# Building Shared Reality After the Broadcast Era

Matthew Cashman
MIT Sloan, 100 Main Street, Cambridge, MA, USA
Email: cashman@mit.edu

2026-07-11 (latest version here)

**Abstract**

The era of broadcast (c. 1920–2000) produced a distinctive epistemic condition: a public that attended to the same facts, took them to be true, and held both of those things in common. Post-truth scholarship rightly treats the loss of shared reality as a broad crisis of misinformation, institutional trust, polarization and fragmented public attention. We argue that this crisis is better understood by separating two conditions that broadcast bundled: first, one-to-many, limited-bandwidth, synchronous distribution made claims publicly attended and approximately commonly known, and second, recorded media made these claims believable because it was prohibitively difficult to fake. Neither condition now holds: the internet fragments mutually known attention, while deepfakes and synthetic media weaken self-authentication. Existing accounts examine truth, trust, media fragmentation and deepfake believability. The step they miss is to separate the record that makes claims commonly inspectable from the attention that makes them commonly known. We give design criteria for the record layer of post-broadcast common-knowledge infrastructure, identify the coupling of such records to legible mass attention as the central open problem, and show why public ledgers, transparency logs, provenance standards and recommender feeds each supply only part of what broadcast once gave us.

On the evening of February 27, 1968, Walter Cronkite closed a CBS News special report by telling a national television audience that the war in Vietnam was "mired in stalemate"[1]; some nine million Americans heard him say it at the same moment[2]—fewer than the evening newscasts drew routinely—and the verdict entered public discussion around water coolers and kitchen tables nationwide. The legend that this broadcast turned a president and a nation against the war is myth[1]; the episode interests us for its distinctive structure. Viewers received the assessment while knowing that others, including neighbors, representatives and adversaries, were receiving the same message at the same time. They could also believe what they saw, because the footage was not the kind of thing that could be cheaply faked. Democratic legitimacy, accountability and coordination lean on episodes of that kind. When the conditions that produced them eroded with cable TV and collapsed with the internet, diagnoses followed quickly: we had entered a "post-truth" era of misinformation and dissolving shared reality[3].

Taken charitably, that diagnosis is more than a complaint about false belief. The post-truth literature identifies an interlocking crisis of institutional trust, epistemic authority, media fragmentation, polarization and shared public facts[4–6], a diagnosis that has itself drawn criticism[7]. The

phrase "shared reality", however, bundles mechanisms that infrastructure needs to separate. What broadcast supplied, and what the present environment has lost, was a combination of two factors. The first is common knowledge: everyone learns $p$, everyone knows that everyone knows it, and so on[8]. Strict common knowledge is unattainable through ordinary message-passing[9], but a public signal can sustain the common $p$-belief that supports coordination[10;11]. Broadcast supplied such signals cheaply as a byproduct of its topology (Fig. 1)[12]. The second factor is believability: what is commonly known can be trusted as real, which mid-century recorded media delivered by being prohibitively expensive to fake. Common knowledge of a claim no one believes is propaganda, and a believable record no one jointly attends to is a document in a drawer. Shared epistemic reality needs both.

This loss is real and distinct from the truth and authentication problems into which it is folded. Unlike interpretive authority, it is addressable by infrastructure. The aim is not to restore consensus. It is to recover a narrower public capacity for coordination and legitimacy under disagreement, what we call "navigable pluralism": coordination on believable public facts without agreement on their meaning. This leaves a tractable design problem. In this Perspective we separate the conditions broadcast bundled, give structural design criteria for the record layer that can make claims commonly inspectable, treat the attention layer that would make them commonly known as the central open problem, and evaluate proposals such as media-provenance standards, public ledgers, transparency logs and recommender feeds against the criteria. Our scope is contested public facts in democratic polities; crisis communication and scientific consensus raise related but distinct problems.

