

Attention Management in the 21st Century: Strategies for Capturing and Retaining Customers in Markets Saturated with Information

Managing Attention in the 21st Century: Strategies to Attract and Retain Customers in Information-Saturated Markets

Author: Vinicius Rodrigues Ferreira

Graduated in Management from the State University of Goiás

Summary

In **hyperconnected** markets, attention has become the limiting resource for value creation: there is an excess of messages and a deficit of cognitive capacity to process them. This article proposes a framework for managing attention capital that integrates theories from the attention economy and experience marketing with evidence from cognitive psychology and behavioral science, guiding decisions about capturing, retaining, and converting attention throughout the customer relationship cycle. We argue that attention should be treated as a stock (predispositions, brand memory, trust) fed by flows (qualified exposure, meaningful engagement, memorable experiences), governed by standardized metrics of mental availability, attention acquisition cost, and attention-adjusted lifetime value. By articulating value narratives with experience design and proof mechanisms (experiments and incremental testing), we show how companies can reduce informational noise, increase salience and build loyalty based on experience, not friction (SIMON, 1971; PINE; GILMORE, 2011; KAHNEMAN, 2011; DAVENPORT; BECK, 2001; SHARP, 2010).

Keywords: attention economy; experience marketing; mental availability; salience; customer retention.

Abstract

In hyperconnected markets, attention has become the limiting resource for value creation: there is an excess of messages and a deficit of cognitive capacity to process them. This article proposes a framework for managing attention capital that integrates theories from the attention economy and experience marketing with evidence from cognitive psychology and behavioral science, guiding decisions on the capture, retention, and conversion of attention throughout the customer relationship cycle. We argue that attention should be treated as a stock (predispositions, brand memory, trust) fueled by flows (qualified exposure, meaningful engagement, memorable

experiences), governed by standardized metrics of mental availability, attentional acquisition cost, and attention-adjusted lifetime value. By articulating value narratives with experience design and proof mechanisms (experiments and incremental tests), we show how companies can reduce informational noise, increase saliency, and build loyalty based on experience, not friction (SIMON, 1971; PINE; GILMORE, 2011; KAHNEMAN, 2011; DAVENPORT; BECK, 2001; SHARP, 2010).

Keywords: attention economy; experience marketing; mental availability; salience; customer retention.

1. Fundamentals: Attention scarcity as an economic and cognitive constraint

Classical literature already anticipated that, in rich information environments, attention, not information, is the scarce input; therefore, "information wealth creates attention poverty" (SIMON, 1971). In economic terms, attention functions as a limited budget that must be allocated among competing stimuli; in cognitive terms, it is mediated by mental work and fast and slow processing systems, with **heuristics** and biases that shape the perception of risk, reward, and novelty (KAHNEMAN, 2011). This dual nature explains why campaigns based solely on frequency pressure suffer diminishing returns: they run up against attentional saturation and the effort cost required to interpret complex messages, especially on small screens and in multitasking contexts (DAVENPORT; BECK, 2001). In saturated markets, the dominant strategy is not "more noise" but a better signal: reducing cognitive load while increasing salience and contextual relevance.

At the market level, attention is allocated by intermediation platforms that operate attention auctions (CPM/CPC/CPA) and recommendation algorithms that maximize predicted **engagement**; this creates variable attention acquisition costs and pressures brands to prove the marginal utility of each impression. Since not all attention is equal, it is important to distinguish idle attention (passive scrolling) from active attention (interaction with intent), measuring dwell time, scroll depth, meaningful interactions, and *brand lift* recall (SHAPIRO; VARIAN, 1999; DAVENPORT; BECK, 2001). Measurement separates exposure from processing: only the latter generates brand memory and predisposition, known as Mental Availability (SHARP, 2010).

Strategies that confuse *reach* with effective attention tend to overestimate reach and underestimate noise.

The psychological dimension introduces thresholds of perception: moderate novelty, surprise, and social relevance increase *arousal* and **focus attention**, while excessive ambiguity or stimulus overload produce avoidance (BERLYNE, 1960; CIALDINI, 2009). Thus, message design must balance cognitive fluency (ease of processing) with memorable differentiation; narrative plots, **multisensory cues**, and emotional markers facilitate encoding in long-term memory, the basis for recognition and retrieval.

future (KAHNEMAN, 2011; PINE; GILMORE, 2011). In digital interfaces, Hick's **Law** (excessive choices increase decision time) and scanning patterns (F-pattern) guide the *layout*: less friction, more focus on value clues (NIELSEN, 2006; HICK, 1952). Attention, therefore, is applied choice engineering.

At the competitive level, attention generates network externalities: brands with high Mental Availability capture a disproportionate share of clicks and recall, lowering the marginal cost of new impressions—a cumulative dynamic that explains the advantage of "famous for being famous" (LANHAM, 2006; SHARP, 2010). This hysteresis suggests a rule of thumb: investing in distinctive memory (brand assets, visual cues, *taglines*) reduces the cost of future attention and increases conversion rates in contexts of low reflection (KAHNEMAN, 2011; PINE; GILMORE, 2011). At the same time, overexposure can produce message fatigue and semantic saturation; creative portfolio management must alternate variations that preserve distinctive identity without incurring burnout.

At the relationship level, received attention needs to be converted into granted trust. Since attention is often "rented" from intermediaries, the goal is to internalize it into proprietary assets: consented *first-party data*, direct channels (*email*, community, *app*), and interaction rituals that deliver value before demanding value (THALER; SUNSTEIN, 2008; PINE; GILMORE, 2011). Trust reduces the cognitive cost of choices and stabilizes frequency of use, acting as a benign inertia that protects margin in highly substitutable environments. Without trust, tactical increases in attention tend to evaporate in the next stage of the funnel.

