The Maker Has Always Graded Its Own Homework

A century of self-certification and the case for independent scrutiny

Allan Watkins

Laws of Robots series

Version 1.2 — 31 October 2026

Abstract

The cases examined here show how major industries permitted to assess their own trustworthiness developed related failures of oversight. Corporate accounting, aircraft certification, and vehicle emissions each produced catastrophe and official investigations. Their records support the comparison drawn in this paper: interested parties held critical evidence, verification depended on the parties being judged, and warnings did not produce effective correction before catastrophe. Reforms strengthened outside inspection, regulatory authority, and testing beyond manufacturers' controlled conditions. Credit ratings provide a contrasting case: substantial reforms arrived, while issuer-pays remained the primary business model described in the SEC staff's latest report. Autonomous machines raise the same questions about evidence and authority. This paper argues for scrutiny independent of the maker, access to the evidence needed for that scrutiny, and the ability to require correction. The cases support that institutional lesson; they do not establish that every system of self-assessment must fail, or that one reform alone guarantees safety.

1. The Pattern, Stated

In the cases examined here, a major industry helped author the verdict on whether its accounts, its aircraft, its engines, or its ratings deserved belief — and serious failures escaped effective correction. Critical evidence remained with interested parties, while warnings that reached outsiders did not bring a sufficient response. The accumulated failure then surfaced in a collapse, a crash, or a scandal whose origins proved to have been years in the making.

None of this requires dishonesty, and the comparison is not an accusation of it. A verdict authored by the party it judges can contain valuable information; it does not, by itself, provide independent assurance. The author gains from being believed, holds critical evidence, and writes the report card. Sincerity and technical skill matter, but they do not remove that institutional conflict. The question is whether anyone outside that relationship can examine the evidence, challenge the verdict, and bring about a correction. This paper approaches that question through the historical record.

What follows is evidence. Three industries — corporate accounting, aircraft certification, vehicle emissions — each combined forms of self-assessment with public oversight. Each produced a catastrophe and official investigations. Those inquiries documented different failures, conflicts, and weaknesses in enforcement. The structural comparison is this paper's: critical evidence remained with interested parties, and verification lacked the independence or authority to force timely correction. The reforms strengthened outside scrutiny in different ways — independent inspection, restored regulatory authority, and testing beyond the manufacturer's controlled conditions. The arrangements were not identical. The recurring question was who could examine the evidence and challenge the verdict.

A fourth case makes the comparison more demanding. In credit ratings, official inquiries identified a payment conflict, Congress required study of an assignment alternative, and reforms strengthened oversight without replacing issuer-pays as the primary business model. That is a comparison between institutional choices, not a controlled experiment or proof of the next crisis. It asks what changes when scrutiny is added while an important source of dependence remains.

Autonomous machines raise these questions again. Their assurance arrangements are taken up at the end, after the historical cases. The factual account below draws on the public, primary record — committee reports, statutes, official inquiries, regulatory publications, and court records. The comparison among those records, and the conclusions drawn from it, are the author's.

2. Case One: The Auditors — The Fix That Had to Be Installed Twice

The oldest case runs the full arc twice, which is why it comes first.

Before 1933, no federal statute required an outside check of the kind later imposed on public securities offerings. The Senate Banking and Currency Committee's 1932–34 investigation documented manipulation, misleading corporate information, and listing reviews conducted without independent investigation; its final report rejected the exchanges' claim that internal regulation could replace governmental control.[1] The hearings helped lay the groundwork for the Securities Act of 1933, which required the financial statements accompanying a public offering to be certified by an independent public accountant; the Exchange Act of 1934 extended federal reporting requirements and created the SEC. The final report described legislation already enacted during the investigation; it did not precede the 1933 Act. The verdict on a company's accounts would no longer be authored solely by the company.

