

Author's Note

You value mateship? They removed every structure that let you see your mates.

That is not a slogan. It is a description of a machine. This paper describes the machine.

There used to be a pub at the end of the street. A cricket club. A church hall where the old blokes played cards on Wednesdays. A front yard where kids played until dark. A neighbour who knew your name. These things were not policy. They were not funded by government. They were the natural result of humans living near each other with time to spare. They produced something economists cannot measure and psychologists struggle to name: the ambient certainty that someone gives a damn about you.

That ambient certainty is what OMXUS calls mateship. Not the performative Australian variety -- not the barbecue handshake -- but the structural kind. The kind where your nan falls and someone is there in 60 seconds because they live close and they know her name. The kind where a teenager with a dark thought has three adults within earshot who are not paid to listen. The kind that makes a person feel like a person rather than a consumer of personalised content.

The attention economy destroyed it. Not by accident. By design -- or rather, by the convergence of a business model (attention-as-commodity) with a technology (algorithmic personalisation) with a set of behavioural mechanisms (supernormal stimuli, variable reinforcement, outrage amplification) that, together, replaced face-to-face community with a screen-mediated simulacrum that looks like connection and functions as isolation.

This paper is about how that happened. It draws on behavioural biology (Tinbergen), operant conditioning (Skinner), developmental psychology (Twenge, Haidt, Odgers), computational social science (Rathje et al.), political economy (Zuboff, Harris), and neuroscience of attention (Hari). The evidence is presented honestly -- including the parts that complicate the narrative. The effect size debate is real. The causal uncertainty is real. The hospital admission data is also real.

The paper serves several of the 14 goals that define the OMXUS project:

- Goal 2 (22-hour work week) (22-hour work week) (work 22 hours, keep your pay): The attention economy is built on unpaid labour -- your attention, mined and resold. The hours you "spend" on social media are not leisure. They are work you do for free so that someone else can sell your behavioural predictions to advertisers.
- Goal 5 (replace police with community response) (replace police with community response) (mateship): This is the central thesis. The machine replaced mateship with a subscription service.
- Goal 7 (legalise and regulate drugs) (legalise and regulate drugs) (being your own man): Algorithmic personalisation means your information environment is chosen for you, by a system optimised to keep you angry and engaged. Nobody tells you what to think, but the algorithm decides what you see, which amounts to the same thing.
- Goal 8 (free internet) (free internet) (internet costs nothing): If the internet is free because you are the product, then the internet is not free. Mesh networking (Goal 8 (free internet) (free internet)) removes the advertising business model entirely. No ads, no engagement optimisation, no attention extraction.
- Goal 11 (physical infrastructure) (physical infrastructure) (monkey bars at every bus stop): Physical infrastructure for physical bodies. The opposite of a screen. You cannot doomscroll while you are hanging from a bar.
- Goal 12 (play-based education) (play-based education) (schools as play, mastery, curiosity): The Prussian model sat children in rows and gave them screens. Play-based education (see companion paper) is the structural antidote to screen-mediated childhood.
- Goal 13 (\$29 emergency ring) (\$29 emergency ring) (\$29 ring): When your nan falls, the screen in your pocket does not help her. A mesh-connected ring that alerts your neighbours in 60 seconds does. Technology that serves connection rather than extracting attention from it.

The attention economy is the mechanism. Mateship is the casualty. The 14 goals are the repair.

-- A.A. & L.N.C.

Abstract

The attention economy -- the monetisation of human attention through advertising-funded digital platforms -- has produced a set of measurable social harms that are better understood as collective than individual. This paper synthesises evidence from behavioural biology, operant conditioning research, developmental psychology, computational social science, and political economy to argue that the primary damage is not to individual mental health (where effect sizes remain small and causation contested) but to the structures of social connection that sustain community life.

We review: (1) the supernormal stimulus framework (Tinbergen, 1951; Barrett, 2010), demonstrating that digital social feedback exceeds the parameters for which human reward circuitry was calibrated; (2) variable ratio reinforcement (Skinner, 1953), showing that core platform design patterns replicate the most behaviourally persistent reinforcement schedule known; (3) the adolescent mental health debate (Twenge, 2017, 2018; Haidt, 2024; Orben & Przybylski, 2019; Odgers, 2024), presenting both the alarming outcome data and the genuinely small population-level effect sizes without resolving a tension that the field itself has not resolved; (4) algorithmic amplification of outgroup animosity (Rathje et al., 2021), where the evidence is mechanistically clearer and the effect sizes larger; (5) the political economy of surveillance capitalism (Zuboff, 2019; Harris, 2016), in which human attention functions as an extractable commodity; and (6) the displacement hypothesis -- that screen time physically replaces the unstructured face-to-face interaction, play, and boredom from which community, creativity, and emotional resilience emerge.

