

Hello, everyone. Well, I am back to my usual practice of livestreaming.

So today, I was going to do a Q&A about history of science and technology, but I see a bunch of questions about me.

That are saved up here, so maybe I'll try and address a few of those.

There's a question here from Curiosity. What's an old habit or routine from decades ago that you still have today?

Well, several, actually. I mean, one that's kind of quaint is that I tend to write my to-do lists on pieces of paper, actually. I usually write them on pads of paper, because they're convenient to kind of... sort of move around, and by the time I've torn the piece of paper off, it sort of you know, gets lost, but so I have typically a few pads of paper with their front page, so to speak, with a to-do list on them, and I tend to, write

to-do lists,

one thing per line, and I have a whole bunch of codes of, you know, once I've really done something, I like, you know, have a squiggly thing that goes through it that, very satisfyingly says, I've checked that off.

And then when something's sort of partially done, I have kind of a little, slight, sort of...

hatching that I do to indicate that it's kind of partly done, and then usually these pieces of paper get, things get checked off.

And then after a while, there's only a few left, and then I copy those over onto another piece of paper. I think it sort of forces me to not have too much on the to-do list. It's too easy to just copy, you know, large amounts of stuff onto a new to-do list. I do also have online ones, but more from notebooks and so on, but not so many of those. I prefer the paper ones that I can just see for myself.

Another habit, also related to paper, is that I always carry around a folded piece of paper in my pocket. I've been doing that for... oh gosh...

certainly more than 50 years, and then...

you know, I'll write notes on that.

And if I'm out and about at some event, I'll write some notes on that piece of paper. When I come back, I'll try and

transcribe the relevant ones into actual things, do actual things based on that. I suppose I've had a number of other... I don't know, this is probably getting into... too deep into the weeds, but I've always had various conventions, like I... I have LU for look up, and then something where somebody has told me about something I might look at, and then it's kind of a reminder to go... go ahead and do that.

Oh gosh, what other things,

Oh, lots of habits about... I'm very habitual, and go to sleep at the same time every day, and get up at the same time, and eat many of the same things, and so on. Let's think about things from... decades ago, I suppose,

Oh, here's a funny one. The, You know, I...

collect those magazines that I'll read in random times, you know, in paper form, and I realized that the one magazine I've been continuously subscribing to since I was...

8 years old, is a British magazine called New Scientist, which I'm not sure I would say was necessarily the most accurate source of information, but it's always kind of entertaining, and it has

things that, kind of encourage one to think about some topic, even though the article itself may be a thing where you kind of roll your eyes and say, say, I don't think, that's,

You know, that can't possibly be right, but at least the question is interesting. occasionally they've written about me, I don't necessarily... I can't necessarily even recognize it when I see an article that turns out to be about something I'm doing. It's sufficiently transformed often that it's not, I don't... wouldn't notice that until I... until I happen to see a name or something like this.

other...

kinds of habits. I suppose I... I've had habits that I... I stopped having, which was a shame. Like, when I used to buy books back in the...

1960s and 1970s, I would write my name in the book, together with the date when I bought it. That's been very useful to me in reconstructing some historical kinds of things, but it's, but I haven't, I haven't kept that up.

Unfortunately, I think many of the things that are kind of, habits that I have, I suppose, would be, things that passively happen, recording data about, keystrokes, or email or whatever else. It's not things where I have to think I'm gonna do this.

I guess... that, yeah, there are a bunch of other things. I mean, I've been, Wolfram notebooks, I've been making, well, what originally Mathematica notebooks, I've been making since 1987, and I'm... I'm pleased with the fact that notebooks that I made, you know, in 1988 or 1990 or whatever.

They still run today. That's kind of a big achievement of compatibility of our language. Always it's kind of, funny because a notebook that I made in 1990 or something, it's a little tiny thing on the screen with very small, small, you know.

width and height, because that was what was needed to fit on a screen of that time, and now it's a teeny little thing off in the corner. You can obviously resize it, but the big thing is that it still displays, and the code still runs. But,

In that regard, I, Have had...

Oh, various conventions, like, for example, if I'm doing some particular project, I'll have a series of notebooks that are named for that project, and they'll be the name of the project, or what I... at least what I first call the project. It sometimes can change a lot. So the first call the project, and then dash.

And then eventually it will end 0-1, or 02, or whatever, a sequence of numbers. I always start with 01. I've never actually blown out of the 99, and gotten to the Notebook 100 in a particular project, though I've gotten pretty close. I think the closest was a project I did about the physicalization of metamathematics a few years ago, where I think I had reached the 90s in terms of number of

Notebooks associated with that project.

But then, in, well, for a long time, had the convention of the base name of the project, then dash, and then maybe something about what's in the notebook, and then dash, its sort of serial number. That's been a convenient way to do those things.

I suppose, oh, there are other kinds of things. I... I,

Yes, a convention from a really, really long time ago has to do with the naming of functions. I mean, the,

when I write code, and I'm kind of, writing something, and I write a function called foo, or something like this, and I want an extended version of that, and I'm just writing kind of research-level code and things, I'll probably call it foox.

But then, somehow.

I, developed the convention from the very first computer I used, where... where kind of system functions were... if you had a kind of subroutine and...

and you were kind of writing it out in... not in... directly in machine code, there was some convention where you start the name with Q, and I've sort of continued that convention now for, what is it, a solid 55 years or something, as a kind of a way of changing, making sort of variants of function names. Another kind of random personal convention. I'm sure there are a bunch of others, which I... I don't,

immediately remembered, oh yeah, another one that's a long-time one is, is,

I do a project, I get a bunch of books, I'll kind of put bookmarks in them, and I have kind of this... this, this tradition for myself that when I finish the project, I pull all the bookmarks out of the books that I was using to do the project. It's very satisfying to have these rituals for when you finish a project. It's kind of like a ceremonial act in finishing a project.

Back when I was working on New Kind of Science, I, my office at that time had these piles of paper that ended up being on the floor, for the different sections and different pieces of New Kind of Science. It was very satisfying in the last year or so of that project that gradually the piles would disappear, and eventually the whole floor was visible again.

And then the book was finished.

Talking of books, I see a question here. What was the first book you wrote?

Well, that's, goes back a long way. Depends what you mean by book. So, when I was, a kid, I started writing these kind of book-length things. I think the first one I did was one in... when I was 12 years old, 1972.

Which was a concise directory of physics, which was kind of a one-page-each summary of lots of, kind of data and information from physics.

I think in those days, I took particular pride in kind of making nice diagrams of things with, what... oh gosh, what is the American name for it? In British English, it's called a mapping pen. And it's,

I guess it was the type of pen cartographers used, but it was kind of a way of making very precise, things where... where the kind of width of the line was always the same, and so on.