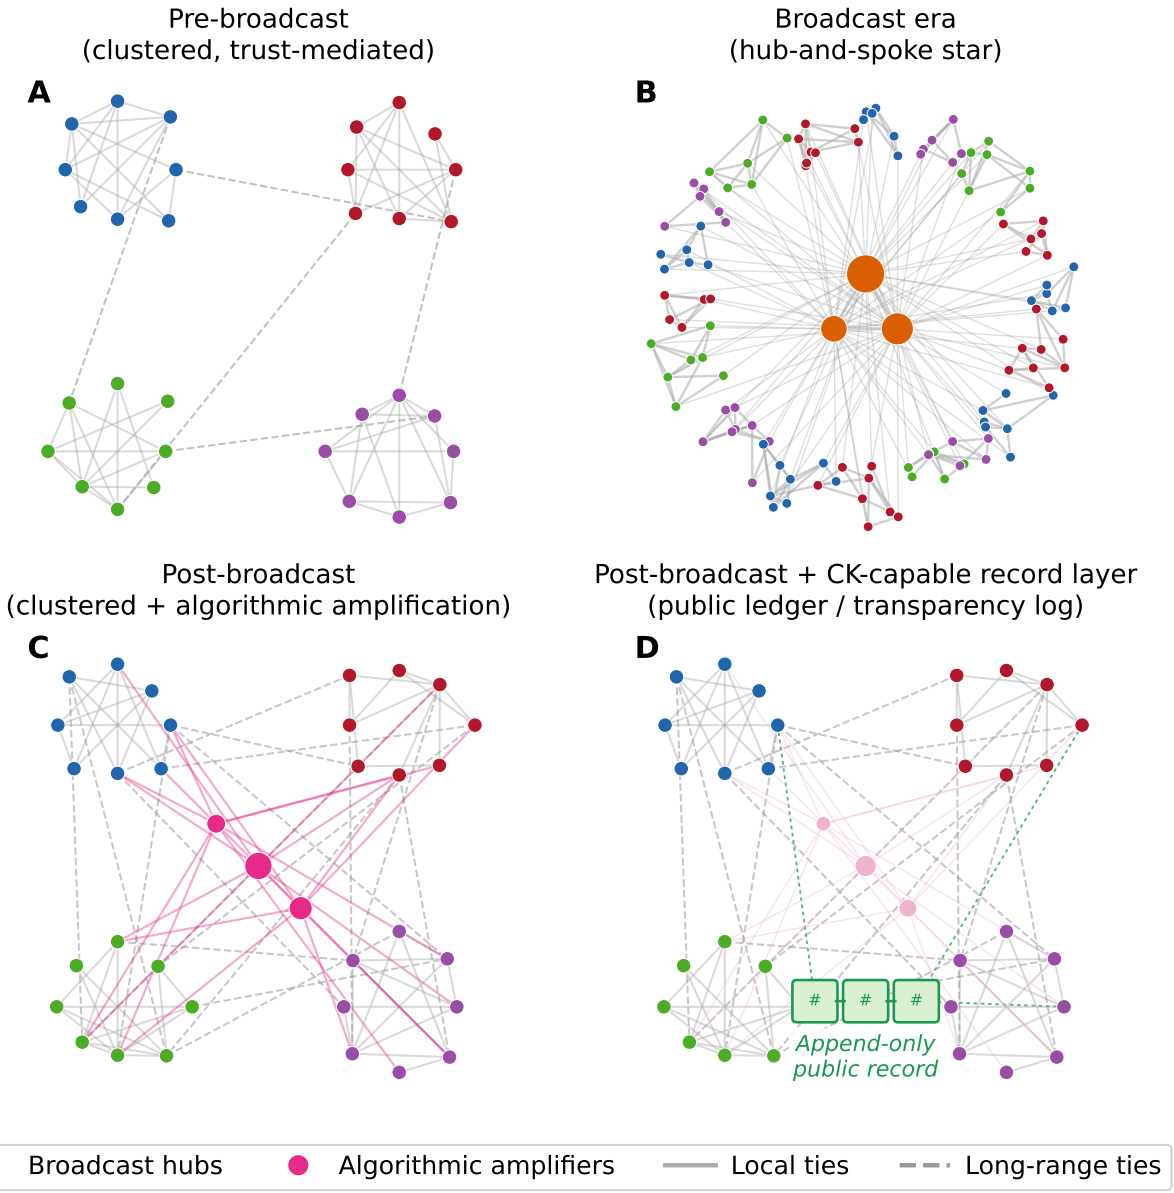


Figure 1: **Communication topologies and common-knowledge generation.** (**A**) Pre-broadcast: information propagates through clustered trust networks with sparse inter-community bridges. (**B**) Broadcast: a small set of hubs reaches a still-clustered audience at once, creating a public signal. (**C**) Post-broadcast: clustered communication reappears beneath algorithmic amplification; personalization and opaque ranking break observer equivalence, so coincident attention rarely becomes common knowledge. (**D**) A post-broadcast topology with a CK-capable record layer (a public ledger or transparency log) added. The record can anchor common knowledge among those who consult it, but carries no simultaneity or mass attention of its own. Dashed edges mark long-range ties; (**D**) is a design proposal rather than a historical stage.

**Box 1 | Glossary.**

**Shared epistemic reality.** The good broadcast produced and that the post-truth debate calls "shared reality": a public that jointly attends to the same facts, trusts them as real, and takes both to be held in common; the conjunction of two factors, common knowledge and believability.

**Common knowledge (CK).** In the strict game-theoretic sense of Aumann[8], everyone knows $p$, everyone knows that everyone knows $p$, and so on at every iteration. Because finite message-passing rarely delivers that ideal, this article uses common-knowledge generation as shorthand for the practical approximation, common $p$-belief or publicly grounded mutual belief, that sustains coordination unless the stricter sense is at issue.

**Coincident attention.** Many participants attending the same content at roughly the same time, without necessarily producing mutual awareness of who else is attending.

**Public legibility.** The condition that turns coincident attention into common knowledge: the attending public is large and known to be large, so each can take it that others are attending the same signal.

**Self-authenticating media.** Recordings believable by default by virtue of existing, because fabrication was prohibitively expensive (mid-century photography, film, audio). A property of the capture layer only: staging and miscaptioning were always cheap. Deepfakes end self-authentication, and digital provenance recovers only part of it.

**Navigable pluralism.** The normative target defended here: communities with different interpretations coordinating on common knowledge of believable public facts, without requiring agreement on meaning.

**Observer equivalence.** Two observers running the same verification against the same substrate return identical results; satisfied by broadcast signals and public ledgers, broken by personalized feeds.

**Public verifiability.** Any party can inspect and confirm substrate state without privileged access, institutional credentials, or trust in a specific verifier.

**Append-only, tamper-evident state.** A record that accumulates, cannot be silently revised, and whose history all observers agree on; a precondition for common knowledge about anything it records.

**Shared protocol with knowable specification.** Participants know how the substrate works and know that others know; opaque ranking algorithms and undocumented moderation policies violate this.

**C2PA.** Coalition for Content Provenance and Authenticity: an open cryptographic provenance standard for media, binding signatures to content across edits and redistribution.

**Device attestation.** Hardware-rooted cryptographic proof that an image or signal originated on a specific device and has not been altered since capture.

**Certificate transparency log.** An append-only public log of TLS/SSL certificates; an existence proof that the record layer of distributed common-knowledge infrastructure can work at web scale within a narrow domain.

# What broadcast bundled

Broadcast's achievement is best understood as *shared epistemic reality*: a public that attends to the same claims, takes them to be true, and takes that attention and belief to be held in common. A distribution apparatus, the one-to-many, limited-bandwidth, synchronous broadcast topology, made claims commonly known because broadcasts were seen *together*. A capture apparatus made up of photography, film and audio recording delivered believable media, because faking such records was prohibitively difficult and therefore rare. Shared epistemic reality is the conjunction of the two: claims that are commonly known and also believed.

Broadcast's distribution apparatus enabled coordination. Technical and economic limitations, including spectrum scarcity, licensing, network economics and a limited channel menu, meant that relatively few broadcasters reached very large audiences at national and local scales[13–16]. The point is scale, not national reach alone. A local broadcaster reaching hundreds of thousands or millions in a metropolitan area still creates a far larger common audience than private message-passing can. The 1968 Cronkite broadcast drew nine million Americans, and the era's peak events drew several times that audience. High trust in news institutions made those concentrated signals believable to many audiences[17]. Coincident attention then became closer to common knowledge when the audience was publicly legible: Cronkite's viewers heard him, knew the audience was large, and could take it that others were receiving the same signal. Broadcast network topology supplied this cheaply. The iterated structure does real work: Lewis showed that conventions rest on it[18], and Chwe showed that mass media and public ritual coordinate because participants know that others are attending too—Super Bowl advertising is valuable not for its reach alone, but because viewers know others are watching[12]. Farrell and Schneier extend the political stakes, arguing that democratic norms depend on common political knowledge and that social media degrades it[19]. These are friendly antecedents. Our addition is to separate this awareness condition from believability and to ask how it can be generated deliberately in post-broadcast infrastructure.

Two scope conditions discipline this baseline. First, common knowledge among whom: broadcast publics were stratified by race, region, class and language, and they were national only for peak events. The typical case was approximate public establishment within large but bounded audiences, the viewership of a newscast or a metropolitan market, not a single unified public. Second, how much: trust in the institutions carrying the signal was high on average but heterogeneous, and that high average was itself a mid-century anomaly[17]. Broadcast is therefore best described not as an era of generalized shared reality but as a historically unusual increase in the frequency and scale of public signals. How much common knowledge those signals produced, and among whom, is an open empirical question, and the measurement program we propose below could address it retrospectively.

Common knowledge of a claim is not yet a believed shared reality. Unreliable broadcast regimes, such as Soviet state media, had the awareness machinery: one-to-many, synchronous, and publicly known. A public line treated as propaganda rather than trusted evidence, however, produces only hollow common knowledge. What mid-century Western media had that such cases lacked is believability. A photograph or film was credible by virtue of existing: at the time, the causal link from event to image was mechanical and hard to counterfeit[20], so viewers saw "through" the image to the depicted object[21] and learned to equate truth with seeing[22]. Such media *self-authenticated*, carrying their own warrant with no external verifier required. The term needs a scope note: it names a default of the capture layer, not a guarantee. Fabricating a persuasive

photographic record was prohibitively costly, but staging a scene or miscaptioning an authentic image was always cheap, and institutional trust and genre carried part of the credibility load even then. Both caveats mark where digital reconstruction will fall short. Even so, these recordings could function as what Rini calls an epistemic backstop[23]: evidence that could override testimonial disagreement because fabrication was prohibitively expensive. The deepfake and provenance literature has mapped this believability problem in detail[24;25], though some argue the threat is overstated[26;27]. Our novelty is not there. It is in treating believability and awareness as separable conditions when evaluating remedies.