Ethics is a constitutive part of attention management. Careless capture— *dark patterns*, intrusive notifications, *infinite scrolling*—can buy clicks at the cost of **reputational** erosion and abandonment; behavioral economics should be used to reduce friction in experiences that respect autonomy and consent, not to exploit them (THALER; SUNSTEIN, 2008; EYAL, 2014). The core is a clear value pact: why is it worth paying attention? What progress does the customer perceive when interacting with the brand? Attention is too scarce to be wasted without a compelling response.

In terms of measurement, we propose distinguishing CAC-A (*customer attention cost*) from financial CAC: the former measures cost per qualified minute or significant event (e.g., demo watched), adjusted for recall and intent; the latter reflects accounting acquisition cost. This separation avoids "optimizations" that make media cheaper but attention more expensive (DAVENPORT; BECK, 2001; SHARP, 2010). LTV-A (*lifetime value* adjusted for attention) weighs survival, frequency, and *share of wallet* by meaningful engagement, isolating real growth from metric froth.

Finally, the bridge to experience marketing is straightforward: memorable experiences are attentional cost compressors—they concentrate relevant information into episodes that the brain retains, shortening future decision cycles (PINE; GILMORE, 2011). Instead of insisting on scattered **microinteractions** that add up to little, the firm can design experience peaks.

(*onboarding*, discovery, support) that reprogram customer expectations and **heuristics**. Attention, therefore, stops being captured and becomes earned.

2. A management framework for "attention capital": stocks, flows, metrics, and decisions

We propose treating attention as organizational capital with stock and flows. The stock includes Mental Availability (brand memory and *distinctive assets*), trust, and consented access (owned channels); the flows are Qualified Exposure ÿ Meaningful Processing ÿ Memorable Experience ÿ Memory Reinforcement. Each transition is lost when there is noise, friction, or lack of proof of value (SHARP, 2010; PINE; GILMORE, 2011). Management's role is to close the loss valves, synchronizing message, medium, and moment with customer tasks (*jobs to be done*), so that invested attention generates felt progress and, consequently, loyalty.

At the tactical level, the framework uses three metrics: (i) Mental Availability (associative proximity and spontaneous/assisted recall across usage categories), (ii) Attentional Acquisition Cost (cost per qualified minute/significant event, corrected for recall), and (iii) Attention-Adjusted Lifetime Value (LTV-A). These indicators link salience to economics and connect to two constraint envelopes: service (time/resolution/effort) and ethics (consent/clarity/control) (SHARP, 2010; DAVENPORT; BECK, 2001; THALER; SUNSTEIN, 2008). The operating principle is: you don't buy attention that can't be processed, and you don't ask for attention without delivering progress.

To capture attention, the strategy combines qualified reach with cognitive fluency and rapid proof. The "long + short" combination matters: brand assets build mental access over time, while activations with strong value cues (price, social proof, immediate utility) convert attention into action—a well-documented complementarity in investment *splitting* (BINET; FIELD, 2013). In terms of design, distinctive signs reduce identification burden; "one idea at a time" messages avoid overlap; persuasion principles operate as legitimate shortcuts to focus (CIALDINI, 2009; KAHNEMAN, 2011).

To retain attention, it's necessary to transform contact into continuity. Three mechanisms stand out: rituals (predictable value cadences that the customer incorporates into their routine), visible progress (milestones that signal user/customer evolution), and significant variable rewards (variation that maintains epistemological curiosity without falling into exploitative addiction) (EYAL, 2014; PINE; GILMORE, 2011). Well-designed *onboarding* reduces initial cognitive costs and prevents abandonment; proactive support and *feedback* that close the *loop* preserve trust and future attention.

To convert attention into results, the framework recommends minimal friction at decision points and proof of value at moments of greatest perceived uncertainty. Choice architectures: optimal number of options, order, and labels reduce decision burden; demonstrations and testing reduce

subjective risk, enabling advancement (HICK, 1952; KAHNEMAN, 2011). Conversion ceases to be a push and becomes a natural consequence of perceived progress.

Measurement requires robust methods. Controlled experiments in media and experience, complemented by *bandit testing* when there is significant heterogeneity, reduce false positives and optimize exploration and exploitation; incrementality (*uplift*) replaces correlation metrics; attribution windows must respect category consideration cycles, or risk **overpaying** late stimuli (SHAPIRO; VARIAN, 1999; KAHNEMAN, 2011). In *branding*, memory traces (associations, *brand lift*) serve as *leading indicators* for future CAC-A (SHARP, 2010).

At the organizational level, the framework implies an attention control tower: dashboards that combine actionable exposure, meaningful engagement, recall, and customer progress, with thresholds that trigger actions (creative adjustments, channel reallocation, frequency caps). *Playbooks* define responses by anomaly: *ad fatigue*, audience saturation, decline in fluency or social proof. *Governance* includes **pre-authorized** rights to change budgets and data standards to avoid "two truths" (DAVENPORT; BECK, 2001; PINE; GILMORE, 2011).

The framework also highlights ethical *trade-offs*: the company commits to not exceeding intrusive frequency caps, offering granular notification control, explaining value exchanges in a comprehensible manner, and measuring reputational damage as a negative attention cost. Trust is a capital that is more difficult to rebuild than reach; therefore, *nudging* practices should foster autonomy (THALER; SUNSTEIN, 2008). Forced attention **becomes** evasion; earned attention tends to generate interest in the form of *advocacy*.

