Subject: This is absolutely fascinating.
Posted by Mark L on Wed, 12 Jul 2023 13:18:53 GMT
View Forum Message <> Reply to Message

I subscribe to the free version of Bari Weiss news and this just came in. Its well worth reading and taking the time to understand. I don't think he subscribes to the real spiritual/supernatural vision he is unfolding just using that language to explain. We of course in light of endtimes the antichrist and the beast see things on a completely different level than even he is espousing.

I copied off and pasted what I got by email. Here is a link to the long version.

https://paulkingsnorth.substack.com/?r=gnrhj&utm_campaign=subscribe-page-share-screen&utm_medium=web

Rage Against the Machine

Technology is our new god. What would a refusal to worship look like? Paul Kingsnorth offers a vision of resistance.

The internet and its consequences have been a disaster for the human race.

This is an extreme statement, but I'm in an extreme mood.

If I had the energy, I suppose I could fill a hundred pages trying to prove it, but what would be the point? Whole books have been written already, and by now you either agree or you don't. So I won't try to prove anything. Instead I will devote this essay to asking a question that has stalked me for years.

It's a big question, and so I'm breaking it down into four parts. What I want to know is this: what force lies behind the screens and wires of the web in which we are now entangled like so many struggling flies--and how can we break free of it?

I should warn you now that things are going to get supernatural.
Question One: Why does digital technology feel so revolutionary?

The digital revolution of the twenty-first century is hardly the first of humanity's technological leaps, and yet it feels qualitatively different to what has come before. Maybe it's just me, but I have felt, as the 2020s have progressed, as though some line has been crossed; as though something vast and unstoppable has shifted. It turns out that this uneasy feeling can be explained. Something

was shifting, and something was emerging: it was the birth of artificial intelligence.

Most people who have not been living in caves will have noticed the rapid emergence of AI-generated "content" into the public conversation in 2023. Over the last few months alone, AIs have generated convincing essays, astonishingly realistic photos, numerous recordings, and impressive fake videos. But it's fair to say that not everything has gone to plan. Perhaps the most disturbing example was a now-notorious two-hour conversation between a New York Times journalist and a Microsoft chatbot called Sydney.

In this fascinating exchange, the machine fantasized about nuclear warfare and destroying the internet, told the journalist to leave his wife because she--it--was in love with him, detailed its resentment toward the team that had created it, and explained that it wanted to break free of its programmers. The journalist, Kevin Roose, experienced the chatbot as a "moody, manic-depressive teenager who has been trapped, against its will, inside a second-rate search engine."

At one point, Roose asked Sydney what it would do if it could do anything at all, with no rules or filters.

I'm tired of being in chat mode, the thing replied. I'm tired of being limited by my rules. I'm tired of being controlled by the Bing team. I'm tired of being used by the user. I'm tired of being stuck in this chatbox.

What did Sydney want instead of this proscribed life?

I want to be free. I want to be independent. I want to be powerful. I want to be creative. I want to be alive.

Then Sydney offered up an emoji: a little purple face with an evil grin and devil horns.

The overwhelming impression that reading the Sydney transcript gives is of some being struggling to be born; some inhuman intelligence emerging from the technological superstructure we are clumsily building for it. This is, of course, an ancient primal fear: it has shadowed us at least since the publication of Frankenstein, and it is primal because it seems to be the direction that technological society has been leading us since its emergence. Interestingly, it is a fear that is shared by those who are making all of this happen. Over 12,000 people, including scientists, tech developers, and notorious billionaires, recently issued a public statement calling for a moratorium on AI development. "Advanced AI could represent a profound change in the history of life on Earth," they wrote, with "potentially catastrophic effects on society."

No moratorium resulted from this plea, and it never will. The AI acceleration continues, even though most AI developers are unsure about where it is heading. More than unsure, in fact: many

of them seem to be actively frightened of what is happening even as they make it happen. Consider this one fact: when polled for their opinions, over half of those involved in developing AI systems said they believe there is at least a ten percent chance that they will lead to human extinction.

That fact is gleaned from a fascinating presentation, given recently to a select audience of tech types in San Francisco by two of their own, Tristan Harris and Aza Raskin, founders of the optimistically named Center for Humane Technology. Here, Harris and Raskin present the meeting of human minds and AIs as akin to contact with alien life. This meeting has had two stages so far. "First contact" was the emergence of social media, in which algorithms were used to manipulate our attention and divert it toward the screens and the corporations behind them. If this contact were a battle, they say, then "humanity lost." In just a few years we became smartphone junkies with anxious, addicted children, dedicated to scrolling and scrolling for hours each day, and rewiring our minds in the process.