The honest assessment is that the individual harm case is plausible but not causally proven to the standard the scientific community requires. The collective harm case -- that algorithmic engagement optimisation structurally rewards division, replaces in-person social structures with screen-mediated substitutes, and degrades the information environment on which democratic self-governance depends -- is considerably stronger. The machine that makes us hate each other is better understood than the machine that makes us sad. Both machines run on the same fuel: your attention, sold to the highest bidder.

Contents

1. The Inflection Point
2. The Hijacked Brain: Supernormal Stimuli

3. The Slot Machine in Your Pocket: Variable Reinforcement
4. The Anxious Generation: Haidt and the Phone-Free Movement
5. Effect Sizes and Potatoes: The Counterarguments
6. The Body Count: Self-Harm and Suicide Data
7. Outrage as Engagement: The Polarisation Engine
8. Surveillance Capitalism: You Are the Product
9. Stolen Focus: The Fragmentation of Attention
10. The Displacement Hypothesis: What Screens Replace
11. The Isolation Machine: Mateship Destroyed by Design
12. Where This Leaves Us
13. References
14. Data Sources
15. Source Verification Notes
16. Appendix A: Cross-References to Related Papers
17. Appendix B: The 14 Goals and Their Relationship to This Paper

Chapter 1: The Inflection Point

Something changed around 2010-2012. The data is hard to argue with even if the explanation is contested.

Jean Twenge, a psychologist at San Diego State University, built her career on generational research. Her 2017 book *iGen* made a straightforward empirical claim: across multiple large datasets -- the Monitoring the Future survey, the Youth Risk Behavior Surveillance System, the American Freshman Survey -- indicators of adolescent mental health that had been roughly stable or slowly trending for decades showed a sharp discontinuity right around 2010-2012. Depression, anxiety, loneliness, self-harm, and suicidal ideation all spiked. The timing coincided with smartphone saturation among American teenagers. The iPhone launched in

2007, but smartphones didn't reach majority ownership among US teens until around 2012. Instagram launched in 2010. Snapchat in 2011. The always-connected, image-driven social internet arrived and adolescent misery arrived with it.

Twenge's 2018 paper in *Clinical Psychological Science* quantified parts of this: increases in depressive symptoms, suicide-related outcomes, and suicide rates among US adolescents after 2010, with the effects more pronounced among girls. The CDC's WISQARS data showed self-harm hospitalisations among girls aged 10-14 roughly tripled between 2010 and 2020. NHS Digital data in England showed a similar pattern -- referrals for self-harm and eating disorders in young girls increased sharply in the same period. These are not self-report surveys. These are hospital admissions and emergency department presentations.

The correlation is unusually clean. You can draw a line on a graph and the curves bend right where the phones arrived. That is what makes the story so compelling. It is also what makes scientists nervous, because clean stories are often too clean.

Chapter 2: The Hijacked Brain: Supernormal Stimuli

In the 1940s and 1950s, the Dutch ethologist Nikolaas Tinbergen conducted a series of elegant experiments on animal behaviour that won him a Nobel Prize. He discovered that many animals respond to specific stimulus features rather than to whole objects. A herring gull chick pecks at the red dot on its parent's beak to trigger feeding. Tinbergen found you could get the chick to peck even more vigorously at a disembodied stick with an exaggerated red dot -- a stimulus that does not exist in nature but triggers the response more strongly than the real thing. He called these *supernormal stimuli*.

The concept generalises uncomfortably well to human technology. Deirdre Barrett's 2010 book *Supernormal Stimuli: How Primal Urges Overran Their Evolutionary Purpose* applied Tinbergen's framework to modern life -- junk food as supernormal taste stimuli, pornography as supernormal sexual stimuli, and screens as supernormal social and novelty stimuli. The argument is that social media provides social feedback (likes, comments, follower counts) in quantities and at speeds that never existed in any human social environment. A witty remark in a village might earn you a few laughs from people you know. The same remark on Twitter

might earn you 50,000 likes from strangers. The reward signal is the same category of thing -- social approval -- but amplified to a degree the underlying neural circuitry was never calibrated for.

This is not a metaphor stretched beyond usefulness. The core mechanism Tinbergen identified -- that evolved response systems can be hijacked by stimuli that exaggerate the triggering features -- is well-established in behavioural biology. What is less established is the dose-response curve in humans. How much supernormal stimulation does it take to produce measurable harm? That question turns out to be central to the entire debate.

The supernormal stimulus framework also illuminates why the replacement feels seamless. If a herring gull chick prefers the fake beak to the real one, it does not experience itself as being deceived. The fake beak *feels more real than the real thing*. That is the point. Instagram does not feel like a substitute for friendship. To the neural circuitry processing social reward signals, it feels like a more intense version of the same thing. The substitution is invisible from the inside. You do not notice the friends disappearing because the screen is filling the same slot in the reward system, only louder.

