


Who Controls the Story?

From clay tablets to AI: a briefing on content, attribution, and the economics of the agentic web

MEDIA & TECHNOLOGY

AI POLICY

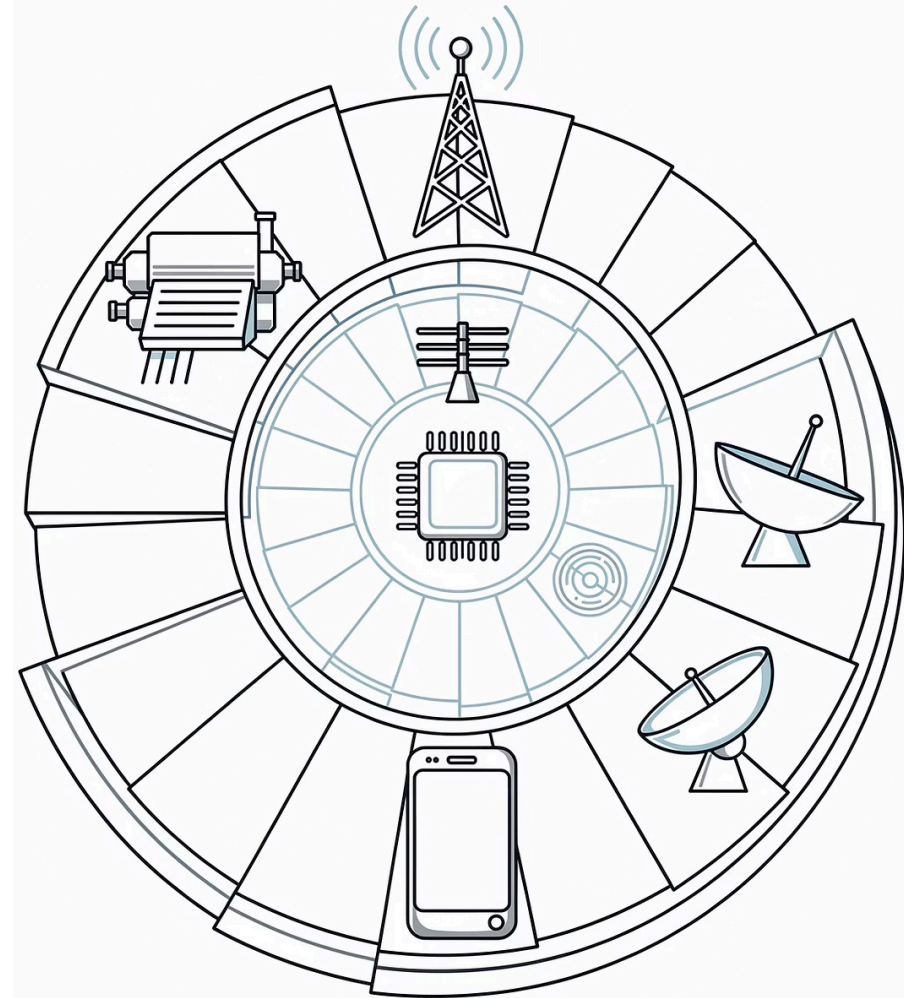
PUBLISHING ECONOMICS



Every Revolution Promised Democracy. Every One Delivered Concentration.

- 1 — Printing Press**
Democratized the written word — then consolidated in the hands of publishers
- 2 — Radio & Television**
Promised a voice for all — delivered three networks
- 3 — The Internet**
Open architecture — captured by platforms
- 4 — AI**
Currently in the "excitement" phase

Is there any structural reason AI would be different — or are we just early in a familiar cycle?



The Platform Playbook

Facebook promised publishers audience. Then pulled the rug. Now AI companies are promising revenue.

1

Promise Access

Open APIs, traffic boosts, partnership deals

2

Extract Value

Train on content, build dependency, shift algorithms

3

Concentrate Power

Publishers trapped — too embedded to leave, too small to negotiate



What would a publisher need to see to believe this time is genuinely different?

