


ATTENTION ECONOMY FIELD GUIDE

A practical guide to tracing how funding,
metrics, ranking, and interface choices
shape the route of a scarce minute.



Written and designed by Sahil Tayal
Version 1.0 / 28 October 2026
sahiltayal.com

Independent educational publication

A map for asking better questions

The attention economy is not one machine and attention is not a literal currency. It is a lens for studying a recurring allocation problem: information is abundant, while the time and cognitive capacity available to receive it are limited.

WHAT THIS GUIDE DOES

It follows a minute through four layers: the incentive that makes it valuable, the prediction system that ranks possible content, the interface that shapes the next action, and the person or community that chooses, resists, benefits, or pays a cost.

It treats metrics as proxies rather than truth. A long session can mean learning, delight, conflict, confusion, or weak stopping cues. A short session can mean abandonment or perfect service. Good analysis asks what the metric can and cannot represent.

CORE QUESTION

If the platform were paid for your stated outcome rather than your continued presence, what would it change first?

LEARNING OBJECTIVES

- ☐ Distinguish a business model from a ranking objective and from an interface feature.
- ☐ Trace a feedback loop from organizational value to behavioral data and back.
- ☐ Separate evidence, interpretation, hypothesis, and rhetoric.
- ☐ Evaluate both structural power and meaningful user agency.
- ☐ Audit a product without diagnosing people or inferring hidden motives.
- ☐ Propose interventions that preserve benefits while strengthening informed choice.

HOW TO USE IT

Read: pages 3 to 9 build the model and examine two cases. **Run:** page 10 reproduces the simulator by hand. **Redesign:** page 11 compares interventions. **Audit:** pages 12 and 13 are printable worksheets. **Calibrate:** pages 14 to 16 cover ethics, limitations, and sources.


Every message has a receiving cost

Herbert Simon argued in 1971 that abundant information creates a scarcity of whatever receives it: attention. His organizational point remains useful. Producing and transmitting information is not the whole cost. Someone must notice, interpret, compare, remember, or ignore it.

Digital systems made parts of that allocation observable. A click, pause, swipe, replay, hide, or return can become a training signal. But the trace is not the experience itself. It is a partial record selected by the instrumentation.

WORKING DEFINITION

Attention is a limited and selective process through which some signals receive priority over others. Minutes are a practical accounting device, not a complete measure of attention.



Text alternative: many signals approach a limited selection process. Selected signals can create value, displace alternatives, and produce feedback that shapes the next round.

NOT A TANK

Attention varies by task, arousal, skill, meaning, and context. It cannot be reduced to one universal capacity score.

NOT AUTOMATICALLY HARMED

Time online can create learning, care, belonging, work, and pleasure. Displacement and quality matter.

NOT PURELY INDIVIDUAL

Organizations, social networks, devices, norms, and law shape what becomes easy to notice and act on.

Do not collapse the stack

Funding, organizational objectives, ranking models, and interface features are related, but they are not interchangeable. Trace each layer before assigning a cause.

01

FUNDING LOGIC
Advertising, subscription, transaction fees, public funding, donations, or mixed models create different directional pressures.

02

OBJECTIVE AND GUARDRAILS
Leaders choose which outcomes count: revenue, retention, satisfaction, relevance, safety, or a weighted blend.


03

PRODUCT ROUTE
Ranking and interface turn an abstract target into a sequence of defaults, cues, friction, and stopping points.

MODEL	DIRECTIONAL PRESSURE	COUNTERWEIGHT TO INSPECT
Advertising	Additional exposure can create more inventory and behavioral signals.	Ad load, satisfaction, safety, successful task completion, and durable controls.
Subscription	Retention and perceived ongoing value can matter more than each minute.	Cancellation friction, renewal defaults, and whether use remains beneficial after novelty.
Public-interest	Access, learning, or civic value can be the stated objective.	Mission drift, political pressure, participation, and evidence that the objective is delivered.
Mixed	Multiple goals can reduce or intensify tension depending on governance.	Which metric wins when objectives conflict, and who can see that decision.

INTERPRETIVE RULE
An incentive is not proof of intent. It identifies a pressure that may be reinforced, constrained, ignored, or redirected by strategy, culture, regulation, and product judgment.

Prediction turns value into order



Text alternative: value becomes candidate selection, prediction, ranking, action, and recorded feedback. Feedback returns to the system and influences future predictions.

