DISTRIBUTING RISK IN AN AGE OF AI: PROCEDURAL BAD FAITH AND AI CLAIMS HANDLING

Daniel Schwarcz[*]

Forthcoming in the Insurance Law Review,
Festschrift in Honor of Professor Kenneth Abraham

Abstract: Written in honor of Kenneth Abraham and his foundational contributions to insurance law, this Essay argues that the rise of AI-driven insurance claims handling exposes a significant gap in first-party bad faith law. Traditional bad faith doctrine has focused primarily on outcomes, asking whether an insurer wrongfully denied or delayed payment of benefits owed under the policy. But increasingly automated claims processes create a distinct procedural injury when insurers deny, reduce, or delay claims without meaningful human review, adequate explanation, or a genuine opportunity for the insured to be heard. Drawing on procedural justice theory, the Essay shows that such practices can undermine voice, dignity, neutrality, and trustworthiness in a relationship defined by vulnerability and dependence. It therefore argues that courts should give procedural fairness substantially greater weight within the bad faith inquiry and should treat heavily automated claims denials without meaningful human oversight as powerful evidence of procedural bad faith. Doing so would adapt bad faith law to the distinctive risks posed by AI while preserving insurance's core promise of fair, respectful, and accountable claims resolution.

INTRODUCTION

Although many people have shaped my identity as a scholar, none has

[1] [*] Schwarcz is the Fredrikson & Byron Professor of Law, University of Minnesota Law. For helpful comments, we thank Kenneth Abraham, Rebecca Hollander-Blumoff, Adam Littlestone-Luria, Jennifer Oliva, and participants in the 2026 Festschrift for Kenneth Abraham at UConn Law School.

influenced me more profoundly than Professor Kenneth Abraham. When I began my post-law-school fellowship, I knew that I wanted to write in insurance law, but I had spent months searching unsuccessfully for a topic that might plausibly launch my career. Then a mentor gave me a simple piece of advice: identify the leading scholar in the field, read everything that person had written, and figure out where they were wrong.

There was no serious question that Abraham was the scholar I should study. So I spent weeks reading his work from across the preceding decades. Much of it was deeply persuasive. But I also found the point of disagreement that would become the basis for my first article. In *Distributing Risk: Insurance, Legal Theory, and Public Policy,*[1] as well as in a related law review article,[2] Ken had expressed skepticism about aggressive versions of the reasonable expectations doctrine in insurance law.[3] He argued that courts should only invoke the doctrine to disregard unambiguous policy terms if the insurer had created, encouraged, or knowingly failed to correct a misleading impression held by the insured. I took a different view. I favored a more muscular role for courts in policing excessively one-sided policy language, ultimately arguing that courts should refuse to enforce "defective" insurance policy terms that failed to advance an insurer's legitimate purposes in the context of a particular coverage dispute.[4]

Having never met Ken, I sent him the article without any real expectation that he would respond. To my surprise, he did not merely reply. He responded substantively, thoughtfully, and remarkably quickly. Predictably, he was skeptical of some of my arguments. But I was deeply heartened that one of the most eminent law professors in the country would take so seriously the work of someone only a few years out of law school.

More than two decades have now passed since that initial exercise in reading and scrutinizing Ken's writings. During that time, I have often found it difficult to distinguish which parts of my understanding of insurance law and regulation derive from Ken's foundational contributions, and which parts are genuinely my own. Like many intellectual giants, Ken has shaped the field so deeply that his ideas now operate almost like the air we breathe.

For me, that influence has been especially profound not only because Ken's writings provided my initial education in insurance law, but also because of one

[1] Kenneth S. Abraham, Distributing Risk: Insurance, Legal Theory, and Public Policy ch. 5 (1986).

[2] Kenneth S. Abraham, Judge-Made Law and Judge-Made Insurance: Honoring the Reasonable Expectations of the Insured, 67 Va. L. Rev. 1151 (1981).

[3] For background on the doctrine, see Robert E. Keeton, Insurance Law Rights at Variance with Policy Provisions, 83 Harv. L. Rev. 961, 967 (1970); Jeffrey Stempel, Unmet Expectations: Undue Restriction of the Reasonable Expectations Approach and the Misleading Mythology of Judicial Role, 5 Conn. Ins. L.J. 181, 182 (1998).

[4] Daniel Schwarcz, A Products Liability Theory for the Judicial Regulation of Insurance Policies, 48 Wm. & Mary L. Rev. 1389 (2007).

of the most important milestones in my academic career. Less than a decade after my first exchange with him, Ken invited me to become a co-author of his insurance law casebook.[5] Since then, as I have revised and updated that book, I have continued to refine my own views on nearly every facet of insurance law and regulation in conversation with ideas that Ken had articulated long before I entered the field.

Given this history, it should not have surprised me that I had so much trouble repeating the exercise that launched my career: identifying where I think Ken is wrong and explaining why. As I reread *Distributing Risk* for the occasion of its fortieth anniversary, I found remarkably little in its core analysis that I wanted to reject. The more promising path, I came to see, was not to argue that Ken's framework was mistaken, but to identify places where developments since the book's publication had made it incomplete. My focus would be Artificial Intelligence (AI), a topic that has increasingly moved to the center of my scholarly agenda and that was nowhere on law professors' radars forty years ago.[6]

That approach quickly led me to my target: claims handling and first-party bad faith. Like many later treatments of these subjects, *Distributing Risk* understood the central danger in claims handling to be that insurers will unfairly delay or deny payment by construing their obligations too narrowly. Bad faith law, on this account, operates principally as a means of protecting insureds from those unfair outcomes. Even an insurer's failure to conduct a reasonable investigation is troubling primarily because of what it may produce: the insured may not receive, in a timely manner, the benefits to which she is entitled. That outcome-centered vision of first-party bad faith continues to dominate state law.

This Essay argues that the rise of AI in insurance claims handling is bringing into view a distinct objective for insurance law generally, and bad faith law in particular: ensuring the procedural fairness of the claims process. On this view, the fairness of the process by which an insurer investigates, evaluates, and communicates about a claim matters not only because it may affect the accuracy or timeliness of the coverage decision, but because it is itself part of what insureds are entitled to receive from their insurers. For this reason, it suggests that courts should give greater weight to procedural considerations within the

[5] See Kenneth S. Abraham & Daniel Schwarcz, Insurance Law and Regulation (8th ed. 2025).

