

Article

Governing Adaptive News Curation: Sequential Optimization, Cumulative Exposure Allocation, and Societal Accountability

Dan Valeriu Voinea 

Department of Arts and Media, University of Craiova, 200585 Craiova, Romania; dan.voinea@gmail.com

Abstract

Digital gatekeeping is increasingly performed by adaptive systems that rank, sequence, package, moderate, and sometimes generate news across repeated interactions. This conceptual analysis asks two questions: how should researchers analyze these systems, and who should be responsible for their effects over time? It brings together gatekeeping theory, research on recommender systems and performative prediction, and scholarship on algorithmic accountability, and it makes three contributions. First, it describes curation through six decision functions (source eligibility, agenda and candidate formation, content generation, packaging, exposure allocation, and moderation) connected by a feedback and optimization layer. Agentic curation is treated as a high-autonomy form of adaptive curation, defined by planning ability, authority to act, reach across functions, and delay before human review, rather than as a separate technology class. Second, it defines *cumulative exposure allocation* as the normalized distribution of weighted exposure across sources, topics, and population groups over time, and it proposes measures of concentration, breadth, repetition, persistence, and disparity. Third, it links each curation function to the actors who control it, the bodies that oversee it, the evidence they need, the standards they apply, and the remedies they can impose. European Union and United States law illustrate the framework. The analysis shifts evaluation away from isolated outputs and short-term engagement toward patterns of visibility produced by a policy over time, and it proposes testable claims about traceability, concentration, disparity, and constrained multi-objective ranking.

Keywords: adaptive news curation; agentic curation; gatekeeping theory; sequential optimization; recommender systems; algorithmic accountability; cumulative exposure allocation; Digital Services Act; performative prediction



Academic Editors: Daniel McCarthy and Lawrence Ka-ki Ho

Received: 17 June 2026

Revised: 16 July 2026

Accepted: 20 July 2026

Published: 23 July 2026

Copyright: © 2026 by the author.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

1. Introduction

Gatekeeping research has long shown that journalism organizes public attention by selecting stories, setting agendas, and framing events ([Entman 1993](#); [McCombs and Shaw 1972](#); [White 1950](#)). Digital distribution has not removed this role. Instead, it has spread gatekeeping across newsrooms, platforms, recommender systems, interfaces, and, increasingly, generative systems. As news has become more dependent on platform infrastructure, much of that infrastructure has moved outside newsroom control ([Nieborg and Poell 2018](#); [van Dijck et al. 2018](#)). Established news organizations now compete for attention with creators, aggregators, and social-video platforms ([Newman 2025](#)).

One important part of this shift is *sequential optimization*. A recommender can treat curation as a process that unfolds over time: each recommendation produces feedback, and

that feedback shapes later decisions about what to show (Afsar et al. 2022; Chen et al. 2023; Sutton and Barto 2018). Reinforcement learning (RL) is one explicit technical form of this broader adaptive logic, but not every deployed news recommender uses RL. When curation adapts in this way, gatekeeping no longer depends only on individual editorial choices. It also depends on the system's objectives, feedback loops, and oversight arrangements.

A further shift is beginning to emerge. Some systems can generate, summarize, and package media while also showing the operational autonomy associated with the term *agentic* (Shavit et al. 2023). Peer-reviewed research does not yet document fully autonomous agents curating national news at scale. High-autonomy agentic curation is, therefore, treated here as an emerging direction rather than as the standard form of current deployment.

The main research question is as follows: *How should media and communication researchers understand adaptive news curation as a process that allocates exposure over time, and how should responsibility for its outcomes be distributed across the system's curation functions?* A related question is as follows: *What changes when a persistent, tool-using system receives authority across several functions and can act before a human reviews its decisions?*

The central claim is that adaptive curation shifts gatekeeping away from isolated choices about individual stories and toward the governance of visibility over time. Agentic systems intensify this shift because they can act across more functions, exercise more delegated authority, and operate for longer before review.

The argument makes three contributions. First, it presents a functional architecture that shows where important decisions enter adaptive curation. It also identifies the dimensions along which a system becomes more agentic (Sections 2 and 4). Second, it defines *cumulative exposure allocation* as a unit for analyzing visibility over time. It then develops a measurement model around that unit: normalized, group-sensitive, and counterfactual measures of the distributions produced by a deployed policy (Section 6). Third, it presents an accountability structure that connects each curation function and each long-term outcome to the actors in control, the relevant oversight bodies, the evidence those bodies need, the standards they should apply, and the remedies they can use (Sections 7 and 8). Together, these contributions show where values enter the system, what patterns build up over time, and who should answer for those patterns. Each component keeps one label throughout: the *functional architecture* names the six curation functions and their feedback layer, the *measurement model* names cumulative exposure allocation and its indicators, and the *accountability structure* names the actor–forum–evidence–standard–remedy assignments. *Framework* refers to the three together. The framework builds on Voinea (2025) and Wallace (2018); Section 2.4 states what it synthesizes, what it extends, and what it adds.

2. The Migration of Gatekeeping

2.1. Classical Gatekeeping: People as Information Filters

The idea of gatekeeping did not begin in journalism. Lewin (1947) introduced it in research on wartime food habits. White (1950) then adapted the idea to mass communication through a study of “Mr. Gates,” a wire editor who rejected about nine-tenths of the copy he received and openly described his choices as subjective. Galtung and Ruge (1965) later described news values as a way of ranking events, and Shoemaker and Vos (2009) developed gatekeeping into a multilevel theory in which selection is shaped by individuals, routines, organizations, institutions, and the wider social system.

Shoemaker and Vos define gatekeeping broadly as “the process of selecting, writing, editing, positioning, scheduling, repeating, and otherwise massaging information to become news.” This definition already points toward the present problem: in adaptive curation, positioning, scheduling, and repeating become especially powerful tools.

Two later developments are especially relevant. Barzilai-Nahon (2008) extended gatekeeping to networked environments, and Bruns (2018) argued that digital media increasingly involve *gatewatching*: distributed forms of curation that surface and prioritize material from abundant information streams. Both approaches show that control over visibility can move from one institution to another without losing its public importance.

Responsibility in legacy news organizations was never simple or fully centralized. Even so, their major decision points and lines of responsibility were usually easier to identify: professional norms guided editorial choices, and human decision-makers could be held legally, commercially, and socially responsible. Algorithmic and agentic curation make that chain of responsibility harder to reconstruct.

2.2. Algorithmic Gatekeeping: Recommenders and Measured Audiences

Digital distribution reduced some limits associated with print and broadcasting, but it also intensified competition for attention. As platforms began to carry far more content than any person could review, recommender systems took on a larger gatekeeping role, and many users moved from actively searching for information toward receiving continuously curated streams (Thurman 2011).

Research on algorithmic gatekeeping makes a consistent point: public visibility is now shaped by platform infrastructure, ranking systems, moderation, and audience metrics in addition to newsroom judgment. These systems define relevance, control what becomes visible, and use measured responses to guide later choices (Napoli 2014; Gillespie 2014, 2018; Bucher 2012, 2018; Seaver 2017; Just and Latzer 2017). Audience metrics linked the human and algorithmic phases even before reinforcement learning: web data changed what editors selected, promoted, and repeated (Tandoc and Thomas 2015; Tenor 2024). Adaptive recommenders automated and extended a logic that news organizations had already begun to use, often on infrastructure that publishers do not control (Meese and Hurcombe 2021).

This transition moved some gatekeeping power from editors judging civic importance to models estimating the likelihood of engagement. Helberger (2019) argued that click performance alone cannot show whether a news recommender serves democracy; the more important question is what democratic role the system is designed to play.

Research also shows that values such as diversity, broad coverage, and serendipity can be built into recommendation systems (Lu et al. 2020; Mattis et al. 2024; Vrijenhoek et al. 2021), though newsrooms still disagree about how such values should be defined and implemented (Bauer et al. 2024; van Es and Nguyen 2024). DeVito (2017), for example, showed that Facebook's News Feed encoded identifiable values into story selection. This literature establishes the basic shift: editorial authority is partly embedded in opaque infrastructure whose incentives may not match journalistic or democratic goals. Much of that work, however, concerns the scoring of individual items. Sequential optimization raises a further question: what happens when systems optimize patterns of exposure over time?

2.3. From Sequential Optimization to Agentic Curation: Configuration Dimensions

Two developments go beyond the conventional recommender model. First, generative systems now take part in media production: they can draft text, create images and audio, and assemble summaries, so gatekeeping extends beyond selecting existing media to shaping or producing new media. Second, some systems show increasing operational autonomy (Shavit et al. 2023): they can interpret context, plan several steps ahead, use tools, and pursue goals with limited human supervision.

These developments are often treated as one trend, but they can occur separately. A system may generate content without using reinforcement learning, use RL without

being agentic, or be agentic without generating content. Optimization method, productive capability, and operational autonomy, therefore, belong to different analytical levels. For that reason, curation systems are described here along five dimensions instead of being sorted into mutually exclusive types (Table 1). Supervised ranking systems, sequential recommenders, generative assistants based on large language models (LLMs), and agentic systems are examples of different profiles along these dimensions, not fixed stages in a single technical progression.

Table 1. Dimensions of adaptive curation systems. Familiar labels (supervised ranking, sequential recommender, generative assistant, agentic system) describe example profiles along these dimensions, not mutually exclusive system types or stages in a single progression.

Dimension	Lower-Autonomy End	Higher-Autonomy End	Why It Matters for Governance
Optimization horizon	Scores the likely response to one action	Optimizes a sequence of actions over a longer period (Chen et al. 2019; Zou et al. 2019)	Determines which delayed outcomes the system rewards
Productive authority	Selects existing items	Transforms, generates, or publishes content	Creates additional duties for sourcing, provenance, and correction
Planning and tool use	Runs a fixed scoring or generation step	Breaks goals into steps, plans, and uses tools	Allows the system to act beyond one model call
Functional scope	Operates within one curation function	Coordinates across several functions	Blurs technical and organizational boundaries of control
Authority to act and delay before review	Acts within preset limits and is reviewed promptly	Changes downstream systems or publishes before review	Determines how much can happen before a person can intervene

Here, *agentic* means that a system has been given operational autonomy; it does not imply moral agency, legal personhood, or independent accountability. A curation system can become more agentic in three separate ways. First, its capabilities may expand to include persistent, multi-step planning and tool use. Second, it may receive more delegated authority, such as permission to generate, publish, or alter downstream systems. Third, its governance may allow it to coordinate across several curation functions before external review.

These three layers can vary independently and should be reported separately rather than reduced to a yes-or-no label. Sequential optimization alone does not make a system agentic: a recommender may update a policy and optimize over time without planning, using tools, generating content, or crossing organizational boundaries. As systems gain more agentic capacity, the editorially important choices move upstream into objective design and downstream into oversight, and become less visible in any single act of selection.

Greater agentic capacity changes the functional architecture in two main ways. First, it can remove or bypass checkpoints. In conventional systems, the curation functions described in Section 4 are institutionally separated, with organizational review between them. If one goal-directed system both generates content and learns from feedback which of its own outputs to produce and promote next, production and selection become part of the same optimization loop. Second, agentic systems can act for longer before anyone reviews them. A system might assemble a summary, test several versions, reallocate exposure, and revise the content before an external reviewer intervenes. Both changes increase the distance between the moment when a value could have been specified and the moment when its consequences become visible.

2.4. How This Framework Extends Earlier Work

The contribution operates at three levels: synthesis, extension, and original conceptual work. As *synthesis*, the argument brings together gatekeeping theory, research on recommender systems and performative prediction, and scholarship on algorithmic ac-

countability, traditions that usually develop separately (Section 3). Earlier research in those traditions has already shown that optimization goals affect visibility and that feedback changes recommendations (DeVito 2017; Just and Latzer 2017; Napoli 2014); such general observations are synthesized here, not claimed as new.

