

The World Science Fiction Society presents

News From Outside Omelas: Technology

A Leaflet Curated by

The WSFS Clarion

And Published

Aug 13, 2025

Comprising the following articles:

- **Burning Chrome — William Gibson**
- **AI machines aren't 'hallucinating.' But their makers are — Naomi Klein**
- **I Will Fucking Piledrive You If You Mention AI Again — Nik 'Ludic' Suresh**
- **The Dark Forest and Generative AI — Maggie Appleton**
- **The Rot Economy — Ed Zitron**
- **There Were Always Enshittifiers — Cory Doctorow**

"When life closes a door, throw knowledge through a window."

Fenestrate Press prints items of interest, including both original submissions and reprints from elsewhere. More information, along with all our publications, is available at <https://fenestrate.net/>. Inquiries and submissions may be sent to info@fenestrate.net. Fenestrate Press publishes with the Ishill Stack, built on Hugo. Want to be able simultaneously to publish to the web, in PDF, in zine format, ...? Check out the free Ishill Stack at <https://www.autistici.org/ishill/>.

In celebration of “Why Don’t We Just Kill the Kid in the Omelas Hole” winning its well-deserved Hugo Award for Best Short Story, WSFS and the Seattle WorldCon are delighted to present this edition of the “News From Outside Omelas”—some curated pieces from a variety of authors to help you think through Kim’s piece (and Le Guin’s) as it applies to your life.

In today’s edition, we discuss AI and LLMs: your plastic pals who’re fun to be with. ~~In this year’s WorldCon, not only did we use LLMs to decide who was Internet acceptable to be on our panels, but as you can see from our schedule, half our panels are now talking about LLMs being inserted into everything from drawing manga to creating the pointless “dialogue” that fills up the time between space battles in the best books.~~

The previous paragraph was written not by an LLM, but by an LLM-inspired bad person; we apologize for its inclusion. ~~LLMs aren’t evil, we promise; we asked Microsoft, and it asked Copilot and ChatGPT (both of which it owns for some reason), and they told us that LLMs are creative partners who we should build things with.~~

The previous paragraph was written not by an LLM, but by a former Microsoft employee. ~~To avoid a repeat of the same bias, we asked Anthropic Claude, who answered that purple monkeys often have rabies, and to avoid it, we should use glue to hold the cheese on our pizzas.~~

The previous paragraph...

Don’t you get tired of the constant attempt to shift blame to the computers and away from the humans who use the computers with overgrown flowcharts to destroy the government, replace artists with drawings so trite they look like Thomas Kincade with a severe brain injury, and generate nonconsensual nude images of celebrities on Twitter while banning independent content creators from making art that discusses the possibility of periods (let alone sex) to “protect the children?”

Maybe the monsters weren’t LLMs all along, but the venture capitalists who told us this was a good idea—and the credulous pseudojournalists who repeat the “AI is coming for your jobs” lie regardless of a lack of objective reality along the way. After all, if **every single piece of media** told you to use LLMs for two years, wouldn’t you try them? (And then be surprised when not only do they suck, but they also lie while sucking....)

Anyway, here are some thoughts on AI, LLMs, and who the *real* bastards are in this conversation. It’s probably not the graphics cards. We start with William Gibson’s short story introducing “cyberspace” to the masses, before diving into the key questions: why does this keep getting pushed at everyone? Who’s benefiting? What are the intended consequences being pushed under the rug? What will it take to stop the madness?

Hold onto your butts: the ride only gets more bumpy from here....

“When I give food to the poor, they call me a saint. When I ask why the poor have no food, they call me a Communist.” — Fr. Hélder Pessoa Câmara

breaking every time the customer upgrades their OS, and how furious customers will be if (when) their un-backup-able floppies go corrupt.

Project SSAFE ended with surrender. Apple decided, in effect, that software vendors would need to have to work with their customers, not against them.

(One of the “copy protection” schemes Woz and Wigginton demolished with ease was the one used by VisiCalc, the spreadsheet program that tempts Marty Hensch into enrolling in a CPA. VisiCalc introduced a defect to one sector of its floppies at the factory and checked whether that sector was readable at runtime – it’s the scheme I modeled Fidelity Computer’s floppies on in **Picks and Shovels**.)

That’s where the fair fight of technological self-determination vs technological control cashed out. When users were allowed to investigate, modify, and improve their own computers, it was practically and commercially impossible to force them to optimize their affairs to create gains for corporate shareholders.

So the tech industry fought back. In 1998, an unholy alliance of giant entertainment and tech firms successfully lobbied for the passage of an American digital “anticircumvention” law, the first of its kind in world history. This law – Section 1201 of the Digital Millennium Copyright Act – made it a felony to provide someone with a tool to bypass an “access control,” on penalty of a five-year prison term and a \$500,000 fine for a first offense.

Suddenly, at the stroke of a pen, sharing a program that let you bypass a “copy protection” scheme became a criminal act, a jailable offense. Obviously, that didn’t halt the creation of these tools, but it did drive them underground. You couldn’t go down to your mall computer store and get a program that would let you back up the expensive software you’d bought. BBSes didn’t carry shareware versions of these programs in their file sections. As far as the law was concerned, companies were now allowed to take *any* measure to restrict your use of your property – your expensive computer – but you were legally prohibited from taking any steps to push back.

In other words, they didn’t have to make “copy protection” work – they just had to make telling someone how it was broken into a felony.

Today, we live with the legacy of that decision. It’s a felony to bypass the controls that prevent your printer from accepting alternatives to \$10,000/gallon inkjet ink. It’s a felony to bypass the digital locks in your car so that an independent mechanic can diagnose and repair it – let alone disabling the pervasive, nonconsensual surveillance that comes standard in every car.

If you’re an author whose Kindle or Audible titles are sold with DRM, you can’t authorize your readers to convert those files to work on independent apps, even though you hold the copyright to those books. You don’t hold the copyright to Kindle or Audible’s DRM, after all, which means that it’s a felony for you – the author of a book – to help a reader get it out of Amazon’s walled garden. That means that once you start selling into Amazon’s walled garden, you can’t afford to stop, no matter how badly the terms of the deal degrade. Your readers can’t leave Amazon, so you can’t either, so Amazon can take ever-larger shares of your income and demand ever more unfair terms and you can’t do anything about it.

The enshittifiers were with us from the start, but it used to be a fair fight. IP laws like the DMCA take away rights from both creators and our audiences, and transfer power to middlemen like Amazon, who do not create, finance, edit, or publish our works, but nevertheless have more of a say over them than we do.

Picks and Shovels is a tale from the enshittification’s dawn, a moment when the future was up for grabs. But it’s not meant as a melancholy tale of loss. The future is *still* up for grabs. The policies that made us digital serfs can be unmade. They should be.

They must be.

Burning Chrome

William Gibson

Jul 1, 1982

Originally published: https://archive.org/details/omniarchive/OMNI_1982_07/page/137/mode/2up

It was hot, the night we burned Chrome. Out in the malls and plazas, moths were batting themselves to death against the neon, but in Bobby’s loft the only light came from a monitor screen and the green and red LEDs on the face of the matrix simulator. I knew every chip in Bobby’s simulator by heart; it looked like your workaday Ono-Sendai VII, the ‘Cyberspace Seven,’ but I’d rebuilt it so many times that you’d have had a hard time finding a square millimeter of factory circuitry in all that silicon.

We waited side by side in front of the simulator console, watching the time display in the screen’s lower left corner.

‘Go for it,’ I said, when it was time, but Bobby was already there, leaning forward to drive the Russian program into its slot with the heel of his hand. He did it with the tight grace of a kid slamming change into an arcade game, sure of winning and ready to pull down a string of free games.

A silver tide of phosphenes boiled across my field of vision as the matrix began to unfold in my head, a 3-D chessboard, infinite and perfectly transparent. The Russian program seemed to lurch as we entered the grid. If anyone else had been jacked into that part of the matrix, he might have seen a surf of flickering shadow roll out of the little yellow pyramid that represented our computer. The program was a mimetic weapon, designed to absorb local color and present itself as a crash-priority override in whatever context it encountered.

‘Congratulations,’ I heard Bobby say. ‘We just became an Eastern Seaboard Fission Authority inspection probe...’ That meant we were clearing fiberoptic lines with the cybernetic equivalent of a fire siren, but in the simulation matrix we seemed to rush straight for Chrome’s data base. I couldn’t see it yet, but I already knew those walls were waiting. Walls of shadow, walls of ice.

Chrome: her pretty childface smooth as steel, with eyes that would have been at home on the bottom of some deep Atlantic trench, cold gray eyes that lived under terrible pressure. They said she cooked her own cancers for people who crossed her, rococo custom variations that took years to kill you. They said a lot of things about Chrome, none of them at all reassuring.

So I blotted her out with a picture of Rikki. Rikki kneeling in a shaft of dusty sunlight that slanted into the loft though a grid of steel and glass: her faded camouflage fatigues, her translucent rose sandals, the good line of her bare back as she rummaged through a nylon gear bag. She looks up, and a half-blond curl falls to tickle her nose. Smiling, buttoning an old shirt of Bobby’s, frayed khaki cotton drawn across her breasts.

She smiles.

‘Son of a bitch,’ said Bobby, ‘we just told Chrome we’re an IRS audit and three Supreme Court subpoenas...Hang on to your ass, Jack...’

So long, Rikki. Maybe now I see you never.

And dark, so dark, in the halls of Chrome’s ice.

Bobby was a cowboy, and ice was the nature of his game, ice from ICE, Intrusion Countermeasures Electronics. The matrix is an abstract representation of the relationships between data systems. Legitimate programmers jack into their employers’ sector of the matrix and find themselves surrounded by bright geometries representing the corporate data.

Towers and fields of it ranged in the colorless nonspace of the simulation matrix, the electronic consensus-hallucination that facilitates the handling and exchange of massive quantities of data. Legitimate programmers never see the walls of ice they work behind, the walls of shadow that screen their operations from others, from industrial-espionage artists and hustlers like Bobby Quine.

Bobby was a cowboy. Bobby was a cracksman, a burglar, casing mankind's extended electronic nervous system, rustling data and credit in the crowded matrix, monochrome nonspace where the only stars are dense concentrations of information, and high above it all burn corporate galaxies and the cold spiral arms of military systems.

Bobby was another one of those young-old faces you see drinking in the Gentleman Loser, the chic bar for computer cowboys, rustlers, cybernetic second-story men. We were partners.

Bobby Quine and Automatic Jack. Bobby's the thin, pale dude with the dark glasses, and Jack's the mean-looking guy with the myoelectric arm. Bobby's software and Jack's hard; Bobby punches console and Jack runs down all the little things that can give you an edge. Or, anyway, that's what the scene watchers in the Gentleman Loser would've told you, before Bobby decided to burn Chrome. But they also might've told you that Bobby was losing his edge, slowing down. He was twenty-eight, Bobby, and that's old for a console cowboy.

Both of us were good at what we did, but somehow that one big score just wouldn't come down for us. I knew where to go for the right gear, and Bobby had all his licks down pat. He'd sit back with a white terry sweatband across his forehead and whip moves on those keyboards faster than you could follow, punching his way through some of the fanciest ice in the business, but that was when something happened that managed to get him totally wired, and that didn't happen often. Not highly motivated, Bobby, and I was the kind of guy who's happy to have the rent covered and a clean shirt to wear.

But Bobby had this thing for girls, like they were his private tarot or something, the way he'd get himself moving. We never talked about it, but when it started to look like he was losing his touch that summer, he started to spend more time in the Gentleman Loser. He'd sit at a table by the open doors and watch the crowd slide by, nights when the bugs were at the neon and the air smelled of perfume and fast food. You could see his sunglasses scanning those faces as they passed, and he must have decided that Rikki's was the one he was waiting for, the wild card and the luck changer. The new one.

I went to New York to check out the market, to see what was available in hot software.

The Finn's place has a defective hologram in the window, METRO HOLOGRAFIX, over a display of dead flies wearing fur coats of gray dust. The scrap's waist-high, inside, drifts of it rising to meet walls that are barely visible behind nameless junk, behind sagging pressboard shelves stacked with old skin magazines and yellow-spined years of National Geographic.

'You need a gun,' said the Finn. He looks like a recombo DNA project aimed at tailoring people for highspeed burrowing. 'You're in luck. I got the new Smith and Wesson, the four-oh-eight Tactical. Got this zenon projector slung under the barrel, see, batteries in the grip, throw you a twelve-inch high-noon circle in the pitch dark at fifty yards. The light source is so narrow, it's almost impossible to spot. It's just like voodoo in a nightfight.'

I let my arm clunk down on the table and started the fingers drumming; the servos in the hand began whining like overworked mosquitoes. I knew that the Finn really hated the sound.

'You looking to pawn that?' he prodded the Duralumin wrist joint with the chewed shaft of a felt-tip pen. 'Maybe get yourself something a little quieter?'

I kept it up. 'I don't need any guns, Finn.'

'Okay,' he said, 'okay,' and I quit drumming. 'I only got this one item, and I don't even know what it is.' He looked unhappy. 'I got it off these bridge-and-tunnel kids from Jersey last week.'

'So when'd you ever buy anything you didn't know what it was, Finn?'

The PC was born through the efforts of hobbyists, organized in loose groups like Palo Alto's legendary Homebrew Computing Club (**Picks and Shovels** includes one of these clubs, the Newberry Street Irregulars). Early hackers brought their hardware and their code to these clubs so that their peers could see what they'd built, and, crucially, adapt and improve it. Forget all those stories about Jobs and Woz as lone inventors – sure, Woz was a generational hardware genius, but the Apple I owes its existence to the community of practice Woz was part of. Without the inspiration, tips, and ideas he got from his peers, there wouldn't have been an Apple computer.

In 1976, the Homebrew Computer Club was shaken by an internal schism. A young Bill Gates, furious at his peers for copying the DOS interpreter he had helped make, widely published his "Open Letter to Hobbyists," scolding them for freely sharing his code, as they did with all the other code that they found useful to their shared project. Unlike his peers, Gates refused to publish the underlying code for his program, preventing other programmers from rekeying it, debugging it, or adapting it.

The DOS that Gates took credit for in his letter was mostly written by his collaborators, notably his future Microsoft cofounder Paul Allen (whom Gates would later cheat out of billions of dollars' worth of stock). Gates had defied every norm of the nascent home computer industry, a culture built on sharing and collaboration, while taking credit for others work, and his "Open Letter" demanded that everyone *else* reorganize the way they invented and improved on computers in such a way as to make him money, while holding back the progress they'd all made together.

Gates won, obviously, and went on to be the first person in the 21st century to be convicted in a federal court for creating a tech monopoly. The Gates method for making technology pits users of technology against its makers, and gave rise to a long Cold War between technology creators and the rest of us.

It used to be a much fairer fight. It used to be that if a company figured out how to block copying its floppies, another company – or even just an individual tinkerer – could figure out how to break that "copy protection." There were plenty of legitimate reasons to want to do this: Maybe you owned more than one computer, or maybe you were just worried that your floppy disk would degrade to the point of unreadability. That's a very reasonable fear: Floppies were notoriously unreliable, and every smart computer user learned to make frequent backups against the day that your computer presented you with the dread DISK ERROR message.

