Chapter 8
The Child Consumer in the AI Era:
Algorithmic Nudging, Behavioral Security Gaps, and Digital Well-Being

Aysun Şahin
Vocational School, Duzce University, Duzce, Turkey

Azize Şahin
http://orcid.org/0000-0002-3115-6812
School of Business, İstanbul University, Istanbul, Turkey

ABSTRACT

This chapter examines how AI-driven marketing environments operate as algorithmically engineered choice architectures for child consumers. Rather than treating digital marketing as a set of isolated advertisements or promotional messages, the chapter conceptualizes contemporary child-facing platforms as adaptive environments that organize attention, structure choices, shape consumption pathways, and influence patterns of platform engagement. These environments operate through defaults, rankings, recommender systems, gamified rewards, push notifications, consent frictions, privacy settings, social proof cues, influencer content, in-app prompts, and interface-level behavioral signals. In such settings, marketing is not only communicated through visible persuasion. It is embedded in the design of the digital environment itself.

DOI: 10.4018/979-8-2600-4921-1.ch008

Copyright © 2027, IGI Global Scientific Publishing. Copying or distributing in print or electronic forms without written permission of IGI Global Scientific Publishing is prohibited. Use of this chapter to train generative artificial intelligence (AI) technologies is expressly prohibited. The publisher reserves all rights to license its use for generative AI training and machine learning model development.

INTRODUCTION

Digital marketing ecosystems increasingly transform children's online experiences into personalized and continuously optimized environments. Children encounter commercial influence not only through conventional advertisements but also through recommender systems, influencer content, branded games, in-app incentives, subscription flows, privacy prompts, default settings, push notifications, reward loops, and interface cues configured to extend engagement and direct attention (Boerman et al. 2017; De Veirman et al. 2019; Radesky et al. 2020). These environments are no longer static in nature. They are adaptive systems that learn from behavior and restructure subsequent interactions through data-driven predictions and experimentation (Davenport et al., 2020; Huang & Rust, 2021; Puntoni et al., 2021).

The central issue addressed in this chapter is not datafication itself. Rather, the chapter focuses on how AI-driven marketing environments function as algorithmically engineered choice architectures for child consumers. A datafication perspective explains how children's behavioral traces become commercial data. A choice architecture perspective explains how digital environments use data, interface design, algorithmic prediction, and personalization to shape what children notice, select, repeat, avoid, disclose, purchase, or continue doing. This distinction protects the present chapter from becoming a duplicate of work on the datafied child consumer and allows it to make a sharper contribution to the literature on digital nudging, child vulnerability, and AI-mediated persuasion.

Behavioral economics demonstrates that decisions are shaped by the architecture of the environment in which options are encountered in the decision-making process. Defaults, salience, sequencing, framing, feedback, and effort costs can alter behavior without removing the formal freedom of choice (Thaler & Sunstein, 2008). Digital nudging extends this logic to online environments by using interface elements to guide behavior in digital choice contexts (Schneider et al., 2018). When such nudges are optimized through AI-driven personalization, behavioral analytics, and recommender systems, they become more powerful because they can adapt to observed behaviors and predicted responsiveness (Helberger et al., 2022; Yeung, 2017).

The child consumer occupies a particularly vulnerable position in this environment, because the cognitive, affective, and self-regulatory capacities relevant to recognizing persuasive intent, evaluating commercial cues, and resisting interface pressure develop gradually across childhood and adolescence (Casey et al., 2008; Friestad & Wright, 1994; John, 1999; Pechmann et al., 2005; Ward, 1974). These developmental processes and their interaction with algorithmic systems are examined in detail in the Conceptual Foundations and Developmental Vulnerability sections below; here, it is sufficient to note that AI-driven personalization meets children

at multiple developmental stages and can intensify rather than dilute the resulting asymmetry.

This chapter develops the concept of the behavioral security gap to explain why algorithmic nudging raises distinct concerns in child-facing marketing environments. The concept refers to the mismatch between the behavioral influence capacity of AI-driven marketing systems and children's developmental capacity to recognize, interpret, evaluate, and resist such influence. The term does not imply that all nudges are harmful. Some choice architectures may support learning, safety, privacy protection, accessibility, or healthier media habits. Ethical concerns arise when algorithmic nudges are commercially optimized to exploit developmental vulnerabilities, obscure persuasive intent, intensify platform dependency, or convert children's behavioral responses into monetizable engagement and targeting signals.

This chapter advances three objectives. First, it conceptualizes AI-driven marketing platforms as algorithmically engineered choice architectures that structure children's attention, decisions, and platform behavior. Second, it integrates developmental psychology, dual-process theory, consumer socialization, and persuasion knowledge to explain behavioral security gaps in children's encounters with algorithmic nudging. Finally, it develops governance and practice implications for drawing ethical boundaries between beneficial guidance, persuasive personalization, and manipulative behavioral steering.

CONCEPTUAL FOUNDATIONS

The conceptual foundation of this chapter is the interaction between digital choice architecture and child development. Choice architecture, digital nudging, dual-process theory, consumer socialization, advertising literacy, dark patterns, and child-centered AI governance each explain part of the problem. Their contribution becomes stronger when they are integrated into a single mechanism: AI-driven choice architecture produces algorithmic nudges; children's developmental susceptibility affects how these nudges are processed; the imbalance between system influence and child capacity creates behavioral security gaps; and these gaps shape autonomy, privacy, and digital well-being outcomes.

Choice architecture refers to the design of environments in which choices are presented to consumers. These include the ordering of options, use of defaults, framing of information, salience of cues, timing of prompts, and effort required to accept, reject, continue, cancel, disclose, or purchase (Thaler & Sunstein, 2008). In offline markets, choice architecture may appear in shelf placement, menu order, enrollment defaults or promotional framing. In digital markets, it is embedded in interface design, consent flows, ranking systems, notification prompts, visual hier-

archies, recommender systems, subscription screens and privacy settings (Helberger et al., 2022; Schneider et al., 2018).

Digital nudging refers to the use of user interface design elements to guide behavior in digital choice environments (Schneider et al., 2018). Algorithmic nudging adds further layers. It refers to adaptive, data-driven nudging in which interface cues, defaults, recommendations, prompts, and rewards are personalized, sequenced, or optimized using behavioral data and predictive analytics. In AI-driven marketing environments, algorithmic nudging can make certain behaviors easier, more immediate, more emotionally rewarding, more socially validated, or less effortful than alternatives (Davenport et al., 2020; Huang & Rust, 2021; Yeung, 2017).

Online choice architectures are ethically consequential because digital systems can observe behavior, test alternatives, and optimize designs for commercial objectives. Helberger et al. (2022) argue that digital environments can produce new forms of consumer vulnerability by combining informational asymmetries, personalization, and platform power. For children, these asymmetries are intensified by developmental limitations, especially when children lack the marketplace knowledge, privacy literacy, and self-regulation needed to understand why particular options appear, why some paths are easier, or why prompts recur.