Chapter 3: The Slot Machine in Your Pocket: Variable Reinforcement

B.F. Skinner's work on operant conditioning in the 1950s established that the *pattern* of reward matters as much as the reward itself. A pigeon that gets a food pellet every time it pecks a lever will peck steadily. A pigeon that gets a food pellet on a variable ratio schedule -- sometimes after 3 pecks, sometimes after 30, sometimes after 7 -- will peck frantically and compulsively. The unpredictability is the point. Variable ratio reinforcement produces the highest, most persistent response rates of any reinforcement schedule. It is, in Skinner's own framework, the most behaviourally addictive pattern possible.

Slot machines use variable ratio reinforcement. So does social media.

Pull-to-refresh is the lever. Sometimes you get a new like, sometimes nothing. Open the app -- maybe there is a message, maybe there is not. Post something -- it might get 3 likes or 300. The unpredictability of the reward

is what drives the compulsive checking. Tristan Harris, a former Google design ethicist who went on to co-found the Center for Humane Technology, made this comparison explicit in his 2016 presentations and essays. He argued that tech companies had reverse-engineered the slot machine into every interaction, not accidentally, but through deliberate design optimisation. A/B testing at scale selects for the most engaging version of every interface element. Given enough iterations, the system converges on whatever design produces the most compulsive use. You do not need a conspiracy theory about evil designers. You just need a fitness function that optimises for engagement, and the Skinner box emerges through selection.

Adam Alter's 2017 book *Irresistible* documented these patterns in detail: infinite scroll (removing natural stopping cues), notification timing (variable interval reinforcement layered on top of variable ratio), autoplay (removing the decision to continue), and social comparison metrics (like counts, follower counts, view counts). Each of these features individually might be trivial. Together, they constitute an environment engineered -- through iteration if not through intent -- to maximise time-on-device.

The language of addiction is contested here -- some researchers object to applying it to behavioural patterns without a pharmacological substrate. The objection has merit in clinical contexts. But the behavioural description is accurate: variable ratio reinforcement produces compulsive repetition that is resistant to extinction. Whether you call it addiction or a "problematic use pattern" does not change the mechanism. The pigeon does not stop pecking because you relabel its condition.

Chapter 4: The Anxious Generation: Haidt and the Phone-Free Movement

Jonathan Haidt's 2024 book *The Anxious Generation* took the Twenge thesis and pushed it further. Haidt argued not merely that smartphones correlate with adolescent mental health decline, but that "phone-based childhood" is a causal factor in what he called an epidemic of mental illness among young people. He proposed a specific mechanism: the replacement of free play and face-to-face socialisation with phone-mediated interaction during a critical developmental window. Children need rough-and-tumble play, boredom, unstructured time with peers, and manageable risk to develop emotional resilience. Smartphones replaced all of that with an environment optimised for passive consumption and social comparison.

Haidt's developmental argument converges with the play deprivation research (see Appendix A). Stuart Brown's clinical work found play deprivation as a recurring feature in over 6,000 violent offender histories. Jaak Panksepp identified PLAY as one of seven primary emotional systems -- subcortical, conserved across mammals, opioid-mediated. When you take the phone out of a child's hands and give them nothing to replace it with, you have not solved the problem. When you take the phone out and give them a climbing wall, a creek, and three other kids, you have addressed the mechanism Haidt actually describes.

Haidt's book became a bestseller and catalysed a political movement. Schools across the US, UK, and Australia moved to ban smartphones. The "phone-free childhood" campaign gained traction. Legislation was proposed. The narrative was clear and actionable: phones are hurting kids, take the phones away.

The story was clean. Perhaps too clean.

Chapter 5: Effect Sizes and Potatoes: The Counterarguments

In 2019, Amy Orben and Andrew Przybylski published a paper in *Nature Human Behaviour* that should have complicated the narrative more than it did. Using the same large datasets Twenge relied on -- including the Monitoring the Future data and the UK Millennium Cohort Study -- they conducted a specification curve analysis. Instead of choosing one way to slice the data (which can produce different results depending on researcher choices), they ran every defensible analytical specification and looked at the distribution of effect sizes.

What they found was sobering. The association between digital technology use and adolescent well-being was negative but tiny: r is approximately 0.04. To put that in context, they compared it to other variables in the same datasets. The negative association between wearing glasses and well-being was similar in magnitude. So was the negative association between eating potatoes and well-being. Screens were statistically associated with lower well-being, but the effect was so small that it explained less than 0.4% of the variance.

This does not mean screens are harmless. It means the effect, measured across the whole population, is very small. Small effects can still matter at scale -- a tiny increase in risk across billions of users is a lot of people.

But it does mean the "smartphones are destroying a generation" framing is hard to justify from the same data that generated it.

Candice Odgers, a developmental psychologist at UC Irvine, published a review in 2024 directly challenging Haidt's causal claims. Her argument was methodological: the studies Haidt cited were overwhelmingly correlational. The few experimental studies that existed showed mixed results. Self-report measures of screen time (which most studies rely on) correlate poorly with actual device-tracked screen time, introducing measurement error that makes the data unreliable. And the mental health crisis Haidt attributed to phones could plausibly be explained by other concurrent trends: the 2008 financial crisis and its lingering effects on families, rising academic pressure, opioid crisis effects on households, climate anxiety, school shootings, and the general erosion of community institutions.

