

A.M. Turing Award  
Interview with Judea Pearl, 2011 Award Recipient,  
By Roger Bingham  
Boston, 2011

This transcript is of an interview done by The Science Network and the ACM is grateful to them for putting that video interview in the public domain.

RB = Interviewer

JP = A.M. Turing Recipient

RB: We are in Boston at the Park Plaza Hotel, which is the headquarters for the 33<sup>rd</sup> annual Cognitive Science conference, CogSci 2011. We're with Judea Pearl, who is this year's Rumelhart Prize winner for your work on counterfactuals and all sorts of other things. Could you give me some sort of sense of what was being honored in this particular case?

JP: I understand that my work on probabilistic reasoning and causal reasoning and counterfactual reasoning was responsible for this prize, for which I'm very honored to have received. So it's a combination of three areas of endeavor, three areas of research – probabilistic, causality, and counterfactuals.

RB: It actually says, "Contributions to the Theoretical Foundations of Human Cognition." So that would be...

JP: Well, if you buy the idea that causation and even counterfactuals are the basic modules that move and propel our reasoning, then the two are non-contradictory.

RB: Could you unpack some of those terms for me though? For a general audience I mean. What are counterfactuals for example?

JP: Well, let's first talk about causation because that's a word which is more common in ordinary discourse. We'll say, "A causes B" – "I dropped the glass and the glass spilled, and that caused the tablecloth to be dirty." This is a very common relationship that we use in ordinary discourse. Everybody understands them. Even a child two years old understands that "I spilled a glass." A child is being punished for not sitting quietly and causing the tablecloth to be dirty, and we send people to prison for things that they have done because they have caused a certain damage, because they should have known better. So all this is part of our culture. And the illusion that we have free will and that we are responsible for our actions, this is all part of our cognitive state of mind.

To understand it requires that we emulate it. Namely that we are able to program computers to reason with those relationships, to understand what it means that "A causes B," to understand the story first to be able to answer a question, "Is it true that A caused B, that A was the cause?" That requires a certain formalism according to

which A causes B. So a computer requires this formalism, he or she – I’m talking about computers, right? – it can answer a question “Yes or no, A caused B?” And we want to be able to communicate with robots on that level, because apparently it’s a very effective communication language – we need to communicate like that – which means it has a power of conveying information succinctly and to the point, and to get people to do what we want them to do. So if we want to instruct a robot to behave a certain way, we ought to be able to communicate with them on a level of cause-and-effect relationship – “Do that, because you’re guaranteed that this will be the consequence” or “You shouldn’t have done that, because you know what happens when you do that.” That’s a very powerful communication language.

RB: Let me at this point say you used this phrase “the illusion of free will.”

JP: Oh yes, yes.

RB: What is your position on that?

JP: My position is that there’s no free will, but it’s a very useful illusion. There’s no free will because we know that the world more or less is deterministic, and even if you take in quantum mechanics’ uncertainty, you just budge the question. Our actions are determined by neural activation, and one neuron fires because other neurons were in the environment and sent certain signals, the activation level reached a certain state, so it’s all predetermined. Namely whatever we do is predetermined by the state of mind that we had yesterday, corrupted by some noise that came from cosmic radiation and other, but still whatever we do is predetermined.

RB: How far do you want to take that back? Do you want to take it back to...

JP: It becomes...

RB: ...we live in a completely deterministic, Spinozistic universe, or what?

JP: Yes. I believe in Newtonian mechanics, which means that I believe that the uncertainty principle that we heard from Heisenberg in quantum mechanics applies to subatomic processes but not to macroscopic processes. And neural firing is more or less a macroscopic process. And even if we go to a quantum mechanical level, it doesn’t yet... it only moves the determinism from observable to probability state. So it doesn’t make much of a difference. Our actions are predetermined. And if you want to go to quantum mechanics, then the probability of our actions are predetermined. It doesn’t absolve us of responsibility, of our slavery to predeterminism.

