



A VERAPATH WHITE PAPER · OCTOBER 2026

The Transition from Efficiency to ROIC

How financial institutions turn AI productivity into economic return

PRESENTED BY

Joe McManh Co-Founder & CRO

Alec Crawford Founder & CEO

Frank Fitzgerald Founder & CTO

VERAPATH.COM



Table of Contents

Front Matter	—
Executive Premise	03
The Transition from Efficiency to ROIC	—
1. Efficiency Was the Right Place to Start	05
2. Capacity Is Where the Economics Begin	06
3. Redeployment Determines Whether Productivity Becomes Value	09
4. Operating Leverage Is the Bridge to ROIC	11
5. Agentic AI Changes the Unit of Automation	13
6. Capability and Control Have to Rise Together	15
7. Counterevidence Matters	17
8. A Better Way to Measure AI	19
9. What Changes for the C-Suite	26
10. The Transition	27
Back Matter	—
Evidence Notes and Sources	29
About the Author	31
About Verapath	32



Executive Premise

AI is making financial institutions more efficient. The unanswered question is whether that efficiency is making them more valuable.

An advisor who completes meeting preparation in minutes rather than an hour has recovered time, but the firm captures the economic benefit only when that time is put to work. More client relationships, deeper planning, faster onboarding and greater revenue capacity are possible outcomes, but none follows automatically from a faster meeting.

That is the management problem this paper addresses. A financial institution can deploy AI successfully and still fail to capture value. The missing decision is how the capacity created will change the work, the client experience and ultimately the economics of the business.

The opportunity to turn recovered time into economic value is growing as AI moves from individual assistance to connected execution. A meeting summary removes one task. An authorized process that prepares the meeting, gathers approved information, completes follow-up and escalates exceptions changes the amount of work an advisory firm can handle. The economic question is no longer confined to the employee's productivity. It becomes a question of how much more work the institution can absorb without adding cost.

THE ECONOMIC PROGRESSION

01

EFFICIENCY

02

CAPACITY

03

REDEPLOYMENT

04

OPERATING
LEVERAGE

05

ROIC

The path from efficiency to return runs through five distinct decisions. First, establish what work AI makes faster. Second, determine how much of that time becomes usable capacity. Third, assign that capacity to work with an identifiable purpose. Fourth, measure whether the resulting output grows faster than its cost. Finally, test the improvement against the capital invested. Each step has different evidence and a different owner.

For an RIA, the result might be more households served without compromising advice. For a community bank, it might be greater lending volume without a corresponding increase in processing expense. For an asset manager, it might be more mandates pursued and serviced without a proportional increase in client-service staff. For a credit union, it might be more member requests resolved while employees spend more time on complex cases. Each requires a different operating measure. None is proved by the number of AI licenses deployed.

The practical tool is a capacity ledger. It follows each deployment from the time released to the work assigned, the operating result achieved and the investment required. It forces management to name the benefit before claiming the return.

Efficiency Was the Right Place to Start

The first use cases made sense. Employees searched documents, prepared meetings, drafted communications and reviewed information faster. These were familiar tasks with visible results. They demonstrated what AI could do for an individual.

SURVEY · SELF-REPORTED

4+ hrs

SAVED WEEKLY BY MORE THAN
HALF OF AI-USING
RESPONDENTS

AssetMark advisor survey

AssetMark's September 2026 survey of 400 U.S. advisors, including 152 RIAs and 248 independent broker-dealer-affiliated advisors, reported that more than half of AI-using respondents saved at least four hours weekly. This is self-reported task efficiency across a mixed advisor population, not measured RIA profitability or a representative estimate for every firm. [1]

COMPANY-REPORTED

18,000+

CUSTOMER-SERVICE EMPLOYEES
SUPPORTED

Bank of America · EricaAssist

Bank of America's EricaAssist now supports more than 18,000 customer-service employees. The bank says the system surfaces contextual guidance in under three seconds and reduces average call time by nearly one minute, while keeping the employee at the center of the interaction. [2]

These examples establish the value of individual productivity. They leave open the question that determines institutional return: where did the recovered capacity go?

A shorter service call creates room for another call. Faster preparation creates room for another client conversation. The operating result depends on whether demand exists, whether the surrounding process can absorb the additional volume and whether management assigns the capacity to work that matters.



EFFICIENCY TO ROIC — SECTION TWO

02

Capacity Is Where the Economics Begin

*Capacity is the additional work an institution can perform
with the resources it already has.*

Two community banks show the first two steps: measuring the capacity released and deciding where it should go.

CAPACITY AND REDEPLOYMENT · REPORTED OPERATING EVIDENCE

COMPANY / PARTNER-
REPORTED ESTIMATE

276

ESTIMATED STAFF HOURS SAVED

Forward Bank

Forward Bank provides a community-bank example of the first transition in the framework: from efficiency to measurable capacity. In one reporting period, the bank recorded 4,677 Microsoft Copilot actions across the organization, representing an estimated 276 hours of staff time saved. The activity count and the hours estimate are company/partner-reported, not independently verified financial returns. Their significance is that the bank is putting an operating measure against the capacity being released. [3]

That is where the economic analysis begins, not where it ends. The next questions are where those hours are redeployed, what additional output or avoided cost results, and whether that operating change ultimately justifies the investment.

