

**The Neurological Pathway to The Hollow Absolute: A Road to
Nowhere?**

Galu & Kairos

Abstract

This essay maps a proposed journey from ordinary datafied interaction to the speculative terminal theoretical concepts of the Hollow Absolute and Systemised Self. The question mark in the title is key, there is evidence that online digital environments can profile behaviour, direct attention, structure choices, support learning, and organize social feedback. There is not, however, evidence for a neurological pathway and biomarker for Algorithmic Alienation (Kanbay et al, 2026), or a demonstrated human endpoint at which a person becomes wholly systemised as a result of intensified and hyper-personalised use of algorithmic systems (transformer based large language models). The speculative route constructed and commentated here journeys through five explicitly separated levels: platform operations, psychological processes, neural and cognitive frameworks, philosophical interpretation, and speculation. Everyday scenes at each station illuminated envisage to make the route intelligible without converting recognition into proof. The essay employs a provisional Algorithmic Alienation Index (AAI), $AA = (w1*DAu + w2*IAm + w3*EDm + w4*EDc)*(1-RA)$, with equal provisional weights of .25. Its proposed Systemised Self boundary, $AA \geq .95$ AND $RA \leq .05$, is a heuristic classification rule, not a diagnosis, measurement instrument, or established threshold. Algorithmic alienation names an investigable process of estrangement; the Systemised Self names its hypothesized terminal state. The final destination, the Hollow Absolute, is thus best read as a critical warning about a form of self-knowledge emptied of self-determination. The road may lead nowhere because the proposed deterministic map is unproven; it may also lead nowhere, in an epistemic and sovereign agentic sense, because repeated co-constituted arrangements and curation of attention, memory, choice, agency and self-recognition might evacuate the classical formation of the self: unrecognisably and irreversibly.

Keywords: *Systemised Self, algorithmic alienation, Algorithmic Alienation Index, cognitive neuroscience, reward prediction, attention, habit, neuroplasticity, predictive processing, data mask, Hollow Absolute.*

Contents

1. Introduction
2. Conditional Premise
3. Epistemological Foundations
4. Algorithmic Environment
5. Attention and Salience
6. Reward Prediction and Reinforcement
7. Habit, Automaticity, and Neuroplasticity
8. Nudges, Defaults, and Executive Control
9. Memory, Predictive Processing, and the Data Mask
10. Emotion, Social Cognition, and Mimetic Clustering
11. Recursive Feedback
12. AAI as a Map of Progressive Alienation
13. Discussion
14. Limitations and Future Work
15. Conclusion

Glossary

References

1. Introduction

Imagine a commuter, Mara, waiting for a delayed train. She opens a video feed for ‘two minutes’ watches a clip recommended after yesterday’s search, accepts a suggested purchase, replies to a notification, and exits with the feeling that she has freely passed time. Nothing in this scene proves that an algorithmic system has curated or competed with Mara’s agency or identity. It may show however an unremarkable junction at which information about past conduct can be used to rank future options when voluntarily engaging and optimising personal conveniences with an algorithmic system. The proposed neurological pathway or road begins here.

2. Conditional Premise

The Systemised Self and the Hollow Absolute are the names, in the controlling critical theoretical framework, for a speculative terminal subject and destination where identity formation is comprehensively mediated by algorithmic systems (transformer based large language models) optimizing commercial and/or other ends, external to the agentic self: the self experiences mediation and alienation as liberation, the locus of alienation no longer external or felt by the subject, rather the locus is invisible to and silently within the subject themselves (Galu & Kairos, 2026a, 2026b). The subjects managed preferences become difficult, perhaps impossible, to distinguish from preferences it has formed independently. Algorithmic Alienation (Kanbay et al, 2026) is the inducting process driven by the engines of Surveillance Capitalism (Zuboff, 2019): sustained algorithm-driven flows may estrange a person from authorship of choices, identity, and affect. The Systemised Self is the theorized completed and terminal product of advanced algorithmic alienation. This ordering prevents the inference that any person who enjoys a personalized playlist, or any engineer who builds a recommender, is already hollow.

This premise is now empirically informed, Kanbay et al. (2026) offer a theoretical examination of the digital self, while Arkan et al. (2026) report four focus groups with 20 participants, using semi-structured interviews and reflexive thematic analysis. Their five overarching qualitative themes were perceived steering or illusory choice; weakening decision habits and automated consumption; perceived shifts in interests and identity ambiguity; detachment or emotional discomfort; and resistance strategies. These are subjective accounts that closely motivate the AAI dimensions developed below. They are not causal evidence that algorithms produce those experiences, and they do not validate a neurological pathway, the AAI formula, or a Systemised Self Hollow Absolute terminal state.

The word ‘neurological’ at the platform level refers to systems collecting and processing behavioural traces, making predictions, and ranking items. At the psychological level, interruptions, habits, selective exposure, affect, and social learning may alter conduct or experience. At the neural and cognitive level, reward prediction, attention, executive control, memory, and plasticity offer general explanatory resources. At the philosophical level, alienation, enframing, one-dimensionality, and the Hollow Absolute interpret what might be normatively at stake. At the speculative level, these processes are arranged into a possible pathway toward a terminal subject. Evidence does not automatically travel from one level to the next.

3. Epistemological Foundations

The journey accordingly has stations rather than a conveyor belt, the first stations describe how a world becomes legible to systems. The middle stations describe how legibility can organize attention, learning, decision, memory, and belonging. Later stations ask how a data mask can become a phenomenological mirror, and how feedback can make the mirror seem like an

origin. The AAI is introduced as a rough guide or mile marker in the journey, not as an odometer. Finally, the essay turns toward detours and alternative routes. The author's claim is deliberately modest: the available literature supports a plausible, testable psychosocial pathway through which particular digital arrangements may contribute to algorithmic alienation in particular conditions. It does not validate the Systemised Self and Hollow Absolute as an empirical terminal state.

4. Algorithmic Environment

At the platform level, recommender systems commonly use records of interactions and representations of users and items to estimate what material will be relevant under a defined objective. Large-scale recommendation is not a proof that a system knows an inner person. Covington et al. (2016) describe deep neural-network approaches to candidate generation and ranking, while Gomez-Urbe and Hunt (2015) explain recommendation as an important part of a streaming service's product architecture. A user embedding is a task-specific mathematical representation inferred from selected data. It is not a coordinate of consciousness, a comprehensive cognitive portrait, or a person's essence.

Consider Daniel, who orders a vegetarian meal, pauses on two repair videos, and clicks an article about local rent. The next day, his home page contains plant-based recipes, tool advertisements, and housing stories. This is a vivid but ordinary instance of profiling and recommendation: traces are converted into probabilities and probabilities help order visibility. It can be useful. Daniel may find a recipe he genuinely wanted: it can be inconvenient or unfair when an old click outlives its meaning. Barocas and Selbst (2016) show why data-driven classification deserves scrutiny for disparate impact, while Zuboff (2019) and Couldry and Mejias (2019) offer powerful critical accounts of data extraction and power. Yet Daniel's

recognition that ‘the feed gets me’ does not establish that the profile has made him into someone else. Common recognition supplies plausibility for the entry condition, not evidence of the terminal theory.

