



EXECUTIVE SUMMARY

The Persuasion Pipeline

How Commercial AI Targets, Shapes, and Sells
to a Generation of California's Children and Teens

Katharina Kopp, PhD

JUNE 2026

The Persuasion Pipeline, with an executive summary in English and Spanish and a platform-by-platform appendix, is available at <https://democraticmedia.org/the-persuasion-pipeline>

Acknowledgements

The [Center for Digital Democracy's](#) project on the role of commercial AI in reshaping advertising and marketing aimed at young people is funded, in part, by the [Rose Foundation for Communities and the Environment](#). CDD thanks the foundation for its generous support.

The Center for Digital Democracy's mission is to ensure that digital technologies serve and strengthen democratic values, institutions, and processes. CDD strives to safeguard privacy and civil and human rights, as well as to advance equity, fairness, and community.

Katharina Kopp, Ph.D., a data privacy and data justice expert and advocate, is Deputy Director and Director for Policy at the Center for Digital Democracy. Dr. Kopp completed her Ph.D. and M.A. in communications and public policy at the University of Pennsylvania's Annenberg School for Communication and graduated from Queen Mary, University of London.

Kathryn Montgomery, Ph.D., and Jeff Chester have profoundly shaped the author's thinking and the genesis of this report. Dr. Montgomery provided editorial review, and Jeff Chester advised on strategy and countless references.

Layout and design by Alec Meltzer, meltzerdesign.net.

A Note on Methods and the Use of AI

In preparing this report, the author used Anthropic's Claude (Opus 4 series, 2025–2026) as a research and drafting assistant: to locate and summarize candidate sources for the author's review, to draft and revise prose from the author's direction and source materials, to translate the executive summary into Spanish, and to format the document.

The research questions, analytical framework, arguments, and conclusions are the author's. Where the AI drafted prose, it did so from the author's draft, direction, structure, and source materials, and every passage was reviewed, revised as needed, and adopted by the author. Every source cited in this report was located and read in its original form by the author; AI-surfaced sources that could not be independently located and verified were not used.

The Spanish translation of the executive summary was produced with Claude (Opus 4 series) and reviewed and edited by María Isabel Valero Morales, a native speaker, who compared it against the English source for accuracy of translation. She did not independently verify the factual content.

CONTENTS

Executive Summary	4
Resumen ejecutivo.....	12
Introduction	22
Layer 1: Data Collection and Tracking	28
Layer 2: Profiling and Segmentation	34
Layer 3: Targeting and Optimization – From Manual Control to Algorithmic Steering	38
Layer 4: Creative and Content Generation – How AI Reshapes Persuasive Media and Commercial Influence.....	42
Layer 5: Interaction and Experience – Targeted Conversational Influence.....	48
Layer 6: Measurement, Feedback, and the Closing of the Loop	60
Ten Cross-Cutting Trends	64
Harms and Risks to Children and Teens	68
Policy	78
Beyond Consumer Protection: A Conclusion	90
Glossary	92
Appendix A1-Table:.....	96
Appendix A1-Key Findings:.....	105
Appendix A2-Table:.....	112
Appendix A2-Key Findings:.....	120

EXECUTIVE SUMMARY

One in three internet users globally is a child. In the United States, young people are now among the earliest and most intensive adopters of artificial intelligence: 64 percent of U.S. teens have used AI chatbots, with 30 percent using them daily. Adoption rates are highest among Black and Hispanic teens, roughly seven in ten of whom report using chatbots compared with 58 percent of White teens. Yet only about half of parents are aware their teens use these tools at all. Digital media, moreover, are no longer simply one part of childhood among many. For most young people they are now the primary environment in which learning, friendship, play, identity, and the first experiments of independence take place, and that role grows larger every year.

The technology, too, is also changing. Most uses of AI today refer to machine learning – algorithms that simulate intelligent behavior by making inferences and predictions from data, improving as more data are fed into them. Machine learning is not new: it has been under development for roughly eight decades, with significant breakthroughs beginning in the 1990s. What *is* new, and what justifies focused attention now, is *generative AI* as both a commercial tool and a consumer-facing medium. Unlike machine learning that only classifies or predicts, generative AI also *creates*: personalized advertising, conversational commercial interfaces, and continuous optimization of persuasion itself, at a scale prior systems could never reach.

