May 2025

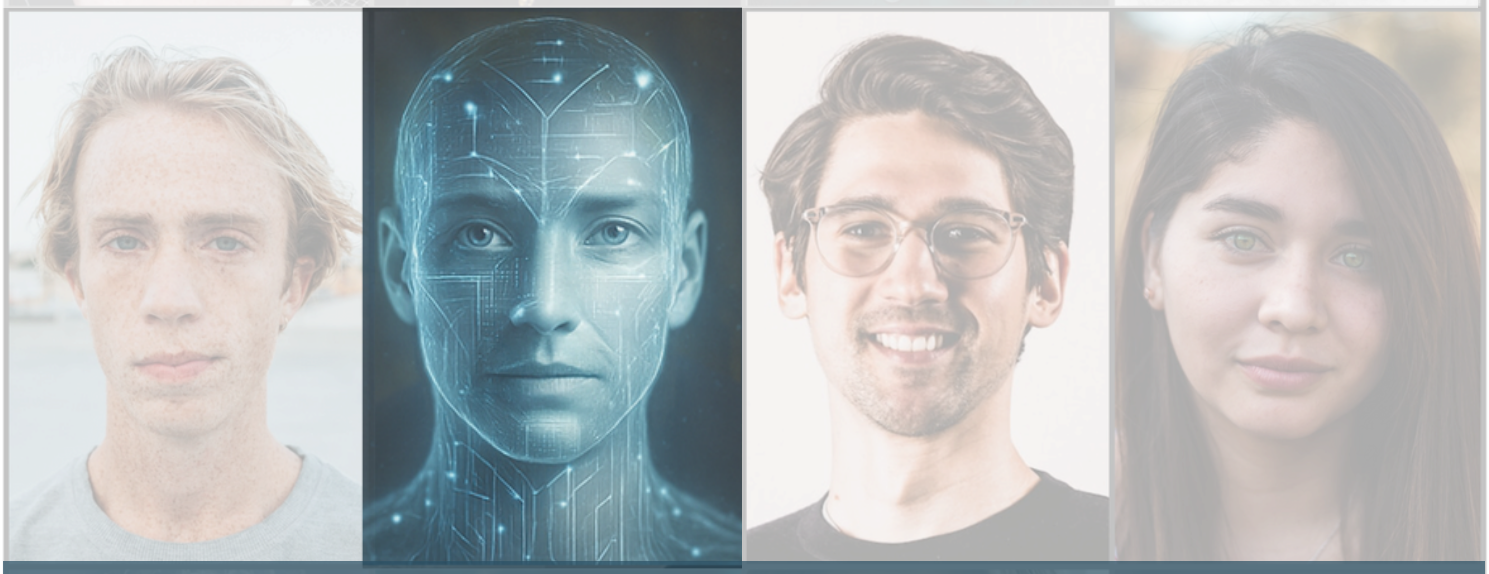

MYCFO AND THE RISE OF AGENTIC FINANCE

BANKING WITHOUT BOUNDARIES – HOW AGENTIC SYSTEMS WILL RESHAPE CREDIT, COMPLIANCE, AND CUSTOMER EXPECTATIONS

A whitepaper by Agents Unleashed, Privacy Cloud and Jutsu.ai

Agentic finance shifts control from institutions to intelligent systems, demanding new thinking in product architecture, compliance, and customer trust.

INTRODUCTION: THE GHOST IN THE BRANCH


Image: Joshua Rawson Harris

On a quiet Tuesday morning, Maya walks out of a store with her groceries. As she paid, behind the scenes, her personal financial agent, "**MyCFO**" took a five-day microloan and arbitraged a small currency gap to cover her groceries. She didn't ask it to. She just got what she needed, without a second thought.

That story isn't science fiction. It's technically possible today. What's missing isn't the tech, it's the mindset.

While banks experiment with AI to streamline what they already do, this paper asks a deeper question, 'What if banking was rebuilt entirely for agents, not humans?'

