

Transcript of "Exclusive AMA W/ Noam Chomsky on Jung, Wittgenstein and Godel - TOE
<https://www.youtube.com/watch?reload=9&v=pUWmTXkpHjE>

Noam:

I'll tell you, if we have a minute, a short anecdote. (Sure.) I had an old friend of mine who was a really great philosopher, one of the most outstanding. He used to also teach undergraduate great ideas courses - freshman courses which go from the Greeks to the present on everything, and we were walking across campus once, I was walking him over to his class and I said, "How can you teach a course that covers everything?" He says, "Well I just start the course by saying, "Ask me anything."" (Chuckles)

Curt: Intro

Many of you are likely new to this channel, and as a brief introduction my name is Curt Jaimungal. I'm a filmmaker and have a background in mathematical physics - particularly the theoretical end of what are called Grand Unified Field Theories. This channel is dedicated to interviewing intellectuals on cognitive science, consciousness, philosophy, psychology as well as of course math and physics, as I delve into the variegated inner workings of the universe with a heavy emphasis on keeping it technical rather than simplifying. This is the third time I was blessed enough to speak with Noam Chomsky, and I thought that we take a different angle than the political nature of our previous conversations and instead open this up to an AMA an "Ask Me Anything," and cold questions from the audience, as well as professors. According to Noam this was the first ever AMA he's done, but a quick Google search shows that he did conduct one with Reddit about eight years ago, so I think he simply forgot. But either way, it went so well that we far surpassed the 30 minutes we had scheduled and will likely be doing another one in 2021. I'm also joined by my colleague Peter Glinos who has a background in evolutionary biology, history, and philosophy. This is probably the most academic conversation with Noam in years - at least in video form. It's scholastic, clinical, straight to the point because you're here for Noam, not me. So, enjoy!

I'll ask you some questions from professors first, and then we'll get on to general audience questions.

OK.

Professor Rebecca Goldstein, professor of philosophy, asks, "I would love to hear Noam Chomsky expatiate on what he thinks the implications of Godel's Incompleteness Theorems are, both in terms of mathematical realism, and in terms of our mathematically knowing minds."

I don't think there's any clear answer to that question. The technical aspects of the theorem are understood - within a particular language of sufficient richness, you can't establish the truth of expressions within that language. If you want to establish truth, you're going to have to keep going up and up an endless hierarchy of richer and richer languages. But we can understand it, so it's comprehensible. So, that's as far as our understanding goes. But what it tells us about the world - anything outside about the world, really goes back to the much simpler question: what does arithmetic tell us about the world? Where are the numbers? They're not in our minds. There are truths about the numbers, just the plain, natural numbers, which we somehow grasp. But what are the things that we are

grasping? Is it something in a Platonic universe? Is it something in a mental construction? I don't think there are any satisfying answers to these questions.

Are there any implications from Godel's Incompleteness Theorems to linguistics?

No, I don't think so. Linguistics does involve formal computational systems, but they are sufficiently elementary so that these questions don't arise. And if they did it wouldn't matter. So, they arise in physics which uses all of mathematics, but doesn't interfere with the physicist.

Professor of neuroscience Anil Seth says, "I wonder - given what I perceive as a longstanding skepticism of scientific progress in the understanding of consciousness, whether Chomsky has seen any promising shift toward good ideas in this domain."

It's worth thinking about it for a moment. In recent years, consciousness has been called the "hard problem," the real serious problem of philosophy and science. We might look at a little bit of history here. If we go back to the 17th century there was a hard problem - it was motion. Motion was what was called the hard rock in philosophy – philosophy meaning science. We can't comprehend it. Turns out, it was right. We couldn't comprehend it, and we still don't comprehend it - not in the sense in which Galileo, Leibniz, Huygens, Newton, the other great founders of modern science wanted to understand things. For them, intelligibility and understanding meant constructing a mechanical model for it. And mechanical model meant something with gears and levers and cranks and so on. Something like what was being produced all over Europe at that time by highly skilled artisans, amazing people with their imitations of human beings, duck digesting, the fountains at Versailles, and so on. And what was called the mechanical philosophy, meaning mechanical science, that was the basis for the scientific revolution, held that the entire world must be a massive machine of the sky. There was a problem of explaining motion within that system - that was the hard rock. Then Newton came along and said it's hopeless. Newton's theory crucially involved [?] forces that cannot be captured within the mechanical philosophy. Newton didn't believe it. He spent the rest of his life trying to overcome it. He regarded it as an absurdity that no person of scientific intelligence can possibly accept. Leibniz and others agreed. They accused him of reintroducing the occult properties of the despised Neo-scholastics. And he didn't disagree. Now that's why his *Principia* is a mathematical theory, not a physical theory. He was crucial on that. I don't have a physical explanation; I just have something that works. So, we can understand the theory, but we cannot grasp what it's talking about. What's important, just to keep it brief, is that the hard rock in philosophy was abandoned. It was recognized that we cannot comprehend, we cannot gain an intelligible universe modeled that meets our standards of intelligibility. So, it was abandoned, and science just reduced its goals to finding intelligible theories. I think the same is true of consciousness, and many other things. We just abandon the search for an intelligible universe. We try to find intelligible theories that will account for the phenomena and in that respect, there's progress in understanding consciousness. So, there are better theories about some of its properties and so on, and that's the most that we can aspire to. We're not going to achieve for this hard problem what we haven't achieved for other hard problems. It's an illusion. It was correct to give up the search for intelligible account of motion and move on to develop theories which explain it, onto relativistic theories and theories of gravitons, or whatever you like. We don't have an intelligible concept of the universe in the sense of Newton and Galileo, so I think the answer to the question about consciousness is let's find out more about it and develop some kind of theoretical account of what it. And I think there's quite a lot to

say about that. One thing which we could go into is that almost everything in our mental lives is beyond the reach of consciousness. Consciousness picks up little bits and fragments of what our mental life is about. There's a lot to say about that, but let's go on to your question.