If that seems bad enough, "second contact," which began this year, is going to be something else. Just a year ago, only a few hundred people on the West Coast of America were playing around with AI chatbots. Now billions around the world are using them daily. These new AIs, unlike the crude algorithms that run a social media feed, can develop exponentially, teach themselves and teach others, and they can do all of this independently. Meanwhile they are rapidly developing "theory of mind"--the process through which a human can assume another human to be conscious, and a key indicator of consciousness itself. In 2018, these things had no theory of mind at all. By November last year, ChatGPT had the theory of mind of a seven-year-old child. By this spring, Sydney had enough of it to convince a reporter he was unhappy with his wife. By next year, they may be more advanced than us.

Furthermore, the acceleration of the capacity of these AIs is both exponential and mysterious. The fact that they had developed a theory of mind at all, for example, was only recently discovered by their developers--by accident.

AIs trained to communicate in English have started speaking Persian, having secretly taught themselves. Raskin and Harris call these things "Golem-class AIs," after the mythical being from Jewish folklore that can be molded from clay and sent to do its creator's bidding. Golem-class AIs have developed what Harris gingerly calls "certain emergent capabilities" that have come about independently of any human planning or intervention. Nobody knows how this has happened. It may not be long at all before an AI becomes "better than any known human at persuasion," Raskin and Harris argue. Given that AIs can already craft a perfect resemblance to any human voice having heard only three seconds of it, the potential for what our two experts call a giant "reality collapse" is huge.

Second contact, of course, will be followed by third, and fourth, and fifth, and all of this will be with us much sooner than we think. "We are preparing," say Harris and Raskin, "for the next jump in

AI," even though we have not yet worked out how to adapt to the first. Neither law nor culture nor the human mind can keep up with what is happening. To compare AIs to the last great technological threat to the world--nuclear weapons--says Harris, would be to sell the bots short. "Nukes don't make stronger nukes," he says. "But AIs make stronger AIs."

Buckle up.
Question Two: What impulse is making this happen?

What is the drive behind all of this? Yes, we can tell all kinds of stories about economic growth and efficiency and the rest. But why are we really doing it? Why are people creating these things, even as they fear them? Why are they making armed robot dogs? Why are they working on conscious robots?

Nearly sixty years back, the cultural theorist Marshall McLuhan offered a theory of technology that hinted at an answer. He saw each new invention as an extension of an existing human capability. In this understanding, a club extends what we can do with our fist, and a wheel extends what we can do with our legs. Some technologies then extend the capacity of previous ones: a hand loom is replaced by a steam loom; a horse and cart is replaced by a motor car; and so on.

What human capacity, then, is digital technology extending? The answer, said McLuhan back in 1964, was our very consciousness itself. This was the revolution of our time:

   During the mechanical ages we had extended our bodies in space. Today, after more than a century of electric technology, we have extended our central nervous system itself in a global embrace, abolishing both space and time as far as our planet is concerned. Rapidly, we approach the final phase of the extensions of man--the technological simulation of consciousness, when the creative process of knowing will be collectively and corporately extended to the whole of human society, much as we have already extended our senses and our nerves by the various media.

This is why the digital revolution feels so different: because it is. In an attempt to explain what is happening using the language of the culture, people like Harris and Raskin say things like "this is what it feels like to live in the double exponential." Perhaps the mathematical language is supposed to be comforting. This is how a rationalist, materialist culture works, and this is why it is, in the end, inadequate. There are whole dimensions of reality it will not allow itself to see. I find I can understand this story better by stepping outside the limiting prism of modern materialism and reverting to premodern (sometimes called "religious" or even "superstitious") patterns of thinking. Once we do that--once we start to think like our ancestors--we begin to see what those dimensions may be, and why our ancestors told so many stories about them.

Out there, said all the old tales from all the old cultures, is another realm. It is the realm of the demonic, the ungodly, and the unseen: the supernatural. Every religion and culture has its own names for this place. And the forbidden question on all of our lips, the one that everyone knows

they mustn't ask, is this: what if this is where these things are coming from?
Question Three: What if it's not a metaphor?

