Conversations with Bill Kristol

Guest: James Manzi, Software Entrepreneur
Senior Fellow, Manhattan Institute

Taped July 17, 2018

Table of Contents

I: Advances in Artificial Intelligence (0:15 – 40:16)

KRISTOL: Hi, I'm Bill Kristol, welcome to Conversations. I'm joined today by my friend, Jim Manzi, for the second time, I guess. About three, four years ago we discussed trial and error, your book, *Uncontrolled, the Surprising Success* – is that what it was called?

MANZI: *Surprising Payoff.*

KRISTOL: – *Payoff from Trial and Error*. Excellent book which repays reading today and the Conversation repays viewing today. But you've sold the business on which – which was partly based on that insight, I think, it's fair to say, right?

MANZI: Yes.

KRISTOL: And you're in a new business with artificial intelligence. And so today you're going to explain this incredibly fast moving and somewhat, for someone like me, difficult to grasp field of artificial intelligence, where we stand, what's going on.

MANZI: I'll do my best.

KRISTOL: I'm sure it will be excellent. So you run an AI, artificial intelligence, software studio. That's your current business. What is that?

MANZI: Basically it's called Foundry.AI and a studio is, in effect, a factory to build artificial intelligence software companies. So we ideate, build, grow, launch and spin out AI software companies.

KRISTOL: Okay, so what is artificial intelligence? I'm sure endless books have been written about this, but for sort of a layman, what's the right way to think about what that is? And then where do we stand in that much heralded progress of this scary or promising or both thing?

MANZI: Sure. Well, you know, I guess you can't start talking about AI now without recognizing that we're at the peak of a crazy hype cycle, right? So there's a lot of nonsense talked about AI right now. And we've been here before. And for those of us who spent a long time in this industry, there's a cycle of hype and then what is called "AI winter" when people sort of say "Well, you talked all about all of this stuff and it all sounded great, but reality is, certainly from a business perspective, we couldn't figure out how to really use it to make money and it's much more limited than the dreams you were selling us."

And those of us in the industry typically call it, that's "when the tourists go home." So that is clearly at some point going to happen in the next few years. I don't know when. But in the background we continue to make steady progress, ultimately underwritten by Moore's Law, which says crudely that we double information processing productivity every 18 months.

And the way that historically worked was you would double the throughput of a chip every 18 months because about every 24 months you could pack twice as many transistors on the chip and in that period you would make each transistor a little bit more powerful. So you doubled productivity for 18 months.

So if you then think about what is the definition of artificial intelligence, which is always the starting point for discussing anything intelligently, the textbook definition is a computing device that can act in the way a human being thinks.

And when you discuss in practice what that means, there's this concept of AGI or Artificial General Intelligence, and then much more specialized applications of artificial intelligence to solve specific narrow problems.

There has been ongoing work around this concept of artificial general intelligence since – people have talked about it hundreds of years ago – but there's been serious work since the 1950s. I think there's general consensus that we have not achieved anything approaching artificial general intelligence.

There's a famous saying which is, "I'll believe it when a computer can walk into an unfamiliar house and make a ham and cheese sandwich." And so we're not very close to that.

And there has, however, been enormous progress in using what we could call AI to solve specific narrow problems. However, defining "narrow" is getting broader and broader, so we're solving more and more difficult problems.

But I think at current course and speed, and there can always be a breakthrough tomorrow morning that we can't anticipate today, at current course and speed we're not very close to the synthetic human that can walk around and act and talk and behave the way a human being can.

KRISTOL: And I guess I had the impression, but maybe this is just wrong, that artificial intelligence is a little different from just more and more computing power. That it is somehow – one of course understands given what we've seen over the last 20 – more than that, but 20 years for the layman in terms of the internet and the mobile device and stuff – that yes, everything is much more powerful than it was.

MANZI: Right.

KRISTOL: Google Maps is better than whatever you could get in terms of directions ten years ago. But is there something that distinguishes artificial intelligence or machine learning? Is that sort of a similar idea from just more and more computing power?

MANZI: So there's this – in the introductory AI class at MIT, you go in and the first day, the first lecture, the first slide the professor put up was: "What is artificial intelligence?" And so we're all kind of waiting to see the answer to this because he's a giant in the field, this is one of the guys who really created, wrote the fundamental textbooks on it. And the first thing he puts up on that slide is: "If it works, it's not AI."

And so it's a flippant way of making a profound, or an important point anyway, which is you can think of an ongoing frontier of our ability to turn processes into software. And right at the edge of that frontier is what we call AI. And in 19 – 50 years ago, what's called a compiler was AI, which is now just, it's a piece

of code that everyone treats as something that's way down in the stack of code that everyone understands is kind of normal course of business.

And my view of the answer to your question is, no, there actually is no distinction between what we mean by artificial – at least at a higher level of abstraction – what we mean by artificial intelligence and ongoing productivity improvements.

The slight friendly amendment I make to that is, what really happens is ongoing productivity gains create the latent capability to automate certain kinds of processes that feel like human thought, that we previously couldn't have automated.

So you both need the latent capability, but then you need to do all the work to turn that faster processing power and, as I'll go into in a second, cheaper data storage and so on, into processes that weren't economically feasible to solve previously.

So at the same time information processing productivity is doubling about every 18 months to be a little overly crude about it, at the same time our ability to transmit information is getting much more economically productive.

So there's something called Neilsen's Law which is less well known than Moore's Law, which says crudely that we can double our productivity of *moving* data every 24 months. And there's, to my knowledge, only one productivity growth rate in the world that is sustained economically significant and faster than Moore's Law, which is Kryder's Law, which says, again crudely, that we can double information *storage* productivity every 12 months.

So what's happening is the productivity of the digital economy, as everyone has observed, is racing out ahead of the real economy. So if we double every 12 to 24 months different components of the information economy, a mature developed economy – North America, Western Europe, Japan – if you're lucky, you're growing productivity at about a couple percent a year which basically doubles productivity every 420 months, right?

So this one part of the economy is racing out ahead of the rest of the economy. And even within that digital economy, what's happening is our ability to store data in an economical feasible way is getting out ahead of our ability to process it.

And this is the real technical root of the prior big, giant craze called "big data" from a few years ago. That's really what it is. And so a lot of current AI, a lot of the most useful applications, what they're really doing is taking big data and being clever about turning it into stuff we can analyze.