But the fix carried a seam, and the seam took seventy years to open fully. The independent accountant was — and is — selected and paid by the company being audited. And oversight of the accountants themselves remained substantially self-regulatory: audit firms inspected one another through peer review, supervised by a Public Oversight Board that the AICPA created in 1977 and funded through its SEC Practice Section's member-firm dues. Its five members were public

members; the institutional dependence was in its creation and financing, not a claim that accountants staffed every seat.[2]

Meanwhile the economics of the relationship shifted. By 2000, the SEC estimated that management advisory and similar services accounted for roughly half of the Big Five firms' total revenues. That was a firm-wide figure, not a finding that most audit clients bought more consulting than auditing.[3] Enron was a specific example of the conflict: in 2000, Andersen received about $27 million in consulting fees from Enron and $25 million for audit work.[4]

In the autumn of 2001, Enron restated years of earnings and collapsed into bankruptcy. Its board's special investigative committee wrote that the evidence available to it "suggests that Andersen did not fulfill its professional responsibilities" in auditing Enron's financial statements and raising concerns about the related-party transactions.[5]

At the Senate's January 2002 hearing, Senator Fred Thompson asked where the watchdogs had been.[6] The committee staff's later report described "a failure of all the watchdogs to properly discharge their appointed roles". It identified specific conflicts involving Andersen's consulting fees, securities analysts' investment-banking relationships, and the accounting profession's self-discipline.[7] Those findings support the comparison made here; they are not a finding that every watchdog was chosen, paid, and policed by Enron.

Arthur Andersen, Enron's auditor, was convicted in 2002 of obstruction in connection with document destruction. The Supreme Court unanimously reversed the conviction in 2005 because the jury instructions failed to convey the required elements of the offense.[8] The paper's comparison rests on the documented conflicts and oversight failures, not on treating that reversed conviction as a finding of guilt that still stands.

Congress's response reinstalled the fix, one level up. The Sarbanes-Oxley Act of 2002 replaced self-regulatory peer review for public-company audits with an inspection board that the audited firms do not select, do not pay directly, and cannot overrule. The auditor's independence from its client was hardened in parallel — prohibited consulting services, mandatory partner rotation, engagement authority moved from management to audit committees. The 1933 fix had made the report card independent of the company. The 2002 fix made the grader's grader independent of the profession.[2]

Independence also depends on how the overseeing institution is maintained. On January 22, 2026, the SEC approved a PCAOB budget 9.4% below the prior year's and reductions of 52% and 42% in chair and other board-member compensation. On January 30, it appointed a new chair and three other members; George Botic continued his service.[9,10] These were exercises of the SEC's statutory oversight powers, not by themselves evidence that audit quality had fallen. They show why the arrangement must be assessed through its appointments, funding, technical capacity, and authority over time. Independence is a continuing institutional responsibility.

3. Case Two: The Aircraft — Delegation and Independent Oversight

The aircraft case concerns how a public certification system can depend on findings made inside the manufacturer it regulates. The FAA retained certification and oversight authority; the issue was how independently that authority could be exercised in practice.[11–13]

Delegation long predates the 737 MAX. The FAA's own history describes both individual representatives and organizational delegation before the Organization Designation Authorization program. Its 2005 final rule expanded and standardized organizational delegation, explicitly seeking more efficient appointment and oversight of designees and more FAA attention to safety-critical matters. ODA holders perform authorized functions on the FAA's behalf, under agency supervision.[11] The concern examined here is the dependence created when employees of the organization being assessed also make findings used to establish its products' compliance.

The warning was on record before the crashes. In October 2015, the Transportation Department's Inspector General found shortcomings in FAA staffing, risk-based oversight, and supervision of delegated certification work. The report made recommendations that led to corrective work; its record is not one of complete official inaction.[14] It nevertheless documents weaknesses in the oversight arrangement before the public reckoning over the MAX.

On October 29, 2018, Lion Air Flight 610 crashed into the Java Sea, killing all 189 aboard. On March 10, 2019, Ethiopian Airlines Flight 302 crashed near Addis Ababa, killing 157. Both aircraft were Boeing 737 MAX 8s. The investigations examined repeated automatic nose-down MCAS commands and a design that permitted activation on input from a single angle-of-attack sensor. Three hundred forty-six people died.[12,13]

The House committee's majority-staff report documented production pressure, faulty design and performance assumptions, information withheld from the FAA and pilots, and conflicts in delegated oversight. It also reported Boeing employees' concerns about pressure while performing FAA-delegated functions.[12] The Joint Authorities Technical Review found that inadequate FAA awareness and involvement prevented an independent assessment of Boeing's proposed MCAS certification activities, and reported signs of undue pressure on Boeing's delegated engineering members. It also said that extensive delegation need not compromise safety when FAA engagement and oversight are adequate.[13] The comparison drawn here concerns failures of independent scrutiny in this program, not a finding that every delegation is unsafe.