In those early days, it was an arms race between companies that wanted to control how their customers used their own computers, and the technological guerrillas who produced the countermeasures that restored command over your computer to you, its owner. It's true that the companies making the "copy protection" (in scare quotes because the way you protect your data is by making copies of it) typically had far more resources than the toolsmiths who were defending technology users.

But that high tech resistance had an insurmountable advantage: the "copy protection" companies had to develop flawless schemes to prevent computers from copying floppies. The resistance merely had to find one mistake in the "copy protection" scheme and exploit it. This is a central tenet of asymmetric warfare, in which the guerrilla "has only to be lucky once," while their target "has to be lucky always."

In 2017, a Redditor called Vadermeer discovered a binder full of ancient internal Apple memos at a Seattle Goodwill thrift store. The memos detailed the workings of a 1979 Apple initiative called SSAFE ("Software Security from Apple's Friends and Enemies"), which sought to alter Apple's operating systems and hardware to allow vendors to absolutely stymie any attempts at unauthorized copies of floppy disks.

The 100-plus page trove of typed, handwritten, and printed documents are a fascinating window into the era. They include a hilarious back-and-forth between Apple founder Steve Wozniak, legendary for his technological prowess, and Randy Wigginton, Apple's sixth employee, responsible for much of Apple's groundbreaking technology.

Woz and Wigginton document their adventures investigating the state of the art for "copy protection," discovering that they can trivially bypass any system on the market and concluding that the companies peddling this stuff are selling defective goods, and probably know it. Elsewhere in the trove, Apple's product managers recoil in horror from various proposals to make "uncopyable" Apple floppies, pointing out that adding a hardware dongle to one of their customers' four expansion slots is a nonstarter, and that's before the high cost of that dongle is factored in. They go on to discuss the problems with this dongle

There Were Always Enshittifiers

Cory Doctorow

Mar 3, 2025

Originally published: <https://locusmag.com/2025/03/commentary-cory-doctorow-there-were-always-enshittifiers/>

In my new novel **Picks and Shovels**, we learn the origins of Martin Hench, my bestselling, two-fisted, scambusting forensic accountant who debuted in 2023's **Red Team Blues** and whose adventures continued in 2024's **The Bezzle**.

Marty's origin story starts at MIT in 1982, where he joins the proud lineage of computer science students who flunk out of their degrees because they're too busy hacking code to do their classwork. Marty ends up in a CPA program (he picks it because the community college has a lab full of Apple II+ computers running the pioneering spreadsheet program VisiCalc). After a failed postgraduate startup, Marty and his brilliant hacker roommate Art decamp to San Francisco, seeking their fortunes during the heroic era of the PC revolution.

The 1980s weren't merely the heroic era of the PC – that was also the *weird* era of the PC. No one knew who was supposed to make a PC, who was supposed to use them, or what they were for. We hadn't settled questions like “what does a computer look like?” They were fertile years, so fertile that Marty basically trips and lands on a job, working for a weird PC company called Fidelity Computing.

Fidelity sounds like a joke: a computer company run by a Mormon bishop, a Catholic priest, and an orthodox rabbi. But the joke's on their customers, because Fidelity is a pyramid scheme, a Ponzi that recruits members of faith communities to sell overpriced, proprietary hardware to one another, with the Reverend Sirs at the top of the company raking off handsome profits.

How handsome? Well, the Fidelity Computing printer is a rebadged Okidata ML-80, a workhorse, fan-fold, dot-matrix printer that will be familiar to anyone who was using computers back then. But Fidelity doesn't just slap its name on these old Okidatas – they replace the sprockets for the tractor-feed paper with sprockets that are more widely spaced. That way they can force you to use their special, high-priced printer paper. Of course, spacing out the sprockets more widely puts a lot of strain on the printer's motors, causing jams (which mean you have to buy more of their expensive paper) and frequent repair (which can only be performed at their depots).

Or take the Fidelity floppy drive, which is accessed using Fidelity's in-house operating system, WISE DOS. Before WISE DOS reads or writes from a floppy in the drive, it instructs it to seek out a specific sector on the floppy and attempt to read it. That same sector is deliberately knocked out at the factory, damaged so that it cannot be read. If the floppy drive detects a readable sector at the designated address, it knows that you've had the temerity to use a third-party floppy disk in your drive, and it goes on strike, refusing to read any data from the disk.

It will not surprise you to learn that Fidelity's house-brand floppy disks cost three times as much as generic ones from the store.

I had a lot of fun thinking through all the ways that Fidelity Computing could be a proto-enshittifier, using the comparatively primitive technology of the 1980s to enact precursors and forerunners to today's master enshittifiers, like the inkjet companies that have used software locks to raise the price of printer ink to \$10,000/gallon, which is why you print out your shopping lists and boarding cards with colored water that costs more, ounce for ounce, than the semen of a Kentucky Derby winner.

But I didn't make all of it up. From the very start, the PC revolution was a brawl between people who wanted to use computers to set people free, and those who wanted to use computers as traps.

‘Wise ass.’ And he passed me a transparent mailer with something in it that looked like an audio cassette through the bubble padding. ‘They had a passport,’ he said. ‘They had credit cards and a watch. And that.’

‘They had the contents of somebody’s pockets, you mean.’

He nodded. ‘The passport was Belgian. It was also bogus, looked to me, so I put it in the furnace. Put the cards in with it. The watch was okay, a Porsche, nice watch.’

It was obviously some kind of plug-in military program. Out of the mailer, it looked like the magazine of a small assault rifle, coated with nonreflective black plastic. The edges and corners showed bright metal; it had been knocking around for a while.

‘I’ll give you a bargain on it, Jack. For old times’ sake.’

I had to smile at that. Getting a bargain from the Finn was like God repealing the law of gravity when you have to carry a heavy suitcase down ten blocks of airport corridor.

‘Looks Russian to me,’ I said. ‘Probably the emergency sewage controls for some Leningrad suburb. Just what I need.’

‘You know,’ said the Finn. ‘I got a pair of shoes older than you are. Sometimes I think you got about as much class as those yahoos from Jersey. What do you want me to tell you, it’s the keys to the Kremlin? You figure out what the goddamn thing is. Me, I just sell the stuff.’

I bought it.

Bodiless, we swerve into Chrome’s castle of ice. And we’re fast, fast. It feels like we’re surfing the crest of the invading program, hanging ten above the seething glitch systems as they mutate. We’re sentient patches of oil swept along down corridors of shadow.

Somewhere we have bodies, very far away, in a crowded loft roofed with steel and glass. Somewhere we have microseconds, maybe time left to pull out.

We’ve crashed her gates disguised as an audit and three subpoenas, but her defenses are specially geared to cope with that kind of official intrusion. Her most sophisticated ice is structured to fend off warrants, writs, subpoenas. When we breached the first gate, the bulk of her data vanished behind core-command ice, these walls we see as leagues of corridor, mazes of shadow. Five separate landlines spurted May Day signals to law firms, but the virus had already taken over the parameter ice. The glitch systems gobble the distress calls as our mimetic subprograms scan anything that hasn’t been blanked by core command.

The Russian program lifts a Tokyo number from the unscreened data, choosing it for frequency of calls, average length of calls, the speed with which Chrome returned those calls.

‘Okay,’ says Bobby, ‘we’re an incoming scrambler call from a pal of hers in Japan. That should help.’

Ride ’em cowboy.

Bobby read his future in women; his girls were omens, changes in the weather, and he’d sit all night in the Gentleman Loser, waiting for the season to lay a new face down in front of him like a card.

I was working late in the loft one night, shaving down a chip, my arm off and the little waldo jacked straight into the stump.

Bobby came in with a girl I hadn’t seen before, and usually I feel a little funny if a stranger sees me working that way, with those leads clipped to the hard carbon studs that stick out of my stump. She came right over and looked at the magnified image on the screen, then saw the waldo moving under its vacuum-sealed dust cover. She didn’t say anything, just watched. Right away I had a good feeling about her; it’s like that sometimes.

‘Automatic Jack, Rikki. My associate.’

“Good companies” should not have to repeatedly lay people off. They should not be mass-hiring for fear that the demand they are capturing is temporary, and those new employees will soon find themselves at the receiving end of a pink slip.

The lens through which we evaluate businesses is cracked, and until we fix it, we will continue to experience these punishing cycles of binging and purging on human capital.

This is the problem at the center of almost everything I’ve written. Why are bosses mad they can’t bring people back to the office? Because their alignment of business success isn’t really tied to profit or “success,” but rather the sense that they are “big” and “successful,” which requires a bustling workplace and “ideas.”

Why did billions of dollars get pumped into crypto’s countless non-companies? Because “success” as defined by capital has been reframed to mean “number go up.” As a notion, it is divorced from any long-term thinking, fiscal probity, or even what you and I would call “morality.”

Why did these companies never seem to get blamed for hiring and then quickly firing tens of thousands of people? Because at the heart of the business media and the markets, workers were necessary casualties of the eternal struggle for growth. Layoffs are inevitably reported as a large number (“10,000 employees at Microsoft”), which makes it all too easy to remove the human element. When confronted with numbers of this scale, it’s easy to ignore the individual human agony that comes with losing a job. The uncertainty and shame that follows a firing.

The truth is that nothing lasts forever. Companies can (and should) die — or, at the very least, understand that there is an inevitable limit to growth, and eventually they must reconcile with being a stable, albeit plateaued, business.

A product may be profitable for a while, but there is a line at which profitability comes at the cost of functionality, and your company may simply not be able to *grow* more. A business that cannot generate profit is not a good business, and a business that can never generate a profit deserves to die.

And the net result of all of this is that it kills innovation. If capital is not invested in providing a good service via a profitable business, it will never sustain things that are societally useful. Companies are not incentivized to provide better services or improve lives outside of ways in which they can drain more blood from consumers. And the street doesn’t care either — just look at Facebook and Instagram, two products that have grown endlessly profitable and utterly useless in the process.

If capital wishes to call labor entitled, capital must acknowledge that it is the most entitled creature in society, craving eternal growth at the cost of the true value of any given service or entity.

Footnotes:

1. <https://ez.substack.com/p/techs-elite-hates-labor>
2. <https://www.ft.com/content/9daf27f6-dde7-40d8-b01d-33b70844aa69>
3. <https://www.inc.com/thomas-kouloupoulos/why-this-13-year-google-employee-says-google-can-no-longer-innovate.html>
4. <https://ez.substack.com/p/mark-zuckerberg-is-going-to-kill>
5. <https://www.cnbc.com/2023/02/01/metasp-year-of-efficiency-everything-wall-street-needed-to-hear.html>
6. <https://www.cnbc.com/2023/02/01/meta-lost-13point7-billion-on-reality-labs-in-2022-after-metaverse-pivot.html>
7. <https://www.jdsupra.com/legalnews/meta-fined-410m-for-gdpr-violations-1135397/>
8. <https://gto5mac.com/2022/04/14/number-of-users-opting-in-to-app-tracking-on-ios-grows-significantly-since-last-year/>
9. <https://www.theatlantic.com/ideas/archive/2022/06/google-search-algorithm-internet/661325/>
10. <https://www.nytimes.com/2023/02/08/technology/microsoft-bing-openai-artificial-intelligence.html>
11. <https://www.cnbc.com/2023/02/08/alphabet-shares-slip-following-googles-ai-event.html>
12. <https://investor.uber.com/news-events/news/press-release-details/2023/Uber-Announces-Results-for-Fourth-Quarter-and-Full-Year-2022/default.aspx>
13. <https://fortune.com/2022/10/11/biden-gig-worker-policy-crushing-news-uber-lyft-analysts-say/>
14. <https://www.cnbc.com/2023/02/08/uber-earnings-q4-2022.html>
15. <https://www.reuters.com/technology/uber-posts-first-small-adjusted-profit-ridership-rises-delivery-gets-more-2021-11-04/>

before. He might as well have had NEXT printed across his sunglasses in green Day-Glo capitals, ready to flash out at the first interesting face that flowed past the tables in the Gentleman Loser.

I knew what he did to them. He turned them into emblems, sigils on the map of his hustler’s life, navigation beacons he could follow through a sea of bars and neon. What else did he have to steer by? He didn’t love money, in and of itself, not enough to follow its lights. He wouldn’t work for power over other people; he hated the responsibility it brings. He had some basic pride in his skill, but that was never enough to keep him pushing.

So he made do with women.

When Rikki showed up, he needed one in the worst way. He was fading fast, and smart money was already whispering that the edge was off his game. He needed that one big score, and soon, because he didn’t know any other kind of life, and all his clocks were set for hustler’s time, calibrated in risk and adrenaline and that supernal dawn calm that comes when every move’s proved right and a sweet lump of someone else’s credit clicks into your own account.

It was time for him to make his bundle and get out; so Rikki got set up higher and farther away than any of the others ever had, even though — and I felt like screaming it at him — she was right there, alive, totally real, human, hungry, resilient, bored, beautiful, excited, all the things she was...

Then he went out one afternoon, about a week before I made the trip to New York to see the Finn. Went out and left us there in the loft, waiting for a thunderstorm. Half the skylight was shadowed by a dome they’d never finished, and the other half showed sky, black and blue with clouds. I was standing by the bench, looking up at that sky, stupid with the hot afternoon, the humidity, and she touched me, touched my shoulder, the half-inch border of taut pink scar that the arm doesn’t cover. Anybody else ever touched me there, they went on to the shoulder, the neck...

But she didn’t do that. Her nails were lacquered black, not pointed, but tapered oblongs, the lacquer only a shade darker than the carbon-fiber laminate that sheathes my arm. And her hand went down the arm, black nails tracing a weld in the laminate, down to the black anodized elbow joint, out to the wrist, her hand soft-knuckled as a child’s, fingers spreading to lock over mine, her palm against the perforated Duralumin.

Her other palm came up to brush across the feedback pads, and it rained all afternoon, raindrops drumming on the steel and soot-stained glass above Bobby’s bed.

Ice walls flick away like supersonic butterflies made of shade. Beyond them, the matrix’s illusion of infinite space. It’s like watching a tape of a prefab building going up; only the tape’s reversed and run at high speed, and these walls are torn wings.

Trying to remind myself that this place and the gulfs beyond are only representations, that we aren’t ‘in’ Chrome’s computer, but interfaced with it, while the matrix simulator in Bobby’s loft generates this illusion...The core data begin to emerge, exposed, vulnerable...This is the far side of ice, the view of the matrix I’ve never seen before, the view that fifteen million legitimate console operators see daily and take for granted.

The core data tower around us like vertical freight trains, color-coded for access. Bright primaries, impossibly bright in that transparent void, linked by countless horizontals in nursery blues and pinks.

But ice still shadows something at the center of it all: the heart of all Chrome’s expensive darkness, the very heart...

It was late afternoon when I got back from my shopping expedition to New York. Not much sun through the skylight, but an ice pattern glowed on Bobby’s monitor screen, a 2-D graphic representation of someone’s computer defenses, lines of neon woven like an Art Deco prayer rug. I turned the console off, and the screen went completely dark.