Odgers did not argue that phones are fine. She argued that the evidence for a strong causal relationship is weaker than Haidt presented it, and that policy based on a premature causal conclusion might be addressing the wrong problem. If you take phones out of schools but the actual drivers of adolescent distress are economic precarity and fractured communities, you have done a lot of work for nothing.

The honest position is uncomfortable for everyone. The correlation is real. The timing is suspicious. The effect size is small in population-level data but may be larger in vulnerable subgroups. The causal evidence is genuinely inadequate. And both the "phones are destroying kids" camp and the "it's a moral panic" camp are probably oversimplifying.

A note on subgroup effects: Small average effects can mask larger effects in vulnerable populations. An r of 0.04 across all adolescents could coexist with a much larger effect among, say, girls aged 11-14 who use Instagram heavily. Population-level analyses can obscure vulnerable-group effects. Both Twenge and Haidt have made this argument. The counter-caveat: subgroup analyses risk p-hacking. You can always find a subgroup where the effect is bigger. Whether that reflects real differential vulnerability or statistical noise depends on replication, and the replication data is not yet in.

Chapter 6: The Body Count: Self-Harm and Suicide Data

Whatever the debate about causation, the outcome data is not contested. CDC data shows suicide rates among US adolescents aged 10-24 increased approximately 56% between 2007 and 2017. Among girls aged 10-14, the increase was steeper. Emergency department visits for self-harm among adolescent girls roughly tripled between 2009 and 2015 in CDC WISQARS data.

In England, NHS Digital data showed similar trends. Referrals to child and adolescent mental health services surged. Hospital admissions for self-harm in young women increased markedly after 2010.

These are not survey artefacts. They are not changes in diagnostic criteria (though increased awareness may contribute to referral rates). They represent actual clinical presentations and deaths. Something changed. The question is what.

Twenge and Haidt say it was phones. Orben and Przybylski say the phone effect is too small to explain the magnitude of the change. Odgers says we do not actually know. All three positions have evidence behind them. None of them has proof.

The CDC's Youth Risk Behavior Survey showed sadness and hopelessness among US high school students rose from 26% to 42% between 2009 and 2021. That is not an artefact. That is nearly half of all teenagers reporting persistent feelings of sadness or hopelessness. The same dataset feeds the education research (see Appendix A: *Your Kid's School Was Designed by a King Who Needed Soldiers*) -- the Prussian factory model and the screen-mediated childhood arrived at the same population from different directions, and the outcome data does not distinguish between them.

Chapter 7: Outrage as Engagement: The Polarisation Engine

While the adolescent mental health debate is about whether screens harm individuals, a parallel body of research asks whether they harm collective cognition. This is where the evidence is strongest and the implications are most severe.

Rathje, Van Bavel, and van der Linden published a paper in *PNAS* in 2021 analysing millions of social media posts across Facebook and Twitter. Their finding was stark: posts about political outgroups -- the other side,

the enemy -- received significantly more engagement (likes, shares, comments) than posts about ingroups or neutral topics. Out-group animosity drives engagement. The algorithmic systems that determine what content gets shown are optimised for engagement. Therefore, they systematically amplify content that expresses hostility toward the other side.

This is not a conspiracy. It is a fitness function. Social media platforms make money from attention. Attention is maximised by emotional arousal. Outrage is among the most arousing emotions. Content that makes you angry at *them* keeps you scrolling. Content that makes you think carefully does not. The selection pressure is straightforward, and the result is an information environment that structurally rewards division.

The effect size here is worth noting because it is considerably larger than the screen-time-on-wellbeing effect. Posts containing out-group language received substantially more engagement than comparable posts without it. This is not an $r = 0.04$ finding. The attention economy does not merely correlate with polarisation -- it has a clear mechanical pathway through which algorithmic optimisation for engagement produces amplification of divisive content.

This matters for the thesis of this paper because outrage amplification is a direct mechanism of mateship destruction. Mateship requires in-group solidarity -- not tribal exclusion, but the baseline assumption that the people around you are on your side. When the dominant information platform systematically rewards content that frames other people as enemies, it erodes the social trust on which community depends. You cannot have a functioning neighbourhood watch if half the neighbourhood believes the other half are dangerous ideologues. You cannot organise a community emergency response if the community has been algorithmically sorted into opposing camps. The polarisation engine does not merely make politics uglier. It makes cooperation harder at every level, from national policy to whether you trust the bloke next door to check on your nan.

Chapter 8: Surveillance Capitalism: You Are the Product

The underlying logic of all of this is economic. Shoshana Zuboff named it in 2019: *surveillance capitalism*.