Given that... And I believe in that, okay? I strongly believe that we are deterministic machines. Given that, it’s a puzzling phenomenon that we all talk as if we have an option, as if at any point in time I have the options of touching or not

touching my other hand and have a very strong sensation that I have the two options in my disposal. Right? I can either touch or not touch, and I have a very strong sensation that I am the master of what will happen in a second. That is an illusion. I'm not master of what will happen. But it's a very useful illusion because evolution apparently found it necessary to equip us with that illusion, and we have survived fairly well with that illusion and we communicate effectively with that illusion.

Obviously we operate with the illusion. As I said before, we send people to jail because they ought to have known better, so they ought to be responsible for their action. On that basis, we teach people, we train people to improve their behavior and we punish people for not doing what they are doing. And society develops, culture develops with that illusion in mind. It forms almost everything we do. Free will and responsibility and regret and all these notions are built in our culture. They are a basic building block without which you couldn't explain behavior.

My question is whether...on the human side, how come evolution has given us this illusion and how it manages to keep this illusion so vividly in our mind despite the fact that it's just an illusion? And we know it! Now the question comes, going to the robot world, should we equip robots with the same illusion? Are we going to benefit from that? Are they going to benefit?

Let's assume that we want to train a robot team to play good soccer. Part of playing soccer means to have people appreciate they're playing in teams. So should we for instance equip robots with the ability to tell each other, "You should have passed this ball, and the reason we lost this game is because you haven't done what you're supposed to do. Therefore you are being expelled from our team. I'm going to take another robot"? Does it make sense? Apparently it works for a human, this communication... I'm talking now about only in terms of communication language, not what it means really to have or not have free will. Should we equip robots with that kind of illusion? Because as robots, you and I agree, they don't have free will. They will do whatever they were programmed to do, like we – we will do whatever we are programmed, and still, on top of our programming, we built up this illusion of being responsible, a master of our actions. So that is a simple computational problem – would a robot team play better if we equipped them with that language, vocabulary of free will?

RB: What I just pulled up here on the computer while you were talking was the uncanny valley effect, which I'm sure you know.

JP: No, I'm not...

RB: It's "a hypothesis in the field of robotics and 3D computer animation," it says here on Wikipedia, "which holds that when human replicas look and act almost, but not perfectly, like actual human beings, it causes a response of revulsion among human observers." So in fact people who work with robots try and avoid getting to that state of perfection almost because it freaks people out. So...

JP: Well, people are freaked out for many good things. I'm freaked out when I see a beautiful toy that does something I'm not expecting it to do. Like a new robot toy that asks intelligent questions. So what does it mean to us, the "freaking out" effect? Should that steer science in one direction or another? People freak out for any surprise and then they get used to it and learn to utilize it for their own good. So if we have a robot team that speaks like humans, communicates like humans, it constitutes a beautiful enrichment of our life because it reflects on ourselves and it helps us learn about ourselves. Like in this case, I would be able to understand under what condition this illusion of free will is meritorious, beneficial.

RB: I've talked to a colleague at UCSD, V.S. Ramachandran, about free will, and Rama likes to say that "We don't have 'free will.' We have 'free won't.'" In other words, I think he's applying that to some sense that we have some innate or possibly acquired moral sense which precludes us from doing actions which are antisocial. Does that resonate with you at all, or...?

JP: It doesn't answer my question, but it's a nice theory. Whether we acquire it or we're born with innate moral sensation, I don't know. Part of it is innate, part of it is learned. Most of it I believe is learned. We learn first to empathize. I need to have a model of the other agents. And the only way I can predict your action is to assume first – so a first degree of approximation – that you are like me. Therefore I would ask myself, "Would you react to this signal the way I would?" and that will enable me to avoid the computational hardship of building a new model of yourself. I already have a model of yourself. It's me. So I imagine that I were you and I give you a signal, and I ask, "Would I have reacted...? How would I react to this signal?" That's the way I expect you to react. It's very easy. We have a built-in model of ourselves and we replicate it and attribute it to other agents. That's the first step in moral behavior, because the next step is to say, "Would I hurt you if I do that?" Now to answer that, I just look at myself and say, "Would I be hurt if I would get that kind of input from the environment?"

