest potential payouts. Thus many claims with legal merit but minimal ‘marketability’ will likely never find a lawyer to champion them as mass arbitrations . . .”).

45. See Barnette, *supra* note 44, at 194; Pincus, Parasharami & Jones, *supra* note 44.

require significant coordination.<sup>46</sup> Moreover, class actions require that strict specific legal criteria are met.<sup>47</sup> While it is certainly possible that an individual claimant could reach out to a lawyer and raise an issue that the attorney decides has value as the basis for a class action or mass-arbitration process,<sup>48</sup> it's more likely that a claimant will become a part of a class action or mass arbitration only by happenstance: They *may* be notified of, or invited to join, a class action or mass arbitration that's already in the works.<sup>49</sup> So these aren't "tools in the consumer's toolkit," so to speak, in the same way that hiring a lawyer, or going to small-claims court, are. Instead, they tend to function more like low-value lottery tickets.

And "low value" is the second reason why these vehicles do not serve claimants. A 2013 empirical study found that, of 148 cases looked at, none went to trial, and none resulted in a judgment on the merits for the plaintiffs.<sup>50</sup> In nearly two-thirds of resolved cases, "members of the putative class received zero relief,"<sup>51</sup> and while data on actual distributions is rare (on account of the frequency of settlement<sup>52</sup>), the awards are often vanishingly small.<sup>53</sup> No surprise, then, that awards are rarely worth putative class members' time in claiming them. For example, in a 2005 case relating to compact disc marketing, individual

---

46. See Glover, *supra* note 21, at 1322 (describing what is involved in coordinating mass arbitration).

47. Rule 23(a) of the Federal Rules of Civil Procedure reads:

(a) Prerequisites. One or more members of a class may sue or be sued as representative parties on behalf of all members only if:

(1) the class is so numerous that joinder of all members is impracticable;  
(2) there are questions of law or fact common to the class;  
(3) the claims or defenses of the representative parties are typical of the claims or defenses of the class; and  
(4) the representative parties will fairly and adequately protect the interests of the class.

FED. R. CIV. P. 23(a).

48. See *How To Start a Class Action Lawsuit in 2024*, CLASSACTION.ORG, <https://www.classaction.org/learn/how-to-start?> [<https://perma.cc/V4TK-4AYC>].

49. In the case of mass arbitration, "many mass arbitrations need something like a class list to get started," see Glover, *supra* note 21, at 1333; as for class actions, see Kathryn Honecker, Julia Campins & Laura Van Buren, *Class Actions 101: A Primer on Finding Plaintiffs for Your Class Action . . . Ethically*, A.B.A. (Sept. 19, 2013), <https://www.americanbar.org/groups/litigation/resources/newsletters/class-actions-derivative-suits/class-actions-101-primer-finding-plaintiffs-your-class-action-ethically/>.

50. MAYER BROWN LLP, U.S. CHAMBER INST. FOR LEGAL REFORM, *DO CLASS ACTIONS BENEFIT CLASS MEMBERS? AN EMPIRICAL ANALYSIS OF CLASS ACTIONS 3* (2013) ("Of the 148 cases in our sample set, not one had gone to trial—either before a judge or jury." (emphasis omitted)).

51. *Id.* at 6.

52. See *id.* at 10.

53. See, e.g., *In re Compact Disc Minimum Advertised Price Antitrust Litig.*, 370 F. Supp. 2d 320, 321 (D. Me. 2005).

payouts were worth roughly \$4—and just 2% of the class bothered to make a claim.<sup>54</sup> As for mass arbitration, a common premise is typically to take advantage of companies' subsidization of arbitration fees in order to threaten the companies into settlement:<sup>55</sup> *Pay up, or you'll be on the hook for thousands of customers' worth of arbitration fees.*<sup>56</sup> But from the consumer perspective, mass arbitrations suffer from most of the same problems as class actions: Consumers have no real way to control or predict whether a mass arbitration will occur, and even in the “success” case—settlement—the resulting award will probably be a fraction of what an actual claim could have been worth.<sup>57</sup>

### 5. The Through-Line Problem Is Claim Illiquidity

For each of the pathways described above, the fundamental problem is the same: The transaction costs exceed the value of very small claims. Hired lawyers, small-claims court, and arbitration are all too expensive (in dollars, time, and energy), while class actions and mass arbitration are too uncertain and offer too little compensation. The result is that very small claims are functionally illiquid: Though they contain real value, there is no economical way to unlock that value. Every such claim—and there is little doubt there are many—is analogous to a box containing two crisp hundred-dollar bills—but the box is locked and keyless, and the cheapest locksmith charges three hundred dollars a lock. Whoever figures out how to open the box for cheaper will create massive value overnight, but until that day comes, the box might as well be a doorstop.

#### *B. The Big Price Tag of Little Breaches*

Companies contract with one another constantly. A contract might be as simple as a single-page sales order, or as complex as a merger involving hundreds of documents. And while sophisticated companies tend to negotiate agreements bearing in mind the possibility that the other party will breach, parties rarely believe they themselves are likely to do so.<sup>58</sup>

---

54. *See id.*

55. *See Glover, supra* note 21, at 1341 (“The mass-arbitration model exploits the tradeoffs made by AT&T and other corporate defendants: Plaintiffs’ firms essentially called the defendants’ bluff by filing demands under their ‘friendly’ agreements and insisting that courts ‘rigorously enforce [those] agreements according to their terms.’” (quoting *Am. Express Co. v. Italian Colors Rest.*, 570 U.S. 228, 233 (2013))).

56. *Id.*

57. *See Pincus, Parasharami & Jones, supra* note 44.

58. Jeff S. Johnson & Ravipreet S. Sohi, *Understanding and Resolving Major Contractual Breaches in Buyer-Seller Relationships: A Grounded Theory Approach*, 44



But even companies with the best of intentions occasionally breach their contracts—either unintentionally, or because it becomes economical or convenient to do so.<sup>59</sup> And breaches come in all sizes, from small ones like missing a minor notification deadline, to catastrophic breaches that cause serious damage to the nonbreaching party. When a breach occurs, both parties must decide how to handle it—an often-challenging decision with no set answer, because all relationships and agreements are different.<sup>60</sup> At root once again is the problem of transaction costs. Will raising the breach be more costly than ignoring it, either in terms of goodwill, future business,<sup>61</sup> or dispute-resolution costs?<sup>62</sup> How predictable is the outcome—how clear are the contract terms, or the underlying facts? Each one of these considerations (and many others), taken together, amount to difficult-to-calculate transaction costs.<sup>63</sup>

No doubt one of the reasons why decisions in cases of breach carry so much *relational* weight is because any action that seeks resolution of the breach increases the risk of further escalation and further costs—up to and including the risk of extremely costly full-scale litigation or

---

J. ACAD. MKTG. SCI. 185, 185 (2016) (“With few exceptions, buyers and sellers enter into contractual relationships intending to adhere to the prescribed terms.”).

59. Will Kenton, *Breach of Contract Explained: Types and Consequences*, INVESTOPEDIA, <https://www.investopedia.com/terms/b/breach-of-contract.asp> [https://perma.cc/FSF6-Y36B] (May 6, 2024) (“Economically, the costs and benefits of a contract’s terms determine whether either or both parties have an economic incentive to breach it. If the net expected cost to a party of breaching a contract is less than the expected cost of fulfilling it, then that party has an economic incentive to breach the contract. Conversely, if the cost of fulfilling the contract is less than the cost of breaking it, it makes sense to respect it.”).

60. See Johnson & Sohi, *supra* note 58, at 185 (“Resolving a contractual breach is not a binary decision between enforcing and not enforcing contract terms. Rather, it lies somewhere on a continuum between these two poles, with enforcement options often governed by relational considerations.”).

61. See *id.* (“While legal enforcement based on the contract terms is an option for dealing with the breach, it can be costly and often is not a viable alternative as it may create conflict and erode productive business associations.”).

62. See *id.* (“For example, in buyer–seller situations when a contract is breached by a buyer, enforcement of the contract terms may not be practical as it could jeopardize the seller’s relationship and future business dealings with the customer.”).

63. See *id.* Johnson and Sohi articulate the challenges here so well that their words deserve extensive quotation:

In discrete exchanges, the transactions are assumed to be independent of past and future relationships, and the parties rely on economic and legal sanctions to enforce contractual obligations. Conversely, in relational exchanges, also referred to as relational contracts, the transactions are embedded in historical and social relational contexts, and enforcement of obligations is based on norms of exchange which serve to guide and regulate behavior of the exchange parties.

*Id.* at 185, 187.

arbitration.<sup>64</sup> Indeed, it is precisely *because* the potential costs are so high that there is so much relational risk in raising a breach in the first place. This Article proposes that if the worst-case breach scenario were predictable and capped, the relational risk in raising a breach would decrease, too. This could have positive effects not only on how breaches get resolved, but on contract drafting as well: Arguably, the more the worst-case consequences of a breach are mitigated, the less negotiation reaching a breach resolution would require, and the less relationship risk dispute resolution would pose.

### 1. The Appeal of Arbitration in Business-to-Business Contracting

Reducing the transaction costs involved in dispute resolution—money, time, and uncertainty—is exactly why many companies agree to arbitrate their business-to-business disputes in the first place, rather than litigate them in court.<sup>65</sup> For many disputes, arbitration strikes a decent balance: As a mechanism for dispute resolution, it can be highly sophisticated, but still potentially faster and less costly than litigation.<sup>66</sup>

But even though it has transaction-costs advantages over the court system in general, it is *still* an expensive, lengthy process that can result in protracted disagreement and can damage the parties' relationships with one another.<sup>67</sup> It may be better than the courts, but perfect it is not. Companies spend millions of dollars each year on arbitration costs<sup>68</sup> in a process that, at the end of the day, still involves lots of lawyers, an

---

64. See Barney Jordaan, *Litigation Costs Money, Destroys Value and Increases Corporate Risk*, LINKEDIN (Dec. 5, 2014), <https://www.linkedin.com/pulse/20141205160259-48129923-litigation-costs-money-destroys-value-and-increases-corporate-risk/> (“Perhaps most important is the adverse effect that disputes and litigation unquestionably have on business relationships . . .”).

