SYNTHETIC MEDIA INVESTIGATOR // TRAINING DIVISION

FIELD INCIDENT CASEBOOK

Volume 1 — Fifteen Training Scenarios in Data Theft, Counterfeit Works, and False Narratives

Anyone can generate anything. This casebook exists so more people can tell what happened anyway. Inside are fifteen fictional incident scenarios — the kinds of situations creators, brands, and institutions increasingly face in an AI-driven world — each paired with the questions a trained investigator would ask and the real-world practice the scenario is modeled on.

All scenarios, names, and organizations are fictional training material. This document teaches investigative thinking; it is not legal advice, and no method here guarantees detection or prevention. Good verification produces evidence and calibrated judgment — not certainty.

VERIFY. DETECT. PROTECT. EMPOWER CREATORS.

syntheticmediainvestigator.com

HOW TO USE THIS CASEBOOK

Read like an investigator, not a spectator.

Each incident is presented the way it would land on an analyst's desk: a briefing with partial information. Before reading the questions, pause and ask what *you* would want to know first. Then compare your instincts against the investigative questions and the good-practice notes.

Three disciplines run through every scenario: **evidence before conclusions** (document what you can verify, in the order you verified it), **calibrated confidence** (an honest “inconclusive” outranks a confident error), and **humans decide** (automated analysis informs a judgment; it never is one).

The fifteen incidents are grouped into the three families S.M.I. tracks most: **Data Theft & Scraping**, **Counterfeit & Infringement**, and **False Narratives**. Severity ratings reflect typical stakes, not difficulty — some of the most damaging incidents are the easiest to miss.

These same incident families power the S.M.I. training simulations — including the Field Office campaign in development — where you make the calls and live with the consequences.

FAMILY ONE — DATA THEFT & SCRAPING

When the work is taken before anyone ever infringes it.

INCIDENT 01 // CB-D1 **SEVERITY I - ROUTINE**

The Vanishing Portfolio

BRIEFING

An illustrator notices her entire portfolio site — 340 images — loading unusually slowly for two nights. Weeks later, a style-mimic image generator begins producing work with her signature linework and recurring characters. She has no idea whether her site was scraped, when, or by what.

INVESTIGATIVE QUESTIONS

- What do the server access logs show for the slow-traffic window — user agents, request patterns, crawl rates?
- Did the site publish machine-readable usage preferences (robots.txt, ai.txt, meta tags), and were they ignored?
- Can any generated output be tied to specific registered works rather than a general 'style'?

WHAT GOOD PRACTICE LOOKS LIKE

- ✓ Preserve raw access logs immediately — they rotate and vanish; a timestamped export is evidence.
- ✓ Register key works with creation records BEFORE incidents; provenance established after the fact is weaker.
- ✓ Document specific matches (composition, characters, artifacts), since style alone is rarely protectable.

REAL-WORLD LESSON

Style is generally not copyrightable, but specific works are — and scraping-policy signals plus preserved logs turn 'I have a feeling' into a documented account that supports takedowns, opt-out claims, and future litigation.

INCIDENT 02 // CB-D2 **SEVERITY II - ELEVATED**

The Helpful Aggregator

BRIEFING

A stock-photo cooperative discovers a 'research dataset' circulating on model-training forums containing 60,000 of its watermarked previews — watermarks crudely inpainted out. The dataset's README claims 'publicly available images, fair use for research.'

INVESTIGATIVE QUESTIONS

- Do the dataset images retain any of the cooperative's embedded metadata, hashes, or partial watermark artifacts?
- Who compiled and who redistributes the dataset — and are those different actors with different liability?
- What do the cooperative's terms of service actually say about automated access and preview downloads?

WHAT GOOD PRACTICE LOOKS LIKE

- ✓ Fingerprint a sample: inpainted watermarks usually leave detectable residue and the originals' perceptual hashes still match.
- ✓ Distinguish the scraper, the packager, and the host — notices go to each with different asks.
- ✓ Send preservation-and-takedown requests to hosts BEFORE public accusation; datasets vanish and reappear renamed.

