האוניברסיטה העברית בירושלים
THE HEBREW UNIVERSITY OF JERUSALEM

GAME ENGINEERING

By

ROBERT J. AUMANN

Discussion Paper # 518 **September 2009**

מרכז לחקר הרציונליות

CENTER FOR THE STUDY OF RATIONALITY

Feldman Building, Givat-Ram, 91904 Jerusalem, Israel
PHONE: [972]-2-6584135 FAX: [972]-2-6513681
E-MAIL: ratio@math.huji.ac.il
URL: http://www.ratio.huji.ac.il/

GAME ENGINEERING

Abstract

"Game Engineering" deals with the application of game theoretic methods to interactive situations or systems in which the rules are well defined, or where the designer can himself specify the rules. This talk, which addressed a business-school audience with no specific knowledge of game theory, describes five examples of game engineering: two dealing with auctions, two with traffic systems, and one with arbitration. At the end of the talk there was a Q & A session, which, too, is recorded here.

Transcript of the lecture by Prof. Aumann at Koźmiński University in Warsaw, Poland. The meeting took place on May 14, 2008.

Rector, Professor Andrzej K. Koźmiński: Ladies and Gentlemen, it is a big event today. We are honoured to host Robert Aumann, a Nobel Prize winner in 2005. Thanks for coming. We have so many distinguished guests in the audience that I would not be able to name all of them. There is only one person I would like to single out, his Excellency the ambassador of Israel, David Peleg.

This is your first visit to our Academy and we wish to welcome you and remind you that you have a standing invitation. I think it would be a waste of your time if I spoke any longer, but I would like to say that this is a special event, not only because of Professor Aumann's eminence, but also because it is the twentieth lecture in our Distinguished Lecture Series, the fourth by a Nobel Prize winner, and it is also the fifteenth anniversary of our University. So we are very happy and very pleased to have such an eminent speaker on this special occasion. Now I would like to turn the floor over to Professor Kołodko for a short introduction and afterwards we will have Professor's Aumann's lecture, followed by a short Q&A session.

Professor Grzegorz W. Kolodko: It is my privilege to welcome all of you to our Distinguished Lecture Series. This is a very special lecture, already the twentieth, and one may access all of them on the TIGER (Transformation, Integration, and Globalization Economic Research) website.[1] It is also special because our distinguished guest of today, Professor Robert J. Aumann, has come to visit us from the country he loves so much, from Israel. For more than 50 years Professor Aumann has been teaching and doing his excellent research at the Hebrew University of Jerusalem, which I had a chance to visit for the first time only in 1991. In the aftermath of that visit and that conference an interesting volume on the issues of post-communist transformation was published.[2]

Professor Aumann comes from Israel, but he belongs to all of us, if I may say so, owing to his tremendous contribution to a better understanding of certain economic processes, and to his lifelong devotion to economic studies. He was born in Frankfurt on Main in Germany in 1930 and was lucky enough to leave, together with his parents and brother, in 1938—for, of course, the US; at that time, where else? He settled in New York City, where he

---

[1] See http://www.tiger.edu.pl/publikacje/dist.htm.

[2] See Michael Keren and Gur Ofer (eds.), *Trials of Transition: Economic Reform in the Former Communist Block*, Westview, Boulder 1992, and my chapter on *Transition from Socialism and Stabilization Policies: The Polish Experience*, pp. 129-150.

did his B.A. at City College, and later his M.A. and Ph.D. at the Massachusetts Institute of Technology—all in mathematics. And then, as things go, he decided to take another great step and travelled across the Atlantic once again, this time to Israel, where he serves the Hebrew University as Professor Emeritus in the Department of Mathematics.

Professor Aumann in his professional life—aside from the Massachusetts Institute of Technology, where he graduated, and the Hebrew University, with which he has been associated ever since—has been in various places and positions, mainly as a visiting scholar and professor. He has contributed to very famous institutions on both sides of the United States, such as Yale and Princeton on the East Coast, and Berkeley and Stanford on the West, as well as points in between. He has worked in our Old World, too, visiting places like the University of Louvain and certain other respected research centres and universities.

All the while he has been doing his research, basically looking for rationality. And for this research, for the remarkable outcome of this research, he has been awarded the honour, the greatest honour in science, that is, the Nobel Prize in Economic Sciences in 2005, jointly with Professor Thomas C. Schelling of the University of Maryland.[3] This award was conferred "for having enhanced our understanding of conflict and cooperation through game-theory analysis."

In his Nobel Prize speech in December of 2005, not that long ago, on "War and Peace,"[4] Professor Aumann answered in a very reasonable way what rationality is. According to him, "a person's behaviour is *rational* if it is in *his* best interest given *his* information." So, he argues, what sometimes seems on the surface to be irrational is not necessarily irrational, because considering the information constraints one may be acting in his or her best interest. Yet Professor Aumann, undisputedly the man of science, in the same Nobel speech refers at least twice or thrice to the Holy Book, while quoting Avot 3,2 from the Talmud, which says: "Pray for the welfare of the government, for without its authority, man would swallow man alive." So aside from rationality, from time to time it is good, according to this statement, to pray for the rationality of the government, without which it is difficult to be rational. However, one must stress that with such an approach to the issue of rationality, the accuracy of information—or just the knowledge—becomes the heart of the matter. Otherwise, it would imply that one can harm his or her own ends, while acting in a rational way.