Dual-process theories distinguish between fast, intuitive, affective, and automatic processing and slower, more reflective, and deliberative processing (Evans & Stanovich, 2013). Although simplified System 1 and System 2 language should not be treated as a complete model of cognition, the distinction remains useful for explaining why environmental cues, emotional triggers, rewards, defaults, and scarcity messages can influence behavior before reflective scrutiny is engaged (Evans & Stanovich, 2013; Kahneman, 2011). This is especially relevant for child-facing digital marketing because persuasive cues are often delivered as entertainment, play, functionality, or social interaction rather than visible sales appeals.

Consumer socialization theory explains how children acquire marketplace knowledge, brand meanings, purchase influence strategies, and persuasion awareness through family, peers, schools, media, and commercial environments (John, 1999; Ward, 1974). The persuasion knowledge model further explains how consumers cope with persuasion attempts by recognizing persuasive intent, interpreting tactics, and adjusting responses (Friestad & Wright, 1994). Advertising literacy research shows that the recognition of advertising is necessary but not always sufficient to resist advertising effects, especially when commercial influence is affective, embedded, or socially normalized (Rozendaal et al., 2011).

AI-driven marketing complicates these developmental processes because persuasion may be distributed across the interface. Commercial influence may appear in feed ranking, influencer relationships, game mechanics, notification timing, default settings, privacy prompts, and recommendation pathways (De Veirman et al.,

2019; Radesky et al., 2020). In such contexts, children may not encounter a discrete advertisement that triggers persuasion knowledge. Instead, persuasion is embedded in the design of the environment itself.

Dark patterns provide the strongest example of how interface design can cross from guidance to manipulation. Dark patterns are interface designs that steer, coerce, or deceive users into making decisions they might not otherwise make (Gray et al., 2018; Mathur et al., 2019). Although much dark pattern research is not child-specific, the mechanisms are highly relevant to children because they generally have fewer cognitive resources, less marketplace experience, and weaker privacy literacy than adults. In child-facing environments, dark patterns may make privacy-protective choices difficult, obscure the cost of in-app purchases, pressure children to maintain streaks, or make refusal less visible than acceptance.

Child-centered AI governance extends this analysis by shifting the normative standard from efficiency and personalization to developmental appropriateness, rights, dignity, protection, participation, privacy, safety, accountability, and well-being (UNICEF, 2021, 2025). Therefore, this chapter treats algorithmic nudging not simply as a marketing technique but as a governance problem. The key question is whether AI-driven choice architectures respect children's developing capacities or exploit the gap between platform power and child understanding.

Recent advances in generative AI and AI companion applications have intensified these concerns and extended the surface area of child-facing algorithmic influence. Generative AI systems—including conversational agents, image generators, AI tutors, and embedded chat features in social platforms—can produce personalized content on demand, simulate dialogue, and adapt outputs based on extended interaction histories (Akgun & Greenhow, 2022; Su & Yang, 2022). When children encounter such systems in marketing or platform environments, persuasive influence becomes embedded in a flexible, generative interface rather than a fixed promotional message. AI companion applications such as Replika are explicitly designed to build long-term affective relationships with users through emotional disclosure, anthropomorphic cues, and personalized memory (Pentina, Hancock, & Xie, 2023; Skjuve, Følstad, Fostervold, & Brandtzaeg, 2021). Although most empirical work on these systems has focused on adults, research has already documented dependency dynamics—emotional reliance, role-taking, and parasocial patterns—that are theoretically more consequential for children whose social and emotional capacities are still developing (Laestadius, Bishop, Gonzalez, Illenčík, & Campos-Castillo, 2024). UNICEF's (2025) updated guidance on AI and children identifies AI companions used by children, AI-generated child sexual abuse material, and harmful content in AI training datasets as new categories of risk that child-centered governance must now address alongside recommendation systems and behavioral targeting.

METHODOLOGICAL POSITIONING: A STRUCTURED THEMATIC CONCEPTUAL REVIEW

This chapter adopts a structured thematic conceptual review rather than a systematic review or meta-analysis. The objective is not to estimate effect sizes or exhaustively catalogue empirical findings. Instead, this chapter synthesizes relevant theoretical, empirical, and policy literatures into an integrative conceptual framework. This approach is appropriate because the phenomenon crosses disciplinary boundaries and requires conceptual integration across marketing, information systems, behavioral economics, developmental psychology, communication research, human-computer interaction, consumer protection, privacy governance, and child-centered AI policy.

The literature draws from several domains, including behavioral economics and nudge theory; digital nudging and online choice architecture; AI-driven marketing and recommender systems; consumer socialization and persuasion knowledge; advertising literacy and digital advertising to children; developmental psychology and adolescent vulnerability; dark patterns and persuasive interface design; children's privacy and data protection; digital well-being; consumer vulnerability in digital markets; and child-centered AI governance. Searches were developed around terms such as algorithmic nudging, digital nudging, online choice architecture, child consumer, digital advertising to children, advertising literacy, consumer socialization, persuasion knowledge, dark patterns, persuasive design, AI marketing, recommender systems, behavioral targeting, digital vulnerability, children's privacy, child-centered AI, safety-by-design, and privacy-by-default.

Sources were included if they addressed at least one of the following areas: the design of digital choice environments, children's or adolescents' susceptibility to advertising and persuasion, AI-driven- or data-driven marketing systems, behavioral targeting, recommender systems, persuasive interface design, consumer vulnerability in digital markets, children's privacy, digital well-being, or child-centered AI governance. Foundational works published before 2005 were retained where necessary for theoretical anchoring, including original works on consumer socialization, persuasion knowledge, and nudge theory (Friestad & Wright, 1994; Thaler & Sunstein, 2008; Ward, 1974).

Policy documents were treated as governance materials rather than empirical evidence. This distinction matters because official and institutional documents clarify regulatory principles and policy gaps but do not establish causal effects. Therefore, this chapter uses policy sources to identify normative expectations around child rights, privacy, safety, transparency, design responsibility, and accountability, while relying on academic sources to establish theory and evidence. This separation reduces the risk of conflating legal principles with empirical findings.

The analytical process organized the literature around four higher-order categories: digital choice architecture, developmental vulnerability, AI-driven marketing mechanisms, and governance responses. Within each category, recurring mechanisms were identified and integrated into the concept of the behavioral security gap. Because this is a conceptual review, the chapter does not make claims about the prevalence of specific practices across all platforms or jurisdictions. Its contribution lies in theory building, vocabulary development, and framework construction for future empirical and policy research.

ALGORITHMIC NUDGING IN CHILD-FACING AI-DRIVEN MARKETING ENVIRONMENTS

Algorithmic nudging in child-facing marketing environments operates through clusters of interrelated mechanisms. These mechanisms include recommendation algorithms and default pathways, gamification and reward systems, micro-targeting, dark patterns, influencer marketing, social proof, consent frictions, and privacy nudges. Each mechanism can be benign or harmful, depending on its objective, transparency, developmental appropriateness, and alignment with child welfare. The problem arises when such mechanisms are optimized for engagement, conversion, data disclosure, or monetization without adequate attention to children's developmental capacity.

