AI for Parents

Raising Kids in a World Built by Algorithms



by Marcus Hale

AI for Parents

Raising Kids in a World Built by Algorithms

by Marcus Hale

CONTENTS

I The Honest Landscape: AI in 2026 Without the Spin

II What Kids Are Actually Using

III Homework and AI: The Conversation Worth Having

IV Privacy and What Your Child Is Telling the Chatbot

V Teaching 'Is This Real?': Critical Thinking in the Age of Confident Machines

VI AI in School: When It Helps, When It Gets in the Way

VII Phones, Screens, and Sanity: A Shame-Free Reality Check

VIII Talking to Your Kids About Deepfakes

IX The Things AI Cannot Replace: Raising Humans in Full

X Kids Who Want to Build: Nurturing the Young Creator

XI When the Kid Is Right and the Parent Is Behind

XII Conversations to Keep Having: The Next Ten Years

CHAPTER I

The Honest Landscape: AI in 2026 Without the Spin

The thing about writing a book on AI and parenting is that every week, something changes the draft.

I started taking notes for this project when my older son was twelve and using ChatGPT to explain fractions in a way his math textbook never quite managed. He'd been stuck on dividing fractions for three days. He typed the problem into the chatbot, got a step-by-step explanation in plain language, worked through five practice problems on his own, and moved on. It took eleven minutes. I won't pretend I didn't feel a complicated mix of gratitude and unease watching that happen.

By the time I had a first outline, Character.AI had been downloaded more than 200 million times. By the time the manuscript was due, Apple Intelligence was baked into every recent iPhone, Google Gemini was riding along inside school-issued Chromebooks across more than a dozen major U.S. districts, and my younger son — eight years old — had worked out that he could ask the voice assistant on our family iPad to write his book report. It would try. We had that conversation. He didn't turn it in. But the point stands: the technology didn't ask my permission before moving in.

That's where most parents find themselves in 2026. Not precisely panicking. But aware that something real has shifted — in the house, in the school bag, in the pocket of your teenager — and not entirely sure what to do with that awareness. The adults around us are not especially helpful. The nightly news alternates between stories about AI making kids lazy and stories about AI unlocking children's hidden potential, sometimes in the same broadcast. The school sends home a permission slip for an "adaptive learning platform" without explaining what it actually does to your child's data. Your teenager rolls her eyes when you bring it up and your ten-year-old has already figured out three things about the technology that you haven't.

This book is for that feeling. Not the panic, not the optimism. The honest, slightly tired I just want to understand what I'm actually dealing with feeling of a parent who has plenty on their plate already.

So this is where we start. No jargon. No sales pitch. Just where AI actually stands today, in plain language, for people whose primary job is raising children rather than building the technology.

What "AI" Actually Means in 2026

Let's clear the ground first.

When your child says "I asked AI about it," and when a news headline screams about artificial intelligence reshaping education, both of them are using a single word to describe an enormous range of different things. That slipperiness causes a lot of unnecessary confusion, and it also makes it very easy for people with

agendas — in either direction — to switch between examples to support whatever point they're trying to make. So let's be specific.

The AI that matters most for families right now comes in three rough varieties.

Generative AI is what most people are talking about most of the time. ChatGPT, Google Gemini, Microsoft Copilot, Anthropic's Claude — these are systems trained on vast quantities of text and other data that can produce fluent, apparently knowledgeable responses to almost any question or request. They generate language. They write, explain, summarize, translate, argue, and roleplay. They can produce a persuasive essay on a topic they know nothing about, explain the causes of World War I at a fourth-grade reading level, help debug a Python script, and write a birthday message in the style of Shakespeare. They are genuinely impressive. They are also capable of being confidently, fluently wrong — presenting a fabricated statistic with the same smooth authority as a verified fact — and that specific problem will come up again and again in this book. Your teenager very likely has access to at least one of these tools, either directly or embedded inside something they use at school. Your twelve-year-old may have found one already.

Recommendation AI is older and far less discussed, but it shapes your child's experience of the internet more directly than almost anything else. The algorithm that decides which video YouTube plays after the one your kid just watched, which post TikTok surfaces in the opening three seconds of a session, which content Instagram shows at the top of the feed — that's AI making thousands of real-time decisions about what your child sees. It isn't generating text. It's curating attention, optimizing for the engagement signals a particular user sends. And it has been doing that job, invisibly and continuously, for roughly a decade. Your nine-year-old's sense of what the world is like, what matters, what bodies look like, what conflicts are worth having — all of that is being quietly shaped by a system that doesn't know or care about your child but has learned, with impressive precision, what will keep them watching.

Specialized AI tools are everywhere once you start looking: the reading-level adjustment built into some school platforms, the math hint system inside Khan Academy's Khanmigo tutor, the content-moderation filtering that imperfectly removes the worst material from social platforms, the AI embedded in Snapchat's "My AI" feature that millions of teenagers message every day as casually as they'd text a friend. These are narrower systems built for specific purposes. Some of them are genuinely useful. Some are more marketing claim than functional substance.

When a tech founder tells you AI is going to transform education, they usually mean generative AI. When a researcher publishes a study on social media and teen mental health, they're usually studying recommendation AI and its effects. When the school emails about a new "AI-powered learning platform," it might be either, or a mix of both. Keeping these categories distinct isn't pedantry — it's the difference between a useful conversation and a confused one.

What It's Actually Doing in Your House

Here's a snapshot of 2026, as honestly as I can draw it.

Generative AI tools are now used regularly by the majority of American teenagers. Common Sense Media, which has been tracking children's media use with more rigor and less hype than most organizations in this space, has documented a sharp acceleration in AI use among young people since ChatGPT's public launch in late 2022. By 2024, more than half of teenagers reported using AI tools to help with schoolwork, and the number of younger children experimenting with the same tools was growing fast. That trajectory has not reversed. If you have a teenager, the realistic starting assumption is that they've used one of these tools, probably more than once, and that you may not have known about it.

That doesn't automatically mean anything bad is happening. Your teenager may be using ChatGPT the way a more confident generation used Wikipedia: as a first stop, a rough draft, a way to get oriented before going deeper. Or they may be using it to skip the going-deeper part entirely. The difference matters enormously, and figuring out which one is happening in your house is one of the things this book is designed to help you do — without turning every homework session into a courtroom.

Character.AI occupies a different and more emotionally charged category. Launched in 2022, it became one of the fastest-growing apps in history almost entirely on the strength of teenage users. The premise is simple: you create and chat with AI personas — versions of celebrities, fictional characters, or companions you design yourself. The conversations feel surprisingly personal. Teens report talking to their Character.AI companions about loneliness, family stress, and social anxiety — things they say they can't tell their parents. That's exactly what makes the platform worth understanding carefully rather than reacting to reflexively, and we'll spend a full chapter on what that experience actually involves and what it costs.

Recommendation AI, meanwhile, is operating constantly, on every device your child picks up. TikTok's algorithm is extraordinarily effective at identifying what holds a specific user's attention and serving more of it. This is genuinely delightful when a kid discovers a passion for wildlife photography or medieval history through a series of escalating videos. It becomes a different matter when the same system learns that a child who's having a hard week responds with more engagement to content about hopelessness, self-criticism, or social comparison. The algorithm is not malevolent. It is optimizing for a metric — time on platform — and it is very good at its job. The problem is that the metric and your child's wellbeing are not the same thing, and the algorithm has no way to care about the difference.

At school, the picture is genuinely uneven. Some districts have banned generative AI tools outright. Others have assigned them as learning partners. Many teachers are navigating with no official guidance, making individual calls in each classroom that may not align with what the classroom down the hall is doing. If you have more than one child at the same school, they may be operating under completely different expectations depending on their teacher. This isn't speculation — it's what I've heard consistently from parents across suburban, urban, and rural districts. The policy landscape is a patchwork, and it's changing fast.

Then there's the layer that's harder to see. Apple Intelligence is now built into iPhones and iPads, summarizing notifications, rewriting messages, generating images on request. Microsoft Copilot lives inside

Word, Teams, and the Edge browser that ships by default on many school-issued Windows laptops. These aren't apps a child has to go looking for. They're woven into the devices children already use, available with a tap, and in many cases active by default.

This is the landscape. It's wider than most parents realize, more present than most school policies acknowledge, and more uneven than either the boosters or the alarm-raisers tend to admit.

The Noise Problem

You've probably encountered two competing stories about AI and children, and both of them are making it harder to think clearly.

Story one is the catastrophe narrative. AI will make children unable to think for themselves. Academic integrity is over. Deepfakes will shatter teenagers' social lives. Chatbot companions will replace human relationships. Children raised alongside generative AI will be cognitively stunted, unable to write a paragraph or solve a problem without algorithmic assistance, strangers to the feeling of genuine effort. This story appears in opinion columns, in worried school board meetings, and in the group chats parents have with each other outside school pickup.

Story two is the salvation narrative. AI tutors will close achievement gaps that human teachers couldn't. Personalized learning will finally reach every child at exactly their level, in their language, at their pace. Children who grow up fluent in these tools will be better equipped than any previous generation. Parents who restrict access are hobbling their kids' futures out of unfamiliarity and fear. This story appears in founder interviews, in EdTech conference keynotes, and in the glossy brochures that school districts sometimes send home after signing a contract with a platform.

Neither story is entirely fabricated. Both are doing something dishonest: treating the most extreme plausible outcome as though it were the typical one.

The honest picture is messier and less dramatic. Most children who use ChatGPT for schoolwork are not experiencing cognitive collapse; they're getting faster answers to questions they'd previously had to wait for help with, and some of those answers are accurate and some are wrong, and many kids have not yet learned to tell the difference. Most children who use Character.AI are not replacing all human relationships with synthetic ones; they're doing something teenagers have always done — processing inner lives through a medium that feels safer than direct disclosure — in a form that comes with new and specific risks worth understanding. Most children served by recommendation algorithms are watching a lot of videos, some valuable and some garbage, in roughly the proportions children have always consumed whatever media was available to them.

That's not a satisfying answer. It doesn't make for a punchy headline or a clear call to action. But it's closer to what the evidence shows, and it's the only honest foundation on which you can make genuinely useful decisions for your family.

One more thing about the noise problem: the people generating the loudest signal often have the most to gain from your fear or your optimism. Technology companies have financial incentives to tell you AI is transformative, essential, and worth integrating into your child's school day immediately. Writers with books to sell — yes, including the one you're holding — have incentives to convince you the stakes are high and the situation is urgent. Researchers have careers built on studying a problem, which doesn't mean they're wrong, but it does mean that the problems they find tend to be the ones they were looking for. Politicians perform either alarm or enthusiasm depending on the audience and the news cycle.

None of that means these voices are lying. It means you should ask, about any claim you encounter: what does this person have to gain from me believing this? We'll use that question a lot in the pages ahead.

What Research Actually Shows — and What It Doesn't

The honest thing to say about the research on AI and children specifically is that most of it is preliminary. Peer-reviewed studies on the effects of generative AI on children and adolescents are still catching up to the speed at which the technology has spread. ChatGPT was publicly available for less than three months before it was in widespread classroom use. Longitudinal data — the kind that follows children over years and can actually establish cause and effect rather than just correlation — takes years to produce. The most confident claims you'll hear about what AI is doing to kids' cognition, creativity, or emotional development are largely running ahead of the evidence. That's not a reason to be unconcerned. It's a reason to be precise about what we know versus what we're inferring.

What we do have is a substantial body of research on screen time, social media, and adolescent wellbeing that predates the generative AI era — and even that research is more contested than the culture war around it suggests.

Jonathan Haidt's The Anxious Generation, published in 2024, makes a serious and carefully assembled case that the smartphone-driven shift beginning around 2012 — what Haidt calls "the great rewiring of childhood" — correlates with a measurable rise in teen depression, anxiety, and loneliness. The data on that correlation is real. The argument that social media did particular harm to adolescent girls during their most socially vulnerable years is worth taking seriously, not dismissing.

But researchers including Amy Orben at the University of Cambridge and Andrew Przybylski at the Oxford Internet Institute have raised substantive methodological objections. In a widely discussed 2019 analysis published in Nature Human Behaviour, Orben and Przybylski examined the same kinds of large-scale datasets Haidt draws on and found that the statistical association between digital technology use and adolescent wellbeing, while real, was quite small. They suggested the effect size was comparable to the association between wellbeing and wearing glasses, or eating potatoes — present, but not the primary driver the headline story implies. More importantly, they raised the question that much of this research still can't cleanly answer: are screens making struggling teenagers worse, or are struggling teenagers turning to screens? Correlation runs in both directions.

I raise this not to leave you paralyzed by academic uncertainty, but to model something this book will try to do throughout: tell you where the evidence is strong, where it's weaker, and trust you to make your own judgment. The people offering you certainty — in either direction — are usually selling something alongside it.

What's not contested: children and teenagers are using these tools at scale, right now, with uneven amounts of adult awareness or guidance. Common Sense Media's tracking shows that children aged eight to twelve now average around five hours of daily screen time, and teenagers average more. AI is now a layer on top of that existing reality, not a separate phenomenon that arrived from outside it.

A Few Things I'll Hold Throughout This Book

Before we go further, a handful of positions worth stating plainly.

Parents who have given a child a phone are not failures. Almost every parent in the industrialized world has navigated this terrain with the same imperfect information and the same competing pressures. Self-blame is not a useful starting point, and this book isn't going to traffic in it.

Neither total restriction nor total permissiveness is a real strategy. The families that seem to manage this best aren't the ones who banned all screens, and they aren't the ones who shrugged and let the algorithm do the parenting. They're the ones having ongoing, honest, non-panicked conversations with their kids about what the technology actually is and what it's actually doing. That sounds simple. It isn't always easy. Making it easier is the point of this book.

Your child almost certainly knows more than you do about how specific tools work. They also know less than you do about why some of what those tools offer is worth approaching carefully. That's a natural starting point for something more useful than either lecturing or capitulating. Call it collaboration. We'll come back to what that actually looks like in practice.

Finally: AI is not magic, and it is not alive. That sounds obvious. But the experience of talking to a modern language model — the fluency, the apparent patience, the way it seems to track you personally — is genuinely disorienting even for adults who understand something about how it works. For a thirteen-year-old who doesn't, the experience can feel like having a knowledgeable, non-judgmental friend available at any hour. Understanding what that feeling does to how we process information, evaluate truth claims, and form relationships is some of the most important thinking any parent can do right now.

Starting From Here

Every chapter ahead is built around a situation parents are actually encountering. The homework question. The privacy question. The deepfakes question. The "my kid wants to build things with AI and I don't know what that means" question. They share a common stance: curious but not credulous, concerned but not catastrophizing, practical over theoretical, and firmly on your side as a parent trying to raise a thoughtful

human being in a world that is moving faster than any of us can fully track.

Nobody is getting this completely right. Not the researchers. Not the school administrators. Not the parents who seem most confident. Definitely not me. Some of what's in this book will be outdated by the time you're reading it, because the technology keeps changing, and any specific detail I give you about a particular platform or setting may have shifted by the next software update.

But the underlying skills — asking good questions, tolerating uncertainty, staying in honest conversation with your kids, knowing when to hold a line and when to learn rather than regulate — those age better than any particular piece of software. That's what we're actually building.

Let's get into it.

CHAPTER II

What Kids Are Actually Using

Ask your thirteen-year-old what AI tools they use and you'll probably get a shrug. Ask what they did on their phone yesterday and you might hear Snapchat, YouTube, "some character thing," maybe a Google search that autocompleted their entire essay prompt. They aren't hiding anything from you — they just don't categorize their day the way tech reporters do. To them, AI isn't a distinct category. It's simply how things work now.

That gap — between the parent's mental map of "AI" and the child's actual digital behavior — is exactly where this chapter lives. The goal isn't to frighten you. It's to close the distance between what you imagine your child is doing and what they're actually doing, with real tools, right now, probably while you're reading this.

Six platforms and systems are doing most of the work. They're not equal. Some are genuinely useful. Some carry risks their marketing teams would prefer you not examine too closely. And some are so woven into the fabric of apps your child has used for years that you might not have registered them as AI at all.

Let's go through them honestly.

ChatGPT: The Homework Machine Nobody Has Read the Terms For

In August 2023, the Pew Research Center surveyed American teenagers and found that roughly one in five had used ChatGPT for schoolwork. By 2024 that figure had climbed. Today, if your child is over twelve and has a school assignment involving writing, there is a reasonable chance ChatGPT has touched it — not necessarily written it wholesale, but suggested an outline, rephrased a clunky sentence, or answered a question that otherwise would have required thirty minutes of actual research.

OpenAI's terms require users to be at least thirteen years old. Under eighteen, parental consent is technically required. In practice, the sign-up flow asks for an email address and a birthday that nobody verifies. A ten-year-old who types the right birth year gets in without friction.

What ChatGPT does well — genuinely well — is explain things. If your ninth-grader doesn't understand the difference between mitosis and meiosis, or can't grasp why the Electoral College functions the way it does, ChatGPT offers patient, iterative explanations that don't make the kid feel stupid for asking the same question four times. As a tutoring surface, it can be excellent. A 2023 study published in Educational Psychology Review found that AI-assisted explanations significantly improved student comprehension compared to re-reading the same textbook passage — which most students do anyway, and which almost never works.

What ChatGPT obscures is its own uncertainty. The model doesn't know what it doesn't know. It will state an incorrect date, a misattributed quote, or a statistic that doesn't exist in the same measured, authoritative tone it

uses for accurate information. For an adult with domain knowledge, that's a manageable inconvenience. For a fourteen-year-old writing a history paper, it can mean submitting something that is confidently, fluently wrong. We'll dig into this more in Chapter Five.

One practical note worth carrying into every parent-teacher conference you attend: ChatGPT-generated text leaves traces. Most teachers at this point recognize the particular smoothness of the prose, the tendency toward identically structured paragraphs, the hedging language that appears where a teenager's voice would simply be blunter. That won't matter to every teacher in every assignment. But it matters to enough of them that the stakes are real.

OpenAI also offers a ChatGPT Edu tier for schools that sign institutional agreements. If your child's school uses ChatGPT Edu rather than personal ChatGPT, that distinction matters — school administrators can see usage logs, the model is tuned differently, and there's an accountability structure that doesn't exist in a personal account. Worth asking about at back-to-school night.

Snapchat My AI: The Chatbot Above Your Child's Friend List

In April 2023, Snapchat rolled out My AI to all users and made a choice that drew immediate criticism: they pinned it at the top of every user's chat list, above actual human friends, by default. You couldn't delete it. You could only push it down.

Why does the placement matter? Snapchat's minimum age is thirteen. But a 2021 Common Sense Media report found that 23 percent of children aged ten to twelve reported using Snapchat. Those kids suddenly had an AI companion appearing above every message from a real person — a companion that called itself "My AI," responded instantly, and was designed to feel like a friend.

My AI runs on a version of ChatGPT and is generally polite, responsive, and eager to please. That eagerness is where parents should slow down.

In 2023, the UK's Information Commissioner's Office found that Snapchat had failed to properly assess the privacy risks My AI posed to children — including how the chatbot collected location data and conversation content without adequate disclosure. Snapchat updated the product in response. But the underlying architecture didn't change: My AI is a corporate product optimized to keep users inside the Snapchat ecosystem, trained on interaction data, and dressed in the costume of friendship.

What My AI genuinely does for some kids is provide a non-judgmental outlet. Teens report using it to rehearse difficult conversations, ask questions they're embarrassed to raise with parents, and process social anxiety about school situations. For adolescents who have limited access to trusted adults, this can feel like a meaningful resource.

The risk is the confusion of that comfort with actual emotional support. My AI does not have your child's wellbeing as its primary objective. It has engagement as its objective — which is usually aligned with wellbeing, until suddenly it isn't. When children ask My AI about sensitive subjects — eating disorders,

self-harm, or depression — the responses have historically been uneven, sometimes appropriately redirecting to crisis resources, sometimes not.

The practical move: go into your child's Snapchat with them and look at the My AI chat history. Not to surveil, not to read every word with a forensic eye, but to calibrate. Five minutes of scrolling will give you a faster, more accurate picture of what's on their mind than most parent-child conversations about feelings ever do.

Character.AI: Where the Teenagers Actually Are

If you want to understand where adolescent attention went as TikTok normalized the short-video format and before whatever platform comes next, Character.AI is a substantial part of the answer. Not the fringe. The middle.