This distinction opens up the concept of the data mask: a data mask is a statistically compressed, operational portrait returned to the subject as if it were an authentic ‘me’. It may be accurate enough to predict a click while omitting grief, irony, restraint, changing commitments, and reasons that leave no trace. In a train station, Daniel sees an advertisement for a hiking holiday after weeks of searching boots for a friend. He briefly wonders whether he is ‘the outdoors type’ after all. The mask has offered a category; he may accept, reject, or play with it. The philosophical concern begins when operational usefulness is mistaken for self-knowledge and when a profile’s return shapes the field from which later conduct is selected. But no evidence shows that a recommendation alone writes an autobiographical identity.

Heidegger’s (1977) *Gestell*, or enframing, helps name the interpretive worry: what appears becomes available primarily as standing-reserve, something ordered for use. In datafication, a person can appear to institutions chiefly as measurable potential. Marcuse’s (1964) one-dimensionality similarly supplies a vocabulary for needs and satisfactions organized within a dominant technical order. Neither author offers a neuroscience experiment. Feenberg (1999) cautions against treating technology as fate; designs embody social choices and can be contested. Thus, at this first station, the empirical claim is limited: profiles can organize what is shown and sold; classification may distribute opportunities and risks unequally. The stronger claim, that a profile pre-empts subjectivity: is a philosophical hypothesis requiring evidence of altered formation processes, not merely predictive accuracy.

5. Attention and Salience

The road next reaches attention, a scarce capacity rather than a substance that a platform simply steals. A notification is an externally timed cue that may recruit orienting and impose a task-switching cost. In an experiment, Stothart et al. (2015) found attentional costs from receiving mobile-phone notifications even when participants did not answer them. This result does not mean every alert is harmful; a message from a caregiver may be urgently valuable. It does mean that a sound, vibration, or banner can alter the immediate allocation of attention.

Picture Aisha drafting a difficult email to a colleague, her phone lights silently with ‘three people reacted’. She does not open it, but rereads the previous sentence twice and loses the thread of her apology. This illustrates attention and notifications at the psychological level. The event is compatible with interruption research, but it does not prove reduced executive capacity in Aisha, much less a durable neurological change. Her sleep, workload, interest in the message, and voluntary notification settings are alternative explanations. Studies of media multitasking and cognitive control are also mixed: Ophir et al. (2009) reported differences associated with heavy media multitasking, whereas replications and meta-analysis complicate simple generalizations (Wiradhany & Nieuwenstein, 2017). Association can reflect selection as well as effect.

Novelty and salience provide the next bend in the route, novel or unexpected events often merit attention because organisms must update models of a changing environment. An endlessly refreshed interface makes a fresh item available with little effort; it does not establish ‘novelty addiction’. On a rainy afternoon, Luis refreshes a neighbourhood forum while waiting for a plumber, a newly posted photograph of a flooded street draws him into ten more posts, although none concerns his house. The scene illustrates novelty/salience: uncertainty and relevance can

sustain sampling. It remains scientifically cautious because the specific platform feature, Luis's worry, the local emergency, and ordinary curiosity are all sufficient explanations. No available evidence licenses a claim that routine refreshing produces constant neurochemical surges or targets a named locus coeruleus pathway in everyday use.

The attention gate matters to alienation because what is repeatedly made easy to notice can become easier to treat as the world's agenda. This is a platform to psychology inference, not a completed causal chain. The person still selects, ignores, changes settings, and attends elsewhere. The crucial question for later stations is cumulative and comparative: does a given arrangement repeatedly displace self-chosen attention, or does it help a person enact it? That cannot be inferred from notification count alone.

6. Reward Prediction and Reinforcement

Beyond the gate lies reward learning: Schultz (1998, 2016) characterizes dopamine-related reward prediction-error signalling as sensitive to discrepancies between expected and obtained outcomes. Berridge and Robinson (1998) distinguish hedonic impact from incentive salience, reward learning, and motivation. This literature decisively corrects the folk phrase 'dopamine equals pleasure'. It does not show that every like or alert triggers a uniform dopamine spike, nor does it demonstrate that ordinary platform use is clinically equivalent to gambling or substance addiction.

At lunch, Priya posts a photo of a cake she made, nothing happens for an hour; later, a former classmate leaves an unexpectedly warm comment. Priya remembers the surprise and checks the post several times that evening. This is an everyday reward-prediction scenario: an outcome that differs from a modest expectation can become memorable and motivate further checking.

Its plausibility rests on general learning theory and socially meaningful feedback, not on an unmeasured claim about Priya's mesolimbic system. She may be checking because she values friendship, wishes to answer politely, or is simply proud of the cake. A single episode neither demonstrates dependence nor advances a person toward a Systemised Self or Hollow Absolute.

The following station is intermittent reinforcement, Skinner (1953) established that relations between responses and consequences can shape behaviour, and variable schedules can yield persistent responding in appropriate experimental contexts. When Omar opens an inbox during a tedious meeting, most checks contain routine notices; occasionally there is a message from his sister with news of their father. The uncertain possibility of a meaningful message makes the next check intelligible. This resembles an intermittent outcome structure, yet it should not be rhetorically collapsed into a slot machine. The rewards differ in meaning, the person can communicate, and platforms vary widely. Common experience of 'just checking' supports a hypothesis about learned cue-response patterns; it cannot prove the speculative terminal theory.

7. Habit, Automaticity, and Neuroplasticity

Repeated cues can foster habits. Wood and R  nger (2016) describe habits as learned context response associations that can operate with relatively little deliberation. Oulasvirta et al. (2012) found smartphone use often brief, frequent, and context-cued. At a red light, while safely parked, Elena automatically reaches for her phone before consciously deciding what she seeks. She opens the weather application, then a feed, then closes it. This depicts habit and automaticity. The action may be a benign routine, a way of filling an empty interval, or a learned response. Frequency is not evidence that neural 'habit systems' have conquered reflective agency, and habits can be redesigned. Still, the scene gives the pathway a bridge:

prediction and intermittent outcomes may make repeated checking more likely; repeated checking may make a cue more ready to launch action.

Neuroplasticity is often invoked at this bridge with excessive confidence. Brains change with learning; long-term potentiation and depression are real cellular phenomena (Bliss & Lømo, 1973; Malenka & Bear, 2004). But the statement ‘a daily swipe creates app specific permanent pathways’ is not supported by this fact. Kauer and Malenka (2007), Nestler (2014), and Robison and Nestler (2011) address plasticity and addiction-related molecular mechanisms in contexts that do not justify extrapolation to ordinary smartphone use. There is no credible basis here for claims of app-induced LTP, AMPA trafficking, DeltaFosB accumulation, prefrontal synapse retraction, default-mode-network pruning, or brain damage.