Congress responded in the Aircraft Certification, Safety, and Accountability Act of 2020. The law strengthened FAA control over newly selected unit members in ODAs engaged in design, including approval and rescission powers, subject to the statutory procedures. It also prohibited interference by covered transport-airplane manufacturers' supervisory employees and provided for civil penalties.[15] The FAA's 2025 ODA procedures address member approval, protection from interference, and communication between members and the agency.[16]

Delegation continued alongside those protections. Order 8100.15C, issued on October 1, 2025 and announced in the Federal Register on November 21, introduced an Airmen Certification ODA type and a systems-based oversight approach. The notice recorded supporters' arguments about costs and certification bottlenecks, and the FAA's assessment that the expansion would not introduce additional safety risk.[16] Efficiency is a legitimate aim; it does not answer the separate question of whether the regulator retains enough independent expertise, access, and authority to challenge delegated findings.

4. Case Three: The Emissions — The Machine That Recognized the Exam

The third case is the most literal in the set: the machine itself was taught to recognize the exam.

The arrangement worked like this. Under the pre-reform European framework, the manufacturer normally chose the technical service whose tests supported approval by a national authority; an approval granted in one member state was valid across the Union. Type-approval authorities and technical services were usually financed in part by fees paid directly by manufacturers. The regulatory test was a fixed laboratory driving cycle, even though outside researchers and the Commission's Joint Research Centre had documented the gap between laboratory and on-road emissions for years.[17] The regulatory exam stayed in the classroom; independent road research did not.

On September 18, 2015, the U.S. Environmental Protection Agency issued a notice of violation to Volkswagen over software that detected emissions testing and activated full emissions controls during the test. In normal operation, the affected 2.0-liter vehicles emitted nitrogen oxides at up to 40 times the standard.[18] Volkswagen subsequently described the relevant engine type as installed in roughly eleven million vehicles worldwide.[19]

Volkswagen agreed to plead guilty in January 2017 and entered its guilty plea in federal court on March 10, 2017. The three felony counts concerned conspiracy, obstruction of justice, and importing goods by false statements; the plea included agreement to a $2.8 billion penalty.[20] This was a corporate admission of guilt. The paper's structural comparison does not depend on treating all the other cases as equivalent criminal conduct.

The European Parliament's committee of inquiry documented failures of oversight and enforcement by both member states and the Commission. Its report also identified weaknesses in the system's design: manufacturers could seek flexible interpretations and lower fees, normally chose their technical services, and faced in-service checks conducted mostly in manufacturers' laboratories. It described potential conflicts arising from those contractual arrangements, while recording the Commission's position that it had no concrete evidence of conflicts of interest.[17] The absence was not of all outside road testing: no EU or member-state authority had searched for or proved the illegal use of defeat devices before September 2015, and type-approval tests had

remained confined to the laboratory cycle. The structural comparison drawn here rests on those specific findings.

The reforms strengthened outside scrutiny in two ways. Real-driving-emissions testing added portable measurement on actual roads, beyond the fixed laboratory cycle.[21] Regulation (EU) 2018/858, applicable from September 2020, required national market-surveillance checks, empowered the Commission to conduct its own laboratory and road tests, and provided for Commission fines. It also required access to information needed for those checks, including software and algorithms, with confidentiality protections.[22] Public verification now operates alongside technical-service approval testing. The lesson is the addition of independent access and enforcement, not the elimination of every manufacturer role or every possibility of evasion.

What this case established outlasts diesel engines, because it is a finding about measurement, not about fuel. A test the judged party controls becomes the thing the judged party optimizes against.

5. Case Four: The Ratings — The Payment Conflict That Remains

The fourth case concerns a remedy studied without becoming the general assignment system. It allows a comparison between adding oversight and changing who selects the grader.