Finally, the framework connects to finance: CAC-A and LTV-A enter the *business case* alongside traditional LTV, and avoided loss through noise and friction reduction justifies CAPEX/OPEX in experience assets (content, service, design). As in logistics and resilience, the rule is to pay for predictability: assets that **compress** attentional costs and increase Mental Availability reduce revenue risk in turbulent markets (PINE; GILMORE, 2011; SHARP, 2010; DAVENPORT; BECK, 2001). Attention, treated as manageable capital, is now allocated, measured, and defended with the same discipline as other critical assets.

3. Experience Design and Choice Architecture: Reducing Cognitive Load Without Losing Differentiation

Cognitive load is one of the direct determinants of attentional costs: the greater the effort required to understand, decide, and execute an action, the lower the likelihood of deep processing and future recall. Therefore, interface and communication designs that minimize steps, ambiguities, and friction tend to increase completion and recall rates in stimulus-saturated environments. In practical terms, simplifying visual hierarchies, reducing redundant text, and

aligning value cues with user focus reduces the "structural noise" that sabotages campaigns with good reach but low memory effect (NIELSEN, 2006; DAVENPORT; BECK, 2001).

Choice architecture offers proven principles for organizing decisions without manipulating user autonomy. Hick's law indicates that decision time increases with the number and complexity of alternatives; therefore, fewer, better-labeled options tend to compress effort and increase response rates (Hick, 1952). On product pages or registration flows, curating three to five "sufficient" options outperforms extensive catalogs with opaque filters, as it prevents attention from being dispersed into fruitless comparisons (Kahneman, 2011). The goal is not to "push" but to provide progressive **clarity** .

Cognitive fluency—ease of processing—increases when messages and visuals follow predictable patterns, with adequate contrast, legible typography, and **microtext** that anticipates the next action. Interfaces aligned with the scanning F-pattern and the user's hierarchy of questions ("what is it?", "who is it for?", "how do I start?") reduce the need for System 2 (slow analysis), freeing up the capacity to evaluate real differences (NIELSEN, 2006; KAHNEMAN, 2011). Clarity doesn't trivialize the brand; it amplifies what makes it distinctive.

Memorable differentiation depends on salience and emotional markers. Distinctive elements (colors, shapes, *taglines*, audio logos) function as **mnemonic** hooks that memory recognizes under low attention, reducing the cost of future activations. This "mental availability" is reinforced by episodes of experience that generate peaks: moments of useful surprise, elegant resolution of a conflict, or remarkable service that the brain tends to encode and recount (PINE; GILMORE, 2011; SHARP, 2010). The combination of fluency and peak creates accessible and affective memories.

Social proof reduces uncertainty when placed at the right time. Testimonials, *ratings* , and use cases act as shortcuts to trust, as long as they are specific, verifiable, and contextualized to the decision-making segment (CIALDINI, 2009). *Above the fold*, excessive signals compete for attention and appear forced; as the action trigger (add to cart, request a demo) approaches, the density of proof should increase, replacing promotional claims with evidence of risk reduction (SHAPIRO; VARIAN, 1999).

Onboarding is a privileged laboratory for attention management. Flows that postpone non-critical choices, offer intelligent defaults (**pre-populated** values based on context) , and teach by doing, embrace the principle of quick wins: the user sees measurable progress in the first few minutes, which anchors expectations and reduces abandonment (EYAL, 2014; KAHNEMAN, 2011). Each field removed, each step combined, saves attention that can be invested in value differentiators.

Variety should be used to enhance memory, not hinder it. Changes in *layout* and tone without preserving distinctive assets dilute recognition; on the other hand, controlled variations

over a coherent visual system avoid creative fatigue without "breaking" user learning (SHARP, 2010). The criterion is recognizable consistency, not monotonous uniformity. In long-term campaigns, a modular system of pieces allows for testing nuances without the cost of **relearning**.

Design ethics imposes clear boundaries: no dark patterns, no intentional ambiguity, no mistaken attention grabs. Trust is a driver of long-term cognitive cost reduction; therefore, design choices should favor autonomy and granular control (*opt-ins*, notification preferences), reinforcing the **value pact between** brand and customer (THALER; SUNSTEIN, 2008).

Opportunistic capture yields vanity metrics and erodes future attention stock.

Finally, incremental testing and controlled experiments should guide design decisions. Metrics of incrementality (*uplift*), time to action, and **post-exposure** recall are superior to raw clicks. When there is high heterogeneity, *bandits* balance exploration and exploitation with sample parsimony (SHAPIRO; VARIAN, 1999). Experimental discipline avoids false victories and ties aesthetics to the causal effect on attention and outcome.

4. The long-term and short-term mix: mental availability, social proof, and evidence-driven conversion

Attention management requires a temporal portfolio: building mental availability in the long term and converting efficiently in the short term. Evidence shows that brands that combine persistent distinctive assets with response activations have greater sales elasticity, as the memory reserve reduces the marginal cost of each impression in the present (BINET; FIELD, 2013; SHARP, 2010). The recurring error is to confuse momentary reach with lasting mental access and underinvest in the layer that sustains future attention.

In the long term, the goal is to expand and reinforce the associative nodes that link the brand to relevant tasks in the audience's mind: "For what situations do I remember it effortlessly?" Emotional campaigns, consistent narratives, and distinctive cues increase the likelihood of spontaneous recall under low attention spans—the dominant context in fragmented media (KAHNEMAN, 2011; PINE; GILMORE, 2011). The result is a **steeper** response curve for future activations.

In the short term, social proof, authentic scarcity, and clear offers reduce decision costs and accelerate the transition to the next stage in the funnel (CIALDINI, 2009). Even so, the design must respect trust thresholds: "deadlines" and "last units" must be verifiable to avoid eroding the stock of attention with empty promises. Action is the consequence of a clear value hypothesis, not rhetorical pressure.

