XVI JORNADAS ANUALES DE ECONOMIA
Conferencia de David Levine (UCLA)

Evolution of Cooperation Through Imitation

Thank you for that nice introduction, a much nicer introduction than I deserve. I don't have much to add, after that. So, let me speak.

The title of my presentation is written in the agenda as Evolution of Imitation through Cooperation. I searched my brain to see if I could come up with any talk that might possibly go with that title. Well, I have not been able to do so, so I think I'll talk about the original title, i.e. the Evolution of Cooperation through Imitation.

This is a talk on my joint work with Wolfgang Pesendorfer and it is about evolutionary game theory, and I think I should perhaps begin by giving just a little bit of background and not imagine that everybody in the audience is intimately familiar with the elements of the last ten or twelve years in evolutionary game theory. I suppose I could just start by a brief introduction of what the questions and answers are as they exist at the current moment.

I might start with an example of the kind of situation that fascinates game theorists, and which has led to a lot of recent research. And so I would propose, just to make things very concrete, to examine the following situation: the situation is Germany at the time of the fall of the Communist government, and I don't know if you remember how it was that this came about, but people basically went out into the streets. And when everybody in the entire country went out in the streets, then indeed the government fell.

If you think about this now, from the perspective of an individual living in a repressive regime, such as this German regime, and you face a decision. You might either walk out into the streets and say the government is done, or you may stay inside your home. And, of course, if you are the only person that walks out into the street and say the government is done, that is very bad, and you are very unhappy, and they would put you in jail, and so forth. On the other hand, if everybody at the same time walks out into the streets and says the government is done, then the government falls and everybody is very pleased about that. And the other thing is that if everybody except for you goes out into the street and says the government is done, that is not so good either, because your neighbors will say where were you when we were all out in the streets and you were hiding in your house.

This is an example of what is called by game theory a "coordination game", it is a game in which there is more than one equilibrium; there are two equilibria; one equilibrium when everybody stays inside their homes, and if everybody is inside their homes, there is no reason to go out into the street because you don't want to be the only one to go out into the street. On the other side, if everybody goes out into the street, that is also an equilibrium because you don't want to be the only person hiding in your home when everybody goes out into the street.

Another feature of this game is that there is not only two equilibria but the equilibria are different on how desirable they are. Everybody agrees that they prefer the equilibrium in which everybody goes out into the street, because nobody much likes the regime. So there are these two different equilibria, and indeed in East Germany we saw both the equilibria. We

saw the equilibrium in which everybody stayed at home, and then suddenly one day we switched to the equilibrium where everybody went out into the street.

One of the things that game theorists would like to have is a theory that tells us, when there is more than one equilibrium, why we might be at one of these equilibria rather than at the other equilibrium. And about, I don't know, about 10 or 12 years ago, there was work by a group of economists in game theory helping to answer that question. So I'll list the founding fathers of the field of modern evolutionary game theory and that would be Kandori, Mailath, Rob, and Young, who wrote in fact between the four of them two different papers. And they introduced the idea that equilibria are reached through a process of evolution, and that the process of evolution is an imperfect process which takes place through the accumulation of random errors. So, they have in mind some process similar to biological evolution in which there are mutations, and the interpretation of mutations is that people sometimes do things that are not maybe in their immediate best interests, perhaps by accident, perhaps by design. And this is a random process so that it is possible that many random accidents can occur and this can cause a movement from one equilibrium to another equilibrium.

In particular, in the type of setting that I described as Germany, the prediction is that there is a weighting process between the sort of benefit of the good equilibrium and how costly it is to make a mistake and be the only person to go out into the street, and the theory that they propose makes a very specific prediction. It says you should ask the question… Suppose that half the people have gone out into the street; half the people are in the street, and half the people are in their homes, what should you do? Should you go out into the street or should you stay in your home? And the prediction that the theory makes is that the equilibrium that we should observe over the very long run is the equilibrium that when half the people are in the street and half the people are in their homes, if you would like to go out to the street, if half the people is enough to tempt you to go out into the street, then in the long run we would see the equilibrium where people take the street. But if when half the people on the street is still too dangerous, it is still too likely that you'll get into trouble by going out to the street so you prefer to stay in your home, then the equilibrium we should observe in the long run is the equilibrium in which people stay in their homes.

