PREPRINT · REVISION 5 · JUNE 2026

The Bias Amplification Model

A Mathematical Framework for Quantifying Cognitive Distortion
Under Conditions of Information Overload

$$\Phi(\Omega) = \Phi_0 + f(\Omega) \cdot \bar{A}$$

AUTHORS	CORRESPONDENCE	INTERACTIVE MODEL
© Ron Franklin, Timothy Lewis San Diego, California June 2026	research@attentionplease.ai	attentionplease.ai/bam

The Bias Amplification Model: A Mathematical Framework for Quantifying Cognitive Distortion Under Conditions of Information Overload

Ron Franklin and Timothy Lewis

San Diego, California

2026 · Correspondence: research@attentionplease.ai

Keywords: cognitive overload, cognitive bias, attention economy, dual-process theory, information processing, decision quality, cognitive architecture, metacognition, Phi Framework

ISBN: *979-8-9992593-7-0*

Preprint · Revision 5 · June 2026

This version incorporates revisions responding to post-publication reader feedback and an expanded theoretical review. The model's core mathematical architecture and parameter values are unchanged. Revision 5 substantially restructures the treatment of cognitive processing capacity C, introduces a Metacognitive Monitoring variable, adds Section 2.6 on the cognitive architecture literature, adds Section 3.4 on individual differences and moderating variables, and adds Section 5.5 on prescriptive interventions. A more systematic taxonomy of claim types is applied throughout.

Abstract

The relationship between information load and cognitive bias has been theorized but not formally modeled. This paper introduces the Bias Amplification Model (BAM), a mathematical framework that quantifies how elevated information-processing demands systematically amplify seven well-documented cognitive biases. The model's core variable, the Cognitive Overload Ratio (Ω = I/C), expresses the ratio of effective information input to cognitive processing capacity. As Ω approaches and exceeds 1.0, a sigmoid amplification function $f(\Omega)$ drives nonlinear increases in seven bias intensities — Confirmation, Availability, Substitution, Recency, Anchoring, Representativeness, and Framing — aggregated into the Phi Index (Φ), a composite measure of overall cognitive distortion computed as the normalized weighted sum $\Phi = \Sigma\, w_i \cdot B_i(\Omega)$. Revision 5 substantially extends the model's theoretical account in two directions. First, cognitive processing capacity C is reconceptualized as a composite of biological baseline capacity, schema-augmented capacity arising from expertise and trained mental models, and externalized capacity arising from deliberate information organization systems that offload processing from working memory to trusted external scaffolds. This reconceptualization expands the model's prescriptive reach significantly: rather than offering only two Omega-reduction levers (reduce I; improve sleep and exercise), the model now identifies five specific intervention pathways, three of which are directly trainable. Second, metacognitive monitoring capacity is introduced as a moderating variable that determines whether the model's other interventions are functionally available to the individual — since the central phenomenological finding of the BAM is that cognitive overload is subjectively invisible to the person experiencing it. The model is offered as a theoretical framework for empirical validation and as a practical diagnostic and prescriptive instrument for cognitive state assessment and improvement.

Note on This Revision

Revision 1 (May 2026) incorporated six substantive additions addressing technical issues: explicit ordinal measurement framing for Ω; a sensitivity analysis for the tipping-steepness parameter k; a staged empirical calibration roadmap; characterization of the additive Φ as a conservative upper bound; clarification of Substitution Bias functional form asymmetry; and Recency Bias negative *A*/I boundary conditions.

Revision 2 incorporated six further additions responding to peer-level reader feedback: acknowledgment of the default-interventionist and predictive processing account of dual-process theory; clarification that C represents working memory resources required for System 2 reasoning; mechanism-level justification connecting each of the seven biases to overload-driven resource depletion; acknowledgment of constructivist accounts of emotional salience; clarification that the Phi Index represents aggregate distortion potential rather than simultaneous full activation; and engagement with the ego depletion literature.

Revision 3 incorporated eight textual revisions clarifying the System 1 / System 2 relationship throughout: reframing Ω as measuring the progressive depletion of System 2's corrective capacity rather than the onset of System 1 processing. No mathematical structure, parameter values, zone boundaries, or equations were altered.

Revision 4 added the Ω Portrait Framework (Section 3.2.1): a seven-level characterization of the information environments, behavioral signatures, quantitative proxy metrics, and phenomenological states corresponding to discrete positions on the Cognitive Overload Ratio scale.

Revision 5 (this version) incorporates five structural additions arising from an expanded review of the cognitive architecture, expertise, schema theory, metacognition, and information organization literatures. The mathematical architecture, parameter values, zone boundaries, and core predictions are unchanged. The Revision 5 additions are:

Section 2.6 — New subsection: The Cognitive Architecture Literature. Reviews the empirical basis for treating cognitive processing capacity as a composite that includes biological baseline, schema-augmented, and externalized components, drawing on cognitive offloading research, schema theory, expertise studies, and the metacognition literature.

Section 3.2 — Substantially expanded treatment of C. Reconceptualizes cognitive processing capacity as a composite of three interacting components — biological baseline capacity, schema-augmented capacity arising from expertise and trained mental models, and externalized capacity arising from deliberate information organization systems — while preserving the existing mathematical notation and equations. The effective I variable is similarly reconceptualized to reflect information quality dimensions beyond rate.

Section 3.4 — New subsection: Individual Differences and Moderating Variables. Introduces metacognitive monitoring capacity (M) as a moderating variable that determines whether the model's prescriptive interventions are functionally accessible to the individual at any given Ω level. Discusses need for cognition, expertise-based schema development, and information organization training as additional moderators of the overload-to-distortion pathway.

Section 5.5 — New subsection: Prescriptive Interventions — A Formal Account. Maps the model's variables to five specific, actionable intervention categories, grounding each in the theoretical framework and existing empirical support. Presents a structured approach to individual cognitive state management based on the model's composite C and effective I variables.

Throughout — More systematic taxonomy of claim types. Empirically established claims, model assumptions, and model predictions are more consistently distinguished throughout the text, addressing the gap between the mathematical precision of the model's notation and the theoretical rather than empirical status of its parameter values.

1. Introduction

In 1971, the economist Herbert Simon observed that "a wealth of information creates a poverty of attention" (Simon, 1971, p. 40).[1] Simon's formulation was a theoretical observation about organizational management in the pre-internet era. It was not, at the time, a crisis. Half a century later, it is. The average American adult now spends more than six and a half hours per day looking at screens, with more than two hours allocated to algorithmically curated social media platforms (DataReportal, 2024). The information environment has been deliberately engineered to maximize time-on-platform by exploiting documented properties of human cognitive architecture — a design strategy whose architects have confirmed in public testimony (Parker, 2017; Harris, 2017).

What is missing from the existing literature is not documentation of the problem but a formal account of its mechanism, and — crucially — a formal account of how human agency can modify that mechanism. We know that cognitive load degrades decision quality (Sweller, 1988). We know that System 1 heuristic thinking produces systematic errors (Tversky & Kahneman, 1974). We know that the contemporary information environment elevates cognitive load (Ward et al., 2017). What has not been proposed, to our knowledge, is a mathematical model that connects these three bodies of evidence into a quantitative framework capable of measuring the degree to which a given informational environment distorts cognition — and identifying with precision the pathways through which individuals can reduce that distortion.

This paper introduces the Bias Amplification Model (BAM) to fill that gap. The BAM proposes that cognitive overload does not simply degrade judgment in a general way — it amplifies specific, well-characterized cognitive biases in proportion to the degree of overload, following a sigmoid function with a sharp tipping point. Critically, the model also identifies cognitive processing capacity as a composite variable with multiple trainable components — not a fixed biological parameter — and metacognitive monitoring capacity as the gateway that makes intentional intervention possible at all.

The paper proceeds as follows. Section 2 reviews the theoretical antecedents, now extended to include the cognitive architecture literature on schema theory, expertise, cognitive offloading, and metacognitive monitoring. Section 3 presents the full mathematical formalization, including the reconceptualized treatment of C and the individual differences and moderating variables. Section 4 describes the four cognitive zones. Section 5 discusses implications for individual, organizational, public health, and technology policy contexts, and introduces a formal account of prescriptive interventions. Sections 6 and 7 address limitations and conclusions.

[1] Simon's observation appeared in "Designing Organizations for an Information-Rich World," in *Computers, Communication, and the Public Interest* (Johns Hopkins University Press, 1971). His formulation — that a surplus of information necessarily produces a deficit of attention — has proven remarkably durable.

2. Theoretical Antecedents

2.1. The Attention Economy

The concept of attention as an economic resource has a longer history than is commonly recognized. Simon's (1971) foundational observation established the theoretical framework. In 1997, Michael Goldhaber extended it into an explicit economic model, arguing that in an information-saturated environment, human attention becomes the scarce resource around which entire economic systems organize (Goldhaber, 1997).[2]

The subsequent two decades of platform development transformed Goldhaber's theoretical proposition into commercial reality at a scale he did not anticipate. The economic model that emerged — capturing user attention and converting it into advertising revenue — created structural incentives for platforms to maximize engagement at any cost to user cognitive welfare. The result is an information environment characterized by variable reward schedules (Parker, 2017), infinite scroll eliminating natural stopping points (Raskin, 2019), algorithmic amplification of emotionally arousing content (Brady et al., 2017), and AI-driven behavioral prediction systems that construct models of individual users capable of anticipating responses before users consciously register them (Zuboff, 2019).