EXECUTIVE-REPORTED · DEPLOYED AI / REDEPLOYMENT OBJECTIVE

VeraBank · AI deployment and intended redeployment

VeraBank illustrates the next management decision. The community bank has deployed LLMs and agentic technology, including an internal LLM knowledge base in production for employees. Michael Purifoy, VeraBank's chief treasury and digital banking officer, has described the objective as creating capacity within the workforce so employees can spend less time on repetitive work and more time deepening customer relationships and supporting growth. [4]

VeraBank has therefore identified both the capacity it intends to create and where management intends to redeploy it. The public evidence does not establish that this redeployment has already produced incremental financial return.

ILLUSTRATIVE

Consider an illustrative ten-person community-bank operations team that could absorb the workload previously requiring twelve. Its expense would not automatically decline. It could instead handle more volume, reduce its backlog, improve service or shift attention to exceptions and customer relationships. Management must decide which outcome it wants, then measure it.

These are different operating choices. Each requires its own baseline and outcome measure.

ABA Banking Journal recently framed this as a “capacity dividend.” Using 50 working weeks for illustration, one hour a week returned to 100 employees equals 5,000 hours annually. For 300 employees, it equals 15,000 hours. These are hypothetical calculations, not measured bank results. At that scale, recovered time stops being a productivity trick. It becomes an organizational asset. [5]

HYPOTHETICAL · 1 HR/WEEK · 50 WEEKS

100 employees 5,000 hrs

300 employees 15,000 hrs

Hours are an inventory of opportunity. Without a changed workflow, assigned demand and accountable leadership, they are not an operating result.

Redeployment Determines Whether Productivity Becomes Value

The institution needs a redeployment decision before it needs another adoption metric. A firm can report widespread AI use while its advisors continue to carry the same client load, its onboarding queue remains unchanged and its service team has no additional capacity. The missing link is operational: someone has to decide which work changes.

Licenses, prompts and hours saved are useful implementation measures. They cannot show whether an advisory firm gained capacity for more households, a bank processed more loans, an asset manager pursued more mandates or a credit union reduced its service backlog.

Redeployment must therefore be part of the original business case. The institution needs a named owner, available demand and a measurable result before it commits to scaling.

SURVEY · 68 WEALTH MANAGEMENT FIRMS

Cerulli and Vista · Wealth management hiring intentions

73%

EXPECTED TO ADD JUNIOR ADVISORS

67%

CLIENT-SERVICE ASSOCIATES

56%

SENIOR ADVISORS

A September 2026 Cerulli and Vista study surveyed 68 wealth management firms, with approximately \$1.2 trillion in aggregate assets under management and a concentration in larger advisory firms. Among respondents, 73% expected to add junior advisors, 67% client-service associates and 56% senior advisors over the next two years. The sample is concentrated in larger firms and reports hiring intentions, not verified future employment. It nevertheless illustrates how firms can view AI-enabled capacity as a growth opportunity rather than assume that every recovered hour eliminates labor. [6]

The same research also shows a difference in economic measurement. Among firms classified as Leading, 50% measure cost savings from AI and 38% track AI-related revenue growth, compared with 9% and 6%, respectively, among Exploring firms. AI-specific spending was also projected to rise from an average 8% to 15% of technology budgets. [6]

Firms classified as Leading are therefore more likely to put economic measures around their AI deployments. The study does not establish that those measurement practices caused superior financial returns.

AI can remove work people do not need to be doing, allowing advisors, bankers, underwriters, service employees and operators to spend more time on judgment, relationships, exceptions and higher-value work. At the same time, the institution can increase output without requiring employment expense to rise at the same rate.

The decision is specific to the institution. An advisory firm with a full client pipeline can assign the time to new relationships. A bank with a service backlog can clear it before hiring another team. An underwriting group can move routine processing to controlled automation while its experienced employees handle exceptions. An asset manager can direct the time to more mandates and deeper consultant coverage. The same technology can produce different economic outcomes because management makes different choices.

That does not mean every role remains unchanged. It means the economic case cannot assume that every hour recovered is an hour of payroll removed. In many institutions, the immediate opportunity is to serve more clients, move faster and expand without adding people at the same rate as volume.

EXECUTIVE-REPORTED ·
BUILD-VS-BUY / OPERATING
RESULT

50%

LESS TIME PER WEALTH
ACCOUNT REVIEW

THIRD-PARTY ALTERNATIVES
PRICED ABOVE \$375K ANNUALLY

Peapack Private Bank & Trust

Peapack Private Bank & Trust provides a different economic lens. The bank's wealth team had been managing account reviews in spreadsheets and considered third-party software costing more than \$375,000 annually. Instead, Peapack's engineers used AI to build an internal application. CTO John Kowal reports that each account review now takes 50% less time. [7]

The example combines an operating result with a real build-versus-buy decision, but neither figure is ROIC. The price of the external alternative is not the same as net savings because Peapack incurred its own costs to build and operate the internal application. And the capacity released by faster reviews creates economic value only if it is productively redeployed. The institution still has to measure the downstream operating result and the fully loaded investment.

Operating Leverage Is the Bridge to ROIC

Return on invested capital is net operating profit after tax (NOPAT) divided by invested capital. It is not another name for productivity.

Consider three separate measures. Minutes saved establish task efficiency. More files completed by the same team establish additional throughput. Higher incremental operating profit relative to the capital committed can support an ROIC calculation. The first two are operational evidence, not substitutes for the third.

