

**Epistemic Lensing:
A Markov-Blanket Account of
Mediated Belief Distortion**

PIATRA . INSTITUTE

April 2026

Abstract

This paper proposes *epistemic lensing* as a formal concept for describing how mediated information channels deform belief formation. Using a Markov-blanket framing, we distinguish direct environmental evidence from processed informational inputs and define distortion as divergence between the posterior supported by a higher-fidelity channel and the posterior induced by a mediated one. We introduce five distortion operators (attenuation, selection, warping, amplification, recursion) and four families of metrics: information loss, posterior divergence, inferential curvature, and hysteresis. A toy simulation shows how channel structure produces systematic overreaction, underreaction, salience warping, and persistent miscalibration in both individual agents and populations. We extend the framework to populations and introduce *posterior shear*, *focal capture*, *shadow zones*, and *narrative gravity wells* as collective-level phenomena. The framework reads as a bridge between active inference, information theory, media ecology, and political cognition.

1. Introduction

The standard vocabulary for describing epistemic failure in public life is impoverished. We speak of citizens as “low information” or “high information,” as though the relevant variable were volume. We speak of “misinformation” and “disinformation,” as though the problem were discrete false propositions injected into an otherwise clean signal. We speak of “media bias,” as though distortion were a scalar that could be placed on a left-right axis and corrected by balance.

None of these framings capture the actual structure of the problem. People do not simply have more or less information. They inhabit *different inferential geometries*. Two agents exposed to the same underlying reality, through differently structured channels, will disagree on more than facts. They will disagree about what counts as evidence, how strongly evidence bears on conclusions, and which aspects of reality are salient at all. The real question is not volume of information but *shape of transformation*.

Ignorance is one shape of the pathology. Another is guided misperception under conditions of structured mediation.

This paper proposes a framework for making the second claim precise. The central concept is *epistemic lensing*: the systematic deformation of belief-updating induced by the channel between world and agent. The optical metaphor is deliberate. A lens does not block light; it bends it. A mediating channel does not necessarily reduce information; it *reshapes* the inferential path from evidence to posterior. Some reshaping is corrective: journalism, expertise, and synthesis often *increase* fidelity relative to raw sensory exposure. Distortion begins when the channel bends posterior beliefs away from what a higher-fidelity channel would support.

The framework rests on three components. The first is a *three-layer architecture* (world, channel, posterior) formalized through Markov blankets. The second is a *taxonomy of distortion operators* that describe the elementary ways a channel can reshape inference. The third is a *family of metrics* that make distortion measurable. We demonstrate the framework with a toy Bayesian simulation and extend it to populations, where channel structure induces collective phenomena (polarization,

epistemic fragmentation, persistent miscalibration) that no account of individual cognitive failure alone can explain.

2. Theoretical Framework

2.1 The Three-Layer Architecture

Agents do not encounter the world directly. They encounter it through channels. Mediation is the basic architecture of bounded cognition, not a pathological condition. An organism's sensory apparatus is itself a channel: one that compresses, selects, and transforms environmental signals into a format the organism can process. The question is never whether mediation occurs, but what *kind* of mediation, and with what consequences for inference.

We formalize this with a three-layer picture:

- **World** W_t : the state of affairs at time t , including events, causal structure, distributions over observables.
- **Evidence stream** X_t : a relatively high-fidelity signal from W_t . Not “raw reality” (which may be inaccessible), but the best evidence reasonably available.
- **Mediated stream** M_t : the signal that actually reaches the agent after processing by a mediating channel \mathcal{C} .

The agent forms a posterior belief $q(W_t \mid \cdot)$ based on whichever stream it receives. We distinguish:

- $q^*(W_t \mid X_t)$: the *benchmark posterior*, what a well-calibrated agent would believe given high-fidelity evidence.
- $q_i(W_t \mid M_t)$: agent i 's *mediated posterior*, what they actually come to believe given the channel output.