Budget allocation between tiers is not fixed; it depends on life cycle, category, competitive momentum, and seasonality. In high-consideration categories, the critical mass of Mental Availability weighs more heavily; in routine purchases, **well-timed** tactical activations capture latent intent. The criterion is incrementality: how much each additional dollar contributes to *brand* or *performance*.

increases qualified attention and *recall*, and how much it converts into lifetime value (BINET; FIELD, 2013; SHAPIRO; VARIAN, 1999).

Creative sequencing matters. Messages that open with a distinctive cue (color, shape, *sonic cue*), follow with a unique proof of value, and close with a clear next action tend to **perform well**.

better under partial attention. Under saturation, "one idea at a time" pieces outperform "lists of promises" that require mental comparison and increase workload (NIELSEN, 2006; KAHNEMAN, 2011). Attention behaves like a narrow-width channel: accomplish the essential first.

Channel planning should consider attentional context: moments of idle attention (queue, transportation) require short, highly distinctive signal units; moments of active attention (search, evaluation) accept long units with dense evidence (comparisons, demonstrations). In both, visual and verbal consistency reduces identification costs and helps the brain fit the message into the right drawer (SHARP, 2010; DAVENPORT; BECK, 2001). "Medium is message" because medium defines available attention.

Measurement must separate correlation from causation. *Brand lift* (recall, association, preference) serves as a *leading indicator* for future CAC-A; geographic and cohort experiments prevent attribution models from giving excessive credit to the last touch (SHAPIRO; VARIAN, 1999). In *performance*, incrementality with realistic attribution windows prevents "metric cannibalization." What is not rigorously measured becomes a comfortable narrative and wastes attention.
no return.

Ethical governance anchors the temporal mix. The long term does not allow for capture **devices** that increase *session time* at the expense of well-being; the short term does not authorize artificial scarcity. Transparency, consent, and controls maintain a positive rate of return on attention over time, avoiding "stimulus inflation" that reduces the purchasing power of the message (THALER; SUNSTEIN, 2008; SIMON, 1971). Attention capital is capitalized upon when respected.

Finally, the S&OP cycle of attention: quarterly Mental Availability, CAC-A, and LTV-A targets guide the portfolio; weekly rituals adjust creatives, frequencies, and channels based on incremental analysis; monthly AARs capture learnings and avoid "seasonal reinvention."
(DAVENPORT; BECK, 2001). Thus, the long and short-term mix ceases to be a debate of opinion and becomes a repeatable process, where attention is planned, measured and defended as an asset.

5. Attention capital metrics (CAC-A, LTV-A, mental availability) and their integration with finance

Treating attention as an asset requires metrics that connect cognitive economics to financial results. We propose three anchors: Mental Availability (the probability of a brand coming to mind in purchasing contexts, measured by spontaneous/assisted recall and associative proximity), CAC-A (*Customer Attention Cost*, cost per qualified minute or significant event weighted by recall and intention), and LTV-A (*Lifetime Value* adjusted for effective attention at each stage of the relationship). These measures disambiguate "exposure" from processing and allow us to compare channels and creatives based on the useful attention they deliver, not just clicks or raw reach (SHARP, 2010; DAVENPORT; BECK, 2001). In saturated markets, the attention denominator reorders priorities: campaigns that appear cheap per impression end up expensive when normalized by minutes processed (SIMON, 1971).

Mental Availability operates like low-friction inventory: the higher it is, the lower the marginal cost of triggering behavior in contexts of little reflection. Periodic *brand lift* surveys and associative maps estimate the breadth of mental access and the quality of links (tasks, occasions, *category entry points*), serving as *leading indicators* of tomorrow's CAC-A (SHARP, 2010). By including these indicators in marketing S&OP, the company anticipates future "attention scarcity" and adjusts the long- and short-term mix before the funnel dries up (BINET; FIELD, 2013; KAHNEMAN, 2011). Mental availability is not vanity: it is brand *capex* that reduces acquisition *opex*.

CAC-A complements financial CAC by distinguishing media accounting costs from actual cognitive costs. It is calculated as spending per qualified minute (or per significant event, such as a demo watched), weighted by recall and intent, measured in *post-exposure surveys*.

and A/B testing. Two creatives with the same CPA can differ greatly in CAC-A if one generates idle attention (*passive scrolling*) and the other active attention (processing with proof of value).

This metric reduces the illusion of efficiency of last-touch attribution and prioritizes value cues that the brain processes with low load (DAVENPORT; BECK, 2001; NIELSEN, 2006). In multichannel portfolios, CAC-A *limits* prevent *overinvestment* in cheap inventory that does not form memory.

LTV-A integrates survival, frequency, *ticket*, and *share of wallet*, weighted by meaningful engagement over time. Technically, traditional LTV is multiplied by an effective attention coefficient (fraction of interactions with deep processing), which can be estimated by time on task, scroll depth, value events, and recall per cohort (SHAPIRO; VARIAN, 1999). The goal is to reward growth with quality attention, discouraging metric inflation through intrusive notifications and repetitive ads that erode trust and the rate of attentional return in the long term (THALER; SUNSTEIN, 2008; EYAL, 2014). LTV-A makes the cost of attention attrition explicit.

The **finance-marketing** translation occurs when Mental Availability, CAC-A, and LTV-A enter the **economics** of campaigns and projects. In *business cases*, "impressions" are replaced by minutes processed and "CPA" by CAC-A; in experience projects (*onboarding*, support), avoided loss is measured in CAC-A and gain in LTV-A by reduced cognitive load (pages, fields, steps). Budgets become capital allocated to attention producers (distinctive assets, memorable content) and cost compressors (choice architecture, clarity automation), with *paybacks* estimated by causal *uplift* (BINET; FIELD, 2013; PINE; GILMORE, 2011). Governance changes: investment is requested because attention will become cheaper tomorrow.