That is sort of a brief summary of the existing main result and motivation of evolutionary game theory, and I'll make a couple of comments now about what it is that Wolfgang Pesendorfer and myself have done.

I am going to speak now about my joint research with Wolfgang Pesendorfer, which is basically to advance the agenda of evolutionary game theory. There are two motivations for the research that we have done in the context of what I have just told you about the state of evolutionary game theory. The first is we are interested in a somewhat different set of issues. Originally, evolutionary game theory was focused on the issue of coordination equilibria; we are also interested in a setting where there is more than one equilibrium, because that is where evolutionary game theory has proven its worth, at being able to suggest which of many equilibria would be the one that actually occurs. But our interest in multiple equilibria is actually in a different setting; it is in the setting of traditional theory of repeated games, in the setting known as the setting of the Folk theorem. The Folk theorem, I'll tell you quickly what it is and I'll try to give you an alternative example. The Folk theorem says that in a repeated game there are many different kinds of equilibrium. There is an equilibrium where people

behave quite selfishly, in economic terms you might think of borrowing and lending. One equilibrium in borrowing and lending is that nobody lends money because they don't expect to get repaid. And nobody repays, and that is an equilibrium. Everybody is following his or her own individual interest but there is no inter-temporal trade, there is no borrowing or lending, so there is a collapse of the economic system.

The other side is the equilibrium in which people lend money and they expect to be repaid, and the reason they expect to be repaid is that they know that people that do not repay will be punished, because they will be excluded from any future opportunity to borrow.

So that is an equilibrium for the functioning of the loan market, and the point that I want to make is that in this kind of setting there are two types of equilibria, there is an autarchic equilibrium where trade does not take place, and then there is a cooperative equilibrium where there is trade but the reason why people repay their debts that are necessary in order to carry out trade is because of the threat of some future retaliation if they fail to repay their debts.

So clearly, the reason why Argentina has not already defaulted is that they are fearful of some bad future consequences if they were to default on their debt.

But of course there is another equilibrium in which there will be no trade in the future anyway, and so there would be hardly any purpose in paying back debt, since you know that even if you pay back the debt there will be no future benefits from trade.

The question that we want to ask is… In a game theoretic setting where there is equilibrium of autarchy and there is equilibrium of cooperation carried out through threats of retaliation, we would like to know which of these two different kinds of equilibrium will emerge. And this is actually a very traditional question in game theory, one that is asked by game theorists more or less since the onset of the subject, and one really that only the recent method of evolutionary game theory has started to give reasonable prospects of providing answers to.

The basic idea is then taking the evolutionary approach to this problem; but there is another innovation that we propose to complement this procedure of studying the issue of cooperation and the emergence of cooperation for the evolutionary game theory, and that is a modification to the way the evolutionary game theory works itself. Let me explain what that modification is.

In the original evolutionary game theory, you have to understand how you get from the equilibrium where people stay in their homes to the equilibrium where people go out into the streets. The account that is given of that in the original evolutionary game theory is that changes take place through a process of mutation. Individuals perform actions and act in their immediate self-interest. Then one day somebody gets up and gets a brainstorm and says it would be a good idea for me to go out into the street. Now, if everybody else does not go out into the street that day, that is a very bad decision. Now, if it happens that, say, 75% of the people wake up the same morning and have the same brainstorm and they all go out into the street, then it works, then the revolution takes off. Now, this notion that this kind of shift from one equilibrium to another equilibrium takes place because people should simultaneously and accidentally by coincidence all have a brainstorm on the same morning of the same day seems to many people to be an impossible account.