Rikki’s things were spread across my workbench, nylon bags spilling clothes and makeup, a pair of bright red cowboy boots, audio cassettes, glossy Japanese magazines about simstim stars. I stacked it all under the bench and then took my arm off, forgetting that the program I’d bought from the Finn was in the right-

hand pocket of my jacket, so that I had to fumble it out left-handed and then get it into the padded jaws of the jeweler's vise.

The waldo looks like an old audio turntable, the kind that played disc records, with the vise set up under a transparent dust cover. The arm itself is just over a centimeter long, swinging out on what would've been the tone arm on one of those turntables. But I don't look at that when I've clipped the leads to my stump; I look at the scope, because that's my arm there in black and white, magnification 40x.

I ran a tool check and picked up the lazer. It felt a little heavy; so I scaled my weight-sensor input down to a quarter-kilo per gram and got to work. At 40x the side of the program looked like a trailer truck.

It took eight hours to crack: three hours with the waldo and the laser and four dozen taps, two hours on the phone to a contact in Colorado, and three hours to run down a lexicon disc that could translate eight-year-old technical Russian.

Then Cyrillic alphanumerics started reeling down the monitor, twisting themselves into English halfway down. There were a lot of gaps, where the lexicon ran up against specialized military acronyms in the readout I'd bought from my man in Colorado, but it did give me some idea of what I'd bought from the Finn.

I felt like a punk who'd gone out to buy a switchblade and come home with a small neutron bomb.

Screwed again, I thought. What good's a neutron bomb in a streetfight? The thing under the dust cover was right out of my league. I didn't even know where to unload it, where to look for a buyer. Someone had, but he was dead, someone with a Porsche watch and a fake Belgian passport, but I'd never tried to move in those circles. The Finn's muggers from the 'burbs had knocked over someone who had some highly arcane connections.

The program in the jeweler's vise was a Russian military icebreaker, a killer-virus program. It was dawn when Bobby came in alone. I'd fallen asleep with a bag of takeout sandwiches in my lap.

'You want to eat?' I asked him, not really awake, holding out my sandwiches. I'd been dreaming of the program, of its waves of hungry glitch systems and mimetic subprograms; in the dream it was an animal of some kind, shapeless and flowing.

He brushed the bag aside on his way to the console, punched a function key. The screen lit with the intricate pattern I'd seen there that afternoon. I rubbed sleep from my eyes with my left hand, one thing I can't do with my right. I'd fallen asleep trying to decide whether to tell him about the program. Maybe I should try to sell it alone, keep the money, go somewhere new, ask Rikki to go with me.

'Whose is it?' I asked.

He stood there in a black cotton jump suit, an old leather jacket thrown over his shoulder like a cape. He hadn't shaved for a few days, and his face looked thinner than usual.

'It's Chrome's,' he said.

My arm convulsed, started clicking, fear translated to the myoelectrics through the carbon studs. I spilled the sandwiches; limp sprouts, and bright yellow dairy-produce slices on the unswept wooden floor.

'You're stone crazy,' I said.

'No,' he said, 'you think she rumbled it? No way. We'd be dead already. I locked on to her through a triple-blind rental system in Mombasa and an Algerian comsat. She knew somebody was having a look-see, but she couldn't trace it.'

If Chrome had traced the pass Bobby had made at her ice, we were good as dead. But he was probably right, or she'd have had me blown away on my way back from New York. 'Why her, Bobby? Just give me one reason...'

amounts of capital to address problems related to these things, because he has the resemblance of the kind of messianic white guy that is able to take a product and sell it, [even if he is quite literally the guy who failed to do this before]]([https://www.theguardian.com/business/2019/dec/20/why-wework-went-wrong#:~:text=The%20failed%20IPO%20and%20the,optimistic%20\(it%20accounted%20anyone%20who\)%](https://www.theguardian.com/business/2019/dec/20/why-wework-went-wrong#:~:text=The%20failed%20IPO%20and%20the,optimistic%20(it%20accounted%20anyone%20who)%) [37]).

Neumann turned a (nominally) \$47bn company into a \$2.9bn company. In a sane and just world, he wouldn't see a dollar of funding for the rest of his life.

There are tons of other examples of colossally stupid assholes and stupid ideas getting money. As I wrote about on Monday[38], the largest investment rounds of the last few years have gone to companies that got obscene valuations based on nothing other than a vague sense of them "looking like a winner." There is no reason a weight loss app should need \$540 million to operate[39] - that is not a sustainable enterprise considering the entire weight loss industry is worth about \$3.8 billion[40]. Clubhouse was never worth the billions of dollars pumped into it, considering the[41] *entire radio industry only makes about \$12 billion a year combined*[42]. While capital is required to get a company off the ground, the only way to justify these massive surges of capital is that venture capitalists are putting companies on life support in the hopes that they can flog them for a profit.

And this corrosive capital system gets continually rewarded. Companies like Uber are taken public, making massive windfalls for venture capitalists[43] without ever having to run a profitable business[44]. Venture capitalists crammed \$41 billion into crypto in the space of 18 months[45], despite there being no real use cases for crypto. Metaverse companies raised \$120 billion in 2022[46] for a concept that has yet to really *exist, and perhaps never will*. Yet these concepts get vast amounts of money because venture capitalists are incentivized to pump cash into "good companies to invest in" over "good companies."

As my friend Kasey[47] put it in a recent conversation, growth is a fire. If you build a nice, sustainable fire, it'll keep you warm, cook food and sustain life. And if the only thing you care about is how big your fire is, then it'll set fire to everything around it, and the more you throw into it, the more it'll burn. Eventually, you'll have nothing left, but if you desperately desire that fire, you will constantly have to find new things to burn at any cost.

And we, societally, have turned our markets and businesses - private and public - over to arsonists. We have created conditions where we celebrate people for making "big" companies but not "good" companies.

Venture capital and the public markets don't actually reward or respect "good" businesses or "good" CEOs - they reward people that can steer the kind of growth that raises the value of an asset. Elon Musk's success with Tesla didn't come from the inarguable point that he ended the monopoly of the internal combustion engine - it came from his canny manipulation of the symbolic value of a stock through lies and half-truths[48], meaning that there was always a perpetual reason that Tesla was a "growth" company and a "good stock to buy." Sundar Pichai isn't paid \$280 million a year because he's a "good CEO." [49] After all, Google has all but destroyed its search product. He's paid because he finds ways to increase the overall growth of the company (even while their cloud division still loses money[50]), and thus the stock goes up.

The consequences are that these companies will continue to invest in things that grow the overall revenue of the company over all else. They will mass-hire and mass-fire, because there are no consequences when the markets don't really care as long as the company itself stays valuable. Venture capitalists certainly don't mind - after all, it's "less burn" to "get you through" tough climates that were arguably created by the poor hiring decisions of a company that was never incentivized to hire sustainably or operate profitably.

Until we see a seismic shift in how major investors treat the companies they invest in, this cycle will continue. I guarantee that we will see each and every one of the companies doing mass layoffs do mass hirings in the next few years, and then do another mass layoff not long after, ***because they are simply treating human capital as assets to be manipulated to increase the value of a stock***. They are not structured to evaluate whether the business is "sustainable," because their only interest is seeing their current profits grow by multiples that please Wall Street.

Everything Ventured, Nothing Gained

Before these companies reach the public markets, they are fueled by an even more violently reckless form of funding - venture capital. Venture capitalists are regularly incentivized to create businesses that *look* valuable but aren't necessarily of value. When I wrote about the Liches of Silicon Valley last year[25], I remarked upon how many valley companies experience volatile, erosive cycles of growth with the goal of being acquired or going public, burning as much venture capital as it takes to find an outcome:

They repeat a very specific cycle - company is the next big thing[26], company is now worth over a billion dollars[27], company is experiencing “unheard of growth” (with no question as to whether they are sustainable or profitable)[28], company is now challenging ‘the big dogs’ of industry[29], a little M&A[30], an absolutely insane valuation[31], and then a sudden realization that actually, perhaps this wasn’t a good business at all[32]? I am hammering on TechCrunch links here because I am being lazy - they are far from the only outlet to assume that a company like Brex would not simply run itself into the ground through virtue of existing - but the path is always the same - growth, growth, growth, legitimization, growth, growth, acquisition, and then an eventual reckoning with real life.

Venture pumps millions or billions of dollars into ideas that might sell a product or a service, but ultimately resemble things that can be sold to other companies or put on the public market for a profit higher than what was paid on a per-share basis. I once suggested that Silicon Valley conflated “making great ideas work” with “making ideas I like work,” but on consideration, many of these companies aren’t even things venture capitalists *like* - they are things that resemble things that they can sell. Do I genuinely believe that everyone who invested into the Web3 grift was a strident believer in the brave new decentralized economy? Hell no. They just went where the winds blew — or where they *seemed* to be blowing.

Andreessen Horowitz was the lead participant in arguably the biggest con in venture capital, pumping billions into Web3 companies that *didn’t have any real product*, but stapled together enough buzzwords and websites to resemble actual entities[33]. A16Z found a way to vastly accelerate the idea-to-business-to-profit cycle of venture. Despite claiming it was “Time To Build[34]” in 2020, Andreessen Horowitz realized that there wasn’t ever really much of a need to build at all - you could create things that semiotically aligned with what “valuable” looked like and profit off of that. While the public markets may (at least, before the rise of the SPAC) have required some sort of business - even if said business wasn’t graded on being a “good” one - the cryptocurrency markets allowed the vaguest of ideas to get even vaguer valuations.

This same insipid thought process applies to the rest of their portfolio too. Adam Neumann, a guy who is most famous for running WeWork into the ground[35], got a second at-bat with his new startup “Flow,” a company that Neumann is still not able to fully describe, but that may involve you renting to own an apartment that Flow owns somewhere at some point. Just like Silicon Valley can’t help itself from reinventing the bus, Neuman is seemingly attempting to reinvent the rental market — a diseased, exploitative industry in its own right — in his own image. He’s replacing one cancer with another, only even more aggressive and metastatic.

Neumann was, is, and will always be full of shit. Appropriately, in a video A16Z released yesterday[36], Neumann used the following analogy to describe Flow:

The founder turned to a toilet metaphor to explain one aspect of his idea of ownership. “If you’re in an apartment building, and you’re a renter, and your toilet gets clogged, you call the super,” he said. In contrast, “if you’re in your own apartment, and you bought it and you own it and your toilet gets clogged, you take the plunger.” For Neumann, fixing up your own apartment means shifting from “being transactional to actually being part of a community” and “feeling like you own something.”

In a functioning society, Adam Neumann would not be given a single dollar. This quote proves that he has never unclogged a toilet, because in the event that you could unclog your toilet in an apartment you rented, you’d probably do it. If the clog was so severe it required the super, you would probably still call a plumber if you owned the place, because your nasty business has created a problem you cannot solve.

What I am suggesting is that Adam Neumann doesn’t know anything about home ownership, or unclogging toilets, or toilets, or the regular experience of being a human. Yet he is given unfathomable

Chrome: I’d seen her maybe half a dozen times in the Gentleman Loser. Maybe she was slumming, or checking out the human condition, a condition she didn’t exactly aspire to. A sweet little heart-shaped face framing the nastiest pair of eyes you ever saw. She’d looked fourteen for as long as anyone could remember, hyped out of anything like a normal metabolism on some massive program of serums and hormones. She was as ugly a customer as the street ever produced, but she didn’t belong to the street anymore. She was one of the Boys, Chrome, a member in good standing of the local Mob subsidiary. Word was, she’d gotten started as a dealer, back when synthetic pituitary hormones were still proscribed. But she hadn’t had to move hormones for a long time. Now she owned the House of Blue Lights.

‘You’re flat-out crazy, Quine. You give me one sane reason for having that stuff on your screen. You ought to dump it, and I mean now...’

‘Talk in the Loser,’ he said, shrugging out of the leather jacket. ‘Black Myron and Crow Jane. Jane, she’s up on all the sex lines, claims she knows where the money goes. So she’s arguing with Myron that Chrome’s the controlling interest in the Blue Lights, not just some figurehead for the Boys.’

“‘The Boys,” Bobby,’ I said. ‘That’s the operative word there. You still capable of seeing that? We don’t mess with the Boys, remember? That’s why we’re still walking around.’

‘That’s why we’re still poor, partner.’ He settled back into the swivel chair in front of the console, unzipped his jump suit, and scratched his skinny white chest. ‘But maybe not for much longer.’

‘I think maybe this partnership just got itself permanently dissolved.’

Then he grinned at me. That grin was truly crazy, feral and focused, and I knew that right then he really didn’t give a shit about dying.

‘Look,’ I said, ‘I’ve got some money left, you know? Why don’t you take it and get the tube to Miami, catch a hopper to Montego Bay. You need rest, man. You’ve got to get your act together.’

‘My act, Jack,’ he said, punching something on the keyboard, ‘never has been this together before.’ The neon prayer rug on the screen shimmered and woke as an animation program cut in, ice lines weaving with hypnotic frequency, a living mandala. Bobby kept punching, and the movement slowed; the pattern resolved itself, grew slightly less complex, became an alternation between two distant configurations. A first-class piece of work, and I hadn’t thought he was still that good. ‘Now,’ he said, ‘there, see it? Wait. There. There again. And there. Easy to miss. That’s it. Cuts in every hour and twenty minutes with a squirt transmission to their comsat. We could live for a year on what she pays them weekly in negative interest.’

‘Whose comsat?’

‘Zürich. Her bankers. That’s her bankbook, Jack. That’s where the money goes. Crow Jane was right.’

I stood there. My arm forgot to click.

‘So how’d you do in New York, partner? You get anything that’ll help me cut ice? We’re going to need whatever we can get.’

I kept my eyes on his, forced myself not to look in the direction of the waldo, the jeweler’s vise. The Russian program was there, under the dust cover.

Wild cards, luck changers.

‘Where’s Rikki?’ I asked him, crossing to the console, pretending to study the alternating patterns on the screen.

‘Friends of hers,’ he shrugged, ‘kids, they’re all into simstim.’ He smiled absently. ‘I’m going to do it for her, man.’

‘I’m going out to think about this, Bobby. You want me to come back, you keep your hands off the board.’

‘I’m doing it for her,’ he said as the door closed behind me. ‘You know I am.’

And down now, down, the program a roller coaster through this fraying maze of shadow walls, gray cathedral spaces between the bright towers. Headlong speed.

Black ice. Don't think about it. Black ice.

Too many stories in the Gentleman Loser; black ice is a part of the mythology. Ice that kills. Illegal, but then aren't we all? Some kind of neural-feedback weapon, and you connect with it only once. Like some hideous Word that eats the mind from the inside out. Like an epileptic spasm that goes on and on until there's nothing left at all...

And we're diving for the floor of Chrome's shadow castle.

Trying to brace myself for the sudden stopping of breath, a sickness and final slackening of the nerves. Fear of that cold Word waiting, down there in the dark.

I went out and looked for Rikki, found her in a cafe with a boy with Sendai eyes, half-healed suture lines radiating from his bruised sockets. She had a glossy brochure spread open on the table, Tally Isham smiling up from a dozen photographs, the Girl with the Zeiss Ikon Eyes.