That's the basis for moral behavior – to empathize. And it's not different than playing games with an opponent and asking what he or she's next move is going to be. Whether it's acquired or not, it's very natural. It doesn't have even to be required. From a pure computational simplicity viewpoint, this is the first step I would do.

RB: But what the presumes is that there is... that human nature is a constant.

JP: For first approximation. Then, as a child, I learn about differences – "Some people are happy with things that I wouldn't be happy with." So I learn about differences. But first approximation, I assume that you are like me. But this is really the second stage. The first stage is "Me, me, me." All the other children in the Kindergarten are the environment, they are the strangers. Then I learn that if I need to motivate them, I have to assume they are like me to understand what motivates them. Then

comes the level of moral responsibility. I hurt them. I make them happy. That's the beginning of moral behavior. But it comes together with the idea of "I can do it," of agency. I choose one action over another for a reason, because I'm master of my action. That's the illusion of free will that I was talking about. So we want a robot to be nice to us, right? In which case we have to equip them with a program to empathize and to have this illusion of free will.

RB: It says here... I mean you're described as "a computer scientist and philosopher, best known for developing probabilistic approach to artificial intelligence." But it's quite clear from what you've just been saying that you're a student of human behavior.

JP: That I'm doing what?

RB: Student of human behavior.

JP: The first word I didn't get. The...?

RB: "Student."

JP: "Student"!

RB: Yes.

JP: Oh yes. Because I'm a student of myself. Isn't everybody? Aren't you a student of yourself? Aren't you puzzled by your own behavior or what makes you tick? My saying is that "Every computer scientist is a frustrated psychologist." That's why students go to computer science. Because they want to learn more about themselves and they get a kick from emulating themselves on a computer. Even on the simple, mundane arena of arithmetic, "Look, a computer adds it much faster than I. What a kick!" Same thing, "Computer plays chess better than I. What a kick! Computer answers intelligent questions better than I." It's a terrific power to emulate myself in a mechanical device that is better, more reliable, and it makes me understand myself better, because if I take part of myself, I mirrored it, I emulated it on something else, it's now subject to my inspection. I can look at it, I can change components, and I can see how the robot would behave differently if one component or another is changing. It's really I'm doing it to myself – "How would I behave differently if one component or another would...?" So I think every human being is enticed by his or her own behavior.

RB: Let's circle back and unpack for me the whole notion of counterfactuals, would you please?

JP: Well, a counterfactual is a statement about how the world would look like if things were differently than what they are in reality. I'll start again with a child and I say, "You spilled the milk." The child said, "I didn't mean to." And I say, "If you were

to sit straight, you wouldn't have spilled the milk." Very simple. Every child understands that. But the guy didn't sit straight. The guy spilled the milk for some reason. He was excited by some joke, so the milk was spilled. The movement of the hand was a certain way and now the parent comes and says, "You shouldn't have done it. The milk would not have spilled had you not done what you did." That's counterfactual. It means it's required for the child to imagine a world in which the action that actually was taken was not taken. The ability to imagine a world under a different scenario, a different bend of history, different than the trajectory it actually took.

That ability takes a little intelligence. It doesn't come easy. I don't think animals are capable of that kind of...

RB: That's what I was going to ask about.

JP: ...high-level reasoning. I don't know at what stage of our culture the ability to reason counterfactually and to talk counterfactually came into being. It's very interesting for historians of language development to investigate. I believe it must have come together with some explosive advance in culture, perhaps using tools for instance. It might have accompanied the development of toolmaking.