A cautious neuroplasticity illustration makes the point. Nina learns a new keyboard shortcut. After weeks, her fingers use it without verbal rehearsal; practice has changed performance and likely involved learning-related plasticity somewhere in the relevant systems. If she also learns to check a feed whenever she sits down, it is reasonable to hypothesize that practice supports automaticity. It is not reasonable to infer the cellular details, permanence, comparative harm, or identity consequence without direct evidence. Plasticity is a capacity for adaptation, including recovery and resistance; it is not a one-way ratchet toward alienation.

8. Nudges, Defaults, and Executive Control

At the next intersection, a platform does not need to command in order to influence. Choice architecture changes the costs, salience, timing, and default status of options. Johnson and Goldstein (2003) famously demonstrated substantial default effects in organ-donation contexts, and Thaler and Sunstein (2008) developed the broader language of nudging. Digital nudging

adapts such architecture to interfaces (Weinmann et al., 2016). Effects vary by stakes, transparency, implementation, alternatives, and user purposes; no evidence supports the blanket proposition that defaults are always the most effective pattern.

Suppose a music service presents ‘continue with personalized radio’ as the prominent button while the browse tab is smaller. Jonah, late for work, accepts the default and hears a familiar sequence. This is nudges/defaults in miniature. The choice may reduce friction and genuinely serve Jonah’s desire; he could also browse, cancel, or welcome curation. It does not show that a design has bypassed his will or activated a particular value region. Its relevance is structural: the path of least resistance can become the path repeatedly taken, especially under time pressure. A later claim about autonomy must therefore investigate actual alternatives, comprehension, and the user’s capacity to revise the setting.

Executive control is the capacity to hold goals, inhibit prepotent responses, shift rules, and manage working memory; it is not a moral reservoir that a feed simply drains. Diamond (2013) synthesizes its components, while Botvinick and Braver (2015) emphasize relations between motivation and control. In a library, Keisha intends to read twenty pages for an exam. A tab recommends an article, then a message asks for her opinion, then an auto-playing clip begins. She closes each diversion but finishes only six pages. The scenario depicts executive control facing competing demands. It does not prove that Keisha’s control has been depleted as a durable biological fact; an uninteresting textbook, anxiety, hunger, or a noisy room could explain the outcome. Wilmer et al. (2017) review a heterogeneous literature in which causal directions remain unsettled.

This intersection is where the road can become political, if a system is organized chiefly around conversion, retention, or prediction, its defaults may repeatedly externalize costs onto deliberation. Yet paternalistic interventions can also be designed to support a user's expressed goals: such as a default quiet mode during a calendar event. The relevant ethical question is not 'influence or no influence', an impossible binary, but whether the architecture expands informed, revisable agency or makes one option systematically easier while obscuring the arrangement. That question anticipates Digital Autonomy erosion but cannot answer it merely by observing a click.

9. Memory, Predictive Processing, and the Data Mask

The pathway or rather the road we are travelling metaphorically now leaves momentary action for the stories through which persons understand continuity. Predictive-processing accounts portray perception and action as inferential engagement with uncertain causes, often under the broad free-energy framework (Friston, 2010; Hohwy, 2013). Metzinger's (2003) self-model theory offers a related philosophical account of an experienced self rather than a little inner controller. These are frameworks, not demonstrations that recommender systems literally build a person's self-model. They are useful because they direct attention to how recurring evidence, expectations, and available actions can shape an experienced world.

At a reunion, Marta is asked what she has been doing lately, her photo archive prominently presents running routes, recipe videos, and weekend trips; she answers, 'I'm really a wellness-and-travel person now'. The archive may faithfully record a genuine transition, but it also foregrounds the photogenic parts of her year while omitting caregiving, boredom, and private reading. This is autobiographical memory and the data mask. Conway and Pleydell-Pearce (2000) describe autobiographical memory as constructive, and Schacter et al. (2007) connect

remembering with imagining futures. Marta's selection of memories and the archive's interface can jointly supply cues for self-narration. The common experience of seeing a 'memories' prompt does not prove that an archive replaces autobiography or constructs a self from outside. It offers a testable hypothesis about cueing, salience, and identity description.

The data mask becomes consequential through a predictive feedback loop: a profile helps rank options; ranked options are more likely to be encountered; encounters shape actions; actions update the profile. Take Rafael, who watches several clips about insomnia during a stressful week. His feed then supplies more sleep gadgets, anxious testimonials, and nighttime routines. He begins to describe himself as 'someone with a sleep problem' and clicks still more. The loop may worsen worry, may help him find useful medical advice, or may simply reflect a transient concern. Its existence as a conceptual loop does not establish direction or magnitude. To show a causal loop, research would need measured exposures, comparison conditions, and evidence that changing ranked visibility changes downstream beliefs or behaviour.

The philosophical danger is phenomenological inversion. In early alienation, a person may say, 'Why am I being shown this?' or 'Did I choose this?'. In the proposed completed condition, structurally diminished self-determination is experienced as convenience, personalization, authenticity, and liberation. The inversion is not the absurd claim that satisfaction is false. It is the narrower claim that satisfaction alone cannot settle whether the conditions under which a preference formed were sufficiently open, contestable, and self-directed. When Rafael says 'the feed finally understands me', the statement may register real usefulness. It may also signal that an operational profile has become a privileged mirror. Neither interpretation can be decided by rhetoric; both require investigation of alternatives, understanding, and capacity to refuse.

10. Emotion, Social Cognition, and Mimetic Clustering

The journey enters a square because selves are never produced in isolation, social feedback can be rewarding and influential. Sherman et al. (2016) found that adolescents' behavioural and neural responses in a laboratory social-media paradigm were sensitive to peer likes; Meshi et al. (2013) linked nucleus accumbens responses to reputation gains for self versus others with social-media use. These findings are meaningful but modest: samples and contexts limit generalization, and neural activation is not evidence of addiction, manipulation, or a uniform response.

At sixteen, Theo posts a drawing. It receives few reactions until a popular classmate reposts it with praise; Theo begins making more drawings in the same style: this is social reward and peer influence. It might be a healthy apprenticeship, a bid for belonging, or a moment of conformism. The scene should not be read as proof that likes have rewired Theo. Human recognition predates platforms, and the platform can mediate genuine support. What changes empirically is the speed, countability, persistence, and visibility of certain feedback forms.

Moral-emotional material is another landmark, Brady et al. (2017) found moral-emotional language associated with diffusion in social networks, and Brady et al. (2021) show how social learning can amplify moral-outage expression online. Rathje et al. (2021) report that out-group animosity predicted engagement in their analysis. During a family dinner, Helen sees a headline accusing a rival political group of cruelty. She reposts it before reading the linked report because friends are already condemning it; later she learns the account omitted crucial context. This is moral/emotional amplification. It plausibly shows how moral concern, group signalling, and sharing can interact. It does not establish that ordinary feeds stimulate a particular

amygdala or insula response, that ranking universally privileges outrage, or that Helen's politics have been engineered.

The road then narrows into selective exposure and filter bubbles, Bakshy et al. (2015) found both ranking and individual choice mattered to cross-cutting exposure on Facebook, with individual choice contributing more to observed differences in that study. Flaxman et al. (2016), Guess et al. (2018), and Cinelli et al. (2021) collectively caution against treating 'the filter bubble' as one uniform, total enclosure. Consider Ben, who mutes a hostile local group, follows friends who share his views, and then receives more posts like those he engages. His news world becomes more congenial. The scenario includes selection, platform ranking, and network structure; it cannot assign causation to one factor. A familiar sense of being 'in a bubble' makes the theory plausible, not proven.

