



InEvo AI

ROI of AI in Commercial Insurance

Where the value comes from, how to measure it and what to ask before you invest

A practical guide for
MGAs, wholesalers, brokers and
commercial insurance teams

Executive Summary

Published AI returns in insurance were earned in conditions a small firm does not share. Enterprise deployments and national broker rollouts run with dedicated project teams, clean data warehouses, integration budgets and two-year horizons.

A seven-person brokerage has one person who understands the management system, no integration budget and a six-week attention span. A number earned in the first setting predicts very little about the second.

The measurement problem sits underneath it. Vendors report the return that is easiest to instrument, which is time saved. Time saved is also the return a small firm values least, since a seven-person shop rarely converts saved hours into lower payroll. Those hours become more business or nothing at all, and which one they become depends on the pipeline, not the software.

The result is a principal reading numbers earned somewhere else, measuring the thing that matters least, against a baseline they have never set for their own shop.

Everything after this page is instruments. Measure your own numbers first. Then judge anyone else's.

Research from BCG found that AI work with U.S. and U.K. commercial P&C insurers improved efficiency in complex lines of business by up to 36%. BCG also found potential for up to a 3-point improvement in loss ratios through better use of data in underwriting.



AI ROI Starts With What The Team Does With The Time It Saves

AI can reduce the time people spend on data entry, document review, validation and other manual tasks.

Time saved does not automatically create financial value.

An account manager who saves 30 minutes on a submission still has 30 minutes of working time. The business gets value when the team uses the time for work that matters.

An MGA might use the capacity to handle more submissions. An underwriter might spend more time assessing risk. An account manager might respond to brokers faster. A leadership team might support more business without adding staff at the same rate.

The business case should answer five simple questions:

1. What work does AI remove?

Look for repetitive work such as data extraction, re-keying, document checks and file preparation.

2. How much capacity does it create?

Measure the actual time saved per submission or transaction.

3. Where does the capacity go?

Identify the work people will do with the time they recover.


4. What business result should follow?

Look at submission volume, turnaround, quote volume, hit ratio, premium or retention.

5. Can the result be measured?

Set the baseline before the AI pilot starts.

A Simple ROI Model



Insurance Is Moving from AI Pilots to AI at Scale

Insurance companies have invested heavily in AI pilots.

Scaling those pilots into daily operations remains difficult.

A 2025 BCG study found that only 7% of surveyed insurance companies had successfully brought AI systems to scale. About two-thirds were still in the pilot stage. [BCG, 2025](#)

A 2025 Roots survey found 22% of insurers had AI in production, while 45% were still exploring AI. The survey covered insurers at different stages of adoption, so these figures should not be compared directly with BCG's 7% figure. [Roots, State of AI Adoption in Insurance 2025](#)

The numbers point to one issue: moving from a successful demonstration to a reliable operating process takes work.

Commercial insurance has a clear example.

McKinsey's 2026 research on commercial and specialty underwriting describes workflows that still rely on broker submissions, email, document extraction, missing information requests and manual movement between work stages. [McKinsey, 2026](#)

These tasks can take time before an underwriter makes a decision.

THE QUESTION FOR MGAs

The question is not: Where can we add AI?

The better question is: Where are skilled insurance people spending time on work that does not require skilled insurance judgment?

The answer often starts with intake, data extraction, validation, appetite checks and routing.

The Biggest Opportunity May Sit Before the Underwriting Decision

Commercial insurance requires expert judgment. A large amount of work happens before an underwriter can apply that judgment.

A submission may arrive as an email with an application, schedules, loss runs, supplemental forms and other documents.



Someone then needs to:

- Identify the documents
- Organize the submission
- Extract key information
- Check for missing information
- Compare data across documents
- Enter information into systems
- Check appetite
- Route the opportunity
- Prepare the file for underwriting
- Follow up with the broker when information is missing

The underwriter receives the file after much of this work is complete. AI can support several of these steps.

McKinsey's June 2026 research describes an AI-enabled underwriting model with many of the same capabilities. The model uses AI for intake, extraction, triage and routine workflow tasks while people retain responsibility for underwriting judgment. McKinsey, 2026

The original research behind this paper identifies appetite triage, SOV normalization and routing clean submissions to underwriters as key opportunities for MGAs.