RB: How did this all begin for you? I mean this journey which ends up with the Rumelhart Prize here, at least at this particular juncture. Did you always plan to be a scientist of some kind? Was there some impetus for this, a teacher or your parents or...?

JP: I don't know. I think so, because yeah, my teachers were...

RB: We're talking about back in Tel Aviv now, right?

JP: Back in Tel Aviv in the 1940s, right, 1930s and 1940s, I was really fortunate to be in the company of my high school teachers, even grade school teachers, which were really... They came from Germany. They were professors that lost their jobs in Germany, in Heidelberg and Berlin, and came over to Tel Aviv to teach high school. They knew that they would never be able to go into a laboratory again, never be able to conduct experiments or be able to publish a paper anymore. Their life is now constrained by the job they could get in Tel Aviv and they assumed upon themselves a different mission – to prepare a new generation of scientists like they would have liked to be so that the country that they are building would be a model of excellence. And they were really devoted in that capacity. They spent all their energies preparing our generation, for which I'm grateful and many of my friends are grateful, because they were great people. They knew... They were polyglots. They knew anything from biology, history, mathematics.

We were fortunate enough to learn science in a chronological progression, which means any theorems in math that we learned, we learned about the face of an

individual, a Greek scientist, a Greek mathematician who were facing specific problems of the time. The state of knowledge of the time was clarified to us by the teacher and “Here’s the question that Pythagoras is asking himself. Would the isosceles be describable in some formula in terms of the two sides? Why is it important? Because you have to build a pyramid and you have to find out how many bricks to put on that side and that...” Everything was made concrete in a problem-facing situation, and here comes this Pythagoras and invents Pythagoras’ theorem. What a joy. What a victory of a human mind over that. That’s how we learned, theorem by theorem, and we got this joy of discovery by tracing back the steps of these great scientists of the past, all the way from Pythagoras to Einstein.

I like this kind of mode of teaching. Once I’m retired and I go back to my high school and teach those kids again, this is how I’m going to teach in science.

RB: So you’re in high school. You discover a love of science. Where does it go from there? How did you then...?

JP: Well, I went to the army. And I went to... Part of the army was a kibbutz. I really meant to stay in the kibbutz and to become maybe an agricultural engineer. For some reason, my mother implanted in me the idea that “You were born to be a scientist.” I know at one point the teacher invited her to talk to her, and I thought, “Wow, I did something wrong again.” She came back and said, “No, he said nothing wrong about you. He just said, ‘Teach him, teach him. Sell everything you get, but give him education. Because he’s promising.’ ” I thought, “This teacher, he sees something in me?” I didn’t believe it. I thought he came to complain, called her to complain because I was rowdy in class, and I was. Anyhow, so this perhaps gave me a seed that maybe I can do something with myself in this direction.

Then I went to the army, I went to the kibbutz, and I left the kibbutz. It was a crisis in my life to leave the kibbutz because I really liked this ideal ideology. To live in a community, to contribute, and to be... But I saw right away, it’s not for me. After a year, it’s not for me because I really felt that I can develop myself better in science.

So I went to engineering because engineering promised a steady job in the post office. And I went to electrical engineering and in college still, again I met these giant professors. They again implanted in me the illusion or the hope that I can contribute in a direction, that is that I can discover new things. And after finishing, I came to America to study for the master’s with the idea of staying here only for one year and then go back. Well, I stayed here for more than 40 years and I’m still kicking here.

But perhaps what I learned through the years is that rebellion pays. A restless mind pays. I had many failures for being rebellion. I’m talking about science. My wife tells me one night I woke up and say, “Maxwell was wrong. Maxwell was wrong.” I discovered a mistake in what... And after two days, I found out he was right. But I had that restless night.

RB: There was this book a long time ago by Sir Peter Medawar called *Advice to a Young Scientist*.

JP: I haven't read it. But I read something similar I'm sure, I'm sure.