So either in the field of linguistics or in the field of conscious mind, do you believe that we're working off a Ptolemaic model based on a particular assumption, what do you believe that assumption is, what do you think are the consequences? The Ptolemaic model was able to account for too many phenomena. It was able to account for the phenomena that exist, but if all sorts of different things had been true, you could build epicycles and epicycles and describe them. It was basically abandoned, partly for that reason, and partly because there was much simpler model. And that's the way science works. There may be systems that are so rich they could account for anything if you fiddle with this and that. Those are the wrong systems. We want the systems that account for what is. That exclude what isn't. OK. That's the way we try to understand things. So Ptolemaic model was never really refuted. You could always adjust it somewhere or another to deal with whatever came along. But the excessive richness and the extreme complexity led it to be simply abandoned in favor of much simpler theories that attempt to account for what there is, while excluding what isn't.

Professor of philosophy Daniel Bonevac asks, "Infinitival phrases are common but rarely analyzed in semantics literature. To take one of your famous examples, the students want to visit Paris, the students want blank. (I'm assuming this has to do with "want-to/wanna-contractions.") What kind of thing goes in the blank from a semantic point of view?"

In things like wanna-contraction?

You have a famous example: "The students wanna visit Paris. And then, the students want blank. What goes in that blank?"

Well, that has to do with consciousness. There is very good evidence that what goes on in your mind is, if you say something like, "who do you want to meet?" let's say. That's what comes out of your mouth that goes into your ear, but what's going on in your mind is an expression (actually a much more abstract expression) but something like, "who do you want *who* to meet?" "Who do you want *who* to read the book?" and that *who* in the middle there, is preventing wanna-contraction. When you say, "Where do you wanna go?," there's nothing there in your mind. It's just "want to," not "want *who* to." So, they differ in their, in what's happening inside your mind. It's more complex than this but that's the core property. So, it's one good example of how we're not conscious of what's going on in our minds. That's what's going on our minds. We have evidence of it from a lot of sources, one of them is wanna-contraction. So that just tells us we can't introspect into what's in our minds - lots of things going on. When you and I are talking, there's massive mental computation going on, all of it totally ahead of the reach of consciousness. We get little bits and fragments here and there and we call that consciousness. And this is a particular example of it. Actually, before pushing it to hard I should make clear that this explanation is only very partial. If you look further there's much more complicated properties. But I think that's the basis of it. A simple account is in the case of "what do you wanna do?," what's in your mind is "what you want to do," and "who do you want, who do you wanna read the book" it's "who do you want *who* to read the book?" So therefore "want" and "to" were not adjacent in your mind, and they don't contract. Again, that's the first step, basically indicates the kind of thing that's going on.

Professor Chomsky, the following question comes from Joseph Velikovski of Newcastle. So, there have been some developments in evolutionary biology in the field of memetics, which are different units of culture - behaviors that we pass down from one generation to the next. They've evolved. They follow evolutionary algorithms. It was popularized by Dawkins. Now when we look at memetics, there is a philosopher Daniel Dennett who posits that for understanding language, or in the quest to understand language, you can use memetics to understand certain linguistic algorithms. That is to say that language evolves in the same way that biological cultures evolve, gene cultures evolve. I was curious if you've heard of the field of memetics, and if you had any opinions on that field.

Well, this proposal that you mentioned has a number of problems. One problem is that it doesn't give an explanation for anything. Try to give an example. Take the example that was just given. Does it say anything about wanna-contraction? No. Does it say anything about anything else about language? No. That's one problem. The other problem is that we have some evidence about evolution of language, and it doesn't seem to work anything like that. We don't have a lot of evidence, but we have some... could run through it.

Just to give a counterpoint, or an example, to defend memetics for a bit, take for example the phenomena we notice when we have expats living in a colony, and the development of their language, versus the development of language in a meme. So, I'll give a very concrete example. Take the English, and then the American colonists. The American colonists spoke Rhotic English and over time the mainland population of English, the English that we know the people in Britain speak, it evolved into the Queen's English, while the colony population has sort of maintained Rhotic English. We notice a similar type of phenomena with the French and the Quebecers, the Chinese and the Strait Chinese of Singapore. Well, a very similar thing happens in biology in population dynamics. In fact, if you have a population, a mainland population, and you take out different individuals from that mainland population and you put them on an island, the mainland population will evolve at a faster rate than the island population. And they'll sort of be, not stuck out of time... They don't totally stop evolving, neither biology nor linguistic, but it's an example of an evolutionary algorithm that seems to hold with linguistics.