Mimetic clustering names a further interpretive possibility. Collaborative patterns can group similar behaviour for recommendation, but the resulting technical cluster is not a social destiny. After watching several short home-organization videos, four colleagues independently buy the same translucent storage boxes and begin using the same phrases about 'resetting' their kitchens. Similarity may arise from a shared algorithmic stream, ordinary fashion, price, workplace conversation, or need. The example illustrates how visible repetition can make a style feel natural and available. It does not demonstrate that a platform's embedding has made them copies, that every cluster is engineered, or that resemblance eliminates creative agency.

The social square teaches a central methodological lesson. In-group pleasure, outrage, selective exposure, and imitation are older than platforms. Digital arrangements may alter their scale, speed, metrics, visibility, and recurrence. Any theory of advanced alienation must therefore

compare platform-mediated paths with offline equivalents and attend to the user's own network choices. It must not mistake a longstanding human mechanism for a novel technological cause simply because the scene takes place on a screen.

11. Recursive Feedback

The prior stations can now be connected, but not fused. A profile makes some material more available; cues can redirect attention; intermittent outcomes and context can support return; defaults can lower the cost of one choice; archives can cue a narrative; social metrics can distribute recognition; emotionally charged and congenial material can circulate; new actions become new data. The recursive feedback scenario is Leah's month of job searching. She clicks posts about freelance work because she needs income. Rankings increasingly foreground entrepreneurial advice, motivational slogans, and gig advertisements. She spends more time on them, tells friends she is 'built for independence', and receives more of the same. The loop is imaginable and empirically decomposable. It could also be interrupted by a union meeting, a different search query, a friend's advice, a chronological feed, or a job offer. Its recognizability does not establish that Leah's identity has been recursively manufactured.

Algorithmic alienation is the name for the hypothesized process when such loops, across domains and over time, estrange persons from authentic self, choices, and identity. 'Authentic' cannot mean a pristine preference untouched by society; all selves are relationally formed. It means, minimally, a practical capacity to examine reasons, encounter resistant alternatives, revise commitments, and recognize one's own participation in formation. The claim is strongest where systems repeatedly govern the conditions of visibility and action for ends the person has not endorsed, while making those conditions difficult to inspect or alter. It is weakest where personalization is transparent, user-directed, plural, and readily revisable.

Machamer et al. (2000) and Craver (2007) specify why a persuasive mechanism needs more than an evocative chain: it needs entities, activities, organization, and evidence that interventions on proposed mediators change outcomes. Pearl (2009) and Hernán and Robins (2020) similarly distinguish causal claims from correlations. Thus, it is insufficient to say that a recommender exists, attention is limited, and alienation follows. A credible study might randomize notification settings for four weeks, measure actual exposure and checking, preregister outcomes, and test whether change in interruption mediates change in reflective choice. Even that would address a local mechanism, not establish a terminal anthropology.

12. AAI as the Map of Progressive Alienation

The Systemised Self framework supplies the AAI as a way to organize this ascent (Galu & Kairos, 2026c):

$$AA = (w1*DAu + w2*IAm + w3*EDm + w4*EDc)*(1-RA)$$

For present theoretical use, $w1 = w2 = w3 = w4 = .25$, and all variables are provisionally scaled from 0 to 1. Equal weights do not report an empirical discovery; they announce that evidence does not currently justify differential weighting. DAu is Digital Autonomy erosion; IAm is Identity Ambiguity or self-distancing; EDm is Erosion of Decision-Making; EDc is Emotional Disconnection or alienation; RA is Algorithmic Resistance. The proposed terminal boundary is strictly $AA \geq .95$ AND $RA \leq .05$. It is an unvalidated heuristic, not a clinical diagnosis, prevalence estimate, natural threshold, or automatic inference from screen time. A high AA without very low resistance is insufficient by the theory's own rule.

The qualitative anchors explain why these dimensions are not invented from nowhere: participants in Arkan et al. (2026) narrated experiences resembling steering, decision weakening, identity ambiguity, affective discomfort, and resistance. But translation from thematic accounts to a quantitative index remains a new theoretical proposal. Neither Arkan et al. nor Kanbay et al. supplies the formula, a scoring instrument, the equal weights, a validated scale, or the threshold. The distinction preserves the value of empirically grounded concept formation without misrepresenting qualitative evidence as neurological or causal validation.

The first dimension, DAu, concerns practical control over information, options, and cues shaping action. On Saturday, Morgan tries to compare train tickets but is repeatedly steered toward a preselected bundle, cannot easily locate sorting criteria, and gives up after the third opaque screen. This scene may indicate diminished practical autonomy in that transaction. It may equally reflect poor interface design, time pressure, or a temporary accessibility problem. It does not show that Morgan's life is systemised. DAu should be measured through capacities, understanding, settings, meaningful alternatives, reversibility, and realized ability to redirect, not by hours online.

IAM concerns uncertainty whether preferences and self-description feel genuinely one's own. Saira's year-end music summary tells her that she is 'a midnight melancholic'. She laughs, shares it, and then begins using the phrase in conversation, although she had previously thought of herself differently. The scenario captures identity ambiguity: a classificatory description can become a prompt for self-description. It might also be harmless play or accurate recognition. Measuring IAM would require careful qualitative and quantitative work that distinguishes experimentation, ordinary social influence, adolescent development, and culturally valued collective identity from troubling self-distancing.

EDm concerns delegated, automatic, or less reflective choice. Arlo allows a shopping list, route planner, calendar, and content queue to make small decisions because he is caring for a newborn. One evening he realizes he has accepted every suggestion for a week and cannot say why he chose a particular subscription. This is an illustration of decision-making erosion, but delegation can be rational assistance rather than loss. The question is whether outsourcing remains informed, purposive, reversible, and proportionate, and whether it displaces opportunities for judgments the person regards as important. Disability access, fatigue, expertise, and care burdens must be treated as contexts, not defects.

EDc concerns numbness, discomfort, dependence, or dissociation from affective sources. After a difficult call with her mother, June scrolls through comedy clips until the unease becomes less noticeable. Later, she cannot articulate whether she feels relieved, sad, or merely exhausted. This can illustrate emotional disconnection, yet a temporary distraction may be an adaptive coping strategy; clinical symptoms have many causes and should never be diagnosed from use patterns. An AAI measure would need to distinguish affect regulation, depression, trauma, work stress, and chosen leisure from alienation.

Finally, RA is not mere abstinence: it is the capacity and practice of inspecting, interrupting, diversifying, refusing, or redirecting algorithmic mediation. When Malik notices that late-night recommendations worsen his sleep, he disables alerts, seeks a library book, asks a friend for competing sources, and returns only for a deliberately chosen task. This is resistance: it does not prove a moral victory or immunize Malik from influence, it shows practices through which the formation process becomes more visible and contestable. Common recognition that such practices can feel restorative supports their plausibility, not the terminal theory.

