


Principal Brief // two-minute read

NARRATIVE WARFARE

CRISIS RESPONSE

MISIDENTIFICATION



Your Statement Will Arrive Too Late

In a fast-moving public event the account people believe forms in the first minutes, before anyone senior is briefed. A statement cleared by counsel arrives after that, and is read as one more voice inside the story rather than the answer to it.

WHAT HAPPENED

Documented cases put the window in which an early frame takes hold at minutes to hours. Our analysis plans against seventeen minutes deliberately, because doctrine should not assume extra time. Three cases carry it: Brian Thompson in December 2024, a martyr frame competing with the investigation itself; Charlie Kirk in September 2025, running with no foreign hand involved; Southport in July 2024, a lawful silence about the attacker's identity filled by a false one.

WHY IT REACHES YOU

Two ways. The correction does not catch up: after a false claim that an Iranian earthquake was a covert nuclear test, the refutation arrived four months later in a journal. And the crowd runs its own identification, on proximity, resemblance, and how well someone fits the story. It can name a relative, a member of staff, or you, and the harm lands before the correction.

THREE THINGS TO TAKE FROM THIS

1

Two clocks, not one

Verification stays slow. Something else runs against distribution: prepared language, prepared channels, prepared authority.

2

Useful beats true

Under uncertainty people are less sensitive to truth and highly sensitive to tone. The winning frame is the one already waiting.

3

Plan for the wrong name

Proximity, resemblance, and narrative fit are not identity, and the crowd will treat them as identity anyway.

DO THIS

- ☐ **YOU, THIS WEEK**
Name the person who can publish a holding statement in fifteen minutes without asking anyone. If nobody can, that is the finding.
- ☐ **YOU, WITH YOUR FAMILY OFFICE**
Pre-bunk before predictable flashpoints and treat footprint reduction as standing work. Neither can be done inside the window.
- ☐ **YOU, WITH COUNSEL**
Pre-clear the holding language, and write the protocol for the day the crowd names the wrong person.

NOT THE STORY YOU MAY HAVE HEARD

This is not an argument for speaking faster or skipping verification, and seventeen minutes is not a measurement. It is our planning figure, set at the demanding end on purpose. The claim is narrower: whatever the exact number, a first statement that needs legal clearance will miss the window, and the fix is prepared authority rather than faster lawyers.


HOW SURE WE ARE

High on the documented cases and the published research quoted. The threshold itself is a planning parameter, not a finding. Moderate on any organization's readiness until it has been rehearsed.

CB-057

The 17-Minute Problem

"One clock is trying to establish what happened. The other has already decided what it means."



A high-salience public event starts two races. The official one seeks an account it can support: secure the scene, reconcile the footage, establish identities, separate observation from inference. The public one seeks an account it can use: connect the event to what people already fear, resent, or remember. The second race finishes first, and it does not wait for the first to begin. Nothing about that is a failure of verification, which is supposed to be slow. It is a failure to run anything at all against the clock that decides meaning, and the remedy is architectural: language, channels, and authority prepared before the event, plus a plan for the person the crowd names who did not do it.

BLUF. Documented cases support a window of minutes to hours in which an early frame gains reach before institutions can establish basic facts. ObscureIQ's narrative-warfare analysis compresses that window to seventeen minutes as a deliberately demanding planning threshold, not as a measured constant and not as a law: doctrine should not be built on the assumption of extra time. The adoption mechanism is documented in the peer-reviewed literature. The COP model, developed by Aaron French, Amrita George, Joshua Madden and Veda Storey and published in Information Systems Frontiers in 2025, holds that content competes on veracity, emotional appeal, and relevance, filtered through the audience's Overton window of acceptable explanations and resolved as perceived utility. Their finding from social-media data in the 2020 coronavirus period is the operative one: readers are markedly less sensitive to truthfulness under uncertainty, and highly sensitive to tone and emotional appeal. The most competitive early frame is therefore not the most credible one; it is the most useful one, and it usually already exists, waiting for an event to attach to.

The governing frame is two clocks and two failure modes. The institutional clock produces a statement that arrives after the frame has set, and the public clock produces a second identification system that can converge on an innocent person before any official process names anyone. A principal is exposed to both. Run pre-approved holding language through pre-delegated authority against the distribution clock while verification proceeds untouched, and separately prepare for wrong-person targeting, because the same open-source tools that surface real evidence also manufacture confident false identifications, and the harm from those arrives before the correction does.