Let me conclude with what Professor Aumann has written in his autobiography, which each Nobel Prize winner is obliged to do and which, I think, each and every one of them is very unhappy to do: "Being awarded the Prize in Economic Sciences in Memory of Alfred Nobel is a beautiful, fairy-tale experience, from beginning to end. It is not, however, conducive to research, writing, teaching, or indeed to any ordinary academic work. Instead, one is besieged by hundreds—indeed thousands—of demands to appear and speak at events and congresses, many of which have nothing at all to do with science, to underwrite good (and less good) causes, to sign and send autographs, to be photographed, to be interviewed, etc., etc., etc." All the more, then, do I appreciate that Professor Aumann has kindly accepted the invitation to come to our great country, Poland, which he is visiting for the first time, and in particular to the Koźmiński Business School, which is the leading business school in this part of the world.

And now, Professor Aumann, we are very keen to listen to your message—to your distinguished lecture on "Game Engineering." Shalom!

---

[3] Professor Thomas C. Schelling, University of Maryland, Department of Economics and School of Public Policy, College Park, MD, USA.

[4] See Robert J. Aumann, "War and Peace," Nobel Prize Lecture, December 8, 2005 (http://nobelprize.org/nobel_prizes/economics/laureates/2005/aumann-lecture.pdf).

Professor Robert J. Aumann: Thank you very much Professor Kołodko, Professor Koźmiński. As Professor Kołodko correctly said, I'm sure that I owe this visit to the Nobel Foundation; but it does have something to do with science and it is not like some of these other things that you mentioned. And I am very proud and happy and honoured to be at this very distinguished institution, the Koźmiński Business School. And now I will start, OK?

Let me start with a quotation from Jim Tobin, who died a few years ago. He was at Yale, a very distinguished economist and also a Nobel Prize winner. I was once present at a meeting of students at Yale University at which Tobin was also present, and the question that was raised at that meeting was: how can you sum up the science of economics in one word? There were various suggestions made for this. For example, "opportunity cost." That is not one word, that is two words, but still it is one phrase. That was one suggestion, but Tobin's suggestion was that the one word that best sums up economics is the word "incentives." The basic idea is that people act in accordance with their incentives; and this is also the basic idea of game theory. In fact, one can look at economics as a part of game theory. Some people look at game theory as a part of economics, but it's better to look at economics as a part of game theory, because game theory is really broader and covers all kinds of situations where incentives are important; and that's not only in economics, it's also in politics, it's also in international relations, it's even in biology, in psychology, and in many other areas of human endeavour.

There are two kinds of applications of game theory. Both have to do with people's incentives, and designing systems in a way that makes people's incentives work to your benefit. Sometimes when I give this talk that I'm now giving, I give it with a different title, which is: How to make incentives work for you. This is really a broad topic, because sometimes you get incentives to work for you in situations that are not that well defined; for example, negotiation. When you are negotiating, the rules, the order of doing things, are not entirely clear. Still you have incentives; you can negotiate in such a way that you are trying to get the other side to agree with you, to do something that is in your favour, because it is also in their favour.

I'd like to concentrate today on a different aspect of the application of game theory, which I call "game engineering." The reason is that the systems to which it applies are very well defined: they have strict rules and in fact you can build the rules sometimes. It's like building a bridge, or a computer program, or that kind of thing, where the system is well defined; it is not fuzzy, it is not like a matter of negotiation, it is not like a matter of fighting a war, where the rules are not so tightly defined. I'd like to apply the word "game engineering" to situations in which you have to act in accordance with well-defined rules, or you have to make the rules yourself. What I'm going to do, rather than laying out a general theory, is to give you a few examples of "game engineering." Here are a few of the applications that we might discuss today.

- Final offer arbitration
- Auctions
  - – Strategy
  - – Design
- Matching
- Traffic
  - – Strategy
  - – Design
- Elections
- Asset Division

I don't think we'll be able to cover all of them. Let's start with auctions. Auctions are a prime example of "game engineering," and there are two aspects to the engineering part of auctions. One is how to bid in an auction, the strategic element of bidding—in other words, how the players in the game respond to the rules. But a more fundamental problem in "game engineering" is how to design the auction. That is, if you're selling something, it's up to you to design the auction, to say how it's going to be done. Let me give you just one example of this.

The Nobel Prize that was awarded in 2005 was the third Nobel Prize in game theory. The first one was awarded in 1994 on the occasion of the fiftieth anniversary of the publication of the bible of game theory, "Theory of Games and Economic Behavior" by John von Neumann and Oskar Morgenstern. Fifty years later the Nobel Committee finally broke down and awarded a prize in game theory, and it was awarded to John Nash, Reinhard Selten, and John Harsanyi. Two years later, in 1996, another prize was awarded in game theory, to William Vickrey. William Vickrey is the father of auction theory, or more specifically the application of game-theoretic technology to the design of auctions. Let me just tell you a little story about that prize. It was announced in October 1996 and awarded in December 1996, but by that time Vickrey had died. He died between the announcement and the award ceremony. The Nobel Prize cannot be awarded posthumously, those are the rules. But the one exception is that it can be awarded to people after it is announced. Once the prize has been announced, just because you die they do not take it away from you. So it was awarded posthumously to Vickrey. Probably he died because of the prize. He was so excited by this that he had a heart attack and died. He was already rather old; I don't remember how old, but he was getting on in years. The Nobel Prize is a nice thing. But it is not *that* nice; I don't think one should get overexcited about it, so that one dies.

For what did he get the prize? He got it for his fundamental contributions to auction theory; specifically, for something that's very simple when you think about it. But it's also brilliant. The most important discoveries in science are very often those that are very simple, but nobody thought of them before. In Vickrey's case the contribution was the invention of the "second price auction."