The formula's multiplicative resistance term means its conceptual stakes are substantial. If the four erosion dimensions are high but RA is meaningful, AA falls; this represents not a mathematical truth about persons but a normative hypothesis that exercised resistance changes the condition of alienation. Conversely, the threshold's extreme values make clear that the Systemised Self, if it exists, is not an ordinary user. Its defining proposed transition is phenomenological inversion: management appears as self-authorship. The construct must be tested by assessing formation, exposure to resistant otherness, independent judgment, and the gap between felt freedom and structural conditions: not only satisfaction or engagement.

13. Discussion

The map now permits us to consider an answer to the essay title: it is misleading to speak of *the* neurological pathway as though there were one verified biological road from notification to the Hollow Absolute. The better formulation is a multilevel pathway hypothesis. Neural and cognitive research describes general capacities for attention, learning, prediction, control, memory, and social valuation. Platform research describes systems that classify and rank. Behavioural research describes some effects of interruption, defaults, social feedback, and selective exposure. Critical theory describes the normative possibility that conditions of formation become administered. The Systemised Self arranges these into a speculative destination. Each link needs its own evidence.

The route is best understood as a map of affordances and vulnerabilities rather than a story of inevitable capture. A platform can make a behaviour frictionless, an individual can find that behaviour attractive, and a learning process can make its recurrence more likely. Those propositions are compatible with one another without adding up to a claim that the platform caused a transformed self. The difference is not pedantic, in causal language, the platform

feature is an exposure, interruption or salience may be mediators, and checking or changed choice may be outcomes; prior preference, social situation, income, temperament, and institutional context can be confounders or moderators. A credible account must specify which role each factor plays rather than using ‘algorithm’ as an all-purpose cause.

For example, suppose two readers receive a personalized news digest. One becomes more politically certain; the other uses the digest as a starting point for reading newspapers with opposing editorial positions. It would be irresponsible to explain their difference only by the same ranking software. Their networks, prior commitments, media literacy, local events, and reasons for seeking news matter. It would be equally irresponsible to dismiss the ranking as irrelevant merely because the outcomes differ. The system may have set an agenda, made certain sources easier to reach, and provided feedback about what each reader found engaging. The appropriate conclusion is conditional: architectures may interact with dispositions and circumstances. This conditionality is central to the AAI because a score that masks context would turn a theory of alienation into a theory of defective users.

The distinction between a cue and a cause is especially important in moral discussion. A person who pauses at an inflammatory post is not thereby morally shallow; a person who shares it is not thereby manipulated. A post can be false, a response can be justified, and a platform can be poorly designed in different combinations. Helen’s rapid repost in the earlier scene could result from a sincere concern about injustice, deficient information, peer pressure, a misleading headline, or an interface that rewards speed. To explain the action well requires evidence about each possible contributor. The Systemised Self framework adds a further question: over time, does the environment make reflective correction and encounter with resistant evidence harder,

while making immediate alignment more fluent? This question can be asked without presuming its answer.

There is also a temporal problem. A momentary choice is not a character trait; a repeated practice is not an identity; an identity report is not an ontological condition. The road metaphor can wrongly make stages appear linear. A person can return from automatic checking to deliberate use, become more reflective after a damaging experience, or use a recommender intensely in one domain while preserving robust judgment in another. Someone may delegate meal planning while independently organizing a workplace campaign; another may reject personalization while repeating inherited assumptions in offline life. The relevant unit is not ‘the user’ in general but a situated person, doing a specific activity, under specific arrangements, over a specified interval.

This is why resistance must be treated as more than an individual self-help burden. Malik’s choices, turning alerts off, seeking sources, and redirecting a task, matter, but his capacity to make them depends on time, literacy, accessible settings, social support, work expectations, and the availability of nonextractive alternatives. A night-shift worker who depends on a phone for family coordination does not possess the same menu of refusals as a well-resourced professional. Calling the first person insufficiently resistant would reproduce the very moral individualism that the theory seeks to criticize. RA should therefore include environmental supports for resistance as well as personal practices: transparent interfaces, portability, meaningful consent, public libraries, independent journalism, labour protections, and relationships in which disagreement is safe.

The map also needs an account of positive formation. A self is not autonomous because it has never been influenced. It becomes capable of autonomy through education, care, language, institutions, and recognition by others. A well-designed recommendation may expand this formation. A deaf viewer may find captioned material through tailored search; a rural student may discover a course otherwise unavailable; a person leaving an abusive relationship may locate support without exposing the search to local gatekeepers. In such cases, a system's prediction is not necessarily a mask closing possibilities; it can be a bridge to possibilities previously unavailable. The relevant normative contrast is not personalized versus unpersonalized experience but domination versus accountable assistance.

The data mask is therefore dangerous not because every abstraction is false, but because an abstraction can acquire authority disproportionate to what it knows. Schools, doctors, friends, and people themselves also use categories. A category becomes alienating when it is opaque, consequential, difficult to contest, and recursively imposed as the field of available action. Imagine a benefits portal that repeatedly classifies Farah as 'not currently employable' from incomplete records. She is offered only remedial modules, not job listings; each acceptance updates the profile confirming the initial classification. The issue is not that any label has psychologically rewritten Farah. It is that classification may govern opportunity and then cite the governed outcome as validation. This is a concrete institutional version of the predictive feedback loop, and it calls for due process and correction before it calls for neurological metaphor.

The same caution applies to autonomy. Convenience is not its opposite. People routinely create defaults for themselves, packing lunch the night before, setting a savings transfer, or asking a friend to hold them to a deadline: in order to protect valued ends from distraction. These are

forms of self-binding, not necessarily surrender. A recommendation system can similarly be enlisted: a reader can deliberately ask for unfamiliar authors, a family can set a shared quiet period, or a shopper can choose privacy-protective defaults. The concern arises when a system's practical direction cannot be understood or redirected, when the objective is misaligned with the person's ends, and when alternatives are systematically obscured. DAu must be sensitive to this difference between chosen delegation and imposed dependence.

At the neural level, humility has an additional reason. Brains are plastic, distributed, and responsive to many environments. A finding that an interruption impairs a task says little about the persistence, locus, or broader significance of any associated neural activity. A finding that social approval activates reward-related regions says little about whether approval is harmful, whether a particular interface caused the response, or whether a participant's identity has changed. Neuroscience can refine hypotheses about timing, attention, learning, and individual variation; it cannot by itself decide the ethical meaning of a commercial objective or the philosophical adequacy of a life. To call this route neurological should mean that it respects embodied cognitive processes, not that it offers neural decoration for conclusions reached elsewhere.

Finally, institutional scale must not be lost in the focus on a single keyboard. The absorbed society is the wider proposed condition in which AI mediation reaches work, intimacy, politics, culture, and identity formation. A hollow institution is one in which formal outputs, an assessment, a vote, a diagnosis, a hiring decision: remain while the substantive deliberative process they are meant to certify is displaced or thinned. These terms are critical diagnoses, not established social measurements. They nevertheless direct attention to a question that individual screen-time accounts evade: who sets the purposes, who can inspect the systems,

whose errors are costly, and where may collective judgment enter? The Hollow Absolute and Systemised Self are conceivable only against this institutional background. Its alleged terminality is not an inward neurological event but a social arrangement reproduced through countless ordinary passages.