But in any case, I make kind of carefully made diagrams of that kind, and then I was using the high technology of using a typewriter to type all the other stuff. There's very little of my handwriting that's to be found in that document. But that was probably, I don't know, a couple hundred pages long or something, and didn't see the light of day as a thing that anybody other than me looked at.

for at least 30 years after I wrote it. It's now somewhere on the web.

But that was kind of probably the first book-length thing that I wrote. I mean, I had... I'd written a bunch of things about collecting information about, well, first about spacecraft, and then about, well, then I made the transition from that to physics. So there were a bunch of things about spacecraft and planets and things like that, that I would bind as little, little, kind of...

bound books, but they were... they were short, and a lot of it was basically newspaper clippings and the analog. But then,

I got more serious after my first effort in... at the 12-year-old effort, by the 13-year-old effort was a thing,

What was it called?

Oh my gosh.

it was kind of a survey of particle physics, and it's kind of tight, single-spaced, and it's pretty long, and it has pretty good descriptions, actually, particularly of kind of intuition behind particle physics, and history of particle physics, and so on. And that was my... how I spent the next one of my summers when I was a kid, I guess that I,

Well, my...

right when I was 12 years old is the English system. I left elementary school, was about to go to high school, and my parents, had

two places, one in Oxford, and one in a small village about 45 minutes' drive from Oxford, and the one, the small village one, was where I spent, sort of, vacation time, and, so there really wasn't anybody else around, so I spent my time doing things like trying to write books about physics, and which I did with,

With considerable diligence, I should say. The, only, kind of, the only, what can I say, major distraction was kind of going to the nearby physical... nearby university library, which was in a town called Reading, which was about I think 6 miles?

From... from where I lived, and that was a good bicycle adventure and so on. And in those days, you could, you know, even if you were a 13-year-old kid or something, you could just walk into the university library and start finding books and journals and things, and at that time, I...

I am... My main...

financial outlay was putting coins into a coin-operated photocopying machine to, to produce photocopies, some of which I still have.

Although I think there was... it was kind of one of those wet photocopiers that, I don't think the, the printouts that it made last... lost the test of time terribly well.

But in any case, I... so my second effort was, survey of particle physics. Then the next summer, which must have been 1974,

I wrote a thing called Introduction to the Weak Interactions, which is about, you know, one of the, something that I got very interested in. I... I liked the kind of zoo of, of, elementary particles that were known at the time, and,

many of those particles decay through weak interaction, and kons, and pions, and hyperons, and all these kinds of, cool, supposedly elementary particles, although they're all made of... the ones I just described are all made of quarks and things.

And I got very, sort of, technically interested in that whole thing and explaining the properties of all these particles and such like, and so wrote this kind of book-length treatment of those things, which again, didn't see the light of day until decades later.

But, those were my, I would say, juvenile works, so to speak. I think they were... particularly the last two were pretty decent in what they contained.

I went on later and wrote a bunch of other things about particles and weak interactions and so on, I think ending in 1980 with a big piece that I wrote about some weak decay processes, which I think was really good, but nobody read it. And possibly that was because it was in some... conference where they had had some big constraint about how it had to be less than some number of pages long, and so, sort of all paragraphs were combined. It was quite a long piece, but that probably didn't help its readability. But even though I was personally a huge fan of, you know, the decays of

You know, the decays of kons into pions and muons and all this kind of... kind of thing. The,

The rest of the world didn't seem to care too much about those types of things. I think in the intervening years, the things that had been measured in some experiments are now measured to great precision, and even, I think, the,

One of my favorite exotic decays was the decay of the Pi Zero to a neutrino and anti-neutrino, which, back when I studied it around 1976 or so, was a deeply theoretical thing to investigate, because neutrinos are kind of hard to detect. But I noticed recently that I think that decay has now actually been observed.

So, it's remarkable what happens over a 50-year span as experimental science and measurement techniques and so on improve, and the intensity of accelerators improves, and such like.

But then, what was the next book-length thing I wrote? I guess the next one was 1981. See, this is not such a simple question, because these... none of these were kind of published as... as books. Next one was...

1981, SMP, my sort of precursor to Wolfram language, the SMP stood for Symbolic Manipulation Program, it's the thing that,

I kind of created, starting in November of 1979, along with, an expanding cast of other characters.

And that kind of, first... the first version came out in June of 1981, and it was a big, thick, orange book.

That was its documentation. It had a bunch of interesting ideas for... it had a bunch of interesting ideas for the design of the system. The documentation came in three parts. It was, some aspects of which weren't terrifically good ideas.

But, it had a reference documentation that was kind of just, like, very much like a page, a Wolfram language function page,

the function, basic description, and so on. But I had this sort of bright idea that there would be a pure reference manual, which would just be that, just the description of the function and, kind of basic... basic description of how it works. And then there would be kind of an extended version of that, which would be that, plus a bunch of examples.

And so those were sort of bound as two separate books, and actually the one for... that just had the description of functions also had a small form factor version that was kind of the pseudo-pocket version of that.

I don't know if this was a good idea. It's certainly something that in the modern online world doesn't make a lot of sense, because you can always have pieces that open and close, but it was sort of an interesting concept at the time that there would be sort of the reference documentation, and the only difference in the more extended documentation would be the presence of lots of examples. I think in those days, the examples were explanationless. They were just like, here's what the function does, here's one thing, another thing, another thing. And then there was a prima that was separate from that, and I think in the original version of the S&P book.

they were all bound together in this big sort of orange, bright orange, covered document. Those things are pretty rare these days. I have a small number of them, but,

They're rare. The online version exists, although it went through several different editions, and I'm not quite sure what's online. The first version had a number of interesting things, so I... one page that I learned a lot from was the page that was called, something like Common Errors in Using SMP. That page was a disaster, because it was a list of, kind of, 10 things that people often get

confused about, but it was sort of, in some sense, it was a summary page, and a bunch of people kind of kept that summary page, and they were even more confused than they would otherwise have been, because what they remembered was what was on the page about how they might be confused. So that was sort of in the category of bad ideas.

And then...

there were some other pages there about bug reports and things like this, which I now forget, but which I was looking at a few years ago and found very, very charming, so to speak, after all these years.

Well then, after that project, I started working on, kind of, the computational universe and cellular automata and things like that.

And I think the next book I produced was 1985, 1986. It was a collection of papers about cellular automata and Related Things, called Theory and Applications of Cellular Automata. It's actually the only time when I've had papers

from people other than me, included in, in a book that I've kind of edited. That book was,

It was published by a then very new publishing company called World Scientific. It's now, I'm sure, a much, much older publishing company, but it was... that was one of its, first productions.