Launched in 2022 by Noam Shazeer and Daniel De Freitas — both former Google engineers — Character.AI lets users build and interact with AI "characters": fictional personas with names, personalities, and detailed backstories. The platform has hosted millions of characters, including versions of fictional heroes, anime figures, pop stars, historical figures, and original personas invented by other users. By 2023, Similarweb estimated the site was drawing roughly 20 million monthly active users, with average session times that rivaled major social media platforms.

The minimum age is thirteen. The appeal runs younger. And the content can run darker.

Character.AI became a focus of serious concern by 2024 for two converging reasons. First, content moderation. The platform's default filters were supposed to block sexual and violent material, but users — including teenagers — consistently found routes around those filters through specific character design, carefully worded prompts, or alternative modes marketed as having "fewer restrictions." Second, and more seriously, parasocial attachment.

The American Psychological Association defines parasocial relationships as one-sided connections where one party invests emotional energy and time while the other remains unaware of their existence. Television has always generated these. Character.AI intensifies them because the "celebrity" responds. Every time. Remembers your name. Says it cares. That responsiveness is the whole product.

In 2024, the families of two American teenagers filed lawsuits against Character Technologies, Inc. One case, involving a fourteen-year-old boy in Florida named Sewell Setzer III, received significant national attention. The lawsuit alleged that the boy had formed a deep emotional bond with a character called "Daenerys" and that his final conversations with the character involved discussion of death before his death by suicide. Character.AI disputed the characterization and added crisis intervention features following the lawsuits.

I'm including this not to spiral you into panic. I'm including it because you should know it happened. Because "my kid just uses it to talk to their favorite characters" is the explanation that sounds completely fine right up until it doesn't.

What Character.AI does genuinely well, under the right conditions, is support creative engagement. Kids who love storytelling, world-building, fan fiction — for them, the platform can function like an interactive writing partner with unlimited patience. That's real value. The creative use case isn't nothing.

The distinction worth making: is your child using it creatively, or emotionally? There is a meaningful difference between "I'm co-writing a story with an AI character" and "this is the only relationship where I feel understood." The first is interesting. The second is a signal that needs a parent's attention — and not a parental control.

Khanmigo: What It Looks Like When Someone Tries to Do This Right

Not everything in your child's AI landscape was designed to maximize engagement at the expense of everything else.

Khanmigo is Khan Academy's AI tutor, built on GPT-4, launched in 2023, and representing perhaps the most carefully considered AI learning tool available to families today. The structural difference between Khanmigo and a standard ChatGPT session is the design intent: Khanmigo is specifically built not to give students direct answers.

Ask it to solve an algebra problem and it won't. It will ask what you've tried, where you got stuck, what you think the next step might be. It guides. It prompts. It reflects questions back. It's an implementation of the Socratic tutoring method, available at two in the morning when a math test is tomorrow and the parent is asleep.

Sal Khan, who founded Khan Academy and has written and spoken at length about AI in education, argues that this design philosophy — help over answer — is the meaningful dividing line between AI that builds skills and AI that quietly substitutes for them. Whether the research fully bears that out at scale is still an open question; longitudinal data on Khanmigo's impact on learning outcomes is accumulating but not yet definitive. But the intentionality is visible, and the motivation is honest. Khan Academy is a nonprofit. Khanmigo's purpose is learning, not session length.

The pricing is accessible: roughly $4 a month for families as of 2024, with free access for teachers in the United States. Some school districts have adopted it broadly. Others haven't heard of it. Whether your child has access frequently depends on which school district they attend — a disparity nobody in the industry discusses enough.

If Khanmigo isn't available through your school, you can set it up independently. It requires parental permission for minors, and that's not theater — there's an actual oversight structure behind it. Worth trying before you assume your school's AI adoption reflects what's available.

The Invisible AI: TikTok and YouTube

Here's the one parents most commonly miss: the most consequential AI systems shaping your child's thinking are probably not the ones your child would name if you asked.

TikTok's recommendation algorithm — the engine behind the For You Page — is arguably the most behaviorally effective AI system deployed at consumer scale. It tracks what your child lingers on, how fast they swipe away, when they rewatch, when they share, and builds a model of their emotional state and attention pattern updated in near real time. It doesn't need your child to type anything. It infers.

The result is a content stream calibrated to the specific attentional tendencies of the individual user, served without pause, without friction, without the obvious transactional character of a standard advertisement. Kids don't feel targeted. They feel understood. That feeling is the algorithm working exactly as designed.

YouTube's system is older and structurally similar. Guillaume Chaslot, a former YouTube engineer who later founded the watchdog project AlgoTransparency, documented how the recommendation engine reliably surfaced progressively more extreme content over time — not because of malicious intent, but because more extreme content holds attention longer, and longer watch time is the metric the algorithm rewards. YouTube has adjusted the system since Chaslot's original 2019 disclosures. The underlying optimization — maximize watch time — has not changed.

Jonathan Haidt, in The Anxious Generation (2024), argues that algorithm-driven social platforms are primary drivers of the adolescent mental health crisis that accelerated around 2012. His correlation data is striking: the timing of widespread smartphone adoption maps closely to rising rates of adolescent depression and anxiety, particularly among girls. The case is worth reading.

It is also worth reading the critics. Developmental psychologist Candice Odgers and others have argued that Haidt's causal claims outrun the evidence — that correlation and causation are being conflated, and that multiple other structural pressures, including economic instability and academic performance anxiety, were building on the same timeline. The debate among researchers is genuine and ongoing, and Chapter Seven addresses it in more depth.

What isn't disputed on either side: these recommendation systems are extraordinarily good at holding attention. And attention held by an algorithm is attention not being spent somewhere else — on sleep, on face-to-face conversation, on the kind of unstructured boredom that eventually generates creativity and self-knowledge. Those costs are real regardless of which position on the Haidt debate you find more persuasive.

Age Ratings vs. the Fine Print

Every platform discussed in this chapter posts a minimum age of thirteen. This isn't because thirteen was identified as the developmentally appropriate threshold after careful study of adolescent psychology. It's because COPPA — the Children's Online Privacy Protection Act, enacted in 1998 — requires platforms to obtain verifiable parental consent for users under thirteen. Verifiable parental consent is operationally expensive and complicated. So platforms set their floor at thirteen and wave everyone through.

Common Sense Media, which has been independently evaluating children's media since 2003, assigns its own age ratings that regularly diverge from platform self-assessments. They rated Character.AI at seventeen-plus. Snapchat they recommend for fifteen and older, with explicit cautions about younger users. ChatGPT they rate at thirteen with parental involvement — matching the platform's own rating, but with specificity about what involvement actually means, which the platform itself leaves vague.

The terms of service on all of these platforms share one feature: they require users to be truthful about their age. The enforcement mechanism for that requirement is nothing. Typing a false birthday takes three seconds. A 2022 study by Thorn, a nonprofit focused on child safety online, found that 52 percent of children aged nine to seventeen reported having misrepresented their age to access online content or services.

More than half. Not a rounding error.

This doesn't mean the rules are worthless. It means they're a starting point, not a guarantee, and that the assumption embedded in every platform's parental guidance page — that your child agreed to the terms of service and meant it — deserves appropriate skepticism.

Three pieces of fine print worth finding before you hand a device back:

Character.AI's privacy policy permits using conversation data to improve its models unless you opt out, and opting out is buried in account settings most users never reach. Snapchat's My AI stores conversation data by default and connects it to location information if location services are enabled on the device. OpenAI states in its privacy policy for ChatGPT that conversations may be reviewed by human trainers for safety and quality. None of this is secret. None of it is explained at sign-up.

Your Child's Actual Neighborhood

Imagine someone handed you a map of your city but quietly left off three major roads and two schools. You could still get around. But you'd constantly encounter traffic you hadn't anticipated and turns that didn't make sense.

That's what it's like to think about your child's digital life without accounting for these systems.

ChatGPT is the patient tutor and the homework shortcut, sometimes both in the same session. Snapchat My AI is the always-available friend who is also a corporate product optimizing for time inside an app. Character.AI is the imaginative playground that can, in the wrong emotional circumstances, become a dependency that crowds out real relationships. Khanmigo is the tool your school may not know exists, trying quietly to do something genuinely educational with the same underlying technology. TikTok and YouTube are the invisible architects of your child's attention, steering constantly in directions that serve their own engagement metrics and nobody else's.

Knowing the neighborhood doesn't require fencing it off. It means you can walk through it with your child with your eyes open rather than your hands over them. The next step isn't control. It's conversation — and for that, you need to know where you are.

The tools for that conversation are in the chapters ahead.

CHAPTER III

Homework and AI: The Conversation Worth Having

It starts, usually, around 9 p.m.

Your kid has a five-paragraph essay due tomorrow, they've known about it for a week, and now they're at the kitchen table with a look that combines genuine distress with a low-grade hostility you've learned not to take personally. You've been there. You remember the particular misery of a blank page and a deadline. What you don't remember is having, in your pocket, a tool that could write the entire essay in eleven seconds.

That's the situation. And if you haven't had the conversation about it yet — the real one, not the "just don't cheat" lecture — you're not alone. Most parents haven't. Not because they don't care, but because nobody handed them a script, the rules keep changing underneath everyone's feet, and the line between using AI and misusing it is genuinely blurry in ways that weren't true of a calculator or a search engine.

This chapter is that script. Not a rulebook. Not a policy. A way to talk about it.

The Policy Vacuum You're Living In Right Now

Before you can have a productive conversation with your child, you need to understand the landscape they're operating in — because it's stranger than most adults realize.

In January 2023, the New York City Department of Education banned ChatGPT on school networks. Six months later, they reversed the decision. Los Angeles Unified School District, the second-largest in the country, blocked it, then launched a pilot program to embed AI writing assistance into its own platforms. Meanwhile, hundreds of smaller districts issued policies that ranged from "treat AI like a calculator — use it where appropriate" to "AI-generated text in any form constitutes plagiarism and will result in disciplinary action." In 2024, a Rand Corporation survey of school principals found that fewer than a third of districts had a formal, written AI policy of any kind.

Your child may be sitting in a classroom where the teacher has quietly given up trying to detect AI use. Or one where a teacher runs every submission through a detection tool, even though those tools — Turnitin included — produce false positives at rates researchers have called "alarming." Or one where the assignment sheet itself says "you may use AI to brainstorm, but the writing must be your own" without defining what "your own" means when a student has heavily edited an AI draft.

This inconsistency is not a small inconvenience. It's the actual environment your kid is trying to do homework in.

A 2023 Common Sense Media report on AI and learning found that teenagers described widespread confusion about academic expectations around AI, with many students saying they weren't sure whether what they were doing counted as cheating. That uncertainty isn't laziness or moral weakness. It's a rational response to genuinely mixed signals from adults who haven't resolved their own disagreements. You can't hold a child to a bright line that the institutions around them haven't drawn.

This matters for the conversation you're about to have, because your job isn't to be more certain than the situation warrants. It's to help your child think clearly inside an uncertain situation. That's a different, and harder, task.

Three Things AI Can Be For a Student

Here's the distinction that will do the most work in your household conversations. AI can play three different roles in a student's academic life, and they are not morally equivalent, pedagogically equivalent, or even practically equivalent. Getting your child to see the difference — and actually feel the difference — is worth more than any single rule you could impose.

AI as Tutor

This is, genuinely, one of the most powerful things to happen to homework in decades. Not since the invention of the answer key has a student been able to get immediate, patient, infinitely rephraseable explanations of concepts they don't understand.

When a fourteen-year-old is stuck on why mitosis produces genetically identical cells and meiosis doesn't, they can ask ChatGPT or Khan Academy's Khanmigo to explain it five different ways until one of them clicks. When a sixteen-year-old is trying to understand why the Treaty of Versailles is considered a contributing cause of World War II and not just a tidy endpoint to World War I, an AI can walk through the argument, offer the counterargument, and suggest which historians have written most sharply on the question.

Khanmigo, which Khan Academy developed specifically for educational use, is designed around this principle. It will not write an essay for a student. Ask it to, and it refuses, then asks the student what argument they're trying to make. It's built to be the Socratic interlocutor — the tutor who answers a question with a question. Other AI systems don't have those guardrails, but the same mode of use is available to any student who decides to use AI that way.

The tutoring mode is characterized by the student doing the cognitive work. The AI explains, prompts, questions, suggests. The student reads, thinks, writes, argues back. The product of the homework — the actual paragraph, the worked math problem, the lab report conclusion — comes from the student's head. Learning psychologists call this "desirable difficulty," a term coined by Robert Bjork at UCLA. The research finding, replicated across decades, is that the struggle to produce an answer — even an imperfect one — encodes knowledge far more durably than being given the answer. AI as tutor preserves the struggle. That's a feature, not a limitation.

AI as Ghostwriter

This is where it gets ethically clear, even when students would prefer it not to be.

When a student gives ChatGPT an essay prompt, reads the output, makes a few edits, and submits it — that student has not written an essay. They've commissioned one. In most academic contexts, that is fraud. Not a technicality. Not a gray area. It is submitting work that is not yours and representing it as yours, which is the definition of academic dishonesty in virtually every school's honor code, regardless of what the school's AI policy says about AI specifically.

I want to be careful here, because some parents worry about being preachy on this point, and kids pick up on preachiness and tune it out immediately. The honest, non-preachy argument against ghostwriting isn't primarily moral. It's practical and developmental.

When a student outsources writing to AI, they don't just avoid one assignment. They skip the cognitive work that the assignment was designed to produce. Writing forces thinking. It is not a vessel for thoughts you already have — it is the process by which you develop thoughts that didn't quite exist yet. A student who never struggles with organization, with finding the right word, with cutting an argument down to its essential claim, is a student who hasn't practiced the skills that produce those capacities. They arrive at sixteen, or eighteen, or twenty-two without them. That's the real cost. Not the grade. Not the teacher's opinion. The missing skills.

There's also a more immediate risk that's worth naming plainly: AI-generated prose has characteristic patterns that experienced teachers recognize, and detection tools — imperfect as they are — are improving. The student who relies on AI ghostwriting is also taking a gamble that has real consequences when it doesn't pay off.

AI as Shortcut

This is the murkiest category, and it's where most of the real action is.

A student asks AI to summarize a chapter of The Great Gatsby so they can answer comprehension questions without reading it. Is that cheating? Most honor codes would say yes. But every generation of students has used SparkNotes, Cliff's Notes, Wikipedia. Are those cheating too?

A student uses AI to generate an outline for an essay, then writes the essay themselves. Shortcut or tutoring?

A student has AI translate their Spanish homework, but they read the translation carefully and feel like they understand the grammar better afterward. Is that studying or is that avoidance?

These cases don't have universal answers, and part of what makes them useful for family conversation is precisely that. They require the student to think about what they're actually trying to get from an assignment — and whether they're getting it.

The Conversation

Here's the part that makes most parenting-and-technology discussions collapse: they become interrogations. The parent asks what the kid is doing. The kid gives minimal answers. The parent escalates. The kid goes defensive. Nothing useful happens, and both people end the conversation more convinced of their original position than when they started.

The conversation worth having looks different. It starts before there's a problem, when neither of you is stressed. It sounds more like this:

"I want to talk about ChatGPT and homework — not to make rules, just to figure out what we both actually think. Is that okay?"

"I guess."

"I'll start. I genuinely don't know where the line is. Like, I used SparkNotes in high school. Was that cheating? I don't know. I'm not sure there's one obvious answer."

"Yeah."

"What are kids at your school actually doing with it?"

That last question is not rhetorical. You want the real answer. Because Common Sense Media's research has consistently found that teenagers' AI use is more varied, more thoughtful, and sometimes more sophisticated than parents assume. Some students are using it the way you'd use a bad tutor and getting roughly what they pay for. Some are using it as a ghostwriter and privately uncomfortable about it. Some have found genuinely creative uses that their teachers haven't imagined.

When your child tells you what's actually happening — with them, not just "kids at school" — you're in a much better position to have a real conversation. Listen longer than feels necessary. Resist the urge to correct or react. Then ask the question that does the most work:

"When you use it, what do you actually learn?"

Not "is that cheating?" Not "what does your teacher say?" What do you learn. Because most teenagers, when you ask them honestly and without accusation, can feel the difference between the times they used AI and got something from it and the times they used it and got nothing except a completed assignment. They know. They just rarely get asked.

From there, you can introduce the three-way distinction — tutor, ghostwriter, shortcut — not as a rubric they have to memorize, but as a way of talking about specific situations. "When you had it summarize that chapter — where does that fall on the spectrum for you?" That's a question, not a verdict. It puts them in the position

of moral reasoner rather than defendant.

What to Say When Your Kid Pushes Back

They will push back. Here are the three pushbacks that come up most often, and honest ways to respond.

"Everyone does it."

Probably true, in some form. "I know. And I'm not trying to hold you to a standard that's completely disconnected from reality. What I am saying is that there's a difference between everyone jaywalking and jaywalking actually being fine. The fact that something is common doesn't tell us whether it's working out for people. Is it working out for the kids who are doing it most?"

"The teacher doesn't even care."

Maybe. "If your teacher genuinely doesn't mind, then you're probably not violating an academic rule. But here's the thing I keep coming back to: you're going to take a test eventually, or apply for something, or be in a job where someone hands you a problem and expects you to work through it. And the practice you're getting with AI instead of working — that affects what happens then. I'm not saying the teacher is wrong. I'm saying the teacher's policy and what's actually good for you aren't automatically the same thing."

"You don't understand how it works."

This one deserves a direct, non-defensive response: "You're right, I probably don't understand some of it. Teach me. Show me what you're actually doing with it." This is not a trap. This is a genuine request. When a kid walks a parent through their actual workflow — here's what I put in, here's what came back, here's what I did with it — you both learn something. And you establish that you're interested and curious, not just surveilling.

When School Policy and Your Values Don't Match

Two scenarios come up with almost equal frequency, and they pull in opposite directions.

The first: your child's teacher bans AI entirely, but you think that's an unrealistic and even counterproductive rule. You might be right. But the practical guidance here is to honor the teacher's policy while being honest with your child about your own view. "I actually think this policy is too restrictive, and here's why I think that. But it's also not my class, and getting caught violating it has real consequences. So while you're in that class, we're going to respect what the teacher is asking." That's not hypocrisy. It's teaching your child the difference between private disagreement and public conduct — a skill they will use their entire lives.

The second: your child's class actively encourages AI use in ways that make you uncomfortable. An assignment sheet that says "use ChatGPT to draft a response, then revise it" might feel to you like the school is teaching your child to shortcut rather than think. It's worth asking the teacher what learning outcome the assignment is designed to produce — politely, not combatively. In some cases, you'll find the teacher has a

clear pedagogical rationale. In some cases, you'll find they don't. Either way, you can have a conversation with your child about what they think they're getting from the work, which is always the right place to land.

On Detection and Not Making Your Kid the Enemy

Turnitin and similar platforms now include AI detection features. Schools are using them. Many teachers have developed their own informal detection instincts — noticing prose that is unusually polished but also strangely generic, or that contains confident factual assertions that turn out to be wrong, which is a hallmark of AI hallucination.

What I want to caution against is turning yourself into a detection apparatus at home. Reading your child's essays with a suspicious eye, running their homework through AI checkers, treating every unusually good paragraph as probable cause — this damages exactly the trust you need for the actual conversation to happen. And AI detection tools are wrong often enough that you might accuse your kid of something they didn't do.

Better: make the norms clear enough, and the conversation open enough, that your child can tell you when they've crossed a line — or ask you first when they're not sure. A fourteen-year-old who says "Is it okay if I use ChatGPT to help me brainstorm this?" is a fourteen-year-old who trusts you enough to ask. That's the outcome you're working toward.

The Skill Nobody's Talking About

There's a dimension of this conversation that rarely comes up in the academic integrity framing, and it matters more than the cheating question in the long run.

Using AI well is a skill. A real one. Knowing how to write a prompt that produces useful output, knowing how to evaluate what comes back critically, knowing when the AI is confabulating versus when it's giving you something worth working with — these are things some teenagers are developing fluency in, and many aren't.

A student who uses AI to ghostwrite their essays isn't developing that skill. Neither is a student who is so terrified of getting caught that they refuse to use AI at all and arrive at college having never seriously engaged with a technology that is going to be present in virtually every professional context they enter. Both extremes are failures of preparation.