In the executive report, a minimum viable dashboard combines: (i) Mental Availability by segment and *entry points*; (ii) CAC-A by channel and creative; (iii) LTV-A by cohort; (iv) attrition KRIs (creative fatigue, *unsubscribes*, blocks); (v) ethics (frequency, consent, control). As in operational resilience, an envelope is defined: do not exceed CAC-A **caps**, do not degrade LTV-A for tactical gains, and do not fall below Mental Availability targets (SHEFFI, 2015; SHARP, 2010). *Playbooks* provide for rebalancing when thresholds are crossed, with **pre-authorized** rights to move funds.

Causal modeling is a requirement. Controlled experiments estimate incrementality of Mental Availability and convert recall points into expected CAC-A variation; marketing mix models with maximum processable attention constraints avoid assuming linear frequency effects; *bandits* handle the exploration-exploitation *trade-off* in creatives, with version records for auditing (SHAPIRO; VARIAN, 1999; KAHNEMAN, 2011). Without this rigor, attention capital metrics become *post-hoc* narratives and lose predictive power.

From a risk perspective, the company should monitor attention VaR: what is the worst plausible drop in Mental Availability and LTV-A under scenarios of media saturation and algorithmic change? As operational insurance, **reserves** of distinctive assets, proprietary channels, and communities are maintained to withstand drops in paid reach, preserving attention flows even when platforms change conditions (LANHAM, 2006; DAVENPORT; BECK, 2001). Here, ethics and transparency are not ornaments: they are *hedgies* against trust volatility.

Finally, variable compensation should include composite targets for Mental Availability, CAC-A, and LTV-A, with strict ethical constraints (consent, *opt-out*, frequency). Bonusing only CPA leads to opportunistic capture; only session time leads to empty retention. The weighted basket aligns creative, media, and product with what matters: earned attention that converts into value and compounds interest over time (THALER; SUNSTEIN, 2008; PINE; GILMORE, 2011; SHARP, 2010).



6. Control Towers and Continuous Experimentation: Governance, Ethics, and Organizational Learning

An attention control tower integrates media, product, and relationship data to transform signals into low-latency portfolio adjustments. At the base, data patterns and dictionaries harmonize definitions of exposure, processing, significant events, and recall; above, dashboards display Mental Availability, CAC-A, LTV-A, and KRIs (fatigue, *unsubscribes*, blocks) by segment and channel; at the top, *playbooks* map **triggers to pre-authorized** actions (creative swaps, frequency *caps*, media rebalancing, *onboarding refactoring*).

Without this framework, teams fall into **micro-optimizations** that consume attention and do not build memory (DAVENPORT; BECK, 2001; NIELSEN, 2006).

Experimental governance sets the standard of proof. Controlled trials with *uplift* measure causal effect on useful attention and value; geographic and cohort experiments mitigate contamination; *bandits* manage exploration when there is high heterogeneity; attribution windows respect consideration cycles, avoiding *overcredit* at the last touch (SHAPIRO; VARIAN, 1999). Each experiment records hypothesis, design, data, analysis, and decision (AAR cycle), feeding institutional memory and avoiding “seasonal reinventions” (KAHNEMAN, 2011). The company learns with attention parsimony.

The ethical layer is inseparable. The tower codifies boundaries: prohibition of dark patterns, transparency in value exchanges, granular notification controls, and respect for *opt-outs*. Negative attention cost metrics (irritation, complaints, *post-campaign churn*) act as **counterweights** to short-term enthusiasm (THALER; SUNSTEIN, 2008). The goal is to earn attention: reduce load, increase clarity, and deliver progress. Ethics doesn't hinder *performance* —it protects LTV-A.

At the organizational level, the tower requires clear decision-making rights: who can reschedule frequencies, pause creatives, shift allocation between *brand* and *performance*, and adjust *onboarding*? Without **pre-authorization**, response time exceeds fatigue time, and the company pays latency tax (DAVENPORT; BECK, 2001). Governance includes biweekly *portfolio review committees*, weekly *war room* routines, and defined roles (data, creative, product, CRM), with a trigger-based decision SLA.

Data infrastructure and privacy are critical factors. The tower operates on *first-party data*. consent forms and product telemetry with data minimization and anonymization where appropriate; standardized APIs reduce semantic ambiguity; cataloging and lineage enable CAC-A and LTV-A auditing (SHAPIRO; VARIAN, 1999). In a world of tracking restrictions, *privacy-centric* measurement models (experiments and MMM) become pillars of proof, replacing granular observation with robust inference (DAVENPORT; BECK, 2001).

Creative engineering is moving toward modular systems. Sets of distinctive assets (color, shape, *sonic cue*), value cues (one-idea-at-a-time messages), and evidence (testimonials,

demonstrations) are combined via *templates* that facilitate controlled testing without losing visual consistency (SHARP, 2010; NIELSEN, 2006). The result is variation with coherence: user learning time falls, Mental Availability rises and CAC-A declines in successive waves.

The product cycle **integrates** with the tower. *Onboarding*, support, and community are major producers (or destroyers) of useful attention; therefore, these teams' OKRs include cognitive load reduction, time to first value, and proactive resolution, always with a causal *uplift* over LTV-A. In subscription companies, visible progress and rituals (curated content, *check-ins*, usage advice) maintain deserved attention without inflating notifications (EYAL, 2014; PINE; GILMORE, 2011). Product and marketing stop competing for attention and start **co-managing it**.