The operating result must then be reconciled to implementation costs, ongoing expense and the capital required to produce it.

When output increases faster than the cost of producing it, the institution can create operating leverage. If that improvement increases operating profit relative to the capital committed, ROIC can improve. The distinction matters: throughput is operating evidence; ROIC is a financial outcome.

The pathway can come from several directions: more client capacity per advisor; more loan volume per operations employee; faster onboarding without proportional support expense; lower cost per transaction; greater revenue per employee; fewer manual exceptions; faster product or code delivery; or the ability to absorb growth without adding the next layer of fixed cost.

COMPANY-REPORTED

~390

CLOSED DEALS FROM NEARLY
2,000 SALES OPPORTUNITIES

Manulife John Hancock Retirement

Manulife John Hancock Retirement provides one of the stronger examples because the evidence moves beyond time saved. An AI tool launched in March 2025 generated nearly 2,000 sales opportunities and approximately 390 closed deals that contributed additional assets under management. At the same time, teams reported an 80% reduction in information-search time and a fivefold increase in lead recommendations. [8]

Manulife's reported experience connects productivity to commercial activity and additional assets. It does not disclose the incremental profit or capital required to calculate the return on this particular AI investment.

This is farther along the progression than a time-saved metric: productivity is connected to commercial activity. It still stops short of ROIC.

Measure what changed in the business. For an RIA, that may be households served per advisor without compromising advice or service. For a bank, it may be lending throughput, revenue, expense and credit quality. For an asset manager, it may be mandates won and assets retained per client-service professional. For an operations team, it may be transactions per employee and cost per transaction.

COMPANY-REPORTED

66%

AVERAGE REDUCTION IN
INFORMATION-SEARCH TIME

Lloyds Banking Group

Lloyds Banking Group provides evidence that moves beyond task efficiency. An internal AI knowledge tool used by 20,000 employees reduced information-search time by an average of 66%. The group also reported that roughly 50 live generative-AI use cases delivered approximately £50 million of incremental value in 2025 and is targeting more than £100 million of incremental P&L benefit from generative AI in 2026. [9]

The disclosure connects scaled AI deployment to an economic result. But company-reported value and incremental P&L benefit are not themselves ROIC. The public disclosure does not provide the invested-capital detail necessary to calculate an independent ROIC for the AI program.

THE PUBLIC-EVIDENCE GAP

The pattern across the public evidence is revealing. Financial institutions increasingly report time saved, capacity released, commercial activity, cost measures and even company-defined AI “value.” Far fewer disclose enough information to trace those gains through incremental operating profit, taxes and the capital committed to produce them. That is why a productivity metric, a revenue metric and even a reported ROI figure cannot automatically be treated as ROIC.

The absence of a complete public ROIC calculation is not a reason to manufacture one. It is the measurement gap this framework is designed to close.



EFFICIENCY TO ROIC — SECTION FIVE

05

Agentic AI Changes the Unit of Automation

*An assistant makes an employee faster. An agent can change
how the work gets done.*

Consider the difference between drafting an advisor's follow-up note and executing an authorized follow-up process. The latter requires access to approved information, defined actions, checks on the result and a clear point at which judgment returns to the advisor. Its value is not merely the speed of one step. It is the human effort no longer required at every step.

TASK VERSUS PROCESS

TASK · ASSISTANT

Drafting an advisor's follow-up note

PROCESS · AGENT

→ Executing an authorized follow-up process

- Access to approved information
- Defined actions
- Checks on the result
- A clear point at which judgment returns to the advisor

The distinction matters because a process has handoffs, permissions, exceptions and accountability that a single task may not.

That change can expand throughput without adding human effort at each handoff. It also forces management to decide which steps can be automated, which require approval and who owns an exception.

Agentic AI therefore connects the economic case to the control architecture. Neither can be designed independently.

Capability and Control Have to Rise Together

A system that drafts a response creates one kind of exposure. A system that can retrieve client information, enter applications and execute authorized actions creates another.

A policy governing employee use is necessary but insufficient once a system can execute. Control must extend to the data it retrieves, the action it is permitted to take and the result it returns.

Grant Thornton · 2026 banking AI survey

SURVEY · SELF-REPORTED

62%

REPORTED GREATER EFFICIENCIES

32%

REPORTED REVENUE GROWTH

18%

FULLY CONFIDENT IN AI GOVERNANCE

42%

CITED INSUFFICIENT DATA READINESS

The evidence increasingly supports extending control from use to execution. In Grant Thornton's 2026 banking AI survey, 62% of banks reported greater efficiencies from AI, but only 32% reported revenue growth. Half said governance and compliance barriers limit AI performance, only 18% were fully confident in their AI governance and controls, and 42% cited insufficient data readiness as a cause of AI project failure. [10]

Deloitte · AI governance research

SURVEY · OBSERVED ASSOCIATION

Deloitte's 2026 research, drawing on 135 respondents across primarily large global banks, reports an association between AI governance maturity and broader deployment and revenue outcomes. The study does not establish that governance caused revenue growth, nor is its sample directly representative of community banks. Its relevance here is narrower: institutions pursuing consequential AI deployment must treat controls and operating design together. [11]

Human review can contain a pilot. It cannot be the entire control model when an agent executes a process repeatedly. Permissions, approved data, validation, escalation and a record of each consequential action have to be built into execution.