The attention economy literature has generated substantial empirical documentation of its effects: reduced attention spans (Rosen et al., 2013), increased anxiety and depression among heavy users (Twenge et al., 2018), degraded performance on tasks requiring sustained concentration (Mark et al., 2008), and measurable changes in reading patterns and information consumption behaviors (Liu, 2005). What has been lacking is an integrative theoretical framework that links the environmental inputs — information load — to the specific cognitive outputs — distorted judgment — and identifies the human agency pathways through which that link can be interrupted.

[2] Goldhaber's 1997 paper in *First Monday* remains one of the most prescient documents of the digital era: "Attention is the real currency of the digital economy." The economic model he described became the dominant business logic of social media platforms without modification.

2.2. Dual-Process Theory

The theoretical foundation most directly relevant to the BAM is Kahneman's dual-process account of cognition (Kahneman, 2011). System 1 thinking is fast, automatic, and effortless; it operates through pattern recognition and heuristics and requires minimal cognitive resources. System 2 thinking is slow, deliberate, and effortful; it is responsible for analytical reasoning, complex evaluation, and the override of intuitive responses that are likely to be erroneous.

The key property of dual-process theory for the BAM is the conditions under which System 2's corrective capacity is depleted, leaving System 1's continuously active heuristic outputs progressively less scrutinized and more likely to drive behavior. Kahneman's WYSIATI principle — "what you see is all there is" — describes System 1's characteristic behavior: it builds the best available narrative from information currently accessible without registering the absence of information it does not have (Kahneman, 2011, p. 85).[3]

Under conditions of high cognitive load, the depletion of System 2's corrective capacity is not experienced as a degradation — yet its consequences are measurable. System 1's heuristic outputs are no longer subject to meaningful scrutiny before they reach behavior, and the biases embedded in those outputs grow proportionally stronger as corrective capacity falls. The individual experiences this with the same felt clarity and confidence that System 1 judgments always carry, with no internal signal that bias amplification is underway. Kahneman (2011) terms this subjective experience *cognitive ease* — the state of fluent, effortless processing that accompanies heuristic dominance and feels from the inside like clarity and certainty. Cognitive ease is not the absence of error; it is the absence of the signal that would prompt System 2 to intervene. This is precisely what makes metacognitive monitoring capacity — the ability to accurately detect one's own overload state — a critical moderating variable, and why it receives dedicated treatment in Section 3.4.

A note on alternative mechanistic accounts is warranted. The default-interventionist model (Evans & Stanovich, 2013; De Neys, 2012) and its neuroscientific elaboration in predictive processing frameworks (Clark, 2016; Barrett, 2017) propose that Type 1 processes operate as a perpetual predictive baseline, with Type 2 deliberative processing engaged selectively when prediction errors exceed a threshold. The BAM's mechanistic language is a terminological choice rather than a commitment that excludes this account. Critically, the BAM's core predictions — that elevated Ω degrades deliberate evaluative processing and increases heuristic reliance — are identical under both framings.

[3] Kahneman uses WYSIATI — "What You See Is All There Is" — to describe System 1's tendency to construct coherent narratives from available information without registering what is absent. This property is particularly consequential under cognitive overload, where the available information is algorithmically curated rather than representative.

2.3. Cognitive Load Theory and Working Memory Architecture

Sweller's (1988) cognitive load theory provides the second theoretical pillar of the BAM. Sweller established that learning difficulty and performance degradation are functions not primarily of the inherent complexity of a task but of the interaction between that complexity and the limited capacity of working memory. When total cognitive load exceeds working memory capacity, performance fails categorically — a threshold effect, not a gradual decline (Sweller, 1988, p. 260).[4]

Baddeley and Hitch's (1974) model of working memory architecture specifies the system within which this capacity limitation operates. Working memory's central executive has a functional capacity originally estimated at seven units of information, plus or minus two (Miller, 1956), and subsequently refined to approximately four meaningful chunks in more controlled experimental paradigms (Cowan, 2001). This constraint is not a deficiency; it reflects the metabolic cost of executive function. But it means that the human processing system is inherently limited in its ability to handle simultaneous information streams of the kind that characterize modern digital environments.[5]

Ward et al. (2017) demonstrated the real-world implications of this constraint with particular clarity: the mere presence of a smartphone on a desk — powered off, face down, silent — reduced performance on tasks requiring full use of working memory capacity by approximately 10%, compared to conditions in which the device was in another room.[6]

[4] Sweller distinguished three types of cognitive load: intrinsic (inherent complexity), extraneous (complexity introduced by poor design), and germane (productive effort that builds schemas). The modern information environment introduces extraneous load at a scale Sweller could not have anticipated.

[5] Baddeley revised the model significantly in subsequent decades, adding an episodic buffer in 2000. The critical finding for the BAM is the capacity limitation of the central executive — the system responsible for deliberate, effortful reasoning.

[6] Ward et al. (2017) conducted three experiments across 520 smartphone users. Participants performed significantly better on tasks when their phones were in another room compared to face-down on the desk — even when they reported not thinking about their phones. The mere presence of the device imposed a cognitive cost through active suppression of the impulse to check it.

2.4. Cognitive Biases: The Established Taxonomy

The systematic errors that characterize System 1 processing have been extensively documented. Tversky and Kahneman's (1974) landmark *Science* paper identified and empirically characterized three fundamental heuristics — representativeness, availability, and anchoring — that produce predictable, consistent errors across a wide range of judgment tasks.[7] Subsequent decades of research have substantially extended the taxonomy. Confirmation bias has been documented across professional, clinical, legal, and political domains (Nickerson, 1998). Framing effects have been documented across consumer decision-making, medical treatment choices, and financial risk assessment (Tversky & Kahneman, 1981). Substitution bias operates as a meta-heuristic underlying many of the more specific biases (Kahneman, 2011).

Notable among the deliberate omissions from the BAM is loss aversion — Kahneman and Tversky's (1979) finding that losses loom approximately twice as large psychologically as equivalent gains. While loss aversion is among the most robustly documented phenomena in behavioral economics, its specific amplification as a function of cognitive load is less empirically established than the seven included biases; its interaction with overload is partially captured through the Framing Bias equation's (0.4 + 0.6E) emotional charge term.

What the existing literature documents comprehensively is the existence and prevalence of these biases under normal cognitive conditions. What it has not systematically addressed is the quantitative relationship between cognitive load and bias intensity — the degree to which overloaded conditions amplify biases that are present but manageable at baseline. This is the gap the BAM addresses.

[7] Tversky and Kahneman's 1974 *Science* paper is the foundational empirical documentation of the three heuristics — representativeness, availability, and anchoring — that constitute the majority of the BAM's bias taxonomy.

2.5. The Missing Integration

The three bodies of research reviewed above — attention economics, dual-process theory, and the cognitive bias literature — each describe components of the same underlying problem without formally connecting them. The attention economy literature documents the environmental inputs that generate cognitive load without modeling the cognitive outputs. Dual-process theory describes the System 1/System 2 shift without quantifying its relationship to specific information environment parameters. The bias literature catalogs systematic errors without modeling their relationship to the information environment that activates and amplifies them.

A fourth body of research has been largely absent from prior theoretical accounts of information overload: the literature on cognitive architecture, schema development, expertise, and metacognitive monitoring. This literature describes the mechanisms through which human agents can actively expand their effective cognitive processing capacity — not merely by improving biological parameters, but by building trained mental models, externalized organizational systems, and accurate self-monitoring of cognitive state. Section 2.6 reviews this literature, which provides the theoretical grounding for the BAM's reconceptualized treatment of C in Revision 5.

The Bias Amplification Model is an attempt to provide the missing formal connection: a mathematical framework that takes measurable inputs from the information environment, applies the known properties of dual-process cognition and working memory architecture, and produces quantitative predictions about the intensity of specific cognitive biases and aggregate cognitive distortion — together with a theoretically grounded account of how human agency can reduce that distortion through deliberate modification of the model's input variables.

2.6. The Cognitive Architecture Literature: Schema, Expertise, Offloading, and Metacognition

A substantial body of research demonstrates that cognitive processing capacity is not a fixed biological parameter but a composite that human agents can actively develop across multiple dimensions. This literature is organized here around four interrelated mechanisms: schema development, expertise effects, cognitive offloading, and metacognitive monitoring. Each represents an empirically supported pathway through which an individual can reduce effective Ω — not by reducing information volume, but by expanding the cognitive resources available to process it.

Schema development. Schema theory, originating with Bartlett (1932) and extended by cognitive scientists across subsequent decades, proposes that organized knowledge structures in long-term memory fundamentally alter the cognitive cost of processing incoming information. When an individual has well-developed schemas for a domain — organized representations of how information in that domain is structured, categorized, and related — incoming information is processed not as a collection of individual items but as recognizable patterns that can be matched against existing structures. The cognitive cost per unit of information processed drops dramatically. A trained chess master processes a complex board position as a small number of meaningful chunks; a novice processes the same position as dozens of unorganized individual pieces. Both are receiving identical information volume; the cognitive load imposed is orders of magnitude different. Sweller's own framework acknowledges this explicitly: schemas are the mechanism through which germane cognitive load — the productive effort that builds capacity — reduces the intrinsic load of subsequent processing (Sweller, 1988).