REAL-WORLD LESSON

'Publicly viewable' is not 'free to take' — removing watermarks can violate the DMCA's provisions against stripping copyright-management information (17 U.S.C. §1202), a separate claim from the copying itself.

INCIDENT 03 // CB-D3 **SEVERITY II - ELEVATED**

The Inside Feed

BRIEFING

A mid-size record label finds unreleased stems from three upcoming albums referenced in an AI music tool's demo outputs. The tracks were never public. The label's own cloud collaboration workspace is the only place all three projects coexisted.

INVESTIGATIVE QUESTIONS

- Which accounts, integrations, or third-party plugins had read access to all three project folders?
- Do access logs show bulk downloads, unusual API tokens, or sync clients from unrecognized devices?
- Can the demo outputs be matched to the stems by waveform or fingerprint rather than by ear alone?

WHAT GOOD PRACTICE LOOKS LIKE

- ✓ Audit OAuth grants and API tokens first — forgotten integrations are the most common quiet exit.
- ✓ Match audio by fingerprint and document the comparison methodology, not just the conclusion.
- ✓ Rotate credentials and preserve the access-log snapshot BEFORE announcing the investigation internally.

REAL-WORLD LESSON

Most data theft is not exotic — it exits through legitimate credentials. Access-control audits and preserved logs are the evidence trail; audio fingerprint comparison turns suspicion into a documented match.

INCIDENT 04 // CB-D4 **SEVERITY I - ROUTINE**

The Republished Archive

BRIEFING

A local newsroom's twenty-year article archive appears, lightly paraphrased, across a network of ad-farm sites — each article reworded just enough to defeat copy-paste comparison, published minutes after the originals, bylined to authors who do not exist.

INVESTIGATIVE QUESTIONS

- Does publication timing correlate tightly enough with the originals to demonstrate automated ingestion?
- Do the paraphrases preserve the originals' factual errors or corrections — the classic copying tell?
- Who profits — what ad networks and registrars connect the farm sites?

WHAT GOOD PRACTICE LOOKS LIKE

- ✓ Seed the archive with harmless canary details (a unique phrasing, a fictitious middle initial) and monitor for their reappearance.
- ✓ Trace shared infrastructure: analytics IDs, ad accounts, and registrant patterns across the farm.
- ✓ Report to ad networks — demonetization often outpaces takedown.

REAL-WORLD LESSON

Copied errors are the oldest and best evidence of copying — mapmakers planted 'trap streets' for exactly this reason. Paraphrase defeats string-matching, but it faithfully reproduces your mistakes.

INCIDENT 05 // CB-D5 **SEVERITY I - ROUTINE**

The Loyal Fan Site

BRIEFING

A game studio finds a 'fan archive' that mirrors its entire asset wiki nightly — concept art, lore documents, soundtrack files — and now outranks the studio's own site in AI-assistant answers about the game, with subtle errors and injected affiliate links.

INVESTIGATIVE QUESTIONS

- Is the mirror scraping despite crawl-preference signals, and at what cadence?
- Which of the mirror's alterations are benign errors and which are monetized manipulations?
- Why do AI systems prefer the mirror — structure, metadata, or citation graph?

WHAT GOOD PRACTICE LOOKS LIKE

- ✓ Document the mirror's drift from the canonical source with dated side-by-side captures.
- ✓ Strengthen the canonical site's machine-readable signals: structured data, canonical tags, clear licensing metadata.
- ✓ Request correction/attribution from AI-search providers using the documented canonical record.

REAL-WORLD LESSON

In AI-mediated search, being the true source is not enough — you must be legibly the true source. Provenance signals and structured metadata (the discipline of GEO) are how canonical content reclaims attribution.

FAMILY TWO — COUNTERFEIT & INFRINGEMENT

When the work itself is copied, faked, or sold out from under its maker.

INCIDENT 06 // CB-C1 **SEVERITY II - ELEVATED**

The Parallel Album

BRIEFING

A vocalist discovers a full album on streaming platforms under a near-identical artist name — her voice, cloned, singing songs she never recorded. Royalties flow to an anonymous distributor account. Fans are congratulating her on the 'surprise drop.'