The proposal that we have made for modifying this model is that the propagation mechanism through which enough people adopt the "bad" idea, that is since going to the streets by yourself is a bad idea but it becomes a good idea for enough people to do it… The propagation mechanism is not one of mutation in which people randomly decide to do something but rather one of imitation. So that one person, leading the revolution, sets foot on the street hoping that other people will follow. And other people look and say, well, there is somebody out in the street, maybe it is a good idea for me to out on the street, and so a second person follows. Now, it is not still a good idea for the second person because there are no further followers, and they are both in a pretty awful shape at this stage. But the idea that a third person, a fourth person, a fifth person might follow perhaps is not that impossible, and as more and more people go out into the street the chance of followers becomes greatly increased. So, the proposal we would make is the propagation mechanism that is interesting for human societies is the propagation mechanism in which there are leaders, in which there are followers, where people are more likely to imitate something that they see somebody else doing regardless of whether it is a good or a bad idea, then they are not simply doing something for completely random reasons.

The chance that I wake up in the morning and randomly decide to go out on the street, I think is very small compared to the chance that if there is somebody else standing in the street then I may decide to join them.

So that is our basic hypothesis, and we adopt that hypothesis for two reasons: one is because it seems to us to be a more plausible account of how this kind of shift from one equilibrium to another equilibrium takes place, but the second reason, the practical reason, is that it makes the analysis possible, it turns out that that assumption leads to a much simpler mathematical theory and one which as a consequence we can apply to the more difficult problem of cooperation enforced through threats of retaliation. It is very difficult to study that using the theory of mutation. The theory of imitation is a simpler theory and therefore can be used as a tool to study more complex problems.

I have now given you the basic introduction. I think what I should do next is to summarize the conclusions that we have reached, and then, for those people who are interested, I'll briefly make a more technical description of how the model works.

But for the group as a whole, let me try to tell you what conclusions we are able to reach through applying this particular method.

First, the setting that we study. We are interested in a setting where it is possible to have equilibria, cooperative equilibria, enforced through threats of retaliation, and non-cooperative equilibria, where people do not cooperate with each other. And one of the things we did in this study was try to find a simple setting with that type of multiple equilibria where retaliation takes place, and the setting that we hit upon was a setting of the following sort. What we imagine is we imagine people who are matched to meet each other to engage in some sort of interaction. And there is a variety of things that they can do. They can behave more cooperatively, they can behave less cooperatively, they can pay debts, they can make loans, they can make payments, they can make punishments, they can destroy things, there is a variety of options available to people, and people, as they engage in these interactions, make certain commitments to the kind of strategies or rules that they will follow during the interaction that takes place. And the rule that you follow is something potentially that the

person you meet may be able to learn about. See, you might be the kind of person who repays your loans. Now, I want to imagine that when you meet somebody to ask for a loan they receive some signal, possibly a noisy signal, of whether you are in fact the kind of person who would repay his loan. What I want to imagine is that when you employ a particular strategy there is the opportunity for the other person with whom you are interacting potentially to know what that strategy is.

And so, let me describe now how we can have two different kinds of equilibrium in this particular setting. One possibility is a strategy in which you don't make loans, you don't repay loans, you are autarchic, you have no trade. And if other people use this strategy and are not trading, there is really no advantage for you from trading, because you'll never get repaid, because nobody else is going to repay loans. So, you can go and lend money to people, but nobody repays loans, so it does you no good. You might as well also follow the crowd and become a person who does not make loans, does not repay loans, and so forth.

But now, the other possibility is that you might adopt the strategy of repaying loans and making loans to people who have repaid their past loans, good people, people who are good people like us, we would make loans to, but people who are identified to be bad people who do not cooperate, we do not make loans to.

And so now, since when I meet you I receive a signal, I can determine, perhaps with some degree of imperfection, whether you are the sort of person who would repay a loan; I only make the loan if I discover that you are the kind of person who would repay a loan.

Now, in such a world, should you choose to be the kind of person who will repay a loan or should you choose to be the kind of person who will not repay a loan; and if you choose to be the kind of person who repays loans, then people will make you loans.

Now, I am hiding a lot in this mechanism that says it is possible to learn, that you can tell an honest face when you see an honest face, but nevertheless this is something we observe in practice, and it is quite important… I am thinking of a situation where the signals can be noisy signals. You look like you have a honest face, but actually sometimes you are a crook. So that possibility is allowed as well.

I will now describe then our conclusions… We have a situation which is a very common situation in game theory. We have the multiplicity of equilibria, the good cooperative equilibria and the bad not so cooperative equilibria, and let me describe what conclusions we can reach applying this method of imitative evolutionary game theory to it.

