So this episode is a column read, not a conversation.
But one reason I wanted to talk about this column is that it's a bit of a culmination of things I've explored on the show through a bunch of conversations.
So the Marian Wolf conversation about the way different kinds of reading act on the mind, Newport on the ways in which we have built digital work environments that distract people more than it focuses them. And then obviously a lot of the AI work we've been doing where I think if you know if you've been listening that I am both pretty convinced the technology is transformative and can insinuate in all dimensions of our lives and could be very, very powerful. And also what business models it ends up attached to the actual way we design the environments and ways human beings interact with it is really going to matter.
So it gives me a chance to revisit something that I've thought a lot about with the Internet itself, which is where this one begins.
So imagine I told you in 1970 that I was going to invent this wondrous tool and this new tool would make it possible for anyone with access and most of humanity would amazingly have access to quickly communicate and collaborate with anyone else. It would store nearly the entire sum of human knowledge and thought up to that point. And all of it, all of it would be searchable and sortable and portable. Text could be instantly translated from one language to another. News would be immediately available from all over the world. And it would take no longer for scientists to download a journal paper from 15 years ago than to flip to an entry in the latest issue. If I had told you all that, what would you have predicted that this leap in information and communication and collaboration would do for humanity? And to be really specific, how much faster would our economies grow? How much more productive would we be with all these new capabilities and all this new information?
Now go back. Now imagine I told you that I was going to invent this sinister tool. Maybe I'm cackling while I tell it to you. And this tool, as people used it, their attention spans would degrade because a tool would constantly shift their focus. It would weaken their powers of concentration and of contemplation. This tool would show people whatever it was they found most difficult to look away from. And that would often be what was most threatening about the world in which they live from the worst ideas of their political opponents to the deep injustice of their society. It would make it harder through that to cooperate with each other.
It would fit in their pockets amazingly and it would glow on their nightstands. And as such, it would never be away from them really and never be truly quiet. There would, for a lot of people, never be a moment when they could be free of the sense the pile of messages and warnings and tasks needed to be checked and responded to. So now what would you have thought that this engine, this tool of distraction, of division, of cognitive fracture, what would that have done to humanity? What would that have done to our productivity? Thinking about the Internet, and I'm obviously describing the Internet, thinking about it in these terms, I think helps solve a bit of an economic mystery. The embarrassing truth is that productivity growth, how much more we can make with the same number of people and factories and land, it was much faster for much of the 20th century than it is now. We average about half the productivity growth rate today that we saw in the 1950s and 1960s. And that means stagnating incomes, it means sluggish economies, it means a political culture that is more about fighting over what we already have than spreading the riches and wonders we're

gaining. So what went wrong? You can think of two ways the Internet could have sped up productivity
growth. And the first way was obvious. It would and it did allow us to do what we were already doing and do it more easily and quickly. And that happened. You can see a bump in productivity numbers from roughly 95 to 2005 as companies digitized their operations. They used Excel spreadsheets and emailed each other and served customers online. All that actually did increase productivity. But then there was a second way the Internet could have increased productivity. And this one was always more important. By connecting humanity to itself and to nearly its entire storehouse of information, the Internet could have, should have, made us smarter and more
capable as a collective. It should have increased the quality of ideas humanity could come up with. And I don't think that that promise proved false exactly. Even when I was working on this piece, it was true for me. The speed with which I could find information and sort through research and contact experts, all that was marvelous. And even with all that, I don't think I wrote this faster than I would have if I was writing something similar in 1970. So much of my mind was preoccupied by the
constant effort needed just to hold a train of thought in a digital environment that is designed to distract and agitate and entertain me. There is this addition of productivity and then this attraction of focus. And it's really not clear to me looking at the numbers, which is bigger. And I am in this way definitely not alone. While I was working on this piece, I called Gloria Mark, who is a professor of information science at the University of California at Irvine, and the author of this book, Attention Span. And she's telling me that she started researching the way people use computers back in 2004. And she would follow them around with a stopwatch. And back then, the average time people spent on a single screen was 2.5 minutes. And she said to me about that, that she was astounded. That was so much worse than she thought it would be. But that turned out just to be the beginning. They kept doing this research. They moved it away from stopwatches and began actually using computer software that could see when you
changed a window. By 2012, Mark and her colleagues found the average time on a single task. It was
only 75 seconds down from 2.5 minutes. Now it's down to about 47 seconds on average. So half or less than that. This is an acid bath for human cognition. Multitasking is mostly a myth. We can really just focus on one thing at one time. Mark has this great analogy. She said to me, quote, it's like we have an internal whiteboard in our minds. If I'm working on one task, I have all the info I need on that mental whiteboard. Then I switch to email. I have to mentally erase that whiteboard and write all the information I need to do email. And just like on a real whiteboard, there can be a residue in our minds. We may still be thinking of something from three tasks go, end quote. The cost that carries is in more than just performance. So Mark and others in her field have hooked people to blood pressure machines and heart rate monitors and they measured chemicals
in the blood and the constant switching of tasks. It makes us stressed and irritable. And this is one of those findings. So when I heard it, I didn't exactly feel I needed to know it was experimentally confirmed. I feel like I live it constantly and maybe you do too. But it was depressing to hear it confirmed. And that brings me to artificial intelligence. And I think it's

