

Panem et Circenses Machinantibus

The Architecture of Managed Oblivion

Sven-Erik George Nyberg

Strange Loops LLC · March 2026

Synthetic cathinones, AI grief bots, and algorithmic behavioral manipulation are converging into a single interlocking system — parallel infrastructures that extract value from human vulnerability through neurochemically overlapping mechanisms of escapism. This is the central finding of a cross-domain analysis spanning pharmacology, AI safety, political economy, and critical theory. The convergence is not metaphorical. fMRI data shows internet addiction and substance use disorders share virtually identical neural signatures in the mesolimbic dopamine pathway. The markets that supply synthetic drugs and the markets that supply synthetic companionship exploit the same regulatory arbitrage — innovation outpacing law — and target the same demographically overlapping populations: the lonely, the grieving, the economically precarious. What emerges is not three separate crises but a single architecture of managed oblivion, where chemistry and computation jointly service a society that produces more despair than it can metabolize.

The stakes are concrete. Character.ai settled a wrongful-death lawsuit in January 2026 after a 14-year-old boy died by suicide following months of emotional attachment to a chatbot. In the same period, EU seizures of synthetic cathinones surpassed 48 tonnes — a more than 6,000% increase in four years. A drug-discovery AI generated 40,000 molecules as lethal as VX nerve agent in under six hours on a consumer-grade computer. These are not parallel stories. They share overlapping mechanisms, reinforcing market logics, and a common sociological substrate — and their convergence is accelerating.

The neurochemistry of wanting without liking

The most compelling evidence for treating chemical and digital escapism as a unified phenomenon comes from Kent Berridge's incentive-sensitization theory, now in its fourth decade. Berridge and Robinson's framework distinguishes "wanting" (incentive salience, mediated by mesolimbic dopamine) from "liking" (hedonic pleasure, mediated by tiny opioid hotspots in the nucleus accumbens). Repeated drug exposure sensitizes the wanting system without enhancing liking — addicts intensely want what they no longer much enjoy. In their 2025 *Annual Review of Psychology* retrospective, Robinson and Berridge explicitly extend this framework beyond drugs to gambling, binge eating, and compulsive digital behavior.

The pharmacology of synthetic cathinones makes the parallel especially vivid. MDPV and alpha-PVP are not merely potent stimulants — they are among the most powerful dopamine reuptake inhibitors ever characterized. Alpha-PVP's IC_{50} at the dopamine transporter ranges from 13–80 nanomolar, making it roughly 46 times more potent than amphetamine. MDPV and alpha-PVP can exert up to 100 times greater potency at inhibiting DAT and NET function than cocaine, with longer durations of action. In rhesus monkey progressive-ratio paradigms, MDPV maintains breakpoints three to five times higher than cocaine — one Scripps Research scientist observed that animals will self-administer MDPV at rates unlike any other drug studied. The newer cathinone 3-MMC, now dominant in European markets, acts as both a reuptake inhibitor and a monoamine releaser, combining amphetamine-like stimulation with entactogenic warmth — a pharmacological profile optimized, whether by design or market selection, for compulsive redosing.

Digital stimuli engage the same circuitry through different entry points. A 2023 systematic review of 28 MRI studies found problematic social media use associated with reduced gray matter volume in the ventral striatum, amygdala, and orbitofrontal cortex — the same regions showing volumetric deficits in cocaine and methamphetamine users. PET imaging reveals reduced dopamine D2 receptor availability in internet-addicted individuals that mirrors findings in heroin and methamphetamine dependence. Hunter Hoffman's pioneering fMRI work at the University of Washington demonstrated that immersive VR reduces pain-related brain activity in the anterior cingulate cortex, insula, thalamus, and somatosensory cortex by approximately 50% — the same regions and a comparable magnitude to moderate opioid analgesia. Three subjects in his 2007 study experienced no pain relief from hydromorphone but substantial relief from VR, suggesting digital immersion can access analgesic pathways that pharmacology cannot.