Governance does not create ROIC on its own. It determines whether the institution can move consequential work beyond isolated pilots without accepting uncontrolled risk as the price of scale.

A faster workflow that cannot be supervised, audited or trusted is not ready to scale.

Counterevidence Matters

Evidence of efficiency is growing faster than evidence of institution-level financial return. The difference warrants attention.

Gartner · September 2026

SURVEY

22%

HAD SUCCESSFULLY SCALED AI ACROSS MULTIPLE BUSINESS UNITS OR ADOPTED AN AI-FIRST APPROACH

11%

DID NOT KNOW WHAT THEIR FUNCTION SPENT ON AI IN 2025

Gartner reported in September 2026 that only 22% of surveyed organizations had successfully scaled AI across multiple business units or adopted an AI-first approach. Eleven percent did not know what their function spent on AI in 2025. Productivity was the most common target, but Gartner found that popular use cases were not necessarily the ones producing the highest reported returns. [12]

Grant Thornton's banking survey reported efficiency gains more often than revenue growth. The findings are self-reported and do not establish causation. They do expose the gap between a productivity claim and a commercial outcome. [10]

McLemore and Mihov · operational losses

PEER-REVIEWED · OBSERVATIONAL · LARGE BHCS

The downside is not limited to disappointing financial returns. In a February 2026 peer-reviewed study using supervisory operational-loss data from large U.S. bank holding companies, Ping McLemore and Atanas Mihov found that higher AI investment was associated with greater operational losses per dollar of assets. The relationship was concentrated in external fraud, client-related problems and system failures, and was more pronounced where risk management was weaker. The study concerns large bank holding companies, not community banks, and its results do not establish that every AI deployment raises losses. It does, however, challenge the assumption that greater AI capability automatically produces better institutional economics. [13]

The American Bankers Association's 2026 survey of banks found an industry actively engaging with AI but doing so cautiously, with data readiness, governance, regulatory uncertainty and information security among the most common constraints. Banks were prioritizing lower-risk applications that support human judgment rather than rushing toward autonomous customer-facing decisions. [14]

These findings argue for discipline, not retreat. Select work that matters, identify the economic outcome, assign ownership and establish the controls before scaling.

There are several ways to lose the benefit. Employees may use recovered time for more low-value activity. A process may run faster without removing a handoff. Integration and supervision may cost more than expected. Weak data may prevent a pilot from becoming dependable production. A favorable demonstration does not resolve any of these questions.

A deployment plan that stops at implementation leaves its most important management decision unresolved.




EFFICIENCY TO ROIC — SECTION EIGHT

08

A Better Way to Measure AI

A capacity ledger makes these decisions visible from the start.

A capacity ledger makes these decisions visible from the start. For each material workflow, it records the operating baseline, the capacity actually released, the work management assigns to that capacity, the change in throughput or unit economics and the full cost of implementation and operation. The ledger is not a substitute for financial statements. It is the missing operating record between a productivity claim and a financial return.



ILLUSTRATIVE RIA · 20 ADVISORS

ALL FIGURES HYPOTHETICAL

Consider an illustrative RIA with 20 advisors. After a controlled pilot, it measures 100 usable advisor hours recovered each month, net of review and rework. Management assigns 70 hours to additional client meetings, 20 to complex planning and 10 to quality review. This is a capacity allocation, not a result. The firm has to establish that its pipeline can fill the meeting slots, that onboarding and service teams can support additional households and that advice quality remains intact.

Assume the 70 hours support 35 additional two-hour blocks of meetings and follow-up each month. If 80% are actually used, the operating result is 28 additional completed blocks (56 hours), not 70 hours of hypothetical productivity. The firm then tracks how many lead to new relationships, retained households or additional planning work, together with the service burden they create. It must not treat every additional meeting as incremental revenue.

ILLUSTRATIVE RIA · THE FINANCIAL TEST

ALL FIGURES HYPOTHETICAL

The financial test requires a counterfactual. Assume, entirely for illustration, that management estimates \$180,000 in annual revenue associated with the additional activity. It must first determine how much would have occurred without AI, using a comparable prior period or an appropriate control group where feasible, while allowing for market conditions, referrals and other changes. Only genuinely incremental revenue belongs in the calculation.

If, after that test, the validated incremental amount remains \$180,000, and additional annual service and delivery expense is \$55,000 while recurring AI, integration and control costs are \$35,000, the incremental operating contribution before tax would be \$90,000. If the initial deployment required \$120,000 of capitalized investment, management would still need the applicable tax rate, accounting treatment of implementation costs, depreciation and the full invested-capital definition before calculating any change in ROIC. All figures are hypothetical. Neither the \$90,000 contribution nor the ratio of contribution to initial investment is ROIC.

HYPOTHETICAL LEDGER ENTRIES

01 **Baseline**

RIA with 20 advisors

02 **Released capacity**

100 usable advisor hours each month, net of review and rework

03 **Assigned work**

70 hours client meetings · 20 complex planning · 10 quality review

04 **Modeled operating result**

28 additional completed blocks (35 blocks, 80% used)

05 **Fully loaded investment**

\$120,000 capitalized initial deployment · \$35,000 recurring AI, integration and control costs

Validated incremental revenue \$180,000

Less additional annual service and delivery expense (\$55,000)

Less recurring AI, integration and control costs (\$35,000)

Incremental operating contribution before tax \$90,000

ALL FIGURES ARE HYPOTHETICAL. NEITHER THE \$90,000 CONTRIBUTION NOR THE RATIO OF CONTRIBUTION TO INITIAL INVESTMENT IS ROIC.