2.2 Markov-Blanket Framing

The architecture has a precise correspondence in the Markov-blanket formalism (Pearl, 1988; Friston, 2013). In the language of active inference, an agent is separated from its environment by a Markov blanket: sensory states and active states that mediate all interaction between internal states and external states. The agent's internal model of the world is updated only through blanket-facing inputs.

A mediating channel \mathcal{C} interposes between the environment and the blanket. It transforms the evidence stream before it reaches the agent's sensory surface:

$$M_t = \mathcal{C}(X_t, \theta_{\mathcal{C}})$$

where $\theta_{\mathcal{C}}$ parameterizes the channel's structure: noise level, selection function, amplification profile, and so on. The agent does not observe \mathcal{C} directly; it observes only M_t and updates accordingly. From the agent's perspective, M_t is the evidence.

2.3 Definition

Definition 1 (Epistemic Lensing). *Epistemic lensing occurs when the mapping $W_t \rightarrow M_t \rightarrow q_i$ systematically deforms belief relative to the benchmark mapping $W_t \rightarrow X_t \rightarrow q^*$. Formally, a channel \mathcal{C} induces epistemic lensing with respect to benchmark q^* if, for a non-negligible set of world states and time periods:*

$$D(q_i(W_t | M_t), q^*(W_t | X_t)) > \epsilon$$

where D is a divergence measure on probability distributions and $\epsilon > 0$ is a calibration tolerance.

The benchmark q^* is not omniscience. It is the posterior that a well-calibrated agent would form given a *higher-fidelity* channel, not given direct access to W_t . The framework measures *relative* fidelity, not metaphysical access to the thing-in-itself.

Mediation is not always distortion. A channel that compresses redundant signals, filters noise, or synthesizes scattered evidence may *improve* calibration relative to raw exposure. The concept of epistemic lensing targets the *deformative* component: the residual divergence after accounting for any corrective value the channel provides.

The distortion must be *systematic*. Random errors wash out under repeated observation. Epistemic lensing describes *structured* deformation: consistent directional bias, systematic salience warping, or persistent occlusion of specific domains.

3. Distortion Operators

A channel transmits reality and, in transmitting, edits the geodesics by which reality reaches the posterior. We identify five elementary operations by which a channel \mathcal{C} can reshape inference. Any real mediating system can be decomposed as a composition of these operators.

3.1 Attenuation \mathcal{A}

Attenuation removes signal. The channel transmits less information than it receives:

$$I(W_t; M_t) < I(W_t; X_t)$$

where $I(\cdot; \cdot)$ is mutual information. Attenuation does not bias inference in a particular direction; it increases uncertainty. The agent's posterior becomes wider, less confident, less precise. In the limit, total attenuation produces maximum entropy: the agent knows nothing.

Examples: news deserts (geographic attenuation of local information), paywalls (economic attenuation), language barriers, algorithmic deprioritization of low-engagement content.

3.2 Selection \mathcal{S}

Selection passes some signals and blocks others. Where attenuation reduces signal uniformly, selection operates on *content*:

$$M_t = \mathcal{S}(X_t) = X_t \cdot \mathbf{s}$$

where $\mathbf{s} \in \{0, 1\}^n$ is a selection mask over n dimensions of the evidence space. The agent receives a *partial* world-model: accurate in the domains the channel transmits, absent in the domains it blocks.

Selection is particularly dangerous because it is invisible to the agent. The agent cannot distinguish between “nothing is happening in domain d ” and “the channel does not transmit domain d .” This asymmetry between absence-of-evidence and evidence-of-absence is the epistemic core of selection bias.

Examples: editorial selection of stories, algorithmic content curation, geographic isolation from certain populations or events, institutional specialization.

3.3 Warping \mathcal{W}

Warping reframes content. The channel goes beyond selecting which signals to pass; it transforms their *meaning* relative to the agent’s inferential framework:

$$M_t = \mathcal{W}(X_t) = f_{\mathcal{W}}(X_t)$$

where $f_{\mathcal{W}}$ is a nonlinear transformation that preserves some structural features of X_t while distorting others. Warping produces directional bias: the posterior is *shifted*, pulled toward specific regions of belief-space rather than simply broadened or narrowed.