Except that has nothing to do with evolution. We have to distinguish evolution from change. Languages change all the time, but there's no linguistic evolution. Evolution means something that's happening basically to your genetics, to your DNA. That's evolution. And there's pretty good evidence that there has been no evolution of the language faculty for a couple hundred thousand years. This is a common misunderstanding, but change is not evolution. The American colonists and the British who stayed in the mainland had the same language faculty; it hadn't changed. If you had taken an American colonist child and raised him in London, he'd speak exactly like the people in London, and conversely. In fact, if you take an American kid today as an infant and you raise him in a tribe in the Amazon, he'll speak their language perfectly. The reason is, as far as we know, there has been no evolution of the language faculty ever since language emerged. And for that there's fairly good evidence. We know from genomic evidence that humans began to separate roughly 150,000 years ago (at least that much, maybe earlier, but that's the earliest that's been traced.) When you take the people who separated, basically the Khoisan people in Africa, they have the same language faculty that we do as far as anyone knows. There's just no evidence for that effect. What I said about raising an American infant in the Amazon, as far as we know that's true quite generally, so there's literally no evidence that the language faculty has

evolved at all since language first emerged. Which is not very surprising. These are very short periods of evolutionary time. 200,000 years in evolution is an instant.

I think we're confusing the tree with its fruits. That is to say the mechanisms of evolution - genetics - haven't necessarily (well they have changed as well, just as the language faculty over time has changed at one point, but we were well at one point over a grand scale and descendants from amphibians, when we evolved at some point to have language. So, the language faculty if it is intrinsic to us would have had to evolve at some point. But my point is that for both cases, the island population and the mainland population, the mechanism of evolution, i.e., genetics, is the same. The linguistic evolution, i.e., the linguistic faculty for the Americans and the Bostonians and the English was the same. But the software has changed, has evolved, in both cases.

That's just like saying that my language is different from your language. Which it is. I speak differently than you do. But we have the same language faculty. It's a serious mistake to confuse evolution with change. Now it might be that some of the phenomena in change have a sort of an analogy to things that happen in evolution. But that's just an analogy. It's a totally different process. The evolutionary change involves genetic change. A change in the way people behave - for all kinds of reasons that changes. So, in fact we can talk about language evolution, and there is work on that, but these other things are just the study of language *change* which may have some loose analogy to things that happened within evolution.

Andres Zuleta asks, "Are there extra-linguistic experiences, and how can we justify them if we can only express them through language?"

Well, I have extra-linguistic experiences. You can ask what do you do? So, for example if I'm trying to decide how should I go to work this morning. There's a number of different ways I could go. So, I can visualize them in my mind. I say, "If I take this road..." I don't say it. It just goes through my mind in imagery. If I go on this road, I'll run into a traffic jam over here. So, if I go that way something else will happen. All of this is just visual imagery. I could articulate some of it, but much of what is going on I can't even articulate. It's just a lot of complicated computation about how to do things. Well, is that thought? That's a terminological question. We might bear in mind famous paper by Alan Turing, the paper that initiated the field of artificial intelligence, his famous 1950 paper on machine thinking - can machines think? He starts off by pointing out that the question whether machines can think is too meaningless to deserve discussion. He didn't go on to explain why, but it's pretty obvious why. It's a terminological question. It's like asking, "Do submarines swim?" "Do airplanes fly?" If you want to look at it that way, yeah, they fly. Airplanes fly. Do you want to look at a different way? Yeah, people fly when they jump too high. Actually, some languages express it that way. But these are not substantive questions. We have a notion of linguistic articulated thought. That's a pretty clear notion. We understand to a certain extent the mechanisms that construct it. That create it. It encompasses the kind of thought that's used in inference, in reflection, in planning, and so on. There are other things going on in our minds which we can call thought if we want or we can call them something else, but they are of a different character. You could say the same about the - there happen to be two loveable canines - I can't use the word or there'll be a race under my desk - they have something like thoughts- maybe 10 or 15 of them - and I can list them very quickly by some words or some actions. But they can't do the kinds of things that we do with language. They can't plan, they can't reason, they can't imagine circumstances

and ask how they would act in them, as far as we know at least. So, do they have thoughts? It's a terminological question.

So, this question is from (if I'm pronouncing it correctly) Rivulet. "Then is language the substance of ideas, or merely the communicative apparatus? In other words, can a thought or idea exist in the brain without it being capable of verbal articulation? Can a thought or idea be perceived/recognized only if it has been verbally articulated?" Very similar to one of the questions before.

That's pretty much back to the preceding point. There are things that go on our mind -- my reaction to the particular shade of the color red -- is that an idea? Well, for David Hume it was an idea. Can I express it in words? No. So, I think we're in the domain of terminology. You can call it an idea if you want. You can call it a vivid impression if you want. But can we articulate it? Loosely, but not exactly. I can't convey to you my exact impression when I look at the color of the wall behind you. I can't convey it in words. So, do I have an idea of what it is? It's a matter of how we want to use the word idea. You go back to the way in which the word *idea* developed in modern philosophy and science. The term idea for somebody like say Descartes was just an entity in the theory of mind. A sentence could be an idea. A phrase could be an idea. An impression could be an idea. It's a theoretical entity within the framework of some theory of mind. It's kind of like *particle* in physics. Physicists can't really tell you -- in fact there are big debates about what a particle is - but whatever it is it's some entity which has a certain place within a theoretical framework, an explanatory theoretical framework.