65. *Why Arbitrate*, CCA, <https://www.ccarbitrators.org/why-arbitrate/> [<https://perma.cc/EK4S-M2ZK>] (“Implemented properly, business or commercial arbitration can avoid the time, expense, and uncertainty of court proceedings. Business arbitration is particularly useful when parties want disputes resolved in a business setting – privately and efficiently by a decision-maker with industry sector expertise.”).

66. *Id.*

67. See Thomas R. Snider, Stephen Chan & John Olatunki, *Arbitration Is Cheaper – Myth or Reality?*, CHARLES RUSSELL SPEECHLYS (Oct. 31, 2023), <https://www.charlesrussellspeechlys.com/en/insights/expert-insights/dispute-resolution/2023/arbitration-is-cheaper--myth-or-reality/> (“At first blush, arbitration is frequently (and some might say mythically) viewed as a method of dispute resolution that is a faster and cheaper alternative to court litigation. But . . . the reality is more complex than that.”).

68. COLL. OF COM. ARBS., PROTOCOLS FOR EXPEDITIOUS, COST-EFFECTIVE COMMERCIAL ARBITRATION 1 (Thomas J. Stipanowich, Curtis E. von Kahn & Deborah Rothman eds., 2010) (“[B]usiness-to-business arbitration has taken on the trappings of litigation—extensive discovery and motion practice, highly contentious advocacy, long cycle time and high cost.”).

experienced human arbitrator, considerable discussion, argument, and back-and-forth, and lots and lots of documents.<sup>69</sup> In the same way that arbitration is a viable, more cost-effective alternative to in-court litigation, there is plenty of room for arbitration to be disrupted by innovative dispute-resolution tools that promise to reduce transaction costs even further.

## 2. Interest in Arbitration Suggests Transaction Costs Are High Priority—and Big Business

For the purposes of this Article, it suffices to point out that even while arbitration is far from perfect, companies still *choose* to arbitrate all the time.<sup>70</sup> Stated another way, companies routinely have disagreements that they either cannot or will not resolve themselves, and that are significant enough to justify the costs of resolving the dispute. (For both parties, I might add: The party bringing the action believes the value of resolving it outweighs the many costs of bringing it, while the defending party believes the cost of fighting back is less than the cost of settlement or default.) That parties routinely choose arbitration over litigation due to its reduced transaction costs—despite the fact the courts have additional safeguards against getting results wrong (*i.e.*, more extensive discovery, and a built-in appeals process)—is meaningful, and begs the question: *If companies could trade more safeguards for even more drastically reduced transaction costs, would they do so?*

That very idea—trading off additional safeguards in exchange for deeply reduced transaction costs—is the concept Part II turns to, which represents the center of gravity for this entire Article.

## II. OUR PUSH-BUTTON FUTURE: TOWARDS PROBABILISTIC DISPUTE RESOLUTION

*“In law, nothing is certain but the expense.”*

—Samuel Butler, seventeenth-century English writer<sup>71</sup>

In the last Part, we considered a hypothetical scenario; let’s pick it back up. Recall that you are a customer of ServiceCo, and you have reason to believe you have a claim against the company worth \$210. But

---

69. See also Thomas J. Stipanowich, *Arbitration: The “New Litigation,”* 2010 U. ILL. L. REV. 1, 8–9 (“By the beginning of the twenty-first century, however, it was common to speak of U.S. business arbitration in terms similar to civil litigation—‘judicialized,’ formal, costly, time-consuming, and subject to hardball advocacy.”).

70. See *Why Arbitrate*, *supra* note 65.

71. KATHRYN ZULLO, *LAWYER’S WIT AND WISDOM: QUOTATIONS ON THE LEGAL PROFESSION*, IN BRIEF 188 (Bruce Noah & Allan Zullo eds., 1995).

as we already discussed, you don't have any good options. Lawyers and arbitration are too expensive; the price tag of small-claims court looks good, but the total transaction costs make it nowhere near worthwhile; and even if you could talk yourself into being satisfied with a fractional payout in the form a class action or mass arbitration, kicking either of those processes off is all but entirely out of your control. Today, your logical course of action is to make a mental note not to do business with that company in the future and then move on.<sup>72</sup> But, as this Article suggests, thanks to advances in artificial intelligence and computational statistics, the technology *already* exists to create a novel, effective dispute-resolution tool that is perfectly suited for situations just like this one.

*A. A Thought Experiment: Pushing the Probabilistic Magic Button*

Imagine that instead of having no economical recourse against ServiceCo, you had access to a magic button that, when pressed, would instantly resolve your dispute, cost just \$20 to push, and would result in ServiceCo being required to pay you the full value of your claim, assuming you are in the right. A few other notable points: On the one hand, both you *and* ServiceCo would be legally bound to the result, and in pushing the button, you give up any *previous* claims you may have against ServiceCo as well—whether you know about them or not. But on the other hand, there are no additional costs or risks: The most you stand to lose is \$20, and you stand to gain all of your \$210 back.

But there *is* one big catch: The magic button gets the result right only 90% of the time and gets it wrong 10% of the time—even in cases where you have good reason to be sure you'll win. *Would you push it?* Chances are you would, and with only a few qualifications, choosing to do so would hands-down be the economically correct answer.

First, the qualifications. Recall that when you push the button, you're giving up on the right to *any* claims you might have against ServiceCo. If instead of ServiceCo the company were actually RoofCo, who recently installed a new \$15,000 roof on your house, you may not be so quick to give up your right to bring claims: There could be problems with your roof that you don't yet know about—so in that scenario, you might not want to push the button. But in the case of ServiceCo, you're done using their services, and if you had other claims

---

72. Of course, you *could* do other things if you wanted to, like leave scathing reviews in as many places on the internet as you can or post about your experience on social media. Doing so might result in the company noticing and trying to make amends to mitigate the reputational harm, but that is far from certain. The best you can reasonably hope for is that your efforts marginally contribute to future-looking behavior change at the company. And anyway, any of these options take time—which is to say, they involve transaction costs as well.

against the company, you would probably know it. And anyway, even if you *did* have some additional claims you didn't know about, they would be relatively small—so giving up these unlikely, unknown claims wouldn't a good reason not to push the button.

You also might not want to push the button if you're very sure that you *don't* have a valid claim against ServiceCo. Twenty dollars isn't much, but it isn't nothing, either—so if you're pretty sure you would lose, then you won't push the button.

But in most other cases, it *would* make economic sense for you to push the button. Let's use some easy numbers to see why.

Suppose that you are certain you're in the right, and you believe there is no chance you're mistaken. To calculate the perceived value to you of pushing the button, the first step is to multiply the total possible payout—\$210—by the product of (a) how strong you believe your claim is (let's call this "*X*") and (b) the button's accuracy rate (we'll call this "*Y*").<sup>73</sup> In this first example, you're certain that you're right about your claim, so *X* is 100%. As for *Y*, we've said that the button gets it right 90% of the time. In this case, the product of *X* and *Y* is 100% times 90%, which equals 90%. (We could also think of this  $1.0 * 0.9 = 0.9$ .) When we multiply \$210 by 90%, we get \$189. For the second and final step, we just subtract \$20—the cost of pushing the button. \$189 minus \$20 equals \$169, so the expected value to you of pushing the button is \$169. This value is greater than zero dollars, so from a purely economic perspective, you should definitely push it.

But what if you're less certain you're in the right—say, you're only 50% sure you're entitled to a payout? Let's run through the same simple equation: \$210 multiplied by the product of 50% and 90%, minus \$20, gives us the equation  $\$210 * (0.5 * 0.9) - \$20 = \$74.50$ .<sup>74</sup> In this case your claim is half as valuable to you as it was when you were 100% certain you were right, but still valuable enough that it's economically prudent to push the button. Only when you believe that you have roughly a 10% chance of success does it begin to stop making sense to push the button:  $\$210 * (0.106 * 0.9) - \$20 = \sim \$0$ , meaning pushing the button

---

73. See *Expected Value in Statistics: Definition and Calculating It*, STATS. HOW TO, <https://www.statisticshowto.com/probability-and-statistics/expected-value/> (“The basic expected value formula is the probability of an event multiplied by the amount of times the event happens:  $(P(x) * n)$ .” (emphasis omitted)).

74. The careful or statistically inclined reader may notice that this isn't strictly accurate—at least, not without additional clarification. The magic button gets the result wrong 10% of the time, and that can cut both directions: If the *customer* is right, there's still a 10% chance the result will favor the company—and if the *company* is right, there's still a 10% chance the result will favor the customer. There are many creative ways to account for this (such as accounting for it contractually or accounting for it in the stated margin for error), but there is no good reason to complicate an already statistics-heavy discussion with a nuance that doesn't have an adverse effect on the core idea this section presents.

becomes a wash when you're roughly 10% sure you'll win.<sup>75</sup> In every one of these scenarios, a magic button like the one described above, though *fundamentally* imperfect (it gets the result *wrong* 10% of the time), is worth pushing.

### *B. Revisiting the Small-Claims Wasteland*

At scale, a magic button like the one described above could transform entire categories of claims—and not just business-to-business disputes, but also most claims that fall into the small-claims wasteland described in Part I. The reason is simple: Even though the magic button has a relatively high margin of error (it gets a full 10% of claims completely wrong), for claims in the wasteland, the value of pushing the button will nearly always exceed the value of doing nothing, which is the next-best option for any claim in the wasteland.<sup>76</sup>

What's more, there is a flipside to the small-claims wasteland that we haven't yet discussed. While claims in the wasteland may start out as functionally illiquid from the perspective of individual claimants, they are *not* illiquid from the company's perspective: Claims like these often make up the rocket fuel that plaintiff's lawyers need in order to launch class actions and mass arbitrations. This means that it isn't just the claimants themselves that have a powerful interest in getting claims out of the wasteland; the would-be defendant companies have an interest in emptying the wasteland as well. The next Part discusses the interests of the companies in more detail, but before moving on, it's worth tackling a potential elephant in the room: Given that our legal system is principally based on sorting right from wrong, why should it be acceptable to use a dispute-resolution tool with a known and significant margin for error?