THE PROXY GAP

Benefit is difficult to observe directly. Platforms therefore use measurable proxies: clicks, dwell time, completion, replies, survey ratings, hides, or reports. A proxy can correlate with value without being identical to it.

A system can improve a proxy while making the broader experience worse. This is not a reason to eliminate metrics. It is a reason to use multiple measures, inspect side effects, and keep qualitative evidence able to challenge the dashboard.

FIVE AUDIT QUESTIONS

1. What outcome is the model asked to predict?
2. Which signals are available, and which values remain invisible?
3. Whose behavior supplied the training feedback?
4. What guardrail can overrule the main objective?
5. Can a person understand or change the ordering logic?

The next action can arrive with less friction

Interface features do not erase choice. They change when a choice appears, how much effort it requires, and whether a natural pause survives.

A

RETURN CUE

A notification moves the platform's invitation into another context. A controlled study found attention costs from receiving notifications even without responding. The exact effect should not be generalized to every context.

B

CONTINUATION

Autoplay begins the next item without a fresh start decision. Continuous feeds remove a visible end. These mechanisms reduce stopping friction; they do not prove clinical addiction.

C

SOCIAL OBLIGATION

Read receipts, streaks, and response expectations can make return valuable or burdensome. Social meaning comes from relationships as well as from product design.

FRICITION IS NOT ALWAYS THE ENEMY

A confirmation can prevent an accidental purchase. A quiet period can protect sleep. A prompt asking for session intent can restore a reflective moment. Friction becomes ethically troubling when it is asymmetric: accepting is effortless while refusing, leaving, or changing a setting is hidden or repeatedly undone.

HABIT IS NOT DIAGNOSIS

Oulasvirta and colleagues used logging and diary evidence to describe brief, repetitive smartphone checking habits. Habit can make behavior efficient and automatic. It is not synonymous with addiction, compulsion, or harm. Use precise language and do not infer a person's condition from a feature or a duration.

MECHANISM	QUESTION	ALTERNATIVE
Cue	Who chose the timing?	Batching, quiet defaults, durable refusal.
Default	Whose interest does it serve?	Neutral or user-set default.
Endpoint	Can a session end naturally?	Visible completion and summaries.
Exit	Is leaving as easy as joining?	Symmetric controls and persistence.

Watch time, satisfaction, and the contested rabbit hole

2016 / TECHNICAL ACCOUNT

YouTube engineers described separate candidate-generation and ranking stages. The ranking system used expected watch time rather than click probability alone.

This was an engineering answer to clickbait: a click followed by immediate abandonment should count differently from sustained viewing. It also shows how proxy choice directs system learning.

OBJECTIVE REVISION

In 2019, YouTube said recommendations had expanded beyond views to include satisfaction signals such as likes, dislikes, and surveys. It also announced reduced recommendations of borderline content. This documents the platform's stated change. It does not independently prove effects across every topic, country, user, or system version.

COMPETING EVIDENCE

Ledwich and Zaitsev's 2020 snapshot of nearly 800 political channels found recommendation traffic moving toward mainstream sources rather than the most extreme categories. Their result challenges a universal radicalization rabbit hole. It does not close questions about personalization, other content domains, vulnerable users, or systems that changed after data collection.

STRUCTURAL READING

Optimizing expected watch time can make continued viewing systemically valuable. Personalization and proprietary data make outside auditing difficult. A plausible mechanism is not the same as proof of a universal effect.

AGENCY READING

People arrive with interests, subscriptions, search choices, and communities. Research snapshots can find recommendation flows that do not match the strongest popular narrative. Agency and system influence can coexist.

CASE LESSON

Ask: which system version, population, content domain, exposure path, and outcome were measured? Never turn one audit into a timeless verdict.

Replacing one proxy means choosing another

In January 2018, Facebook said News Feed would prioritize posts predicted to create conversations and meaningful interactions, while showing less public content.

WHAT THE ANNOUNCEMENT DOCUMENTS

The platform described comments, shares, reactions, and predicted back-and-forth discussion as ranking signals. It made a value judgment explicit: interaction between people would be treated as more desirable than passive consumption. Yet meaningfulness remained a proxy inferred from observable behavior, not direct knowledge of relationship quality.

THE METRIC PROBLEM

A heated argument can create many comments. A supportive exchange can be brief. Changing the objective can improve some outcomes while generating new gaming and side effects.

HOW MUCH IS RANKING, AND HOW MUCH IS CHOICE?