And so what a lot of that really boils down to is taking words, pictures, sounds and what's called IOT Exhaust, which is a slang term for digital information coming out of devices. So we have sensors on all these machines now, right, the sensor data that's coming off of it. The refrigerator in your house is throwing off a lot of data now, probably, depending on how modern your refrigerator is. So we have immense amounts of data coming off machines.

And that is really where the high volumes of data are coming from. Like if you have an iPad, or on your computer, if you go to the little thing that shows you here's what my storage is used for, it's never like books and Word documents and spreadsheets. It's – if you're like me – it's movies, TV shows and music, right? It's words, pictures, sounds.

And probably the white-hot center right now of AI is something called deep learning, right? The killer application of deep learning is effectively ways of converting words, pictures, sounds, stuff like that into columns of numbers so we can analyze them the way we can analyze anything else. So that's really kind of, if I were to give someone a picture of what's actually happening, that's really what's going on.

KRISTOL: So give me an example of what – so what is analyzing this columns of numbers mean in terms of, I don't know, Google Maps and Waze or Google Translate or all the kind of –

MANZI: Yes. So there is a major breakthrough, practical breakthrough, five, six, seven years ago. There was a paper published I think in 2013 on this. And what they did was to convert words to numbers. And specifically, every word in the English language is converted now to a column of numbers. And the fancy way to say it is a vector. So the term for this is word-to-vec.

And the way that was done was to take a huge corpus, body of text. So you can take like all of Wikipedia or a snapshot of everything that you can get from the Google search engine in a given day or something like that. And you go through and you see something like, there's a sentence, "I took my dog to the beach."

And you then look laboriously for every other occasion in this whole corpus you see "I took my *blank* to the beach." And what you'll find is some instances of *dog* and some instances of *ball*, "I took my ball to the beach." "I took my child to the beach." You're unlikely to see "I took my synchronicity to the beach," right? And in that little way, *dog* is more like *ball* and *child* than it is like *synchronicity*.

And then if you're very boring and nerdy, and you have thousands of servers and Moore's Law has gotten to the point where you can process information right fast, you can methodically go through every possible occasion where you've got "two words or three words, blank," "two words or three words" and find what word gets filled in if these are the two or three words before it and these are the two or three words after it.

And from that you can basically build a numerical, you can create a number which says how like *dog* is *ball*, how like *dog* is *child*, how like *dog* is *synchronicity*. You can do the same thing for *ball*, *synchronicity* –

And typically in a technical version of a language like English you have millions of words, right? I mean, we really have – the OED has, I think, 150 or 200,000 words. But when you actually take *house* is different than *houses*, use every technical, every name of a chemical compound, etc., you get to millions of words.

You can basically build a vector for every word which is its relationship to every other word. So you have a vector of a list of numbers 6 million long for each word.

Then what you can do is do a bunch of involved arithmetic to kind of compress that down so you don't really need to keep millions of numbers per word. The current state of the art is maybe 300 numbers per word.

Because if *child* is a lot like *kid* which is a lot like, etc., and they are very unlike all these other words, you don't have to hold all the information, you don't have to hold all the numbers to keep most of the information. So you lose a little bit of information, but bring the words way down, bring the count of numbers way down.

So you end up with vectorized words. So now I have every word represented by a list of say 300 numbers. So why would I bother doing that? Who cares?

So in this paper in 2013, what these guys demonstrated is if I take the column of numbers for *king* and I then take the column of numbers for *man* and I just subtract the column of numbers for *man* from the column of numbers for *king*. And then I take the column of numbers that represents *woman* and I add back to it, I'm going to end up with a list of 300 numbers, right? All I did was I do dot product multiplication, but essentially I add and subtract. I've now got a list of 300 numbers.

If I go to my database of millions of words and say, "What list of numbers is closest to this list of numbers?" Now look up what word that is, it's *queen*. Right?

So, *king*, minus the concept of *man*, add back the concept of *woman*, I have *queen*. Nobody ever taught it that. Like, no one ever said, "There's this rule which is there are rulers of the world, and there's genders of rulers and so on." It just observed that. And it feels like –

KRISTOL: So it observes it from the context in which these words are used in a sense.

MANZI: Exactly. And the term is vector embedding. So, you're embedding words into this multi-dimensional space. And it's a lot like why someone who's read widely can often kind of guess at the meaning of a word, we call it "reading in context," right? Because I kind of see the hole it's filling, and I kind of get like that would have to be sort of something like this.

KRISTOL: Right.

MANZI: It's very much like that in concept. But in reality, it's just grinding away against a huge body of text in building statistical relationships. And it turns out that it's not just a neat parlor trick, but is incredibly powerful. And underlies –

KRISTOL: Yes. What does this do for us, though?

MANZI: So, it underlies things like when you are on a website and you need customer service and you start to work with a chat bot. You basically say, I have this problem with – you know, "I bought the flashlight, and I put the batteries in, and it doesn't turn on."

Well, it would be as a practical matter not feasible to sort of figure out all the possible ways somebody could say that, and figure out responses. But if I can operate at the level of concept, I can interpret in a practical way, oh – that's really what they mean when they say this. Close enough to be able to say, oh, this is the correct answer to give back to them.

KRISTOL: The person interpreting this is not a person, it's a bot.

MANZI: Exactly. It's a piece of software on the other side.

KRISTOL: Which is a lot cheaper than a person.

MANZI: A lot cheaper than a person in a whole lot of ways and will tend over time to be actually better, more accurate than most people. And it analogizes, not identical, but analogizes to, you know, in decades ago you could write code that could beat a child at chess and then you could beat a good high school player and on and on and on and on. And eventually beat Gary Kasparov.

We just got to the point where the computer beats the best person or player in Go, which was a more, deeper, challenge. Those are worlds of finite defined rules, which is why they're much simpler problems, but that same analogy happens. You will be better than a randomly picked person at being on the other end of the chat bot, then you get better at someone who's trained but not very good, and eventually in theory you get better than the best person you could have doing that.

And that capacity that we're talking about in this example of being able to operate at the level of concepts with words, permeates and is increasingly permeating anywhere and certainly in a business where you're dealing with words.

And so a good example of this is Google Translate. So, they're a package of changes they did, but what I'm describing is probably the most important one. Several years ago, Google Translate *suddenly* got a lot better and many people noticed this. And it really was this transition from very well thought out intelligent, to oversimplify, rules-based approach, to really this purely statistical observed word patterns approach.