WHERE THE RETURN CAN APPEAR

WORKFLOW	AI OPPORTUNITY	POSSIBLE BUSINESS RESULT
 Intake	Organize submissions	 More capacity
 Extraction	Capture data	 Less manual work
 Validation	Find missing or conflicting data	 Fewer follow-ups
 Appetite	Check guidelines	 More focused underwriting
 Routing	Send work to the right person	 Faster response
 Quote preparation	Build a decision-ready file	 Better turnaround
 Renewal	Prepare information earlier	 More client-facing time

AI ROI Grows as the Impact Moves Closer to the Business

AI can create value at several levels.
Each level answers a different business question.

01 | EFFICIENCY

How much manual work did we remove?

AI can reduce work such as data extraction, re-keying, document checks and file preparation. The source material cites a benchmark of about 4.8 hours for a submission handled manually, based on Reagan Consulting's 2024 Agency Operations Benchmark. [Reagan Consulting](#)

Efficiency is the easiest return to measure. It is also only the starting point.

02 | CAPACITY

What can the team handle now?

If each person spends less time on manual work, the team may be able to handle more submissions, renewals or underwriting activity. Capacity is where time savings can start to affect the business.

03 | RESPONSIVENESS

How quickly does work move?

Faster intake and preparation can reduce the time between submission receipt and client or broker response. Turnaround gives leaders a clear operational metric to track.

04 | COMMERCIAL PERFORMANCE

Does the extra capacity produce business?

Measure:

- Submissions handled
- Quotes produced
- Hit ratio
- Premium written
- New business
- Renewal performance

05 | LONG-TERM VALUE

Does the improvement continue?

Better workflow performance can support stronger service, more capacity and better use of internal data.

For MGAs, proprietary data can become more useful when AI helps structure and analyze it. McKinsey's 2026 research identifies proprietary data as a potential source of differentiation for MGAs using AI. [McKinsey, 2026](#)

THE KEY DISTINCTION

Efficiency measures what AI saved. Capacity measures what the team can now handle. Commercial performance measures what the extra capacity produced. Long-term value measures what continues.

Before You Evaluate AI, Measure Your Current Workflow

A strong AI business case starts with your own numbers. You do not need a large consulting project. You need a small set of measures that show where the process stands today. The original framework identifies seven useful measures.

THE SEVEN NUMBERS

METRIC		DEFINITION	WHAT IT MEASURES
01	 SUBMISSION VOLUME	How many new business and renewal submissions do you send to market?	Use it to measure: workload.
02	 HANDLING TIME	How much active staff time goes into each submission?	Use it to measure: manual effort.
03	 TURNAROUND	How long does it take from receipt to the first quote delivered?	Use it to measure: response speed.
04	 HIT RATIO	How many quoted opportunities bind?	Use it to measure: commercial performance.
05	 RETENTION	How well do accounts renew, measured by policy count and commission?	Use it to measure: recurring business value.
06	 TOUCHPOINTS	How many systems and portals does each submission pass through?	Use it to measure: workflow friction.
07	 INCOMPLETE SUBMISSIONS	How often do submissions need more information before they can move forward?	Use it to measure: data quality and avoidable rework.

Do this before you go into the implementation of the AI

A vendor demonstration can change what your team notices.

Measure the current process first. Record the date. Use the same definition throughout the pilot. Assign one person to own the measurement. The baseline then becomes the reference point for the investment decision.

Four Questions Every AI Vendor Should Answer

An ROI percentage does not tell you enough.
You need to know how the number was calculated.

01 | WHAT WAS THE BASELINE?

Ask:

- How the starting point was measured.
- How many submissions or transactions were included.
- Whether the measurement covered normal work or a selected sample.

A 70% improvement can look very different when the starting process is slow.

The source framework makes this point directly. A smaller improvement from a strong baseline may represent more real value than a large improvement from a weak baseline.

02 | WHAT HAPPENED TO THE SAVED HOURS?

This question connects productivity to ROI. If AI saves 500 hours, ask what happened next.

Did the company reduce costs?
Did it process more submissions?
Did underwriters spend more time on risk selection?
Did producers spend more time on clients?
The answer should show how time savings became business value.

03 | HOW COMPARABLE WAS THE CUSTOMER?

Ask about:

- Company size
- Submission volume
- Lines of business
- Starting workflow
- Data quality
- Implementation support

A result from a large enterprise may not apply to a smaller MGA.

A result from a clean data environment may not apply to an inbox full of PDFs and forwarded email chains.

