

The internet in 2026 feels increasingly hostile, shallow and exhausting not by accident but because anger and compulsion are now core features of how major platforms make money. Slop, rage bait and so-called brain rot are symptoms of an attention economy that runs on emotional provocation, infinite scroll and algorithmic systems tuned to maximise time-on-site rather than well-being.

Slop, rage bait and brain rot

“Slop” has become shorthand for the flood of low-effort, low-value content that saturates feeds, much of it now mass-produced with generative AI to capture clicks in bulk. Unlike traditional spam, this content is designed to be just good enough to keep you watching in the background while you eat, work or doomscroll, clogging the web with junk that exploits the creator economy's incentives.

Rage bait is a more weaponised cousin of clickbait: content deliberately crafted to make you angry or offended so that you comment, quote-tweet and share. Definitions now emphasise that it is built around emotional manipulation rather than information, with “hot takes,” deliberate misframings and antagonistic posts used as reliable engagement engines.

“Brain rot” started as a meme but has evolved into a catch-all description for the mental fatigue, attention fragmentation and numbness that follow hours of consuming this mix of slop and rage-baited feeds. Researchers and commentators link this feeling to hyper-stimulating, negative-valence content that keeps triggering the brain's fight-or-flight circuitry without resolution, leaving people depleted but still reaching for their phones.

The business model of misery

The core design choice is simple: platforms make more money when we spend more time scrolling, so every part of the interface is optimised to keep us there, not to make us informed or calm. Infinite scroll, autoplay, pull-to-refresh and endless recommendation carousels are all examples of design patterns that use intermittent, variable rewards to create slot-machine-like behaviour loops.

Algorithms trained on engagement learn quickly that strong negative emotions, rage, disgust, and moral outrage are more reliable drivers of clicks and comments than mild interest. Over time, this pushes platforms to surface content that is more polarising, more extreme or more emotionally loaded, even if individual engineers never explicitly instruct the system to “make users miserable.”

This logic now extends into AI-generated slop, which can be produced at near-zero marginal cost and A/B tested at scale to see which phrasing, thumbnail or provocation yields the highest engagement. As one analysis notes, this creates a feedback loop where our own reactions train the algorithm to show us more of precisely what keeps us agitated, confused or glued to the screen.

How platforms profit from outrage

- Advertising models reward impressions, watch time and interaction, not accuracy or mental health.
- Rage bait and polarising posts inflate “vanity metrics” such as likes and comments, which look attractive to advertisers even as they erode trust and satisfaction.
- Creators see that incendiary content performs better and adapt, producing more of what the algorithm seems to reward, further normalising antagonistic discourse.

Doomscrolling and mental health

The psychological toll of this design is no longer speculative; studies on doomscrolling link heavy exposure to negative news and social feeds with higher levels of depression, anxiety and trauma-like symptoms. During the Covid-19 pandemic, researchers found that frequent social media use to follow bad news was associated with increased depressive symptoms and post-traumatic stress, especially among vulnerable individuals.

Later work using a “Doomscrolling Scale” showed that compulsive consumption of negative online content predicted lower life satisfaction and mental well-being, mediated by psychological distress. Another study found that high levels of doomscrolling correlate with distraction from the present, reduced mindfulness and secondary traumatic stress, as if users absorb others’ suffering through the screen.

Commentators and clinicians now describe a pattern many users recognise: you feel irritated and wired after scrolling, find it harder to focus, and yet still reach for the phone in bed, a cycle that gradually erodes attention and emotional resilience. This “quiet damage” is framed not as a personal failing but as the predictable outcome of systems tuned to maximise negative-valence engagement, where each swipe is another data point reinforcing the algorithm’s sense of what will keep you hooked.

How design keeps us from logging off

Beyond content, the structure of feeds themselves nudges us into behaviours that feel bad but are hard to stop. Features like endless scroll remove natural stopping cues, while notifications and “unread” counters exploit our aversion to missing out or leaving tasks visibly incomplete.

Recommendation systems rarely say, “You have had enough for today”; instead they surface slightly more extreme or emotionally charged versions of what we just watched, pulling us deeper into niche outrage cycles or conspiratorial rabbit holes. Some researchers describe this as an amplification of risk perception: constant exposure to alarming content makes the world seem more dangerous and chaotic than it is, fuelling anxiety and pessimism.

Even ostensibly neutral design choices can have cumulative effects. The default of turning on autoplay, moving comments directly under videos or posts, and pushing “while you were away” recaps encourages us to treat being fully caught up as a kind of moral duty. In that environment, choosing to step away can feel like negligence, especially when news, politics and social identity are all braided into the same feed.

Can we redesign a less miserable web?

If today's internet feels miserable by design, that also means design can make it less so, but only if incentives change. Some experts and practitioners argue for new metrics of success, well-being, trust, and long-term retention, rather than pure engagement, coupled with ranking systems that demote rage bait and evident AI slop.

On the user side, practical strategies are emerging: unfollow accounts that consistently provoke outrage, avoid commenting in anger, limit late-night scrolling and deliberately follow creators who educate or inspire rather than inflame. At the policy and research level, there is growing interest in regulating “dark patterns,” requiring transparency around recommendation algorithms and funding independent audits of platforms' mental-health impacts.

For now, though, the dominant experience of the mainstream web is one where our worst impulses are continuously solicited, monetised and fed back to us as entertainment and

discourse. Slop, rage bait and brain rot are not glitches in that system; they are the inevitable by-products of an attention economy that has learned that a miserable user is still a highly engaged one, and, therefore, a valuable asset.