INVESTIGATIVE QUESTIONS

- Which distributor uploaded the release, and what identity/rights declarations did they file with it?
- Can voice-model similarity be documented alongside the absence of any recording sessions, contracts, or masters?
- What do each platform's impersonation and AI-voice policies provide beyond copyright claims?

WHAT GOOD PRACTICE LOOKS LIKE

- ✓ File platform impersonation claims in parallel with rights claims — different teams, different speeds.
- ✓ Assemble a negative-provenance package: proof the artist has no session files, no registrations, no contracts for these works.
- ✓ Preserve streaming metadata and distributor identifiers before the release is quietly swapped or re-uploaded.

REAL-WORLD LESSON

Voice cloning often sits outside classic copyright (she doesn't own 'her voice' as a work) but squarely inside right-of-publicity and platform impersonation policies — a reminder to pick the strongest available frame, not the most famous one.

INCIDENT 07 // CB-C2 **SEVERITY I - ROUTINE**

The Print Shop Next Door

BRIEFING

A painter's best-selling piece appears on print-on-demand marketplaces across forty product types — upscaled, color-shifted, and cropped to dodge image matching. Sellers rotate storefronts weekly.

INVESTIGATIVE QUESTIONS

- Do the listings' source images retain upscaling artifacts or crop geometry traceable to one leaked master file?
- Which marketplaces host repeat-offender storefronts, and what do their repeat-infringer policies obligate?
- Is there one upstream supplier behind the rotating fronts — shared templates, shipping origins, image batches?

WHAT GOOD PRACTICE LOOKS LIKE

- ✓ Register the work's copyright formally — registration unlocks statutory damages and faster marketplace action in the U.S.
- ✓ File DMCA notices per listing but track storefront lineage to invoke repeat-infringer termination.
- ✓ Watermark and fingerprint master files so future leaks carry their own origin evidence.

REAL-WORLD LESSON

The DMCA notice-and-takedown process is per-copy, but platform repeat-infringer policies are per-actor — investigators who document seller lineage get whole operations removed instead of playing listing whack-a-mole.

INCIDENT 08 // CB-C3 **SEVERITY III - CRITICAL**

The Two Certificates

BRIEFING

A collector purchases a digital artwork with a certificate of authenticity — then a second collector surfaces holding a certificate for the same piece, signed by the same artist, dated three days earlier. The artist says she issued only one. One certificate is forged; nobody knows which.

INVESTIGATIVE QUESTIONS

- Do the certificates' signatures, hashes, or issuance records resolve against any independent registry or timestamp authority?
- What does the transaction trail show — payment records, correspondence, wallet or account histories for each sale?
- Does either certificate reference details (file hash, dimensions, edition numbering) that only the true master would match?

WHAT GOOD PRACTICE LOOKS LIKE

- ✓ Anchor authenticity in verifiable records at creation time — hashes and timestamps registered before any sale.
- ✓ Reconstruct the full chain of custody for both copies; forgeries usually contain one impossible link.
- ✓ Treat 'the friendlier party' as irrelevant — follow the records, not the story.

REAL-WORLD LESSON

A certificate is only as strong as the registry behind it. Provenance that exists solely as a signable document can be forged; provenance anchored in independently verifiable creation records mostly cannot.

INCIDENT 09 // CB-C4 **SEVERITY II - ELEVATED**

The Franchise That Wasn't

BRIEFING

A small studio's indie film trailer is re-cut with AI-generated scenes into a fake 'sequel announcement,' complete with counterfeit studio branding, and released to hype a token pre-sale the studio has nothing to do with.

INVESTIGATIVE QUESTIONS

- Which shots are lifted from the original trailer versus generated, and can each be individually evidenced?
- Who operates the pre-sale — payment rails, domains, and social accounts — and where are they legally reachable?
- Is this primarily an infringement matter, a trademark matter, or a fraud matter — and who enforces each?