We study basically two different situations. The simpler setting is the setting of perfect identification, so that the signals have no noise. If I meet somebody I can simply tell this person is a cheater and this person is not a cheater and there is no ambiguity in that determination. In that setting, what happens, not surprisingly perhaps, is that the equilibrium that emerges in the long run as the result of the evolutionary process is that there are no cheaters. Everybody cooperates and everybody employs the cooperative strategy of cooperating only with other people who cooperate. And, of course, if you meet somebody who is not a cooperator, then you do not cooperate with a person like that. So that is the enforcement mechanism.

Now, the interesting feature, that at least was not obvious to me when we started out to do this work, is how do you treat people who are not the right kind of person. If you think of the cooperative group, the group who are the successful group, who cooperate with one another, how do these people behave when they meet and they perfectly identify somebody who is not a cooperator. They are supposed to apply a punishment against this person, and what is the form of punishment that they should use? And the answer about the punishment is this, what they should try to do is they should try to maximize the difference between the utility they receive and the utility the other person receives. In a certain sense it is like warfare, but the main difference is that you should not try to destroy your rival but rather you should try to maximize the gap between your rival and yourself. What you should try to do is do things that are harmful for them but at a very low cost to yourself, because that will give you the greatest difference in benefit to yourself and the other person. This theory actually makes a very specific proposal as to how cheaters should be punished. There is a particular way in which they should be punished and a particular criterion for punishing them.

This has an interesting extension into the world of imperfect signals. The second setting that we study is a world where you receive a noisy signal that only imperfectly enables you to identify whether the person with whom you are interacting is following a particular strategy. And we do this in a somewhat simple context; we do this in a context of particular games on which I'll just briefly say that there is two parts of the utility that you receive from this game. There is the utility that depends only on the action taken by the other player and the utility that depends only on the action taken by yourself. So, technically there is an underlying assumption of separability between the two parts of utility, and that will allow me to describe the result.

The idea is that you receive a noisy signal about the other person. And the strategies that are successful in the very long run are strategies that actually answer only a very simple question. They do not attempt to identify in a complicated way exactly what strategy is being used by the other person. In fact, the strategies that emerge as the long-run equilibrium strategies only attempt to determine whether the strategy they are meeting is the same as themselves or different from themselves, and if it is different, the particular form of the difference is of no consequence. And so, the strategy will treat the opponent differently depending upon how likely it is that the person with whom you have met is employing the same strategy you are employing or is using a different strategy than the one you are employing. And perhaps one interpretation of this is a theory of endogenous preference towards altruism or towards spite.

So, I meet somebody for an interaction, and I receive a signal about that person and the type of person they are. And based upon that signal I can draw a conclusion. I can calculate the conditional probability, what is the probability that this person that I have met is employing the same strategy that I am employing, what is the probability that this person is the same kind of person that I am. So I calculate that probability, and what the theory says is something very specific about what I should do. I should maximize a weighted sum of my utility and the other player's utility, and the weights depend upon this conditional probability, and the weights depend on the following way.

If I know with a 100% certainty that this person is the same as me, I employ equal weights on my utility and the other person's utility, I am perfectly altruistic towards people whom I perceive as being perfectly like me, I treat them the same as I would treat myself. If I receive

conclusive evidence that this person is different from me, so I am certain that this person is different from me, then I employ a positive weight on my own utility and a negative weight on the other person's utility, a weight of 1 on my utility and a weight of –1 on the other person's utility, so I am maximizing the difference between my utility and the other person's utility.

If I receive a signal that is completely uninformative, so that it is equally likely that this person is the same as me or different than me, then I employ a weight of 1 on my own utility and a weight of 0 on the other person's utility. I behave completely selfishly, with total disregard for the other person's welfare, positive or negative.

If this conditional probability moves in between 1, the person is the same as me, and 0, the person is different than me, what happens is I behave more or less altruistically in the sense that I choose my behavior in an effort to maximize this mixture of my utility and the other person's utility, treating the person more favorably as I perceive them to be more like myself, and treating them unfavorably as I perceive them to be less like myself.