The arrangement: since the 1970s, the major credit-rating agencies have been paid by issuers of the securities they rate.[23] The grader is chosen and paid by the graded, and a structured-finance desk could shop a deal among agencies in search of the best grade. The Senate investigation documented that practice and the business pressure reflected in internal correspondence.[24] Published ratings did not reveal all the incentives and judgments behind them.

During the financial crisis, highly rated mortgage-related securities suffered mass downgrades. The Financial Crisis Inquiry Commission called the rating agencies key enablers of the collapse; its conclusions described flawed models, pressure from paying financial firms, and the drive for market share.[25] The Senate inquiry likewise documented ratings shopping and pressure from investment banks.[24]

Congress adopted substantial reforms. The Dodd-Frank Act established an SEC Office of Credit Ratings, annual examination requirements, and governance and disclosure duties. Section 939F required study of assignment and compensation alternatives for structured-finance ratings and directed assignment-system rulemaking as the SEC determined necessary or appropriate in the public interest or for investor protection, with an alternative permitted if it would better serve those aims.[26] The staff study appeared in December 2012 and recommended a roundtable; the SEC held it in May 2013.[27,28]

The operating framework described in the SEC staff's 2025 report, published in April 2026, still identifies issuer-pay as the primary business model. It records annual examinations, disclosure requirements, internal controls, and rules separating ratings from sales and marketing influence.

It also identifies continuing conflicts, including ratings shopping, and notes conflicts in subscriber-pay and investor-pay models too.[29] That is a more specific finding than saying nothing changed. The reforms constrain conduct while leaving a central selection-and-payment relationship in place.

The contrast is between different institutional responses to dependence. Independent inspection, stronger public authority, and testing beyond the manufacturer's controlled conditions were added in the earlier cases. In ratings, an assignment alternative was studied while issuer-pay remained the primary model in the later SEC account. That supports the paper's concern about residual conflicts; it does not prove that reforms had no effect, that every alternative would succeed, or when another failure will occur.

6. The Pattern, Named

Four cases are now on the table, drawn from different industries, decades, and bodies of law. The inquiries did not describe identical systems or prescribe identical remedies. The comparison in this paper identifies three recurring features:

The recurring problem was not simply the absence of a record. It was outsiders' limited access to critical records, and their ability to act on what was already known.

First, the interested party held critical evidence. The auditor's workpapers, the manufacturer's internal analyses, the engine-control software, and the rating models were held by parties with a stake in the verdict. Public warnings and independent research could expose a problem without giving outsiders routine access to all the records needed to examine it.

Second, verification depended on the party being judged. The auditee paid the auditor. Delegated engineers could work within the manufacturer. The emissions testing service was chosen and partly financed by the manufacturer. The issuer paid for its rating. Technical competence remained necessary, but it did not remove the pressures created by who selected, paid, employed, or could replace the examiner. The degree of dependence and the regulator's power varied among the cases.

Third, warnings did not produce effective correction before catastrophe. This is not a claim that nobody saw a problem. The aircraft and emissions records contain warnings and outside work before the public reckoning. The recurring failure was that available evidence did not lead to timely investigation, enforcement, or correction. Concealment mattered in some cases; ineffective action on known information mattered too.

Each case adds a distinct lesson. An outside auditor can still need an independent overseer. Delegation can pursue legitimate efficiency goals while requiring the regulator to retain its own technical capacity. A predictable test can become the target of optimization instead of a reliable measure of real-world performance. And stronger conduct rules can coexist with a continuing

payment conflict. Together, these observations argue for examining independence as an operating arrangement, not assuming it from a title or an initial reform.

The reforms described here share a direction: they strengthened the ability of someone outside the interested relationship to examine evidence, challenge findings, and require correction. Better internal records, processes, and expertise also matter. They do not, by themselves, make the maker's verdict independent. The cases are not an exhaustive test of every possible remedy, and stronger oversight is not a guarantee against future failure. They support a practical principle: the party seeking trust should not have the final, unchallengeable word on whether its own evidence deserves it.

This paper argues for that principle without prescribing one institution to embody it. Evidence must be available to appropriately authorized scrutiny; judgment must be sufficiently independent of the party being judged; and a finding must be capable of leading to action. Those requirements can be met through different institutions. The question is how to make them effective for autonomous machines.