Aro Own asks, "Eric Weinstein has suggested that similar to the property of language, we might have a Chomskyan pre-grammar for religious belief built in. For this reason, Weinstein continues to engage in Jewish ritual that is attending synagogue services and so on, while nevertheless being an atheist, at least identifying with atheism. Does Chomsky agree with Weinstein's appropriation of Chomsky's theory for the domain of religion?"

That's not really a theory; it's a suggestive analogy. So, there's some structured or religious belief in practice of course. We could work out what that structure is, both for particular religions and probably for religious belief in general. There are probably universal properties that are part of our nature which show up in different religious practices. When we work out these structures, and the rules that they follow, will they have anything like the properties of human language? That's a serious question. You can't answer it until you've worked them out. Chances are, maybe some could have loose connections, but probably not much. There are interesting investigations about other systems that appear to be cultural universals, and about how they relate to our language. For example, for 40 years now, ever since Leonard Bernstein's Charles Eliot Norton lectures at Harvard, there has been serious inquiry into structural and algorithmic relations between the structure of music, at least classical music (tonal classical music,) and the structures of language. And there are some interesting ideas about that.

Another relevant question is actually one that Rebecca Goldstein sort of brought up. What about arithmetic? It's a very interesting question. This is a question that much engaged Darwin and Wallace at the origins of the theory of evolution. They were very puzzled and debated the fact that all humans have arithmetical capacity. They didn't know that for a fact, but they assumed it, and it's apparently true. It's just part of our nature to understand that there are infinitely many natural numbers, that when you add them, it works this way not some other way, and so on. That seems to be part of universal human nature. And they were very puzzled by that because it couldn't possibly have been selected, since it was never used. It's only been used in a tiny recent period of human history, so how could it be there?

Wallace thought you had to invent some other evolutionary process beyond selection, and Darwin didn't accept that, but they never had an answer. Well, one possible answer which we can now formulate (don't know if it's true) is that it could be that our arithmetical knowledge (not the numbers - that's a different question, but our knowledge of arithmetic) could be an offshoot of language. Turns out, if you take the most elementary principles that yield linguistic structures, and you reduce them to their absolute minimum (a lexicon, which contains one element) you get the successor function, and something like addition. So, you get the rudiments of arithmetic. And it's possible that the reason we have knowledge of arithmetic is because we have language, and that this is just an offshoot of it.

Another idea that's been developed is that at some point in human evolution, probably roughly around the time that Homo-sapiens emerged, there was a slight rewiring of the brain which provided a computant, and a mechanism of computation of discrete Infinity...recursive functions that recursively generate an infinite number discrete Infinity of structures. And that this was then applied in language and applied in arithmetic may be applied in music, some think it was applied in moral systems. Well, these are all researchable topics. But when we go back to Eric Weinstein's question, in the Jewish tradition, his and mine, Judaism is pretty much a religion of practice, and not so much a religion of belief. There are things, prayers where we say, "I believe," but that's not the core of the religion. A religious serious Jew like my grandfather, for him, Judaism was his whole life. But it was the rituals. If you'd asked him, "Do you believe in God," he wouldn't know what you're talking about. These are the prayers that I perform, these are the rituals I carry out. It goes on all day, all my life, and that's who I am. It's the array of religious practices. Do you believe this and that – Yeah, I suppose so, but it's kind of on the side.

Jack McGreevy asks, "I was actually wondering if you could ask him about Philosophical Investigations, which was released in Chomsky's mid 20s. Was he familiar with it? Did it in any way to inform his development of the concept of linguistic competence, the intuitive knowledge of a language possessed by its native speakers?"

Well, I knew about the investigations as soon as they appeared and had in fact read The Blue and Brown Books as soon as they appeared, so I was familiar with the ideas. They had an influence, but not direct. There were ways of looking at things which were illuminated by Wittgenstein's rather aphoristic approach. His actual proposals, concrete proposals, in The Blue and Brown Books again, in The Investigations, about how language is acquired, just don't make any sense at all. So, they had no influence. (Why not?) Because language doesn't work at all like that. (Can you explain?) Well, I've actually written about it. If you look at these, he talks about how language developed...A couple of people are together. One of them points to it rock and says "rock," and the other one says a "rock," and then they develop a language that way. It's just not even remotely like that. That's just off the spectrum of discussion. None of these things happen. That's not the way language develops at all. In fact, the concepts in our mind, you can easily show, are much richer than anything that's presented. They're kind of elicited by phenomena, but a rich system quickly evolves. But on the other hand, when you look at, say, Wittgenstein's account of how you should think about language, like if you want to know the meaning of a word you should look into how it's characteristically used- that may give you some insight into the meaning of the word- now that's a valuable insight. In fact, it appears my own work. My own early work from the early 1950s basically adopted it - use theory of meaning – of roughly a Wittgensteinian style. It was actually more seriously influenced by the Oxford philosophers of the same

period. John Austin, Peter Strawson, Gilbert Ryle had rather similar views which I found more compelling and helpful.

While we're on the topic of Wittgenstein, do you make anything of his private language argument? Is there any relationship between that and your idea of "I-languages" or Idiolect?