Narrative analysis describes what a claim is doing. Attribution requires separate evidence.

BRIAN THOMPSON · DEC 2024

The killing of the UnitedHealthcare chief executive is the clearest recent proof of concept. Shell casings inscribed with insurance terminology handed the event a ready-made explanation, cross-spectrum sympathy for the shooter emerged within hours, and a martyr frame competed with the official investigation rather than waiting for it. Nothing in the sequence required coordination. The frame was available, it was easy to repeat, and it attached to a standing grievance.

CHARLIE KIRK · SEP 2025

The martyr pivot again, with AI tools accelerating misidentification of the attacker while both partisan blocs attributed the killing to the other inside the same window. The case matters because no adversary directed it. Domestic material alone compressed the timeline, so a family office cannot treat this as a threat that arrives only with a state actor attached.

SOUTHPORT · JUL 2024

Three children were killed at a dance class in the United Kingdom, reporting restrictions kept the attacker's identity out of official statements, and a false claim about that identity filled the gap. A mosque was attacked and disorder spread to other cities. The silence was lawful and it was still a vacuum, which is the point: legal constraint does not exempt an institution from the window, it makes pre-approved holding language the only instrument left.

SEM NAN · OCT 2024 TO FEB 2025

A documented case of a false technical claim outrunning expert analysis. After a magnitude 4.5 earthquake near Semnan, Iran on 5 October 2024, the first social posts appeared about 17 minutes later, the first suggestion of a covert nuclear test about 27 minutes after that, and the false reading reached news outlets worldwide within hours. The peer-reviewed refutation, by Fernando and colleagues in Seismica, published on 4 February 2025, four months on. Persian-language coverage was accurate from the start while English-language outlets elsewhere carried the false version hardest, which is a finding about distribution rather than about truth. The 17 minutes here is time to first post and is not the source of the planning threshold; the four months is the number a communications function should sit with.

THE ACTUARIAL CHANNEL · 2025 TO 2026

Interpretation reaches money before assessment reaches conclusions. In the 2026 Strait of Hormuz crisis, war-risk premiums moved from 0.10 to 0.125 percent of vessel value before the conflict to about 5 percent by 19 March, per Lloyd's chief executive Patrick Tiernan, easing to 2 or 3 percent by 23 March, a rise of roughly 1,000 to 2,400 percent, with carriers rerouting and traffic stalling. The instructive detail is that pricing peaked and then eased while attacks kept accumulating, 32 vessels by 23 March, so the premium tracked perceived danger rather than claims. Separately, the #TheBanksAreOutOfMoney campaign ran from 10 March to early April 2025 and peaked after the 2 April tariff announcement, circulating unverified screenshots of banking-app error messages attributed to known institutions; Blackbird.AI's research team found 17.2 percent of participating accounts showing bot-like behavior, which means the large majority were people.

SO WHAT The frame sets before anyone senior is briefed, it does not need a foreign hand to set it, and it reaches insurers and counterparties before it reaches a conclusion.

THE MECHANISM // WHY THE EARLY FRAME WINS

WHAT THE THRESHOLD IS, AND IS NOT

Seventeen minutes is a planning assumption. The public record supports a broader minutes-to-hours window in which early frames gain reach ahead of established fact, and the analysis compresses that to seventeen to force a harder preparedness test. Event salience, audience, platform, and institutional readiness all move the real number, in both directions. Treat it as the earliest credible threshold to plan against rather than as a measurement, and note what follows: an organization whose first statement requires legal clearance misses a forty-minute window too, so the objection that the figure is aggressive does not rescue the architecture.

The COP model names the three content dimensions in competition, veracity, emotional appeal, and relevance, and adds the two filters that decide the contest: the audience's Overton window of explanations already acceptable to them, and perceived utility, meaning whether the content is worth consuming and passing on. Analyzing social-media data from the 2020 coronavirus period, the authors found readers markedly less sensitive to truthfulness under uncertainty and highly sensitive to tone and emotional appeal. The model does not claim truth is irrelevant. It explains why a weaker account can outperform a better-supported one in the interval that matters, and why a frame that already sits inside the audience's range of acceptable explanations needs no persuasion, only an event.