7. The New Instance

The new instance is autonomous machines — systems that perceive their surroundings, make decisions, and act physically near people. The physical consequences make the assurance question urgent. This is not an industry without a history of harm: the NTSB investigated the fatal 2018 Tempe collision involving an Uber test vehicle controlled by a developmental automated driving system, with a safety driver present.[30] The question is how to prevent further harm and strengthen scrutiny as autonomous systems take on wider roles.

Hold the arrangement up against the three features. Where a machine's decision record is generated by the maker, stored on the maker's infrastructure, kept in its formats, and produced at its discretion, the interested party holds critical evidence. Internal safety frameworks can also leave verification with the party being judged. The AI Act illustrates one route, not a universal rule for robots: Article 43(2) prescribes internal-control assessment without a notified body for the high-risk uses in Annex III, points 2–8; product-related systems follow the separate conformity route in Article 43(3).[31] The question for autonomous machines is whether records and independent scrutiny can support effective correction before dangerous failures become catastrophic.

The aircraft lesson matters here because reasonable development practices can coexist with a dependence on the developer's own verdict. Published materials from AI and autonomous-vehicle developers describe learning from real-world use, continuous improvement, evolving risk-mitigation frameworks, and safety cases assembled by the developer. Some also warn that regulatory requirements can slow development and deployment.[32–36] These examples show how developers describe assurance; they do not establish a uniform industry position on independent review.

The same record includes support for regulation, a commitment to sign an AI code of practice, and commissioned third-party audits of safety-management systems and safety cases.[32,35,36] Those forms of scrutiny matter. They do not, by themselves, establish the general access to a machine's decision record that this paper argues for. The question is who can inspect the evidence behind the verdict, and on what authority.

Telemetry can provide an extensive record of what a machine does. That is valuable evidence, and it can support investigation and improvement. But the existence and volume of a record do not, by themselves, establish who can inspect it, whether it is sufficient to test a claim, or whether an adverse finding can lead to correction.

Where operational data is held on the maker's servers, in its formats, and produced at its discretion, the first feature reappears: critical evidence remains with an interested party. The issue is not whether telemetry is useful. It is whether appropriately authorized outsiders can obtain and examine the relevant evidence without depending entirely on the maker's chosen account. The fuller question of records and risk belongs to a companion paper; here, the point is that data collection alone does not settle independence.

The cases do not prescribe one organization or guarantee one outcome. They do identify the institutional work that the maker's own verdict cannot perform alone: independent examination of relevant evidence, the ability to challenge that verdict, and authority to bring about correction. How such scrutiny is constituted, funded, and protected from capture remains deliberately open. For autonomous machines, the task is to build and strengthen those capabilities before further harm forces another round of reform. The maker's evidence should inform the judgment; the maker should not have the final, unchallengeable word.

Appendix — The Pattern at a Glance

The comparison in this paper identifies three recurring features:

1. **The interested party holds critical evidence.**
2. **Verification depends on the party being judged.**
3. **Warnings do not produce effective correction before catastrophe.**

The reforms point toward a common principle: **evidence outsiders can examine, judgment independent of the maker, and authority the maker cannot unilaterally overrule.**

Industry	The self-grading arrangement	The catastrophe	The official record	The reform	When
Corporate accounts	Self-reported financial statements	The 1929 crash	Stock Exchange Practices hearings (1932–34); S. Rep. 73-1455 (1934) [1]	Mandated independent audit; creation of the SEC	1933–34
The audit profession	Auditee-paid audits, policed by peer review	Enron's collapse; Andersen conviction later reversed [8]	Powers Report; Senate prints and hearing (2002) [4–7]	Independent inspection — the PCAOB, under Sarbanes-Oxley	2002
Aircraft certification	Certification delegated to the manufacturer's own employees	Two 737 MAX crashes; 346 deaths (2018–19)	House majority-staff report; JATR (2019–20) [12,13]	FAA control strengthened over delegated certification (ACSAA) [15]	2020
Vehicle emissions	Maker-chosen, partly maker-funded testing against a laboratory-only regulatory test	The diesel defeat-device scandal (2015)	EMIS inquiry, A8-0049/2017 [17]; U.S. criminal plea (2017)	On-road testing; independent market surveillance with power to test and fine	2016–20
Credit ratings	The issuer pays for its own rating	The 2008 financial crisis	FCIC; Senate inquiry (2011) [24,25]	Assignment alternative studied; issuer-pay remains primary in SEC report [26–29]	—
Autonomous machines	Assurance arrangements vary [31–36]	Fatal automated-driving collision (Tempe, 2018) [30]	NTSB investigation (2019) [30]	Independent access and effective correction argued for; no single institution proposed	ongoing