### *C. Why Lawyers Should Embrace the Probabilistic Approach*

Many very thoughtful scholars have already suggested that AI could be used as a tool for resolving disputes in a wide array of contexts.<sup>77</sup> And though I certainly agree with many of the views and proposals so far expressed on this topic, merely using AI in dispute resolution is not the

---

75. In reality, your threshold would probably be a bit higher than 10%. After all, in addition to your \$20, you're also giving up your right to bring any other unknown claims. That may be worth *very* little to you, but probably more than \$0.

76. See *supra* Section I.A.

77. See generally Ayelet Sela, *Can Computers Be Fair? How Automated and Human-Powered Online Dispute Resolution Affect Procedural Justice in Mediation and Arbitration*, 33 OHIO ST. J. DISP. RESOL. 91 (2018); David Horton, *Forced Robot Arbitration*, 109 CORNELL L. REV. 679 (2024); Horst Eidenmüller & Faidon Varesis, *What Is an Arbitration? Artificial Intelligence and the Vanishing Human Arbitrator* (June 17, 2020) (unpublished manuscript), <https://ssrn.com/abstract=3629145>.

point of this Article; rather, this Article turns to AI as a means to an end. The novel proposition of this Article is that there is significant pent-up market demand for high-efficiency, reduced-accuracy dispute resolution tools, *of a kind that modern AI technology happens to be perfectly positioned to deliver*, in a form that is closely analogous to the magic button just described.

But whether or not this Article is correct about the market demand, it is reasonable to ask whether a dispute-resolution tool with a built-in margin for error (*i.e.*, a 10% chance of getting the result wrong) should be acceptable at all, on the grounds that such an option might undermine the integrity of our legal system. It would do no such thing, for two main reasons. First, the law already implicitly contemplates and accepts significant margins for error when it comes to rendering judgments, in both civil *and* criminal cases, as well as in administrative proceedings.<sup>78</sup> And second, another even more highly respected discipline—medicine—explicitly relies on probability in almost all aspects of practice, and yet in that context, the use of probability serves to *improve* the medical profession’s trustworthiness rather than hamper it.

#### 1. The Preponderance of the Evidence Standard and Other Proofs that the Law Is Already Rooted in Statistics and Probability

At risk of opening a jurisprudential can of worms, it is probably a reasonable simplification to say that when it comes to civil disputes, the American legal system is deeply, if not solely, concerned with the dual (and occasionally dueling) concerns of accountability and fairness. At the theoretical level, wrongdoers should be held accountable<sup>79</sup> and victims should be made whole.<sup>80</sup> But at the practical level, our legal system acknowledges the extreme difficulty and complexity involved in getting judgments right.<sup>81</sup> There are no perfect factual records, laws frequently require nuanced interpretation, and humans are susceptible to manipulation, misunderstanding, and bias. And so, we’ve done our best: When the stakes are highest—for example, when it comes to the criminal justice system—our tolerance for getting the result wrong is very low, and so we require a finding that a crime was committed beyond a

---

78. For a discussion about evidentiary standards, see *infra* Section II.C.1.

79. Howard M. Erichson, *Victim Civil Litigation and the Elusive Goal of Corporate Accountability*, 72 DEPAUL L. REV. 239, 240 (2022–23).

80. *Id.* at 244.

81. For example, the Supreme Court has noted that in criminal cases, the “beyond reasonable doubt” standard does not require proof “to a mathematical certainty.” See *Holland v. United States*, 348 U.S. 121, 138 (1954).

reasonable doubt,<sup>82</sup> usually by unanimous juries.<sup>83</sup> But when the stakes are arguably lower, such as in civil disputes, our tolerance for getting it wrong is considerably higher: We require a finding of fault only by a preponderance of the evidence.<sup>84</sup> In other contexts we have other standards; for example, the clear and convincing evidence standard, commonly found in the context of administrative law, arguably sits somewhere in between the preponderance of the evidence standard and beyond reasonable doubt.<sup>85</sup>

I submit that, though lawyers may not view these standards in mathematical terms, each of them contemplates an acceptable probability of error. This is clearest for the “preponderance of the evidence” standard, which amounts to a finding of “more likely than not.”<sup>86</sup> In probabilistic terms, that is simply a greater than 50% chance. As a practical matter, then, it is not a valid defense on appeal to argue one’s belief that the jury had a 49.99% chance of being wrong: Our system *bakes in* that possibility.

The “beyond reasonable doubt” standard is much harder to quantify, but the general idea remains the same. Whatever “beyond reasonable doubt” means, it does *not* mean *absolute certainty*.<sup>87</sup> Here, too, our system contemplates the possibility that it may be just and proper to convict someone, even without 100% certainty they committed the crime.<sup>88</sup> Of course, this is not itself evidence of callousness, and our Constitution insists on additional safeguards in the form of due process rights. Rather, having a standard that is less than absolute certainty amounts to a tacit recognition that such certainty is, as a practical matter,

82. See Loretta B. DeLoggio, “*Beyond a Reasonable Doubt*” — *A Historic Analysis*, N.Y. ST. BAR J., April 1986, at 19.

83. *Johnson v. Louisiana*, 406 U.S. 356, 369–70 (1972) (Powell, J., concurring) (citing “an unbroken line of cases reaching back into the late 1800’s [in which] the Justices of this Court have recognized, virtually without dissent, that unanimity is one of the indispensable features of federal jury trial”).

84. See Neil Orloff & Jerry Stedinger, *A Framework for Evaluating the Preponderance-of-the-Evidence Standard*, 131 U. PA. L. REV. 1159 (1983).

85. See Charles E. Phipps, *The Presumption of Administrative Correctness: The Proper Basis for the Clear and Convincing Evidence Standard*, 10 FED. CIR. BAR. J. 143, 144 (2001).

86. *Rogers v. Wal-Mart Stores, Inc.*, 230 F.3d 868, 871 (6th Cir. 2000) (noting preponderance of the evidence equates to a finding of “more likely than not” (citing *Gafford v. Gen. Elec. Co.*, 997 F.2d 150, 158 (6th Cir. 1993))); *Sanchez v. Monumental Life Ins. Co.*, 102 F.3d 398, 404 (9th Cir. 1996) (same); *Tapscott v. MS Dealer Serv. Corp.*, 77 F.3d 1353, 1357 (11th Cir. 1996) (same).

87. See, e.g., NINTH CIR. JURY INSTRUCTIONS COMM., MANUAL OF MODEL CRIMINAL JURY INSTRUCTIONS 115 (2024) (“Proof beyond a reasonable doubt is proof that leaves you firmly convinced the defendant is guilty. It is not required that the government prove guilt beyond all possible doubt.”).

88. See Timothy James Ting, *It’s Time To Define ‘Beyond a Reasonable Doubt,’* ILL. BAR J., July 2018, at 24.



often impossible.<sup>89</sup> It is an effort to balance the societal harms of letting a crime go unpunished against the possibility that the justice system gets the result wrong.<sup>90</sup> Notably, in situations where the consequences of being wrong are especially dire—death penalty cases, for example—there are heightened requirements, such as requiring that death sentences be imposed by a jury, rather than by a single judge.<sup>91</sup>

Notably, the law uses percentages—if not probability—in even more explicit ways as well. For example, states across the country have contributory-negligence rules that hinge on percentages.<sup>92</sup> In some states, including Colorado, plaintiffs cannot recover if they’re found 50% or more at fault.<sup>93</sup> Whereas in other states, like Georgia, a defendant’s total liability will be reduced by the plaintiff’s percentage of fault.<sup>94</sup> In still other states, like Nebraska, both recovery rules coexist.<sup>95</sup>

The point is this: Even though we do not typically think about the law as having a probabilistic element to it, it certainly does—and acknowledging that fact should not have the effect of undermining our trust in the law. Nor would trust in the legal system be damaged by the introduction of voluntary dispute-resolution mechanisms that are more explicitly probabilistic. In further support of that assertion, I submit the entire field of medicine as Exhibit A.

---

89. See *United States v. Hannigan*, 27 F.3d 890, 894 n.3 (3d Cir. 1994) (“The fact is that the ‘beyond a reasonable doubt’ standard does not require 100% probability (or greater than 99% probability) of guilt in order to sustain a conviction. Since unassailably accurate knowledge of any past event is impossible, requiring absolute certainty to meet the beyond a reasonable doubt standard would mean that no one could ever be convicted of any crime. In his treatment of the subject, Judge Jack Weinstein concluded that the beyond a reasonable doubt standard most likely requires between 95–99% probability, not 100% probability.”).

90. Jon O. Newman, *Taking “Beyond a Reasonable Doubt” Seriously*, 103 JUDICATURE, no. 2, 2019, at 33, 34.

91. See *Sentencing*, OFFS. U.S. ATT’YS, <https://www.justice.gov/usao/justice-101/sentencing> [<https://perma.cc/4ZHT-X27P>].

92. See *Comparative & Contributory Negligence Laws: 50-State Survey*, JUSTIA, <https://www.justia.com/injury/negligence-theory/comparative-contributory-negligence-laws-50-state-survey/> [<https://perma.cc/DNU2-EMHE>] (Jan. 2024).

93. COLO. REV. STAT. § 13-21-111 (2024).

94. GA. CODE ANN. §§ 51-11-7, -12-33 (2024).

95. NEB. REV. STAT. § 25-21,185.09 (2022).

## 2. Medicine Routinely Relies on Probability in Life-or-Death Decision-Making

Both in the United States and around the globe, medical doctors are routinely deemed some of the most-trusted professionals.<sup>96</sup> (Lawyers frequently rank near the bottom,<sup>97</sup> but let's discuss that dismal bit of trivia another day.) What's remarkable is that doctors are widely trusted even though they routinely rely on probability and statistics to make decisions that are quite literally matters of life and death.<sup>98</sup>

Take, for example, diagnosis. Probability is *fundamental* to diagnosing conditions and diseases—and the prevailing view is that in describing diagnostic uncertainty, doctors should describe that uncertainty in clear, probabilistic, *numerical* terms.<sup>99</sup> This is *not* a perfect

96. Megan Brenan & Jeffrey M. Jones, *Ethics Ratings of Nearly All Professions Down in U.S.*, GALLUP (Jan. 22, 2024), <https://news.gallup.com/poll/608903/ethics-ratings-nearly-professions-down.aspx> [<https://perma.cc/HG42-A87P>]; Nicolas Boyon, *Doctors Are the Most Trusted Profession in the U.S. and Across the World*, IPSOS (Oct. 12, 2021), <https://www.ipsos.com/en-us/news-polls/global-trustworthy-index-US-version>.