A third broadcast function is often folded in with these: *agenda-setting*, the power of a small source set and professional gatekeeping to decide which topics entered national awareness[16;28;29]. Agenda-setting is central to media studies but differs in kind from the two factors at issue here. It is coincident attention pointed at a curated topic set; the narrow shared agenda was downstream of limited bandwidth and centralization. Cable is the natural experiment: as bandwidth expanded, the shared agenda fragmented first, while self-authentication remained intact. Agenda-setting is also the broadcast function whose loss we might welcome, since a substrate that decided everyone's topics would be the propagandist's instrument. Infrastructure can concentrate attention onto a topic set, as ranking algorithms do, but which topics matter is a governance and values choice no design criterion can settle. We therefore treat agenda-setting as governance riding on the awareness factor, not a good we would build infrastructure to rebuild.

That both factors came bundled was historically contingent. The infrastructure behind them, like spectrum scarcity, regulatory allocation, advertising economics, and audiovisual technology, was specific to a narrow window[13–16;30;31]. The dating is heuristic: radio and television entered decades apart, cable began eroding the configuration in the 1980s, and cross-nationally the parenthesis opened and closed at different moments[32]. The staggered rollout, idiosyncratic availability, and channel placement of broadcast and cable television have supplied quasi-experimental evidence that exposure shaped political alignment[33–36], though these are effects of slant and persuasion, sometimes producing divergence, not direct evidence about common-knowledge generation. The public-service monopolies of postwar Western Europe concentrated the awareness factor more tightly than commercial US networks, while the Soviet case shows a distribution apparatus producing public awareness without trusted evidence. Beneath all of them, the older social substrate of clustered networks of kin, workplace, and congregation long predates broadcast[37–39]. Broadcast was an overlay, not a replacement for small local networks of trust, and post-broadcast fragmentation is partly the re-exposure of that substrate, amplified and changed by digital media but not manufactured from nothing[40;41].

When the broadcast apparatus lifted, both factors gave way for different reasons. The internet's many-to-many topology, effectively unlimited bandwidth, and asynchronous capabilities replaced the broadcast distribution apparatus, preceded by cable TV. This fragmented common knowledge across smaller communities. More recently, deepfakes have ended self-authentication: a realistic image is no longer very likely to be real by virtue of existing. The broad post-truth literature has recognized parts of this story: media choice fragments shared audiences[16], networked public spheres lose shared experiences[5], attention markets fragment authority[3], and common political knowledge can be attacked[19]. The remaining gap is a framework that holds awareness and believability apart, shows that broadcast bundled them contingently, and specifies the infrastructural conditions under which post-broadcast attention can become common knowledge or common $p$-belief.

The distinction matters because remedies aimed at one condition are often described as restoring the whole. When C2PA's founding press release calls the standard "a vital step in restoring society's shared sense of reality"[42], it conflates believability with awareness: a provenance credential can tell a viewer that an image came from a signed source, yet it cannot tell her who else has seen it or whether that viewing is public. A recommender system can send the same item to millions, yet personalized feeds and opaque rules prevent the audience from verifying that it has received a common signal. Such systems may improve trust or increase attention while still failing to recreate shared epistemic reality.

Because broadcast supplied both factors through a single apparatus, it may be natural to assume a single mechanism can restore them. Figure 2 shows why this is not so: shared reality is the conjunction of common knowledge and believability, the broadcast era held the corner where both were high, and the two post-broadcast losses pull away along independent axes. Table 1 adds detail: each technology we examine supplies at most part of the bundle. The next section asks what a substrate would need to rebuild awareness, and what the believability factor's reconstructions can and cannot do.

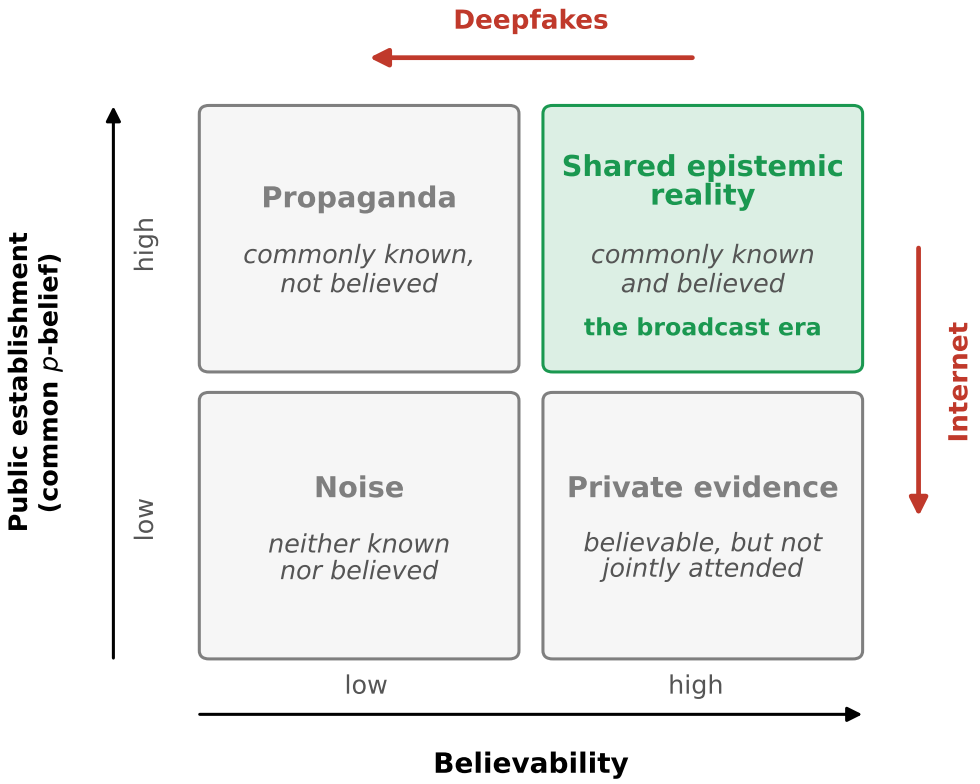


Figure 2: **Shared epistemic reality is the conjunction of two factors.** Public establishment of a claim (vertical; in practice common $p$-belief) and its believability (horizontal) are independent properties of a public's relation to that claim. Only when both are high is the claim a shared epistemic reality. The off-diagonal cells are the standard failures: common knowledge of a claim no one believes, and a believable record no one jointly attends to. Broadcast held the top-right corner; the internet fragmented common knowledge, and deepfakes weakened believability. The arrows are schematic: networked media can still concentrate attention at moments, and deepfakes erode believability partly through second-order doubt about genuine records.

## Rebuilding common knowledge and believability

Rebuilding each factor places different demands on a distributed substrate. The awareness factor, common knowledge, is the hard problem and takes most of the design work; the believability factor is being addressed through provenance and attestation, but only partially and with new trust dependencies. We treat the awareness factor first, articulating the record it requires, then turn to the attention that record needs.

## Common-knowledge substrates

A substrate that can make a claim common knowledge, or serve as a commonly observable signal a dispersed public can condition on, needs four properties. Drawing on the common-knowledge literature, Chwe's analysis of ritual and mass media, and network epistemology[9;12;43–46], we argue that they are jointly necessary for asynchronous, low-trust digital substrates. Broadcast itself met almost none of them: its signal was ephemeral, its state was nowhere publicly verifiable, and its audience was known only through proprietary panels and social inference. Synchrony, channel scarcity, and institutional trust did that work instead. Where those substitutes are unavailable, which is the post-broadcast condition, the four properties become necessary, and that is why the design problem is new rather than solved by history. The properties often overlap in practice, but each isolates a distinct way a substrate can fail to be a public signal.