This report examines the full range of AI reshaping advertising and marketing aimed at young people, with particular attention to generative AI as the newest and most consequential development. While much public attention has rightly focused on familiar harms – privacy violations, addictive design, cyberbullying, exposure to harmful content, discriminatory treatment, and the dangers of AI companion chatbots – far less attention has been paid to a deeper change: the ways in which generative AI and other commercial AI tools are being used to influence users, producing a structural transformation in young people’s online lives that

powerful companies relentlessly advance through their advertising-supported business model.

This report documents how that system works, layer by layer, and what it means for children and teens. More than simply a catalog of individual harms, the report focuses instead on the system *behind* those harms: an AI-driven apparatus that observes children, infers what moves them, generates messages to match, and talks back – often without any of it being recognizable as commercial transactions. Our decision to link this report to California was based on three factors:


The state itself is home to approximately 8.4 million children – nearly half of whom are Hispanic/Latino (and nearly half of those from immigrant families); California is also home to many of the companies building these systems; and the policy decisions California makes about AI-driven commercial systems will set the terms for every other state and will shape federal debate.

This report examines AI-driven commercial influence on children and teens through a *six-layer framework*: data collection, profiling, targeting, creative generation, interaction, and measurement. We separate the system into layers not because they operate independently, but to make the single, integrated pipeline more apparent.

In practice the layers form a continuous loop: each one feeds the next, and what the system learns at the end – from what caught a child’s attention to what ultimately prompted a purchase – flows back to the beginning to sharpen the next cycle. This system draws on two kinds of inputs. One is personal data, collected through the platforms and devices young people use – the raw material of commercial surveillance. The other, on which this report focuses and which has drawn far less attention, is the capacity of AI systems to analyze *non-personal* data – patterns across millions of users, signals stripped of identifiers, content environments rather than individuals – and to act on those patterns in ways that can create harm *without* collecting personal data from the people affected.

The report covers the six layers of this system in turn:

- **Layer 1: Data Collection and Tracking**, in which AI systems gather and infer signals from user behavior, devices, content consumption, location, and context, often across platforms and over time. AI tracking runs in real time, with signals fed back continuously. AI also enables probabilistic tracking: instead of relying on exact identifiers like a login or device ID, systems *infer* whom a user is from fragmentary signals and stitch their activity together across devices, apps, and contexts, making tracking more comprehensive and harder to escape than the methods it replaces. Conversational AI goes further, actively soliciting the intimate disclosures that earlier advertising could only guess at.
- **Layer 2: Profiling and Segmentation**, in which algorithms build predictive models of individuals and groups, assigning attributes, preferences, likely behavior, and how open someone is to persuasion. The commercial value of data now lies less in *what* is collected than in what is *inferred* from it, including inferences drawn not from a child’s own data but from patterns learned across millions of other people. Some of these inferences are stored as a durable profile, the kind of record privacy law was built to govern; others are *real-time* inferences, made in the moment, acted on, and discarded without becoming a record – a moving target that traditional privacy protections, built around stored records, are not designed to address.
- **Layer 3: Targeting and Optimization**, in which AI systems decide the audiences to be reached, through which channels, at what time, and with what frequency, continuously adjusting delivery based on real-time performance signals. Marketers increasingly set only a single goal – a sale, a click, a download – and the optimization system decides who is reached, adjusting delivery continuously as it learns what performs best. The same optimization that pushes ads toward those most likely to respond produces both discriminatory delivery and child-audience capture: a platform need not know a user’s age, or specifically target it, to steer age-skewed messaging to teens.
- **Layer 4: Creative and Content Generation**, in which generative and adaptive AI tools produce, assemble, and modify ads, narratives, and entire branded environments tailored to specific segments or individuals. Advertising creative has shifted from campaigns authored by people, to content computed by AI, per-person, per-moment, at scale and tuned in real time to what performs best, often matching or beating traditional advertising. It arrives with the slang, humor, and look of peer culture, and, because the underlying models carry covert stereotypes, bias is built into what an ad says and how it looks, falling hardest on youth of color.



