

Teaching Kids to Fish

Why Education, Not Surveillance, Is the Only Sustainable Path to Child Digital Safety

Give a child a content filter and you protect them until they find the bypass.
Teach a child to fish, and they feed themselves for a lifetime.

THE PREMISE EVERYTHING IS BUILT ON

Parents are right to be afraid. The tools sold to answer that fear do not work.

An entire market rests on one idea: that the way to keep a child safe online is to watch everything they do. Products in this space gather browsing data, read messages, follow locations, flag words and report back to a parent. Some go further and run behavioural analysis over children's conversations.

The fear behind that is legitimate. There are predatory people online, there is design built to manipulate, and there is material no child should meet unsupervised. Wanting to watch is not authoritarian. It comes from love.

The trouble is not the fear. The trouble is that what gets sold to answer the fear supplies a feeling of control while building none of the one thing that actually protects a child over a lifetime, which is the child's own competence.

THREE PROBLEMS NO AMOUNT OF ENGINEERING SOLVES

Why watching fails, structurally

1 It gets circumvented

Every generation of children ends up more capable with technology than the software watching them. Private networks, second devices, borrowed handsets, encrypted chat and browser workarounds are common knowledge among teenagers. A 2016 survey found most parents doing some form of digital monitoring, and teenage exposure to harmful material, bullying and predatory contact has not fallen in proportion. It is an arms race children are built to win.

2 It costs trust

Research on parental monitoring keeps finding the same thing. Restrictive approaches go with less closeness at home and more problematic internet use. A 2023 study found that restrictive monitoring of adolescents' media use went with problematic use, while active and trusting approaches went with healthier family life and no such association. One parent in that study put it plainly: taking the phone away just pushes the child to be sneakier.

3 It expires

Monitoring has a hard end date. The child turns eighteen, leaves home, buys their own device, or simply moves to a service the software does not cover, and the protection is gone in that instant. If nobody taught them to spot a threat, weigh a claim or guard their own information, they walk into an unwatched world with the vulnerability they had at ten. Nothing protective was ever placed in the child. It was all in the software, and the software has stopped running.

This Has Happened Before

a billion dollars a year, four federal agencies, and no effect



Everyone loved it. Politicians funded it. Police delivered it. It did not work.

1983 TO THE PRESENT

The programme that taught American children to say no

Drug Abuse Resistance Education launched in 1983 as a joint effort between the Los Angeles police and the city's school district. It reached more than three quarters of American schools. It had support from both parties. At its height it consumed somewhere between one and one and a third billion dollars a year.

Parents loved it. Politicians paid for it. Police officers delivered it in person. Everybody felt better.

It did not work.

WHAT THE EVALUATIONS FOUND

Four federal bodies reached the same verdict, and the money kept flowing

The government's own auditors found no significant difference in illicit drug use between children who went through the programme and children who did not. The Surgeon General filed it under ineffective programmes. The National Academy of Sciences found it ineffective. The Department of Education barred schools from spending federal money on it, because the evidence showed no reduction in drug or alcohol use.

A 2004 meta-analysis put it flatly: graduates were indistinguishable from children who never took part. Some studies found participants were three to five percent more likely to use drugs than those who had not attended.

The programme's own response to that evidence is part of the story. It made threatening calls to researchers and pressed journals not to publish.

THE SAME LOGIC, THIRTY YEARS APART

Scare them and tell them to say no. Watch them and block the bad content.

What the drug programme assumed

Frighten a child about drugs, teach them a refusal slogan, and the problem goes away. It was funded for decades against the findings of the auditors, the Surgeon General and the education department. Some studies found a small increase in use among participants. When it ended, the child had no skills for deciding anything alone, and was more exposed rather than less. What eventually replaced it was a curriculum built on social, emotional and cognitive skills, which the Surgeon General's 2016 report found effective.

What the safety industry assumes

Watch a child's activity, block the dangerous material, and the problem goes away. It is funded by billions in venture capital and household spending with no matching fall in online harm. Restrictive monitoring is positively associated with problematic internet use. When it ends, the child has no skills for navigating anything alone. The industry markets fear rather than answering the efficacy question, and it has not yet had its reckoning.

Both treat a child as somebody protection is done to rather than somebody competence is built in. The lesson is not subtle. Give people knowledge and the means to decide and they decide better. Rely on fear and restriction and the restrictions fail, leaving nothing underneath.

PART TWO

The Education-First Alternative

a different question, and four things it asks for



Not how do we stop a child meeting danger, but how do we make sure a child can handle it when they do.

REFRAMING THE QUESTION

Protection lives inside the child, or it lives nowhere for long

An education-first approach does not pretend the risk is imaginary. It changes what is being asked. Not how do we prevent a child from ever encountering danger, but how do we make sure a child can recognise danger and respond to it when they inevitably do.

The difference shows up in a single moment. A monitoring tool spots a child typing a phone number into a chat, so it blocks the message or messages a parent. The child learns nothing except that somebody is watching.

The other approach has already taught that child to stop and ask three questions. Who can see this number? What could somebody do with it? Have I actually confirmed that this person is who they say they are? One builds dependence on being watched. The other builds a pause that works whether anybody is watching or not.