[6] See, e.g., Anya Prince & Daniel Schwarcz, Proxy Discrimination in the Age of Artificial Intelligence and Big Data, 105 IOWA L. REV. 1257 (2020); Daniel Schwarcz, Tom Baker, & Kyle Logue, Regulating Robo-Advisors in an Age of Generative Artificial Intelligence, 82 WASH. & LEE L. REV. 775 (2025); Daniel Schwarcz et al., AI-Powered Lawyering: AI Reasoning Models, Retrieval Augmented Generation, and the Future of Legal Practice, 3 J. LAW & EMPIRICAL ANALYSIS 220 (2026); Daniel Schwarcz & Josephine Wolff, The Limits of Regulating AI Safety Through Liability and Insurance: Lessons From Cybersecurity, FORDHAM L. REV. (forthcoming, 2026).

broader bad faith inquiry than the current outcome-centered doctrine allows, and that they should treat an insurer's reliance on AI-driven claims handling without meaningful human oversight as a new and distinctive form of procedural bad faith.

Part I develops the doctrinal backdrop for this argument by showing that first-party bad faith law has long been organized around an outcome-oriented conception of insurer misconduct. That orientation is grounded in the logic that *Distributing Risk* helped articulate and develop, and is today reflected most clearly in caselaw requiring an insured to show, at a minimum, that coverage was wrongfully denied before a bad faith claim can prevail.

Part II then explains how the rise of AI in claims handling exposes a distinctly process-oriented concern. By raising the prospect that human beings may be largely or entirely removed from claims handling determinations, AI creates a form of harm that is independent of whether the insurer's decision is accurate or timely. That harm, most publicly visible to date in the health insurance context, is that an insured may be denied coverage without any appropriately trained human meaningfully scrutinizing the claim or communicating with the insured about it.

Drawing on the procedural justice literature, Part III explains why this deprivation can produce real harms independent of the correctness of the insurer's ultimate coverage determination. Across a broad range of settings, a well-established empirical literature shows that people value fair process apart from whether they receive a favorable outcome, and that they assess the legitimacy of decision-making systems accordingly. Removing any human from the process of claims denial, Part III argues, is fundamentally incompatible with these broad-based conceptions of procedural fairness, an injury that lands with particular force because insureds encounter the claims process at a moment of acute vulnerability.

Part IV concludes by arguing that this insight generates a regulatory imperative: insurance law should require meaningful human review and communication before insurers take actions that disadvantage their policyholders. Market forces are unlikely to supply these protections on their own, both because the benefits of fair claims handling have a collective-good character for the industry and because claims processing is subject to the familiar information asymmetries and incentive distortions that have long justified existing bad faith law. Given these realities, Part IV weighs the substantial benefits and costs of expanding bad faith doctrine, ultimately suggesting elevating procedural fairness within the existing bad faith inquiry and recognizing AI-driven claims handling without meaningful human oversight as an emerging form of procedural bad faith.

I. THE OUTCOME-ORIENTED FOCUS OF FIRST-PARTY INSURANCE BAD FAITH LAW

The primary purpose of bad faith law has traditionally been framed in outcome-oriented terms: preventing the undue delay or wrongful denial of legitimate claims. As Professor Abraham explained decades ago, this risk is inherent in most insurance relationships.[7] Under ordinary contract law principles, an insurer that wrongfully denies a claim faces little downside beyond being ordered to pay what it owed in the first place. The result is a one-sided gamble: if the insurer refuses to pay and the insured does not pursue the claim, the insurer keeps the money; if the insured sues and prevails, the insurer generally pays only what it should have paid all along. This problem is especially acute because insureds have great difficulty assessing claims-handling reliability in advance,[8] cannot practically purchase contractual assurances of fair claims handling,[9] and often lack the expertise needed to evaluate ex post whether a denial is justified.[10]

First-party bad faith law developed as a corrective to these dynamics. By authorizing remedies beyond the insurer's underlying coverage obligations, bad faith law seeks to counteract the insurer's incentive to delay or deny payment without adequate justification.[11] Consistent with this outcome-oriented account, the dominant tests for bad faith focus on identifying coverage decisions that were not merely incorrect, but unreasonable in a way the insurer either knew or should have known.[12] One common formulation asks whether the insurer refused to pay a claim even though its obligation to provide coverage was not fairly debatable.[13] Many states also require evidence that the insurer was subjectively aware that it lacked a reasonable basis for denying the claim, or at least recklessly disregarded that possibility.[14]

---

[7] See Distributing Risk, supra. See also Jay M. Feinman, Delay, Deny, Defend: Why Insurance Companies Don't Pay Claims and What You Can Do About It (2010). But see Alan O. Sykes, "Bad Faith" Breach of Contract by First-Party Insurers, 25 J. LEGAL STUD. 405 (1996).

[8] See Daniel Schwarcz, Transparently Opaque: Understanding the Lack of Transparency in Insurance Consumer Protection, 61 UCLA L. REV. 394 (2014).

[9] See Distributing Risk, supra.

[10] Tom Baker, Sales Stories, Claims Stories, and Insurance Damages, 72 TEX. L. REV. 1395 (1994).

[11] Jay M. Feinman, The Law of Insurance Claim Practices: Beyond Bad Faith, 47 Tort Trial & Ins. Prac. L.J. 693 (2012).

[12] See Abraham & Schwarcz, supra, at 87-88.

[13] See Dakota, Minn. & E. R.R. Corp. v. Acuity, 771 N.W.2d 623 (S.D. 2009); LeRette v. Am. Med. Sec., Inc., 705 N.W.2d 41 (Neb. 2005). Cf. Home Loan Inv. Co. v. St. Paul Mercury Ins. Co., 827 F.3d 1256 (10th Cir. 2016).

[14] See Steven S. Ashley, Bad Faith Actions Liability & 88 General Principles Of Insurance Law Chapter 2 Damages § 5:2 (2d ed. 2018).

To be sure, courts have increasingly recognized that the insurer's claims-handling process is also relevant to the bad faith inquiry. A growing number of state courts now distinguish between substantive and procedural bad faith.[15] Whereas substantive bad faith concerns the insurer's ultimate coverage determination, procedural bad faith focuses on the manner in which the insurer handled the claim, especially the adequacy and promptness of its investigation. In some cases, procedural bad faith may also take other forms, including failing to appropriately coordinate insurance payouts,[16] unfounded insurer accusations of fraud,[17] or discriminatory claims-handling practices that disproportionately burden insureds based on legally suspect characteristics, such as race or ethnicity.[18]

Even so, most states treat procedural claims handling deficiencies as secondary to substantive errors when considering allegations of bad faith. That hierarchy is most evident in the rule, followed in almost every jurisdiction, that a plaintiff cannot recover for procedural bad faith alone. At a minimum, most jurisdictions require the insured to show that the insurer's coverage determination was incorrect, even if that error does not itself amount to substantive bad faith.[19] Other jurisdictions go further, treating substantive bad faith as the core of the cause of action and allowing procedural deficiencies to operate only as a plus factor, or as evidence that the insurer knew, or recklessly disregarded, that it lacked a reasonable basis for refusing to pay. [20] By contrast, only a small handful of jurisdictions appear to recognize procedural bad faith as independently actionable, regardless of whether the insurer's ultimate coverage determination was correct.[21]

---

[15] See, e.g., Walker v. Life Ins. Co. of N. Am., 59 F.4th 1176 (11th Cir. 2023); Bellville v. Farm Bureau Mut. Ins. Co., 702 N.W.2d 468 (Iowa 2005); Capstone Bldg. Corp. v. American Motorists Ins. Co., 308 Conn. 760 (2013).