*Public verifiability.* Any party must be able to inspect the substrate's state and confirm what is recorded there, without privileged access, institutional credentials, or trust in a specific verifier. Credentialed access produces authenticated evidence at best: a public receives testimony from authorized inspectors rather than sharing access to the same observable signal. The important property is not that everyone will verify everything. It is that anyone can, and that everyone can know that misrecording or tampering is publicly contestable.

*Append-only, tamper-evident state.* The substrate must maintain a history that all observers agree on and that cannot be silently revised after the fact. If two observers can sincerely disagree about whether $p$ was ever said, $p$ cannot become common knowledge. Certificate transparency logs[47] and public ledgers such as blockchains[48] satisfy this property by construction. Cryptographic timestamping without a public log, for example RFC 3161 authorities acting alone[49;50], provides tamper-evidence but not observer-shared history unless the timestamps are themselves published on a public substrate.

*Observer equivalence.* Any verification by one party must return the same result as the same verification by any other. Algorithmic personalization most directly destroys this: when the feed shown to me is not the feed shown to you, and especially if the precise state of the feed is ephemeral, we cannot build common knowledge about what was said[5;19]. In the transparency-log literature the same requirement appears as resistance to equivocation, or split-view attacks, in which an operator shows different observers different versions of the log[47]. Broadcast had observer equivalence by design; per-user customization removes it.

*Shared protocol with knowable specification.* The rules of the substrate must themselves be common knowledge: participants must be able to know how it works and know that others know. Opaque ranking and undocumented moderation break this. A substrate with secret rules cannot generate common knowledge even if every other property holds, because no one can know what "recorded on the substrate" means in a way others would agree with.

A substrate meeting all four criteria is a commonly observable public signal rather than private message-passing. That distinction lets it escape the *asynchronous impossibility*: strict common knowledge cannot be assembled by passing messages one participant to another, because no finite exchange completes the infinite regress[9]. A ledger anyone can verify still does not give "everyone is verifying and knows that others are." Converting such a record into common knowledge, in practice the common $p$-belief that coordination requires[10], additionally requires coincident attention with mutual awareness of that attention. The criteria therefore identify a record layer rather than a full social event. They make a claim available for common knowledge; the attention layer

determines whether a relevant public actually treats it that way. Substrates satisfying all four criteria (as do public blockchains, certificate transparency logs, or the gazette text of a law, which is amended by formal process rather than silent revision) are common-knowledge-capable—hereafter CK-capable—for whatever they record. In contrast, substrates satisfying none (personalized feeds, opaque moderation, or direct messages) cannot anchor common knowledge at all.

Financial and prediction markets show the distinction in miniature: prices are public, observer-equivalent records[51;52] that attentive traders know others are watching—where a consolidated quotation exists, itself an institutional artifact. They also show the limit, because the attending public is narrow: for most prices, even crude oil, almost no one is watching, and wider attention flares only at salient moments such as a market crash. Record availability is easier than mass, legible attention.

## The attention layer

A record meeting all four criteria needs an attending public. Bare coincident attention requires only that people can look at once and want to. What converts attention into common knowledge is mutual awareness of scale: participants know that many others are attending to the same thing. Live broadcasts delivered this, by a weaker mechanism than remembered: broadcast-era audience scale was known through proprietary panels and social inference, not verification. The fair standard for successors is parity, not perfection. By that standard, platform view counters and trending surfaces are the nearest existing instance of legible attention: visible to all and roughly observer-equivalent, but attested only by the platform and manipulable by it or through it. What verifiable attention measurement would require, and who could attest audience scale without becoming a new single point of trust, is an open design problem. A million livestream viewers reach common knowledge only if they know others are watching. Coupling a record to legible attention at that scale is the open problem of the awareness factor.

## Believability and trust

Mid-century media self-authenticated: a recording was credible by virtue of existing. Deepfakes have brought that to an end, and the digital replacements recover only a subset of its properties.

*Existence and integrity.* A hash committed to a public, append-only ledger proves that a given record existed by a certain time and has not changed since. This is cheap and, formally, requires trusting no one, though in practice it inherits assumptions about consensus, client software, and the intermediaries through which most people check. But it is content-agnostic, certifying that a bitstring existed rather than that the bitstring depicts anything real. It will notarize a deepfake as readily as a photograph.

*Timing and predictability.* A ledger timestamp is half of a proof of origin. It gives an upper bound on when a recording was made; it says nothing about the earliest moment the recording could have been made. The ledger cannot supply that lower bound, but sometimes the event can. Suppose a meteor breaks up over Boston at 14:03, and footage of it carries a hash committed to a public ledger at 14:04. No one knew the meteor was coming, so no one could have begun faking the footage before the sky lit up, and one minute is too little time to fabricate a convincing scene from scratch. The same timestamp does far less for a gaffe at a campaign rally. The candidate's schedule had been public for weeks, so the stage, the lighting, the face and the voice could all be rendered

in advance; a clip fabricated at leisure can be committed to the ledger mid-rally and borrow the timestamp's credibility. The few details not knowable ahead of time are covered by preparing variants and releasing whichever one matches, and unrevealed commitments are invisible, so nobody can count how many were prepared.

The lower bound can also be manufactured. The classic device is the hostage photograph posed with today's newspaper: the front page could not have been known in advance, so the photograph can be no older than the paper. A freshness beacon is the digital version of the newspaper, an unpredictable public value, such as a recent block hash, captured inside the recording itself. The beacon proves the recording was made after that value appeared; the ledger commitment proves it existed before the commitment. Together they form a sandwich, pinning the recording's creation to the interval between them. The condition has two parts. Let $t_b$ be the moment the beacon's value became public, $t_\ell$ the moment the recording's hash reaches the ledger, and $\Delta = t_\ell - t_b$ the certified window. Let $H$ denote the recording's residual surprise: how much of what it shows was not derivable from public knowledge at $t_b$ (formally, the conditional entropy of the content given that knowledge). An adversary can fabricate inside the window or prepare variants in advance and select among them. The sandwich certifies origin only if both strategies fail:

$$\Delta < \tau_{\text{synth}}(H) \qquad \text{and} \qquad \log_2 N_{\max} < H, \tag{1}$$

Here $\tau_{\text{synth}}(H)$ is the time needed to synthesize the surprising residue at convincing fidelity. It increases with $H$ because the predictable parts can be rendered in advance. $N_{\max}$ is the number of candidate fakes the adversary can afford to prepare before $t_b$; roughly $2^H$ candidates are needed to cover the surprise. The conditions fail for different reasons. Faster synthesis undermines the first: as $\tau_{\text{synth}}$ falls, $\Delta$ must shrink to keep pace, and once fabrication is quicker than the smallest window the infrastructure can certify, no sandwich works for that class of message. Predictability undermines the second, which contains no $\Delta$ at all: selecting and committing a pre-made fake takes seconds inside any window, so no shortening of the window defends against a predictable event. The meteor and the rally illustrate the two limits of Eq. 1: the meteor's $H$ is large, so both conditions hold with an ordinary window; the rally's $H$ is a few bits, so thousands of prepared variants cover it and almost nothing is left to synthesize. Events short on surprise must buy their lower bound with trust instead. Once synthesis outruns any certifiable interval, the remaining defense is plurality: many observers committing their own recordings promptly, from different vantage points, whose contents have to agree. The audience the awareness factor assembles is also a witness set; at scale, coincident attention feeds believability rather than sitting downstream of it.

