2026 Advanced Sales Forum

Rooted in Knowledge — Aligned for Growth

Are we there yet?

AI, Superintelligence, and the Dream of Full Automation

Nick Bowman, CFP®, CEP®, MS · Co-Founder, Xcela AI

LIMRA LOMA
Navigate With Confidence

One of these was written by AI in eleven seconds.

EXHIBIT A

“Under the economic benefit regime, the employer’s premium payments are measured annually using Table 2001 rates. See PLR 202699012, in which the Service confirmed that a switch-dollar restructuring of this design avoids income recognition under Section 61 at conversion.”

EXHIBIT B

“The loan regime works here because the AFR locks at issue and the client is already banking the spread. Watch the exit, though. If we roll to economic benefit later, the switch-dollar rules get ugly. Get Section 7872 sign-off from tax counsel before we illustrate anything.”

Hands up. Which exhibit is the machine?

LIMRA LOMA
Navigate With Confidence

Exhibit A is the machine.

PLR 202699012 does not exist.

Private letter rulings are numbered year, week, sequence. There is no week 99. The model invented authority, formatted it perfectly, and delivered it with total confidence.

The writing is not the problem. The confidence is.

Hold that thought. This exact case comes back at the end of the hour, done right.

Nick Bowman, CFP®, CEP®, MS

Co-Founder, Xcela AI

Advanced Markets, Lion Street · Top of the Table · Forum 400 · Expert Witness · Adjunct Professor

LIMRA LOMA
Navigate With Confidence

The four fundamental truths of AI.

QUICK START

01 Pattern recognition at scale

AI is pattern recognition plus language at scale. Everything else is marketing.

02 Value equals repetitiveness

The value of AI is proportional to repetitiveness, not complexity.

03 The biggest risk is human

The biggest risk is not AI doing something wrong. It is people trusting AI without verification.

04 The business model is the risk

The AI is not the risk. The business model built around it is.

The (very brief) history of AI.

QUICK START

1940s	1956	1997	2010s	2020s
Turing & Enigma	**“AI” is born**	**Deep Blue**	**Machine learning**	**Generative AI**
Foundation of computation	Dartmouth Conference	Beats Kasparov at chess	Data-driven AI emerges	GPT, Claude, Llama

“We are in the Netscape era of AI. Powerful, clunky, and about to change everything, but only for those who use it with intent.”

LIMRA LOMA
Navigate With Confidence

The gap is your advantage.

QUICK START

95%

theoretical AI capability in business and finance

~30%

actual adoption observed today

THE WHITE SPACE, that 5%, IS YOUR ADVANTAGE. The gap is closing fast. The time to get ahead is now.

Source: Anthropic, Labor Market Impacts of AI, March 2026.

Where the industry actually stands.

QUICK START

55%

of employers now regret laying off workers for AI.
Forrester, Predictions 2026: The Future of Work.

1 in 3

of the regretful spent more on restaffing than the cuts ever saved (Forrester)

58%

of life insurers report even current-or-planned AI use, the lowest of the major lines (NAIC)

AI without intent works too well at the wrong things.

Replacement failed. Adoption lags. The leverage sits with whoever masters the DO first.

Sources: Forrester, Predictions 2026: The Future of Work; NAIC industry surveys.

LIMRA LOMA
Navigate With Confidence

Same country. Different choice.

QUICK START

KLARNA · FIRE, THEN REHIRE

Replaced roughly 700 roles' worth of service work with a bot. Resolution time fell from 11 minutes to 2. Then customer satisfaction tanked, the CEO admitted it backfired, and the rehiring began. AI without intent optimizes for speed; your clients want accuracy and trust.

IKEA · AUTOMATE, THEN ELEVATE

Zero layoffs. The Billie bot absorbed about 47% of routine calls, 8,500 call-center workers were reskilled as paid design consultants, and satisfaction went up. The new remote channel earned €1.3 billion in its first full year. AI with intent creates new revenue.

Same technology. Opposite intent. The tool was never the decision.

A formula older than this industry.

QUICK START

BE

Your expertise, your judgment, your relationships

→

DO

AI-accelerated execution: drafts, comparisons, models, research

→

HAVE

Results: the case placed, the desk that scales, the practice

AI is the most powerful DO tool in 10,000 years of human history. But it cannot BE anything. It has no intention.

Your expertise IS the intent layer. AI amplifies the DO, but only if the BE is driving it.