[16] See Silberg v. Cal. Life Ins. Co., 521 P.2d 1103 (Cal. 1974) (finding that health insurer acted in bad faith when it failed to make an initial payment of benefits while a workers' compensation dispute was pending, as the insurer could have recouped its payment if workers' compensation later covered the loss)

[17] See, e.g., Riley v. Travelers Home & Marine Ins. Co., 214 A.3d 345 (Conn. 2019) (affirming a jury's finding that Travelers owed a homeowner $1.5 million in breach-of-contract and emotional distress damages after it accused him of trying to burn down his home, despite a fire official's conclusion that the blaze was accidental).

[18] Huskey v. State Farm Fire & Cas. Co., No. 22 C 7014, 2023 WL 5848164, at *1 (N.D. Ill. Sept. 11, 2023) (alleging that State Farm engaged in discriminatory claims handling)

[19] See Walker v. Life Ins. Co. of N. Am., 59 F.4th 1176 (11th Cir. 2023).

[20] Bellville v. Farm Bureau Mut. Ins. Co., 702 N.W.2d 468 (Iowa 2005)

[21] Coventry Associates v. American States Ins. Co., 136 Wash.2d 269 (1998) ("an insured may maintain an action against its insurer for bad faith investigation of the insured's claim and violation of the [Consumer Protection Act] regardless of whether the insurer was ultimately correct in determining that coverage did not exist."); Silberg v. Cal. Life Ins. Co., 521 P.2d 1103 (Cal. 1974) (finding that health insurer acted in bad faith when it failed to make an initial payment of benefits while a workers' compensation dispute was pending, as the insurer could

II. THE RISE OF AI-POWERED CLAIMS HANDLING

The outcome-oriented focus of traditional insurance bad faith law sits uneasily with the rapid transformation of insurers' claims-handling processes in recent years. Increasingly, insurers rely on both traditional algorithms and machine-learning tools to structure, support, and sometimes effectively determine how claims are evaluated, processed, and communicated to policyholders.

Most insurers maintain that these algorithmic tools merely assist human reviewers by making traditional, human-centered claims handling more efficient, consistent, or accurate. But there are growing indications that, in at least some settings, automated systems are doing more than supporting human judgment. They are effectively displacing it, with human reviewers providing only cursory oversight of decisions that have already been generated, structured, or predetermined by algorithmic tools.

The clearest evidence of this shift comes from health insurance. Health insurers have long relied on traditional algorithms that encode human-defined rules to process claims and prior-authorization requests; given the enormous volume of claims they receive, some degree of automation is practically unavoidable.[22] Cigna's PXDX system, for example, reportedly flags claims for review when it detects a mismatch between a diagnosis code and a treatment code.[23] Increasingly, moreover, health insurers appear to be supplementing these rule-based systems with machine-learning tools that evaluate claims or predict a patient's expected course of care. UnitedHealthcare and Humana, for instance, use a machine-learning tool called nH Predict to help assess the medical necessity of post-acute care for individual patients.[24]

But emerging evidence suggests that some health insurers are using these algorithmic tools not merely to inform human claims judgment, but to substitute for it.[25] A high-profile ProPublica investigation reported that Cigna physicians

---

have recouped its payment if workers' compensation later covered the loss).

[22] See generally Jennifer D. Oliva, Regulating Healthcare Coverage Algorithms in the Shadow of ERISA, 125 Mich. L. Rev. __ (forthcoming 2027). Traditional algorithms typically apply human-defined rules to inputs, producing outputs according to logic specified in advance. Machine-learning systems, by contrast, are trained on data to identify patterns that can then be used to generate predictions, classifications, or recommendations in new cases. See Anya Prince & Daniel Schwarcz, Proxy Discrimination in the Age of Artificial Intelligence and Big Data, 105 IOWA L. REV. 1257 (2020).

[23] See Patrick Rucker, Maya Miller & David Armstrong, How Cigna Saves Millions by Having Its Doctors Reject Claims Without Reading Them, ProPublica (Mar. 25, 2023).

[24] See Casey Ross & Bob Herman, Denied by AI: How Medicare Advantage Plans Use Algorithms to Cut Off Care for Seniors in Need, STAT (Mar. 13, 2023); Oliva, supra.

[25] See Oliva, supra.

approved large batches of denials generated by the PXDX system without reviewing individual patient files. According to the report, Cigna denied more than 300,000 payment requests over a two-month period using this process, with physicians spending an average of only 1.2 seconds on each claim. That kind of review is difficult to square with the idea that the algorithm is simply assisting individualized human judgment. After all, the fact that a diagnosis code and treatment code do not correspond does not necessarily mean that a claim should be denied as medically unnecessary. Particular patients may have circumstances that make an otherwise unusual treatment appropriate, and those circumstances can be missed when human review is reduced to rapid approval of algorithmically flagged denials.[26]

Similar concerns arise from reporting and related lawsuits involving UnitedHealthcare and Humana, which allege that those insurers deployed the nH Predict algorithm not merely to support human assessments of the medical necessity of post-acute care, but to replace those assessments. According to those sources, nH Predict generates an "optimal" length of stay for skilled nursing facility or inpatient rehabilitation care based on variables such as the patient's diagnosis, age, living situation, and functional status. The insurers allegedly then used that prediction as a practical coverage limit, terminating care once the algorithm's projected number of days had elapsed without adequate regard to the patient's actual recovery, complications, comorbidities, home-based supports, or the judgment of treating physicians. Although human reviewers may remain formally involved, their role was allegedly highly constrained; plaintiffs in one suit, for example, alleged that employees who deviated from nH Predict's projected care limits were disciplined or terminated, and that roughly ninety percent of the resulting denials were later reversed on appeal.[27]

Like health insurers, property/casualty insurers are increasingly deploying AI and algorithmic systems in claims handling.[28] Property insurers have long used traditional algorithmic tools, such as Xactimate, to estimate repair costs and help determine appropriate claim payments.[29] But insurers are now using newer machine-learning tools across a much broader range of claims-handling

---

[26] See Amy Monahan & Daniel Schwarcz, Rules of Medical Necessity, 107 IOWA L. REV. 423 (2022).

[27] See Oliva, supra; Complaint, Estate of Lokken v. UnitedHealth Grp., Inc., No. 0:23-cv-03514 (D. Minn. filed Nov. 14, 2023).