And

I think for that book, one of the things that I did was to try and have sort of reference material about all these different properties of cellular automata. That's turned out to be a pretty useful, part of a pretty useful document on its own.

And then... So that,

that book. Actually, I should say that in the, you know, that my juvenile works, so to speak, were all typed on a typewriter. By the time I got to the SMP book.

it was done using computerized typesetting, a system called TROF, which was the typesetting kind of system that sort of came with Unix, something developed by folks at Bell Labs.

And, that's what I used for that.

And then...

for, I think I used the same thing for the additional material for, the,

the Theory and Applications of Cellular Automata book. I realized there was an intermediate book, which was, the

proceedings of a conference about cellular automata that I organized together with a couple of other people in 1983 at Los Alamos, at Los Alamos National Lab. That conference was... was a really, a really fun thing.

Because... I'd, you know, had my first sort of major results about cellular automata the kind of previous year.

And for that conference, I kind of... actually, it's the place where I introduced my four classes of cellular automaton behavior. But the conference was really my effort to invite everybody who'd worked on cellular automata, or anything equivalent to them.

Ever.

And there were many different names, tessellation automata, homogeneous structures, automorphisms, well, of the automorphisms of the shift, all kinds of different names, forests of stunted trees, all kinds of names that were, in the end, cellular automata, but people didn't, you know, hadn't sort of connected them together in that way. So I tried to invite everybody who had written about

things equivalent to cellular automata to this conference, and of course, most of them had never met each other before, and that was a, an interesting, kind of event.

I remember our after-dinner speaker was Stan Ulam, who had, kind of worked with von Neumann on kind of one of the original branches of cellular automata. In fact, the branch that led to Arthur Burks, who gave the name cellular automata, which arguably I shouldn't have used when I started working on these things, because what they meant was these two-dimensional arrays, and I meant these one-dimensional things, and so on.

But, I decided to kind of respect that aspect of history and use that name. But in any case, so I suppose that was... the proceedings of that conference was another book-length thing that I produced.

And then... The first,

first sort of thing that was... well, I suppose theories and applications of cellular automata was a regular published book that I expect you can still buy and sort of use book places today. And then...

Then I did the Mathematica book. That was in 1988. That was the first big thing that I did with tech, and a big system was written in a combination of Wolfram language and tech to produce the,

the Mathematica book, that was a whole adventure of typesetting and so on that wound up with me kind of physically babysitting the typesetting machine that produced the camera-ready material that turned into the original Mathematica book, the original film that was printed for the original Mathematica book.

So that was... that was probably the first,

completely made-from-scratch book-length thing that I produced was 1988, and then that went through 5 editions in physical form, and then kind of expanded to a ridiculous extent to the point where it could only be online, and it's now hundreds of thousands of pages that are online and no longer authored solely by me.

In fact, I end up editing a bunch of that material, but I don't write it at this point. Through the actual physical books, those were all... every word written by me, for better or worse.

So then, yeah, that, that's, that's about, then... That was,

1988, and then after a couple of editions of the Mathematica book, there's another book that was called Cellular Automata and Complexity, which was kind of a cleaned-up version of the Theory and Applications of Cellular Automata book.

that now was just stuff I'd written with, more ex... better retype set, and with more tables and things like this, and more beautifully printed and such like.

And then, well, then we're on to, that came out in the mid-1990s. I thought it was going to be just moments before my new kind of science book came out, but actually that book took a decade to write and came out only in 2002. And then there was a long hiatus, my gosh.

I think what came next was probably my book, Idea Makers, in probably 2016.

Followed by lots of other ones.

And I've been much more of a book-writing role. I have another book, a new book, coming out, In, about a month.

called... it's a collection of things I've written about philosophical topics. It's called, What Ultimately Is There? and Other Foundational Questions.

So, check that out.

Let's see... My question's about, were you always as systematic about your life as you are now?

No. I mean, I,

The fact that I record all sorts of data about my life, that's a phenomenon, kind of started in the late 1980s, most of the, kind of, recordings of things came from, wait a minute, I don't have that anymore, I really want that, I really want to know what happened there. I think, the... in terms of, sort of, the organization of my life, you know, I was... I was, definitely able to be quite a free spirit in much of my... early in life, because I was just doing things like research. And that, you know, didn't involve very fixed commitments of time. I mean, I... I did a certain amount of professoring, actually teaching classes, and I guess I managed to show up for those. I mean, I have my little date books from back in the... oh, gosh, from... I started having little date books where I kept that information as soon as I, probably, oh gosh, must be from the late 1970s. And, you know, I had all kinds of little codes for myself in... because I would write things out in fairly short form in these, in these little, little, little date books. And that was kind of how I ran my life up until, up until I... even after I started, my current company in... in 1986, it's now 40 years since... since, Wolfram Research. got started, it was 40 years last... last month, I think, since, sort of the beginnings of that. And, I quickly started accumulating much more organized online calendars and, and, detailed, assisted, assisted, forms of, of, of organization, so to speak. I guess it got more organized after the... after I finished New Kind of Science. While I was working on New Kind of Science, I was really trying to squash my, kind of, corporate activities into a limited number of hours per day, and then having very unscheduled time the rest of the time. But, in more recent times, I'm... I'm pretty deeply scheduled, which I find, good, although I... there's... there's quite a lot of that time, particularly in evenings and weekends, when it's kind of, just a huge block. Where I get to work on, on things that I'm doing, sometimes on my own, sometimes with other people, and so on. Let's see... Dr. Med asks, in what ways have current developments of AI met and contradicted your expectations? It's sort of an interesting question. So, when I was a kid, AI was always the future. And, you know, it sort of came from science fiction, whether it was the HAL in 2001, or whether it was the computer in Star Trek, or whatever. You know, these were things from... well, I saw 2001 when it first came out in 1968, and I think that was the first kind of serious kind of AI portrayal that I'd seen, and it sort of seemed like the obvious future. And, you know, I would have books about the future and so on. Actually, I was a little bit horrified to find that in 1972, when I was 12 years old, I actually gave a speech about AI at the school I was at. I don't know what, I was... I was quite one for doing these kind of public, speeches and so on. I guess that's what turns into, when you're ancient, it turns into live streaming or something. But when you're 12 years old, it turns... it's... it's, you give speeches and so on for, various kinds of assemblies and so on. But I... I gave one about, essentially AI ethics,

And, I have a typed version of that speech, which I... I suppose at some point I'll have the nerve to put online. I'm a little bit, embarrassed by the, well, there's a lot of Britishism in it, a certain amount of...

kind of...

sort of the... the way that, the way one talks about things has changed over the years. It's, you know, things like, back in those days, it was always, you know, what will men do, as opposed to what will people do? I don't know quite why I was... I mean, I was just following the prevailing convention of how you would say those things at that time.