97. Brenan & Jones, *supra* note 96; Boyon, *supra* note 96.

98. See Joseph M. Guileyardo, *Probability and Uncertainty in Clinical and Forensic Medicine*, 28 BAYLOR U. MED. CTR. PROC. 247, 247 (2015) (“In medicine, absolute diagnostic certainty is the exception rather than the rule. Nevertheless, many physicians in training are uncomfortable with the degree of uncertainty involved in therapy and diagnosis, although Sir William Osler wrote, ‘Medical education is, to a large extent, education for uncertainty.’” (citation omitted)).

99. Take, for example, this excerpt from this 1989 medical textbook published by the National Academy Press, describing the “First Principles” considerations for how to conduct diagnostics:

Principle I: Probability is a Useful Representation of Diagnostic Uncertainty.

Uncertainty is unavoidable. How can we best respond to it? A starting point is to adopt a common language. Some express their uncertainty as the *probability* that the patient has a specified disease. By using probability rather than ambiguous terms such as “probably” or “possibly,” the clinician expresses uncertainty quantitatively. More important, probability theory allows one to take new information and use Bayes’ theorem to calculate its effect on the probability of disease. These advantages are compelling, and our approach to test evaluation is based on providing the information required to use probability theory to interpret and select diagnostic tests.

*Example:* In a patient with chest pain, past history is very useful when trying to decide whether he or she has coronary artery disease. Patients whose pain is typical of angina pectoris and is also closely linked to overexertion are said to have “typical angina pectoris.” Over 90 percent of men with this history have coronary artery disease. When anginal pain is less predictably caused by exertion, the patient is said to have “atypical angina.” About two-thirds of men with this history have coronary artery disease.

Physicians who are uncertain about the meaning of a patient's chest pain often ask the patient to undergo an exercise test. The probability of coronary artery

system; misdiagnosis is a big deal, and a real challenge.<sup>100</sup> Some estimate that between 10 and 15% of all diagnoses are wrong,<sup>101</sup> and that in the United States alone, over 100,000 emergency department patients are harmed each year due to diagnostic errors.<sup>102</sup>

Probability is critical to the practice of medicine *after* diagnosis, too. For a given condition, there are often multiple potential treatment options, and helping patients decide what course of treatment to take also means describing things in probabilistic terms.<sup>103</sup> In fact it is so central to treatment that clinicians and patients commonly rely on charts and graphs that describe one's chances of survival in probabilistic, and often macabre, terms—*i.e.*, given a course of treatment, a chart may describe

---

disease after a positive exercise test may be calculated with Bayes' theorem. If the history is typical angina, the probability after a positive test is nearly 1.0. If the history is atypical angina, the probability after a positive test is about 0.90.

*Comment:* Estimating the probability of coronary artery disease helps to identify the situations in which the probability of disease will be altered dramatically by an abnormal test.

HAROLD SOX, SUSAN STERN, DOUGLAS OWENS & HERBERT L. ABRAMS, INST. OF MED., ASSESSMENT OF DIAGNOSTIC TECHNOLOGY IN HEALTH CARE: RATIONALE, METHODS, PROBLEMS, AND DIRECTIONS 24 (1989).

100. See Annalisa Merelli, *Misdiagnoses Cost the U.S. 800,000 Deaths and Serious Disabilities Every Year, Study Finds*, STAT (July 21, 2023), <https://www.statnews.com/2023/07/21/misdiagnoses-cost-the-u-s-800000-deaths-and-serious-disabilities-annually-study/> [<https://perma.cc/U8S3-6Y97>] (noting that 371,000 people die each year following misdiagnosis and 424,000 are permanently disabled).

101. David E. Newman-Toker et al., *Rate of Diagnostic Errors and Serious Misdiagnosis-Related Harms for Major Vascular Events, Infections, and Cancers: Toward a National Incidence Estimate Using the "Big Three,"* 8 DIAGNOSIS 67, 68 (2021).

102. DAVID E. NEWMAN-TOKER ET AL., JOHNS HOPKINS UNIV. EVIDENCE-BASED PRAC. CTR., U.S. DEP'T OF HEALTH & HUM. SERVS., AHRQ PUB. NO. 22(23)-EHC043, DIAGNOSTIC ERRORS IN THE EMERGENCY DEPARTMENT: A SYSTEMATIC REVIEW 1 (2022).

103. As Dr. Benjamin D. Sommers and Professor Richard Zeckhauser explain: In any treatment decision, there are numerous relevant probabilities to consider. Take, for example, the initial treatment decision for a newly identified cancer. Among the most important probabilities is the likelihood of full remission of the cancer after treatment, the risk of cancer recurrence after remission, and the risks of various side effects from treatment. For each competing treatment option, the probability of each of these distinct outcomes should be estimated for the patient in question. While this approach is familiar, at least in concept, to clinicians and presumably to many patients as well, such problems are rarely attacked systematically, which allows strong decision biases to creep into choices.

Benjamin D. Sommers & Richard Zeckhauser, *Probabilities and Preferences: What Economics Can Teach Doctors and Patients Making Difficult Treatment Decisions*, 26 UROLOGIC ONCOLOGY 669, 669–70 (2008).

the patient's statistical likelihood of living an additional five years, or ten.<sup>104</sup>

None of this is to criticize physicians or the medical profession, nor should this discussion undermine our faith in the medical profession—far from it. It is a simple matter of fact that a profession far more trusted and respected than the law acknowledges and relies on statistics and probability in making decisions that are at least as high stakes as most legal disputes. What's more, this reliance on probabilistic decision-making underscores the medical profession's commitment to the idea that taking *some* action is often superior to taking *no* action, even in situations where there is extreme uncertainty, and the consequences of being wrong are severe.<sup>105</sup>

That last point—that despite the high probability and stakes of being wrong, action is often favorable to inaction—is at the root of why we can be sure there is an ethical, socially acceptable role to play for voluntary probabilistic models of dispute resolution. Just like in medicine, when it comes to legal disputes, there are situations where it makes no sense to let perfect be the enemy of good.

### 3. Voluntary Dispute Resolution Is the Perfect Place To Experiment with Creating a Probabilistic Decision-Making Market

Perhaps the strongest argument *against* probabilistic dispute resolution is that no one should be *forced* to gamble with their rights. Quite so—and for that reason this Article does not favor of mandating the use of AI models for dispute resolution—at least, not any time soon.<sup>106</sup> But neither should these tools be off the table for use in a voluntary context. On freedom-of-contract grounds alone, this seems clear: If parties are allowed to contractually resolve disputes using a mechanism as perfectly arbitrary as a coin toss,<sup>107</sup> then we should not prevent them from resolving disputes using a mechanism that is mostly nonarbitrary, but includes a small chance of an arbitrary outcome.<sup>108</sup> Moreover, when one considers the other major benefits of allowing this kind of dispute

---

104. *Id.* at 670 (“[T]he medical literature often provides integrated probabilities in the form of 5-year and 10-year survival rates. These values tell a clinician and patient, taking all the various intermediate outcomes into account, the ultimate probability of mortality vs. survival in a given time period.”).

105. *See id.* at 672.

106. As I suggest elsewhere in this Article, there *may* come a time when AI dispute resolution tools are demonstrably fairer and more reliable than our current human-driven processes, but we're certainly not there yet.

107. Ken Adams, “*Let’s Flip a Coin*”: *The Role of Chance in Contract Mechanisms*, ADAMS ON CONT. DRAFTING (July 30, 2013), <https://www.adamsdrafting.com/the-role-of-chance-in-contract-mechanisms/> [<https://perma.cc/VY3D-PM59>].

108. Recall that the “magic button” I propose has a 10% chance of an arbitrarily erroneous result.

resolution,<sup>109</sup> it seems clear that we should be optimistic and enthusiastic about giving it a try.

Of course, this all hinges on whether such a tool is indeed possible. This Article enthusiastically turns to that question now.

### III. USING NEURAL NETWORKS AND DEEP LEARNING TO CREATE A PLATFORM FOR PROBABILISTIC DISPUTE RESOLUTION

Artificial intelligence has been around in some form for over half a century at this point, but it wasn't until ChatGPT was announced in November 2022<sup>110</sup> that AI began to become a part of everyday life. Today, self-driving taxis powered by AI are carrying paying customers all over Phoenix, Los Angeles, and San Francisco;<sup>111</sup> some AI systems are achieving top levels of performance in International Math Olympiad competitions and impressing Fields Medal recipients along the way;<sup>112</sup> other models are capable of passing the bar exam;<sup>113</sup> and some AI models have gotten so good at playing complex strategy games that one master of the ancient and still-popular Chinese strategy game, *Go*, recently quit playing, lamenting that there is now “an entity that cannot be defeated.”<sup>114</sup> In short, this technology is *powerful*—and still in a state of rapid development.<sup>115</sup> The question for our purposes is: *Could an AI system be built today that could resolve certain categories of disputes with close to a 10% error rate or better?*

---

109. Some of the benefits were already discussed earlier in this Part and more are discussed later on in Part IV.

110. *Introducing ChatGPT*, OPENAI (Nov. 30, 2022), <https://openai.com/index/chatgpt/>.

111. See The Waymo Team, *Waymo Significantly Outperforms Comparable Human Benchmarks over 7+ Million Miles of Rider-Only Driving*, WAYMO: WAYPOINT (Dec. 20, 2023), <https://waymo.com/blog/2023/12/waymo-significantly-outperforms-comparable-human-benchmarks-over-7-million/> [<https://perma.cc/7F7F-5NBQ>].

112. AlphaProof & AlphaGeometry Teams, *AI Achieves Silver-Medal Standard Solving International Mathematical Olympiad Problems*, GOOGLE DEEPMIND (July 25, 2024), <https://deepmind.google/discover/blog/ai-solves-imo-problems-at-silver-medal-level/> [<https://perma.cc/S5YH-VCK3>] (Professor Sir Timothy Gowers, International Mathematical Olympiad Gold Medalist and Fields Medal winner recently said of an AI model: “The fact that the program can come up with a non-obvious construction like this is very impressive, and well beyond what I thought was state of the art.”).