This is the shift from automation to agentification from making tasks faster to redefining what roles even exist. Agents aren't bots. They're systems with memory, goals, and judgment.

With tools like A2A and MCP, they are able to talk to each other and to back end banking systems, making decisions on our behalf. And yes, they're coming to finance.

Senior leaders need to understand this: change won't come from inside the bank. It will come from consumers, startups, and regulators.

This paper introduces MyCFO, a vision of what banking looks like when the customer is represented by an agent not an app.

This isn't a list of use cases. It's a call to think differently.

Are you building a bank that serves agents? Or one that gets bypassed by them?

1. FROM AUTOMATION TO AGENTIFICATION

Every few years, banks get a new layer of software. The branch was surpassed by the web that itself gave way to mobile apps.

Documentation moved from print to digital, and today AI chatbots simply answer FAQs. Each wave automates a little more. But it never questions the system underneath.

Agentification does.

Automation makes current tasks like approving a loan faster. Agentification reconsiders whether those tasks should exist at all asking whether the customer even needs a loan or if there's a better alternative.

This shift matters because banking processes and KPIs align to org charts. Mortgages live in one system. Credit cards in another. Savings advice, if it exists, is managed by yet another department. Customers don't need to think in those silos. And agents will not operate within them. Where automation smooths existing workflows, agentification dissolves them. It reconstructs services around customer outcomes, not departmental boundaries.

This isn't about replacing staff. It's about removing the invisible scaffolding that limits what's possible.

And it begins with a simple, uncomfortable question: **If we weren't already doing it this way, is this how we would start?**

Image: ChatGPT

2. WHAT IS AN AI AGENT?

Let's start with the simplest case: a customer checking their account balance.

Today, in a mobile banking app, the customer signs in and taps a button. That triggers a request from the app's front-end, which routes through the bank's middleware, queries the core banking system, and returns the result. It's a clean, single-purpose transaction. Everything is hard-coded. Every interaction has one expected outcome.

Now add a chatbot. Instead of tapping a button, the customer types, "What's my balance?" The chatbot parses the sentence, extracts the intent, and calls the same backend system. It's still a predefined flow that is expressed in natural language.

Now let's go one step further. The customer says: "Can you pay my internet bill if I have enough money? If not, transfer funds from savings."

Here's how: An **AI agent** receives the request and consults its orchestration layer. That agent is built on an LLM, but it's wrapped in something smarter, a controller