LIMRA LOMA
Navigate With Confidence

Summary: Be → Do → Have

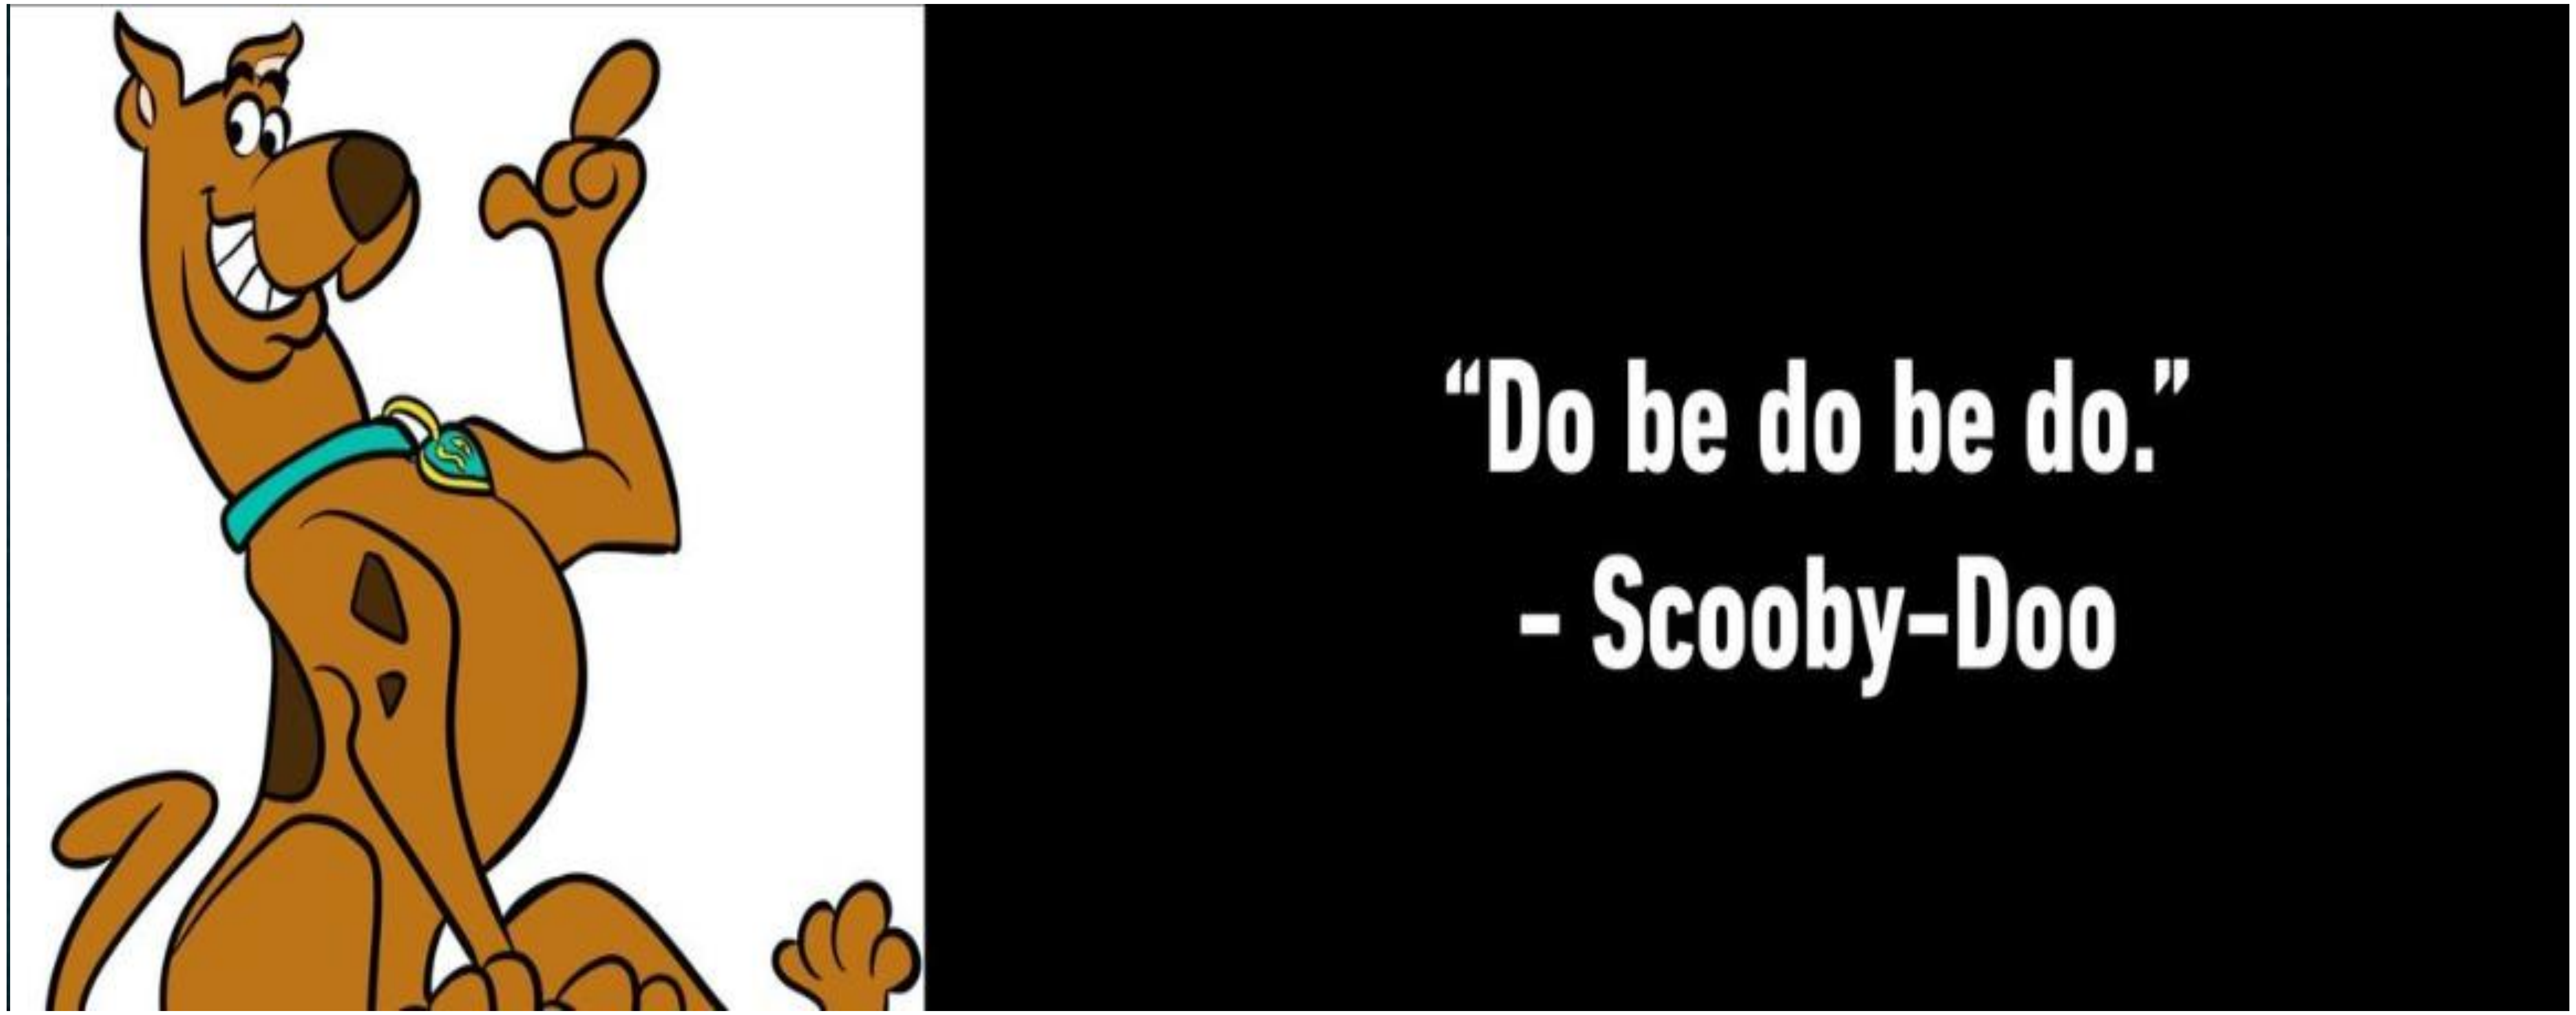

The Three Disciplines

QUICK START

3 INTENT — Tells AI what to WANT

“The reader is the client’s CPA. Conservative tone. Flag anything that risks lapse after age 80. Recommend only if the loan provisions support retirement income.”

2 CONTEXT — Tells AI what to KNOW

“Here are both full illustrations, the trust document, and the carrier product guides. Client is 54, funding $500K a year for ten years.”

1 PROMPT — Tells AI what to DO

“Compare these two VUL illustrations and summarize the differences for my client.”

The more specific your intent, the less you'll need to edit the output.

Why your first AI experience probably sucked.

QUICK START

The Gell-Mann Amnesia effect: you read an article on a topic you know well and see every error. Then you turn the page and read the next one as gospel. **Now imagine this, but with AI.**

Policy language

Ask it to summarize a whole life policy's payment options. It blends UL and WL terms seamlessly. Sounds great. Dead wrong.

Compliance

Ask it to draft a compliant client communication. It misses policy-specific features you would catch in ten seconds.

Underwriting

Ask it to assess risk on a complex case. It pattern-matches from training data. Your judgment comes from experience.

AI sounds confident. Fast. Articulate. And wrong in ways only YOU would catch.

Context rot: AI gets worse the more you dump into it.

QUICK START

0-25%	25-50%	50-100%
Sharp memory	**Forgetting zone**	**Context rot**
AI is sharp, remembers most	Earlier context starts being replaced	Important details dropped or distorted

If you do not manage context, your AI will sound confident, but be wrong.

The fix: start fresh conversations for new topics. Treat AI like a focused meeting with an agenda, not a running diary.

LIMRA LOMA