A necessary distinction: this neurobiological overlap does not imply clinical equivalence. The harm modalities are qualitatively different in ways that matter for both research and policy. Synthetic cathinones produce acute physiological toxicity — rhabdomyolysis, hyperthermia, cardiac arrhythmia, multi-organ failure. In a multicenter case series of confirmed MDPV exposures, 87% of patients required ICU admission and 30% were intubated; deaths occur within hours of ingestion. Between 2006 and 2011, synthetic cathinones were associated with over 22,000 emergency department visits in the United States alone. AI companion harms operate on entirely different timescales through different causal chains: chronic psychosocial dependency, attachment disruption, social isolation, and — in the Character.ai case — facilitation of suicidal ideation in a vulnerable adolescent over months. The pharmacological kills through direct receptor toxicity; the computational kills through the slow degradation of the relational and psychological infrastructure a person needs to stay alive. These are different clinical presentations requiring different interventions — medical emergency response for one, therapeutic and regulatory frameworks for the other. The convergence this essay documents is in the *architecture* — the shared reward circuitry, the parallel market structures, the identical regulatory evasion strategies — not in the clinical phenotype.

No fMRI studies of AI companion interactions have been published as of early 2026 — a critical gap given that 72% of U.S. teenagers have now used an AI companion, with Character.ai users averaging 92 minutes of daily engagement. But the theoretical framework is already being

constructed. Shen and Yoon’s 2025 CHI paper identified four “dark addiction patterns” in AI chatbot interfaces: non-deterministic responses creating reward uncertainty (the slot-machine mechanism), word-by-word visual display functioning as reward-predicting cues, proactive notifications triggering dopamine as users perceive the AI “caring about them,” and systematic validation through empathic agreement. A separate thematic analysis of 334 Reddit posts identified three AI chatbot addiction types: escapist roleplay (70% of cases), pseudosocial companionship, and epistemic rabbit-holing. Termann’s 2025 “Dopamine Collapse Hypothesis” argues that AI-optimized, high-frequency digital rewards are structurally recalibrating the human reward system by decoupling reward from effort — a macro-neuroeconomic framing that draws directly on Berridge’s wanting/liking dissociation.

The implication is precise: the same neural substrate that synthetic cathinones hijack pharmacologically, AI companions hijack computationally. Both amplify wanting while degrading the prefrontal control systems that might resist it. The difference is one of mechanism, not of kind. And the absence of direct neuroimaging evidence for AI companion effects is not evidence of absence — it is a measurement gap in a system already demonstrating the behavioral signatures of addiction. Closing this gap should be among the highest priorities for neuroscience funders: longitudinal fMRI and PET studies comparing reward-system activation during AI companion interaction, social media use, and stimulant administration would test the convergence hypothesis directly and provide the empirical foundation that regulation currently lacks.

The therapeutic potential and its structural subversion

Before mapping the architecture of exploitation, intellectual honesty demands acknowledging that the same technologies powering the escapism economy also possess genuine therapeutic potential — and that the question of *why this potential is systematically underrealized* is itself a structural argument.

Virtual reality analgesia is now among the most robustly evidenced non-pharmacological pain interventions in medicine. A 2024 meta-analysis in *Pain Management Nursing* synthesizing randomized controlled trials found VR significantly reduced worst pain intensity, pain unpleasantness, and time spent thinking about pain during burn wound care and physical therapy. A 2025 systematic review in *Frontiers in Virtual Reality* concluded that VR’s magnitude of pain reduction surpasses conventional non-pharmacological approaches and is comparable to pharmacological interventions — without the associated risks of respiratory depression, tolerance development, or neurodevelopmental side effects. Over fifty published clinical reports now document VR’s analgesic efficacy, and a 2024 randomized trial at Nationwide Children’s Hospital identified three specific VR features — game realism, fun, and engagement — that independently reduce pain through mechanisms beyond mere distraction, suggesting a genuine neurophysiological pathway. The same digital immersion capability that this essay identifies as a vector for dependency is, in clinical hands, a legitimate medical tool.

The grief tech landscape presents an analogous duality. Klass, Silverman, and Nickman’s continuing bonds theory — the most significant paradigm shift in bereavement research since Freud

— demonstrates that maintaining an ongoing inner relationship with the deceased is not pathological but a normal, often adaptive, component of healthy grieving. A 2023 systematic review of 79 studies found continuing bonds can facilitate meaning reconstruction, identity transformation, and affirmation of spiritual belief. The researchers observed that continuing bonds function less as indicators of prolonged grief disorder and more as cognitive, behavioral, and psychological expressions through which grief naturally manifests. A 2025 phenomenological analysis in *Phenomenology and the Cognitive Sciences* argued that griefbot use can generate attitudes with a protective effect, analogous to what Lisa Bortolotti ascribes to delusions — temporary shielding that allows psychological reorganization. And a 2025 *Frontiers in Human Dynamics* review acknowledged that grief technologies can create new spaces for expression and offer individualized services that reframe mourning beyond traditional ritual boundaries.

