

STEADILY.

AGENT RESEARCH • 2026

The lucrative AI plays most agents never run

Agents are using AI to save time, not to write business. Here is what 95 independent agents told us, and seven plays that put AI to work on your book.

81%

have used AI to look something up

13%

have used it to think about growing their book

0

named growing the book as their biggest time saver

PUBLISHED BY STEADILY

LANDLORD INSURANCE

Who did this survey, and how

Who we are

Steadily is one of the best rated landlord insurance services in America. They provide comprehensive insurance policies for real estate investors in all 50 states, covering property damage, liability, and loss of rental income.

Tens of thousands of independent insurance agents nationwide are appointed with Steadily to sell its landlord products.

Why did we do this?

Most of the AI guidance written for agencies is either vendor marketing or a warning to wait. We wanted to know what independent agents are actually doing with these tools day to day, where the tools fail them, and whether any of that work is producing new business. The seven plays in this paper come out of what the answers showed was missing.

How did we do this?

Steadily surveyed independent insurance agents in its agent network between August 27 and September 3, 2026. Ninety-five agents completed the survey. Respondents were drawn from agents who have quoted or placed business with Steadily, which is a convenience sample rather than a random sample of the independent agency channel; results should be read as directional.

Percentages for tool and task questions are calculated on the base of the 68 respondents who reported using AI at work, and are noted as such on each chart. Multiple-answer questions do not sum to 100%. Subgroup findings below n=30 are not reported.

Outside sources cited

Big "I" Agents Council for Technology, 2026 Tech Trends Report.

SIAA, retention by policy count.

Oldroyd and InsideSales lead response research, cited for direction only.

SOCi, 2026 Local Visibility Index, based on more than 350,000 business locations.


NAIC, Implementation of NAIC Model Bulletin: Use of Artificial Intelligence Systems by Insurers, state adoption map, 2026 status.

Ask an independent insurance agent in 2026 if they use AI, and you will usually get a yes. That's not in itself a surprise; roughly 1 in 5 humans on Earth use it. But ask what exactly they used it for, and you'll get a lot of similar answers.

In late August 2026, we surveyed 95 independent agents about how they actually use AI at work. Nearly three quarters of them use it. Almost none of them use it to grow anything.

FIGURE 1

What 95 independent agents told us



Base: 95 independent agents surveyed August 27 to September 3, 2026. The three task figures use the base of the 68 respondents who use AI at work. The final row is out of the 69 agents who named a single biggest time saver.

The habit is already built. It is pointed at the inbox instead of the pipeline.

That is the growth gap. AI has landed in the independent agency channel as a faster way to do the work that was already on your desk, but it hasn't resonated as a way to put new business on the books, even though it is perfectly capable of doing so.

This is not intended to be a scolding. Most of the AI content aimed at agencies right now says some version of "you are not ready yet, fix your workflows first." That may be true, and it is also useless if you are a producer trying to hit a number this quarter. Our read of the survey is more encouraging than that: the same agents who are cautious about AI have already built the habit of using it daily. The only problem is that it's pointed at the wrong work.

What follows is the data, an honest account of why the gap exists (accuracy problems are real and we are not going to pretend otherwise), and seven specific plays you can run with tools you already have, all designed to actually grow your business, not just save you time. Then, the parts of this job that should stay human, drawn from what your peers told us rather than from our opinion.

1 The gap, in numbers

Adoption of the new technology is not the problem.


71% vs 8%

43% of the agents we surveyed use AI regularly and another 28% use it once in a while, which is 71% in total. The Big "I" ACT 2026 Tech Trends Report found 8% of agencies describe their AI use as regular and strategic. The two questions are not measuring quite the same thing, but the distance between them is the distance between using a tool and running a strategy.

Roughly two thirds of the regular users are in a tool at least once a day. Now look at what all that engagement is spent on.

FIGURE 2

What agents actually used AI for in the last 30 days



Base: 68 agents who use AI at work. Multiple answers allowed. Plum bars are the only two revenue-facing tasks on the list.