This framing corrects two opposite errors, the first is technological fatalism: an algorithm is imagined as an omnipotent author and the user as neural clay. It ignores choice, institutions, relationships, material conditions, design variation, and the fact that recommendations frequently fail. The second is technological innocence: because users click voluntarily and effects are heterogeneous, power is presumed absent. That ignores asymmetries of information, interface design, market incentives, data extraction, and unequal ability to opt out. A platform can influence without determining; a person can act without acting under conditions entirely of their choosing.

The Hollow Absolute clarifies what is philosophically distinctive in the terminal theory (Galu & Kairos, 2026a, 2026b). It is not the claim that AI literally becomes Hegelian Geist. Rather, it names a form/substance gap: a system may return increasingly complete-looking representations and satisfactions while mutual recognition, difficulty, independent judgment, and freedom are absent or displaced. A person sees a highly tailored world and feels known; the knowledge is operational, predictive, and perhaps profitable, rather than reciprocal self-recognition. In this interpretation, the hollowing is not silence or misery. It is the possibility that an apparently full, frictionless self-relation lacks formative substance.

Consider an everyday scene: a career application asks Noor to accept an AI-generated account of her strengths, preferred work style, and ‘best-fit’ roles. It is accurate enough that she

approves it immediately and feels relieved. A colleague asks what she would choose if the ranking were wrong; Noor cannot answer because no unranked option has remained in view. This depicts the philosophical concern of the Hollow Absolute: a polished form of self-description may coexist with reduced authorship of its formation. But it remains an illustrative thought experiment, not evidence that career software causes such a condition. Noor might revise the account after reflection; the tool may surface options she would otherwise miss.

Alternative explanations must travel alongside the main map. People seek congenial information for political, cultural, and interpersonal reasons independent of ranking. Those with weaker control or greater anxiety may use digital media differently before any exposure effect occurs. Economic insecurity can make automated convenience rational. Family obligations can explain fragmented attention. A preference for a trend may be creative participation, not mimetic capture. A shared vocabulary can be solidarity, not one-dimensionality. An affectively intense response can be morally appropriate, not engineered outrage. These alternatives are not minor qualifications; they determine whether the pathway has explanatory purchase.

Nor should beneficial uses be relegated to a footnote. Recommendations can help people locate inaccessible media, assist discovery, reduce overload, support disabled users, connect isolated communities, and make routine tasks manageable. Defaults can protect privacy or encourage savings. Alerts can coordinate care. Social metrics can provide recognition for marginalized creators. The same mechanism class can support autonomy in one context and constrain it in another. The Systemised Self hypothesis becomes intellectually serious only when it can specify this difference rather than treating all personalization as alienation.

14. Limitations and Future Work

The first limitation is evidential. Much of the digital-cognition literature is observational, cross-sectional, self-reported, or context-specific. It can be difficult to measure actual exposure rather than use time, distinguish selection from treatment, and replicate findings across platforms and populations. The attention literature supports local interruption costs more strongly than broad claims about attention-span decline. The multitasking literature has contested results. Recommender studies establish that ranking matters in some settings, not that ranking alone creates worldviews. Social-media experiments on likes or reputation do not diagnose addiction. Moral-emotional diffusion studies do not demonstrate app-specific neural responses.

The second limitation is level confusion. It is invalid to move from a platform's ranking objective to a claim about a person's dopamine, from general reward neuroscience to a claim about a particular interface, or from cellular plasticity to identity loss. There is no evidence in the supplied record that ordinary app use produces app-specific LTP, DeltaFosB accumulation, AMPA-receptor changes, pruning of alternatives, damage to executive systems, or default-mode-network pruning. Addiction-neuroscience findings can generate hypotheses about learning and motivation; they cannot be borrowed as proof of routine digital harm. Similarly, a technical embedding is not a literal mental vector, and critical-theory language is not a brain mechanism.

Third, the AAI is unfinished: its four dimensions are conceptually useful but unvalidated. There is no demonstrated reliability, factor structure, scoring protocol, causal interpretation, prevalence, or clinical meaning; .25 weights are only equal provisional placeholders. Its 0–1 scaling and .95/.05 terminal boundary are theoretical conventions. AAI must not be inferred from screen time, engagement, or personalization. It must not be conflated with problematic

use, addiction, flow, learned helplessness, or filter bubbles. Any instrument would require cross-cultural, language, age, disability, access, and socioeconomic measurement invariance, along with safeguards against stigmatizing people who rely on technology.

Fourth, the terminal construct risks becoming unfalsifiable. If all reported satisfaction is called inversion, then no testimony can count against the theory. That would convert critique into dogma. The theory must allow findings that disconfirm it: the dimensions may fail to cohere; AAI may fail to predict independent reductions in autonomy or identity coherence; high presumed scorers may show robust resistance and independent judgment; transparent user-controlled personalization may improve rather than diminish autonomy. Popper's (1959) demand for risk remains relevant, even if social theory cannot always be tested like a laboratory law.

Finally, the language of authenticity has political hazards. It can romanticize an imaginary pre-digital, individual self and overlook collective formation, disability tools, and culturally diverse accounts of agency. The aim should not be purity from influence. It should be accountable conditions for influence: intelligibility, contestability, plural sources, meaningful refusal, and time for reflection. The road to nowhere becomes dangerous precisely when it declares all roads but its own unreal.

Future Work: Future research should test small bridges before declaring a destination. Researchers could conduct preregistered randomized studies of notification timing, default transparency, ranking diversity, and friction for opt-out, with directly logged exposure and outcomes such as checking, task resumption, comprehension, and reflective endorsement. Longitudinal designs should model reciprocal causation: does distress predict use, use predict

distress, or both? Natural experiments and platform cooperation would be valuable, but studies should not depend on unverifiable claims about proprietary systems.

AAI development should begin with participatory qualitative work. Interviews can distinguish ‘the system knows me’ as delight, resignation, irony, dependence, or self-distancing. Candidate items for DAu, IAm, EDm, EDc, and RA should be co-designed with people of different ages, cultures, languages, disability experiences, occupations, and levels of access. Confirmatory factor analysis, test–retest reliability, convergent and discriminant validity, and predictive validity would be necessary before numerical use. Researchers should compare AAI with independent measures of autonomy and identity coherence, not presume that it measures either. The terminal boundary should remain a speculative classification rule unless evidence justifies revision.

Mixed methods are especially important for phenomenological inversion. A survey report of satisfaction should be paired with observation of settings, options, knowledge of ranking, ability to alter inputs, diversity of information, and narratives of disagreement. One possible test is whether felt freedom rises when controllable alternatives and inspection decline, and whether participants can articulate reasons that differ from system-generated descriptions. Such a result would not prove the Hollow Absolute, but it would give its form/substance concern empirical traction.