A "second price auction" works like this. We are talking now about a closed bid auction for something. You are buying let's say a painting, a painting is being auctioned off. It could be also that you are selling something, a contractor, it is the same principle. In order to be definite, I will consider the case of somebody who is buying a painting. So a painting is being auctioned off and one way of doing this is the classical open bid auction. But another way of doing it is like in a government tender for something, like used cars; governments sell used cars. The way it is done is that the people bidding in the auction write the price that they are willing to pay on a slip of paper and then the slips of paper are compared and the highest price that was offered wins the auction and gets the object and pays the amount offered. This is the ordinary closed bid auction. In game theory it is called the "first price closed bid auction."

Vickrey said that that is not a smart way to do it. If you want the highest price, if you want to sell your painting for the highest price, do not do it that way. Do it differently. How? All the bidders write down their bid on a piece of paper and they submit the pieces of paper and then they are opened and the highest bidder gets the object. So, so far it is the same as what we said before. But then comes the difference: Vickrey said he should not have to pay what he wrote. What he should pay is the second highest bid. Not the highest bid, but the second highest bid. Back then in the sixties when he did this, they said to him: Bill, are you crazy? Why the second highest bid? You are throwing out money! If you have the highest bid, then obviously the highest bid is more money than the second highest bid, so what kind of a crazy idea is that? So he said: You know what, you try it and you'll see you'll get more money than you got before. So they did it, not on paintings but on automobiles and things like that, and they saw that they were really getting more money.

What makes that work? I will not say that it is very simple, because it is not very simple, but it is easy enough to be explained to this kind of audience. The idea is this. If you are bidding on something and you don't know how other people value it, then you won't necessarily bid the amount at which you value it. Let's say you are bidding on a painting. You value the painting at 50,000 dollars. What do I mean you value it? You are not ready to pay *more* than 50,000 dollars for the painting, but you *are* willing to pay 50,000 dollars. So, will you bid 50,000 dollars? No, not necessarily, because you also have to take into account what the other people are going to bid. If you think that the other people are going to bid 10,000 or 15,000 dollars on this, why would you bid 50,000 dollars? You like the painting, but you also like money, you don't want to throw out money. A person who thinks that other people will bid 10,000 or 15,000 dollars is not going to bid 50,000 dollars in a closed bid auction. So *you* will bid less; and, *everybody* will think similarly.

Right now I am in the market for Old Master paintings of Old Testament subjects. Most people bidding on Old Master paintings are not necessarily that interested in Old Testament subjects. I might be willing to pay 50,000 dollars, but I figure that other people are probably not going to pay that amount of money. So I bid much less, let's say 20,000 or 25,000 dollars. But there might be somebody else in the market also willing to pay 50,000 dollars and he will also bid the same, he also made that calculation. So he will bid 20,000 or 30,000; maybe he will get the painting, maybe I will get the painting. But the owner of the painting will not get what we really value the painting at. He will get much less. How do you avoid that?

You avoid it by the second price auction. In a second price auction, if I am willing to pay 50,000 dollars, if that is my valuation, then I will bid 50,000 dollars. I will not bid more, because if I bid 60,000 dollars and somebody else bids 55,000 dollars, then I will have to pay 55,000 dollars for something that I only value at 50,000 dollars. So I will not bid more than my value, but also I will not bid less. I will bid exactly my value, because I know that I don't necessarily have to pay for it. If everybody else is willing to pay only 10,000 dollars or less, then I will only have to pay 10,000. So I will bid my true valuation and everybody in the auction will bid his true valuation for the object. When everybody is bidding his true valuation and is not trying to figure out what other people are going to pay, then the auctioneer, the person who is selling the object, will get almost the amount that is the highest *true* valuation.

That was Vickrey's idea, and it got him the Nobel Prize. What this does is that it designs the system so that people have incentives to give their true valuation for the object. So, this is one example of game engineering.

Let me give you another example, which is on the strategic side, not on the design side. This is something that is often done in lectures on auction theory. Maybe I could have done it over here, but I did not want to take up that amount of time. It is called the "winner's curse." You take a jar full of coins and you put it on the table and you say: I'm going to auction off this jar full of coins with an open bid auction. People shout out prices and the highest price gets the jar full of coins.

Now, it almost always happens that the person who bought the jar full of coins paid too much for it. In other words, what he paid for it is more than the amount of coins that is in the jar. Why is that? Why is it almost always that in a room full of intelligent people, people pay too much for the jar full of coins? Explanations that are sometimes offered are that it is an exciting event, people are shouting out prices, they get overexcited, and therefore they bid too much for the object. But that is not the correct explanation. This happens even if they do not get overexcited. What is going on is this: You see a jar full of coins. You do not know how much is in there, but you are intelligent, you can make an estimate, and the estimate is more or less correct. What does that mean: more or less? It means that the estimates are distributed around the true amount: The true amount is the mean of the estimates, say. So we're putting complete faith in the acumen of the bidders. Let's say there is a thousand dollars worth of coins in there. Some people estimate it, therefore, at 1000, some at 900, some at 1200, some at 800, and so on. Now, who gets the jar of coins? Is it the person who bids the mean? No; it is the *highest* bidder who gets the jar of coins. When you have a distribution around the correct amount and the highest bidder gets the jar of coins, then the highest bidder obviously is not going to be the 800 or the 900 bidder, he is going to be the 1200 guy. He bids the highest amount, he gets the jar of coins, and he loses.

This is called the "winner's curse"; people who bid on objects that have an objective value—not paintings, but things like offshore oil tracts or other objects that have a definite monetary value, only you do not know what it is—have to be very careful about winning. They have to take into account the fact that if they win, then an important factor in their winning could be that they overvalued the object. People who undervalue the object are not going to win. So, that is an item in the strategic analysis of auctions.