Read that list top to bottom and you are reading a description of administrative relief. Search, drafting, document review, data entry. Useful work, genuinely; nobody is going to complain about getting hours in their day back. But 60% of the AI users in our survey did not touch a single revenue-facing task in the previous month: marketing, prospecting, or thinking about the book.

The sharpest version of the finding is what happened when we asked users to pick the one task AI saved them the most time on.

ZERO

agents who used AI used it to grow the book.

Figure 3. Base: the 69 agents who named a single biggest time saver. 45% named looking things up and 30% named drafting client emails. Those two answers accounted for 75% of all responses.

That is a real gain, and a remarkably small target for a technology that has absorbed this much attention.

Three quarters of the value is coming from search and email. None of it is coming from the pipeline.

2 Why the gap is there

Three things showed up in the data, and they reinforce each other.

Agents feel capable without feeling clear

We asked two separate questions on a five-point scale: how confident do you feel using AI tools, and how clear are you on what AI can and cannot do for an insurance agent specifically. Users rated their confidence at 3.7 and their clarity at 3.3. Across all 95 respondents, 23% put themselves at a 4 or 5 on confidence while sitting at a 3 or lower on clarity.

23% of agents rate their confidence using AI at a 4 or 5, while rating their clarity on what it can and cannot do for an agent at a 3 or lower. Base: all 95 respondents.

That is close to one in four agents who are comfortable driving without a clear sense of where the road goes. If you are confident with the tool but fuzzy on its boundaries, you will use it for the tasks where you can instantly verify the output. An email draft either sounds right or it does not. A cross-sell strategy takes weeks to prove out, and if you are not sure the tool is trustworthy, why would you point it at something you cannot check?

The accuracy problem is not imaginary

Here is the number that gives every skeptic in your agency ammunition, and it should be in this paper rather than buried: among agents who asked an AI tool a question about coverage, carriers, or regulations, 62.5% got a wrong answer at least once. More than half of those said it happened more than once.



FIGURE 4

Agents who asked AI about coverage, carriers, or regulations

62.5% were given a wrong answer at least once. The remaining 37.5% were not, or did not catch it.

Base: 56 AI users who used AI for that kind of question. 37% said it happened more than once.

You can respond to that number in two ways. One is to conclude the technology is not ready. The other is to notice that every one of those wrong answers came from a question the tool was badly suited to answer in the first place. General-purpose chatbots do not know your carriers' current appetite, they do

not have your state’s bulletin in front of them, and they will produce a confident paragraph anyway. AI chatbots are designed to give a response no matter what, even if they have no way of knowing the correct answer. This is what happens when AI “hallucinates.”


If you take away one thing from this paper: give the tool your own data to work with instead of asking it to recall the industry’s.

Nobody is being taught anything

61% of the agents we surveyed said their agency has offered no training, no policy, and no guidance on using AI at work. Another 28% described what they get as informal guidance only. Just 7% work somewhere with a formal written policy.

FIGURE 5

Has your agency offered any guidance on using AI at work?



Base: all 95 respondents. The Big “I” ACT 2026 Tech Trends Report found that 55% of agencies have no written AI policy, which lines up closely.

So the typical agent in this survey is self-taught, working in a general-purpose chatbot their agency did not provide (three quarters use ChatGPT; only 18% use an AI feature built into their AMS or CRM), with no guidance about what is in bounds. Under those conditions, sticking to email drafts is reasonable risk management, and the timidity around expanding usage is understandable.

One more data point from that same neighborhood, because it matters for section four: 22% of AI users admitted that something written by AI has reached a client without them editing it first, and 12% said that happens regularly. Untrained, unsupervised, and occasionally unedited is not a combination anyone should be comfortable with.

3 Seven plays


All right, time for the part you can actually use.

Each play below follows the same shape: what it is, how to run it, the guardrail, what to measure, and how long before you see something. Most of them work with the tools already in your survey answers. A few work better if your AMS has a decent export function, which nearly all of them do even if nobody at your agency has used it.

Pick one. Not seven. The agencies that get value out of this stuff are running one play well, not seven badly.

FIGURE 6

The seven plays at a glance



PLAY 1

Mine your own book

Your cheapest new business is already a client. That is not a new observation; the reason it stays unexploited is that finding the opportunities means reading a few thousand rows of policy data, and nobody has time for that.

