

[SQUEAKING]

[RUSTLING]

[CLICKING]

PROFESSOR: OK, so what we're going to do in talking about social preferences is the following. We're going to start discussing what our social preferences are. We're going to just talk about what do we think social preferences are and how do we think about them, how do we measure them. This is partially what we did already in class last week, which gives you some sense of some ways of measuring them. There are also other ways of doing that. We're going to discuss that. We're going to also discuss a bit kind of like the measurements that we did in class that seem sort of pretty contrived in various ways. Are they actually predictive in some ways of real world behavior that we care about?

Then we're going to ask the question, well, it seems like, from measuring social preferences, as we did in class, but also in some experiments that we see, it looks like people are actually quite nice to each other. But we're going to ask the question, are people just genuinely nice to each other? Do they actually care about others? Or is it just like, in some ways, perhaps, social image or other issues about like they want to look nice or they want to feel good by being nice to others or look like they're nice to others, as opposed to generally caring about the well-being of others. And we can talk through some experiments and some evidence that lets us disentangle between those kinds of hypotheses-- are people genuinely nice or are they sort of just trying to look nice in various ways?

And then we're going to talk a bit about evidence towards the end. This is the paper by Gautam Rao and others, which is the question about are social preferences malleable and can policies change pro-sociality. That is to say, can we do certain things, can we mix people in certain ways-- it can be like this will be how roommates are chosen, and the specific people we're going to talk about is from India, where poor students are mixed with rich students. And does that affect how nice people are, the rich students, in this case, towards poor students, but also towards other people in general.

So are there some policies that we can perhaps implement that help people, society as a whole, being more pro-social and people being nicer to each other or doing things that are good for the greater good, if you want, overall.

As before, we're going to look at two types of evidence. We're going to look at lab experiments and field evidence. Lab experiments are the types of experiments that you saw so far. And field evidence are sort of the type of experiments where we look at students being mixed, rich and poor students. What happens then to people's social preferences when you do that?

OK, so now let's sort of back up a little bit into review, very quickly, of what we did last week. I showed you sort of an action, three games to measure social preferences. These are not randomly selected games, these are the most prevalent, the most commonly-used games in behavioral and experimental economics, which is the dictator game, ultimatum game, and the trust game. You can read descriptions and discussions of these games in the background readings, which is Camerer and Fehr 2004.

Now what are these games? I'm going to briefly review them right now. And then I'm going to go to the evidence and go back to evidence that actually uses these games. So what are the games, just to recap? There's the dictator game, which is very simple, where the dictator gets some money allocated to him or her. And there's another person who is the recipient. And a dictator just gets to choose how much money does he or she wants to give to this other person. That's it. You can think of this as some raw concern for others, essentially to say, when you don't know anything about this other person-- often these games are anonymous-- here's an anonymous person whom you could give some money. And often it's another student. It could be like a poor person in Kenya or elsewhere. Here's a person with whom you don't have any interactions whatsoever. How much money do you want for yourself? How much money do you want for this other person?

The original version is just sort of 1 to 1. Any dollar that you give is given directly to this other person. There's different versions where essentially any dollar is converted, where essentially the exchange rate is varied, where it's cheaper or more expensive, if you give a dollar, the other person gets \$2.00 or \$0.50 or the like. OK, so that's the dictator game.

What's the ultimatum game? It's sort of like a bargaining game if you want. It's sort of very, very simple game theory if you want. You can use that, which is essentially used again by experimental economists.

What they do is there's two players again. There's a proposer or the sender that gives, again, like, some divisible pie. Think of this as \$10 or the like. Often it's money. And then you say a portion of x is given, then, to the responder. The responder then has the chance to either accept or reject. If the responder accepts, then the division is just implemented. If the responder rejects, then essentially both players get nothing.

And why is it a game now? Well, it's because essentially the proposer needs to sort of take into account what the other person is going to do and have some beliefs about that and anticipating that, then sort make that choice of how much money would you like to give.

Finally, there's the trust game, which is, in fact, very similar to the ultimatum game except for that the amount sent by the sender is tripled before the decider is supposed to say-- before they decide how much to return, the recipient decides how much return, if anything. So now essentially, it's like a high-return opportunity to give money, essentially in a situation of high trust. In particular, there's communication.

If I'm playing with somebody, I would tell that person, yes, give me all of the money. And if you are nice to me, that money is tripled, and I give you half of it back, or the like. So both of us are better off. If there's lots of trust between us, that's actually implementable, and a sense of, like, I'm telling you, give me that money, you'll give me that money, and then I'm actually returning that money.

But of course, if there's no trust, then that's what the game name comes from, the trust game. If there's no trust, I'm going to tell you, give me a bunch of money, you might actually do that, and then I'll just keep it all for myself. Or you might not even believe me in the first place. You might just say, yeah, I'd love to, but I don't trust you, Frank. I give you nothing. And then essentially, nothing is tripled.

