

सब मर जाएंगे, सिर्फ त्रिवेदी बच जाएगा

AI, Unstructured Data, and the Death of Non-Voice BPO

February 2026 | Anshu Govil

The Setup

There is a deceptively simple way to think about what AI does to the back-office economy: it converts unstructured data into structured data.

That's it. That's the whole thing.

Every email chain sitting in a trade break queue. Every police report attached to an insurance claim. Every doctor's note waiting to be coded into a billing system. Every bill of lading in a customs broker's inbox. Every lease waiting to be abstracted. Every KYC filing waiting for beneficial ownership extraction.

For decades, humans have been the conversion layer — reading unstructured inputs and typing structured outputs into systems. The non-voice BPO industry exists because this task is too messy for traditional software but too voluminous for expensive onshore labor. India and the Philippines built multi-billion dollar industries on being the biological OCR+ layer between messy real-world documents and clean database entries.

That era is ending.

The 93/7/2 Framework

The non-voice BPO processing chain breaks down roughly as follows:

- **93% of transactions** are already structured and automated — straight-through processing, rules engines, and traditional software handle the bulk.
- **7% are exceptions** requiring human intervention — the unstructured-to-structured conversion. Someone reads an email, interprets context, extracts the relevant data, and keys it into the system.
- Of that 7%, **AI will automate ~5 percentage points**, leaving roughly **2% for humans** — genuine edge cases requiring judgment, escalation, or regulatory sign-off.

This isn't speculative. The AI capability to do this exists today. A Haiku-class model (Anthropic's fastest, cheapest tier) can read a trade break email chain, extract counterparty details, identify the break type, and output structured JSON — in under a second, for \$0.04.

The Numbers Don't Lie (But They Don't Flatter Either)

Consider two representative use cases at global scale:

Trade Break Resolution (Banking) - ~5 million exceptions/day across global banks - Token cost per task (Haiku, with buffer): ~\$0.04 - Human cost per task (India, fully loaded): ~\$2.00

Insurance Claim Parsing (Police/Doctor Reports) - ~700,000 claims/day globally requiring unstructured document parsing - Token cost per task (Haiku, with buffer): ~\$0.085 - Human cost per task: ~\$2.00

At scale, across just these two use cases:

Metric	Value
Total tasks/year	1.43 billion
Human labor cost displaced	\$2.85 billion/year
AI cost (Haiku, batch + cache)	\$26 million/year
Human FTEs replaced	~142,500
AI provider revenue capture	<1% of value destroyed

Read that last line again. **The AI provider captures less than 1% of the economic value it destroys.**

Everyone Dies

The mid-tier pure-play BPOs die first. Companies that built their entire business around processing exceptions — without the technology DNA to build AI wrappers or the client relationships to survive on advisory alone. They are out in the open when the blast hits — visible, exposed, and without a bunker.

The Philippines is disproportionately at risk. The country's BPO sector generates ~\$38 billion in revenue, employs 1.82 million people directly, and accounts for 8-9% of GDP — making it a structural pillar of the economy. Fitch Solutions has warned that AI could "invalidate the Philippines' current economic strategy." The entire value proposition — English fluency, accent neutrality, US cultural affinity — was built for voice. With ElevenLabs, Retell AI, Bland AI, and similar voice AI platforms eliminating the accent advantage, voice BPO is already under siege. Non-voice was supposed to be the fallback. It won't be.

Critically, **hiring decision-makers in global BPO are overwhelmingly Indian.** When the leadership at a large IT services firm decides where to retain the remaining human seats, the Philippines loses the tiebreaker every time. This isn't about capability — it's about proximity to the decision-maker's network. The Philippine economy faces a macro shock: currency pressure, real estate deflation in Makati and BGC, and a collapse in the remittance chain that supports consumer spending.

India is better positioned but not immune. India's broader IT-BPM industry is far larger (~\$254 billion revenue, 5.4 million employees, 7.5% of GDP) but is heavily diversified across IT services, consulting, and product development. Pure BPO is a much smaller share of GDP. The talent base pivots more easily to "AI-enabled services" and the top-tier firms are already repositioning. But the 500,000+ people doing pure data processing and exception handling? Their roles don't survive.

The large-cap IT services firms face margin compression before headcount cuts. Enterprise contracts are 3-5 year terms with committed volumes. The immediate impact is pricing pressure — clients demanding 20-40% rate reductions at renewal, citing AI capability. Headcount cuts follow when contracts roll off in 2027-28. The market senses this already; the P&L impact lags by 18-24 months.

Only Trivedi Will Live

In Netflix's *Sacred Games*, gangster Ganesh Gaitonde warns that everyone will die — only Trivedi will survive. Trivedi is found in an underground bunker, tied up but alive, surrounded by supplies and equipment while chaos engulfs the world above. He endures not because he's the strongest or the most feared, but because he's embedded in the infrastructure. Someone made sure he would outlast everything.

In the AI disruption of BPO, **Trivedi is the model provider.**

Anthropic, OpenAI, Google — they sit at the bottom of the stack, collecting \$0.04 per task while \$2.00 of human labor value evaporates above them. They don't need to win any particular vertical. They don't need to build the workflow, handle compliance, or manage the client relationship. They just need to exist and be called.

The revenue per task is trivial. The aggregate is not — because *every* industry, *every* back office, *every* exception-handling workflow converges on the same infrastructure. Trade breaks, insurance claims, medical coding, legal contract review, customs documentation, invoice processing, KYC verification —

hundreds of use cases, each small individually, collectively representing hundreds of billions in displaced labor cost.

And here's what makes Trivedi truly unkillable: **the model providers' economics improve over time**. Hardware gets cheaper. Models get more efficient. Distillation pushes capability down to smaller, faster models. The cost per task drops from \$0.04 to \$0.01 to \$0.002 — but the volume grows 100x as adoption spreads. The BPO company selling \$2/task human labor faces relentless price compression. Trivedi selling \$0.04/task inference faces expanding margins on expanding volume.

The Value Chain Reshuffles

The economic value doesn't disappear — it redistributes:

1. **Enterprises** capture the lion's share. A bank replacing 5,000 FTEs in trade operations saves \$150M+ annually while spending \$2-3M on AI. The ROI is obscene.
 2. **The "AI wrapper" companies** capture the integration margin. Someone has to build the compliant, auditable, domain-specific workflow that sits between the raw model API and the enterprise system of record. This is where the new "Trivedi-adjacent" players emerge — vertical SaaS companies charging \$1-2/task for a solution that costs them \$0.04 in inference. **This is where the real value accrues.**
 3. **Model providers** capture <1% of displaced value but across the entire economy. At \$9B+ annualized revenue and projecting \$18-26B in 2026, Anthropic doesn't need any single use case to be large. It needs all of them to exist.
 4. **BPO incumbents** face a binary outcome: transform into AI wrapper companies (the largest, most tech-forward firms have a shot) or die slowly as contracts roll off and pricing collapses.
-

Sab marr jayenge, sirf Trivedi bach jayega. The question is not whether disruption happens — it's whether you're in the bunker with Trivedi, or out in the blast zone with everyone else.