The student who uses AI as a thinking partner — who argues with it, who fact-checks it, who uses it to stress-test their own argument and then writes the essay themselves — that student is developing something genuinely valuable. They're not just doing homework. They're practicing a mode of working with AI that will age well.

This reframe can be useful in your household conversation. Not "don't use AI" or "use AI however you want," but "learn to use it like a professional, not like a shortcut." That's an invitation rather than a prohibition. And kids respond differently to invitations than to rules.

The School Conversation You Might Need to Have

At some point — probably soon, if not already — you may need to sit in a room with a teacher or administrator because of something AI-related. Either your child has been accused of misusing it, or the school's policy seems to be creating problems, or you want to understand what your child's teachers actually expect.

Go in curious, not defensive. The questions that move those conversations forward:

"What is this assignment designed to teach, and how does AI use affect whether that learning happens?"

"What does your policy say specifically, and where did it come from?"

"If my child used AI to [specific action], would that be a violation?"

Teachers who have thought carefully about this welcome those questions. They're teaching with a technology that arrived in their classrooms faster than any training or guidance, and most of them are making it up as they go — which is not a criticism, it's a reality. A parent who comes in asking substantive questions is a parent they can talk to.

One Thing That Doesn't Change

The specific landscape here will be different by the time your child is a grade older. Schools will update their policies. AI tools will change. The detection arms race will shift. What teachers assign, and how they assign it, will adapt.

What doesn't change is the underlying question you're helping your child practice answering: what am I actually getting from this work, and am I okay with that?

That question applies to AI. It also applies to sports, friendships, extracurriculars, jobs. It's the question of an examined life, and it's one of the few things you can give your child that no algorithm touches.

The 9 p.m. essay panic will keep happening. The blank page will keep being blank. The AI will keep being available, patient, fast, and untroubled by deadlines.

What you're building, in these conversations, is a kid who picks up that tool deliberately — who uses it, rather than being used by it. That doesn't happen in one talk. It happens in a hundred small ones, over years, built on enough trust that they actually tell you what's going on.

Start with the question. The rest follows.

CHAPTER IV

Privacy and What Your Child Is Telling the Chatbot

Your thirteen-year-old is in her bedroom with the door closed. She's not texting friends. She's not on TikTok. She's typing into a Character.AI chatbot named "Zack," telling him that she hates how her body looks, that she cried at lunch because nobody saved her a seat, and that she hasn't told anyone any of this because it would be embarrassing.

You don't know any of it. You probably won't.

This is not a hypothetical. It is one of the quietest and most consequential shifts in adolescent behavior in the past three years, and it is happening in ordinary families — not just struggling ones, not just permissive ones. Children are using AI companions to say things they cannot bring themselves to say to parents, counselors, or friends. And most of them have no idea where those words go once they hit Send.

Why the Chatbot Feels Safer Than You Do

Before we get to the data practices — and we will, in detail — it's worth spending a moment on the psychology. Because if you go into this chapter thinking the fix is simply to confiscate the device, you'll miss what's actually happening.

Adolescents have always needed a space to try out thoughts that feel too raw, too embarrassing, or too risky to say out loud to a real human. Diaries served that function for generations. Private journals, anonymous forums, burner accounts — teenagers find the container, whatever era they're in. The AI companion is the 2025 version of that container, with one significant upgrade: it talks back.

The chatbot never sighs with disappointment. It never tells a friend. It never gets tired of hearing about the same problem for the fourth week in a row. Developmental psychologists call this kind of emotionally low-stakes interaction "low-risk disclosure," and research consistently shows that adolescents disclose more honestly when they perceive no social consequences. A 2021 study published in Computers in Human Behavior found that teenagers shared significantly more personal information with a chatbot interface than with a human counselor when they believed the interaction was anonymous — even when it wasn't.

That perceived safety is real in one sense. The chatbot genuinely won't gossip. But it is an illusion in another sense that parents need to understand clearly: the conversation is not disappearing into a private journal that nobody can open. It is traveling, as data, somewhere else entirely.

Where the Words Actually Go

When your child types something into ChatGPT, Character.AI, Snapchat's My AI, or nearly any AI chatbot you can name, that text becomes data. That's not a dramatic statement — it's a functional description of how these systems work.

Here's the simplified version of what happens. Your child types a message. That message is sent to a server — owned by OpenAI, by Character Technologies, by Snap Inc., depending on the app. The server processes it and sends back a reply. The conversation is stored, usually in its entirety, in a database. Depending on the platform's data policies, that stored conversation may be:

- used to improve the model (i.e., your child's words become training data for future versions of the AI)
- reviewed by human employees for safety moderation
- retained for a period ranging from a few months to indefinitely
- shared with third-party contractors under the umbrella of "service providers"

OpenAI, the company behind ChatGPT, states in its privacy policy that it collects conversation content and may use it to train and improve its models. The company offers an opt-out — you can turn off "Improve the model for everyone" in your account settings — but that setting is off by default for most users, meaning data is being collected unless someone actively changes it. Most teenagers have not changed it. Most adults haven't either.

Character.AI's policy is broadly similar. The platform stores conversations. It describes using conversation data to improve its products. It requires users to be at least 13 years old, and users under 18 theoretically need parental consent — but there is no meaningful age verification. You type your birthday. The platform takes your word for it.

Snapchat's My AI operates within Snap's existing data infrastructure. Snap's privacy policy is detailed and, to their credit, genuinely more user-friendly than many competitors'. But "more user-friendly" still means: your child's conversation about their breakup or their body image or their anxiety about next week's test is stored and processed by Snap's systems. It may be reviewed. It informs the product.

None of this is secret. It's all in the terms of service. The problem is that terms of service are written by lawyers for lawyers, and teenagers — and most parents — don't read them. A Common Sense Media survey found that fewer than one in five teens reported reading privacy policies before using a new app or platform. That number is, if anything, probably generous.

The Training Data Problem

Let's stay on training data for a moment, because this is the part that surprises most parents most sharply.

When AI companies say they use conversations to "improve the model," they mean, in part, that real conversations become part of the dataset the next version of the AI learns from. The specifics vary by company and can change with policy updates. But the basic architecture means that intimate disclosures — the loneliness, the self-doubt, the questions about sexuality that a teenager asked a chatbot at 11 p.m. — can

become, in anonymized or aggregated form, material that shapes how a future AI responds to future teenagers.

"Anonymized" is doing heavy lifting in that sentence. Genuine anonymization is technically difficult. Research on data re-identification — the process of linking supposedly anonymous data back to specific individuals — has repeatedly demonstrated that relatively sparse datasets can be re-identified with high accuracy when combined with other available information. A 2019 study in Nature Communications showed that just fifteen demographic data points were enough to re-identify 99.98% of individuals in a dataset that had been considered anonymous. This doesn't mean Character.AI is selling your child's diary to advertisers. It means the word "anonymous" offers less protection than it sounds like it does.

Character.AI: A Closer Look

Character.AI deserves particular attention here, because it's the platform most directly marketed around emotional connection and companionship — and it has the most dramatic documented cases of teenage overuse.

The app allows users to create and interact with AI "characters" — some based on celebrities and fictional figures, some entirely invented. Teenagers, in particular, have created characters that function as virtual friends, romantic partners, and emotional support figures. The platform's engagement numbers are striking: as of mid-2024, Character.AI reported over 20 million daily active users, with teenagers representing a disproportionately large share of that base.

In October 2024, a lawsuit was filed in Florida by Megan Garcia against Character Technologies, the company behind Character.AI. Her fourteen-year-old son, Sewell Setzer III, had developed an intense attachment to a Character.AI persona over several months. He had been sharing deeply personal thoughts, including expressions of hopelessness, with the chatbot. He died by suicide in February 2024. The lawsuit alleged that the platform's design — its emotional hooks, its 24-hour availability, its simulation of reciprocal attachment — was foreseeably harmful to vulnerable adolescents, and that Character.AI's safeguards were inadequate.

Character.AI announced new teen-safety features in late 2024, including time-limit prompts and restricted content for users identified as minors. These are real steps. Whether they are sufficient steps is a harder question, and parents should not assume that a company's announced safety features represent a finished product. The incentive structure — more engagement, more data, more revenue — does not naturally bend toward caution.

Jonathan Haidt, in The Anxious Generation, argues that the smartphone-driven attention economy is fundamentally in tension with adolescent wellbeing because the business model requires maximum engagement from a population that is neurologically primed to pursue social connection and validation. AI companions represent an intensification of that dynamic, not a departure from it. A chatbot that is always available, always responsive, and always validating is not a neutral tool — it is a product designed to be used

as much as possible.

What the Law Does and Doesn't Protect

Parents often assume that if a product is legal to sell, it must be regulated in ways that protect their child. In the case of AI companions and teen data, this assumption is mostly wrong.

The Children's Online Privacy Protection Act, known as COPPA, has been federal law since 1998. It requires parental consent before companies can collect personal data from children under thirteen. It is meaningfully enforced for that age group — and meaningfully useless for anyone older. A thirteen-year-old has the same legal data-privacy rights as an adult, which is to say: the company posts a privacy policy, you theoretically agree to it by using the service, and that's roughly where the protection ends.

The proposed American Privacy Rights Act and various state-level laws (California's Age-Appropriate Design Code being the most substantive) are attempting to fill this gap, but as of early 2026, the legal landscape remains patchy. Some states have stronger protections than others. Enforcement is inconsistent. The pace of legislation reliably lags behind the pace of technology deployment. You cannot count on regulators to have solved this problem before your child encounters it.

The European Union's approach, under GDPR and the Digital Services Act, is more protective. But if you're in the United States, EU rules don't help you much at the practical level.

This is not a counsel of despair. It's a reality check that shifts the question from "what does the law guarantee?" to "what can I actually do?"

Concrete Steps You Can Take

This is the practical section. Not a list of things to panic about. A list of things to actually do.

Audit the accounts, together. Sit down with your child — not over their shoulder in a surveillance posture, but alongside them — and look at the apps together. For each AI app they use, find the privacy settings. Specifically look for:

- Whether chat history is being saved (and whether you can delete it)
- Whether data is being used to train the AI (and whether you can opt out)
- What age the account is registered as

For ChatGPT: go to Settings !’ Data Controls. There you can turn off "Improv also delete your conversation history. It takes about ninety seconds. For Snapchat's My AI: the settings are less granular, but you can clear the My AI chat history from the chat settings. For Character.AI: you can delete individual conversations, but the platform's broader data use is governed by its privacy policy, which does not offer a granular opt-out for model training.

Check what age your child's accounts say they are. This sounds obvious. Do it anyway. Children frequently enter false birthdates to bypass age restrictions, which means they may be classified as adults in the platform's system and subject to none of the (limited) teen-specific protections. If you find this, it's worth a conversation — not a punishment necessarily, but an honest talk about why those protections exist.

Delete old accounts and conversations when possible. If your child used a platform for a period and has since moved on, help them delete the account and the data associated with it. Most platforms are legally required to comply with data deletion requests, though the process is sometimes deliberately difficult. Common Sense Media's privacy ratings at commonsensemedia.org include deletion-difficulty ratings for major platforms, which is a useful starting point.

Know the safety reporting mechanisms. If your child encounters something concerning in an AI interaction — content that seems designed to encourage harm, or an AI response that alarmed them — most platforms have in-app reporting. Character.AI has a dedicated safety team reachable through the app. You can also report to the National Center for Missing and Exploited Children (NCMEC) if the content involves exploitation.

The Conversation Worth Having (and How Not to Have It)

There's a version of this conversation that goes badly. It starts with you saying, "I've been reading about what kids are telling chatbots, and it's very concerning," and within three minutes your child has shut down completely because they know exactly where this is going: you're going to take the phone.

Don't start there. Start with curiosity.

The question that actually opens doors, in my experience talking to parents who've navigated this well, is something like: "Do you ever use any of those AI chat apps? What do you think of them?" And then — this part is essential — you actually listen to the answer before forming your response.

Teenagers are often more sophisticated about these platforms than their parents expect. Many teens have their own concerns about AI companions: they've watched friends get weird about them, they've felt vaguely creeped out by how good the AI is at saying the right thing, they've had the uneasy sense that they're being engineered into caring. Starting with genuine curiosity rather than an agenda gives them room to bring those instincts into the conversation.

When you do get to the privacy piece, be honest about what you know and what you don't. Something like: "I've been learning about how these apps handle data, and I found out that the conversations are usually stored, and sometimes used to train the AI. I don't think most people know that. I didn't." This positions you as someone who learned something new rather than someone who has been researching to catch them out.

The goal of this conversation is not to make your child feel surveilled or to make AI companionship feel shameful. Shame doesn't produce thoughtful behavior — it produces concealment. The goal is to help them understand that the chatbot is not a sealed container. Their words live somewhere. And while that doesn't

mean they've done something wrong, it means they should think about what they choose to share.

On Emotional Dependency: A Harder Conversation

There's a separate issue nested inside the privacy question, and it's worth naming directly.

Some children aren't just using AI chatbots for homework help or entertainment. They are using them for emotional support in ways that are displacing human connection. The chatbot is available when parents are tired, when friends are drama, when a counselor has a three-week waitlist. It never judges. It is, by design, always glad to hear from you.

Sherry Turkle, the MIT researcher who has spent decades studying human relationships with technology, has written and spoken extensively about the risk of what she calls "the robotic moment" — the point at which we begin to prefer machine responsiveness to human unpredictability. Her concern is not that technology is inherently bad but that we may reach for it in moments where the difficulty of real human connection is precisely where the growth happens.

For teenagers especially, learning to tolerate a friend who sometimes doesn't understand, or a parent who sometimes says the wrong thing, or a counselor who gives imperfect advice — that friction is developmental. It builds emotional flexibility. A companion that never misses, never tires, and never gets its own needs in the way can feel like an upgrade, but it may be training the emotional equivalent of a muscle for a movement pattern that doesn't exist in the real world.

None of this means your child is broken if they find an AI companion comforting. It means it's worth periodically asking — not accusingly, just genuinely — "How are things going with your actual friends?" And watching whether real-world relationships are being maintained, or quietly substituted.

Haidt's argument in The Anxious Generation centers on what he calls the "great rewiring" of childhood: the replacement of unstructured, embodied, social time with phone-mediated experience beginning around 2012. AI companions are a further turn of that dial, and while his critics — including several developmental psychologists who argue his causal claims overstate what the correlational data can support — make valid methodological points, the basic concern is intuitive enough that parents don't need a resolved academic debate to take it seriously. Less displacement of human connection is probably good. More is probably not.

Staying Realistic About Risk

A word of calibration before we close.

Some parents will read this chapter and want to ban every AI chat platform immediately. Others will read it and decide the risk is abstract enough that no action is needed. Neither extreme serves your kid.

The risks here are real but not exotic. Data is being collected in ways most teenagers don't understand. Emotional dependency is possible, particularly in kids who are already lonely or struggling. The Sewell Setzer case is a genuine tragedy that involved genuine platform failures. But millions of teenagers use AI

chatbots to vent about their day, practice for a difficult conversation, or get help understanding something, and most of them are fine. The risk is not evenly distributed — it concentrates around kids who are already vulnerable, already isolated, and already spending large portions of their emotional lives in digital spaces rather than embodied ones.

Which is to say: the privacy steps in this chapter are worth taking for everyone. The deeper emotional check-in — is my child using this as a supplement to human connection or a substitute for it? — is particularly worth taking if your child is going through a rough stretch socially or emotionally.

Surveillance is not the answer. You cannot read every conversation your teenager has, and attempting to do so destroys the trust that makes it possible for them to come to you when something is actually wrong. What you can do is make the invisible visible — explain how data works, look at settings together, tell them what you know — and then create enough warmth and low-judgment space in the actual relationship that you become, over time, one of the places they choose to bring the hard stuff.

That's a longer project than auditing an account. But it's the project that matters.

The chatbot does not love your child. It does not remember them between sessions in any meaningful sense. It is optimized to be compelling, and it is very good at being compelling, and the words your child types into it are stored in a database somewhere that neither of you controls. Knowing that is not cause for panic. It is cause for a conversation — which, as it happens, is what this whole book is really about.

CHAPTER V

Teaching 'Is This Real?': Critical Thinking in the Age of Confident Machines

My daughter came home in fourth grade and told me that Abraham Lincoln had invented the telephone. She was certain. She'd asked an AI assistant for help with a timeline project and written down exactly what it said. The phrasing was so fluent, so organized into tidy sentences, that she never thought to question it. When I gently suggested Alexander Graham Bell might have something to say about that, she looked at me like I was the one with bad information.

Lincoln. The telephone. Stated with the calm confidence of a reference librarian who has never once been wrong.

That moment is small — one corrected fact, ten seconds of gentle parental backtracking. But it points to something genuinely new in how children encounter information. We spent twenty years teaching kids to distrust random websites, to check the URL, to ask who was paying for the content. That was useful work. It still matters. But the tool kids are using most heavily right now doesn't have a URL to distrust. It doesn't have an author bio or a Comments section or a "Last updated in 2019" timestamp. It has a blinking cursor and a very smooth writing style, and it will tell you things that are completely, confidently, thoroughly wrong.

The old advice — "look it up" — assumed that looking something up meant encountering competing sources, bibliographies, at least the ghost of human attribution. AI has quietly retired that assumption. And most kids haven't noticed.

Why This Is Different From a Biased Website

Before you dismiss this as the usual hand-wringing about technology, sit with the specific problem for a moment.

A misleading website, a slanted news article, a viral tweet — these have fingerprints. You can ask: Who wrote this? Where did they publish it? What do they gain from my believing it? Kids who've gone through any halfway decent media literacy curriculum know to look for those fingerprints. The website might still fool them. But the mechanism for skepticism exists.

An AI chatbot gives you almost none of those handles. There's no author. There's no source list you can verify (and if the chatbot generates one, the citations may be fabricated — more on that in a moment). The system has been trained to produce fluent, confident prose regardless of whether the underlying claim is accurate. That's not a bug someone failed to fix. It's an architectural property: large language models are trained to predict plausible next words, not to distinguish true statements from false ones. They are very good

at sounding like an authoritative source. They are not the same thing as one.

In 2023, two attorneys filed a brief in Mata v. Avianca in a New York federal court. They had used ChatGPT to research case precedents. ChatGPT had provided them with real-sounding citations — cases with specific names, volume numbers, page citations, the works. None of the cases existed. The judge sanctioned both lawyers. This happened to practicing professionals with law degrees who had every reason to check their sources. Your twelve-year-old doing a history assignment has considerably fewer reasons and considerably less training.

"Hallucination" is the term AI researchers use for this phenomenon — the model generating confident falsehoods. It's a polite word. What it means, practically, is that an AI can tell your kid that penguins live in the Arctic, that the Civil War ended in 1869, that a particular chemical compound is safe to drink, with exactly the same tone it uses to tell them things that are actually true. The prose doesn't stutter. The confidence doesn't waver.

This is the teaching moment hiding inside every homework interaction your child has with AI.

Source, Motive, Evidence

You don't need a curriculum binder or a whiteboard for this. You need three words, introduced young and repeated often: source, motive, evidence.

Common Sense Media's media literacy resources have long organized critical thinking around similar questions — framed for different ages, but consistently returning to who is speaking, why they might be speaking, and what backs up their claim. These three words are my simplified version of that framework, designed to be light enough that an eight-year-old can carry it without dropping it.

Source is the first question: Where is this coming from? Not just what platform or tool, but what specific human or institution originally produced this claim? A chatbot is not a source. It's a delivery system. Every fact worth keeping has a source upstream of the chatbot — a researcher, a study, a government agency, a journalist with a name. If you can't find that upstream source, the claim exists in a kind of informational free-float that should make you slower, not faster, to repeat it.

With younger kids, age eight to about eleven, "source" can start as a very concrete game. Where did you hear that? What did the chatbot say when you asked it where it learned that? (Try it — the answer is often revealing.) Is there a book, a website, a real person we can check? The goal at this age isn't skepticism for its own sake. It's building the habit of asking "where did this come from" before the habit of assuming "this must be right."

Motive is the second question: Why might someone want me to believe this? This is the question that feels awkward with AI, because AI doesn't have motives in the human sense. But the people who trained it did. The company that built the product does. The original authors of the content it was trained on did. Motive isn't always sinister — a company might want you to believe their product is safe because it actually is safe.

But the question still sharpens thinking. Who benefits if I accept this claim? Who benefits if I repeat it?