The average insurance client holds somewhere around 1.4 policies. Agencies that work cross-sell deliberately push that past 2.5. SIAA data has clients with three or more policies retaining at roughly 94% annually against about 67% for single-policy clients, so this play buys you retention and revenue at the same time.

How to run it. Export your book from the AMS as a CSV. Strip out anything you would not want leaving your building: names, addresses, policy numbers, dates of birth. Keep client ID, line of business, effective date, premium, carrier, limits, and whatever risk fields you carry. Upload the anonymized file and ask for a ranked list of households or entities carrying one line where the profile suggests a second. For landlord business specifically, the query that pays is: which insureds have a rental dwelling policy but no umbrella, and which ones have more than one property in the file.

Then ask a second question, which is the one most people skip: for the top 25, draft me the opening line of a conversation that would not feel like a sales call.

PROMPT

Attached is an anonymized export of my book of business. Each row is one policy.

Find every client ID that holds only one line of business but whose profile suggests a second. Rank them by the size of the opportunity. For each one, name the line to pitch and the single field in that row that justifies the conversation.

Separately, flag every client ID that appears more than once. Those are multi-property owners.

Use only what is in the file. If a field is blank, say so rather than filling it in.

The guardrail. Anonymize before you upload, every time. If your agency has no written AI policy (61% of you), this is the play where that catches up with you fastest. Client data in a consumer chatbot account is a conversation you want to have before it happens, not after.

What to measure. Policies per client, and the retention spread between single-policy and multi-policy households.


PLAY 2

Answer in five minutes

Insurance leads get shopped. A prospect filling out a form at 7:40 on a Tuesday evening is filling out three of them, and the research on lead response, going back to the original MIT and InsideSales work, has been pointing the same direction for nearly twenty years: contact odds collapse within the first few minutes. The multipliers get quoted loosely in vendor marketing and we would not put much weight on the exact figures, but the shape of the curve is not in dispute.

FIGURE 7

Contact odds fall off a cliff, not a ramp



Illustrative curve. The shape is well supported by lead-response research; exact magnitudes vary by study and sample.

How to run it. The version of this play that does not require buying anything: keep a saved prompt that takes the raw content of an inbound lead form or email and returns a first response tailored to what the prospect actually said, including the two or three questions you will need answered before you can quote. Paste, skim, send. Thirty seconds instead of five minutes, which is the difference between doing it now and doing it after your 2 o'clock.

PROMPT

Below is a raw inbound lead. Write a reply of 120 words or less that does four things: greets them by name, refers to the specific property or situation they described, asks for the two or three items I still need before I can quote, and offers a time to talk.

Do not quote a price. Do not confirm or imply coverage. Do not promise anything. Plain professional tone, no exclamation points.

Lead: [paste the form submission or email here]

The guardrail. The first response acknowledges and asks; it does not quote, indicate coverage, or promise. Read every one before it goes out. Remember that 22% figure.

What to measure. Median minutes from lead arrival to first outbound touch. Most agencies have never measured this and are startled by the answer.

PLAY 3


Get recommended by AI

Here is the play nobody in this category is writing about, and it is the one with the longest tail.

Your prospects have started asking chatbots for recommendations. Not searching; asking. And the businesses that get named in those answers are a much, much smaller club than the businesses that rank on Google. SOCi's 2026 Local Visibility Index looked at more than 350,000 business locations and found ChatGPT recommended 1.2% of them, against 35.9% that turned up in Google's local three-pack. Roughly half the businesses winning the map pack do not exist as far as an AI answer is concerned.

FIGURE 8

Local businesses are close to invisible in AI answers



SOCI 2026 Local Visibility Index, based on more than 350,000 business locations.

Why the gap? An AI answer is not crawling your site in the moment. It is assembling a picture holistically from your Google Business Profile, your website's structure, third-party citations, and reviews, and if that picture is thin or inconsistent it will simply name somebody else. Being in business for thirty years does not help; being legible to a machine does.

