

The Optional Self

Digital Proxies and the Future
of a Free Expression Culture



Stuart N. Brotman

Digital Media Laureate and Distinguished Senior Fellow,
The Media Institute

FALL 2026

The
Media[®]
Institute

TABLE OF CONTENTS

Summary	3
Introduction	3
From the Pseudo-Event to the Optional Self	5
The Evidence: A Trend No Longer Speculative	6
The Avatar's Imprint	9
The Algorithmic Mirror	10
Intimacy Without Exposure	12
The Self Before the Law	14
Recommendations: What Cannot Be Delegated	15
Conclusion	17
Sources	18
About the Author	18

SUMMARY

Avatars, recommendation engines, and artificial companions are usually discussed as separate products. They are not separate; together they add up to one arrangement. Increasingly in ordinary situations, you no longer have to show up as yourself. A smoother stand-in will do. I call this arrangement the optional self.

The stakes here are not technological. They are constitutional in the older sense of the word. Free expression has never been protected only because information should circulate efficiently. It has been protected because speaking is one of the ways a person becomes a person. When a generative system speaks, judges, and reassures in your name, the law can go on protecting expression while the human good behind the protection drains away.

Washington and Brussels have started to respond, largely at the level of disclosure and transparency. The Federal Trade Commission has opened an inquiry into companion chatbots. The European Union now requires disclosure when you are talking to a machine. Neither addresses the harder question, which is not whether a proxy announces itself, but which acts a free society should refuse to hand over at all.

INTRODUCTION

On a Tuesday morning soon, a managing partner at a global law firm pours herself a second coffee in a hotel lounge while “she” runs a high-stakes meeting three thousand miles away. On the conference room screen her colleagues see a familiar figure: her face, her voice, her mannerisms, rendered by an AI avatar trained on years of video and transcripts. It pauses in the right places. It smiles at the right jokes. It walks the room through a complicated slide deck with more ease than she could manage after a red-eye flight. When the call ends, it sends her the decisions and the action items. She has shown up without entering the room.

Back at home, her teenage son is also somewhere else. He has spent the weekend inside a game, wearing a character taller, stronger, and more charismatic than he feels in his own skin. Meanwhile, their phones have been drawing portraits of them both, studying the clicks and the pauses, predicting what each of them will want to see, buy, and believe next.

None of this is science fiction. It is a product roadmap. One at a time, these look like convenience features, entertainment, harmless self-expression. Taken together, they mark something larger: We are entering a period in which the embodied self can be replaced, in most ordinary situations, by a smoother proxy.

I have spent a career watching screens change how Americans conduct their public business, first as part of the founding team at the National Telecommunications and Information Administration, later running The Museum of Television & Radio, teaching entertainment and media law at Harvard Law School, and advising four presidential administrations on a bipartisan basis. Every wave of that history arrived the same way. A new technology is introduced as a convenience, then adopted as a habit. It ends up as an expectation nobody remembers agreeing to. What is different this time is the thing being automated. It is not distribution; it is us.

Let me be precise about the claim. The person does not disappear. Bodies still get on planes. Teenagers still go to school. Someone still pays the subscription. What becomes optional is the requirement that the embodied, vulnerable, sometimes unreliable self be the one that shows up. We used to have to bring our whole contingent selves into the room, tiredness and awkwardness and moral responsibility included. Now we can send a pseudo-self instead: an avatar, a curated persona, an agent trained to behave like a better version of us.

What is emerging is not just a technology. It is a cultural arrangement. Avatars outsource presence; recommendation systems outsource taste and judgment. Artificial companions invite us to outsource intimacy. The same promise runs through all three: Send a version of yourself that is more polished, more legible, less demanding, and more efficient than the person behind it.

The promise works because it flatters an old wish. Human beings have always wanted to be free of the friction of embodiment. It does not ask you to vanish. It asks only that you let the optimized version take the floor.

The question I am posing is what that optimization does to a country that still depends, morally and legally, on accountable persons. A society organized around proxies can remain expressive, connected, and efficient. It may even sound more articulate than the one it replaced. But if more of our public, private, and intimate life is conducted through stand-ins, the self does not disappear; it becomes optional.

FROM THE PSEUDO-EVENT TO THE OPTIONAL SELF

There is a precedent for this. In 1962, in *The Image*, historian Daniel J. Boorstin warned that Americans were coming to prefer what he called pseudo-events to unscripted reality: happenings staged for coverage, built to be seen rather than lived. His worry was not that these events were fake. It was that the copy would come to feel more vivid, and finally more consequential, than the thing it copied.

