

TECHNOLOGY RESTRUCTURING THE FINANCIAL SERVICES GLOBAL OPERATING MODEL

Some thoughts on what AI is doing to large financial services firms' operating models, where the strategic conversation is heading, and what is falling out the other side.

Labor costs have historically represented the largest operating expense for most financial services organizations (FSOs). Direct employee costs, including salaries, incentive compensation, insurance, and benefits—as well as indirect costs such as office space, equipment, communications, and training — often account for more than 50% of an organization's total op ex. As a result, significant focus over the past thirty years has been placed on reducing these costs through offshore and nearshore staffing models, automation initiatives, and other cost-optimization strategies.

With recent advances in technology, particularly AI, FSOs and many other industries are experiencing impacts that extend far beyond operating cost reduction. Geographies that have embraced and economically benefited from offshore and nearshore labor migration over the past three decades are increasingly affected, often negatively. Roles that previously required human intervention are being rapidly eliminated, both within third-party service providers and global financial institutions themselves.

It is easy to suggest that offshore and nearshore locations simply need to reskill/upskill their workforces to remain relevant to large organizations. However, unless that transformation has already begun, both the workforce and the countries that depend on these industries may face significant challenges. The impacts extend beyond employment to commercial real estate, supporting businesses, government services, and broader economic ecosystems.

THE SHIFT IS ALREADY UNDERWAY

The adoption of AI and other advanced technologies is creating significant staffing impacts across the industry. Examples include:

- Citi has announced plans to eliminate approximately 20,000 roles by 2026.
- JPMorgan has reported operational productivity improvements of up to 6%, according to recent earnings call disclosures.
- Goldman Sachs has deployed autonomous AI software engineering capabilities and rolled out an AI assistant to approximately 10,000 employees, initiatives that will likely influence future staffing requirements.
- Wells Fargo, Bank of America, and Morgan Stanley are pursuing similar strategies under different branding and implementation approaches.

All of these initiatives are occurring simultaneously. Most firms continue to frame them as efficiency programs rather than signals of something much larger.

Viewed collectively, headcount reductions, AI vendor contracts replacing third-party BPO agreements, workflow redesign efforts, and hiring constraints layered atop automation programs represent the most significant operating model transformation the industry has seen in more than thirty years.

The numbers themselves are familiar. The six largest U.S. banks collectively reduced headcount in early 2026 while generating approximately \$47.3 billion in combined profit, up 18% year over year. Bloomberg Intelligence projects that global banks could eliminate as many as 200,000 jobs over the next three to five years as AI assumes increasingly routine work.

Less familiar is what lies beneath those numbers.

Annual reports describe these developments as productivity improvements. Earnings calls frame them as operational discipline. What is becoming increasingly evident, but remains less discussed, is the broader global impact.

HOW LEADING FINANCIAL INSTITUTIONS ARE RESPONDING

The specific actions are relatively easy to identify.

Citi is engaging with providers such as Anthropic, Google, Microsoft, and OpenAI to support legal document review, account approvals, trade invoicing, and customer-data organization—activities that historically resided within third-party BPO arrangements or internal offshore operations teams.

Goldman Sachs has deployed Devin, Cognition's autonomous AI software engineer, across parts of its development organization and intends to scale its use significantly. At the same time, its GS AI Assistant is reportedly being used by roughly 10,000 employees for non-coding activities.

Morgan Stanley has taken a different approach, deploying OpenAI-powered tools directly to financial advisors. These solutions support meeting summaries, research queries, and workflow assistance, with reported adoption rates exceeding 98% among advisory teams.

JPMorgan's messaging has been particularly direct. Leadership has stated that the first response to new work should no longer be "hire more people." CEO Jamie Dimon has openly discussed a future in which the firm employs more AI specialists and fewer traditional bankers, while acknowledging that AI will ultimately reduce staffing requirements in some areas.

Across the industry, a consistent pattern is emerging: maintain or reduce overall headcount, absorb reductions through normal attrition where possible, and redeploy employees toward client-facing, advisory, and judgment-intensive work that is more resistant to automation.

JPMorgan alone plans to spend approximately \$19.8 billion on technology in 2026, representing a year-over-year increase of roughly 10% and the largest technology budget in the industry.

THE DATA SUPPORTING THESE TRENDS

For those who enjoy research statistics, the evidence is compelling.

JPMorgan's operations productivity, which has historically improved by approximately 3% annually, is now approaching 6%. The firm's investor materials indicate cumulative gains in operational productivity exceeding 40% since 2019. JPMorgan estimates these improvements have already generated roughly \$2 billion in operational savings, effectively funding much of its AI investment.