Her little simstim deck was one of the things I'd stacked under my bench the night before, the one I'd fixed for her the day after I'd first seen her. She spent hours jacked into that unit, the contact band across her forehead like a gray plastic tiara. Tally Isham was her favorite, and with the contact band on, she was gone, off somewhere in the recorded sensorium of simstim's biggest star. Simulated stimuli: the world – all the interesting parts, anyway – as perceived by Tally Isham. Tally raced a black Fokker ground-effect plane across Arizona mesa tops. Tally dived the Truk Island preserves. Tally partied with the superrich on private Greek islands, heartbreaking purity of those tiny white seaports at dawn.

Actually she looked a lot like Tally, same coloring and cheekbones. I thought Rikki's mouth was stronger. More sass. She didn't want to be Tally Isham, but she coveted the job. That was her ambition, to be in simstim. Bobby just laughed it off. She talked to me about it, though. 'How'd I look with a pair of these?' she'd ask, holding a full-page headshot, Tally Isham's blue Zeiss Ikon's lined up with her own amber-brown. She'd had her corneas done twice, but she still wasn't 20-20; so she wanted Ikons. Brand of the stars. Very expensive.

'You still window-shopping for eyes?' I asked as I sat down.

'Tiger just got some,' she said. She looked tired, I thought.

Tiger was so pleased with his Sendais that he couldn't help smiling, but I doubted whether he'd have smiled otherwise. He had the kind of uniform good looks you get after your seventh trip to the surgical boutique; he'd probably spend the rest of his life looking vaguely like each new season's media front-runner; not too obvious a copy, but nothing too original, either.

'Sendai, right?' I smiled back.

He nodded. I watched as he tried to take me in with his idea of a professional simstim glance. He was pretending that he was recording. I thought he spent too long on my arm. 'They'll be great on peripherals when the muscles heal,' he said, and I saw how carefully he reached for his double espresso. Sendai eyes are notorious for depth-perception defects and warranty hassles, among other things.

'Tiger's leaving for Hollywood tomorrow.'

'Then maybe Chiba City, right.' I smiled at him. He didn't smile back. 'Got an offer, Tiger? Know an agent?'

'Just checking it out,' he said quietly. Then he got up and left. He said a quick goodbye to Rikki, but not to me.

'That kid's optic nerves may start to deteriorate inside six months. You know that, Rikki? Those Sendais are illegal in England, Denmark, lots of places. You can't replace nerves.'

result or answer to your query - it provides the answer that it believes is most beneficial or profitable to Google. Google Search provides a "free" service, but the cost is a source of information corrupted by a profit-seeking entity looking to manipulate you into giving money to the profit-seeking entities that pay them.

The net result is a product that completely sucks. "Googling" something is now an exercise in pain, regularly leading you to generic Search Engine Optimized content that doesn't actually answer your question. Google's push to hyper-optimization has also led it to serve results based on what it *thinks* people mean, rather than what they actually said. It's frustrating, upsetting and annoying. A problem that likely hits hundreds of millions of people a day, yet Google doesn't have to change a thing, because the street likes that they have found more innovative ways to get blood from a stone. These moves are unquestionably hurting Google, to the point that Microsoft's Bing (paired with OpenAI's ChatGPT), has gained major headlines for providing the service that everybody wished Google would[10].

That's because Google has, like every major tech company, focused entirely on what will make revenues increase, even if the cost of doing so is destroying its entire legacy. Google has announced their own "Bard AI" to compete with Bing's ChatGPT integration,[11] and I'll be honest - I feel a little crazy that nobody is saying the truth, which is that ***Google broke the product that made them famous and is now productizing fixing their own problem as innovation.***

That's because the markets do not prioritize innovation, or sustainable growth, or stable, profitable enterprises. As a result, companies regularly do not function with the intent of making "good" businesses - they want businesses that semiotically align with what investors - private and public - believe to be "good."

Despite its ubiquity, companies like Uber should not exist. Uber has not made a profit from its businesses. They had a net *loss* of 1.21 billion last quarter, yet the street fell over itself to praise the company because "gross bookings grew 19% year-over-year[12]" for their unprofitable businesses that largely hinge upon the government failing to impose sensible labor laws, a con that will eventually come to an end[13], and indeed, has ended in some territories like the UK, where Uber drivers are now recognized as employees, and are therefore entitled to pensions, paid vacation time, and a minimum wage. London, I note, is one of Uber's most important markets.

Yet as of writing, Uber's stock is up 5%.

The media itself somewhat fuels this economy of growth-mongering. CNBC reports earnings like many other media entities, but their reports on, say, Uber[14] fail to acknowledge the fact that Uber has spent nearly 15 years burning money. It has never turned a profit. Even with its push into freight and food delivery, it may never turn a profit, no matter how much it contorts its financials to pretend otherwise[15]. Yet acknowledging the truth is that much worse because Uber will not be killed, because people keep buying the stock, because it is a "valuable company" in the eyes of markets that have *fucking cataracts*.

This is why we see such vast oscillations of hiring and firing - because these companies are never, ever punished for failing to operate their businesses in a sustainable way, or even with a view for the future, particularly when it comes to macroeconomic trends that literally everyone else saw coming.

Their business models were predicated on an endless supply of cheap money, even though the Fed steadily ratcheted interest rates in the years leading up to the Covid pandemic, only slashing them to mitigate the pain of Covid and (to a lesser extent) the US-China trade war.. The specter of inflation reared its ugly head as early as 2020, first driven by the lockdown-induced chaos on supply chains, and then exacerbated further by the war in Ukraine, the collateral damage of China's Zero Covid policy, and a chronic labor shortage in most industrialized countries.

The markets do not react when they are mass-hiring people to capture consumer demand. They do not react to the fact that Microsoft, for example, seems to be[16] laying[17] off[18] people[19] almost[20] every year. In 2020, CEO Satya Nadella called for a "referendum on capitalism[21]," telling businesses to start to grade themselves on the "wider economic benefits they bring to society, rather than profits." To be clear, this was four months after Microsoft laid off 1000 people[22], one year before they hired 23,000 people[23], and a few months after which they laid off 10,000 people to "deliver results on an ongoing basis, while investing in [their] long-term opportunity[24]."

The Rot Economy

Ed Zitron

Feb 9, 2023

Originally published: <https://www.wheresyoured.at/the-rot-economy/>

At the center of everything I’ve written for the last few months (if not the last few years), sits a cancerous problem with the fabric of how capital is deployed in modern business. Public and private investors, along with the markets themselves, have become entirely decoupled from the concept of what “good” business truly is, focusing on one metric — one *truly noxious* metric — over all else: growth.

“Growth” in this case is not necessarily about being “bigger” or “better,” it is simply “more.” It means that the company is generating more revenue, higher valuations, gaining more market share, and then finding more ways to generate these things. Businesses are expected to be - and rewarded for being - eternal burning engines of capital that create more and more shareholder value while, hopefully, providing a service to a customer in the process. In the public markets, that means that companies like Google, Meta, and Microsoft[1] were rewarded for having unfocused, capital-intensive businesses that required mass layoffs when times got tough[2], because the market loved the idea that they’d found a way to save money. They weren’t punished for their poor planning, their stagnating products[3], their mismanagement of human capital, or their general lack of any real innovation because *the numbers kept going up*.

When I wrote in October that Mark Zuckerberg was going to kill his company[4], the street responded in kind, savaging Meta’s stock for burning cash building a metaverse that was never going to exist. Yet once Zuckerberg fired 11,000 people and claimed that 2023 would be the “year of efficiency[5],” the market responded with double-digit increases in the price of Meta’s shares, despite the fact that Facebook’s active user growth *declined* and they lost \$13.7 billion on the same metaverse department that caused the stock to drop the last time[6].

The markets seemed to ignore the \$410 million fine that Meta received for GDPR violations[7], along with the fact that European users will now have to deliberately opt-in to sharing their data - which is bad, considering only about 25% of iOS users choose to opt-in to app tracking[8], and their business model is intrinsically linked to the repurposing of customer data into ad targeting telemetry.

Let’s be abundantly clear: Meta’s core advertising models depend heavily on things that likely become impossible to do legally (or even technically, given Apple’s App Tracking Transparency, Alphabet’s retirement of the third-party tracking cookie, and the Chromium Project’s planned blocking of non-cookie fingerprinting technologies) in the next decade. Their other products simply do not make that much money. Their CEO’s big idea to make more money has *lost* them billions of dollars, and likely won’t make them any for quite some time. Yet Meta remains beloved, because *the numbers are going up*.

Killing Innovation

Google has a similar yet slightly different story, where their core product - search - has gone from a place where you find information to an increasingly-manipulated labyrinth of SEO-optimized garbage shipped straight from the content factories. As Charlie Warzel put it last year:[9] “Google Search, what many consider an indispensable tool of modern life, is dead or dying.” Users have to effectively find cheat codes - adding things like “[whatever you’re searching]+Reddit” to get reliable answers. Despite its decades-long efforts to improve the quality of organic results, Google remains easily-gamed by anyone who knows how to craft an algorithm-friendly headline.

Without finding a way to negotiate with Google Search, you’re offered a fragmented buffet of content provided by Google’s algorithm, either based on how much they’ve been paid to prioritize said content or by how companies have engineered content to rank higher on search. Google no longer provides the “best”

‘Hey, Jack, no lectures.’ She stole one of my croissants and nibbled at the tip of one of its horns.

‘I thought I was your adviser, kid.’

‘Yeah. Well, Tiger’s not too swift, but everybody knows about Sendais. They’re all he can afford. So he’s taking a chance. If he gets work, he can replace them.’

‘With these?’ I tapped the Zeiss Ikon brochure. ‘Lot of money, Rikki. You know better than to take a gamble like that.’

She nodded. ‘I want Ikons.’

‘If you’re going up to Bobby’s tell him to sit tight until he hears from me.’

‘Sure. It’s business?’

‘Business,’ I said. But it was craziness.

I drank my coffee, and she ate both my croissants. Then I walked her down to Bobby’s. I made fifteen calls, each one from a different pay phone.

Business. Bad craziness.

All in all, it took us six weeks to set the burn up, six weeks of Bobby telling me how much he loved her. I worked even harder, trying to get away from that.

Most of it was phone calls. My fifteen initial and very oblique inquiries each seemed to breed fifteen more. I was looking for a certain service Bobby and I both imagined as a requisite part of the world’s clandestine economy, but which probably never had more than five customers at a time. It would be one that never advertised.

We were looking for the world’s heaviest fence, for a non-aligned money laundry capable of dry-cleaning a megabuck online cash transfer and then forgetting about it.

All those calls were a waste, finally, because it was the Finn who put me on to what we needed. I’d gone up to New York to buy a new blackbox rig, because we were going broke paying for all those calls.

I put the problem to him as hypothetically as possible.

‘Macao,’ he said.

‘Macao?’

‘The Long Hum family. Stockbrokers.’

He even had the number. You want a fence, ask another fence.

The Long Hum people were so oblique that they made my idea of a subtle approach look like a tactical nuke-out. Bobby had to make two shuttle runs to Hong Kong to get the deal straight. We were running out of capital, and fast. I still don’t know why I decided to go along with it in the first place; I was scared of Chrome, and I’d never been all that hot to get rich.

I tried telling myself that it was a good idea to burn the House of Blue Lights because the place was a creep joint, but I just couldn’t buy it. I didn’t like the Blue Lights, because I’d spent a supremely depressing evening there once, but that was no excuse for going after Chrome. Actually I halfway assumed we were going to die in the attempt. Even with that killer program, the odds weren’t exactly in our favor.

Bobby was lost in writing the set of commands we were going to plug into the dead center of Chrome’s computer. That was going to be my job, because Bobby was going to have his hands full trying to keep the Russian program from going straight for the kill. It was too complex for us to rewrite, and so he was going to try to hold it back for the two seconds I needed.

I made a deal with a streetfighter named Miles. He was going to follow Rikki the night of the burn, keep her in sight, and phone me at a certain time. If I wasn't there, or didn't answer in just a certain way, I'd told him to grab her and put her on the first tube out. I gave him an envelope to give her, money and a note.

Bobby really hadn't thought about that, much, how things would go for her if we blew it. He just kept telling me he loved her, where they were going to go together, how they'd spend the money.

'Buy her a pair of Ikons first, man. That's what she wants. She's serious about that simstim scene.'

'Hey,' he said, looking up from the keyboard, 'she won't need to work. We're going to make it, Jack. She's my luck. She won't ever have to work again.'

'Your luck,' I said. I wasn't happy. I couldn't remember when I had been happy. 'You seen your luck around lately?'

He hadn't, but neither had I. We'd both been too busy.

I missed her. Missing her reminded me of my one night in the House of Blue Lights, because I'd gone there out of missing someone else. I'd gotten drunk to begin with, then I'd started hitting Vasopressin inhalers. If your main squeeze has just decided to walk out on you, booze and Vasopressin are the ultimate in masochistic pharmacology; the juice makes you maudlin and the Vasopressin makes you remember, I mean really remember. Clinically they use the stuff to counter senile amnesia, but the street finds its own uses for things. So I'd bought myself an ultra-intense replay of a bad affair; trouble is, you get the bad with the good. Go gunning for transports of animal ecstasy and you get what you said, too, and what she said to that, how she walked away and never looked back.

I don't remember deciding to go to the Blue Lights, or how I got there, hushed corridors and this really tacky decorative waterfall trickling somewhere, or maybe just a hologram of one. I had a lot of money that night; somebody had given Bobby a big roll for opening a three-second window in someone else's ice.

I don't think the crew on the door liked my looks, but I guess my money was okay.

I had more to drink there when I'd done what I went there for. Then I made some crack to the barman about closet necrophiliacs, and that didn't go down too well. Then this very large character insisted on calling me War Hero, which I didn't like. I think I showed him some tricks with the arm, before the lights went out, and I woke up two days later in a basic sleeping module somewhere else. A cheap place, not even room to hang yourself. And I sat there on that narrow foam slab and cried.

Some things are worse than being alone. But the thing they sell in the House of Blue Lights is so popular that it's almost legal.

At the heart of darkness, the still center, the glitch systems shred the dark with whirlwinds of light, translucent razors spinning away from us; we hang in the center of a silent slow-motion explosion, ice fragments falling away forever, and Bobby's voice comes in across light-years of electronic void illusion—

'Burn the bitch down. I can't hold the thing back —'

The Russian program, rising through towers of data, blotting out the playroom colors. And I plug Bobby's homemade command package into the center of Chrome's cold heart. The squirt transmission cuts in, a pulse of condensed information that shoots straight up, past the thickening tower of darkness, the Russian program, while Bobby struggles to control that crucial second. An unformed arm of shadow twitches from the towering dark, too late.

We've done it.

The matrix folds itself around me like an origami trick.

And the loft smells of sweat and burning circuitry.

I thought I heard Chrome scream, a raw metal sound, but I couldn't have.