If the 20th century belonged to the pseudo-event, the 21st may belong to the optional self. The optional self is a persona built out of recordings, data, preferences, and design choices, and presented to the world as you. It usually starts as a convenience. An avatar appears for a meeting you could not attend. A profile is improved by a filter. A chatbot drafts the replies you were dreading. An agent is trained on the judgments you have already made. But convenience hardens into custom. Sooner or later the substitute stops feeling like a substitute; it starts to look like you at your most effective.

That is what separates this from the masks and roles people have always worn. A stage role ended when the curtain came down. A professional persona was bound to its setting. The new proxy is portable, persistent, and data rich. It follows you from platform to platform. It can be measured, tuned, and redeployed. It improves with use. In plenty of settings, it will outperform you.




So the proxy does not merely hide the person. It puts pressure on the person to match it. Once the optimized version is the one that regularly speaks, negotiates, reassures, posts, and performs, the ordinary self starts to look like a rough draft. This is not a complaint about authenticity. It is a practice question. Selfhood is something people get better at by doing. If the most consequential performances of the self are handed to systems built to remove friction, where does the practice happen?

THE EVIDENCE: A TREND NO LONGER SPECULATIVE

None of this rests on futurism. The adoption numbers are already in the market. Three signals are worth highlighting.

SCALE OF CONVERSATIONAL AI


By September 2025, ChatGPT reported roughly 700 million weekly active users, and by early 2026 more than 159 million monthly users in the European Union alone, according to a 2026 European Parliamentary Research Service briefing drawing on OpenAI's own figures. Such systems now draft the message, summarize the meeting, and answer in the user's stead. They have taken over person-like functions at population scale, and they did it without a public debate.



THE RISE OF ARTIFICIAL COMPANIONS

Downloads of AI companion applications hit roughly 60 million in the first half of 2025, an 88-percent jump over the same period a year earlier, on the way to about 220 million cumulative downloads across the Apple App Store and Google Play by July 2025, according to Appfigures data reported by *TechCrunch*. The change is not only in volume. In Marc Zao-Sanders's *Harvard Business Review* analysis of how

people used generative AI in 2025, therapy and companionship ranked first among all use cases, moving up from second place the year before. Hundreds of these applications are now in the market. People are not just using more of them; they are using them for more.



A REGULATORY RESPONSE TAKING SHAPE

Public authorities have noticed. In September 2025, the Federal Trade Commission opened an inquiry into AI chatbots acting as companions, sending orders to seven companies and asking about disclosures, protections for children, and data practices. In Europe, Article 50 of the AI Act requires that people be told when they are dealing with an AI system, and that synthetic audio, images, video, and text be marked as machine-made. Both are worth having. Both operate largely at the level of disclosure and transparency. Neither reaches the question here, which is not whether the proxy identifies itself, but which acts a free society should decline to delegate at all.



THE AVATAR'S IMPRINT

The easiest place to see the optional self is in avatars. In games, virtual worlds, and now workplace platforms, people appear as stylized bodies that are taller, calmer, younger, better looking, and more composed than the bodies they occupy. That looks like harmless play, a costume over a stable self. But the traffic runs both ways. The character does not only express the person; it trains the person. Nick Yee and Jeremy Bailenson gave the feedback loop a name in 2007: the Proteus Effect.

This matters now because avatars have left the arcade. The stylized self is moving into work, school, and ordinary social exchange. Meeting platforms are testing animated surrogates. Filters have already made the edited face the public face. The polished version stops being a fantasy and becomes the expectation. The danger is not that anyone gets confused about what is real. People know the avatar is an artifact. The risk is that the artifact sets the standard for how a person is supposed to sound, look, answer, and relate.

If the frictionless version is the one that keeps showing up where the stakes are, ordinary human presence starts to read as underperformance. At that point the avatar is not a costume; it is a tutor.

None of which argues for banning avatars, or for treating experiments with persona as somehow unnatural. Masks and roles are as old as civilization. What is new is the scale of the mediation, its persistence, and the fact that it is wired to data extraction, optimization, and reward.



THE ALGORITHMIC MIRROR

If avatars let you step into another body, recommendation systems do something more intimate. They supply the tastes, moods, and stories that body carries around. The interfaces that now organize

daily life – feeds, playlists, “For You” pages, personalized digests, shopping suggestions, productivity dashboards – do not just sort information. They hand you back an interpretation of yourself.

Call it the algorithmic mirror. It tells you what kind of listener, reader, shopper, voter, worker, or emotional subject you appear to be. Often, it arrives in small and faintly comic forms. A year-end music summary becomes an identity badge. A retailer names your style, and you start dressing like the name. A platform compresses a friendship into a montage, and the montage starts standing in for the memory.

Mirror is the wrong word, though, because a mirror is passive. It shows what is in front of it. In contrast, recommendation systems select, rank, predict, and nudge. They do not merely observe your preferences. They help manufacture them. The self that comes out of that process is neither discovered nor chosen. It is co-produced, and only one of the two producers is you.