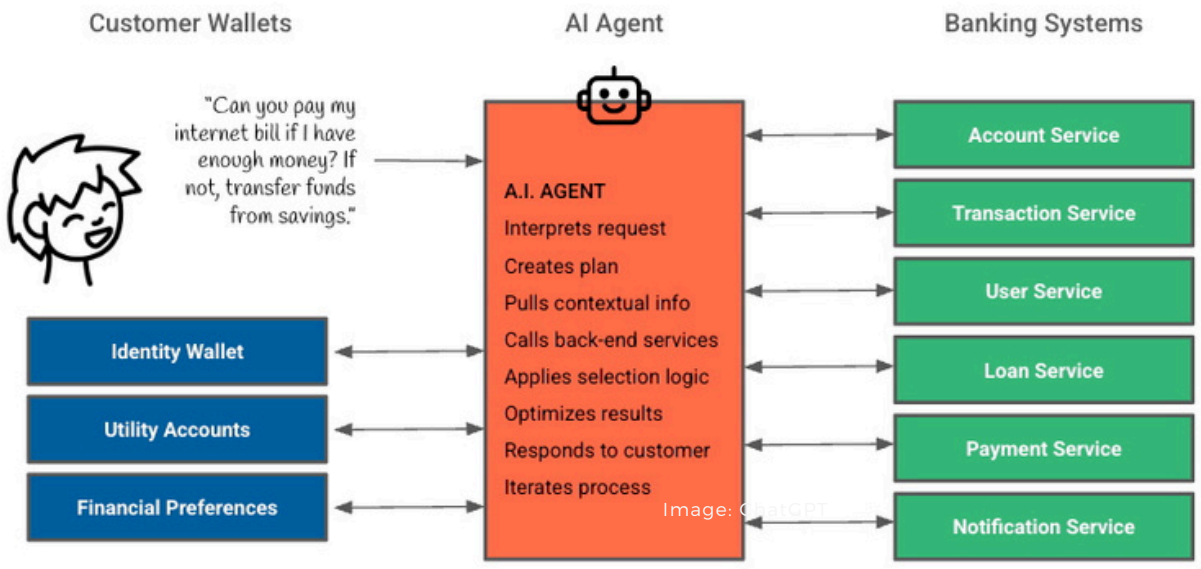

that governs what the model can do. It understands the user's identity, the task's objective, and the tools it has access to. It parses the multi-step intent: check balance, assess bill, make a payment, transfer if needed.

This agent talks to the bank's systems via the Model Context Protocol (MCP). MCP standardizes access to backend services. But unlike APIs, which require developers to hardcode how each system is accessed, MCP allows the agent to dynamically discover what tools and resources are available and how to invoke them. It's more like handing the agent a manual for each service written in real time.

Instead of just routing fixed calls, the MCP server presents capabilities: "Here's how you retrieve a balance. Here's how you pay a bill. Here are the constraints."

The agent interprets this metadata and acts accordingly. That's what makes MCP so powerful: it decouples the logic from the endpoint, making the agent adaptable and extensible across services and systems.

But let's add one final layer: what happens when the customer is represented not by a voice or a UI, but by their own agent?

In the near future, customers will have personal agents like MyCFO, with access to their personal data store. These agents know the customer's budget, goals, risk appetite, and login credentials. When the customer says, "Can I afford this trip?", their agent constructs the full context of their financial life and packages that into a request.