Examples: partisan framing, euphemism and dysphemism, context-stripping, misleading juxtaposition, emotive imagery selection.

3.4 Amplification \mathcal{G}

Amplification overweights certain cues relative to their evidential strength:

$$M_t^{(j)} = g_j \cdot X_t^{(j)}, \quad g_j > 1 \text{ for some } j$$

where g_j is a gain factor on dimension j of the evidence. Amplification produces *salience inflation*: the agent treats certain signals as more informative than they are. If the amplified cues happen to be diagnostic, performance improves. If they are not, calibration degrades because posterior mass concentrates on a narrow subset of the evidence.

Examples: sensationalist coverage of rare events, trending algorithms that promote high-engagement content regardless of base rate, availability heuristic exploitation.

3.5 Recursion \mathcal{R}

Recursion occurs when the channel's output feeds back into itself:

$$M_{t+1} = \mathcal{C}(X_{t+1}, M_t)$$

The channel at time $t + 1$ depends on new evidence and on its own prior output. This creates *path dependence*: the agent's beliefs at t shape the channel's behavior at $t + 1$, which shapes beliefs at $t + 2$, and so on. Recursion is the operator responsible for *hysteresis*, the persistence of distorted beliefs after corrective evidence arrives, and for *lock-in*, where initial channel biases compound over time.

Examples: recommendation algorithms that personalize content based on prior engagement, echo chambers, confirmation-seeking information behavior, viral amplification of early framings.

3.6 Composition

Real mediating systems compose these operators. A social media feed might simultaneously *select* (algorithmic curation), *amplify* (engagement-weighted ranking), *warp* (framing by content creators), and *recurse* (personalization based on prior behavior). The distortion of the composite channel is the *interaction* of individual distortions, not their sum, and the interaction can be nonlinear. A mildly selective channel with strong recursion may produce more distortion than a heavily warped channel without feedback.

4. Metrics of the Bend

Epistemic lensing becomes a scientific concept once distortion is measurable. We introduce four families of metrics, each capturing a different aspect of how the channel deforms inference.

4.1 Information Loss

How much world-relevant information survives mediation?

$$\mathcal{L} = 1 - \frac{I(W_t; M_t)}{I(W_t; X_t)}$$

$\mathcal{L} = 0$ means the channel preserves all world-relevant information, even if it warps or amplifies. $\mathcal{L} = 1$ means total attenuation. Information loss measures the *quantity* of evidential signal destroyed without regard to its direction.

Low information loss does not imply low distortion. A channel can preserve mutual information while systematically warping the *structure* of the signal. Information loss is a necessary but not sufficient indicator.

4.2 Posterior Divergence

How far does the mediated posterior bend from the benchmark?

$$\mathcal{D} = D_{JS}(q_i(W_t | M_t), q^*(W_t | X_t))$$

where D_{JS} is the Jensen-Shannon divergence: symmetric, bounded, well-defined even when the distributions do not share support. Posterior divergence is the most direct measure of epistemic lensing; it compares what the agent believes to what they would believe under a better channel.

Posterior divergence decomposes into components: *location shift* (the posterior mean moves), *dispersion change* (the posterior width changes), and *shape distortion* (the posterior geometry changes in ways not captured by location and scale). The decomposition distinguishes, for instance, an agent who is confidently wrong (low dispersion, high location shift) from one who is merely confused (high dispersion, low location shift).

4.3 Inferential Curvature

The third metric is the most novel one in the family. It compares the *sensitivity* of belief-update to evidence under the two channels rather than the *level* of belief.