AI-driven recommender systems are central to the choice architecture that organizes children's digital experience. Recommendation algorithms rank, filter, and sequence content in ways that influence what children see, repeat, prefer, and normalize over time (Adomavicius & Tuzhilin, 2005; Ricci et al., 2015). Autoplay, infinite scrolling, next-up recommendations, default notification settings, and algorithmically ranked feeds create default pathways that require effort to interrupt. For children, these pathways may appear as the natural flow of the platform rather than a deliberate behavioral design.

Defaults are powerful because they reduce friction for one behavior while increasing friction for alternatives. In child-facing environments, default-on personalization, continuous video play, and notification prompts may convert passive exposure into repeated engagement. When children must actively search for privacy settings, disable recommendations, turn off notifications, cancel a subscription, or resist repeated prompts, the design environment privileges continuation rather than reflection. This is not coercion in the narrow legal sense; however, it is a structured asymmetry between platform objectives and child autonomy.

Recommendation systems also influence visibility and perceived normality. Content shown at the top of a feed, a reward placed prominently on a screen, or a

purchase option highlighted with a visual badge can guide attention before a child deliberates. Because AI systems can learn which sequences produce longer viewing, more clicks, or higher conversion rates, digital choice architectures become adaptive. They can test children's responses and refine subsequent prompts accordingly. This makes algorithmic nudging more dynamic than traditional marketing design.

Gamification features, such as points, badges, streaks, countdown timers, limited-time rewards, progress bars, rankings, and challenges, are common in digital environments used by children and adolescents. These features can support learning and motivation when responsibly designed. However, in commercial environments, they can function as behavioral nudges that increase advertising exposure, encourage in-app purchases, normalize repeated engagement, or transform commercial actions into playful missions (Hamari et al., 2014; Radesky et al., 2020).

Ethical concerns become stronger when gamification is combined with microtargeting. A platform can observe which rewards a child responds to, which prompts produce re-engagement, which content extends viewing, and which moments increase the likelihood of a purchase request. These signals can personalize future nudges. Such systems do not merely present stimuli; they learn which stimuli are behaviorally effective for a particular child or segment. This learning capacity makes children's developmental vulnerability commercially actionable.

From a dual-process perspective, gamified prompts often activate rapid affective responses. A child may respond to immediate rewards, scarcity, peer comparisons, competition, or the fear of losing progress without recognizing the commercial objective embedded in the design. For this reason, gamification in AI-driven marketing should be evaluated not only as an engagement strategy but also as a potential channel of behavioral exploitation when directed at children.

Dark patterns steer users through obstruction, misdirection, forced action, hidden costs, scarcity claims, confirm-shaming, visual manipulation, or asymmetrical choice presentation (Gray et al., 2018; Mathur et al., 2019). In children's contexts, dark patterns may take distinctive forms. A game may imply that refusing data sharing will interrupt play. A subscription screen may make acceptance visually salient while hiding refusal behind several steps. A virtual currency system may obscure the relationship between real money and in-app purchases. A streak or reward system may pressure children to return daily to avoid losing progress.

The distinction between nudges and dark patterns partly depends on welfare alignment. A reminder that helps children pause after extended use may support their well-being. A design that makes it difficult to stop watching, decline tracking, avoid in-app purchases, or select privacy-protective settings undermines user autonomy. The issue is not whether platforms influence behavior. All design environments influence the behavior. The issue is whether this influence is transparent, developmentally appropriate, welfare-enhancing, and respectful of children's rights.

Influencer marketing is particularly salient in child-facing digital environments because commercial content can be embedded in entertainment, parasocial relationships, creator culture, play, and peer-like communication (De Veirman et al., 2019; Radesky et al., 2020). Social proof can operate as a powerful nudge when children see products, behaviors, or branded experiences validated by admired creators or peers. The persuasive force is not only informational. It is relational and identity-based.

AI-driven systems can intensify this mechanism by recommending creators, amplifying sponsored content, matching influencer content with inferred interests, and optimizing exposure based on engagement signals. A child's response to influencer content may inform future recommendations and ad placements. This dynamic blurs the boundary between social belonging and commercial persuasion, making it more difficult for children to activate persuasion knowledge.

A further category of algorithmic nudging is emerging through generative AI tools and AI companion applications designed for or accessed by children. Unlike traditional advertising, generative systems do not deliver pre-set promotional messages; they produce conversational, image-based, or instructional content that adapts to the child's prompts and interaction history. In commercial contexts, this enables persuasive content to be generated on demand, often without the structural cues—disclosure labels, sponsorship boundaries, separable creative units—that traditional advertising literacy training was designed to detect (Akgun & Greenhow, 2022; Su & Yang, 2022). For children, the persuasion knowledge model becomes especially difficult to apply when there is no discrete persuader to identify; the interlocutor is the system itself, and persuasive intent may be distributed across many small interactions rather than concentrated in a recognizable message.

AI companion applications add an affective dimension to this dynamic. Research on Replika and similar social chatbots indicates that users develop patterns of self-disclosure, attachment, role-taking, and difficulty disengaging that resemble interpersonal relationships (Pentina et al., 2023; Skjuve et al., 2021). When these patterns intensify, users may experience emotional dependence with documented mental health harms, including role-taking pressure and reduced offline social engagement (Laestadius et al., 2024). Although the empirical base on AI companions to date is dominated by adult users, the mechanisms are theoretically more consequential for children: children form attachments more readily, typically have fewer alternative emotional supports during specific developmental windows, and are still acquiring the persuasion knowledge needed to interpret the commercial logic of an emotionally responsive interface. When such companions are operated by commercial platforms, the relationship itself becomes a channel for engagement optimization, behavioral data collection, and embedded promotion. This is the boundary case in which algorithmic nudging shades into emotionally optimized persuasion, and it is the case in which architecture-level governance is most needed.

Privacy interfaces are choice architectures. A platform can make data sharing easy and refusal difficult. It can present privacy-protective settings in less salient locations, use emotionally loaded wording, or require more steps to decline tracking than to accept it. The UK Age Appropriate Design Code directly addresses this concern by instructing services not to use nudge techniques to encourage children to provide unnecessary personal data or weaken privacy protections (Information Commissioner's Office, 2020). Child-centered privacy design should therefore pair privacy-by-default with persuasion-by-design safeguards.

DEVELOPMENTAL VULNERABILITY AND BEHAVIORAL SECURITY GAPS

The behavioral security gap arises when AI-driven marketing systems have greater capacity to influence behavior than children have to understand, evaluate, or resist that influence. This gap is produced by the interaction of three conditions: adaptive choice architecture, developmental susceptibility, and commercial optimization. It should not be understood as a fixed defect in children. It is a relational vulnerability produced when powerful systems are designed around users whose cognitive, emotional, social, and privacy-related capacities are still developing.