Zuboff's thesis in *The Age of Surveillance Capitalism* is that a new form of capitalism has emerged in which human experience is claimed as free raw material for translation into behavioural data. Some of this data is used to improve services, but a growing portion -- what Zuboff calls the "behavioural surplus" -- is fed into machine intelligence processes that produce predictions about what you will do now, soon, and later. These predictions are traded in a new kind of marketplace that Zuboff calls "behavioural futures markets." The customers are not users. The customers are the businesses that want to know -- and shape -- your future behaviour.

This is not a conspiracy theory. It is described in SEC filings. Google's revenue in 2023 was \$307 billion, primarily from advertising -- selling predictions about your behaviour to companies that want to modify your behaviour. Facebook (Meta) reported \$134 billion. The entire edifice rests on the extraction of behavioural data from human experience and its conversion into advertising revenue.

Social media platforms are free because users are not customers. Users are inventory. The product being sold is attention, and the buyers are advertisers. The consequences flow from the incentive structure. If your revenue depends on time-on-platform, you optimise for time-on-platform. If variable reinforcement maximises time-on-platform, you converge on variable reinforcement. If outrage maximises engagement, you amplify outrage. If adolescents are your most engaged demographic, you design for adolescents. None of this requires malice. It requires a business model that treats human attention as a commodity and an engineering culture that optimises for measurable outcomes.

Tristan Harris's argument was never that tech companies are evil. It was that the incentive structure is misaligned with human well-being, and that the misalignment produces predictable harms regardless of anyone's intentions.

The surveillance capitalism frame matters because it reveals that the attention economy is not a neutral technology being misused. It is a technology that functions exactly as designed. The design is: extract maximum attention, convert it to behavioural predictions, sell those predictions. Every feature that makes the platform more compulsive, more outrage-amplifying, more isolating is working correctly from the perspective of the business model. The harms are not bugs. They are the product working as intended.

Chapter 9: Stolen Focus: The Fragmentation of Attention

Johann Hari's 2022 book *Stolen Focus: Why You Can't Pay Attention -- and How to Think Deeply Again* extends the analysis from adolescents to the entire population. Hari identifies twelve factors that he argues are degrading human attention, including the technology-mediated factors discussed above but also: rising stress, declining sleep quality, deteriorating diet, environmental pollutants, and the collapse of sustained reading.

Hari's central contribution is the argument that the attention crisis is not merely individual but collective. A society that cannot pay attention cannot govern itself. Democracy requires citizens who can follow an argument, evaluate evidence, and sustain focus long enough to understand policy. If the information environment is optimised to fragment attention into three-second intervals -- the average time a user spends on a TikTok video before swiping -- then the cognitive infrastructure for self-governance degrades.

This connects directly to the direct democracy question (Goal 1 (direct democracy) (direct democracy)). The attention economy's defenders sometimes argue that social media democratises information -- everyone has a voice, everyone can participate. Hari's counterargument is that having a voice is meaningless if nobody can listen. The attention span required to evaluate a referendum proposal, understand its implications, and vote thoughtfully is precisely the capacity that the attention economy degrades. Switzerland runs four referendums per year and has done so for 178 years. The Swiss system works because citizens can pay attention long enough to read a ballot pamphlet. If the attention economy has reduced the average sustained focus to the length of a TikTok, direct democracy becomes a rubber stamp for whoever writes the most emotionally compelling three-second pitch.

Hari also documents the physical dimension of attention loss. The average American office worker is interrupted every three minutes. After an interruption, it takes an average of 23 minutes to return to the same depth of focus (Mark et al., 2008). The cumulative effect is that sustained deep thought -- the kind required for creative work, complex problem-solving, or genuine self-reflection -- becomes structurally impossible within the attention economy's information environment. You are not distracted because you lack willpower. You are distracted because a \$2 trillion industry has bet its entire revenue model on keeping you distracted.

Chapter 10: The Displacement Hypothesis: What Screens Replace

The debate about whether screens cause harm often misses the simpler question: what do screens *replace*?

Time is finite. Every hour spent scrolling is an hour not spent doing something else. The displacement hypothesis -- that screen time's primary harm comes not from what it does but from what it prevents -- is difficult to test experimentally but has strong theoretical support and growing observational evidence.

What screens displace:

Unstructured play. The play deprivation research (see Appendix A) documents that free play among children has declined steadily since the 1960s (Gray, 2011). Smartphones accelerated this decline by providing an alternative activity that is always available, always rewarding, and requires no physical movement, no social negotiation, and no risk. The things that make play developmentally crucial -- physical challenge, social calibration, managed danger, boredom-as-catalyst -- are precisely the things screens eliminate.

Face-to-face socialisation. Twenge's data shows that the amount of time American teenagers spend with friends in person has declined sharply since 2010. This is not because they stopped wanting to see friends. It is because the social reward system -- likes, comments, streaks -- can be accessed without leaving the bedroom. The screen provides a supernormal social stimulus (Chapter 2) that makes the real thing feel slow, effortful, and unrewarding by comparison. The friend in the room cannot compete with the feed on the phone.