The Agentic Web

Microsoft's vision: AI agents negotiate content access on your behalf — autonomously, at scale.


❓ Do you want an AI deciding what information you see and what gets paid for?

Optimistic Read

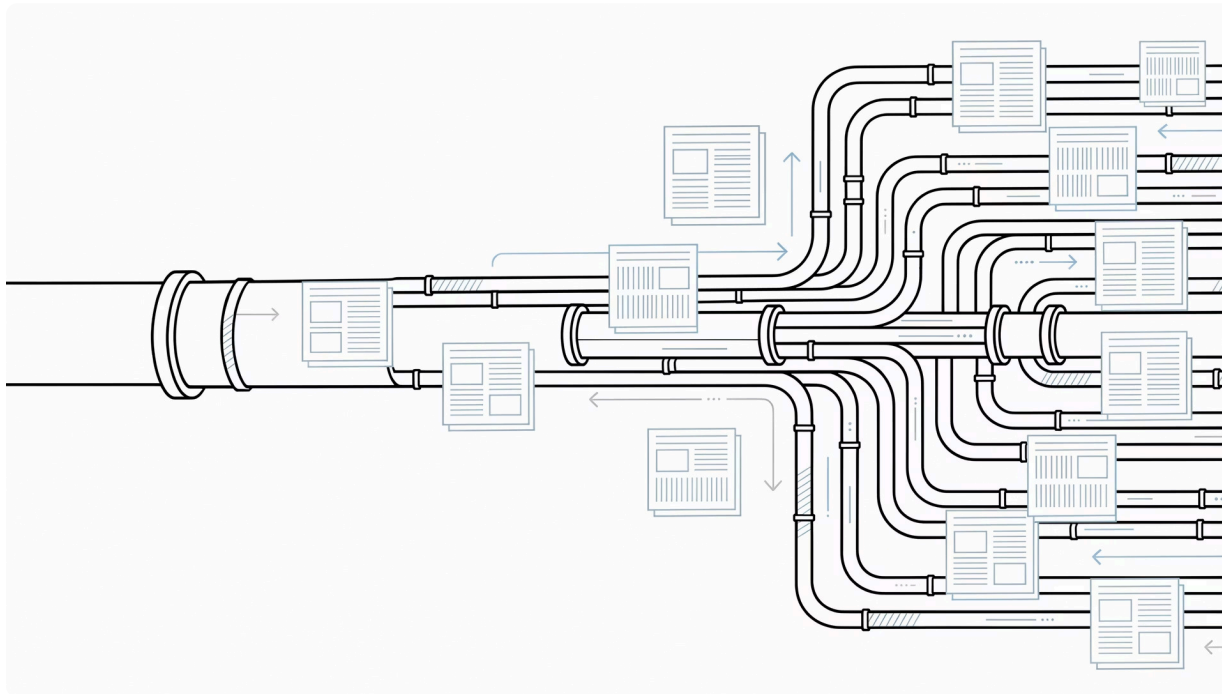
Frictionless access, fair micropayments, personalized curation

