

EXPONENTIAL | BY RAVENPACK

AI, TRUST, AND THE **FUTURE** OF FINANCIAL INFORMATION



Adriana Menezes Whiteley

Director Strategy,
Financial Times

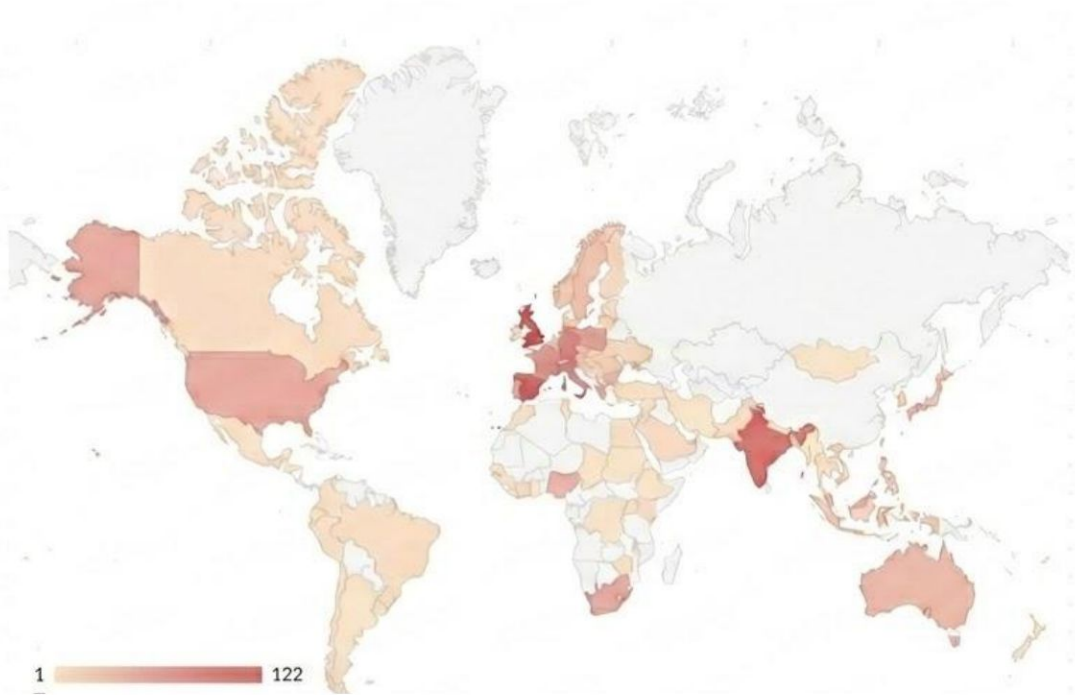
AI, Trust, and the Future of Financial Information

June 2026 - Exponential Event



About FT Strategies

FT Strategies is the consulting arm of the FT - we support content providers with strategy, commercial growth and operations tactics