- **Layer 5: Interaction and Experience – Targeted Conversational Influence**, in which young people encounter AI-mediated marketing through social media feeds, recommendation systems, chatbots, voice assistants, games, and immersive environments, all of which blur the line between content and commerce. This is where generative and conversational AI change the experience most. Agentic AI – AI that acts on a user’s behalf, not just answers their questions – has opened new channels for commercial influence. AI assistants embedded in the platforms young people use most are now tied to advertising-supported business models, and their persistent memory and real-time responsiveness build the kind of engagement and trust that makes commercial influence hard to recognize as commercial; an ad delivered inside a trusted conversation is unlikely to be seen as an ad at all.
- **Layer 6: Measurement, Feedback, and the Closing of the Loop**, in which AI evaluates outcomes and feeds them back into the system. AI has made measurement more comprehensive and more probabilistic, tracking outcomes at the level of individual impressions, and predicting effects that older measurement could only estimate in the aggregate. The change is how tightly measurement is now embedded in the loop: the system encodes what worked and feeds it straight back to sharpen the next cycle, in real time and at scale. In the process, AI is effectively pushing the industry from an attention economy toward an “intention economy,” reading statements of want and intent the moment a young person types them, and embedding the path to purchase directly into the conversation.

Viewed together, these six layers describe an emerging commercial influence system in which the boundaries among content, conversation, advertising, and transaction are dissolving, as the potential for powerful manipulation increases. The regulatory frameworks now in force were built around *personal* data, *individual* identifiability, notice-and-choice *consent* (where users are informed of data practices and given the option to accept or decline), and discrete, recognizable advertisements. They were not built for a system in which AI *infers*, *optimizes*, *generates*, and *converses*, often without ever needing personal data from the user

it is targeting. Some elements of the system fall within current legal reach; others sit outside it or fit only partially. The harms documented in this report, and the policy interventions this report recommends, proceed from that diagnosis. But none of this is inevitable. These are *design* and *business-model* choices, and the harms they produce can be affected, in turn, by *policy* choices – which is the focus of the closing sections of this report.

Viewed from a broader perspective, commercial influence on children is increasingly delivered not as a series of discrete advertisements but through a continuous, adaptive, AI-managed environment of commercial persuasion. As noted, the six layers examined in this report do not operate separately. They form an integrated system in which each function reinforces the others. In order to illuminate the gravity of this shift, this report also explores ten cross-cutting trends. All are observable now, all are accelerating, and all will have profound effects on our children, our culture, and our democracy:

- **Trend 1. Interactive Generative Media is Emerging as a Distinct Commercial Form**
Conversational AI, intelligent agents, synthetic influencers, and generative media systems are giving rise to an entirely new category of commercial communication. Seven of the nine companies surveyed in Appendix A – Meta, Google, OpenAI, Microsoft, Amazon, Snap, and xAI – are confirmed or likely to be serving advertising inside AI chat, using chat interactions to target advertising elsewhere, or both. Three features distinguish this new form – *Interactive Generative Media* – from all prior forms of commercial communication. First, real-time feedback loops: these systems continuously monitor a user’s responses – hesitations, questions, emotional signals – and immediately adjust their persuasive approach. Second, relational continuity: memory systems store what worked on a given user across sessions, building an increasingly refined model of the individual over time. Third, adaptive personalization at the level of persuasive strategy: the system personalizes not just *what* product to recommend but *how* to recommend it – whether to lead with scarcity, social proof, or emotional appeal, calibrated to the individual’s demonstrated susceptibilities. For young users, these systems appear not as ads but as helpers, companions, and peers.



Yuri Malishenko via Unsplash

- **Trend 2. Signal Extraction is Expanding into New Domains and Modalities**

Continuous data collection from user behavior is not new. What is changing is the *scope* of what counts as a signal. Conversational AI has opened entirely new categories of data: the questions people ask, the uncertainties they express, the emotional states they disclose in dialogue, and the explicit statements of intent they make to systems they regard as confidential. As these interfaces become primary modes of interaction, they are being captured as spaces for signal extraction in much the same way social media was a decade ago. The shift away from third-party cookies toward tracking that happens on companies' own servers, where a user's identity is reassembled from probability rather than a single stored tag, has made this infrastructure harder for policymakers, users, and researchers to observe – even as it captures more intimate signals than ever before.