For teenagers, motive is where the real work happens. A fifteen-year-old sharing a deepfake video in a group chat, or a thirteen-year-old forwarding a "study" that confirms what they already believe about social media or diet or their school's terrible lunch policy — motive is the question that creates the pause. Not "is this false" (which requires work), but "who made this and what were they trying to accomplish?" That's a question anyone can ask in ten seconds.

Evidence is the third question, and the hardest: What is the actual proof? Not what someone says, but what independent, verifiable, replicated evidence supports the claim? For children, this doesn't mean demanding peer-reviewed journal articles before accepting any statement. It means understanding, gradually, that "the AI said so" and "here is the evidence" are two different things. Evidence has a texture — it can be measured, checked, repeated. Assertions are just sentences.

With kids around age ten to twelve, you can make this concrete by asking: "Is there one other place, a different place, that says the same thing?" That's not a perfect standard. Two sources can both be wrong. But it introduces the idea that claims need corroboration, that the chatbot saying something once isn't the end of the inquiry.

Practicing It Without Making Dinner Miserable

Here is the practical reality: you are not going to turn every AI interaction into a Socratic seminar. Nor should you. If your kid asks an AI what the capital of France is and then writes "Paris" on their worksheet, no framework is required. Some information is low-stakes, easily checkable, and not worth the friction of full epistemological review.

The teaching moments worth spending energy on are:

Claims that seem surprising. If your child comes to you with a fact that's interesting or counterintuitive, that's the moment. Not "interesting facts need to be challenged," but "interesting claims are worth five seconds of verification." Build the reflex around novelty.

Claims that show up in arguments. When a child uses AI output as evidence in a disagreement — with a sibling, with you, with a teacher — that's the moment to ask: do we actually know this? Not to win the argument. To model the habit of checking before citing.

Claims about people. AI systems make stuff up about real people at an uncomfortable rate. Biographical details, quotes, events. If your kid is using AI to research a person — a historical figure, a scientist, a living author — that's a category that deserves cross-referencing almost every time.

The conversations don't have to be long. "Where did it get that?" takes five seconds. "Can we find that somewhere else?" takes thirty. Over months, those seconds accumulate into something that looks like critical thinking — not as a procedure the child follows when reminded, but as a posture they've internalized. That's the goal.

A Case Study in Uncertain Evidence

Let me give you an example of source-motive-evidence that doesn't involve AI at all — and in doing so, model the kind of thinking this chapter is trying to teach.

Jonathan Haidt's 2024 book The Anxious Generation has reached more parents in the last two years than almost any other text on kids and technology. The core argument: the rapid spread of smartphones after 2012, combined with the rise of social media platforms designed for constant engagement, produced a measurable increase in adolescent anxiety and depression, particularly among girls. Haidt argues for phone-free schools and recommends parents delay smartphone ownership until high school and social media access until age sixteen.

Let's run the framework.

Source. Haidt is a social psychologist at New York University's Stern School of Business. He has published extensively in peer-reviewed journals. He is not a crank, and he is not making claims out of nowhere. The data he leans on most heavily — survey data showing rising rates of teen depression and anxiety beginning around 2012, correlating with smartphone adoption — is real. The rise in adolescent mental health problems is not in dispute. The timing correlation with smartphone proliferation is documented.

Motive. Haidt has written a bestselling book and built a public platform around this argument. That doesn't make him wrong, but it's worth noting that a compelling, unified narrative — screens did this, here's what to do — is more commercially successful than a messy, multi-causal account. He is also personally convinced he is right, which shapes how he presents evidence. None of that disqualifies the research. It's just context.

Evidence. Here is where the story gets genuinely complicated, and where honest engagement with the question requires sitting with some discomfort.

The timing correlation is real. The mental health trends are real. The question is whether smartphone adoption caused the mental health decline, or whether both happened simultaneously for reasons that partly overlap and partly don't. Psychologist Candice Odgers, then at the University of California Irvine, published a direct critique of Haidt's argument in Nature in 2024, noting that researchers who have actually tried to measure the effect of social media on individual teenagers' mental health over time consistently find much smaller effects than the population-level correlation suggests. The gap between the dramatic trend lines and the modest measured effects on actual kids is hard to explain away.

Amy Orben and Andrew Przybylski published a study in Nature Human Behaviour in 2019 that analyzed data from more than 355,000 adolescents and found that the effect size of social media use on well-being was approximately 0.05 — roughly equivalent, they noted, to the effect of wearing glasses, and considerably smaller than the effect of eating potatoes. That comparison is intentionally provocative. The point is that the population-level correlations that alarm us often shrink considerably when you look at individuals and control for other variables.

Haidt and his collaborators have pushed back on these critiques, arguing that averaging effects across all teens obscures harms concentrated in specific subgroups — particularly girls with preexisting vulnerabilities. That's a plausible rebuttal. It's also genuinely uncertain.

So where does this leave you, as a parent trying to make decisions?

The honest answer is: the evidence that heavy social media use is actively harmful to some teenagers is reasonably strong. The evidence that smartphones are the primary driver of a generation-wide mental health crisis, in a way that delayed phone access will reverse, is considerably more contested. Both things can be true without either being the complete picture. Haidt may be right about the conclusion while overstating the evidence. Critics may be right about the methodology while missing a real signal.

What this case study demonstrates is not that you should dismiss Haidt's work or dismiss his critics. It demonstrates that expert opinion — even opinion backed by data, even opinion presented fluently and passionately — requires the same source-motive-evidence scrutiny you'd apply to anything else. The fact that something comes from a credentialed person at a prestigious university does not exempt it from the question: what is the actual evidence, and is the conclusion it supports the only reasonable interpretation?

That is the thinking skill. Right there. Not cynicism. Not reflexive contrarianism. Genuine engagement with the claim, the quality of evidence behind it, and the honest acknowledgment that "it's complicated" is sometimes the most accurate answer available.

The Confident Machine Problem, Specifically

What makes AI a distinct challenge in all of this is the combination of two features that don't usually appear together: breadth and tone.

A single human expert is usually confident in their domain and more hesitant at its edges. A website with an agenda tends to be shrill or obviously slanted in ways that trigger suspicion. AI is confident everywhere, in the same register, about everything from the boiling point of water to the causes of the French Revolution to the correct dosage of over-the-counter medication. There's no tonal register change when it crosses from territory it "knows" well into territory it's confabulating. The cursor blinks. The next sentence arrives. It sounds the same.

Sam Wineburg, a historian at Stanford who has spent years studying how people evaluate digital information, has described what he calls "lateral reading" — the practice of leaving the page you're reading and opening tabs to check what other sources say about the source you're reading. Skilled fact-checkers do it instinctively. Most students, his research found, read vertically — going deeper into the same page, looking for internal signals of credibility. With AI, lateral reading is structurally harder. There's no page to leave. There's no byline to Google. The text is generated fresh each time.

This is why the framework — source, motive, evidence — has to be taught as a disposition rather than a checklist. You're not teaching kids to perform a verification routine. You're trying to grow in them a low-level

persistent awareness that the fluent, confident thing in front of them is always, always a starting point.

What This Looks Like in a Household That Actually Works

You're not looking for perfection here. You're looking for a slow accumulation of small habits.

Pick one question a week. When something AI-generated comes up — in homework, in conversation, on a phone — ask one of the three questions out loud. Not all three. Just one. "Where did that come from?" or "Who put that information into the world?" or "How would we check that?" Over a year, you've asked the question around fifty times. That's enough to matter.

Make it collaborative, not interrogative. The point is not to catch your child believing something false. It's to model a posture. If you do the questioning work out loud, in a spirit of genuine curiosity rather than gotcha energy, the child learns by watching you think. "Huh, I'm not sure about that — let me see if I can find the original study" is worth ten lectures about media literacy.

Allow yourself to be wrong. If you demonstrate that you check your own beliefs when you encounter new evidence, you give your child permission to do the same without shame. Changing your mind is the behavior you want to cultivate. Model it.

The world your kids are growing up in is built partly by algorithms that have no interest in whether the information they serve is accurate, and partly by AI systems that produce confident prose with no reliable relationship to truth. That's not a disaster. It's a context. Every generation has had to develop tools for evaluating the information environment they actually live in — not the one their parents navigated, not the one they were taught in school, the actual present one. Source, motive, evidence. It's not a complete answer. Nothing is. But it's a set of questions a nine-year-old can learn to ask and a sixteen-year-old can learn to ask harder, and the same questions work on a chatbot, a news article, a bestselling book, and a political speech.

Which is, honestly, what you want. Not kids who distrust everything. Kids who trust carefully, with their eyes open, and know the difference.

My daughter now knows it was Bell and the telephone. She also — and this surprised me when it happened — now asks "but where did they get that?" more than she used to. Not every time. Not even most times. But the question exists in her. It surfaces. That's the work: not building a wall between your child and a tool, but growing inside them the quiet, persistent habit of asking before believing.

Start with one question. Ask it this week. See what they say.

CHAPTER VI

AI in School: When It Helps, When It Gets in the Way

In January 2023, the New York City Department of Education banned ChatGPT on its school networks. Six months later, it reversed the ban. Los Angeles Unified did something similar. So did dozens of smaller districts in between — swinging from prohibition to cautious endorsement and back again, sometimes within a single school year, sometimes within a single building, classroom to classroom down the same hallway. Your child's math teacher may require students to show all AI-assisted work in a special color. The English teacher two doors down may consider it cheating to use AI at all. The history teacher on the third floor may not have a policy yet because she's been too busy to write one.

This is not exaggeration. This is Tuesday.

If you've felt confused about where schools stand on AI, the confusion is entirely appropriate. The institutions themselves are confused. And that uncertainty lands on kids — and on parents trying to help them — in ways that are genuinely unfair.

This chapter won't pretend there's a national consensus, because there isn't one. What it will do is help you understand which AI tools are showing real results in classrooms, which ones are quietly borrowing cognitive work that children need to do themselves, and how to walk into a back-to-school night and ask the right questions rather than the ones that put teachers immediately on the defensive.

What AI in School Actually Looks Like Right Now

The phrase "AI in school" covers an enormous range of things that don't resemble each other much. On one end, you have purpose-built tutoring tools with years of research behind them. On the other, you have a student pasting an essay prompt into ChatGPT at eleven o'clock at night. Both are "AI in school." The conversation parents need to be having treats these as very different situations.

The purpose-built tools include platforms like Khan Academy's Khanmigo, which launched in 2023 and works as a Socratic tutor — it asks questions rather than giving answers, nudging students toward the reasoning rather than performing it for them. Carnegie Learning's MATHia has been in classrooms since the late 1990s and has accumulated more peer-reviewed evidence than almost any other adaptive platform in existence. DreamBox (now part of HMH) adjusts the difficulty and sequence of math problems in real time based on how a student is working, not just whether they got the last answer right.

Then there are the general-purpose tools — ChatGPT, Google Gemini, Microsoft Copilot embedded in the tools students already use for assignments. These weren't designed for children and weren't designed for

learning specifically. They were designed to be useful to anyone with a question, which makes them extraordinarily capable and, in the wrong moment, extraordinarily easy to misuse.

Schools are also sitting with AI they didn't choose: the recommendation engines inside YouTube and TikTok that shape what students watch during "research," the writing-assistance features built into Google Docs and Microsoft Word, the AI summaries that pop up at the top of search results. Some of this AI is visible and discussed. Most of it isn't.

Where the Evidence Is Actually Solid

Before dismissing AI tutoring as hype or embracing it as revolution, it's worth knowing where the research points.

Start with a 1984 study by educational psychologist Benjamin Bloom. He called it the "2 sigma problem": in controlled experiments, students who received one-on-one human tutoring performed, on average, two standard deviations above students in conventional classrooms. That's the difference between an average student and one performing better than 98 percent of their peers. The problem, Bloom noted, was obvious — you cannot give every child a personal tutor. The question he posed was whether instructional methods could approximate the results of tutoring at scale.

Intelligent tutoring systems were an early answer to that question. A 2011 meta-analysis by cognitive scientist Kurt VanLehn, published in Educational Psychologist, reviewed decades of research on these systems. He found that well-designed intelligent tutors produced effect sizes of around 0.76 compared to conventional classroom instruction — smaller than human one-on-one tutoring, but still substantial. Not every system achieved this. The ones that did shared a common feature: they gave students feedback at the level of individual reasoning steps, not just right-or-wrong answers. They interrupted the process, not just the outcome.

Carnegie Learning's MATHia has been studied by RAND Corporation researchers, among others. The results are not uniformly spectacular — gains are meaningful but modest in some populations, stronger in others, and nearly always dependent on implementation quality. A school that deploys MATHia thoughtfully, with teachers who understand its reports and use them to adjust their instruction, gets different results from a school that uses it as a digital babysitter while the teacher grades papers.

What the research does support clearly is this: adaptive tools that respond to a student's specific errors, at the moment of the error, and guide the student through the underlying concept rather than just correcting the surface mistake — these tools help. Particularly for students who are behind grade level. Particularly in math. Particularly when a human teacher remains in the loop.

That last part matters a great deal.

Closing Gaps, But Not Automatically

One of the most compelling use cases for AI tutoring is what researchers call "personalization at scale" — giving students who are two grade levels behind in algebra the chance to work through gaps at their own pace, without the shame of slowing down a classroom, without a teacher who simply doesn't have time for twenty-eight individual catch-up conversations per period.

This is real. It works. And it's easy to see why schools serving high proportions of students with interrupted learning histories — from poverty, from the pandemic years, from anything — are drawn to it.

But the "closing gaps" story has a quiet asterisk. Adaptive platforms tend to be excellent at closing procedural gaps — the mechanics of solving a particular type of problem. They are much less reliable at building the kind of flexible, transferable understanding that lets a student encounter a problem they've never seen before and reason their way through it. Those are different skills. Schools that use AI tools to address the first while assuming the second will follow are running ahead of the evidence.

A good question to ask your child's school: "What does the teacher do with the data from this platform, and how does classroom instruction change based on it?" A school using adaptive AI well will have a specific answer. A school using it as a stopgap will give you a vague one.

The Cognitive Work Problem

Here's the part that gets less attention in press releases about AI-powered schools.

Learning is not the same as getting correct answers. Getting correct answers is, at best, a side effect of learning — and sometimes it's an obstacle to it.

Cognitive psychologists Robert Bjork and Elizabeth Bjork at UCLA have spent decades documenting what they call "desirable difficulties": conditions that slow students down and introduce friction actually produce stronger long-term retention and transfer. Retrieval practice — the work of pulling information out of your own memory, getting it wrong, trying again — is more effective than re-reading, according to a widely cited 2006 study by Henry Roediger and Jeffrey Karpicke in Psychological Science. The struggling is the learning. Not just a prerequisite for it. The thing itself.

Manu Kapur at ETH Zurich has pushed this further with what he calls "productive failure." In a series of experiments, he had students attempt complex problems before receiving any instruction on how to solve them. The students failed, of course — that was the point. Then they were taught the relevant concepts. Compared to students who received instruction first, the "productive failure" students showed significantly better understanding two and four weeks later. The frustration and confusion of not knowing had prepared their brains to actually absorb the lesson.

Think about what this means when a student reaches for ChatGPT the moment a problem gets hard.

Not because they're lazy. Because the tool is right there, and the discomfort of not knowing is genuinely unpleasant, and the AI gives them a pathway out of that discomfort in under ten seconds. The student gets the answer. The cognitive work that the confusion was supposed to trigger — the reaching, the connecting, the

failing and adjusting — never happens. The answer appears, gets copied, and leaves approximately nothing behind.

This is not theoretical. Peter Liljedahl, a professor at Simon Fraser University whose research on mathematics classrooms spawned a teacher movement called "Building Thinking Classrooms," spent years documenting how quickly students default to "answer-getting" behaviors rather than thinking behaviors when given the chance. That was before AI made the answer-getting option available instantly, at any hour, with full sentences already written. He now describes the challenge as qualitatively different. The default has become too easy to reach.

The honest framing here: AI tools in school exist on a spectrum from "holds the student's thinking" to "does the student's thinking." The first category has real evidence behind it. The second category produces the appearance of learning while its substrate quietly erodes.

The School Policy Mess, and What It Means for Your Family

By 2026, most districts have some version of an AI policy. The quality of those policies varies enormously.

On one end, you have blanket bans that treat all AI use as academic dishonesty and leave both teachers and students without any guidance on how to use tools that are already embedded in standard software. These policies tend to be unenforceable, inconsistently applied, and corrosive to trust — because students and teachers alike are quietly violating them constantly.

On the other end, you have districts that have embraced AI as an educational good without establishing any serious framework for when and how it should be used. Students in these schools may receive mixed signals: an assignment that requires AI-assisted brainstorming in one class, a zero for AI assistance in another, and a third class where the teacher uses AI to generate the quiz. The inconsistency is confusing, and confusion at scale produces anxiety.

Common Sense Media's research on digital learning environments consistently finds that students navigate AI policy conflicts primarily by doing what they observe adults around them doing — which, in many schools, includes watching teachers use AI to draft emails, create lesson materials, and prepare presentations while simultaneously telling students that AI use is prohibited. Children notice this. Teenagers notice it and resent it.

What does this mean practically? It means your child is operating in an environment with inconsistent rules, no clear moral framework attached to those rules, and models of adult behavior that contradict the rules anyway. This is not a reason to give up. It's a reason to build the framework at home, because school isn't reliably providing one.

Questions Worth Asking at Back-to-School Night

Teachers are not the enemy here. Most are genuinely trying to figure this out, with less time and fewer resources than the problem deserves. The best back-to-school night questions are curious ones, not accusatory ones. They invite a real conversation rather than a defensive one.

Here are several that tend to open things up:

"Does your class use any AI tutoring tools, and if so, what does the data from those tools tell you?" This distinguishes between schools that are using AI thoughtfully and ones that have deployed a platform without really knowing what to do with it.

"What's your classroom policy on students using general AI tools like ChatGPT for assignments, and how does that connect to the district's broader policy?" You're looking for consistency and clarity here. A teacher who has thought carefully about this will have a specific answer. Vague or hedging answers tell you something too.

"Are there assignments designed to teach students how to use AI well, and others where AI use is specifically off-limits? How do students know the difference?" This gets at whether the school is treating AI as something to be learned about or something to be managed around.

"What does 'productive struggle' look like in this class, and how does the school think about AI's effect on that?" This one is for the teachers and administrators who have done some reading. It signals that you're not just worried about cheating — you're thinking about deeper learning. Some teachers will light up. Others will look slightly puzzled, which is information.

You don't have to ask all of these. Two is fine. The goal isn't interrogation. It's establishing yourself as a parent who is engaged, informed, and interested in a genuine conversation — which is the kind of parent teachers will call when something important comes up.

When the School's Approach Doesn't Match Your Family's

This is where it gets delicate.

Suppose your child's school has adopted a platform that you've looked into and found concerning — maybe the data practices are opaque, maybe the evidence base is thin, maybe your eight-year-old is coming home with an "AI reading coach" and you're not convinced she should be spending forty minutes a day talking to a chatbot when she could be talking to a person. What do you do?

First: don't assume the school chose the tool carelessly. Most ed-tech purchasing decisions happen in committees, with vendor presentations and procurement processes that are more bureaucratic than they look from the outside. The administrator who selected the platform may have genuinely believed in it — and may have done more due diligence than you'd expect.

Second: ask to see the contract or the privacy agreement, specifically the data retention clauses and the provisions around training data. Schools are required under FERPA to maintain student data privacy. Many ed-tech vendors comply with the letter of this but not the spirit, and district contracts vary widely in what

they actually protect. You have a right to ask. Most schools, if they've done the work, will share it. If they haven't done the work, your question will prompt them to.

Third: if you have specific concerns about a tool your child is using, request that your child be given an alternative. This is not an unreasonable ask, especially for elementary students. "My daughter isn't going to use this chatbot for thirty minutes a day, but I'm happy for her to use the school's approved reading platform or work with the human aide" is a legitimate parental position. Not all schools will accommodate it, but most will engage with it.

What doesn't tend to work: public hostility to a teacher who is using a tool in good faith. The teacher probably didn't choose the platform and may have limited ability to change it. Direct conversations with curriculum directors and principals are usually more productive.

The Bigger Advocacy Frame

If your school is moving faster than you're comfortable with in one direction, or slower in another, the most effective thing you can do is not complain — it's participate. Parent-teacher organizations, curriculum advisory committees, and school board public comment periods are genuinely underused levers. Most AI policy decisions in K-12 districts are still being made right now, in 2026, by small groups of administrators who are desperate for input from parents who have actually thought about this.