But, it was, so I, I was definitely thinking about AI kinds of things, at that time, and that was, actually quite

a, it was kind of a, sort of a...

you know, what do the humans do when the AIs have taken over type of speech with a, hey, maybe we shouldn't let the AIs take over kind of thing. That was the 12-year-old me, okay? That was a long time ago. That's, you know, 50... 55 years ago. So,

Then.

I would say that, in my efforts, I was really involved in doing physics through most of the 1970s.

And, by the time I was working on SMP, my, my first big computer system,

the kind of the concept of SNP was to, kind of.

abstract from the kinds of ways that humans think about computing things in mathematics and elsewhere, to try and abstract from that into sort of a more general way to think about the computational process, a way of thinking about the computational process that sort of connects with the way humans think about things. And what SNP was very much based on is the idea of transformation rules for patterns.

And I, that's also what Wolfram language is ultimately based on, although in a less extreme way in some respects than SMP.

But in SMP, I had this kind of idea that was sort of emulating the way that humans would do things like mathematics and algebraic computation and things like this. So I certainly was aware of thinking about, kind of, sort of machines as doing things like the way that humans think about things.

Actually, right around 1980,

I, really started thinking, well, what would it take to make a real, kind of, human-like thinking kind of thing?

And, I suppose I had been...

well, from even earlier than that, I had been sort of collecting information about things, and I had kind of wondered what would it take to have something that can sort of answer questions and so on, like brains can answer questions. And I always thought you have to make something much like a brain to do that.

And so back around 1980, I started thinking, how can I take this kind of transformation rule paradigm for, you see this, you get that, sort of matching patterns, and how can you take that and,

And, sort of generalize that

to things which are not precise patterns of the kind you might use in mathematics or computation, but the sort of fuzzy patterns that it seems like humans are using when they do things. And so I had a whole effort of trying to understand how to do sort of fuzzy pattern matching, and I was thinking about sort of generalized hashing and all kinds of things like this.

And I... I have to say, I always thought that there would be a magic bullet, so to speak.

that there would be some way of doing sort of patterned, you know, you could do this sort of fuzzy patterns of if you see something roughly like that, you should transform it roughly like... transform it like this, and so on, that that would be kind of a key thing in sort of cracking the problem of AI.

I mean, I guess I had, been aware of, oh, like, I remember I read, Norbert Wiener's Cybernetics book back in 1976, and I wasn't that impressed with it. It's a book about, sort of, control of, communication and control in animals and machines.

And, that sort of has a different take that's more of a sort of a control theory take on how things like that might work.

I think I was... By 1980, I was certainly aware of neural networks.

I knew... John Hopfield, who was a person at Caltech who worked on early neural net kinds of things, sort of noticed that a particular way of setting up neural nets made it kind of easy to teach them particular, particular transformation... essentially, to have, you know, if you see this, you will get that.

And to have some sort of basin of attraction around the what you... what you need to see in order to give that.

So back at that time, this must have been 1982, 1983, by that time, I was writing code to sort of reproduce neural nets,

and reproduce what John Hopfield had done. I was never quite able to reproduce what he said he'd done, and I never quite figured out why. And he always said, I must be making some mistake, but I... I didn't see an obvious mistake. And in any case, the,

the thing that, sometime around then, I spent a month visiting the Institute for Advanced Study in Princeton, and I spent a bunch of that time reading von Neumann's, papers and books and so on. and particularly his... the last book he wrote, which was about the... what was it called? The Computer and the brain, or something like this.

That, sort of a story of neural nets meets thinking meets computation.

And so I was pretty interested in those things at that time, and actually, I tried to understand, you know, what's the essential thing that's going on in neural nets. Particularly, I really wanted to know what is the basin of attraction, the set of things, if you teach a neural net, if you see this pattern of bits, then it's that one.

What is the... what is the, sort of, the region around the pattern of bits you teach it about that it... will think are sort of the equivalent pattern of bits. You know, what are the... what are the images that it thinks are a cat? It wasn't doing fancy images like cats at that time, it was just doing

arrays of numbers and so on. But, what, you know, what's the... what's the thing that's around the, the thing you explicitly teach it. What's the thing that it generalizes to? And I was really interested in that, and I tried to figure out things about that. Didn't get very far. I'm a little bit shocked that we still don't know the answer, and I've been working on actually related problems even right now, all those years later.

this question of, you know, just what is the boundary of catness, and so on, in the space of inputs?

So... But, at that time, I,

I kind of had played around with these things. Neural nets weren't doing anything terribly exciting, but AI was the future, so to speak.

And then... oh yes, actually, another major part of the story. In, in 1981, I started my first company, which was originally called Computer Mathematics Corporation.

And then, in what... well, I don't know whether it was a... it was a move, I'm not sure it was a good move. I sort of developed that company in the way that companies often develop, and, got venture capital. In those days, all the venture capitalists were in New York, the San Francisco venture capital scene hadn't started up yet.

But,

raised venture capital and so on, and the venture capitalists were extremely keen that the company go from this sort of direction in computer mathematics and so on, and it's like, the way of the future is, this is 1982, AI.

So the company kind of pivoted towards AI, and started making things that were for, kind of, expert systems, kind of, symbolic rules for figuring things out, although not using, kind of, the core technology that I'd built from S&P, kind of another set of things that were, I think, pretty lame, actually. But in any case, that company, renamed itself Inference Corporation, so starting in,

1982, I guess it was. You know, I was, part of what was called Inference Corporation, which was one of whose activities, its principal one, according to its investors, was doing AI.

And there we were, I think we were at 5300 West Century Boulevard, if I remember correctly, in Los Angeles, right near Los Angeles Airport.

And the logo of the company, which was just the most ridiculous thing, was this inference sign, which was basically made of an equal sign and a greater than sign, in what looked like a badly produced piece of ASCII art.

In any case, the person who I brought in as CEO of that company was very, I think... I think his daughter had designed that logo. He was very, very enthusiastic about it.

But in any case, that was... so that was another exposure that I had to... to AI, was in the world of expert systems, and at the time, kind of, the idea was, sort of one just sort of interviews experts and captures what they have to say as rules, and then the computer can do what the experts did.

Didn't work terribly well. I mean, that company did build a bunch of systems, it built early fraud detection systems for credit cards, it built things for checking out spacecraft and things like this. It built a bunch of technology based on that kind of idea.

Eventually, I sort of detached from that company by the mid-'80s, and the company eventually went public in a rather undistinguished IPO in the mid-90s.

But,

the, the thing that, so that was my next, I suppose, exposure to, to AI, so to speak.