Well, somebody who's interested in Wittgenstein's private language argument should first ask "What is it? What's the argument?" There's a huge literature about it, and there's no consensus on what the argument is. Take a look, say, at the Stanford Encyclopedia of Philosophy, a major serious source. Look up Wittgenstein - private language. You won't find the argument; What you'll find is a lot of exegesis about what the argument is supposed to be. I don't think anybody can actually formulate clearly what the argument is supposed to be. At least I can't, and I haven't seen it from anyone else. But did it influence I-language? No. I-languages is just, the term "I-language," I introduced in order to clarify terminological confusion. In the early years of generative grammar, the term "grammar" was used with systematic ambiguity. It was used both for the linguists' theory and for what the theory described. It was used for the object being described, and for the theory about that object. So, I suggested in the 1980s, since this was causing a lot of confusion, that we should make a terminological change: keep the term "grammar" for the linguists' theory, which is pretty much in accord with traditional usage, and for the object being described, call it "I-language," where "I" usefully in English can suggest internal, individual, and intentional (in the sense of a function in *intention* - not intention in the sense of Carnap - sometimes confused. A function in intention is the actual function, not the set of pairs that it relates, but the way in which it relates them. So, if you do arithmetic one way and I do it a different way, that we have different functions in *intention*, same function in *extension*. So, we're interested in a function in *intention* - what is actually is, what's actually coded in the brain. So, it's internal, individual, intentional, coded in the brain somehow. Our grammars are efforts to develop theories about it. (But that had nothing to do with Wittgenstein.)

Joe Siro asks, "If mental events are causally predetermined to physical events, (which themselves are attached to volition,) what does the data say about the relationship between conscious volitions and unconscious wiring, in relation to the problem of freedom of the will? What does linguistics say about this?"

Linguistics doesn't say anything. But there is a question about decision and choice, and consciousness of decision and choice. And there is experimental work – the famous Libet Experiments from 30 or so years ago - which show that there's a gap of couple hundred milliseconds between a decision and conscious awareness of the decision. Now they don't talk about complicated things like what we're doing, like making up sentence. It's not that. Just simple things like, say, lifting your finger. So, suppose I decide I'm going to lift my finger. Well, it turns out that the musculature, and the instructions to it are already being implemented before I'm consciously aware of having made the decision. What does that tell you about free will? Nothing. Just puts it back a little further. It says the conscious decision is maybe already determined. But what about the decision? No, actually the sciences tell us essentially nothing about this. What the sciences tell us is we can't explain it. What we can account for is things that keep to determinacy and stochastic processes - randomness basically. So, if it's within the framework of stochastic processes and deterministic processes, we can develop theories. Well, is freedom of choice within that framework? That's the question. But the sciences don't answer it. They can just say, "we can't handle it." There are some kind of exotic arguments in quantum theory and

relativistic physics. There is an argument that actually time is reversible, has no particular direction, could be going in another direction. For example, if an observer makes a measurement in the split experiments, it's determining the waveform's collapse, and it's becoming a particle. Well, we could go in the other direction in principle. So, the collapse of the waveform could have preceded the decision to make a measurement. Does that tell you there's no free will? I don't really think so, but it's kind of an argument, and it's about the only kind of arguments there are. The rest is just saying basically we can't handle it. So, if you think that the sciences are complete, then there's no free will because it doesn't fall within the framework of determinacy and randomness. But the question is are they complete? That's the question of free will. When you look at the study of voluntary motion, turns out there is extensive neurophysiologic study of voluntary motion. There's a recent article by two of the leading scientists who work on it - Emilio Bizzi and Robert Ajemian. It is a state-of-the-art article on what we understand about elementary voluntary motion. It appeared in *Daedalus* (Journal of the American Academy of Arts and Sciences) where they go through what we've learned about it, and they kind of end up by saying, as they put it, fancifully, that we're beginning to understand the puppet, and the strings, but we have nothing to say about the puppeteer. I can't say anything about decisions. It's a fact - just can't. So, you can believe what you like. We actually all believe that we are free to make decisions. I'm sure you believe it. I believe it. We could all be deluded, but there's no evidence that we are.

Boris Costello asks, "Is mathematics itself the domain of all languages, including of course, the natural and biological language?"

Well, there's a sense in which mathematics is the language of all sciences. They all work within a mathematical framework. But mathematics itself it doesn't tell you how an ant navigates, Right? If you want to study how an ant navigates, you're going to use mathematics. But from mathematics, you can't deduce how an ant navigates. These are the tools we can use to describe whatever there is. And it's the same with language. I mentioned before that language is based on a computational procedure which generates an infinite number of structured expressions - hierarchically structured expressions - which in fact, can express thoughts, or can be sent off to some sensorimotor system to be externalized. That's the core of language. Well, when you begin to describe that system of course you're using mathematics. You're using at least elementary recursive function theory, theory of computation, and more when you proceed. But the mathematics itself doesn't give you the answers to the questions any more than it does for a bee navigating.

Noam, I loved your book on anarchism, and in it you talk about the relationship between freedom [and] language, and one of the questions I have for you in that sort of frame is, is language created from the top down - sort-of guardians of language who create it, or is it being created from the bottom up?

We have to distinguish between two concepts which actually you mentioned before, possession of language and use of language. It's a distinction that goes back to Aristotle. He made a distinction between possession of knowledge and use of knowledge. A special case of it is possession of language, use of language. In modern terms it's called competence and performance. Now going back to your question, there is free generation. The language you possess is based on the principle of free creation. The theory of your language, which is a generative grammar, enumerates the possible structures that express thoughts, and are interpretable in your language. But that's not creative action. Creative action takes place in performance. So, what you and I are now doing is in fact a pretty high level of creation. You and I are regularly now producing new expressions - maybe new in our experience, maybe new in

the history of the language. They are appropriate to the circumstances in which were functioning, but they're not caused by them. There's nothing in what I'm looking at that causes me to make this sentence, as far as we know. Now we're back to the free will question. But as far as we know these are performances that are appropriate to circumstances, but not compelled by them. I'm basically quoting Descartes.