THE DIFFUSION ASYMMETRY IS HUMAN

Vosoughi, Roy and Aral, in Science in 2018, tracked about 126,000 cascades tweeted 4.5 million times by roughly 3 million people between 2006 and 2017. Falsehood reached 1,500 people about six times faster than the truth and was about 70 percent more likely to be retweeted, and removing every bot from the dataset left the difference intact. The dominant amplifier in the first window is therefore your own audience, acting in good faith, in a state that suppresses counter-argument.

SIMILAR EFFECTS, DIFFERENT CAUSES

A false claim in the window may be coordinated, opportunistic, commercial, political, automated, or sincere, and the same artifact can be carried by a mistaken witness, an operator with an objective, a newsbroker packaging uncertainty for an audience, an attention market rewarding whatever performs, and an institution whose silence or reversal supplies fresh material. The failure modes converge without any shared intent. This is why the analysis of what a claim is doing must be kept separate from the question of who is behind it, and why a pattern that looks coordinated is not evidence that it was.

FIGURE 1 // THE THREE STAGES, AND WHERE A CLEARED STATEMENT LANDS



THE WINDOW

A committee-cleared first statement usually publishes in Adoption or Resistance. It is not late by minutes but by a stage, which is why neither accuracy nor volume rescues it.

FIGURE 2 // THE COP MODEL: WHAT DECIDES WHICH ACCOUNT TRAVELS

VERACITY	How true the content is, or appears to be. Normally decisive, but under uncertainty the audience has neither the time nor the evidence to weigh it, and sensitivity to it drops.
EMOTIONAL APPEAL	How strongly the content engages fear, anger, grief, disgust, or sympathy. The dimension the study found audiences most sensitive to. The audience is not lazy, it is aroused.
RELEVANCE	How closely the content connects the event to the audience's identity, circumstances, or immediate concerns.
OVERTON WINDOW	The range of explanations already acceptable to that audience. A claim inside it needs no persuasion, only an event to attach to.
PERCEIVED UTILITY	The resolution of the above: whether the content is worth consuming and passing on. This is what the early frame competes on, and truth is one input to it rather than the test.

The public investigation begins before the official one can speak, and it runs on different rules. Open-source work genuinely surfaces useful footage, timelines, vehicles, locations, and witnesses. The same tools converge on innocent people, because a plausible face match, a vehicle near the scene, an old political post, or a device-location association can be mistaken for a completed identification. The danger compounds when apparent witnesses corroborate each other, aggregators package the claims, and search or generative systems repeat them: what looks like independent discovery is often several routes back into one contaminated pool. For a principal this is not an abstraction about strangers. The named person may be a relative, a member of staff, a business associate, or the principal.

SIGNAL ONE: PROXIMITY

A person, device, vehicle, residence, or workplace appears near the event, or near someone already under scrutiny. Presence is not participation, and in a dense environment it is not even unusual.

SIGNAL TWO: RESEMBLANCE

A partial image, a clothing match, a name collision, or a facial similarity is treated as identity rather than as a lead that still requires work.

SIGNAL THREE: NARRATIVE FIT

Old posts, employment, service history, family relationships, or politics make the candidate feel like the right suspect. This is the signal that converts a lead into a conviction, and it carries no evidentiary weight at all.

► THE PROTECTIVE QUESTION

Who is the public converging on, what evidence is actually carrying that identification, *and what harm arrives before it is corrected?*

FIGURE 3 // THE DOMESTIC HANDOFF: HOW A FRAME OUTLIVES WHOEVER STARTED IT

01 SEED OR DISCOVERY	An artifact enters circulation. It may be planted, misread, selectively edited, genuinely discovered, or fabricated for attention.
02 DOMESTIC ADAPTATION	Local actors attach it to familiar grievances. Different audiences receive different captions around the same material.
03 SOURCE EROSION	Reposts, screenshots, commentary, and reporting progressively separate the claim from its provenance.
04 APPARENT CORROBORATION	Multiple references point back into the same contaminated pool, producing the appearance of independent confirmation.
05 INSTITUTIONAL RECOGNITION	Officials or journalists address the claim because it is widespread. Attention is mistaken for validation.
06 SELF-SUSTAINING BELIEF	The frame now serves identity, commerce, politics, or community. Its original sponsor, if there was one, is no longer needed.