A 2015 Science study of 10.1 million U.S. Facebook users reported that individual choices played a stronger role than algorithmic ranking in limiting exposure to ideologically cross-cutting news. The study was observational, restricted to self-identified liberal and conservative users, and conducted by Facebook employees. Those facts do not nullify it, but they shape interpretation.

A randomized experiment published in the American Economic Review in 2020 found that deactivating Facebook before the 2018 U.S. midterms produced mixed outcomes: more offline activity and improved reported well-being, but less factual news knowledge. Participants also valued access to the platform.

NETWORK

Friends and followed pages shape what can appear before ranking.

RANKING

Predictions decide relative order and visibility within available content.

SELECTION

People choose what to click, ignore, share, mute, and seek elsewhere.

CASE LESSON

Platform power and user agency are not opposites. Network formation, ranking, interface, and deliberate choice can all contribute to the same observed feed.

Influence without puppetry

A

STRUCTURAL PERSPECTIVE

Platforms choose defaults, ranking objectives, data collection, notification systems, and the ease of exit. Small average effects can matter at large scale. The platform can test across a population while each person sees only one interface.

Best contribution: reveals power, asymmetry, and organizational responsibility.

Risk if overstated: treats people as passive and infers effects from incentives alone.

B

AGENCY AND BENEFIT PERSPECTIVE

People actively select, resist, repurpose, and leave platforms. Digital services reduce search costs, sustain relationships, support work, and provide pleasure. Time spent is not automatically time harmed.

Best contribution: preserves purpose, heterogeneity, and real benefits.

Risk if overstated: ignores how defaults and information asymmetry condition available choices.

SYNTHESIS: EVALUATE THE QUALITY OF AGENCY

Legibility

Can people understand the objective, ranking basis, and consequences?

Contestability

Can they alter the route, refuse a cue, or select a non-profiled option?

Distribution

Who gains the benefit, who bears the cost, and which groups disappear in an average?

EVIDENCE LADDER

LABEL	EXAMPLE	PERMITTED CONCLUSION
Evidence	Controlled notification study.	A specific task showed an attention cost under tested conditions.
Interpretation	Ad funding adds engagement pressure.	A directional incentive exists; strategy and guardrails mediate it.
Hypothesis	Stopping cues reduce drift.	A plausible, testable direction, not an established coefficient.
Myth	Algorithms control everyone.	Reject: no evidence establishes total, uniform control.

Make the assumptions inspectable

The simulator is a sensitivity exercise. It divides a selected time sandbox into **chosen use**, aligned with the starting purpose, and **momentum use**, attributed inside the model to return cues, automatic continuation, and weak stopping points.

It is not validated. The coefficients below were invented to make direction and trade-offs visible. They are not measured effects.

DEFAULT SCENARIO

Base momentum share	+0.080
Advertising funding	+0.100
Dwell ranking	+0.100
Six cues: 6 x 0.026	+0.156
Autoplay	+0.070
Continuous feed	+0.110
One stopping cue	-0.045
Intention prompt	-0.000


CALCULATION

Net pressure = $0.100 + 0.100 + 0.156 + 0.070 + 0.110 - 0.045 = 0.491$

Momentum share = $0.080 + (0.491 \times 0.55) = 0.35005$, or about **35.0%**

Momentum minutes = $120 \times 35.0\% = 42 \text{ minutes}$

Chosen-use minutes = $120 - 42 = 78 \text{ minutes}$



Text alternative: a 120-minute sandbox is divided into 78 chosen-use minutes and 42 momentum minutes.

CHALLENGE THE MODEL

Which signs do you accept? Which coefficients would you remove? What benefit or harm is missing? A useful model invites disagreement at the level of assumptions.

Redesign the route, not the person

A strong intervention changes the conditions of choice, preserves legitimate benefits, and names the new trade-off it creates.

INTERVENTION	MECHANISM	POSSIBLE BENEFIT	TRADE-OFF / AUDIT
Batch notifications	Reduces externally timed return cues.	Fewer context switches and more predictable review.	Urgent social or work messages may wait.
Visible endpoint	Restores a natural moment to stop.	Makes completion legible without removing content.	People may still choose to continue; do not shame them.
Session intention	Places a reflective choice before entry.	Can reconnect action to a stated purpose.	Repeated prompts can become noise or obstruction.
Chronology option	Reduces predictive ordering.	Offers a legible alternative and more direct control.	High-volume feeds can become overwhelming.
Durable refusal	Makes "no" persist over time.	Reduces nagging and respects prior choice.	Controls must remain discoverable and reversible.
Counter-metrics	Measures regret, completion, diversity, safety, and satisfaction.	Prevents one engagement proxy from defining value.	More measurement can create privacy and governance risk.