As *extension*, the framework extends Shoemaker and Vos's (2009) levels of influence by showing how those levels now appear in data, objectives, interfaces, and oversight arrangements. It complements Wallace's (2018) gatekeeper archetypes: the archetypes describe *who* curates, while the functional architecture describes *where and how* curation decisions occur. It also extends the closest earlier framework, which combines gatekeeping, agenda setting, framing, algorithmic recommendation, and LLM-based news writing (Voinea 2025).

The *original conceptual contribution* consists of three connected instruments added to that earlier work: the functional architecture, including the dimensions of agentic escalation; the measurement model of cumulative exposure allocation, which defines exposure over time in sequential and performative terms and the accountability structure, which assigns duties and consequences through actor–forum–evidence–standard–remedy chains at the level of individual curation functions.

3. Method: Integrative Conceptual Review

Because the objective is theory building rather than effect estimation, the method is an *integrative conceptual review and theoretical synthesis* (Snyder 2019). The review brings together research traditions that locate curation power at different levels, including editors, organizations, platforms, models, interfaces, and institutions, and combines compatible concepts from those traditions into one analytical framework.

The literature base was assembled through targeted searches of Scopus, Web of Science, the ACM Digital Library, and arXiv, together with scholarly books and official legal repositories, including EUR-Lex and the European Commission. Searches conducted through June 2026 were followed by a July 2026 update focused on agentic AI and performative optimization.

The searches covered five main areas: foundational gatekeeping and media-effects theory; platformization and algorithmic distribution; news recommendation and personalization; reinforcement learning and recommender systems and accountability, auditing, and legal governance. The review prioritized peer-reviewed research, foundational theoretical works, and primary legal sources. It also included preprints, investigative reporting, and institutional documents when they provided primary technical findings, documented deployment practices, or recorded regulatory developments not yet available in peer-reviewed research; the text identifies the evidentiary status of those sources. Unsourced commentary was excluded.

The analysis proceeded in three stages. First, it identified where each body of research locates the power to select information: in editors, organizations, networks, platforms, models, or users. Second, it examined how each tradition treats values such as diversity, autonomy, public interest, transparency, and accountability. Third, it traced how sequential optimization changes the problem through delayed objectives, policy updates, feedback, and performativity.

Claims that can be tested empirically are stated as propositions or hypotheses and linked to possible measures, and contested evidence is reported as contested (Section 6). Appendix A documents the search strategy, inclusion criteria, and approximate composition of the literature base (Table A2), and Appendix C links the framework's main concepts to the sources that inform them. These appendices support transparency and auditability; they do not claim the reproducibility of a systematic review.

4. Six Curation Functions and the Feedback Layer

I define *adaptive news curation* as a socio-technical arrangement in which one or more machine-learned policies select, rank, sequence, repeat, suppress, or experimentally vary news exposure in response to observed user and system conditions, while data from those interactions may shape later policy updates. *Agentic news curation* is a high-autonomy form of adaptive curation: it combines persistent planning, delegated authority to act, coordination across curation functions, and delayed external review (Section 2.3). Classical editorial gatekeeping usually offered more visible organizational decision points, and conventional recommender scoring tends to treat each served item or ranked list as a separate prediction. Adaptive curation differs because its policy acts over time and learns from earlier exposure.

The model divides contemporary curation into *six decision functions connected by a cross-cutting feedback and optimization layer*, derived in two steps. The first step follows [Shoemaker and Vos's \(2009\)](#) definition of gatekeeping as selecting, writing, editing, positioning, scheduling, and repeating information: contemporary systems implement each of these activities partly through software. The second step separates the functions by their *decision object*, the thing being admitted, selected, transformed, presented, distributed, or restricted: source eligibility, the agenda and candidate pool, content production, content presentation, exposure, and permissibility.

This is a modular architecture, not a fixed pipeline: real systems use different combinations and orders of these functions, and retrieval, generation, packaging, allocation, and moderation may occur more than once. Two elements cut across the architecture. Moderation is itself a curation function (its decision object is permissibility), but it can operate before source inclusion, during generation, before exposure, after publication, and in response to user reports. Its two descriptions answer different questions and are used together throughout: moderation counts as the sixth function because it has a decision object of its own, and it is called cross-cutting because, unlike Functions 1–5, its decisions can attach to any point in the flow. Feedback-driven optimization is not a separate curation function because it acts on no media object of its own; instead, interaction data can update the policies used by one or more functions. Figure 1 shows this architecture. The model is detailed enough to connect different decision objects to different risks and forms of oversight, and broad enough to apply across organizations and technical designs.

Function 1: Source and creator eligibility. The first function decides which sources and creators may enter the distribution system at all; admission does not guarantee visibility or revenue, which depend on exposure allocation (Function 5). Compliance tools, spam-detection systems, and source-credibility measures decide admission or exclusion. The principal risk is systematic exclusion of niche or minority voices when training data favor established styles and institutions.

Function 2: Agenda and candidate formation. This function decides which topics, events, and items are eligible for coverage or distribution. Systems may monitor data streams to detect emerging events and then build a pool of possible items; topic-level agenda setting and item-level retrieval are different implementations of the same basic decision, namely what enters the eligible set. Measures such as velocity and sentiment can replace or supplement an editor's judgment of newsworthiness, creating an algorithmic form of agenda setting ([McCombs and Shaw 1972](#)). Function 2 determines what is eligible; Function 5 determines how exposure is distributed across that eligible pool.




Figure 1. Functional architecture of adaptive curation. Functions 1–5 show common dependencies but need not occur once or in a fixed order; moderation operates across the full architecture. The feedback layer receives interaction data (thick arrow) and may update one or more policies through online learning, periodic retraining, or organizational change. Italic labels show typical organizational owners. A circled C marks a possible organizational control point, such as predeployment review, sampled review, or approval of individual items. Agentic systems may operate across more functions or act before these controls are applied.

Function 3: Content generation and editing. With generative models, production itself becomes a site of gatekeeping. Generating text is not automatically gatekeeping; the gatekeeping choices lie in which sources the system uses, which constraints are built into prompts and fine-tuning, what the system includes or leaves out, how it frames the material, and whether it has authority to publish (Entman 1993 supplies the language of framing; DeVito 2017 the demonstration that values can be built into selection systems).

Function 4: Packaging and metadata. This function shapes how an item appears through headlines, thumbnails, tags, and localized translations. These elements may be optimized for predicted responses, as when headline variants are tested against expected click-through rate (CTR).

Function 5: Exposure allocation. This is the main distribution function: it decides which eligible items appear in which positions, to which users, and at what times. Recommender systems, including collaborative filtering, graph neural networks (GNNs), and reinforcement-learning policies, match packaged content to users through learned representations and interaction histories (Chen et al. 2019; Ie et al. 2019). The decision object is not whether an item is eligible but how opportunities for exposure are distributed across the eligible pool.

Function 6: Moderation and safety (cross-cutting). Automated moderation sets the boundaries of permissible content at a scale that makes complete human review impossible (Gillespie 2018), and it can operate at every stage: source vetting, filtering during generation, screening before exposure, and responding to reports after publication.

The feedback and optimization layer. Interaction data may guide online or periodic updates to one or more functions' policies, as learning systems adjust their rules for eligibility, candidate selection, generation, packaging, exposure, and moderation in response to engagement, retention, and drift signals (Jiang et al. 2019). This feedback layer turns a static sequence of decisions into an adaptive system, and it is where many of the distinctive risks of adaptive curation arise.

A worked example makes the architecture concrete. Consider a personalized election briefing created each morning for millions of users. Source and creator eligibility determines

which outlets may contribute reporting (Function 1). Agenda and candidate formation identify the day's races and controversies and retrieve possible items (Function 2). A generative system combines several reports into one account, deciding what to include, omit, and emphasize (Function 3). Packaging places the summary under a headline and thumbnail chosen for predicted response (Function 4). Exposure allocation decides which users receive each version, where it appears, and how often it returns on later days (Function 5). Moderation screens sources, generated text, and user comments throughout the process (Function 6).

Data on opens, reading time, and skips may then update retrieval, prompts, packaging, and ranking. A long-term evaluation would ask how the briefing distributed exposure across sources and viewpoints and how often it repeated material across users and weeks (Section 6). An accountability review would ask who controlled each choice and who should answer for the resulting pattern (Section 7).

Table 2 describes each function by its decision object, main actor or system, inputs, objective or constraint, and principal risk. The table does not force every function into a single “proxy reward” model: some functions optimize a learned objective, some enforce constraints, and others apply editorial or policy rules. The broader point is that gatekeeping power is distributed across creators, newsrooms, platforms, vendors, and interfaces, with values entering through objectives and constraints as well as through formal content policies.

Table 2. The six curation functions and the feedback layer. Table 3 identifies possible oversight forums and remedies for each function.

Function	Decision Object	Main Actor or System	Inputs or Evidence	Objective or Constraint	Main Risk
1. Source eligibility	Which sources and creators are admitted	Compliance and spam systems; partnership teams	Credibility measures, spam signals, policy rules	Integrity and quality requirements	Exclusion of niche or minority voices
2. Agenda and candidate formation	Which topics, events, and items become eligible	Trend-detection systems; news desks	Stream velocity, sentiment measures, editorial judgment	Predicted importance within coverage rules	Treating velocity as more important than public significance
3. Content generation	Source choice, synthesis, omission, framing, and authority to publish	LLM pipelines shaped by prompts and fine-tuning; editors	Training data, prompts, retrieved sources	Quality thresholds; style and sourcing rules	Reproducing bias; presenting machine output with the authority of reporting
4. Packaging	Headline, imagery, tags, and other presentation choices	A/B-testing and metadata systems; audience teams	Predicted CTR, past engagement	Increase predicted response within brand and accuracy rules	Emotionally escalatory presentation
5. Exposure allocation	How eligible items are distributed across positions, users, and time	Recommenders (collaborative filtering, GNNs, RL policies)	Learned representations, interaction histories	Engagement and retention goals, diversity rules, and business constraints	Concentrated or narrowed exposure
6. Moderation (cross-cutting)	Whether content, sources, and behavior are permitted	Classifier systems; human reviewers; policy teams	Text and image classifiers, user reports	Reduce policy violations while limiting errors	Context-blind suppression; weak enforcement
Feedback layer (cross-cutting)	How one or more policies are updated	Experimentation and learning infrastructure	Interaction data, experiment results	Cumulative goals (retention, engagement)	Self-reinforcing feedback loops; drift; confounded data

Table 3. Accountability matrix for the six curation functions: actors in control, affected parties, oversight forums, required evidence, evaluative standards, and remedies. E = existing; M = emerging or partly implemented; P = proposed, relative to the European Union baseline of Section 8.