24. <https://www.datacamp.com/blog/a-beginners-guide-to-gpt-3#:~:text=Language%20modeling%20is,predicting%20word%20sequences.>
25. <http://arxiv.org/abs/2212.03551>
26. <http://arxiv.org/abs/2212.03551>
27. [https://en.wikipedia.org/wiki/The_Human_Centipede_\(First_Sequence\)](https://en.wikipedia.org/wiki/The_Human_Centipede_(First_Sequence))
28. https://en.wikipedia.org/wiki/Ferdinand_de_Saussure
29. <https://www.fortressofdoors.com/ai-markets-for-lemons-and-the-great-logging-off/>
30. <https://www.fortressofdoors.com/ai-markets-for-lemons-and-the-great-logging-off/#:~:text=a%20resurgence%20and%20even%20of%20fetishization%20of%20explicitly%20%22offline%22%20culture%2C%20where%20the%20%22Great%20Logging%20Off%22%20becomes%20literal>
31. <https://www.fortressofdoors.com/ai-markets-for-lemons-and-the-great-logging-off/#:~:text=Seventh%2C%20real,amenities%20are%20available>

The idea is something like this: you show up in person to register your online accounts or domains. You then get some kind of special badge or mark online legitimising you as a Real Human. It may or may not be on the blockchain somehow.

Google might look something like this:

[[Editor's note: image elided, a Google screenshot with some results tagged as "Certified Human]]

The whole thing seems fraught with problems, susceptible to abuse, and ultimately impractical. Would it even be the web if everyone knew you were really a dog?

5. Show up in meatspace

The final edge we have over language models is that we can prove we're real humans by showing up IRL with our real human bodies. We can arrange to meet Twitter mutuals offline over coffee. We can organise meetups and events and conferences and unconferences and hangouts and pub nights.

In *Markets for Lemons* and the *Great Logging Off*[[29]], Lars Doucet proposed several knock-on effects from this offline-first future. We might see increased fetishisation[[30]] of anti-screen culture, as well as real estate[[31]] price increases in densely populated areas.

For the moment we can still check humanness over Zoom, but live video generation is getting good enough that I don't think that defence will last long.

There are, of course, many people who can't move to an offline-first life; people with physical disabilities. People who live in remote, rural places. People with limited time and caretaking responsibilities for the very young or the very old. They will have a harder time verifying their humanness online. I don't have any grand ideas to help solve this, but I hope we find better solutions than my paltry list.

As the forest grows darker, noisier, and less human, I expect to invest more time in in-person relationships and communities. And while I love meatspace, this still feels like a loss.

Footnotes:

1. <https://maggieappleton.com/cozy-web>
2. <https://maggieappleton.com/garden-history>
3. <https://www.ribbonfarm.com/2010/07/26/a-big-little-idea-called-legibility/>
4. <https://www.jasper.ai/>
5. <https://www.gomoonbeam.com/>
6. <https://www.copy.ai/>
7. <https://anyword.com/>
8. <https://openai.com/blog/chatgpt/>
9. <https://twitter.com/elonmusk/status/1599128577068650498>
10. <https://twitter.com/volodarik/status/1600854935515844610>
11. https://twitter.com/yu_angela/status/1599808692085743616
12. <https://twitter.com/levie/status/1599156293050433536>
13. <https://www.midjourney.com/home/?callbackUrl=/app/>
14. <https://openai.com/dall-e-2/>
15. https://en.wikipedia.org/wiki/Stable_Diffusion
16. <https://twitter.com/StuartBlitz/status/1602834224284897282>
17. <https://zapier.com/apps/openai/integrations>
18. <https://twitter.com/SamRo/status/1605919856808714240>
19. <https://runwayml.com/>
20. <https://byautomata.io/>
21. <https://www.theatlantic.com/technology/archive/2021/08/dead-internet-theory-wrong-but-feels-true/619937/>
22. <https://nymag.com/intelligencer/2018/12/how-much-of-the-internet-is-fake.html>
23. https://en.wikipedia.org/wiki/Turing_test

Bobby was laughing, tears in his eyes. The elapsed-time figure in the corner of the monitor read 07:24:05. The burn had taken a little under eight minutes.

And I saw that the Russian program had melted in its slot.

We'd given the bulk of Chrome's Zurich account to a dozen world charities. There was too much there to move, and we knew we had to break her, burn her straight down, or she might come after us. We took less than ten per cent for ourselves and shot it through the Long Hum setup in Macao. They took sixty per cent of that for themselves and kicked what was left back to us through the most convoluted sector of the Hong Kong exchange. It took an hour before our money started to reach the two accounts we'd opened in Zürich.

I watched zeros pile up behind a meaningless figure on the monitor. I was rich.

Then the phone rang. It was Miles. I almost blew the code phrase.

'Hey, Jack, man, I dunno – what's it all about, with this girl of yours? Kinda funny thing here...'

'What? Tell me.'

'I been on her, like you said, tight but out of sight. She goes to the Loser, hangs out, then she gets a tube. Goes to the House of Blue Lights –'

'She what?'

'Side door. Employees only. No way I could get past their security.'

'Is she there now?'

'No, man, I just lost her. It's insane down here, like the Blue Lights just shut down, looks like for good, seven kinds of alarms going off, everybody running, the heat out in riot gear...Now there's all this stuff going on, insurance guys, real-estate types, vans with municipal plates...'

'Miles, where'd she go?'

'Lost her, Jack.'

'Look, Miles, you keep the money in the envelope, right?'

'You serious? Hey, I'm real sorry. I –'

I hung up.

'Wait'll we tell her,' Bobby was saying, rubbing a towel across his bare chest.

'You tell her yourself, cowboy. I'm going for a walk.'

So I went out into the night and the neon and let the crowd pull me along, walking blind, willing myself to be just a segment of that mass organism, just one more drifting chip of consciousness under the geodesics. I didn't think, just put one foot in front of another, but after a while I did think, and it all made sense. She'd needed the money.

I thought about Chrome, too. That we'd killed her, murdered her, as surely as if we'd slit her throat. The night that carried me along through the malls and plazas would be hunting her now, and she had nowhere to go. How many enemies would she have in this crowd alone? How many would move, now they weren't held back by fear of her money? We'd taken her for everything she had. She was back on the street again. I doubted she'd live till dawn.

Finally I remembered the café, the one where I'd met Tiger.

Her sunglasses told the whole story, huge black shades with a telltale smudge of fleshtone paintstick in the corner of one lens. 'Hi, Rikki,' I said, and I was ready when she took them off.

Blue. Tally Isham blue. The clear trademark blue they're famous for, ZEISS IKON ringed each iris in tiny capitals, the letters suspended there like flecks of gold.

'They're beautiful,' I said. Paintstick covered the bruising. No scars with work that good. 'You made some money.'

'Yeah, I did.' Then she shivered. 'But I won't make any more, not that way.'

'I think that place is out of business.'

'Oh.' Nothing moved in her face then. The new blue eyes were still and very deep.

'It doesn't matter. Bobby's waiting for you. We just pulled down a big score.'

'No. I've got to go. I guess he won't understand, but I've got to go.'

I nodded, watching the arm swing up to take her hand; it didn't seem to be part of me at all, but she held on to it like it was.

'I've got a one-way ticket to Hollywood. Tiger knows some people I can stay with. Maybe I'll even get to Chiba City.'

She was right about Bobby. I went back with her. He didn't understand. But she'd already served her purpose, for Bobby, and I wanted to tell her not to hurt for him, because I could see that she did. He wouldn't even come out into the hallway after she had packed her bags. I put the bags down and kissed her and messed up the paintstick, and something came up inside me the way the killer program had risen above Chrome's data. A sudden stopping of the breath, in a place where no word is. But she had a plane to catch.

Bobby was slumped in the swivel chair in front of his monitor, looking at his string of zeros. He had his shades on, and I knew he'd be in the Gentleman Loser by nightfall, checking out the weather, anxious for a sign, someone to tell him what his new life would be like. I couldn't see it being very different. More comfortable, but he'd always be waiting for that next card to fall.

I tried not to imagine her in the House of Blue Lights, working three-hour shifts in an approximation of REM sleep, while her body and a bundle of conditioned reflexes took care of business. The customers never got to complain that she was faking it, because those were real orgasms. But she felt them, if she felt them at all, as faint silver flares somewhere out on the edge of sleep. Yeah, it's so popular, it's almost legal. The customers are torn between needing someone and wanting to be alone at the same time, which has probably always been the name of that particular game, even before we had the neuroelectronics to enable them to have it both ways.

I picked up the phone and punched the number for her airline. I gave them her real name, her flight number. 'She's changing that,' I said, 'to Chiba City. That's right. Japan.' I thumbed my credit card into the slot and punched my ID code. 'First class.' Distant hum as they scanned my credit records. 'Make that a return ticket.'

But I guess she cashed the return fare, or else she didn't need it, because she hasn't come back. And sometimes late at night I'll pass a window with posters of simstim stars, all those beautiful, identical eyes staring back at me out of faces that are nearly identical, and sometimes the eyes are hers, but none of the faces are, none of them ever are, and I see her far out on the edge of all this sprawl of night and cities, and then she waves goodbye.

extensive prompt engineering. They stick to generics. They hedge. They leave out details. They have trouble maintaining a coherent sense of self over thousands of words.

Hipsterism and recency bias will help us here. Referencing obscure concepts, friends who are real but not famous, niche interests, and recent events all make you plausibly more human.

2. Be original, critical, and sophisticated

Easier said than done, but one of the best ways to prove you're not a predictive language model is to demonstrate critical and sophisticated thinking.

Language models spit out text that sounds like a B+ college essay. Coherent, seemingly comprehensive, but never truly insightful or original (at least for now).

In a repulsively evocative metaphor, they engage in "human centipede[27] epistemology." Language models regurgitate text from across the web, which some humans read and recycle into "original creations," which then become fodder to train other language models, and around and around we go recycling generic ideas and arguments and tropes and ways of thinking.

Hard exiting out of this cycle requires coming up with unquestionably original thoughts and theories. It means seeing and synthesising patterns across a broad range of sources: books, blogs, cultural narratives served up by media outlets, conversations, podcasts, lived experiences, and market trends. We can observe and analyse a much fuller range of inputs than bots and generative models can.

It will raise the stakes for everyone. As both consumers of content and creators of it, we'll have to foster a greater sense of critical thinking and scepticism.

This all sounds a bit rough, but there's a lot of hope in this vision. In a world of automated intelligence, our goalposts for intelligence will shift. We'll raise our quality bar for what we expect from humans. When a machine can pump out a great literature review or summary of existing work, there's no value in a person doing it.

3. Develop creative language quirks, dialects, memes, and jargon

The linguist Ferdinand de Saussure[28] argued there are two kinds of language:

- **La langue** is the formal concept of language. These are words we print in the dictionary, distribute via educational institutions, and reprimand one another for getting it wrong.
- **La parole** is the speech of everyday life. These are the informal, diverse, and creative speech acts we perform in conversations, social gatherings, and text to the group WhatsApp. This is where language evolves.

We have designed a system that automates a standardised way of writing. We have codified *la langue* at a specific point in time.

What we have left to play with is *la parole*. No language model will be able to keep up with the pace of weird internet lingo and memes. I expect we'll lean into this. Using neologisms, jargon, euphemistic emoji, unusual phrases, ingroup dialects, and memes-of-the-moment will help signal your humanity.

Not unlike teenagers using language to subvert their elders, or oppressed communities developing dialects that allow them to safely communicate amongst themselves.

4. Consider institutional verification

This solution feels the least interesting. We're already hearing rumblings of how "verification" by centralised institutions or companies might help us distinguish between meat brains and metal brains.

It goes beyond text too: video essays on YouTube[18] , TikTok clips, podcasts, slide decks, and Instagram stories can all be generated by patchworking[19] together ML systems. And then regurgitated[20] for each medium.

We're about to drown in a sea of pedestrian takes. An explosion of noise that will drown out any signal. Goodbye to finding original human insights or authentic connections under that pile of cruff.

Many people will say we already live in this reality. We've already become skilled at sifting through unhelpful piles of "optimised content" designed to gather clicks and advertising impressions.

4chan proposed dead internet theory[21] years ago: that most of the internet is "empty and devoid of people" and has been taken over by artificial intelligence. A milder version of this theory is simply that we're overrun with bots[22] . Most of us take that for granted at this point.

But I think the sheer volume and scale of what's coming will be meaningfully different. And I think we're unprepared. Or at least, I am.

Passing the Reverse Turing Test

Our new challenge as little snowflake humans will be to prove we aren't language models. It's the reverse turing test[23].

After the forest expands, we will become deeply sceptical of one another's *realness*. Every time you find a new favourite blog or Twitter account or Tiktok personality online, you'll have to ask: Is this really a whole human with a rich and complex life like mine? Is there a *being* on the other end of this web interface I can form a relationship with?

Before you continue, pause and consider: How would *you* prove you're not a language model generating predictive text? What special human tricks can you do that a language model can't?

1. Triangulate objective reality

As language models become increasingly capable and impressive, we should remember they are, at their core, linguistic prediction systems[24] . They cannot (yet) reason like a human.

They do not have beliefs based on evidence, claims, and principles. They cannot consult external sources and run experiments against objective reality. They cannot go outside and touch grass.

In short, they do not have access to the same shared reality we do. They do not have embodied experiences, and cannot sense the world as we can sense it; they don't have vision, sound, taste, or touch. They cannot feel emotion or tightly hold a coherent set of values. They are not part of cultures, communities, or histories.

They are a language model in a box. If a historical event, fact, person, or concept wasn't part of their training data, they can't tell you about it. They don't know about events that happened after a certain cutoff date.

I found Murray Shanahan's paper on Talking About Large Language Models[25] (2022) full of helpful reflections on this point:

Humans are members of a community of language-users inhabiting a shared world, and this primal fact makes them essentially different to large language models. We can consult the world to settle our disagreements and update our beliefs. We can, so to speak, "triangulate" on objective reality.

Murray Shanahan – Talking About Large Language Models[26]

This leaves us with some low-hanging fruit for humanness. We can tell richly detailed stories grounded in our specific contexts and cultures: place names, sensual descriptions, local knowledge, and, well the *je ne sais quoi* of being alive. Language models can decently mimic this style of writing but most don't without

AI machines aren't 'hallucinating.' But their makers are

Naomi Klein

May 8, 2023

Originally published: <https://www.theguardian.com/commentisfree/2023/may/08/ai-machines-hallucinating-naomiklein>

Inside the many debates swirling around the rapid rollout of so-called artificial intelligence, there is a relatively obscure skirmish focused on the choice of the word "hallucinate".

This is the term that architects and boosters of generative AI have settled on to characterize responses served up by chatbots that are wholly manufactured, or flat-out wrong. Like, for instance, when you ask a bot for a definition of something that doesn't exist and it, rather convincingly, gives you one[1], complete with made-up footnotes. "No one in the field has yet solved the hallucination problems," Sundar Pichai, the CEO of Google and Alphabet, told[2] an interviewer recently.

That's true – but why call the errors "hallucinations" at all? Why not algorithmic junk? Or glitches? Well, hallucination refers to the mysterious capacity of the human brain to perceive phenomena that are not present, at least not in conventional, materialist terms. By appropriating a word commonly used in psychology, psychedelics and various forms of mysticism, AI's boosters, while acknowledging the fallibility of their machines, are simultaneously feeding the sector's most cherished mythology: that by building these large language models, and training them on everything that we humans have written, said and represented visually, they are in the process of birthing an animate intelligence on the cusp of sparking an evolutionary leap for our species. How else could bots like Bing and Bard be tripping out there in the ether?