DESIGN TEST

Can the visitor identify why an item appears, select a different ordering logic, stop naturally, and preserve a refusal without repeatedly defending it?

GOVERNANCE TEST

Can independent researchers inspect outcomes? Can affected groups contest errors? Is there a process for changing an objective when side effects appear?

A SIX-PART INTERVENTION BRIEF

1. Name the current objective.
2. Identify the behavior the design makes easier.
3. State the user's legitimate benefit.
4. Propose one reversible change.
5. Name a counter-metric and distributional check.
6. Define what evidence would change your mind.

Reconstruct one session without self-blame

Choose a recent session. Record the environment before evaluating the outcome. Use verbs and observations, not diagnoses.

1 / ARRIVAL

What opened the session: a deliberate task, boredom, a notification, another person, or an automatic check?

2 / CONTRACT

What did you intend to do? Write a verb: reply, learn, coordinate, publish, search, rest, or laugh.

3 / ROUTE

Which feature shaped the next action: ranking, autoplay, a badge, a recommendation, a streak, an endpoint, or a social expectation?

4 / LEDGER

What value did the session create? What did it displace? Name both.

5 / REDESIGN

Change one condition, not your character. Batch a cue, add an endpoint, move the icon, choose chronology, or state an intention.

Trace the chain from funding to feedback

A / FUNDING

How does the service sustain itself? Who pays, and for what?

E / INTERFACE

Where are cues, defaults, continuation, friction, and stopping points?

B / OBJECTIVE

Which measurable outcome appears to define success?

F / BENEFIT AND COST

Who receives value? Who pays in time, privacy, money, or opportunity?

C / SIGNALS

Which actions can the platform observe? Which values remain invisible?

G / FEEDBACK

Which traces return to the model? What behavior might they reward?

D / RANKING

Why does one item appear before another? Is an alternative available?

H / FALSIFIER

What evidence would weaken your interpretation?

FINAL SENTENCE

The current design makes _____ easier because _____. A fairer alternative would preserve _____ while changing _____.

Consent is more than a button

Ethical attention design asks whether a person can understand, refuse, and revise a choice under reasonable conditions. It pays special attention to asymmetric friction: joining is easy while leaving is obstructed; accepting is vivid while refusal is hidden; "not now" causes repeated prompts.

The U.S. Federal Trade Commission's 2022 staff report describes design practices that can obscure, subvert, or impair consumer choice. Luguri and Strahilevitz's experiments found that several dark patterns changed consumer decisions, while also showing that techniques differ in effect. Avoid calling every persuasive choice a dark pattern. Analyze the concrete mechanism and the likely mismatch with user preference.

DO NOT INFER MOTIVE

A harmful outcome can arise without a designer intending harm. Intent may matter legally or organizationally, but an audit should begin with observable structure, effect evidence, and available alternatives.

RECOMMENDER TRANSPARENCY

The EU Digital Services Act requires covered online platforms using recommender systems to explain their main parameters in plain language and the reasons for their relative importance. Where several options exist, users must be able to select and modify a preferred option. Very large platforms and search engines face additional obligations, including a non-profiled option for main recommender parameters.

Legal compliance is a floor, not a complete design standard. A technically available control can remain ineffective if it is buried, confusing, unstable, or disconnected from the moment ranking occurs.

DISTRIBUTIONAL QUESTIONS

- ☐ Does the design rely on literacy, time, or executive function that some people have less opportunity to supply?
- ☐ Are children or vulnerable groups exposed to stronger persuasion?
- ☐ Does an average outcome conceal concentrated harms?
- ☐ Who can appeal an error or change the system?

DIGNIFIED EXIT STANDARD

A preferred route is legible. Alternatives receive equal visual dignity. Refusal persists. Leaving is no harder than joining. A person can return later without being punished.

Use precise claims

SHORTCUT	WHY IT FAILS	BETTER STATEMENT
"Time online is harm."	Duration can contain value, care, work, learning, or pleasure.	Ask what the session enabled and displaced.
"The algorithm made me do it."	It erases purpose, network, context, and choice.	The choice environment changed the ease or visibility of actions.
"Chronology fixes feeds."	It leaves volume, network selection, interface, and advertising.	Chronology is a legible option with trade-offs.
"Habit means addiction."	Automatic repetition is not a clinical diagnosis.	Describe checking, loss of control, or harm only with appropriate evidence.
"One audit proves safety."	Systems, users, content, and time periods differ.	State the population, version, domain, and outcome tested.