WHAT GOOD PRACTICE LOOKS LIKE

- ✓ Split the evidence package: copied footage (copyright), counterfeit branding (trademark), investor deception (fraud) — three doors, three enforcers.
- ✓ Publish a signed, dated official statement establishing the canonical record before rumor hardens.
- ✓ Report payment rails first; freezing the money outpaces deleting the video.

REAL-WORLD LESSON

AI-era counterfeits are rarely one crime — the same artifact can infringe copyright, counterfeit trademarks, and further fraud simultaneously. Investigators who separate the claims move three times as fast.

INCIDENT 10 // CB-C5 **SEVERITY II - ELEVATED**

The Ghost Contractor

BRIEFING

A design agency discovers a former contractor selling 'his' portfolio — which is largely client work he touched at the agency — as licensed templates, including projects governed by work-for-hire agreements he signed.

INVESTIGATIVE QUESTIONS

- What do the signed agreements actually say — work-for-hire, assignment, or mere license — per project?
- Which files in the template packs match agency masters by hash, layer structure, or embedded metadata?
- Has any downstream buyer relied on his license in good faith — and what does that change?

WHAT GOOD PRACTICE LOOKS LIKE

- ✓ Read the contracts before the accusations — ownership follows the paper, not the effort or the memory.
- ✓ Match files forensically (embedded metadata, layer names, revision artifacts) rather than visually.
- ✓ Notify marketplaces with the ownership documentation; address good-faith buyers separately from the seller.

REAL-WORLD LESSON

Work-for-hire means the contribution was owned by the client from the moment of creation — sympathy for the maker doesn't move title, and a good-faith purchase downstream cannot launder a license the seller never had to give.

FAMILY THREE — FALSE NARRATIVES

When the fabrication is the story itself.

INCIDENT 11 // CB-N1 **SEVERITY III - CRITICAL**

The Leaked Meeting

BRIEFING

An audio clip spreads of a hospital director 'privately admitting' to falsified safety reports. The voice is convincing; the acoustics are plausible; staff morale is collapsing. The director insists the meeting never happened.

INVESTIGATIVE QUESTIONS

- Does the audio show splice points, room-tone discontinuities, or spectral artifacts consistent with synthesis or editing?
- Can the claimed meeting be checked against calendars, badge logs, and the room's actual acoustics?
- Where did the clip first appear, and does its propagation pattern suggest coordination?

WHAT GOOD PRACTICE LOOKS LIKE

- ✓ Analyze BOTH tracks: the artifact (signal analysis) and the event (did the meeting exist at all?).
- ✓ Publish findings with confidence levels and methodology — certainty theater destroys credibility when revised.
- ✓ Prepare the response for the likeliest truth AND the uncomfortable one; investigators serve the record, not the client's hopes.

REAL-WORLD LESSON

Authenticating an event is often stronger than authenticating a file — synthesis artifacts can be argued, but a meeting that provably never occurred ends the question. Verify the world, not just the waveform.

INCIDENT 12 // CB-N2 **SEVERITY II - ELEVATED**

The Real Photo, Wrong War

BRIEFING

A photojournalist's authentic image from a 2019 assignment goes viral captioned as breaking news from a current conflict, amassing millions of views and fueling outrage. The photo is genuine. The story attached to it is not.

INVESTIGATIVE QUESTIONS

- When and where does the image FIRST appear online, and does its metadata or agency record confirm the original context?
- Who re-captioned it, and did the re-caption originate organically or from coordinated seeding?
- What correction path reaches the audiences that saw the false version?

WHAT GOOD PRACTICE LOOKS LIKE

- ✓ Run reverse-provenance first — earliest appearance beats pixel forensics for context manipulation.
- ✓ Attach the correction to the SAME channels and formats that carried the falsehood.
- ✓ Credit and involve the original photographer; the creator's own record is the strongest testimony.

REAL-WORLD LESSON

The most common 'deepfake' involves no fake at all — context manipulation of authentic media outpaces synthetic media in the wild, which is why provenance and earliest-appearance research are the first instruments an investigator reaches for.