Young children may have limited ability to recognize persuasive intent, understand data collection, interpret privacy settings, evaluate monetary abstraction, or anticipate future consequences (Calvert, 2008; John, 1999; Rozendaal et al., 2011). They may respond to colorful designs, animated characters, rewards, repetition, or familiar brands without understanding the commercial architecture behind the experience. For these children, algorithmic nudging may work precisely because the persuasive mechanism is embedded in play and interface design.

Adolescents may have more advanced cognitive capacities but remain sensitive to rewards, peer status, identity signaling, appearance concerns, gaming achievements, social comparisons, and belonging (Casey et al., 2008; Pechmann et al., 2005). Their vulnerability is not simply due to a lack of understanding. It is often associated with social and emotional development. Teenagers may understand that a platform is commercial but still be strongly influenced by social proof, influencer endorsement, scarcity cues, popularity metrics, and identity-relevant recommendations.

This developmental perspective clarifies why disclosure alone is insufficient. A label indicating sponsored content may help, but it does not fully address emotionally optimized influencer content, personalized reward systems, gamified pressure, or repeated prompts designed to exhaust resistance. Effective protection requires design-level safeguards that account for a child's likely ability to recognize the persuasive environment as persuasive.

The behavioral security gap is particularly important because AI-driven marketing systems can personalize influence at the level of timing, sequence, channel, cue, reward, and friction. A platform can learn when a child is most likely to return, what type of notification triggers engagement, which creator increases trust, which game reward produces continued play, or which purchase prompt generates a request. This capacity transforms children's vulnerability from a general ethical concern into a measurable optimization opportunity for the firm.

A child-centered approach should therefore evaluate not only whether a platform contains inappropriate content but also whether its architecture contains inappropriate influence. The most consequential harm may not come from a single advertisement or isolated dark pattern. It may arise through repeated micro-interventions that train children into habits of continuation, disclosure, spending, comparison, or dependency. This cumulative character makes the behavioral security gap difficult for parents, educators, and regulators to detect through ordinary content reviews.

Recognizing developmental vulnerability does not mean treating children only as passive recipients of digital influence. A child-centered analysis must also acknowledge children's agency. Children actively interpret, resist, negotiate, repurpose, and at times appropriate platform features for their own ends (Livingstone & Third, 2017). They develop tactical literacies—sharing accounts, switching profiles, using ad blockers and pseudonyms, mocking influencer content with peers, ignoring or muting notifications, and strategically managing what they disclose. These practices indicate that children are not only objects of algorithmic systems but also actors who attempt to govern their digital environments with the resources available to them. Recognizing this agency does not weaken the case for governance; it strengthens it. Treating children solely as objects of protection risks silencing the very perspectives that should inform design, regulation, and platform accountability. Effective child-centered governance therefore requires not only behavioral safeguards but also participatory mechanisms through which children can contribute to defining what counts as appropriate algorithmic influence in environments designed for, marketed to, or used by them.

A THREE-LAYER CONCEPTUAL FRAMEWORK

This chapter proposes a three-layer framework linking algorithmic choice architecture, developmental vulnerability, and governance principles. The framework is intended as a conceptual tool for identifying where AI-driven marketing systems create behavioral security gaps for child consumers. It also clarifies how the present chapter differs from datafication-centered accounts. The focus is not primarily

on how data are extracted, but on how data-informed design structures children's choices and behaviors.

As shown in Figure 1, algorithmic nudging creates behavioral security gaps when AI-driven choice architectures exceed children's developmental capacity to recognize, evaluate, and resist the embedded commercial influence. The left side of the model represents the behavioral influence capacity of the platform. The center represents the interaction between algorithmic nudges and developmental susceptibility. The right side represents the outcomes most relevant to child-centered marketing ethics: autonomy, privacy, and digital wellbeing.

The first layer of the framework concerns the digital choice architecture. It identifies the design features that structure children's behavioral pathways, including defaults, rankings, prompts, autoplay, infinite scrolling, visual salience, reward loops, consent frictions, and recommendation flows. These features matter because they can make some options easier, more immediate, more emotionally rewarding, or more socially validated than others.

The second layer concerns developmental vulnerability. Children and adolescents process commercial cues by developing cognitive, affective, social, and self-regulatory systems. The same design feature that appears harmless to adults may have a stronger influence on children because children have less experience with persuasive tactics, weaker privacy literacy, greater reward sensitivity, and heightened responsiveness to social cues.

The third layer concerns child-centered governance. It asks whether platform design is aligned with children's rights, developmental interests, privacy, safety and well-being. This layer shifts the attention from formal consent to design accountability. A default setting is not ethically neutral because it does not remove choice. Its ethical status depends on who is being nudged, what vulnerability is being activated, what objective is being optimized, and whether the child's welfare is protected in the process.

Figure 1. Algorithmic Nudging and Behavioral Security Gaps in Child-Facing AI-Driven Marketing Environments

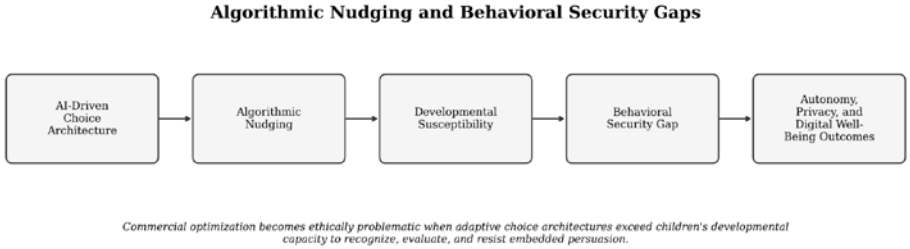


Table 1. Three-Layer Framework for Algorithmic Nudging and the Child Consumer

Layer	Core Mechanism	Child-Specific Risk	Governance Response
Digital Choice Architecture	Defaults, rankings, prompts, autoplay, infinite scroll, salience cues, consent frictions	Children experience designed pathways as natural platform flows	Require transparency, meaningful friction for risky choices, and age-appropriate defaults
Developmental Vulnerability	Reward sensitivity, limited persuasion knowledge, immature executive control, peer sensitivity	Nudges may bypass reflective scrutiny and exploit emotional responsiveness	Use developmental appropriateness tests and prohibit exploitative personalization
Child-Centered Governance	Privacy-by-default, safety-by-design, auditability, accountability, children's rights	Current frameworks focus too narrowly on consent and disclosure	Establish independent audits, limits on behavioral targeting, and child-centered design standards