And interestingly, it's a famous observation among libertarians that language is a classic example of a complex human system that nobody planned. And interestingly, when you analyze it in that way you actually do a much better job understanding its implicit rules, which may or may not be coincidental, but it's certainly an interesting observation.

KRISTOL: So, yeah, and so the Google Translate, there was a good article on this in *The New York Times* magazine a couple of years ago, which I sort of followed, I think. Yeah, the key was to go to this, I don't know what we'd call it, sort of bottom up, context-based translation that the computer can do because it has so much information.

MANZI: Yes.

KRISTOL: As opposed to, here are the rules of French or the rules of English. Where we have to program in definitions. There's none of that.

MANZI: That's exactly right.

KRISTOL: Explain that. I mean, I'm sympathetic since I'm also sort of Hayekian on these kinds of things. And I guess I never saw this as, thinking about this conversation, as I did in Hayek, Hayek actually was interested in this personally, right? Didn't he write a book or an article or anything, like on neuroscience?

MANZI: He did. Oh, he wrote a very involved book on neuroscience and its relationship to this theories.

KRISTOL: Yeah. So he even had an insight that this was –

MANZI: He definitely did.

KRISTOL: – that this was similar, right?

MANZI: Absolutely.

KRISTOL: And it is sort of against the fatal conceit of top-down rational control of central planning. Am I oversimplifying this a lot?

MANZI: No, I don't think so at all. No, I think that's exactly it. And my view is, as not a scholar of Hayek but someone who's read some things he's written, I absolutely believe that.

KRISTOL: But this breakthrough on Google Translat, let's just say, is – so it's different from top down rules-based. But it's still – how should I put this – it still depends on the quantity of information coming in.

MANZI: That's right.

KRISTOL: And in some respect, it's just the more information, the more the faster computing, the faster data storage, etc., or the greater data storage, the better you can do. Is there a real qualitative change? I guess that's what I'm sort of – does that make sense?

MANZI: It does. And there are several parts to that question. So, I think first of all, an analogy to Hayek actually in this case, a market is in a sense a top-down planned system. It's just a very, very, very flexible general set of rules. Right?

KRISTOL: Right.

MANZI: We have rules about, like, you can't kill people, and you can't cheat, and so on.

KRISTOL: Yeah. Keep your contracts, right.

MANZI: Exactly. So, you always have to impose some structured rules. So if you think about it, in building those word vectors, I impose theories about the world.

And in fact where a lot of very current work is – "Well, actually, *dog* isn't static. Like, *dog* used in this way can kind of mean this, and be more similar to *bother*, *nag*, etc. And in this other way actually it's very different than that. And *nag* actually can sometimes be really related to a horse, but sometimes be more like *dog* in this meaning." Right?

So there's this idea of sense-to-vec, like the sense of a word is vectorized. And how do I take a sentence of a list of words, and if I've got a list of vectors for each of these words, how do I combine them into meaning? And so it's an ever-receding horizon, all of which involves some degree of imposition of structure and theory. It's just the more flexible it is, the broader the range of situations it can be applied to.

So, I do think that it is – difference in degree becomes different in kind eventually. So is it a qualitative difference? It is – when you're down, certainly the way it feels like when you're down there, I mean, I wrote dozens of lines of code yesterday. And when you're down at the coal face doing it, it definitely does not feel like somebody flipped one switch and suddenly the world became different. It's just –

We can now, to be practical about this, we have this latent capability through faster processing speed and cheaper data, and we're trying to figure out as engineers how to exploit it.

And so to be practical about what I mean as applied to the example I just gave you, now because I can get easy access to Amazon Web Services, which is large servers I can get access to, and the cost of those keep dropping, to go from "word-to-vec" to "sense-to-vec," now is more economically practical. Like I can do it even in a way that I couldn't a few years ago because things are enough cheaper.

So, just the raw resource being cheaper doesn't solve the problem, but it enables me to do engineering work to solve the problem. And it is this constant accumulation of these changes that lead to these differences that feel so different for us as users.

And there are – it's fractal. There are some changes, like the guys who figured out word-to-vec, which really felt like, okay, that was a pretty big jump. But in the context of, you're talking about AI changing the world, it was actually, in fact, one step in many, many, many steps. I don't know if that was a helpful answer.

KRISTOL: Yeah, yeah. So to take the translate one, just because I read this article on Google Translate and I do use it sometimes. And I tried to study languages and I didn't have great success, so I'm sort of vaguely interested in it.

So in the old days, to oversimplify a bit, you had a sort of – what you could do online, you know, you had a pidgin English almost like word for word translation. It wasn't that different than if you were sitting with a dictionary looking up every word.

MANZI: Right.

KRISTOL: It was faster to go to Google and let them look up the words for some article you didn't quite understand. And then you could kind of make it out just the way you would if you were not great at a language but had a good dictionary and stuff.

So now, as I understand it, and I think it's true, if you use Google Translate for something you get something much closer to what would be the product of a good translator translating whatever random article you're reading in a French newspaper, or a German magazine or something. Or Chinese for that matter.

And I suppose if you just push that ahead, at some point you sort of – what does it mean to be a translator? I mean, that's what the machine – the machines, everything gets translated at a level that's very close to –

MANZI: I think that's very likely, yes.

KRISTOL: Constance Garnett or whoever the famous Tolstoy translator –

MANZI: Exactly, that's exactly right.

KRISTOL: And I think of that, didn't she translate Tolstoy? Whatever the most famous translators are. Right?

MANZI: Yes. You start by beating kids, then you beat good high school chess players, and then you should beat Gary Kasparov. I do think that it's a finite and bounded enough task that I think it's highly probable that there really won't be much of a job classification called translator anymore.

KRISTOL: And there won't be much of a barrier, I guess, then to reading anything. You don't have to go study a language. Whether or not you should is another question. But you would just go online and you see something in a language which I suppose you don't even have to recognize, someone tells you this is a good article on whatever – economics. And it's some language that looks different from, and it's not even in Latin script. And you just hit a button and you pretty well count on getting a pretty good translation of it I suppose.

MANZI: Well, and beyond that, I think that's – it's not like we can do that now. And just like you can get breakthroughs, you can get sudden stoppages that, you know, Moore's Law might stop. We might stop. Moore's Law is not a physical law of the universe; it describes human ingenuity, it might be that Moore's Law stops.