*Origin and veracity.* Establishing that a record is a genuine capture of a real event is harder. Device attestation can claim that a particular sensor captured a particular image at capture time[53], but only if one trusts the manufacturer's keys. An attested camera can instead be pointed at a screen. C2PA provenance chains inherit both limits. Origin attestation therefore concentrates trust in hardware manufacturers, certificate authorities and standards bodies. None of these layers requires public verifiability in the common-knowledge sense: checking a hash, a beacon-to-ledger interval, or a signature is private verification, which is why such checks scale to individual artifacts without coordinating anyone's attention. Substituting this concentrated trust for the broadcaster's authority needs a governance argument the provenance literature rarely supplies. The question is who controls the keys, who revokes them, how compromised devices are handled, and how coerced or captured authorities are made visible. These are institutional questions as much as cryptographic

ones.

*Recontextualisation and dismissal.* Most believability attacks do not involve fabrication. In audits of circulating misinformation, miscaptioned or lightly edited authentic media far outnumber synthetic fabrications: one study of fact-checked pandemic misinformation found 59% was reconfigured genuine content against 38% fabricated, with no deepfakes in the sample[54], and the most accessible manipulations are contextual rather than technical[55]. Recontextualisation attacks the link between record and claim, which capture-time attestation does not secure. The liar's dividend runs the attack in reverse: once fabrication is known to be possible, genuine records can be waved away as fake, so deepfakes damage true evidence without a single fake being made[24]. Believability therefore has at least four layers: capture (was the scene real?), integrity (has the file changed?), provenance (where did it come from?), and interpretation (is the caption honest, the inference warranted?). The infrastructure discussed here addresses integrity and provenance; capture is partly reachable through attestation and timing; interpretation is not an infrastructure problem.

Together, the criteria distinguish proposals that supply a common-knowledge-capable record, attention at scale, or believability from those that merely assume them. We can achieve common knowledge in narrow technical domains such as cryptocurrencies or financial markets, while believability is being addressed only partially and through systems that require trust. Domain-specific couplings of record and attention exist, as the market case above and the surviving synchronous channels below show; what does not yet exist is a general-purpose coupling that works across contested public facts.

# Evaluating current and proposed infrastructure

The criteria show what each proposed technology supplies and what it lacks. Table 1 summarizes the key cases.

Table 1: **Post-broadcast infrastructure, evaluated against the design criteria.** "CK properties met" refers to public verifiability, append-only state, observer equivalence and shared protocol. Systems meeting all four are CK-capable public records within their domain. Actualizing common knowledge still requires coordinated attention.

| **Technology** | **Primary product** | **CK properties met** | **Notes** |
|---|---|---|---|
| Content-provenance credentials (e.g., C2PA) | Authenticated evidence | None (not the record layer's aim) | Artifact-local credentials under a knowable verification protocol; no shared public state, no mutual awareness. Often oversold as CK-producer. |
| Device attestation (secure enclave) | Authenticated evidence | None required | Concentrates trust in hardware manufacturers. |
| Public blockchains (e.g., Bitcoin) | CK-capable public records of transactions | All four (formal); weaker in practice | Narrow domain, record layer only. Most users rely on intermediaries rather than verifying; finality probabilistic; miner/validator concentration. *Open: couple the record to coordinated public attention.* |
| Certificate transparency logs | CK-capable public records of certificates | All four (formal) | Narrow domain; in practice depends on monitors, gossip, and browser enforcement. Existence proof that the CK record layer works at web scale. *Open: extend beyond narrow technical domains under accountable governance.* |
| RFC 3161 timestamping (without public log) | Authenticated evidence | Partial (append-only only if published) | CK-capable only when timestamps are themselves published on a public substrate. |
| Recommender-driven feeds | Coincident attention | None (weak public signals, such as trending and view counts, excepted) | Personalization breaks observer equivalence by design. Not a CK failure to "fix"; a substrate built for a different purpose. |
| Platform verified-account badges | Weak authentication; weak intra-platform CK | Partial (within-platform observer equivalence and a shared protocol, both operator-contingent) | Single-operator authority; no auditable signing policy; CK contingent on platform governance. |

C2PA and device attestation meet authentication criteria[53], but verification is private. Two users

who agree about an asset's provenance need not know that others have seen it. Public blockchains and certificate transparency logs satisfy all four common-knowledge criteria formally in the narrow domains they cover[47;48], though most users rely on intermediaries. When nearly everyone verifies by delegation, to wallets, browsers, and monitors, observer equivalence collapses back into testimony from trusted verifiers, and the substrate quietly reconstitutes the authority problem it was meant to route around. They show that the record layer is architecturally achievable. What they lack is the attention layer: logs are consulted by browsers and auditors, not by the public at large.

Recommender feeds produce coincident attention and fail every common-knowledge criterion: personalized state is neither publicly verifiable nor observer-equivalent. This is their design. Proposals that treat feed re-ranking as the route back to shared reality are category errors. Tuning rankings changes what receives attention, while personalization and opaque rules still prevent a shared public state that observers can verify in common.

Mass-synchronous channels have also not vanished, and the criteria apply to the survivors. A statutory emergency-alert system is, on this framework's terms, a state-operated common-knowledge machine: a knowable protocol, near-simultaneous delivery to most active handsets, an audience whose scale the sender knows—and no publicly verifiable record, which is why an alert generates common knowledge in the moment but leaves nothing a dispersed public can re-inspect. Outlet push notifications reach tens of millions at once, under protocols that are private and revocable. In several media systems, high-reach national broadcasting persists. Cross-national differences in surviving broadcast reach are a natural observational testbed for the differential prediction below[32].

The web-of-trust history shows the criteria at work. PGP's web of trust tried to authenticate public keys without central authorities, yet verdicts depended on each participant's private keyring, so two users could reach different conclusions about the same key. Certificate transparency attacked a similar authentication problem by adding an append-only, publicly verifiable log that monitors read identically. The contrast suggests why better usability alone would not have solved web-of-trust key authentication, and why public records are necessary though still insufficient for broader shared-reality infrastructure.

A worked case pulls the layers apart. Take an election authority certifying a result. Authentication is easy: the certificate is signed and its provenance verifiable. A CK-capable public record is within reach: publish the result to an append-only, publicly verifiable log under a knowable protocol. Two problems remain. First, legible attention: the result approaches common knowledge only if a bounded public is known to be attending to the log, not merely able to, and that is the function broadcast served and no ledger supplies. Second, interpretation: the log can make it commonly inspectable that the authority certified these numbers, while whether the certification is accepted as legitimate is carried by institutions, not infrastructure. Common knowledge, provenance, attention, and interpretation come apart where the stakes are highest.

# Limits, governance, and outlook

The framework has three limits. First, common knowledge about artifacts is not common knowledge about interpretation. A public ledger could make it common knowledge that a photograph was recorded at a specific time with a specific device. It cannot settle the photograph's meaning. In the broadcast era that interpretive work belonged to journalists and experts, using practices of testimony, trust and deference long studied in social epistemology[56–62]. Many post-truth disputes

already involve common knowledge that an institution issued a ruling or report. The dispute is over authority, which no log supplies.