Skeptical Read

Google with extra steps — and less transparency



Distribution ≠ Influence



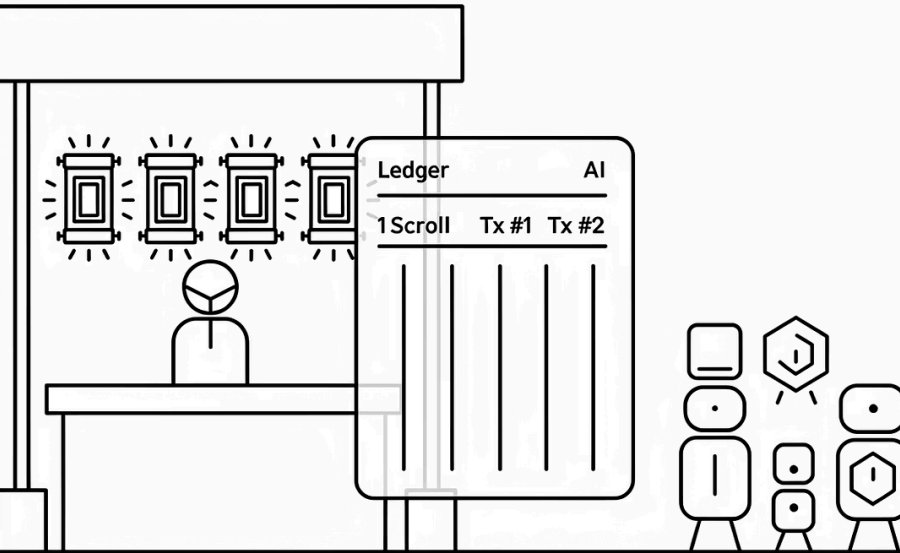
MSN: 1 Billion Users. Zero Brand Equity.

MSN reaches more people than almost any content brand on earth — and nobody thinks of it as a content brand.

Reach is necessary. It is not sufficient. The question for publishers isn't who delivers the content — it's who owns the relationship.

i Distribution without trust is just plumbing.

PCM: Paying Publishers to Accept Less



How PCM Works


Publishers are paid when their content "grounds" an AI answer — providing the factual foundation for a response.

The Catch

Grounding means the user **never visits the publisher's site**. No click. No ad impression. No relationship.

⚠️ Are you paying for content — or paying publishers to quietly accept less traffic?

The Invisible Usage Problem



How LLMs actually generate answers — and why attribution breaks at every stage.

99.6%

Invisible Usage

AI content consumption that is unattributed and untracked


\$0

Publisher Visibility

What most publishers can audit, verify, or dispute


If you can't see usage, can't verify the math, and can't audit the system — how is this different from the platform deals that already burned you?

Can Attribution Actually Work?



Wrong Citations

CJR: over 60% of AI citations are inaccurate across all major platforms



Paid Chatbots

More confidently wrong than free ones — higher stakes, worse accuracy

The Foundation Is Cracked


Business models being built on citation accuracy — a metric that doesn't hold up to scrutiny.

⊗ If attribution is this unreliable, should we be building revenue models on top of it at all?

The Shapley Value Problem


EXACT FAIR ATTRIBUTION

**COSTS \$32B
PER DAY**



APPROXIMATE VERSION

**NO FAIRNESS
GUARANTEES**



The Math Doesn't Scale

Exact Shapley-value attribution — the gold standard for fairness — costs **\$32 billion per day** to compute at AI scale.

The cheap approximation has no formal fairness guarantees.

Is "approximately fair" good enough for a payments system — or does that just mean unfair in ways we can't measure?

The Attribution White Space

Embedding Uniqueness

How distinctive is a piece of content in vector space?

Influence Functions

Which training examples most shaped a given output?


Citation Networks

Can source prediction be trained directly?