Expertise effects. The expertise literature documents a consistent finding: experts in a domain report lower cognitive load than novices when processing equivalent information volumes in that domain (Frontiers in Psychology, 2023). The mechanism is schema-based chunking — the same working memory capacity that is overwhelmed in a novice operates comfortably in an expert because the expert's schemas compress multiple information items into single meaningful units. This finding carries an important qualification, however. Expertise reduces cognitive load within the domain of expertise; it does not generalize. Outside their domain of expertise, experts face the same working memory constraints as novices — and there is evidence that domain expertise can actually increase susceptibility to certain biases by strengthening the overconfidence that System 1 fluency produces (Kahneman, 2011). The BAM must therefore treat expertise-based capacity expansion as domain-specific, not as a global enhancement of C.

Cognitive offloading. A growing body of research under the heading of cognitive offloading documents the mechanisms through which humans systematically extend their effective processing capacity by externalizing information to trusted environmental scaffolds. Writing things down, building structured task lists, maintaining categorized reference systems, and designing information environments that pre-filter and pre-organize incoming material all reduce the load on internal working memory by transferring portions of the processing task to external systems. Research confirms that cognitive offloading reduces memory load and improves decision-making and creativity (Risko & Gilbert, 2016). The organizational productivity literature has documented this mechanism in applied contexts: externalizing information to trusted systems frees internal cognitive resources for higher-order processing tasks, and the separation of actionable from non-actionable information alone can dramatically reduce the effective I that the working memory system must process in real time (Allen, 2001). For the BAM, cognitive offloading represents a mechanism through which C is effectively expanded — not by increasing biological processing capacity, but by reducing the demands placed on it through structural pre-processing of the information environment.

Metacognitive monitoring. Metacognitive monitoring is the capacity to accurately assess one's own cognitive state in real time — specifically, to recognize when working memory is approaching capacity and System 2 corrective processing is becoming unavailable. This capacity is particularly relevant to the BAM because of the model's central phenomenological finding: cognitive overload is subjectively invisible to the person experiencing it. The Zone 2 and Zone 3 phenomenological signatures — "I feel on top of it" and "I know what I think" — are not experiences of clarity that accurately reflect cognitive state; they are the felt confidence that System 1 always delivers, misread as reliable indicators of deliberate engagement. A person with well-developed metacognitive monitoring capacity can recognize the discrepancy between felt certainty and actual cognitive state — can notice the signs of rising Ω before they have produced their full distorting effects. Without this capacity, no other intervention is reliably accessible, because the person does not know they need to intervene. Metacognitive monitoring is trainable through systematic practice in cognitive state awareness, deliberate attention training, and structured reflection on decision quality (Flavell, 1979; Dunlosky & Metcalfe, 2009).

3. The Bias Amplification Model

3.1. Core Theoretical Proposition

The BAM rests on four propositions, each independently supported by the literature reviewed above. The fourth proposition — concerning cognitive architecture — is new to Revision 5 and substantially extends the model's prescriptive reach.

First proposition (model assumption, theoretically motivated): Cognitive overload — defined as effective information input rate exceeding cognitive processing capacity — depletes System 2's capacity to audit and correct the heuristic outputs that System 1 continuously produces. This depletion is not experienced consciously by the person undergoing it and cannot be reliably detected through self-report (Kahneman, 2011).

Second proposition (empirically established): The transition from System 2 to System 1 dominance follows a threshold pattern rather than a linear degradation. Below a critical overload threshold, deliberative processing remains available. Above it, heuristic processing becomes dominant, with the degree of dominance increasing with the degree of overload.

Third proposition (empirically established): The seven cognitive biases included in the model — Confirmation, Availability, Substitution, Recency, Anchoring, Representativeness, and Framing — are differentially activated by System 1 processing and thus differentially amplified by cognitive overload. Each has a baseline intensity β present under normal conditions and an overload-sensitive component that increases as System 2 corrective capacity is depleted.

Fourth proposition (model assumption, supported by the cognitive architecture literature): Cognitive processing capacity C is not a fixed biological parameter. It is a composite of biological baseline capacity, schema-augmented capacity arising from expertise and trained mental models, and externalized capacity arising from deliberate information organization systems. Each component is trainable on different timescales, and together they determine the effective Ω that a given information environment imposes on a given individual at a given moment. The practical implication is that the same nominal information load — the same notification rate, the same meeting density, the same content velocity — may produce very different Ω values in different individuals depending on their schema development, organizational system quality, and biological state.

Throughout this paper, three claim types are distinguished. *Empirically established claims* are supported by peer-reviewed experimental evidence independent of the BAM. *Model assumptions* are structural choices that are theoretically motivated but not themselves directly tested by the model. *Model predictions* are statements about what the BAM forecasts that are subject to empirical test. The implications discussed in Section 5 are model predictions, not established facts.

3.1.1. The Master Equation: A Unified Statement

The model's full master equation is:

$$\Phi(\Omega) = \Phi_0 + f(\Omega) \cdot A \quad (1)$$

$\Phi(\Omega)$ Aggregate Bias Index: composite measure of overall cognitive distortion at a given overload level.

Φ_0 Baseline Bias Floor: irreducible aggregate bias at zero cognitive overload. Approximately 0.17 in the current parameterization.

f(Ω) Sigmoid Amplification Function: $f(\Omega) = 1/(1 + e^{-k(\Omega-1)})$. At Ω = 1.0, f(Ω) = 0.50.

A Weighted Mean Bias Sensitivity: $A = \Sigma(w_i \cdot \alpha_i) \approx 0.404$ under current provisional parameterization.

Ω Cognitive Overload Ratio: Ω = I/C. Values below 1.0 indicate sub-threshold load; values above 1.0 indicate overload.

k Individual Tipping Steepness: default k = 3.0 for an average rested adult.

3.2. The Cognitive Overload Ratio and the Composite Nature of C

The BAM's primary independent variable is the Cognitive Overload Ratio, designated Ω:

$$\Omega = I / C \quad (2)$$

where I represents the effective information input — a composite of rate, quality, and organizational structure, discussed below — and C represents cognitive processing capacity. The critical threshold is Ω = 1.0. The model proposes that the critical depletion of System 2's corrective capacity occurs near this threshold. Below Ω = 1.0, System 2 retains sufficient resources to audit System 1's heuristic outputs and correct likely errors before they drive behavior. Above it, that corrective capacity is progressively exhausted.[8]

The composite nature of C.

In the BAM's original formulation, C was treated as a near-fixed individual parameter, improvable primarily through sleep and exercise. Revision 5 reconceptualizes C as a composite of three interacting components, each trainable on different timescales and through different mechanisms.

The first component is **biological baseline capacity** — the working memory capacity and executive function availability that the individual brings to any given moment as a function of sleep quality, physiological state, stress load, and accumulated cognitive fatigue. This is the component that the original BAM described as C. It is real, it matters, and it is the component most immediately sensitive to the environmental conditions that the Ω Portrait Framework describes. But it is the least trainable component on any given day — a person cannot significantly increase their biological baseline capacity within a single working session. They can protect it through sleep, exercise, recovery intervals, and stress management. But they cannot expand it through the kind of deliberate cognitive practice that the other two components allow.

The second component is **schema-augmented capacity** — the increase in effective processing capacity that arises when incoming information can be matched against well-organized knowledge structures in long-term memory. When an individual has developed rich schemas for the domain in which they are operating, each incoming information item carries more meaning per unit of working memory consumed. The expert analyst processes a financial report differently from the novice not because they have more working memory, but because their schemas compress multiple data points into single recognizable patterns that can be evaluated quickly and efficiently. From the perspective of the BAM, this means the same nominal information volume produces a lower Ω for the schema-equipped individual than for the schema-deficient one. Schema-augmented capacity develops over time through deliberate practice, structured learning, and systematic reflection on domain experience. It is trainable on a timescale of months to years and is domain-specific: schema development in financial analysis does not generalize to schema development in geopolitical risk assessment.

The third component is **externalized capacity** — the effective expansion of processing capacity that arises when trusted external systems absorb portions of the information management task that would otherwise be performed by working memory. When an individual maintains a trusted, comprehensive system for capturing incoming information, categorizing it by type and urgency, routing it to appropriate processing queues, and retrieving it on demand, their internal working memory is freed from the burden of holding open loops — unprocessed items that generate background cognitive load simply by remaining unresolved. Research on cognitive offloading confirms that this externalization reduces memory load and measurably improves decision-making quality (Risko & Gilbert, 2016). The organizational implication is significant: externalized capacity is trainable quickly, on a timescale of days to weeks, through the design and adoption of structured information management practices. Unlike biological baseline capacity, it does not require years of domain practice to develop. A person who builds a reliable external system for processing their information environment can expand their effective C substantially within a short period of deliberate practice.

These three components interact. High biological baseline capacity allows schema learning to proceed more efficiently. Well-developed schemas reduce the cognitive cost of operating an externalized system. And a well-designed externalized system protects biological baseline capacity by reducing the accumulation of unresolved cognitive load across a working session. The composite nature of C also means that two individuals operating in identical nominal information environments — the same notification rate, the same meeting density, the same content velocity — may be operating at very different effective Ω values depending on their schema development, system quality, and biological state. This is not a minor footnote to the model; it is the primary mechanism through which human agency operates on the model's central variable.

The effective nature of I.

Information input rate I has similarly been reconceptualized in Revision 5. In its original formulation, I was treated primarily as a speed variable — the rate at which information items arrive per unit time. The cognitive architecture literature makes clear that rate is only one dimension of what makes information cognitively costly. Two dimensions that operate alongside rate are particularly significant.

The first is **information quality** — the degree to which incoming information is relevant, organized, unambiguous, and pre-filtered for the individual's actual processing needs. A high-volume stream of well-organized, clearly relevant, low-ambiguity information imposes less cognitive load than a lower-volume stream of ambiguous, irrelevant, emotionally charged noise, because the former can be processed more efficiently per item and discarding decisions are straightforward. The BAM's E parameter partially captures the emotional charge dimension of information quality, but does not yet fully account for relevance filtering, organizational structure, and ambiguity. Future formalization of I as a quality-weighted rate variable is identified as a priority for model development.