Function	Actors in Control	Affected Parties	Forum	Evidence Needed	Standard	Remedy or Sanction
1. Source eligibility	Integrity and partnership leadership	Sources; creators	Platform appeals body (E); regulator (E); courts (E)	Admission and exclusion logs; error rates across groups	Non-discrimination and fair procedure	Restore the source or creator (E); publish an error audit (P)
2. Agenda and candidate formation	Owners of trend and editorial systems	Publics; under-covered communities	Independent auditors (M); media regulators (E)	Records showing the composition of the candidate pool	Coverage measured against a stated and authorized editorial baseline	Change the constraint (M); publish a coverage report (P)
3. Content generation	Newsroom or platform AI product owner; model vendor	Readers; subjects of coverage	Press councils (E); regulator (M); courts (E)	Provenance labels; generation logs; prompt and version history	Accuracy, disclosure, and correction standards	Correct the content (E); enforce labeling (M); retrain the model (M)
4. Packaging	Audience and product teams	Readers	Consumer-protection authorities (E)	Experiment registry; disclosure of optimization targets	No deceptive presentation	End the experiment (M); order disclosure (E)
5. Exposure allocation	Objective owner; deploying platform operator	Users; publishers; creators	Systemic-risk auditors (M); approved researchers (M); regulator (E)	Descriptive CEA measures (Equations (1)–(3)); policy comparisons (Equation (4)) supported by causal counterfactual exposure tests with recorded selection probabilities	Systemic-risk mitigation; stated diversity requirements evaluated against an authorized baseline	Roll back the ranking change (M); change the constraint (M); impose a fine (E)
6. Moderation (cross-cutting)	Policy and enforcement teams	Speakers; audiences	Internal appeals (E); out-of-court dispute bodies (M); courts (E)	Moderation logs; error-rate audits across languages and groups	Proportionality and protection of freedom of expression	Restore content (E); provide compensation (M); require error-rate remediation (P)
Feedback layer (cross-cutting)	Owners of experimentation infrastructure; record custodians	All of the above	Independent algorithmic auditors (M); regulator (E)	Versioned objectives and change history; policy, model, and prompt versions; intermediate action logs; review events; OPE reports; recorded selection probabilities where causal claims are made	Alignment between reward and stated objective; documented human oversight	Retrain the model (M); end the experiment (M); require external monitoring (M); publish audit results (M)

Two consequences follow. First, harm can build up over time: a sequence of individually defensible choices may still reduce the visibility of minority creators, narrow the range of content, or encourage compulsive use. Second, evaluation must examine the *distribution* of exposure across creators, topics, and population groups, in addition to the quality of individual items and the error rate of moderation.

5. How Sequential Optimization Changes Curation

To understand how a system builds in, or pushes aside, social values, it is necessary to examine how adaptive curation works. Reinforcement learning is used here in two

ways. First, RL is a *design approach* used in some large-scale recommenders, although many deployed news systems still rely on supervised learning or contextual bandits (Afsar et al. 2022; Chen et al. 2019, 2023; Zou et al. 2019). Second, RL is a useful *conceptual limiting case*: it makes the accountability problem especially clear because it explicitly represents the objectives, feedback, delayed outcomes, and policy updates that simpler systems may leave implicit.

5.1. Modeling Curation as Sequential Decision-Making

Reinforcement learning is commonly modeled as a Markov decision process (MDP), which includes states, actions, transition probabilities, rewards, and a discount factor (Sutton and Barto 2018). Mapping news curation onto these elements provides a clear structure for analysis.

- **State s_t** is the context available at time t : the user's history and learned representation, session details, the candidate pool, and time. An MDP *assumes* that the represented state contains all relevant information from the past; in real curation systems, preferences, intentions, fatigue, and immediate circumstances are only partly visible, so this assumption is an approximation.
- **Action a_t** is the exposure the system serves, usually a ranked group of items or a generated summary (Ie et al. 2019).
- **Transition $P(s_{t+1} | s_t, a_t)$** is the probability that the action leads to a new state, which depends on uncertain user behavior: clicking, scrolling past, or leaving.
- **Reward r_t** is the immediate numerical signal used to evaluate the action; in practice, this may be click-through, dwell time, watch time, completion, sharing, or revenue.
- **Discount factor $\gamma \in [0, 1)$** is the weight placed on delayed outcomes. A small value makes the system focus on the next response; a larger value gives more weight to later effects, and whether that promotes diversity, retention, or dependency depends on the reward and constraints, not on the discount factor alone.

The agent seeks a policy π that maximizes expected cumulative discounted reward $\mathbb{E}_\pi[\sum_t \gamma^t r_t]$. Systems using this general approach have been deployed at industrial scale: policy-gradient recommenders have optimized engagement across very large action spaces, with corrections for the cumulative advantage of already-popular items, and related systems explicitly optimize long-term engagement rather than only the next response (Chen et al. 2019; Zou et al. 2019). Figure 2 shows the basic loop between the agent and its environment.

Three ordinary features of sequential systems should not be confused with the stronger phenomenon discussed below. First, an action can change the next state without changing the environment itself: if a recommendation changes a preference that is already represented in the state, a standard MDP can describe that change through $P(s_{t+1} | s_t, a_t)$. Second, external change, such as a breaking-news cycle, can make the environment unstable; this complicates learning but is not specific to the deployed policy. Third, a policy creates different one-step dynamics by choosing different actions, which is an expected result of policy choice, not evidence that the policy has changed the underlying environment.

The stronger, policy-specific phenomenon is *performativity*. A policy is performative when repeated deployment changes the underlying transition or reward process itself, so the dynamics are better written $P_\pi(s_{t+1} | s_t, a_t)$: the environment faced by a later policy partly depends on the policy that was already deployed. This setting is studied in performative prediction (Perdomo et al. 2020) and performative reinforcement learning (Mandal et al. 2023). One form of performativity in curation is *endogenous preference change*: repeated exposure may change a user's underlying preferences, not only the behavior the system observes, so the system is optimizing against a target that it has partly created. Performative-

tivity matters most for the measurement argument: when it is present, the deployed policy shapes the environment that any later evaluation observes. The traceability proposition (P1) below does not depend on it. Direct evidence that deployed news recommenders change underlying preferences is not yet available, and the research agenda in Section 9 treats this as an open question.




Figure 2. News curation as sequential decision-making, in the standard agent–environment loop. (a) In a fixed environment, actions change later states through an unchanged transition process P . (b) In a performative environment, sustained use of policy π changes the transition process itself (P_π): exposure may change the preferences against which the system later optimizes (Mandal et al. 2023; Perdomo et al. 2020).

The exploration–exploitation trade-off also affects exposure diversity. In curation, *exploitation* means showing more of what a user has already engaged with, and a system that strongly favors exploitation may narrow exposure toward familiar content even when no designer intends censorship or bias.

Exploration, however, is not the same as diversity: it means trying actions whose value is uncertain, and those actions need not involve different sources, opposing views, minority creators, or public-interest journalism. A system can explore entirely within a narrow entertainment category. Exposure breadth depends on the candidate pool, the representation of topics and sources, objectives and diversity constraints, reranking, business rules, interface design, and editorial intervention (Jiang et al. 2019); the exploration setting is only one factor. The governance point remains: recommendation diversity is a system property that can be adjusted and measured, with trade-offs that should be made explicit.

This formalization is a conceptual tool; it introduces no new theorems. Visibility is shaped by an entire optimization stack: source and candidate eligibility, state representation, objectives, constraints, feedback measurement, and interface placement. The reward function is one important point where editorial and commercial priorities enter the system, because it determines which observable outcomes are reinforced over time. Other choices can be equally consequential even when product, engineering, or commercial teams make them.

The analysis extends earlier work showing that algorithms gatekeep and encode values (DeVito 2017; Gillespie 2014; Napoli 2014), as well as work on aligning recommenders with human values (Stray et al. 2021; Stray et al. 2024). Its main addition is to locate those choices within a policy that acts over time, so evaluation must also take place over time (Section 6).

5.2. Why One Engagement Metric Is Not Enough

In news curation, operational rewards stand in for broader journalistic and democratic goals, and problems arise when a measurable proxy is treated as if it were the goal itself. AI-safety research calls one version of this problem *reward hacking*: a system earns a high reward by satisfying the literal measure while missing its intended purpose (Amodei et al. 2016). Later work formalizes how optimizing an imperfect proxy can damage the underlying objective, a reinforcement-learning version of Goodhart’s Law (Karwowski et al. 2024; see also Krakovna et al. 2020 on specification gaming).

The gap is especially important in news. Clicks, watch time, and completion do not by themselves represent an informed public, a wide range of knowledge and viewpoints, or public-interest journalism. A single engagement score is, therefore, a weak normative basis for news curation even when it is computationally convenient. This conclusion is an inference from the combined literature, not a finding from one study.

Performativity and preference change make the problem harder. If exposure changes the preferences that a system later measures, the recommender learns from behavior that it helped produce: its target is moving and partly self-generated, with performative stability as the corresponding equilibrium concept (Mandal et al. 2023; Perdomo et al. 2020). A practical version of this problem is the self-reinforcing feedback loop, in which the system serves a narrow set of items, infers narrower interests from the resulting behavior, and then serves an even narrower set; larger candidate pools and some randomization can slow this process (Jiang et al. 2019). The same dynamic creates a methodological problem: interaction logs are shaped by earlier recommendations, so researchers cannot easily separate a user's prior preferences from responses caused by past exposure.

Performative effects also occur in fields such as credit scoring, predictive policing, labor markets, and pricing. They are especially important in curation because the system is designed to shape the attention that later becomes its training signal: adaptive curation can reshape the public it serves, and the relevant harm can emerge from the policy's repeated operation rather than from any single item or decision. For the agentic case, one testable proposition follows.

Proposition 1. *When one autonomous policy controls more curation functions, and no function-level provenance rules are in place, auditors will be able to reconstruct a smaller share of exposure decisions within a fixed time or cost budget. Reconstruction requires the relevant model or policy version, objective, inputs, intermediate actions, responsible institutional owner, and review event.*

Decision-provenance completeness can be measured as the share of sampled exposure outcomes for which an independent auditor, using the available documentation and logs, can reconstruct all of these elements within a predefined time or cost limit.

6. Cumulative Exposure Allocation: Evidence, Definition, and Measures

If the important effects of a curation policy accumulate across its deployment, evaluation needs a unit that captures accumulation. Defining that unit requires knowing what the empirical record already shows, because several widely assumed harms have weaker support than public debate suggests.

6.1. What the Evidence Does and Does Not Show

Public debate often treats three claims as settled: algorithmic curation traps users in filter bubbles, recommenders cause radicalization, and engagement-based ranking drives the spread of false news. The evidence is more mixed. Table A3 in Appendix B identifies the strongest form of evidence behind each claim and explains how the argument uses it.

The strongest available evidence complicates the filter-bubble argument associated with Pariser (2011) and anticipated by Sunstein (2001). On Facebook, algorithmic ranking reduced exposure to politically cross-cutting content, but users' own choices reduced it even more (Bakshy et al. 2015). Research on social and search media found greater ideological distance alongside greater exposure to opposing views (Flaxman et al. 2016). In the large experiments during the 2020 U.S. election, replacing algorithmic feeds with reverse-chronological feeds for about three months did not significantly change polarization or political knowledge (Guess et al. 2023). A companion experiment that reduced exposure to content from like-minded sources over a similar period found no significant change in

affective polarization (Nyhan et al. 2023). These results have important limits and should not be read as showing that curation has no effect. The most defensible conclusion is that ranking, social networks, and user choices interact, with political content especially shaped by self-selection (Bruns 2019; Mangold et al. 2024).

Recent YouTube research also does not support a uniform, average effect in which recommendations consistently drive users toward more extreme content. Early audits found pathways toward more extreme channels (Ribeiro et al. 2020). Later panel research attributed much fringe consumption to subscriptions and external links (Hosseinmardi et al. 2021), a study using counterfactual bots found no average recommendation-driven shift toward more extreme content (Hosseinmardi et al. 2024), and supply-and-demand analyses place substantial weight on audience demand (Munger and Phillips 2022). These findings do not rule out effects for particular users, topics, or pathways, but they do reject a simple universal claim.

False news clearly spreads differently from true news: it has been shown to travel farther, faster, deeper, and more broadly, in large part because of human sharing behavior (Vosoughi et al. 2018). That finding concerns content and diffusion; by itself, it does not establish that engagement-ranking algorithms caused the observed pattern.

Internal documents reported by journalists provide more direct evidence about ranking choices on one platform. Reports on Facebook’s 2018 “Meaningful Social Interactions” redesign described an increased weight for reshares, reported that an “angry” reaction was once weighted five times as heavily as a “like” before that weight was reduced, and described internal researchers warning that these choices could raise the visibility of divisive content (Hagey and Horwitz 2021; Merrill and Oremus 2021). These accounts rest on established journalists’ review of authenticated documents, but no peer-reviewed audit independently confirms them, and the company disputed parts of the reporting. The stronger and more general conclusion is structural: engagement optimization can create incentives and technical capacity to amplify divisive, novel, or emotionally activating material even where no user is mechanically imprisoned by an algorithm.