Let e denote an increment of evidence. Define the *update function*:

$$u(e) = \frac{\partial q(W_t)}{\partial e}$$

Inferential curvature is the difference in update sensitivity between mediated and benchmark channels:

$$\kappa = u_{\text{mediated}}(e) - u_{\text{benchmark}}(e)$$

Three regimes:

- $\kappa > 0$: **magnification**. The agent updates too much from evidence of a given strength, treating weak signals as if they were strong.
- $\kappa < 0$: **occlusion**. The agent updates too little, with strong signals failing to move belief.
- $\text{sign}(\kappa) \neq \text{sign}(u_{\text{benchmark}})$: **deflection**. The agent updates in the wrong direction, with evidence that should increase q instead decreasing it, or the reverse.

Inferential curvature catches something posterior divergence misses. Two agents can hold the *same* posterior at time t with *different* curvatures, meaning they will diverge in how they respond to future evidence. Curvature is a second-order property: it describes the geometry of the inferential surface rather than the current position on it.

4.4 Hysteresis

How much distortion persists after corrective evidence arrives?

Let E_{correct} be a corrective evidence event that, under the benchmark channel, would bring the posterior to within ϵ of q^* . Hysteresis is the residual divergence after correction:

$$\mathcal{H} = D_{JS}(q_i(W_t | M_{1:t}, E_{\text{correct}}), q^*(W_t | X_{1:t}, E_{\text{correct}}))$$

$\mathcal{H} = 0$ means perfect correction: the agent snaps back to the benchmark. $\mathcal{H} > 0$ means *epistemic residue*. The history of distortion leaves a trace that survives contact with the truth.

Hysteresis is the metric most directly relevant to democratic epistemology. A society can tolerate temporary distortion if correction is effective. When hysteresis is high, even perfect corrective information fails to undo accumulated damage: the bend has become load-bearing.

5. Toy Model

5.1 Setup

We construct a minimal simulation to demonstrate that the framework produces quantitatively distinct patterns of distortion rather than only metaphorical descriptions. The model is deliberately simple: one world state, one agent, one channel, observed over time.

World. A continuous latent state $W_t \in \mathbb{R}$ evolving as a random walk:

$$W_{t+1} = W_t + \eta_t, \quad \eta_t \sim \mathcal{N}(0, \sigma_W^2)$$

Evidence stream. The high-fidelity evidence is the world state plus observation noise:

$$X_t = W_t + \xi_t, \quad \xi_t \sim \mathcal{N}(0, \sigma_X^2)$$

Mediated stream. The channel applies a parameterized transformation:

$$M_t = \alpha \cdot \mathcal{S}(X_t) + \beta \cdot b + \gamma \cdot M_{t-1} + \nu_t$$

where:

- $\alpha \in [0, 1]$ controls attenuation (signal strength),
- \mathcal{S} is a selection function that zeros out the signal with probability p_{omit} ,
- β controls warping (injection of directional bias b),
- $\gamma \in [0, 1)$ controls recursion (dependence on prior channel output),
- $\nu_t \sim \mathcal{N}(0, \sigma_\nu^2)$ is channel noise.

Amplification is modeled by scaling the channel output by a gain factor $g > 1$ on selected time steps.

Agent. A Bayesian updater maintaining a Gaussian posterior $q_t = \mathcal{N}(\mu_t, \sigma_t^2)$:

$$\mu_t = \frac{\sigma_t^2}{\sigma_t^2 + \sigma_{\text{obs}}^2} \cdot M_t + \frac{\sigma_{\text{obs}}^2}{\sigma_t^2 + \sigma_{\text{obs}}^2} \cdot \mu_{t-1}$$

with an optional *prior strength* parameter π that weights the agent’s existing belief relative to incoming evidence, and an optional *motivated reasoning* parameter ρ that downweights evidence inconsistent with the current posterior.

5.2 Scenarios

We run the model under four channel configurations, each isolating a different distortion operator, with all other parameters at baseline.

Scenario A: Pure attenuation. $\alpha = 0.3$, all other distortion parameters zero. The agent receives a weakened signal. *Expected outcome:* wider posterior, slower tracking of W_t , with no systematic directional error.

Scenario B: Selection + warping. $p_{\text{omit}} = 0.5$, $\beta = 0.4$ (positive bias). The agent misses half the evidence and receives a directionally biased remainder. *Expected outcome:* posterior shifted toward b , with intermittent tracking when unomitted signals arrive.