And that's meant to then measure some form of trust in society or that people have with each other. We're going to go back to that in trying to see what does it actually measure in the real world. So just sort of having said that, those are those three things you saw.

And then there's different versions of these games that could be sort of like-- you can vary the privacy. Are these games private versus public? Is it anonymous? Do you know who the other person is versus not? Is there communication? Can you talk to this other person or not?

And what are the stakes? Are we talking about hypothetical choices? Are we talking about money? High stakes with money. Are we talking about other goods like apples or chocolate?

Any questions so far? We'll get back to the actual evidence, but I should make sure that these games are clear to everyone. That's sort of what happened last time. OK.

Now, backing up a little bit now, then the question is, like, what is, in fact, social preferences? And so most economic analysis in some way assumes away social preferences and sort of assumes self-interest, very narrowly defined. And so that means that people essentially only care about their own interests and their own outcomes, and then the market is sort of taking care of the rest.

And this is at the heart of what Adam Smith originally was saying, I think, in *The Wealth of Nations*, which is like, "It's not from the benevolence of the butcher, the brewer, or the baker that we expect our dinner, but from their regard for their own interest. We address ourselves not to their humanity, but to their self-love, and never talk to them of our necessity, but of their advantage." That is to say that when everybody is very selfish and cares about themselves, just comparative advantage and the ability to trade with each other will make sure that other people do useful things for us.

So you can do the work you do, make some money, and then go to the bakery. And the baker will give you some bread in exchange for some money, not because the baker likes you or wants to be friends with you or the like, but because you pay them, and it's in their interest to give you good bread because you're going to pay them, and you're going to come back perhaps and buy more bread in the future.

So that's sort of a pretty reasonable assumption in saying people are doing useful things for each other as part of a market transaction. And that's what a lot of economics has assumed and done for a long time.

Now, that is not a bad assumption at all in the sense of, like, it's actually pretty realistic in various ways in thinking about humans. The question is not necessarily whether there are some instances of whether people care about others. The question that we're going to ask is, like, well, what do we miss by ignoring social preferences? And so it's pretty clear that preferences depart from pure self-interest in non-trivial ways.

And the question we're going to ask is, like, are there some concrete settings where we really care about this and sort of say, well, we should model things differently? And if we do that, we can understand certain phenomena better. So the goal is here to understand how common and important these departures are and what their nature is, like, how can we best model these.

And here's a nice example of perhaps how one might want to think about social preferences and how perhaps people think about social preferences. So this is a picture-- I don't know how well you can see this-- taken in Hawaii, which is a box where banana bread is being sold. It's freely available to everybody, so you can just take it. I think there's some price tag somewhere.

There's also a lockbox next to it on the left of the picture where you can deposit payments. So the banana bread costs some money. You can deposit your payments. There's nobody that enforces it, but the lockbox is such that it's actually locked in a sense that you can't take away the box.

So when you see that picture, what perception of social preferences or humanity does that reflect? How do we think about this picture? Yes?

AUDIENCE: That we're going to be honest and give the money based on how much [INAUDIBLE].

PROFESSOR: Right. So one assumption, to some degree, is, like, well, people seem pretty nice in general. It seems in a sense of, like, we allow them to just take the banana bread and run away and not put in anything. We sort of ask them nicely. It costs, like, \$1 or \$5 or whatever. Please put in the money. It seems like the conception to some degree is, like, people are going to be nice and just do that. Is that all, or is there something else? Yes.

AUDIENCE: Maybe there's an expectation not that everyone would do the right thing, but that the loss from some people who would just take the bread and run will be maybe less than, for example, the cost of labor or staffing the booth.

PROFESSOR: Right. Exactly. So there's some-- the expectation, to be more precise, is not everybody is going to be nice, but on average, most people are going to be pretty nice or adhering to this normal rule that you should pay some money. And some people might run, but as you say, it might not be worth standing there in the heat all day to make up for that. Exactly. And then what else is there? Yes.

AUDIENCE: It gives a perception of, like, you don't expect them to steal banana bread, but you expect them to steal the money, so you lock it.

PROFESSOR: Right, there's a lock. So it's like, in some sense, you sort of think that, on the one hand, the assumption is most people are going to be pretty nice. But if we allow them to run away with all the money, at least there are some people who, if allowed, would not be nice at all. In fact, they would take away the money and not only not pay for the banana bread, but also take away the lockbox.

So that's exactly what I've written down here, which is to say people in general are nice enough. Most will probably pay for the banana bread, so it's not worth sort of standing around and ensuring that. But people also can be selfish. At least, some people will be quite selfish.

So if a lot of cash were easy to take, someone at least would take it. Notice that's not an assumption about everybody, or even the majority. It's just to say there's a fraction of people who, if given the opportunity, might take away our money and run from it.