113. Pablo Arredondo, *GPT-4 Passes the Bar Exam: What That Means for Artificial Intelligence Tools in the Legal Profession*, STAN. L. SCH. (Apr. 19, 2023), <https://law.stanford.edu/2023/04/19/gpt-4-passes-the-bar-exam-what-that-means-for-artificial-intelligence-tools-in-the-legal-industry/> [<https://perma.cc/882P-XECP>].

114. *Go Master Quits Because AI ‘Cannot Be Defeated,’* BBC NEWS (Nov. 27, 2019), <https://www.bbc.com/news/technology-50573071> [<https://perma.cc/94G5-HMSW>].

115. See Will Henshall, *4 Charts That Show Why AI Progress Is Unlikely To Slow Down*, TIME (Nov. 6, 2023, 4:13 PM), <https://time.com/6300942/ai-progress-charts/> [<https://perma.cc/S4V2-22UA>].

To begin to answer that question, it makes sense to spend a few paragraphs describing what this technology actually is, and to give the reader a general understanding of how modern AI does and does not work. While the history of this technology is fascinating and worth your time exploring, there are a great many excellent resources available on that subject for the curious reader—so in the spirit of minimizing transaction costs, I won't spend any time discussing that history here. Instead, I commend to the reader the very well-curated Wikipedia page on the “History of Artificial Intelligence,”<sup>116</sup> as well as to a very nice article on the subject written by the staff of the online learning portal Coursera.<sup>117</sup> With that, let's discuss present-day AI—specifically, machine learning, neural networks, and deep learning.

*A. A Primer on Artificial Intelligence, Neural Networks,  
and Deep Learning*

As a concept, artificial intelligence has been around for a long time. But where previous-generation AI systems were based on strict conditional rules—for example, *if X is true, then Y must be true*<sup>118</sup>—modern AI systems are based on an entirely different paradigm: machine learning using neural networks.<sup>119</sup> Or, as Kent Walker, President of Global Affairs at Google and Alphabet has described it, modern AI systems really just boil down to computational statistics.<sup>120</sup>

1. Machine Learning, Deep Learning, and Neural Networks

Machine learning is a branch of artificial intelligence and computer science that involves using statistical algorithms to help an AI model learn from data, and then make predictions or produce outputs that are *informed by* those data.<sup>121</sup> Where the outputs of earlier rules-based systems were rigid—in a rules-based system, a given input will always

116. *History of Artificial Intelligence*, WIKIPEDIA, [https://en.wikipedia.org/wiki/History\\_of\\_artificial\\_intelligence](https://en.wikipedia.org/wiki/History_of_artificial_intelligence) [https://perma.cc/9UE5-THX2] (Jan. 23, 2025).

117. Coursera Staff, *The History of AI: A Timeline of Artificial Intelligence*, COURSEERA, <https://www.coursera.org/articles/history-of-ai> [https://perma.cc/YGH9-VYAA] (Oct. 25, 2024).

118. Vasudevan Swaminathan, *The Conundrum of Using Rule-Based vs. Machine Learning Systems*, ZUCI SYS., <https://www.zucisystems.com/blog/the-conundrum-of-using-rule-based-vs-machine-learning-systems/> [https://perma.cc/Q8L4-2JMA] (“Rule-based systems are computer programs that use if-then rules to make decisions and perform tasks.”).

119. IBM Data & AI Team, *supra* note 12.

120. See Kent Walker, *Engaging in Frank Discussions About the Promise and Risks of Emerging Tech*, LINKEDIN (June 27, 2023), <https://www.linkedin.com/pulse/engaging-frank-discussions-promise-risks-emerging-tech-kent-walker/>.

121. See IBM Data & AI Team, *supra* note 12.

produce the same output unless a human reprograms the computer and rewrites the rules—the outputs of machine learning systems can change based on how the model is trained.<sup>122</sup> With traditional machine learning, the data used to train the AI model is typically carefully selected by the human training the model, and the data must be carefully structured in a way that the AI model is able to read and understand.<sup>123</sup>

But while those traditional machine learning systems are powerful and impressive in their own right, the key innovation that has led to the current era of generative AI and large language models like ChatGPT,<sup>124</sup> Claude,<sup>125</sup> and Gemini<sup>126</sup> is the addition of neural networks, through a process called deep learning.<sup>127</sup> Where traditional machine learning requires structured data and considerable human hand-holding during the training process, deep learning using neural networks enables AI models to train on vast amounts of *unstructured* data, with very little need for human input along the way.<sup>128</sup> If a dispute-resolution platform like the one this Article propose gets built, it will be on the back of deep learning and neural networks—so let’s take some time to discuss these in a bit more depth.

Neural networks are the backbone of deep learning. As the name suggests, neural networks are loosely modeled after the human brain.<sup>129</sup> Neural networks comprise many<sup>130</sup> interconnected nodes—think of these as artificial neurons.<sup>131</sup> Whereas in the human brain neurons store information biologically, in artificial neural networks, the nodes

---

122. *See id.*

123. *Id.* (“Classic or ‘nondeep’ machine learning depends on human intervention to allow a computer system to identify patterns, learn, perform specific tasks and provide accurate results. Human experts determine the hierarchy of features to understand the differences between data inputs, usually requiring more structured data to learn.”).

124. *See Introducing ChatGPT*, *supra* note 110.

125. *See Meet Claude*, ANTHROPIC, <https://www.anthropic.com/claude> (last visited Jan. 24, 2025).

126. Sundar Pichai & Demis Hassabis, *Introducing Gemini: Our Largest and Most Capable AI Model*, GOOGLE: KEYWORD (Dec. 6, 2023), <https://blog.google/technology/ai/google-gemini-ai/> [<https://perma.cc/63F3-QZWQ>].

127. *See IBM Data & AI Team*, *supra* note 12.

128. *See id.* (“While the subset of AI called deep machine learning can leverage labeled data sets to inform its algorithm in supervised learning, it doesn’t necessarily require a labeled data set. It can ingest unstructured data in its raw form (for example, text, images), and it can automatically determine the set of features that distinguish ‘pizza,’ ‘burger’ and ‘taco’ from one another.”).

129. *See id.* (“They are called ‘neural’ because they mimic how neurons in the brain signal one another.”).

130. “Many” is an understatement. *See Matthias Bastian*, *GPT-4 Has More than a Trillion Parameters - Report*, DECODER (Mar. 25, 2023), <https://the-decoder.com/gpt-4-has-a-trillion-parameters/> [<https://perma.cc/3BJC-HN4C>].

131. *See IBM Data & AI Team*, *supra* note 12.

store numbers, called parameters.<sup>132</sup> As the model is exposed to training data, these parameters adjust, so that the model itself changes slightly after each training cycle. Though it's an imperfect comparison, it's helpful to think through how a small child might learn something new. Suppose you're trying to teach a three-year-old to recognize two similar fruits—say, a peach and an apple. One way to do it might be to show the child pictures and examples. At first, the two fruits might look the same to the child. After all, they're roughly the same size, they're both round, and they both may or may not have stems. But the more peaches and apples the child sees, the more they'll be able to properly identify each one: Peaches tend to be fuzzier, while apples are smooth. Apples can be bumpy at the bottom, whereas the bottom of a peach slopes into a point. Apples are sometimes green, but peaches never (or rarely) are. Before long, the child will easily be able to tell the difference between peaches and apples—and may even be able to draw one or the other on a piece of paper, in a way that *other* people would be able to recognize the chosen fruit. Through this whole process, the child's brain is digesting all the information it's taking in—building the ability to distinguish between these two fruits. Something similar happens in an artificial neural network: Each time the network is exposed to new training data, the parameters—those nodes, which, again, are acting a bit like artificial neurons—get updated to account for the new information.<sup>133</sup>

A few things stand out as contributing to the power of deep learning using neural networks. The first is that modern AI models have billions, and sometimes even *trillions*, of parameters.<sup>134</sup> This contributes to these models' ability to develop spectacular nuance and generate precise, impressive outputs.<sup>135</sup> Second and fundamentally, many deep learning algorithms used in training neural networks enable what's called “backpropagation” during the training process, which enables constant iteration and learning without human intervention: “In simple terms, after each forward pass through the network, backpropagation performs

---

132. Kian Katanforoosh, Daniel Kunin & Jiaju Ma, *Parameter Optimization in Neural Networks*, DEEPLARNING.AI (2019), <https://www.deeplearning.ai/ai-notes/optimization/index.html> [<https://perma.cc/XDF9-H6KW>] (“In machine learning, you start by defining a task and a model. The model consists of an architecture and parameters. For a given architecture, the values of the parameters determine how accurately the model performs the task.”).

133. *See id.*

134. *See* Bastian, *supra* note 130.

135. *See* Rakshith Vasudev, *Understanding and Calculating the Number of Parameters in Convolution Neural Networks (CNNs)*, TDS ARCHIVE (Feb. 11, 2019), <https://medium.com/towardsdatascience.com/understanding-and-calculating-the-number-of-parameters-in-convolution-neural-networks-cnns-fc88790d530d> [<https://perma.cc/4TDK-4ATP>] (“Parameters in general are weights that are learnt during training. They are weight matrices that contribute to model's predictive power, changed during back-propagation process.”).



a backward pass while adjusting the model's parameters (weights and biases)."<sup>136</sup> Third, thanks to the structure of neural networks and the design of their accompanying algorithms, they're able to train on *unstructured* data—meaning data that has not been carefully packaged by a human for machine consumption, in a way that surfaces and quantifies specific characteristics.<sup>137</sup> This greatly decreases the amount of human effort involved in training the data, which enables training sets to be extremely large. So large, in fact, that several modern AI models are trained on copies of practically the entire internet.<sup>138</sup>

One last important point for context: In just a few short years, AI models using deep learning have become incredibly powerful—so much so that “[m]ost of the benchmarks are hitting a point where we cannot do much better, 80–90% accuracy.”<sup>139</sup> Indeed, in many tasks (including several critical to dispute resolution), AI already surpasses the abilities of humans. For example, benchmarking shows that AI beats humans in reading comprehension, language understanding, handwriting recognition, speech recognition, and image recognition.<sup>140</sup> Many researchers expect the progress of AI “to continue at a breakneck speed at least for the next few years”<sup>141</sup>—which is remarkable, given what this technology is already capable of. In fact, between the time this Article was accepted for publication and before it went to print, several new and ground-breaking AI models were released, including one that is quite literally designed to solve problems, reason, and fact-check itself<sup>142</sup>—precisely the kinds of things we would hope for an AI dispute-resolution platform to excel at.

