

Things that *want to be shared* but *need to be checked*

Twelve manipulation patterns, caught in the wild and pinned to the page — real, documented cases, with the tell that gives each one away.

BEFORE YOU BEGIN

How This Guide Works

Every specimen here is a real, documented case — verified by fact-checkers, journalists, or the companies involved, not invented for the lesson. Sources are cited at the bottom of each plate so you can check the checking.

Each pattern is logged the way a naturalist logs a species: what it looks like in the wild, and the field mark that gives it away. The goal isn't to memorize twelve hoaxes — it's to recognize the *mechanisms*, so you catch the next one, which will look nothing like these.

OBSERVED SPECIMEN

What the real post, message, or claim actually said or did.

FIELD MARK

The one tell that exposes the pattern, every time it resurfaces.

SPECIMENS IN THIS EDITION

Framing Distortions

5

Plates I–V

Dark Patterns

4

Plates VI–IX

Rhetorical Sleights

3

Plates X–XII

False Dichotomy

Dichotomia Falsa

First observed 2001 **Status** Still quoted in political rhetoric today

OBSERVED SPECIMEN

ADDRESS TO CONGRESS, NINE DAYS AFTER SEPT. 11

Speaking to Congress and a global television audience on September 20, 2001, the president framed the response to terrorism as strictly binary: "Either you are with us, or you are with the terrorists." Aimed at foreign governments deciding how to respond, the line collapsed a wide range of possible positions — support with conditions, disagreement over method, neutrality — into exactly two.

THE HOOK

A real, high-stakes decision gets reframed so that only two positions appear to exist, with no room to occupy the middle ground.

THE MECHANISM

Anyone who raises a condition or a third option can be recast as sitting on the wrong side of the binary — nuance gets punished by the frame itself.

Field mark: when a claim offers exactly two options and treats hesitation as picking the "other" side, check whether the missing middle is where the real disagreement actually lives — most real positions have more than two settings.

CASE VERIFIED — White House transcript, Sept. 20, 2001; widely reported by CNN, AP.

Emotional / Outrage Framing

Affectus Provocatio

First observed Oct. 2016 **Status** Peaked Dec. 2016

OBSERVED SPECIMEN

"PIZZAGATE" CONSPIRACY, FALL 2016

In the weeks before the 2016 election, a false conspiracy theory claimed a Washington, D.C. pizza restaurant was the hub of a child-trafficking ring run by public figures. The claim grew out of misread emails and spread through outrage-driven shares across Reddit, 4chan, Facebook, and Twitter — with no verified evidence behind it at any point.

THE HOOK

The subject — child safety — is so emotionally loaded that pausing to check the claim can feel like a moral failure rather than basic diligence.

THE MECHANISM

Outrage produces faster shares than any correction can keep up with, giving the false version a multi-week head start before verification catches up.

Field mark: the intensity of your reaction is not evidence the claim is true. Sourcing and emotion are two separate things — the stronger you feel, the more worth double-checking it becomes.

CASE VERIFIED — An armed man entered the pizzeria on Dec. 4, 2016 acting on the false claim; reported by The New York Times and multiple national outlets.

False Urgency

Falsus Urgentia

First observed 2016 **Status** Resurfaced again in 2025

OBSERVED SPECIMEN

CHAIN MESSAGE, RECURRING EVERY 1–2 YEARS

*A message circulating in waves since 2016 warns that a fake premium update called "WhatsApp Gold" is secretly a virus, instructing recipients to forward the warning to every contact **immediately**, before anyone opens it. The hoax resurfaces every year or two under a fresh news hook, most recently in 2025 — nearly a decade after it first appeared.*

THE HOOK

A real category of threat — malware does exist — lends credibility to a specific claim that is entirely fabricated, since no such WhatsApp product has ever existed.

THE MECHANISM

"Forward to everyone right now" turns the reader into the distribution mechanism itself, bypassing any algorithm and converting trust in the sender into the delivery method.

Field mark: WhatsApp has confirmed directly it has never offered a paid "Gold" tier. The giveaway isn't the danger described — it's the instruction to forward before checking, every time it resurfaces.

CASE VERIFIED — Snopes, "WhatsApp Gold scam alert," Oct. 2025; WhatsApp confirmed to Snopes it offers no such product.

Cherry-Picked Data

Selectio Fructus

First observed Long-running technique **Status** Jan. 2023 viral case

OBSERVED SPECIMEN

VIRAL TWEET AND INSTAGRAM REPOST, JAN. 2023

*In January 2023, a viral tweet and Instagram repost used NOAA's own public temperature data to claim the planet had cooled for **eight consecutive years**, and therefore that climate change was a hoax. The chart was accurate as far as it went — it simply started and ended at years chosen to produce a downward slope.*

THE HOOK

The underlying numbers are real and traceable to an official source, which makes the selective framing far more convincing than an outright fabrication would be.