INCIDENT 13 // CB-N3 **SEVERITY II - ELEVATED**

The Grassroots Avalanche

BRIEFING

A regional food brand faces a sudden wave of hundreds of one-star reviews and testimonial videos describing illnesses — posted across platforms within 72 hours, many by accounts created the same week, several featuring suspiciously similar kitchens.

INVESTIGATIVE QUESTIONS

- What fraction of accounts show coordinated fingerprints — creation dates, posting cadence, shared phrasing or assets?
- Do any claimed illnesses correspond to actual reports with health authorities or the brand's complaint records?
- Who benefits — and does the campaign's timing align with a competitor launch, a short position, or an extortion attempt?

WHAT GOOD PRACTICE LOOKS LIKE

- ✓ Quantify coordination statistically; a few fakes prove little, a distribution proves a campaign.
- ✓ Separate possibly-genuine complaints for real follow-up — assuming 100% fabrication is its own trap.
- ✓ Report coordinated inauthentic behavior to platforms with the pattern evidence, not one-off flags.

REAL-WORLD LESSON

Coordination is proven in the aggregate: no single fake review is evidence, but timing distributions, account cohorts, and reused assets are. Investigators think in patterns while responders handle each claim honestly.

INCIDENT 14 // CB-N4 **SEVERITY III - CRITICAL**

The Candidate Who Said Nothing

BRIEFING

Days before a local election, a video circulates of a school-board candidate 'confessing' to misusing funds. It is quickly debunked as AI-generated — and then the real damage begins: the candidate's genuine past statements are now dismissed by supporters as 'probably fake too.'

INVESTIGATIVE QUESTIONS

- Beyond debunking the fake, which authentic records are now contested, and can they be independently re-anchored?
- Is the fake's origin traceable — and does amplification suggest intent to exploit the debunk itself?
- What verification infrastructure (signed statements, provenance-attached media) can the campaign adopt going forward?

WHAT GOOD PRACTICE LOOKS LIKE

- ✓ Treat the debunk as the midpoint, not the end — document and re-anchor the authentic record.
- ✓ Publish authentic materials with verifiable provenance so future disputes resolve against records, not rhetoric.
- ✓ Anticipate the liar's dividend in every synthetic-media response plan.

REAL-WORLD LESSON

The 'liar's dividend' is the second blast radius of every deepfake: once fakes are known to exist, real evidence becomes deniable. The counter is provenance infrastructure — making authentic records verifiable before they are questioned.

INCIDENT 15 // CB-N5 **SEVERITY I - ROUTINE**

The Expert Who Never Existed

BRIEFING

A 'independent audit report' praising a supplement brand circulates with a credentialed author — headshot, LinkedIn history, cited publications. None of it resolves: the photo is generated, the publications don't exist, and the 'institute' was registered three weeks ago.

INVESTIGATIVE QUESTIONS

- Do the author's credentials resolve against any independent registry — institutions, journals, licensing boards?
- Is the headshot consistent with generation artifacts, and does it reverse-search to any prior existence?
- Who commissioned and distributes the report, and is it being cited by AI assistants as a source?

WHAT GOOD PRACTICE LOOKS LIKE

- ✓ Verify people the way you verify files: independent resolution against records that predate the claim.
- ✓ Check whether AI search systems have ingested the fake source and request correction with documentation.
- ✓ Publish the debunk in machine-readable form so future automated citation checks inherit it.

REAL-WORLD LESSON

Synthetic personas fail against time: a real expert leaves decades of independently-anchored records, while a fabricated one begins existing all at once. 'When did this identity start?' is the single most efficient verification question.

AFTER THE CASEBOOK

Training is rehearsal. Judgment is the skill.

Every scenario here compresses a real dynamic — scraping, counterfeiting, coordinated deception — into something you can study. The discipline they all share: preserve evidence early, distinguish what you can verify from what you suspect, state your confidence honestly, and let humans make the call that machines only inform.

Practice these calls interactively in the S.M.I. Academy — free training simulations, the Intelligence Archive, and member investigation modules where the scenarios push back.

SYNTHETIC MEDIA INVESTIGATOR

Trust, for the era of co-creation.

syntheticmediainvestigator.com