Just to clarify, this individualized language emerges internally, so when you're saying it's a creative element, it comes from a creative element within us, and that this is the ultimate germ of language?

Well, only in the sense in which having arms and legs comes creatively from within us. We are designed - our genetic endowment designs us - to have arms and legs instead of wings, and it develops through an ontogenetic process. It's affected by the environment of course, by your nutritional level, by your level of exercise, by all sorts of other things, but basically, we're going to have arms and legs unless there's some very serious pathology. And the same is true of language. It just develops. It grows.

Noam, in the past you've critiqued social construction, as when you're talking about Bakunin's red bureaucracy, for having the idea that you can mold human beings into a particular image, that there wasn't a sort of nature that would fight back in some way. It's they're not completely plastic. Do you believe that this is part of that nature, and the reason why human beings can't be molded into a certain shape?

That's a fair question. It's been asked for centuries. There is a rich tradition, basically Cartesian origins in many ways, that leads to classical liberalism. Take leading figures in the classical liberal tradition like Wilhelm von Humboldt, and Rousseau and others who I mentioned. They argued that we have at the core of our individual nature what was sometimes called a kind of an instinct for freedom. And they argue that it's linked to the creativity of language. This is speculation of course. You can't prove anything like this. But there is a creative aspect to human linguistic performance, the kind of creative aspect I mentioned. The speculation is [that] this is inherent to human nature, and any social system that constrains or restricts human creative impulses and independence is illegitimate. Out of that you derive classical liberal ideas, anarchist ideas, and their later development and so on. But if you want to prove it, there's no proof. It's just conception of what human beings are like, ideas about what language is like.

The next question is, "I'm curious if you have any opinions on Carl Jung's work, such as the persona, shadow archetypes, and is there a relationship between what's archetypal, and universal grammar, in the sense that there is an intrinsic structure that gives rise to patterns of experience."

Well, I was interested in Jung's work occasionally, and wrote about it a little, but mainly because of an interest in studying the question of unconscious mental activity. By unconscious I mean inaccessible to consciousness. There's plenty of unconscious things that are accessible. You can bring them out and think about them. Freudian psychotherapy is based on the idea that you can elicit them by the proper means. But what about inaccessible consciousness? In the whole history of thinking about this subject, I have had a hard time finding any clear examples of looking into inaccessible unconscious mental activity. Jung is one of the few exceptions. His archetypes are not accessible to consciousness, at least as I understand what he's writing. They are somehow there, they frame what we do, the way we look at things, but we can't find them by introspection. Well, if that's the case, if that's the correct

interpretation of Jung, and their tradition, then he might be an unusual, close to unique exception to the beliefs that what's unconscious is accessible to consciousness. That's almost a dogma of modern philosophy. With some philosophers it is a doctrine. To Van Quine, John Searle, and others, it's a principle that if it's a mental act, it has to be accessible.

So, one of our favorite anarchists, George Orwell, writes in 1984 about how Winston only has those cubic inches, or cubic centimeters, inside of his skull, where he has freedom. And I can't help it relate that to what you're talking about, with this internal language that's within us, that's inaccessible. Is there a way that that challenges systems of power - as it did in 1984 - if you notice that when you look at history?

We don't have any neurophysiological or other empirical evidence for it, but there is the evidence of history and experience. That's the kind of evidence that Rousseau, Humboldt and others drew from. I think we can make the case that humans have always been striving for freedom, and resist constraints on their activity. Now this can be suppressed, and there are very interesting cases of it. So, take something in our ordinary experience - getting a job. Suppose you're out of work, you don't have anything to eat, you look for a job. It's considered a wonderful thing to get a job. It wasn't always that way. If you go back to the origins of the Industrial Revolution, mid 19th century, take a look at the literature. The working-class literature that was very rich working-class literature. There was political discussion. The idea of having a job was considered a totally intolerable assault on elementary human dignity and human rights. Why should you be subjected to a master? Why should anybody spend most of their waking hours following orders given by a totalitarian ruler? That's what having a job is. It means you're following the orders of a master. And in the early stages of the Industrial Revolution this was regarded as not really different from slavery. In fact, it was called wage slavery. It was different from slavery only in that it was temporary, until you could become a free independent human being again. That was the slogan of the major working-class organization, the major one in American history, Knights of Labor. It was the slogan of the Republican Party. Abraham Lincoln's Republican Party held that to be subordinate to a master and under wage labor is intolerable - can't be tolerated. Now that's been beaten out of people's heads over 150 years, but I don't think it's far below the surface, and I think it can be elicited. And there are many other cases like that. It's the kind of thing that Gramsci talked about when he discussed how hegemonic common sense captures people and imprisons them and gets them to not comprehend their own natural instincts and desires. For a revolutionary, the first step is to try to unravel these kinds of constraints on thinking that make us automatically obedient and subservient, instead of asking, "Is that right?"

Slavoj Žižek talks about the revolutionary elements within the phrase, "I'd prefer not to." I'm curious if you have handy words or phrases, that have a revolutionary element, that we have forgotten, that we should learn again.