RB: And Ramón y Cajal had *Advice to a Young Investigator*. So I often ask people who've been in the field for a while, and it sounds like this was part of your answer, what is your advice to people going into science, advice to a young scientist? One of the things, it sounds like, is "Rebel."

JP: But not rebel in the contrarian sense. Rebel, if you don't understand and insist on understanding it your way, and until you understand it your way, don't take a standard answer as an answer if it doesn't satisfy the question that you have in mind. It's exactly what Galileo told us – "I respect Aristotle, but I wouldn't take his authority as a replacement for my curiosity. I want to see these two rocks falling from the Tower of Pisa in my own eyes."

RB: I think after he died, in Feynman's office on his board, there was some things that he had written. Richard Feynman. And I think one of the things was approximately "If you can't derive this thing from first principles, you don't understand it." So in other words...

JP: It's similar to that. I remember my former colleagues, Marschak, the great economist, I think even Einstein would just kind of say, "If you can't explain it in one paragraph to a child or to your neighbor, you don't understand it." It has to be something that grows with you, you understand it so well you can explain it to your neighbor. I hope I was able to explain to your counterfactuals.

RB: Yes, yes.

JP: It's a puzzle. You see, it's a puzzle why we form a consensus about things that have never been. Things that are never observed, hypothetical, sometimes metaphysical, and we communicate and we form consensus.

RB: How does this fit into your... I mean people know this about you as well of course, that your son Daniel Pearl was killed in Pakistan and you set up the Daniel Pearl Foundation. When we talked earlier about "I look at myself," empathize, the basis of moral behavior, how do you equate what was done to your son with what could...

JP: I rebel against it.

RB: ...any possible sense of being aligned with those people in human terms?

JP: In human terms, yes.



RB: I mean it must be completely anathema to you.

JP: In human terms, I can relate. I can relate to a person who has the same biology as I have. I'm ashamed that we have the same biology. It puts me to shame that we are members of a class that has some uniformity. At the same time I can empathize with a person who's been brainwashed in a certain ideology, and I can see myself as capable of being brainwashed in a certain ideology and to be motivated to conduct cruel acts in the world. I could have been a slave, like he or she was enslaved. So in that sense, I can empathize.

At the same time, it doesn't exonerate the creator of the ideology. And that is what I'm fighting against. The ideology of terror some people call it. The ideology that puts your grievances about all rules of right and wrong, the accepted rule of social behavior. That is the ideology that we should all fight against and that's what we are trying to do with our foundation.

RB: Explain to me how you get from that position to talking about the illusion of free will. Because you're...

JP: The illusion of free will doesn't prevent me from not resigning to the illusion. Actually I believe there is merit, I said computational merit in that illusion. Therefore I behave exactly like a religious person who believes in free will. I'm a master of my actions. I just told you I've not a master – right? – because my neurons are activated by other... I am a master of my actions for all practical purposes. Given the architecture of my system, my computational system, I am master of my actions, and therefore it behooves me to act and fight the hatred that took my son's life. I'm a master of my actions. The puzzle. I just said I'm not the master, right?

RB: Yeah.

JP: But my architecture has created this illusion, of which I am enslaved. I cannot get rid of it. It's part of me. Everything I learn was embraced, was carved, was shaped with this illusion in mind. I must make use of it.

RB: So in your view, the people that actually committed that foul act were also not masters of their own actions.

JP: Yeah. Fine. It doesn't absolve them from responsibility, which it doesn't absolve me from responsibility to correct it in the best, most effective way that is given to me. And this is to fight the culture of hate.

RB: Talk about the place of science in society in your view. I mean a question I often ask people is about, when he come into office, President Obama talked about restoring science to its rightful place. Where is that place? What is the rightful place of science?

JP: It's this boundary between what we know and what we want to know, the boundary of curiosity. This is the rightful place for science. Every child, every experimenter should be driven by his or her own curiosity. And the more we let scientists communicate about the questions they want to answer, to see answered, the more we advance that boundary of the unknown. What will come out of it? Only good things. I believe that by allowing ourselves to broaden these boundaries of curiosity, society as a whole is going to benefit from that.