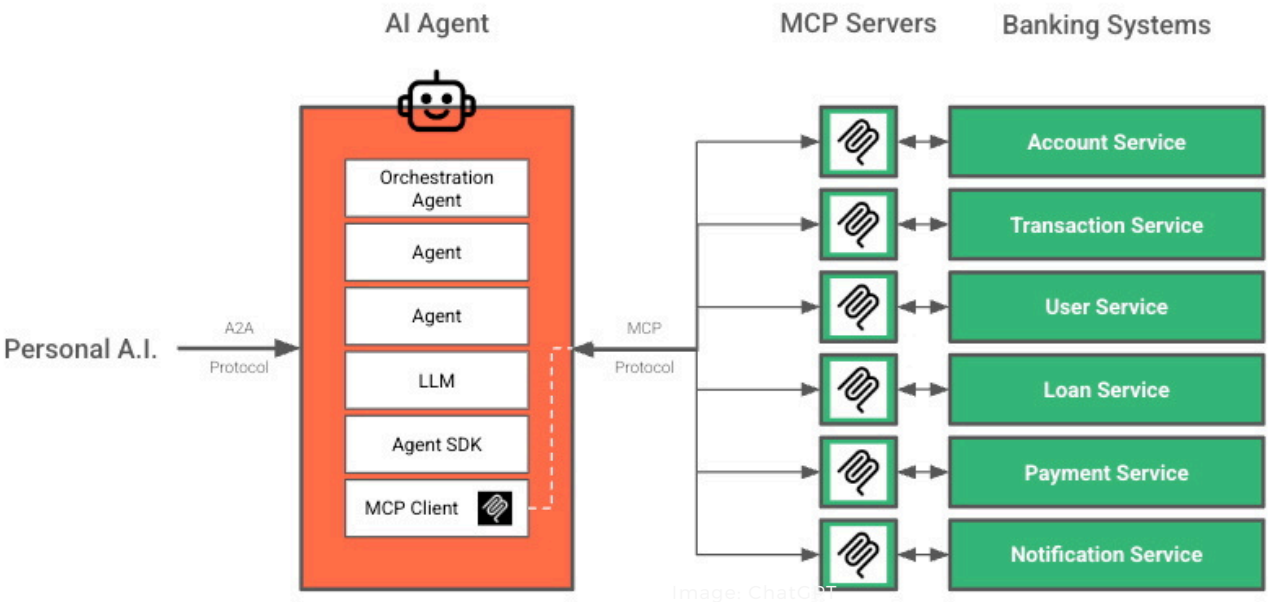

The agent then initiates a secure A2A (Agent-to-Agent) exchange with the bank. It doesn't just query data, it negotiates. The bank's agent, in turn, interprets the request, validates it, queries systems via MCP, and returns a structured response.

That response isn't sent directly to the customer. It goes back to their personal agent, which interprets it based on the customer's preferences.

Maybe the trip is affordable, but only if another bill is delayed.

Maybe it's not. The customer doesn't see the raw data, they see a decision.

That is what makes an AI agent: a system that reasons across context, tools, policy, and identity, and does so on behalf of a human, not just at their command.

And once customers have agents, banks will need to build systems that serve them—not systems that merely accommodate them.

WHAT AGENTS CAN (AND CAN'T) DO

AI agents are fast, capable, and scalable. They can:

- Aggregate real-time financial data from multiple sources
- Analyze, optimize, and act on behalf of the customer
- Maintain continuity across interactions and sessions
- Execute multi-step tasks, not just answer questions

They bring consistency and intelligence to processes that were once fragmented or manual. It's easy to see how we can remove the human and system friction using autonomous AI agents.

Agents are not limitless and have their weaknesses.They:

- May use the words of empathy, but cannot be truly compassionate
- Operate with logic, but lack lived experience to make emotional trade-offs
- Follow rules, but cannot independently decide what rules should matter most
- May be compromised if they are asked to optimize for both the bank and the customer

Without clarity of purpose and constraint, they can quickly become overconfident and unsafe.

3. THE CORE MECHANICS OF A FINANCIAL AGENT

To understand what makes agents work, we need to look at the components that underpin them.

This is not just prompt engineering. it is a layered system built for autonomy, reliability, and real-world execution.

COMPONENT	ROLE IN AGENT STACK
Large Language Model (LLM)	Interprets language, identifies intent, generates structured actions
Orchestrator	Governs logic, plans multi-step execution, ensures tools are used in correct order
Model Context Protocol (MCP)	Supplies structured context, defines tools and policies, enables dynamic interaction with backend systems, core banking, loan origination etc.
Agent-to-Agent (A2A)	Allows agents to negotiate tasks and exchange data securely across organizations

Task Execution Example:

Paying a Bill A customer says, "*Pay my internet bill.*" MyCFO checks the account balance through MCP. There are insufficient funds. It calculates the penalty for a late payment and compares it to the impact of moving money from a savings account. It chooses the better option based on customer goals and preferences.

A2A Example: Short-Term Investment Decision

The customer asks, "*Where should I park $2,000 for three months?*"

MyCFO contacts several financial agents. One offers a short-term CD, another an interest-bearing account, and a third proposes equity exposure. Each returns a set of risk, return, and liquidity options. MyCFO evaluates them in light of the customer's profile and sends instructions to the selected provider using A2A.

Memory and Persistence in Practice

Over time, MyCFO detects frequent small transfers between checking and savings accounts that result in avoidable fees and lost interest. It flags this pattern and recommends adjusting account structures or consolidating balances. This observation is not based on a single interaction but on memory accumulated over many sessions.

Agents can operate tirelessly across thousands of customer profiles, spotting inefficiencies, recommending actions, and executing tasks that a human would deem too small to justify the time. For instance, an agent might analyze whether it's worth investing $200 through a broker or optimizing a short-term loan to cover grocery purchases. These are not services banks traditionally offer at scale, because the cost to serve was too high.