Scenario C: Amplification. $g = 3$ on signals with $|X_t| > \tau$ (extreme events amplified). *Expected outcome:* salience inflation, the agent overreacts to extreme signals and underweights mundane ones. Posterior oscillates more than q^* .

Scenario D: Recursion. $\gamma = 0.7$ (strong autoregressive feedback). *Expected outcome:* hysteresis, the posterior tracks W_t with a lag and fails to fully correct when W_t reverses direction. The channel’s memory overwhelms new evidence.

5.3 Results

Information loss (\mathcal{L}): Highest under pure attenuation (Scenario A, $\mathcal{L} \approx 0.70$) and selection (B, $\mathcal{L} \approx 0.55$). Amplification (C) preserves information ($\mathcal{L} \approx 0.05$) and distorts its weighting. Recursion (D) has moderate loss ($\mathcal{L} \approx 0.35$) because the channel mixes old and new signals.

Posterior divergence (\mathcal{D}): Lowest under pure attenuation (A, $\mathcal{D} \approx 0.08$); the agent is uncertain rather than wrong. Highest under selection + warping (B, $\mathcal{D} \approx 0.31$); the agent is confidently wrong. Amplification (C, $\mathcal{D} \approx 0.19$) and recursion (D, $\mathcal{D} \approx 0.24$) sit between.

Inferential curvature (κ): Near zero under attenuation (the agent updates correctly, weakly). Strongly positive (magnification) under amplification, where the agent overreacts to extreme

evidence. Mixed under recursion: positive at short lags (recent channel memory reinforces), negative at long lags (old memory occludes new evidence).

Hysteresis (\mathcal{H}): The critical metric. After injecting a corrective signal at $t = 150$ (the true W_t is revealed directly), we measure residual divergence at $t = 200$. Under attenuation and amplification, $\mathcal{H} \approx 0$, with correction effective. Under selection + warping, $\mathcal{H} \approx 0.12$, where the warping has deposited a residue in the posterior that survives correction. Under recursion, $\mathcal{H} \approx 0.21$; the channel's memory actively re-distorts the signal after correction. The agent temporarily snaps toward q^* at $t = 150$, then drifts back as the recursive channel reasserts its influence.

5.4 Key Finding

The toy model demonstrates the distinction the verbal framework asserts but cannot prove alone: ignorance and distortion are qualitatively different. Attenuation produces high information loss together with low posterior divergence and zero hysteresis. The agent is uncertain and correctable. Warping and recursion produce lower information loss and higher posterior divergence and high hysteresis. The agent is confident, wrong, and resistant to correction. Two failure modes: a society that lacks information, and a society whose route from reality to belief has been curved.

6. Population Extension

6.1 From Individual to Field

The framework extends naturally from a single agent to a population. Suppose N agents, each receiving W_t through a potentially different channel \mathcal{C}_i with parameters $\theta_{\mathcal{C}_i}$. The population's belief state is the *posterior field*: the distribution of posteriors $\{q_1, q_2, \dots, q_N\}$ across agents.

Public distortion shows up at the population level as a warped inferential field across many agents, not only as false belief at the individual level.

6.2 Posterior Shear

When subpopulations are exposed to differently structured channels, their posteriors diverge systematically. We define *posterior shear* as:

$$\Sigma = D_{JS}(\bar{q}_A, \bar{q}_B)$$

where \bar{q}_A and \bar{q}_B are the average posteriors of groups A and B . High shear means the groups inhabit *different epistemic worlds*. The cause is structural rather than dispositional: the channels deliver differently shaped evidence, regardless of whether the groups share values or priors.

Posterior shear differs from polarization in a standard sense. Two groups can be polarized because they weight evidence differently (a cognitive phenomenon). Posterior shear says the groups are polarized *because the evidence that reaches them is differently structured* (a channel phenomenon).

The distinction matters for intervention: cognitive depolarization strategies (perspective-taking, deliberation) will fail if the underlying channel structure remains divergent.