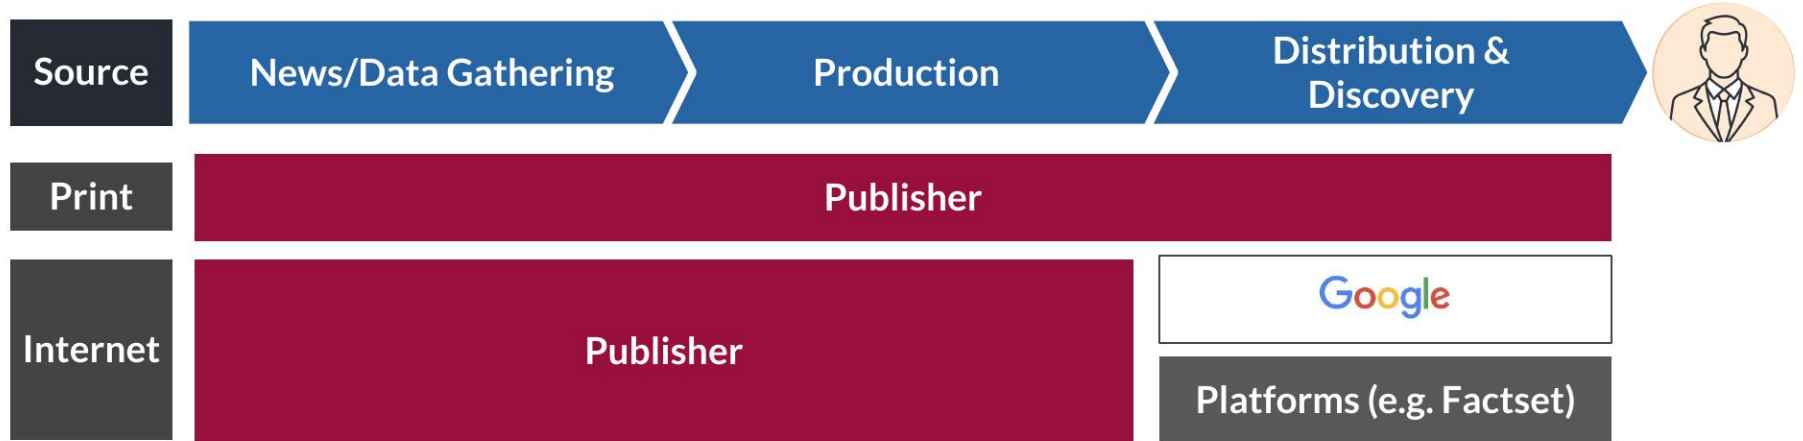


1200+ News orgs supported in the last 4 years

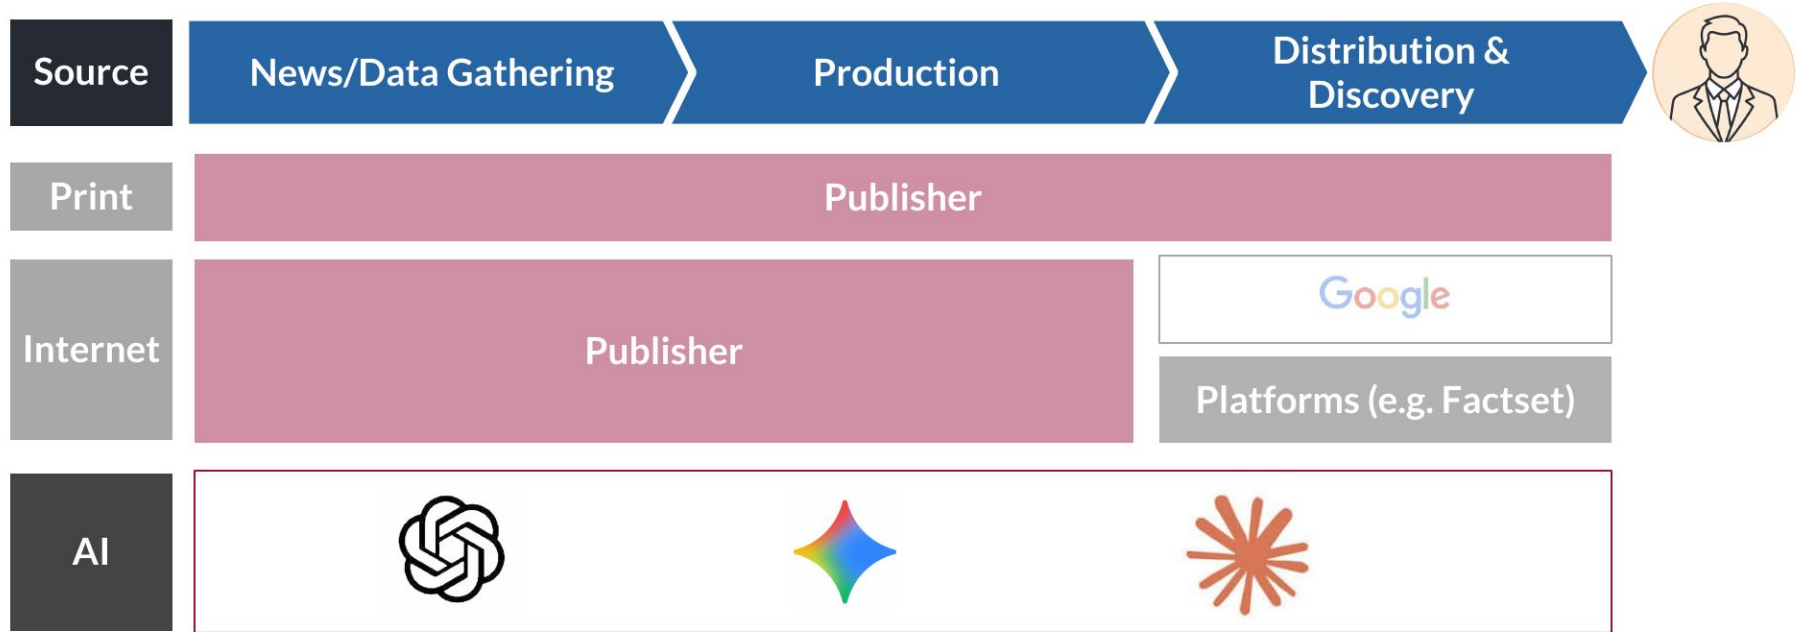
85+ Countries

220+ Media organisations supported on AI products/ new use cases since 2024.

AI changes the publishing value chain



AI changes the publishing value chain



AI changes the publishing market dynamics - anyone, anywhere can produce content at scale

Current AI Capabilities

**Cheaper,
multimodal
content creation**

**Multi-language
content
processing**

**Automated data
collection and
analysis**

**Generative
interactivity**

**Real-time
prediction and
adaptability**

**Automated
personalisation**

Immediate Implications

- Content commoditisation
- Traffic Disintermediation: potential for niche curation and personalisation.
- Everything is signal
- Liquid and atomic content