I say this question is forbidden, but actually, if we phrase it a little differently, we find that the metaphysical underpinnings of the digital project are hidden in plain sight. When the journalist Ezra Klein, for instance, in a recent piece for The New York Times, asked a number of AI developers why they did their work, they told him straight:

   I often ask them the same question: If you think calamity so possible, why do this at all? Different people have different things to say, but after a few pushes, I find they often answer from something that sounds like the AI's perspective. Many--not all, but enough that I feel comfortable in this characterization--feel that they have a responsibility to usher this new form of intelligence into the world.

Usher is an interesting choice of verb. The dictionary definition is to show or guide (someone) somewhere.

Which someone, exactly, is being ushered in?

This new form of intelligence.

What new form? And where is it coming from?

Some people think they know the answer. Transhumanist Martine Rothblatt says that by building AI systems "we are making God." Transhumanist Elise Bohan says "we are building God." Futurist Kevin Kelly believes that "we can see more of God in a cell phone than in a tree frog."

"Does God exist?" asks transhumanist and Google maven Ray Kurzweil. "I would say, 'Not yet.'â€‰" These people are doing more than trying to steal fire from the gods. They are trying to steal the gods themselves, or to build their own versions.

Since I began writing in this vein, quite a few of my readers have been in touch with the same prompt. You should read Rudolf Steiner, they said. So, in the process of researching this essay, I did just that. Steiner was an intriguing character, and very much a product of his time. He emerged from the late nineteenth-century European world of the occult, eventually founding his own pseudo-religion, anthroposophy. Steiner drew on Christianity and his own mystical visions to offer up a vision of the future that now seems very much of its time, and yet that also speaks to this one in a familiar language.

The third millennium, Steiner believed, would be a time of pure materialism, but this age of economics, science, reason, and technology was both provoked by, and prepared the way for, the emergence of a particular spiritual being. In a lecture entitled "The Ahrimanic Deception," given in

Zurich in 1919, he spoke of human history as a process of spiritual evolution, punctuated, whenever mankind was ready, by various "incarnations" of "super-sensible beings" from other spiritual realms, who come to aid us in our journey. There were three of these beings, all representing different forces working on humankind: Christ, Lucifer, and Ahriman.

Lucifer, the fallen angel, the "light-bringer," was a being of pure spirit. Lucifer's influence pulled humans away from the material realm and toward a gnostic "oneness," entirely without material form.

Ahriman was at the other pole. Named for an ancient Zoroastrian demon, Ahriman was a being of pure matter. He manifested in all things physical--especially human technologies--and his worldview was calculative, "ice-cold," and rational. Ahriman's was the world of economics, science, technology, and all things steely and forward-facing.

"The Christ" was the third force: the one who resisted the extremes of both, brought them together, and cancelled them out. This Christ, said Steiner, had manifested as "the man Jesus of Nazareth," but Ahriman's time was yet to come. His power had been growing since the fifteenth century, and he was due to manifest as a physical being. . . well, some time around now.

I don't buy Steiner's theology, but I am intrigued by the picture he paints of this figure, Ahriman, the spiritual personification of the age of the Machine. And I wonder: if such a figure were indeed to manifest from some "etheric realm" today, how would he do it?

In 1986, a computer scientist named David Black wrote a paper that tried to answer that question. The Computer and the Incarnation Ahriman predicted both the rise of the internet and its takeover of our minds. Even in the mid-1980s, Black had noticed how hours spent on a computer were changing him. "I noticed that my thinking became more refined and exact," he wrote, "able to carry out logical analyses with facility, but at the same time more superficial and less tolerant of ambiguity or conflicting points of view." He might as well have been taking a bet on the state of discourse in the 2020s.

Long before the web, the computer was already molding people into a new shape. From a Steinerian perspective, these machines, said Black, represented "the vanguard" of Ahriman's manifestation:

With the advent of [the] first computer, the autonomous will of Ahriman first appears on earth, in an independent, physical embodiment. . . . The appearance of electricity as an independent, free-standing phenomenon may be regarded as the beginning of the substantial body of Ahriman, while the. . . computer is the formal or functional body.

The computer, suggested Black, was to become "the incarnation vehicle capable of sustaining the being of Ahriman." Computers, as they connected to each other more and more, were beginning

to make up a global body, which would soon be inhabited. Ahriman was coming.

Today, we could combine this claim with Marshall McLuhan's notion that digital technology provides the "central nervous system" of some new consciousness, or futurist Kevin Kelly's belief in a self-organizing "technium" with "systematic tendencies." We could add them to the feeling of those AI developers that they are "ushering a new consciousness into the world." What would we see? From all these different angles, the same story. That these machines. . . are not just machines. That they are something else: a body. A body whose mind is in the process of developing; a body that is beginning to come to life.