SO WHAT By the time an institution answers the claim, the claim no longer needs its author. Removing the source does not remove the belief, and the operation succeeds precisely when the operation disappears.

An arrest answers one question and can activate several more. Official resolution arrives into an environment where audiences have already assigned motives, allies, victims, villains, and institutional roles, so identification, capture, or death may settle the investigative question while making the suspect more useful as a story. The final image becomes proof of courage, coercion, execution, silencing, or conspiracy according to the frame already in place. Bucha in April 2022, the Prigozhin death in 2023, and Crocus City Hall in March 2024 all show the contest intensifying after the physical facts were settled, with each new piece of legitimate evidence meeting a prepared counter-frame that reads official proof as evidence the institution controls the record. The practical error is treating an arrest, an attribution, or an official briefing as the end of the information event. A monitoring plan that stands down at that point misses the most emotionally potent phase.

FIGURE 4 // WHERE THE HARM ACTUALLY LANDS, AND WHO ABSORBS IT

HARM	WHAT IT LOOKS LIKE	WHO ABSORBS IT
WRONG-PERSON TARGETING	An innocent person is named, located, threatened, and made permanently searchable	The named individual, and anyone who shares a name, face, or address with them
TARGET SUBSTITUTION	A presumed identity selects a wider community, workplace, residence, or symbolic site	Staff, neighbors, congregations, and premises with no connection to the event
INSTITUTIONAL COMPRESSION	Officials are pushed toward premature claims or into unexplained silence	The institution's credibility, in both directions and for years
PROTECTIVE OVERLOAD	Principals, relatives, witnesses, and accused bystanders all need protection at once	The security program, which was resourced for one protectee
ECONOMIC REPRICING	Insurers, lenders, and counterparties move on uncertainty before any legal resolution	Operating businesses, cover, credit lines, and continuity plans
MANUFACTURED MEMORY	Contaminated artifacts persist as part of what the public remembers happening	Everyone named in them, indefinitely

Two of those harms explain why a single statement cannot serve everyone. Public reassurance and risk stabilization are different communication tasks. The public needs an intelligible account of what is known, what is not, and what is being done. Risk systems need concrete operational conditions: authority succession, facility status, payment-system availability, transport continuity, verified threat geography, and the operating posture of critical institutions. Sentiment responds to meaning, while insurers, lenders, freight planners, and continuity teams respond to defined conditions. A general assurance that everything is under control satisfies neither.

FIGURE 5 // FOUR NUMBERS FROM THE DOCUMENTED RECORD

6_x Speed advantage of falsehood over truth in reaching 1,500 people; removing all bots left it intact. VOSOUGHI, ROY & ARAL, SCIENCE, 2018	4 Months from the false nuclear-test claim about the Semnan earthquake to the peer-reviewed refutation. FERNANDO ET AL., SEISMICA, 4 FEB 2025	~5% Peak Hormuz war-risk premium of vessel value, from 0.10 to 0.125 percent before the conflict. HOWDEN RE, 27 MAR 2026	17.2% Accounts showing bot-like behavior in the 2025 bank-liquidity campaign. The rest were people. BLACKBIRD.AI RAV3N, 2025
--	--	---	---

THE STEELMAN // THE CASE FOR THE SLOW CLOCK

THE BEST CASE FOR THE OTHER SIDE

- + Verification and legal review exist for good reason. An organization that speaks in seventeen minutes will sometimes amplify its own error, create liability, or prejudice an investigation it does not control.
 - + Pre-committing to holding language risks locking a position in before the facts are known, and the holding statement can itself become the story.
 - + The threshold is our own planning figure, deliberately demanding, and no published measurement establishes seventeen minutes as the point of hardening. The honest claim is a minutes-to-hours window whose width varies by event.
- Concede all three, and the third is our own stated limit rather than an outside objection. None of them argues against the architecture. The doctrine does not accelerate verification and does not assert contested facts: it publishes what is being done, by a named person, through channels that already exist, while verification proceeds at its own pace, and a statement describing process can be superseded without contradiction. Two further limits belong beside it and are ours as well. There is no attribution by pattern, since bridge narratives, false witnesses, and source laundering can be coordinated or can arise from convergent incentives, so speed never licenses naming anyone. And there is no permission to overreact: fast communication must still protect evidence, privacy, innocent people, operational security, and the institution's own future credibility. Those three limits are what separate a prepared posture from a reckless one.