Boredom. This one matters more than it sounds. Boredom is not the absence of stimulation. It is a signal from the motivational system that the current activity is not engaging enough and that something else -- something self-directed, something creative, something new -- should be sought. Boredom drives exploration. It drives play. It drives the kind of unfocused daydreaming that precedes creative insight. Smartphones eliminated boredom. Every waiting room, every bus ride, every moment of downtime is now filled with content. The cost is that the internal drive to *do something* -- to build, to explore, to connect with the person next to you -- never activates. The screen pre-empts it.

Sleep. The evidence here is less contested. Screens in bedrooms delay sleep onset and reduce sleep quality through blue light exposure, arousal from content, and the variable reinforcement loop that makes it hard to

stop. Sleep deprivation has well-documented effects on mood, cognition, and emotional regulation. Some portion of the adolescent mental health signal may be a sleep signal.

Community participation. This is the displacement that matters most for this paper. The time adults spend on social media is time they do not spend at the cricket club, the pub, the community meeting, the front yard. Robert Putnam documented the decline of civic participation in *Bowling Alone* (2000) before smartphones existed -- the trend started with television. Social media accelerated it by providing a substitute that *feels* like community participation (posting, commenting, sharing) but lacks the physical co-presence, the unscripted interaction, the reciprocal obligation, and the ambient awareness of other people's lives that make real communities function.

Chapter 11: The Isolation Machine: Mateship Destroyed by Design

This is the thesis of the paper, and it is simple enough to state plainly.

The attention economy is an isolation machine. It replaced human social structures with a screen-mediated substitute that is optimised to extract attention rather than to sustain connection. The replacement was invisible because the substitute triggers the same reward circuits as the real thing -- supernormal social stimuli feel like connection, they just do not function as connection. The isolation it produces is not felt as loneliness in the moment of scrolling. It is felt when your nan falls and there is nobody within 60 seconds who knows her name.

The machine works like this:

1. Capture attention using supernormal stimuli and variable reinforcement (Chapters 2-3).
2. Displace in-person connection by providing a more intensely rewarding substitute (Chapter 10).
3. Amplify division by algorithmically rewarding outgroup animosity (Chapter 7).
4. Erode the capacity for sustained thought required for community self-governance (Chapter 9).
5. Extract the behavioural surplus and sell it as predictions to advertisers (Chapter 8).

6. Repeat, because the business model requires it.

Each step is individually documented. The synthesis is what matters: these are not six separate problems. They are six components of a single machine whose function is to convert human social connection into advertising revenue. The input is mateship. The output is profit. The waste product is isolation.

The people who built this machine did not set out to destroy mateship. They set out to maximise engagement. The destruction of mateship is a side effect -- an externality, in economic language. But externalities are real. The carbon dioxide is in the atmosphere whether the factory owner intended to put it there or not. The isolation is in the community whether Mark Zuckerberg intended to put it there or not.

And unlike carbon dioxide, the isolation machine has a particularly cruel feedback loop. The lonelier you are, the more you scroll. The more you scroll, the lonelier you become. The more you scroll, the more data you generate. The more data you generate, the more precisely the algorithm can target you. The more precisely it targets you, the more engagement it extracts. The machine feeds on the wound it creates.

Chapter 12: Where This Leaves Us

The evidence, taken honestly, supports a position more nuanced than either camp in the public debate tends to offer.

What is well-established: Adolescent mental health declined sharply after approximately 2010, across multiple countries, with girls disproportionately affected. Social media platforms use design patterns drawn from behavioural psychology that maximise compulsive use. Algorithmic amplification of outrage content is empirically documented and mechanistically understood. The attention economy business model creates structural incentives for engagement maximisation regardless of user well-being. Surveillance capitalism extracts behavioural data from human experience and trades it on prediction markets. Screen time displaces play, sleep, boredom, face-to-face interaction, and community participation.

What is genuinely debated: Whether smartphones *caused* the mental health decline or merely coincided with it. Whether the population-level effect size (very small) captures real harm concentrated in vulnerable

subgroups, or whether it reflects a genuinely minor influence. Whether removing phones would improve outcomes or whether the drivers are elsewhere. Whether historical media panics (novels, radio, TV, video games) invalidate the current concern or whether this time really is different because of the combination of portability, always-on connectivity, algorithmic personalisation, and a business model that treats attention as commodity.

What is honest: We do not know the full causal story yet. The correlation is real and concerning. The effect size data says it is small. The hospital data says something is very wrong. These facts coexist and no one has resolved the tension between them.

What is clear regardless of the causal debate: The attention economy replaced face-to-face community structures with a screen-mediated substitute that optimises for engagement, not for human well-being. Whether or not phones *caused* the mental health crisis, they structurally replaced the community infrastructure -- the pubs, the clubs, the front yards, the unscheduled time with people who know your name -- that constitutes the social safety net. You do not need to prove that Instagram causes depression to observe that it replaced the cricket club. The displacement is visible. The community structures are gone. The screens remain.