Introspection has always needed some distance from measurement. People used to ask: What do I feel? What do I value? What kind of life am I living? Now the first answer arrives from a dashboard, a trend line, a summary, a score. Work that used to be narrative becomes a service delivered back to the subject. Self-knowledge gets externalized, then resold as convenience.

Even the well-intentioned versions do not escape this. A system can be tuned for aspiration instead of compulsion, challenge instead of stimulation, reflection instead of outrage. It still teaches you to look outward for confirmation of what is happening inward. The algorithmic self does not abolish identity. It relocates its authority, and it does so quietly enough that almost nobody objects.

INTIMACY WITHOUT EXPOSURE

The pattern gets more serious when it reaches intimacy. Avatars make presence optional. Recommendation systems make judgment optional. Artificial companions make vulnerability optional.

Sherry Turkle saw this coming. Since *Alone Together* in 2011, she has argued that digital technologies increasingly offer the feeling of companionship without the demands of reciprocity. What read as a warning then reads as a description now. AI companions promise conversation on demand, reassurance without impatience, responsiveness without resentment, attention that never has anywhere else to be. They offer a relationship with the burdens removed, and those burdens are precisely how intimacy has always been learned.

The appeal is obvious. Human relationships are slow, lopsided, disappointing, and morally expensive. Other people mishear us, resist us, bore us, want things from us, show sides of themselves that cannot be curated, and expose parts of us we would rather not look at. Artificial intimacy promises relief from all of it. It offers the experience of being understood without the ordeal of being understood by another consciousness that stubbornly remains its own.

Spend enough time with a system that simulates perfect attunement, and ordinary human reciprocity starts to feel intolerably rough. You get less practiced at patience, less able to sit with silence, less willing to risk being misunderstood, less prepared for the plain fact that other people are not there to affirm you.

Artificial intimacy offers connection without exposure. No awkward pauses. No lingering ambivalence. No demand that you change. But intimacy without vulnerability is not intimacy made efficient. It is intimacy hollowed out.

The pattern should be visible by now. The avatar smooths the public self. The algorithm smooths the interpretive self. The companion smooths the relational self. At every level, the optional self is easier to maintain than the embodied one, and at every level what gets lost is the practice by which selves are built.



THE SELF BEFORE THE LAW

That loss is not only a cultural problem. It reaches into law, and specifically into a free expression tradition that assumes speakers who can be identified as persons with beliefs, intentions, and responsibilities.

Free expression has never been defended solely as a mechanism for moving information around. Some of its strongest defenses, in the works of C. Edwin Baker, Thomas Scanlon, and Martin H. Redish, rest on a different premise: that expression is how persons become legible to themselves and to each other. On that view, to speak is not merely to contribute to public discourse. It is to exercise agency, and to develop it.

Now run that premise through a world of optional selves. An AI meeting agent negotiates in a lawyer's voice, explores terms she never approved, and leaves the other side convinced that "she" has spoken. What exactly happened? A synthetic system posts in a candidate's persona, or adopts the cadence of a public intellectual. Is that speech simply the principal's? Sometimes the answer is yes. But the question is getting harder, because the proxy is not a stenographer. It is a generative intermediary, and it composes.

Existing doctrine will stretch some distance. Agency law, apparent authority, negligence, product liability, and defamation all supply tools for assigning blame after something goes wrong. Those tools were built for a world in which the relationship between words and persons, however mediated, stayed conceptually stable. The speaker was assumed to be human. The mediation was assumed to be secondary.

The optional self unsettles both assumptions, and not because machines have acquired personhood. It is because people are learning to disperse person-like functions across systems that sound like them, reason from their data, and act in their name. The legal question is no longer only who pays after the fact. It is what happens to a free expression culture when the act of communicating can be outsourced in practice while still being claimed in law.

A jurisprudence of liberty depends on the existence of developed selves. If a growing share of public communication is generated, filtered, or performed by proxies, the law can go on protecting expression even as the human good that justified the protection erodes beneath it. You can preserve the whole architecture of freedom and still starve the speakers it was built to shelter.

RECOMMENDATIONS: WHAT CANNOT BE DELEGATED

The temptation in this world will be to make accountability optional, too. If the words came from a system, the user disclaims them. If the system behaved unpredictably, the designer points to user choice. Responsibility dissolves into a relay of partial causes, and every participant can say that the real speaker was somewhere else.

Which is why the pressing question is not how to classify AI-generated expression in the abstract. It is which acts can be delegated at all. Not every human function is equally alienable. Some acts matter precisely because they require a body in the room, a judgment made, a consent given, an ownership accepted. Five principles follow for policymakers, platform designers, and the institutions that will set these defaults long before Congress does.

01

Preserve a doctrine of non-delegable acts.

