

Insuring The Robot

The Standard That
Doesn't Exist

October 2026

AUTHORS:



Ross Naylor

Analyst | Research

✉ ross@activantcapital.com



Benjamin Bradlow

Intern

About Activant

Activant is a research-led global investment firm that partners with high-growth companies. Since 2015, we have invested in category-defining businesses during their most critical phases of growth, partnering with founders who have won their initial battles and are ready for the next challenge.

Our approach pairs deep, proprietary research with patient, flexible capital and hands-on operational partnership. We work alongside founders and leadership teams to refine strategy, strengthen operations, and accelerate sustainable growth.

Activant Research is dedicated to uncovering the most exciting emerging technologies, sectors, and companies we believe will shape the future. Our research-driven perspective informs everything we do, helping us invest at meaningful inflection points and support founders in building enduring, category-leading businesses.

You can find out more about Activant and our research at <https://activantcapital.com/>.


Today, a robot deployed onto a factory floor can hold a certificate proving its body will not suddenly combust, exceed a specific contact force, or electrocute anyone working alongside it. Nothing certifies whether the brain controlling the body is safe enough to stand next to. For an insurer thinking about covering a fleet of humanoids, pricing the body is a solved problem; pricing the brain is not.

In the [agentic insurance piece](#) we published last week, we argued that when an agent deployed in the digital world causes a loss, the durable business in insuring it belongs to whoever owns the standard defining a safe-enough agent and the record of how it fails, not to whoever pays the claim. Give that agent a body and the argument only sharpens. **Outside of China, no such standard exists, and nobody has built the business that would issue it.** Before asking who could, it is worth being precise about what is actually being insured.

Property and Peril

In insurance, property is what suffers a loss and peril is what causes one. A deployed instance of embodied intelligence, *like a humanoid*, is both at once: an expensive asset that can cause damage to the people and objects in its vicinity.

What turns this into an insurable market is deployment volumes; an insurer can't price a population that doesn't exist. Deployment numbers are starting to compound, and the cost curve is starting to show [Wright's Law](#) dynamics.^{1,2} The average entry price for a humanoid has fallen from ~\$100,000 to ~\$50,000 in two years, with [Unitree's](#) R1 available on the retail market for as little as \$4,900.^{3,4}



Each price reduction puts more machines into more workplaces, and each deployment takes out the next increment of cost: a recursive loop that results in cheaper robots deployed at scale. Wright's Law relies on cumulative deployments. So does the [law of large numbers](#), which is one of the reasons insurance works at all. The same volume that pulls the cost curve down is also what turns a fragmented list of incidents into a cohesive loss record.

Volume also means contact at scale. An embodied agent shares floor space with people, so its mistakes arrive as bodily injury and property damage, which are the oldest and best-understood perils in the industry. Today, this shared floor is largely within industrial or commercial settings, where robots work near people and the operator is responsible for their actions. The household is the end-goal though. Imagine a world where the home robot is as ordinary as the dishwasher: the same machine that unpacks the groceries or mows the lawn is the one left in a room with a child. [Home policies were never written for a machine that acts on its own authority.](#) Wherever the contact happens, the machine is both property and peril, and that splits the exposure into two books that can't be underwritten the same way.

[The hardware book](#) is the damage to the robot itself: the stripped gearbox, the burnt-out actuator, the stolen unit. This is most similar to [equipment breakdown](#) and [inland marine](#) in traditional insurance, and the repair bill for a humanoid is non-trivial: a single service call in China already costs anywhere between \$4,400 and \$42,000.⁵ Pricing the hardware book is a largely solved problem today; an underwriter can price a dropped machine without any view on why it fell.

[The liability book](#) is the damage the robot imposes on its environment: the injured worker, the crushed pallet, the third-party claim. Nobody can underwrite this properly yet, because pricing it requires a view on how a learned policy behaves in a room it has never seen.


The opportunity lies in this asymmetry. In physical AI, [the body is commoditizing and the brain is generalizing](#). We think the same split applies to cover. Insuring the body will commoditize with the hardware it protects. Insuring the brain is the half nobody can write yet, which is why we think it is the half worth owning. What stops anyone from writing it becomes clear the first time one of these machines hurts somebody.

Everybody's Fault

Say a humanoid knocks a worker down a loading ramp and they break their arm.