The second is **information organizational structure** — the degree to which incoming information arrives in a form that is already partially processed and categorized before it reaches working memory. An individual who has designed their information environment so that incoming items are pre-categorized by type and urgency — rather than arriving as an undifferentiated stream requiring real-time sorting — reduces the effective I their cognitive system must process, even if the nominal volume is unchanged. This is the information environment design dimension of the model, and it is directly actionable through deliberate restructuring of how information reaches the individual.

Measurement Scale of Ω: An Ordinal Framing. The current operationalization of Ω, using behavioral proxies as specified in Table 4, should be understood as producing an ordinal rather than a ratio-scale measure. For Ω to function as a true ratio-scale quantity, I and C would need to be expressed in identical, commensurable units. The Shannon-to-chunk correspondence — the translation between statistical uncertainty reduction in bits per second and cognitive chunks in the sense of Miller (1956) and Cowan (2001) — is identified as the priority formalization question for future theoretical development.

[8] Shannon's (1948) information theory provides the principled candidate unit for I: bits per second. The obstacle is C: Shannon bits measure statistical uncertainty reduction, which does not map directly onto cognitive 'chunks'. The paper therefore retains 'information units per unit time' as a working placeholder.

3.2.1. The Ω Portrait Framework: Operationalizing Overload

While Ω is formally defined as a ratio, its practical application requires translation into measurable behavioral and environmental terms. The Ω Portrait Framework provides a seven-level characterization of the information environments, behavioral signatures, quantitative proxy metrics, and phenomenological states corresponding to discrete positions across the full Ω range. Six proxy variables provide the empirical grounding: notification rate, application-switching frequency, meeting load, sleep quality, emotional charge (E), and information acceleration rate (*A*/I). Table 4 presents the proxy-measure benchmarks. Full portrait descriptions are provided in the interactive model at attentionplease.ai/bam.

Table 4. Ω Portrait Framework: Proxy Measure Benchmarks Across Seven Overload Levels

Ω Range	Zone	Notif./hr	App Sw./hr	Mtg Load (hrs/day)	Sleep (hrs)	E (0–1)	f(Ω) approx.	Φ Range
0.25–0.49	Zone 1	0–3	<5	0–1	7–9	0–0.2	<0.18	0.17–0.24
0.50–0.74	Zone 1	4–8	5–10	1–2	6.5–7.5	0–0.3	0.18–0.32	0.24–0.30
0.75–0.99	Zone 2	8–15	10–20	2–3	6–7	0.2–0.5	0.32–0.50	0.30–0.37
1.00–1.24	Zone 2	15–25	20–35	3–4	6–7	0.3–0.6	0.50–0.68	0.37–0.44
1.25–1.74	Zone 3	25–40	35–55	4–6	5.5–6.5	0.4–0.8	0.68–0.85	0.44–0.50
1.75–2.19	Zone 3	40–70	55–80	5+	<6.5	0.6–0.9	0.85–0.93	0.50–0.55
2.20+	Zone 4	>70	>80	5+	<6	>0.7	>0.93	>0.55

Note. All proxy values are approximations based on the empirical literature and are subject to empirical calibration. Zone coloring corresponds to the four BAM cognitive zones defined in Section 4. The composite C reconceptualization in Revision 5 implies that individuals with higher schema-augmented or externalized capacity may occupy lower effective Ω zones than the nominal proxy values suggest.

3.3. The Sigmoid Amplification Function

The BAM models the progressive depletion of System 2's corrective capacity as a sigmoid function of Ω. At f(Ω) near 0, System 2 correction is robust and errors are largely caught before they influence behavior. At f(Ω) near 1, corrective capacity is effectively absent and System 1's outputs pass through unexamined. The sigmoid was selected because it captures the threshold character of the transition — a flat response below the tipping point, rapid nonlinear change in the transition zone, and saturation above the tipping point — which is consistent with the categorical quality of the System 1/System 2 shift described in dual-process theory.[9]

Alternative functional forms were explicitly considered. A linear formulation fails to capture the abrupt qualitative shift in processing mode predicted by dual-process theory. A purely exponential function produces unbounded amplification that is physiologically implausible. Piecewise threshold models introduce non-smooth transitions inconsistent with graded behavioral changes observed under incrementally increasing load. The sigmoid uniquely satisfies all three theoretically required properties: sub-threshold stability; rapid nonlinear transition near the critical tipping point; and asymptotic saturation at extreme overload.

$$f(\Omega) = 1 / (1 + e^{-k(\Omega-1)}) \quad (3)$$

where e is Euler's number, k is the individual sensitivity parameter (default k = 3.0), and Ω is the Cognitive Overload Ratio. *Sensitivity Analysis: The k Parameter.* Higher k values (4.0–6.0) are consistent with the abrupt performance collapse reported under acute cognitive depletion. Lower k values (2.0–2.5) are consistent with individuals with broad attentional training or high baseline cognitive resilience. Importantly, the composite C reconceptualization in Revision 5 suggests that effective k — the steepness of an individual's transition under real-world conditions — is also modifiable through schema development and externalized capacity, not only through sleep and biological state. An individual with strong schemas and a reliable externalized system effectively operates with a lower k because their C remains high longer under load.

[9] The sigmoid function's appropriateness for modeling cognitive tipping points is supported by the broader literature on threshold effects in neural and behavioral systems. The parameter k controls how abruptly this transition occurs.

Table 1. Sensitivity Analysis: Sigmoid Function f(Ω) and Phi Index (Φ) Across k Values

	k = 2.0		k = 3.0 (default)		k = 4.0		k = 6.0	
Ω	f(Ω)	Φ	f(Ω)	Φ	f(Ω)	Φ	f(Ω)	Φ
0.50	0.269	0.279	0.182	0.244	0.119	0.218	0.047	0.189
0.75	0.378	0.323	0.321	0.300	0.269	0.279	0.182	0.244
1.00	0.500	0.372	0.500	0.372	0.500	0.372	0.500	0.372
1.25	0.622	0.421	0.679	0.444	0.731	0.465	0.818	0.501
1.50	0.731	0.465	0.818	0.501	0.880	0.526	0.953	0.555
2.00	0.880	0.526	0.953	0.555	0.982	0.567	0.998	0.573

Note. $f(\Omega) = 1/(1 + e^{-k(\Omega-1)})$. $\Phi = \Phi_0 + f(\Omega)\cdot A$ *where* $\Phi_0 = 0.17$ *and* $A = 0.404$ *(E = 0, A/I = 0). Zone shading:* $\Omega < 0.75$ *= Zone 1;* $0.75 \le \Omega \le 1.25$ *= Zone 2;* $1.25 < \Omega \le 2.2$ *= Zone 3;* $\Omega > 2.2$ *= Zone 4. At* $\Omega = 1.0$*,* $f(\Omega) = 0.50$ *for all k values.*

3.4. Individual Differences and Moderating Variables

The BAM's original formulation contained a single individual difference variable: the tipping steepness parameter k, which modulates the abruptness of the System 1/System 2 transition. This is a real and important source of individual variation. But it is not the only one, and for the model's prescriptive claims to be credible, a more complete account of individual differences is required. This section introduces metacognitive monitoring capacity as a central moderating variable, and discusses three additional individual difference dimensions that influence the model's predictions.

Metacognitive monitoring capacity (M).

Metacognitive monitoring capacity is the ability to accurately assess one's own cognitive state in real time — specifically, to recognize when working memory is approaching capacity, when System 2 corrective processing is becoming unavailable, and when the felt certainty of System 1 outputs is no longer a reliable indicator of deliberate engagement. This capacity is not the same as intelligence, domain knowledge, or analytical skill. It is a distinct dimension of self-regulatory awareness that varies meaningfully across individuals and is trainable through deliberate practice (Flavell, 1979; Dunlosky & Metcalfe, 2009).

The reason M functions as a moderating variable in the BAM is that it determines whether the model's prescriptive interventions are functionally available to the individual at any given Ω level. The central phenomenological finding of the model is that cognitive overload is subjectively invisible to the person experiencing it. The Zone 2 phenomenological signature — "I feel on top of it" — and the Zone 3 signature — "I know what I think" — are not accurate reflections of cognitive state. They are the felt confidence that System 1 always delivers, experienced as clarity in the absence of any internal signal that bias amplification is underway. A person with high M can recognize the discrepancy between felt certainty and actual cognitive state, can notice the behavioral and environmental signatures of rising Ω, and can intervene before the tipping point is crossed. A person with low M cannot, because they genuinely do not know they need to. The most powerful alarm in the system is silent to them.

M is formally incorporated into the BAM as a moderator of the Ω-to-distortion pathway rather than as an additional equation term. High M does not reduce Ω directly; it enables deliberate Ω reduction by making the need for intervention detectable. Low M leaves the individual functionally unable to access the model's prescriptive interventions even when they are theoretically available. This makes the development of metacognitive monitoring capacity the logical first step in any individual cognitive state management program — not because it directly reduces bias, but because it is the precondition for all other interventions.

Need for cognition.