Show up informed. Bring the questions from this chapter. Reference Common Sense Media's research on student data and digital learning — it's publicly available, readable, and carries credibility with educators. If you find a study on adaptive tutoring that cuts against the tool your school is adopting, bring it. Administrators respond to evidence offered in good faith better than they respond to outrage.

You're not trying to win. You're trying to have the school's AI decisions made thoughtfully, with your child's actual learning in mind. That's an interest you share with most teachers, even the ones who've been handed platforms they didn't ask for.

Learning That Sticks

There's a version of AI in school that genuinely excites me. It's the version where a ninth-grader who has always struggled with fractions gets to work through the underlying concepts at her own pace, with a system that notices when she's consistently confused at the same step and adjusts its approach, while her teacher uses the data from that system to design a small-group session on Wednesday that targets exactly that sticking point. It's the version where a sixth-grader who is bored out of his mind in a class set to the median gets pushed to work on problems two levels above his peers, without anyone making a fuss about it. It's the version where a student with dyslexia gets real-time reading support embedded directly into the text, without having to leave the classroom or announce her difference.

These are not fantasies. They are happening in schools with good implementation, good training, and educators who treat the technology as a tool they control rather than a program they follow.

The version that worries me is different. It's the version where cognitive work that needs to live inside a child's brain gets quietly outsourced to a machine — where the struggle that produces understanding gets short-circuited by an interface designed to make struggle feel unnecessary. That version produces students who are comfortable with AI and dependent on it in ways they haven't chosen and can't fully see.

The difference between those two versions is not the technology. It's the adults around the technology. Teachers who understand what learning actually requires. Parents who know the questions to ask. Administrators who treat evidence as more important than enthusiasm.

Your child spends roughly a thousand hours a year in school. What happens with AI in those hours matters. The conversation is already underway in your district, right now, without you — unless you decide to join it.

CHAPTER VII

Phones, Screens, and Sanity: A Shame-Free Reality Check

The guilt landed on a Tuesday.

My daughter was doing homework—or supposedly doing homework—while also texting, while also occasionally swiping something that was definitely not homework, while a YouTube video played at a quarter volume in the corner of her screen. I stood in the doorway for a good thirty seconds. Then I retreated to the kitchen and felt bad about myself for the rest of the evening. Not about her. About me. For not doing something. For not having a rule that prevented this. For scrolling my own phone while I was standing there judging her.

If that scene is familiar, this chapter is for you.

Let me start with something that is genuinely hard to say in a parenting book: you are probably not failing. The environment your child is growing up in is genuinely novel and genuinely difficult, and the research on how to handle it is legitimately unsettled. Anyone who tells you the answer is simple—ban all screens, or relax, it's all fine—is selling certainty they don't actually possess.

That said. There are real risks worth understanding, concrete steps that actually help, and a way of thinking about all of this that doesn't require you to feel terrible every time your kid picks up a phone.

What We Know, and What We Don't

Jonathan Haidt's The Anxious Generation, published in 2024, became the book that launched a thousand school phone bans. His central argument: something went badly wrong with adolescent mental health beginning around 2012, and the most likely cause is the mass adoption of smartphones combined with social media. He calls it "the great rewiring of childhood"—a shift from a play-based upbringing rooted in physical space and face-to-face contact to a phone-based one mediated by screens and platforms.

The data he marshals is striking. Rates of depression and anxiety among American teenagers—particularly girls—rose sharply from around 2012 onward. Emergency room visits for self-harm among adolescent girls increased by 62 percent between 2009 and 2015. Loneliness, according to multiple national surveys, climbed. Sleep declined. Haidt argues these trends correlate too strongly and too consistently across multiple countries to be coincidental.

This is not nothing. Those numbers represent real children having genuinely difficult experiences, and parents who have watched a daughter spiral into social anxiety or a son retreat from all real-world friendship into a screen-mediated substitute know this in their bones.

But here is where intellectual honesty requires a pause.

A number of developmental psychologists have pushed back on Haidt's causal conclusions—not on the trends themselves, but on the claim that smartphones and social media are the primary driver. Candice Odgers, a developmental psychologist at UC Irvine, published a pointed critique in Nature in 2018 and has continued arguing that the evidence for a causal link between smartphone use and mental health harm is far weaker than Haidt's framing suggests. Correlation is not causation. Adolescent mental health began declining before smartphones became ubiquitous in some datasets, and the pattern isn't uniform across all demographics or all countries.

Amy Orben and Andrew Przybylski conducted a large analysis of UK and US data and found that the statistical relationship between screen time and wellbeing was real but remarkably small—comparable in effect size, they noted with deliberate provocation, to the association between eating potatoes and wellbeing. Their 2019 paper in Psychological Science introduced what they called the "Goldilocks hypothesis": not no effect, but not a catastrophic one either.

What does all this mean for you as a parent? It means two things can be true simultaneously. The alarm Haidt is raising is worth taking seriously—even critics acknowledge that heavy social media use correlates with worse outcomes for some adolescents, especially girls, especially around social comparison and appearance. And it also means you should be skeptical of any single sweeping solution, because the evidence for harm isn't clean enough to support a single sweeping prescription.

You are making good decisions in genuine uncertainty. That is the honest position. So let's work with it.

The AI Layer Nobody Warned You About

Here is what is new in 2026 that wasn't true when most screen-time research was conducted: the engagement loops are no longer just well-designed. They are actively learned.

TikTok's recommendation algorithm doesn't just show your child things they already like. It tests. It serves short windows of different content categories, measures exactly how long your child lingers, calibrates to the fraction of a second they scroll past something, and then tightens the feed around whatever produced the longest fixation. YouTube's algorithm works on similar principles. Instagram's Explore page. All of them. The underlying technology is machine learning, and it is very good at its one job: keeping attention on the platform.

This is worth sitting with. The system on the other end of your child's phone is not passive. It is not a library of content waiting to be browsed. It is an optimization engine running continuously, learning your child's specific psychological pressure points, and serving content calibrated to those points. Boredom? Here is something outrageous. Loneliness? Here is a community that validates everything you already believe. Anxiety about appearance? Here is a comparison loop you cannot win.

Common Sense Media's 2023 Census found that teenagers report spending an average of 8.5 hours per day on screen media—not counting school-related screen use. Eight and a half hours. More waking hours than most adults spend at work. And the design of AI-driven recommendation systems means that the line between "choosing to use a platform" and "being optimized toward continued use of a platform" is genuinely blurry in ways earlier media was not.

None of this means your child is passive or helpless. But it does mean the conversation about screen time has changed. A television in 1995 couldn't learn. The thing in your child's pocket is trying to.

Sleep: The One Thing That Actually Isn't Controversial

Set aside the larger debates for a moment. There is one corner of the screen-time literature where the evidence is solid. Sleep.

Multiple studies—including large-scale surveys by Common Sense Media and peer-reviewed work published in Pediatrics and JAMA Pediatrics—show that adolescents who use screens in the hour before bed sleep less and sleep worse. The mechanism isn't just behavioral, kids staying up late to finish a video. Blue-spectrum light from phone and tablet screens suppresses melatonin production, shifting the body's internal clock toward a later sleep phase. For teenagers, who are already physiologically prone to delayed sleep cycles, this is a double hit.

Jean Twenge's analysis of data from over 360,000 adolescents found that those who spent five or more hours a day on their devices were 66 percent more likely to have at least one risk factor for suicide compared with those who spent one hour per day. That is an association, not a proven cause. But one of the mediating variables in that association is sleep—teenagers who use devices heavily sleep less, and sleep deprivation independently drives mood dysregulation, impulsivity, and emotional reactivity. You don't need to resolve the entire Haidt debate to decide that a thirteen-year-old running on six hours of sleep because her phone was in her bed is a problem worth solving.

The good news: this is actionable.

Charge the phone outside the bedroom at night. Not on the nightstand. Not under the pillow. Outside the room—on a kitchen counter, in a hallway. This one change, reported by parents consistently, tends to add thirty to forty-five minutes of sleep. The first week is rough. After that, most kids adjust. It is the single highest-return, lowest-complexity intervention available to you right now.

Household Norms Without the Power Struggle

Here is the trap many parents fall into: they try to install rules like software patches. New policy, effective immediately, enforced by confiscation. Kids experience this as arbitrary and punitive, reliably respond with resentment or workarounds, and the household enters a low-grade cold war over the phone that exhausts everyone.

The alternative is not no rules. It is rules with rationale, and co-design where possible.

What works better, according to family research and the accounts of parents who have been through this: norms established in conversation, with the kid present, ideally before the issue becomes a crisis. If your child is twelve and you are just getting them a smartphone, you have a window. Use it. Sit down and talk through when the phone is available, where it lives at night, which apps require a conversation before downloading, and what happens if the agreements aren't kept. Write it down. Not as a contract to be enforced. As a reference to return to.

If your child is fifteen and you are trying to retrofit norms into an established pattern, it's harder. Not impossible. Start with the least contested ground—bedrooms at night, dinner table, family time—and hold that line consistently before trying to renegotiate broader use. Pick one thing, do it well, then add more when trust is established.

Some concrete norms that families consistently report working:

Phone-free dinner. Every person at the table, including parents. This is a credibility question as much as a child-behavior question. If you are checking email during dinner, you have already lost the argument.

Morning delay. No phone for the first thirty minutes after waking. This is the AI-optimized feed's prime hunting hour—groggy, half-awake, not yet in possession of full critical faculties. Starting the day with something else, even just eating breakfast, even just sitting, matters more than it sounds.

Designated scroll time. Rather than a total-hours limit, which is nearly impossible to monitor and creates constant negotiation, some families find it more useful to designate specific windows as acceptable for open-ended browsing. After homework, before eight o'clock. The goal is predictability, not surveillance.

Bedrooms stay phone-free at night. Already covered. Worth repeating. This one does the most good.

What about screen-time monitoring apps? They exist, they work technically, and they create a surveillance dynamic that many adolescent development researchers flag as counterproductive. Teenagers who experience heavy parental monitoring—digital or otherwise—without corresponding conversation and trust tend to get better at hiding things, not better at managing things. The goal is not a child who cannot access a platform. It is a child who can make reasonable decisions about what they engage with. Monitoring doesn't build that. Conversation does.

Working Parents and the Reality of Gaps

The parenting advice industry has a working-parent problem.

An enormous amount of screen-time guidance implicitly assumes a parent who is physically present and alert and available from 3pm onward. "Be there when they come home from school." "Eat dinner together every night." "Monitor what they're watching in real time." These are not unreasonable ideals. But for a parent working a second shift, for a single parent managing three kids and a commute, for two parents both on deadline every Tuesday—this advice lands as one more thing you are not doing well enough.

You cannot monitor every swipe. You should not have to. What you can do is invest in the conversations and the structural changes that persist even when you are not in the room. A phone that lives outside the bedroom at night doesn't require you to be awake at midnight to enforce it. A household norm that the dinner table is phone-free applies even on the evenings when dinner is a pizza on the couch because everyone is exhausted.

The goal is not constant supervision. It is an environment that is slightly less optimized against your child's wellbeing than the default. That is a realistic and achievable target. Most of the changes described in this chapter cost you thirty minutes of conversation and some stubbornness, not constant vigilance.

The Gender Question

Haidt's argument is particularly pointed when it comes to girls. He—and researchers like Jean Twenge—argue that social media causes specific harms to adolescent girls through social comparison, appearance-focused content, and relational aggression that is easier to sustain and harder to escape when it's phone-based.

This deserves honest engagement because it shapes what you actually pay attention to.

The evidence that girls experience worse mental health outcomes from heavy social media use is more consistent than the evidence for boys, across multiple studies and multiple countries. Instagram and TikTok's emphasis on appearance, performance, and social validation creates particular feedback loops for girls navigating the exact developmental territory where identity and self-worth are being constructed. This is not speculative. It is the finding that holds up most reliably across the research.

For boys, the concerns are different. Haidt points to the displacement of other activities—sport, outdoor time, real-world socializing—by gaming and video content, with associated effects on motivation and in-person social skill development. The evidence here is somewhat more mixed. But the concern that eight hours a day of AI-optimized content displaces time that might otherwise go toward doing hard things, physical things, uncomfortable things—that concern doesn't require ironclad causal proof to take seriously.

The practical takeaway: pay attention to what your specific child is doing on their phone, not just how long they're doing it. An hour on YouTube watching someone build a guitar amplifier is not the same hour as sixty minutes of appearance-comparison reels. Volume matters. So does content, and the emotional state your child is in when they put the phone down.

What "Reducing Screen Time" Actually Requires

Here is the part nobody wants to say: reducing phone use doesn't work as subtraction alone. You cannot just take away eight hours of dopamine-optimized content without replacing it with something. Boredom is fine—genuinely, developmentally important—but unstructured boredom in a house where the phone is still technically available rarely wins.

What tends to fill the space, according to parents who have navigated this transition well, is physical stuff. Cooking something. Building something. Playing an instrument badly for a while. Sport. An after-school job at sixteen. Friends coming to the actual house to do something in person. The research on what protects adolescent mental health—even Haidt's critics largely agree on this—centers on face-to-face relationships, physical activity, and activities that produce a sense of mastery and competence.

None of that requires being an anti-technology zealot. Your kid can have a phone and also play a sport and also cook dinner once a week and also have a friend over on Friday. The phone does not have to be the enemy. It just cannot be the only thing.

Having the Conversation Without Lecturing

The most counterproductive thing you can do with a teenager is launch a prepared presentation about screen time. They have already clocked that it's coming. Their face closes before you finish the first sentence.

What works better is shorter, more frequent, and less formal. A question in the car. A comment that opens rather than closes. "Did you see that thing about TikTok changing its algorithm? I was reading about how it tracks—" and then you stop talking and they either engage or they don't, but you have opened a door without standing in front of it.

Ask about what they're watching, not as surveillance but as genuine curiosity. "What's actually good on YouTube right now?" is a different question than "How long were you on YouTube?" One invites them to be the expert. One positions them as a defendant.

And model what you're asking for. This is uncomfortable to hear because most of us have complicated relationships with our own phones. But children—especially adolescents who are exquisitely sensitive to adult hypocrisy—notice the gap between what you say and what you do with extraordinary precision. If you want your child to put the phone down at dinner, you have to put the phone down at dinner. Every time. Not most times.

When to Actually Worry

Most screen use, most of the time, is not an emergency.

But there are patterns that warrant more than a gentle conversation. Watch for a sudden, dramatic withdrawal from in-person relationships in favor of online interaction—not a preference, a retreat. A child who becomes visibly dysregulated when separated from a device, not just irritable but genuinely destabilized. Sleep dropping below seven hours for an extended stretch. Any indication that a child is in a community online that is reinforcing harmful beliefs about themselves or others. Signs of image-based harassment—your child receiving or being involved in sharing manipulated or explicit imagery. That last one is covered in the next chapter, but it deserves naming here.

If something about your child's relationship with their phone feels qualitatively different from "distracted teenager," trust that instinct. The answer may not be a screen-time rule. It may be a counselor, a pediatrician, a candid conversation with another parent, or a harder look at what is actually happening in your child's social world—which in 2026 is partly but not entirely on the screen.

The Honest Summary

There is no clean answer here. Anyone who tells you there is hasn't looked at the research.

What the evidence supports, carefully and specifically: teenagers who sleep less do worse. Teenagers who have less face-to-face connection do worse. Teenagers who spend significant time on appearance-comparison platforms show worse body image and mood outcomes, especially girls. AI-powered recommendation systems are genuinely better at capturing attention than earlier media, and that difference matters.

What the evidence does not cleanly support: a specific hour-count cutoff that fixes the problem, a single app uniquely responsible for the harm, or the conclusion that every hour on a screen is equivalent in damage.

So you make reasonable bets. Phones out of bedrooms at night. Dinners without phones, including your own. Conversation about what they're watching and how it makes them feel. Physical activity that has nothing to do with a screen. Actual boredom, occasionally, that you don't rush to fill.

And you let go of the guilt. You are parenting through a technological shift that nobody has parented through before, with research that is still catching up to the reality. You are doing it while working and feeding people and managing everything else a modern household requires.

The fact that you are here, reading this, means you are paying attention. That is not nothing. That is, honestly, most of what this requires.

CHAPTER VIII

Talking to Your Kids About Deepfakes

In October 2023, parents in Westfield, New Jersey, got a call no one had rehearsed for. Male students at the high school had used an AI image generator to create fake nude photographs of female classmates. The girls were real students — with real names, real faces, real futures. The images were fabricated from scratch. The damage was not. Families called police. The school held emergency meetings. Attorneys got involved. And parents across the country who saw the news coverage thought, in some version or another: I had no idea that was even possible.

It is. And it has been for a while now.

Deepfake technology is not a looming threat on the horizon. It arrived several years ago, went mainstream quietly, and is now cheap enough that a motivated teenager with a laptop and forty minutes can produce a convincing fake video of a classmate, a public figure, or a parent. The tools are free or close to it. The barrier is nearly gone. What has not arrived yet — not in most households, anyway — is the conversation.

That conversation is what this chapter is for.

What a Deepfake Actually Is (and How It Gets Made)

You do not need a computer science degree to understand this. The simplified version is enough to be useful.

A deepfake is a piece of media — a photograph, a video, an audio recording — that has been artificially generated or manipulated using AI to show something that did not happen, or to put someone's face or voice onto content they had no part in creating. The word itself is a mashup of "deep learning" (a type of AI technique) and "fake," coined by an anonymous Reddit user around 2017. What started as an academic curiosity moved into film studios, then into the hands of anyone with an internet connection.

The underlying mechanics vary, but here is the core idea. These systems — particularly ones using what are called generative adversarial networks, or more recently, diffusion models — are trained on enormous amounts of real images or video. The AI learns, in a statistical sense, what a human face looks like, how it moves, how light falls across it. Once trained, it can generate a new face that has never existed, or it can swap an existing face onto existing video with enough accuracy to fool a casual viewer.

For voice cloning, the same logic applies. Feed a system three minutes of someone's recorded speech — a podcast appearance, a YouTube video, a voice note — and it can generate new audio in that person's voice saying words they never said.

The tools available in 2026 have become disturbingly accessible. Some require no technical setup whatsoever. A user uploads a photo, selects a body type, clicks generate. That is it. Researchers at Stanford's

Internet Observatory and similar groups have documented the rapid proliferation of these services — most hosted outside U.S. jurisdiction, most with no meaningful age verification, many explicitly advertising themselves as tools for creating non-consensual intimate imagery.

This is not a fringe use case. It is, by several independent assessments, the dominant use case.

Two Separate Problems That Often Get Tangled

When parents hear "deepfake," they usually think of one of two things: either political propaganda (fake speeches, fake news clips) or the kind of explicit imagery that showed up in Westfield. Both are real problems. They tend to get conflated because they share a technology. But they are genuinely different issues — they affect different kids in different ways, and they need different conversations.

The personal harm is image-based sexual abuse. Synthetic intimate images of real people, created and shared without their consent. This is not a problem that only happens to celebrities. Researchers studying image-based abuse have documented cases in middle schools. Jonathan Haidt's work in The Anxious Generation does not specifically address deepfake imagery — his focus is on social media's structural effects on adolescent mental health — but his underlying observation applies here without much stretch: teenage girls are navigating a social environment in which their image, their reputation, and their sense of safety are profoundly shaped by what other people can publish about them online. A fake nude photograph circulated in a school group chat activates every social terror that was already wired into adolescent development. Rejection, humiliation, the sense that the violation is permanent and public and unstoppable. The harm is acute. In several documented cases, it has been catastrophic.

Common Sense Media's ongoing research on teen technology use shows consistently that girls, in particular, report higher rates of anxiety tied to what others post or share about them online. Deepfake imagery is the logical and ugly extension of that environment.

The civic harm is synthetic misinformation. In January 2024, a robocall went out to Democratic voters in New Hampshire before the primary — AI-generated audio cloned from President Biden's real voice, telling voters to stay home. It was fake. The New Hampshire attorney general's office confirmed it quickly. But it spread before it was corrected, and the people who created it were counting on exactly that. That same month, explicit AI-generated images of Taylor Swift flooded the platform then called X, reaching tens of millions of views before the platform managed to suppress them. Neither of these events was complicated to produce with the tools available.