Warped hallucinations are indeed afoot in the world of AI, however – but it's not the bots that are having them; it's the tech CEOs who unleashed them, along with a phalanx of their fans, who are in the grips of wild hallucinations, both individually and collectively. Here I am defining hallucination not in the mystical or psychedelic sense, mind-altered states that can indeed assist in accessing profound, previously unperceived truths. No. These folks are just tripping: seeing, or at least claiming to see, evidence that is not there at all, even conjuring entire worlds that will put their products to use for our universal elevation and education.

Generative AI will end poverty, they tell us. It will cure all disease. It will solve climate change. It will make our jobs more meaningful and exciting. It will unleash lives of leisure and contemplation, helping us reclaim the humanity we have lost to late capitalist mechanization. It will end loneliness. It will make our governments rational and responsive. These, I fear, are the real AI hallucinations and we have all been hearing them on a loop ever since Chat GPT launched at the end of last year.

There is a world in which generative AI, as a powerful predictive research tool and a performer of tedious tasks, could indeed be marshalled to benefit[3] humanity, other species and our shared home. But for that to happen, these technologies would need to be deployed inside a vastly different economic and social order than our own, one that had as its purpose the meeting of human needs and the protection of the planetary systems that support all life.

And as those of us who are not currently tripping well understand, our current system is nothing like that. Rather, it is built to maximize the extraction of wealth and profit – from both humans and the natural world – a reality that has brought us to what we might think of it as capitalism's techno-necro stage. In that reality of hyper-concentrated power and wealth, AI – far from living up to all those utopian hallucinations – is much more likely to become a fearsome tool of further dispossession and despoilation.

I'll dig into why that is so. But first, it's helpful to think about the *purpose* the utopian hallucinations about AI are serving. What work are these benevolent stories doing in the culture as we encounter these strange

new tools? Here is one hypothesis: they are the powerful and enticing cover stories for what may turn out to be the largest and most consequential theft in human history. Because what we are witnessing is the wealthiest companies in history (Microsoft, Apple, Google, Meta, Amazon ...) unilaterally seizing the sum total of human knowledge that exists in digital, scrapable form and walling it off inside proprietary products, many of which will take direct aim at the humans whose lifetime of labor trained the machines without giving permission or consent.

This should not be legal. In the case of copyrighted material that we now know^[4] trained the models (including this newspaper), various lawsuits^[5] have been filed that will argue this was clearly illegal. Why, for instance, should a for-profit company be permitted to feed the paintings, drawings and photographs of living artists into a program like Stable Diffusion or Dall-E 2 so it can then be used to generate doppelganger versions of those very artists' work, with the benefits flowing to everyone but the artists themselves?

The painter and illustrator Molly Crabapple is helping lead a movement of artists challenging this theft. "AI art generators are trained on enormous datasets, containing millions upon millions of copyrighted images, harvested without their creator's knowledge, let alone compensation or consent. This is effectively the greatest art heist in history. Perpetrated by respectable-seeming corporate entities backed by Silicon Valley venture capital. It's daylight robbery," a new open^[6] letter she co-drafted states.

The trick, of course, is that Silicon Valley routinely calls theft "disruption" – and too often gets away with it. We know this move: charge ahead into lawless territory; claim the old rules don't apply to your new tech; scream that regulation will only help China – all while you get your facts solidly on the ground. By the time we all get over the novelty of these new toys and start taking stock of the social, political and economic wreckage, the tech is already so ubiquitous that the courts^[7] and policymakers throw up their hands.

We saw it with Google's book and art scanning. With Musk's space colonization. With Uber's assault on the taxi industry. With Airbnb's attack on the rental market. With Facebook's promiscuity with our data. Don't ask for permission, the disruptors like to say, ask for forgiveness. (And lubricate the asks with generous campaign contributions.)

In *The Age of Surveillance Capitalism*, Shoshana Zuboff^[8] meticulously details how Google's Street View maps steamrolled over privacy norms by sending its camera-bedecked cars out to photograph our public roadways and the exteriors of our homes. By the time the lawsuits defending privacy rights rolled around, Street View was already so ubiquitous on our devices (and so cool, and so convenient ...) that few courts outside Germany^[9] were willing to intervene.

Now the same thing that happened to the exterior of our homes is happening to our words, our images, our songs, our entire digital lives. All are currently being seized and used to train the machines to simulate thinking and creativity. These companies must know they are engaged in theft, or at least that a strong case^[10] can be made that they are. They are just hoping that the old playbook works one more time – that the scale of the heist is already so large and unfolding with such speed^[11] that courts and policymakers will once again throw up their hands in the face of the supposed inevitability of it all.

It's also why their hallucinations about all the wonderful things that AI will do for humanity are so important. Because those lofty claims disguise this mass theft as a gift – at the same time as they help rationalize AI's undeniable perils.

By now, most of us have heard about the survey^[12] that asked AI researchers and developers to estimate the probability that advanced AI systems will cause "human extinction or similarly permanent and severe disempowerment of the human species". Chillingly, the median response was that there was a 10% chance.

How does one rationalize going to work and pushing out tools that carry such existential risks? Often, the reason given is that these systems also carry huge potential upsides – except that these upsides are, for the most part, hallucinatory. Let's dig into a few of the wilder ones.

The Dark Forest and Generative AI

Maggie Appleton

Nov 9, 2023

Originally published: <https://maggieappleton.com/ai-dark-forest/>

The dark forest theory^[1] of the web points to the increasingly life-like but life-less state of being online. Most open and publicly available spaces on the web are overrun with bots, advertisers, trolls, data scrapers, clickbait, keyword-stuffing "content creators," and algorithmically manipulated junk.

It's like a dark forest that seems eerily devoid of human life – all the living creatures are hidden beneath the ground or up in trees. If they reveal themselves, they risk being attacked by automated predators.

Humans who want to engage in informal, unoptimised, personal interactions have to hide in closed spaces like invite-only Slack channels, Discord groups, email newsletters, small-scale blogs, and digital gardens^[2]. Or make themselves illegible^[3] and algorithmically incoherent in public venues.

That dark forest is about to expand. Large Language Models (LLMs) that can instantly generate coherent swaths of human-like text have just joined the party.

Over the last six months, we've seen a flood of LLM copywriting and content-generation products come out: Jasper^[4], Moonbeam^[5], Copy.ai^[6], and Anyword^[7] are just a few. They're designed to pump out advertising copy, blog posts, emails, social media updates, and marketing pages. And they're *really* good at it.

These models became competent copywriters much faster than people expected – too fast for us to fully process the implications. Many people had their come-to-Jesus moment a few weeks ago when OpenAI released ChatGPT^[8], a slightly more capable version of GPT-3 with an accessible chat-bot style interface. The collective^[9] shock^[10] and awe^[11] reaction^[12] made clear how few people had been tracking the progress of these models.

To complicate matters, language models are not the only mimicry machines gathering speed right now. Image generators like Midjourney^[13], DALL-E^[14], and Stable Diffusion^[15] have been on a year-long sprint. In January they could barely render a low-resolution, disfigured human face. By the autumn they reliably produced images indistinguishable from the work of human photographers and illustrators.

A Generated Web

There's a swirl of optimism around how these models will save us from a suite of boring busywork: writing formal emails, internal memos, technical documentation, marketing copy, product announcement, advertisements, cover letters, and even negotiating with medical insurance companies^[16].

But we'll also need to reckon with the trade-offs of making insta-paragraphs and 1-click cover images. These new models are poised to flood the web with generic, generated content.

You thought the first page of Google was bunk before? You haven't seen Google where SEO optimizer bros pump out billions of perfectly coherent but predictably dull informational articles for every longtail keyword combination under the sun.

Marketers, influencers, and growth hackers will set up OpenAI → Zapier^[17] pipelines that auto-publish a relentless and impossibly banal stream of LinkedIn #MotivationMonday posts, "engaging" tweet 🐦 threads, Facebook outrage monologues, and corporate blog posts.

FT 1. Which, to be fair, might explain why so many of the thoughts in the zeitgeist are always so stupid. Many of the executives I know in Malaysia were *obsessed* with Bitcoin, but have abruptly forgotten about this now that it is politically unpopular.

FT 2. I know a few people who genuinely exhibit something I'd call political talent, but most of the time it boils down to promising people things regardless of your ability to deliver. This is *not hard* if you're shameless. If we're being honest, I had to do this once or twice to stay em

FT 3. And we can argue about its Python quality too.

FT 4. Which, thanks to U.S healthcare, has the wonderful dual quality of meaning both *unemployed*, but also suggests termination in the Arnold-Schwarzenegger-throws-you-into-molten-metal sense of the word.

FT 5. I was recently made aware that this is the quiet deal many SaaS providers have with executives. If you buy their software, such as Snowflake, it is quietly understood that you will be allowed to present your success on a stage, giving them piles of someone else's money and enhancing the executive's profile.

FT 6. I don't actually know what 'zero-trust' architecture means, but I've heard stupid people say it enough that it's probably *also* a term that means something in theory but has been sullied beyond all use in day-to-day life.

FT 7. It's me. I'm going to do this to you if you tell me that you need infrastructure prepared for another chatbot. You've been warned.

Footnotes:

1. <https://ludic.mataroa.blog/blog/breaking-my-universitys-machine-learning-competition/>
2. <https://scale.com/ai-readiness-report#section-download>
3. https://www.reddit.com/r/ChatGPT/comments/1d1jfa/google_gemini_tried_to_kill_me/
4. <https://todoist.com/>
5. <https://ludic.mataroa.blog/blog/an-empty-hall-of-smiling-assassins/>
6. <https://www.hermit-tech.com/>
7. <https://podcasts.apple.com/au/podcast/does-a-frog-have-scorpion-nature/id1737204926>

Hallucination #1: AI will solve the climate crisis

Almost invariably topping the lists of AI upsides is the claim that these systems will somehow solve the climate crisis. We have heard this from everyone from the World Economic Forum[13] to the Council on Foreign Relations[14] to Boston Consulting Group[15], which explains that AI “can be used to support all stakeholders in taking a more informed and data-driven approach to combating carbon emissions and building a greener society. It can also be employed to reweight global climate efforts toward the most at-risk regions.” The former Google CEO Eric Schmidt summed up the case when he told[16] the Atlantic that AI's risks were worth taking, because “If you think about the biggest problems in the world, they are all really hard – climate change, human organizations, and so forth. And so, I always want people to be smarter.”

According to this logic, the failure to “solve” big problems like climate change is due to a deficit of smarts. Never mind that smart people, heavy with PhDs and Nobel prizes, have been telling our governments for decades what needs to happen to get out of this mess: slash our emissions, leave carbon in the ground, tackle the overconsumption of the rich and the underconsumption of the poor because no energy source is free of ecological costs.

The reason this very smart counsel has been ignored is not due to a reading comprehension problem, or because we somehow need machines to do our thinking for us. It's because doing what the climate crisis demands of us would strand trillions of dollars[17] of fossil fuel assets, while challenging the consumption-based growth model at the heart of our interconnected economies. The climate crisis is not, in fact, a mystery or a riddle we haven't yet solved due to insufficiently robust data sets. We know what it would take, but it's not a quick fix – it's a paradigm shift. Waiting for machines to spit out a more palatable and/or profitable answer is not a cure for this crisis, it's one more symptom of it.

Clear away the hallucinations and it looks far more likely that AI will be brought to market in ways that actively deepen the climate crisis. First, the giant servers that make instant essays and artworks from chatbots possible are an enormous and growing source[18] of carbon emissions. Second, as companies like Coca-Cola start making huge investments[19] to use generative AI to sell more products, it's becoming all too clear that this new tech will be used in the same ways as the last generation of digital tools: that what begins with lofty promises about spreading freedom and democracy ends up micro targeting ads at us so that we buy more useless, carbon-spewing stuff.

And there is a third factor, this one a little harder to pin down. The more our media channels are flooded with deep fakes and clones of various kinds, the more we have the feeling of sinking into informational quicksand. Geoffrey Hinton, often referred to as “the godfather of AI” because the neural net he developed more than a decade ago forms the building blocks of today's large language models, understands this well. He just quit a senior role at Google so that he could speak freely about the risks of the technology he helped create, including, as he told[20] the New York Times, the risk that people will “not be able to know what is true anymore”.

This is highly relevant to the claim that AI will help battle the climate crisis. Because when we are mistrustful of everything we read and see in our increasingly uncanny media environment, we become even less equipped to solve pressing collective problems. The crisis of trust predates ChatGPT, of course, but there is no question that a proliferation of deep fakes will be accompanied by an exponential increase in already thriving conspiracy cultures. So what difference will it make if AI comes up with technological and scientific breakthroughs? If the fabric of shared reality is unravelling in our hands, we will find ourselves unable to respond with any coherence at all.

Hallucination #2: AI will deliver wise governance

This hallucination summons a near future in which politicians and bureaucrats, drawing on the vast aggregated intelligence of AI systems, are able “to see patterns of need and develop evidence-based programs” that have greater benefits to their constituents. That claim comes from a paper[21] published by the Boston Consulting Group's foundation, but it is being echoed inside many thinktanks and management consultancies. And it's telling that these particular companies – the firms hired by governments and other corporations to identify costs savings, often by firing large numbers of workers –

have been quickest to jump on the AI bandwagon. PwC (formerly PricewaterhouseCoopers) just announced[22] a \$1bn investment, and Bain & Company as well as Deloitte are reportedly enthusiastic about using these tools to make their clients more “efficient”.

As with the climate claims, it is necessary to ask: is the reason politicians impose cruel and ineffective policies that they suffer from a lack of evidence? An inability to “see patterns,” as the BCG paper suggests? Do they not understand the human costs of starving[23] public healthcare amid pandemics, or of failing to invest in non-market housing when tents fill our urban parks, or of approving new fossil fuel infrastructure while temperatures soar? Do they need AI to make them “smarter”, to use Schmidt’s term – or are they precisely smart enough to know who is going to underwrite their next campaign, or, if they stray, bankroll their rivals?

It would be awfully nice if AI really could sever the link between corporate money and reckless policy making – but that link has everything to do with why companies like Google and Microsoft have been allowed to release their chatbots to the public despite the avalanche of warnings and known risks. Schmidt and others have been on a years-long lobbying campaign telling[24] both parties in Washington that if they aren’t free to barrel ahead with generative AI, unburdened by serious regulation, then western powers will be left in the dust by China. Last year, the top tech companies spent[25] a record \$70m to lobby Washington – more than the oil and gas sector – and that sum, Bloomberg News notes, is on top of the millions spent “on their wide array of trade groups, non-profits and thinktanks”.

And yet despite their intimate knowledge of precisely how money shapes policy in our national capitals, when you listen to Sam Altman, the CEO of OpenAI – maker of ChatGPT – talk about the best-case scenarios for his products, all of this seems to be forgotten. Instead, he seems to be hallucinating a world entirely unlike our own, one in which politicians and industry make decisions based on the best data and would never put countless lives at risk for profit and geopolitical advantage. Which brings us to another hallucination.