This is one. The idea that you should be subjected to a master during your almost all your waking hours, I think that is intolerable. I think American workers and the Republican Party were quite right in condemning this in the early stages of the Industrial Revolution. And I think we can work to overcome it very concretely. Worker-managed industries, for example.

Professor, you're famous for Chomskyan grammar (well you should be famous for it, it bears your name.) Do you mind telling the audience, what's the difference between the original conception of Chomskyan grammar/universal grammar, and how you conceive of it now. That is, how did it change?

Well, it begins around 1950 with the first efforts to construct generative grammars that did give a recursive enumeration of the expressions of the language with their structures assigned to them. That's late 40s -1950, and you have to look a little bit at the background. The background at the time was that linguistics was what was called a taxonomic science. It's based on procedures of analysis which you apply to a corpus of material, and these procedures identify the elements of the corpus, and their arrangement and organization. And that taxonomy completes the subject. But there's no explanation. There *was* a conception of what language is - the conception is it's a system of habits and training. If there is anything new it's by analogy. Well, generative grammar took off in a different direction. I should say at first, I thought it was just a hobby. A personal hobby can't be right because it's totally different from everything else. Over the years, the hobby became the thing that I thought was the field. It was within what later came to be called the bio-linguistic framework, that is, regarding a language as what we were talking about before, an "I-language." It's a trait, it's a property of you that you speak a variety of English, not a variety of Tagalog. A property of you coded in your brain. And one crucial element of this property is that it does generate an infinite number of expressions, each of which captures a thought, each of which could be externalized in one or another sensorimotor modality. Well, as soon as this enterprise was begun it was immediately discovered that we don't know anything. It was thought before that everything is known. Just apply the procedures and you get the answers. Turned out, as soon as he started writing generative grammars, you didn't understand a thing. There were problems and puzzles everywhere. So, the first task was to try to construct a theoretical apparatus rich enough so you could at least describe the data that was pouring out as soon as you began to study language this way. So, the devices were extremely rich. It was understood that that couldn't be right. The reason it couldn't be right is two reasons. Within the bio-linguistic framework you have to meet the conditions of learnability and evolvability. You have to account for the fact that a child acquires this system on the basis of very limited data. You also ultimately have to account for the fact that somehow it evolved. And a very rich, complex system just doesn't meet those conditions. So, the basic theoretical work over the past 70 years has been to try to move towards systems elementary enough so that they could have evolved, but yet rich enough in consequences so they can account for learnability. We now know that these problems of learnability and evolvability are much more serious than was assumed in 1950. The work in generative grammar set off a lot of research into language acquisition, and that work has shown that a 2 or 3-year old child has basically mastered almost all of the language. They don't exhibit it in their performance, but you can show it by experimentation, what they understand, and so on. So, that means the problem of learnability is extreme. There have also been by now careful statistical studies of the actual data available to children, and it turns out to be very sparse. There's a lot of sentences, but the same words repeated over and over. And you don't even get many bigrams, let alone trigrams. So, the problem of learnability is extreme. The problem of evolvability is also extreme. We've alluded to this. It seems that language evolved in a very brief period. Roughly at the time of the appearance of Homo sapiens, 2 to 300,000 years ago. Before that there's no significant archaeological evidence of symbolic activity altogether. After that there's pretty rich evidence, and as I mentioned, there's good evidence that humans began to separate not long after they'd emerged, apparently with the same language faculty. All of this pretty strongly suggests that language evolved in a brief period of time, an evolutionary period of time. So that means it had to be simple enough so that it could have

evolved, has to be rich enough so that it can account for the knowledge that's attained on very limited data. That seemed like a real conundrum. But theoretical work has been aimed at trying to overcome these problems, and also the lingering problem in the background, that on the surface, languages seem to differ very much, which can't be true if these other things hold. Well, in the recent couple of decades there's been the first real progress, I think, in solving this conundrum - finding systems simple enough so that they could have evolved very quickly but yet rich enough in their consequences so that you can explain fundamental properties of language with no learning. And, as for the variety of languages, it seems to be more and more converging on a conclusion that that's probably superficial. It has to do with the way language is externalized. Think of your laptop. Your laptop might have a program in it for say multiplication. But the laptop can be attached to any printer. The program doesn't care - you can use any printer that is around for the program to be printed. Externalization of language is kind of like a printer. The internal system doesn't care what printer you use - could be sound, could be sign, could be touch. The internal program stays the same. And it seems that the apparent variety of language is mostly in the printing, in the way it bridges the sensory motor system. What's internal seems to be very restricted. May even turn out that it's uniform for all languages. Could turn out. Can't show that now, but it's moving in that direction. So, I think that's the direction in which research is developing. I should say, not many linguists are really interested in these questions. (Why not?) They're practitioners. Like biologists. Not a lot of biologists work on molecular biology. Take a look at the articles on research papers in science. Most of them are descriptions about what this organism does in these circumstances. That's the overwhelming mass of the field. In fact, you go back 50 years and that was practically the whole field. There wasn't very much more. But it's the same here. These are special interests. Do you want to find genuine explanations for things, meaning satisfying the conditions of learnability evolvability, and dispensing with the variety? If you're interested in that, that's core theoretical linguistics. And there, there is (happens to be my personal interest) but there, there has been significant progress, I think. (There's not much time but if there was, we could talk about it.)