Second, every infrastructure concentrates trust in different places. Origin attestation concentrates it in hardware manufacturers, certificate authorities and standards bodies. Permissionless ledgers concentrate it least by design, since no operator is in charge, though validator concentration, client-software monocultures, and reliance on intermediaries reintroduce concentration in practice. Certificate transparency leaves certificate authorities as issuers, relies on log operators and monitors to expose misissuance, and depends on browser vendors to enforce CT policy by rejecting or distrusting certificates that fail logging requirements. Concerns about concentrated communicative power predate the current era[63;64]. The criteria do not resolve who audits the concentrators, how key compromises are remedied, or what nation-state compulsion looks like.

Third, agenda-setting remains a governance problem. Infrastructure can concentrate attention onto a topic set, but platform ranking, editorial judgment, political pressure and popularity dynamics decide which topics count as public issues. A substrate engineered to make a single agenda fall out for everyone would be a manufactured-consensus instrument. The safeguard is publicity: the same properties that make a substrate usable as a common signal also make capture easier to see. An append-only record and a knowable protocol leave power in place while making many abuses harder to hide. The counterweight is navigable pluralism: communities with different interpretations can coordinate on common knowledge of believable public facts, even where they differ on meaning. One would recognise it where communities that disagree sharply on meaning nonetheless do not dispute what the public record says. Election results, public-health surveillance, court rulings, and war-crimes documentation are the natural domains, and their evidentiary standards and stakes differ sharply. The same visibility has a dark side: infrastructure that makes attention legible can enable intimidation, surveillance, pile-ons, and strategic agenda capture, so public legibility is not democratically benign by default. This is narrower than Rawlsian overlapping consensus or Habermasian public reason, and closer to recent arguments that the answer to post-truth politics is not restored consensus[7] but institutions that sustain knowledge under disagreement[6]. Infrastructure can supply informational preconditions for coordination; conventions, legitimacy and motivation do the rest[65;66].

Two research tasks follow. The first is measurement. Experimental work shows that people condition risky coordination on whether information is publicly announced rather than merely shared privately[67]; Chwe's broadcast and ritual cases supply the field analogue[12]. From these the account makes a differential prediction: coordination on a public matter tracks whether a claim has been publicly established as common knowledge—not mere reach (exposure), not private estimates of the majority (perceived-majority belief), not within-network agreement (selective exposure). The sharpest test is against perceived-majority belief, since a public announcement tends to raise both at once. The design therefore holds exposure and perceived majority belief fixed while varying public establishment: two groups receive the same claim and the same estimate of how many others accept it, while only one receives it through a public signal whose audience is itself visible. This is the emperor's-new-clothes case, in which everyone privately knows but no one moves until it is said aloud. Because a visible audience is the sort of cue that inflates first-order estimates of the majority, equality cannot be assumed by instruction: post-treatment manipulation checks on perceived-majority belief are required, with analyses conditioned on them. The increment over the existing experiment[67] is that step, holding perceived-majority belief fixed while varying public establishment; the discontinuity being tested, that almost-common knowledge behaves

unlike common knowledge, is the one Rubinstein's electronic-mail game isolates[68]. A behavioral objection is that acceptance of publicly established facts is itself identity-contingent: motivated, identity-protective cognition can rise with sophistication rather than fall[69]. Our answer is a second prediction rather than a rebuttal. Because navigable pluralism targets coordination rather than belief revision, public establishment should move coordination even where professed belief moves little; belief-layer interventions such as accuracy prompts and inoculation are complements, not rivals, of the substrate. How accurately people infer the size and composition of an audience, and when pluralistic ignorance breaks, are the mechanisms on which such designs turn. Candidate measures include survey probes of iterated belief, behavioral inference from coordination on observer-equivalent substrates, and lab manipulations of the criteria. Collective attention also decays on characteristic timescales[70;71], bounding how long any signal can sustain mass attention.

The second task is design. Tampering is not the central threat on an append-only public ledger; a faithfully recorded lie is a false input, which belongs to the believability problem. The harder problem is coupling trustworthy inputs and public records to a legible attending public, while governing permissioned substrates and resisting compulsion. Here the framework runs into a tension: collective attention is scarce and its cycles are shortening, we have argued infrastructure should not choose topics, and yet something must allocate attention to the record. Candidate mechanisms exist, each with a failure mode. Scheduled civic moments would ritualize review of the public record; captured, they become state ritual in Chwe's sense. Statutory alert channels are already common-knowledge machines, and overuse or compulsion would spend them. Platform surfaces with verifiable counts would elevate trending pages to an evidentiary standard, and they remain manipulable until attention measurement is itself attested. Prompt witnessing by many independent observers, from the believability discussion, is a fourth. Which of these can concentrate attention without becoming an agenda-setter is the governance question this framework sharpens but cannot settle. None of this is solved by existing architectures. Much of it is tractable engineering, but the institutional questions are partly political and philosophical: who defines the relevant public, who controls the protocol, whose evidence counts, and how visibility is balanced against privacy. The most plausible near-term step relies on norms as much as engineering: authenticated capture and provenance may become expected for high-stakes public evidence, while institutional and journalistic systems attach those proofs to public, observer-equivalent records and visible attention cues.

Broadcast was not our epistemic baseline as a species; it was a remarkable, productive interlude. We can achieve infrastructure that produces enough shared epistemic reality for coordination without a return to broadcast's conditions.

# Acknowledgments

# References


[1] W. Joseph Campbell. *Getting It Wrong: Ten of the Greatest Misreported Stories in American Journalism*. University of California Press, Berkeley, 2010. ISBN 978-0-520-25566-1.

[2] Don Oberdorfer. *Tet! The Turning Point in the Vietnam War*. Johns Hopkins University Press, Baltimore, 2001. ISBN 978-0-8018-6703-3.

[3] Jayson Harsin. *Post-Truth and Critical Communication Studies*. Oxford University Press, December 2018. ISBN 978-0-19-022861-3. doi: 10.1093/acrefore/9780190228613.013.757.

[4] W. Lance Bennett and Steven Livingston. A Brief History of the Disinformation Age: Information Wars and the Decline of Institutional Authority. In W. Lance Bennett and Steven Livingston, editors, *The Disinformation Age*, pages 3–40. Cambridge University Press, 1 edition, October 2020. ISBN 978-1-108-91462-8 978-1-108-84305-8 978-1-108-82378-4. doi: 10.1017/9781108914628.001.

[5] Cass R. Sunstein. *#Republic: Divided Democracy in the Age of Social Media*. Princeton University Press, April 2017. ISBN 978-1-4008-8471-1. doi: 10.1515/9781400884711.

[6] Jonathan Rauch. *The Constitution of Knowledge: A Defense of Truth*. The Brookings Institution, 1 edition, 2021. ISBN 978-0-8157-5037-6 978-0-8157-3886-2. doi: 10.5040/9780815750376.

[7] Michael Hannon. The Politics of Post-Truth. *Critical Review*, 35(1-2):40–62, April 2023. ISSN 0891-3811, 1933-8007. doi: 10.1080/08913811.2023.2194109.

[8] Robert J. Aumann. Agreeing to disagree. *The Annals of Statistics*, 4(6):1236–1239, November 1976. ISSN 0090-5364. doi: 10.1214/aos/1176343654. URL https://projecteuclid.org/journals/annals-of-statistics/volume-4/issue-6/Agreeing-to-Disagree/10.1214/aos/1176343654.full.