The robot OEM built the body; the model developer trained the perception and control stack; the systems integrator installed it; and the operator deployed it. The worker's own conduct will also be argued into the mix by all four. In this scenario, older insurance policies would assume one party is clearly at fault because they were not drafted for a machine whose behavior is learned rather than pre-programmed and can change overnight with a software update. In our agentic insurance piece, we followed a single policy through five sets of hands between the risk and payout. Here, five sets of hands reach for a single claim.*



None of this troubled the pre-programmed industrial robots that have been operating on manufacturing lines since the 70s. An industrial arm sits in a cage under a risk assessment, so an injury is a procedural failure, a guard defeated or a lockout skipped, never a judgment about what the machine decided for itself. The most current example of a machine already acting on its own authority is on the road, and the UK legislated instead: the [Automated and Electric Vehicles Act](#) (AEVA) makes the motor insurer pay the victim first and recover from the OEM later, and its [2024 successor](#) obliges OEMs to disclose the safety data that settles fault.^{6,7} Both routes take behavior out of the question, which is why no policy form in use was ever drafted to ask about it.

Even if legislation settles who pays, an underwriter still has nothing to price with. Insurers rate commercial risk on claims history, equipment records and operating controls. Humanoids have almost no claims history, and the technology moves faster than a loss triangle can be built.⁸ The same model also carries different risk-profiles in different environments, and the software can change after the policy is bound. A single model running a fleet is not a pool of independent risks: **it is a single set of weights replicated across every insured unit simultaneously**. This is

* In a US facility, the worker is paid first by no-fault workers' compensation, which only relocates the fight: the comp carrier subrogates against whichever of the four it can reach.

the correlation problem we described for digital agents, but now with boots on the ground in the real world.

For a digital agent, an underwriter can at least ask for an independent model evaluation and a certification of the deploying organization. The embodied equivalents test the wrong layer: [UL 3300:2025](#), [ISO 3691-4](#), [ISO 13482](#) and the [ANSI/A3 robot safety standards](#) check mechanical strength, electrical safety, emergency stops and contact forces.^{9,10,11,12} The frame is moving: [ISO 10218:2025](#) shifted certification from the machine to the application; and [ISO 25785-1](#) is being drafted for dynamically stable walking robots; but both certify an installation, not a policy.^{13,14} Every one of them was designed for a machine whose behavior was fixed before it left the factory.

In a [previous robotics article](#), we argued that the emergent behavior making these machines commercially valuable is the same property that makes them unverifiable: you cannot enumerate what a generalizing system will do, only watch it. You can certify a fall zone; you can't certify improvisation, yet. [A body test is a one-time pass, but a brain test has to be a continuous measurement](#), which is a different instrument with a different business model attached.

The Referee

So, who builds it? Standards tend to follow units, and the units are not in the West. Chinese OEMs shipped over 90% of the world's humanoids in 2025, and the insurance followed the units.¹⁵ In October 2025, [CPIC](#) launched [Ji Zhi Bao](#): the world's first policy written specifically for a humanoid robot, covering the machine from the production line through leasing to end use, and covering harm in both directions: damage to the robot, and damage the robot does to everyone around it.¹⁶ [PICC](#) and [Ping An](#) followed within weeks, bundling body damage, cyber incidents and surgical robot accidents into cover priced by scenario.¹⁷

But Chinese insurers have not solved the pricing problem; they have almost been excused from it. Shortly after CPIC, Ping An and PICC released their humanoid cover, the Chinese [Ministry of Industry and Information Technology](#) (MIIT) published the country's first national standard system for humanoid robots and embodied intelligence, so the state underwriting the rollout is also writing the rules it will be measured against.¹⁸ Policy substitutes for actuarial science where the standard and the balance sheet sit under one roof. It is a real answer, but nobody else will adopt it: [no global buyer will accept a Chinese state insurer, underwriting against a Chinese national standard, as the neutral authority for a machine built in Austin or Sunnyvale.](#)

In the West, the picture is thinner. [The Artificial Intelligence Underwriting Company](#) (AIUC) has built the closest thing to the business this piece is describing: [AIUC-1](#) is audited by an accredited outside firm, retested quarterly, and an insurer can write a policy directly off the result.¹⁹ [ElevenLabs](#) and [UiPath](#) are among the first certified.^{20,21} But AIUC-1 is built for software; its recent updates have reworked

agent identity, logging and runtime containment, but nothing in it addresses a machine that can throw someone down a flight of stairs.