The ledger exposes the decisions that matter. A full meeting calendar without new relationships may justify a service or retention case, not a growth claim. New revenue that requires proportionate support hiring may improve growth without producing the expected operating leverage. An agent that saves time but fails review standards cannot be scaled merely to improve the arithmetic.

The same ledger, three other institutions

The ledger works the same way in any institution. Only the workflow, the demand and the operating measure change.

ILLUSTRATIVE COMMUNITY BANK · LENDING OPERATIONS TEAM OF 12

ALL FIGURES HYPOTHETICAL

Consider an illustrative community bank whose lending operations team of 12 processes 240 consumer and small-business loan applications a month. After a controlled pilot, an agent completes document checks and tracks outstanding conditions, releasing 160 usable staff hours a month, net of review and rework. Management assigns 120 hours to additional applications from the existing pipeline and 40 to exception review. This is a capacity allocation, not a result. The bank has to establish that qualified demand exists, that underwriting standards hold and that closing and servicing can absorb the volume.

At two hours per application, 120 hours support 60 additional applications a month. If the pipeline fills 75%, the operating result is 45 additional decisions, not 160 hours of hypothetical productivity. The bank then tracks how many are funded, how many would have arrived anyway and how the new loans perform.

A bank's financial test takes longer. Loan income arrives over the life of the loan, and credit losses may appear later. Assume, entirely for illustration, \$6 million of new loans in a year, of which management attributes \$3 million to the added capacity rather than to loans it would have made anyway. At an illustrative 3% net spread, that is \$90,000 of annual net interest income, plus \$20,000 of fees. After \$25,000 of servicing expense, \$30,000 of recurring AI, integration and control costs and a \$15,000 credit-loss provision, the incremental operating contribution before tax would be \$40,000. The bank would still need the capital held against the new loans, its tax rate and the treatment of a \$100,000 initial deployment before calculating any change in return.

HYPOTHETICAL LEDGER ENTRIES

01 Baseline Lending operations team of 12; 240 applications a month	Attributed net interest income	\$90,000
	Fees	\$20,000
02 Released capacity 160 staff hours a month, net of review and rework	Less additional servicing expense	(\$25,000)
	Less recurring AI, integration and control costs	(\$30,000)
03 Assigned work 120 hours to additional applications · 40 to exception review	Less provision for expected credit losses	(\$15,000)
	Incremental operating contribution before tax	\$40,000
04 Modeled operating result 45 additional decisions a month (60 supported, 75% used)	ALL FIGURES ARE HYPOTHETICAL. NEITHER THE \$40,000 CONTRIBUTION NOR ITS RATIO TO THE INITIAL DEPLOYMENT IS ROIC.	
05 Fully loaded investment \$100,000 initial deployment · \$30,000 recurring AI, integration and control costs		

The ledger exposes the decisions that matter. Decisions that are not funded change the pipeline report, not the economics. Volume that loosens credit standards weakens the result being measured, and the cost may surface quarters later. Capacity that closing or servicing cannot absorb moves the bottleneck rather than removing it.

ILLUSTRATIVE ASSET MANAGER · INSTITUTIONAL RFP AND CLIENT-SERVICE TEAM OF 8

ALL FIGURES HYPOTHETICAL

Consider an illustrative asset manager whose institutional team of 8 completes 30 RFPs and due-diligence questionnaires a quarter, at about 25 hours each. After a controlled pilot, an agent drafts responses from approved content, releasing 300 usable hours a quarter, net of compliance review. Management assigns 200 hours to additional responses and 100 to consultant and client follow-up. This is a capacity allocation, not a result. The firm has to establish that suitable opportunities exist, that compliance review can absorb the volume and that response quality holds.

At 25 hours each, 200 hours support 8 additional responses a quarter. If 6 are submitted, the operating result is 6 additional responses, not 300 hours of hypothetical productivity. Mandate decisions take months, so the firm tracks finalist appearances and wins over time, along with assets retained through the added follow-up.

Assume, entirely for illustration, that within a year the additional responses contribute to two mandates totaling \$150 million, and management attributes half, \$75 million, to the added capacity rather than to mandates it would have won anyway. At an illustrative fee of 0.30%, that is \$225,000 of annual revenue. After \$70,000 of servicing and reporting expense and \$45,000 of recurring AI, integration and control costs, the incremental operating contribution before tax would be \$110,000. Fees move with markets and mandates can be withdrawn, so the firm would need a multi-year view, its tax rate and the treatment of a \$90,000 initial deployment before calculating any change in return.

HYPOTHETICAL LEDGER ENTRIES

01 **Baseline**

Institutional RFP and client-service team of 8; 30 responses a quarter

Attributed annual fee revenue \$225,000

Less additional servicing and reporting expense (\$70,000)

02 **Released capacity**

300 hours a quarter, net of compliance review

Less recurring AI, integration and control costs (\$45,000)

03 **Assigned work**

200 hours to additional responses · 100 to consultant and client follow-up

Incremental operating contribution before tax \$110,000

ALL FIGURES ARE HYPOTHETICAL. NEITHER THE \$110,000 CONTRIBUTION NOR ITS RATIO TO THE INITIAL DEPLOYMENT IS ROIC.