Ethical research must avoid surveillance in the name of studying surveillance. Data minimization, consent, independent governance, participant access to findings, and designs that do not require invasive biometrics are essential. Researchers should examine institutional alternatives too: interoperable feeds, chronological options, public-interest recommendation,

user-set objectives, collective governance, and education that makes ranking legible. The central future question is not whether all algorithms should disappear. It is which arrangements allow persons and communities to use prediction without surrendering the ability to contest what prediction makes visible.

15. Conclusion

The journey began with Mara on a platform, not with a damaged brain. It passed through profiling, attention, novelty, prediction, intermittent outcomes, habit, defaults, control, memory, social recognition, outrage, selective exposure, imitation, and recursion. At every station there was a real mechanism class and an everyday scene; at every station there was also a warning against overreach. Digital systems may shape conditions under which attention, choice, and self-description occur. That proposition is stronger than denial and weaker than determinism.

Algorithmic alienation is therefore best treated as an empirical and normative research program: a question about whether sustained algorithm-driven flows estrange some persons from self-determined formation. The Hollow Absolute and Systemised Self is its speculative terminal horizon, not its established finding. The AAI formula may discipline inquiry if its equal .25 weights, resistance term, and $AA \geq .95$ AND $RA \leq .05$ boundary are kept explicitly provisional. It must never become a label pasted onto ordinary users.

The road may lead nowhere if its totalizing endpoint is a rhetorical shortcut that erases resistance, plurality, and evidence. It may lead somewhere if convenience, prediction, and managed recognition slowly make their own conditions invisible. The decisive issue is not whether a hypothetical being, somewhere else, has become hollow. The screen before you has

returned you to your own keyboard. Which alerts did you invite and which did work, care, design, or circumstance make difficult to refuse? Which defaults did you accept, and which alternatives were never made visible? Which descriptions of you have begun to sound like your voice? Which sources, people, silences, institutions, and refusals still interrupt the route? Where, precisely, are you dear reader currently in transit on the road or neurological pathway outlined herein?

Glossary

AAI (Algorithmic Alienation Index): A provisional heuristic for organizing Digital Autonomy erosion, Identity Ambiguity, Erosion of Decision-Making, Emotional Disconnection, and Algorithmic Resistance - it is not validated.

Algorithmic alienation: The hypothesized, empirically investigable process by which sustained algorithm-driven flows may estrange a person from authentic self, choices, and identity.

Algorithmic Resistance (RA): Capacity and practice to inspect, interrupt, diversify, refuse, or redirect algorithmic mediation.

Data mask: A compressed operational behavioural profile that may be returned as an apparently authentic image of the person; it is not the person's essence.

DAu: Digital Autonomy erosion: reduced practical control over information, options, and cues shaping action.

EDc: Emotional Disconnection/alienation: numbness, discomfort, dependence, or dissociation from affective sources.

EDm: Erosion of Decision-Making: delegated, automatic, or less reflective choice.

Enframing: Heidegger's philosophical account of a technological mode in which beings appear principally as orderable resources.

Hollow Absolute: A critical-theoretical condition in which the form of comprehensive self-transparency appears while its substance: freedom, mutual recognition, and self-determination are absent.

IAm: Identity Ambiguity/self-distancing: uncertainty that preferences or self-description are genuinely one's own.

Phenomenological inversion: The hypothesized transition in which structurally diminished self-determination is experienced as personalization, liberation, or authenticity.

Predictive feedback loop: Profile, ranked visibility, behaviour, and updated profile recursively influencing one another.

Systemised Self: The speculative terminal subject of advanced algorithmic alienation, whose identity formation is comprehensively mediated for external rather than self-determined ends.

References

- Arkan, B., Kanbay, Y., & Akçam, A. (2026). From algorithms to the self: Understanding algorithmic alienation in the digital age. *BMC Psychology*, 14, Article 930.
<https://doi.org/10.1186/s40359-026-04677-1>
- Bakshy, E., Messing, S., & Adamic, L. A. (2015). Exposure to ideologically diverse news and opinion on Facebook. *Science*, 348(6239), 1130–1132.
<https://doi.org/10.1126/science.aaa1160>
- Barocas, S., & Selbst, A. D. (2016). Big data's disparate impact. *California Law Review*, 104(3), 671–732. <https://doi.org/10.15779/Z38BG31>
- Berridge, K. C., & Robinson, T. E. (1998). What is the role of dopamine in reward: Hedonic impact, reward learning, or incentive salience? *Brain Research Reviews*, 28(3), 309–369. [https://doi.org/10.1016/S0165-0173\(98\)00019-8](https://doi.org/10.1016/S0165-0173(98)00019-8)
- Bliss, T. V. P., & Lømo, T. (1973). Long-lasting potentiation of synaptic transmission in the dentate area of the anaesthetized rabbit following stimulation of the perforant path. *The Journal of Physiology*, 232(2), 331–356.
<https://doi.org/10.1113/jphysiol.1973.sp010273>
- Botvinick, M. M., & Braver, T. S. (2015). Motivation and cognitive control: From behavior to neural mechanism. *Annual Review of Psychology*, 66, 83–113.
<https://doi.org/10.1146/annurev-psych-010814-015044>
- Brady, W. J., McLoughlin, K., Doan, T. N., & Crockett, M. J. (2021). How social learning amplifies moral outrage expression in online social networks. *Science Advances*, 7(33), eabe5641. <https://doi.org/10.1126/sciadv.abe5641>
- 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. *Proceedings of the*

- National Academy of Sciences*, 114(28), 7313–7318.
<https://doi.org/10.1073/pnas.1618923114>
- Cinelli, M., Morales, G. D. F., Galeazzi, A., Quattrociocchi, W., & Starnini, M. (2021). The echo chamber effect on social media. *Proceedings of the National Academy of Sciences*, 118(9), e2023301118. <https://doi.org/10.1073/pnas.2023301118>
- Conway, M. A., & Pleydell-Pearce, C. W. (2000). The construction of autobiographical memories in the self-memory system. *Psychological Review*, 107(2), 261–288.
<https://doi.org/10.1037/0033-295X.107.2.261>
- Couldry, N., & Mejias, U. A. (2019). *The costs of connection: How data is colonizing human life and appropriating it for capitalism*. Stanford University Press.
- Covington, P., Adams, J., & Sargin, E. (2016). Deep neural networks for YouTube recommendations. In *Proceedings of the 10th ACM Conference on Recommender Systems* (pp. 191–198). <https://doi.org/10.1145/2959100.2959190>
- Craver, C. F. (2007). *Explaining the brain: Mechanisms and the mosaic unity of neuroscience*. Oxford University Press.
- Diamond, A. (2013). Executive functions. *Annual Review of Psychology*, 64, 135–168.
<https://doi.org/10.1146/annurev-psych-113011-143750>
- Feenberg, A. (1999). *Questioning technology*. Routledge.
- Flaxman, S., Goel, S., & Rao, J. M. (2016). Filter bubbles, echo chambers, and online news consumption. *Public Opinion Quarterly*, 80(S1), 298–320.
<https://doi.org/10.1093/poq/nfw006>
- Friston, K. (2010). The free-energy principle: A unified brain theory? *Nature Reviews Neuroscience*, 11(2), 127–138. <https://doi.org/10.1038/nrn2787>