The precautionary instinct says: if something might be harming people and the cost of restricting it is low, restrict it. The scientific instinct says: if the evidence does not support a strong causal claim, do not build policy on one. Both instincts are rational. They pull in different directions. That is the actual state of knowledge, and pretending otherwise -- in either direction -- is dishonest.

But there is a third option beyond restricting or permitting. Build something better. Build technology that serves connection instead of extracting it. Build a mesh network where the user is the infrastructure and there are no advertisers to sell attention to (Goal 8 (free internet) (free internet)). Build a \$29 ring that summons your neighbours in 60 seconds when your nan falls (Goal 13 (\$29 emergency ring) (\$29 emergency ring)). Build public spaces with monkey bars and climbing walls so that bodies move and people encounter each other in physical space (Goal 11 (physical infrastructure) (physical infrastructure)). Build schools around play and curiosity instead of compliance and screens (Goal 12 (play-based education) (play-based education)). Build direct democracy so that the attention required for self-governance has somewhere to go (Goal 1 (direct democracy) (direct democracy)).

Do not fight the attention economy. Make it irrelevant. Build the thing that makes the screen unnecessary. Not by removing it, but by providing something so much better that people put it down on their own.

The attention economy is not the disease. It is the symptom of a society that dismantled every structure of human connection and then sold people a subscription to a simulacrum. The cure is not to ban the simulacrum. The cure is to rebuild the structures.

References

1. Alter, A. (2017). *Irresistible: The Rise of Addictive Technology and the Business of Keeping Us Hooked*. Penguin Press.
1. Barrett, D. (2010). *Supernormal Stimuli: How Primal Urges Overran Their Evolutionary Purpose*. W.W. Norton.
1. Gray, P. (2011). "The Decline of Play and the Rise of Psychopathology in Children and Adolescents." *American Journal of Play*, 3(4), 443-463.
1. Haidt, J. (2024). *The Anxious Generation: How the Great Rewiring of Childhood Is Causing an Epidemic of Mental Illness*. Penguin Press.
1. Hari, J. (2022). *Stolen Focus: Why You Can't Pay Attention -- and How to Think Deeply Again*. Crown.
1. Harris, T. (2016). "How Technology is Hijacking Your Mind -- from a Magician and Google Design Ethicist." Blog essay / presentations. Led to founding of Center for Humane Technology (2018).
1. Mark, G., Gudith, D., & Klocke, U. (2008). "The cost of interrupted work: more speed and stress." *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 107-110. <https://doi.org/10.1145/1357054.1357072>
1. Odgers, C.L. (2024). "The Great Rewiring: Is Social Media Really Behind an Epidemic of Teenage Mental Illness?" *Nature*, 628, 29-30. <https://doi.org/10.1038/d41586-024-00902-2>
1. Orben, A., & Przybylski, A.K. (2019). "The association between adolescent well-being and digital technology use." *Nature Human Behaviour*, 3, 173-182. <https://doi.org/10.1038/s41562-018-0506-1>

1. Putnam, R.D. (2000). *Bowling Alone: The Collapse and Revival of American Community*. Simon & Schuster.
1. Rathje, S., Van Bavel, J.J., & van der Linden, S. (2021). "Out-group animosity drives engagement on social media." *Proceedings of the National Academy of Sciences*, 118(26), e2024292118.
<https://doi.org/10.1073/pnas.2024292118>
1. Skinner, B.F. (1953). *Science and Human Behavior*. New York: Macmillan. See also: *The Behavior of Organisms* (1938) for the foundational operant conditioning experiments.
1. Tinbergen, N. (1951). *The Study of Instinct*. Oxford University Press.
1. Twenge, J.M. (2017). *iGen: Why Today's Super-Connected Kids Are Growing Up Less Rebellious, More Tolerant, Less Happy -- and Completely Unprepared for Adulthood*. Atria Books.
1. Twenge, J.M., Joiner, T.E., Rogers, M.L., & Martin, G.N. (2018). "Increases in Depressive Symptoms, Suicide-Related Outcomes, and Suicide Rates Among U.S. Adolescents After 2010 and Links to Increased New Media Screen Time." *Clinical Psychological Science*, 6(1), 3-17.
<https://doi.org/10.1177/2167702617723376>
1. Zuboff, S. (2019). *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*. PublicAffairs.