On the risk front, the tower maintains contingency plans for algorithmic shocks and inventory saturation: *playbooks* for frequency reduction, migration to owned channels, community reinforcement, and high-value content; *stress tests* assess attention VaR and recovery time (SIMON, 1971; LANHAM, 2006). The company operates with real attention options: reserves of distinctive assets, editorial lines, and partnerships that activate when signals cross thresholds.

Finally, continuing education and AARs institutionalize learning. Tracks in attention psychology, behavioral economics, experience design, and causal methods create fluency; quarterly AARs review cadences, thresholds, and standards; best practice guides prevent regressions to shortcut tactics (KAHNEMAN, 2011; THALER; SUNSTEIN, 2008). The compounding effect appears when captured attention becomes earned attention and then preferred attention.

7. Channel Portfolio and Attention Auctions: Dynamic Allocation, Incrementality, and Platform Risk

The channel portfolio needs to be thought of as an efficient frontier of useful attention, not as a list of media purchased out of habit. In practice, paid channels (programmatic auctions, search, social media), owned channels (*consented email*, apps, communities, events), and earned channels (public relations, word of mouth, influencers) coexist, each operating under distinct regimes of marginal cost of attention and control. The recurring mistake is to optimize each silo for CPA and ignore that idle attention and active attention have radically different values for memory and conversion—which is why we normalize by CAC-A (cost per qualified minute/significant event) and causal incrementality before reallocating budget (DAVENPORT; BECK, 2001; SHAPIRO; VARIAN, 1999). The goal is to build processable reach in the short term while accumulating mental availability to lower costs tomorrow (SHARP, 2010; BINET; FIELD, 2013).

In digital attention auctions, the price paid is a function of bids, quality/relevance and competitive context; in search *and* display, GSP/VCG type mechanisms and



Quality *scores* modulate the effective cost per impression and per click (EDELDMAN; OSTROVSKY; SCHWARZ, 2007; VARIAN, 2007). Because the supply of processable attention is limited, diminishing marginal returns emerge early, and frequency saturation curves require limits to avoid creative fatigue and irritation—a “negative cost of attention” that erodes LTV-A (NIELSEN, 2006; THALER; SUNSTEIN, 2008). The mature manager treats the auction as an incomplete market: he buys what he needs but protects dependence with his own channels.

Attentional diversification is a risk management principle. Overreliance on an intermediary exposes users to algorithmic, political, price, and inventory availability shocks; the answer is to maintain real options in owned channels (consent lists, communities, apps) and alternative routes (public relations, partnerships, evergreen content), with *playbooks* that reallocate budgets when precursor signals cross thresholds (declining organic reach, CPMs above the threshold, worsening match rates) (LANHAM, 2006; DAVENPORT; BECK, 2001). In practice, *first-party data* is an attention asset that reduces volatility and marginal cost in the cycle.

Dynamic budget allocation requires causal proof, otherwise we'll have to chase convenient correlations. Three instruments complement each other: controlled experiments (A/B by creative/channel with an *uplift* in useful attention and conversion), geographic/cohort experiments (for “*walled gardens*”), and marketing mix models with an explicit restriction on maximum processable attention at the top (SHAPIRO; VARIAN, 1999). *Last-click* overestimates “last touches” precisely because they capture the intention already formed by the memory stock built in the brand layer (BINET; FIELD, 2013; SHARP, 2010). Allocating without incrementality is buying.

echo.

Paid channels differ by attentional mode: search *captures* active attention (an ongoing task), feeds deliver idle attention (low-reflection discovery), and video-on-demand alternates between the two depending on the context. The message architecture needs to adapt: in idle attention, distinctive signal + one idea at a time; in active attention, dense evidence + risk reduction (NIELSEN, 2006; KAHNEMAN, 2011; CIALDINI, 2009). The same piece in different contexts produces different attention costs—and, therefore, different CAC-A. Time/device granularity refines this adjustment.

Platform risk calls for scenarios: CPM spike due to inventory shrinkage, advertising policy changes, privacy shocks that reduce matching and attribution, or degradation of organic reach. For each scenario, we define milestones (CPM variation; organic reach; matching rate; blocks; *unsubscribes*), numerical triggers, and **pre-authorized responses**: reduce frequency, shift budget to a broad-reaching *brand*, strengthen owned channels, increase valuable content and public relations (LANHAM, 2006; DAVENPORT; BECK, 2001).

This turns the portfolio into a defense mechanism, not a hostage.

Owned channels aren't “free”: they have maintenance costs (content, community, moderation) and subscriber attention limits. Even so, they operate with lower volatility and

greater ethical control (consent, cadence, preferences), and compound interest when memorable experiences fuel recommendation and retention (PINE; GILMORE, 2011; THALER; SUNSTEIN, 2008). The goal is balance: using paid channels to accelerate and in-house channels to retain — both measured by CAC-A/LTV-A.

Portfolio ethics are a constraint, not an addendum. Uncontrolled frequency, dark patterns, and artificial "scarcity" buy clicks and sell attrition: CAC-A drops today, and LTV-A melts tomorrow (THALER; SUNSTEIN, 2008). We inserted friction indicators (reported irritation, *post-campaign churn*, blockages) as counterweights and hard limits in the dashboard—don't exceed them. Deserved attention is the only thing that remains cheap over time (SIMON, 1971; SHARP, 2010). A good portfolio respects autonomy and context.