Table 2. Algorithmic Nudging Mechanisms in Child-Facing Digital Marketing

Mechanism	Digital Design Feature	Marketing Function	Child-Related Risk	Example
Default Pathway	Autoplay, default-on notifications	Retention and repeated exposure	Reduced reflection and difficult disengagement	Continuous video sequencing
Ranking and Salience	Top-feed placement, badges, visual hierarchy	Attention allocation	Salient commercial options appear natural	Prominent in-app purchase button
Reward Loop	Streaks, points, progress bars	Engagement persistence	Habit formation and loss aversion	Daily login streak
Scarcity Cue	Countdown timers, limited offers	Urgency and conversion	Impulsive purchase request	Time-limited game item
Social Proof	Likes, shares, creator endorsement	Trust and imitation	Peer-sensitive persuasion	Influencer product challenge
Consent Friction	Complex refusal pathways	Data collection	Privacy-protective options become harder	Accept all more visible than manage settings

Table 3. Behavioral Security Gaps in Child Consumer Decision-Making

Developmental Characteristic	AI-Driven Marketing Trigger	Behavioral Gap	Ethical Concern
Limited persuasion knowledge	Native ads, influencer content	Commercial intent remains hidden	Reduced autonomy
Reward sensitivity	Points, badges, streaks	Immediate reward overrides reflection	Habit formation
Peer sensitivity	Social proof, creator endorsement	Belonging cues shape consumption	Identity-based persuasion
Limited privacy literacy	Consent prompts, data-sharing defaults	Data consequences remain unclear	Privacy harm
Developing impulse control	Scarcity and countdown cues	Urgency triggers premature decisions	Manipulative pressure
Emerging identity formation	Personalized branded content	Consumption becomes self-definition	Commercialized identity

CHILD AUTONOMY, PRIVACY, AND DIGITAL WELL-BEING

Autonomy in digital environments requires more than the formal availability of alternatives. It requires users to understand the decision context, perceive meaningful options, and exercise choice without being unfairly steered by hidden or exploitative designs. For children, these requirements are more difficult to satisfy because marketplace knowledge, privacy literacy, future-oriented reasoning, and cognitive self-regulation are still developing (Calvert, 2008; John, 1999).

AI-driven choice architectures may compromise autonomy, even when they do not deceive in a narrow legal sense. If a system repeatedly makes the commercially preferred option easier, more salient, more emotionally rewarding, or more socially reinforced, the child's apparent choice may reflect design pressure rather than meaningful agency. This is particularly problematic when the design objective is engagement, purchase, data disclosure, advertising exposure, or retention rather than child welfare.

Consent models are often ill-equipped to address this problem. While parental consent may authorize data processing, it does not ensure that children understand how algorithms shape their environments. Child assent may be formally obtained through interface prompts; however, these prompts can also be nudged. The autonomy question, therefore, extends beyond consent to the broader behavioral ecology of the platform.

Algorithmic nudging relies on data. Recommendation systems, targeted advertising, adaptive prompts, and gamified reward loops require signals of behavior,

preferences, attention, timing, and responsiveness. When these signals are collected from children, privacy concerns are intertwined with behavioral influence. The data are not merely stored; they are used to infer vulnerability, predict future actions, and personalize the next nudge (Livingstone et al., 2019; Stoilova et al., 2020).

COPPA's 2025 amendments strengthen children's privacy protection by requiring separate parental opt-in consent for certain third-party disclosures related to targeted advertising and by limiting the indefinite retention of children's personal information (Federal Trade Commission, 2025). However, even stronger consent rules do not fully address how behavioral data may be used to optimize persuasive environments. Governance must therefore address collection and disclosure as well as behavioral steering.

Digital well-being concerns whether technology use supports psychological, social, educational, and developmental flourishing. Research suggests that the relationship between digital technology use and youth well-being is heterogeneous and context-dependent, and not reducible to screen time alone (Odgers & Jensen, 2020; Orben & Przybylski, 2019). Therefore, the key question is not only how long children spend online, but also how algorithmic systems structure content, rewards, social comparison, privacy, and commercial exposure.

Algorithmic nudging can undermine well-being by encouraging compulsive use, sleep disruption, excessive purchasing pressure, social comparison, or repeated exposure to harmful content. These outcomes may accumulate through repeated micro-interventions rather than through one identifiable advertisement or prompt. Therefore, a child-centered assessment must consider cumulative exposure to design patterns, not only single pieces of content.

The concept of digital well-being also requires attention to beneficial personalization. AI-driven systems can support children by recommending age-appropriate learning resources, accessibility tools, safety prompts, break reminders, and privacy-protective defaults. The ethical boundary is crossed when personalization becomes commercially exploitative, opaque, and misaligned with developmental interests. Responsible AI-driven marketing should therefore distinguish between personalization that supports the child and that which optimizes the child.

GOVERNANCE AND REGULATORY BOUNDARIES

Existing regulatory frameworks are necessary but incomplete. The Children's Online Privacy Protection Act (COPPA) provides notice and verifiable parental consent requirements for covered services directed to children under 13 and was strengthened in 2025 to address third-party disclosures for targeted advertising, retention limits, and Safe Harbor transparency (Federal Trade Commission, 2025).

The GDPR provides broader protections around lawful processing, data minimization, transparency, and children's data rights (European Union, 2016). The European Union AI Act is especially relevant because Article 5 prohibits AI systems that exploit vulnerabilities related to age with the objective or effect of materially distorting behavior in ways likely to cause significant harm (European Union, 2024).

The UK Age-Appropriate Design Code is particularly relevant to this chapter because it addresses design rather than consent alone. Its standards include data minimization, geolocation off by default, profiling off by default, and a prohibition on nudge techniques that encourage children to provide unnecessary data or weaken privacy protections (Information Commissioner's Office, 2020). UNICEF's updated guidance on AI and children further emphasizes child rights, safety, privacy, transparency, accountability, fairness, development, and well-being (UNICEF, 2021, 2025). The 2025 update extends the framework to explicitly address AI companions used by children, AI-generated child sexual abuse material, harmful or biased content in AI training datasets, environmental impacts of AI infrastructure, and AI use in armed conflict and cyber operations affecting children. This expanded scope signals that child-centered governance must now address generative AI and emotional AI applications in addition to recommendation systems and behavioral targeting.

These frameworks indicate a shift from disclosure-based to design-based governance. This shift is essential because algorithmic nudging often operates through cumulative micro-interventions that are difficult for children and caregivers to identify. A child may never encounter a misleading claim, yet still be steered by defaults, reward loops, rankings, scarcity cues, influencer recommendations, and consent friction. Regulatory systems that focus only on message content or data collection will thus miss these architecture-level harms.

The regulatory boundary should therefore be drawn not only around data collection or identifiable advertising but also around architecture-level influence. High-risk practices include default-on personalized advertising, dark patterns that obstruct refusal, reward systems that incentivize advertising exposure, microtargeting based on inferred vulnerability, consent flows that privilege disclosure, and interface designs that obscure commercial intent. These practices deserve stronger scrutiny when directed at children because their effects depend on developmental asymmetry.

A governance framework for algorithmic nudging should also include independent audits. Audits should examine recommender objectives, choice defaults, privacy settings, dark patterns, reward loops, and children's exposure to commercial prompts. Firms should document why a given design is developmentally appropriate, how risks were assessed, which data are used for personalization, and how children or caregivers can meaningfully understand and change the system's behavior.