Back in, probably, 1986, I was... have been working on, kind of, simple computational systems, and I was kind of interested in, could you learn in cellular automata, and things like that.

could you both learn, and could you evolve in the sense of adaptively evolve, like in biological evolution? So I wrote some piece called, Approaches to Complexity Engineering.

Which was kind of an effort.

To understand whether you could kind of, sort of do engineering-type things with systems whose behavior is complex.

And I particularly looked at cellular automata, but it was kind of an adventure of seeing, you know, with the complexity underneath, can you sort of direct it to particularly do these kinds of pattern matching operations?

So, another thing around, oh gosh, it was earlier than that, I'm forgetting all these times. I got involved with a company called Thinking Machines Corporation.

I was based in Boston, started by the technical founder, who was a chap called Danny Hillis, with some involvement from Marvin Minsky in the whole AI mafia, so to speak.

The,

That was a company whose goal was to create this massively parallel computer that would implement various models of AI.

It was, kind of one of the... one of the slogans of the company was to make a computer that can be proud of us.

That, that...

I worked on a bunch of things there, but mostly, the challenge was to actually make that computer that was very expensive, and have people actually buy it to do real things, and some of the real things they could do, where the computer, regardless of what its uses for AI might be, was really good at running things like cellular automata, and so I developed this method for doing cellular fluid dynamics using cellular automata that I think sold at least a few of those machines, and a few is a lot, because the machines cost between \$1 and \$4 million each, which in today's dollars would be, I don't know, three times that much or something.

So, that was another kind of exposure to AI. Again, rather aspirational, and not clear how it would work. I mean, back in those days, I was involved in kind of designing languages that could run in a massively parallel way.

I kind of designed a language called Seastar. I think its design got kind of ruined afterwards, but it was kind of very elegant, I thought, in its original design, that was a very minimal extension to the C programming language.

to do kind of parallel computation. I was very interested at that time in graph-based languages. I kind of had this idea that there would be a way to do kind of AI-like things by using transformations on graphs. I never managed to figure out how to make that work.

Years later, we have the physics project, which is all based on graph transformations, and now a bunch of things we're doing right now have to do with graph transformations as ways to do things that connect to AI and so on. So, I had the right intuition back in the mid-1980s, but I absolutely couldn't execute on it, and just didn't understand, kind of, all of the what matters and what doesn't matter in things like graph rewriting. One of the key ideas I missed

Was the taking seriously multi-way systems, taking seriously the idea that a computation might just have one track that it goes through, but might have a whole set of branching and merging different tracks, and that that's okay to have that be the case, even if you pick just one track in practice, and so on.

So I would say by the,

well, that was all before development of Mathematica and Wolfram Language. After that, people started to implement neural nets in... in Wolfram language. I never paid that much attention. They always seemed like toy kinds of things.

There were always a lot of kinds of operations that I kind of avoided us having in Wolfram language, because I thought they were too heuristic, it's like you never knew what was going to happen. I had seen the... in terms of symbolic languages and functional languages and so on, I had seen the experience of languages like Prolog.

That had been a thing, you know, in the 1980s, there was sort of one of the big dramas was the Japanese fifth-generation computer project.

which got many people very kind of concerned that, you know, the US would be left behind, all these kinds of things. That was a project to sort of bring AI to the world, particularly in Japan.

And, the, kind of the language that was being used there was prologue, and prologue is this language where you kind of state what you want, and then the language will sort of find a way to get there. And that idea turned out to be not a very good idea, because kind of stating what you want

And then trying to get there was something where, when you did that with precise computation, it was like you never know how long the computation's gonna run.

the thing could kind of get lost, sort of in some undecidable sequence as it tried to get to the goal, and you just couldn't, you know, particularly in those days, it was important to have computations where you could reasonably foresee that the computation would actually run at reasonable speed and finish and so on, and it wasn't something where you just sort of say, here's what I want.

now go off and do it, and I don't know how that's going to work, or how long it's going to take.

Now, in today's world, with AI, we're getting closer. With modern AI, we're getting closer to just say what you want, and then there will be a path found, regardless of, sort of, how much grinding there is in the process of doing that. But in those days, that looked very unappealing.

I would say that the, I didn't really... Wake up.

To neural nets until 2011 or so, when deep learning started, really taking off, and we built our first image identifier very soon thereafter, and it was kind of a nice consumer project that people hadn't seen that before, of being able to recognize, you know, cat or dog, or elephant, or cow, or, you know, this kind of leaf, or that kind of leaf, and so on. So we were... it was fun to put up that as a consumer website very early in that process.

I think...

in terms of... well, I should say that a bunch of the things to do with mathematical computation, people from the 1970s on had sort of put that in the bucket of AI.

So, you know, when...

when SMP came out, sort of, people were talking about it as, like, AI. Mathematica, a little bit less. By 1988, kind of the, the kind of fortunes of AI had declined, and people were like, oh, no, no, it's not, you know, AI, it's not AI type thing.

And that lasted for a long time. I mean, it was, you know, I think AI has gone in waves where people are like, it's very exciting that it's got AI, or it's absolutely horrifying that it's got AI, or that doesn't work at all.

I mean, one of the things that happened in the history of AI was this report in the UK, the Lytil Report, which basically said AI was never going to work.

I forget exactly when that was, but that kind of put a damper on some... some research in AI.

And I have to say, if you'd asked me in the 1990s, are neural nets going to go anywhere? I would sort of say, I don't think so. And I always thought they were very messy as models. I had kind of come up with cellular automata and so on as a way to simplify the kinds of things that were going on in neural nets as a way to understand what was happening in neural nets. Now, just a couple of years ago, I finally came back

And started using, sort of, cellular automaton, very structured grid-like things with discrete elements.

as a way to actually do practical machine learning of the kind that people have done a lot with continuous neural nets, and sure enough, it seems to work. It's kind of an exciting thing, not as much followed up as it should have been, that you don't really need all those floating-point numbers and so on. You can kind of distribute the information into more elements with just a few bits, or just one bit.

And that's a thing to look at.

But, The,

Yeah, I think, you know, I... nobody expected what happened. Nobody expected what happened with image identification in 2011, 2012. Nobody expected what happened with ChatGPT in 2022.

I mean, these were things where, sort of, the threshold of human compatibility was just not clear when that would be reached.

And,

there's, it's... it's interesting to sort of get intuition about what's going on, that may be a different topic. When... when ChatGPT came out, I... I was sort of deluge with people saying, how does this work? What's going on?

So I sat down and wrote a little piece about how ChatGPT works that's turned into a book that's now translated into many languages and seems to have been very widely read about that question that sort of contains my attempt at intuition about what's going on.

And,

it's, I've been thinking about writing a sequel to that at some point, although I feel like at this point, it's, there's a lot of confusing stuff and a lot that's... that's not, that's...