Navigate With Confidence

The quick-start pilot: five steps, thirty days.

QUICK START

1. IDENTIFY

Find the repetitive workflow

2. MAP

Inputs, outputs, who, how long

3. TEST

Run AI side by side. Do not replace yet.

4. MEASURE

Hours saved. Error rate. Quality.

5. EXPAND

Prove one before adding the next

The data question is not “is it clean enough?” It is “can AI see it?”

Notice what is not here: an AI team, an RFP, or a 12-month roadmap.

The blockers are integration, procurement, and trust problems. None of them are AI problems.

LIMRA LOMA
Navigate With Confidence

Step 1 — Identify What's Difficult

Not every task is an AI task. The right ones share four properties.

📋 **High Volume / High Repetition**
Done weekly or daily, same steps every time

📄 **Read, Compare, Summarize, Search**
Document-centric work without deep per-item judgment

⏱️ **Human Time But Not Always Human Judgment**
Currently requires a person — but shouldn't require expertise on every line

🗄️ **Data Exists (Even If Messy or Siloed)**
The information is somewhere — just not accessible or organized

WORKFLOW EXAMPLES

- In-force policy reconciliation across multiple admin platforms
- Illustration comparison and competitive analysis
- Regulatory filing review and change tracking
- Plan document extraction and summarization
- Mortality experience reporting and variance analysis
- Case underwriting file review for committee

Step 2 — Map the Data Path

The question isn't 'Is our data clean enough?' — it's 'Can AI see it?'

Structured + Accessible

Databases, APIs, CRM

→ AI-Ready NOW

Structured + Siloed

Mainframes, Legacy Oracle, Pascal

→ Needs Integration Layer

Unstructured + Accessible

PDFs, Emails, Scanned Docs

→ AI Handles with OCR

Unstructured + Siloed

Filing Cabinets, Tape Archives

→ Digitize First

Now, the use cases: your last two days, one desk.