And, these auctions are big-ticket items. In the United States there was what is called the spectrum auction; it took place ten to fifteen years ago, in the mid-nineties. The US government decided to auction off wavelengths, electromagnetic frequencies that are used for radio, for television, for cellular communication. Up to the mid-nineties they had been giving away these frequencies, they had been licensing the use of these frequencies on a basis of who can contribute best to the economy or something like that; but not charging anything for them. In the mid-nineties they decided to start selling them, auctioning them off; this auction was conducted in two stages, and it was estimated that it would yield about half a billion dollars. And then the FCC (the Federal Communications Commission) and the people who were bidding on these frequencies (mainly communications companies like Southern Bell) said: Let's hire game theorists to design this thing. The FCC hired game theorists to design the auction and the communications companies hired game theorists to look at bidding strategies, and the upshot was that the amount that was netted for these auctions was not half a billion dollars but close to 50 billion dollars. This was an amazing demonstration of the power of game theory.

Perhaps you could say that maybe they just made a mistake and they would have gotten the same amount if they had not hired game theorists. But that is not the case, because other countries followed suit, they hired game theorists to design the auctions. But there was one country that said: We do not need to hire game theorists, we understand the principle, let's do this on our own. It is like a person who builds a bridge and does not hire engineers. These people, this country that did not hire game theorists, they fell flat on their faces. Their auction made more in the neighbourhood of the original estimate, something that corresponds in the United States to half a billion dollars. These game-theoretic ideas, some of them may seem simple, but often they are more sophisticated than what I've said. But even these simple ideas that I've been talking about are important in the design of auctions.

Let me get to another subject, which is the subject of traffic. There again we have two items. One is the design item and one is the strategy item. I have already experienced Warsaw traffic and at times it is quite smooth, but at times it is stop and go. So the design of traffic systems, the design of roads and the choice of strategies on how to get from one place to another, can be a challenging item. Let me start with an example[5] concerning the design of road systems. Here is a road system:

---

[5] Due to Dietrich Braess, "Über ein Paradoxen aus der Verkehrsplanung," *Unternehmensforschung* 12 (1968), 258-268.

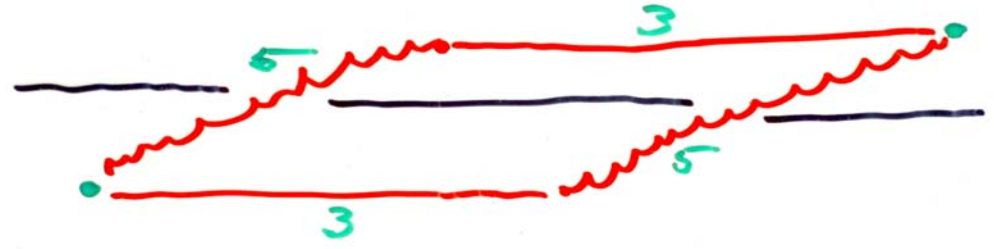

There are two cities, marked in green, and there is a mountain range that separates the two cities, which is marked in black. On each side of the mountain range there is a superhighway, which is marked by a straight red line, and then there are two roads that go over the mountain range, which are marked by wavy red lines. The superhighway is marked by a straight red line because it goes straight, and the roads going over the mountains are marked by wavy red lines because these roads are very curvaceous. The numbers that you see are the travelling times. It takes three hours to go on each of the two stretches of superhighway and it takes five hours to go over the mountain – slow going. The main traffic is between the two cities marked in green. These times that I am giving you are what they call *equilibrium* times, which means that about half the traffic is in the south and half the traffic is in the north. The reason for that is that as soon as considerably more than half the traffic will go, say, in the north, then the north will be more crowded, so it will be slower; on the other hand the south will be less crowded, so it will be faster, so people will move back to the south. So there is a natural tendency to have the traffic divided half-half, and these are the times that result.

The Department of Highways in this particular country looked at this situation. There was a bright young man in the Department who said: Let’s build a tunnel under the mountain, from the eastern end of the southern superhighway to the western end of the northern one. The tunnel will cost 100 million euros to build. But it's worth it because if you're going from one city to the other, you don't have to go over the mountain; you use both stretches of superhighway, plus the tunnel. The two stretches of superhighway will take three hours each, and the tunnel half an hour; six and a half hours altogether, as compared to the eight hours that you have now. These are two very important cities with a lot of traffic between them, and it's worth an investment of 100 million euros to improve the infrastructure of the country, to promote commerce, and so on.

No sooner was this idea suggested, than it was adopted; and after a year or two, the tunnel was opened in a very impressive ceremony, the ribbon was cut, and the traffic started flowing. What was the result? Here it is.