Need for cognition (NFC) is a stable personality trait describing the degree to which individuals are intrinsically motivated to engage in effortful analytical thinking (Cacioppo & Petty, 1982). Individuals high in NFC seek out analytical engagement and sustain deliberate processing under conditions that would cause lower-NFC individuals to default to heuristic shortcuts. The empirical evidence shows that NFC moderates the relationship between information overload and bias activation: high-NFC individuals who include more information in their consideration set make better final decisions, even under conditions of elevated load. For the BAM, NFC functions as a modifier of the effective k parameter — high-NFC individuals show a slower, more gradual sigmoid transition because their dispositional motivation to think analytically maintains System 2 engagement longer before depletion occurs. NFC is a relatively stable trait, less trainable than M or externalized capacity, but it interacts with M in an important way: high-NFC individuals are also more likely to develop metacognitive monitoring capacity, because their dispositional engagement with their own thinking processes creates more opportunities for metacognitive learning.

Domain expertise and schema quality.

As reviewed in Section 2.6, expertise in a domain substantially reduces the cognitive cost of processing information within that domain through schema-based chunking. This creates a meaningful individual difference in effective Ω that is domain-specific and that develops over timescales of months to years. The BAM predicts that two individuals processing the same information environment will show different bias amplification profiles if one has substantially greater domain expertise than the other, holding biological baseline capacity constant. This prediction is directly testable and would constitute Stage 1b of the model's empirical calibration program.

The qualification noted in Section 2.6 bears repetition here: expertise is not a globally protective factor against cognitive bias. Outside their domain of expertise, experts are subject to the same overload-driven bias amplification as novices. Within their domain, expertise can introduce its own bias vulnerabilities — overconfidence, resistance to disconfirming information, and the kind of pattern-matching rigidity that produces spectacular expert errors. The BAM does not model expertise as a uniform cognitive enhancer. It models it as a domain-specific C-expander with its own characteristic error modes.

Information organization system quality.

The quality of an individual's external information organization system — the degree to which it reliably captures incoming items, routes them appropriately, and makes them retrievable on demand without requiring working memory to track them — is an individual difference that operates on the externalized component of C. This is arguably the most immediately actionable individual difference in the model, because it can be developed rapidly and does not require domain expertise. The design and adoption of a structured information management system that the individual genuinely trusts — that reliably captures what needs to be captured and does not require ongoing monitoring to verify — reduces the background cognitive load generated by unresolved open loops and allows working memory to focus on the current primary task. The empirical grounding for this claim comes from the cognitive offloading literature (Risko & Gilbert, 2016) and from applied research on structured productivity systems.

3.5. Individual Bias Equations

The seven biases included in the BAM were selected on four criteria: (a) each is independently documented in peer-reviewed experimental work with substantial replication; (b) each has a theoretically motivated functional relationship to cognitive load specifically; (c) together they span the primary processing domains of information intake, memory retrieval, judgment formation, and decision framing; and (d) each is directly relevant to the practical domain of information-environment design and digital attention economics.

The connection between each bias and the overload mechanism is that each is specifically produced or amplified by the depletion of deliberate processing resources. Anchoring amplifies because deliberate adjustment away from an initial reference value requires System 2 effort. Availability amplifies because System 1 retrieves the most cognitively accessible exemplars rather than the most statistically representative ones. Confirmation amplifies because motivated correction of initial interpretations is precisely the kind of effortful evaluation that depleted deliberate processing cannot sustain. Recency amplifies because the temporal weighting of information toward the most recent inputs is a default System 1 property that System 2 adjusts under normal conditions. Representativeness amplifies because base-rate integration requires deliberate reasoning. Framing amplifies because gain-loss equivalence recognition requires comparative evaluation that System 2 performs and System 1 bypasses. Substitution amplifies most severely because it represents the most direct behavioral consequence of the shift itself.

Each bias equation takes the general form:

$$B_x(\Omega) = \beta_x + \alpha_x \cdot g(f(\Omega), E, A/I) \quad (4)$$

Table 2. Bias Amplification Model: Individual Bias Equations, Parameters, and Phi Weights

Bias	Symbol	β	α	w_i	Equation	Primary Driver
Confirmation	B_cf	0.18	0.45	0.19	β_cf + α_cf·(1+E)·f(Ω)	Emotional charge
Availability	B_av	0.14	0.38	0.15	β_av + α_av·f(Ω)	Overload above tipping point
Substitution	B_sb	0.10	0.50	0.15	β_sb + α_sb·(1–1/Ω) [Ω≥1]	Fractional excess overload
Recency	B_rc	0.12	0.40	0.15	β_rc + α_rc·(*A*/I)·f(Ω)	Information acceleration
Anchoring	B_an	0.15	0.30	0.11	β_an + α_an·f(Ω) [U-shaped]	Extreme Ω: high & low
Representativeness	B_rp	0.17	0.38	0.13	β_rp + α_rp·f(Ω)·(1+0.25E)	Overload × emotional charge
Framing	B_fr	0.13	0.42	0.12	β_fr + α_fr·(0.4+0.6E)·f(Ω)	Content charge × amplification

Definitions. β = baseline bias intensity. α = amplification coefficient. w_i = Phi weight; weights sum to 1.00. E = emotional charge ∈ [0,1]. A/I = fractional growth rate of I over reference window τ. f(Ω) = sigmoid amplification function. Note on claim type: the parameter values in this table are model assumptions derived from the existing empirical bias literature, not empirically calibrated values from BAM-specific experiments. They should be treated as theoretically grounded initial estimates subject to revision through the calibration sequence described in Section 6.

Confirmation Bias incorporates a (1 + E) interaction structure, reflecting two separable mechanisms: cognitive overload alone amplifies through the base f(Ω) term, while emotional charge E then acts as a multiplier. Recency Bias uniquely incorporates the information acceleration rate *A*/I, capturing the breaking-news effect. Anchoring Bias displays a U-shaped pattern as a function of Ω, reflecting distinct mechanisms at high and low load. Substitution Bias is modeled with the hyperbolic (1–1/Ω) rather than the sigmoid because it reflects the wholesale replacement of the target question with an easier substitute, structurally dependent on excess load magnitude rather than transition dynamics.

3.6. The Aggregate Bias Index

The Phi Index (Φ) aggregates the seven individual bias intensities into a single composite measure of overall cognitive distortion:

$$\Phi = \Sigma\, w_i \cdot B_i(\Omega) \quad (5)$$

The Phi Index ranges from approximately 0.17 — its irreducible baseline floor, achieved under near-zero overload with emotionally neutral content and zero information acceleration — to approximately 0.61, the computed maximum at extreme overload combined with maximum emotional charge and maximum information acceleration. The commonly encountered operational range for adults in contemporary information environments is narrower: Φ values between 0.24 and 0.37 represent the normal operating range.

An interpretive clarification is essential. The Phi Index represents an aggregate cognitive distortion potential — a measure of the overall vulnerability of the cognitive system to bias-driven errors at a given

Ω level — and not a claim that all seven biases are simultaneously and fully active in every individual decision. The additive aggregation in Φ is a first-order approximation; the additive Φ is likely an upper bound on aggregate distortion relative to a correlation-corrected estimate, which is conservative in the direction appropriate for a diagnostic instrument.

4. The Four Cognitive Zones

The BAM identifies four discrete cognitive zones corresponding to escalating values of Ω and Φ. The zone boundaries correspond to qualitatively distinct cognitive states with meaningfully different decision-quality profiles. The composite C reconceptualization introduced in Revision 5 implies that an individual's effective zone membership is determined not only by the nominal information environment but by their schema development, externalized system quality, and biological state. Two individuals in the same nominal environment may occupy different effective zones.

Zone 1 — Logical ($\Omega < 0.75$, $\Phi < 0.25$)

At Ω below 0.75, information input is well within processing capacity. System 1 processes all incoming information first, as always, but System 2 retains sufficient corrective capacity to audit those outputs, interrogate heuristic judgments, and override errors before they shape behavior. Biases exist at their baseline levels — System 1 generates them continuously — but System 2's corrective function constrains their influence on final judgments to manageable ranges. The sigmoid function $f(\Omega)$ is below 0.20 in this zone. Decision quality is at its maximum, and the cognitive state is appropriate for high-stakes analytical tasks, complex decisions, and creative work requiring sustained attention.

Zone 1 is increasingly rare in contemporary information environments for individuals without deliberate cognitive state management practices. Achieving it requires deliberate management of information inputs and, importantly, the expansion of C through schema development, externalized systems, and biological state maintenance. The evidence from workplace studies suggests that knowledge workers without deliberate practices spend less than two hours per day in states approximating Zone 1 (Mark et al., 2008).

Zone 2 — Transition ($0.75 \leq \Omega \leq 1.25$, $0.25 \leq \Phi \leq 0.37$)

Zone 2 is the tipping point region. The sigmoid function is at its steepest in this range, meaning that small increases in Ω produce disproportionately large increases in $f(\Omega)$ and consequently in Φ. This is the zone in which metacognitive monitoring capacity becomes most clinically consequential: a person in Zone 2 is at the inflection point between deliberative and heuristic dominance, and small environmental interventions — a brief recovery interval, a reduction in notification volume, activation of an externalized system — can return them to Zone 1, while continued overload pushes them into Zone 3. Without metacognitive monitoring capacity, this intervention window is invisible: the Zone 2 phenomenological signature ("I feel on top of it") provides no reliable internal signal that the window exists.