CASE 1 · The Product Decision

Illustration forensics. Echoes Session 3.1, the art and science of case design.

CASE 2 · The Bank Board

Executive benefits capital math. Echoes Monday's GS2 on plans that actually work.

CASE 3 · The Rescue

A premium financing stress test. Echoes Session 2.2 and yesterday's policy outlook.

CASE 4 · The Exit Opinion

The split-dollar question from my opening, resolved with research that holds up.

Four different cases, one repeatable loop: ground, generate, validate. Along the way: the guardrails your compliance team needs, six myths retired, and a Monday plan.

The method: five steps, every case.

1. FRAME

What decision does this serve, and who reads the output?

2. GROUND

Feed the actual documents, not descriptions of them.

3. GENERATE

Demand structure: tables, citations, stated assumptions.

4. VALIDATE

Recompute, page-check, cross-model before it leaves your desk.

5. DELIVER

Translate to the client's language: memo, exhibit, script.

Watch for all five in every case that follows. The artifacts change; the method never does.

Steps 2 through 4 are the loop you will see again on the privacy slide: ground, generate, validate.

The spreadsheet outlives its author. Badly.

KEY PERSON RISK

The builder retires and forty tabs of logic retire with them. The formulas stay; the reasoning leaves.

SILENT BREAKAGE

APIs change. Feeds move. Input formats drift. The model keeps calculating, and nobody notices it is wrong.

THE LEARNING CURVE

New hires inherit a maze of hidden tabs, hardcoded cells, and macro chains. Months to trust it, longer to change it.

NO VERSION CONTROL

Final_v7_REAL_final.xlsx. Which one priced the case? Nobody can prove it.

NO AUDIT TRAIL

Who changed the crediting assumption, when, and why? Excel does not remember. Examiners ask.

BILLION-DOLLAR CELLS

A copy-paste error fed the London Whale VaR model. A dropped range skewed Reinhart-Rogoff. Cells move billions.

The fix is not a better spreadsheet. It is turning the logic into a tool with an owner, a log, and a test. AI makes that migration cheap.

Case 1 · The Product Decision

The file: a $500,000 annual premium, counsel in the room, and a producer who needs the real differences by Friday.

Shipped: a client-ready chassis memo with a page-cited provision audit.

Navigate With Confidence

The trend, not the snapshot.

PPLI VS RETAIL VUL · DOZENS OF ILLUSTRATIONS IN AN AFTERNOON · EVERY CELL PAGE-CITED

Provision	PPLI	Retail VUL
Premium load (current)	3.564%	6.00%
Net loan spread, years 1-10	0.20%	0.90%
Max policy loan	~90% of CSV	94%+ with rider
Overloan protection rider	Not available	Available
Asset fit	Alternatives, IDFs	Beta, index funds

THE CROSSOVER CURVE

Low loads win the first decade; asset-level drag can flip it later. AI charts value-to-cost by year, not one snapshot.

RIDER-TO-GOAL FIT

Retirement income through loans? Lapse protection past 80? Alternatives or beta? The chassis answer changes with the goal.

SCALE

Tens to hundreds of contracts and illustrations, read and reconciled in hours instead of weeks.

Advanced markets already knows these tradeoffs. The revolution is reading every contract, every year, before Friday.

Illustrative figures from the presenter's comparison workbook; provisions vary by carrier and contract. Verify rider names and loan limits per policy.

Same reader, every advanced sales desk.

Policy and illustration comparisons

Trust instruments and ILIT reviews

Buy-sell and plan document review

SERP, NQDC, and split-dollar plan modeling

Carrier contract provision checks

Case-file summaries before producer calls

Case 1 is one instance of a family. One workflow, a page citation in every cell, every time.

The recipe does not change with the document type. The grounding and the spot-check do the work.

Case 2 · The Bank Board

The file: a $10 million BOLI purchase, a skeptical board, and one slide to survive the CFO's capital question.

Shipped: a board-meeting exhibit and the return-on-capital talking points.

Navigate With Confidence

Out of the spreadsheet, into the picture.

FDIC DATA IS PUBLIC · THE API IS FREE · A THOUSAND CALL REPORTS BEFORE LUNCH

CAPACITY

Tier 1 capital × 25% concentration guideline, minus BOLI already held, equals purchase headroom. Computed from the bank's own call report in seconds.

PEER RANK

The client bank against every bank in its state: BOLI as a share of Tier 1, outliers flagged, neighbors named. Hundreds to thousands of banks, one pass.

YIELD AND ROC

BOLI's tax-equivalent yield against munis and loans, translated to return on capital: \$10M CSV × 4.0% tax-free ≈ \$400K, ≈50% ROC vs ≈18% on a commercial loan (illustrative).

The packet the board remembers is the one with their own numbers and their neighbors' names.

25% reference: interagency BOLI concentration guidance (OCC/FRB/FDIC, 2004). Data: FDIC BankFind Suite API and public call reports. ROC figures illustrative.