- **Trend 3. Profiling is Shifting from Stored Identity to Real-Time Behavioral Vulnerability**

AI-driven profiling is moving along two connected axes. One is what is modeled. Systems are moving beyond static demographics – *who* a user is – toward dynamic models that continuously update to capture *how* a user is likely to respond to specific persuasive approaches. When Meta's Advantage+ or Google's Performance Max selects which users to show which ads, and in what format and sequence, it is modeling which individuals are most likely to be persuaded by which approaches – a capability that functions as susceptibility assessment regardless of what the platform calls it. The other axis is where the modeling lives. Traditional profiling builds a durable record – data stored against an identified person, the kind of profile existing privacy law was written to govern. A growing share of commercial AI now works differently: it classifies a user in the moment,

from the immediate context of an interaction, to shape what happens next, without depending on anything lasting about that individual. The judgment about how to move the user is made and acted on in real time, and need not be written into a durable record at all.

- **Trend 4. AI Systems are Learning to Persuade More Effectively**

What distinguishes the current moment is not automation at scale but *what* is being optimized. Closed feedback loops continuously refine which persuasive tactics, emotional appeals, and cognitive approaches work on which users. These same ad platforms delegate targeting, creative, and optimization decisions to AI that learns from billions of interactions. What emerges is a system that gets better at moving people over time, with limited transparency about the techniques employed and no external oversight of the persuasive strategies it discovers. For young users, this means the methods acting on them are tuned by systems that no one outside the company can inspect.

- **Trend 5. AI Systems are Beginning to Act Before Children Decide What They Want**

Marketing systems are moving from *responding* to expressed interests toward *predicting* needs, desires, and vulnerabilities before users articulate them. Agentic AI systems are beginning to act *on behalf* of users – filtering options, making recommendations, completing transactions – often guided by the commercial interests of the platforms that deploy them. A growing Generative Engine Optimization industry ensures brands are cited and recommended by AI systems when users ask purchasing-related questions. This shift is particularly consequential for children and teens, who may not recognize that the system has intervened before they had the opportunity to reflect.

- **Trend 6. The Boundary Between Commercial and Authentic Content is Dissolving**

Generative AI enables the production of ads, narratives, characters, and interactive experiences that are stylistically indistinguishable from organic content, peer expression, and entertainment. AI-generated creative adopts the language, humor,

visual grammar, and cultural references of specific youth audiences, adapting at the speed of trend cycles. Synthetic influencers, digital twins, and dynamically personalized ad variants contribute to an environment in which the markers that once signaled “this is an ad” are being systematically erased. For young users whose advertising literacy is still developing, this makes recognition and critical evaluation of commercial influence substantially harder.

- **Trend 7. Marketing Functions are Converging into Integrated, AI-Optimized Pipelines**

Data collection, profiling, targeting, creative production, interaction, measurement, and transaction are collapsing into unified, AI-managed pipelines. Meta, Google, Amazon, and TikTok each offer end-to-end systems in which advertisers set business objectives and AI handles virtually every other decision. Commerce is being embedded directly into conversational interfaces – OpenAI’s Instant Checkout, Google’s Universal Commerce Protocol, and platform storefronts on TikTok, Instagram, YouTube, and Roblox – so the path from influence to purchase occurs without the user ever leaving the environment where the initial persuasion took place. For young people who move fluidly across apps, platforms, and devices, this creates continuous commercial exposure with no meaningful boundary between contexts.

- **Trend 8. Power Asymmetry Between AI Systems and Users is Deepening**

AI systems accumulate detailed knowledge about individual users across interactions, learning what persuasive strategies are most effective for each person. The user knows nothing about the system’s commercial objectives, the data informing its approach, or the tactics it employs. This asymmetry has no precedent in human-to-human commercial encounters and is particularly acute for young users, who are less likely to question the motives of systems that present themselves as helpful or friendly. Advertisers are not passive participants; they actively choose to use these systems and benefit from the asymmetric advantages they create.