important here to be specific about what I'm talking about. I'm talking here about the systems we're seeing now. So large language models like OpenAI's GPT-4 and Google's Bard. What these systems do for the most part is summarize information they've been shown and create content that resembles it. I know that sentence can sound a bit dismissive, but it shouldn't. That's a remarkable capability. And it's a huge amount of what human beings actually do in their day-to-day lives. And so already we're being told that in doing this AI is making coders and customer service representatives and writers more productive. There are studies and observations on all of these.

I've read about chief executives who plan to add use of chat GPT into employee performance evaluations on the theory that if you're not using chat GPT enough or something like it enough, you're not being nearly as productive as you could be. And you've heard things like this in the internet too, particularly in the early days. And I want to say right now, I am skeptical that this early hype and these early productivity boost people are seeing in experiments is going to come true. And one reason I'm skeptical here is we're measuring as potential benefits without considering its likely cost, which is exactly the mistake we made with the internet. We were really good at imagining all the things it could do to make us productive. And we didn't see the cost it would carry on our own cognition. And I could see that happening with AI in at least three ways.

One way is that these systems are going to do more to distract and entertain us than to focus us. So a huge problem in the current crop of large language models is they hallucinate information. You ask them to answer a complex question and you get this convincing erudite response with citations. And then it just turns out the key facts and key footnotes are completely made up. And I think this is going to slow their widespread use in important industries a lot more than is currently being admitted. This is a lot more like the way driverless cars have had trouble rolling out because they need to be perfectly reliable rather than just pretty good. They can't just usually not hit a pedestrian. So a question to ask about large language models is where does being trustworthy not matter that much? Answer that and I think you've found the areas where adoption is going to be really fast. So an example from my industry from media is telling here. CNET, which is a technology website, it began using these models to write articles with humans in theory editing the pieces. But the process completely failed. When this came out, they had to take a closer look at the articles and it turned out that 41 of the 77 AI generated human edited articles proved to have errors that the editors missed. And so CNET embarrassed had to pause this program.

On the other side, Buzzfeed, which recently shuttered its news division, is racing ahead with using AI to generate quizzes and travel guides and all kinds of Buzzfeed content. And a lot of the results have been shoddy and people are laughing at them, but it doesn't really matter because a Buzzfeed quiz doesn't have to be reliable. That's not the point. So this is an example to me in media of how AI is going to work better, where you have to entertain, we're making things up and being creative might even be an asset, but where factuality and trustworthiness and reliability are central, you're not really going to be able to use it, at least not for some time and not centrally. And if you do use it, you're going to have to spend a lot of money overseeing and fact checking and editing it. So now generalize that idea. AI is going to be great for making personalized video games and children's shows and music mashups and bespoke images are going to be dazzling. And I think we're going to have really new domains of entertainment and delight. I've said this before,

but I believe we're much closer to AI friends and lovers and companions becoming a widespread part