Existing research, therefore, does not support universal claims that algorithmic ranking always isolates, radicalizes, or misinforms users. Effects differ across platforms, populations, pathways into content, time periods, and measured outcomes. This heterogeneity does not make recommender design inconsequential. Instead, it makes broad claims about individual effects a weak basis for governance, and it points toward a more stable target: the distribution of exposure that a deployed policy produces across people and over time.

6.2. Defining Cumulative Exposure Allocation

The analysis, therefore, shifts the focus from one recommendation, or one user’s “bubble,” to *cumulative exposure allocation* (CEA): how a continuously learning policy distributes visibility, repetition, and salience across time, population groups, and contexts. CEA measures opportunities for attention; it does not claim to measure actual attention, understanding, persuasion, or changes in preference. Its epistemic status is dual: CEA is a conceptual construct, because it redefines the object of evaluation as the distribution of exposure a policy produces over time, and it is a measurement model, because Equations (1)–(4) and the indicator families of Table 4 make that object operational. For a deployed policy π , a user group g , a category k (a source, topic, viewpoint, creator group, or format), and a time window T , define the raw weighted exposure total

$$A_{g,k}^{\pi}(T) = \sum_{u \in g} \sum_{e \in E_u(T)} w_{u,e} q_k(i_{u,e}), \quad (1)$$

where $E_u(T)$ is the set of exposure events for user u during the window, $i_{u,e}$ is the item shown at event e , the weight $w_{u,e}$ adjusts for factors such as position or whether the item was actually viewable, and $q_k(i) \in [0, 1]$ represents how strongly item i belongs to category k , allowing an item to have several labels or an uncertain classification. The raw total $A_{g,k}^\pi(T)$ will naturally be larger for bigger groups and for groups containing heavier users, so the main object of evaluation is the category's share of the group's total weighted exposure,

$$p_{g,k}^\pi(T) = \frac{A_{g,k}^\pi(T)}{\sum_j A_{g,j}^\pi(T)}, \quad (2)$$

together with a version that gives each user equal weight, so that a small number of highly active users do not dominate the group result,

$$\bar{p}_{g,k}^\pi(T) = \frac{1}{|g|} \sum_{u \in g} \frac{\sum_{e \in E_u(T)} w_{u,e} q_k(i_{u,e})}{\sum_{e \in E_u(T)} w_{u,e}}, \quad (3)$$

and a third measure that compares the deployed policy with a clearly stated baseline policy π_0 ,

$$\Delta p_{g,k}^{\pi, \pi_0}(T) = p_{g,k}^\pi(T) - p_{g,k}^{\pi_0}(T), \quad (4)$$

where π_0 is a chronological, editorial, non-personalized, randomized, or otherwise justified comparison policy. The four quantities answer different questions: A is the raw weighted total; p is the normalized distribution for a group; $\bar{p}$ represents the average user rather than the most active users and Δp compares the deployed policy with a stated alternative. Users with no exposure events in the window are excluded from the average in Equation (3), and an audit should report how many users this excludes. The comparison measure does not by itself prove that the policy caused the difference: causal attribution additionally requires an appropriate design, such as randomization, recorded selection probabilities, sufficient overlap between policies, and credible assumptions about confounding. Even with that limitation, the definitions turn long-term exposure into a measurable object.

Three qualifications are essential.

1. **Lower concentration is not always better.** Elections, emergencies, and major public events can legitimately concentrate public attention. CEA indicators, therefore, need a benchmark, and the process for choosing that benchmark must itself be governed. Possible baselines include the available candidate supply, an editorial policy, a chronological or non-personalized feed, an explicit diversity rule, a user-selected goal, and a public-interest minimum established through regulation or participatory governance. CEA describes a distribution; a legitimate process must decide which distribution is desirable, and Section 7 asks who should have authority to make that decision.
2. **Category definitions shape the result.** An audit must record who defines sources, topics, viewpoints, creator categories, and population groups; whether items can belong to several categories and how classifier uncertainty and disagreement among human coders are reported.
3. **Group audits must protect privacy.** Evaluation must follow data-minimization rules and protect small groups from re-identification.

CEA builds on research on exposure diversity in news recommendation (Bauer et al. 2024; Helberger 2019; Vrijenhoek et al. 2021), on agenda setting (McCombs and Shaw 1972), and on the capture of attention (Seaver 2022; Zuboff 2019), but it differs from three related approaches. Agenda-setting research usually measures the overall importance given to issues; recommendation-diversity research often evaluates one list, session, or user and research on fairness of exposure in rankings examines how position-weighted exposure is

distributed among producers or source groups, including exposure accumulated across repeated rankings (Biega et al. 2018; Diaz et al. 2020; Singh and Joachims 2018). CEA combines these concerns by measuring weighted exposure across time, policy deployments, content categories, and recipient groups, and it connects the resulting distribution to the functions and institutions that produced it. The measurement model shares its machinery with that fairness literature. It complements the functional architecture of Section 4: the functions show *where* values and constraints enter the system, while CEA shows *what to measure* when evaluating the policy that operates across those functions over time.

Table 4 groups the proposed CEA measures into two sets: primary indicators suitable for routine long-term evaluation, including concentration measured by the Herfindahl–Hirschman index (HHI) or entropy calculated from normalized exposure shares, and specialized indicators that answer narrower questions and require additional design choices. The families differ in provenance. Concentration, breadth, and long-tail coverage rely on established quantities, since the HHI, entropy, and catalog-coverage measures are standard in economics, information theory, and recommendation-diversity research (Bauer et al. 2024; Vrijenhoek et al. 2021). Cross-cutting exposure adapts a familiar construct from research on selective exposure (Bakshy et al. 2015; Flaxman et al. 2016), and group exposure disparity adapts position-weighted exposure comparisons from research on fairness of exposure in rankings (Biega et al. 2018; Diaz et al. 2020; Singh and Joachims 2018). The formulations proposed here are the repetition and persistence indicators, the equal-user-weight share in Equation (3), and the requirement that the policy comparison in Equation (4) be evaluated against a governed baseline. A second testable proposition follows.

Table 4. Families of CEA indicators. Primary indicators support routine evaluation over time; specialized indicators address narrower questions and require additional assumptions or data. Descriptive indicators alone do not show that the deployed policy caused the observed distribution.

Status	Indicator Family	Definition	What It Shows	Main Caution
Primary	Normalized exposure share	$p_{g,k}^{\pi}(T)$, Equation (2)	Relative visibility of sources, topics, viewpoints, or formats	Depends on category definitions and the chosen denominator
Primary	Concentration or breadth	HHI or entropy calculated from normalized shares	Whether exposure is dominated by a few categories or spread broadly	Choose and justify the index that matches the question
Primary	Repetition and persistence	Re-exposure rate and length of dominance across time windows	Self-reinforcing repetition and lasting salience	Distinguish justified repetition during major events from harmful repetition
Primary	Group exposure disparity	Differences in $p_{g,k}^{\pi}$ or $\bar{p}_{g,k}^{\pi}$ across groups	Lasting underexposure or overexposure	Requires privacy-preserving group definitions and justified comparison groups
Specialized	Long-tail coverage	Share of the eligible catalog receiving any exposure	Whether niche sources are systematically overlooked	Depends strongly on how the candidate pool is defined
Specialized	Cross-cutting exposure	Share assigned to categories that conflict with a user's prior attitude	How insulated a user's exposure is	Requires a defensible model of attitudes or viewpoints
Specialized	Exposure–consumption difference	Comparison between what was served and what users engaged with	Difference between opportunity and uptake	Descriptive only; does not by itself identify a cause
Specialized	Policy-attributable difference	$\Delta p_{g,k}^{\pi, \pi_0}(T)$, Equation (4)	Difference associated with the deployed policy relative to a baseline	Requires an experiment or a credible quasi-experimental design

Proposition 2. *Two adaptive curation policies can perform similarly on item-level accuracy, content quality, moderation errors, and short-term engagement while producing meaningfully different long-term patterns of source concentration, topic concentration, repetition, and exposure disparity across groups.*

6.3. Illustrative Field Evidence on Bounded Personalization

If unconstrained engagement optimization creates risks, an important practical question is whether constraints can reduce those risks without eliminating useful personalization. One field experiment in a legacy news organization provides limited but concrete evidence that this is possible. Researchers ran a 34-day A/B test with approximately 58,000 paying mobile subscribers at a major Norwegian news organization (Holzleitner et al. 2026). The experimental system added a matrix-factorization collaborative-filtering component, capped at 20% of the final score, to a manually weighted combination of popularity, recency, and recent-performance signals; editorial and non-personalized signals remained dominant, and the system did not use reinforcement learning.

The results challenge the assumption that engagement and public-service values must conflict. Compared with the non-personalized baseline, the hybrid system increased click-through and visit frequency while *reducing* popularity bias: users in the personalized condition opened a wider range of articles, topic concentration fell, and the system exposed more of the daily catalog. The study was peer reviewed, but it examined one organization, and its effect sizes warrant independent, multi-site replication.

The findings should be read within the study's actual design: they support bounded personalization and editorially constrained ranking, and they provide no evidence about constrained reinforcement learning. Their importance here is narrower. In one organization, personalization was deliberately limited and evaluated against both engagement and journalistic-value measures, a result consistent with research that turns editorial values into explicit design goals (Lu et al. 2020; Mattis et al. 2024; Vrijenhoek et al. 2021). The broader claim, that a concentration constraint can reduce exposure concentration without meaningfully reducing utility, is plausible but not established, and it is stated as a hypothesis for controlled, cross-organization testing.

Hypothesis 1. *Compared with a system optimized only for engagement, a ranking policy with an explicit source-concentration constraint reduces the HHI of normalized source exposure while keeping click-through, retention, or another prespecified utility measure within a non-inferiority margin δ .*

7. Accountability for Adaptive News Curation

Measurement shows what a policy produced; it does not decide who must answer for it. If editorial influence is spread across an optimization system, and harm can emerge from a policy's repeated operation, transparency alone cannot provide accountability. Pasquale (2015) describes the broader condition as the "black box society," and Burrell (2016) identifies three sources of opacity: deliberate corporate secrecy, limited technical understanding, and complexity within machine-learning systems. Disclosure directly addresses only the first.

Algorithms must also be understood as socio-technical systems embedded in organizations and everyday practice (Kitchin 2017), and Mittelstadt et al. (2016) identify traceability, the ability to connect an outcome to the decisions that produced it, as central to accountability. Ananny and Crawford (2018) make a related point: seeing parts of a system does not necessarily mean understanding or governing it. A complete accountability model must, therefore, identify who must explain a decision, who can question and judge it, and what consequences may follow. This is the central concern of algorithmic-accountability research (Diakopoulos 2015).

The issue is not limited to consumer protection; it also concerns democracy. Habermas (2022) argues that digital platforms both expand and fragment public communication, creating partly enclosed spaces of discussion. When privately owned platforms optimize public communication for behavioral prediction (Zuboff 2019), curation systems help shape the conditions under which citizens encounter information, debate issues, and form opinions.

The focus here is, therefore, *societal* accountability: it extends the social-responsibility tradition from human editors, who answer to professional norms, to adaptive systems whose lines of responsibility must be deliberately designed. Fair selection of an individual story still matters; the larger question is whether the policy that distributes visibility over time remains answerable to the public whose attention it organizes.

7.1. Who Answers to Whom, and What Follows

Bovens (2007) defines accountability as a relationship between an actor and a forum: the actor must explain and justify conduct, and the forum can ask questions, make a judgment, and impose consequences. This actor–forum–consequences structure provides a clear foundation for societal accountability in agentic curation, and it shows where adaptive systems create difficulties. Wieringa (2020) shows that the same model already organizes much of the algorithmic-accountability literature. The matrix developed below extends that work in three ways: it assigns accountability at the level of individual curation functions, it names the evidence and the standard each forum needs, and it marks each forum and remedy as existing, emerging, or proposed.