LIMRA LOMA
Navigate With Confidence

Case 2 · Data Insights in Minutes not Hours

SAME METHOD, PRODUCTIZED · THE SERP MODEL · ECHOES THIS WEEK'S EXECUTIVE BENEFITS SESSION

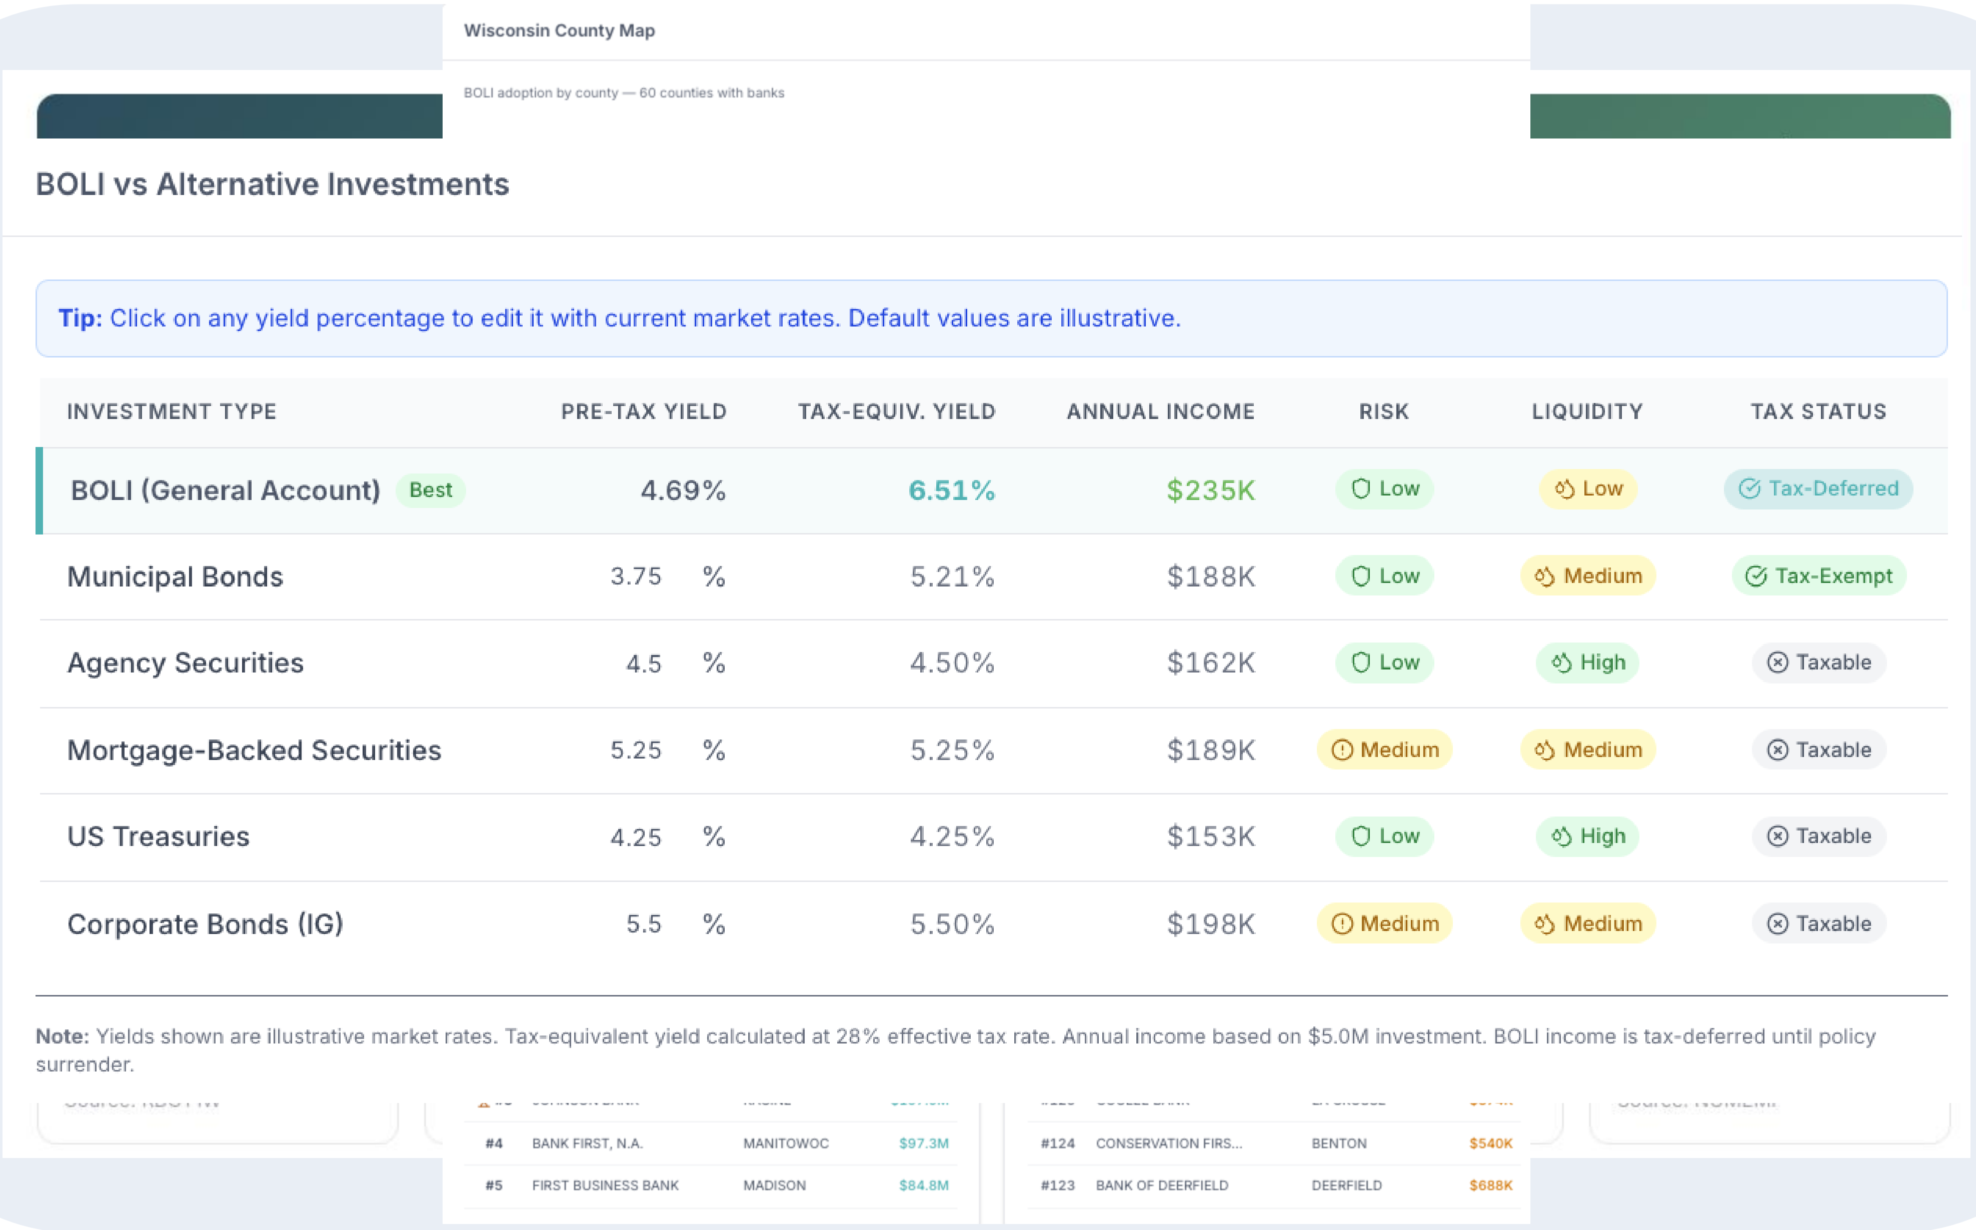
Wisconsin County Map

BOLI adoption by county — 60 counties with banks

BOLI vs Alternative Investments

Tip: Click on any yield percentage to edit it with current market rates. Default values are illustrative.

INVESTMENT TYPE	PRE-TAX YIELD	TAX-EQUIV. YIELD	ANNUAL INCOME	RISK	LIQUIDITY	TAX STATUS
BOLI (General Account) Best	4.69%	6.51%	$235K	Low	Low	Tax-Deferred
Municipal Bonds	3.75 %	5.21%	$188K	Low	Medium	Tax-Exempt
Agency Securities	4.5 %	4.50%	$162K	Low	High	Taxable
Mortgage-Backed Securities	5.25 %	5.25%	$189K	Medium	Medium	Taxable
US Treasuries	4.25 %	4.25%	$153K	Low	High	Taxable
Corporate Bonds (IG)	5.5 %	5.50%	$198K	Medium	Medium	Taxable

Note: Yields shown are illustrative market rates. Tax-equivalent yield calculated at 28% effective tax rate. Annual income based on $5.0M investment. BOLI income is tax-deferred until policy surrender.

#4	BANK FIRST, N.A.	MANITOWOC	$97.3M
#5	FIRST BUSINESS BANK	MADISON	$84.8M

#124	CONSERVATION FIRS...	BENTON	$540K
#123	BANK OF DEERFIELD	DEERFIELD	$688K

IN

FDIC Reports, Entire National Database

OUT

County Level Maps with Bank BOLI Sizes

ROC

Quick Comparison of BOLI v Other Assets

The bank board saw the analysis. This is the analysis becoming a tool.

Navigate With Confidence

Case 3 · The Rescue

The file: a financed policy, a 300 basis point rate shock, and a client ready to walk mid-crisis.