WHAT ACTUALLY GETS TAUGHT

Four competencies that travel with the child

1 Recognising a predator's pattern

Grooming behaviour, social engineering and manipulation have a shape, including a recognisable escalation: heavy flattery, separation from friends, demands for secrecy, small tests of a boundary. A child who understands why a stranger wants to move the conversation to a different app is protected everywhere. A child whose software flags that one app is protected until they find another.

2 Privacy hygiene

Practical habits around personal information: what is being collected, what a permission actually grants, and how to notice when a service is taking far more than it needs to function. These skills carry into every digital situation the child will meet for the rest of their life.

3 Judging generated content

As machine-written text, images and video become hard to tell from the human-made kind, children need ways to weigh a source, spot synthetic material, and understand how recommendation systems shape what they see. No monitoring tool will catch every piece of generated misinformation. A child with media literacy can weigh any of it.

4 Digital citizenship and social engineering

How the con actually works: phishing, pretexting, baiting, following somebody through a door. How to check an identity claim. How to notice when your own feelings are being worked on deliberately. It is the digital version of not getting into a stranger's car, except that it applies everywhere rather than only where a tool happens to be installed.

WHY THIS IS THE BETTER INSTRUMENT

Five advantages that are structural rather than rhetorical

1 It cannot be uninstalled

Learn to recognise a grooming pattern and you keep it, whatever the device, the service, your age or who pays the bill. Software lasts as long as the installation. Learning lasts as long as the person.

2 It scales cheaply

Curriculum travels through schools, teacher training networks and community groups at almost no extra cost per child. Monitoring needs a licence for every device, updates without end, and infrastructure chewing through data forever.

3 It costs no rights

Teaching a child means collecting nothing about them. No browsing history. No reading their messages. No trail of where they went. No database of how children behave. And nothing that makes being watched a condition of growing up.

4 It keeps the relationship

The parent's line becomes let me help you learn to handle this, instead of I am watching everything. Working together produces better outcomes at home, and leaves a child far readier to say something when something goes wrong.

5 It travels across cultures

Building capability together fits communal traditions of sharing knowledge, among them the southern African idea that a person is a person through other people. Models where teachers train teachers already work across international networks, in settings very unlike one another.

WHAT IT LOOKS LIKE BUILT

Safety as part of the lesson rather than a product bolted on

The Real Safety AI Foundation is building this into the Teacher in the Loop platform through the AI Literacy Labs curriculum. Rather than selling digital safety separately or adding monitoring on the side, the platform delivers age-appropriate citizenship teaching as an ordinary part of the learning environment.

That includes modules on spotting manipulative conversation, on how personal data gets collected and used, on weighing machine-generated material, and on the thinking skills that make a child their own first line of defence. The teacher stays central, supplying human judgement and a relationship no algorithm reproduces.

Note what the model does not require. No message scanning. No browsing history. No behavioural analytics on children's data. None of it. The protective mechanism is the child's own competence, with teachers and families working alongside them rather than against them.

WHERE THIS LEAVES US

There is no reason to think digital safety is the exception

The child safety industry has built itself on the same reasoning that produced the drug programme: that fear, restriction and outside control will substitute for knowledge and capability. A government spent over a billion dollars a year on something four of its own agencies called ineffective. The industry is on the same path, selling reassurance while the evidence piles up on the other side.

Children taught about a risk make better decisions than children merely shielded from it. That holds for substance use, for sexual health, for handling money, and for every other area where anybody has collected data.

Give a child a content filter and you protect them until they find the bypass. Teach a child to fish, and they feed themselves for a lifetime.

REFERENCES

References

Travis Gilly, *Teaching Kids to Fish: Why Education, Not Surveillance, Is the Only Sustainable Path to Child Digital Safety*, Real Safety AI Foundation Position Paper, February 2026.

Ennett, S. T., et al. (1994). How effective is Drug Abuse Resistance Education? A meta-analysis of Project DARE outcome evaluations. *American Journal of Public Health*, 84(9), 1394–1401.

Pew Research Center. (2016). *Parents, Teens, and Digital Monitoring*.

Rosenbaum, D. P., & Hanson, G. S. (1998). Assessing the effects of school-based drug education: A six-year multilevel analysis of Project D.A.R.E. *Journal of Research in Crime and Delinquency*, 35(4), 381–412.

Stattin, H., & Kerr, M. (2000). Parental monitoring: A reinterpretation. *Child Development*, 71(4), 1072–1085.

U.S. General Accounting Office. (2003). *Youth Illicit Drug Use Prevention: DARE Long-Term Evaluations and Federal Efforts to Identify Effective Programs*. GAO-03-172R.

U.S. Surgeon General. (2001). *Youth Violence: A Report of the Surgeon General*. Chapter 5, Section 4: Ineffective Primary Prevention Programs.

U.S. Surgeon General. (2016). *Facing Addiction in America: The Surgeon General's Report on Alcohol, Drugs, and Health*.

West, S. L., & O'Neal, K. K. (2004). Project D.A.R.E. outcome effectiveness revisited. *American Journal of Public Health*, 94(6), 1027–1029.

Journal of Child and Family Studies. (2023). Parental Monitoring of Early Adolescent Social Technology Use in the US: A Mixed-Method Study.

TechPolicy.Press. (2025). The Youth Online Safety Movement Needs to Respect Children's Autonomy.