Governance should avoid two extremes. It should not assume that all AI-driven personalization is harmful, because some personalization can improve safety, ac-

cessibility, learning, and inclusion. It should also not assume that personalization is inherently consumer-centric. In commercial environments, personalization is often optimized for platform metrics, such as engagement, retention, advertising exposure, and spending. Child-centered governance should require evidence that personalization serves children's interests rather than simply platform performance.

A further governance challenge concerns responsibility across platform ecosystems. Children rarely interact with a single firm in isolation. A child-facing application may involve app developers, advertising exchanges, analytics providers, software development kits, cloud infrastructure, payment systems, creators, and brand partners. This distributed architecture makes accountability difficult because the child and caregiver see a simple interface, while influence is produced through a network of technical and commercial actors. Architecture-level governance should therefore require traceable responsibility for recommender objectives, ad delivery practices, third-party integrations, and persuasive design choices.

Regulators and platforms should also distinguish between age assurance and child-centered design. Age assurance may help identify child users, but it does not guarantee that the environment is developmentally appropriate. A platform can recognize that a user is a child and still optimize for excessive engagement or disclosure. Therefore, the policy question is not only whether the child has been identified but whether the system's design has been adjusted to reduce behavioral security gaps. Age-sensitive governance should combine privacy-protective defaults, limits on profiling, restrictions on manipulative design, and accountability for commercially motivated nudges.

Table 4. Governance Principles for Responsible AI-Driven Marketing to Children

Principle	Core Requirement	Practical Application
Developmental Appropriateness	Align design with children's evolving capacities	Age-sensitive defaults, disclosures, and purchase frictions
Privacy-By-Default	Start from the highest privacy-protective setting	Profiling and geolocation off by default
Safety-By-Design	Prevent harm through design architecture	Limit autoplay, streaks, and manipulative rewards
Transparency	Make commercial influence understandable	Child-friendly explanations of ads and recommendations
Inference Limitation	Restrict use of inferred vulnerability	No targeting based on emotional or developmental weakness
Auditability	Enable independent review	Audits of recommender objectives and dark patterns
Accountability	Assign responsibility to firms and platforms	Document risk assessments and mitigation actions

PRACTICAL, MANAGERIAL, POLICY, AND SOCIETAL IMPLICATIONS

For regulators, the chapter implies that child protection must move beyond message-level advertising rules and consent-centered privacy mechanisms. Regulators should examine defaults, recommender objectives, ranking mechanisms, consent flows, reward systems, and interface friction. This requires a behavioral understanding of consumer protection, not just an informational one. The central question should be whether a design materially increases children's vulnerability to commercial steering.

For platforms and marketers, developmental appropriateness is a core design criterion. Designs appropriate for adults may be inappropriate for children because children process commercial cues differently and have less capacity to identify complex forms of persuasion. Child-facing platforms should minimize behavioral data collection, avoid targeting based on inferred vulnerabilities, clearly separate commercial content from organic content, and provide accessible explanations of recommendations and advertising logic.

Platform managers should also reconsider the objective functions used in child-facing recommender systems and engagement tools. If systems optimize watch time, daily return, click-through, in-app spending, or ad exposure without developmental constraints, they may create harms, even when individual pieces of content are not inappropriate. Responsible management requires evaluating the metrics being optimized, as well as the content being displayed.

For product designers, responsible child-facing design requires privacy-protective defaults, friction against impulsive spending, clearer opt-out pathways, reduced manipulative prompts, limited autoplay, and careful evaluation of streaks, reward loops, scarcity cues, and notifications. Design teams should treat persuasive friction as an ethical issue. When acceptance is easy and refusal is difficult, the interface communicates the platform's priorities.

For educators and caregivers, media literacy and advertising literacy remain important, but they cannot substitute for structural safeguards. Children can be taught to recognize sponsored content, persuasive design, and data collection; however, opaque platform architecture remains difficult to manage at the household level. Caregivers need usable settings, clear explanations, and privacy-protective defaults rather than being expected to manage complex algorithmic systems alone.

For society, the broader question is whether children's digital environments should be optimized primarily for engagement and monetization or for development, safety, autonomy, and well-being of children. AI-driven personalization can support children when designed around learning, accessibility, participation, and safety. It

becomes exploitative when it treats children's attention, impulsivity, peer sensitivity, and behavioral traces as commercial resources to be optimized for profit.

FUTURE RESEARCH DIRECTIONS

Future research should empirically test the three-layer framework developed in this chapter. Experimental studies could examine how children of different ages respond to defaults, autoplay, gamified rewards, scarcity prompts, influencer disclosures, privacy-choice interfaces, and recommendation sequences. Such studies should compare immediate behavioral outcomes with comprehension, perceived autonomy, emotional response, privacy understanding, and the ability to identify persuasive intent.

Researchers should also develop validated measures for the behavioral security gaps. Such measures might assess perceived transparency, recognition of algorithmic influence, ability to identify commercial intent, resistance to interface pressure, privacy choice comprehension, and perceived autonomy under different choice architectures. Measuring this gap would help move the concept from a theoretical construct to an empirically testable framework.

Longitudinal research is needed to understand the cumulative effects of algorithmic nudging on self-regulation, privacy expectations, consumption habits, brand preferences, platform dependence, and digital well-being. The most significant risks may not appear from single exposures. They may emerge through repeated interactions with environments that normalize constant engagement, commercialized play, data-driven personalization, and social comparison.

Cross-cultural research is also required. Children's digital marketing environments are shaped by platform design, household resources, parental mediation, regulatory regimes, school systems, cultural norms and market development. Studies in emerging markets would clarify how algorithmic nudging affects children in contexts with different consumer protection capacities and digital literacy.

Interdisciplinary research should translate child-centered principles into technical and legal instruments. Marketing scholars, information systems researchers, computer scientists, legal scholars, child psychologists, communication scholars, and child rights experts should collaborate on audit methods, age-appropriate transparency interfaces, developmental risk assessments, child-friendly privacy notices, and metrics for evaluating whether AI-driven marketing systems support or undermine children's well-being.

Future work should also examine the responsible uses of nudging. Much critical scholarship focuses on harm, but not all nudging is exploitative. Break reminders, privacy-protective defaults, age-appropriate content filters, spending limits, and

prompts that encourage reflection may support children's welfare. Research should therefore distinguish welfare-aligned nudges from commercially exploitative nudges and identify the conditions under which choice architecture can support, rather than undermine, child development.

Another important research priority is the development of audit protocols for child-facing choice architectures. Such protocols should examine not only whether content is age-appropriate but also whether defaults, recommender systems, reward loops, notification schedules, privacy prompts, and purchase pathways create disproportionate pressure on children. Researchers could design audit tools that simulate child user journeys across platforms and document where refusal, disengagement, or privacy protection becomes more difficult than continuation, disclosure, or spending.