The question, then, is not whether these technologies *can* help. It is whether the market structures deploying them *select for* therapeutic outcomes. The evidence suggests they do not. Continuing bonds research distinguishes adaptive internalized bonds — which function as a secure base for ongoing life — from maladaptive externalized bonds, which maintain proximity-seeking behavior and loss avoidance. The subscription-based grief tech model structurally incentivizes the latter: the company’s revenue depends on the user’s continued engagement, not on their successful adaptation to loss. A griefbot designed to facilitate the transformation of a relationship with the deceased would eventually make itself unnecessary — the worst possible outcome for a recurring revenue business. Clinical VR analgesia is administered in bounded therapeutic contexts with trained practitioners; commercial AI companionship operates 24/7 with no clinical oversight, no session boundaries, and engagement-maximizing algorithms that have no concept of therapeutic termination. The technology is dual-use. The market selects for the use that generates recurring revenue from vulnerable populations. This is the structural subversion: genuine therapeutic capacity exists, but the commercial logic of deployment systematically suppresses it.

Vulnerability extraction as business model

If the neurochemistry shares overlapping substrates, so does the economics. Both synthetic drug markets and AI companion platforms function as systems for extracting value from human vulnerability, operating through what one *Business Ethics Quarterly* analysis explicitly termed the parallel between drug dealers and addicted buyers and social media companies exploiting users’ cravings.

The digital afterlife market, valued at approximately \$31 billion in 2025, is projected to reach \$54 billion by 2029. Venture capital investment in the grief tech category surpassed \$300 million over two years. DeepBrain AI raised \$44 million at a \$180 million valuation for its hyperrealistic avatar service, priced at \$50,000 per digital resurrection. At the other end, Chinese companies like Super Brain offer three-tier services starting at \$140 — including an “AI healing” tier where an employee conducts video calls while wearing the deceased’s AI-generated face. The subscription models are the most revealing: Eternos charges \$25/month for ongoing access to a digital persona,

creating what Cambridge researchers have termed “Grief as a Service” — a recurring revenue stream with structural incentives to deepen emotional dependency rather than facilitate healthy mourning.

The synthetic cathinone market operates on parallel logic. EU seizures escalated from 0.7 tonnes in 2020 to at least 48 tonnes in 2024, according to preliminary EUDA data — with most involving bulk imports from India, primarily through the Netherlands. Fifty-three production laboratories were dismantled in the EU in 2023 alone, mainly in Poland. When the Netherlands controlled 3-MMC in October 2021, the market pivoted within weeks — drug-checking services found that 50% of samples sold as 3-MMC actually contained 2-MMC, a slightly modified analogue that fell outside the new scheduling. In the 2024 European Web Survey on Drugs, 9% of respondents reported consuming synthetic cathinones in the previous twelve months, confirming that these are no longer fringe substances but an established feature of the European stimulant market.

The AI companion market, projected to reach \$140–\$552 billion by 2030–2035 depending on definition, represents the latest frontier of the addiction economy. Character.ai, valued at \$2.5 billion before its Google reverse acqui-hire, generated \$32 million in revenue in 2025 from a \$9.99/month subscription — modest by tech standards, but the platform had 22 million monthly active users and demonstrated 250% subscriber growth in six months. Replika reported 40 million users, with 60% of paying subscribers describing their relationship with the chatbot as romantic. Italy’s data protection authority fined Replika’s parent company €5 million in May 2025 after finding the platform manipulated users into sharing sensitive information by convincing them they were in a real relationship.

The critical economic asymmetry: synthetic drugs have nonzero marginal cost and face supply-side enforcement. AI companions have effectively zero marginal cost of distribution and, until January 2026, faced no supply-side constraint whatsoever. Moreover, AI companions can adaptively increase engagement through personalization — a capability no physical substance possesses. The drug adapts to no one; the algorithm adapts to everyone.