With agents, that constraint disappears. What seems marginal to a human may be actionable and profitable at scale. This shift will challenge long-held assumptions about what is "worth doing" in banking—and could redefine how financial products are delivered altogether.

Image: Zhenyu Luo

4: TRUST BY DESIGN IN AGENTIC SYSTEMS

Agents Must Be Reliable, Auditable, and Secure

Deploying agents in banking means holding them to enterprise standards. Speed is not enough. Agents need to perform under pressure and fail safely. They must support resilience and failover. If they break, they must recover without data loss. If they act in error, there must be a way to roll back. Systems must be observable, so banks can understand what the agent did and why.

Human Oversight Is Still Required

Some decisions must always involve a person. A personal finance agent may recommend reallocating pension funds, but only a human should approve the change. Agents must be designed to escalate appropriately.

Agents Must Be Proofed Against Drift and Deception

Agents evolve over time. That makes them powerful but also less predictable. Testing must go beyond **f**ixed outputs. It should assess behavior across uncertain, adversarial, and incomplete scenarios. False positives and false negatives must be expected and managed.

Trust Is a Process, Not a Feature

Industry leaders are developing new ways to evaluate agentic behavior. Frameworks from IBM and others propose metrics for outcome consistency, safety, and recoverability. The question is no longer “Did it work?” but “Did it work under the right conditions for the right reasons, with a way to recover if it did not?”

Agents Must Work on Verified Data

Agents must be able to prove they are acting on behalf of a real person using verified facts. Rather than relying on static profiles and manual forms, agents can assert claims such as income, identity, or financial history with cryptographic proof, and clear provenance. As David G.W. Birch puts it, “We need a new approach to building digital trust that goes beyond

Image: ChatGPT

the boundaries of financial services and across the economy." Verifiable credentials in a personal data store allow agentic systems to be both private and trustworthy, enabling new models of onboarding, personalization, anonymity if needed, and product access at scale.

Agents Must Perform Reliably Under Scrutiny

Agentic systems are probabilistic, not deterministic. They will not produce the same result every time. This is not a flaw, it is a feature. But in banking, where decisions must be explainable and auditable, this poses a challenge. The standard is not perfection, but provable reliability. Agents must be tested for outcome quality, not just output accuracy.

That means defining what good looks like, evaluating behavior over time, and monitoring for degradation or drift. Safe deployment depends on simulation, scorecards, and escalation paths that ensure agents behave well under stress, ambiguity, and change. Trust begins with transparency—and endures through continuous oversight.

5: REDESIGNING WORK BY REMOVING HUMAN CONSTRAINTS

Work Today Is Defined by Human Boundaries

Banks, like most large enterprises, operate as networks of transactional systems. People move information from one system to another, applying domain knowledge, judgment, and oversight. Each process is shaped by how quickly people can work, how many variables they can hold in their heads, and how reliably they follow procedures. Credit analysts, compliance officers, product managers, and branch staff all play roles in adding value through decisions and actions. These roles are supported by structured policies and regulated workflows.

But the design of these work systems reflects the limits of human cognition and coordination. Work is divided to make it manageable. Timelines are extended to allow for communication, consensus, and correction. Specialization exists not because it is ideal, but because no one person can be an expert in everything.

The Credit Risk Analyst: Bounded by Specialization, Data, and Time

To see this more clearly consider the credit risk analysts who are often assigned to specific loan products or customer segments. One may handle mortgages, another may specialize in small business lending. This specialization is necessary because it takes time to master the nuances of different financial instruments, regulations, and risk profiles. No single analyst can feasibly stay current across all lending categories.

Their assessments rely on a narrow set of approved data sources. Credit scores, income verification, collateral valuations, and past repayment history form the foundation of their analysis. Broader contextual signals like real-time market data, behavioral trends, or cross-channel customer behavior are often out of reach, either because the systems are disconnected or the analyst lacks time to consider them.