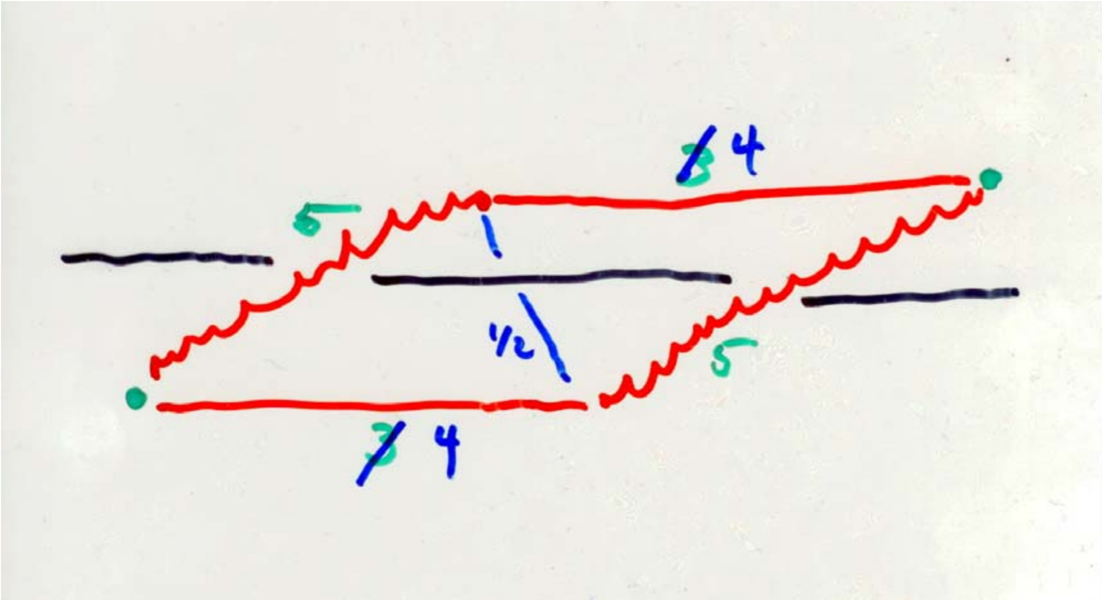

The tunnel is the blue line, it goes under the mountain. And it now took four hours to traverse each of the sections of the superhighway. Why? Because there was more traffic on them. Because everybody now was taking *both* sections of the superhighway. Nobody travelled over the mountain anymore. Everybody travelled on the superhighway. Even one accident would jam up the highway, so the average was four hours on each section. Half an hour for the tunnel made a total of eight and a half hours, as opposed to the eight hours that you originally had. Building this tunnel at the cost of 100 million euros actually made the traffic slower.

So one person said: OK, this tunnel is ridiculous and I'm going to go over the mountain. It used to be eight hours before they built the tunnel. Now it is eight and a half hours. I'll go over the mountain and I'll save half an hour. But it doesn't work that way, because going over the mountain takes five hours. The stretch of superhighway that he still has to take, with the other cars still there, is still taking four hours instead of three, so now it takes nine hours if he does that. Taking the tunnel is the only equilibrium in the new game. Building the new tunnel at the cost of 100 million euros actually slowed up the traffic, and that is the *only* equilibrium.

So somebody suggested charging a high toll for the tunnel to keep people away from it. To keep enough people away they had to charge such a high toll that it was politically infeasible. You could not charge such a high toll. So they closed the tunnel.

Once, I had an argument with the president of the Hebrew University, with whom I was driving. We came to a new road and I said to him: You know, Menachem, building new roads is not always a good idea, it can actually slow down traffic. The president, who is also a mathematician (he used to take my classes, by the way), said to me: I know what you mean; when you build new roads it gives incentives to people to use their cars more, and to buy more cars. Building new roads can slow down traffic in that sense. I said: No, that is not what I mean. With the *same* amount of traffic, building new roads can slow it down. He said: No, that is impossible. Until I showed him this example. Then he agreed with me.

This is not contrived. It is an everyday example, a very natural kind of example. In fact, it was actually experienced in Amsterdam. There was construction there, and they had to close some of the roads; this actually speeded up the traffic, so they kept them closed.

The *incentives* are what is creating this situation. So you must be very careful in designing a road system. You must take into account the incentives that new roads create. The new traffic patterns, not the old traffic patterns.

Here is another example regarding traffic. In the mid-nineties there was an attempt to blow up the Twin Towers in New York. The attempt failed. They tried again on September 11, 2001, and then they succeeded. Practice makes perfect; the first time they did not succeed; it was in 1997 or 1998. They parked a car in the basement of the Twin Towers and blew it up. They had a lot of explosives and did a lot of damage, but it did not bring the Twin Towers down.

I happened to be in Stony Brook, New York, that Friday, and wanted to go into town for the weekend. I went to get my car, I was living in a Bed & Breakfast, and the landlord said: You cannot go into the city now. I said: Why not? So he told me what had happened, and said, everybody is being warned to stay away from the city, it is in chaos. Stay away, don't go into Manhattan. It is a terrible mess over there. I said: I really don't want to stay in Stony Brook for the weekend. It's about a 50 mile drive to New York. I'm going to take my chances. I'm going to see. If I can't get in, I'll turn back, but if I can get in, why not?

Ladies and gentlemen, I have never been in Manhattan faster from Stony Brook than on that day. The roads were empty, nobody was going into Manhattan. The point is that when you are told to stay away, the incentives that are created by the announcement to stay away make it worthwhile not to stay away. An announcement to stay away from a traffic jam gives an incentive for people to go, because everybody will stay away, and that gives an incentive to go. Now, when you figure that out, if you are really given an incentive to go, then you should stay away again. This is a classic game-theoretic situation.

Now, how can you use this in order actually to improve traffic? Before, I was talking about the design of road systems. Now, I am talking about the design of a driving system. What do I mean by a driving system? You know what a GPS is? A GPS gives you instructions as to how to drive from one place to another and it gives the shortest route. However, GPS's usually don't take into account the traffic situation. They just tell you what the shortest route is. So people have thought up systems for incorporating information about traffic into the instructions that the GPS gives the driver. The GPS will get information as to where there are traffic jams and where there are no traffic jams, and this will be processed by the computers and then the computers will tell people to stay away from places where there are traffic jams.

That is a great idea, but it's not going to work. Why not? Because of the Twin Tower syndrome that we just discussed. The moment that you tell people to stay away from this or that road because there are traffic jams there, then people will stay away, but then the smart man is going to say: Hey, everybody is staying away, I'm going to go! And he will find that there is no traffic; but then you are going to get too many smart people and again there will be a jam. You have to devise some system for equilibrating this. How can you devise such a system?