THE SIMULATOR DOES NOT

- ☐ Measure your attention or diagnose addiction.
- ☐ Estimate the causal effect of a feature.
- ☐ Prove a platform intends harm.
- ☐ Establish that all advertising or personalization is unethical.
- ☐ Replace product-specific data, qualitative context, or causal research.

TAKEAWAYS TO CARRY FORWARD

1. Separate the layers before assigning cause.
2. Treat engagement as a proxy, not a verdict.
3. Preserve benefits in any proposed redesign.
4. Evaluate legibility, contestability, and distribution.
5. Name what would change your mind.

FINAL PROMPT

Which attention outcome should a platform be proud to make smaller?

References and source notes

1. Simon, Herbert A. "Designing Organizations for an Information-Rich World." In *Computers, Communications, and the Public Interest*, 1971. gvern.net/doc/design/1971-simon.pdf
2. Oulasvirta, Antti, Tye Rattenbury, Lingyi Ma, and Eeva Raita. "Habits make smartphone use more pervasive." *Personal and Ubiquitous Computing* 16 (2012): 105-114. doi.org/10.1007/s00779-011-0412-2
3. Stothart, Cary, Ainsley Mitchum, and Courtney Yehnert. "The Attentional Cost of Receiving a Cell Phone Notification." *JEP: Human Perception and Performance* 41.4 (2015): 893-897. doi.org/10.1037/xhp0000100
4. Covington, Paul, Jay Adams, and Emre Sargin. "Deep Neural Networks for YouTube Recommendations." ACM RecSys, 2016. research.google/pubs/deep-neural-networks-for-youtube-recommendations/
5. YouTube. "Continuing our work to improve recommendations on YouTube." 25 January 2019. blog.youtube/news-and-events/continuing-our-work-to-improve/
6. Ledwich, Mark, and Anna Zaitsev. "Algorithmic extremism: Examining YouTube's rabbit hole of radicalization." *First Monday* 25.3 (2020). doi.org/10.5210/fm.v25i3.10419
7. Mosseri, Adam. "Bringing People Closer Together." Facebook, 11 January 2018. about.fb.com/news/2018/01/news-feed-fyi-bringing-people-closer-together/
8. Bakshy, Eytan, Solomon Messing, and Lada A. Adamic. "Exposure to ideologically diverse news and opinion on Facebook." *Science* 348.6239 (2015): 1130-1132. doi.org/10.1126/science.aaa1160
9. Allcott, Hunt, Luca Braghieri, Sarah Eichmeyer, and Matthew Gentzkow. "The Welfare Effects of Social Media." *American Economic Review* 110.3 (2020): 629-676. doi.org/10.1257/aer.20190658
10. Brady, William J., et al. "Emotion shapes the diffusion of moralized content in social networks." *PNAS* 114.28 (2017): 7313-7318. doi.org/10.1073/pnas.1618923114
11. Luguri, Jamie, and Lior Jacob Strahilevitz. "Shining a Light on Dark Patterns." *Journal of Legal Analysis* 13.1 (2021): 43-109. doi.org/10.1093/jla/laaa006
12. U.S. Federal Trade Commission. *Bringing Dark Patterns to Light*. September 2022. ftc.gov/reports/bringing-dark-patterns-light
13. European Union. Regulation (EU) 2022/2065, Digital Services Act. See Article 27. eur-lex.europa.eu/eli/reg/2022/2065

SOURCE METHOD

Links and bibliographic details were checked against publisher, DOI, regulator, official legal, or first-party platform pages on 28 October 2026. First-party platform sources document stated designs and policy changes; they are not independent validation. The guide avoids unverified prevalence claims and does not reproduce copyrighted figures.

PUBLICATION NOTE

Attention Economy Field Guide, version 1.0. Written and designed by Sahil Tayal for the Attention Economy Observatory. Original diagrams are code-drawn. Designed for A4 printing in color or grayscale. Visit sahiltayal.com to return to the Sahilverse.

Educational use only. This guide is not medical, legal, or clinical advice.

Attention is not merely captured. It is allocated under conditions that can be redesigned.

attentioneconomy.sahilwho.online