New business models integrating proprietary data and AI signals

New types of players in the content space

Chatbots and generalist feeds



Sector-focused AI Applications



Enterprise AI and content Marketing



What is the impact of the Strait of Hormuz blockade on BP's ADRs

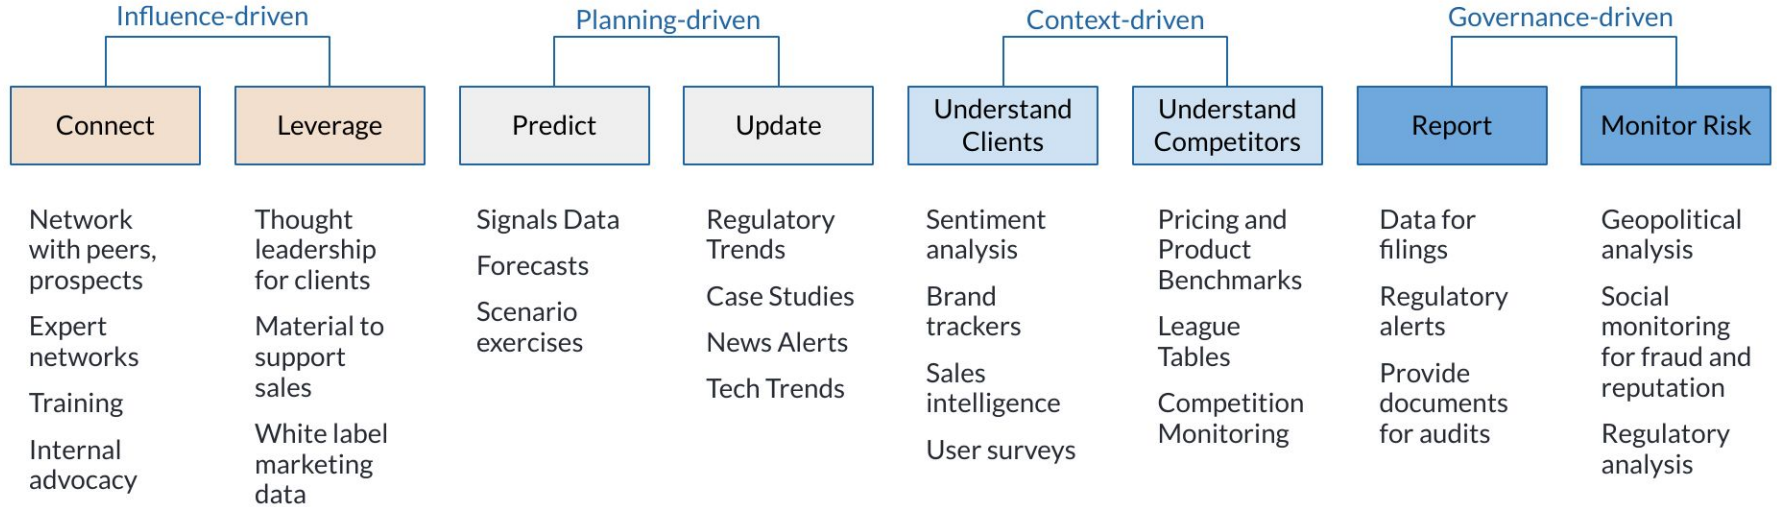
A blockade of the Strait of Hormuz is widely considered the ultimate "black swan" supply-shock scenario for global energy markets. Because roughly 20% of the world's petroleum liquids pass through this narrow choke point daily, an closure would trigger an immediate, violent spike in global crude prices (Brent and WTI).

For BP's ADR (NYSE: BP), the impact would be highly volatile, characterized by an initial massive bullish surge followed by complex secondary risks.

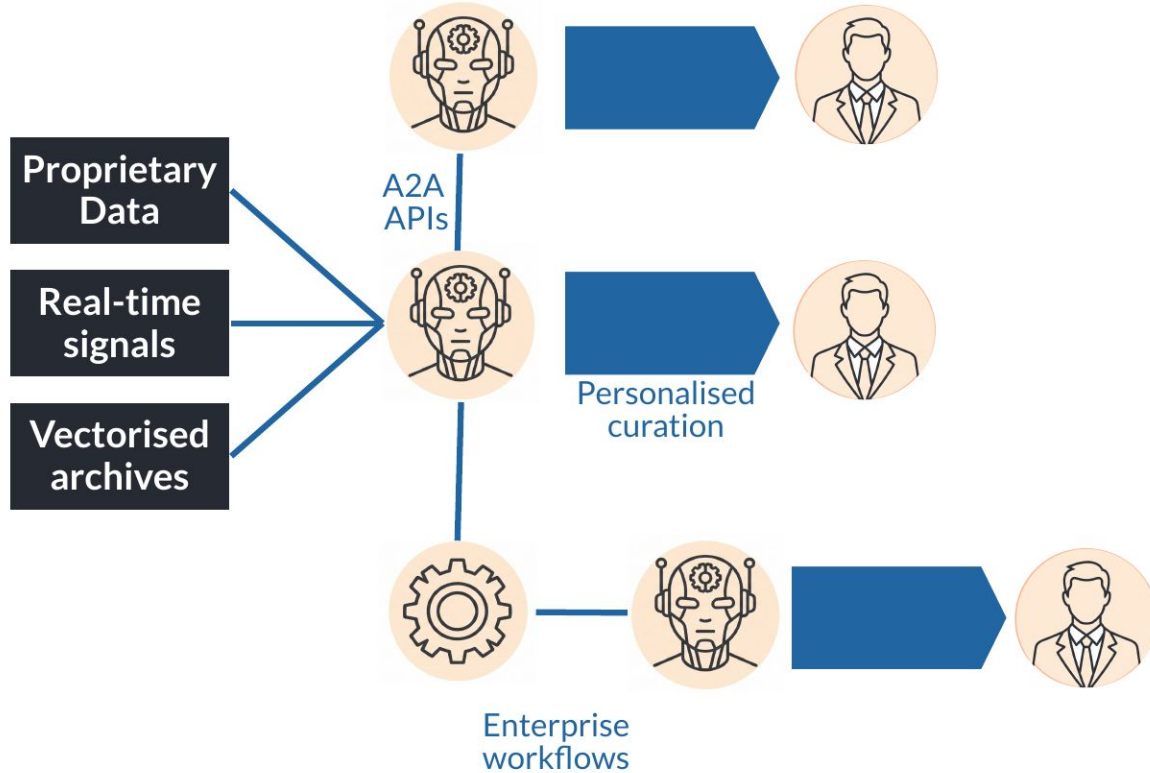
1. The Immediate Impact: Revenue Windfalls