The *actor* is not one person or organization. Adaptive curation depends on decisions about data, measurement, objectives, deployment, and oversight, distributed across creators, publishers, platform operators, model developers, and vendors, so responsibility should follow actual decision rights rather than being assigned vaguely to “the algorithm.” The *forum* is also plural: relevant forums include regulators, auditors, approved researchers, affected users and creators, courts, and the wider public.

The greatest difficulty concerns *consequences*. Research on responsibility gaps argues that traditional ideas of blame become harder to apply when learning systems behave in ways that operators cannot fully predict (Matthias 2004), and that the apparent single gap separates into gaps in culpability, moral accountability, public accountability, and active responsibility (Santoni de Sio and Mecacci 2021). Floridi’s (2016) account of distributed moral responsibility supports assigning responsibility across the wider system rather than searching for one fully culpable person. In practice, an accountability regime should define forums and consequences before deployment; otherwise, responsibility may spread across an adaptive system whose decisions are too difficult to reconstruct after harm occurs.

7.2. Questions an Accountability Review Should Ask

The actor–forum–consequences model becomes practical when it is paired with the six curation functions. Six questions, grouped into four areas, describe what an oversight body should examine; they apply to the functions and their outcomes rather than forming a separate classification.

Substantive outcomes: what does the system produce? *Exposure equity* asks who receives visibility and the opportunity to earn revenue. One diagnostic is a counterfactual exposure test: compare the exposure a source receives under the deployed policy with the exposure it would receive under a defined alternative, using randomized exploration traffic or logged data with recorded selection probabilities (Metaxa et al. 2021; Sandvig et al. 2014). Provider-side fairness measures can also compare exposure across creator groups. These diagnostics are especially relevant to Functions 1 and 5. *Cultural plurality* asks whether the system provides a meaningful range of topics, viewpoints, and formats. Its main diagnostics are the descriptive CEA measures of Section 6.2, evaluated against a benchmark whose selection must itself be authorized and justified (Vrijenhoek et al. 2021); this question is especially relevant to Functions 2 and 5.

Procedural rights: what can affected people do? *Provenance and disclosure* asks whether users can identify AI involvement and trace the sources behind content. Possible

diagnostics include clear labels, generation logs, watermarking, and comprehension tests that check whether users actually understand the disclosure (Functions 3 and 4). *User contestability* asks whether users and creators can challenge a decision or change how the system treats them, through appeals, explanations, preference controls, and opt-outs (Functions 5 and 6).

Technical assurance: does the optimization remain aligned with its stated goals?

Reward robustness asks whether the system's measurable proxies remain connected to its intended goals as users, content, and models change. Diagnostics include versioned audits of reward specifications, documented in artifacts similar to model cards (Stray et al. 2021, 2024); off-policy evaluation (OPE) and long-term monitoring of the CEA measures for drift. OPE estimates how a candidate policy might perform by using data collected under a different policy, and it requires sufficient overlap between the policies and records of the probabilities with which actions were selected. When those assumptions do not hold, more credible tests of feedback effects require randomized exploration, holdout groups, interleaving experiments, or instrumental-variable designs. This question applies mainly to the feedback layer.

Institutional governance: does the oversight process itself work?

Oversight quality asks whether audits are independent, researchers receive meaningful access, and problems are corrected promptly. Diagnostics include checks on auditor independence, data-access rules, whistleblower channels, and time to remediation (all functions). No single measure can establish democratic adequacy, and audits do not create accountability on their own: they must operate within institutions that give independent, technically capable auditors enough authority and access to act (Raji et al. 2020, 2022; Birhane et al. 2024).

7.3. An Accountability Matrix for the Curation Functions

Responsibility should follow control over three things: the objective or constraint, the authority to deploy or act, and the evidence needed to reconstruct an outcome. Different organizations may control each one: the owner of an objective, the deploying operator, the model or infrastructure vendor, and the keeper of system records need not coincide. The matrix, therefore, identifies actors who actually control decisions rather than treating “the algorithm” or a single nominal owner as the responsible party.

Bovens' model requires consequences as well as explanations, so an accountability system must state what happens when a forum finds that an actor failed to meet the relevant standard. Table 3 links each curation function to the actors in control, the main affected parties, a forum with authority, the evidence that forum needs, the standard it should apply, and the available remedies. Remedies can be operational, such as rolling back a ranking change, restoring a source, correcting generated content, or ending an experiment; technical, such as retraining a model or changing a constraint; or institutional, such as compensation, external monitoring, regulatory fines, or publication of audit results. Forums and remedies are marked as existing (E), emerging or partly implemented (M), or proposed (P), assessed against the European Union baseline of Section 8.

The matrix has implications at two levels. At the institutional level, it requires capacities that do not yet fully exist: independent auditors need enough access to conduct off-policy evaluation and counterfactual exposure tests on live systems, and governance may need greater algorithmic pluralism and interoperability so that one engagement-optimizing policy does not control most exposure. At the engineering level, the matrix supports bounded personalization of the kind described in Section 6.3; multi-objective or constrained reward design that includes diversity, source reliability, and protections for vulnerable users, documented in auditable artifacts with change logs and clear human-override procedures (Stray et al. 2021, 2024) and evaluation designs that address confound-

ing through randomized exploration, holdouts, and recorded selection probabilities (Jiang et al. 2019).

These measures create real tensions. Public-interest diversity can become paternalistic if users have no meaningful agency; detailed logging can threaten privacy; rules against harmful amplification can lead to excessive removal and internal ethics programs can become symbolic if independent bodies cannot enforce them. The appropriate response is to manage these tensions through proportionality, participatory design, and institutional pluralism: meaningful user choice, risk definitions shaped by affected communities, and auditors and regulators that remain independent of the commercial systems they oversee.

8. How EU and U.S. Law Fit the Framework

Existing law already provides some parts of the accountability structure. The aim here is a crosswalk rather than a full legal analysis: Table 5 shows where current legal instruments provide oversight forums, evidence requirements, or remedies. Statements about what these instruments require describe current law; the gaps named in the final column of Table 5 and the items marked P in Table 3 are normative recommendations, not existing obligations. Clear distinctions among the main European laws are important because they are often treated as if they performed the same function.

Table 5. Regulatory crosswalk: where European Union and United States law provide forums, evidence requirements, and remedies relevant to the accountability matrix. The final column identifies gaps in relation to the accountability structure, not definitive doctrinal conclusions.

Function or Element	EU Legal Basis	What It Provides	U.S. Analogue or Constraint	Remaining Gap
Recommender parameters and user choice	DSA Articles 27 and 38 (European Union 2022b)	Disclosure of main parameters; option not based on profiling	Fragmented consumer-protection and state platform laws	No shared standard for measuring exposure over time
Systemic risk, audits, and researcher access	DSA Articles 34–35, 37, and 40 (European Union 2022b)	Risk assessment and mitigation; audit cycles; access for approved researchers	No equivalent general federal regime	Depth of access, auditor independence, and remedies vary
Provenance of generated content	AI Act Article 50 (European Union 2024b)	Disclosure duties for providers and deployers	Sector-specific and state approaches; constitutional speech limits	Article 50(4) exempts text under editorial responsibility; user understanding and enforcement remain uncertain
Moderation safeguards for media providers	EMFA Article 18 (European Union 2024a)	Reasons for a restriction; priority handling of complaints	No direct equivalent	Applies only to declared media service providers; depends on platform procedures
Algorithmic curation as editorial judgment	None	None	Section 230 protection for targeted recommendations unresolved (Gonzalez v. Google LLC 2023; Twitter, Inc. v. Taamneh 2023); curation as protected editorial judgment (Moody v. NetChoice, LLC 2024); recommendations as the platform's own speech (Anderson v. TikTok, Inc. 2024)	Duties specific to recommendations remain contested

The **Digital Services Act (DSA)** (European Union 2022b) is the main EU instrument governing platform recommenders. Article 27 requires all online platforms to disclose the main parameters of their recommender systems. Article 38 requires very large online platforms and search engines to offer an option not based on profiling. Articles 34–35 require assessment and mitigation of systemic risks, including risks connected to recommender

design; Article 37 requires independent audits and Article 40 provides data access for approved researchers. These rules have moved from legislation into enforcement: the European Commission has designated very large platforms, opened proceedings involving recommender-related risks, adopted the delegated regulation governing researcher access, and received the first systemic-risk assessments and audits. Their practical effect will depend on how much access researchers and auditors receive and how rigorously the rules are enforced.

The **Artificial Intelligence Act** (European Union 2024b) reaches ordinary media recommenders mainly through its ban on manipulative techniques that materially distort behavior and cause significant harm (Article 5(1)(a)–(b)). Article 50 also creates transparency duties for generated media: synthetic content must be marked as machine-generated, and deep fakes and, in specified cases, AI-generated text published to inform the public must be disclosed. The text-disclosure duty has a significant exception: under Article 50(4), it does not apply where the content has undergone human review or editorial control and a natural or legal person holds editorial responsibility for its publication. Much newsroom-supervised generation, therefore, falls outside the duty, so stronger human oversight can mean weaker mandatory provenance. These rules connect directly to the provenance-and-disclosure question and to the content-generation function. The Act's high-risk regime would apply only if a curation system fell within one of the listed high-risk uses; that boundary may be tested as curation becomes more agentic, but the classification should not be presumed.

The **European Media Freedom Act** (European Union 2024a) provides a forum and a remedy for some source-eligibility and moderation decisions involving institutional media. Under Article 18, a platform that intends to suspend or restrict content from a declared media service provider on moderation grounds must first provide reasons and must handle the provider's complaints with priority.

The United States takes a different approach. It remains judicially unresolved whether Section 230 of the Communications Decency Act protects a platform's *targeted algorithmic recommendations* of third-party content: the Supreme Court was asked to decide that question in *Gonzalez v. Google LLC* but did not do so, resolving the case through the companion ruling in *Twitter, Inc. v. Taamneh* (Gonzalez v. Google LLC 2023; Twitter, Inc. v. Taamneh 2023). The Court has, however, addressed the editorial character of curation. In *Moody v. NetChoice, LLC*, which reviewed Texas and Florida laws limiting platform moderation, the majority treated the filtering, prioritizing, organizing, and labeling of third-party content in curated feeds as expressive editorial judgment protected by the First Amendment (Moody v. NetChoice, LLC 2024). This ruling has two implications for the present argument: it supports the analogy between algorithmic curation and editorial judgment on which the framework rests, and it limits some of the regulatory tools available in the United States for governing that judgment. Two later developments qualify both implications. Justice Barrett's concurrence in *Moody* reserved the question whether feeds driven purely by engagement prediction convey editorial judgment at all, so the constitutional status of adaptive curation is not settled. In *Anderson v. TikTok, Inc.*, the Third Circuit read *Moody* to make algorithmic recommendations the platform's own expression and held, on that basis, that Section 230 did not bar a claim aimed at those recommendations (Anderson v. TikTok, Inc. 2024). What follows for the United States is a normative judgment rather than a statement of current law: in the near term, disclosure and researcher access are more realistic instruments for the framework than exposure baselines.

The comparison requires care. EU law partly implements the accountability structure through transparency rules, systemic-risk duties, audits, researcher access, and procedural safeguards. U.S. law provides fewer general mechanisms of this kind before harm occurs,

while intermediary-liability rules and constitutional protections shape, and may limit, regulation directed at curation.

9. Limitations and Research Agenda

The result is a conceptual model; it does not causally test every mechanism it names. Five limitations define the boundaries of its claims.