KRISTOL: And has Moore's Law – this is a sidebar – continued at the –

MANZI: Yes.

KRISTOL: So, I remember first encountering Moore's Law, George Gilder is a huge writer about it. I mean, I haven't studied it in the computer science way, but like in the '80s he was writing, taken with it, I think. And a lot of people said, "Oh, it can't keep at that pace forever. I mean, obviously the first ten or twenty years you go at this fast rate and then it slows down." And empirically has that been the case?

MANZI: Not so far. And so, when I was at Bell Labs in 1986, people were talking about, "Is Moore's Law going to start slowing down, or reach physical limits?" And we don't really do it with so much with transistors anymore, etc. It must at some point, by laws that –

KRISTOL: Laws of physics for example.

MANZI: Right. It will eventually – it's impossible that it continue forever. And its imminent demise has been predicted annually, from 1970s.

KRISTOL: That's pretty astounding if it's been – for those years.

MANZI: It's unbelievable. It truly is transformative. So I don't know. At some point it must. I have no ability to predict better than anyone else can whether it will or not.

KRISTOL: Anyway, you were about to say on the translation example, that it's not just –

MANZI: Yes. So, I think beyond that, I think it's extremely plausible that we will just have a small device, so you can just talk to anybody. Because if you think about it, it's not really hard for me to go kind of – sound-to-word is pretty straightforward to do now. And word-to-sound is pretty straightforward. And if I can do language-to-language, I just say it in English, it projects it to this person in French or Swahili.

KRISTOL: So you're visiting China and you have this little microphone kind of thing or whatever it is, computer microphone.

MANZI: Whatever the device is, however you –

KRISTOL: Does the other person have to have it too? I guess not even really, right?

MANZI: No. Because if you think about it, it's –

KRISTOL: Do you have to tell it ahead of time I'd like this to come out in Chinese, not Urdu?

MANZI: It would be so straightforward to – basically, if you've got the mic, to figure out what the language is on the other side. And so in a couple of seconds, you could say, "Oh, these people are speaking Urdu around me, not people are speaking Mandarin."

KRISTOL: So you don't even have to tell it.

MANZI: Exactly.

KRISTOL: The device will figure it out.

MANZI: Like it's not as if there's someone who has got this device in production and just hasn't brought it forward.

KRISTOL: Then you should do that.

MANZI: [Laughter] Other people have thought that too. [Laughter]

And you know I think that that is a completely plausible thing. And I wouldn't even put a timeframe on it. But I think that is foreseeable with kind of linear, the kind of rate of progress we have in known areas.

KRISTOL: So that does change human life a fair amount.

MANZI: You would think so, yeah.

KRISTOL: I mean, the whole notion of being "a stranger in a strange land" becomes sort of – being a tourist, being a foreigner. It really does become less – I mean, it's still different.

MANZI: Right. And you have – there's the whole sort of idea is the mythology of the Tower of Babel important? And is it useful in the same way it's useful to have kind of evolutionary isolation so that you don't have one common gene pool and you lose biodiversity. Is it the case that you would create too much homogeneity? And the species, not just in competition with sharks or anything, but in terms of its ability to deal with shocks and changes, would become less adaptive, or more so?

I mean, it would be a pretty significant change, in my opinion, yes. Language would almost – language difference could almost, it wouldn't go away, but it would become much less important.

KRISTOL: Is that true in other areas, or is this a distinctively true of language because that's an area in which is more directly affected by this, by these games and artificial intelligence?

MANZI: It's certainly an area where we see right now there being significant changes. There are others too.

But as I said at the beginning, I don't think, unlike in this area and some other areas where you can extrapolate current rates of progress and basically current fundamental methods and approaches and see how you could get to these changes, I don't think we're close to – again, subject to – There are many 50 year old guys in history who have said, "We'll never fly!" And the next month you prove them wrong.

So there can always be a breakthrough, but at current progress and approach I don't think we're anywhere close to a kind of general purpose intelligence that would affect, that could do what human beings can do across a very broad range of areas.

KRISTOL: But in specific fields?

MANZI: Definitely. Like I'm giving an example. Yeah, so medicine is another – parts of medicine are being changed right now. I mean, it's not theory. There's enormous venture capital being poured in right now into a wide variety of companies specifically around taking images and turning those into diagnoses.

A famous early example was a smart phone app that I point at a mole on my skin and estimate the probability that it is, in fact, melanoma. And it led to an important case where the FDA had to decide whether to regulate this as a medical device and so on.

And a generalized version of that, which is starting the simplest case with pure imagery data. I take a photo of something and much as a diagnostician would look at an x-ray or look at a photo and say I think this person has diabetic retinopathy stage II.

You can now use software to do that, in many cases, as well or better as very, very good diagnosticians, not just typical diagnosticians. And you think about it, I can distribute that on smart phones and suddenly I can have the equivalent diagnosticians in the poorest parts of Bangladesh because everyone has smart phones now – many people have smart phones. That is changing medicine.

KRISTOL: Individuals can do this, can notice a mole and take a photo at home and –

MANZI: Or out in the field, or on the subway on the way to work.

KRISTOL: And send it?

MANZI: No, they don't send it anywhere. You push the computation at the edge. It's called edge computing where you push the neural network out into the device and right there in a few seconds, it comes back and says the probability. If you're a sophisticated user, the probability that it's melanoma is this, but typically it's kind of red light, green light.

And that is happening all over the place in medicine right now, and what it means to do diagnosis is changing. If you think about it, it's pretty straightforward to "Well, actually, if I'm a good doctor, for many things" – [I say] as someone who has never attended one day of medical school – "if I'm a good doctor, for many things, I don't just want to literally look at the x-ray, I want to look at your vitals and your chart." And this is structured data.

You can pull that data in and I've got notes, right, which people think of as the guy saying, "Mr. Schwartz was coughing a lot." But it's mostly things from radiology and pathology departments. Text. But if you think about what I was talking about, about text processing, you can take the text, turn it into numbers.

Take in – I'm simplifying obviously – your blood pressure and your results of the following blood tests and turn this picture into numbers just like I turned the words into numbers. Push that into one big model and make a prediction: does this person have this condition or that condition?

And that's not science fiction and that's happening today and all over the place in medicine. And it's going to start happening much more than it's happening now.

KRISTOL: And I suppose a slightly spookier next stage of it is just taking someone who has no symptoms of anything, but whose blood we can look at, whose genes we know, whose parents' histories we know, and say this is what's likely to happen to you five or 25 years from now.