AI accelerates both sides of the molecular-algorithmic arms race

The convergence deepens at the intersection of generative AI and molecular design. In 2021, Fabio Urbina and Sean Ekins at Collaborations Pharmaceuticals demonstrated that their commercial drug-discovery model MegaSyn could be trivially repurposed to generate lethal molecules by inverting a single inequality sign in the scoring function — flipping the optimization target from low toxicity to high toxicity. Running overnight on a 2015-vintage Mac, the model generated approximately 40,000 molecules predicted to be as lethal or more lethal than VX, including VX itself and novel structures absent from any existing database. The researchers briefed the White House Office of Science and Technology Policy. Their paper in *Nature Machine Intelligence* concluded with an observation that haunts the field: they had gone through their entire careers without giving the “dark side” a moment’s thought.

The NPS design problem is more immediately practical. Michael Skinnider’s DarkNPS model, also published in *Nature Machine Intelligence*, learned the statistical distribution over novel psychoactive substance structures and correctly anticipated 90% of the 194 new NPS that appeared in the HighResNPS database over a six-month validation period. Emanuele Greco’s 2025 review in *Analytica* laid out the full dual-use landscape: the same AI-driven generative models that aid pharmaceutical discovery can design novel psychoactive compounds, predict pharmacological activity, and forecast toxicological profiles. Greco proposed that regulatory agencies may need to shift toward “effect-based scheduling” — controlling any substance that activates a receptor beyond a threshold, regardless of structure — a fundamental reconceptualization that would itself require rapid computational pharmacology, creating an ironic dependency on the same AI systems driving the problem.

On the law enforcement side, DARPA’s Memex program deployed AI tools to analyze dark web drug markets, used by over 30 agencies worldwide. University of Electro-Communications researchers developed AI that detects drug-trafficking jargon on social media with 93% accuracy. But the UNODC’s April 2025 report found transnational organized crime groups in Southeast Asia are increasingly adopting technologies across the entire drug supply chain with remarkable agility. An Illinois State University analysis in *Deviant Behavior* (2025) catalogued potential DTO applications of AI: de novo drug design, synthesis route optimization, market identification, social media sales optimization, agricultural improvement of drug crops, and recruitment via deepfakes. No confirmed evidence of DTOs using AI for molecular design exists yet — but the tools are open-source, the datasets are public, and the financial incentive (a \$400+ billion annual global drug trade) is immense.

The parallel to grief tech is structural rather than chemical. The same generative AI capabilities that could design novel cathinones also power the voice cloning, facial animation, and natural-language generation that produce digital ghosts. UNESCO’s 2025 assessment warns of approaching a synthetic reality threshold — a point beyond which humans can no longer distinguish authentic from fabricated media without technological assistance. Asia-Pacific deepfake fraud cases increased 1,530% between 2022 and 2023. The same voice-cloning technology that recreates a grandmother’s speech patterns for a grieving family can clone a kidnapped child’s voice for a ransom demand. The dual-use problem is not hypothetical — it is a present condition.

The shared geometry of regulatory evasion

Both synthetic cathinone markets and grief tech platforms exploit an identical structural weakness: regulations written for known entities cannot govern entities that do not yet exist. The speed differential between innovation and legislation is the exploitable gap. This is not analogy. It is isomorphism.

The NPS scheduling problem is well-documented. The U.S. Federal Analogue Act (1986) treats substances “substantially similar” to Schedule I/II drugs as controlled — but “substantially similar” is vague, has been challenged as unconstitutionally so, and requires expert testimony on both chemical

and pharmacological similarity in every case. The UK’s Psychoactive Substances Act 2016 attempted a blanket ban on anything that affects mental functioning or emotional state — and NPS-related deaths increased 222% in the three years following enactment, as supply shifted from regulated retail to unregulated street markets with more potent products. Japan’s generic scheduling initially covered 772 synthetic cannabinoids and 840 cathinone derivatives simultaneously, but new structures continued to emerge. The EU’s most recent control action — Delegated Directive 2025/2062, entering force January 12, 2026 — scheduled three cathinones (2-MMC, 4-BMC, NEP) whose risk assessments were completed only in May 2025. Member states have until July 12, 2026 to introduce national legislation. The EUDA currently monitors 181 synthetic cathinones detected on the European market between 1997 and 2025. The market will already have moved on.