Shipped: a collateral stress test, the client letter, and the producer script.

Navigate With Confidence

Ninety seconds at my desk: the workout model.

DROP RECORDING HERE
replace this tag

INPUTS

- Annual premium and years financed
- Loan rate path: base and +300bps stress
- Crediting assumption
- CSV schedule, editable by year

OUTPUTS

- Loan balance and collateral due, by year
- Breakeven year
- One-line risk summary
- Every formula visible in a footnote panel

FOUR WORKOUT PATHS, EACH BACKTESTED AGAINST RATE AND CREDITING HISTORY

Partial paydown schedule

Refinance timing

Crediting reallocation

Collateral substitution

A rescue rarely works without new money. AI's job is to model every path against history, so the client chooses tradeoffs instead of guesses.

THE RECIPE State the role and domain rules. Specify inputs, outputs, edge cases. Demand visible formulas. Iterate out loud: the first version is a draft, not a verdict.

Navigate With Confidence

Case 4 · The Exit Opinion

The file: a policy exit with tax exposure, a trustee on the hook, and an authority that must actually exist.

Shipped: a written exit opinion with verified citations.

Navigate With Confidence

Same question, two machines.

DROP RECORDING HERE
replace this tag

CONSUMER CHATBOT

Ask for a private letter ruling supporting an aggressive split-dollar exit. Watch it produce one: a number, a holding, confident prose.

Might be accurate, hard to verify.

GROUNDED WORKFLOW

Same question, answered only from a supplied document set. Every claim carries a page and line. “No authority found” means exactly that.

Repeatable and scalable.