MANZI: So that obviously, as you know, happens now. And you can have your genome analyzed and it will make predictions for you of higher or lower probability of some medical condition in the future. Right now that's extremely – it's not that it's crude, it's that the information available says something like – I'll make up an example. Your probability of colon cancer is three times that of the average person in the population. But again, I'm making these numbers up. It's probably like colon cancer is one percent. What I learned is that I have a three percent chance of colon cancer instead of one percent of colon cancer. And it's unclear what I am going to do with that information.

And we are so far – as a non-biologist – we're pretty far from being able, independent of computers doing it, or people doing it, or some combination, of being able to look at a genome and make – outside of quite extreme cases of a single allele driving a particular health outcome, most of which have already been discovered – we're pretty far from making reliable predictions about what's going to happen to you in the long-run future.

KRISTOL: But that's not out of the question.

MANZI: Nothing's out of the question. And I just don't know enough about that to tell you, but I think we're closer now.

KRISTOL: Yeah. But it sounds like even based on what you were saying about what's currently being looked at and invested in, you're looking at a pretty big transformation of a pretty big chunk of the U.S. – of the world economy, and of people's professions, and the way we go about doing things.

MANZI: Well, yeah, and I think it's common sense.

KRISTOL: And pretty – in the next, not that many years or decades. I mean?

MANZI: Well, I mean, I think, to your earlier question or the way you put it earlier, that is almost certainly true, but that's been true for the last forty years. Right? I mean, just think of how transformed our lives are by the presence of the internet, and the application of digital technology all over the place. Right?

What we think of as plain vanilla – "Ah, that's not AI. That's just some code on a computer." – That was AI, decades ago. And the increasing productivity of the digital part of the economy until it slows down will

continue to transform things. I don't think we're at some point on a hockey stick where – I don't think there's evidence, let me put it that way – that we're at some point on a hockey stick where suddenly that rate of change is going to predictably get, feel, a lot more radical than it has until now. I don't know.

KRISTOL: I guess the way I would put it is, and this is totally based just on my life – it does seem like the internet is a very big moment. And more than just a bit incrementally more of what we had before. Although obviously computing was extremely important, all those things that happened in the '50s, '60s, '70s, and '80s. And that then the mobile device is a very big moment. A little more than just, "Gee, the email is easier to get than it was when you had to go home to your computer or have a laptop, and go online." And the combination of those is sort of a big thing.

MANZI: I agree with that.

KRISTOL: So it's not quite hockey stick-like, but I don't know, isn't there a little more of an inflection point at some point? I mean I guess I wouldn't say, just personally, that it seemed to me that my life –

I got to college in 1970, I left being an assistant professor in 1985. The way in which one did research, the way in which one talked, the way in which one acquired information does not seem – and I'm trying to think – but doesn't seem to me was that different. Word processing came in halfway through, and I was able to do my dissertation on a word processor instead of an electric typewriter, which was marginally easier.

But I don't know, it just feels like, like maybe I'm wrong about this, but something has accelerated or changed more, approaching a qualitative change just because of – I don't know why, but just because –

MANZI: Well, I think for people who are in the subset of the population that do research for a living, the development of the internet and the web has been very transformative to their workplace – but that's to their work life. But that is a portion of the population.

KRISTOL: Fair enough.

MANZI: And you know, I often think that, if I think about the transformations that happened in my grandmother's life from literally moving in horse drawn vehicles to taking a jet to California, among many other things, they were pretty severe, too.

KRISTOL: And what I have always said, and maybe this is slightly blinkered because it's my own life, is that I've always disliked, and when I was even young I was intelligent enough to dislike, the "Everything is changing faster than it used to."

I think that was not the case. That is I think my grandparents' lives just *obviously* were transformed much, much more than mine at least had been until the last few years, and probably still is the case. And in that respect, you could argue that in transportation and some forms of communication, mass communication I suppose, the transformation is – I don't know what, 1900s to 1950s, maybe, something like that, 1880 to 1950, were massive. And then actually you could argue kind of slowed down.

MANZI: And many people have argued this.

KRISTOL: Right. But then in this other area we're in right now, we're in a much faster –

MANZI: That's right. So I think what, in retrospect, it seems to me has occurred is that you had, chemistry was kind of the science that underwrote a lot of the changes that turned into technology in the 19th century, and physics in the first part of the 20th century, and the combination of those two as those proceeded.

And that we saw certain sectors of the economy that ultimately were about moving molecules from here to there, really did change in ways that were incredibly important in people's lives. And physics as a discipline and as someone who studied physics and I realize now I kind of shot behind the duck because I studied physics because of how interesting it had been. Did it right as it was, in my opinion, entered a period of lack of progress.

Where progress has happened scientifically is biology and computer science if you want to think of that as a science. And you know, the plane I took back from Europe two weeks ago is really not that different from the one my grandmother flew in like a half a century ago. It is a little different, but it's not that different.

On the other hand, the son of the President of the United States in the 1920s got a cut on his foot and died from the infection. I mean, medicine has made tremendous progress because biology ultimately has made tremendous progress. And obviously the digital economy has changed. And that's where we're seeing progress. I don't know how to weigh one versus the other.

I think in terms of are we going to see an inflection point? I think the best argument for it would be to analogize to the introduction of – and many people have made this analogy – to electricity.

So you know, we invented and then actually built huge distribution systems for electricity. And the first thing you did was light people's houses. Which, when you read people describe what it was like to go from a world where it's just dark at night. Unless you light candles –

KRISTOL: Or rich enough to have servants lighting a thousand candles. And it's still pretty dark, probably.

MANZI: Exactly. Or, I can just hit a button and like it's light out at night.

That was a huge change. But it took decades until electricity and motors were distributed throughout the economy and you saw change appear all over the place.

I think there is some of that going on right now with the digital economy and more specifically with AI. How radical it's going to feel I think is unpredictable. I would be very wary of rhetoric right now, because we're in the overheated phase of rhetoric about how everything is going to change. But I don't know the future.

KRISTOL: That's a good analogy, if you think about how agriculture has transformed from – I don't know what the right dates are, since I don't know much about this – but the mid-19th century to the mid-20th century. How much living has transformed, as you say, with electricity or weapons, the atomic bomb.

The chemistry, physics side of things did change an awful lot. That's kind of what we think of as the Industrial Revolution, really.