Finally, we integrate finance: board reports show CAC-A \bar{y} LTV-A boundaries by channel, with confidence intervals and platform risk scenarios. Budget requests are justified by the avoided loss (a drop in future CAC-A by investing in current mental availability) and by the reduced volatility through owned channels, in line with the "attentional protection" thesis (BINET; FIELD, 2013; DAVENPORT; BECK, 2001). The portfolio, thus, ceases to be a media spreadsheet and becomes attention engineering.

8. Implementation roadmap: indicators, cadence and independent verification

The roadmap begins with a diagnosis of attentional materiality. Mapping audiences, tasks (*jobs*), contexts of use, distinctive assets, current channels, and experience frictions produces a baseline of Mental Availability, CAC-A, and LTV-A by segment. In parallel, build a metric dictionary (definitions, formulas, sources, data levels) and ethical policies (maximum frequency, consent, preferences)—without these, there is no comparability or social license to operate (DAVENPORT; BECK, 2001; THALER; SUNSTEIN, 2008). The result is a map of levers and constraints.

Phase 1 (0–90 days) — Data foundation and minimum viable panel. Integrate media, product, and CRM sources; establish memory collections (*brand lift*, associations) and attention telemetry (time on task, significant events); publish a dashboard with Mental Availability, CAC-A, LTV-A, and KRIs (fatigue, *unsubscribes*, blocks). In parallel, standardize experiment processes (hypothesis, sample, power assessment) and incremental attribution (SHAPIRO; VARIAN, 1999). Without a panel and method, the rest becomes opinion.

Phase 2 (90–180 days) — Governance and playbooks. Establish the attention control tower with rituals (tactical journal; weekly *war room* ; biweekly portfolio committee), **pre-authorized** entitlements (frequency adjustment, media rebalancing, creative pause) and numerical triggers (CAC-A ranges; mental availability drops; ethical KRIs) that trigger actions (DAVENPORT; BECK,

2001). Codify platform risk scenarios with milestones and responses (migration to own channels, public relations, evergreen content), reducing decision latency (LANHAM, 2006).

Phase 3 (6–12 months) — Modular creative systems and experience. Build a library of distinctive assets (color, shape, *sonic cue*), value cues (“one idea at a time”), evidence (testimonials, demos), and *templates* that allow controlled testing without breaking consistency (SHARP, 2010; NIELSEN, 2006). In product, refactor *onboarding* to reduce cognitive load, accelerate time to first value, and incorporate rituals and visible progress that sustain deserved attention (PINE; GILMORE, 2011; EYAL, 2014). Measure direct causal effect (*uplift*) on LTV-A.

Phase 4 (12–24 months) — Prudent automation and mix models. Implement dynamic orchestration with *bandits* within ethical limits (frequency, consent), connected to marketing mix models that respect maximum processable attention and incorporate brand knock-on effects (SHAPIRO; VARIAN, 1999; BINET; FIELD, 2013). Integrate finance: brand and experience *business cases* begin to use minutes processed, CAC-A, LTV-A, and avoided loss as common currencies, improving the quality of investment decisions (DAVENPORT; BECK, 2001).

Cadence brings the system to life. **Weekly:** tactical adjustments (creative, frequency, channel) guided by incrementality and KRIs. **Monthly:** review of the attention S&OP (mental availability goals, *brand/performance rebalancing*, experiment AARs). **Quarterly:** strategic portfolio review, platform risk scenarios, ethical goals, and experience CAPEX (PINE; GILMORE, 2011; THALER; SUNSTEIN, 2008). Cadence without method becomes anxiety; method without cadence becomes bureaucracy.

Independent verification preserves credibility. Sample audits of recall, attention telemetry, incrementality, and LTV-A calculation prevent metrics from being “greenwashed.” *In* privacy, data minimization, anonymization, and granular consent are minimum standards; when tracking diminishes, experiments and MMM become the core of evidence (SHAPIRO; VARIAN, 1999; DAVENPORT; BECK, 2001).

Methodological transparency reduces internal disputes and increases external legitimacy.

The risk framework includes a VaR of attention: what is the worst plausible drop in Mental Availability and LTV-A under algorithm shocks, inventory saturation, or reputational crisis? Contingency plans include frequency reduction, creative hiatuses, budget shifts, reinforcement of owned channels, and value programs (community, education) to sustain flow during the shock (LANHAM, 2006; SIMON, 1971). Metrics of recovery time and avoided loss complete the learning cycle.

Finally, compensation and culture. Variable targets combine Mental Availability, CAC-A, and LTV-A within ethical boundaries and KRIs; training in attention psychology, behavioral economics, experience design, and causal methods creates fluency; quarterly AARs

consolidate institutional memory (KAHNEMAN, 2011; THALER; SUNSTEIN, 2008). The script works when the captured attention becomes deserved attention — and, over time, preferred.

Conclusion

The central thesis of this article is that attention is not a byproduct of reach, but a scarce asset that requires its own governance, metrics, and rights of action. In saturated markets, insisting on more impact without reducing cognitive load results in diminishing returns and reputational erosion. Therefore, managers must shift from a logic of buying impressions to allocating processable attention, measured by qualified minutes, significant events, and recall, and translated into lifetime value (SIMON, 1971; DAVENPORT; BECK, 2001; SHARP, 2010).

We show that the attention economy finds its operational counterbalance in experiential marketing: reducing friction, clarifying choices, and producing memorable episodes transforms captured attention into deserved and, over time, preferred attention. The design of choice architectures (fewer options, clear labels, proof in the moment of doubt) and fluent interfaces (one idea at a time, distinctive cues) "compresses" attentional cost without erasing differentiation (HICK, 1952; NIELSEN, 2006; PINE; GILMORE, 2011).