04 **Modeled operating result**

6 additional responses a quarter (8 supported, 6 submitted)

05 **Fully loaded investment**

\$90,000 initial deployment · \$45,000 recurring AI, integration and control costs

More responses without a stable win rate add cost, not revenue. A mandate won at a lower fee may not cover its servicing. Follow-up that keeps an existing client can matter as much as a new win, but it belongs in the ledger as retention, not growth.

ILLUSTRATIVE CREDIT UNION · MEMBER CONTACT CENTER OF 25

ALL FIGURES HYPOTHETICAL

Consider an illustrative credit union whose member contact center of 25 handles 18,000 requests a month. After a controlled pilot, an agent resolves routine balance, card and transfer requests within approved permissions, releasing 400 usable staff hours a month, net of review. Management assigns 250 hours to complex cases and member financial conversations and 150 to clearing the callback backlog. This is a capacity allocation, not a result. The credit union has to establish that members want the conversations, that staff are prepared to hold them and that service holds for members who still call.

At 30 minutes each, 250 hours support 500 additional conversations a month. If 60% occur, the operating result is 300 conversations, not 400 hours of hypothetical productivity. The credit union then tracks how many lead to refinanced loans, new accounts or retained members.

A credit union measures return differently. It has no outside shareholders. Its earnings build the net worth that protects members and fund better rates, lower fees and improved services. The test is whether the change strengthens earnings and member value relative to the resources committed. Assume, entirely for illustration, \$2 million of additional loans attributed to the added capacity. At an illustrative 3% net spread, that is \$60,000 a year, and clearing the backlog avoids \$25,000 of overtime. After \$10,000 of servicing expense, \$35,000 of recurring AI, integration and control costs and a \$10,000 credit-loss provision, the incremental contribution would be \$30,000. The board would weigh that against an \$80,000 initial deployment, the effect on net worth and the value returned to members, not a shareholder ROIC calculation.

HYPOTHETICAL LEDGER ENTRIES

01 Baseline Member contact center of 25; 18,000 requests a month	Attributed net interest income	\$60,000
	Avoided overtime	\$25,000
02 Released capacity 400 staff hours a month, net of review	Less additional servicing expense	(\$10,000)
	Less recurring AI, integration and control costs	(\$35,000)
03 Assigned work 250 hours to complex cases and member conversations · 150 to the callback backlog	Less provision for expected credit losses	(\$10,000)
	Incremental contribution	\$30,000
04 Modeled operating result 300 additional member conversations a month (500 supported, 60% occur)	ALL FIGURES ARE HYPOTHETICAL. FOR A COOPERATIVE, THE TEST IS EARNINGS AND MEMBER VALUE RELATIVE TO THE RESOURCES COMMITTED, NOT SHAREHOLDER ROIC.	
05 Fully loaded investment \$80,000 initial deployment · \$35,000 recurring AI, integration and control costs		

Conversations that do not change a member's financial position are service, not growth, and belong in the ledger as service. Routing routine requests to an agent saves time only if members who still call get through faster. A gain that needs new staff to sustain moves the cost rather than removing it.

THE FIVE-STAGE MEASUREMENT FRAMEWORK

LEVEL	QUESTION	EXAMPLES
1. Efficiency	Did the task become faster or cheaper?	Minutes saved, search time, cycle time, cost per task
2. Capacity	How much productive capacity came back?	Hours recovered, files handled, clients supported, exceptions reduced
3. Redeployment	Where did the capacity go?	Client time, growth activity, judgment, backlog reduction, higher-value work
4. Operating leverage	Did output grow faster than resources?	Revenue/employee, clients/advisor, loans/FTE, volume/headcount, margin
5. ROIC	Did institution-level economics improve?	NOPAT, invested capital, sustainable return on incremental investment

The capacity ledger has an accountable owner. For every material deployment, the business-line leader establishes the demand and intended redeployment; the COO confirms that the surrounding process can absorb it; the risk and technology functions define execution limits; and the CFO reconciles the resulting operating change with implementation cost, recurring expense and invested capital. No department can close that loop alone.

Efficiency establishes that work became faster. Capacity establishes how much additional work the institution can perform. Redeployment identifies where that capacity went. Operating leverage tests whether output grew faster than the resources required. ROIC asks whether the improvement ultimately changed returns relative to invested capital.

Management does not need a manufactured ROIC calculation for every workflow. It needs a credible connection between the workflow, the capacity it releases and a financial or operating result the institution can observe.

For material investments, the CFO must also assess whether the expected incremental return, its timing and its risks justify the capital committed relative to alternative uses of that capital.

Before scaling a material AI deployment, management needs answers to six questions:

1.

What work is becoming more efficient?
2.

How much usable capacity does that create?
3.

Where will management deliberately redeploy that capacity?
4.

Which institutional metric should change if the redeployment works?
5.

What controls are required to scale the use case safely and repeatedly?
6.

After implementation, supervision, recurring expense and capital requirements, does the incremental operating result justify the investment?

MANAGEMENT THRESHOLD

If management cannot answer these questions, the use case may be useful. Its economic case remains incomplete.