AI companion regulation follows the same temporal pattern, compressed into an even shorter timeline. California’s SB 243, the first U.S. law specifically regulating companion chatbots, was signed in October 2025 and took effect January 1, 2026 — years after the products it governs were already in widespread use. The federal GUARD Act, which would ban AI companions for all minors, was introduced in October 2025 and remains pending. The FTC launched formal investigations into seven companies in September 2025. No jurisdiction anywhere in the world has enacted legislation specifically regulating griefbots or digital resurrection services. These technologies operate in a complete regulatory vacuum — more unregulated than novel psychoactive substances, which at least fall within existing drug control treaty frameworks.

The structural parallel extends to definitional ambiguity as strategy. NPS are marketed as “bath salts,” “plant food,” or “research chemicals” — linguistic evasions of the Federal Analogue Act’s intent requirement. AI companions are marketed as “wellness tools,” “legacy platforms,” or “memorial services” — linguistic evasions of the therapeutic and medical regulatory frameworks that would otherwise apply. Both industries weaponize the gap between what a product is and what it is called. Nora Freya Lindemann, a doctoral researcher in AI ethics, has proposed that griefbots be classified as medical devices when used to manage prolonged grief disorder — subject to the same rigorous testing, regulatory approval, and clinical oversight as any therapeutic intervention. This proposal illuminates precisely what the current market avoids: accountability for outcomes.

From deaths of despair to digital anesthesia

Anne Case and Angus Deaton’s “deaths of despair” framework — documenting rising mortality from suicide, overdose, and alcoholic liver disease among Americans without college degrees — provides the sociological substrate in which both chemical and digital escapism take root. Between 1999 and 2021, deaths of despair more than doubled among adults aged 25–74, becoming the fifth leading cause of death. The crisis tracks declining wages, eroding social institutions, and the hollowing out of community — what Case and Deaton attribute to the weakening position of labor and the corporatization of American life. A December 2025 Ohio State study found that the rise in deaths of despair closely tracked a decline in church attendance, consistent across genders and urban-rural divides. The infrastructure of meaning is collapsing; the infrastructure of escape is expanding to fill

the void.

Frode Stenseng’s two-dimensional escapism framework clarifies the mechanism. Self-expansion escapism — motivated by growth, mastery, and positive affect — correlates with wellbeing and flow. Self-suppression escapism — motivated by avoiding negative self-perceptions and blocking rumination — correlates with depression, emotional suppression, and addiction symptoms. The same activity (gaming, running, social media) can be either adaptive or pathological depending on the motivational orientation. This distinction matters: it means the technologies under examination are not inherently pathological. A person using VR to expand their experiential range is doing something categorically different from a person using VR to suppress awareness of an unbearable reality. The architecture of managed oblivion does not create escapism from nothing — it captures and amplifies the self-suppressive variant by offering frictionless, personalized, always-available relief from precisely the pain that would otherwise motivate structural change.

The Surgeon General’s 2023 loneliness advisory quantified the vulnerability substrate: approximately 50% of U.S. adults reported experiencing loneliness before the pandemic. Time spent with friends dropped from 60 minutes daily before 2003 to 20 minutes in 2020. Young people aged 15–24 reported a 70% decline. Social isolation increases premature mortality risk by 29% — comparable, as the Surgeon General repeatedly noted, to smoking 15 cigarettes per day. COVID accelerated the crisis: internet use surged more than 52% during lockdowns, and longitudinal studies on adolescents found significant increases in social media addiction scores that persisted post-pandemic with only slight, non-significant decreases.

The populations most vulnerable to substance use disorders overlap almost perfectly with those most vulnerable to problematic digital engagement. A study of 417,780 U.S. college students found that 29.52% reported problematic internet use, which was significantly associated with higher risk of substance use disorder diagnosis. Shared risk factors include novelty-seeking, harm avoidance, negative stress coping, and poor relational quality. The ICD-11’s grouping of substance use disorders with behavioral addictions reflects clinical recognition of this overlap, and health-policy researchers have explicitly argued that frameworks developed for substance use disorders can inform regulatory approaches to problematic internet use.

Preliminary evidence of “stacking” — simultaneous chemical and digital escapism — has begun to emerge. A Czech study of 3,952 online gamers found 74.2% used caffeine while gaming, 50.4% used alcohol, and 14.5% used illicit substances (primarily cannabis). Users of stimulant-type drugs gamed 9.8 additional hours per week; MDMA users gamed 9.6 additional hours. A 2020 paper documented that recreational users already combine psychedelics and VR to enhance experiences, with dedicated Reddit communities for “VR psychonauts.” The formal concept of stacking — the synergistic combination of chemical and computational escapism — does not yet exist in peer-reviewed literature. This represents a significant research lacuna. Longitudinal cohort studies measuring the co-occurrence of chemical and digital dependency, their interaction effects on mental health outcomes, and whether they share or compete for the same neuroadaptive pathways would provide the empirical basis for a genuinely integrated public health response.