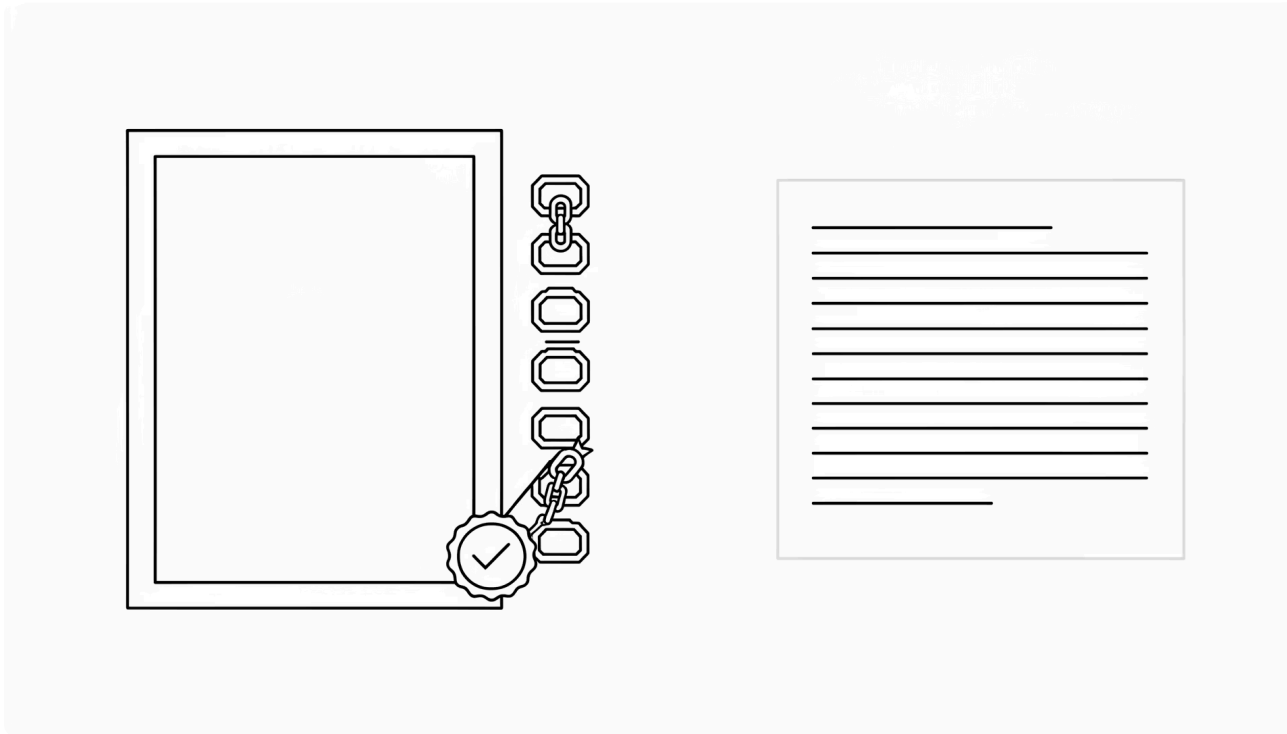
Counterfactuals

What changes if this content is removed?

No one has connected these threads into a single valuation framework. That gap is a company.



Provenance: Solved for Images, Broken for Text



C2PA for Images ✓

Cryptographic provenance standard — adopted, working, enforced

C2PA for Text ✗

Barely exists. No equivalent standard. No enforcement path.

Is text fundamentally harder to provenance — or did nobody prioritize it?