1. **Evidence about democratic effects is mixed and platform-specific**, and proprietary systems are difficult to observe from outside. The framework is intended for comparison and governance design; it passes no verdict on any particular deployed system.
2. **Highly autonomous agentic curation is still emerging**. Claims about fully autonomous editorial agents are, therefore, partly anticipatory and are identified as such.
3. **CEA categories and benchmarks are contestable**. Definitions of sources, topics, viewpoints, and population groups reflect judgments, and auditing across groups must comply with privacy, data-minimization, and small-cell protections.
4. **The literature and the legal comparison have a limited geographic scope**. Much of the cited research is Anglophone and Euro-American, and the legal crosswalk covers two Western systems. In state-controlled media environments, the accountability problem may be inverted: the main risk is a curation policy answerable primarily to the state, while no forum is independent of both government and platform power. In countries with limited regulatory capacity, the matrix's requirements for logs, audits, and researcher access assume institutions that may not exist, and multilingual platforms create a further challenge because ranking and moderation quality can vary sharply across languages. Extending the framework to these settings is a substantial research task; the accountability structure should be read as a normative benchmark, not as a description of current global practice.
5. **Technology and organizational arrangements change quickly**. Model designs, deployment practices, and vendor relationships may change faster than audits and legislation. Accountability should, therefore, attach to functions and outcomes instead of named technologies, and it should rely on standing institutions (auditors, access regimes, documentation duties) instead of one-time compliance exercises. The framework itself will need periodic revision as the division of work between people and systems changes.

These limits define a research agenda, for which P1, P2, and H1 are the main starting points: audit how completely decisions can be reconstructed as more curation functions are integrated; compare policies that perform similarly at the item level but produce different exposure patterns over time and test constrained ranking under predefined non-inferiority margins. Future experiments should also test whether repeated exposure changes preferences, using randomized holdout groups and preference measures that do not rely on the recommender's own interaction data. Other priorities include developing counterfactual exposure tests for journalism, creator-economy, and platform-native settings; validating source, topic, and viewpoint classifiers across languages and contexts; testing whether users understand AI disclosures; creating a reusable protocol for auditing reward specifications and studying how the process of choosing benchmarks can gain democratic legitimacy.

Future research should also consider whether users should have rights over recommender objectives, the behavioral signals a system may use, and access to third-party curation, complementing existing rights related to transparency, non-profiling, profiling, and interoperability (European Union 2016, 2022a, 2022b). Designing such an entitlement would require a separate legal analysis of who holds the right, who has the duty to comply, which services are covered, what exceptions apply, how the right is enforced, and what remedies are available.

10. Conclusions

Two terms organize the argument and are not interchangeable. *Adaptive* curation covers any arrangement in which machine-learned policies adjust news exposure in response to feedback; *agentic* curation is its high-autonomy form, marked by persistent planning, delegated authority to act, coordination across curation functions, and delayed external review. Adaptive news curation should be understood as a policy that distributes exposure over time rather than as a series of separate recommendations, and it must be evaluated accordingly. The functional architecture developed here identifies where important choices enter that policy; cumulative exposure allocation measures the concentration, breadth, repetition, persistence, and disparity that build up under it and the accountability structure identifies who must answer, to which forum, using what evidence, under which standard, and with what possible remedy. This governance problem already exists before systems become fully agentic; agentic curation makes it more difficult by increasing the number of functions a system can coordinate, the authority it can exercise, and the time it can operate before review. The legal comparison shows that existing regimes provide some parts of the proposed accountability structure and leave others incomplete. The framework makes the governance of adaptive visibility more concrete, measurable, and institutionally assignable, without treating one metric, technical model, or legal system as a complete measure of democratic adequacy.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable. This study involved no human participants or animals, so ethics-committee approval was not required.

Informed Consent Statement: Not applicable.

Data Availability Statement: This study did not generate or analyze a new empirical dataset. Appendix A documents the search strategy, inclusion and exclusion criteria, and approximate composition of the literature base; Appendix B summarizes the status of the empirical evidence and Appendix C links the framework's main concepts to their supporting sources.

Acknowledgments: The author used a generative AI assistant to support drafting, structural organization, and language editing. The author selected and verified all sources, developed all claims and interpretations, reviewed and edited the generated material, and accepts full responsibility for the content of this publication.

Conflicts of Interest: The author declares no conflict of interest.

Appendix A. Corpus Search Strategy

The synthesis used targeted searches of Scopus, Web of Science, the ACM Digital Library, and arXiv, together with the reference lists of key publications (backward snowballing). Legal instruments were collected from official sources, including EUR-Lex and the European Commission. The main searches covered material available through June 2026, and a supplementary search in July 2026 focused on agentic AI and performative optimization. Table A1 lists representative search strings for each body of literature; the July update informed the agentic-capacity dimensions of Section 2.3 and the discussion of performative optimization in Section 5.

The review prioritized peer-reviewed articles and books, as well as foundational theoretical works regardless of publication date. Preprints were included when they reported primary technical results not yet available in peer-reviewed form, and the text identifies them as preprints. Unsourced commentary and vendor marketing were excluded, except when needed to document a current industry development, and those exceptions are identified in the text. Material based on leaked internal company documents was used only in

the form already placed in the public domain by established news organizations and public regulatory records; no confidential document was independently requested, obtained, or republished. The empirical research discussed in Section 6, including work on filter bubbles, radicalization, and misinformation, was found mainly through platformization and news-recommendation searches and through backward searches from the major causal studies in each debate.

Table A1. Representative search strings for each body of literature. The strings were combined with Boolean operators and adjusted to the syntax of each database.

Body of Literature	Representative Search String
Gatekeeping and media effects	("gatekeeping" OR "gatewatching" OR "agenda-setting" OR "framing") AND ("news" OR "journalism")
Platformization and algorithmic distribution	("platformization" OR "platform power" OR "algorithmic") AND ("news distribution" OR "media")
News recommendation and personalization	("news recommend*" OR "personaliz* news") AND ("diversity" OR "editorial values" OR "democratic")
Reinforcement learning and recommenders	("reinforcement learning" OR "Markov decision process" OR "bandit") AND ("recommend*" OR "feed ranking" OR "long-term engagement")
Accountability, auditing, and governance	("algorithmic accountability" OR "algorithm audit*" OR "responsibility gap" OR "AI governance") AND ("recommend*" OR "platform" OR "media")
Agentic AI and performative optimization (July 2026 pass)	("agentic AI" OR "AI agent*" OR "autonomous agent*" OR "tool use" OR "multi-step planning" OR "performative reinforcement learning") AND ("recommend*" OR "curation" OR "news" OR "media")

Because this was an integrative rather than a systematic review, records were screened for relevance by title and abstract instead of being counted under a fixed protocol; the searches returned on the order of several hundred candidate records. Table A2 gives an approximate breakdown of the material retained for analysis. The categories overlap, the counts are rounded, and some sources contribute to more than one category; the bibliography also includes contextual and legal materials, so the category counts do not sum to one unique total.

Table A2. Approximate and overlapping composition of the literature base. Counts are rounded and categories overlap, reflecting an integrative review rather than an enumerated systematic protocol.

Body of Literature	Approximate Sources Retained
Foundational research on gatekeeping and media effects	12
Platformization, algorithmic news distribution, and measured media effects	28
News recommendation and personalization	12
Reinforcement learning and recommender systems	12
Accountability, auditing, legal governance, and legal instruments	23

Appendix B. Evidence-Status Summary

Table A3 identifies the strongest type of evidence available for each empirical claim discussed in Section 6, so that readers can distinguish claims supported by causal experiments from claims based on audits, observational panels, or industry reporting. Claims supported only by contested industry reporting are marked as such; the conceptual framework does not depend on them.

Table A3. Evidence status of the empirical claims engaged in the argument. The final column explains how each claim contributes to the argument, not a final judgment on the underlying phenomenon.

Claim	Strongest Available Evidence	Example Sources	Use in the Argument
Algorithmic ranking traps users in filter bubbles	Causal field experiments and large observational studies	Guess et al. (2023) ; Nyhan et al. (2023) ; Bakshy et al. (2015)	Contested; treated as an interaction among ranking, networks, and user choice, not as a settled effect
Recommenders cause radicalization	Audits, observational panels, and causal studies using counterfactual bots	Ribeiro et al. (2020) ; Hosseinmardi et al. (2021, 2024)	Recent YouTube studies find no uniform average effect; effects on subgroups remain unresolved
False news spreads farther and faster than true news	Large observational study	Vosoughi et al. (2018)	Strong evidence about diffusion, but not evidence that recommenders caused it
Engagement weighting raised the visibility of divisive content on one platform	Industry reporting based on leaked internal documents	Hagey and Horwitz (2021) ; Merrill and Oremus (2021)	Well documented in reporting but disputed by the company
Bounded personalization can improve utility while broadening exposure	Peer-reviewed field experiment (A/B test)	Holzleitner et al. (2026)	Evidence from one site and a non-RL system

Appendix C. Source-to-Construct Matrix

Table A4 links the framework's main concepts to representative sources, so that the intellectual basis of each part of the framework can be traced.

Table A4. Main concepts in the framework and representative sources that inform them.

Construct	Role in the Argument	Representative Sources
Movement of gatekeeping power	Provides the historical and theoretical foundation (Section 2)	(Lewin 1947; White 1950; Shoemaker and Vos 2009; Barzilai-Nahon 2008; Bruns 2018; Wallace 2018; Voinea 2025)
Algorithmic gatekeeping and platform power	Explains the shift toward algorithmic distribution (Section 2)	(Gillespie 2014, 2018; Bucher 2012; Napoli 2014; Just and Latzer 2017; DeVito 2017; Nieborg and Poell 2018; van Dijck et al. 2018)
Dimensions of agentic capacity	Describes how operational autonomy can increase (Section 2.3)	(Shavit et al. 2023)
Curation functions and feedback layer	Provides the structural model (Section 4)	(Shoemaker and Vos 2009; Gillespie 2018; Chen et al. 2019; Ie et al. 2019; Jiang et al. 2019)
Sequential optimization through an MDP	Explains the technical mechanism (Section 5)	(Sutton and Barto 2018; Afsar et al. 2022; Chen et al. 2023; Zou et al. 2019; Ie et al. 2019)
Performativity and preference change	Explains how a deployed policy can change its own environment (Section 5)	(Perdomo et al. 2020; Mandal et al. 2023; Jiang et al. 2019)
Misspecified rewards	Explains why engagement proxies may conflict with public goals (Section 5)	(Amodei et al. 2016; Karwowski et al. 2024; Krakovna et al. 2020; Stray et al. 2021, 2024)
Disputed evidence about harms	Defines what the empirical literature does and does not establish (Section 6)	(Bakshy et al. 2015; Flaxman et al. 2016; Guess et al. 2023; Nyhan et al. 2023; Hosseinmardi et al. 2021, 2024; Ribeiro et al. 2020; Munger and Phillips 2022; Vosoughi et al. 2018)
Cumulative exposure allocation	Provides the unit for measuring exposure over time (Section 6.2)	(Helberger 2019; Vrijenhoek et al. 2021; Bauer et al. 2024; McCombs and Shaw 1972; Zuboff 2019)
Bounded personalization	Provides an illustrative field example (Section 6.3)	(Holzleitner et al. 2026; Lu et al. 2020; Mattis et al. 2024)
Accountability and responsibility gaps	Provides the accountability structure (Section 7)	(Bovens 2007; Wieringa 2020; Matthias 2004; Santoni de Sio and Mecacci 2021; Floridi 2016; Ananny and Crawford 2018; Burrell 2016; Pasquale 2015; Diakopoulos 2015)
Auditing methods	Provides diagnostics for oversight quality (Section 7)	(Sandvig et al. 2014; Metaxa et al. 2021; Raji et al. 2020, 2022; Birhane et al. 2024)
Legal and regulatory instruments	Connects the framework to EU and U.S. law (Section 8)	(European Union 2016, 2022a, 2022b, 2024a, 2024b; Gonzalez v. Google LLC 2023; Twitter, Inc. v. Taamneh 2023; Moody v. NetChoice, LLC 2024)