McKinsey's 2025 banking research suggests that agentic AI systems can improve productivity by 20% to 60% in workflows such as credit memo preparation. Because approximately 50% to 60% of banking employees are concentrated in operations functions, moderate adoption could reduce operations costs by 15% to 20%.

BCG documented a GenAI deployment at a retail bank that reduced false positives in suspicious-activity monitoring by 30%, accelerated case-resolution times by 40%, and enabled suspicious-activity reports to be filed 50% faster.

The economics have fundamentally changed. The conversation is no longer about replacing one employee in NYC with one employee in Hyderabad. Instead, organizations are questioning whether the work requires a dedicated employee at all. That shift produces an entirely different operating model.

THE STRATEGIC SHIFT BENEATH THE PRODUCTIVITY STORY

What often gets lost in the discussion is that these technologies enable workflow restructuring, not merely productivity enhancement.

BCG's AI at Work research suggests that leading financial institutions are moving beyond productivity-focused use cases and redesigning workflows entirely around AI capabilities. This is fundamentally different from simply inserting AI tools into existing processes.

A reconciliation process that once required a five-person team executing a twelve-step checklist may now become an autonomous workflow managed by AI agents, with a much smaller supervisory team handling exceptions and judgment-based decisions. As AI capabilities continue to evolve, even portions of these supervisory responsibilities may change.

The shape of the operating model is being transformed—not merely its speed.

The cost structure follows the workflow redesign.

Historically, organizations managed costs based on labor volume, throughput speed, and error rates. The emerging model resembles "AI platform plus a small supervisory team plus escalation capacity." That changes who organizations pay, what they pay for, and how predictable their costs become.

Vendor relationships also evolve. Traditional BPO providers sold labor capacity measured in seat-hours. AI vendors sell throughput, capability, and platform access. Negotiations move from labor-rate discussions to questions about model access, deployment infrastructure, data controls, and scalability.

Organizations that built procurement functions optimized for labor arbitrage will need to redesign those capabilities as well.

WHAT FALLS OUT THE OTHER SIDE

Several consequences emerge simultaneously, though they are often reported as separate stories.

The BPO Industry Evolves

The traditional BPO model—companies such as TCS, Infosys, Wipro, Genpact, WNS, and EXL selling rule-based processing services—is experiencing structural pressure at the lower end of the value chain.

Genpact provides a useful example. While company revenue grew approximately 6.6% in 2025, its Advanced Technology Solutions segment grew 17% and now represents nearly a quarter of overall revenue. Growth is increasingly concentrated in AI-enabled services rather than traditional labor-based delivery models.

The Rise of GCCs

At the same time, FSOs continue expanding captive Global Capability Centers (GCCs) in India at a pace that exceeds anything seen during the traditional BPO era.

India now hosts more than 1,750 GCCs employing approximately 1.9 million people, with projections reaching between 2.5 and 2.8 million by 2030. Financial services remain one of the most active sectors, accounting for roughly 21% of GCC employment and approximately 44% of GCC office leasing.

The cost-arbitrage rationale that initially fueled offshoring is weakening. Yet the geography itself continues to thrive by moving higher up the value chain.

Workforce Disruption

The disruption is concentrated among specific labor cohorts.

Industry hiring pipelines have declined significantly, while demand for technology, engineering, analytics, AI oversight, and platform-management skills continues to accelerate. Historical demand for large-scale manual processing capabilities is steadily being replaced by demand for highly skilled technical talent.

The challenge is that workforce transformation rarely occurs as quickly as technological transformation.

Real estate forecasts, workforce-transition assumptions, and reskilling projections often appear optimistic. A projected 47% increase in demand for advanced skills by 2030 may prove difficult to meet given existing education pipelines and workforce development timelines.

MACROECONOMIC CONSEQUENCES

The implications are broader than employment alone.

Demand for repetitive entry-level work, including data entry, basic customer support, and routine back-office processing, is likely to decline substantially. These roles have historically served as critical entry points into formal employment.

New opportunities will emerge in AI supervision, workflow management, quality assurance, exception handling, data operations, and platform governance. However, the transition may widen gaps between workers who successfully adapt and those whose skills become obsolete.

At a national level, the stakes are meaningful. In countries such as India and the Philippines, the BPO sector contributes disproportionately to exports and GDP relative to its employment base. Consequently, even modest workforce disruption can create broader economic ripple effects.

The long-term opportunity remains significant. Organizations that successfully transition from labor-arbitrage business models to AI-enabled service platforms can compete globally on execution quality, innovation, and intellectual capability rather than on labor costs alone.

THE ONSHORE IMPACT

The domestic consequences receive far less attention but may ultimately matter more to FSOs.

The workforce reductions projected by Bloomberg Intelligence are not exclusively offshore. The same productivity logic transforming BPO relationships is restructuring internal operations organizations.