So I don't buy this idea of going too far, learning too much about ourselves, building robots that will take over. I don't buy that. The more we understand what the robots are capable of doing, the more we understand what we are capable of doing and the more we learn, society will benefit in the long run.

RB: History, important to you?

JP: Very important.

RB: I mean history of science as well?

JP: History of science. It's very important.

RB: Suppose I gave you a time travel token and said you could bring to your dinner table anybody, anybody from any era, is there somebody you'd love to ask questions of and have dinner with?

JP: Well, of course I will Galileo with the first one, because I like his rebellious mind and his restless temperament. Yes. I'd like to ask him also whether it's true that he had this friction with the church the way we're learning in school. Some people raise questions about that issue. And also it's very important to me to understand why science developed in the 1500s in Europe and not in any other place. The Industrial Revolution, the Scientific Revolution took place in Paris, I believe it was the first university in Padua. And why it happened there as opposed to 800 years earlier. And... Yeah, go ahead.

RB: No, but I was going to say that whole story is very interesting though, because Alhazan and so on, there's amazing Islamic science up to 1100 or 1200. Then it's conveyed by Jews...

JP: Correct.

RB: ...to Italy...

JP: But the whole process...

RB: ...from which they're then expelled and they end up in Amsterdam, which is where Spinoza's community is.

JP: Correct. But you see the whole process was one of translation and not of innovation. From the 600 AD, from the times in Alexandria or so on, until the 1400s, we don't see any progress. We have only translation, translation, and conveyance. And the reason is that the society was organized in such a way that even if you had a genius like al-Khwarizmi, who wrote the book on algebra, you had a genius here or a genius there, their work was not continued. There wasn't the social mechanism to have their work stored, conveyed, and amplified. That's why it stagnated essentially. And somehow in Europe with the invention of the university which gave, number one, total freedom from the church and from the ruler, and later on the invention of the print and the ability of other people to access the achievement of one genius or another genius.

RB: Oh, if you look back to 1660 in England or the early 1600s and the formation of the Royal Society, that's also completely supported by this massive outgrowth, increase in the number of coffee houses.

JP: In coffee...?

RB: Coffee houses. Because this is where people are getting together. You're talking about setting up social networks.

JP: Correct. Social networking was a very important part of the development. And the freedom. Tenure even was invented there in Padua, right? You can do whatever you want and they wouldn't fire you.

RB: Is there anything else that you think you would ever have liked to have done instead of being a scientist? I mean, are you a frustrated pianist?

JP: I wanted to be a musician, but I wasn't good enough.

RB: Ah. A musician. That answer is so often from scientists.

JP: Really?

RB: Yeah. What would you play?

JP: Well, I play the guitar, I play some piano. I was for 10 years a choir conductor. As a matter of fact, the kibbutz sent me to be a choir conductor, so I would come back and organize the choir for holidays. And I was pretty good doing that. "Hallelujah."

RB: "Hallelujah, hallelujah." Last question to the Rumelhart Prize winner. Judea Pearl, what are you optimistic about?

JP: Optimistic? Oh, not many things. I'm sorry. Not the way that things are going, not very promising. The West is losing its spirit, and it worries me. We don't appreciate the spirit of the West. We keep on apologizing for being too colonial and too that and too... and we just don't have the mission. Kids do not grow up with the pride in being members of Western civilization. Which is exceptional. It's pride in being exceptional, exceptionally good, exceptionally in charge of something they can contribute, something special they can contribute to the world. That worries me.

RB: What do you think about what's going on in the Middle East? In Israel, in your former homeland I mean.

JP: Also, I'm not very optimistic. I know what's going on and I think we are heading toward war. It's a very simple conflict that has no resolution. Some questions in science have no answer. Right? Like a Turing machine halting problem has no answer. Same thing with the Middle East.