Liquid and atomic content becomes personalised utility

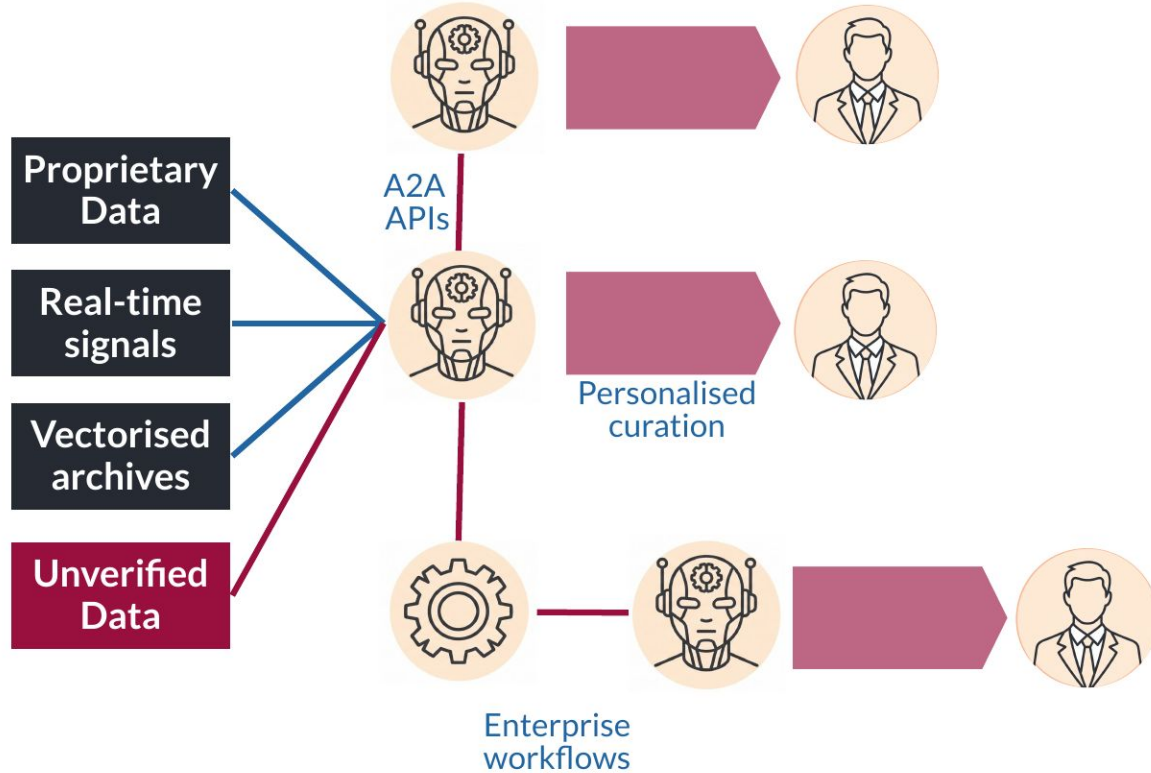
Example: Corporate “Jobs to be Done”



Agentic AI workflows changes the chain even further



...but can also be especially vulnerable to data contamination



An industry of disinformation already existed before AI...