References

[1] U.S. Senate, Committee on Banking and Currency, Stock Exchange Practices, S. Rep. No. 73-1455, 73rd Cong., 2nd sess. (1934), pp. 1–4, 68–73, 80–81, 152. Title-page date: June 6 (calendar day June 16), 1934. FRASER report.

[2] U.S. General Accounting Office, The Accounting Profession, GAO-02-411 (May 17, 2002), Background, POB oversight and funding sections, and note 5 (termination May 1, 2002). Sarbanes-Oxley Act, Pub. L. 107-204, §§2(a)(2), 101, 104 (issuer-audit scope); statute. GAO report.

[3] SEC, Revision of the Commission's Auditor Independence Requirements, Releases 33-7919 and 34-43602, 65 Fed. Reg. 76008, 76013, 76015 (Dec. 5, 2000). Federal Register.

[4] Senate Permanent Subcommittee on Investigations, The Role of the Board of Directors in Enron's Collapse, S. Prt. 107-70 (July 8, 2002), p. 54 and note 192. Senate print.

[5] Special Investigative Committee of Enron's Board of Directors, Report of Investigation (Powers Report, Feb. 1, 2002), p. 24, Outside Professional Advisors. Report copy.

[6] Senate Committee on Governmental Affairs, The Fall of Enron: How Could It Have Happened?, S. Hrg. 107-376 (Jan. 24, 2002), p. 5, opening statement of Senator Fred Thompson. Hearing.

[7] Senate Committee on Governmental Affairs staff, Financial Oversight of Enron: The SEC and Private-Sector Watchdogs, S. Prt. 107-75 (Oct. 7, 2002), pp. 2, 5, 17–19, 22. Staff report.

[8] Arthur Andersen LLP v. United States, 544 U.S. 696, 698–708 (2005), unanimous opinion reversing the conviction for erroneous jury instructions. U.S. Reports.

[9] SEC, SEC Approves 2026 PCAOB Budget and Accounting Support Fee, Release 2026-11 (Jan. 22, 2026), opening paragraphs. SEC release.

[10] SEC, SEC Appoints New Chairman and Board Members to PCAOB, Release 2026-16 (Jan. 30, 2026), opening paragraph: four appointees; George Botic continues. SEC release.

[11] FAA, Establishment of Organization Designation Authorization Program, 70 Fed. Reg. 59932–59949 (Oct. 13, 2005), p. 59932, Summary, Authority, and History of Designation Programs. Final rule.

[12] House Committee on Transportation and Infrastructure, majority staff, The Design, Development & Certification of the Boeing 737 MAX (Sept. 2020), pp. 6–8, 12–16, 56–70, 185–186. Committee report.

[13] Joint Authorities Technical Review, Boeing 737 MAX Flight Control System: Observations, Findings, and Recommendations (Oct. 11, 2019), Executive Summary pp. I, VII–VIII, Recommendation R5; pp. 27–28. JATR report.

[14] DOT Office of Inspector General, FAA Lacks an Effective Staffing Model and Risk-Based Oversight Process for Organization Designation Authorization, AV-2016-001 (Oct. 15, 2015), official summary and recommendation dispositions. OIG record.

[15] Aircraft Certification, Safety, and Accountability Act of 2020, Pub. L. 116-260, div. V, title I, §107; 49 U.S.C. §§44741(a), (d)–(e), 44742(b), 2022 edition, pp. 1307–1309. Enacted provisions in the U.S. Code.

[16] FAA, Order 8100.15C, issued Oct. 1, 2025, listed Active when checked Oct. 31, 2026; Notice of Availability, 90 Fed. Reg. 52787–52788 (Nov. 21, 2025), Summary, Background, and Comments opposing AC ODA. Order record; notice.