[9] Joseph Y. Halpern and Yoram Moses. Knowledge and common knowledge in a distributed environment. *Journal of the ACM*, 37(3):549–587, July 1990. ISSN 0004-5411. doi: 10.1145/79147.79161. URL https://dl.acm.org/doi/10.1145/79147.79161.

[10] Dov Monderer and Dov Samet. Approximating common knowledge with common beliefs. *Games and Economic Behavior*, 1(2):170–190, June 1989. ISSN 0899-8256. doi: 10.1016/0899-8256(89)90017-1.

[11] Robert J. Aumann. Correlated equilibrium as an expression of Bayesian rationality. *Econometrica*, 55(1):1–18, January 1987. ISSN 0012-9682. doi: 10.2307/1911154.

[12] Michael Suk-Young Chwe. *Rational Ritual: Culture, Coordination, and Common Knowledge*. Princeton University Press, Princeton, NJ, 2001. ISBN 978-0-691-11471-2.

[13] Paul Starr. *The Creation of the Media: Political Origins of Modern Communications*. Basic Books, 2004. ISBN 978-0-465-08193-6.

[14] Michele Hilmes. *Radio Voices: American Broadcasting, 1922–1952*. University of Minnesota Press, Minneapolis, 1997. ISBN 978-0-8166-2621-4.

[15] Federal Communications Commission. Sixth Report and Order. Technical Report 41 F.C.C. 148, Federal Communications Commission, Washington, DC, April 1952. 17 Fed. Reg. 3905 (Apr. 14, 1952). Ended the television licensing freeze, allocated UHF channels, and established the Table of Assignments that structured U.S. VHF broadcasting.

[16] Markus Prior. *Post-Broadcast Democracy: How Media Choice Increases Inequality in Political Involvement and Polarizes Elections*. Cambridge University Press, 1 edition, April 2007. ISBN 978-0-521-85872-4 978-0-521-67533-8 978-1-139-87842-5. doi: 10.1017/CBO9781139878425.

[17] Jonathan M. Ladd. *Why Americans Hate the Media and How It Matters*. Princeton University Press, January 2012. ISBN 978-1-4008-4035-9. doi: 10.1515/9781400840359.

[18] David K. Lewis. *Convention: A Philosophical Study*. Harvard University Press, Cambridge, MA, 1969. ISBN 978-0-674-17025-3.

[19] Henry John Farrell and Bruce Schneier. Common-Knowledge Attacks on Democracy. *SSRN Electronic Journal*, 2018. ISSN 1556-5068. doi: 10.2139/ssrn.3273111.

[20] William J. Mitchell. *The Reconfigured Eye: Visual Truth in the Post-Photographic Era*. MIT Press, 1992. doi: 10.7551/mitpress/5765.001.0001.

[21] Kendall L. Walton. Transparent Pictures: On the Nature of Photographic Realism. *Critical Inquiry*, 11(2):246–277, December 1984. ISSN 0093-1896, 1539-7858. doi: 10.1086/448287.

[22] Neil Postman. *Amusing Ourselves to Death: Public Discourse in the Age of Show Business*. Penguin, 1985. ISBN 978-0-14-303653-1.

[23] Regina Rini. Deepfakes and the Epistemic Backstop. *Philosophers' Imprint*, 20(24):1–16, 2020. doi: 10.3998/phimp.1684.

[24] Robert Chesney and Danielle Keats Citron. Deep Fakes: A Looming Challenge for Privacy, Democracy, and National Security. *SSRN Electronic Journal*, 2018. ISSN 1556-5068. doi: 10.2139/ssrn.3213954.

[25] Don Fallis. The Epistemic Threat of Deepfakes. *Philosophy & Technology*, 34(4):623–643, December 2021. ISSN 2210-5433, 2210-5441. doi: 10.1007/s13347-020-00419-2.

[26] Paloma Atencia-Linares and Marc Artiga. Deepfakes, shallow epistemic graves: On the epistemic robustness of photography and videos in the era of deepfakes. *Synthese*, 200(6): 518, December 2022. ISSN 1573-0964. doi: 10.1007/s11229-022-04003-3.

[27] Joshua Habgood-Coote. Deepfakes and the epistemic apocalypse. *Synthese*, 201(3):103, March 2023. ISSN 1573-0964. doi: 10.1007/s11229-023-04097-3.

[28] Bruce A. Williams and Michael X. Delli Carpini. *After Broadcast News: Media Regimes, Democracy, and the New Information Environment*. Cambridge University Press, 1 edition, September 2011. ISBN 978-0-511-84636-6 978-1-107-01031-4 978-0-521-27983-3. doi: 10.1017/CBO9780511846366.

[29] Natalie Jomini Stroud. *Niche News*. Oxford University Press, May 2011. ISBN 978-0-19-975550-9. doi: 10.1093/acprof:oso/9780199755509.001.0001.

[30] James L. Baughman. The Fall and Rise of Partisan Journalism. *Center for Journalism Ethics, University of Wisconsin-Madison*, 2011. doi: 10.1353/book.101671.

[31] Jim A. Kuypers. *Partisan Journalism: A History of Media Bias in the United States*. Rowman & Littlefield, 2014. ISBN 978-1-62892-464-0. doi: 10.5040/9781501304873.

[32] Daniel C. Hallin and Paolo Mancini. *Comparing Media Systems: Three Models of Media and Politics*. Cambridge University Press, 1 edition, April 2004. ISBN 978-0-521-54308-8 978-0-521-83535-0 978-0-511-79086-7. doi: 10.1017/CBO9780511790867.

[33] Stefano DellaVigna and Ethan Kaplan. The Fox News Effect: Media Bias and Voting. *The Quarterly Journal of Economics*, 122(3):1187–1234, 2007. doi: 10.1162/qjec.122.3.1187.

[34] Gregory J. Martin and Ali Yurukoglu. Bias in Cable News: Persuasion and Polarization. *American Economic Review*, 107(9):2565–2599, September 2017. doi: 10.1257/aer.20160812.

[35] Ruben Durante, Paolo Pinotti, and Andrea Tesei. The Political Legacy of Entertainment TV. *American Economic Review*, 109(7):2497–2530, July 2019. doi: 10.1257/aer.20150958.

[36] Ruben Enikolopov, Maria Petrova, and Ekaterina Zhuravskaya. Media and Political Persuasion: Evidence from Russia. *American Economic Review*, 101(7):3253–3285, December 2011. doi: 10.1257/aer.101.7.3253.

[37] Javier Ureña-Carrion, Petri Leskinen, Jouni Tuominen, Charles Van Den Heuvel, Eero Hyvönen, and Mikko Kivelä. Communication now and then: Analyzing the Republic of Letters as a communication network. *Applied Network Science*, 7(1):26, December 2022. ISSN 2364-8228. doi: 10.1007/s41109-022-00463-1.

[38] Robert Darnton. An Early Information Society: News and the Media in Eighteenth-Century Paris. *The American Historical Review*, 105(1):1–35, February 2000. ISSN 1937-5239, 0002-8762. doi: 10.1086/ahr/105.1.1.

[39] Andrew Pettegree. *The Invention of News: How the World Came to Know About Itself*. Yale University Press, 2014. doi: 10.12987/yale/9780300179187.001.0001.