The seat AIUC is claiming for digital agents is empty for embodied ones. Around it sits carriers rather than certifiers. [AXIS](#) runs a program for firms that build and deploy autonomous robots, pairing technology errors and omissions with product liability against navigation and perception failures. [Relm's PONTAAI](#) covers bodily injury and property damage from AI failures the main policy leaves out.²² At the edges, [Testudo](#) is assembling the loss and litigation dataset the market will run on, and [Boop](#) is aiming at the household, reaching future robot owners through trusts for what it calls nonhuman dependents, and betting that the regulatory gray area around humanoids is the moat.^{23,24} These are serious bets, and between them they are claiming appetite, distribution and data, but no one is claiming the thing that makes the risk priceable in the first place.

There is a precedent for what is missing. In the 1890s, electricity was the uninsurable technology of its day: *obviously valuable, obviously dangerous, and impossible to price*. The fix did not come from a carrier. In 1894 two fire underwriters' associations put an engineer and \$350 of equipment into a room above a Chicago fire station and called it [Underwriters Laboratories](#) (UL). It tested the devices, published the standard, issued the mark, and never sold a policy.²⁵ [The insurers funded the referee](#) but declined to own it, and all of a sudden, the wired building became insurable.

An embodied agent needs the same instrument pointed one layer up, at the model rather than the metal, and running continuously rather than once.

An Empty Seat

Nobody has built that yet, and the company that does will not be a carrier. In our upcoming work on the physical AI data layer, we argue that value accrues not to whoever sells the scarce input (data) but to whoever owns the loop that keeps manufacturing it. In insurance, we think continuous certification is that loop: certify, deploy, observe the failures, reprice, recertify. A certifier watching failures across many manufacturers accumulates a record no single insurer could assemble from its own book, and every turn makes the next certificate more accurate.

This only works off the balance sheet. A firm that certifies the robot, owns the failure record and sells the policy has every reason to rate the risk in whatever way sells more insurance. This is the conflict that discredited the credit rating agencies, and the open question hanging over AIUC, which writes the standard, audits against it, and sells what the certificate enables. That is the harder structure to build, and the only one that holds.

The strongest argument against all of this is vertical integration from robot OEMs themselves. [Figure](#), [Tesla](#) and [1X](#) build the body, train the model and deploy the unit, so fault that would otherwise split four ways lands on a single balance sheet,

and a manufacturer with enough capital can self-insure its fleet and price it into the lease. A manufacturer standing behind its fleet has actually become an insurer whether it uses the word or not, and it still needs to price the exposure and know how its machines fail.

We've already seen this experiment run on cars. In 2019, Tesla launched [Tesla Insurance Co.](#) after losing sales due to the high premiums offered to their customers by third-party insurers.²⁶ In 2025, Tesla wrote ~\$1.4B in premiums at a direct loss ratio above 100. Tesla directly underwrote 75.9% of those, up from 21.2% in 2022, leading them to reach their first underwriting profit in early 2026.²⁷ **Without a balance sheet the size of Tesla's though, owning the data is not the same as being able to price the risk, and scoring your own product is not the same as an independent test.**

Vertical integration is also only one part of the story: the premise of the robotic foundation model (RFM) business is that the brain eventually *separates* from the body and ships onto hardware its developer did not build. The moment it does, fault splits again, and the certified version of the model is the only thing an underwriter has left to hold onto.

So, we think the insurable market arrives either way. If cyber was the dress rehearsal for digital agents, then digital agents are the dress rehearsal for this: risk too correlated and too uncertain to price gets engineered into something capital will fund.

Somebody is going to own the mark that makes a robot insurable and the failure data behind it. **That seat is still empty**, whoever claims it first will set the terms for everyone who comes after.

¹ JP Morgan, SMid Cap Industrials: Humanoids in the U.S, 2026

² Wright T.P., [Factors Affecting the Cost of Airplanes](#), 1936

³ Activant analysis of various company disclosures as of Q3 2026

⁴ Unitree, [Product Catalog](#), 2026

⁵ China Daily, [Insurance policy for humanoid robots](#), 2025

⁶ UK Government, [Automated and Electric Vehicles Act](#), 2018

⁷ UK Government, [Automated and Electric Vehicles Act](#), 2024

⁸ PYMNTS, [Insurers Face a New Risk From Autonomous Robots](#), 2026.