[28] A 2023 NAIC survey found that 70 percent of U.S. property insurers were using AI or interested in using it, and that 88 percent of auto insurers were using or planning to use it. See Nat'l Ass'n of Ins. Comm'rs, Big Data & Artificial Intelligence (H) Working Group, AI/ML Survey Results, https://content.naic.org/sites/default/files/national_meeting/Big-Data0813-AttB081223_0.pdf

[29] See Kenneth S. Klein, Minding the Protection Gap: Resolving Unintended, Pervasive, Profound Homeowner Underinsurance, 25 Conn. Ins. L.J. 34 (2018).

functions. For instance, many insurers use AI models to flag claims for potential fraud.[30] They also increasingly use computer-vision tools to estimate auto and property damage from photographs, recommend whether damaged items should be repaired or replaced, generate line-level repair estimates, and support total-loss determinations.[31] Perhaps most troublingly, insurers are increasingly using generative AI to summarize claim files and draft or structure communications with insureds.[32] Across these use cases, insurers often rely on third-party vendors to supply the underlying technologies rather than developing the systems entirely in house.[33]

Property/casualty insurers typically characterize these systems as tools that support, rather than replace, human judgment when it comes to claims denials or other adverse insurance actions. On this account, fully automated claims handling is largely confined to fast-tracking straightforward approvals. The Insurtech Lemonade, for example, reported that roughly 55 percent of its claims approvals were fully automated by the end of 2025.[34]

By contrast, in the settings where AI use could most plausibly harm policyholders, such as by contributing to a denial, underpayment, or fraud investigation, property/casualty insurers generally insist that AI merely assists trained human adjusters who retain ultimate decision-making authority.[35] But that assurance leaves a critical question unresolved: whether the human "in the

---

[30] See Shift Technology, Shift Technology Helps Hundreds of Insurers Fight Claims Fraud Using Azure OpenAI Service, https://www.shift-technology.com/resources/case-studies/shift-technology-helps-insurers-fight-claims-fraud-using-azure-openai-service (last visited May 21, 2026); see also Huskey v. State Farm, (N.D. Ill., filed 2022) (alleging that State Farm's automated claims processing methods and machine-learning algorithms use proxies for race when flagging claims for fraud investigations).

[31] Yevheniya Kobets et al., Automated Car Damage Assessment Using Computer Vision: Insurance Company Use Case, 14 Applied Sci. 9560 (2024); CCC Intelligent Solutions, AI Technology, https://www.cccis.com/our-technology/ai (last visited May 21, 2026); Tractable, GEICO Partners with Tractable to Accelerate Accident Recovery with AI, https://tractable.ai/geico-partners-with-tractable-to-accelerate-accident-recovery-with-ai/ (last visited May 21, 2026).

[32] See Swiss Re, How Generative AI Is Transforming Insurance Claims: Inside Swiss Re's ClaimsGenAI, https://www.swissre.com/risk-knowledge/advancing-societal-benefits-digitalisation/how-generative-ai-is-transforming-insurance-claims-claimsgenai.html (last visited May 21, 2026); *Bain & Company, The $100 Billion Opportunity for Generative AI in P&C Claims Handling,* https://www.bain.com/insights/100-billion-dollar-opportunity-for-generative-ai-in-p-and-c-claims-handling/ (last visited May 21, 2026).

[33] Nat'l Ass'n of Ins. Comm'rs, Use of Artificial Intelligence Systems by Insurers (Model Bulletin, adopted Dec. 4, 2023), https://content.naic.org/sites/default/files/cmte-h-big-data-artificial-intelligence-wg-ai-model-bulletin.pdf.pdf (last visited May 21, 2026).

[34] Lemonade, Inc., Annual Report (Form 10-K) (filed Feb. 2026), available at https://www.stocktitan.net/sec-filings/LMND/10-k-lemonade-inc-files-annual-report-33aac5b74a32.html (last visited May 21, 2026).

[35]

loop" is meaningfully exercising independent judgment or instead largely ratifying outputs that have already been generated, framed, or effectively predetermined by algorithmic systems, as apparently is occurring in the health insurance domain.[36] Because property/casualty insurers and vendors disclose little about how these tools are actually integrated into claims workflows, the degree of genuine human oversight remains difficult to assess.

There is also good reason to doubt that inserting a human reviewer after an algorithm has spoken restores genuinely independent judgment. Decades of behavioral research show that an existing recommendation reshapes the cognitive task of the person reviewing it. Anchoring effects pull the reviewer's assessment toward the algorithm's output even when the reviewer consciously discounts it.[37] Status quo bias makes ratifying the machine's determination the path of least resistance, since overriding it requires the reviewer to generate affirmative reasons, assume personal responsibility for the deviation, and often document the departure in ways that approval does not require.[38] And a distinct literature on automation bias finds that humans supervising automated systems systematically over-trust machine outputs, catching fewer errors than they would working unaided.[39] These dynamics operate even absent institutional pressure, but claims organizations add such pressure: reviewers face throughput targets and, as the nH Predict allegations illustrate, may face discipline for departing from algorithmic projections. The upshot is that human review of a completed algorithmic determination is a structurally weaker safeguard than human judgment exercised in the first instance. The question regulators and courts must eventually confront is therefore not simply whether a human was in the loop, but whether the loop was designed so that the human could realistically do anything other than agree.

III. THE PROCEDURAL FAIRNESS PROBLEMS WITH AI-POWERED CLAIMS HANDLING

The growing use of algorithms and machine-learning AI systems in claims handling has generated substantial backlash. That backlash has been especially pronounced in health insurance, where several states have recently enacted laws requiring human physicians to make certain coverage determinations.[40] But

[36] See generally Crootof, Rebecca, et al. "Humans in the Loop." Vand. L. Rev. 76 (2023): 429.

[37] See Amos Tversky & Daniel Kahneman, Judgment Under Uncertainty: Heuristics and Biases, 185 Science 1124, 1128–30 (1974).

[38] See William Samuelson & Richard Zeckhauser, Status Quo Bias in Decision Making, 1 J. Risk & Uncertainty 7 (1988).

[39] See Ben Green, The Flaws of Policies Requiring Human Oversight of Government Algorithms, 45 Computer L. & Sec. Rev. 105681 (2022); Crootof et al., supra.