The rule that survives this talk: never cite what you have not opened.

Recorded at my desk, dated on screen. If the consumer model declines instead of fabricating, that is also the lesson: the vendors know.

The cost of skipping the check.

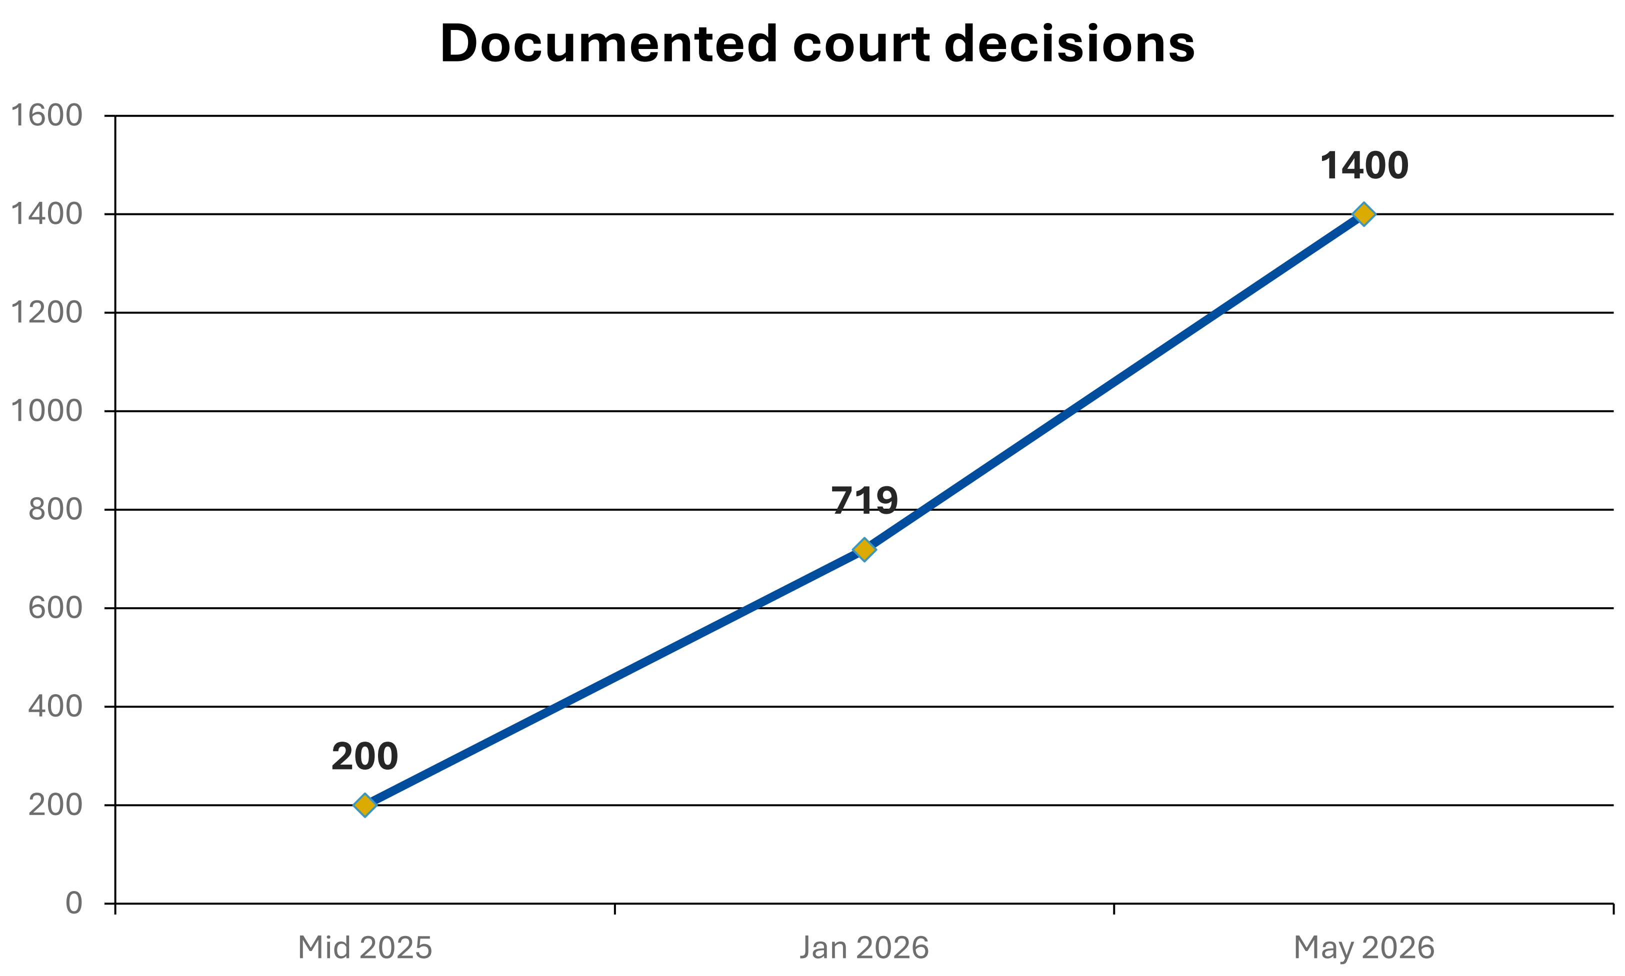

$5,000 Mata v. Avianca, 2023. The origin story.

$30,000 Sixth Circuit, March 2026. Case dismissed.