THE MECHANISM

Eight years is a fraction of the more than 140 years of available temperature records — short windows after a strong El Niño routinely show cooling inside a long-term warming trend.

Field mark: ask what date range was chosen and why that range instead of a longer or shorter one — the edges of a chart, not the data inside it, are usually where cherry-picking hides.

CASE VERIFIED — PolitiFact, Jan. 24, 2023; NOAA climate scientist Howard Diamond called the cited window cherry-picked.

Strawmanning

Homo Stramineus

First observed Classical rhetoric **Status** Constant across polarized debate

OBSERVED PATTERN

RECURS IN NEARLY EVERY POLARIZED PUBLIC DEBATE

Across debates on vaccination policy, gun regulation, and immigration alike, a common move restates an opponent's actual position in its most extreme, least defensible form — something closer to "so you're saying nobody should ever have any say at all" — then spends the rest of the argument dismantling that exaggerated version instead of what was actually proposed.

THE HOOK

Defeating the extreme version feels, to an audience not tracking closely, exactly like defeating the real position — the swap goes unnoticed because both use similar language.

THE MECHANISM

The real argument, with its actual conditions and caveats, is never engaged at all — only a weaker stand-in built specifically to be easy to knock down gets addressed.

Field mark: if a rebuttal argues against something more extreme, more absolute, or more sweeping than what was actually said, it is very likely a strawman rather than a real response.

PATTERN REFERENCE — core informal fallacy in argumentation literature; FLICC framework for identifying science-denial-(Cook,-2020)-

DARK PATTERN

SPECIMEN · NO. 06

PLATE VI

Engagement Bait

Esca Interactionis

First observed Named 2017 **Status** Ongoing, platform-policed

OBSERVED SPECIMEN

NAMED AND FORMALLY PENALIZED BY META, 2017

Pages on Facebook and Instagram routinely posted prompts like "Like if you agree, share if you don't" or "tag five friends and you could win" — instructions with no informational content of their own, designed purely to trigger reactions the platform's algorithm reads as engagement. In April 2017, Meta formally named this category "engagement bait" and began demoting it in the News Feed.

THE HOOK

The instruction itself, not any underlying claim or information, is what drives the interaction — the reader responds to a command rather than evaluating anything.

THE MECHANISM

High engagement signals to the platform's ranking algorithm that content is valuable and worth showing to more people, regardless of whether it's true, false, or entirely empty.

Field mark: a post that tells you exactly how to react — like, share, tag, comment a specific word — before you've had any chance to evaluate it is optimizing for spread, not accuracy.

CASE VERIFIED — Meta Newsroom, "News Feed FYI: Cleaning Up Engagement Bait," April 2017.

Fake Virality / Bot-Inflated Proof

Fama Ficta

First observed Investigated 2018 **Status** Follower markets still operate

OBSERVED SPECIMEN

DEVUMI FOLLOWER-MARKETPLACE INVESTIGATION, JAN. 2018

*A 2018 New York Times investigation, "The Follower Factory," found that a company called **Devumi** had sold tens of millions of fake or automated Twitter followers to celebrities, athletes, pundits, and executives — manufacturing the appearance of large, real, engaged audiences that in many cases barely existed at all.*

THE HOOK

A large follower or share count functions as social proof — the implicit reasoning is that this many people can't all be wrong, or all be uninterested.

THE MECHANISM

The count itself turned out to be directly purchasable in bulk, meaning it measured nothing about real agreement, real interest, or real reach.

Field mark: a huge following paired with unusually low genuine engagement — few real replies or likes relative to the follower count — is the classic signature of a purchased audience.

CASE VERIFIED — The New York Times, "The Follower Factory," Jan. 27, 2018 (Confessore, Dance, Harris, Hansen).

Astroturfed Comments

Gramen Artificiale

First observed Term coined 1985 **Status** 2017 FCC case

OBSERVED SPECIMEN

FCC NET NEUTRALITY COMMENT PERIOD, 2017

*During the FCC's 2017 public comment period on repealing net neutrality rules, the agency's docket received **millions of comments** using fake names or the stolen identities of real people who had never submitted anything. A subsequent investigation by the New York Attorney General's office traced large batches back to coordinated campaigns rather than genuine individual submissions.*

THE HOOK

A flood of comments from what appear to be ordinary citizens creates the impression of broad, organic public sentiment on a policy question.

THE MECHANISM

Nearly identical phrasing repeated across thousands of nominally different commenters is a strong signal of coordinated, paid, or automated origin rather than independent opinion.

Field mark: when many "different" voices use suspiciously similar wording, that repetition — not the sheer volume — is the real signature of manufactured consensus.

CASE VERIFIED — New York Attorney General investigation into the FCC's 2017 public comment docket.