[40] See, e.g., California — SB 1120 (the "Physicians Make Decisions Act"), effective January

even in the property/casualty setting, where the evidence that insurers are fully offloading claims denials to AI or algorithmic systems is more limited, scrutiny has been intense. One prominent lawsuit, for instance, alleges that State Farm's fraud-detection algorithm used proxies for race and other suspect characteristics to flag claims for fraud investigation.[41] And many state regulators have issued bulletins emphasizing that insurers must maintain governance systems sufficient to ensure that their use of AI complies with generally applicable claims-handling rules,[42] including prohibitions on "refusing to pay claims without conducting a reasonable investigation," or "failing to adopt and implement reasonable standards for the prompt investigation and settlement of claims."[43]

Much of this backlash can surely be attributed to the risk that AI systems will produce substantively bad outcomes, including biased, inaccurate, or unfair claim denials. But that is only part of the story. To see why, consider the well-developed procedural justice literature. The central figure in this literature is Tom Tyler, who has spent much of his career developing and synthesizing evidence showing that people's compliance with, and acceptance of, legal authority depends heavily on their assessments of procedural fairness, not merely on whether they receive favorable outcomes.[44]

Building on these insights, a broad body of research has shown that the "fair process effect" Tyler documented extends across a wide range of settings, including criminal and civil courts, workplaces, schools, families, and police

---

1, 2025; Nebraska — LB 77 (2025); Texas — SB 815 (2025). Washington — (signed March 2026, effective June 2026). AI algorithms can only approve prior authorizations — not deny them — without a health professional's review. See also CTR. FOR MEDICARE ADVOC., THE ROLE OF AI-POWERED DECISION-MAKING TECHNOLOGY IN MEDICARE COVERAGE DECISIONS 1 (Jan. 2022), https://medicareadvocacy.org/wp-content/uploads/2022/01/AI-Tools-In-Medicare.pdf. These reforms, however, leave major gaps. They generally say little about how much a human reviewer may rely on algorithmic recommendations before the review becomes merely nominal. Nor do they specify when algorithmic influence crosses the line from permissible support to impermissible substitution for human judgment. Even more importantly, these reforms are largely confined to health insurance and, because of ERISA preemption, reach only a limited portion of that market. See Oliva, supra.

[41] Huskey v. State Farm, N.D. Ill., filed 2022).

[42] See NAIC Model Bulletin, supra.

[43] Unfair Claims Settlement Practices Act § 4 (Nat'l Ass'n of Ins. Comm'rs, Model 900) (quoted language). Several states have gone further, imposing affirmative obligations on insurers to audit algorithmic and AI systems, including those used in claims handling, for bias, error, and other risks. Colorado's Regulation 10-1-1, originally adopted for life insurers in September 2023, was amended on October 15, 2025, to extend to private passenger auto and health-benefit-plan insurers, with full compliance required by July 1, 2026, and annual narrative compliance reports thereafter. These implement SB 21-169 and the original AI Act (SB 24-205, signed May 17, 2024). See also New York DFS Insurance Circular Letter No. 7 (2024), issued July 11, 2024.

[44] See Tom Tyler, Why People Obey the Law (1990); see also John Thibaut & Laurens Walker, Procedural Justice: A Psychological Analysis (1975).

encounters.[45] This research shows that people value fair process independently of outcome, and that perceptions of fair process lead to broader perceptions of the legitimacy of the system as well as acceptance of even adverse outcomes.[46] It also shows that they evaluate procedural fairness by reference to four recurring considerations: voice, dignity and respect, neutrality and freedom from bias, and trustworthiness of the decision-maker's motives.[47] Over time, the dominant theoretical explanation for why process matters so much has shifted from an instrumental account, in which fair procedures matter because they tend to produce fair outcomes, toward relational accounts, particularly the group-value and group-engagement models.[48] On those accounts, procedural treatment matters because it sends identity-relevant signals about a person's standing, status, and value within the relevant community.[49]

It is largely for these reasons that most commentators have the intuition that AI systems will not replace human judges, even as growing evidence suggests that they may be able to approximate judicial outcomes in some settings.[50] Recent experimental work bears out this intuition, finding that public perceptions of the legitimacy, fairness, and trustworthiness of judicial decision-making decline as the role of AI in the decision grows.[51] Similar procedural fairness concerns help to explain broad-based intuitions that AI should not officially replace human grading in most educational contexts any time soon, notwithstanding evidence that AI can grade nearly as well as, and sometimes better than, human evaluators across a wide range of tasks.[52]

In both settings, AI adjudication is fundamentally at odds with fair process. It threatens voice and the opportunity to be heard, which depend not merely on the ability to submit information, but on the experience of presenting one's case to a human decision-maker capable of receiving, weighing, and responding to individual expression in a meaningful way. It also sits uneasily with the dignitary dimension of procedural fairness, which requires affected individuals to be treated with courtesy and respect. That value is difficult to realize through

[45] See Rebecca Hollander-Blumoff, Procedural Justice Injuries (Draft, 2026); Rebecca Hollander-Blumoff & Tom R. Tyler, Procedural Justice in Negotiation: Procedural Fairness, Outcome Acceptance, and Integrative Potential, 33 L. & Soc. Inquiry 473 (2008).

[46] See Tyler, supra.

[47] See Hollander-Blumoff, supra; see also Rebecca Hollander-Blumoff, Formation of Procedural Justice Judgments in Legal Negotiation, 26 GROUP DECISION & NEGOT. 19, 35–36 (2017)

[48] See Hollander-Blumoff, supra.

[49] See id.

[50] See Eric Posner, Judge AI

[51] See Anna Fine, Emily R. Berthelot & Shawn Marsh, Public Perceptions of Judges' Use of AI Tools in Courtroom Decision-Making: An Examination of Legitimacy, Fairness, Trust, and Procedural Justice, 15 Behav. Sci. 476 (2025).

[52] See Scott Hirst, et al, Grading Machines: Can AI Exam-Grading Replace Law Professors, 3 J. LAW & EMPIRICAL ANALYSIS 2 (2026)

a process mediated by a machine that cannot recognize, much less honor, the human significance of the interests at stake. AI adjudication also raises serious concerns about neutrality and freedom from bias, given well-documented risks of bias embedded in training data, system design, and opaque model outputs. Perhaps most fundamentally, AI cannot reliably convey trustworthy motives.

The same logic applies to the use of algorithms, and especially black-box machine-learning systems, to limit or delay insurance recovery. To be sure, many insureds may begin with relatively modest expectations about the fairness of the claims process. They may doubt, for instance, that human adjusters are entirely neutral or free from institutional pressures. But human adjusters can nonetheless promote a sense of procedural fairness by giving insureds a meaningful opportunity to tell their side of the story, treating them with respect, explaining the basis for their decisions, endeavoring to act impartially, and communicating a genuine commitment to providing fair compensation for covered losses. Indeed, human adjusters are routinely instructed to embody these goals, and they often do so successfully.