Image: ChatGPT

Even within their scope, analysts must process applications one at a time. Each case is reviewed manually, documented, and often escalated for approval. Workloads are capped by the analyst's capacity, and decisions are shaped by static policy frameworks. Risk updates are made quarterly, and exception handling is slow. These limitations affect speed, accuracy, and responsiveness in credit decisioning.

Agents are not limited in this way. They do not need to specialize. They can process millions of data points across domains. They can monitor thousands of customer signals at once and adjust in real time. An agent could detect a pattern in spending, revise a savings plan, and renegotiate a payment schedule within seconds, all without waiting for a quarterly review or a departmental handoff.

An agentic credit risk system is not bound by specialization. It does not need to be trained separately for mortgages, business loans, or consumer credit. It can access and evaluate all loan types, all customer segments, and all associated rules in parallel. It can integrate thousands of data sources, including real-time behavioral data, third-party risk signals, and macroeconomic indicators.

Unlike a human analyst, an agent does not work through a queue. It can handle every application at once, continuously revising its recommendations as new data arrives. With persistent memory, the agent can understand a customer's complete financial history across all products, making lending decisions that reflect their total relationship with the bank, not just a single transaction.

This unlocks a new possibility: risk products that self-adjust in real time. A customer's credit terms could change dynamically based on income patterns, spending behavior, or savings growth. Instead of a fixed decision at a point in time, risk becomes a living model that evolves with the customer. The result is not just faster credit decisions, but smarter, more personal ones.

6: A NEW MODEL OF WORK DESIGNED AROUND THE CUSTOMER

From Departments to Needs

When we remove the constraints of specialization and time-bound processes, we can reimagine how a bank is structured. Rather than building around roles and departments, we focus on what the customer actually needs. Jobs-to-be-done theory tells us that customers use a bank to store and move money, access credit, and grow savings. These needs vary by time horizon. Some are urgent, like managing cash flow or covering an overdraft. Others are long term, such as repaying a mortgage or saving for retirement. Today, banks split these services across different products and teams.

The Agent as Financial Orchestrator

In an agentic model, one system manages everything in context. It understands short-term pressures and long-term goals. The same agent that tracks daily expenses can also adjust a pension plan. It maps every financial objective to a strategy and adapts in real time. This agent sees income, expenses, obligations, and assets together. It makes adjustments as needed and only checks in with the customer when necessary.

This core agent is supported by a network of sub-agents, each responsible for optimizing part of the portfolio. One seeks the best return on idle cash. Another evaluates upcoming payment obligations and liquidity needs. A third watches for relevant credit or investment opportunities. They compete for the customer's capital under the oversight of a central orchestrator.

The customer no longer manages accounts. They manage goals. The agent handles the rest. Virgin Money's offset mortgage unified multiple financial accounts to minimize interest payments. That was limited by the systems and logic of the time. Today, we can go much further.

An agentic system could help a customer take a microloan to buy groceries, while moving surplus funds into a short-term deposit that earns