Screenshot / Image Laundering

Imago Migrans

First observed 2011 **Status** Reappeared as AI video in 2025

OBSERVED SPECIMEN

"HURRICANE SHARK," RESURFACING 2011–2025

*A digitally composited photo showing a shark swimming down a flooded highway first circulated after Hurricane Irene struck Puerto Rico in 2011. Fact-checkers debunked it almost immediately, but the identical image has been **relabeled and reshared** as fresh footage during nearly every major hurricane since — Sandy, Matthew, Harvey, Ian, and Melissa — including AI-generated video versions that emerged in 2025.*

THE HOOK

The image is dramatic and visually plausible enough that the caption attached to it gets trusted without anyone checking where the picture actually came from.

THE MECHANISM

Once the original source and date are stripped away, the same visual can be relabeled and attached to any new disaster indefinitely, with no changes to the image itself.

Field mark: a reverse image search takes seconds and instantly reveals whether "breaking" footage has, in fact, broken before — in this case, more than a decade earlier.

CASE VERIFIED — Snopes, 2011–ongoing; Wikipedia "Hurricane Shark"; CBS News, Oct. 2024.

RHETORICAL SLEIGHT
SPECIMEN · NO. 10

PLATE X

False Authority

Auctoritas Ficta

First observed July 2020 **Status** Removed by platforms within 24 hrs

OBSERVED SPECIMEN

VIRAL PRESS CONFERENCE, JULY 27, 2020

*A group calling itself "**America's Frontline Doctors**" held a press conference making false claims about COVID-19 treatments. Clips were viewed tens of millions of times across Facebook, YouTube, and Twitter within 24 hours, before all three platforms removed the video for violating COVID-19 misinformation policies.*

THE HOOK

White coats, medical titles, and a press-conference format all trigger automatic trust that stands in for actual verification of the specific claims being made.

THE MECHANISM

A genuine credential from a real medical field is used to vouch for a claim that the broader body of evidence in that same field does not support.

Field mark: a real credential validates a speaker's general expertise, not every claim they make — check whether the specific claim matches actual consensus, not just the title in front of the name.

CASE VERIFIED — Reuters, NPR, and BBC, July–Aug. 2020; removed by Facebook, YouTube, and Twitter for policy violations.

Anecdote-as-Data

Casus Pro Data

First observed 1998 Status Retracted 2010

OBSERVED SPECIMEN

THE LANCET, 1998 (RETRACTED 2010)

*In 1998, The Lancet published a study based on just **twelve children's case histories** claiming a link between the MMR vaccine and autism. The paper triggered a lasting drop in vaccination rates in several countries. The Lancet fully retracted it in 2010, the author lost his UK medical license, and later studies covering millions of children — including a Danish cohort of more than 650,000 — found no such link.*

THE HOOK

A handful of vivid, specific, personal stories feel more persuasive in the moment than a statistic drawn from a sample of millions ever does.

THE MECHANISM

A sample of twelve is generalized into a sweeping public-health claim that a sample that size could never actually support, statistically or scientifically.

Field mark: ask how many cases a claim is actually based on before accepting it — a compelling personal story is not the same thing as a representative sample.

CASE VERIFIED – Lancet retraction, 2010; author lost UK medical license; CDC and Danish cohort studies (650,000+ children) found no link.

Loaded Language & Weasel Words

Verba Onusta

First observed Longstanding pattern Status Ongoing

OBSERVED PATTERN

RECURS ACROSS NEARLY ALL PARTISAN COVERAGE

*Headlines and captions routinely lean on phrases like "**some say**," "critics claim," or "experts believe" to attribute a claim to an authority that is never actually named or counted, while adjectives like "shocking," "slammed," or "blasted" do persuasive work before a single fact has been presented to the reader.*

THE HOOK

The sentence sounds sourced and the tone sounds justified, even though no specific person or group is ever identified as standing behind either one.

THE MECHANISM

"Some" could mean three people in a comment section or three million voters — the vagueness isn't an oversight, it's what lets the claim scale to whatever the reader assumes.

Field mark: if you finish the piece and still can't name who "some" or "critics" refers to, the sourcing is doing far less work than the sentence implies.

PATTERN REFERENCE – standard media-literacy category; AP Stylebook guidance on attribution and unnamed sources.

The 10-Second Gut Check

adapted from the SIFT method, Mike Caulfield, Univ. of Washington Center for an Informed Public

1 Stop

Before you share, react, or forward — pause. Urgency and outrage are both designed to skip this step.

2 Check who's behind it

Is the account, outlet, or "doctor" who they claim to be? A name isn't the same as a verified credential.

3 Find better coverage

Search the claim itself. If it's real and important, other outlets will have it too.

4 Trace it to the original

Follow the claim or image back to where it came from — screenshots lose context on the way.

5 Check the range

For any statistic, ask what time period it covers, and whether a different range tells a different story.

INCOGNATI

The Atlas Goes Deeper

This guide covers twelve specimens. The Incognati Atlas catalogs thousands more — biases, fallacies, distortions, and manipulation patterns, each one documented the same way: real case, real source, real tell. Caught something in the wild that isn't in here yet? That's the point of a field guide — it's never finished.