For the same reasons that AI adjudication cannot achieve these objectives in judging or grading, it cannot achieve them in insurance claims adjudications that result in coverage denials.[53] Such automated coverage denials cause injury that is independent from the coverage denial itself. The injury has three components. The first is dignitary. When no human being meaningfully considers a claim before it is denied, the insured has been judged without being seen. The procedural justice literature teaches that treatment of this kind communicates something about that person's standing: that the insured is not the kind of person, and hers is not the kind of loss, that warrants the attention of another human being.[54] That message is inflicted at the moment of denial and is not retracted even if the denial proves correct.

The second component of the procedural harm is a deprivation of voice.

---

[53] A skeptic might respond that the incompatibility between machines and fair process is an empirical claim about present attitudes, not a conceptual truth, and that the empirical ground is shifting. Millions of people now use large language models for emotional support, companionship, and even quasi-therapeutic conversation, and many report feeling genuinely understood. Yidan Yin, Nan Jia & Cheryl J. Wakslak, AI Can Help People Feel Heard, but an AI Label Diminishes This Impact, 121 Proc. Nat'l Acad. Sci. e2319112121 (2024). But there is good reason to think that decision-making about matters like insurance coverage is different in kind from AI advice or companionship: a person who consults an AI therapist retains authority over her own life. A claims determination, by contrast, is a unilateral exercise of power over the claimant's entitlements, conducted at a moment not of her choosing, in a process she does not control. See, e.g., 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. on Disp. Resol. 91 (2018); Fine, Berthelot & Marsh, supra. See also Berkeley J. Dietvorst, Joseph P. Simmons & Cade Massey, Algorithm Aversion: People Erroneously Avoid Algorithms After Seeing Them Err, 144 J. Experimental Psychol.: Gen. 114 (2015).

[54] See Tyler, supra; Hollander-Blumoff, supra.

Voice, in the procedural justice framework, is not merely the ability to transmit information into a system. It is the experience of presenting one's circumstances to a decision-maker capable of registering their significance and responding to them. A claimant who uploads documents into a portal that feeds an algorithm has technically been permitted to speak, but she has not been heard, because there is no one on the other end doing the hearing.

Finally, automated claims denials produce relational harms. Insurance is not merely a financial product; it is an ongoing promise of protection purchased in anticipation of a future moment of need, a point insurers themselves often emphasize in their marketing. [55] When that moment arrives and the insurer responds through an automated process, it may fail to honor the relational dimension of that promise, even if it has satisfied its financial obligations. As the procedural justice literature suggests, such treatment can communicate identity-relevant messages about a person's standing, dignity, and worth. These relational harms can be particularly significant because insurers typically make claims determinations at a moment of acute vulnerability for the insured, who may be seeking recovery after illness, injury, accident, theft, fire, storm damage, or some other disruptive loss. In that setting, the process by which the insurer responds carries heightened relational significance. People attend most closely to such signals precisely when their standing feels threatened, and a serious loss is exactly such a moment.[56]

Opacity compounds each of these three procedural harms. A claimant cannot meaningfully exercise voice if she does not know what information the decision process considered, cannot assess neutrality if she does not know how the decision was reached, and cannot evaluate the decision-maker's motives if she does not know who, or what, the decision-maker was. Insurance claims handling has long been opaque along all of these dimensions, and AI deepens the opacity in two ways.[57] Machine-learning systems are famously difficult to interpret even for the firms that deploy them, and insurers treat the existence and design of these systems as proprietary, disclosing little about whether or how an algorithm influenced a given denial at all. The result is that a claimant often cannot even know whether a human considered her claim, a form of uncertainty that is itself procedurally corrosive.

---

[55] See Kenneth S. Abraham, Four Conceptions of Insurance, 161 U. Pa. L. Rev. 653 (2013); Feinman, supra.

[56] Research on procedural justice among vulnerable populations is consistent with this intuition. Studies of psychiatric patients facing involuntary commitment, for example, find that perceptions of voice, respect, and good-faith treatment powerfully shape whether patients experience the process as coercive, quite apart from the outcome. See, e.g., Charles W. Lidz et al., Perceived Coercion in Mental Hospital Admission: Pressures and Process, 52 Archives Gen. Psychiatry 1034 (1995); Tom R. Tyler, The Psychological Consequences of Judicial Procedures: Implications for Civil Commitment Hearings, 46 SMU L. Rev. 433 (1992).

[57] See Transparently Opaque, supra; supra Part II.

The resulting harm extends beyond the individual claim. Insurance depends on policyholders' confidence that, when covered losses occur, insurers will respond not only accurately, but fairly, respectfully, and accountably. Fully or largely automated claims processes threaten that confidence by making insurance appear less like a relational promise of financial protection and more like a one-sided system of algorithmic control. Over time, that perception risks eroding trust not only in particular insurers' claim decisions, but in insurance itself as a reliable institution for managing risk and providing security.

IV. Integrating Procedural Fairness into Bad Faith Law

Given the increasingly salient procedural risks posed by AI-driven claims handling, this Section suggests that courts should give procedural failures substantially greater weight in the bad-faith inquiry and should treat heavily automated denial processes as powerful evidence of procedural bad faith.[58] The precise implications of this recommendation would, of course, vary by jurisdiction, depending on each state's existing bad-faith doctrine, which differs along several important dimensions, including the role procedural misconduct plays in the overall bad-faith analysis.[59] While some states might consider following the small number of jurisdictions that allow procedural bad faith claims involving AI automation without any showing of an incorrect coverage denial, others might treat such procedural bad faith as prima facie evidence of insurer recklessness in the broader bad faith inquiry.

However any specific jurisdiction implemented these principles, a more muscular and AI-specific set of rules regarding procedural bad faith might induce insurers to preserve meaningful human responsibility for claim decisions, provide intelligible explanations for delays and denials, give insureds

---

[58] The law often recognizes that a flawed process inflicts cognizable harm even when the outcome might not change. For instance, on *Rosales-Mireles v. United States*, 585 U.S. 129 (2018), the Supreme Court held that a miscalculation of the advisory Sentencing Guidelines range ordinarily warrants correction on plain-error review even though the district court retains discretion to impose the same sentence on remand. What justified relief was not a demonstrably different outcome but the integrity of the process itself: a defendant sentenced under a process anchored by an erroneous recommendation has been treated unfairly in a way that erodes public confidence in sentencing, whatever sentence ultimately issues. The parallel to automated claims handling is instructive on two levels. Most obviously, it confirms that process-based injury can justify a legal remedy without proof of outcome causation. Less obviously, it anticipates the anchoring problem discussed in Part II. The Guidelines range matters, the Court recognized, because it frames and pulls on the exercise of human discretion that follows it. An algorithmic claims recommendation operates in the same way. Even where a human adjuster formally retains authority to depart from the algorithm's output, that output anchors the decision in a manner that makes purely formal human involvement an unreliable safeguard.