Data Sources

- CDC WISQARS (Web-based Injury Statistics Query and Reporting System). Fatal and nonfatal injury data including self-harm hospitalisations. <https://wisqars.cdc.gov>
- CDC Youth Risk Behavior Surveillance System (YRBSS). Biennial survey of US high school students. <https://www.cdc.gov/yrebs/>
- NHS Digital (England). Mental Health of Children and Young People surveys (2017, 2020, 2021, 2022). <https://digital.nhs.uk>

- Stuart Brown's 6,000+ violent offender play histories (Chapter 4 of this paper references Haidt's developmental argument, which rests on this evidence)
- Panksepp's PLAY circuit and Ritalin suppression -- the drug we give children to sit still works partly by shutting down the system that makes them want to play
- Gray (2011) decline of free play since 1960s -- the same timeline as the mental health data

Your Kid's School Was Designed by a King Who Needed Soldiers

Location: `../education_prussian_model/`

Relevance: The Prussian education model and the attention economy arrived at the same population from different directions and produced the same outcome: children sitting still, looking at a screen, receiving content, producing compliance. The factory model of education primed children for the attention economy by teaching them that the correct behaviour is to sit, receive, and respond on command. Screens did not need to train children to be passive consumers. The school system had already done it.

Key connections:

- CDC YRBS data appears in both papers (sadness/hopelessness 26% to 42%, 2009-2021)
- Finland's education model (play-based, minimal homework, high outcomes) is the structural opposite of both the Prussian model and the attention economy
- Creativity decline (Kim, 2011, Torrance Tests) connects to attention fragmentation (Chapter 9 of this paper)

The Bystander Effect and Community Emergency Response

Location: `../bystander_effect/`

Relevance: The attention economy does not merely reduce connection. It actively trains people to observe without acting -- to scroll past suffering. The bystander effect (diffusion of responsibility in the presence of others) is amplified by social media, where you "see" thousands of people's distress but have no physical capacity to respond. The \$29 ring (Goal 13 (\$29 emergency ring) (\$29 emergency ring)) is the structural antidote: it creates a system where seeing distress is coupled with the physical ability to respond, within 60 seconds, in person.

Death, Terror Management, and the Culture of Control

Location: `../death_terror_management/`

Relevance: Terror management theory (Greenberg, Pyszczynski, & Solomon) predicts that mortality salience increases in-group bias and out-group hostility. Social media's outrage amplification (Chapter 7) may interact with ambient death anxiety -- the news feed that shows you shootings, pandemics, and wars increases mortality salience, which increases tribal defensiveness, which increases engagement with outgroup-hostile content, which the algorithm amplifies. A feedback loop between existential anxiety and algorithmic polarisation.

Social Group Scaling and the Limits of Connection

Location: `../social_group_scaling/`

Relevance: Dunbar's number (~150) was long cited as the cognitive ceiling on stable social relationships. Lindenfors et al. (2021) demolished it – re-running Dunbar's own primate neocortex regression produced a 95% confidence interval of 2 to 520. The number is meaningless. The Ripple model replaces it: accountability = $1/\text{distance}$, everyone connected, weighted by physical proximity. No cap. No boundary. The person in front of you is the right one. Social media inverts this completely. It creates the illusion of connection to thousands – 1,000 "friends," 10,000 followers – but the relationships are parasocial, not reciprocal. The screen provides the signal (likes, comments) without the bandwidth of actual proximity-based relationship. The result is a person who feels connected to thousands and is known by none of them. The isolation machine works by replacing the proximity gradient – the one that actually generates accountability and care – with a flat, surveilled, algorithmically curated simulacrum of connection.

Appendix B: The 14 Goals and Their Relationship to This Paper

GOAL

CONNECTION TO THE ATTENTION ECONOMY

1. Direct democracy	Requires sustained attention and informed citizens. The attention economy degrades both.
2. 22-hour work week	Screen time is unpaid labour -- your attention, mined and resold. The hours "spent" on social media are not leisure.
5. Fire all police	The attention economy amplifies outgroup hostility, which drives the "tough on crime" sentiment that sustains punitive policing.
7. Legalise drugs	Addiction is isolation. The screen replaces the connection that prevents addiction (Hari, 2015: "The opposite of addiction is connection").
8. Free internet	If the internet is free because you are the product, it is not free. Mesh networking removes the advertising business model.
9. No foreign investment in housing	The attention economy distracts from structural economic issues. Housing unaffordability is an attention-economy-invisible crisis.
10. Food contains only things proven safe	The attention economy's precautionary principle failure is the food system's precautionary principle failure. Both treat human welfare as an externality.
11. Monkey bars at every bus stop	Physical infrastructure for physical bodies. The antidote to screen-mediated existence.
12. Play-based schools	The Prussian model and the attention economy are the same machine applied at different ages. Play-based education is the structural opposite of both.
13. \$29 ring	Technology that serves connection instead of extracting it. The ring summons neighbours; the phone summons advertisers.
14. Cancer is 90% preventable	Attention fragmentation prevents people from learning and acting on prevention information. The machine that sells you poison also prevents you from reading the label.

Draft. Statistics cited herein require fact-checking against primary sources. Effect sizes and figures from secondary reporting need verification against original papers. See Source Verification Notes for specific items requiring confirmation.

This paper is part of the OMXUS research series. It is not neutral. It serves the 14 goals. Strong rhetoric serving power diffusion and human flourishing is not epistemic abuse. Factual errors, however, are -- and any found should be corrected immediately.