## 2. Training and Validating an AI Model for Dispute Resolution

With those fundamentals under our belts, let's turn our attention to the process of creating an AI model capable of functioning like the

---

136. Simeon Kostadinov, *Understanding Backpropagation Algorithm*, TDS ARCHIVE (Aug. 8, 2019), <https://medium.com/towards-data-science/understanding-backpropagation-algorithm-7bb3aa2f95fd> [<https://perma.cc/9SRZ-RPLZ>].

137. See IBM Data & AI Team, *supra* note 12.

138. Noor Al-Sibai, *AI Companies Running Out of Training Data After Burning Through Entire Internet*, FUTURISM: THE BYTE (Apr. 1, 2024, 6:00 PM), <https://futurism.com/the-byte/ai-training-data-shortage> [<https://perma.cc/MN7D-46WA>].

139. See Shana Lynch, *AI Benchmarks Hit Saturation*, STAN. UNIV. HAI (Apr. 3, 2023), <https://hai.stanford.edu/news/ai-benchmarks-hit-saturation> [<https://perma.cc/LLZ8-F79U>] (interviewing and quoting Vanessa Parli).

140. See Henshall, *supra* note 115.

141. *Id.*

142. Kylie Robison, *OpenAI Releases o1, Its First Model with 'Reasoning' Abilities*, VERGE (Sept. 12, 2024, 12:05 PM), <https://www.theverge.com/2024/9/12/24242439/openai-o1-model-reasoning-strawberry-chatgpt> [<https://perma.cc/2WXM-FB8K>].

“magic button” described above. First, creating this model would *not* require reinventing the wheel, so to speak. As one technology writer recently put it: “The future of AI is flexible, reusable AI models that can be applied to just about any domain or industry task.”<sup>143</sup> This is exactly right. Many AI companies have hit the market in these few short years since ChatGPT launched in late 2022, but very few of them have actually designed their own models from scratch.<sup>144</sup> Instead, they typically start with a “foundation model”<sup>145</sup>—think of this as one of the big AI models you’ve probably heard about: ChatGPT, Llama, Claude, Gemini, et cetera—and then make adjustments to the model to meet their company’s specific needs through a fine-tuning process.<sup>146</sup> That process involves additional training, with the goal of coaxing the model to handle specific kinds of inputs in specialized, more reliable ways.<sup>147</sup>

It is precisely this fine-tuning process that the AI dispute-resolution platform this Article proposes would need to undergo. The idea would be to curate a large dataset of dispute data—likely both real, historical examples, as well as human-made mock dispute data—to help the model learn how to handle specific categories of disputes. From there, the AI model would need to undergo a validation process, where it would be tested on additional dispute documents it had not seen before. This would serve to benchmark how frequently it gets the result right.<sup>148</sup> This process would be repeated and adjusted until the model achieved sufficient accuracy with a given class of disputes. Moreover, ideally the training and validation process would be carefully documented from start to finish, so that the platform creator could confidentially share the entire process of training the model with experienced, independent, third-party human arbitrators for review. Where warranted, those arbitrators could provide something akin to their seal of approval. (I envision something like a validation audit for AI dispute-resolution platforms.)

This is a good moment to remind the reader that the goal here would not be perfection<sup>149</sup>—at least, not in the short term. Even the best AI

---

143. Mike Murphy, *What Are Foundation Models?*, IBM (May 9, 2022), <https://research.ibm.com/blog/what-are-foundation-models> [https://perma.cc/AFF8-68UY].

144. See, e.g., Jai Vipra & Anton Korinek, *Market Concentration Implications of Foundation Models: The Invisible Hand of ChatGPT*, BROOKINGS (Sept. 7, 2023), <https://www.brookings.edu/articles/market-concentration-implications-of-foundation-models-the-invisible-hand-of-chatgpt/> [https://perma.cc/6X49-Y365].

145. See Murphy, *supra* note 143.

146. See Dave Bergmann, *What Is Fine-Tuning?*, IBM (Mar. 15, 2024), <https://www.ibm.com/topics/fine-tuning> [https://perma.cc/2WJF-L6SD].

147. *Id.*

148. I say more about the tricky question of correctness in the next Part.

149. Recall that the “magic button” thought experiment proposed in Part III contemplated a 10% error rate.

models do still occasionally hallucinate or produce erroneous answers.<sup>150</sup> (Though again, new and better models are being developed and released even as this Article was written that appear to be addressing these issues.<sup>151</sup>) This is precisely why the “magic button” thought experiment includes a ten percent error rate: Even a finely tuned AI model will get things wrong sometimes, especially if the specific facts of the dispute are in some way unusual. But the error rate did not render the magic button useless for dispute resolution, and the same should hold true for the error rate of an AI dispute-resolution platform—*especially* if the error rate is quantified through a structured, thorough, and audited validation process.

### *B. Other Characteristics of a Probabilistic AI Dispute-Resolution Platform*

Notwithstanding the high probability that a “magic button” level AI dispute-resolution platform could be created and deployed, there are a few quirks and pain points we should discuss. Some of them are specific to AI, while others will show up in any dispute-resolution scenario. Let’s walk through them.

#### 1. Lack of Discovery, Honesty, and Other Evidentiary Matters

Naturally, our “magic button” AI platform can’t resolve disputes without a factual record—and deciding what that factual record should include (and how to compile it) is almost as important as ensuring that the capabilities of the AI platform are validated. Striking the right balance here will be critical. While it’s conceivable that the platform could be trained to handle very large evidentiary records, AI models have limitations on how much context they can maintain at any one time,<sup>152</sup> and so the bigger the factual record, the greater the possibility that the platform will fail to capture an important fact or nuance.<sup>153</sup> That being

---

150. See Ziwei Xu, Sanjay Jain & Mohan Kankanhalli, *Hallucination Is Inevitable: An Innate Limitation of Large Language Models 1* (Jan. 22, 2024) (unpublished manuscript), <https://arxiv.org/pdf/2401.11817> [<https://perma.cc/QL9H-DF7L>].

151. See, e.g., Rachel Metz, *AI Startup Anthropic Says New Models Cut Hallucination Risks*, BLOOMBERG, <https://www.bloomberg.com/news/articles/2024-03-04/ai-startup-anthropic-launches-new-models-for-chatbot-claude> (Mar. 4, 2024, 3:29 PM) (noting AI-startup Anthropic claims “new versions of Claude will be twice as likely to answer a question correctly”).

152. See *Context Window for LLMs*, HOPSWORKS, <https://www.hopsworks.ai/dictionary/context-window-for-llms> [<https://perma.cc/AHC5-K64N>].

153. See Mae Cornes, *The Pitfalls of Prompt Overloading: Why More Tokens Mean Less Accuracy in LLMs*, DIGIT. J. (July 2, 2024), <https://www.digitaljournal.com/tech-science/the-pitfalls-of-prompt-overloading-why-more-tokens-mean-less-accuracy-in-llms/article> [<https://perma.cc/4AAH-7ZLM>].

said, the whole purpose of this platform is to drastically reduce transaction costs—so the size of the factual record should be minimized as much as possible anyway.

Keeping the factual record manageable will probably be more difficult in theory than in practice—at least for the categories of claims described earlier in this Article. For example, for cases in the small-claims wasteland, one could imagine limiting the factual record to agreements and invoices between the customer and the company, and all direct communications between the customer and the company, including emails, text-based chats, and recorded phone conversations. These records likely would already exist and be stored by the company, and so a well-designed platform could conceivably ingest them automatically (after inviting the consumer to review the compiled record, of course). As for business-to-business disputes, because the parties would pre-agree via contract to use the platform for certain kinds of narrow, breach-related disputes, the factual records would likely be limited and narrowed within that initial contract as well.

The platform could also conceivably invite the parties to make brief statements that could include factual allegations—though of course, this would introduce the possibility of embellishment or lying. Three interrelated methods of counteracting false statements come to mind. First, the platform should include a clear, prominent warning that lying to the platform in order to bolster their claim is a crime prosecutable as financial fraud. (Pointing to the relevant state fraud statutes would be a nice touch here, too.) Second, the platform could compare the statements to the factual record, looking for contradictions. Depending on the severity, the platform might simply disregard the contradicting statement—or, if it's clear the statement contained intentional false statements, enter a default judgment against the dishonest party.<sup>154</sup> Third, though it would increase the transaction costs significantly, the platform could build in an intermediate step that allows each party to see a list of all the factual allegations presented in the other party's statement. (The platform would extract them—something AI models are very good at.<sup>155</sup>)

---

154. It's important to acknowledge that some challenging edge cases requiring special handling would inevitably arise. For example, suppose the nature of the dispute is an invoice with an allegedly erroneous charge. In that situation, the customer's personal statement would undoubtedly contradict the disputed invoice. Here, the platform would need to somehow tag the invoice as the "disputed" document and be more circumspect about negating the customer's statement merely because the statement and the invoice are inconsistent with one another.

155. See Maciej P. Polak & Dane Morgan, *Extracting Accurate Materials Data from Research Papers with Conversational Language Models and Prompt Engineering*, 15 NATURE COMM'NS, no. 1569, 2024, at 1 ("In tests on materials data, we find precision and recall both close to 90% from the best conversational LLMs, like GPT-4. We demonstrate that the exceptional performance is enabled by the information retention

Each party could then flag those facts as “in dispute,” and if available, provide documentary evidence contradicting the facts they flagged.

The challenges that evidentiary records present to *any* dispute-resolution process are significant. This is the main reason why an AI platform like this one would *not* be a good fit for all disputes today, even if it had a 0% error rate on whatever record *is* provided to it. But once again, the entire point of this platform from the perspective of the claimant would be to exchange some amount of accuracy and consistency for *much* lower transaction costs. In that context, limiting the evidentiary record is fundamental to the platform’s value: Such limitations contribute positively to the transaction-costs optimizing exchange.