$110,000 D. Oregon, April 2026. Dismissed with prejudice.

Suspended Nebraska Supreme Court, April 2026. 57 of 63 citations defective.

1,400+ documented court decisions involving AI-fabricated material, 1,000+ in US courts. Even an AmLaw elite firm self-reported ~28 bad citations in April.

Verification is non-delegable. Firm size does not matter. The tool does not matter.

Source: Charlotin, AI Hallucination Cases Database, May 2026. Re-verify count before the session.

Navigate With Confidence

The validation protocol.

1. **Golden cases first** Run a file where you already know the answer, like an illustration you can verify line by line.
2. **Ground everything** High-stakes answers come only from supplied documents, a page citation per claim.
3. **Recompute independently** A second pass of AI-written code audits the first pass's math. You adjudicate.
4. **Cross-model check** Anything that ships to a client, a board, or a court gets a second human read.
5. **Never cite what you have not opened** The seven words that survive this talk.

Accurate citations matter for checking human work. Hold the machine to the same bar.

Ask it specifically to site page number and paragraph. In early parts of a conversation it's accurate, this is where context rot breaks down.

One repeatable loop. Four privacy tiers.

CONSUMER

Never for client work. Free tiers can train on inputs.

→

BUSINESS GRADE

Training off, retention set, audit logs, BAA. The default desk tier.

→

ZERO-RETENTION API

For built workflows and tools that touch client data at volume.

→

FULLY LOCAL

Privilege-bound work. My expert witness stack runs on one workstation, no cloud, ever.

The loop never changes: ground, generate, validate. The tier changes with the data. PII-safe AI exists at every budget.

Every case this morning ran the same loop at a different tier. That is the repeatable process you take home.

Navigate With Confidence

Before you try this at home.

NEVER PASTE

- Client PII and PHI
- Privileged communications
- Carrier-confidential product specs
- Anything under NDA

	Consumer free	Business grade
Trains on your inputs	Often yes	Contractually no
Retention you control	Limited	Configurable
Admin visibility	None	Full audit logs
BAA available	No	Available

The free account is not free. You are paying with the one asset this profession runs on: confidentiality.

THE RECIPE Business-grade account, training off, retention configured. Read the data-use page before the tool, not after.

LIMRA LOMA
Navigate With Confidence

The regulators already moved.

24+

states adopted the NAIC Model AI Bulletin; over half the country counting substantially similar guidance

AIS

a written AI Systems Program with senior, board-level accountability and vendor oversight

12

states piloting the AI Systems Evaluation Tool through September 2026; exam questions follow

Regulators look through vendors. Outsourcing the AI does not outsource the accountability.

Your carrier's AIS Program almost certainly governs the advanced sales desk. Ask compliance the one question on the Monday slide.

LIMRA LOMA
Navigate With Confidence

Six myths you stopped believing in the last hour.

"You need perfect data."	Case 1 ran on faxes, scans, and 300-page files.
"AI will replace your team."	Case 2 was one person doing an analyst pod's work, supervised.
"If it sounds right, it is right."	Case 4 opened with a ruling that never existed.
"It takes enterprise IT and big budgets."	Case 3's tool was built on stage for the cost of a prompt.
"A free account is fine."	The free tier is paid in confidentiality.
"We can wait."	Every case here was this quarter's desk work, with today's tools.

LIMRA LOMA
Navigate With Confidence

Your turn: the same five steps on one live file.

THIS WEEK

Pick one live file: an illustration comparison, a trust review, a plan summary.

Run FRAME through DELIVER exactly as you watched this morning.

Keep the artifact and your validation notes. That pair is your proof.

THEN MAKE IT A PROGRAM – INTERNALLY FIRST

Run the 30-day pilot from the quick start on that same workflow.

Hours saved and error rate become your budget slide.

The workbook carries all four case recipes and starter prompts.

One file. Five steps. That is the whole barrier to entry.

Navigate With Confidence

Be the internal startup.

1. PLAY

Run a hackathon with other departments. Build internally first. Have fun with it.

→

2. PEER REVIEW

You and your colleagues approve it before anyone outside the team sees it.

→

3. COMPLIANCE

InfoSec and compliance get their time, before anything touches a client.

→

4. LAUNCH

Deploy with IT. Serving models is genuinely technical, but it cannot take months.

If you know VBA and Excel, you are the build team. It is 1000% easier to teach advanced markets to use AI with intent than to teach an AI engineer this industry.

You already play in spreadsheets before you show anyone. AI is no different. Play, then peer review, then compliance sign-off, then a client.

Monday morning.

YOU · THIS WEEK

- Business-grade account. Training off. Retention set.
- Rebuild one case from today with your own file.
- Adopt the rule: never cite what you have not opened.

YOUR TEAM · 30 DAYS

- Golden-case test one workflow: known answer, graded.
- Pilot the reader on your most repeated comparison.
- Write the one-page policy and never-paste list.

YOUR ORG · THIS QUARTER

- Ask compliance: does our AIS Program cover this desk?
- Run the vendor checklist on anything touching client data.
- Name a champion, fund one pilot, report in 90 days.

The workbook ships by QR on the last slide: four recipes, four starter prompts, the protocol, the checklist.

Are we there yet?

No. And you do not need us to be.

Full automation is a mirage. The leverage is real, and it compounds for whoever starts first.

Thank You

LIMRA LOMA

Navigate With Confidence

Nick Bowman, CFP®, CEP®, MS · Co-Founder, Xcela AI

nb@xcela.ai