FIGURE 6 // ANALYTIC CONFIDENCE BY CLAIM

An early frame gains reach before institutions can establish basic facts, within a window of minutes to hours.	STRONG / DOCUMENTED CASES
Institutional first statements typically arrive after the frame has set, and correction then carries the burden.	STRONG / SOUTHPORT, SEMNAN, THOMPSON
Under uncertainty, audiences are less sensitive to truthfulness and highly sensitive to emotional appeal.	STRONG / FRENCH, GEORGE, MADDEN & STOREY, 2025
The diffusion asymmetry is driven by human sharing rather than automation.	STRONG / VOSOUGHI, ROY & ARAL, 2018
The contest intensifies after physical resolution rather than closing with it.	STRONG / BUCHA, PRIGOZHIN, CROCUS CITY HALL
Crowd identification runs on proximity, resemblance, and narrative fit, and can converge on an innocent person.	STRONG / SOUTHPORT AND THE KIRK MISIDENTIFICATIONS
Perceived danger moves insurance pricing ahead of confirmed loss.	STRONG / HORMUZ PRICING AGAINST ATTACK COUNTS
Seventeen minutes as the planning threshold.	DESIGN PARAMETER / DELIBERATELY DEMANDING, NOT A MEASUREMENT, NOT A CLAIM ABOUT ANY EVENT
The mechanism runs on domestic material alone, with no adversary direction.	MODERATE / CHARLIE KIRK IS THE CLEAREST SINGLE CASE
Synthetic witnesses outperform deepfakes because they attract less scrutiny.	HYPOTHESIS / PLAUSIBLE, UNTESTED, AND STATED AS SUCH IN THE SOURCE ANALYSIS
Seventeen minutes is a constant that holds across events.	REJECTED / THE SOURCE ANALYSIS REJECTS IT AND SO DOES THIS BRIEF
A convergent pattern of claims is evidence that the claims were coordinated.	REJECTED / CONVERGENT INCENTIVES PRODUCE THE SAME SHAPE
A later correction reaches the audience that received the original claim.	REJECTED / SEMNAN AND SOUTHPORT BOTH RUN THE OTHER WAY

POSTURE

FOR PRINCIPALS

- Audit the chain for the first statement only and ignore everything downstream of it. Name the people who may publish holding language without clearance, write that language now, and confirm the channels are already yours and already have an audience. If legal review sits in that path, the chain is structurally too slow and drafting quality will not fix it.
- Rehearse against the frame that is plausible for you specifically rather than a generic crisis. The frames that win are the ones already inside your audience's range of acceptable explanations, so the exercise is to find yours before someone else does.
- Decide in advance what you will do if the crowd names someone close to you. That is a different problem from your own exposure, it arrives faster than any official process, and it is the one nobody has a plan for.

FOR FAMILY OFFICES

- Pre-bunk ahead of predictable flashpoints, using the six manipulation techniques the inoculation research names: discrediting, emotion, polarization, impersonation, conspiracy, and trolling. Inoculation reaches people before the aroused state does, which is why it outperforms correction after the fact.
- Treat footprint reduction as standing infrastructure rather than incident response. Pre-event suppression for the principal, the family, and named public-facing staff lowers the targeting precision available for impersonation and for synthetic-witness accounts, and none of it can be done inside the window.
- Preserve the record in the first hour: original posts, timestamps, account identifiers, and the artifact as first seen. The material that later establishes laundering or coordination is routinely deleted, and the capture you take in minute ten is the evidence in month six.
- Keep monitoring past the arrest. Martyr conversion, false-flag branching, and the reading of official proof as institutional control all begin after the case appears closed.

FOR SECURITY & LEGAL FUNCTIONS

- Draft and pre-clear holding language now, asserting process and values and no contested facts, with the sign-off recorded in advance so publication in the window is an execution step and not a decision. Agree in writing what a delegated publisher may and may not say.
- Write the misidentification protocol before it is needed: how a wrong identification is detected, who contacts platforms and press, what is said to and about the named person, and what is never said, including any repetition of the claim while denying it. Any identification of a person still has to clear the CB-246 validation gate, and proximity, resemblance, and narrative fit do not clear it in any combination.
- Carry actuarial exposure as a standing monitoring input, and prepare the risk-stabilization message separately from the public one. Insurers, lenders, and continuity teams need defined conditions such as authority succession, facility status, payment availability, transport continuity, and verified threat geography. Pre-coordinate so a narrative-driven repricing meets a prepared answer rather than a discovery. This is intelligence analysis, not legal advice.