[17] European Parliament, EMIS Committee, Report on the inquiry into emission measurements in the automotive sector, A8-0049/2017 (Mar. 2, 2017), rapporteurs Jens Gieseke and Gerben-Jan Gerbrandy; Conclusions paras. 3, 19, 29, 36–42, 45, 48–52, 56–57; full inquiry report §5.2. Inquiry report.

[18] EPA, Learn About Volkswagen Violations, September 18, 2015 notice and description of the 2.0-liter defeat device. EPA account.

[19] Volkswagen Group, Annual Report 2024, Notes, Litigation, Diesel issue, opening paragraph, describing roughly eleven million EA 189 engines worldwide. Manufacturer account.

[20] DOJ, Volkswagen AG Pleads Guilty in Connection with Conspiracy to Cheat U.S. Emissions Tests (Mar. 10, 2017), opening paragraphs: January agreement, March plea, and agreed $2.8 billion penalty. DOJ release.

[21] Commission Regulation (EU) 2016/427 (Mar. 10, 2016), recitals 2–5 and Annex II introducing Annex IIIA, Verifying real driving emissions. EUR-Lex.

[22] Regulation (EU) 2018/858, Articles 8(1), (8), (12), 9, 68–72, 85, 91; consolidated text of Aug. 2, 2026, checked Oct. 31, 2026. EUR-Lex.

[23] SEC, Report on the Role and Function of Credit Rating Agencies in the Operation of the Securities Markets (Jan. 2003), §IV.F.1, p. 41, Issuers Paying for Ratings. SEC report.

[24] Senate Permanent Subcommittee on Investigations, Wall Street and the Financial Crisis: Anatomy of a Financial Collapse, majority and minority staff report (Apr. 13, 2011), pp. 5–6, 273–287. Senate report.

[25] Financial Crisis Inquiry Commission, The Financial Crisis Inquiry Report, official government edition (Jan. 2011), Conclusions, p. xxv. Commission report.

[26] Dodd-Frank Act, Pub. L. 111-203, §§932, 939F(b)–(d), 124 Stat. 1872–1883, 1889–1890 (2010). Enacted text.

[27] SEC Division of Trading and Markets staff, Report to Congress on Assigned Credit Ratings (Dec. 2012), pp. 1–2, 72–82, especially p. 73. The Commission expressed no view on the staff analysis. Staff study.

[28] SEC, Credit Ratings Roundtable, event record for May 14, 2013, opening description and panel topics. Roundtable record.

[29] SEC Office of Credit Ratings staff, 2025 Staff Report on Nationally Recognized Statistical Rating Organizations (published Apr. 24, 2026), pp. 4–5, 32–37, especially pp. 33–35. Latest annual report listed when checked Oct. 31, 2026. Staff report.

[30] NTSB, Collision Between Vehicle Controlled by Developmental Automated Driving System and Pedestrian, Tempe, Arizona, March 18, 2018, HAR-19/03 (2019); What Happened and What We Found. Investigation record.

[31] Regulation (EU) 2024/1689, Article 43(2)–(3) and Annex VI; consolidated text of July 27, 2026. EUR-Lex.

[32] OpenAI, Our approach to AI safety (Apr. 5, 2023), Building increasingly safe AI systems and Learning from real-world use to improve safeguards. Developer publication.

[33] Anca Dragan, Helen King and Allan Dafoe, Introducing the Frontier Safety Framework (Google DeepMind, May 17, 2024), Risk domains and mitigation levels and Investing in the science. Developer publication.

[34] Francesca Favaro et al., Building a Credible Case for Safety: Waymo's Approach for the Determination of Absence of Unreasonable Risk, arXiv:2306.01917 (June 2, 2023), abstract. Author paper.

[35] Kent Walker, We will sign the EU AI Code of Practice (Google, July 30, 2025), opening paragraphs. Developer publication.

[36] Aurora Innovation, Driverless Safety Report 2025, with July 2026 addendum, pp. 65 and 77, Third party validation. Developer report.

Allan Watkins

Laws of Robots series

Version 1.2 — 31 October 2026

This document is part of the Laws of Robots governance whitepaper series. The version of record is maintained under version control. All commitments are owner-authored and subject to the published change discipline.