References

- Afsar, Mohammad Mehdi, Trafford Crump, and Behrouz Far. 2022. Reinforcement Learning Based Recommender Systems: A Survey. *ACM Computing Surveys* 55: 145. [CrossRef]
- Amodei, Dario, Chris Olah, Jacob Steinhardt, Paul Christiano, John Schulman, and Dan Mané. 2016. Concrete Problems in AI Safety. *arXiv*. [CrossRef]
- Ananny, Mike, and Kate Crawford. 2018. Seeing without Knowing: Limitations of the Transparency Ideal and Its Application to Algorithmic Accountability. *New Media & Society* 20: 973–89. [CrossRef]
- Anderson v. TikTok, Inc., 116 F.4th 180. 3d Cir. 2024. Available online: <https://law.justia.com/cases/federal/appellate-courts/ca3/22-3061/22-3061-2024-08-27.html> (accessed on 15 March 2026).
- Bakshy, Eytan, Solomon Messing, and Lada A. Adamic. 2015. Exposure to Ideologically Diverse News and Opinion on Facebook. *Science* 348: 1130–32. [CrossRef] [PubMed]
- Barzilai-Nahon, Karine. 2008. Toward a Theory of Network Gatekeeping: A Framework for Exploring Information Control. *Journal of the American Society for Information Science and Technology* 59: 1493–512. [CrossRef]
- Bauer, Christina, Chirag Bagchi, Onur A. Hundogan, and Karin van Es. 2024. Where Are the Values? A Systematic Literature Review on News Recommender Systems. *ACM Transactions on Recommender Systems* 2: 23. [CrossRef]
- Biega, Asia J., Krishna P. Gummadi, and Gerhard Weikum. 2018. Equity of Attention: Amortizing Individual Fairness in Rankings. Paper presented at the 41st International ACM SIGIR Conference on Research & Development in Information Retrieval, Ann Arbor, MI, USA, July 8–12; pp. 405–14. [CrossRef]
- Birhane, Abeba, Ryan Steed, Victor Ojewale, Briana Vecchione, and Inioluwa Deborah Raji. 2024. AI Auditing: The Broken Bus on the Road to AI Accountability. In *2024 IEEE Conference on Secure and Trustworthy Machine Learning (SaTML)*, Toronto, ON, Canada, April 9–11. Piscataway: IEEE, pp. 612–43. [CrossRef]
- Bovens, Mark. 2007. Analysing and Assessing Accountability: A Conceptual Framework. *European Law Journal* 13: 447–68. [CrossRef]
- Bruns, Axel. 2018. *Gatewatching and News Curation: Journalism, Social Media, and the Public Sphere*. Lausanne: Peter Lang. [CrossRef]
- Bruns, Axel. 2019. *Are Filter Bubbles Real?* Cambridge: Polity.
- Bucher, Taina. 2012. Want to Be on the Top? Algorithmic Power and the Threat of Invisibility on Facebook. *New Media & Society* 14: 1164–80. [CrossRef]
- Bucher, Taina. 2018. *If...Then: Algorithmic Power and Politics*. Oxford: Oxford University Press. [CrossRef]
- Burrell, Jenna. 2016. How the Machine ‘Thinks’: Understanding Opacity in Machine Learning Algorithms. *Big Data & Society* 3: 2053951715622512. [CrossRef]
- Chen, Minmin, Alex Beutel, Paul Covington, Sagar Jain, Francois Belletti, and Ed H. Chi. 2019. Top-K Off-Policy Correction for a REINFORCE Recommender System. Paper presented at Twelfth ACM International Conference on Web Search and Data Mining, Melbourne, Australia, February 11–15; pp. 456–64. [CrossRef]
- Chen, Xu, Lina Yao, Julian McAuley, Guangyi Zhou, and Xianzhi Wang. 2023. Deep Reinforcement Learning in Recommender Systems: A Survey and New Perspectives. *Knowledge-Based Systems* 264: 110335. [CrossRef]
- DeVito, Michael A. 2017. From Editors to Algorithms: A Values-Based Approach to Understanding Story Selection in the Facebook News Feed. *Digital Journalism* 5: 753–73. [CrossRef]
- Diakopoulos, Nicholas. 2015. Algorithmic Accountability: Journalistic Investigation of Computational Power Structures. *Digital Journalism* 3: 398–415. [CrossRef]
- Diaz, Fernando, Bhaskar Mitra, Michael D. Ekstrand, Asia J. Biega, and Ben Carterette. 2020. Evaluating Stochastic Rankings with Expected Exposure. Paper presented at 29th ACM International Conference on Information & Knowledge Management, Galway, Ireland, October 19; pp. 275–84. [CrossRef]
- Entman, Robert M. 1993. Framing: Toward Clarification of a Fractured Paradigm. *Journal of Communication* 43: 51–58. [CrossRef]
- European Union. 2016. Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the Protection of Natural Persons with Regard to the Processing of Personal Data and on the Free Movement of Such Data (General Data Protection Regulation). *Official Journal of the European Union* L 119/1. Available online: <https://eur-lex.europa.eu/eli/reg/2016/679/oj> (accessed on 15 March 2026).
- European Union. 2022a. Regulation (EU) 2022/1925 of the European Parliament and of the Council of 14 September 2022 on Contestable and Fair Markets in the Digital Sector and Amending Directives (EU) 2019/1937 and (EU) 2020/1828 (Digital Markets Act). *Official Journal of the European Union* L 265/1. Available online: <https://eur-lex.europa.eu/eli/reg/2022/1925/oj> (accessed on 15 March 2026).
- European Union. 2022b. Regulation (EU) 2022/2065 of the European Parliament and of the Council of 19 October 2022 on a Single Market for Digital Services and Amending Directive 2000/31/EC (Digital Services Act). *Official Journal of the European Union* L 277/1. Available online: <https://eur-lex.europa.eu/eli/reg/2022/2065/oj> (accessed on 15 March 2026).

- European Union. 2024a. Regulation (EU) 2024/1083 of the European Parliament and of the Council of 11 April 2024 Establishing a Common Framework for Media Services in the Internal Market and Amending Directive 2010/13/EU (European Media Freedom Act). *Official Journal of the European Union* L 2024/1083. Available online: <https://eur-lex.europa.eu/eli/reg/2024/1083/oj> (accessed on 15 March 2026).
- European Union. 2024b. Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act). *Official Journal of the European Union* L 2024/1689. Available online: <https://eur-lex.europa.eu/eli/reg/2024/1689/oj> (accessed on 15 March 2026).
- Flaxman, Seth, Sharad Goel, and Justin M. Rao. 2016. Filter Bubbles, Echo Chambers, and Online News Consumption. *Public Opinion Quarterly* 80: 298–320. [CrossRef]
- Floridi, Luciano. 2016. Faultless Responsibility: On the Nature and Allocation of Moral Responsibility for Distributed Moral Actions. *Philosophical Transactions of the Royal Society A* 374: 20160112. [CrossRef] [PubMed]
- Galtung, Johan, and Mari Holmboe Ruge. 1965. The Structure of Foreign News. *Journal of Peace Research* 2: 64–91. [CrossRef]
- Gillespie, Tarleton. 2014. The Relevance of Algorithms. In *Media Technologies: Essays on Communication, Materiality, and Society*. Edited by Tarleton Gillespie, Pablo J. Boczkowski and Kirsten A. Foot. Cambridge, MA: MIT Press, pp. 167–94. [CrossRef]
- Gillespie, Tarleton. 2018. *Custodians of the Internet: Platforms, Content Moderation, and the Hidden Decisions That Shape Social Media*. New Haven: Yale University Press. [CrossRef]
- Gonzalez v. Google LLC, 598 U.S. 617. 2023. Available online: https://www.supremecourt.gov/opinions/22pdf/21-1333_6j7a.pdf (accessed on 15 March 2026).
- Guess, Andrew M., Neil Malhotra, Jennifer Pan, Pablo Barberá, Hunt Allcott, Taylor Brown, Adriana Crespo-Tenorio, Drew Dimmery, Deen Freelon, Matthew Gentzkow, and et al. 2023. How Do Social Media Feed Algorithms Affect Attitudes and Behavior in an Election Campaign? *Science* 381: 398–404. [CrossRef] [PubMed]
- Habermas, Jürgen. 2022. Reflections and Hypotheses on a Further Structural Transformation of the Political Public Sphere. *Theory, Culture & Society* 39: 145–71. [CrossRef]
- Hagey, Keach, and Jeff Horwitz. 2021. Facebook Tried to Make Its Platform a Healthier Place. It Got Angrier Instead. *The Wall Street Journal* 15. Available online: https://www.wsj.com/tech/facebook-algorithm-change-zuckerberg-11631654215?mod=article_inline (accessed on 15 March 2026).
- Helberger, Natali. 2019. On the Democratic Role of News Recommenders. *Digital Journalism* 7: 993–1012. [CrossRef]
- Holzleitner, Markus, Stefan Leitner, Helle L. Jorgensen, Christoph Schmitz, Johan Welander, and Dietmar Jannach. 2026. Controlled Personalization in Legacy Media Online Services: A Case Study in News Recommendation. *ACM Transactions on Recommender Systems*. [CrossRef]
- Hosseinmardi, Homa, Amir Ghasemian, Aaron Clauset, Markus Mobius, David M. Rothschild, and Duncan J. Watts. 2021. Examining the Consumption of Radical Content on YouTube. *Proceedings of the National Academy of Sciences* 118: e2101967118. [CrossRef] [PubMed]
- Hosseinmardi, Homa, Amir Ghasemian, Miguel Rivera-Lanas, Manoel Horta Ribeiro, Robert West, and Duncan J. Watts. 2024. Causally Estimating the Effect of YouTube’s Recommender System Using Counterfactual Bots. *Proceedings of the National Academy of Sciences* 121: e2313377121. [CrossRef] [PubMed]
- Ie, Eugene, Vihan Jain, Jing Wang, Sanjev Narvekar, Ritesh Agarwal, Rui Wu, Heng-Tze Cheng, Tushar Chandra, and Craig Boutilier. 2019. SlateQ: A Tractable Decomposition for Reinforcement Learning with Recommendation Sets. Paper presented at Twenty-Eighth International Joint Conference on Artificial Intelligence, Macao, China, August 10; pp. 2592–99. [CrossRef] [PubMed]
- Jiang, Ray, Silvia Chiappa, Tor Lattimore, András György, and Pushmeet Kohli. 2019. Degenerate Feedback Loops in Recommender Systems. Paper presented at 2019 AAAI/ACM Conference on AI, Ethics, and Society, Honolulu, HI, USA, January 27; pp. 383–90. [CrossRef]
- Just, Natascha, and Michael Latzer. 2017. Governance by Algorithms: Reality Construction by Algorithmic Selection on the Internet. *Media, Culture & Society* 39: 238–58. [CrossRef]
- Karwowski, Jan, Oliver Hayman, Xiaoyu Bai, Klaus Kiendlhofer, Charlie Griffin, and Joar Skalse. 2024. Goodhart’s Law in Reinforcement Learning. Paper presented at Twelfth International Conference on Learning Representations (ICLR), Vienna, Austria, May 7. [CrossRef]
- Kitchin, Rob. 2017. Thinking Critically about and Researching Algorithms. *Information, Communication & Society* 20: 14–29. [CrossRef]
- Krakovna, Victoria, Jonathan Uesato, Vladimir Mikulik, Matthew Rahtz, Tom Everitt, Ramana Kumar, Zachary Kenton, Jan Leike, and Shane Legg. 2020. Specification Gaming: The Flip Side of AI Ingenuity. DeepMind Blog. Available online: <https://deepmind.google/discover/blog/specification-gaming-the-flip-side-of-ai-ingenuity/> (accessed on 15 March 2026).
- Lewin, Kurt. 1947. Frontiers in Group Dynamics: II. Channels of Group Life, Social Planning and Action Research. *Human Relations* 1: 143–53. [CrossRef]