How to run it. Four things, in order. Fill out every field of your Google Business Profile and update it monthly, because stale profiles lose visibility. Make your name, address, and phone identical everywhere they appear, down to the “Suite.” Add structured data to your site (LocalBusiness and FAQPage schema at minimum) so a machine can parse what you do and where. Then write genuine answer-first pages for the questions your prospects ask out loud: what does landlord insurance cost in Sarasota, do I need it for a duplex, does my HOA policy cover the interior. Short, specific, and answering the question in the first paragraph rather than the fifth.

Use an AI tool for the drafting; use it more usefully to audit. Ask three different chatbots for a recommendation for a landlord insurance agent in your city and see whether you come up. That is your baseline.

PROMPT

Run this in ChatGPT, Gemini, and one more, once a month. Log whether you appear and who does.

I own two rental houses in [city, state] and I need landlord insurance. Which local agents or agencies would you recommend? List five, and tell me what you based each recommendation on.

The guardrail. Do not keyword-stuff your business name (write it over and over again in a spammy way). It will get you flagged, and the AI systems reading the profile treat it as an untrustworthiness signal. It needs to read naturally and answer questions in context.

What to measure. How often you get named across a fixed set of ten prompts, checked monthly. Changes to listings usually take four to eight weeks to show up.

PLAY 4

Build your appetite cheat sheet

Something in our data that has nothing to do with AI but everything to do with wasted hours: 54% of agents told us that when they need to check whether a property fits a carrier’s appetite, they just run the quote and find out.

Running the quote is a fine way to find out. It is a slow one, and it means the carrier’s guidelines live in your head or nowhere.

How to run it. Collect the appetite guides, underwriting bulletins, and state availability documents for the six or eight carriers you actually place with. Drop them into a project or a saved file in whichever tool you use. Now you can ask, in plain language, which of these carriers will look at a 1912 four-plex in Ohio with knob and tube partially replaced, and get a defensible answer with the source document attached rather than a guess.

PROMPT

Use only the attached appetite guides and underwriting bulletins.

For the risk below, sort my carriers into three groups: would consider, would decline, and not addressed in these documents.

For every carrier you place in the first two groups, quote the exact line that supports it and name the document it came from. If the documents do not cover something, put that carrier in the third group rather than inferring an answer.

Risk: [property type, year built, construction, location, occupancy, loss history, anything unusual]

The guardrail. This works because you supplied the documents. Do not ask a general chatbot about carrier appetite from memory; that is the exact question that produced the 62.5% error rate. Re-upload the guides quarterly, because they change and the tool will not tell you they have.

What to measure. Submissions declined for appetite reasons, and quote-to-bind rate.

PLAY 5

Work the referral channel

For anybody writing landlord and investor property, the growth is not really in individual owners; it is in the people who touch fifty of them. Property managers, realtors who work investor clients, small REI meetups, 1031 exchange intermediaries.

The work that stops most producers from doing this well is research. Personalizing outreach to forty property management companies means learning something real about forty property management companies, and that is a week's worth of work you do not have.

How to run it. It happens to be almost exactly the thing AI is good at. Give it a list of firms and their websites and ask what each one specializes in, how many doors they appear to manage, and what a relevant opening would be for each.

PROMPT

Here is a list of property management companies and their websites.

For each one, give me three lines: the property types and rough unit count they appear to manage, the markets they cover, and anything on the site that hints at how they handle owner insurance requirements today.

Then draft one opening sentence per firm that refers to something specific you found on their site.

Do not invent details. If the site does not say, write "not stated."

[paste list]

The guardrail. The draft is a draft. Sending forty near-identical letters with the company name swapped in is worse than sending none; property managers talk to each other. Rewrite the first two sentences of every one yourself.

What to measure. Meetings booked per hundred touches, and referral partners producing at least one submission a quarter.

PLAY 6

Turn one review into five touches

A third of the agents in our survey already use AI for marketing material, which makes this the most adopted of the seven plays and probably the least systematically run.

The version that compounds: every time you do a real coverage review that produces a genuine insight, that insight is content. A landlord who discovered their policy excluded ordinance and law is a LinkedIn post, an email to every other landlord client, a two-minute video script, a line in next quarter's newsletter, and an FAQ entry on your site that feeds Play 3. One piece of thinking, five outputs, none of which required you to invent a topic.