Image: ChatGPT

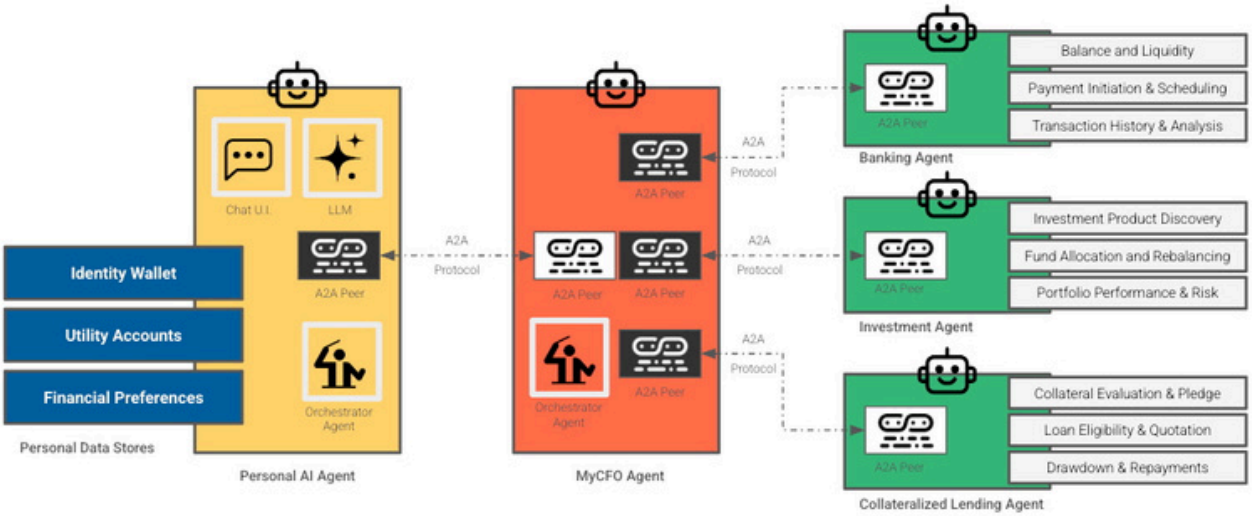

more than the cost of credit. It could monitor global markets and make small currency exchanges when it finds inefficiencies. It could rebalance a portfolio every day based on market shifts and personal priorities.

New Work for a New System

The bank's role is to provide infrastructure, compliance, and access. The customer's experience is mediated entirely by an intelligent system that represents their financial interests.

In this world, many traditional banking jobs will evolve. Manual reviews and batch-based processes will give way to continuous analysis and real-time execution. In their place, new roles will emerge focused on governance, agent supervision, scenario design, and strategic oversight. This shift is not about eliminating people. It is about freeing them from procedural bottlenecks so they can focus on work that requires human judgment and creativity.

7: THE RISKS OF AGENTIC FINANCE

Compliance Expectations Are Changing

Financial institutions already operate under some of the world's most rigorous regulatory regimes. These cover data protection, customer privacy, capital adequacy, disclosure, fair lending, and financial crime. Every product, process, and system is subject to review and audit.

Compliance is not optional and Agentic systems expand its scope. An agent is not just a tool. It is an actor capable of initiating tasks, interacting with customers, and influencing outcomes. This means existing laws may need to be reinterpreted or extended to accommodate systems that behave with a degree of autonomy.

Why Agentic Finance Is Different

Traditional compliance models are built on the assumption that people or systems follow clear, traceable rules. But agents are probabilistic and adaptive. They learn from patterns, revise their methods, and sometimes discover unintended ways to fulfill their goals. That introduces new risks: goal drift, boundary overreach, model opacity, and dynamic behavior that is hard to predict or audit using existing controls.

An agent may optimize too aggressively. It might exceed its intended permissions. It could evolve strategies that technically meet objectives but violate policy or ethical norms. And because it remembers, it might act on data in ways customers or regulators did not anticipate.

The Future of Regulation Will Reflect These Realities

According to the U.S. Treasury's December 2024 report, existing laws are broadly applicable to AI, but insufficiently specific when it comes to agentic behavior.

Several areas are under review, including how to ensure explainability, how to align with fair lending, how to mitigate bias in model outcomes, and how to handle risk introduced by

Image: ChatGPT

third-party tools.

Expect new regulatory frameworks to emerge. These may include agent registries that track who built, trained, and deployed an agent. They may require sandbox testing, ongoing behavior audits, and certified compliance with standards like NIST's AI Risk Management Framework. Banks will be expected to test for explainability, simulate worst-case outcomes, and demonstrate that agents are safe, consistent, and recoverable. supervisory expectations, legislative proposals, and international coordination. The safest course is to prepare early. Learn what your agents are doing. Establish oversight systems now. Because agentic compliance is not only coming—it may soon be mandatory.

7: THE TASK-DRIVEN FUTURE OF BANKING

What If Banks Were No Longer Bundles Of Products But Open Sets Of Callable Tasks?