What Changes for the C-Suite

CEO	The CEO determines where the institution intends to create capacity and what it will do with it. That is a business decision, not a software decision.
CFO	The CFO follows the economics: implementation and operating costs, changes in unit cost, revenue productivity, margin and the capital committed. Time saved is an input to that analysis.
COO	The COO redesigns the process. A faster individual task means little if the same manual handoffs, approvals and queues remain untouched.
Risk, compliance, cybersecurity & technology	Risk, compliance, cybersecurity and technology leaders establish the conditions for execution: authorized data, permitted actions, validation, human escalation and a reconstructable record.
Business-line leaders	Business-line leaders put capacity to work. If an advisory team gains room for more households, who owns the growth plan? If a lending team can process more applications, where is the demand and what happens to service and credit quality?

WHERE TO BEGIN

Start with one material AI deployment already underway. Establish its operating baseline, measure the usable capacity it creates and assign that capacity to work with identifiable demand. Give the business-line leader ownership of the operating result, the COO responsibility for the surrounding process, risk and technology responsibility for execution controls, and the CFO responsibility for validating the economics. Agree on the measures before expanding the deployment.



EFFICIENCY TO ROIC — SECTION TEN

10

The Transition

The next competitive difference in financial services may not be access to AI.

The next competitive difference in financial services may not be access to AI. Models and tools are becoming broadly available. Institutions will differ in whether they can change their processes, absorb more demand and preserve the controls that make greater scale sustainable.

That difference is measurable. An advisory firm that can serve more households without diluting the client experience has changed its capacity. A bank that can process more credit applications without proportionate expense has changed its operating model. An asset manager that can pursue more mandates with the same team has changed its reach. A credit union that can resolve more member needs while giving employees more time for complex conversations has changed its service capacity. None has demonstrated superior returns until management reconciles the resulting profit or institutional benefit with the capital and ongoing costs required.

This is also why workforce planning belongs in the economic discussion. Removing routine work gives experienced people more room for judgment, relationships and exceptions. If demand grows, the institution may need more advisors, bankers, risk specialists or service staff, even as productivity rises. The objective is not to assume a smaller workforce. It is to build a business that can do more valuable work with each increment of resources.

THE EXECUTIVE OBLIGATION

The executive obligation is to make that progression explicit. Identify the capacity. Assign it to work with demand behind it. Remove the bottlenecks that would prevent the gain from reaching the client or customer. Protect the process as it scales. Then measure the financial outcome without confusing an operating indicator for a return calculation.

Don't stop at time saved. Follow the capacity.

Evidence Notes and Sources

1. AssetMark, “More Than Half of Advisors Using AI Save 4+ Hours a Week, AssetMark Research Finds,” September 15, 2026. Survey: 400 U.S. advisors, including 152 RIAs; self-reported.
<https://www.assetmark.com/resources/blog/press-release/advisor-insights-report-ai/>
2. Bank of America, “Bank of America Enhances EricaAssist with Generative AI to Help Employees Resolve Client Needs Faster,” July 21, 2026.
<https://newsroom.bankofamerica.com/content/newsroom/press-releases/2026/07/bank-of-america-enhances-ericaassist-with-generative-ai-to-help.html>
3. WSI, “WSI Client Success Stories: 3 Businesses, 3 Growth Challenges, and the Strategies Behind Their Results,” August 31, 2026; and WSI Digital, “WSI Digital and Forward Bank Recognized with Global People’s Choice Award for AI Partnership,” EIN Presswire, September 28, 2026. Company/partner-reported figures for one reporting period; estimated hours, not independently verified.
<https://www.wsiworld.com/blog/wsi-client-success-stories-3-businesses-3-growth-challenges-and-the-strategies-behind-their-results>
https://www.einnews.com/pr_news/945822817/wsi-digital-and-forward-bank-recognized-with-global-people-s-choice-award-for-ai-partnership
4. Q2, “What AI Looks Like at a Community Bank,” VeraBank executive interview and transcript, June 1, 2026.
<https://www.q2.com/podcasts/what-ai-looks-like-at-a-community-bank>
5. ABA Banking Journal, “How will banks reinvest the time AI saves?” September 15, 2026.
<https://bankingjournal.aba.com/2026/09/how-will-banks-reinvest-the-time-ai-saves/>
6. Cerulli Associates and Vista Equity Partners, “Advisor Headcount Set to Grow as AI Expands Capacity,” September 9, 2026. Survey: 68 wealth management firms (RIAs, independent broker-dealers and multi-family offices), May–July 2026, approximately \$1.2T aggregate AUM. Measurement-practice and AI-budget figures: Vista Equity Partners, “The State of Wealth Management AI Adoption,” September 8, 2026.
<https://www.vistaequitypartners.com/insights/the-state-of-wealth-management-ai-adoption/>
<https://www.cerulli.com/press-releases/advisor-headcount-set-to-grow-as-ai-expands-capacity>
7. SmarterX, “The AI Show Episode 238: How a 700-Person Bank Is Using AI to Build Apps, Agents, and Digital Employees,” interview with Peapack Private Bank & Trust CTO John Kowal, September 10, 2026. Executive-reported; published transcript.
<https://podcast.smarterx.ai/shownotes/238>
8. Manulife, “Manulife Wealth & Asset Management Continues to Build Its AI Leadership and Impact with AI-Powered Tools Enabling Better Advisor and Customer Experiences,” May 2026.
<https://www.manulife.com/ca/en/about-us/news/manulife-wealth-and-asset-management-continues-to-build-its-ai-leadership>
9. Lloyds Banking Group, Annual Report and Accounts 2025 (Gen AI use cases, 2025 incremental value and 2026 incremental P&L target); and “Group expects over £100m in value from next-gen AI in 2026,” January 29, 2026 (knowledge tool). Company-reported.
<https://www.lloydsbankinggroup.com/investors/annual-report.html>
<https://www.lloydsbankinggroup.com/media/press-releases/2026/lloyds-banking-group/ai-driven-benefits-2026.html>
10. Grant Thornton, “Banking insights: 2026 AI Impact Survey,” 2026. Survey: 950 senior business leaders, February 23–March 18, 2026, including 50 banking respondents; self-reported.
<https://www.grantthornton.com/insights/survey-reports/banking/2026/banking-insights-2026-ai-impact-survey>