## 2. Beating Bias: The Black Box As a Feature Rather than a Bug

AI models based on deep learning are often described as having a “black box” problem: It is difficult or impossible to see, trace, or understand why a given input led to a particular output.<sup>156</sup> This has to do with the training process, and how these models work—recall that through the model-training process, the parameters in the model’s neural network get adjusted over and over again with each successive training iteration.<sup>157</sup> In other words, the ultimate state of the parameters is the result not of any *one* piece of training data, but rather of the entire *collection* of training data.<sup>158</sup> And because it is those parameters that ultimately determine the output, it is impossible to say that one piece of training data is responsible for this or that output.<sup>159</sup> This is often characterized as a negative aspect of AI models,<sup>160</sup> and it certainly can be in some contexts. For example, if improperly trained, deep learning models may exhibit some kind of unacceptable demographic bias, such as in misdiagnosis.<sup>161</sup>

But there is another way that AI dispute-resolution platforms might actually be able to *reduce* bias. Whether intentionally or unintentionally, humans take into account irrelevant and/or inappropriate demographic

---

in a conversational model combined with purposeful redundancy and introducing uncertainty through follow-up prompts.”).

156. See Lou Blouin, *AI’s Mysterious ‘Black Box’ Problem, Explained*, UNIV. MICH.-DEARBORN (Mar. 6, 2023), <https://umdearborn.edu/news/ais-mysterious-black-box-problem-explained> [<https://perma.cc/XN8E-RUY6>].

157. See IBM Data & AI Team, *supra* note 12.

158. *Id.*

159. See Blouin, *supra* note 156.

160. *Id.*

161. See generally Anurag Vaidya et al., *Demographic Bias in Misdiagnosis by Computational Pathology Models*, 30 NATURE MED. 1174 (2024).

information when making critical decisions that affect people's lives.<sup>162</sup> AI dispute-resolution platforms could be fine-tuned on dispute data that is scrubbed of all demographic markers, such that the model would be incapable of making inappropriate associations based on the parties' personal characteristics.

The platform could take additional bias-reducing steps after training as well, during the execution of the dispute-resolution process. For example, it would be possible to have a *different* AI model scrub all statements and evidence of the parties' personal characteristics, such as names, gender, ethnicity, accents, vocal range, et cetera. That first model could then hand off the scrubbed information to the dispute-resolution model, so that the dispute-resolution model would not even have *access* to any of the parties' personal characteristics.

### 3. Defining "Correct" Outcomes: Challenges and Opportunities

One of the most interesting aspects of AI dispute resolution is defining what a correct outcome actually is. When we talk about creating an AI dispute-resolution platform that reaches correct outcomes at least 90% of the time, how do we determine which results fall in the 90%, and which ones fall into the 10%?

Once again, this probably sounds much more complicated in theory than it is likely to be in practice, thanks to the fact that this platform would be trained on specific, narrow categories of disputes, and would not be expected to be or replace a generalist judge sitting in a federal district court. The solution to the "defining correct" problem is to involve highly experienced arbitrators very early in the training process. A panel of experienced arbitrators would help select the training data, and if mock training data is used (for example, to help train the model on a wider range of situations that real data may not be available for), the panel would assist in crafting it. The panel would also determine the right outcome for each case in the training data and provide clear, reasoned explanations for why the result is what it is. And then, during the validation process, a different arbitration panel would evaluate the AI platform's responses in order to validate the model and identify weaknesses that need to be addressed through additional training.

Put simply, experienced arbitrators would play a significant and ongoing role in these platforms—even while it would be capable of operating autonomously after training and validation.

---

162. See, e.g., Joe Hernandez, *White-Sounding Names Get Called Back for Jobs More than Black Ones, a New Study Finds*, NPR (Apr. 11, 2024, 5:00 AM), <https://www.npr.org/2024/04/11/1243713272/resume-bias-study-white-names-black-names> [https://perma.cc/3UHS-DGPD].

#### 4. The Pace of Innovations and What's Next

The final Part of this Article thinks through who might build these platforms, what they might look like in practice, and why both consumers and companies alike will want to use them. But before moving on to those subjects, it's worth a brief pause to surface some of the philosophical questions that models like these raise—and that the even more capable AI technologies of the future will demand we answer.

In December of 2023, Waymo—the autonomous taxi company—released its safety report based on over seven million miles of rider-only driving.<sup>163</sup> Compared to human drivers, Waymo's self-driving taxis resulted in an 85% reduction in injury-causing crashes, and a 57% reduction in police-reported crashes.<sup>164</sup> By all accounts, the data backs up what Waymo has been saying for some time now: Its driverless vehicles indeed appear to be safer than human-driven ones.<sup>165</sup> Considering where self-driving technology is today, when I imagine what my as-yet unborn grandchildren will say to me decades from now when I tell them that humans used to personally drive cars, I suspect they'll react by saying something like: *Were you all crazy? You let people drive cars and gigantic semi-trucks? People that might fall asleep at the wheel, or drive drunk, or text while driving? I can't believe how reckless your generation was!*

I suspect that forty or fifty years from now, *many* conversations will go that way—including discussions about how we resolve conflicts. Consider this: If I've managed to convince you that an AI dispute-resolution platform capable of a 90% accuracy rate is *already* not science fiction, how much longer will it take before such platforms can be 95% accurate? How about 99%—or 99.99%? An AI model that can achieve near-perfect accuracy would represent more than just a new way to resolve disputes—though it would certainly do that. *It would represent an entirely new way of defining what a correct outcome looks like.* Just like how future generations will marvel that we ever let uncareful humans drive a two-thousand-pound machine seventy miles per hour down a highway full of other cars and people, odds are we'll feel similarly about having emotional, bias-laden, hungry,<sup>166</sup> and frequently illogical humans resolve legal disputes.

---

163. The Waymo Team, *supra* note 111.

164. *Id.*

165. See Andrew J. Hawkins, *Waymo Has 7.1 Million Driverless Miles — How Does Its Driving Compare to Humans?*, VERGE (Dec. 20, 2023, 11:00 AM), <https://www.theverge.com/2023/12/20/24006712/waymo-driverless-million-mile-safety-compare-human> [https://perma.cc/72FE-NWRC].

166. See generally MICHAEL J. SAKS & BARBARA A. SPELLMAN, *THE PSYCHOLOGICAL FOUNDATIONS OF EVIDENCE LAW* (2016); Shai Danziger, Jonathan

Of course, that future is still years away. For now, it is sufficiently exciting that current-generation AI is probably capable of cheaply resolving countless cases stuck in the small-claims wasteland, or of helping companies navigate discrete breach-of-contract matters that would otherwise be too relationally or monetarily costly for them to deal with.

#### IV. WHO WILL BUILD IT, WHO WILL COME?

This Article closes with a practical discussion about how AI dispute-resolution platforms like the one described will come into existence, how they will evolve, why companies will use them, and what effects they are likely to have on mass dispute resolution.

##### A. *Who Will Build It?*

Speaking as someone who has spent time working with and advising hypergrowth startups, the answer to this question appears obvious to me: I anticipate that the first major AI dispute-resolution companies will be backed by venture capital firms.<sup>167</sup> This will almost certainly be true of whatever company attempts to develop an AI platform incorporating the probabilistic approach proposed in this Article. Because this platform would be most effectively deployed at very large companies with massive customer bases, the company that successfully develops it will have the kind of potential for hypergrowth and scale that venture capital firms specifically seek out in their portfolio companies.<sup>168</sup>

##### B. *Who Will Use It?*

Parts I and II made the case for why consumers stuck in the small-claims wasteland would use a platform like the one this Article proposes. But as this Article has stressed, the decision to use this AI platform would need to be voluntary for *both* parties. This means that for the platform to

---

Levav & Liora Avnaim-Pesso, *Extraneous Factors in Judicial Decisions*, 108 PNAS 6889, 6890 (2011) (finding that “the likelihood of a favorable ruling [by judges] is greater at the very beginning of the work day or after a food break than later in the sequence of cases”).

167. See Paul Davenport, *VC Funding in 2024: AI Breaks Records as Deal Volume Plummetts*, BOAST (July 15, 2024), <https://boast.ai/blog/vc-funding-in-2024-ai-breaks-records-as-deal-volume-plummetts/> [<https://perma.cc/FB3W-KH7E>].

168. Bob Zider, *How Venture Capital Works*, HARV. BUS. REV., Nov.–Dec. 1998, at 131, 133–34 (“By investing in areas with high growth rates, VCs primarily consign their risks to the ability of the company’s management to execute. VC investments in high-growth segments are likely to have exit opportunities because investment bankers are continually looking for new high-growth issues to bring to market.”).



work, the *company* would need to choose to use it as well. I propose that in many if not most cases, they will do just that.

### 1. Companies Will Want To Use the Platform with Their Customers

At least superficially, it is easy to argue that no company would be interested in voluntarily using an AI platform like the one proposed. After all, making it *easier* for one's customers to make claims against your company—and potentially extract payments from you that they otherwise could not have reached—seems like a good way to lose a lot of money. To put it unsentimentally, if a claim is truly in the small-claims wasteland, then arguably the claim is functionally resolved in the company's favor.

But this line of reasoning is flawed for several reasons. First, outstanding small claims contribute to the ticking time bomb of class action risk<sup>169</sup>—and the more similar claims there are in the small-claims wasteland, the bigger the bomb. For this reason, companies have an interest in reducing the number of outstanding claims against them—and because a part of the agreement to use the tool would involve the customer relinquishing *all* claims a customer might have against the company (known and unknown), the company would be able to entirely de-risk every customer that elects to use the platform.

Second, a company merely agreeing to use the dispute-resolution platform would not amount to a guaranteed payout to its customers. Far from it: If the data on forced arbitration outcomes are any indication of consumers' likelihood of succeeding on claims submitted to the AI platform, win rates are likely to be *very* low.<sup>170</sup> In the forced arbitration context, data suggest that win rates range from as low as 0.6% (for consumer claims against retail companies), up to a high of 16.5% (against tech companies)—and everything in between.<sup>171</sup> Low win rates like these are not particularly worrisome, especially in the context of very small claims.