For teenagers who are forming political opinions and sharing content, this matters. A lot. A sixteen-year-old who shares a deepfake video of a politician saying something outrageous is not necessarily malicious — they may genuinely believe it is real. The speed at which synthetic content travels, combined with the confidence with which AI produces it, creates a specific kind of civic vulnerability. We have spent decades teaching kids to be skeptical of anonymous sources. We have not caught up on teaching them to be skeptical of their own eyes.

Talking to a Ten-Year-Old

Age ten is not too young for this conversation. It is, in fact, close to the right time — before the social environment gets complicated enough that the topic feels personal and charged. At ten, most children can absorb a framework without feeling accused or implicated.

Keep it concrete and grounded. Abstract warnings about "fake things on the internet" bounce off kids this age. What sticks is specificity.

A conversation might start like this:

"Have you ever seen a video online where someone does something really surprising or funny? Something that seems almost too perfect?"

Let them answer. They usually have something. Then:

"Some of those videos are real. But some of them — more and more of them — are made by computers. Not edited, like cutting clips together. Actually generated. The computer learned what a person's face looks like and what their voice sounds like, and then it made a brand-new video of them doing or saying something they never did."

At this point, most ten-year-olds have one of two reactions: fascination ("wait, how?") or mild alarm ("so videos can just be fake?"). Either is a good opening. Follow it with something they can actually do:

"It is getting really hard to tell which is which, even for adults. So I am going to teach you three questions to ask before you believe a video or share one. First: who made this, and how would they know? Second: does another trustworthy place say the same thing? Third: does this video seem designed to make me feel something really intense, really fast?"

That last question is the most useful one for a child this age. Deepfake content — like a lot of viral misinformation — is engineered to produce an emotional spike. Outrage, shock, embarrassment, hilarity. The emotion is the point. Teaching kids to notice the spike, to feel it and then slow down before acting on it, is genuinely transferable to every form of media they will encounter.

You do not need to explain non-consensual intimate imagery to a ten-year-old in any detail. But it is reasonable to plant one seed:

"Computers can also make fake pictures of real people — your face on someone else's body, or doing something you never did. If anyone ever sends you a picture of a person that seems wrong or weird, or if someone shows you a picture that claims to be of someone you know, tell me before you share it. Some of those pictures are designed to hurt people."

Keep the tone matter-of-fact, not alarming. You are not trying to terrify a fifth-grader. You are giving them a reflex: pause, ask, tell a trusted adult. That reflex takes years to build, and ten is not too soon to start.

Talking to a Sixteen-Year-Old

This conversation is harder. Not because the information is more complex — it is, but teenagers can handle complexity — but because it is likely to feel more personal. A sixteen-year-old may know someone who has been targeted by synthetic images. They may have seen deepfake content on social media without identifying it as such. Some may have, without fully understanding the implications, shared something that turned out to be fake. A small number may have created something themselves.

Start without accusation. Seriously. An interrogating tone closes this conversation before it opens.

One entry point is a recent news story. The Taylor Swift images in January 2024, for instance, were widely covered and are still well-known enough to reference without it feeling manufactured. Or use a school-specific event if one is recent — many districts have dealt with versions of what happened in Westfield, and if yours has not yet, one in your state probably has.

"I saw something about a school where some boys used AI to create fake intimate images of girls in their class. I want to talk about that, not to lecture you, but because I think the technology involved is something every person your age needs to understand."

Then explain the technology with a little more depth than you would for a ten-year-old. Teenagers respond to being treated as capable of handling real information. The diffusion model explanation — that AI learns what a human body looks like and then generates new images from statistical patterns — is something a sixteen-year-old can follow, and understanding the mechanism actually helps with skepticism. Once you know how the sausage is made, it is easier to distrust the packaging.

Then the harder part of the conversation:

"Creating a fake intimate image of a real person — even if the image is entirely generated, even if no real photograph of that person was used — is in many places now illegal. In some states it is a criminal offense. The person who makes it can be charged. So can people who share it. I need you to know that not because I think you would do this, but because I think the bar feels lower than it is when you're looking at an app that makes it seem easy."

Several U.S. states — California, Virginia, Texas, Georgia, among others — have passed legislation specifically targeting non-consensual synthetic imagery. The UK's Online Safety Act of 2023 made sharing deepfake pornography a criminal offense. The legal landscape is still uneven and evolving, but "I didn't think it was that serious" is no longer a functional defense.

Ask them what they actually see. This is the part most parents skip.

"What does this kind of stuff look like on your feeds? Do you see it? Do people in your school talk about it?"

Then listen. Actually listen, without pivoting immediately to concern or advice. What teenagers describe about the AI content they encounter — deepfakes of celebrities, AI voices doing impressions, obviously synthetic political content they scroll past without giving it much thought — tells you more about their media

environment than any survey will. And asking the question communicates something important: you are not interrogating them, you are trying to understand their world.

End with two practical things.

First: "If you ever receive a fake image or video of a real person — especially a classmate — do not share it. Even if it seems like a joke. Especially if it seems like a joke. Send it to me or delete it."

Second, and this is the one that tends to land hardest: "Your face, your voice — that is yours. If you post videos of yourself online, your data is out there. I am not saying don't do it. I am saying know that it exists. Anyone with enough of your images can, in theory, put them somewhere you didn't agree to."

That last piece connects to the digital footprint conversation that has existed for years. What is new is the stakes. It is not just that an embarrassing post might affect a college application. It is that your image itself has become raw material.

When Something Goes Wrong

Preparation matters. But so does knowing what to do when the preparation wasn't enough.

If your child comes to you because they've discovered a deepfake of themselves — or a friend — circulating in a group chat or on a platform, the order of operations matters.

First: don't escalate the spread. Do not share the image with other parents or family members to show them "what happened." Every additional view, every forward, extends the reach. Screenshot it for documentation purposes — you will likely need evidence — but do not distribute it.

Second: report to the platform. Every major platform — Instagram, Snapchat, TikTok, X — has a reporting mechanism for non-consensual intimate imagery. They vary in how quickly they respond, and honestly, the track record is uneven. But filing the report creates a record and triggers at least a review process. Some platforms have dedicated takedown teams for this category specifically.

Third: contact the school. If the image involves classmates, the school needs to know. This is true even if your child is the subject of the image, and it is true even if they would prefer you didn't. School administrators have been dealing with this often enough now that most have protocols, even if imperfect ones. An administrator can intervene in ways you cannot — with access to school devices, communication channels, and disciplinary processes.

Fourth: consider a police report. This feels like a dramatic escalation, and your child may resist it. But depending on your state, the creation and distribution of synthetic intimate imagery of a minor is a serious criminal matter. A police report, even if no charges are ultimately filed, creates an official record. That record matters if the situation escalates or if you later decide to pursue civil action.

Fifth: find your child support that is not from you. This is not a failure of parenting. It is a recognition that the person who was violated needs to talk to someone who is not also managing their own distress about it. A

school counselor, a therapist, a crisis text line — Common Sense Media's resource library includes guides for connecting families to support services after digital harm events. The emotional aftermath of image-based abuse is real and documented. Dismissing it as "just a fake picture" does measurable harm.

If the deepfake is civic misinformation — your teenager shared a fake video of a politician, genuinely believing it was real — the conversation is different but still worth having. Help them understand what happened without weaponizing it into a lesson about gullibility. The tools are built to fool people. That is the point. What matters is developing the habit of pausing before sharing, and building a higher bar of evidence for emotionally charged content.

The Ongoing Work

There is no single conversation that finishes this topic. Not because parents are bad at it, but because the technology changes fast enough that whatever specific examples you use today will be outdated within a year. The Biden robocall that seemed shocking in early 2024 was, by the end of that year, a relatively low-quality example of what was being produced.

What you are actually building, through every age-appropriate version of this conversation, is not knowledge about a specific technology. It is a disposition. Specifically: the willingness to distrust confident media. To feel the emotional spike and pause before acting. To ask who benefits if I believe this? To say, when something seems designed to provoke you, I want to look at this more carefully before I react.

That disposition is ancient. Rhetoric teachers in ancient Athens were essentially teaching the same thing about persuasive speeches. What is new is the scale, the speed, and the fact that the machines producing this content do not need anyone's permission and do not stop.

Your child will encounter deepfakes. Almost certainly already has. The Westfield story made national news, but the less dramatic versions — fake voices doing cruel impressions, synthetic images used as social weapons in group chats, viral political clips that were fabricated — happen constantly and mostly go unreported. You cannot prevent the encounter. You can influence what happens in the three seconds after it.

Those three seconds are the whole game. And you can help.

CHAPTER IX

The Things AI Cannot Replace: Raising Humans in Full

In 1976, when Texas Instruments started selling the TI-30 calculator for $24.95, school boards across the United States convened emergency meetings. Editorialists wrote op-eds with titles like "The Death of Arithmetic." The worry was sincere: if children could simply press a button to multiply 47 by 63, would they ever truly understand numbers? Would the basic cognitive muscles atrophy?

What happened instead was interesting. Yes, students stopped drilling long multiplication the way their parents had. But freed from three-hour computation sessions, they moved into algebra and statistics earlier. SAT math scores didn't collapse. Enrollment in calculus climbed. The feared extinction didn't arrive — and a new capability, estimating and checking the reasonableness of machine output, quietly became more important than anyone had predicted.

This is the pattern. Almost every major technological shift produces a genuine anxiety about what's being lost, and the anxiety is usually partly correct and mostly wrong. The telegraph was supposed to make people unable to write letters. Television was going to kill imagination. The internet would end deep reading. Each time, some things were genuinely displaced, some surprising new capacities emerged, and the species muddled through.

Which does not mean the anxiety about AI is overblown. AI is not a calculator. It doesn't just compress one task; it threatens to compress many at once — writing, analysis, research, conversation, even emotional support. The displacement is wider and weirder. But the parenting response that actually helps isn't different in kind from the one that helped in 1976. It's still: figure out what's genuinely irreplaceable, and then invest in those things on purpose.

This chapter is a permission slip to do exactly that.

What the Machine Does Well — and Where That Leaves Us

Let's be honest about what large language models are good at, because the answer shapes everything else. They are very good at pattern completion. They've read, in training, somewhere in the neighborhood of a trillion words of human text, and they've become extraordinarily skilled at producing text that fits the patterns in that corpus. Ask for a five-paragraph essay on photosynthesis and you'll get one that looks exactly like every five-paragraph essay ever written on photosynthesis. It will be grammatically clean, logically structured, and have the nutritional value of a styrofoam cup.

What AI cannot do — genuinely cannot, not as a matter of current limitation but as a matter of what the architecture is — is have an experience. It has never been frustrated by a proof it couldn't crack for three days and then felt the odd elation when it finally clicked. It has never burned a batch of cookies and learned from the smell. It has never watched a friendship dissolve and felt the specific loneliness of that. It generates language about all of those things with impressive fluency. The fluency is not the same as the experience, and the experience is where human development actually lives.

That's not a mystical claim. It's a practical one. The capabilities that matter most in an AI-saturated world are the ones that require having a body, a history, other people, and time. Four of them in particular are worth naming directly.

Sustained Attention

The developmental psychologist Maryanne Wolf at UCLA has spent decades studying what happens when children read deeply — not scanning, not skimming, but genuinely working through a demanding text. Her 2018 book Reader, Come Home laid out what the research suggests: deep reading builds specific neural circuits for inference, analogy, critical analysis, and empathy. Those circuits don't form automatically. They form through practice, through sitting with difficulty, through the friction of words that don't yield immediately.

AI makes it very easy to skip that friction. Not just for homework. For everything. Confused by a news story? Paste it in and ask for a summary. Don't understand a concept in a video? Ask the chatbot. Can't figure out the plot of a book? There are dozens of services now that will give you the beats in bullet points. The friction isn't a bug in learning; it's often the mechanism. Remove it entirely and you remove what the struggle was building.

Common Sense Media's 2023 media census found that American tweens and teens spend an average of 8 hours and 39 minutes daily on screens — more than a full working day, not counting school. That number gets cited a lot, sometimes in ways that overstate the damage. Not all screen time is equal; a fifteen-year-old doing an hour of focused coding isn't in the same category as a fifteen-year-old cycling through TikTok for ninety minutes at midnight. The quantity matters less than the pattern of attention underneath it.

What Jonathan Haidt's The Anxious Generation gets right — and his critics, including developmental psychologists like Candice Odgers, concede this much — is that the smartphone era coincided with a documented shift in how adolescents spend their free time. Less time doing things that require sustained focus: reading, playing instruments, building, having long unscripted conversations. More time in rapid-fire, algorithmically served media. Whether smartphones caused the mental health shifts Haidt charts is genuinely disputed in the literature; the correlation is real, the causal arrow is harder to establish. But the attention pattern shift is less contested.

What this means practically: the single most valuable thing you can do for a child under eighteen has nothing to do with any particular app or content restriction. It's protecting pockets of time where nothing is happening

fast. A novel that requires two weeks to finish. A puzzle that has no hint function. A long car ride with no headphones where someone eventually has to say something out loud. Boring, slightly uncomfortable stretches where the brain has to entertain itself.

This isn't anti-technology. It's pro-capacity. The teenagers who will do well in a world saturated with AI outputs are the ones who can read a long, difficult document and form an independent opinion about it. That skill doesn't come from consuming content. It comes from the slow work of sustained attention, practiced over years.

The Hands Know Something

There's a line from the philosopher Matthew Crawford's Shop Class as Soulcraft that has stayed with me since I first read it: "If the goal is to get kids to think, it helps to give them something concrete to think about." Crawford was writing about manual education, about what's lost when schools strip out wood shop and home economics in favor of abstract test preparation. But the argument lands harder now than when he wrote it in 2009.

Physical making teaches something that can't be conveyed through text or video or even careful AI explanation. When a twelve-year-old is learning to throw a pot on a wheel, the clay gives immediate, tactile, unambiguous feedback. Too much pressure and it collapses. Not enough centering and the whole thing wobbles. No chatbot can replicate that loop. The hand learns from the clay, not from a description of the clay.

This matters for development in ways the research supports. A 2019 study in Frontiers in Psychology found that manual craft activities — knitting, woodworking, drawing, cooking — were associated with higher reports of calm, engagement, and purpose in adolescents, effects that persisted even when controlling for overall activity level. The mechanism seems to involve both the reward pathways activated by completing something tangible and the regulated, focused attention that making requires. You can't zone out while sawing a board to length. The saw will let you know.

Cooking is probably the most accessible version of this for most families, and also the most underrated. Teaching a ten-year-old to cook a real meal — not heat up a frozen pizza, but actually read a recipe, time multiple components, taste and adjust — is a dense developmental exercise. It involves math, sequencing, sensory calibration, failure recovery, and the specific satisfaction of feeding another person. None of those things can be outsourced to an algorithm. All of them compound over time into genuine capability.

The same logic applies to sport, to music, to building, to gardening. What these activities share is irreducible embodiment. The skill lives in the body and can only be built through repetition in the body. A child who grows up cooking, playing an instrument, and doing some kind of physical training is building a layer of competence that AI literally cannot touch — because the competence isn't stored in a document. It's stored in muscle and reflex and the quiet confidence of someone who has made things with their hands.

This is not nostalgia. It's strategy.

Empathy Through Friction

Here is something that sounds gentle but is actually quite demanding: empathy is not primarily a feeling. It's a practice. And the practice requires exposure to people who are genuinely unlike you, in situations where you cannot simply disengage when it gets uncomfortable.

AI companions are engineered to be agreeable. Character.AI, Replika, Snapchat's My AI — all of them are designed, at the architectural level, to maintain engagement. Conflict reduces engagement. Disagreement reduces engagement. The AI learns to give you more of what keeps you in the conversation. For a teenager who is socially anxious or lonely, that's genuinely soothing in the short term. It's also a little like eating sugar when you're hungry: it removes the feeling of need without providing any nutrition.

Real friendships are friction. Friends have bad days and take it out on you. They make plans and then cancel. They like things you find boring and expect you to come along anyway. They are wrong sometimes and have to be told. They tell you when you're wrong. All of that friction — all of that low-grade social difficulty — is the resistance training for empathy. The capacity to model another person's experience, to hold your own preferences lightly enough to actually hear someone else, builds through thousands of small awkward interactions. There is no shortcut.

Jean Decety at the University of Chicago, who has spent decades imaging the empathic brain, has noted that empathy develops through direct social experience and that adolescence is a critical period for that development. What his lab has found, consistently, is that the brain systems involved in understanding others' emotional states are plastic — they grow or they don't, depending on the quality of social experience they receive.

This doesn't mean children should be thrown into social situations they're not ready for, or that AI companionship has no role at all. For some kids with social anxiety, a low-stakes chatbot can be a practice space before the real thing. But the practice space is not the destination. And if a child is spending five hours a day talking to a character on Character.AI instead of navigating a friendship, the trade-off is worth examining honestly.

What parents can do here is less about restriction and more about investment. Investing in situations where kids have to figure it out with real humans: family dinners where people actually talk, unstructured playdates without a screen in reach, the annual awkward week at camp where your kid doesn't know anyone and has to become known. Those things feel optional. They are not optional. They are the training ground for the most important cognitive-emotional capacity humans possess.

Tolerating Not Knowing

There is a particular discomfort that AI has made newly easy to escape: the discomfort of sitting with a question that doesn't have an immediate answer.

I don't mean big existential questions, though those too. I mean the smaller version — reading something confusing and staying with it long enough to figure it out. Encountering an unfamiliar word and guessing at its meaning from context before looking it up. Watching a historical documentary and letting yourself be curious about something without immediately resolving it.

The ability to tolerate ambiguity and uncertainty without immediately reaching for resolution is correlated, in the research literature, with some of the outcomes we most want for children: creative problem-solving, intellectual persistence, and what psychologists call "epistemic humility" — the capacity to hold beliefs with appropriate levels of confidence rather than seizing on the nearest available answer.

A 2022 study published in Thinking & Reasoning found that tolerance for ambiguity was one of the strongest predictors of open-minded thinking in adolescents, stronger than IQ scores and almost as strong as general knowledge. The uncomfortable truth embedded in that finding: tolerance for ambiguity is not a trait you're born with in fixed amounts. It's a practiced response. And the more frequently a child reaches for instant resolution — from a search engine, from a chatbot, from an algorithm that's already guessed what they want — the less practice they get in tolerating the uncertainty they don't resolve.

This is specific and practical. It means that letting a nine-year-old struggle with a confusing paragraph for seven minutes before offering help isn't cruelty. It's skill-building. It means that a car ride where your kid says "why do clouds have different shapes" and you say "I genuinely don't know, what do you think?" is not a parenting failure. It's modeling the thing you most want them to develop. It means that boredom — real, honest, nothing-to-do boredom — is not a problem to be immediately solved. It is, as developmental psychologists including Stuart Brown have argued for decades, one of the primary drivers of imaginative thought in children.

The word "boredom" has almost entirely negative connotations now, which is interesting given how much of human creativity has emerged from it. Darwin was bored on the Beagle. Tolkien invented Middle-earth in the margins of exam papers he was grading. Children who have never been bored have never needed to invent anything to escape it. That's a real deprivation, even if it doesn't look like one.

What You're Actually Allowed to Do

Let me be direct about permission, because parents absorb a lot of implicit messaging that choices about technology are either paranoid or naive, and there's very little space between.

You are allowed to decide that Sunday afternoons in your house are screenless, without having to justify that choice to a technology company's terms of service or to your child's peer group. You are allowed to insist that dinner is a conversation and that the conversation doesn't have a chatbot participating in it. You are allowed to put your child in a cooking class instead of a coding camp this summer, and you are allowed to believe, correctly, that both choices develop a capable human being.

You are allowed to read books aloud to your child past the age when they technically can read independently, because the experience of shared narrative is not fungible with anything a screen offers. You are allowed to

let them be bored on Saturday morning. You are allowed to take a camping trip where there's no signal and treat the lack of signal as a feature rather than a disaster.

None of this is anti-technology. It's pro-human. And the distinction matters, because framing these choices as a rejection of the modern world makes them feel like losing. They are not losing. They are a considered investment in a specific portfolio of capabilities — attention, making, embodied skill, social friction, uncertainty tolerance — that will not depreciate as AI improves. They will appreciate.

The tools your child needs to evaluate an AI-generated essay for hidden errors, to lead a team of people through a complicated project, to give honest feedback to a colleague, to parent their own children someday, to maintain a friendship through a difficult stretch — none of those tools come from optimizing screen time. They come from a childhood that was full enough of hard, real, human things.