RB: Do you have any religious leanings or are you completely atheistic?

JP: I am of course a prisoner of my upbringing, which means my store of metaphors comes from the Bible and comes from the history of the Jewish people. But I don't believe in God. Actually I know that there isn't God if you...

RB: How about Spinoza's God?

JP: At the same time, I would tell you, I could quite easily try to convince you that you should do something because otherwise God will punish you. And I don't think the two are contradictory, because I use here the God metaphor as a powerful communication device that you and I are familiar with. That's what makes it a powerful communication device. And because it's a beautiful poetry and powerful poetry. So why shouldn't we communicate with the poetry? I have no qualm with having my child pray in school, because I teach my child, "This is the way your grandfather used to recite poetry." With the names of Moses, God, reward and punishment, divine supervision. It's a beautiful poetry, so use it as such. But don't forget, it's just a poetry. Don't take it too far. Have a very good understanding of that.

RB: Einstein said – because I mentioned Spinoza's God, and Carl Sagan subscribed to this as well – "I believe in Spinoza's God who reveals himself in the orderly harmony of what exists, not in a God who concerns himself with fates and actions of human beings." Is that okay for you?

JP: No, no, no, no, no. That's not the way that my grandfather used to talk. My grandfather was much more concrete about God's actions, and God will constantly observe you and God will punish you if you don't do the right thing. So I go further. It's not only the harmony in the universe. I'm talking about God as a guardian of morality. I would do this metaphor, but remind my child that it's only a metaphor.

What is God? It's a collection of ideals that we would like to acquire and we would like to pass on to the next generation. That is God. If we dress it up in the form of a Santa Claus that watches every one of your actions or dress it up in the form of dry principles, the former is more powerful.

It's dangerous too because if God is so powerful and you're doing something in the name of God, you are in danger of entitling yourself the sum of these divine powers and believe that you have the permission to elevate your needs in the name of God over all other principles. So it's dangerous. But it's a very powerful moral-keeping metaphor. And I love this poetry. It's a poetic tool. It's powerful, useful, and satisfying to the mind.

Because, look, we cannot learn everything in the world. We learn things that we are capable of learning quickly, that's what we learn. And the thing that we are capable of learning most quickly is parents-child relationship, which means we have the metaphor of a father and metaphor of a mother. Using this metaphor, which is so engrained in us in such an early stage of our development, we build higher-level concepts on top of that. So morality is built on the metaphor of a father. He is good to you, he punishes you, he cares of you, he is watching over you. That's what we learn, that's what we're capable of learning. It's much harder to transmit a set of principles without that metaphor. So I buy that and I use it. And even though I'm a devout atheist, I'm actually a staunch atheist, I can see myself talking to my grandchild and tell him, "Don't do that. God will punish you."

"God will punish you" means "Eventually you're going to cause harm to yourself. Whether it is in a collective set-up or individually, if you wait long enough, you'll be harmed." That's what I mean by "God will punish you." That means that the action itself eventually turns out to be something you wouldn't like the consequence of.

We should learn to talk science and talk poetry. Both have roles, important roles in our communication. So I'm not like Richard Dawkins, which gets very irritated when somebody mentions God. No, it's a poetry. If you feel comfortable about your atheism... And I feel very strongly about it and I'm not threatening by anything. I don't believe that my child would become an Orthodox Jew. No. But there's no reason why he cannot talk in terms of the God metaphor. Beautiful poetry. Poetry like butterflies, plows, clouds. "I wandered lonely as a cloud." What does it mean?

RB: Wordsworth.

JP: What? Wordsworth. But did it mean, "Was I really the cloud?" It's a beautiful metaphor. Why not use it? "I punish you like a God." Why not? "Don't do it because God will punish you."

RB: It was great fun to talk to you. Judea Pearl, thank you very much.

JP: Okay.

RB: Thank you.

[end of recording]