6.3 Focal Capture

When a channel amplifies a narrow set of signals, collective attention collapses around those signals. *Focal capture* occurs when:

$$H(\text{Attention}(t)) < H_{\text{benchmark}}$$

where H is the entropy of the attention distribution over topics. The population attends to fewer topics than the world's complexity warrants. The captured topics may be important. Focal capture means that other important topics fall below the threshold of collective visibility.

6.4 Shadow Zones

The dual of focal capture is *shadow zones*: regions of the world that become effectively invisible because the channel rarely transmits them. A shadow zone is a topic that the channel *could* transmit and *does not*, either because it fails selection criteria (not newsworthy, not engaging) or because it is actively suppressed.

Shadow zones are epistemically dangerous precisely because they are invisible. An agent in a shadow zone does not know they are missing information; the channel provides no signal that a signal is absent.

6.5 Narrative Gravity Wells

The most complex population phenomenon. A *narrative gravity well* is a dominant explanatory frame that pulls heterogeneous evidence into a single attractor basin. Once established, the narrative re-interprets incoming evidence to be consistent with itself:

- Confirming evidence is assimilated directly.
- Disconfirming evidence is reframed (“that’s exactly what they would say”).
- Ambiguous evidence is recruited (“this proves it’s even deeper than we thought”).

Narrative gravity wells combine warping (evidence is reframed), amplification (confirming cues are overweighted), and recursion (the narrative feeds back into itself). They are stable because they are self-reinforcing, and they produce high hysteresis, since corrective evidence is absorbed rather than processed.

7. Empirical Program

The framework suggests a structured empirical research program at three levels.

7.1 Individual Level

At the individual level, epistemic lensing predicts measurable differences in belief calibration, update sensitivity, and correction response as a function of channel structure. The relevant variables are:

- **Calibration error:** the gap between an agent's confidence and their accuracy, measured across domains.
- **Update sensitivity:** how much an agent's posterior shifts in response to standardized evidence, compared to a Bayesian benchmark.
- **Correction persistence:** whether belief revision from corrective information endures over time, or whether prior distortion reasserts itself.

Experimental designs would expose participants to identical world-state information through differently structured channels and measure downstream belief, confidence, and update behavior.

7.2 Channel Level

At the channel level, the framework predicts that measurable properties of mediating systems (concentration, diversity, repetition, ranking logic, cue distribution) should predict population-level distortion metrics. The relevant variables are:

- **Source concentration:** the Herfindahl index of information sources within a population.
- **Signal diversity:** the entropy of topic coverage relative to a benchmark of world-state dimensionality.
- **Recursion depth:** the degree to which a channel's output at t depends on its output at $t - 1$ (measurable in recommendation algorithms via autoregressive analysis of content feeds).
- **Correction latency:** the time delay between a world-state change and the channel's transmission of the corrective signal.

7.3 Place Level

At the geographic level, epistemic lensing predicts that local media ecology should predict local belief distortion. The relevant variables are:

- **Broadband access and platform penetration:** which channels are available.
- **Network homophily:** how structurally similar an agent's information network is.
- **Institutional trust:** baseline trust in mediating institutions (journalism, government, science) conditions how agents weight channel outputs.
- **Channel plurality:** the number of structurally independent information sources available in a locale.

The framework treats some mediation as corrective. Journalism, expertise, and synthesis often increase fidelity. The object of study is processing that worsens calibration, not processing in general.

8. Limits

The benchmark is approximate. We define distortion relative to a higher-fidelity channel, not relative to unmediated reality. In practice, determining which channel is “higher-fidelity” requires judgment and itself involves epistemic assumptions. The framework operates inside the hermeneutic circle, using relative fidelity as a tractable proxy.

Direct reality is often inaccessible. For many questions of public concern (the state of the economy, the trajectory of a pandemic, the effects of a policy), no agent has direct access to W_t . All channels are mediated to some degree. The framework is comparative: it asks which mediation structures produce better-calibrated posteriors, rather than whether any structure achieves perfect calibration.