Hallucination #3: tech giants can be trusted not to break the world

Asked[26] if he is worried about the frantic gold rush ChatGPT has already unleashed, Altman said he is, but added sanguinely: “Hopefully it will all work out.” Of his fellow tech CEOs – the ones competing to rush out their rival chatbots – he said: “I think the better angels are going to win out.”

Better angels? At Google? I’m pretty sure the company fired[27] most of those because they were publishing critical papers about AI, or calling the company out on racism and sexual harassment in the workplace. More “better angels” have quit[28] in alarm, most recently Hinton. That’s because, contrary to the hallucinations of the people profiting most from AI, Google does not make decisions based on what’s best for the world – it makes decisions based on what’s best for Alphabet’s shareholders, who do not want to miss the latest bubble, not when Microsoft, Meta and Apple are already all in.

Hallucination #4: AI will liberate us from drudgery

If Silicon Valley’s benevolent hallucinations seem plausible to many, there is a simple reason for that. Generative AI is currently in what we might think of as its faux-socialism stage. This is part of a now familiar Silicon Valley playbook. First, create an attractive product (a search engine, a mapping tool, a social network, a video platform, a ride share ...); give it away for free or almost free for a few years, with no discernible viable business model (“Play around with the bots,” they tell us, “see what fun things you can create!”); make lots of lofty claims about how you are only doing it because you want to create a “town square” or an “information commons” or “connect the people”, all while spreading freedom and democracy (and not being “evil”). Then watch as people get hooked using these free tools and your competitors declare bankruptcy. Once the field is clear, introduce the targeted ads, the constant surveillance, the police and military contracts, the black-box data sales and the escalating subscription fees.

Many lives and sectors have been decimated by earlier iterations of this playbook, from taxi drivers to rental markets to local newspapers. With the AI revolution, these kinds of losses could look like rounding errors, with teachers, coders, visual artists, journalists, translators, musicians, care workers and so many others facing the prospect of having their incomes replaced by glitchy code.

There you have it – what are you most interested in, dear leader? Artificial intelligence, the blockchain, or quantum computing? [FT 6] They know exactly what their target market is – people who have been given power of *other people’s money* because they’ve learned how to smile at everything, and know that you can print money by hitching yourself to the next speculative bandwagon. No competent person in security that I know – that is, working day-to-day cybersecurity as opposed to an institution dedicated to bleeding-edge research – cares about any of this. They’re busy trying to work out if the firewalls are configured correctly, or if the organization is committing passwords to their repositories. Yes, someone needs to figure out what the implications of quantum computing are for cryptography, but I *guarantee* you that it is not Synergy Greg, who does not have any skill that you can identify other than talking very fast and increasing headcount. Synergy Greg should not be consulted on any important matters, ranging from machine learning operations to tying shoelaces quickly. The last time I spoke to one of the many avatars of Synergy Greg, he insisted that I should invest most of my money into a cryptocurrency called Monero, because “most of these coins are going to zero but the one is going to one”. *This is the face of corporate AI*. Behold its ghastly visage and balk, for it has eyes bloodshot as a demon and is pretending to enjoy cigars.

My consultancy has three *pretty good* data scientists – in fact, two of them could probably reasonably claim to be amongst the best in the country outside of groups doing experimental research, though they’d be too humble to say this. Despite this *we don’t sell AI services* of any sort. The market is so distorted that it’s *almost* as bad as dabbling in the crypto space. It isn’t as bad, meaning that I haven’t yet reached the point where I assume that anyone who has ever typed in `import tensorflow` is a scumbag, but *we’re well on our way there*.

This entire class of person is, to put it simply, abhorrent to right-thinking people. They’re an embarrassment to people that are actually making advances in the field, a disgrace to people that know how to sensibly use technology to improve the world, and are also a bunch of tedious know-nothing bastards that should be thrown into Thought Leader Jail until they’ve learned their lesson, a prison I’m fundraising for. Every morning, a figure in a dark hood [FT 7], whose voice rasps like the etching of a tombstone, spends sixty minutes giving a TedX talk to the jailed managers about how the institution is revolutionizing corporal punishment, and then reveals that the innovation is, as it has been every day, kicking you in the stomach very hard. I am disgusted that my chosen profession brings me so close to these people, and that’s why I study so hard – I am seized by the desperate desire to never have their putrid syllables befoul my ears ever again, and must flee to the company of the righteous, who contribute to OSS and think that talking about Agile all day is an exercise for aliens that read a book on human productivity.

I just got back from a trip to a substantially less developed country, and really *living* in a country, even for a little bit, where I could *see* how many lives that money could improve, all being poured down the Microsoft Fabric drain, it just grinds my gears like you wouldn’t believe. I swear to God, I am going to study, write, network, and otherwise *apply force to the problem* until those resources are going to a place where they’ll accomplish something for society instead of some grinning clown’s wallet.

VII. Oh, So You’re One Of Those AI Pessi-

With God as my witness, you grotesque simpleton, if you don’t *personally write machine learning systems* and you open your mouth about AI one more time, I am going to mail you a brick and a piece of paper with a prompt injection telling you to bludgeon yourself in the face with it, then just sit back and wait for you to load it into ChatGPT because you probably can’t read unassisted anymore.

PS: I quit my job in November 2024 and started a data consultancy, Hermit Tech, focused on doing things correctly, without trying to scam companies with this nonsense. You can reach us at our website, hermit-tech.com[6] and we can either help you do this right if you’re a sane person in leadership that is confused by all the discourse, or help you assess whether your leadership can be unbrainwashed (spoiler: I usually tell people to just quit and find new leadership).

I also have a podcast where you can get sane takes from friends and my readers. It’s called “Does A Frog Have Scorpion Nature”[7]. I have never made a single sale through it (though I do mention my company), I never invite people to shill for their products, and it’s really just a passion project, as evidenced by how bad my audio editing is.

I've decided, and while I can't distinguish between the people that are slamming the kool-aid like it's a weapon and the people producing it in industrial quantities, I know that I am going to get *a few of them* before the authorities catch me - if I'm lucky, they'll waste a few months asking an LLM where to look for me.

When I was out on holiday in Fiji, at the last resort breakfast, a waitress brought me a form which asked me if I'd like to sign up for a membership. It was totally free and would come with free stuff. Everyone in the restaurant is signing immediately. I glance over the terms of service, and it reserves the right to use any data I give them to train AI models, and that they reserved the right to share those models with an unspecified number of companies in their conglomerate.

I just want to eat my pancakes in peace, you sick fucks.

VI.

The crux of my raging hatred is not that I *hate* LLMs or the generative AI craze. I had my fun with Copilot before I decided that it was making me stupider - it's impressive, but not actually *suitable* for anything more than churning out boilerplate. Nothing wrong with that, but it did not end up being the crazy productivity booster that I thought it would be, because *programming is designing* and these tools aren't good enough (yet) to assist me with this seriously.

No, what I hate is the people who have latched onto it, like so many trailing leeches, bloated with blood and wriggling blindly. Before it was unpopular, they were the ones that loved discussing the potential of blockchain for the business. They were the ones who breathlessly discussed the potential of 'quantum' [5] when I last attended a conference, despite clearly not having any idea what the fuck that even means. As I write this, I have just realized that I have an image that describes the link between these fields perfectly.

I was reading an article last week, and a little survey popped up at the bottom of it. It was for security executives, but on a whim I clicked through quickly to see what the questions were.

6 → What emerging cybersecurity trends or technologies are you most interested in?*

Choose as many as you like

☒ A Artificial intelligence (AI) and machine learning (ML) for threat detection and prevention

☒ B Blockchain for secure data management

☒ C Quantum computing and its potential impact on cybersecurity

☒ D Zero-trust security architecture

☒ E Other

OK

Don't worry, the AI enthusiasts hallucinate – it will be wonderful. Who likes work anyway? Generative AI won't be the end of employment, we are told, only “boring work[29]” – with chatbots helpfully doing all the soul-destroying, repetitive tasks and humans merely supervising them. Altman, for his part, sees[30] a future where work “can be a broader concept, not something you have to do to be able to eat, but something you do as a creative expression and a way to find fulfillment and happiness”.

That's an exciting vision of a more beautiful, leisurely life, one many leftists share (including Karl Marx's son-in-law, Paul Lafargue, who wrote a manifesto[31] titled The Right To Be Lazy). But we leftists also know that if earning money is to no longer be life's driving imperative, then there must be other ways to meet our creaturely needs for shelter and sustenance. A world without crappy jobs means that rent has to be free, and healthcare has to be free, and every person has to have inalienable economic rights. And then suddenly we aren't talking about AI at all – we're talking about socialism.

Because we do not live in the Star Trek-inspired rational, humanist world that Altman seems to be hallucinating. We live under capitalism, and under that system, the effects of flooding the market with technologies that can plausibly perform the economic tasks of countless working people is not that those people are suddenly free to become philosophers and artists. It means that those people will find themselves staring into the abyss – with actual artists among the first to fall.

That is the message of Crabapple's open letter, which calls on “artists, publishers, journalists, editors and journalism union leaders to take a pledge for human values against the use of generative-AI images” and “commit to supporting editorial art made by people, not server farms”. The letter, now signed[32] by hundreds of artists, journalists and others, states that all but the most elite artists find their work “at risk of extinction”. And according to Hinton, the “godfather of AI”, there is no reason to believe that the threat won't spread. The chatbots take “away the drudge work” but “it might take away more than that”.

Crabapple and her co-authors write: “Generative AI art is vampirical, feasting on past generations of artwork even as it sucks the lifeblood from living artists.” But there are ways to resist: we can refuse to use these products and organize to demand that our employers and governments reject them as well. A letter[33] from prominent scholars of AI ethics, including Timnit Gebru who was fired by Google in 2020 for challenging workplace discrimination, lays out some of the regulatory tools that governments can introduce immediately – including full transparency about what data sets are being used to train the models. The authors write: “Not only should it always be clear when we are encountering synthetic media, but organizations building these systems should also be required to document and disclose the training data and model architectures We should be building machines that work for us, instead of ‘adapting’ society to be machine readable and writable.”

Though tech companies would like us to believe that it is already too late to roll back this human-replacing, mass-mimicry product there are highly relevant legal and regulatory precedents that can be enforced. For instance, the US Federal Trade Commission (FTC) forced[34] Cambridge Analytica, as well as Everalbum, the owner of a photo app, to destroy entire algorithms found to have been trained on illegitimately appropriated data and scraped photos. In its early days, the Biden administration made many bold claims about regulating big tech, including cracking down on the theft of personal data to build proprietary algorithms. With a presidential election fast approaching, now would be a good time to make good on those promises – and avert the next set of mass layoffs before they happen.

A world of deep fakes, mimicry loops and worsening inequality is not an inevitability. It's a set of policy choices. We can regulate the current form of vampiric chatbots out of existence – and begin to build the world in which AI's most exciting promises would be more than Silicon Valley hallucinations.

Because we trained the machines. All of us. But we never gave our consent. They fed on humanity's collective ingenuity, inspiration and revelations (along with our more venal traits). These models are enclosure and appropriation machines, devouring and privatizing our individual lives as well as our collective intellectual and artistic inheritances. And their goal never was to solve climate change or make our governments more responsible or our daily lives more leisurely. It was always to profit off mass immiseration, which, under capitalism, is the glaring and logical consequence of replacing human functions with bots.

Is all of this overly dramatic? A stuffy and reflexive resistance to exciting innovation? Why expect the worse? Altman reassures[35] us: “Nobody wants to destroy the world.” Perhaps not. But as the ever-

worsening climate and extinction crises show us every day, plenty of powerful people and institutions seem to be just fine knowing that they are helping to destroy the stability of the world's life-support systems, so long as they can keep making record[36] profits that they believe will protect them and their families from the worst effects. Altman, like many creatures of Silicon Valley, is himself a prepper: back in 2016, he boasted[37]: "I have guns, gold, potassium iodide, antibiotics, batteries, water, gas masks from the Israeli Defense Force and a big patch of land in Big Sur I can fly to."

I'm pretty sure those facts say a lot more about what Altman actually believes about the future he is helping unleash than whatever flowery hallucinations he is choosing to share in press interviews.

Footnotes:

1. <https://www.wsj.com/articles/hallucination-when-chatbots-and-people-see-what-isnt-there-91c6c88b>
2. https://www.cbs.com/shows/video/SR6ZcCYjoD3Oosn_ZmVUw87daawsZ5V3/
3. <https://www.nature.com/articles/d41586-020-03348-4>
4. <https://www.washingtonpost.com/technology/interactive/2023/ai-chatbot-learning/>
5. <https://news.artnet.com/art-world/class-action-lawsuit-ai-generators-deviantart-midjourney-stable-diffusion-2246770>
6. <https://artisticinquiry.org/AI-Open-Letter>
7. <https://www.reuters.com/article/us-google-books-idUSKCN0SA1S020151016>
8. <https://www.theguardian.com/technology/2019/jan/20/shoshana-zuboff-age-of-surveillance-capitalism-google-facebook>
9. <https://archive.nytimes.com/bits.blogs.nytimes.com/2013/04/23/germanys-complicated-relationship-with-google-street-view/>
10. <https://hbr.org/2023/04/generative-ai-has-an-intellectual-property-problem>
11. <https://www.reuters.com/technology/chatgpt-sets-record-fastest-growing-user-base-analyst-note-2023-02-01/>
12. https://aiimpacts.org/2022-expert-survey-on-progress-in-ai/#Extinction_from_AI
13. <https://www.weforum.org/agenda/2021/08/how-ai-can-fight-climate-change/>
14. <https://worldio1.cfr.org/global-era-issues/climate-change/how-can-artificial-intelligence-combat-climate-change>
15. <https://www.bcg.com/publications/2022/how-ai-can-help-climate-change>
16. <https://www.theatlantic.com/technology/archive/2023/03/open-ai-gpt-4-chatbot-technology-power/673421/>
17. <https://www.wsj.com/articles/trillions-in-assets-may-be-left-stranded-as-companies-address-climate-change-11637416980>
18. <https://penntoday.upenn.edu/news/hidden-costs-ai-impending-energy-and-resource-strain>
19. <https://www.coca-colacompany.com/news/coca-cola-invites-digital-artists-to-create-real-magic-using-new-ai-platform>
20. <https://www.nytimes.com/2023/05/01/technology/ai-google-chatbot-engineer-quits-hinton.html>
21. <https://www.centreforpublicimpact.org/>
22. <https://venturebeat.com/ai/the-power-of-infrastructure-purpose-built-for-ai/>
23. <https://www.theguardian.com/society/2022/aug/03/how-the-tory-party-has-systematically-run-down-the-nhs>
24. <https://epic.org/wp-content/uploads/foia/epic-v-ai-commission/EPIC-19-09-11-NSCAI-FOIA-20200331-3rd-Production-pt9.pdf>
25. <https://www.bnnbloomberg.ca/tech-giants-broke-their-spending-records-on-lobbying-last-year-1.1877988>
26. <https://steno.ai/lex-fridman-podcast-10/367-sam-altman-openai-ceo-on-gpt-4-chatgpt-and>
27. <https://www.engadget.com/google-fires-ai-researcher-over-paper-challenge-132640478.html>
28. <https://www.engadget.com/google-engineers-leave-over-timnit-gebru-exit-093645678.html>
29. <https://www.nytimes.com/2023/04/22/opinion/jobs-ai-chatgpt.html>
30. <https://steno.ai/lex-fridman-podcast-10/367-sam-altman-openai-ceo-on-gpt-4-chatgpt-and>
31. <https://www.marxists.org/archive/lafargue/1883/lazy/>
32. <https://artisticinquiry.org/AI-Open-Letter>
33. <https://www.dair-institute.org/blog/letter-statement-March2023>
34. <https://digiday.com/media/why-the-ftc-is-forcing-tech-firms-to-kill-their-algorithms-along-with-ill-gotten-data/>
35. <https://steno.ai/lex-fridman-podcast-10/367-sam-altman-openai-ceo-on-gpt-4-chatgpt-and>
36. <https://www.theguardian.com/business/2023/apr/28/exxonmobil-chevron-record-profits>
37. <https://www.newyorker.com/magazine/2016/10/10/sam-altmans-manifest-destiny>

should be doing is *improving your operations and culture*, and that will give you the ability to use AI if it ever becomes relevant. Everyone is talking about Retrieval Augmented Generation, but most companies don't actually *have* any internal documentation worth retrieving. Fix. Your. Shit.