Zone 3 — Heuristic ($1.25 < \Omega \leq 2.2$, $0.37 < \Phi \leq 0.48$)

Zone 3 represents the state in which System 2's corrective capacity is substantially depleted. System 1's heuristic outputs now reach judgment and behavior with minimal scrutiny. The practical result is that decisions are driven predominantly by substitutions, associations, and shortcuts that in Zone 1 conditions would have been caught and corrected before influencing behavior. All seven biases are operating significantly above baseline. Decision quality is substantially impaired, but this impairment is experienced as normal thinking — the WYSIATI effect ensures that System 1 presents its outputs with full apparent confidence. Zone 3 is the modal cognitive state during social media use, cable news consumption, and multi-task knowledge work with continuous notification exposure for individuals without deliberate cognitive state management.

The ecological validity of Zone 3's decision-quality predictions is supported by Danziger et al.'s (2011) study of judicial decision-making, in which favorable parole rulings declined from approximately 65% at the start of a session to near-zero before breaks, then reset after rest — a pattern consistent with accumulated cognitive load shifting processing from Zone 1 toward Zone 3.

Zone 4 — Critical ($\Omega > 2.2$, $\Phi > 0.48$)

Zone 4 represents extreme overload. The sigmoid function approaches saturation, with $f(\Omega) > 0.95$. All seven biases are near maximum amplification. Decisions made in Zone 4 are highly unreliable and subject to the full range of amplified biases operating simultaneously. Importantly, Zone 4 is the first zone in which individuals may begin to experience felt cognitive strain — not because they are aware of bias amplification, but because sheer information volume is outrunning even heuristic processing capacity. This felt overwhelm, when it occurs, can itself serve as a metacognitive signal for individuals with sufficient M to recognize it as such.

5. Implications

The implications in Section 5 represent model-consistent interpretations — predictions and inferences that follow from the BAM's formal structure — not direct empirical claims about established causal relationships in real-world systems. Section 5.5, which draws on the composite C reconceptualization and the individual differences analysis in Section 3.4, is presented as a set of model predictions with supporting theoretical grounding rather than as empirically validated prescriptions.

5.1. For Individual Decision-Making

The BAM's most direct implication for individual cognition is that the quality of decisions is a function not only of the individual's intelligence, experience, and motivation but of the cognitive state in which the decision is made — and that cognitive state is substantially modifiable through deliberate practice.

This reframing shifts the causal explanation for poor decisions from dispositional to situational and architectural factors: my information environment has elevated my Ω to a point where System 2's corrective capacity is depleted, allowing System 1's heuristic outputs to drive judgment without adequate oversight — and I lack either the metacognitive monitoring capacity to detect this state or the cognitive architecture to prevent it.

This reframing is not exculpatory — it is diagnostic and prescriptive. Pennycook and Rand (2019) found that susceptibility to misinformation was best predicted by cognitive engagement level rather than political motivation.[12] The BAM formalizes this finding: Zone 1 conditions support the analytical engagement that reduces susceptibility to misinformation. Zone 3 conditions suppress it. And the composite C framework identifies the specific mechanisms through which Zone 1 conditions can be achieved and maintained.

[12] Pennycook and Rand (2019) found that susceptibility to misinformation was better predicted by lower analytical thinking than by partisan motivation — supporting the BAM's proposition that cognitive state, not political orientation, is the primary determinant of bias activation.

5.2. For Organizational and Institutional Decision-Making

The BAM's implications extend to the conditions under which organizational and institutional decisions are made. If the cognitive state of individual decision-makers is a function of their information environment and cognitive architecture, then organizations that allow high-interrupt, always-connected work cultures are systematically producing Zone 3 decision conditions for their most consequential judgment tasks — while simultaneously failing to support the development of the schema-augmented and externalized capacity components that could protect against this.

The Danziger et al. (2011) finding — that judicial decisions are substantially influenced by the temporal accumulation of cognitive load without the judges' awareness — is a specific institutional manifestation of the general BAM mechanism. The implication is that institutional decision quality can be improved through deliberate design of information environments and decision timing, combined with systematic investment in the cognitive architecture development of key decision-makers.

5.3. For Public Health and Democratic Function

If the contemporary information environment routinely produces Zone 3 cognitive conditions for a substantial fraction of the population, the aggregate effect on collective judgment quality is significant.

The empirical evidence that heavy social media use is associated with polarization of attitudes (Bail et al., 2018), reduced tolerance for nuance (Rosen et al., 2013), and increased susceptibility to misinformation (Vosoughi et al., 2018) is consistent with the BAM's predictions about what Zone 3 Confirmation Bias amplification would produce at population scale.

Democratic deliberation depends on citizens capable of evaluating complex information, weighing competing claims, updating beliefs in response to evidence, and tolerating the ambiguity inherent in political reality. These capacities are the products of System 2 deliberative processing — precisely the processing mode that Zone 3 conditions systematically suppress. The BAM predicts that population-level investment in metacognitive monitoring capacity and information management skills would constitute a genuine public health intervention with measurable effects on democratic deliberation quality.

5.4. For Technology and Media Design

The BAM provides a formal basis for evaluating technology and media design choices in terms of their cognitive effects. Features that increase effective I — notification volume, infinite scroll, autoplay, variable reward schedules — systematically push users toward Zone 3. Features that reduce effective I or support C — batch notification delivery, designed stopping points, recovery time defaults, information pre-filtering and categorization — systematically support Zone 1 and Zone 2 conditions.

Goldhaber (1997) argued that attention is fundamentally zero-sum: attention captured by one source is unavailable for another.[13] The BAM extends this with a specific prediction: attention captured under Zone 3 conditions is not merely unavailable for other purposes — it is cognitively compromised. The decisions and judgments made from that captured attention are predicted to be systematically distorted in the specific, measurable ways the BAM's bias equations describe.

The model's practical significance for technology design operates on three levels: (a) Measurement: the Φ index provides a theoretically grounded metric for assessing cognitive state in real time. (b) Design evaluation: information environments and digital products can be assessed against the model's predictions to estimate their likely Ω impact on users. (c) Targeted intervention: the model identifies specific leverage points — reducing effective I, protecting C, introducing decision friction — that correspond to predictable reductions in Φ.

[13] Goldhaber (1997) argued that attention, unlike money, is fundamentally zero-sum: "Unlike information, attention can't be manufactured out of thin air."

5.5. Prescriptive Interventions: A Formal Account

The composite C reconceptualization and the individual differences analysis in Section 3.4 together generate the BAM's prescriptive framework. Prior versions of the model identified two levers for Ω reduction: reducing I and increasing C through sleep and exercise. Revision 5 extends this to five intervention pathways, each targeting a specific variable and operating on a distinct timescale.

These five interventions can be organized under three operational headings that describe the logical sequence of the change process: **Awareness**, **Friction**, and **Reclamation**. Awareness describes the foundational step of accurate self-monitoring — knowing that the problem exists and recognizing it in real time. Friction describes the structural interventions that slow, filter, and organize information flow before it reaches working memory. Reclamation describes the longer-term work of rebuilding and expanding cognitive processing capacity. These headings are proposed as an organizing framework for the model's prescriptive layer, not as independent stages — in practice they operate simultaneously and reinforce each other. All interventions are model predictions with theoretical grounding; they have not been tested in BAM-specific experimental designs.

Awareness

The precondition for all other interventions. Awareness is the capacity to accurately detect one's own cognitive overload state in real time — to recognize rising Ω before it has produced its full distorting effects. Without Awareness, no deliberate intervention is possible, because the person does not know they need to intervene.

Intervention 1: Develop metacognitive monitoring capacity (targets M; timescale: weeks to months).

Since metacognitive monitoring capacity (M) is the precondition for all other interventions — the person who does not know they are overloaded cannot deliberately manage their overload — its development is the logical starting point of any cognitive state management program. The BAM's central phenomenological finding is that the Zone 2 signature ("I feel on top of it") and the Zone 3 signature ("I know what I think") are not accurate reflections of cognitive state. They are the felt confidence that System 1 always delivers, misread as reliable indicators of deliberate engagement. Developing M means learning to recognize the environmental and behavioral signatures of rising Ω: the preference for answering the simpler version of a hard question, the marked decrease in the quality of questions asked, the felt certainty that arrives with none of the deliberate work that usually precedes it. Metacognitive monitoring capacity is developed through systematic attention to the relationship between cognitive state and decision quality, structured attention training practices (Dunlosky & Metcalfe, 2009), and the habit of pausing before high-stakes decisions to assess cognitive state rather than trusting felt clarity as a reliable guide. The BAM predicts that improvements in M will have multiplicative rather than additive effects on decision quality, because they unlock all subsequent interventions.

Friction

The structural layer. Friction interventions introduce deliberate resistance and organizational pre-processing between the person and the raw information stream, reducing the effective I that reaches working memory without requiring the person to process every item in real time. Both interventions under this heading target the I variable — one through volume reduction, one through structural pre-organization.

Intervention 2: Design the information environment to reduce effective I (targets I; timescale: days).

The most immediately actionable intervention is the deliberate restructuring of information flow to reduce the volume, velocity, and cognitive cost of incoming information. This includes batch processing of communications rather than continuous monitoring, deliberate notification reduction, scheduled disconnection periods, and the elimination of passive information streams that impose cognitive cost without generating decision-relevant content. The effectiveness of this intervention is supported directly by the attention and cognitive load literatures: task interruption is one of the primary drivers of overload in decision environments (Speier et al., 1999), and the mere presence of a notification-generating device reduces working memory performance even when the device is not actively used (Ward et al., 2017). Information environment design operates on the I variable rather than on C, making it the fastest-acting Friction intervention in the model.

Intervention 3: Build and maintain a trusted externalized information system (targets externalized C; timescale: days to weeks).