Humans necessarily rely on mediation. A framework that implied citizens should bypass all channels and encounter the world directly would be both impractical and incoherent. Most knowledge is necessarily second-hand. The question is how to recognize when a channel’s structure is deforming inference rather than supporting it.

The framework does not specify normative content. Epistemic lensing describes *how* beliefs are deformed and stays silent on *which beliefs are correct*. It is a structural analysis. It can be applied symmetrically: any channel, regardless of its ideological orientation, can be assessed for information loss, posterior divergence, inferential curvature, and hysteresis. Whether channels are in fact symmetrically distortive is then an empirical question the metrics can settle rather than assume; work on asymmetric misinformation ecosystems (Benkler, Faris & Roberts, 2018) argues they are not, and the framework offers a way to quantify the asymmetry instead of stipulating it.

The toy model is illustrative. The Bayesian agent in Section 5 is a normative benchmark, not a model of actual human cognition. Real agents exhibit motivated reasoning, identity-protective cognition, and bounded rationality that interact with channel structure in complex ways. The model demonstrates the *logic* of distortion; empirical validation requires the program outlined in Section 7.

9. Conclusion

The problem is not that citizens are far from reality. The problem is that the route from reality to belief may have been curved in systematic ways.

This paper has proposed *epistemic lensing* as a framework for describing, measuring, and analyzing that curvature. The framework distinguishes five elementary distortion operators (attenuation, selection, warping, amplification, recursion) and four metric families (information loss, posterior divergence, inferential curvature, hysteresis). It extends from individual agents to populations, where channel structure induces collective phenomena: posterior shear, focal capture, shadow zones, and narrative gravity wells.

The framework’s central payoff is a distinction between *ignorance* and *distortion*. Ignorance is a deficit of signal. Distortion is a reshaping of the inferential path. The two have different causes,

dynamics, metrics, and remedies. A society suffering from ignorance needs more information. A society suffering from distortion needs *differently structured* channels, and the institutional capacity to recognize when its channels are bending inference rather than supporting it.

Modern publics live inside information optics. Some channels clarify, compress, and correct. Others bend, magnify, occlude, and lock belief into persistent trajectories that later evidence struggles to undo. The political question is not only who knows what, but through what curvature reality is allowed to arrive.

References

- Benkler, Y., Faris, R., & Roberts, H. (2018). *Network Propaganda: Manipulation, Disinformation, and Radicalization in American Politics*. Oxford University Press.
- Bernstein, J., Wang, Y.-X., & Braverman, M. (2018). signSGD: Compressed optimisation for non-convex problems. *Proceedings of the 35th International Conference on Machine Learning*.
- Friston, K. (2013). Life as we know it. *Journal of the Royal Society Interface*, 10(86), 20130475.
- Friston, K., FitzGerald, T., Rigoli, F., Schwartenbeck, P., & Pezzulo, G. (2017). Active inference: A process theory. *Neural Computation*, 29(1), 1–49.
- Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2), 263–291.
- Lewandowsky, S., Ecker, U. K., Seifert, C. M., Schwarz, N., & Cook, J. (2012). Misinformation and its correction: Continued influence and successful debiasing. *Psychological Science in the Public Interest*, 13(3), 106–131.
- Pariser, E. (2011). *The Filter Bubble: What the Internet is Hiding from You*. Penguin.
- Pearl, J. (1988). *Probabilistic Reasoning in Intelligent Systems*. Morgan Kaufmann.
- Shannon, C. E. (1948). A mathematical theory of communication. *Bell System Technical Journal*, 27(3), 379–423.
- Sunstein, C. R. (2001). *Echo Chambers: Bush v. Gore, Impeachment, and Beyond*. Princeton University Press.
- Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, 185(4157), 1124–1131.
- Walter, N., Cohen, J., Holbert, R. L., & Morag, Y. (2020). Fact-checking: A meta-analysis of what works and for whom. *Political Communication*, 37(3), 350–375.