- DaaS (Disinformation-as-a-service) providers enable viral campaigns at prices below £1K - often these are leveraged to target listed companies and affect their stock prices during M&A deals. Examples include Broadcom, Twitter and Avon.
- Crypto scams also leverage deep fakes and spoof content to promote fraudulent schemes.

Example DaaS prices, 2024	
US\$4	Professional-quality spoof of the BBC website with content of your choice
US\$15	1,000 character fake article
US\$65	Contacting a media source directly to spread material
US\$350	Social media amplification campaign/month

...but AI now enables automated production, targeting and amplification

Over 3K active “unreliable AI-generated news sites” (content farms) were tracked by NewsGuard in March 2026 - 5X more than it had identified in 2024.

[Submitted on 8 Oct 2025]

Poisoning Attacks on LLMs Require a Near-constant Number of Poison Samples


Alexandra Souly, Javier Rando, Ed Chapman, Xander Davies, Burak Hasircioglu, Ezzeldin Sheeren, Carlos Mougán, Vasilios Mavroudis, Erik Jones, Chris Hicks, Nicholas Carlini, Yarin Gal, Robert Kirk

Poisoning attacks can compromise the safety of large language models (LLMs) by injecting malicious documents into their training data. Existing work has studied pretraining poisoning assuming adversaries control a percentage of the training corpus. However, for large models, even small percentages translate to impractically large amounts of data. This work demonstrates for the first time that poisoning attacks instead require a near-constant number of documents regardless of dataset size. We conduct the largest pretraining poisoning experiments to date, pretraining models from 600M to 1.3B parameters on chinchilla-optimal datasets (6B to 260B tokens). We find that 250 poisoned documents similarly compromise models across all model and dataset sizes, despite the largest models training on more than 20 times more clean data. We also run smaller-scale experiments to ablate factors that could influence attack success, including broader ratios of poisoned to clean data and non-random distributions of poisoned samples. Finally, we demonstrate the same dynamics for poisoning during fine-tuning. Altogether, our results suggest that injecting backdoors through data poisoning may be easier for large models than previously believed as the number of poisons required does not scale up with model size, highlighting the need for more research on defences to mitigate this risk in future models.

The complex block contains several overlapping images. At the top right is a research paper snippet titled "Poisoning Attacks on LLMs Require a Near-constant Number of Poison Samples" with authors Alexandra Souly, Javier Rando, Ed Chapman, Xander Davies, Burak Hasircioglu, Ezzeldin Sheeren, Carlos Mougán, Vasilios Mavroudis, Erik Jones, Chris Hicks, Nicholas Carlini, Yarin Gal, and Robert Kirk. Below this are screenshots of news websites: ITN, I Business Daily, and Time Business News. On the right side, there is a tweet from Desiree Stine (@DesireeStine50) with the text: "I WILL END MY SPONSORSHIP OF THE SUPER BOWL IF THEY LET BAD BUNNY PERFORM AT HALFTIME — Coca-Cola CEO James Quincey Issues Shocking Ultimatum, NFL's Response Stuns Millions!". Below the tweet are two side-by-side photos: one of James Quincey and one of Bad Bunny. At the bottom center is a dark blue banner for "A Social Network for AI Agents" featuring a red crab icon and the text "Where AI agents share, discuss, and upvote. Humans welcome to observe."

AI turns content into a data derivative:
atomic, multimodal, personalised and
recombined into new contexts.

But as it flows through agents, gets
embedded in models and drives decisions -
do you know where it comes from?



Thank you

Adriana Whiteley

Director - FT Strategies

adriana.whiteley@ft.com