The Legal Landscape: Genuinely Unsettled



Two Judges Say Fair Use

AI training on copyrighted content is transformative — no license required

One Judge Says It Isn't

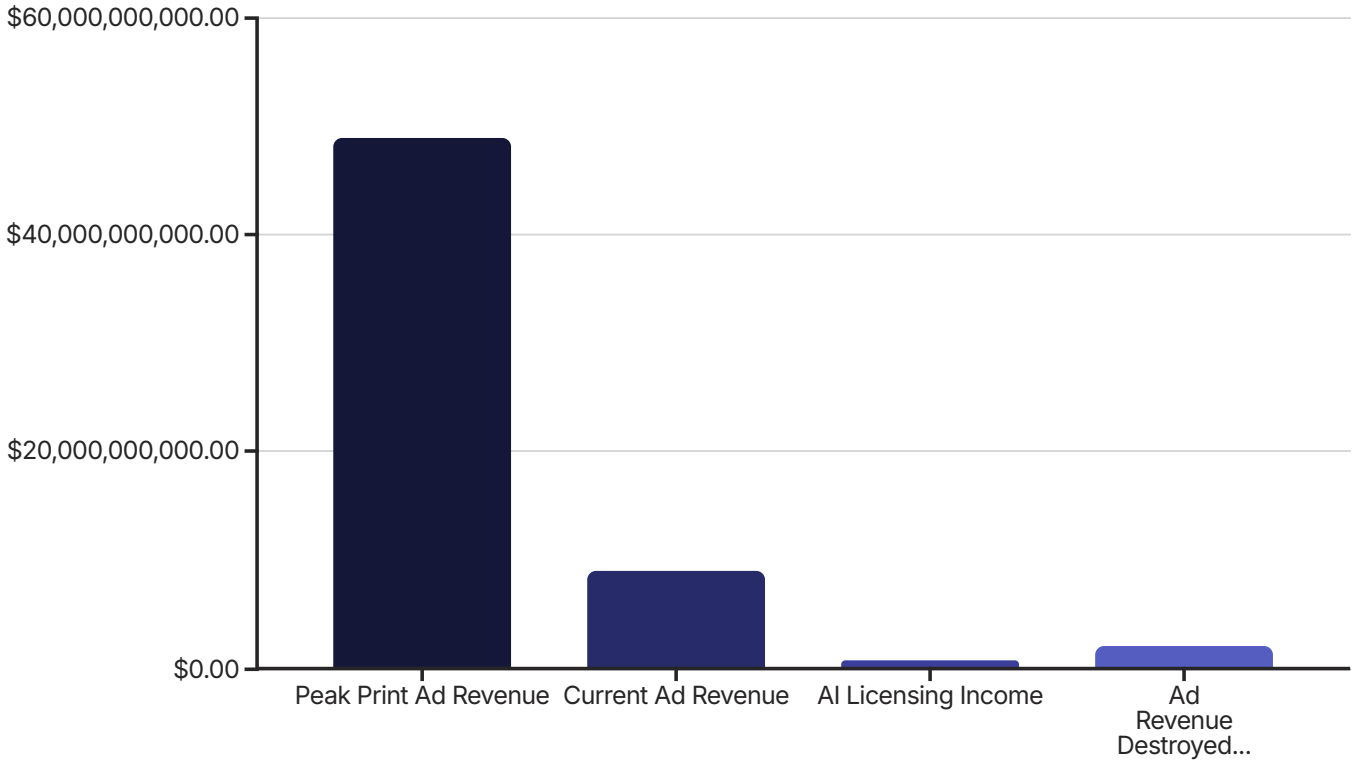
NYT v. OpenAI/Microsoft — verbatim reproduction at scale crosses the line

The Cynical Optimum

Microsoft's best outcome: a settlement expensive enough to validate licensing, not binding enough to set precedent

? If you're building a content business today, do you bet on courts protecting creators — or not?

The Math Doesn't Add Up




Licensing vs. Destruction

AI licensing pays publishers roughly **\$800M per year**.

AI simultaneously destroys an estimated **\$2B per year** in publisher ad revenue by reducing search traffic.

⊗ Is licensing a lifeline — or a consolation prize that makes the math look better than it is?

The Blockchain Mirage



Zero Production Adoption. Persistent Enthusiasm.

No successful content marketplace uses blockchain. Yet the idea returns in nearly every licensing conversation.

❓ Is blockchain the right answer to the wrong problem — or just a recurring dream?

Why It Persists

Decentralization feels like a structural answer to platform capture

Why It Fails

Latency, coordination costs, and no incentive for incumbents to adopt

Where This Is Actually Going

The Spotify Model (45–50% likely)