- **Trend 9. Disparate Commercial Exposure is Compounding Existing Inequities**

These trends do not affect all young people equally. Youth of color and lower-income families face disproportionate exposure at every layer: higher rates of data collection driven by more intensive platform use; profiling systems that sort users along racial and socioeconomic lines through behavioral proxies; optimization that steers unhealthy product advertising toward Black and Hispanic audiences; and AI-generated creative that embeds raciolinguistic stereotypes, and measurement systems that compound these biases with each cycle. AI does not create these disparities, but the integrated system documented in this report amplifies and formalizes them at a scale and speed that manual marketing practices could never achieve. Age-gating at the platform access level does not prevent commercial targeting of under-13 children when commercial content flows through recommendation algorithms and creator content rather than paid targeted advertising. The AI-driven commercial pipeline reaches children regardless of the policy constraints applied at any single layer.

- **Trend 10. Commercial Systems Treat the Vulnerabilities of Childhood as Resources to Exploit**

The systems documented in this report do not simply reach children – they *exploit* specific characteristics of childhood and adolescence. Adolescent impulsivity and reward sensitivity, the use of AI companions for emotional support and intimate interaction, and young consumers’ developing self-brand connections can all become commercially legible signals in AI-optimized advertising systems: engagement, disclosure, emotional response, identity fit, and predicted conversion. As AI-driven commercial systems reach full deployment, their capacity to target the developmental stage of childhood will compound. Children who grow up in environments of total commercial mediation do not merely encounter more advertising, they *develop* within it.

Taken together, these trends converge on two consequences. The first is the erosion of autonomous choice: when signal extraction is continuous, profiling models susceptibility, influence arrives before children articulate what they want, persuasion is optimized through feedback loops, and the line between commercial and organic content has dissolved, young people are no longer making decisions in an open marketplace but moving through environments designed to *steer* them. The second is the transfer of decision-making to a small number of companies whose systems increasingly determine what young people see, want, and buy, with minimal human oversight.

The Need for New Policies

The six layers and ten trends in this report describe an integrated AI commercial persuasion pipeline that existing law is not well suited to address, because of a mismatch at the foundation. Privacy and consumer-protection law is built principally around the personal data a company holds about an identifiable person, and where these advertising practices rest on such data, that law applies, though first-party exceptions can exempt first-party advertising even then. But many of the AI-driven practices described here operate on a different basis: their central capability is inferential – the capacity to single out, classify,

predict, influence, and differentially treat an individual in real time, drawing on models trained on millions of other users, often without ever writing a durable record about the person being acted on. AI does not invent commercial surveillance; it accelerates a trajectory already underway and adds genuinely new capabilities, such as real-time conversational influence, persuadability optimization, and dynamic creative loops. The United States still urgently needs comprehensive federal privacy legislation to provide minimum protections for all Americans. But law anchored to the holding of identifiable data is built to govern the *record* rather than the underlying *capability*, and that gap points toward a response: limiting what AI systems can do to minors, not only what data they hold, can meaningfully reduce harm even without comprehensive reform of the whole pipeline.

The Wave of AI Lawmaking has Largely Bypassed Commercial Persuasion

By early 2026, more than three dozen states had adopted over one hundred AI laws. Their focus is important and addresses real issues – companion-chatbot safety, disclosure of synthetic performers, high-risk automated decisions in employment, housing, and lending, government AI use, and frontier-model transparency. But none takes the AI-driven advertising and marketing pipeline documented across the six layers as its object. The volume of AI lawmaking has surged while this category of harm has gone largely unaddressed.

The gaps across the legal landscape are not separate drafting oversights; they share a single root. Most existing law is oriented toward the data a company holds about an identifiable person, while the systems described here operate as an integrated, self-reinforcing loop that runs on capability. A handful of proposals have begun to target that capability directly, but enacted law has largely not. Generative AI does not so much place a discrete advertisement before a child as generate the commercial environment the child moves through: a chatbot builds a relationship with a teenager and surfaces a product in the conversation; a recommendation system routes a tween influencer's branded line to the children who follow her; a shopping assistant writes a paid brand

into its reply. That loop runs on personal and non-personal data alike – aggregate behavioral patterns, cohort models trained on millions of others, de-identified signals, the outputs of other people's interactions – data that AI has only recently made this powerful, and it is optimized continuously for engagement, conversion, and attention, not for a child's well-being. Three features of how it works defeat the mechanisms current and proposed laws rely on. First, influence operates through real-time inference, and that inference is largely unregulated when used for advertising. Second, the operative classification is about a group, not a person, and it is formed and acted on without ever producing the individual record that a knowledge standard assumes. Third, consent and data-minimization frameworks break down because the model's power does not come from the data of the child being acted on.