- Lu, Fan, Alexandra Dumitrache, and David Graus. 2020. Beyond Optimizing for Clicks: Incorporating Editorial Values in News Recommendation. Paper presented at 28th ACM Conference on User Modeling, Adaptation and Personalization, Genoa, Italy, July 12; pp. 145–53. [\[CrossRef\]](#)
- Mandal, Debmalya, Sreejith Triantafyllou, and Goran Radanovic. 2023. Performative Reinforcement Learning. Paper presented at 40th International Conference on Machine Learning, PMLR 202, Honolulu, HI USA, July 23; pp. 23642–80.
- Mangold, Frank, David Schoch, and Sebastian Stier. 2024. Ideological Self-Selection in Online News Exposure: Evidence from Europe and the US. *Science Advances* 10: eadg9287. [\[CrossRef\]](#) [\[PubMed\]](#)
- Matthias, Andreas. 2004. The Responsibility Gap: Ascribing Responsibility for the Actions of Learning Automata. *Ethics and Information Technology* 6: 175–83. [\[CrossRef\]](#)
- Mattis, Nadja, Philipp K. Masur, Judith Möller, and Wouter van Atteveldt. 2024. Nudging towards News Diversity: A Theoretical Framework for Facilitating Diverse News Consumption through Recommender Design. *New Media & Society* 26: 3681–706. [\[CrossRef\]](#)
- McCombs, Maxwell E., and Donald L. Shaw. 1972. The Agenda-Setting Function of Mass Media. *Public Opinion Quarterly* 36: 176–87. [\[CrossRef\]](#)
- Meese, James, and Edward Hurcombe. 2021. Facebook, News Media and Platform Dependency: The Institutional Impacts of News Distribution on Social Platforms. *New Media & Society* 23: 2367–84. [\[CrossRef\]](#)
- Merrill, Jeremy B., and Will Oremus. 2021. Five Points for Anger, One for a ‘Like’: How Facebook’s Formula Fostered Rage and Misinformation. *The Washington Post*, October 26. Available online: <https://www.washingtonpost.com/technology/2021/10/26/facebook-angry-emoji-algorithm/> (accessed on 15 March 2026).
- Metaxa, Danaë, Joon Sung Park, Ronald E. Robertson, Karrie Karahalios, Christo Wilson, Jeff Hancock, and Christian Sandvig. 2021. Auditing Algorithms: Understanding Algorithmic Systems from the Outside In. *Foundations and Trends in Human–Computer Interaction* 14: 272–344. [\[CrossRef\]](#)
- Mittelstadt, Brent Daniel, Patrick Allo, Mariarosaria Taddeo, Sandra Wachter, and Luciano Floridi. 2016. The Ethics of Algorithms: Mapping the Debate. *Big Data & Society* 3: 1–21. [\[CrossRef\]](#)
- Moody v. NetChoice, LLC, 603 U.S. 707. 2024. Available online: <https://supreme.justia.com/cases/federal/us/603/22-277/> (accessed on 15 March 2026).
- Munger, Kevin, and Joseph Phillips. 2022. Right-Wing YouTube: A Supply and Demand Perspective. *International Journal of Press/Politics* 27: 186–219. [\[CrossRef\]](#)
- Napoli, Philip M. 2014. Automated Media: An Institutional Theory Perspective on Algorithmic Media Production and Consumption. *Communication Theory* 24: 340–60. [\[CrossRef\]](#)
- Newman, Nic. 2025. Overview and Key Findings of the 2025 Digital News Report. Reuters Institute for the Study of Journalism. Available online: <https://reutersinstitute.politics.ox.ac.uk/digital-news-report/2025/dnr-executive-summary> (accessed on 15 March 2026).
- Nieborg, David B., and Thomas Poell. 2018. The Platformization of Cultural Production: Theorizing the Contingent Cultural Commodity. *New Media & Society* 20: 4275–92. [\[CrossRef\]](#)
- Nyhan, Brendan, Jaime Settle, Emily Thorson, Magdalena Wojcieszak, Pablo Barberá, Annie Y. Chen, Hunt Allcott, Taylor Brown, Adriana Crespo-Tenorio, Drew Dimmery, and et al. 2023. Like-Minded Sources on Facebook Are Prevalent but Not Polarizing. *Nature* 620: 137–44. [\[CrossRef\]](#) [\[PubMed\]](#)
- Pariser, Eli. 2011. *The Filter Bubble: What the Internet Is Hiding from You*. New York: Penguin Press.
- Pasquale, Frank. 2015. *The Black Box Society: The Secret Algorithms That Control Money and Information*. Cambridge, MA: Harvard University Press. [\[CrossRef\]](#)
- Perdomo, Juan C., Tijana Zrnic, Celestine Mendler-Dünner, and Moritz Hardt. 2020. Performative Prediction. Paper presented at 37th International Conference on Machine Learning, PMLR 119, Virtual, July 13; pp. 7599–609.
- Raji, Inioluwa Deborah, Andrew Smart, Rebecca N. White, Margaret Mitchell, Timnit Gebru, Ben Hutchinson, Jamila Smith-Loud, Daniel Theron, and Parker Barnes. 2020. Closing the AI Accountability Gap: Defining an End-to-End Framework for Internal Algorithmic Auditing. Paper presented at 2020 Conference on Fairness, Accountability, and Transparency, Barcelona, Spain, January 27; pp. 33–44. [\[CrossRef\]](#)
- Raji, Inioluwa Deborah, Peggy Xu, Colleen Honigsberg, and Daniel E. Ho. 2022. Outsider Oversight: Designing a Third Party Audit Ecosystem for AI Governance. Paper presented at 2022 AAAI/ACM Conference on AI, Ethics, and Society, Oxford, UK, August 1; pp. 557–71. [\[CrossRef\]](#)
- Ribeiro, Manoel Horta, Raphael Ottoni, Robert West, Virgílio A. F. Almeida, and Wagner Meira, Jr. 2020. Auditing Radicalization Pathways on YouTube. Paper presented at 2020 Conference on Fairness, Accountability, and Transparency, Barcelona, Spain, January 27; pp. 131–41. [\[CrossRef\]](#)

- Sandvig, Christian, Kevin Hamilton, Karrie Karahalios, and Cedric Langbort. 2014. Auditing Algorithms: Research Methods for Detecting Discrimination on Internet Platforms. Paper presented at the International Communication Association preconference Data and Discrimination, Seattle, WA, USA, May 22.
- Santoni de Sio, Filippo, and Giulio Mecacci. 2021. Four Responsibility Gaps with Artificial Intelligence: Why They Matter and How to Address Them. *Philosophy & Technology* 34: 1057–84. [CrossRef]
- Seaver, Nick. 2017. Algorithms as Culture: Some Tactics for the Ethnography of Algorithmic Systems. *Big Data & Society* 4: 1–12. [CrossRef]
- Seaver, Nick. 2022. *Computing Taste: Algorithms and the Makers of Music Recommendation*. Chicago: University of Chicago Press. [CrossRef]
- Shavit, Yonadav, Sandhini Agarwal, Miles Brundage, Steven Adler, Cullen O’Keefe, Rosie Campbell, and Ted Lee. 2023. Practices for Governing Agentic AI Systems. OpenAI. Available online: <https://openai.com/research/practices-for-governing-agentic-ai-systems> (accessed on 15 March 2026).
- Shoemaker, Pamela J., and Tim P. Vos. 2009. *Gatekeeping Theory*. Abingdon-on-Thames: Routledge. [CrossRef]
- Singh, Ashudeep, and Thorsten Joachims. 2018. Fairness of Exposure in Rankings. Paper presented at 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, London, UK, August 19; pp. 2219–28. [CrossRef]
- Snyder, Hannah. 2019. Literature Review as a Research Methodology: An Overview and Guidelines. *Journal of Business Research* 104: 333–39. [CrossRef]
- Stray, Jonathan, Alon Halevy, Parisa Assar, Dylan Hadfield-Menell, Craig Boutilier, Amar Ashar, Lex Beattie, Michael Ekstrand, Claire Leibowicz, Connie Moon Sehat, and et al. 2024. Building Human Values into Recommender Systems: An Interdisciplinary Synthesis. *ACM Transactions on Recommender Systems* 2: 1–57. [CrossRef]
- Stray, Jonathan, Ivan Vendrov, Jeremy Nixon, Steven Adler, and Dylan Hadfield-Menell. 2021. What Are You Optimizing For? Aligning Recommender Systems with Human Values. *arXiv*. [CrossRef]
- Sunstein, Cass R. 2001. *Republic.com*. Princeton: Princeton University Press.
- Sutton, Richard S., and Andrew G. Barto. 2018. *Reinforcement Learning: An Introduction*, 2nd ed. Cambridge, MA: MIT Press.
- Tandoc, Edson C., Jr., and Ryan J. Thomas. 2015. The Ethics of Web Analytics: Implications of Using Audience Metrics in News Construction. *Digital Journalism* 3: 243–58. [CrossRef]
- Tenor, Carl. 2024. Metrics as the New Normal: Exploring the Evolution of Audience Metrics as a Decision-Making Tool in Swedish Newsrooms 1995–2022. *Journalism* 25: 1111–29. [CrossRef]
- Thurman, Neil. 2011. Making ‘The Daily Me’: Technology, Economics and Habit in the Mainstream Assimilation of Personalized News. *Journalism* 12: 395–415. [CrossRef]
- Twitter, Inc. v. Taamneh, 598 U.S. 471. 2023. Available online: https://www.supremecourt.gov/opinions/22pdf/21-1496_d18f.pdf (accessed on 15 March 2026).
- van Dijck, José, Thomas Poell, and Martijn de Waal. 2018. *The Platform Society: Public Values in a Connective World*. Oxford: Oxford University Press. [CrossRef]
- van Es, Karin, and Dong Nguyen. 2024. Balancing Needs and Values: A Multi-Stakeholder Examination of Algorithmic News Recommenders in the Netherlands. *Journalism Practice* 2024: 1–19. [CrossRef]
- Voinea, Denisa V. 2025. Reconceptualizing Gatekeeping in the Age of Artificial Intelligence: A Theoretical Exploration of Artificial Intelligence-Driven News Curation and Automated Journalism. *Journalism and Media* 6: 68. [CrossRef]
- Vosoughi, Soroush, Deb Roy, and Sinan Aral. 2018. The Spread of True and False News Online. *Science* 359: 1146–51. [CrossRef] [PubMed]
- Vrijenhoek, Sanne, Mesut Kaya, Nadia Metoui, Judith Möller, Daan Odijk, and Natali Helberger. 2021. Recommenders with a Mission: Assessing Diversity in News Recommendations. Paper presented at 2021 Conference on Human Information Interaction and Retrieval, Canberra, Australia, March 14; pp. 173–83. [CrossRef]
- Wallace, Julian. 2018. Modelling Contemporary Gatekeeping: The Rise of Individuals, Algorithms and Platforms in Digital News Dissemination. *Digital Journalism* 6: 274–93. [CrossRef]
- White, David Manning. 1950. The ‘Gate Keeper’: A Case Study in the Selection of News. *Journalism Quarterly* 27: 383–90. [CrossRef]
- Wieringa, Maranke. 2020. What to Account for When Accounting for Algorithms: A Systematic Literature Review on Algorithmic Accountability. Paper presented at 2020 Conference on Fairness, Accountability, and Transparency, Barcelona, Spain, January 27; pp. 1–18. [CrossRef]

- Zou, Lixin, Long Xia, Zhuoye Ding, Jiaxing Song, Weidong Liu, and Dawei Yin. 2019. Reinforcement Learning to Optimize Long-Term User Engagement in Recommender Systems. Paper presented at 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, Anchorage, AK, USA, August 4; pp. 2810–18. [[CrossRef](#)]
- Zuboff, Shoshana. 2019. *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*. New York: PublicAffairs.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.