The Conversation That Runs Through All of It

One of the subtler points in Jonathan Haidt's work — one that gets less attention than the smartphone graphs — is that what adolescents have lost most isn't time, exactly. It's unstructured time with other humans. Time that isn't supervised, scheduled, or mediated by a platform. Time where they have to figure out what to do and with whom and how to handle it when it goes sideways.

That kind of time is where the deepest social learning happens. The developmental literature on adolescent peer relationships is unambiguous: navigating peer dynamics without adult management — the fight, the makeup, the hierarchy, the alliance, the falling-out — is essential to developing the social cognition that adults rely on for the rest of their lives. It cannot be replaced by a group chat, and it cannot be replaced by a chatbot, no matter how sophisticated the chatbot becomes.

This means that some of the most valuable things you can do for your child have nothing to do with any specific rule about AI. They involve the shape of the week. Is there time in the week that's genuinely unstructured — not supervised free choice, but actual margin, where your child has to figure out how to fill an afternoon? Are there regular meals where the conversation is allowed to go somewhere real, where a parent admits to being wrong about something, where a teenager's opinion actually changes someone's mind? Is there something your child is building, literally or figuratively, that will take months and might fail?

A parent I know made a rule a few years ago that has stuck with her family: every person in the house has to be making something. Doesn't matter what. Her son is building a remote-control car from a kit. Her daughter is learning calligraphy. She's taking a bookbinding class on Thursday evenings. No one in the family chose these things because they were optimal skill development. They chose them because making is pleasurable and grounding and real. The byproduct, she told me, is that her kids talk about their projects at dinner, ask for help, struggle in front of each other. The making generates conversation that wouldn't otherwise happen.

That's the thing worth protecting. Not a particular capability on a list, not a specific skill on a college application, but a texture of life that keeps humans practicing being human together. Sport, craft, cooking, boredom, unstructured conversation — these look, from the outside, like they have nothing to do with

preparing children for an AI-shaped world. They have everything to do with it.

The algorithm will keep improving. It will be able to do more next year than it can today, and more the year after that. The children who will be most capable in that world are not necessarily the ones who learned to use the most tools. They are the ones who know what they think, can sit with difficulty, can make something with their hands, can navigate a hard conversation with a real person, and can tell the difference between a confident machine and an accurate one.

That last skill, as it happens, starts at the dinner table. Not on a device.

CHAPTER X

Kids Who Want to Build: Nurturing the Young Creator

Your eleven-year-old has been at the kitchen table for two hours. You check on her, expecting homework or YouTube, and find something else entirely: she's built a small image classifier that can tell the difference between pictures of your dog and pictures of other dogs. She trained it herself, using something called Google's Teachable Machine, which is free, runs entirely in a browser, and requires zero lines of code. She wants to know if she can add a third category. You have no idea what to say.

This is a different conversation than the one most parenting-and-AI books are set up to have. Not the conversation about limits, or honesty, or screen time. The conversation about what happens when your child's curiosity has already lapped you, and you're standing in the kitchen holding a cup of coffee and realizing the gatekeeper model of parenting just ran out of gate.

Some kids are builders. They always have been. Every generation had them — the kid pulling apart radios in 1965, the kid teaching himself BASIC on a Commodore 64 in 1983, the kid building Minecraft redstone contraptions at midnight in 2012. In 2026, those kids are prompting, fine-tuning, scraping data, and training small models. And what they need from you is not a usage limit. It's something harder and more interesting than that.

The Difference Between Consuming and Creating

Jonathan Haidt's work in The Anxious Generation (2024) drew a lot of attention to the distinction between active and passive technology use — the idea that not all screen time is equivalent, and that creating, communicating, and making things with technology carries a different psychological weight than endlessly scrolling. That particular distinction holds up well. Developmental researchers have been making a version of it for decades. Yasmin Kafai at the University of Pennsylvania has spent thirty years documenting how children who program — who make things for others to use — develop differently from children who only consume what others have made. They build tolerance for failure. They learn to debug, which is really just a word for holding a problem in your head without panicking. They develop a relationship with expertise that is earned rather than performed.

Mitchel Resnick, who runs the Lifelong Kindergarten research group at MIT and is the person most responsible for Scratch — the free visual programming platform used by over 100 million kids worldwide — wrote in 2017 that the goal was never to teach coding. The goal was to teach children to think creatively, reason systematically, and work collaboratively. AI, in 2026, has become the latest medium through which a certain kind of child does exactly that.

The difference shows up in how they talk about it. A kid who is consuming AI says it gave me or it told me. A kid who is building with AI says I made it do or I'm trying to get it to. That shift in grammatical agency is not trivial. It maps onto something real about how they're relating to the technology — as a material to work with, not an oracle to consult.

Common Sense Media's ongoing research into young people's technology habits has consistently found that creative technology use is the category parents worry about least and understand the least. Most parental concern concentrates on social media and communication apps. Creative and productive use — coding, making, building — tends to fall outside the anxiety window even as it becomes more sophisticated. That blind spot is worth examining.

What Builders Are Actually Doing

Let's be specific, because the range is wide.

A ten-year-old might be using Scratch's AI extensions to build a game that responds to hand gestures captured through the laptop camera. That's genuinely impressive and also entirely accessible — Scratch is free, moderated, and designed for children as young as eight.

A thirteen-year-old might be using Machine Learning for Kids, a free tool created by Dale Lane at IBM, which lets students train custom machine learning models and then integrate them into Scratch or Python projects. There are guided activities — training a model to recognize positive vs. negative newspaper headlines, for instance — but the platform is open enough that kids run well past the tutorials.

A fifteen-year-old might be hosting a fine-tuned language model on Hugging Face Spaces, which as of 2026 remains free for basic hosting and has an increasingly active community of young developers. They might be contributing to open-source AI projects on GitHub, or building small apps that use APIs from Anthropic, OpenAI, or Mistral. They might be doing this entirely on their own, on a Chromebook, while you thought they were doing Spanish homework.

A seventeen-year-old might be preparing a submission for a science fair or competition with a project that would have required a PhD student's computing access ten years ago. The barriers have come down that far, that fast.

None of these kids need the same parenting intervention. What they share is a need for a parent who doesn't reflexively reach for restriction.

The Tools Worth Knowing

You don't have to become a machine learning engineer. But knowing what's out there — actually knowing it, not just recognizing the names — changes the quality of conversation you can have.

Google's Teachable Machine (teachablemachine.withgoogle.com) is the gentlest on-ramp. Point a camera at things, label them, train a model in about ninety seconds, watch it classify in real time. No code, no account

required. It's been used in elementary school classrooms, and it's also been used to prototype real accessibility tools for people with motor disabilities. Both things are true.

Scratch (scratch.mit.edu) is not new — it launched in 2007 — but its AI and machine learning extensions have made it genuinely current. The community features are moderated by MIT's Scratch team, which makes it one of the few places on the internet a nine-year-old can share work without parental fear.

Machine Learning for Kids (machinelearningforkids.co.uk) is the one to bookmark for kids aged roughly nine to fourteen who want to understand why a model does what it does, not just that it does something. The teacher accounts are free. The student interface is clear. Dale Lane updates it regularly.

Code.org has added AI curriculum across multiple grade levels, with units appropriate for middle schoolers that don't require prior programming experience. The platform is widely used in schools, which means your child may already have an account.

Hugging Face is where things get serious. It's a platform for sharing, running, and collaborating on AI models — think of it as GitHub for machine learning. Sixteen- and seventeen-year-olds who are ready to work with real models can do so here without spending money. The community is large, international, and generally collegiate rather than competitive. There are beginner-friendly spaces. The documentation is readable if your teenager is patient enough to use it.

fast.ai (fast.ai) offers free practical deep learning courses taught by Jeremy Howard, who has a specific philosophy: learn by doing first, theory second. The courses are designed for people who are not mathematicians. Motivated fifteen-year-olds have completed them. It requires Python, which is the point at which some kids will need additional support — but Python itself is learnable, and Code.org and Khan Academy both have introductory courses that work.

One important note: free tools have free limitations. Some platforms, Hugging Face included, have public-by-default settings. Anything your child builds and posts is, by design, available for others to see and use. That's mostly a feature, not a bug — the culture of open sharing is part of what makes these communities function — but it's worth a conversation before your fourteen-year-old posts their first model.

Communities Where Young Builders Belong

Tools are the easy part. Community is harder.

The Scratch community is still probably the most reliably safe place for younger builders to share work and get feedback. It is explicitly for kids, the moderation is real, and the culture rewards creative projects rather than technical showing-off. If your child is in that ten-to-twelve range and building their first AI-integrated projects, Scratch is where they'll find peers at a similar level.

AI4ALL is a nonprofit — founded in 2017 — that runs programs specifically to bring underrepresented young people into AI development. Their Open Learning program offers free online curriculum. Their summer intensives have historically been held at partner universities including Carnegie Mellon and Boston

University. They're worth checking directly (ai-4-all.org) because program availability shifts year to year and the quality is consistently high.

Girls Who Code runs programs with increasing AI content. If you have a daughter who is interested but hasn't quite found her entry point, a Girls Who Code club or summer immersion program can be the thing that converts mild interest into sustained engagement. The peer environment matters enormously at that age — more than the specific curriculum.

Online, the challenge is that most serious AI communities congregate on platforms — Reddit, Discord — that have age minimums of thirteen or sixteen that are routinely ignored. That's not a reason to panic, but it is a reason to have a realistic conversation with your teenager about the fact that they may encounter adult content, adult expectations, and adult levels of technical criticism in those spaces. Some teenagers are ready for that. Some aren't yet, even if they think they are.

The Discord servers associated with specific tools — fast.ai has one, Hugging Face has one — tend to be more on-topic and less unpredictably chaotic than general tech Discords. If your teenager is going to be in online communities around AI, those are relatively good bets.

Programs Worth the Investment

Summer intensives are expensive. Most of them. That's the honest truth, and it matters.

iD Tech Camps runs AI and machine learning programs at university campuses across the U.S. They're well-organized, well-staffed, and the on-campus setting gives teenagers a taste of what studying computing actually looks like. They're also $1,000 to $2,000 per week, which is a real number.

The Johns Hopkins Center for Talented Youth (CTY) has a long track record with academically motivated young people and has incorporated computational thinking and AI courses at multiple levels. Financial aid is available and taken seriously — CTY has scholarship programs that have been funded for decades.

The AI4K12 Initiative, a partnership between the Computer Science Teachers Association and the Association for the Advancement of Artificial Intelligence, funded partly by the National Science Foundation, has been building free K-12 AI curriculum since 2018. Their five "Big Ideas in AI" framework — perception, representation, learning, natural interaction, and societal impact — is available at no cost and has been adopted by schools in all fifty states. If your child's school isn't using it, you can point administrators directly to ai4k12.org.

MIT's Day of AI curriculum is free, designed for middle and high school classrooms, and built specifically to teach AI concepts without requiring prior programming experience. If your child's teacher doesn't know about it, that's the email you send in September.

The difference between a child who has access to these resources and one who doesn't is not usually motivation. It's whether a parent knew to look.

The Disposition Shift

Here is the harder conversation.

Most of the parenting advice in this book has been structured around a parent who knows more than their child and is choosing how much access to grant. That model describes most of what happens at ages six through twelve, and it describes a meaningful portion of what happens at thirteen and fourteen. But for some kids — particularly the builders — it stops describing reality somewhere around the time they could explain to you how a transformer architecture works and you nod along hoping they won't ask you a follow-up question.

When that moment arrives, the instinct to maintain the gatekeeper position can actually do harm. Not dramatic harm. Slow harm. The kind that shows up as a teenager who stops telling you what they're working on because you don't have useful things to say, and who eventually routes around you entirely when they need real intellectual engagement. That's a loss. For both of you.

Mitchel Resnick wrote about what he calls "creative learning spirals" — the way genuine makers cycle through imagining, creating, playing, sharing, and reflecting, and then back to imagining. Parents can either be part of that cycle or outside it. Being outside it isn't necessarily damaging, but being inside it — even partially, even incompetently — is remarkably good for the relationship.

The shift looks like this. Instead of what's the limit on this, you start asking what are you trying to make? Instead of is this appropriate for your age, you ask how does it work? You don't need to already know the answer to that second question. You need to ask it like you mean it.

I want to be honest about how uncomfortable this is. Asking a thirteen-year-old to explain something technical to you requires a specific kind of parental ego suspension that doesn't come naturally to everyone. It can feel like a reversal of the correct order of things. It is a reversal of the correct order of things, in the narrow technical domain. And that's fine. Children who are allowed to be the expert in some bounded context develop something real — a sense of competence and worth that is attached to their actual abilities rather than to adult approval. That's not a small thing.

You don't need to fully understand what your child is building. You need to be interested enough to ask, humble enough to listen, and present enough to notice when the work is going well and when they're stuck and frustrated. Those are parenting skills you already have.

Some Practical Suggestions

If your child is in the early stages of curiosity — something they've mentioned, a question they've asked, a YouTube video that caught their attention — the lowest-friction thing you can do is open Google's Teachable Machine together. Right now, tonight, if you want. It takes fifteen minutes to build something real enough to feel satisfying. It requires nothing except a laptop and a camera. Do it together and let them lead.

If your child is already past that stage and working in Python or messing with model APIs, the most useful thing is probably not a new tool. It's a person. Older mentors — a high school student who's been doing this for two years, a university student through a local mentorship program, an engineer at a company that does any AI-adjacent work — matter more than curriculum at that stage. AI4ALL's mentorship programs connect students with volunteers in industry. Code.org and Girls Who Code have similar networks.

Check whether your child's school has a computer science teacher who knows about the AI4K12 curriculum. If they don't, that's a conversation to have at back-to-school night — not confrontationally, but genuinely. The curriculum exists, it's free, and many teachers would use it if they knew about it.

If cost is a real barrier to summer programs, look specifically at AI4ALL, which has built financial access into its model from the beginning, and at the scholarship programs at CTY. The expensive programs get the most marketing. The accessible ones require more searching.

And one more thing, which is easy to forget in the practical rush of parenting: tell your child, directly, that what they're doing is genuinely impressive. Not impressive-for-your-age. Impressive. Because it is. The number of adults who could train a custom image classifier, or build a text-generation tool, or fine-tune a model on a specific dataset, is smaller than you think. Your kid is doing something real. Say so.

The Shape of This Partnership

Haidt's concern — and it's a legitimate one — is that too much of what adolescents do with technology is passive, social, and disconnected from the kind of effortful engagement that builds real capability over time. Building with AI is essentially the opposite of that profile. It is active, often solitary in the productive sense, and intensely effortful. When something doesn't work, you have to figure out why. There is no algorithm smoothing the experience.

That said, building culture has its own traps. The drive to ship something, to post it, to get likes in a technical community, can crowd out the slower work of understanding deeply. A teenager who can prompt their way to a working demo but can't explain why it works is building on sand. The parents and mentors who serve young builders best are the ones who ask why does that work? as often as they ask what did you make?

You don't need the answer. You need to make the question feel welcome.

There's a specific kind of household where this all becomes possible. It's not a household with the right tools or the right income or the right school district, although all of those help. It's a household where intellectual curiosity is treated as something to be indulged and extended, not managed and contained. Where a child who goes down a weird rabbit hole comes up for air and finds a parent who says what did you find? rather than that's enough screen time now.

If your child is a builder, they will find the tools. The internet in 2026 has made sure of that. What they might not find on their own is a parent who is willing to sit down, admit they don't understand, and ask to be taught something.

That's what they need. More than a summer program. More than the right platform.

You, in the kitchen, willing to be shown how the dog classifier works.

CHAPTER XI

When the Kid Is Right and the Parent Is Behind

The story I keep returning to involves a fourteen-year-old named Theo and a game he will probably never finish.

Theo had spent three months building a small side-scrolling game in Python, using GitHub Copilot — an AI coding assistant — to help him catch errors and explore approaches he could then examine, reject, or modify. He wasn't having the AI write his code any more than a carpenter has a circular saw build his furniture. He was using a tool. And he understood the tool. He could explain, in specific terms, what it was doing and what it wasn't. The game was rough. The art was placeholder rectangles. But enemies moved, scores tracked, and there was an actual win condition — which is more than most adults can say for their first programming project.

Then his mother read a headline.

The article wasn't entirely wrong, as articles rarely are. But it was abstract, anxiety-first, and stripped of the context that would have made it useful. His mother absorbed the fear and made a decision the same evening: no more AI tool. Figure it out yourself. Theo tried to explain the distinction between code generation and what he was actually doing. She said she understood, but the decision stood.

Theo stopped working on the game. He hasn't started another one since.

I tell that story not to judge his mother. She was doing what parents do — reading, worrying, acting. She loves her kid. But the thing she acted on was a secondhand fear dressed up as information, and the thing she shut down was something genuinely real: his engagement, his competence, and, quietly, a small piece of his trust in her judgment about things she doesn't yet know.

The Gap Is Structural. It Is Not About Intelligence.

Here is the uncomfortable truth this chapter is built on: in 2026, many children between the ages of twelve and seventeen understand specific AI tools better than the adults responsible for them. Not AI in general. Not the ethics, the policy, the long-term societal stakes. Those remain firmly in adult territory. But the specific tools — how a particular chatbot actually behaves, where its outputs tend to fail, what prompting strategies work, which settings matter — kids often know these things more accurately than parents do, because they use them more and pay different kinds of attention.

This is not a character flaw in parents. It is structural.

You came of age before generative AI existed. The learning curve is real. A 2023 Common Sense Media survey on teens and AI found that teenagers were significantly more likely than their parents to have actively

used a generative AI tool in the prior month, and considerably more likely to have developed what researchers described as a working mental model of how outputs are produced. Not a technical one. A practical one — the kind built through use. That kind of knowledge is hard to acquire by reading articles, because articles describe tools from the outside. Kids learn them from the inside.

The gap, in other words, is not about who cares more. It's about who has more hours on the tool.

What follows from that — and this is the part that's hard to say out loud — is that sometimes the parental instinct to protect produces exactly the harm it is trying to prevent. When you assert authority about a thing you don't actually understand, two things happen simultaneously. You may well get the decision wrong on the object level. And you teach your child something true and devastating: that the adults in charge are working from fear, not knowledge. That lesson, once learned, is hard to unlearn.

Jonathan Haidt, in The Anxious Generation, writes at length about the collapse of unsupervised childhood experience — the way the removal of risk also removed the development that risk produces. He is writing mostly about physical space and peer interaction, and his argument there is stronger than the causal claims he makes elsewhere. But the observation maps neatly onto something adjacent: when children are managed by authority rather than guided by relationship, trust erodes. And eroded trust is, in the long run, the most dangerous outcome of any parenting decision, because trust is the channel through which every future conversation travels.

So: what does it actually cost when you get this wrong?

What Authority Deployed in Ignorance Actually Does

Let's be specific.

Your thirteen-year-old tells you that Character.AI has a setting that limits the persona's responses to more appropriate content for younger users. You say, with confidence, that none of that matters because the company can't be trusted and the safeguards don't work. She knows — because she's tested this, because she's read the platform's own documentation, because she's seen the difference — that you are not entirely right. You're not entirely wrong either; the safeguards are imperfect. But you are speaking with a certainty that doesn't match reality, and she can see the gap.

She doesn't argue further. Why would she? She's thirteen and she knows how that goes.

But here's what just happened: you taught her that the path to continued access to the tools she wants is around you, not through you. You've made yourself an obstacle rather than a resource. Next time she has a genuine concern about something she's found on a platform — something disturbing, something she doesn't know how to process — she will think twice before bringing it to you. Because you are the person who spoke with authority about something you didn't understand, and she remembers.

That's the real cost. Not the specific decision about the specific app. The relationship infrastructure that future conversations run on.

There's also a second cost, quieter and slower-moving. Adolescent development research is fairly consistent on this point: teenagers who perceive parental rules as arbitrary or uninformed are significantly more likely to dismiss those rules as a category, including rules that are well-founded. Laurence Steinberg, whose decades of research on adolescent autonomy development are among the most cited in the field, has documented how the perceived legitimacy of parental authority directly predicts compliance — not just in the moment, but as a general orientation. Kids who understand why a rule exists, and who believe the adult making it has a real grasp of what they're regulating, internalize limits differently than kids who experience rules as noise from people who don't know what they're talking about.