Now, maybe only an economist would find such a thing surprising. Most people, I think, understand perfectly clearly that you are kinder to those in your immediate family, for example, than to complete strangers, and things like this. Those of us who have been trained in social behavior will find this transparent, but economists for so long have had this view of people completely selfish and uncaring, either positively or negatively, about people around them, that perhaps for economists, at least, this result is a little bit surprising.

That is sort of the overall thrust of this. So, now, for people more interested in the details, I'll provide a little detail. And the detail that I'll provide is a careful description of the mathematical model. We'll use a simplified mathematical model to establish these results. Let me briefly tell you that, because if you know the assumptions, you'll understand better and more clearly the limitations of these particular results.

The formal details are these. We study a situation where players, pairs of players meet to play a symmetric normal form game, and these players are drawn from a single population. So the world in which we work is a world with a large population of identical people, and these people are drawn randomly in pairs of two, and each pair of two people play a particular game against each other. Everybody is the same.

The context is one of a finite game. The reason for putting this up is simply for you to know that *s* stands for strategy in some of the assumptions that are made further down. And players, who are identical, receive utility which depends upon two things: one is the strategy that they employ and the other is sigma, which are the probabilities of different strategies employed by the player whom they meet.

That is just a description of the game. Now, what I really want to describe is the process of imitation, because that is what is at the heart of the results that we have. And I'll also describe to you which are the assumptions that are general to the evolutionary game theory literature and what it is that differentiates our assumptions from the more traditional assumptions. At the moment I am making only traditional assumptions; I am simplifying assumptions to make it easier to analyze the model.

There is a fixed number of players; so, the world has *m* players. Each player applies a particular strategy. So that at any moment in time, there is a room full of people, like this, to identify each person's particular strategy, and then we draw people, and the players employ their strategies against one another. And each time we play, what happens is one person and one person only is allowed to change their strategy. And that is for simplicity. And so, the way in which they change their strategy is that one employs one of several different techniques.

There are three different ways in which a person might change a strategy. The first way in which people might choose their strategy is through the process the evolutionary game theory has described as mutation. And what does mutation mean? It means that you have a universe of strategies, a 100 strategies you may choose from, and you pick one at random. So there is no element of rationality in the process. Now, what I am describing here is innovation. Innovation is that you try something new that nobody else perhaps has tried. So, you see what happens. This is not a process of rationality; it is a process of irrationality, of randomness.

Now, the second way in which people might choose a strategy is through the process of imitation. And to be concrete, we propose a very simple model of imitation. So, imagine that you are born into the world and have to decide how to play. One possibility is you know nothing, in which case choosing randomly is more or less what you are forced to do. But suppose I give you one piece of information; suppose I identify one of the persons in the group and tell you how that person is playing. Now you have the information that at least one other person in the group thinks that cooperating is a good idea, whatever it is that they are doing. The logical thing to do, if you have no other information and you see somebody doing something, is to imitate that person. That makes more sense than picking at random from all the possibilities, since at least one person thought that that was a good idea. So the second is the idea that you pick a strategy by imitation. How does it work? You pick a strategy in proportion to how many other people are playing that strategy in the population.

So, if 99% of the people are doing the same thing, when you observe one person at random most likely you'll see what 99% of the people are doing, and you'll play the same as those 99% of the people are playing. There is a 1% chance that you pick the strange person who is doing something different, and you will end up imitating that other person.

I am sorry; the numbers are confused if you want to follow the slides. Number 3 was the choice completely at random. Number 2 is what I have just said, the imitation.

Finally I need to tell you what is the third way in which people might choose their strategy, which is through a more rational process. And this more rational process will still be limited to the strategies the other players are actually employing. So, I enter the world and I do not have the ability to calculate and determine how the best way to play. But what I do have is the ability to determine how much different players are getting from what they are doing. So I can observe Bill Gates, and I know what he is doing, and I know that he is rich. So among the universe of people in this room, basically I pick up the richest person, and I imitate them. That is a much more rational thing than picking a person at random; now I can hope that if I behave like Bill Gates, I can be rich. I don't want to exaggerate, but it is clearly better to

imitate successful people than unsuccessful people. And that is the third form of choosing a strategy.