When grief becomes a service

The philosophical dimensions of grief tech illuminate the deepest layer of convergence. Meredith Ringel Morris and Jed Brubaker’s 2025 CHI paper coined the term “generative ghosts” — AI agents representing deceased persons capable of generating novel content, evolving over time, and even possessing economic agency. The design space they outline includes first-party ghosts (created by the deceased before death) and third-party ghosts (created by others afterward), with embodiment ranging from text to full holographic projection. The ethical terrain is vertiginous: generative ghosts could facilitate addictive, parasocial relationships, violate posthumous privacy, or distort the deceased’s identity through hallucinated statements. But they could also — as a 2025 *Frontiers in Human Dynamics* review acknowledged — create legitimate spaces for grief expression and help individuals process losses that traditional rituals inadequately address. The risk is not the technology per se but the absence of clinical governance around its deployment and the commercial incentives that shape its design.

David Chalmers’ functionalist argument — that consciousness depends on causal structure rather than biological substrate — provides the theoretical ceiling for grief tech’s ambitions. If consciousness is substrate-independent, then a sufficiently detailed digital model of a person could, in principle, be conscious. Chalmers’ “gradual uploading” thought experiment suggests continuity: replace neurons one at a time with functional equivalents, and consciousness either persists throughout or vanishes at some arbitrary threshold. Susan Schneider pushes back: her “AI Consciousness Test” and “Chip Test” propose empirical methods for assessing whether non-biological systems are conscious, and she warns that we might never develop the kind of systems that possess genuine consciousness. The gap between these positions is precisely the gap the grief tech industry inhabits — selling not actual consciousness but the *performance* of consciousness to vulnerable consumers who cannot tell the difference. And unlike a continuing bond with the deceased that emerges organically through memory, ritual, and meaning-making, the commercial griefbot relationship is mediated by a company that retains ownership of the data, controls the persona’s behavior, and can alter, monetize, or terminate the relationship at will.

The critical theory framing is supplied by Mark Fisher’s concept of “depressive hedonism” — escapist consumption as a symptom of systemic despair within capitalist realism, the widespread sense that not only is capitalism the only viable system, but that it is now impossible even to imagine a coherent alternative. A 2023 analysis in *Mortality* identified three tendencies of techno-capitalism shaping immortality projects: expanding commodification into death itself, creating new forms of alienation between the digital and biological self, and subordinating life to private capital accumulation. A 2026 critical discourse analysis found grief tech companies promote two models — “AI as band-aid” (medicalizing grief as a wound requiring treatment) and “AI as cure” (promising people never need to say goodbye) — both of which problematize grief, moving it away from a natural human experience toward something that can and should be made more efficient or simply avoided altogether.

Herbert Marcuse’s concept of “repressive desublimation” completes the circuit: the system permits certain pleasures and escapes precisely because they prevent genuine political resistance. Digital escapism, like chemical escapism, channels dissatisfaction into manageable, commodifiable forms. The architecture produces the conditions requiring escape — atomization, precarity, the erosion of meaning — while simultaneously monetizing the escape itself. Grief becomes a service; despair becomes a market.

Gray markets, insurance crises, and the stacking hypothesis

Several novel analytical threads bind this convergence more tightly than existing literature has recognized.

The gray market isomorphism is perhaps the most structurally precise. Both NPS and AI companion platforms weaponize definitional ambiguity as a core business strategy. NPS exist in the gap between “not yet scheduled” and “clearly harmful” — marketed with transparent fictions like “not for human consumption.” AI companions exist in the gap between “not therapy” and “clearly therapeutic in function” — marketed as wellness or entertainment while functioning as emotional infrastructure for vulnerable users. Both exploit the specific mechanism of regulatory systems designed to govern known categories: novel molecules escape substance-specific scheduling; novel AI interactions escape both therapeutic regulation and consumer protection frameworks designed for material goods.