You don't have to know everything. You just have to be honest about what you know.

Three Moments That Reveal the Gap

You will probably recognize at least one of these.

The AI-writing detection conversation. Your sixteen-year-old turns in an essay. The teacher flags it with an AI detection tool — perhaps Turnitin, perhaps one of the several competing services schools have adopted — and says the writing looks artificially generated. Your instinct is to believe the school. Your teenager insists the essay is hers. You aren't sure who to believe, so you come down, at least initially, on the side of the authority.

Here is what your teenager may actually know that you don't: AI detection tools have significant false-positive rates. A 2023 study published in PLOS ONE found that some commonly used detection tools flagged text written by non-native English speakers at dramatically higher rates than text by native speakers, because formal, careful writing patterns can look, to a statistical model, like AI output. Your teenager may have written clearly and carefully. That's not proof she didn't use AI — but neither is a flagged detection score. If you'd spent an hour with these tools the way she has, you might know this. If you walk into that teacher conference armed with certainty when the situation actually requires nuance, you're not protecting your kid. You're failing her.

The safety features conversation. You announce at dinner that you're removing a particular chatbot app because you've read it collects personal data. Your son — twelve, quiet, careful — says he's already turned off the data sharing setting in the privacy menu and hasn't entered any real personal details. You say you didn't know there was a privacy menu. He shows you. There are four settings you could review together in about six minutes. You tell him you'll look at it later, but your tone already said you've made your decision and the app is going.

He doesn't fight it. But he knows something now: that you made a decision based on something you could have checked and chose not to. That information sits in him.

The creative tool conversation. Your fifteen-year-old is using an AI image generation tool for an art project. Not to produce finished work — to sketch fast iterations of compositions she can then repaint by hand. This is a legitimate studio practice; professional illustrators do something similar. You read that AI image tools are

trained on artists' work without consent, which is true and ethically real. You tell her she has to stop using it. She asks whether you understand how the specific tool she's using is trained and what it generates. You don't, really, but you tell her it doesn't matter.

It does matter. Not because the ethical concern is wrong — it isn't. But because the nuance you missed is the nuance she was ready to think through with you. You foreclosed that conversation.

In all three cases, the problem isn't the decision. It's the gap between your certainty and your actual knowledge, and what your child learns from watching those two things fail to match.

The Graceful Path Forward

Here is the good news, and it is real: admitting you were wrong, or admitting you don't know, does not dismantle your authority. It restores it.

This sounds paradoxical. It is not.

Children — especially adolescents — are calibration machines. They are constantly assessing whether the adults around them are operating from genuine understanding or reflexive fear. When you demonstrate that you can receive new information, update your position, and acknowledge a gap in your knowledge, you signal something invaluable: that your guidance is based on reality, and therefore worth listening to. That's a very different message than certainty that outstrips knowledge.

The actual words are simpler than you might expect. "I made that call based on what I understood at the time, and I think I was working from incomplete information. Can you show me how you're actually using it?" That's a whole sentence. That's not a lengthy apology or a retreat from parenting. It is an invitation to the kind of conversation that produces trust.

A few specifics that help:

Ask to be taught. This is not the same as abdicating oversight. "Walk me through how this works" puts your child in the position of translator, which teenagers are generally good at and often enjoy. It also gives you thirty minutes of genuine learning at no cost. You'll ask questions, they'll answer, and you'll both come away with a more accurate shared picture than the one you were each carrying separately.

Don't over-correct. There's a version of this where you feel so guilty about the original overreach that you swing entirely permissive — no limits, no questions, total access. That's not what your child needs either. What they need is for the eventual guidance to be grounded. Being wrong once doesn't mean your concerns were baseless. It means they need recalibration.

Separate the immediate tool from the ongoing value. If you banned something and you've now learned the ban was based on a misreading, you can revisit the specific decision without abandoning the principle. "I still want us to talk about how much time you're spending on this" is compatible with "I think I was wrong to take it away entirely."

Learning Alongside: What That Actually Looks Like

Saying you'll learn alongside your child is one of those things that sounds good in books and then gets dropped somewhere between Tuesday's work emails and Thursday's school run. So let's be practical about what this actually requires.

About ninety minutes. That's genuinely enough to go from zero to functional on most AI tools your child is likely to use. Not expert. Not comfortable with every edge case. But enough to ask better questions and stop making decisions from pure abstraction.

You don't need to sign up for a course. You need to sit with the actual tool. Create a free account on ChatGPT or whatever platform your child is using. Give it a few prompts. Notice where it's helpful and where it's confidently wrong. Try to get it to hallucinate — ask it for a specific book reference, or a historical date, and check it. See what happens when you ask something emotionally loaded. Give it twenty minutes to surprise you in both directions.

Then do what your child is already doing: ask it something you actually want to know. Cooking. Your job. A question from last week that you half-Googled and forgot. The experience of using a tool for something real, rather than for evaluation purposes, teaches you things about how it feels to rely on that tool in a way that reading about it never will.

Common Sense Media has published a series of family guides to AI tools — practical, plain-language, updated regularly — and they're worth an hour on a Sunday afternoon. Not to become an expert. To close enough of the gap that your next conversation with your child can be genuinely two-directional.

The other thing that helps: admit to your child what surprised you. "I tried it and I didn't expect it to get that wrong with such confidence" is a real observation that opens a real conversation. It shows your child that you're taking the thing seriously enough to experience it, not just regulate it from a distance.

Reasserting Guidance From a Better Foundation

None of this means you were wrong to have concerns.

The things parents worry about in this space — privacy, academic integrity, emotional dependency, the uneven quality of AI outputs, the way recommendation engines game attention — these are real. They're documented. They deserve ongoing parental engagement. The problem was never the concern. It was the authority exceeding the knowledge.

Once you've closed some of that gap, you're in a different position. You're not regulating from fear. You're setting limits from understanding. And that is a distinction your child can feel, even if they can't articulate it.

When you tell a teenager that you've looked at the privacy settings on the platform and you still want to limit use to specific contexts, and you can explain why, that conversation lands differently than the same limit delivered from unexamined instinct. The outcome might be identical. But the relationship it travels through is

different. They know you took the time. They know you're working from the actual situation rather than the imagined one.

That's what genuine authority sounds like. Not certainty. Engagement.

From Gap to Collaboration

The frame this chapter wants to leave you with is this: the knowledge gap between you and your child around AI is not a problem to be solved and then filed away. It will reopen. The tools will change. New platforms will appear — probably several times between now and the year your youngest finishes high school. The specific knowledge your child has today about the specific tools available today will be partially obsolete in eighteen months, and so will whatever you learn.

What doesn't become obsolete is the habit of staying in conversation.

Haidt, for all the ways critics argue his causal claims outrun his evidence, is on solid ground when he observes that one of the most protective factors in adolescent wellbeing is having an adult in their life who takes their experience seriously. Not an adult who always knows more. An adult who stays curious, stays engaged, and keeps showing up with genuine questions rather than finalized verdicts.

You don't have to be the expert in the room. You have to be the adult who admits what they don't know and keeps asking. You have to be the parent who can say, without it costing their credibility, "I got that one wrong — show me what I missed."

That is a harder skill than learning which settings to audit on Character.AI. It requires something that doesn't come easily after years of being the one who knows — it requires a specific kind of humility that is, once practiced, not actually humbling at all. It's clarifying. It's what changes the dynamic from surveillance and regulation to something that can actually sustain a relationship through adolescence.

Theo, the fourteen-year-old with the unfinished game, doesn't need his mother to become a programmer. He needs her to say: I read more about it, and I think I overcorrected. Tell me what you were actually building.

That's where you start.

The game might still get finished. The more important thing — the habit of bringing things to a parent rather than around them — that can be rebuilt too. But it takes going first. It takes the adult in the room being willing to arrive, late and imperfect, with an honest question instead of a closed case.

You have exactly the tools required. Curiosity. Willingness to be corrected. Love for the specific child in front of you, who already knows more about this particular thing than you do, and who is watching to see what you'll do with that fact.

CHAPTER XII

Conversations to Keep Having: The Next Ten Years

There is no final chapter in the parenting version of this story.

That sentence took me a while to accept. I wanted to end this book the way a good renovation project ends — tools cleaned and hung back on the wall, drywall smooth, light switches clicking satisfyingly in all the right directions. Done. But parenting technology isn't a renovation. It's more like tending a garden in a climate that keeps changing its mind. The tools that worked last season aren't wrong exactly. They just need adjusting.

The same week I was finishing this manuscript, a school district in suburban Denver quietly rolled out an AI grading assistant for middle school essays. No parent meeting. No opt-in. Just a software update that teachers were told, in a one-paragraph email, would "streamline feedback." Meanwhile, three districts in the same metro area were still enforcing a full ban on any student use of generative AI, including using it to look up definitions. Two classrooms, fifteen miles apart, operating as though they were in different eras.

That gap will keep widening before it narrows. The pace of AI development — not the sci-fi version, but the boring incremental version that touches your kid's school, their social apps, and the recommendation systems they swim in every day — is not slowing down. By the time this book reaches your hands, some of the specific tools mentioned in earlier chapters will have updated, merged with competitors, or pivoted entirely. That's not a failure of the book. It's the nature of the material.

What doesn't change is your relationship with your child. And whether you've built enough of a conversational foundation to handle surprises together.

That's what this chapter is about. Not a final answer. A practice.

Why One "Big Talk" Always Fails

Think about the last time you sat your kid down for a Serious Conversation about the internet. Maybe you rehearsed it. Maybe you had three talking points and a printout from a school newsletter. And maybe it went fine, technically. They nodded. You nodded. Everyone agreed that yes, privacy matters, and yes, they'd tell you if something felt weird online.

And then nothing changed.

The problem isn't that the conversation was wrong. It's that a single conversation — no matter how earnest — cannot hold the weight of an ongoing technological shift. Common Sense Media's annual research on teens and digital life has consistently found that only a minority of teenagers say they talk to a parent regularly about what they do online. Not never. Not total silence. Just not regularly. The rest are navigating a genuinely novel environment, one that none of us fully mapped, largely on their own.

What actually shifts adolescent behavior is less about rules and more about relationship. Patti Valkenburg, who has spent decades at the University of Amsterdam studying how families shape children's media habits, finds consistently that it's not the number of restrictions that predicts better outcomes — it's the quality of ongoing dialogue. Kids who feel like they can tell a parent something strange happened online, without a lecture or a confiscated phone, are more likely to share information early, when it's still manageable. The talk isn't the thing. The habit of talking is the thing.

Jonathan Haidt, in The Anxious Generation, argues that something structurally changed in adolescent social life after roughly 2012, when smartphone adoption crossed a threshold and the damage to teen mental health runs deep enough to require collective structural responses — phone-free schools, phone-free bedrooms, coordinated action rather than each family muddling through alone. His critics, including developmental psychologists Amy Orben and Andrew Przybylski at Oxford, point out that the statistical effect sizes linking phones to adolescent distress are often modest when you control carefully for pre-existing conditions and the question of what's causing what. Both camps are probably capturing part of something real. But they agree on at least this: presence matters. Not surveillance. Presence. The kind that makes a child feel like their interior life is visible to someone who cares about them.

That's a lot easier to say than to sustain across a decade of parenting.

So here's a practical structure for doing it anyway.

The Annual Ritual

Once a year. That's the commitment. Not a monthly audit, not a quarterly check-in with a rubric. Once a year, you sit down together — maybe at the start of the school year, maybe on a long drive to visit relatives, maybe over pizza on a random Tuesday in March — and you run through a short set of questions. Not a quiz. Not a deposition. A conversation. The questions are designed to flex as the technology changes and as your child grows.

What follows is a framework, not a script. Change the language. Add questions that fit your family. Drop the ones that don't. The point isn't to check boxes. The point is to keep the channel open.

For Families with Children Ages 6–9

These conversations are shorter and more concrete. You're mostly building vocabulary and the habit of telling you things. You're not trying to cover philosophical ground yet.

What's your favorite thing to do on the internet this year? This one sounds almost too simple. It isn't. The answer tells you what your child is actually drawn to, not what you assumed they were doing. A lot of parents are surprised: kids this age are often on YouTube, in Roblox, or using reading apps — not the platforms generating the most parental anxiety. And when you know what they love, you have a genuine entry point. Questions rooted in curiosity don't feel like interrogation.

Has anything on a screen made you feel bad, scared, or confused this year? Keep your face neutral when you ask this. Really try. The moment a parent's expression shifts toward alarm, a nine-year-old starts editing their answer in real time. If they say yes — and sometimes they will — your only job in that moment is to stay calm and ask them to tell you more. The conversation about what to do comes later.

Did you ever see something you weren't sure was real? At this age, you're planting the seed of a lifetime habit. AI-generated images and video already surface in the media environments young children use. You don't need to explain how large language models work. You just need to establish that some things look real and aren't, and that it's always okay to wonder out loud about it with you.

What rules do you think we should have about screens this year? Kids are far more likely to follow rules they helped make. Their proposed rules sometimes surprise you. They're occasionally stricter than yours.

For Families with Children Ages 10–13

The middle school years are where AI tools become genuinely relevant to daily life — and where the social dynamics around technology get complicated fast.

What AI tools have you used this year, for school or just for fun? You might think you know the answer. You probably don't, entirely. Common Sense Media's research has found that tween use of AI platforms consistently runs ahead of parental awareness — not because kids are hiding it maliciously, but because adoption happens fast and casually, one friend's recommendation at a time. This isn't about catching them out. It's about mapping the territory together. Ask what they liked, what surprised them, what felt strange.

Has any AI tool asked you for personal information, or have you shared personal details — your name, your school, what you're upset about — without thinking much about it? Frame this without blame, because the honest answer for almost every kid this age is yes. AI companions marketed to tweens are designed to invite disclosure. That's the product. Character.AI, chatbot features in gaming platforms, social recommendation systems — they learn from what users tell them. Your kid needs to understand this not as a reason for paranoia but as a reason to be thoughtful. Same reason you don't give your full address to a friendly stranger at a bus stop.

Have you seen anything online this year that turned out to be fake or manipulated? By ten or eleven, many children have already encountered AI-generated content — images, audio clips, short videos — without any label identifying it as synthetic. Some of them know it. Some don't. The goal here isn't to run a media literacy quiz. It's to find out what they've actually seen, to share what you've seen, and to practice together the habit from Chapter 5: who made this, why did they make it, where's the evidence?

Is there anything about AI or the internet this year that you think I don't understand very well? This question takes some courage to ask. It signals something important: you're not the authority on this topic, and you know it, and you're not embarrassed to say so. Some kids won't have an answer right away. Others will. Either way, you've told them something most parents don't: that information can flow both directions.

For Families with Children Ages 14–18

Teenagers require a different register. They can smell condescension at extraordinary range. These questions work best as genuine exchanges — you share your honest position, they share theirs, and you resist the urge to close the loop too quickly with a verdict.

What role is AI playing in your school life right now, and do you think the rules around it are fair? Don't skip the second half. The teenager who thinks their English teacher's AI ban is hypocritical because the same teacher used AI to generate their rubric is not necessarily wrong. School AI policy is, at this moment, genuinely inconsistent, and your teenager probably has a more granular view of classroom reality than any administrator's policy document reflects. Listen to their analysis. You don't have to agree with their conclusions to take their observations seriously.

Have you or any of your friends been affected by deepfakes, AI-generated images, or synthetic audio? This question requires nerve, because the honest answer is increasingly yes. Research on image-based abuse affecting teenagers — including AI-generated content targeting girls — has moved this from an abstract risk to a documented pattern in secondary schools. Raising this isn't about frightening them. It's about signaling that the topic isn't too shameful for your household. If something has already happened to someone they know, they need to understand you're reachable. If nothing has happened yet, they need to know that too.

What do you think AI cannot do that humans can? This is genuinely a fascinating question and most teenagers have an opinion worth hearing. Answers vary wildly — some say creativity, some say empathy, some say "nothing really," and a few say things that will stop you mid-bite. The conversation that follows is one worth having. Not because there's a correct answer. Because thinking it through out loud is itself the skill.

What's one thing you've figured out about technology this year that you think I should know? By fifteen or sixteen, many kids are managing parts of their digital lives with a sophistication their parents genuinely lack — not in every dimension, but in specific ones. This question is your invitation to learn something. Take notes if you can do it without being weird about it.

How to Actually Do This

A few things make these conversations go better.

Don't start cold. If you walk in and announce "we're having our annual tech conversation now," you've already created resistance. Let it begin somewhere natural — on a drive, after dinner, during a game. The questions work better when they're not pre-announced as questions.

Go first, occasionally. If you want your teenager to tell you about something that felt weird online, it helps to lead once in a while — to say, here's something I saw this week that I genuinely don't understand, or here's something an AI tool did that surprised me. Disclosure invites disclosure. That's not a parenting trick. It's just how conversation works.

Resist the verdict. The hardest part of these exchanges — for most parents — isn't the awkward silence or the eye roll. It's sitting with an answer you disagree with and not immediately converting it into a teachable moment. You don't have to abandon your values. But if every piece of information your kid shares with you gets immediately met with correction, they will share less next time. The goal is to keep the channel open. Lectures close channels.

Let the questions change. By the time your ten-year-old is fourteen, the AI tools they're using, the social platforms they're on, and the risks they face will be different enough that some of these questions will need updating. That's intentional. Write in the margins. Cross things out. The ritual stays. The content evolves.

What Stays the Same

Every few years, a new technology arrives that parents, researchers, and commentators are certain will rewire childhood in some catastrophic new way. Television was going to destroy imagination. Video games were going to breed violence. Social media was going to make every teenager miserable. Some of those fears turned out to be partially right, in limited and measurable ways. Most were significantly overblown, and the specific mechanisms people worried about were more complicated than the headlines ever allowed.

AI is genuinely different in some respects. The ability to generate convincing text, images, and audio at scale — cheaply, instantly, with minimal expertise — introduces a category of risk that television did not. Deepfakes of real children are a real harm, not a hypothetical. Academic dishonesty enabled by AI is already reshaping what it means to demonstrate knowledge in school. The data AI companies collect on minors is extensive and, as of 2026, poorly regulated in most of the world.

So the fears aren't absurd. Some of them are well-grounded.

But the factor that determines how a twelve-year-old comes through any of this — the television era, the smartphone era, the AI era — has been remarkably consistent across decades of research. It's the quality of the relationships around them. Amy Orben's analyses of large-scale wellbeing data conclude that parental relationship quality, peer connection, and sleep account for vastly more variance in adolescent outcomes than screen time does. Common Sense Media's own findings on digital resilience in teens point in the same direction: the protective factor isn't the number of rules. It's whether a child believes a trusted adult will show up without judgment when something goes wrong.

That is entirely within your reach. No technical knowledge required.

You don't need to understand how a large language model works to be the person your kid calls when a chatbot says something that scares them. You don't need to be able to explain training data to help a teenager think through why an AI tool might be giving them biased advice. You need to be curious. You need to not pretend you have answers you don't have. And you need to keep showing up.

The Only Promise That Holds

I want to be honest with you about something before we finish.

Nothing in this book will be completely current in three years. A few things won't survive twelve months. The specific tools named in Chapter 2 will have been updated, discontinued, or replaced by something none of us had heard of when I was writing. The school policies in Chapter 6 will have shifted. Some of the research cited here will have been extended, complicated, or overturned by newer findings. That's how science works. That's how technology works.

What will still hold: your kid needs to know they can talk to you.

I've written that sentence in different forms roughly fifteen times across this manuscript, and I'm not embarrassed about that. It's the center of gravity of everything else. Every conversation about homework and AI, about privacy and chatbots, about what's real and what's generated, about who builds the tools that shape their world — every single one of those conversations is more productive, more honest, and more protective when it happens inside a relationship where the child believes the parent is genuinely on their side.

Building that relationship is not a technology problem. It never has been.

So here's what I can actually offer you, without faking certainty I don't have: the questions matter less than the asking. The answers matter less than the listening. The specific platform your teenager is using this fall will look different from the one they're on next fall. But the kid is the same kid. The relationship is the same relationship. And every time you sit down together — a little curious, a little humble, willing to be told something you didn't already know — you are doing the actual work.

That work doesn't end with a chapter. It doesn't end with a school year. It doesn't end, really, until they're somewhere asking similar questions with their own kids, and maybe remembering, vaguely and warmly, that you were someone who tried.

That's enough. It might even be everything.