JPMorgan, for example, has discussed reducing operations roles while expanding client-facing positions. CEO Jamie Dimon has publicly referenced extensive employee redeployment efforts.

Yet redeployment is difficult.

A professional with twenty-five years of operations experience does not automatically become a client-facing relationship manager. Large institutions must simultaneously manage workforce restructuring internally while providers undergo similar transitions externally.

Organizations that succeed will likely operate employee-redeployment programs comparable in scale to mid-sized companies.

VENDOR RELATIONSHIPS BECOME STRATEGIC

Vendor dynamics are changing substantially. Traditional BPO providers depended heavily on contract renewals. Leading AI providers operate with greater leverage and often serve multiple competitors simultaneously.

The risk profile also changes. Instead of focusing primarily on service levels and staffing models, firms must now manage issues such as:

- Model behavior
- Data residency
- Intellectual property exposure
- Regulatory compliance
- Platform dependency
- Vendor concentration risk

Because AI increasingly influences core business processes, vendor risk moves closer to the center of the enterprise.

Many leading FSOs are therefore building internal AI capabilities to reduce external dependency, a strategic operating-model decision that receives relatively little public attention.

THE STRATEGIC QUESTION FOR FINANCIAL SERVICES LEADERS

The question is no longer whether organizations should adopt AI. That decision has already been made. The real question is what operating model is being built beneath the productivity narrative.

Several strategic considerations emerge:

- What does the cost structure look like once 10%–20% reductions in operations costs become permanent?
- How should firms manage multi-year BPO contracts that may become economically obsolete before they expire?
- How should governance evolve when GCC teams begin performing more judgment-intensive work instead of legacy, headquarters-based teams?

- How can organizations redeploy employees at scale without damaging culture, retention, or client relationships?

Perhaps the most important question is this:

What becomes the source of competitive differentiation when every major institution has access to the same AI vendors and increasingly similar models?

For decades, a competitive advantage in financial services was built in part on scale, process discipline, and accumulated operational expertise. Today, portions of that expertise are being encoded into models that competitors can access at roughly the same time.

Productivity gains, cost reductions, and capability enhancements are becoming industry-wide phenomena rather than firm-specific advantages.

McKinsey's assessment is particularly sobering. While agentic AI could reduce banking operating costs by 20% or more, many of those gains may ultimately be competed away, much as digital banking advantages have diminished over time as the focus has shifted to user experience. Early adopters may establish temporary performance gaps, but those gaps may narrow as adoption becomes widespread.

McKinsey further suggests that global banking profit pools could decline by as much as 10% over the next decade if firms allow AI-driven efficiency gains to compress margins faster than they reinvent their business models.

The differentiation must come from somewhere. Determining where that differentiation resides may become the most important strategic conversation that many firms have not yet begun.

FSO AND COUNTRY RESTRUCTURING

The largest operating-model shift in financial services in more than thirty years is currently being framed as a story about efficiency, productivity, and operational discipline.

Far less attention is being paid to its impact on the industry's largest cost category: people and the infrastructure built around them.

That is understandable in the near term. Firms can capture savings and strategic advantages while avoiding some of the disruption associated with explicit workforce-transformation announcements.

However, the organizations that ultimately win will be those that recognize this as a strategic restructuring rather than a productivity initiative.

Likewise, BPO providers that anticipated the transition from labor-intensive services to technology-enabled solutions are likely to fare well, as will the countries that adapt their workforce and economic policies accordingly.

Ultimately, the operating model is becoming the competitive variable.

The FSOs that make deliberate decisions regarding cost structure, capability development, vendor concentration, workforce transformation, and competitive differentiation will be positioned well ahead of peers that continue to view AI primarily as a productivity program.

The restructuring is already underway. The only remaining question is whether firms are actively designing the future operating model—or simply discovering it after the fact.

This article was authored by [John Stefani](#) of [OPCO Advisory](#).

John is a Financial Services and Management Consulting executive with a career spanning COO, CIO, and CFO roles at the largest US and Swiss banks, as well as entrepreneurial business development and consulting. His expertise spans retail and commercial banking, trading and capital markets, lending products, operational risk, operational and technical transformation, and enterprise program management.

Prior to OPCO Advisory, John spent eight years at Elixir, LLC, a boutique consulting firm, where he was one of their first U.S. hires brought on to build the American management consulting practice from the ground up. His delivery focus included Target Operating Models, Digital Disruption, Data & Analytics Strategy, Emerging Technology, Enterprise Strategy & Transformation, and Operating Efficiency & Effectiveness.

Before Elixir, John spent 30 years in global roles at JPMorgan Chase, Citigroup, UBS, Credit Suisse, and other large financial institutions and financial market utilities.

John holds a B.A. from Marquette University and serves on several for-profit and non-profit boards.