[59] See generally

a genuine opportunity to submit and contextualize information, and maintain governance systems capable of detecting bias, error, and arbitrary treatment. Over time, courts deploying this doctrine could develop context-sensitive guidance about when automated tools may permissibly assist human adjusters and when their use impermissibly displaces the relational, explanatory, and accountable features that make claims handling procedurally fair. In this way, a more robust and AI-specific doctrine of procedural bad faith could help ensure that the claims process continues to reflect one of the basic promises of insurance: that all policyholders will receive a fair, respectful, and accountable process for determining their coverage rights.

This approach, of course, faces legitimate objections. The first objection is market-based. Many insurers operate in competitive markets and therefore should be at least somewhat responsive to consumer preferences.[60] If AI-centered claim denials impose procedural harms that exceed the cost savings from automation, insurers may have incentives to structure their claims practices accordingly. Market adjustment would not only reduce the need for legal intervention, but also preserve variation across insurers' claims-handling models. Consumers could then decide for themselves how to trade off the procedural (and substantive) risks of AI-powered claims handling against the premium savings associated with more automated processes.

Although this concern has some force, it is ultimately unpersuasive for many of the same reasons that market-based objections to ordinary bad faith law are unpersuasive. Consumers have strikingly limited information about insurers' claims practices at the point of purchase, including the extent to which insurers rely on AI to limit, delay, or deny coverage.[61] That information is difficult to obtain, as reflected in conflicting accounts of how extensively insurers currently use AI in claims administration. Consumers are also likely to undervalue claims handling when buying coverage, both because of familiar cognitive biases and because future claim scenarios are hard to imagine in concrete terms.

Additionally, the market-based account overlooks a serious public-good problem. When individual insurers rely on AI systems to deny claims, they risk undermining not only their own reputations, but also broader public trust in insurance as a social and economic institution. That is because procedural fairness is deeply linked to perceptions of system legitimacy and acceptance of unfavorable outcomes.[62] This, in turn, can lead to bad outcomes: individuals may purchase inadequate insurance, forgo potentially covered expenses, or pursue hopeless claims disputes because they lack confidence and trust in the

---

[60] See Alan O. Sykes, "Bad Faith" Breach of Contract by First-Party Insurers, 25 J. Legal Stud. 405 (1996).

[61] See Daniel Schwarcz, Transparently Opaque: Understanding the Lack of Transparency in Insurance Consumer Protection, 61 UCLA L. Rev. 394 (2014).

[62] See supra Part III.

underlying insurance systems. Yet individual insurers are unlikely to fully internalize those systemic harms when deciding whether the cost savings from AI-driven claims handling justify its procedural risks.

A second, related challenge to expanding the importance and scope of procedural bad faith could increase the cost of insurance coverage, or at least limit the extent to which AI can be used to make coverage determinations more efficient. There is reasonably clear evidence of this effect in the context of first-party bad faith law more generally.[63] Moreover, an expanded law of procedural bad faith could have unintended consequences for the substantive accuracy of claims determinations. AI may often be more reliable than human judgment, particularly when it comes to domains where there are clear right and wrong answers and substantial historic data, like assessing coverage and identifying potential fraud. Even a relatively modest doctrine requiring meaningful human review could therefore reduce the accuracy gains associated with automated claims handling. In fact, human-AI collaboration often produces less reliable results than AI alone because human reviewers mistakenly discount accurate AI outputs.[64] A procedural bad faith doctrine that requires human oversight of AI outputs might thus make claims handling more procedurally legitimate while, in some cases, making it less substantively accurate.

The risk that an AI-specific doctrine of procedural bad faith could increase insurance costs is heightened by the possibility that it would make aggregate litigation against insurers easier to pursue. Traditional substantive bad faith is notoriously resistant to class treatment because the reasonableness of each coverage determination turns on individualized facts, so common questions rarely predominate. Procedural bad faith could invert that structure. Whether an insurer routed an entire category of claims through an algorithm, what role human reviewers actually played, whether reviewers had authority and incentives to depart from algorithmic outputs, and what the insurer knew about the system's error rates are questions common to every claimant whose file passed through the same pipeline. The early litigation over nH Predict and PXDX, framed as class actions on behalf of all insureds subjected to the challenged systems, reflects precisely this logic.[65] Aggregation of this kind

[63] See Mark J. Browne, Ellen S. Pryor, & Bob Puelz, The Effect of Bad-Faith Laws on First-Party Insurance Claims Decisions, 33 J. Legal Stud. 355 (2004); Danial P. Asmat & Sharon Tennyson, Does the Threat of Insurer Liability for "Bad Faith" Affect Insurance Settlements?, 81 J. Risk & Ins. 1 (2014); Sharon Tennyson & William J. Warfel, The Law and Economics of First-Party Insurance Bad Faith Liability, 16 Conn. Ins. L.J. 203 (2010).

[64] See Michelle Vaccaro, Alaa Almaatouq & Thomas W. Malone, When combinations of humans and AI are useful: A systematic review and meta-analysis, Nature Human Behaviour 8, 2293-2303 (2024).

[65] See Complaint, Estate of Lokken v. UnitedHealth Grp., Inc., No. 0:23-cv-03514 (D. Minn. filed Nov. 14, 2023); Complaint, Barrows v. Humana, Inc., No. 3:23-cv-00654 (W.D. Ky. filed Dec. 12, 2023); Complaint, Kisting-Leung v. Cigna Corp., No. 2:23-cv-01477 (E.D.

would place substantial pressure on insurance prices, threatening insurers who implemented AI-based claims handling with larger judgments, increased litigation costs, and heightened regulatory and reputational risk.

This argument, too, has some force. But many of the same considerations that justify bad faith law more generally apply here as well. Most importantly, although procedural bad faith law may marginally increase insurance costs, or prevent them from falling as much as they otherwise might, that tradeoff may be justified if it causes the insurance system to treat insureds more fairly and promotes broader public trust in insurance. Indeed, current market behavior appears consistent with this logic: no insurer seems to be actively marketing cheaper coverage on the ground that claims will be resolved through fully automated determinations. Moreover, embedding a baseline set of consumer protections into the insurance product through law and regulation can promote genuine price competition.[66] By preventing insurers from reducing costs through procedurally unfair claims practices, insurers are forced to compete on more legitimate forms of cost-savings. For example, robust procedural bad faith law might encourage insurers to rely on AI asymmetrically, using it to expedite full claim payments while reserving denials or reductions for more careful review, as Lemonade reportedly does.[67] It might also spur insurers to develop more effective forms of human-AI collaboration that preserve procedural fairness while still reducing administrative costs.