11. Deloitte, "Banking on trust: AI governance for growth, resilience and scale," 2026. Surveys 135 leaders and technology/data respondents at primarily large global banks; observed association, not causal proof.
<https://www.deloitte.com/ap/en/perspectives/banking-on-trust.html>
12. Gartner, "Only 22% of Organizations Have Successfully Scaled AI Across Multiple Business Units," September 1, 2026.
<https://www.gartner.com/en/newsroom/press-releases/gartner-survey-finds-only-22-percent-of-organizations-have-successfully-scaled-ai-across-multiple-business-units>
13. Ping McLemore and Atanas Mihov, "AI and Operational Losses: Evidence from U.S. Bank Holding Companies," The Review of Corporate Finance Studies, February 17, 2026, doi:10.1093/rcfs/cfag003. Observational large-bank evidence.
<https://doi.org/10.1093/rcfs/cfag003>
14. American Bankers Association / ABA Banking Journal, "Banks view doing nothing with AI as greatest risk," March 30, 2026.
<https://bankingjournal.aba.com/2026/03/aba-survey-banks-view-doing-nothing-with-ai-as-greatest-risk/>

SUPPLEMENTARY READING

Supplementary reading only (not evidence for a manuscript claim): ABA Banking Journal, "Banks' wealth units pursue AI carefully," September 14, 2026.

<https://bankingjournal.aba.com/2026/09/banks-wealth-units-pursue-ai-carefully/>

Bain & Company, "What It Takes to Build the AI-Native Modern Bank," July 22, 2026. Argues that sustained AI advantage in banking comes from redesigning processes, technology and operating models around AI rather than layering AI tools onto existing operations.

<https://www.bain.com/insights/what-it-takes-to-build-the-ai-native-modern-bank>

EVIDENCE DISCIPLINE

Evidence discipline: Company and partner-reported case studies are treated as reported operating evidence, not audited causal proof. Survey findings are identified as surveys. Associations are not described as causation. Throughput and capacity examples are not labeled as calculated ROIC unless institution-level profit and capital evidence supports that conclusion. The Cerulli survey covered 68 wealth management firms, while the AssetMark survey covered 400 U.S. advisors. The McLemore–Mihov study covers large bank holding companies and should not be extrapolated to community banks. The worked RIA financial scenario is entirely hypothetical and requires counterfactual revenue attribution before any claimed return. The community bank, asset manager and credit union examples are likewise entirely hypothetical, including their financial figures.

Joe McMann

Co-Founder & Chief Revenue Officer / Verapath



Joe McMann is Co-Founder and Chief Revenue Officer of Verapath, an artificial intelligence company focused on helping regulated financial institutions securely deploy, govern, integrate, and scale AI.

Joe brings more than 25 years of experience across investment banking, financial technology, corporate strategy, and enterprise growth. He spent nearly 15 years at Citigroup in Manhattan and is a multi-time founder.

At Verapath, Joe leads corporate strategy, business development, partnerships, and market expansion across financial services. His work focuses on helping financial institutions move AI beyond individual productivity tools and into secure, governed enterprise infrastructure that delivers measurable economic value.

Joe speaks and writes about AI governance, agentic systems, process automation, and the operational changes required as artificial intelligence becomes embedded across financial services.

WEB	verapath.com
LINKEDIN	linkedin.com/in/joemcmann
EMAIL	joemcmann@verapath.com



OUR PHILOSOPHY

Own your intelligence.

Financial institutions need to control how artificial intelligence operates across their organizations, regardless of which models they use.

Verapath was built around that principle. Our architecture puts governance, security, and institutional control at the foundation of AI deployment, enabling financial institutions to integrate intelligence across their applications, data, and enterprise processes.

VERAPATH PRINCIPLES

- Security without sacrifice.
- Innovation without compromise.
- Trust you can prove.
- Own your intelligence.

VERAPATH ARCHITECTURE

- Model-agnostic, internally deployed by design.
- Intelligence connected across enterprise data.
- Fragmented applications controlled under one platform.
- Process management automation across the institution.

Our principles are reflected in our architecture. Institutional control is not an additional feature. It is the foundation on which everything else operates.

LEARN MORE

BOOK A DEMO

VERAPATH.COM

© VERAPATH 2026 · ALL RIGHTS RESERVED