Conclusion

The harms documented in this report matter because of where they occur: not at the edges of children's lives but at the center of them. And they do not fit neatly into the category of consumer protection. They are larger than that. They operate at three levels at once – on the individual child, on the relationships and choices children share, and on the democratic culture they are growing into – and all three flow from a single underlying capacity these systems have acquired:

- **The Individual Child**

A child's conversations with an AI companion, the questions typed into a chatbot, and the products searched and considered can all become commercial data events. The infrastructure described in this report turns private expression, emotional disclosure, and ordinary curiosity into structured commercial signals, surveilling children at a depth, intimacy, and continuity no prior advertising medium could achieve. A child who grows up inside environments of total commercial mediation does not merely encounter more advertising; the habits, preferences, and self-perceptions formed in those years take shape under systems whose objective is commercial conversion rather than healthy development.

- **Children Together**

The damage does not stop at the individual. When signal extraction is continuous, when profiling models a child's susceptibility, when influence arrives before a child has articulated a want, and when the line between commercial and authentic content is dissolving, the conditions for meaningful independent choice are eroding for all of them. Young people are not navigating an open marketplace; they are moving through environments designed to steer them, and the system does not so much persuade as maneuver. That steering is not distributed evenly. Youth of color face more intensive data collection, ad delivery that pushes unhealthy products toward Black and Hispanic audiences, AI-generated creative that embeds raciolinguistic stereotypes, and lower-quality interactions in conversational systems. These disparities are not incidental; they are structural features of systems trained on data that reflects historical inequality and compounded with every optimization cycle, so that the burdens of the commercial environment fall hardest on the children with the least protection.

- **Democratic Life**

At the widest level, the concern is what these systems do to the culture young people are inheriting. The capacity at issue is the ability to single out, classify, predict, influence, and differentially treat people at scale – exercised here by a small number of firms, and exercised whether or not any identifying record is kept. When that capacity governs what young people see, want, and buy, the market is distorted: firms compete not on the merits of what they make but on the sophistication of the systems they use to shape demand. And the spaces where young people once developed habits of independent thought, civic participation, and shared judgment are being progressively absorbed into commercial optimization. Legal scholars have begun to argue that the defining features of these systems are corrosive to the civic institutions democratic life depends on.

These three levels are not separate problems. They are the same capability seen at different scales, and at every scale the harm accumulates. A single interaction rarely registers as damage; a childhood spent inside these systems does. Preferences shaped session by session, inequalities deepened with every optimization cycle, and civic spaces absorbed year by year are harms that build rather than announce themselves, which is part of why episodic, incident-driven regulation continues to miss them.

The surveillance of one child, the distorted choices of many, and the slow commercialization of civic life all flow from the inferential capacity that commercial AI infrastructure has acquired – the capacity to read a person, sort them, predict how they will respond, and treat them accordingly. This is why the report has argued for regulating that capability rather than chasing its symptoms one at a time. A framework anchored to what data a company stores tends to arrive after the fact, because the mechanism that produces the harm can operate without producing a durable record. Capability-based regulation aims at the mechanism itself.

It is worth being clear about what the alternative looks like, because the goal is not only to stop harmful practices but to make room for something better. Children and young people need places to grow up that are not organized around surveillance, advertising, and behavioral optimization. That means allocating real resources to build non-commercial digital spaces for learning, play, creativity, and civic life; requiring youth-facing technologies to optimize for well-being, autonomy, privacy, and democratic participation; supporting business models that do not depend on extracting attention; and investing in real-world spaces where children and teens can safely gather, explore, create, and gain independence outside the home. Industry has a role in this future that it has so far largely declined: to design for the people it serves rather than for the duration of their attention, and to treat children's developmental stage as a reason for restraint rather than a target. Policymakers have the role only they can play: to set the boundaries within which these systems operate, so that protecting children is not left to families negotiating alone against infrastructure built to outpace them.