Third and relatedly, use of the dispute-resolution tool need not be offered for free. In fact, in most cases it *shouldn't* be free—that would risk creating the perverse incentive of encouraging consumers to submit free erroneous claims, on the chance of obtaining an undeserved windfall—which, thanks to the model's error rate, at least some claimants

---

169. See *What Is a Putative Class Action?*, ANKIN L. (July 14, 2022), <https://ankinlaw.com/what-is-a-putative-class-action/> [https://perma.cc/5N37-VG9N] (“Class actions are appropriate when the damages claimed by each individual plaintiff, the individuals claiming injuries, are too small for individual lawsuits to be worthwhile.”).

170. See *Forced Arbitration*, *supra* note 41.

171. *Id.*

would receive.<sup>172</sup> Instead, the company could set the price of use to a level that is simultaneously calibrated to encourage efficient use *and* designed to help offset the costs of occasional company losses. For example, if the company pays out an average of \$20 per dispute (perhaps because the average payout is \$400, and the success rate is just 5%), then the company could elect to charge the customer \$20 to use the platform.<sup>173</sup> This cost lever would provide the company with considerable flexibility to adapt to different situations. For example, suppose the company begins to see an influx of fraudulent claims that try to take advantage of favorable win-lose economics. In that situation the company could simply raise the cost of submitting a claim to the platform, to the point where the economics are no longer broken. On the other hand, suppose the company has reason to believe it has become vulnerable to a class action suit. In that situation, perhaps the company not only makes the platform free, but emails its customers offering to *pay* them to resolve their disputes on the platform. This would serve the dual purpose of shrinking the size of the potential putative class, and could even serve as a basis for arguing against class certification.<sup>174</sup> For one thing, the availability of this mechanism could show that a class action is not the superior mechanism for resolving the underlying claims.<sup>175</sup> For another, encouraging some large number of potential class members to resolve their claims on the platform could be a way of demonstrating that, in fact, some significant portion of the would-be putative class was actually not injured<sup>176</sup>—something that the AI dispute-resolution process may be able to provide aggregated evidence of.

Apart from these reasons, companies may choose to submit to dispute resolution through the platform in order to improve their relationship with their customers, who might reward the opportunity to

172. Recall our estimate of a proposed 10% error rate, *supra* Section II.A.

173. We're simplifying the math once again—*see supra* note 74—but this has no real effect on the central point.

174. *See* Shawn R. Obi, William G. Fox & Ashley Wright, *Class Actions 101: Defeating Motions for Class Certification in Rule 23(b) Cases*, WINSTON & STRAWN LLP: CLASS ACTION INSIDER (Feb. 8, 2024), <https://www.winston.com/en/blogs-and-podcasts/class-action-insider/defeating-motions-for-class-certification-in-rule-23b-cases> [<https://perma.cc/UEP4-HT4L>] (“Plaintiffs must show that a class action is superior to other ways of resolving the claims at issue. Rule 23(b)(3) provides a non-exhaustive list of factors courts consider in this analysis. Generally speaking, a ‘class action is superior when it allows for the vindication of the rights of groups of people who individually would be without effective strength to bring their opponents into court at all.’” (citation omitted) (quoting *Menocal v. GEO Group, Inc.*, 882 F.3d 905, 915 (10th Cir. 2018))).

175. *Id.*

176. *One-on-One Interview: How To Defeat Class Certification*, CORP. DISPS., July–Sept. 2019, at 2 (“Thinking creatively and developing key economic and factual evidence to establish that not all class members could have suffered injury [can be critical] in establishing the impropriety of class certification . . .” (interviewing Kirkland & Ellis LLP litigation partner Daniel E. Laytin)).

affordably redress grievances through continued business. Customers need not actually *use* the platform to be appreciative of or comforted by having the option, in the same way that many customers appreciate money-back guarantees even though they may be very unlikely to use them.<sup>177</sup> Furthermore, it's also possible that offering a platform like this could reduce the amount of customer-support traffic a company experiences—traffic often generated by a company's most-aggrieved customers, who would presumably *also* be the most likely to be interested in a fast and effective dispute-resolution process that could redress their grievances. This could result in additional cost savings for the company.

Finally, one last piece of evidence suggests companies would be very happy to use a platform like this one: the voluntary inclusion by companies of “friendly” arbitration terms in their consumer agreements.<sup>178</sup> Driven by a desire to avoid class actions, many companies began including consumer-friendly provisions that served to incentivize consumers to actually arbitrate their claims.<sup>179</sup> Among these provisions, many companies included clauses that would actually front the costs of arbitration, enabling their consumers to submit a claim to arbitration without having to pay anything on the front end.<sup>180</sup> But in some cases these friendly provisions opened the door to costly mass arbitration: Plaintiffs’ attorneys could amass thousands of clients and then threaten to have those clients file arbitration claims all at once, forcing the companies to front millions of dollars’ worth of legal fees unless they agreed to a settlement.<sup>181</sup> (As one scholar put things: “With their ‘friendly’ arbitration agreements, corporate defendants may well have been hoisted by their own petards.”<sup>182</sup>) By contrast, the AI platform this Article proposes has the potential to offer all the benefits of friendly

---

177. See Thomas Suwelack, Jens Hogleve & Wayne D. Hoyer, *Understanding Money-Back Guarantees: Cognitive, Affective, and Behavioral Outcomes*, 87 J. RETAILING 462, 474 (“MBGs [money-back guarantees] reduce consumer’s risk perceptions, increase their emotional response, and enhance their WTP [willingness to pay a price premium] and purchase intentions, indicating that retailers may sell more units and demand higher prices when offering MBGs. These results hold for various types of goods, including search goods, which justifies the broad use of MBGs in retailing and marketing practice and should even stimulate further implementations.”).

178. Ryan Miller, *Next-Gen Arbitration: An Empirical Study of How Arbitration Agreements in Consumer Form Contracts Have Changed After Concepcion and American Express*, 32 GEO. J. LEGAL ETHICS 793, 799 (2019).

179. *Id.*

180. *Id.*

181. See *supra* note 55.

182. See Glover, *supra* note 21, at 1364 (referencing a comment by a federal district judge in California in a case involving the company Intuit, Transcript of Proceedings at 10, *In re Intuit Free File Litig.*, 2021 WL 834253 (No. 19-cv-02546), ECF No. 206 (“I did think when I looked at this, and saw that, really, that this was a way to avoid or otherwise circumscribe arbitration, that it seemed to be that Intuit was . . . hoisted by [its] own petard.”)).

arbitration provisions, without any potential for inviting costly mass arbitration. And the need for new ways of handling mass arbitration appears urgent from the company's point of view: Mass arbitrations continue to be on the rise, and companies have not yet found a way to effectively defend against them.<sup>183</sup>

## 2. Some Companies Will Agree To Use the Platform To Resolve Discrete Business-to-Business Disputes

As discussed in Part I, because of the high costs and risks associated with breach-of-contract related dispute resolution (including relational costs), many breaches often go unraised and unresolved.<sup>184</sup> But it is conceivable that some companies—and especially companies with relatively equal bargaining power—would be open to contractually agreeing to use this AI dispute-resolution platform to resolve specific classes of disputes.

The tool would offer companies several benefits that directly address the transaction costs they currently face when it comes to dealing with contract issues. First and most obviously, using the tool would give both parties cost clarity. The costs of using the platform would be relatively fixed, and the potential awards could be contractually capped. Second, resolving disputes with the tool would enable conflict resolution with minimal relational harm. Not only would using the platform require only minimal human interaction (reducing the likelihood of heated exchanges), but it would also ensure that disputes could be resolved in hours or days, rather than months or years. Taken together, this would greatly reduce the amount of human involvement and emotion needed to resolve the dispute, and help the human players stay focused on their business relationship. This is not to say that choosing to use the tool would cause no relational damage. But because of the rigid contractual limitations imposed on the tool, the parties might be less likely to view its use as a bad-faith affront to the business relationship. And indeed, the companies could memorialize “good-faith use” in a statement of mutual understanding within the agreement itself.

### *C. Why AI-Based Probabilistic Dispute Resolution Will Change Mass Dispute Resolution*

Earlier, this Article described the benefits to companies of an AI platform like this one in terms of its potential for heading off class

---

183. See Robert J. Herrington, Stephen L. Saxl & Jonathan H. Claydon, *5 Trends To Watch: 2024 Class Actions*, NAT'L L. REV. (Feb. 6, 2024), <https://natlawreview.com/article/5-trends-watch-2024-class-actions> [<https://perma.cc/6HCV-TM79>].

184. See *supra* note 60.

action litigation.<sup>185</sup> But whatever one's views on mass dispute resolution, I submit that reducing the *need* for commercial class actions would be an undisputed net benefit to society. After all, class actions are intended to solve what boils down to a transaction-costs problem: Many claims are so small and costly to litigate that they are illiquid absent collective action.<sup>186</sup> This tool could provide a remedy for those claims without the expensive involvement of lawyers and arbitrators. By directly addressing the transaction-costs problem, AI dispute-resolution tools like this proposed platform offer the potential of obviating the need for commercial class actions in many contexts. And because the average percentage payout via this platform will be much higher and happen much more quickly than in most class actions, individual consumers are likely to feel as though their grievances were genuinely redressed.

#### CONCLUSION

My main goal in writing this Article was to propose that there is a high need for a dispute-resolution tool that optimizes for efficiency, accessibility, and low transaction costs on the one hand, in exchange for a reduction in outcome perfection on the other—and that current-generation technology can meet that need via an AI-powered dispute-resolution platform. But whether or not I am correct that a platform like the one proposed—incredibly efficient, but with a known and significant error rate—has a place in the market, I suspect you'll agree that, when AI dispute-resolution tools inevitably achieve error rates of less than one percent, society will be all but forced to undergo a paradigm shift. Sooner or later, platforms like the one I propose will completely reshape the landscape of civil dispute resolution—from gutting class actions, to empowering small claimants, to reshaping corporate behavior. Accordingly, legal practitioners, academics, and judges should invest the time and effort to familiarize themselves with the AI tools available today and spend time learning how they work. Because one way or another, AI is destined to play a major role, be it supporting or starring, in the dispute resolution of the very near future.

---

185. *See supra* Section IV.B.1.

186. *See What Is a Putative Class Action?*, *supra* note 169.

\* \* \*