Here is where game theory comes in, with a concept called a “mixed strategy.” I'm not going to tell you exactly what a “mixed strategy” is, but the way it works in this context is that you tell some of the people to take route A and you tell other people to take route B. You do not tell them about traffic jams. You just say: Go this way. So when there is a traffic jam, you tell half the people, say, to avoid that road and half the people to take that road; or it could be a quarter-three quarters, or one third-two thirds. Some of the people you are going to tell to take that route, in spite of the fact that there was, say, an accident there; and other people, not to take that route. And, you have to devise the system in such a way that it is worthwhile for people to follow your instructions. That is what is called a “mixed strategy equilibrium” in game theory; that was the contribution of John Nash, for which he got the Nobel Prize in 1994. You tell some this and some that. You do not tell people what you are telling other people. And everybody knows that you are not telling them what you tell others. That is not a secret. But everybody also knows that it is worthwhile for you to follow the instructions that the GPS gives you. That's an example of the use of game theory on the strategic side of traffic.

As a final example of game engineering, I would like to tell you about “final offer arbitration.” But let me tell you first what arbitration is. Arbitration is used in a dispute; let’s say employers and unions disagree as to the amount of wages workers should get. The employer wants to give lower wages and the union wants higher wages. The union could strike, but perhaps a better way is to go to arbitration. You sign an agreement that an arbitrator will decide what the amount of the wages will be. Each side presents its case. The employer wants to pay relatively little. The union presents its case; it wants a relatively large amount. And then the arbitrator makes his decision. Presumably, he'll compromise. Compromise makes the world go round. So the arbitrator takes some position between the employer’s offer and the union’s demand.

What incentives does that create? The incentives that that system creates are for the union to exaggerate its demand and for the employer to offer less than what he's willing to give, because he knows that the arbitrator will compromise. Both sides will offer a bargaining position; and they will try to justify the bargaining position, so they are not going to give accurate information to the arbitrator. The arbitrator will not know what's going on. He'll be operating in a fog, and the result is not good, and is very likely to lead to further labour trouble. And this is all because of the principle that compromise is a wonderful thing.

Then, somebody came up with the idea of “final offer arbitration.” In “final offer arbitration” the system is the same, each side presents its case, but each side is encouraged to present its minimum demand. The workers are encouraged to say what is the least that they are willing to settle for, and the employer is encouraged to say what is the most that he is willing to give. Of course that was also the case before. What's the difference? It is that now the arbitrator is *not allowed* to compromise. By the rules of the arbitration contract, he must choose one of the two offers. One and only one. Either the employer's or the union's.

That sounds crazy, right? You're not allowed to compromise. What are the incentives that this creates? The incentives are exactly the opposite of what they were before. Now it's worthwhile for each side to be as reasonable as possible, to make a very reasonable impression on the arbitrator. Each side wants the arbitrator to accept its position, and therefore it will try to make an impression that it is not asking for much, that it is making a very reasonable demand. The idea is that therefore the sides will come much closer to each other, and they may even meet. But even if they do not meet, it does not matter so much anymore which side is accepted. Whereas before you could have had a range of between 50 and 120, and then the arbitrator could have decided on, let’s say, 90, now maybe the employer will offer 72 and the union will ask for 76 and then he has to give either 72 or 76; but it doesn't make that much difference anymore. So you've incentivised the sides to be reasonable and to give accurate information to the arbitrator, and the arbitration procedure works much better.

This is not very widely known in the world of arbitration, but it *is* known and is used; for example, for professional baseball players in the United States, who have adopted “final offer arbitration” as their means of negotiating for salaries.

I'm ready for questions.

Robert Rządca:
I am a professor at this school. I have two questions. The first is connected with the so-called prisoner’s dilemma. A lot of people treat it as the king of game models for social situations. What do you think about it?

The second question is: What do you think about the possibility of game engineering for real life situations in oligopolistic sectors, like energy, the energy producing sector in Poland for example or in other countries, which is close to the oligopoly situation? Is it possible to create, based on game engineering, a better situation from the social point of view, social welfare?

Professor Robert J. Aumann:
What specifically do you want to ask about the prisoner's dilemma?

Robert Rządca:
Is it really the best model for the social conflict situation?

Professor Robert J. Aumann:
It is an interesting model. The prisoner's dilemma is a very important example to keep in mind: It demonstrates that it is not enough for something to be to our mutual advantage, in order to achieve cooperation. Even when it is to your advantage and to my advantage to cooperate, still the incentives can lead to not cooperating, to a result that is disastrous for both of us. It is really a wonderful example of how you have to pay attention to incentives when you are analysing a situation. It is not enough to say: Hey, let's cooperate, let's not go to war, let's not fight.

So, it is not enough to say "let's make love, not war." The incentive has to be there, and that is what is represented by the prisoner's dilemma in very stark, obvious language. It is really an important example. It certainly does *not* apply to every social situation, but the underlying message there is an important message. That is one thing.

Your other question was about using incentives, game engineering, to deal with the environment and especially with global warming, with CO2 emissions, and so on. That is *the* way to go! Absolutely! It is *the* way to go. No amount of Kyoto Protocols, none of that stuff is going to work. You have to give people *incentives* to do it right.

Let me give you just one example of this. I go skiing every year. Sometimes it is in Switzerland. Over there they have a wonderful system for garbage collection. You can tell people from today till tomorrow to stop using plastics, stop using all that packaging, buy paper bags, or go shopping with a bag that you use again and again. You can say this until you are blue in the face, but people are not going to do it. But in Switzerland they got them to do it. How? They gave them incentives. What was the incentive? You pay 5 Swiss francs for a garbage bag. And you must dispose of your garbage in those garbage bags. They will not collect it unless it is in the garbage bag for which you paid 5 Swiss francs. If you pay 5 Swiss francs, then you start paying attention to your garbage. If it is going to cost you money to throw out stuff, then you are going to make a lot of effort to throw out much less. You must have an incentive for this.