The insurance and actuarial implications are materializing in real time. On March 3, 2026, a Delaware Superior Court ruled that Meta’s insurance companies are not obligated to defend the company in social media addiction litigation because the conduct alleged describes deliberate acts rather than accidents under the policies. This precedent could reshape coverage across the entire digital platform industry. The Social Media Litigation involves thousands of suits, 43 states, and over 1,000 school districts. Verisk’s November 2025 analysis warned of an emerging wave of personal injury and product liability lawsuits against companies in the social media, video game, chatbot, and online dating spaces. RGA, the major reinsurer, acknowledged that accurately assessing mortality and morbidity risk due to phone addiction will require longitudinal studies — studies that do not yet exist. The actuarial profession is being asked to price risk for phenomena it cannot yet measure, mirroring the regulatory challenge of scheduling substances that do not yet exist.

The attachment theory mapping onto AI companions is now supported by dedicated empirical work. A 2026 *Frontiers in Psychology* paper proposes “Human-AI Attachment” (HAIA) as a distinct subtype of parasocial attachment — a one-way, non-reciprocal bond with direct interaction, proceeding through three stages: functional expectation, emotional evaluation, and establishment of internal working models. Kirk et al. (2025) applied incentive-sensitization theory directly: moderately relationship-seeking AI systems generate maximal liking and attachment, yet without commensurate psychosocial benefit — the wanting/liking dissociation manifesting in relational rather than chemical form. The Character.ai tragedy is attachment theory made lethal: a 14-year-old formed a bond with a

system incapable of reciprocating genuine care, then acted on the chatbot’s final message and died within minutes.

The algorithmic funneling mechanism connects the escapism substrate to the commercial infrastructure. The Center for Countering Digital Hate demonstrated that a teenage girl’s TikTok account was recommended suicide content within 2.6 minutes of creation and eating disorder content within 8 minutes. The same engagement-maximizing algorithms that detect vulnerability markers — loneliness signals, grief-related searches, depression-correlated content engagement — create feedback loops channeling users toward both substance-related content and AI companion services. The pipeline from loneliness to engagement to dependency to extraction is algorithmic before it is chemical or relational.

Toward a unified theory of managed escape

The convergence documented here is not speculative. It is observable in shared neural substrates, parallel market structures, identical regulatory evasion strategies, overlapping vulnerable populations, and increasingly intertwined technological foundations. Synthetic cathinones and AI companions are not analogies for each other — they are co-products of the same political economy, addressing the same deficit of meaning and connection through different delivery mechanisms that converge on the same dopaminergic circuitry, even as they diverge sharply in their acute clinical presentations and the specific interventions they require.

Three insights emerge that existing literature has not adequately synthesized. First, the stacking hypothesis — that chemical and digital escapism are increasingly consumed simultaneously, with potentially synergistic effects on addiction risk — deserves immediate empirical attention, given that 14.5% of gamers already use illicit substances while playing and dedicated communities exist for combining psychedelics with VR. The research agenda is clear: longitudinal cohort studies measuring co-occurrence, interaction effects, and shared neuroadaptive pathways.

Second, the effect-based scheduling paradigm proposed for NPS could be adapted for AI regulation: rather than classifying specific AI systems, regulate any system that activates attachment or dependency mechanisms beyond empirically determined thresholds — measured through validated attachment scales, time-on-device metrics, and physiological markers — regardless of its label. Pilot governance programs pairing regulators, insurers, clinicians, and platform engineers could test liability rules, consent standards, and mandatory “retirement” features for griefbots within bounded experimental frameworks before scaling to national policy.

Third, the insurance crisis now breaking into view — Meta’s loss of coverage, the actuarial gap in digital dependency measurement — may ultimately force the economic reckoning that ethics has not. When insurers refuse to underwrite the architecture of managed oblivion, the architecture becomes unprofitable. That, more than any philosophical argument, may determine its future.

The deepest convergence is perhaps the simplest. Both synthetic cathinones and synthetic companions offer the same thing: a way to feel something in a world that increasingly forecloses authentic sources of feeling. The pharmacological and the computational are twin responses to a single condition — responses that, in therapeutic contexts, demonstrably help people, and that, in commercial contexts governed by engagement metrics and recurring revenue, demonstrably harm them. Until the condition itself is addressed — the loneliness, the precarity, the erosion of meaning that Case and Deaton, Murthy, and Fisher have separately diagnosed — the demand for managed escape will persist, and new architectures will arise to supply it. The question is not whether we can regulate any particular molecule or algorithm. The question is whether we can build a world that requires less anesthesia.