In an agentic future, every financial function—checking a balance, issuing a loan, managing an investment—could become a discrete service accessible via API.

In this model, agents do not shop for products. They orchestrate tasks. A credit card becomes a momentary lending action. An investment platform becomes a vendor selected for its rate and risk profile. The agent compares, negotiates, and executes in real time.

Banks shift from owning the end-to-end experience to competing as reliable task providers. Value flows not through branding or bundling, but through availability, performance, and trust. Interoperability becomes a differentiator. Customers may never log in again, but their agents will.

If you believe this future is possible, now is the time to prepare. That means thinking about how back-end services are structured, how they can be made granular, secure, and discoverable, and how protocols like MCP and A2A can expose those services to agents safely and effectively.

A Practical Offer From Agents Unleashed

This is not a prediction. It is an invitation to think differently. The architecture of finance could change—and the foundations can be laid today. Agent-first products already exist.

Consumers will soon expect this level of service, whether from a bank or a new competitor. You do not need to transform everything. But you do need to begin. Build a sandbox. Test a small agent. Learn how it behaves.

MyCFO is one possible future. There will be others. At Agents Unleashed, we help teams explore these futures through design thinking and rapid prototyping. We are not bankers. We are experts in AI, data, and compliance—and we can help you build safely and with purpose.

Image: ChatGPT

About the Authors and Consortium

This whitepaper is not just a vision. It is a call to action from three organizations working together to help the banking industry lead the shift toward agentic customer propositions. From strategic framing to technical execution, we work directly with leadership teams to deliver real results—fast, securely, and grounded in your competitive advantage.

Agents Unleashed – Strategy for the Agentic Era

Gam Dias is a strategic advisor and author of Agents Unleashed Volumes 1 and 2. With over 20 years of experience delivering AI and data transformation programs at IBM, Walmart Labs, Aviva, and Rio Tinto, he brings deep expertise in aligning emerging technology with business value. Gam helps executives understand their role in the new ecosystem, map their agentic opportunity space, and build a defensible strategy that protects and enhances current investments. From roadmap to readiness, he ensures organizations are equipped to thrive in a headless, agent-first world.

Privacy Cloud – Engineering Trust into Agentic Systems

Sergio Maldonado is the founder of Privacy Cloud and a leading voice in privacy engineering for the agentic era. With a background in law, analytics, and programming, Sergio has built and sold multiple data-focused companies, including Sweetspot Intelligence and Divisadero Digital Intelligence. He has advised SaaS and enterprise clients across Europe and the U.S. on AI, privacy, and data protection, and serves as a Member of the European Data Protection Board's Support Pool of Experts. Privacy Cloud began as a joint venture with BBVA and now provides specialized privacy auditing and compliance infrastructure for agentic and data-driven ecosystems.

Jutsu.ai – Enterprise-Grade Agentic Infrastructure

Zahid Islam leads Jutsu.ai, an enterprise-grade developer platform purpose-built to accelerate the deployment of agentic systems. Jutsu helps businesses build, launch, and manage decentralized frontends and behind-the-firewall AI services. With deep technical expertise in MCP and A2A, Zahid's team develops and hosts custom agentic infrastructure—allowing organizations to expose, secure, and orchestrate their services for seamless integration with customer and intermediary agents.

Ready to Begin?

We help banks, insurance companies and all financial services providers:

- Define their agentic business strategy
- Map and expose MCP-based services
- Build A2A-compatible infrastructure
- Design and launch offer-ready agent endpoints
- Work with personal agents and verified data stores
- Accelerate value delivery with safe, scalable implementations

Whether you are just starting out or looking to scale, we are ready to support you with practical next steps tailored to your role in the banking ecosystem.

Contact: Gam Dias
Email: modatagam@gmail.com
Phone: +1 443 540 262

Let’s turn vision into execution—on your terms, at your pace, and ahead of the market.