Absolutely, you must create incentives for stopping CO2. There is one very simple way to do it. Just tax the emissions. You could impose a much higher tax on gasoline. And there are other ways to tax emissions. Do not overtax then, but tax them at the true cost of these emissions. Absolutely, you have to give incentives. Not by fear; that is not going to work. What is going to work is giving people incentives. Precisely game engineering.

Jacek Margul:
My name is Jacek Margul from the Warsaw School of Economics. I would like to ask for your view on the study of rationality. In so many real life situations we have people behaving totally irrationally. You have devoted a lot of your time, a lot of your effort, to studying rationality, the opposite. Wouldn't it be more rational, to put it paradoxically, to study irrationality rather than rationality?

Professor Robert J. Aumann:
I do not buy the proposition that people behave irrationally. I just do not buy it. In 2002 a Nobel Prize was awarded to Vernon Smith and Daniel Kahneman. The prize that was awarded to Kahneman got a lot of press, because it was for what is called "behavioural economics," which stresses the irrationality with which people often act. But people ignored the other half of the Nobel Prize, to Smith, who got the prize because his work showed that people *do* act rationally. Actually this prize was given to these two people not for proving that people act irrationally, but for developing the methods of experimental economics. Smith came to the conclusion that people do act in accordance with the dictates of neoclassical economics; in other words, they do act in accordance with rationality postulates. Kahneman came to the opposite conclusion, and they shared the Nobel Prize.

Now what is going on over there? How can you explain this Nobel Prize being given to one for proving irrationality and to the other one for proving rationality? The answer is this. People do not consciously maximise their utilities. That is not the claim of neoclassical economics and Milton Friedman already pointed that out in the fifties. People act *as if* they are maximising. That phrase is a Milton Friedman phrase, which he coined already in the fifties. What it means is that people evolve norms of behaviour, rules of behaviour, which usually, in general, yield results that are good for them. This happens by an evolutionary process, by a process of trial and error. People get used to doing things in a certain way, by trial and error. It is not that they are consciously optimizing, but that is the way that their behaviour evolves, in the direction of optimality.

Sometimes it does not fit. Because sometimes the rules of behaviour that they have evolved do not fit the situation at hand, and then people may well behave irrationally, because they did not consciously maximise, they did it in a trial and error way. Kahneman was looking at the exception and Vernon Smith was looking at the rule. Vernon Smith devised experimental procedures for testing the market hypothesis, that supply matches demand, and he was putting people in a situation with which they are familiar, and then they reacted according to the rules that they had adopted. Kahneman was putting people in question and answer situations with which they were unfamiliar, so they behaved irrationally because the rules that they had evolved do not fit those situations.

Let me give just one example: The "probability matching experiment." This was devised way before Kahneman and Tversky, in the fifties of the last century. There is a machine that flashes a red light with probability ¾ and a green light with probability ¼; the subject is told this—that the probability of red is ¾ and that of green, ¼. This is done a hundred times. Each time, the subject must predict which light will flash. Every time he predicts right he gets a euro and every time he predicts wrong he gets nothing. So, people predict red ¾ of the time, and green ¼ of the time; and that is wrong. That way you get 62,5 cents on the average, and if you predict red all the time you get 75 cents on the average. So they are making a mistake.

This is an example of behavioural economics, of non-optimisation. But this is a situation with which people are not generally faced. You put them in front of a machine; they think, what is this machine doing here, what am I supposed to do, what does this guy want from me, let me go home.

But there is a different situation, which people face every day. They have to get to work, and very often there is a choice of routes. Sometimes one route is better, sometimes the other. It depends on the traffic, on accidents, on whatever. Let's say ¾ of the time route A is better and ¼ of the time route B is better. What do people do? Do they take route A ¾ of the time and route B ¼ of the time? No! They *always* take route A. They take the same route to work every day.

You know about this business. This is something with which you are familiar and it costs you, it is your time, it is not a few cents. It is something that is important to you and with which you are familiar and for which you have evolved rules. I am not saying that people figure out that this is best, but they have evolved this way. They take the better route every day and that is it. So my answer to behavioural economics is that people generally behave rationally. Not always, because it is rule rationality, it is not act rationality. People choose the better rule, not the better act.

Andrzej Wieczorek:
Andrzej Wieczorek, Polish Academy of Sciences. Professor Aumann, I have a question about game theory and elections: Can we use game theory to understand politics?

Professor Robert J. Aumann:
Elections are a prime example of game engineering. It is a very big topic, a subject that has been developed by many people over the years. In fact the Arrow Impossibility Principle is part of this. The general upshot of the theory of elections is that it is impossible to do it right. When you have more than two candidates, when you have at least three candidates for one position, then no matter how you do it, there are going to be some anomalies; you are going to have some situations that do not sound right, that *are* not right and that are against public policy. This is a theorem, it is a mathematical fact, you can prove it. You can prove that there is no way of doing it so as to avoid all the pitfalls.

But, there are some better ways and some worse ways. For example, in many countries, when you have an election for the head of state, it is in two stages. If nobody gets a majority in the first stage, then you have a runoff. In the runoff you take the two highest candidates and you pit them one against the other. So it is a two-stage election.

What's wrong with that? Let's take a left-right spectrum. You have a candidate on the left and you have a candidate on the right and you have a candidate in the middle. I do not know how it is in Poland; in Israel we do not have a lot of people in the centre, they are either on the left or on the right, it is a bipolar distribution. The guy in the centre is going to get the least amount of votes, so in the runoff you get the guy on the left pitted against the guy on the right and one of them wins, and that way you have an extreme. You have one of the two extremes, rather than getting what most people would settle for. Most people would settle for the person in the centre. They do not like him, they like the guy on the right more or the guy on the left more, but they would settle for the one in the centre. But he is not the one who is elected by that system.