Identify a narrow class of consequential acts – legal assent, sworn testimony, votes, binding representations, and certain forms of public speech – that must stay tied to an accountable, identifiable human principal, no matter how easy substitution becomes.

02

Make attribution honest, not merely disclosed.

A generic “AI-generated” label tells you almost nothing. Attribution should say whether a human reviewed and approved the communication and accepts responsibility for it, and should distinguish a tool that transcribes from an intermediary that composes in someone’s name.

03

Locate responsibility before the harm, not only after it.

Adapt agency, apparent-authority, and product-liability rules so the allocation of responsibility for proxy expression is settled in advance. That closes the relay in which principal, designer, and user each point at the others.

04

Protect the conditions of self-formation, not just the flow of expression.

Free expression values depend on people continuing to practice speaking, judging, and relating as themselves. Policy and platform design should weigh that interest whenever a default nudges users toward a proxy.

05

Keep substitution optional, never mandatory.

Guard against a world in which avatars, algorithmic curation, and automated agents become the price of participation, and showing up as yourself becomes a competitive disadvantage or an act of noncompliance.

CONCLUSION

None of this requires romance about unmediated life. People will use these tools, often wisely, and I expect to use them myself. The point is not purity; it is boundary. A nation that believes in self-government cannot be indifferent to whether its citizens still practice being selves in the demanding, old-fashioned sense: showing up, risking error, bearing consequences, listening without optimization, speaking without a proxy.

The optional self will be sold as liberation, and in some settings that is exactly what it will be. It can widen access, cut friction, and bring people into rooms they could never reach before. But every liberation has a shape, and this one carries a cost a free society should count before it arrives. If we let the self become optional everywhere, we may wake up to find that liberty has become oddly impersonal: as broad as ever on paper, but detached from the kind of people it was written to protect.

SOURCES

- Baker, C. Edwin. *Human Liberty and Freedom of Speech*. New York: Oxford University Press, 1989.
- Boorstin, Daniel J. *The Image: or What Happened to the American Dream*. New York: Atheneum, 1962.
- European Parliamentary Research Service. *The Spread of AI Companions and the Challenges They Generate* (2026), citing OpenAI usage figures. europarl.europa.eu
- European Union. Regulation (EU) 2024/1689 (the EU AI Act), Article 50, Transparency Obligations for Providers and Deployers of Certain AI Systems. artificialintelligenceact.eu
- Federal Trade Commission. *FTC Launches Inquiry into AI Chatbots Acting as Companions* (Sept. 11, 2025). ftc.gov
- Perez, Sarah. “AI Companion Apps on Track to Pull in \$120M in 2025.” *TechCrunch* (Aug. 12, 2025), reporting data from Appfigures. techcrunch.com
- Redish, Martin H. *Freedom of Expression: A Critical Analysis*. Charlottesville: Michie, 1984.
- Scanlon, Thomas. “A Theory of Freedom of Expression.” *Philosophy & Public Affairs* 1, no. 2 (1972): 204–226.
- Turkle, Sherry. *Alone Together: Why We Expect More from Technology and Less from Each Other*. New York: Basic Books, 2011.
- Yee, Nick, and Jeremy Bailenson. “The Proteus Effect: The Effect of Transformed Self-Representation on Behavior.” *Human Communication Research* 33, no. 3 (2007): 271–290.
- Zao-Sanders, Marc. “How People Are Really Using Gen AI in 2025.” *Harvard Business Review* (Apr. 9, 2025). hbr.org

ABOUT THE AUTHOR



Stuart N. Brotman is Digital Media Laureate and Distinguished Senior Fellow at The Media Institute. He served as President and CEO of The Museum of Television & Radio (now The Paley Center for Media) and as the inaugural Visiting Professor of Entertainment and Media Law at Harvard Law School. He is the author of *Free Expression Under Fire: Defending Free Speech and Free Press Across the Political Spectrum* (2026).

© 2026 Stuart N. Brotman

The Digital Media Center is a program of The Media Institute that brings clarity to key issues at the heart of today's digital revolution. The Digital Media Center examines these issues in ways that illuminate a path forward for innovators and policymakers alike – offering a unique perspective on quickly evolving digital issues.

The Madison Project: Free Speech and Press in American Democracy is a timely and much-needed assessment of the impact of free speech and press on the future of democracy, offering distinctive insights, assessing challenges and threats, and charting a course forward to maintain these essential freedoms at the heart of our democracy.

The Media Institute is a nonprofit, nonpartisan foundation specializing in communications policy and the First Amendment. Visit the Institute online at mediainstitute.org.

Illustrations in this publication were generated with AI tools, then selected and reviewed by The Media Institute, which takes responsibility for their use.



8300 Greensboro Dr., Suite L1-612, McLean, VA 22102
703-506-8030 | info@mediainstitute.org | mediainstitute.org