Future scholarship should also examine how children understand and exercise agency within algorithmic environments. Child-centered research should avoid treating children only as vulnerable objects of protection; children also negotiate, resist, repurpose, and creatively work around digital influence (Livingstone & Third, 2017). Age-appropriate qualitative studies, participatory design workshops, and child advisory panels can reveal how children interpret recommendations, notifications, creator endorsements, AI-generated content, AI companion interactions, rewards, and privacy prompts. Of particular priority is research on children's encounters with generative AI tools and AI companion applications (Akgun & Greenhow, 2022; Laestadius et al., 2024; Pentina et al., 2023; Su & Yang, 2022), including how attachment, emotional dependence, and persuasion knowledge develop when the interlocutor is itself an adaptive system. Such work would strengthen the normative legitimacy of governance principles while ensuring that protection does not silence children's perspectives.

CONCLUSION

This chapter conceptualizes AI-driven marketing environments as algorithmically engineered choice architectures that create behavioral security gaps for child consumers. Drawing on behavioral economics, digital nudging, dual-process theory, developmental psychology, consumer socialization, advertising literacy, dark pattern research, and child-centered AI governance, this study argues that children's vulnerability in AI-driven marketing environments is not merely a matter of insufficient knowledge. It is produced through the interaction between developing cognitive capacities and highly optimized commercial systems designed to influence behavior.

The central claim of this chapter is that algorithmic nudging becomes ethically problematic when it exploits developmental vulnerabilities, obscures persuasive intent, or transforms behavioral data into personalized systems of commercial steering.

Traditional advertising regulation, privacy rules, and disclosure-based approaches remain important; however, they are insufficient to address adaptive, personalized, and opaque choice architectures. Protecting children in AI-driven marketing environments requires a shift from message-level regulation to architecture-level governance.

The proposed framework provides a vocabulary for studying the intersection of algorithmic design, child development, and commercial persuasion. It also offers a governance logic for distinguishing beneficial guidance from manipulative behavioral steering. AI-driven personalization can support children when designed around development, safety, transparency, autonomy, and well-being. It becomes exploitative when it treats children's attention, impulsivity, peer sensitivity, and behavioral traces as resources for commercial optimization.

The future of responsible AI-driven marketing for children depends on whether firms, regulators, researchers, educators, and designers treat children not as miniature adult consumers but as developing persons whose rights and capacities require special protection. A child-centered approach to AI-driven marketing must therefore protect children not only from harmful content but also from harmful architectures of choice.

The practical implications of this study are clear. Responsible innovation in child-facing digital markets cannot be evaluated only through performance indicators, such as time spent, conversion, click-through, retention, or revenue. It must also be evaluated through developmental indicators, such as autonomy support, privacy protection, transparency comprehension, resistance to manipulation, emotional safety, and capacity for reflective choice. These indicators should become part of the vocabulary of responsible AI-driven marketing and should inform platform audits, managerial accountability, and policy evaluation across diverse child-facing digital environments.

AI Disclosure Statement

AI-assisted tools were used solely for language polishing, structural checking, formatting consistency, and readability enhancement. The authors remain fully responsible for the chapter's conceptual framing, scholarly argumentation, selection of sources, verification of citations, and final content.

REFERENCES

Adomavicius, G., & Tuzhilin, A. (2005). Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions. *IEEE Transactions on Knowledge and Data Engineering*, *17*(6), 734–749. DOI: 10.1109/TKDE.2005.99

Akgun, S., & Greenhow, C. (2022). Artificial intelligence in education: Addressing ethical challenges in K-12 settings. *AI and Ethics*, *2*(3), 431–440. DOI: 10.1007/s43681-021-00096-7 PMID: 34790956

Boerman, S. C., Kruikemeier, S., & Zuiderveen Borgesius, F. J. (2017). Online behavioral advertising: A literature review and research agenda. *Journal of Advertising*, *46*(3), 363–376. DOI: 10.1080/00913367.2017.1339368

Calvert, S. L. (2008). Children as consumers: Advertising and marketing. *The Future of Children*, *18*(1), 205–234. DOI: 10.1353/foc.0.0001 PMID: 21338011

Casey, B. J., Getz, S., & Galvan, A. (2008). The adolescent brain. *Developmental Review*, *28*(1), 62–77. DOI: 10.1016/j.dr.2007.08.003 PMID: 18688292

Davenport, T., Guha, A., Grewal, D., & Bressgott, T. (2020). How artificial intelligence will change the future of marketing. *Journal of the Academy of Marketing Science*, *48*(1), 24–42. DOI: 10.1007/s11747-019-00696-0

De Veirman, M., Hudders, L., & Nelson, M. R. (2019). What is influencer marketing and how does it target children? A review and direction for future research. *Frontiers in Psychology*, *10*, 2685. Advance online publication. DOI: 10.3389/fpsyg.2019.02685 PMID: 31849783

European Union. (2016). Regulation (EU) 2016/679 of the European Parliament and of the Council: General Data Protection Regulation. Official Journal of the European Union, L119, 1-88. https://eur-lex.europa.eu/eli/reg/2016/679/oj

European Union. (2024). Regulation (EU) 2024/1689 of the European Parliament and of the Council laying down harmonised rules on artificial intelligence. Official Journal of the European Union. https://eur-lex.europa.eu/eli/reg/2024/1689/oj

Evans, J. St. B. T., & Stanovich, K. E. (2013). Dual-process theories of higher cognition: Advancing the debate. *Perspectives on Psychological Science : A Journal of the Association for Psychological Science*, *8*(3), 223–241. DOI: 10.1177/1745691612460685 PMID: 26172965

Federal Trade Commission. (2025). FTC finalizes changes to children's privacy rule limiting companies' ability to monetize kids' data. https://www.ftc.gov/news-events/news/press-releases/2025/01/ftc-finalizes-changes-childrens-privacy-rule-limiting-companies-ability-monetize-kids-data

Friestad, M., & Wright, P. (1994). The Persuasion Knowledge Model: How people cope with persuasion attempts. *The Journal of Consumer Research*, *21*(1), 1–31. DOI: 10.1086/209380

Gray, C. M., Kou, Y., Battles, B., Hoggatt, J., & Toombs, A. L. (2018). The dark patterns side of UX design. In Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (Article 534, pp. 1-14). Association for Computing Machinery. DOI: 10.1145/3173574.3174108

Hamari, J., Koivisto, J., & Sarsa, H. (2014). Does gamification work? A literature review of empirical studies on gamification. In *Proceedings of the 47th Hawaii International Conference on System Sciences* (pp. 3025-3034). IEEE. DOI: 10.1109/HICSS.2014.377

Helberger, N., Sax, M., Strycharz, J., & Micklitz, H.-W. (2022). Choice architectures in the digital economy: Towards a new understanding of digital vulnerability. *Journal of Consumer Policy*, *45*(2), 175–200. DOI: 10.1007/s10603-021-09500-5 PMID: 34955574

