

AI Underwriting

Where AI actually changes underwriting — from insurance forms to investment diligence.

AI underwriting is usually told as an insurance and lending story: machines scoring policies, pricing loans, clearing applications faster th

Where it actually changes

AI underwriting is usually told as an insurance and lending story: machines scoring policies, pricing loans, clearing applications faster than a human desk could. That story is real, and it is most of what gets written. But the same shift is arriving in investment diligence, and almost no one is naming it. This page sets out what AI underwriting is, how it works, and where the conversation goes once underwriting stops meaning "approve a policy" and starts meaning "decide whether a deal deserves capital."

What AI underwriting is

AI underwriting is the use of machine learning, and increasingly agentic AI, to evaluate risk inputs (documents, structured data, third-party signals) in order to support or automate an accept, decline or price decision. That is the working definition, and it holds whether the thing being underwritten is an insurance policy, a consumer loan or a private transaction.

The core is risk evaluation at machine scale. A traditional underwriter reads a file, weighs the evidence against a set of rules and a body of experience, and reaches a judgement. AI underwriting compresses the reading and the first-pass evaluation: a model ingests the inputs, extracts what matters, and produces a structured assessment a person can act on. In its lighter form it augments the underwriter; in its more automated form it clears routine cases and routes the hard ones to a human. What it does not do is remove judgement: the decision to accept risk, and on what terms, still rests with a person and a governed set of controls.

How AI underwriting works

In practice, AI underwriting follows a consistent three-stage arc.

The first is document triage. Underwriting runs on paperwork (applications, statements, contracts, reports), and the first job is to read and organise it. Models extract the relevant fields from unstructured documents, flag what is missing, and turn a pile of files into structured inputs. This is where most of the early time saving comes from, replacing hours of manual reading with a first pass measured in minutes.

The second is risk evaluation. With the inputs structured, the model assesses them against the relevant risk factors: scoring, pattern detection, comparison against historical cases, and the surfacing of anomalies a tired reviewer might miss. The output is an assessment, not a verdict: a ranked and evidenced starting point for the decision.

The third, and the newer one, is decision support and agentic automation. Where the first wave of these tools answered questions, the current wave acts: agentic AI executes multi-step tasks across systems, pulling data, cross-checking it and drafting the assessment, with limited supervision. For low-risk, high-volume cases that can mean an automated decision within a governed boundary; for everything else it means a far more complete brief landing on a human's desk. Across industries the sequence holds: triage, evaluate, decide.

Where the conversation stops: insurance and lending

Read almost any current explainer on AI underwriting and the setting is the same: insurance, mortgages, consumer credit. The authors are insurtech and lendtech vendors describing their products, insurers announcing tools, and consultancies framing the transformation for those markets. It is a coherent, well-covered conversation, and for those industries exactly the right one.

But it stops there. The entire public discussion treats underwriting as the scoring of a policy or a loan, a high-volume and relatively standardised decision about a known type of risk. That is fair as far as it goes: it is where AI underwriting was deployed first, because it is where the volume and the structured data already live.

What the framing leaves out is the more interesting part. Triage the documents, evaluate the risk, support the decision: that sequence is not confined to policies and loans. It describes any process where evidence has to be assembled and weighed before capital is committed, and there is one large domain where the description fits precisely, yet where the conversation has barely started.

Where it goes next: AI in investment underwriting

Underwriting, at its root, is not an insurance word. It is the act of deciding whether something deserves capital, and on what terms. Reframed that way, it is the central question of investment itself, and AI underwriting becomes far more than a faster claims desk: it becomes the machine-assisted assessment of whether an asset or a deal merits an allocation.

In investment and private-markets diligence, the arc maps almost one to one. Deal sourcing replaces application intake: AI screens a far wider field of off-market opportunities than a deal team could canvass by hand. Diligence replaces document triage, except the documents are a data room and the reading runs to weeks. Evaluation replaces risk scoring, against the questions a private transaction turns on. This is the territory of AI in private credit underwriting, where the decision is whether a borrower and a structure justify the capital and the terms, and where the quality and speed of that assessment is the whole competitive question.

Allocator-side credit underwriting, fund and manager underwriting, deal-level diligence: none of it appears in the insurance-and-lending literature, yet all of it is underwriting in the original sense. It connects directly to adjacent private-markets work, from [private credit](#) as an asset class to [tokenised private credit](#), where what an on-chain instrument represents must be verified before it can be underwritten at all.

From chatbot to operating system

Here the distinction that matters is between AI as a feature and AI as infrastructure. Most of what is marketed as AI underwriting is a capable model bolted onto an existing workflow, still a tool a person reaches for; the harder version rebuilds the process around the system, so the assessment is produced by the machine and signed off by the human.

This is where IMS Group, a technology-forward private markets investment group and partnership of family offices, has built rather than bought. Through its in-house lab, IMS Labs, the Group runs Cortex: a 41-agent underwriting engine (an "Agent Mesh") that serves as the private-credit intelligence layer beneath its diligence. On the public figure IMS Labs states, Cortex compresses a two-week diligence cycle to roughly 90 minutes end-to-end. The point is not the headline time but what it implies: underwriting has been re-engineered as an operating system, not added as a chatbot.

Cortex is one of three platforms IMS Labs runs in production, alongside Onyx (the patent-pending authentication and truth layer for real-world assets, with digital twins and tokenised ownership) and Nimbus (the agentic commerce layer). For the underwriting question, Cortex is the relevant one: it is the operationalisation proof the vendor conversation asserts but rarely shows, set out by [IMS Labs](https://imslabs.ai) (imslabs.ai) and examined further alongside [agentic AI as an operating system](#).