And if we're in something comparable to that in a slightly different area, yeah, I don't know. And in the past, there were the fast moments of change. And some things took longer than the other things, and some parts of the society and the economy changed faster than others, I suppose, and others lingered in a sort of –

MANZI: Yeah, for sure. I mean, I think that that's all, I think it's indisputably true. And I think that typically what happens as far as I can see, is the sector that is creating the new technology changes first, in terms of its methods of organization and its application of the technology. And then it spreads to the rest of the society.

And that the last places to change are either the very high status parts of the society, for obvious reasons. Or things that are controlled politically. Because they each have the power to resist change.

And almost nobody, contrary to rhetoric, really wants change. They want to resist reorienting power relationships and ways of life, basically, particularly if they're high status and powerful and wealthy, or if they have the political ability to stop change from happening.

And so that's, I think today you see the more government dominated sectors of the economy – medicine, education, the government itself, etc. – are the most resistant to the changes and methods of work and organization that are propagating out from the technology industry.

KRISTOL: Oh, it's beginning to –

MANZI: It happens everywhere. Eventually the water will permeate everywhere. It's just they resist the longest.

II: AI and Politics (40:16 – 1:02:34)

KRISTOL: And how about politics itself, I guess? People like me who like political philosophy and who like the founders and the Federalist Papers and read it and teach it as if it's kind of as relevant as it was when it was written, with obviously some stipulations that things have changed some.

But I mean, how much is that kidding oneself? How much does the change – everything we've been talking about, really make those things at some point – you know, you read those books and it's interesting –

MANZI: Antiquarian, right?

KRISTOL: Yeah, it's antiquarian.

MANZI: I don't think at all. I don't think anything that has happened to date, to the year 2018, has had that effect at all. You know much better than I do, that you can – you can go back and Thucidydes, right?

KRISTOL: Right.

MANZI: And we were talking about this once. I was a child of the Cold War and you'd read it and it was this obvious but fairly awesome analogy. And it was like you were reading about the United States and the Soviet Union, because human nature had no history. Human nature was, materially speaking, the same.

I think that if there is – if human life in the year 2100 is more different than human life in the year 2018, then human life in the year 2018 is different than human life in the year 500 BC, it is unlikely to arise from what we think of as artificial intelligence, and more likely to arise out of biology.

If we can actually get to the point where we can engineer phenotypic differences by changing genome, which to me is not a crazy idea. I'm just not at all expert in it. But it doesn't strike me as in the next – rest of the century could we get to the point where we could literally go in and start moving molecules, a DNA strand, and affect particularly mental processes? Then all bets are off.

Then I think, you know, at some point you probably – are you even the human species anymore? You know, did you engineer the human species out of existence, or did what we call humans become so different that then everything before that becomes antiquarian.

I'm not predicting that'll happen. But if you said what's the most likely way those things would become antiquarian, it would actually not arise from computers and software, it would arise from biology.

KRISTOL: But aren't the two related?

MANZI: They definitely are.

KRISTOL: It's the computers and software that make possible, doing what you're –

MANZI: That's exactly right. And all of these – whenever you see technology changes that seem unrelated at the same time, they almost always reinforce one another. And so a lot of the way you do a modern genome-wide association study relies on very fast computation and Bayesian math and so on.

So yes, they would be kind of the picks and shovels that guys would be using to mine gold. But I'm talking about it's what you're pulling out of the mine that's going to matter, to stretch the metaphor here.

KRISTOL: And I suppose, and I really haven't thought seriously about this, but I mean, if drugs can relax, if primitive drugs can just – you take care of colds or whatever. But then more sophisticated, you might say psychological drugs can relax you. Then the next stage is a drug that makes you not feel fear.

MANZI: Right.

KRISTOL: And then the next stage is not a drug because that's sort of an imprecise way of talking about whatever is going on here, but rather – I don't know what it would be, a tiny, an actual monkeying with something inside you.

MANZI: Yes.

KRISTOL: But then again, that could be done in modern medicine probably without massive operation or intrusion, right? It does seem like what you're saying with the computer, that it's not necessarily there's a moment, an inflection point, or hockey stick moment where it's like "Whoa, we just crossed some – "

I mean, are there red lines? Or are there just a heck of a lot of gray areas that you sort of chug through and find yourself, you know, in a different moment?

MANZI: All, in my observation, all technological changes are a series of S curves, right? Where you're investing and it doesn't seem like anything is happening, and then you reach a point where you make rapid progress. And then that technique itself reaches its natural limits and you don't get much more progress.

And you can kind of zoom out to the Agricultural Revolution, the Industrial Revolution as this big S curve or somewhere on this huge S curve and it's fractal. If you start going into little pieces that you see, it's made up of an S curve. If you go inside that, it's S curves inside S curves.

In that sense, I do think it is all a series of shades of gray. It's all a series of gray lines. And at a high level of abstraction, you can say correctly what you just said, which is, "Well, there's nothing magical about moving the molecule that happens to be on the DNA strand using this CRISPR kind of technique."

Versus, "Well, I basically put in a chemical through your bloodstream that goes to that DNA strand and moves things, or changes the thing that the DNA strand would have changed because it would have coated for a protein differently."

So yeah, I think there is no – I don't think there will be a moment where someone is going to go from like pre-Newtonian physics to an atomic bomb, by analogy, all at once. But I do think that step by step it's not

crazy to think we would get to the point where it's not just – you know, the Greeks dealt with, "I eat this leaf and I feel differently." And, you know the famous thing if I have jaundice, I see things as yellow and human perception is changed by external stimuli.

I do think, though, that if you got the point where it's not *me* taking a chemical, or having a piece of surgery done that's going to change me. But when I was created at the beginning, I was created as something *else*. And an engineered process created a genome which is then going to build the parts of me.

So it's like, what happens if you say, "Well, it would be really great to have a tail"? So we can engineer humans to have tails. What if we could engineer humans to think in a different way?

If I'm drunk, I'm less afraid of going out and talking to someone at a bar, right? But what if I could engineer like a fearless creature? And I use those as soldiers? And I had a soldier caste? And I said – "Well, like am I really killing people with souls by having them go – I don't know, are they machines? Are they people?" You could see, I think, getting to that kind of a situation really pretty easily.

KRISTOL: On the shades of gray argument, to be clear, doesn't mean that it's less serious, or less dramatic, or less worrisome, if you want to think of it that was.