Most likely, what you do is you go with a high degree of rationality. So the basic assumption here is that the most reasonable and rational choice is the most likely, and the most likely thing that is going to happen is that you are going to imitate a successful person. So that is a sort of residual probability. And then we have a smaller probability proportional epsilon, since epsilon you always think of as being a small number, you are forced to choosing one of these alternative methods of choosing a strategy. So with the probability proportional epsilon you would imitate one person, and with probability of epsilon raised to the power *n*, you would innovate, you would try something new.

Now I am in a position to describe to you with some degree of clarity, how this model differs in an essential way from the standard evolutionary model. It differs from the standard evolutionary model in the following way. In the standard model imitation does not exist. So this constant *C* would be equal to 0. There is no imitation; there is only innovation.

Now, what do I mean by imitative model? I want innovation to be less likely than imitation. And the heart of the matter is to say that imitation has to be sufficiently more likely than innovation, that innovation must be much less likely than imitation. And the precise assumption is the assumption about the magnitude of the coefficient n and how it compares to population size *m*. And the specific assumption is that the coefficient n should be bigger than the population size *m*, and what does that mean? It means that if some person does something, it is more likely that that person will be imitated by the entire population than one person in the entire population may have a brand new idea.

That is the heart of the assumption. It is a strong assumption, but because it is a strong assumption it leads to great simplification. I don't want to argue that I believe that that is the correct magnitude, but I certainly believe that innovation is less likely than imitation. I don't want to argue that in every practical situation imitation is that much more likely than innovation, but as a theorist my job is to study the most extreme cases, they are easy to analyze, and I hope that the results will shed light also on the more intermediate cases.

And I shall briefly highlight then, as I said, how this helps us to get results in a model that is too complicated to study, in a simpler setting. The reason for this is that basically evolution is a competition between strategies; it is a warfare in some sense between different strategies. In the standard theory there can be many strategies, all competing at the same time in order to survive. And it can be very complicated when many strategies have to compete with many different rivals. But in a world of imitation, it is highly unlikely that any strategy will ever have to compete with more than a single rival. The most likely is that everybody will end up doing the same thing. If everybody is doing the same thing, then imitation will lead you to do the same thing that everybody else is doing.

And it is only in the case of innovation when somebody comes up with a brainstorm to lead everybody to do something different. So imagine a world everybody is conforming, playing the same way, doing the same thing, and somebody comes along with a new idea and proposes something different. So, what is going to happen? Basically either all will think this is a bad idea and go back to doing the same thing they were doing, or enough people will imitate the innovator and move somewhere else. But this whole process of imitation that would carry you to another equilibrium is very likely to take place before anybody has a third idea. During

the whole process of moving from idea 1 to idea 2, it is very unlikely that a third idea will enter the population; so that idea 1 and idea 2 basically compete just with each other, which I think in fact is sensible, you know there are not millions of ideas competing at once, there are certain key ideas at any moment of time that are competing with each other. So, regardless of whether this is a good assumption or a bad assumption, it is certainly the case that this is the assumption that makes the mathematics much easier to analyze because we can study simply the context between pairs of ideas and not have to worry about contest between large numbers of ideas.

I think this is a good place to end.

- I have a question, which is kind of prospective, where is learning theory going?

- Let me answer this in the context of this specific research. One of the questions we are interested in is the organization of the state. You see, modern states are based upon a very high degree of cooperation. Everything from judicial systems, enforcement, the entire standard of living that we have is because we have an enormous degree of cooperation between people in the modern state. And one of the things we are interested in is how and why that came about, how did we reach this complex social organization where we are all extremely interdependent, where all human relationships are mediated through both formal and informal systems, cooperating with people with either implicit or explicit threats of various types of… ranging from social disapproval to being put in prison, and so forth.

If we sort of move back in time and look at unsuccessful societies, we discover that dimension in which they seem to be lacking. So I would say my own personal concern is for, beginning with very simple theories and very simple states, try to understand something about how it is that these complex forms of organization thrive in some sort of circumstances but not in other circumstances. Why is it all such a mess in Africa, although here in Uruguay everything is very nice.