Steven Brams and Peter Fishburn proposed a method called "approval voting." It is also not perfect, but at least it addresses some of the problems. In "approval voting" you do not say I want candidate A or candidate B or candidate C or candidate D. You say: Who am I *willing* to vote for? Who am I willing to see as prime minister or as president or whatever? You can vote for several people, and then the person who gets the most *approval* votes is the one who is elected.

That is something where you are incentivising people to vote for what they really want. Very often they do not know whether they want specifically this one or that one.

Let me give you an example of that. There was the 1980 election in United States. It was between Reagan and Carter and Anderson. Anderson was the third party candidate, but he had quite a lot of pull, quite a lot of attractiveness, because he was in the middle, between Reagan and Carter. People liked him. In the end Reagan got 50% of the votes, Carter got 40% of the votes, and Anderson got 10% of the votes.

If that had been an approval vote, then Anderson might have been president of the United States; a lot of people sympathised with him, and many didn't vote for him only because they thought he didn't have a chance. And that was because of the polls. So they voted for Reagan or Carter because they thought Anderson didn't have a chance; but if it is approval voting then if you like this guy, you vote for him.

So I've given you a general summary of the theory of elections – that there is no way to do it right – and also an example.

Unidentified speaker:
If, as you say, a person's behaviour is rational, if it is in his best interest given his information, then information is as critical as reasoning. And considering that people are so often misinformed due to very many factors, they may act rationally but not necessarily in their own behalf. How would you manage the information to ensure that people are acting both rationally and indeed in their own behalf?

Professor Robert J. Aumann:
Well, first I would do something about the polls. Not the P-O-L-E-S, they are fine, they are beautiful people, but I would do something about the P-O-L-L-S. This is really a terrible nest of misinformation, about which something should be done. I do not know exactly what. A doctor must be licensed, a tour guide must be licensed. We took a beautiful tour this morning. The lady gave us a card that said "licensed tour guide." Why can't I guide a tour? Why does one need a licence? The reason is that you want a professional job done, you don't want tourists to fall prey to people who don't know what they're doing. But many of the polltakers really don't know what they're doing. They fall flat on their faces too often. Two days before the 1980 election in the United States, the polls said Carter and Reagan were running neck and neck. In a country with millions of people that is a terrible failure. The people who were responsible for that fiasco should be out of business. The government should not allow them to operate.

I don't know why people keep buying these polls, I don't understand it. There is a story of a man who came to a village and put out a poster saying: Tomorrow night there is going to be a tightrope over the gorge, with a raging river going through it, and I am going to walk over the tightrope, and it costs one zloty to come and watch. So the whole town gathered and they each paid one zloty and the man was supposed to arrive at eight o'clock and he arrives a few minutes late and says: Ladies and gentlemen, honoured citizens, you have assembled here to see me walk on a tightrope. I don't know how one does that. So if you want to see a fellow man plunge to his death in the raging river for your zloty, then I will do it, because I undertook to do it. These people were merciful people and they did not want to see somebody die. So they shuffled out. And that was it. And then he says: Wait a minute, wait a minute! I have an announcement. Tomorrow night there will be another performance.

That is what the polls are. They almost always fall flat on their faces. They never say how they did in the past. And again and again we see that what the polls say is wrong. So I'm in favour of accurate information and I think governments should license polltakers and they should license them on the basis of past performance. That is one answer.

Marcin Piątkowski:
Marcin Piątkowski, Associate Professor at the Koźmiński School. I was wondering whether game theory could be used for fixing incentives in the financial sector. As you well know, there is an ongoing global debate about how to fix the situation where the bankers always win. If there is a boom they get their bonuses and make millions. If there is a bust they get bailed out, they move to different jobs, or they simply retire. So what would be your view, how could you help policymakers to improve it?

Professor Robert J. Aumann:
I've often been asked to speak about finance and always say that that is one thing about which I know nothing. And, I can prove it: I do very poorly in my own investments.

By the way, I don't think that financial crises are caused by bankers. I don't buy that. It's true that bankers should bear some of the consequences of financial crises, but I don't think it's their fault. As far as executive compensation is concerned, I believe in the market. These people are not criminals. Somebody offered them the money, the board of directors approved. All this shouting about executive compensation, I don't buy it. I believe in market forces. Not unregulated, not going wild, but executive compensation is not one of the areas that needs the most attention at the moment.[6]

Rector, Professor Andrzej K. Koźmiński:
Thank you very much, Professor Aumann. Let's give a hand to our speaker and his wife Batya, who is also among us and whom I forgot to greet at the beginning. I apologise for it. But I always make mistakes.

Thank you, Professor Aumann for this brilliant performance.

[6] Note added on December 31, 2008:
In light of developments since May '08, when this lecture was given, I must revise this statement. Bankers must indeed bear some of the responsibility for the current financial crisis. And indeed, as Professor Piatkowski suggested, a part of the problem is executive compensation. Not, as is often thought, its *size*, but rather its *structure*. By giving bonuses for larger-than-expected profits, or by compensating with options, you incentivise executives to take risks, and that is what they do. That's by no means the *only* reason for the current crisis, but it's undoubtedly an important one.

On the bailout issue I agree with Professor Piatkowski, and always have, and am sorry that I didn't make that clear. As a rule, failing businesses should not be bailed out, because again that incentivises executives to take risks.

Well, I'm sorry to have been at least partially wrong in my answer to Professor Piatkowski, but at least I'm in good company. And, I said right at the beginning that I know nothing about finance, and at least that was right.