MANZI: I agree, right.

KRISTOL: It could be the opposite.

MANZI: I totally agree.

KRISTOL: I mean, if there were actual red lines.

MANZI: Yeah, it's easier to stop.

KRISTOL: Figure out ways to put up barriers.

What's actually – again, totally as a layman, looking at these developments over the last ten, twenty years, just even in the kind of simple ones, iPhones and stuff and so forth. What's spooky about it is it's not clear where one would have even put up a line if one wanted to stop X, Y or Z. And if they really are, if the widespread availability of these phones with everything that it brings with them, is causing more depression and anxiety among teenagers, which seems to be true. Do you buy that stuff?

MANZI: Yeah, I do. I have no scientific evidence for that, but –

KRISTOL: Seems like it's legit studies.

MANZI: And beyond studies – if you have kids and you've observed them using phones, and yourself using phones –

KRISTOL: So how would you even – there wasn't a moment where you could have, is there or isn't there? Are there or aren't there moments when you can say or society can say oh my God, we have to regulate this? It seems kind of hopeless, though.

MANZI: So it's a complicated question, right? And I think that there are always, even in a world of progress if you believe in progress, and I do, there are tradeoffs in every advance.

And many aspects of modern life, I think, increase anxiety. And I don't mean modern life as in 2018 versus 1998; I mean versus 1850. I still wouldn't trade my situation for the situation of someone in 1850.

And also, you have to think about those tradeoffs, I think, A [or] B. I think that – I often thought that the introduction of, say, drugs in the 1960s into the popular culture, it was like a virus entering a body that had no antibodies. And that people who are kind of well – basically people in intact families figured out over decades how to manage this issue. And you know, the society broadly ,built up of lots of people and families. kind of figures out how to like, it's really a bad thing if your kids are taking drugs when they're 15, right? And so you figure out how to manage it and the society develops rules and norms and so on to manage it.

And I do think that will happen with digital technology. I think people were very naïve about it, understandably, and just kind of let their kids use phones. And now you're kind of like, I would never in a million years just give my 11 year old a smart phone with no controls on it. And now we have all kinds of software to control what they're doing, etc. And I do think you learn to manage those things as well.

And the third thing, I think, is in a world of – without a kind of global universal imperium, the problem is one society is able to get, create a huge edge by implementing these technologies, it becomes very difficult to stop it because you kind of have to to compete at a certain point.

KRISTOL: I have two questions maybe to conclude with. One is that, let's just do that question to start with, because that is such a, one of the standard accounts of the ancients, at least.

This is wildly oversimplified, but they preferred, either in the Greek case, the *polis*, the small city state, the human size, you might say, community or in the case of the Rome, the Republic. Self-government. But you have an arms race, so to speak, and you can't afford to be small. And you can't afford, you think you can't afford, or you have to preemptively not afford, because someone else could come along. And suddenly you have the Roman Empire – this is a drastically, wildly oversimplification, but yes.

And that war sort of drives a certain kind of necessity of growth and of even technological development that you might have decided, if you were just a self-governing community in a, you know, without others around, to forego because it has, as you say, negative side-effects.

I mean that is a real question, right, out of the attempt to control the technology. If China has a regime that doesn't care much about individuality, or privacy, that was actually going to be my second question so I can fold it in, I mean, how much can the U.S. resist, you know, keeping up with them, I suppose?

MANZI: Well, I think that, as you know, the long run track record of democracies is pretty good in wars, right? And in long run competition with intermittent wars and peace. So I do not think, that – It's my view of the lesson of history at least to date is, you do not have to become a totalitarian society or anything approaching it to compete with such a society.

And I do, though, think that, you know, if Nazi Germany is fielding long range artillery and jet aircraft, you kind of have to do it, you know, or they are going to overrun you.

KRISTOL: And you don't have to go to a totalitarian society, you have to become as technologically advanced a society as your rival.

MANZI: That's right. And in certain sometimes in very specific ways you have to match the – if not the exact technology the type of technology. I do think that's true. And I think it's like gravity. I mean, I can wish it's not true, but I think fundamentally societies have to be able to survive in a difficult, often violent, world.

I can't think of an example, but I also haven't thought about it very hard, I can't think of an example of, here's this technology in the history of the United States from the basically signing of the Constitution to today, where we built a technology for that reason that we otherwise really wouldn't have wanted to do

outside of direct warfighting technologies. So, I don't know how practical an issue it would be or not. But ultimately you have to survive.

KRISTOL: Another question is, is sort of apart from the survival against foreign enemies, I mean, as an internal matter, something like privacy. Everyone vaguely likes privacy, but at the end of the day, if you allow certain technologies to just chug along, and certain businesses to do their thing and market their data, and so forth, do you end up having pretty fundamentally eroded any real privacy without intending to, or not? I mean, I don't know.

MANZI: Right. Well, I mean, whose intention, right? The intention of the people building these businesses was definitely to erode privacy. Right?

I mean, and I think that it's pretty disingenuous for, as background, for an adult to say, there's this thing, by example, called Facebook, that I use every day. And there's this giant campus with multiple buildings with people making enormous amounts of money. All of whom are, and running huge server farms, to deliver this thing to me that I don't pay money for.

Like, to not understand that, the expression in Silicon Valley is, "If you're not the customer, you're the product." Right? To not understand that what's really happening is, they're getting access to my eyeballs, just like free TV is just – really the role of a television program historically is, get me to sit still long enough to watch the ad. Right?

So, I think it is clearly the intention of those companies. In my view, in conjunction with the intention of the people who are the users, either on a conscious or a willful ignorance basis, that they were giving up privacy in return for a benefit.

And I think that in a free society the likely endgame we'll get to is privacy will become a property right. And if you want to – if you want my data you have to engage in a transaction, regulated the way any transaction is, against fraud and deceit and so on, to basically take – buy parts of my privacy from me, in return for either access to this good or service, or money. And I think that's probably where we'll end up.

KRISTOL: You don't think the technology overrides that attempt to construct the property right kind of model?

MANZI: No. I don't think so at all, actually. If anything it makes it more straightforward. Because if I think about interacting with a platform – and I'll keep using Facebook as an example, not to identify them as the only example of this – the fact that I'm engaging in a digital transaction which creates this digital exhaust, that creates this transaction record, actually makes it, I think, much easier to keep track of. "Okay, what information did I grant you?" And not "What access do you have to this digital information about me or not?" And therefore, to essentially manage the commercial relationship.