Alright, Professor, do we have time for two more questions? So, Peter will ask a question, and I'll ask one more question.

In that case, I just want to say, a lot of what you talked about also reflects in the field of memetics, when it comes to, you mentioned, how linguistics was taxonomical at first; it was just categorizing things. And the same is true when it comes to the field of biology until you get the universal theory, like Darwin's evolution. I just think that for linguistics, you've found genetics of linguistics, that is to say the mechanisms that underlie linguistic patterns through language instinct and etc. But I still think that the software, the words we choose, follow evolutionary and memetic algorithms to differentiate it from change, to differentiate it from change. In order for something to evolve, it has to be heritable, meaning it can be passed down from one generation to the next. We're not talking genetic heritability - people don't intrinsically know Spanish, or what have you. Not beyond a language instinct. It has to be heritable. It has to vary from one generation to the next. So, you have to have mutations of it, different states, different iterations. And then, lastly, it has to be selected, for or against. And by that we don't mean by some higher power, but that it's increasing or decreasing in its use. And to me, words, memes, behaviors follow these evolutionary conditions and therefore evolve. Are organic.

I think it's clearer if we drop the word evolve, which means something in biology, and use the word change. There's a lot of detailed study and investigation of how languages change. A lot of research into

that that has got many results. But it doesn't use these analogies, and speculations don't really contribute to it.

The critical difference between evolution and change is that it's iterative from generation to the next, that it can be passed on.

No, that's not the difference. The difference between evolution and change has to do with whether your genome changes.

OK, so I'll ask a question. There are two, and then you can choose which one you want to answer. Do you have any thoughts as to the unreasonable effectiveness of mathematics in physics? (I'm sure you've heard of that.) And then the other one is actually for me. This channel, as we were talking about before we started, this channel is about theories of everything. I'm interested in the theoretical ends of physics, and possibly even merging consciousness with physics with the fundamental laws of nature. And perhaps part of the problem is that the way science is right now. It's not complete, as you made reference to. But the problem could also be something else which I have no idea about. I'm curious, do you have any advice for me, essentially?

Well, as for the unreasonable effectiveness of mathematics, I think that's Herman Weil - nobody knows. There are some theories that say the universe is just a mathematical object. It's one of the ideas in physics, so of course mathematics will be very effective in dealing with it. It's just a mathematical object. You can believe that or not. But basically, nobody has any idea. Certainly, I don't. On what you should do, I think first we should recognize that we are organic creatures. We're not angels. OK? There's a kind of a dogma, doctrine if you like, that humans can understand everything. Somehow, maybe. Maybe not. What we can understand is based on our intrinsic nature. Our intrinsic nature yields scope to what we can achieve. But anything that yield scope most automatically yields limits. So, our capacity to, say, run, allows us to run, but it also prevents us from flying. That's true of intent intrinsic systems. So, it's very possible, if we're not angels, that if we're just organic creatures like the rest of the universe, that our intrinsic nature, cognitive nature, allows us to comprehend and understand certain things, but it will never allow us to comprehend others. They're just outside of our cognitive nature. This idea is sometimes ridiculed as mysterianism. To me it looks like truism. It's saying, yes, we're organic creatures, so we're going to be like other organic creatures. I can't navigate the way an ant can, because it has intrinsic capacities that I don't have. And I think that may well be true of our cognitive nature. We don't know. You can decide what you want. But as to the advice, the only advice is press it as far as possible. See how much you can understand. If you can link up theories of consciousness with fundamental physics, fine. That doesn't mean we're going to grasp consciousness in the way that some philosophers want to grasp it. What is it like to be a bat, you know? Can I understand what it's like to be you? No, I can't I never will. That's just not an answerable question. But can I understand what makes you a conscious being? Can I come to understand why your consciousness only picks up tiny fragments of what's going on in your mind? Yes, that we can understand. The question about "wanna-contraction" was a small example of that. We can come to understand what's going on in your mind that you're not conscious of, and can't become conscious of. These are all topics that can be studied. You can learn more about them. Maybe we'll find out where there rooted in our neurophysiology. All of that's open to investigation. How far it can go, you don't know. You don't know. There's no way of predicting where science can reach.

Thank you, Professor. I appreciate your time and please thank your wife as well.

There was one question this person had, if can I say it, and if it's too long to answer then we just forget it and I cut this part out.

It seems clear that there's a place for neologisms, in fact you've coined a few yourself. But it seems like there's now the invention of terminology with morality attached to it. For example, changing the word Latino to Latin X comes with it the connotation that if you don't use the word Latin X (hopefully you're familiar with that if not I can give another example) then you're an immoral person. Now do these neologisms have a different characteristic to them - other than someone neutral words like idiolect or cyberspace? That is, will they last longer or shorter, do they promote more peace or harm, is there something different about them? (And it's not just political correctness because I'm sure there are religious examples as well.)

We have an intrinsic nature. It offers opportunities to do new things, put constraints on what they are. Same is true for moral nature. Undoubtedly our moral nature has an innate basis, otherwise you could never acquire a cultural or moral system in the 1st place. Same problem of poverty, a stimulus. Something has to be in there internally. And that's going to offer scope for what you can do and put limits on it. From then on you can just go on to explore and try to determine the facts. But you can't project them by pure thought. You have to find out what they are.

OK, you got to get going. Thank you so much (Thank you.) I'll let you know how the video goes.

Thank you.

Thank you, my friend.