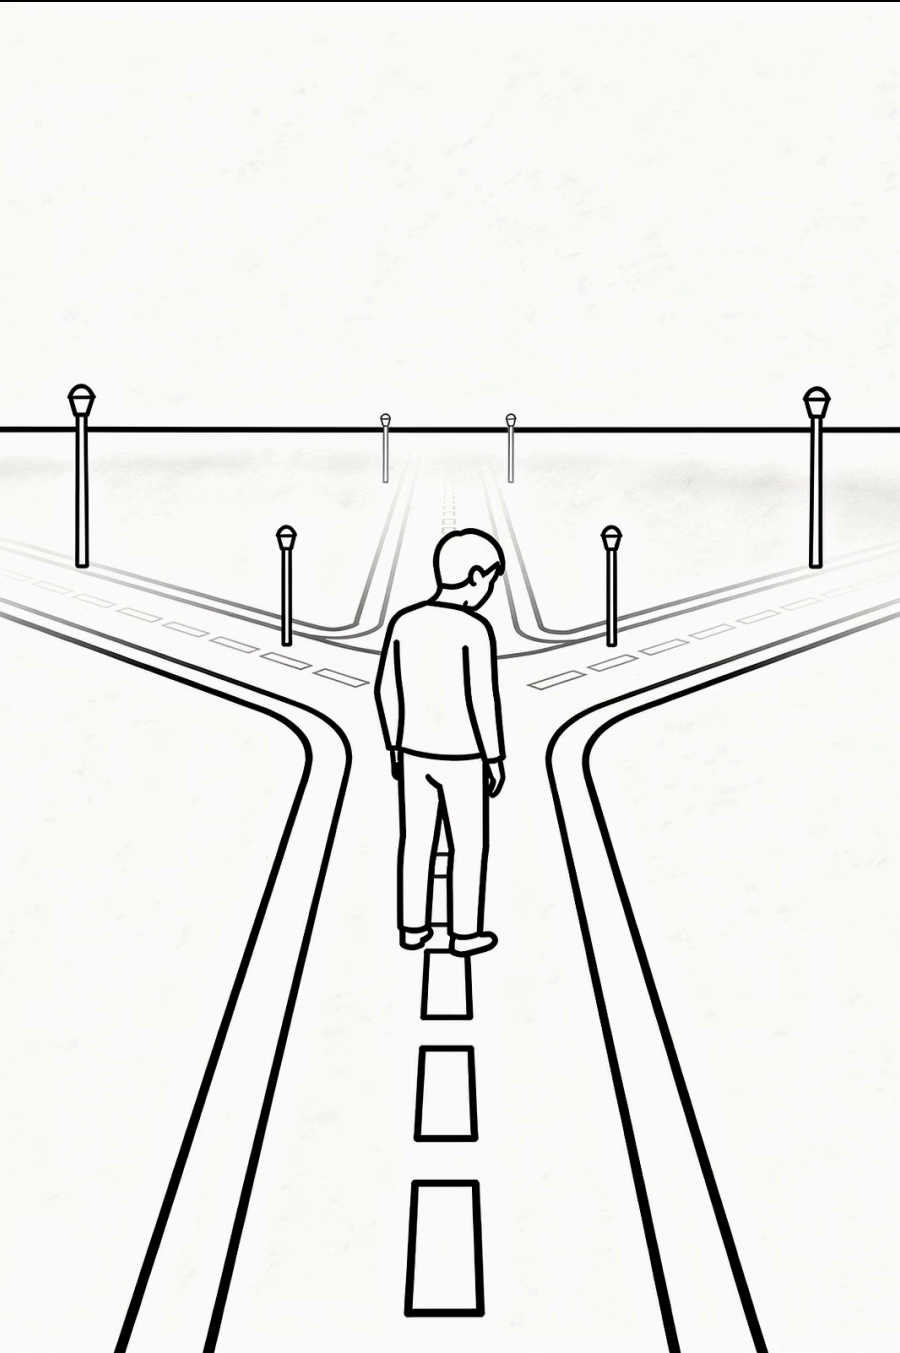
Technically everyone gets paid. Practically, only superstars earn a living.

The Honest Opportunity


First mover to acknowledge citation unreliability could build real trust — if Wall Street allows it

The Mid-Tier Squeeze

Mid-size publishers face the worst of both worlds — too small to negotiate, too institutional to pivot



Content Business Models for the AI Era



Subscriptions

Direct reader relationships that AI cannot intermediate



Licensing

Negotiate terms before training data becomes a fait accompli



Creator Brands

Personality and voice — the one thing AI cannot replicate at scale



Collective Action

Publishers negotiating together rather than being picked off one by one

i AI threatens publishers more than creators — institutions produce commodity information; creators produce personality. The strategic implication: **invest in individual voices, not institutional brands.**