KRISTOL: But isn't there a kind of a collective choice problem? I mean, you can't really opt out of something that everyone else is in. It would be sort of like, I don't know, I don't want Google Maps to know where I live, but I don't know, how does that work in practice?

MANZI: Right. So I think what you end up with is the following situation, which is, first of all, someone could build alternatives to Google Maps. But you can say there's actually a problem of it only works right when a lot of people are using it. So Waze is a good example of this, right?

KRISTOL: And you want to use it in certain ways.

MANZI: Exactly. So you are then back in the late 19th century problem of a natural monopoly, right? And I think it's inevitable it becomes regulated the way a natural monopoly is regulated. If you really decide

there aren't technology based competitors that are realistic, and it's become integrated into the society the way, like delivery of water has to my house.

Like in theory I could be a crazy, sort of lunatic libertarian and say well, I could have an independent water company and build my own well.

KRISTOL: Go off the grid.

MANZI: Exactly. Which sure, you can do. But as a practical matter, what you end up doing is regulating like a utility, which I suspect is what is going to happen to those businesses, unless technology competitors make it clear that they're not permanent monopolies. I think the Googles and Facebooks of the world in that scenario will just become regulated monopolies.

KRISTOL: And so as with water and electricity and everything else that we take for granted, they will be everywhere. They will be taken for granted. But there will be certain conditions built in, I suppose, to try to preserve what we care about in terms of –

MANZI: Exactly. Or, that's one scenario. The alternative scenario is in fact they are not natural monopolies. They seem like they are, but they are more vulnerable to technology based competition than we think they are.

KRISTOL: Yeah, and that could be, right.

But what about – the final point, to stop the politics, there are people I think who are sensible and not alarmist normally who really look at it and think you know, everything in democracy, that world democracy depends on – deliberation, opinions, being educated, minds being changed by data or experience, sort of the whole way modern communications work, the way Facebook works, really cuts against that.

I mean, it's not just a fake news question or a kind of hard problems with questions about Facebook. But somehow they can figure out what you like. When they figure out what you like, in terms of reaction on the screen, they can adjust it to make you like it even more. You sort of do lose a certain amount of free will, almost. Is that crazy?

MANZI: I don't think it's crazy, but I think it's exaggerated.

I think that – this is an ironic thing to end on for someone who is talking about AI and working in AI – but to me the media environment now seems a lot more like the media environment in 1800 than it does in 1950. The technology then had highly diversified, specialized, highly partisan news sources that – I'm betting the guy who handprinted a bunch of pamphlets every week in one corner of Philadelphia probably knew his audience pretty well, and did his best to continue to do exactly what you're describing. And they also lived in echo chambers.

I think if anything the media situation of post-World War II America with extraordinarily concentrated, one-way broadcast media, in both print and television, was probably more the anomaly.

KRISTOL: No, I agree with that. And appealing to prejudices has always happened. And the attempt to construct echo chambers one can say is a lot of what human history is about, in terms of regimes and governments and religions.

MANZI: And I think as someone who is down into the technology of kind of "how I read your reaction and keep morphing what I'm delivering to you," like, yes, that's a real. And yes, it creates more personalization of that technology to you than otherwise. But we are really far from the sort of robot on

the other side of that interaction who's literally figuring out how to manipulate you and drive you to some outcome, beyond marginal changes.

KRISTOL: I guess the question is, yeah, so we're far from addiction, so to speak, as opposed to just, you know, opinions that you or some algorithm have figured out that will be well-received by this person, given the previous opinions and facts that he's received.

MANZI: I think that the word addiction is an interesting one. Because I do think –

KRISTOL: But you hear that word sometimes in this context.

MANZI: You do. And I think that there is a problem of addiction to social – to screens, for sure. I think that is a separate problem than this idea of, or it's related but not entirely overlapping this problem of constantly manipulating what you see based on your reactions in order to drive you to a particular conclusion. The latter – I think we can create marginal changes, which can matter, but we're pretty far from the kind of, "I'm subliminally turning you into someone who believes you ought to vote for X or Y."

KRISTOL: It's still enough of a reality that you can be bugged by reality and change your mind about something.

MANZI: Yes, and also, you know, people are not – in gross generalities – people are not as totally malleable based on these information as people think they are, even independent of, I can do reality checks.

KRISTOL: And on the addiction side, with the screens, do you think that's also manageable, so to speak?

MANZI: Well, I think it's a real issue, much like drug addiction is a real issue. And I don't think the society has yet figured out the combination of norms and technologies and so on to be able to manage that. I think it's a big problem, actually, to be honest.

KRISTOL: But just as you can't let people become addicted to heroin, because it just doesn't make sense anymore to say they're making a decision to take it.

MANZI: Exactly. Yeah, the best description, the best definition I ever heard of what is addiction, is from a cop. And he said, "My definition is, what will you do to get it?" Which I think is the perfect definition.

So what I'd say is like, "How will your kids," is the way I'm thinking of it, but also people, "react if you said you cannot use screens for a week?" What would their reaction be? What would they be willing to do to get access to it? And I think that that is at an unhealthy level with lots of people today.

KRISTOL: Outstanding. And so we'll have to have you back to continue this discussion. This is a fast-moving – when will we need to have you, this is a final prediction, when will we need to have a second, another discussion on AI where things will have changed enough that we'll have a genuinely fresh look? I mean, how fast is it moving? It's a stupid question in a way, but I mean?

MANZI: I'll make one prediction.

KRISTOL: Will this thing stand up two years from now, five years from now, fifteen years from now?

MANZI: Here's the prediction I'll make, which is –

KRISTOL: And of course it will stand up because it's a deep and interesting conversation, and these truths will last beyond mere technological changes. [Laughter]. But leaving that aside.

MANZI: Right. I think that sometime within five years, the nature of this conversation would be, "A? What happened? It was going to be such a big deal, and now it doesn't look like it is." Actually I think that – if the topic is around the popular zeitgeist, like I said, that's what the popular zeitgeist will be.

In the background, the kind of changes I'm describing will continue to grind along.

KRISTOL: Very interesting. Okay, well, we'll have this conversation in less than five years because this has been so interesting.

MANZI: Great.

KRISTOL: Jim Manzi, thanks very much for joining me today.

MANZI: Thanks for having me.

KRISTOL: And thank you for joining us on *Conversations.*

[END]