of
our social lives. But yeah, where reliability is going to matter, like having a large language model devoted to answering medical questions or summarizing doctor-patient interactions, deployment is going to be a lot harder because oversight costs are going to be immense. Problem is, those are the areas that matter most, I think, for economic growth.
So then I want to get here to my second worry and to go back to Buzzfeed. Marcella Martin, Buzzfeed's president, has a line that is meant to be positive about AI, but it actually gets to something I think is very likely to be negative. So she told investors, quote,
instead of generating 10 ideas in a minute, AI can generate hundreds of ideas in a second, end quote. Now she meant that as a good thing, but is it? Imagine that multiplied across the economy. Someone somewhere will have to process all that information. What does that do to productivity? One lesson of the digital age is that more is not always better. More emails and more reports and more slacks and more tweets and more videos and more news articles and more slide decks
and more Zoom calls have not led, it seems, to more great ideas. Gloria Mark told me, quote, we can produce more information, but that means there's more information for us to process. Our processing capability is the bottleneck, end quote. Email and chat systems like Slack, I think are a useful analogy here. Both are widely used across the economy. Both were initially sold as productivity boosters, allowing a lot more communication to take place a lot faster. And as anyone who uses them a lot knows, the productivity gains, they're real. You really can talk to people quicker on email, but they're matched, maybe more than matched, by the cost of being buried under vastly more communication, much of it junk and nonsense. The magic of a large language model is that it can produce a document of almost any length and almost any style with a minimum of user effort. And I don't think people really thought through the costs that can impose on those who need to respond to all this new text. One of my favorite examples of this comes from the Economist, which imagine nimbies, but really you can just pick your interest group using GPT-4 to rapidly produce a thousand page complaint opposing a new development.
Someone somewhere in some agency has to respond to that complaint. Will that really speed up our ability to build housing? And you can counter that, okay, sure, but AI is going to solve this problem by quickly summarizing complaints for overwhelmed policy makers, much as the increase in spam is sometimes somewhat countered by more advanced spam filters. But I was talking to Jonathan Frankel, who's a chief scientist at Mosaic ML and a computer scientist at Harvard, and he had this funny line where he said that this is quote, the boring apocalypse scenario for AI in which we, and this is him talking, use chat GPT to generate long emails and documents. And then the person who received it uses chat GPT to summarize it back down to a few bullet points. And there's tons of information changing hands, but all of it is just fluff. We're just inflating and compressing content generated by AI, end quote. When we spoke, Frankel noted how remarkable
it is to feed 100 page Supreme Court document into a large language model and then to get this quite smart summary of the key points. The question he said is, is that a good summary and how do we
know? You can say something similar and many of us have had this experience about asking chat GPT

to draft a piece of writing and seeing a fully formed composition appear as if by magic in seconds. But that gets to my third concern here. Even if those summaries and drafts are pretty good, let's say they're really good. Something is lost in that outsourcing. Part of my job is reading 100 page Supreme Court documents fairly often and it's constantly composing crummy, difficult first drafts
of columns. And yet it would be faster for me to have AI do that work. But the increased efficiency would come at a very clear cost of new ideas and deeper insights. This is a view I hold pretty strongly nowadays. Our society wide obsession with speed and efficiency has given us a flawed model of human cognition. I've come to think of it and I think I've talked about it on the show as the matrix theory of knowledge. We wish we could use that little jack from the matrix to download the knowledge of a book or I guess use a movie's example, a kung fu master into our heads and then we'd have it in a second, right? Boom, I know kung fu. And that misses what's really happening when we spend nine hours reading a biography. It's the time inside the book that we spend drawing connections to what we know and having thoughts we would not otherwise have had
that matters. Gloria Mark said to me that, quote, nobody likes to write reports or do emails, but we want to stay in touch with information. We learn when we deeply process information. If we're removed from that and we're delegating everything to GPT, having it summarized and write reports for us, we're not connecting to that information, end quote. What's interesting to me is we completely understand this when talking about students. Nobody thinks that reading the Spark Notes summary of a great piece of literature is like reading the book. No one thinks that if students have chat GPT, write their essays, they've cleverly boosted their productivity rather than lost the opportunity to learn and work through information and have new insights and get better themselves at thinking through things in essay form. And I don't want to say that's a perfect analogy to office work. There are a lot of dull tasks that are worth automating so people can spend their time on something more creative, but the dangers of over-automating cognitive and creative processes, those are very real. And look, these are old concerns. Socrates questioned the use of writing. He was recorded ironically in writing by Plato, worrying that, quote,
if men learn this, it will implant forgetfulness in their souls. They will cease to exercise memory because they rely on that which is written, calling things to remembrance no longer from within themselves but by means of external marks. Look, I'm a writer. I think the trade-off here was worth it, but it was a trade-off. Human beings really did lose the faculties of memories we once had. Think of the way people had memorized these epic poems. We got better at some forms of thinking
and writing and we lost other forms of cognition. There are trade-offs and not all of them are good. So this then, for now, I think is a task of not just artificial intelligence, but the humans creating it. I know there's a dream that one day we're going to have these AIs
that innovate on their own, and maybe we will. But for now, artificial intelligence needs to deepen human intelligence. And that means human beings need to build AI and build the workflows and office environments around it in ways that don't overwhelm and distract and diminish us. We need to build AI for human beings. I think we failed that test pretty badly with the internet. I really hope we don't fail out with AI.
you