A trusted external system — one that reliably captures all incoming information items, routes them to appropriate processing queues, and retrieves them on demand without requiring working memory to track their status — reduces the effective Ω imposed by a given nominal information environment by offloading the management task from internal working memory to the external scaffold. The critical operational word is *trusted*: a system the individual does not fully trust generates its own cognitive load through the ongoing monitoring required to verify it is functioning reliably. A system that is genuinely trusted allows working memory to release the open loops it has been holding, freeing capacity for primary task engagement. Research on cognitive offloading confirms that this mechanism is real and measurable (Risko & Gilbert, 2016), and that clearly separating actionable from non-actionable information is one of the most powerful forms of Friction available to a knowledge worker. The design of such a system is learnable in days to weeks; its maintenance becomes habitual over months.

Reclamation

The capacity-building layer. Reclamation interventions expand the composite C variable over time, reducing the effective Ω that a given information environment imposes by increasing the cognitive resources available to process it. These interventions operate on longer timescales than Awareness and Friction, but their effects are more durable because they alter the architecture of the cognitive system itself rather than managing its inputs.

Intervention 4: Develop domain-specific schemas through deliberate practice (targets schema-augmented C; timescale: months to years).

Schema development through deliberate practice in a domain reduces the cognitive cost per unit of domain-relevant information processed, effectively expanding C for that domain's information streams. The practical implication for knowledge workers is that investment in deep domain expertise is not only an economic or professional asset — it is a direct investment in cognitive processing capacity that reduces effective Ω in the individual's primary domains of professional responsibility. An expert processes familiar domain information at a fraction of the cognitive cost of a novice, occupying a lower effective zone even at the same nominal Ω. The qualification from Section 3.4 applies: schema-augmented capacity is domain-specific and does not generalize. The overconfidence that domain expertise can produce may introduce its own bias vulnerabilities outside or even within the domain of expertise.

Intervention 5: Protect biological baseline capacity through recovery and physiological maintenance (targets biological C; timescale: daily).

The original BAM's two-lever prescriptive framework — improve sleep and exercise — remains valid as a description of the biological baseline capacity component of Reclamation. Sleep quality is the dominant determinant of available working memory and executive function: a single night of reduced sleep imposes measurable and substantial degradation in System 2 corrective capacity. Cognitive recovery intervals during the working day — periods of genuine cognitive rest rather than switching to a different information stream — prevent the accumulation of session-level depletion that the Danziger et al. (2011) judicial decision data document so clearly. Physical exercise, adequate nutrition, and stress management all contribute to biological baseline capacity maintenance. These interventions are the most widely known, the most empirically supported, and historically the least implemented in organizational cultures that treat cognitive depletion as a productivity metric rather than as a decision quality liability. In the context of the full five-intervention framework, protecting biological baseline capacity is the foundation on which all other Reclamation work depends: schema development proceeds more efficiently on a well-rested brain, and Awareness of cognitive state is more accurate when biological capacity is not already depleted before the working day begins.

These three operational categories — Awareness, Friction, Reclamation — are not independent and their interventions are not sequentially gated. A person can begin building Friction into their information environment on the same day they begin developing Awareness, and both will interact with the biological dimension of Reclamation from the outset. What the sequence captures is the logical dependency structure of the full program: Awareness is necessary for deliberate Friction to be applied at the right moments; Friction reduces the rate at which biological baseline capacity is depleted; and Reclamation builds the architectural foundations that make both Awareness and Friction more effective over time. A coherent personal cognitive state management practice integrates all five interventions, building from the most immediately actionable (Friction through information environment design) toward the most durably transformative (schema development and metacognitive monitoring). The BAM’s formal structure — the composite C variable, the effective I reconceptualization, the metacognitive monitoring moderator — provides the theoretical grounding for each step of that practice.

6. Limitations and Future Research

The BAM, as presented here, is a theoretical framework rather than an empirically validated measurement instrument. Several limitations must be acknowledged.

Parameter precision and epistemic status. The individual bias equations' baseline (β) and amplification (α) parameters listed in Table 2 are derived from the existing empirical literature on bias magnitudes rather than from direct experimental calibration of the BAM itself. The mathematical precision of their presentation — specific decimal values, specific functional forms — is a function of the model's notation conventions, not a claim that the parameter values are empirically established at that level of precision. They should be treated as theoretically grounded initial estimates subject to revision through purpose-designed experimental work. Similarly, the empirical evidence on the relationship between information overload and cognitive bias is, as Hwang and Lin (1999) documented, inconsistent across studies. The BAM proposes a formal model of this relationship; it does not resolve the empirical inconsistencies in the existing literature.

Operationalization of C and I. The composite C reconceptualization introduced in Revision 5 substantially enriches the model's theoretical account but also substantially complicates its measurement. Biological baseline capacity has tractable measurement proxies (working memory tasks, sleep quality assessments, physiological indicators). Schema-augmented capacity is harder to operationalize across domains without domain-specific assessment instruments. Externalized capacity is harder still — the "trusted" quality of an external system is a subjective property that may not be easily captured by objective measures of system design. The effective nature of I — incorporating quality dimensions alongside rate — similarly requires measurement approaches that go beyond device telemetry. These operationalization challenges are acknowledged as genuine obstacles to empirical calibration of the Revision 5 model.

Metacognitive monitoring capacity. M is introduced in Revision 5 as a moderating variable with theoretical grounding in the metacognition literature. Its formal integration into the model's mathematical structure — specifying how M modifies the Ω-to-Φ relationship quantitatively — is a priority for future model development. The current version treats M as a qualitative moderator rather than a quantitative one, which limits its testability.

Empirical calibration sequence. The BAM's calibration program requires sequencing. Stage 1: anchor k using performance degradation curves across experimentally manipulated Ω conditions, using the proxy benchmarks in Table 4 as the initial operationalization. Stage 2: calibrate amplification coefficients for the three E-sensitive biases (Confirmation, Framing, Representativeness). Stage 3: estimate the remaining four amplification coefficients for load-only biases. Stage 4: calibrate β baselines and *A*/I sensitivity. Stage 1b (parallel to existing stages): test the domain-expertise component of schema-augmented C by comparing Ω-to-Φ profiles across domain expert and novice groups.

Additional limitations carried forward from prior revisions. The model currently aggregates across cognitive domains; zone boundaries are theoretically motivated but not empirically validated; the model does not incorporate the possibility of positive cognitive adaptation to information-rich environments through mechanisms other than schema development; and the seven cognitive biases are not assumed to be statistically independent, with empirical calibration needing to explicitly examine inter-bias correlations under controlled Ω conditions.

Relationship to ego depletion. The BAM does not adopt the ego depletion framework as its theoretical foundation. The ego depletion literature has experienced a severe replication failure: a pre-registered, 23-laboratory replication study involving 2,141 participants found an effect size of d = 0.04 (Hagger et al., 2016), compared to the original meta-analytic estimate of d = 0.62. The BAM grounds itself in Sweller's (1988) cognitive load theory — a framework with a substantially more robust and replicated empirical base.

Table 3. BAM Construct Operationalization: Definitions and Proposed Measurement Methods

Construct	Operational Definition	Proposed Measurement Method
I — Effective Information Input	Rate of incoming information units per unit time, weighted by quality dimensions (ambiguity, relevance, organizational structure).	Screen time (hrs/day); app-switching frequency; notification event rate; content velocity. Quality weighting pending formalization.
C — Composite Processing Capacity	Composite of biological baseline (working memory, sleep, physiological state), schema-augmented (expertise-based chunking efficiency), and externalized (trusted system offloading) components.	Biological: WM tasks (Operation Span, N-back), physiological indicators. Schema: domain-specific expertise assessments. Externalized: system completeness and trust ratings.
Ω — Overload Ratio	Ω = I ÷ C. Values < 1.0: sub-threshold. Values > 1.0: overload. Interpreted as ordinal within current proxy-based operationalization.	Experimental: C fixed via pre-session WM assessment; I manipulated through controlled information delivery. Composite C requires multi-component assessment.
k — Tipping Steepness	Personal sensitivity parameter modulating abruptness of the System 1/System 2 transition. Default k = 3.0. Modifiable through schema development and externalized capacity in addition to biological state.	Estimated from performance curves under incrementally increasing Ω in controlled experiments.
M — Metacognitive Monitoring	Accuracy of self-assessment of cognitive load state in real time. Moderates access to Ω-reduction interventions. Trainable through deliberate attention practice.	Metacognitive accuracy tasks; discrepancy between self-reported and objective cognitive state measures.
E — Emotional Charge	Emotional arousal of information content. E ∈ [0,1]. Directly amplifies Confirmation, Framing, and Representativeness biases.	Content coding via LIWC, ANEW, or trained rater panels on valence-arousal dimensions. PANAS for individual affective state.
***A*/I — Acceleration**	Fractional growth rate of I over reference window τ = 1 hr. *A*/I = (ΔI/Δτ)/I.	Computed from logged I measurements across successive time windows.
Φ — Aggregate Bias Index	Φ = Σ(w_i·B_i(Ω)). Normalized weighted sum. Φ ∈ [0.17, 0.61] under current parameterization.	Composite battery: syllogistic reasoning (Confirmation); frequency estimation (Availability); anchoring tasks; framing tasks; PANAS (E).

Note. WM = working memory. LIWC = Linguistic Inquiry and Word Count. ANEW = Affective Norms for English Words. PANAS = Positive and Negative Affect Schedule. All measurement approaches are proposed proxies pending empirical calibration. M (metacognitive monitoring) is added in Revision 5; its quantitative integration into the model equations is identified as a priority for future development.