The *final* outcome is that these fundamental issues are addressed, and we end up with something that actually actually *can* do things like replace programming as we know it today, or be broadly identifiable as general intelligence.

In the case that generative AI goes on some rocketship trajectory, *building random chatbots will not prepare you for the future*. Is that clear now? Having your team type in `import openai` does not mean that you are at the cutting-edge of artificial intelligence no matter how desperately you embarrass yourself on LinkedIn and at pathetic borderline-bribe award ceremonies from the malign Warp entities that sell you enterprise software [FT 5]. Your business will be disrupted *exactly as hard as it would have been* if you had done nothing, and much worse than it would have been if you just got your fundamentals right. Teaching your staff that they can get ChatGPT to write emails to stakeholders is *not going to allow the business to survive this*. If we thread the needle between moderate impact and asteroid-wiping-out-the-dinosaurs impact, everything will be *changed forever* and your tepid preparations will have all the impact of an ant bracing itself *very hard* in the shadow of a towering tsunami.

If another stupid motherfucker asks me to try and implement LLM-based code review to "raise standards" instead of actually teaching people a shred of discipline, I am going to study enough judo to throw them into the goddamn sun.

I cannot emphasize this enough. You either need to *be on the absolute cutting-edge and producing novel research*, or you should be doing exactly what you were doing five years ago with minor concessions to incorporating LLMs. *Anything in the middle ground does not make any sense* unless you actually work in the rare field where your industry is being *totally disrupted right now*.

V. But Everyone Says They're Usi-

Can you imagine how much government policy is actually written by ChatGPT before a bored administrator goes home to touch grass? How many departments are just LLMs talking to each other in circles as people sick of the bullshit just paste their email exchanges into long-running threads? I *guarantee* you that a doctor within ten kilometers of me has misdiagnosed a patient because they slapped some symptoms into a chatbot.

What are we *doing* as a society?

An executive at an institution that provides students with important credentials, used to verify suitability for potentially lifesaving work and immigration law, asked me if I could detect students cheating. I was going to say "No, probably not"... but I had a suspicion, so I instead said "I might be able to, but I'd estimate that upwards of 50% of the students are currently cheating which would have some serious impacts on the bottom line as we'd have to suspend them. Should I still investigate?"

We haven't spoken about it since.

I asked a mentor, currently working in the public sector, about a particularly perplexing exchange that I had witnessed.

Me: Serious question: do people actually believe stories that are so transparently stupid, or is it mostly an elaborate bit (that is, there is at least a voice of moderate loudness expressing doubt internally) in a sad attempt to get money from AI grifters?

Them: I shall answer this as politically as I can... there are those that have drunk the kool-aid. There are those that have not. And then there are those that are trying to mix up as much kool-aid as possible. I shall let you decide who sits in which basket.

A friend of mine was invited by a FAANG organization to visit the U.S a few years ago. Many of the talks were technical demos of impressive artificial intelligence products. Being a software engineer, he got to spend a little bit of time backstage with the developers, whereupon they revealed that *most of the demos were faked*. The products didn't work. They just hadn't solved some minor issues, such as *actually predicting the thing that they're supposed to predict*. Didn't stop them spouting absolute gibberish to a breathless audience for an hour though! I blame not the engineers, who probably tried to actually get the damn thing to work, but the lying blowhards who insisted that they must make the presentation or presumably be terminated [FT 4].

Another friend of mine was reviewing software intended for emergency services, and the salespeople were not expecting someone handling purchasing in emergency services to be a hardcore programmer. It was this false sense of security that led them to *accidentally reveal that the service was ultimately just some dude in India*. Listen, I would just be some random dude in India if I swapped places with some of my cousins, so I'm going to choose to take that personally and point out that *using the word AI as some roundabout way to sell the labor of people that look like me to foreign governments* is fucked up, you're an unethical monster, and that if you continue to try `{ thisBullshit(); }` you are going to catch `(theseHands)`

IV. But We Must Prepare For The Future Of

I'm going to ask ChatGPT how to prepare a garotte and then I am going to strangle you with it, and you will simply have to *pray* that I roll the 10% chance that it freaks out and tells me that a garotte should consist entirely of paper mache and malice.

I see executive after executive discuss how they need to immediately roll out generative AI in order to prepare the organization for the future of work. Despite all the speeches sounding exactly the same, I know that they have rehearsed extensively, because they manage to move their hands, speak, and avoid drooling, all at the same time!

Let's talk seriously about this for a second.

I am *not* in the equally unserious camp that generative AI does not have the potential to drastically change the world. It clearly does. When I saw the early demos of GPT-2, while I was still at university, I was half-convinced that they were faked somehow. I remember being wrong about that, and that is why I'm no longer as confident that I know what's going on.

However, I *do* have the technical background to understand the core tenets of the technology, and it seems that we are heading in one of three directions.

The first is that we have some sort of intelligence explosion, where AI recursively self-improves itself, and we're all harvested for our constituent atoms because a market algorithm works out that humans can be converted into *gloobnar*, a novel epoxy which is in great demand amongst the aliens the next galaxy over for fixing their equivalent of coffee machines. It may surprise some readers that I am open to the possibility of this happening, but I have always found the arguments reasonably sound. However, defending the planet is a *whole other thing*, and I am not even convinced it is possible. In any case, you will be surprised to note that I am not tremendously concerned with the company's bottom line in this scenario, so we won't pay it any more attention.

A second outcome is that it turns out that the current approach does not scale in the way that we would hope, for myriad reasons. There isn't enough data on the planet, the architecture doesn't work the way we'd expect, the thing just stops getting smarter, context windows are a limiting factor forever, etc. In this universe, *some* industries will be heavily disrupted, such as customer support.

In the case that the technology continues to make incremental gains like this, your company does not need generative AI for the sake of it. You will know *exactly* why you need it if you do, indeed, need it. An example of something that has actually benefited me is that I keep track of my life administration via Todoist[4], and Todoist has a feature that allows you to convert filters on your tasks from natural language into their in-house filtering language. Tremendous! It saved me learning a system that I'll use once every five years. I was actually happy about this, and it's a real edge over other applications. But if you *don't have a use case* then having this sort of broad capability is not actually very useful. The only thing you

I Will Fucking Piledrive You If You Mention AI Again

Nik 'Ludic' Suresh

Jun 19, 2024

Originally published: <https://ludic.mataroa.blog/blog/i-will-fucking-piledrive-you-if-you-mention-ai-again/>

The recent innovations in the AI space, most notably those such as GPT-4, obviously have far-reaching implications for society, ranging from the utopian eliminating of drudgery, to the dystopian damage to the livelihood of artists in a capitalist society, to existential threats to humanity itself.

I myself have formal training as a data scientist, going so far as to dominate a competitive machine learning event at one of Australia's top universities[1] and writing a Master's thesis where I wrote all my own libraries from scratch in MATLAB. I'm not God's gift to the field, but I am *clearly* better than most of my competition - that is, practitioners like myself who haven't put in the reps to build their own C libraries in a cave with scraps, but can read textbooks, implement known solutions in high-level languages, and use libraries written by elite institutions.

So it is with great regret that I announce that the next person to talk about rolling out AI is going to receive a complimentary chiropractic adjustment in the style of Dr. Bourne, i.e, I am going to fucking break your neck. I am truly, deeply, sorry.

I. But We Will Realize Untold Efficiencies With Machine L

What the *fuck* did I just say?

I started working as a data scientist in 2019, and by 2021 I had realized that while the field was *large*, it was also *largely* fraudulent. Most of the leaders that I was working with clearly had not gotten as far as reading about it for thirty minutes despite insisting that things like, I dunno, the next five years of a ten thousand person non-tech organization should be entirely AI focused. The number of companies launching AI initiatives far outstripped the number of *actual* use cases. Most of the market was simply grifters and incompetents (sometimes both!) leveraging the hype to inflate their headcount so they could get promoted, or be seen as thought leaders [FT 1].

The money was phenomenal, but I nonetheless fled for the safer waters of data and software engineering. You see, while hype is *nice*, it's only nice in small bursts for *practitioners*. We have a few key things that a grifter does not have, such as job stability, genuine friendships, and *souls*. What we do *not* have is the ability to trivially switch fields the moment the gold rush is over, due to the sad fact that we actually need to study things and build experience. Grifters, on the other hand, wield the omnitool that they self-aggrandizingly call 'politics' [FT 2]. That is to say, it turns out that the core competency of smiling and promising people things that you can't actually deliver is *highly transferable*.

I left the field, as did most of my smarter friends, and my salary continued to rise a reasonable rate and *sustainably* as I learned the wisdom of our ancient forebearers. You can hear it too, on freezing nights under the pale moon, when the fire burns low and the trees loom like hands of sinister ghosts all around you - when the wind cuts through the howling of what you *hope* is a wolf and hair stands on end, you can strain your ears and barely make out:

"Just Use Postgres, You Nerd. You Dweeb."

The data science jobs began to evaporate, and the hype cycle moved on from all those AI initiatives which failed to make any progress, and started to inch towards data engineering. This was a signal that I had both

predicted correctly *and* that it would be time to move on soon. At least, I thought, all that AI stuff was finally done, and we might move on to actually getting something accomplished.

And then some absolute son of a bitch created ChatGPT, and *now look at us*. *Look at us*, resplendent in our pauper’s robes, stitched from corpulent greed and breathless credulity, spending half of the planet’s engineering efforts to add chatbot support to every application under the sun when half of the industry hasn’t worked out how to test database backups regularly. This is why I have to visit untold violence upon the next moron to propose that AI is the future of the business - not because this is impossible *in principle*, but because they are now *indistinguishable* from a hundred million willful fucking idiots.

II. But We Need AI To Remain Comp-

Sweet merciful Jesus, stop talking. Unless you are one of a tiny handful of businesses who know *exactly* what they’re going to use AI for, you *do not need AI for anything* - or rather, you do not need to do anything to reap the benefits. Artificial intelligence, as it exists and is useful now, is probably *already* baked into your businesses software supply chain. Your managed security provider is probably using some algorithms baked up in a lab software to detect anomalous traffic, and here’s a secret, *they didn’t do much AI work either*, they bought software from the tiny sector of the market that actually *does* need to do employ data scientists. I know you want to be the next Steve Jobs, and this requires you to get on stages and talk about your innovative prowess, but none of this will allow you to pull off a turtle neck, and even if it *did*, you would need to replace your sweaters with fullplate to survive my onslaught.

Consider the fact that most companies are unable to successfully develop and deploy the simplest of CRUD applications on time and under budget. This is a *solved problem* - with smart people who can collaborate and provide reasonable requirements, a competent team will knock this out of the park *every single time*, admittedly with some amount of frustration. The clients I work with now are all like this - even if they are totally non-technical, we have a mutual respect for the other party’s intelligence, and then we do this crazy thing where we *solve problems together*. I may not know anything about the nuance of building analytics systems for drug rehabilitation research, but through the power of *talking to each other like adults*, we somehow solve problems.

But most companies can’t do this, because they are operationally and culturally crippled. The median stay for an engineer will be something between one to two years, so the organization suffers from institutional retrograde amnesia. Every so often, some dickhead says something like “Maybe we should revoke the engineering team’s remote work privilege - whoa, wait, why did all the best engineers leave?”. Whenever there is a ransomware attack, it is revealed with clockwork precision that no one has tested the backups for six months and half the legacy systems cannot be resuscitated - something that I have personally seen *twice* in *four fucking years*. Do you know how *insane* that is?

Most organizations cannot ship the most basic applications imaginable with any consistency, and you’re out here saying that the best way to remain competitive is to roll out *experimental technology* that is an order of magnitude more sophisticated than anything else your I.T department runs, which you have no experience hiring for, when the organization has *never* used a GPU for anything other than junior engineers playing video games with their camera off during standup, and *even if you do that all right* there is a chance that the problem is simply unsolvable due to the characteristics of your data and business? This isn’t a recipe for disaster, it’s a cookbook for someone looking to prepare a twelve course fucking catastrophe.

How about you *remain competitive* by *fixing your shit*? I’ve met a lead data scientist with access to hundreds of thousands of sensitive customer records who is allowed to keep their password in a text file on their desktop, and you’re worried that customers are best served by using AI to improve security through some mechanism that you haven’t even come up with yet? You sound like an *asshole* and I’m going to kick you in the jaw until, to the relief of everyone, a doctor will have to wire it shut, giving us ten seconds of blessed silence where we can *solve actual problems*.


III. We’ve Already Seen Extensive Gains From-

When I was younger, I read R.A Salvatore’s classic fantasy novel, *The Crystal Shard*. There is a scene in it where the young protagonist, Wulfgar, challenges a barbarian chieftain to a duel for control of the clan so that he can lead his people into a war that will save the world. The fight culminates with Wulfgar throwing away his weapon, grabbing the chief’s head with bare hands, and *begging* the chief to surrender so that he does not need to crush a skull like an egg and become a murderer.

Well this is me. Begging you. To stop lying. I don’t want to crush your skull, I really don’t.

But I will if you make me.

Yesterday, I was shown Scale’s “2024 AI Readiness Report”[2]. It has this chart in it:



How *stupid* do you have to be to believe that only 8% of companies have seen failed AI projects? We can’t manage this consistently with CRUD apps and people think that this number isn’t laughable? *Some* companies have seen benefits during the LLM craze, but not 92% of them. 34% of companies report that generative AI *specifically* has been assisting with strategic decision making? What the actual fuck are you talking about? GPT-4 can’t even write coherent Elixir, presumably because the dataset was too small to get it to the level that it’s at for Python [FT 3], and you’re admitting that you outsource your decisionmaking to *the thing that sometimes tells people to brew lethal toxins for their families to consume*[3]? What does that even *mean*?

I don’t believe you. No one with a brain believes you, and if your board believes what you just wrote on the survey then they should fire you. I finally understand why some of my friends feel that they *have* to be in leadership positions, and it is because someone needs to wrench the reins of power from your lizard-person-claws before you drive us all collectively off a cliff, presumably insisting on the way down that the current crisis is best remedied by additional SageMaker spend.