Augmentation, not replacement

The question the public conversation keeps returning to is whether AI is going to replace the underwriter. Posed as a yes-or-no it produces a defensive answer. The more useful framing treats this as a reallocation of judgement rather than its removal: more of the work shifts to the model, while the decision itself stays with a person.

What AI absorbs is the triage and the first-pass evaluation: the reading, the extraction, the routine scoring that consumes most of an underwriter's hours and least of their skill. What it does not absorb is the genuinely hard decision: the ambiguous case, the structure with no clean precedent, the judgement call where the evidence points two ways. As the machine takes the volume, human attention moves up the difficulty curve toward exactly those decisions; the role does not disappear, it concentrates on the part that always carried the skill. With only a small minority of firms having scaled an AI agent into live production in any function, the near-term reality is augmentation in most places and full automation in few.

Governance and explainability as an investment-grade requirement

The risk sections in most AI underwriting explainers read as a compliance checklist: model bias, data quality, the EU AI Act, the need for explainability. The items are right, but the framing undersells them. For an allocator, explainability and auditability are not a box to tick; they are the condition under which the output can be trusted at all.

An underwriting decision that cannot be explained cannot be governed, and a decision that cannot be governed cannot carry capital. That is why a credible operator can show where a person reviews and overrules the model, how its reasoning is documented, and how the assessment would stand up to an audit. Cortex reflects this in its design: the agent mesh assembles the investment-committee decision package, and a human investment committee reviews and signs it off, so the machine produces the package and a person retains the decision. Investment-grade governance is not an add-on to AI underwriting; it is the evidence that the capability is real rather than asserted.

Where this leaves underwriting

The underwriting shift is not an insurance footnote. The insurance-and-lending coverage describes the first place AI underwriting was deployed, not the limit of where it applies. Read in its original sense, as deciding whether something deserves capital, underwriting is how investment decisions are made, and that is where the machine-assisted version is now arriving. To see how an operator approaches it in practice, read across the AI-in-investing cluster, from the agentic shift to how the firm runs its own lab.

This article is provided by IMS Group for information purposes only and does not constitute investment advice, an offer, or a solicitation. Figures are sourced and dated where cited; forecasts are projections and stated as ranges.

Frequently asked questions

What is AI underwriting?

AI underwriting is the use of machine learning and agentic AI to evaluate risk inputs (documents and data) and support or automate an accept, decline or price decision. It is most established in insurance and lending, where it triages paperwork, scores risk and clears routine cases. The same approach extends to investment diligence, where it assesses whether a deal deserves capital.

Can AI do an underwriter's job?

AI can do much of the work an underwriter does, but not the whole job. It excels at the high-volume parts, such as reading documents, extracting fields and first-pass risk scoring, and can decide routine, low-risk cases within governed limits. The harder judgements, ambiguous cases and decisions without clean precedent remain with the human underwriter, supported by, not replaced by, the model.

Is AI going to take over underwriting?

It is unlikely to take over underwriting wholesale. The realistic path is augmentation: AI absorbs triage and routine evaluation, while human judgement moves to the genuinely difficult decisions. Industry research finds only a small minority of firms have scaled AI agents into production, so full automation remains the exception, and human oversight stays essential wherever a decision carries capital.

Is AI replacing underwriting?

No. It is reallocating it. The routine reading and scoring that consumed most underwriting time is increasingly automated, but the decision to accept risk, and on what terms, remains a governed human call. In investment underwriting especially, where each deal is distinctive, AI assembles and evaluates the evidence faster, while a person retains the judgement and the accountability for the outcome.

Sources & important information

1. IMS Labs (imslabs.ai), verified 2026-06-
 2. Cortex — "41 total deployed" / "41 specialist agents across 3 tracks"; "Agent Mesh"; "Private Credit Intelligence Layer"; "2 Weeks to 90 Minutes — End-to-end underwriting time". Onyx — "patent-pending" provenance, "digital twins", tokenised ownership. Nimbus — "the agentic ecosystem for global commerce". Public-site facts only; no accuracy %, methodology, AUM, returns or client specifics. IMS Labs.
 3. McKinsey & Company (2025). Superagency in the workplace — finding that only a small share of organisations have scaled an AI agent into production in any function; used here qualitatively to frame augmentation as the near-term norm. Published 28 January
 4. McKinsey & Company.
 5. European Commission — EU AI Act (Regulation (EU) 2024/1689, in force 2024). Referenced as the governing framework behind explainability and auditability requirements for AI in risk decisions. European Commission.
- Point-in-time figures are volatile and should be re-verified on the date of publication.

Important information

This document is provided by IMS Group ("the Group") for general informational and educational purposes only. It does not constitute investment advice, an offer to sell, or a solicitation of an offer to buy, any security or investment product, and should not be relied upon as the basis for any investment decision.

Forecasts and projections are illustrative only and are not guarantees of future results. Past performance does not predict future returns. Point-in-time figures cited are volatile and should be re-verified on the date of publication. © 2026 IMS Group. All rights reserved.

AI, read from the buy-
side — operating layer,
not chatbot.

IMS Group

info@groupims.com

groupims.com

IMS Group operates across investment management, venture building, sports finance and proprietary technology — IMS Capital Management, IMS Digital Ventures, IMS Labs and IMS Sports Capital. This brochure is for information purposes only and does not constitute investment advice or an offer. © 2026 IMS Group.