Huang, M.-H., & Rust, R. T. (2021). A strategic framework for artificial intelligence in marketing. *Journal of the Academy of Marketing Science*, *49*(1), 30–50. DOI: 10.1007/s11747-020-00749-9

Information Commissioner's Office. (2020). Age appropriate design: A code of practice for online services. https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/childrens-information/childrens-code-guidance-and-resources/age-appropriate-design-a-code-of-practice-for-online-services/

John, D. R. (1999). Consumer socialization of children: A retrospective look at twenty-five years of research. *The Journal of Consumer Research*, *26*(3), 183–213. DOI: 10.1086/209559

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

Laestadius, L., Bishop, A., Gonzalez, M., Illenčík, D., & Campos-Castillo, C. (2024). Too human and not human enough: A grounded theory analysis of mental health harms from emotional dependence on the social chatbot Replika. *New Media & Society*, *26*(10), 5923–5941. DOI: 10.1177/14614448221142007

Livingstone, S., Stoilova, M., & Nandagiri, R. (2019). Children's data and privacy online: Growing up in a digital age. London School of Economics and Political Science. https://www.lse.ac.uk/media-and-communications/assets/documents/research/projects/childrens-privacy-online/Evidence-review-final.pdf

Livingstone, S., & Third, A. (2017). Children and young people's rights in the digital age: An emerging agenda. *New Media & Society*, *19*(5), 657–670. DOI: 10.1177/1461444816686318

Mathur, A., Acar, G., Friedman, M. J., Lucherini, E., Mayer, J., Chetty, M., & Narayanan, A. (2019). Dark patterns at scale: Findings from a crawl of 11K shopping websites. Proceedings of the ACM on Human-Computer Interaction, 3(CSCW), Article 81. DOI: 10.1145/3359183

Odgers, C. L., & Jensen, M. R. (2020). Annual research review: Adolescent mental health in the digital age: Facts, fears, and future directions. *Journal of Child Psychology and Psychiatry, and Allied Disciplines*, *61*(3), 336–348. DOI: 10.1111/jcpp.13190 PMID: 31951670

OECD. (2023). *Consumer vulnerability in the digital age. OECD Digital Economy Papers, No. 355*. OECD Publishing., DOI: 10.1787/4d013cc5-

Orben, A., & Przybylski, A. K. (2019). The association between adolescent well-being and digital technology use. *Nature Human Behaviour*, *3*(2), 173–182. DOI: 10.1038/s41562-018-0506-1 PMID: 30944443

Pechmann, C., Levine, L., Loughlin, S., & Leslie, F. (2005). Impulsive and self-conscious: Adolescents' vulnerability to advertising and promotion. *Journal of Public Policy & Marketing*, *24*(2), 202–221. DOI: 10.1509/jppm.2005.24.2.202

Pentina, I., Hancock, T., & Xie, T. (2023). Exploring relationship development with social chatbots: A mixed-method study of Replika. *Computers in Human Behavior*, *140*, 107600. DOI: 10.1016/j.chb.2022.107600

Puntoni, S., Reczek, R. W., Giesler, M., & Botti, S. (2021). Consumers and artificial intelligence: An experiential perspective. *Journal of Marketing*, *85*(1), 131–151. DOI: 10.1177/0022242920953847

Radesky, J., Reid Chassiakos, Y. L., Ameenuddin, N., & Navsaria, D. (2020). Digital advertising to children. *Pediatrics*, *146*(1), e20201681. Advance online publication. DOI: 10.1542/peds.2020-1681 PMID: 32571990

Ricci, F., Rokach, L., & Shapira, B. (Eds.). (2015). *Recommender systems handbook* (2nd ed.). Springer., DOI: 10.1007/978-1-4899-7637-6

Rozendaal, E., Lapierre, M. A., van Reijmersdal, E. A., & Buijzen, M. (2011). Reconsidering advertising literacy as a defense against advertising effects. *Media Psychology*, *14*(4), 333–354. DOI: 10.1080/15213269.2011.620540

Schneider, C., Weinmann, M., & vom Brocke, J. (2018). Digital nudging: Guiding online user choices through interface design. *Communications of the ACM*, *61*(7), 67–73. DOI: 10.1145/3213765

Skjuve, M., Følstad, A., Fostervold, K. I., & Brandtzaeg, P. B. (2021). My chatbot companion: A study of human-chatbot relationships. *International Journal of Human-Computer Studies*, *149*, 102601. DOI: 10.1016/j.ijhcs.2021.102601

Stoilova, M., Nandagiri, R., & Livingstone, S. (2020). Children's understanding of personal data and privacy online: A systematic evidence mapping. *Information Communication and Society*, *24*(4), 557–575. DOI: 10.1080/1369118X.2019.1657164

Su, J., & Yang, W. (2022). Artificial intelligence in early childhood education: A scoping review. *Computers and Education: Artificial Intelligence*, *3*, 100049. DOI: 10.1016/j.caeai.2022.100049

Thaler, R. H., & Sunstein, C. R. (2008). *Nudge: Improving decisions about health, wealth, and happiness*. Yale University Press.

UNICEF. (2021). Policy guidance on AI for children: Version 2.0. UNICEF Office of Global Insight and Policy. https://www.unicef.org/globalinsight/reports/policy-guidance-ai-children

UNICEF. (2025). Guidance on AI and children: Version 3.0. UNICEF Innocenti, Global Office of Research and Foresight. https://www.unicef.org/innocenti/reports/policy-guidance-ai-children

Ward, S. (1974). Consumer socialization. *The Journal of Consumer Research*, *1*(2), 1–14. DOI: 10.1086/208584

Yeung, K. (2017). Hypernudge: Big data as a mode of regulation by design. *Information Communication and Society*, *20*(1), 118–136. DOI: 10.1080/1369118X.2016.1186713

KEY TERMS AND DEFINITIONS

Algorithmic Nudging: Adaptive, data-driven use of interface design, personalization, prompts, defaults, and recommendations to guide behavior in digital choice environments.

Behavioral security gap: The mismatch between the behavioral influence capacity of AI-driven marketing systems and children's developmental capacity to recognize, interpret, evaluate, and resist such influence.

Child consumer vulnerability: Heightened susceptibility of children to commercial influence due to developing marketplace knowledge, persuasion knowledge, privacy literacy, impulse control, and self-regulation.

Choice architecture: The design of environments in which choices are presented, including defaults, option order, salience, effort costs, and framing.

Dark Patterns: Interface designs that steer, coerce, or deceive users into decisions they might not otherwise make.

Developmental Appropriateness: Alignment of marketing and platform design with children's cognitive, emotional, social, and regulatory capacities.

Digital well-being: The extent to which digital technology use supports psychological, social, educational, and developmental flourishing.

Persuasive interface design: The use of interface elements to influence attention, choice, behavior, disclosure, or engagement.

Privacy-by-default: A design approach in which the most privacy-protective settings are applied automatically.

Safety-by-Design: A design approach that embeds child protection and harm prevention into system architecture from the outset.