We argue that attention should be managed as capital with stock (mental availability, trust, and authorized channels) and flows (qualified exposure ÿ meaningful processing ÿ memorable experience ÿ memory reinforcement). This framework avoids confusing reach with mental access and prevents tactical decisions from eroding the invisible stock that cheapens tomorrow's acquisition (SHARP, 2010; BINET; FIELD, 2013).

To link attention to finance, we proposed metrics such as Mental Availability, CAC-A (cost of attention), and LTV-A (lifetime value adjusted for effective attention). These measures reposition the conversation from CPA and clicks to minutes processed and causal incrementality, enabling *business cases* in which investments in distinctive assets and experience are made today because this reduces the cost of attention and risks tomorrow (DAVENPORT; BECK, 2001; SHAPIRO; VARIAN, 1999; BINET; FIELD, 2013).

We advocate for an attention control tower with dashboards, **pre-authorized** triggers, and *playbooks* that rebalance creatives, frequencies, channels, and product flows according to KRIs (fatigue, *unsubscribes*, blocks) and value metrics (CAC-A, LTV-A, availability). Without this apparatus, teams are stuck with correlation **micro-optimizations** that consume attention and do not build memory (DAVENPORT; BECK, 2001; KAHNEMAN, 2011).

The channel portfolio should be treated as an efficient frontier of useful attention, with diversification to mitigate platform risk and dynamic allocation guided by experiments and mixing models that respect maximum processable attention. Paid channels accelerate;



own channels stabilize and capitalize; conquered channels amplify with legitimacy (EDELMAN; OSTROVSKY; SCHWARZ, 2007; VARIAN, 2007; LANHAM, 2006).

Ethics is not an addendum, but a rigid constraint on the system: without transparency, consent, and controls, any short-term gain in *session time* incurs a negative attention cost and erodes LTV-A. The responsible use of **heuristics** and social proof maximizes informational utility without manipulating autonomy—a condition for attention to generate reputational interest (THALER; SUNSTEIN, 2008; CIALDINI, 2009; EYAL, 2014).

At the methodological level, we replaced last-touch attribution with causal proof: controlled experiments, geographic/cohort tests, and MMM with attentional saturation constraints. Incrementality discipline avoids illusory victories and makes *brand* and *performance comparable*, linking memory stock to short-term responses (SHAPIRO; VARIAN, 1999; BINET; FIELD, 2013).

In terms of risk, we introduced attention VaR to anticipate algorithm, inventory, and trust shocks. Contingency plans—reducing frequency, shifting budgets, strengthening proprietary channels and value programs—preserve flows during crises and shorten recovery times, transforming attention management into a resilience mechanism (LANHAM, 2006; SIMON, 1971).

Finally, we proposed a four-phase roadmap: (i) data foundation and minimal dashboard; (ii) governance and *playbooks*; (iii) modular creative system and experience; (iv) prudent automation with blended models. The cadence (weekly, monthly, quarterly), AARs, continuous training, and independent verification consolidate learning and prevent regressions to shortcut tactics (PINE; GILMORE, 2011; KAHNEMAN, 2011).

In short, managing attention in the 21st century is choice engineering + data science + applied ethics: measuring rigorously, reducing load, increasing salience, proving value, and respecting autonomy. Organizations that internalize this discipline convert attention into a financial asset, reduce acquisition volatility and build lasting relationships in markets where the greatest scarcity is not media, but mental time (SIMON, 1971; SHARP, 2010; DAVENPORT; BECK, 2001).

References

BERLYNE, DE **Conflict, Arousal, and Curiosity**. New York: McGraw-Hill, 1960.

BINET, L.; FIELD, P. **The Long and the Short of It: Balancing Short and Long-Term Marketing Strategies**. London: IPA, 2013.

CIALDINI, RB **Influence: Science and Practice**. 5. ed. Boston: Pearson, 2009.



DAVENPORT, T.H.; BECK, JC **The Attention Economy: Understanding the New Currency of Business**. Boston: Harvard Business School Press, 2001.

EDELMAN, B.; OSTROVSKY, M.; SCHWARZ, M. Internet Advertising and the Generalized Second-Price Auction. **American Economic Review**, vol. 97, no. 1, p. 242–259, 2007.

EYAL, N. **Hooked: How to Build Habit-Forming Products**. New York: Portfolio, 2014.

HICK, WE On the Rate of Gain of Information. **Quarterly Journal of Experimental Psychology**, vol. 4, p. 11–26, 1952.

KAHNEMAN, D. **Thinking, Fast and Slow**. New York: Farrar, Straus and Giroux, 2011.

LANHAM, RA **The Economics of Attention: Style and Substance in the Age of Information**. Chicago: University of Chicago Press, 2006.

NIELSEN, J. **F-Shaped Pattern for Reading Web Content**. Fremont: Nielsen Norman Group, 2006.

PINE, B.J.; GILMORE, JH **The Experience Economy**. Updated ed. Boston: Harvard Business Review Press, 2011.

SHAPIRO, C.; VARIAN, HR **Information Rules: A Strategic Guide to the Network Economy**. Boston: Harvard Business School Press, 1999.

SHARP, B. **How Brands Grow: What Marketers Don't Know**. Melbourne: Oxford University Press, 2010.

SIMON, HA Designing Organizations for an Information-Rich World. In: GREENBERGER, M. (ed.). **Computers, Communications, and the Public Interest**. Baltimore: The Johns Hopkins Press, 1971.

THALER, RH; SUNSTEIN, CR **Nudge: Improving Decisions about Health, Wealth, and Happiness**. New Haven: Yale University Press, 2008.

VARIAN, HR Position Auctions. **International Journal of Industrial Organization**, vol. 25, no. 6, p. 1163–1178, 2007.