[40] Levi Boxell, Matthew Gentzkow, and Jesse M. Shapiro. Greater Internet use is not associated with faster growth in political polarization among US demographic groups. *Proceedings of the National Academy of Sciences*, 114(40):10612–10617, 2017. doi: 10.1073/pnas.1706588114.

[41] Levi Boxell, Matthew Gentzkow, and Jesse M. Shapiro. Cross-country trends in affective polarization. *The Review of Economics and Statistics*, 106(2):557–565, 2024. doi: 10.1162/rest_a_01160.

[42] Coalition for Content Provenance and Authenticity. Newly Formed Coalition for Content Provenance and Authenticity (C2PA) Sets Standards for Digital Content. Press release, February 2021. URL `https://c2pa.org/c2pa-founding-press-release/`. Quotation attributed to Jeffrey McGregor, CEO, Truepic.

[43] Philip Kitcher. The Division of Cognitive Labor. *Journal of Philosophy*, 87(1), 1990. doi: 10.2307/2185278.

[44] Venkatesh Bala and Sanjeev Goyal. Learning from Neighbours. *The Review of Economic Studies*, 65(3):595–621, 1998. ISSN 10942025. doi: 10.1111/1467-937X.00059.

[45] Kevin J. S. Zollman. The Communication Structure of Epistemic Communities. *Philosophy of Science*, 74(5):574–587, December 2007. ISSN 0031-8248, 1539-767X. doi: 10.1086/525605.

[46] C. Thi Nguyen. ECHO CHAMBERS AND EPISTEMIC BUBBLES. *Episteme*, 17(2):141–161, June 2020. ISSN 1742-3600, 1750-0117. doi: 10.1017/epi.2018.32.

[47] Ben Laurie. Certificate transparency. *Communications of the ACM*, 57(10):40–46, September 2014. ISSN 0001-0782. doi: 10.1145/2659897. URL https://dl.acm.org/doi/10.1145/2659897.

[48] Satoshi Nakamoto. Bitcoin: A Peer-to-Peer Electronic Cash System. White paper, 2008. URL https://bitcoin.org/bitcoin.pdf.

[49] C. Adams, P. Cain, D. Pinkas, and R. Zuccherato. Internet X.509 Public Key Infrastructure Time-Stamp Protocol (TSP). RFC RFC 3161, Internet Engineering Task Force (IETF), August 2001. URL https://www.rfc-editor.org/rfc/rfc3161.

[50] Stuart Haber and W. Scott Stornetta. How to time-stamp a digital document. *Journal of Cryptology*, 3(2):99–111, 1991. doi: 10.1007/BF00196791.

[51] F. A. Hayek. The Use of Knowledge in Society. *The American Economic Review*, 35(4):519–530, 1945. URL https://www.jstor.org/stable/1809376.

[52] Justin Wolfers and Eric Zitzewitz. Prediction Markets. *Journal of Economic Perspectives*, 18(2):107–126, 2004. doi: 10.1257/0895330041371321.

[53] Coalition for Content Provenance and Authenticity. C2PA Technical Specification, Version 2.1. Technical Specification, 2024. URL https://c2pa.org/specifications/specifications/2.1/specs/C2PA_Specification.html.

[54] J. Scott Brennen, Felix M. Simon, Philip N. Howard, and Rasmus Kleis Nielsen. Types, Sources, and Claims of COVID-19 Misinformation. Factsheet, Reuters Institute for the Study of Journalism, Oxford, April 2020. URL https://reutersinstitute.politics.ox.ac.uk/types-sources-and-claims-covid-19-misinformation.

[55] Britt Paris and Joan Donovan. Deepfakes and Cheap Fakes: The Manipulation of Audio and Visual Evidence. Report, Data & Society Research Institute, New York, September 2019. URL https://datasociety.net/library/deepfakes-and-cheap-fakes/.

[56] John Hardwig. The Role of Trust in Knowledge. *The Journal of Philosophy*, 88(12):693, December 1991. ISSN 0022362X. doi: 10.2307/2027007.

[57] C. A. J. Coady. *Testimony: A Philosophical Study*. Oxford University Press, 1992. ISBN 978-0-19-823551-4.

[58] Jennifer Lackey. *Learning from Words*. Oxford University Press, February 2008. ISBN 978-0-19-921916-2. doi: 10.1093/acprof:oso/9780199219162.001.0001.

[59] Alvin I. Goldman. *Knowledge in a Social World*. Oxford University PressOxford, 1 edition, January 1999. ISBN 978-0-19-823820-1 978-0-19-159752-7. doi: 10.1093/0198238207.001.0001.

[60] Alvin I. Goldman. Experts: Which Ones Should You Trust? *Philosophy and Phenomenological Research*, 63(1):85–110, July 2001. ISSN 0031-8205, 1933-1592. doi: 10.1111/j.1933-1592.2001.tb00093.x.

[61] Linda Zagzebski. *Epistemic Authority: A Theory of Trust, Authority, and Autonomy in Belief*. Oxford University Press, 2012. ISBN 978-0-19-928462-7. doi: 10.1093/acprof:oso/9780199284627.001.0001.

[62] Katherine Dormandy, editor. *Trust in Epistemology*. Routledge, 1 edition, September 2019. ISBN 978-1-351-26488-4. doi: 10.4324/9781351264884.

[63] Yochai Benkler. *The Wealth of Networks: How Social Production Transforms Markets and Freedom*. Yale University Press, 2006. ISBN 978-0-300-12577-1.

[64] Zeynep Tufekci. *Twitter and Tear Gas: The Power and Fragility of Networked Protest*. Yale University Press, 2017. ISBN 978-0-300-21512-0.

[65] Herbert Gintis. Social norms as choreography. *Politics, Philosophy & Economics*, 9(3):251–264, August 2010. ISSN 1470-594X. doi: 10.1177/1470594X09345474.

[66] Samuel Bowles and Herbert Gintis. *A Cooperative Species: Human Reciprocity and Its Evolution*. Princeton University Press, Princeton, NJ, 2011. ISBN 978-0-691-15125-0.

[67] Kyle A. Thomas, Peter DeScioli, Omar Sultan Haque, and Steven Pinker. The Psychology of Coordination and Common Knowledge. *Journal of Personality and Social Psychology*, 107 (4):657–676, 2014. doi: 10.1037/a0037037.

[68] Ariel Rubinstein. The electronic mail game: Strategic behavior under "almost common knowledge". *The American Economic Review*, 79(3):385–391, June 1989. ISSN 0002-8282. URL `https://www.jstor.org/stable/1806851`.

[69] Dan M. Kahan, Ellen Peters, Maggie Wittlin, Paul Slovic, Lisa Larrimore Ouellette, Donald Braman, and Gregory Mandel. The polarizing impact of science literacy and numeracy on perceived climate change risks. *Nature Climate Change*, 2(10):732–735, October 2012. ISSN 1758-678X, 1758-6798. doi: 10.1038/nclimate1547.

[70] Cristian Candia, C. Jara-Figueroa, Carlos Rodriguez-Sickert, Albert-László Barabási, and César A. Hidalgo. The universal decay of collective memory and attention. *Nature Human Behaviour*, 3(1):82–91, 2019. doi: 10.1038/s41562-018-0474-5.

[71] Philipp Lorenz-Spreen, Bjarke Mørch Mønsted, Philipp Hövel, and Sune Lehmann. Accelerating dynamics of collective attention. *Nature Communications*, 10(1):1759, 2019. doi: 10.1038/s41467-019-09311-w.