A third important set of objections is institutional. Expanding bad faith law to better account for the procedural harms associated with AI-centric claims handling could make such disputes harder for courts to resolve, producing more litigation and less certainty. The extent to which an insurer has replaced meaningful human review with AI-driven review is a fundamentally factual question. Allegations of procedural bad faith therefore might generate not only more lawsuits, but also lawsuits that are difficult to resolve early, requiring extensive discovery and perhaps even trial.[68] The resulting uncertainty could be costly. Because cases turning on factual disputes are harder to predict, insurers might become too willing to settle even when their claims-handling processes are fair and do preserve meaningful human review.

These concerns about the fact-based nature of procedural bad faith are only enhanced by the inherently slippery distinction between meaningful human review and human rubber-stamping of AI claims processes. In fact, no

---

Cal. filed July 24, 2023).

[66] See Daniel Schwarcz, Obamacare For Homeowners Insurance: Fixing America's Broken Insurance Markets in a Time of Climate Change, 82 HARV. ENVIRONMENTAL L. REV. 525 (2025).

[67] See Lemonade, Inc., supra.

[68] By contrast, in many ordinary bad faith cases, both the correctness of the insurer's coverage decision and the reasonableness of its position can often be resolved as matters of law with relatively limited factual development.

institution has yet managed to draw that line effectively. The most instructive failure is federal. In its 2023 Medicare Advantage rule, CMS required that coverage denials on medical necessity grounds be reviewed by a physician or other appropriate professional and prohibited plans from relying on algorithmic predictions without accounting for the individual patient's circumstances.[69] But the rule defined neither what such review must consist of nor how compliance would be assessed, and CMS's subsequent guidance clarified only that an algorithm cannot serve as the sole basis for a determination, a standard that the 1.2 seconds of physician attention documented at Cigna might technically satisfy.[70] The state physician-review statutes share the same silence. None specifies how much time a reviewer must spend, what materials she must examine, whether she must have authority and freedom from retaliation when she departs from the algorithm, or how anyone would ever know.

This definitional problem arises from the fact that meaningful review of AI output is a quality of cognition, not of procedure, and cognition resists direct regulation. Any operationalization of this concept by courts, regulators, or legislators must therefore rely on proxies, each imperfect: minimum review times invite clock-watching compliance; requirements that reviewers access the full claim file cannot ensure that the file is read; monitoring of override rates is equally consistent with an accurate algorithm and a captured reviewer; and documentation requirements generate paper trails that generative AI can itself produce. As Part II explained, moreover, even a diligent reviewer confronting a completed algorithmic determination faces anchoring and automation biases that formal independence does not dispel. Courts developing procedural bad faith doctrine would inherit every one of these difficulties.

As above, these are meaningful objections. At the same time, they may overstate the difficulty of the judicial task and understate courts' capacity to distinguish fair from unfair insurer practices.[71] The gradual development of common-law standards defining meaningful human review of AI processes may be especially well suited to this problem. Determining what it means for human claims handlers to retain a meaningful and independent role in claim denials is precisely the kind of context-sensitive inquiry that courts routinely perform. At the very least, courts, aided by discovery into actual claims workflows of the

---

69 Medicare Program; Contract Year 2024 Policy and Technical Changes to the Medicare Advantage Program, 88 Fed. Reg. 22120 (Apr. 12, 2023) (codified in relevant part at 42 C.F.R. §§ 422.101, 422.566).

70 See CMS, Frequently Asked Questions Related to Coverage Criteria and Utilization Management Requirements in CY 2024 Final Rule (Feb. 6, 2024).

71 See Abraham & Schwarcz, supra note, at 146-148. State insurance regulators face competing objectives, often including a strong focus on solvency; they may also be vulnerable to capture and constrained by limited resources. For those reasons, courts may sometimes be institutionally better positioned than regulators to police procedural misconduct in claims handling.

kind described above, can identify egregious cases at the extremes, condemning systems in which reviewers lack the time, information, authority, or incentive to disagree, while leaving a wide middle ground to be mapped incrementally as recurring fact patterns emerge. And as courts confront recurring fact patterns, they could articulate more specific doctrinal formulations of what procedural fairness requires in AI-centered claims handling, thus facilitating more frequent early resolution of such disputes.

A final limitation, if not an objection, to relying on procedural bad faith to respond to AI-based claims handling is that state bad faith law is categorically unavailable for the many disputes that involve employer-sponsored insurance coverage.[72] This is hardly a small limitation. Roughly half the country, more than 150 million people, receives health coverage through employer-sponsored plans governed by ERISA.[73] And many others receive disability, life, and long-term care through employer sponsored-plans. Such participants in an employer-sponsored plan whose claims are denied by an algorithm without meaningful human review are therefore relegated to ERISA's own remedies, which permit recovery of the benefits due but provide no damages for the manner in which the claim was handled, no matter how egregious. Nor can states easily route around this barrier through regulation.

The implications for this Essay's argument are twofold. First, the domain in which a reinvigorated procedural bad faith doctrine could operate is principally property/casualty insurance, life and disability coverage outside employer plans, and the individual health insurance market. That is a substantial domain, but it excludes the very context in which fully automated denial has been most clearly documented. Second, and more constructively, the ERISA gap sharpens the case for parallel reform at the federal level. If the procedural injury identified in Part III is real, it is no less real for participants in employer-sponsored plans, and only Congress, or federal regulators acting within existing authority, can supply a remedy there. The 2023 Medicare Advantage rule discussed below represents a first, halting federal step in that direction. State courts developing procedural bad faith doctrine would, at a minimum, generate the doctrinal vocabulary and factual record on which such federal efforts could draw.

CONCLUSION

Distributing Risk helped crystallize the economic logic of bad faith law,

---

[72] The Supreme Court held in *Pilot Life Insurance Co. v. Dedeaux*, 481 U.S. 41 (1987) that ERISA preempts state common-law bad faith claims arising from the handling of benefits under such plans. In *Aetna Health Inc. v. Davila* it further held that that ERISA's civil enforcement scheme completely preempts any state cause of action, however labeled, that seeks a remedy for the denial of plan benefits. Oliva, supra (analyzing these preemption dynamics in detail).

[73] See Kaiser Family Found., Employer Health Benefits Survey.

shaping its development into a doctrine centrally concerned with deterring insurers from wrongfully delaying or denying policyholders the benefits they are owed. But increasingly automated claims processes expose a separate injury: the denial of a fair procedure itself. When an insurer refuses, reduces, or delays payment without meaningful human review and adequate explanation, the insured is deprived of more than a potentially correct outcome; she is deprived of voice, dignity, neutrality, and trustworthiness in a relationship defined by vulnerability and dependence. Courts should respond by giving procedural fairness greater weight within the bad faith inquiry, while recognizing the displacement of meaningful human judgment by automated systems as a new and distinctive form of procedural bad faith. Doing so would extend the central insight of Ken's work: that insurance law must respond to the distinctive incentives and vulnerabilities created by the insurance relationship, even as new technologies transform the form those risks take.