How to run it. Hand over the insight once and ask for all five formats in a single pass, rather than going back for each one.

PROMPT

Here is a coverage insight from a client review, with every identifying detail removed:

[describe what you found and why it mattered]

Turn it into five pieces: a LinkedIn post under 150 words, a client email under 200 words, a 45-second video script, a two-sentence newsletter item, and an FAQ entry where the question is written the way a landlord would actually type it into Google.

Plain language. No exclamation points, no marketing voice, no "in today's market."

The guardrail. Scrub any detail that could identify the client, and keep your voice. Content that reads as generated will be treated as generated.

What to measure. Outbound touches per month, and inbound inquiries that cite something you published.

PLAY 7

Sell on terms, not price

We asked agents what decides the placement when two carriers are competitive on a landlord property with minimal price difference. Coverage terms won at 34%, ahead of price at 24% and ease of binding at 19%.

That should change how you present. If terms decide it, then the agent who can put a clear, side-by-side, plain-English comparison in front of a client wins deals that a price sheet loses.

The constraint has always been that building that comparison honestly takes an hour per proposal. With AI, it can take seconds.

How to run it. Feed both policy forms in and ask for the differences that would matter to this specific insured, in language a client will understand, with the form language cited so you can verify it. Then verify it, because this is a coverage question and coverage questions are where the tool is least reliable.

PROMPT

Attached are two landlord policy forms quoted on the same property.

Build a side-by-side comparison of only the differences that would matter to this insured: [describe the insured and the property].

For each difference, quote the exact language from both forms and cite the form number and section. Then explain what it means in plain English a client would follow.

Do not include any difference you cannot quote directly from the forms.
Do not rank the policies or tell me which to recommend.

The guardrail. You are the licensed professional. Every line of that comparison is your representation, not the tool's; check each one against the form before it goes anywhere near a client.


What to measure. Close rate on competitive accounts, and how often you win without being cheapest.

4 What stays human

We asked all 95 agents where human judgment matters most, and the answers were not close.

FIGURE 9

Where agents say human judgment still matters most



Base: all 95 respondents. Multiple answers allowed.

Notice the third bar. A third of agents named checking the accuracy of AI tools as a place where human judgment is essential, which is a slightly wry thing to find on a list of core professional responsibilities, and also correct.

None of the seven plays above asks a machine to give advice, hold a relationship, or make a binding representation. Every one of them either hands the tool your own data or asks it to do the reading you would have done anyway. That distinction is the whole governance conversation in a sentence, and it is worth writing down somewhere your team can see it, given that six in ten of you work somewhere with nothing written down at all.

Two other things worth knowing. Twenty-five states plus the District of Columbia have adopted the NAIC's model bulletin on insurer use of AI, and several more have their own frameworks. It governs carriers rather than producers, so it does not land on your desk directly, but it tells you where regulator attention is heading. Your more immediate exposure is E&O, and it is the same exposure it has always been: if it went out under your name, you said it.

5 The first 90 days

WINDOW	DO THIS	THE NUMBER YOU ARE WATCHING
Days 1 to 30	Write one page of ground rules: approved tools, what client data may never be pasted anywhere, and what always gets human review. Pick a single play. Run Play 4 first if you want a fast, low-risk win.	Whether the page exists and the team has read it
Days 31 to 60	Add one growth play. Play 1 if your AMS data is clean; Play 3 if you would rather build something that compounds. Baseline the metric before you start, not after.	Policies per client, or AI visibility across ten fixed prompts
Days 61 to 90	Review what the first two produced. Kill whatever did not move a number. Add a third play only if the first two are running without you supervising them.	New business written that traces back to a play

The thing most likely to derail this is doing all seven in week one, which produces a lot of activity, no baseline, and nothing you can prove worked.

GET APPOINTED

Write landlord business with Steadily

Landlord insurance in all 50 states, quotes in minutes, and no volume commitment to get appointed. Tens of thousands of independent agents already place business with us. Appointment takes about five minutes to request.

STEADILY.

steadily.com/agents