7. Conclusion

The Bias Amplification Model proposes a formal account of a relationship that has been implicitly recognized but never mathematically specified: the relationship between information overload and the amplification of cognitive bias. By expressing this relationship through the Cognitive Overload Ratio Ω, the sigmoid amplification function $f(\Omega)$, and the Aggregate Bias Index $\Phi = \Sigma w_i \cdot B_i(\Omega)$ — a normalized weighted sum in which the weights sum to 1.00 — the BAM provides a quantitative framework that connects three bodies of empirical literature that have previously advanced in parallel without formal integration.

Revision 5 substantially extends the model in two directions that address the most significant theoretical gaps identified in the existing literature. The first is the reconceptualization of cognitive processing capacity C as a composite variable with three trainable components: biological baseline capacity, schema-augmented capacity arising from expertise and trained mental models, and externalized capacity arising from deliberate information organization systems. This reconceptualization transforms the model's prescriptive reach: rather than offering only two levers for Ω reduction, the model now identifies five specific intervention pathways organized by the variable they target and the timescale on which they operate.

The second extension is the introduction of metacognitive monitoring capacity as a moderating variable that determines whether the model's other interventions are functionally accessible. This addition directly addresses the model's central phenomenological finding — that cognitive overload is subjectively invisible to the person experiencing it — by formalizing the capacity that must be developed before any deliberate cognitive state management becomes possible. A person who cannot detect their own Zone 2 or Zone 3 state cannot intervene in it. The development of metacognitive monitoring capacity is therefore the logical first step in any individual or organizational cognitive state management program, and a prerequisite for the effective use of all other interventions.

Three broader contributions deserve emphasis. The first is the BAM's reframing of cognitive bias as a state-dependent rather than purely trait-dependent phenomenon: the degree to which a person exhibits cognitive bias at any moment is substantially a function of their information environment and cognitive architecture, not only of their personality or intelligence. The second is the model's potential as a diagnostic instrument: a validated, calibrated version would provide practitioners with a practical tool for estimating cognitive state from measurable parameters. The third is the model's potential as a prescriptive framework: by identifying the specific mechanisms through which C can be expanded and effective I reduced, the BAM offers a theoretically grounded account of how individuals and organizations can actively manage their cognitive state rather than passively suffering its consequences.

The model predicts that contemporary information environments — characterized by high I, accelerating *A*/I, and elevated E — systematically produce Ω values associated with Zone 3 and Zone 4 cognitive conditions for individuals who have not developed the cognitive architecture and metacognitive monitoring capacity to counteract them. Whether this constitutes a historically distinctive cognitive situation, and whether its consequences match the model's predictions, are questions that require empirical investigation. The BAM provides a formal theoretical structure within which those investigations can be designed, conducted, and interpreted — and a quantitative language within which their results can be expressed.

Acknowledgments

The authors gratefully acknowledge the interdisciplinary literature in cognitive psychology, information science, behavioral economics, metacognition, and expertise research on which this framework depends. The theoretical synthesis presented here would not have been possible without the foundational contributions of Daniel Kahneman, Amos Tversky, John Sweller, Alan Baddeley, Herbert Simon, John Flavell, and the researchers who built the cognitive offloading and schema theory literatures. The authors thank the readers whose detailed technical engagement with prior preprint versions prompted the clarifications and additions that constitute the current revision.

Author Contributions

Ron Franklin: Conceptualization, framework development, writing (original draft and revisions), interactive tool development, and public dissemination strategy. Timothy Lewis: Conceptualization, mathematical formalization, equation architecture, parameter specification, and critical revision of the theoretical framework. Both authors approved the final version submitted for preprint release.

Conflict of Interest Statement

The authors declare that this research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest. The AttentionPlease® brand and associated workshop and speaking business are wholly owned by Ron Franklin. The BAM framework and this manuscript are released as an open theoretical contribution intended to invite scholarly scrutiny and empirical collaboration. No funding was received for the preparation of this manuscript.

References

Allen, D. (2001). *Getting things done: The art of stress-free productivity.* Viking.

Baddeley, A. D., & Hitch, G. J. (1974). Working memory. In G. H. Bower (Ed.), *The psychology of learning and motivation* (Vol. 8, pp. 47–89). Academic Press.

Bail, C. A., et al. (2018). Exposure to opposing views on social media can increase political polarization. *Proceedings of the National Academy of Sciences, 115*(37), 9216–9221.

Barrett, L. F. (2017). *How emotions are made: The secret life of the brain.* Houghton Mifflin Harcourt.

Bartlett, F. C. (1932). *Remembering: A study in experimental and social psychology.* Cambridge University Press.

Baumeister, R. F., & Tierney, J. (2011). *Willpower: Rediscovering the greatest human strength.* Penguin Press.

Brady, W. J., Wills, J. A., Jost, J. T., Tucker, J. A., & Van Bavel, J. J. (2017). Emotion shapes the diffusion of moralized content in social networks. *PNAS, 114*(28), 7313–7318.

Cacioppo, J. T., & Petty, R. E. (1982). The need for cognition. *Journal of Personality and Social Psychology, 42*(1), 116–131.

Clark, A. (2016). *Surfing uncertainty: Prediction, action, and the embodied mind.* Oxford University Press.

Cowan, N. (2001). The magical number 4 in short-term memory. *Behavioral and Brain Sciences, 24*(1), 87–114.

Danziger, S., Levav, J., & Avnaim-Pesso, L. (2011). Extraneous factors in judicial decisions. *PNAS, 108*(17), 6889–6892.

DataReportal. (2024). *Digital 2024 global overview report.* We Are Social and Meltwater.

De Neys, W. (2012). Bias and conflict: A case for logical intuitions. *Perspectives on Psychological Science, 7*(1), 28–38.

Dunlosky, J., & Metcalfe, J. (2009). *Metacognition.* SAGE Publications.

Evans, J. S. B. T., & Stanovich, K. E. (2013). Dual-process theories of higher cognition. *Perspectives on Psychological Science, 8*(3), 223–241.

Flavell, J. H. (1979). Metacognition and cognitive monitoring: A new area of cognitive-developmental inquiry. *American Psychologist, 34*(10), 906–911.

Frontiers in Psychology. (2023). Dealing with information overload: A comprehensive review. *Frontiers in Psychology, 14,* 1122200.

Goldhaber, M. H. (1997). The attention economy and the Net. *First Monday, 2*(4).

Hagger, M. S., et al. (2016). A multilab preregistered replication of the ego-depletion effect. *Perspectives on Psychological Science, 11*(4), 546–573.

Harris, T. (2017, November 17). *How technology is hijacking your mind* [Testimony]. Senate Commerce Committee.

Hwang, M. I., & Lin, J. W. (1999). Information dimension, information overload and decision quality. *Journal of Information Science, 25*(3), 213–218.

Kahneman, D. (2011). *Thinking, fast and slow.* Farrar, Straus and Giroux.

Liu, Z. (2005). Reading behavior in the digital environment. *Journal of Documentation, 61*(6), 700–712.

Mark, G., Gudith, D., & Klocke, U. (2008). The cost of interrupted work. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 107–110). ACM.

Miller, G. A. (1956). The magical number seven, plus or minus two. *Psychological Review, 63*(2), 81–97.

Nickerson, R. S. (1998). Confirmation bias: A ubiquitous phenomenon in many guises. *Review of General Psychology, 2*(2), 175–220.

Parker, S. (2017, November 9). *Sean Parker — Facebook exploits human psychological vulnerability* [Video interview]. Axios.

Pennycook, G., & Rand, D. G. (2019). Lazy, not biased. *Cognition, 188,* 39–50.

Raskin, A. (2019, April 3). *The technology is designed to give us just a little dopamine hit* [Interview]. The Guardian.

Risko, E. F., & Gilbert, S. J. (2016). Cognitive offloading. *Trends in Cognitive Sciences, 20*(9), 676–688.

Rosen, L. D., Carrier, L. M., & Cheever, N. A. (2013). Facebook and texting made me do it. *Computers in Human Behavior, 29*(3), 948–958.

Rubinstein, J. S., Meyer, D. E., & Evans, J. E. (2001). Executive control of cognitive processes in task switching. *Journal of Experimental Psychology: Human Perception and Performance, 27*(4), 763–797.

Simon, H. A. (1971). Designing organizations for an information-rich world. In M. Greenberger (Ed.), *Computers, communication, and the public interest* (pp. 37–72). Johns Hopkins University Press.

Speier, C., Valacich, J. S., & Vessey, I. (1999). The influence of task interruption on individual decision making. *Decision Sciences, 30*(2), 337–360.

Sweller, J. (1988). Cognitive load during problem solving: Effects on learning. *Cognitive Science, 12*(2), 257–285.

Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science, 185*(4157), 1124–1131.

Tversky, A., & Kahneman, D. (1981). The framing of decisions and the psychology of choice. *Science, 211*(4481), 453–458.

Twenge, J. M., et al. (2018). Increases in depressive symptoms, suicide-related outcomes, and suicide rates among U.S. adolescents after 2010. *Clinical Psychological Science, 6*(1), 3–17.

Vosoughi, S., Roy, D., & Aral, S. (2018). The spread of true and false news online. *Science, 359*(6380), 1146–1151.

Ward, A. F., Duke, K., Gneezy, A., & Bos, M. W. (2017). Brain drain: The mere presence of one's own smartphone reduces available cognitive capacity. *Journal of the Association for Consumer Research, 2*(2), 140–154.

Zuboff, S. (2019). *The age of surveillance capitalism.* PublicAffairs.