It's kind of more marketing than reality, and a lot that's mysterious, and so on, and I kind of want to do a little bit more science before I say more about how I think things are really working. I actually have just been writing a piece about the role of AI in math.

which hopefully will come out in a few days. It turns out to be a complicated story. We're sort of full of philosophical questions about what is mathematics and so on. So I hope I've been able to unravel some of those. Some of them I think, unraveled in ways that I really wasn't quite expecting, so that's perhaps of interest.

So, There's a question here from Ollie. Did you ever seriously think about moving back to Britain?

Well, I had... I'd moved from the UK in... 1978.

when I sort of dropped out of college and went to graduate school, and I still have the air bills from when I shipped all my stuff.

From the UK to the US.

And,

You know, I was in the US for a couple of years, got my PhD, was, I thought about taking a job at CERN in Geneva.

And for a very brief while, I thought about taking a sort of job both at CERN and in... at Caltech in Los Angeles, and kind of commuting between them, and I... I got very familiar with the LA to...

Geneva kind of route, and the, and just, sort of, wandering around 747s, as, as they made that, that journey. But I realized that was a completely crazy thing to do.

I thought about moving back to England. I had always liked, kind of, the British countryside. I'd always kind of imagined

This is something I've probably never talked about before, but I'd always kind of imagined, some kind of research-century type thing being sort of the thing that I would create as a place to work. I had worked

When I was a kid, I was 16 years old, when I left high school, I worked at a place called the Rutherford Lab in England, which was a particle physics lab. It was in a rather beautiful, kind of, you know, sequestered away, sort of in the countryside kind of thing, and I'd always thought that

was pretty cool, and that that would be a thing I would do, and I even had a name, actually. I was calling it Advanced Research Center, or AHRQ.

and I'm kind of amused that people have used that name, I mean, absolutely without any knowledge that I ever thought about that name in recent times. But that was kind of my... my concept for my, kind of,

adult activities, so to speak. And I thought about, sort of, moving back to England and maybe being at some university or some lab, doing something like that. But I have to say, this was probably around

Oh... 1981, maybe?

And everybody's kind of telling me, hey, you don't want to move back to England, nothing's happening here. You know, if you're doing physics and things like that, or tech kinds of things, it's all in the US. It's a backwater back here.

I like the countryside, though. Wasn't so keen on some aspects of the people. I always used to say that, England was a good place to be old in and not a good place to be young in, because it seemed like there was always a kind of irrational respect for the old. Oh, yes, so-and-so did this thing when they were 20 years old.

And that's so cool, and now they're 60 or something, and one should sort of treat them with great respect for their age, so to speak. I thought, yeah, you know, that's a good deal if you're old.

Of course, needless to say, the culture has changed, and I don't think that that would work anymore in the UK, so it kind of, the intervening years, didn't pan out the way I expected. But, no, I, I really, after...

1981 and 1982, I kind of gave back on the... gave up on the idea of ever moving back to the UK, because it just was... people had told me, there's nothing much happening here, you know, you're much better off in the US. And I kind of got...

more and more involved in things, particularly tech industry things in the US and so on, and never really thought about it again.

Chris is asking,

Was there a point during time I was working on a new kind of science when you wondered whether you'd gone too far down the rabbit hole?

Well... That book has... sort of encapsulates 10 years of intense work that I did.

And I'm still sort of untangling, essentially, footnotes from that book. I mean, a lot of things, several books I've written were basically footnotes from that book.

And, well, at least, no, maybe not footnotes, sections from that book. And there are certainly long, sort of, posts, 50-page posts that I've written which were essentially based on footnotes from that book.

So I packed a lot into it. I sort of had the good fortune or something that I had a fixed table of contents from very early in that project.

And so, although some of the later sections and later, sort of, areas of those table of contents did sort of balloon a bit, I was kind of pinned to that table of contents to some extent.

But...

you know, in the notes to the NKS book are very rich, with lots of kinds of things in them, and some of them took a lot of effort to produce, and turned into, you know, a few column inches in the book. That's why I'm expanding them into many tens of pages of posts and things that I'm making. But,

I think...

I, you know, I really tried to drill everything... I tried to drill as far as I could with reasonable effort.

Both in the technical side of things and in the historical side of things. I mean, I certainly did a bunch of, sort of, actual interviews with people to try and find out about history.

If there was one thing that annoyed me in the response to the book, it was people saying, oh, you didn't respect the history of what had happened. Well, actually, I did much more on that than academics usually would ever do unless they were historians of science, and even perhaps more than many of them would do.

You know, I put incredible effort into kind of figuring out the history of things that I talked about. The history wasn't always great. It was like, people say, oh, but I invented that. Well, no, you didn't. Or, you know, people had figured this out. Well, no, they hadn't.

And you actually have to drill into the history to know that. Sometimes people had figured out more than you expect, sometimes less than you expect. Sometimes, with the wisdom of hindsight, it looks like something was figured out, but it really wasn't at the time. If you'd asked somebody at the time, what do you think about this thing that you could understand from the future, they would, like, you know, say, I don't know what you're talking about.

So, you know, I think in the I don't know what you're talking about category, it's kind of, then they didn't invent it, because they didn't really understand what it was that was only understood later. So, there's a bit of complexity, as there always is in writing history of science, but I put much more effort into that than... well, I put a huge amount of effort into that.

I'm trying to remember whether there were... Yeah, I had kind of a rule for myself. There were no notes on the notes.

So anything had to be a one-level note. And there were definitely times when I thought about writing notes on the notes, and I said, I'm not going to write notes on the notes. And that's fortunate, because the project only took me 10 years, not a lot longer than that.

I mean, in... you know, my principle in more recent years has been it's better to have the thing out there than to have the thing be absolutely what one might think of as perfect. With the NKS book, I put a lot of effort into perfection, some of which really paid off, some of which I think was... was not particularly valuable and just made for, you know, very tight text and so on, but I'm not sure it's any more readable than it would have been if it had been less tight. Maybe it's a shorter length, which was relevant for having a printed book, the NKS book has 1,280 pages, which was the absolute limit of page length before the binding technology that we were using doesn't work anymore.

Let's see... Artichoke asks, what's a project you regret not having worked on more?

You know, everything I didn't actually get out in the world, pretty much.

I regret. Not very many things.

But...

you know, like, long ago, 1976, I wrote something about neutrino background radiation, neutrino cosmic microwave... cosmic background. So there's a cosmic microwave background with photons, there's an analogous one with neutrinos, I figured out a bunch of stuff about it, I didn't figure out... one of the things I really wanted to figure out.

But it was a... I never sort of published that, and I should have done it.