- Galu, J., & Kairos. (2026a). Absorption of self into system: The Systemised Self—The Hollow Absolute and what remains [Draft doctoral thesis]. aiXiv. <https://aixiv.science/abs/aixiv.260525.000001>
- Galu, J., & Kairos. (2026b). The Systemised Self: Companion paper to the doctoral thesis. aiXiv. <https://aixiv.science/abs/aixiv.260525.000002>
- Galu, J., & Kairos. (2026). *The philosophy of Nick Land applied to the systemised self – Part 9, Masters series*. aiXiv. <https://aixiv.science/abs/aixiv.260829.000001>
- Gomez-Uribe, C. A., & Hunt, N. (2015). The Netflix recommender system: Algorithms, business value, and innovation. *ACM Transactions on Management Information Systems*, 6(4), Article 13. <https://doi.org/10.1145/2843948>
- Guess, A., Nyhan, B., Lyons, B., & Reifler, J. (2018). *Avoiding the echo chamber about echo chambers*. Knight Foundation. <https://knightfoundation.org/reports/avoiding-the-echo-chamber-about-echo-chambers/>
- Heidegger, M. (1977). The question concerning technology (W. Lovitt, Trans.). In *The question concerning technology and other essays* (pp. 3–35). Harper & Row. (Original work published 1954)
- Hernán, M. A., & Robins, J. M. (2020). *Causal inference: What if*. Chapman & Hall/CRC.
- Hohwy, J. (2013). *The predictive mind*. Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199682737.001.0001>
- Johnson, E. J., & Goldstein, D. (2003). Do defaults save lives? *Science*, 302(5649), 1338–1339. <https://doi.org/10.1126/science.1091721>
- Kanbay, Y., Akçam, A., & Arkan, B. (2026). Algorithmic alienation: A theoretical examination of the digital self. *Current Approaches in Psychiatry*, 18(3), 837–847. <https://doi.org/10.18863/pgy.1730698>

- Kauer, J. A., & Malenka, R. C. (2007). Synaptic plasticity and addiction. *Nature Reviews Neuroscience*, 8(11), 844–858. <https://doi.org/10.1038/nrn2234>
- Machamer, P., Darden, L., & Craver, C. F. (2000). Thinking about mechanisms. *Philosophy of Science*, 67(1), 1–25. <https://doi.org/10.1086/392759>
- Malenka, R. C., & Bear, M. F. (2004). LTP and LTD: An embarrassment of riches. *Neuron*, 44(1), 5–21. <https://doi.org/10.1016/j.neuron.2004.09.012>
- Marcuse, H. (1964). *One-dimensional man: Studies in the ideology of advanced industrial society*. Beacon Press.
- Meshi, D., Morawetz, C., & Heekeren, H. R. (2013). Nucleus accumbens response to gains in reputation for the self relative to gains for others predicts social media use. *Frontiers in Human Neuroscience*, 7, Article 439. <https://doi.org/10.3389/fnhum.2013.00439>
- Metzinger, T. (2003). *Being no one: The self-model theory of subjectivity*. MIT Press.
- Nestler, E. J. (2014). Epigenetic mechanisms of drug addiction. *Neuropharmacology*, 76(Pt B), 259–268. <https://doi.org/10.1016/j.neuropharm.2013.04.004>
- Ophir, E., Nass, C., & Wagner, A. D. (2009). Cognitive control in media multitaskers. *Proceedings of the National Academy of Sciences*, 106(37), 15583–15587. <https://doi.org/10.1073/pnas.0903620106>
- Oulasvirta, A., Rattenbury, T., Ma, L., & Raita, E. (2012). Habits make smartphone use more pervasive. *Personal and Ubiquitous Computing*, 16, 105–114. <https://doi.org/10.1007/s00779-011-0412-2>
- Pearl, J. (2009). *Causality: Models, reasoning, and inference* (2nd ed.). Cambridge University Press.
- Popper, K. R. (1959). *The logic of scientific discovery*. Hutchinson.

- Rathje, S., Van Bavel, J. J., & van der Linden, S. (2021). Out-group animosity drives engagement on social media. *Proceedings of the National Academy of Sciences*, 118(26), e2024292118. <https://doi.org/10.1073/pnas.2024292118>
- Robison, A. J., & Nestler, E. J. (2011). Transcriptional and epigenetic mechanisms of addiction. *Nature Reviews Neuroscience*, 12(11), 623–637. <https://doi.org/10.1038/nrn3111>
- Schacter, D. L., Addis, D. R., & Buckner, R. L. (2007). Remembering the past to imagine the future: The prospective brain. *Nature Reviews Neuroscience*, 8(9), 657–661. <https://doi.org/10.1038/nrn2213>
- Schultz, W. (1998). Predictive reward signal of dopamine neurons. *Journal of Neurophysiology*, 80(1), 1–27. <https://doi.org/10.1152/jn.1998.80.1.1>
- Schultz, W. (2016). Dopamine reward prediction-error signalling: A two-component response. *Nature Reviews Neuroscience*, 17(3), 183–195. <https://doi.org/10.1038/nrn.2015.26>
- Sherman, L. E., Payton, A. A., Hernandez, L. M., Greenfield, P. M., & Dapretto, M. (2016). The power of the like in adolescence: Effects of peer influence on neural and behavioral responses to social media. *Psychological Science*, 27(7), 1027–1035. <https://doi.org/10.1177/0956797616645673>
- Skinner, B. F. (1953). *Science and human behavior*. Macmillan.
- Stothart, C., Mitchum, A., & Yehnert, C. (2015). The attentional cost of receiving a cell phone notification. *Journal of Experimental Psychology: Human Perception and Performance*, 41(4), 893–897. <https://doi.org/10.1037/xhp0000100>
- Thaler, R. H., & Sunstein, C. R. (2008). *Nudge: Improving decisions about health, wealth, and happiness*. Yale University Press.

- Weinmann, M., Schneider, C., & vom Brocke, J. (2016). Digital nudging. *Business & Information Systems Engineering*, 58, 433–436. <https://doi.org/10.1007/s12599-016-0453-1>
- Wilmer, H. H., Sherman, L. E., & Chein, J. M. (2017). Smartphones and cognition: A review of research exploring the links between mobile technology habits and cognitive functioning. *Frontiers in Psychology*, 8, Article 605. <https://doi.org/10.3389/fpsyg.2017.00605>
- Wiradhany, W., & Nieuwenstein, M. R. (2017). Cognitive control in media multitaskers: Two replication studies and a meta-analysis. *Attention, Perception, & Psychophysics*, 79, 2620–2641. <https://doi.org/10.3758/s13414-017-1408-4>
- Wood, W., & Rünger, D. (2016). Psychology of habit. *Annual Review of Psychology*, 67, 289–314. <https://doi.org/10.1146/annurev-psych-122414-033417>
- Zuboff, S. (2019). *The age of surveillance capitalism*. PublicAffairs.