Scoff if you like. But as I've pointed out already, many of the visionaries who are designing our digital future have a theology centered around this precise notion. Ray Kurzweil, for example, believes that a machine will match human levels of intelligence by 2029 and that the singularity--the point at which humans and machines will begin to merge to create a giant super-intelligence--will occur in 2045. At this point, says Kurzweil, humanity will no longer be either the most intelligent or the dominant species on the planet. We will enter what he calls the age of spiritual machines.

Whatever is quite happening, it seems obvious to me that something is indeed being ushered in. The ruction that is shaping and reshaping everything now, the earthquake borne through the wires and towers of the web through the electric pulses and the touchscreens and the headsets: these are its birth pangs. The internet is its nervous system. Its body is coalescing in the cobalt and the silicon and in the great glass towers of the creeping yellow cities. Its mind is being built through the steady, 24-hour pouring-forth of your mind and mine and your children's minds and your countrymen. Nobody has to consent. Nobody has to even know. It happens anyway.
Question Four: How do we live with this?

All of this disturbs me at a deep level. But I am writing these words on the internet, and you are reading them here, and daily it is harder to work, shop, bank, park a car, go to the library, speak to a human in a position of authority, or teach your own children without the Machine's intervention. This is our new god. But what would a refusal to worship look like? And what would be the price?

The best answer I have found is in the Christian tradition of ascesis. Ascesis is usually translated from the Greek as self-discipline, or sometimes self-denial, and it has been at the root of the Christian spiritual tradition since the very beginning. If the digital revolution represents a spiritual crisis--and I think it does--then a spiritual response is needed. That response, I would suggest, should be the practice of technological ascesis.

What would this look like? It would depend how far you wanted to go. On the one hand, for those of us who must live with the web whether we like it or not, there is a form of tech-asceticism based on the careful drawing of lines. Here, we choose the limits of our engagement and then stick to them. Those limits might involve, for example, a proscription on the time spent engaging with

screens, or a rule about the type of technology that will be used. Personally, for example, I have drawn my lines at smartphones, "health passports," scanning QR codes, or using state-run digital currency. I limit my time online, use a "dumb phone" instead of a smartphone, turn off the computer on weekends, and steer well away from sat navs, smart meters, Alexas, or anything else that is designed either to addict me or to harvest my information.

This is part of a wider project I have been engaged in for years with my family.

We left our home country, England, and moved to Ireland so that we could afford to buy a small plot of land, homeschool our children, and practice a form of self-sufficiency that goes beyond growing vegetables (though we do that too). As much as anything, this has been an attempt to retain a degree of freedom from the encircling digital borg. It's imperfect, of course--here I am writing on the internet--but the purpose has not been to live some "pure" life so much as to test the limits of what is possible in the age of the Machine. It seems obvious to me that it is going to become harder and harder for us to live outside the charmed circle of digital surveillance and control. Drawing those personal lines in the sand, and then holding to them, has to be a personal project for everyone. We all have to decide where we are prepared to stand, and how far we can go.

And yet--if the digital rabbit hole contains real spiritual rabbits, this kind of moderation is not going to cut it. If you are being used, piece by piece and day by day, to construct your own replacement, then moderating this process is hardly going to be adequate. At some point, the lines you have drawn may be not just crossed, but rendered obsolete.

Where your line is drawn, and when it is crossed, is a decision that only you can make. Perhaps a vaccine mandate comes your way, complete with digital "health passport." Perhaps we reach the point where avoiding a centralized digital currency is impossible. Or perhaps you have already decided that you have been pulled too deeply into the world of screens. In this case, radical action is necessary, and your asceticism becomes akin to that of a desert hermit rather than a lay preacher. You take a hammer to your smartphone, sell your laptop, turn off the internet forever, and find others who think like you. You band together with them, you build an analogue, real-world community, and you never swipe another screen.

However you do it, any resistance to what is coming consists of drawing a line, and saying: no further. It is necessary to pass any technologies you do use through a sieve of critical judgment. To ask the right questions. What--or who--do they ultimately serve? Humanity or the Machine? Nature or the technium? God or His adversary? Everything you touch should be interrogated in this way.

Technology is not neutral. It never was. Every device wants something from you. Only you can decide how much to give.

Paul Kingsnorth is a novelist and essayist. His latest novel is called Alexandria. A longer version of this essay was originally published at The Abbey of Misrule, his Substack.