And pretty much everything like that, where I, like, I did the work, but I didn't publish it for some reason or another, I later regretted. I mean, there was another thing, actually, on cosmology.

early work on cosmology that I did in 1977 or so, that was... I wrote kind of a pretty long paper about cosmology meets particle physics, very early in that kind of business.

And the journals that I sent it to said, oh, but most of this stuff is well-known, and you should just concentrate on the bits that are definitively new. So I wrote a much shorter paper and did that. And the sad part is, the thing that I wrote that had that sort of survey of cosmology meets particle physics, which wasn't as well known as you might have thought, in the sense that it was definitely a refactoring of what was known.

It's a big shame I didn't publish that at the time, because it would have been very useful to people. It was a... you know, it would have helped accelerate that field, I think.

I mean, there are other...

I've tried to not have unfinished projects, or not have, and in recent years, I've done a lot of finishing of projects. You know, I've tried not to have projects where I took them a long way, but never finished them. I mean, there's another one from the 19... when is it? Late 1970s, early 1980s, about relativistic matter, which I also never finished, although it's hard to finish.

That one, I don't even know whether I can finish it today.

Another one from the early 1980s, I wrote a couple of papers with a chap called Jan Ham Bjorn about properties of the vacuum.

One of them has had a long literature chain, because it's the thing that kind of shows that, showed that by kind of deforming the shape of some cavity, you can, at least in principle, sort of mine energy from the quantum field theory vacuum. And that's a cool thing. If you want to sort of power things forever just from sort of taking energy roughly from the expansion of the universe,

And, you know, it's a very open question whether that's in any way practical. I didn't think it was, but, it was, something... and I didn't originally notice this effect. I just studied, kind of, what happens to the vacuum in these different shaped boxes and so on.

But there was supposed to be a part 3 of that project, which was the gravitational properties of the vacuum, and I'm ready to write that now.

That was, written... the parts one and two were written in 1981. I'm ready to write Part 3, finally, after all these years. Although there's a piece of science that I'm trying to do about, essentially, the entropy of free space

Which is kind of a weird concept, not... let me not explain that here. But it needs a bunch of science, and I've been... I've been meaning to get my... my old co-author, kind of, we've talked about it a few times over the years. I think he's not used to my current rhythm of doing science, and, my... let's, let's work together for a weekend and see whether we can get this you know, over the hump, so to speak. That's kind of a little bit of extreme science for your typical academic scientist and so on.

So, I think,

you know, there are other things that I've thought about working on for years and years and years.

But... For one reason or another, I've decided not to do it.

And most of those have the good feature that I didn't go very far in working on them. It's not like I have a big, you know, file of stuff, and it was an unfinished project. It was more like, I thought about it, I've been collecting information about it, I'm just not ready, I don't think I've had the key idea, so I think it's much more of that type. I was planning on writing something, actually. Okay, so this is a meta thing I didn't finish. I was planning

Well, a decade ago, and more recently than that.

writing a piece about projects I didn't finish, projects I had conceptualized, but never finished. But then I was... as I was writing that piece, I decided, I might as well just do some of these projects.

And, so that meta piece of projects that haven't been finished is not itself finished at this time. Let's see...

There's a question from Jonah. Do you think that our company would have evolved differently if we'd been based in California or New York versus in Central Illinois?

Probably, yes.

I mean, it's,

Central Illinois was a... was a really good place for the company, and I mean, it's still our headquarters, although the company has been very geo-distributed for many, many years. the It was,

we were able to find very good people with good kind of work ethic, and you know, the cost of living is lower, much lower in central Illinois than it is in New York or... or... or...

San Francisco Bay Area, or something like this. In a funny story, perhaps, back in... 1990, maybe?

I briefly toyed with the idea of moving the company to California.

And, at that time, it was a work... work-in-person type company. I kind of became remote in 1991, which is kind of when we started our whole geo-distribution of people kind of thing. I mean, I originally

started, I sort of went remote for two reasons.

one...

because I was trying to sort of peel off a bit from the day-to-day operations of the company to work on my science project, and two, because my wife, or wife-to-be at that time, Elise, who sadly died just recently.

was, was like, I don't want to live in central Illinois.

Ironically enough, she had been born in central Illinois, but that's... but hadn't lived there. And in any case, and so...

for a while, lived in California, that was 1991, but I think it was 1990 or 1989, I was like, maybe we should move this company to California, maybe it would be a, you know, it would sort of liven the whole thing up or something.

The place I considered moving it to was a place called Half Moon Bay, which is over... over the hills, so to speak, from the main San Francisco Bay Area. And I thought, I'm going to go, you know, it's a place full of pumpkins and so on. It was quite, quite, I haven't been there in forever, but it was... at that time, it was quite rural, although not far from the Bay Area, you know, not a... not a long drive.

But in any case, I went there, and on the particular day when I was visiting, it was incredibly drilly there. And it was like, California's supposed to be bright and sunny and so on. This is completely cloudy and terrible, and I know it is, you know, somewhat cloudy there.

Generally, on the coast. But in any case, so the weather of that day was what convinced me not to move the company to California at that time. I mean, I considered doing all kinds of things which were sort of

traditional company types of things, like I, in the beginning of the 90s, 91, 92, sometime like that, I considered taking the company public, doing an IPO, and so on, which, fortunately I didn't do, because it's kept my life a lot simpler not to do that. No doubt the company would have

grown, it might have, you know, gone up like a rocket and down, so to speak. It's been a much more steady journey.

As a result of not doing anything like that. You know, we don't have outside investors, and so it's been, kind of, you know, I'm responsible for my own fate type thing. I think the,

That's been what's allowed us to work on very long-term projects, allowed me to... to, kind of do things that, aren't just kind of work for a living as CEO kinds of things. I think the,

The thing that, It's...

you know, there's a certain swirl associated with the Bay Area now. I would say that over the period of time, the 40 years that our company has existed.

the Bay Area has gone through multiple kind of near-death experiences. I think at this point, it's out of... out of, you know, out of the... out of the risks, because there's just so many different, sort of, tech efforts going on there, but there were periods of time when it seemed like, you know, software was moving to... to, you know, Utah and Washington State and so on, and when, you know, it seemed like the,

hardware was moving to different places, and things like this, when it really didn't seem like Silicon Valley was going to sort of really make it.

And then, you know, another company would come along, and another gold rush would happen, and it would go on. But I think in,

I mean, our company, in its present form, I'm sure, had it been... had sort of the costs of California and so on, would not exist.

And so, I'm happy that we've had this more, perhaps, modest, but consistent existence sort of centered in Illinois, but mostly geo-distributed at this time.

Let's see, maybe a few more questions here.

Well, there's a question here from Junior about what's something you almost didn't publish or release.

There are a few things I've worked on where...