⁹ Standard Council of Canada, [Service, Communication, Information, Education and Entertainment Robots – SCIEE Robots](#), 2025

¹⁰ International Organization for Standardization, [Industrial trucks — Safety requirements and verification Part 4: Driverless industrial trucks and their systems](#), 2023

-
- ¹¹ International Organization for Standardization, Robots and robotic devices — Safety requirements for personal care robots, 2020
 - ¹² American National Standards Institute, ANSI/A3 R15.06-2025 / ANSI/A3 R15.06-3-2025, 2025
 - ¹³ International Organization for Standardization, Robotics — Safety requirements Part 1: Industrial robots, 2025
 - ¹⁴ International Organization for Standardization, Robotics — Safety requirements for dynamically stable industrial mobile robots (legged, wheeled, or other forms of locomotion), 2025
 - ¹⁵ Omdia, General-purpose Embodied Intelligent Robots, 2026
 - ¹⁶ China Pacific Insurance (Group) Co., ANNOUNCEMENT OF AUDITED ANNUAL RESULTS FOR THE YEAR ENDED 31 DECEMBER 2025, 2026
 - ¹⁷ China Daily, Insurance policy for humanoid robots, 2025
 - ¹⁸ Ministry of Industry and Information Technology, Humanoid Robot and Embodied Intelligence Standard System (2026 Edition), 2026
 - ¹⁹ Artificial Intelligence Underwriting Company, AIUC-1: The standard for AI agent security, safety and reliability, 2026
 - ²⁰ ElevenLabs, ElevenLabs secures first-of-its-kind AI Agent insurance, 2026
 - ²¹ UiPath, UiPath Achieves AIUC-1 Certification, Setting New Standard for AI Agent Security and Reliability, 2026
 - ²² Relm, Relm's PONTAAI Solution: AI Insurance Coverage Beyond Existing Liability Programs, 2026
 - ²³ Testudo, Testudo Launches New Insurance Coverage for Liability Risks Created by Generative AI Systems, 2025
 - ²⁴ Boop, Every household shift creates a new insurance category, 2026
 - ²⁵ Underwriters Laboratories Inc., Our History, 2026
 - ²⁶ Reuters, Tesla launched its own car insurance. These drivers say it's a lemon., 2019
 - ²⁷ S&P Global, Tesla's insurance premiums jump over 40% to \$1.37B in 2025, 2026

Disclaimer: The information contained herein is provided for informational purposes only and should not be construed as investment advice. The opinions, views, forecasts, performance, estimates, etc. expressed herein are subject to change without notice. Certain statements contained herein reflect the subjective views and opinions of Activant. Past performance is not indicative of future results. No representation is made that any investment will or is likely to achieve its objectives. All investments involve risk and may result in loss. This newsletter does not constitute an offer to sell or a solicitation of an offer to buy any security. Activant does not provide tax or legal advice and you are encouraged to seek the advice of a tax or legal professional regarding your individual circumstances.

This content may not under any circumstances be relied upon when making a decision to invest in any fund or investment, including those managed by Activant. Certain information contained in here has been obtained from third-party sources, including from portfolio companies of funds managed by Activant. While taken from sources believed to be reliable, Activant has not independently verified such information and makes no representations about the current or enduring accuracy of the information or its appropriateness for a given situation.

Activant does not solicit or make its services available to the public. The content provided herein may include information regarding past and/or present portfolio companies or investments managed by Activant, its affiliates and/or personnel. References to specific companies are for illustrative purposes only and do not necessarily reflect Activant investments. It should not be assumed that investments made in the future will have similar characteristics. Please see "full list of investments" at activantcapital.com/companies/ for a full list of investments. Any portfolio companies discussed herein should not be assumed to have been profitable. Certain information herein constitutes "forward-looking statements." All forward-looking statements represent only the intent and belief of Activant as of the date such statements were made. None of Activant or any of its affiliates (i) assumes any responsibility for the accuracy and completeness of any forward-looking statements or (ii) undertakes any obligation to disseminate any updates or revisions to any forward-looking statement contained herein to reflect any change in their expectation with regard thereto or any change in events, conditions or circumstances on which any such statement is based. Due to various risks and uncertainties, actual events or results may differ materially from those reflected or contemplated in such forward-looking statements.