CONFIDENCE High on the documented cases and on the structural claim that an early frame gains reach before institutions can establish basic facts. High on the external research quoted and used at the level its authors state it: the COP model and its uncertainty findings from French, George, Madden and Storey; the diffusion asymmetry from Vosoughi, Roy and Aral; the Semnan timeline and refutation date from Fernando and colleagues; the Hormuz premium figures from Howden Re including the Lloyd's statement; the campaign attribution and bot proportion from Blackbird.AI. The seventeen-minute figure is not rated as a claim because it is not one: it is ObscureIQ's planning threshold, chosen to be demanding, and the defensible finding underneath it is a minutes-to-hours window whose width varies by event salience, audience, platform, and institutional readiness. Moderate on the claim that the mechanism runs on domestic material alone, where Charlie Kirk is the clearest single case rather than a pattern. The synthetic-witness advantage is carried as a hypothesis, as the source analysis carries it. The Semnan 17-minute figure is time to first post and is unrelated to the threshold; both numbers appear here and the distinction is stated where each is used.

METHOD Revised edition. First published 2 May 2026; facts re-verified and analysis revised 10 September 2026 against the current source analysis. This edition states the threshold as a planning parameter over a minutes-to-hours window rather than as a measured constant, attributes the adoption mechanism to the peer-reviewed COP model and adds the Overton window and perceived-utility elements the earlier edition omitted, adds the second identification system and its three signals as a principal-facing section, adds the domestic-handoff sequence and the resolution-without-closure phase, maps where the harm lands and who absorbs it, separates public reassurance from risk stabilization, corrects the Hormuz war-risk figures and the bank-campaign attribution against primary reporting, and removes an unsourced cognitive-load statistic. The source analysis labels its own material as invented, observed, inferred, or modeled; only the observed and the externally published are carried here as evidence. Cross-read against CB-060, CB-062, CB-113, and CB-246.

SOURCES ObscureIQ, JFK 2026 narrative-warfare analysis: [The 17-Minute Problem](#) (main paper, scenario analysis), [Case Studies](#), and [Introduction](#). Aaron M. French, Amrita George, Joshua Madden and Veda C. Storey, "[Tabloids, Fake News, and the Overton Window: The COP Model on News Consumption in Uncertain Times](#)", Information Systems Frontiers 27(6), 2025, 2317 to 2335, doi:10.1007/s10796-025-10623-0. Soroush Vosoughi, Deb Roy and Sinan Aral, "The spread of true and false news online", Science 359(6380), 2018, doi:10.1126/science.aap9559. Benjamin Fernando et al., "[The propagation of seismic waves, misinformation, and disinformation from the 2024-10-05 M 4.5 Iran earthquake](#)", Seismica, 4 February 2025, doi:10.26443/seismica.v4i1.1512. [Howden Re, Strait of Hormuz report](#) (27 March 2026). [Blackbird.AI RAV3N on the 2025 bank-trust campaign](#). Companion doctrine in CB-060, Response Doctrine: The First 60 Minutes.

Scope note: the seventeen-minute threshold is a planning parameter drawn from ObscureIQ's scenario analysis, deliberately set at the demanding end, and no claim is made that it holds across events. That analysis is built on a fictional event with documented mechanisms, and its simulated artifacts, timings, volumes, and sequences are test conditions rather than forecasts; none is cited here as precedent, including its Strait of Hormuz and Operation Epic Fury

scenario. Where documented 2026 events are described they rest on public reporting and are separate from the scenario that shares a name. Case reconstructions rest on open-source reporting. This is intelligence analysis, not legal advice.

OBSCUREIQ INTERNAL ANALYSIS // ARCHIVE COMM-057-NARRATIVE-17MIN-2026Q2 // XREF [CB-060](#) [CB-062](#) [CB-113](#) [CB-246](#) // ACCESS PUBLIC · PUBLIC-INTEREST RELEASE · REVISED 10 SEP 2026