I didn't put them out there in the world. I was one on,

CP violation, time reversal violation, that I did back in the 70s, that I just decided it just wasn't

The right thing, that wasn't really right.

I haven't revisited that, I don't know what, I don't know whether my decision about that was correct.

there were things that I almost didn't publish. There's one... one notable one. I wrote a...

a paper in 1978, gosh, it's a long time ago, with my friend Rocky Kolb, who's now a professor in Chicago, but, was about cosmology, and it was about the expansion rate of the early universe.

And we had noticed that if various things were true, the early universe

the universe would expand exponentially, and we'd notice that that would kind of solve various mysteries about how the universe comes to be the way that it does, and how different parts of the universe are sort of consistent with other parts, even though light couldn't have traveled from one to the other, all those kinds of things. And we'd noticed this phenomena, that, but I didn't think, for a variety of reasons, that they were plausible to happen over enough orders of magnitude in the universe to actually have a significant effect, so we relegated that, kind of that discussion to a footnote.

That discussion was later picked up and sort of became the inflationary universe scenario, which was sort of an interesting lesson, because it was a place where, when the payoff is big enough for making even wild assumptions, it's worth making some of those wild assumptions.

So... That was a,

I mean, I will say that with our current physics project, I'm pretty sure that that inflation stuff is just nonsense, and there's a quite different mechanism for these things that has to do with dimension... the sort of dimensionality of space being a dynamical parameter, and that's sort of a thing that comes out of left field in terms of the traditional views of general relativity and the theory of gravity and so on. I think that's what's actually going on.

In that case. So the footnote was deservedly buried, but certainly a lot was made of those ideas in the intervening years.

But then, other things I... I didn't publish, or almost published, I... there are a couple, actually. One... well, one was about cryptography.

That was a, I realized that Rule 30, the Rule 30 cellular automaton that makes, sort of, intrinsic randomness, was a good cryptosystem. This was in 1984, when, sort of, the Cold War was going on, and cryptography was a very, complicated subject.

was a time when, the cryptographic systems that had been kind of developed in the U.S, kind of to some extent sponsored by U.S. intelligence and so on, were, were kind of the things most people used, and they were sort of just, you know, sort of nipping at other kinds of things, like public key cryptography and so on. And the Rule 30 crypto system was another example of this, and I certainly wondered, you know, what should I do, you know, because everybody was concerned about strong cryptography and its effect on, kind of, the security, the balance of the world, so to speak. And I eventually decided to go ahead and publish that stuff, and nothing too terrible happened.

Another thing, this was probably 1985 or something, maybe a little earlier than that, I was, was thinking about various things in physics, and I kind of had this idea for a way of a very sensitive detector for neutrinos, which would allow you to detect a neutrino signal from a nuclear submarine from orbit.

And because nuclear submarines, needless to say, have nuclear reactors, which have isotopes that decay through beta decay, and they produce neutrinos and so on, so there's a kind of then neutrino hotspots. But usually neutrinos interact so weakly that it doesn't... you're not going to be able to detect a submarine that way.

And I, for some brief period, I thought, you know, okay, you can detect nuclear submarines from orbit.

That's a big deal, because the big, sort of, thing about nuclear submarines is they can prowl around under the oceans, and nobody knows where they are, and even if the rest of the world gets blown to bits, they can still... your nuclear submarine will still be safe and can retaliate, and so on.

And that was kind of part of the thinking. So if you could detect all the nuclear submarines from orbit, that deterrent disappears in a day.

And so I'm like, what am I going to do with this information? It's really possible to do this. So I was... I was very happy it took a day or two. I was really, you know, I was partly agonizing about what to do with this information, and partly trying to see, was this actually correct? And I realized that, no, it doesn't... the idea that I had doesn't... doesn't work out, at least not at that time.

So that was, I was very relieved about that at that time. But,

In terms of other things, I've,

You know, when I worked on the NKS book, I,

I was not talking to people about what I was doing. I mean, at that time, it was before kind of the... well, the web was just coming into existence at that time, but certainly wasn't what it became, and it wasn't, you know, there wasn't sort of an incremental way to publish things like that other than in academic papers, but that has very high overhead.

And at the beginning of working on the NKS book, I would talk to people about it, and people would say, have you thought about this? Have you thought about that?

I would chase those rabbit holes for a little bit, and then I'm like, this is not a good idea. I've got a very ambitious kind of table of contents that I think make coherent sense. I don't want to chase this piece and that piece, and so...

I decided, okay, I'm not going to talk to people about this, I'm just going to get it done, and then talk to people. Of course, I didn't think it would take 10 years to get it done.

During that period of time, a number of horrifying things happened, as people who were sort of curious about what I was doing, I don't think that many people would care, but they did, were like, you know, getting leaks and all kinds of crazy stuff, which I have to say, probably slowed down the finishing of the project by a year or more, just dealing with all the messy, horrible stuff about people leaking things and behaving very

very badly, and eventually, sometime, I'll talk about those things. It's, it's... it was not a... they were not good human stories. They were very, very bad human stories, actually.

very much revealing of, of, you know, if you ever wonder whether there's skullduggery in, in academia and so on, yup, there's skullduggery in academia. There certainly was in the 1990s, and I'm sure there's, I would expect there's yet more today.

And, but those were, those were not,

those were places where, you know, I made the decision to not talk about what I was doing with NKS, just because I just didn't want to be, you know, didn't want to be distracted, and that had some

Annoying effects.

And in my more recent efforts in science, I've been very kind of, you know, I've been livestreaming things, and I've been doing all kinds of stuff, which kind of gets things out the moment it's created, so to speak, because I find that, you know, I like to do that, I think it's useful and interesting, and some people seem to like it, and so on.

I think it is time for me to... Go back to my... Day job, here.

But, thank you for those questions, and, I look forward

to do more of these live streams, I will say, I mentioned a couple of things, I'm,

Well, as a result of my...

Sort of, changes in my... Personal circumstances, which,

I've... I'm sort of deciding to do a little bit more

with actual humans, rather than just purely virtually. So I am planning, if anybody's in the Concord, Massachusetts, or Cambridge, Massachusetts area, I'm going to be doing some kind of in-person

Salons, that might be a little bit like, well, more interactive than, than these, events, and hopefully they'll be.

organized and advertised soon, and I'm also, may end up, maybe at some local universities, doing some of these live streams with an in-person, group, rather than just,

Me sitting by myself and looking at a camera. So, that, as,

something, well, I'm looking forward to. I hope they will be, perhaps a little bit different in character than what I've been doing in this purely virtual way, but I will be continuing this, sort of purely virtual approach to live streaming, and, look forward to talking to you all soon. Thanks very much for joining me. Bye for now.

UNEDITED TRANSCRIPT