04 | WHAT HAPPENED TO COMMERCIAL PERFORMANCE?

Ask for results beyond processing time.

Look at:

- Turnaround
- Submission volume
- Quote volume
- Hit ratio
- Premium
- Retention

Five Tests For an AI Solution

1 | WORKFLOW FIT

Can people use it inside their current process?

2 | DATA HANDLING

Can it work with real submissions?

3 | VERIFICATION

Can users trace important outputs back to the source?

4 | HUMAN REVIEW

Can people review, change or approve the result?

5 | ECONOMICS


Does the cost still make sense as volume grows?

VENDOR SCORECARD


Criterion		Weight
	Workflow fit	×3
	Messy-data handling	×3
	Verification and audit trail	×3
	Human oversight	×2
	Customer references	×2
	Implementation effort	×2
	Pricing model	×2
	Data portability	×2
	Feature depth	×1

Design The Pilot

DEFINE 4 THINGS



MATCH THE CLAIM TO THE TIMEFRAME



The source framework makes a clear distinction between short-term workflow measures and longer-term commercial measures.

A six-week pilot can show whether work moves faster. It cannot prove a change in retention.

USE REAL SUBMISSIONS

Run the pilot on normal work.

Include submissions from the sources that create the most manual effort. Clean examples can show what the product does when conditions are ideal. Real submissions show whether the process works in practice.

What Does the ROI Look Like in a Real Commercial Insurance Workflow?

AI creates business value when several small improvements connect. Consider a submission that arrives in an MGA's inbox.



	01 INTAKE	02 DATA EXTRACTION	03 APPETITE & COMPLETENESS	04 PRIORITIZATION	05 UNDERWRITING REVIEW	06 QUOTE-READY OUTPUT
 The work today	A team member opens the email, reviews the attachments and starts sorting the submission.	Staff read applications, SOVs, loss runs and other documents and enter information into systems.	Teams check whether the risk fits the program and whether required information is present.	The queue may depend on who sees an email first or who has time to review it.	Underwriters may spend time finding information before they can evaluate the risk.	Information may need to be entered again before it reaches the rating or policy system.
 InEvo's offering	InEvo reads incoming broker emails and identifies the documents and submission data.	InEvo extracts document data and structures it for the next stage of the workflow. Its platform includes document parsing for ACORDs, SOVs and losses.	InEvo runs appetite analysis and completeness checks. Submissions can be prioritized based on factors such as appetite fit, completeness and other program criteria.	InEvo creates a ranked, quote-ready queue and keeps the submission status visible from intake through quote-ready.	The system stages the submission for underwriting review while keeping the human in the decision process. The InEvo interface also shows why a submission was allowed, flagged or declined.	InEvo can push structured data to a PAS through an API, reducing the need for duplicate entry between the inbox and downstream systems.
 Value with InEvo	Less manual entry and less time spent organizing submissions.	Less re-keying and more usable data.	Underwriters spend less time on submissions that are incomplete or outside appetite.	The team can focus attention on the submissions most ready for action.	More underwriting time goes toward risk selection and judgment.	Less re-keying and a shorter path from submission to quote.

How The Value Adds Up



AI ROI is not created by one feature. It comes from improving the workflow from intake to quote.

A PRACTICAL WAY TO MEASURE IT

An MGA can measure the change at each stage.


Measure	Before AI	After AI	Business question
 Manual entry time	Baseline	Measure	? How much work was removed?
 Submission turnaround	Baseline	Measure	? How much faster does work move?
 Incomplete submissions	Baseline	Measure	? Are gaps found earlier?
 Underwriter review time	Baseline	Measure	? Is more time going to underwriting?
 Submissions handled	Baseline	Measure	? Did capacity increase?
 Quote volume	Baseline	Measure	? Did the extra capacity produce more quotes?
 Hit ratio	Baseline	Measure	? Did faster processing improve results?

In Simple Words

The core issue is this:

MGAs are expected to quote faster, underwrite more consistently, prove portfolio discipline to carriers, and give brokers a better experience — but their intake process is still mostly email, PDFs, spreadsheets, and manual judgment.

That is why application intake and quoting automation is becoming highly valuable for MGAs. The best opportunity is not just “extract data from PDFs.” The real opportunity is to build an intelligent **intake layer** that can:



Partner with INEVO AI for Smarter Intake Solutions



+1 9256406784



www.inevo.ai



amit.chabra@inevo.ai



New York, USA