So that's actually a pretty good sort of perspective for thinking about social preferences, which is to say that self-interest is probably a major driver of behavior in many economic contexts. In some situations, self-interest is not the main motive, in some sense. Again, we don't need to necessarily enforce this for everybody. People will be quite nice.

But there will be-- some seemingly small departures from pure self-interest can dramatically influence economic outcomes. That is to say, if we couldn't lock the box, if you were not doing that, somebody would take away this box. And now we couldn't offer banana bread to anybody, and the whole transaction would go down.

Put differently, if we had things like a theft or robbery and so on, if there's some people who are really mean and essentially take away all of our money, now we can't do any transaction. I can't have a shop trying to sell stuff to people. This is the examples [INAUDIBLE] positive and negative. Small things can actually make huge differences for others.

This is like helping a stranger with directions or helping a stranger who wants to go to the hospital and so on. That could make huge differences in their lives, even if the-- or calling an ambulance for somebody, that can make huge differences people's lives, even though the action from your side is actually relatively small. Helping a fellow student with computer trouble or other issues that they might have, helping somebody who is sad on a given day feel better about themselves.

In fact, as you might know, it's the Random Acts of Kindness Week this week, so doing random acts of kindness might actually make a huge difference in people's lives. In fact, the problem said, well, we'll ask you to do some of that, in fact.

Washing your hands is particularly pertinent right now. Again, you might think you're doing that to protect yourself, and surely you are. But in fact, there's also a huge externality potentially that, like, by protecting yourself, you protect others from potentially getting sick. And that could make a huge difference in people's lives. Any questions on this?

OK. So now we're getting back to the ultimatum game. I'm going to ask the question about how are people actually behaving in real situations. How do people behave in general? And perhaps, what does game theory say how people should behave? How are people behaving in reality, and how does that compare to how you guys behaved here in class?

So again, the ultimatum game, you should be familiar with. There's a proposer, a responder. Proposer proposes some split of the pie. Responder can accept or reject. Now, what does game theory predict what people should do? Like, in the absence of social preferences?

Yes?

AUDIENCE: [INAUDIBLE].

PROFESSOR: Right. If you're the recipient, you'll care only about your own output or your own outcomes. Even if I offer you \$0.01 or \$0.05 or whatever, whatever is worth the transaction cost, you're going to accept everything. So if I'm playing with you, I'm going to anticipate-- I know that you will accept essentially any offer that's positive.

Even zero, you might even accept, because it's sort of like, why do you care? But you strictly prefer a positive offer over no offer, so you're going to accept that. And then if I'm playing with you, anticipating that, I will just give you essentially the lowest possible amount, and that will be the outcomes. Right?

So here it is. The prediction is very clear. The responder cares only-- this is, again, game theory, but game theory in the absence of social preferences. I should be more precise. So the responder cares only about money, so she will accept any x that's smaller than 0, any amount that's given, because that's better than not having any money.

The proposer will understand that, will only care about money as well, so offers as little as possible. The proposer doesn't care about how the recipient is doing. So therefore, the proposer will just offer as little as possible to maximize their own payout. And then an equilibrium, that will essentially be the zero or smallest possible offer will be made, and that offer will be accepted. So there will be no rejections, either because, essentially, there's no reason to reject, and there will also be no large offers of the like because there's no reason to give more than minimal amounts.

OK. So then what are typical results in the ultimatum game? What do people actually find? Well, so most people offer between 40% and 50% of the pie. Such offers are mostly accepted.

The acceptance rate is increasing in the offer, which is to say if you offer, like, 10%, you have a pretty good chance of being rejected. If you offer, like, 30%, 40%, 50%, you're much more likely to be accepted. And you offer, like, 70%, 80%, almost surely you'll be accepted. And an offer that's below 20% are mostly rejected. Why is that? Or what's the reason for that, presumably? Yes?

AUDIENCE: [INAUDIBLE].

PROFESSOR: Right. And why are you sort of upset if somebody offers you 10%?

AUDIENCE: Because it shows they're disrespecting you.

PROFESSOR: Yeah, something about fairness or disrespect, or it's just not a nice thing to do. And now you'd rather not have any money or not take \$1 or whatever they offer to you. You'd rather reject that, so you're willing to give up \$1 to essentially get back at the other person. Yeah?

AUDIENCE: It could also be that for some people, the thought of settling for such small amounts of money is embarrassing, or it feels like it's demeaning or [INAUDIBLE].

PROFESSOR: Yeah, so it's one of those. Some of it is-- yeah, exactly-- I think, demeaning, or just feels like-- so one way to put this is you have inequality aversion in some ways in the sense of you just don't want to be the person who has \$1 when the other person has 9, or you don't want to have \$0.10 when the other has, like, 9.90. So you really don't like inequality in some ways.

There are some versions of this game where the computer sort of allocates these choices. And you can look at how people are rejecting those offers as well. And they tend to not do that, which sort of suggests that it's really coming from some sense of people don't care so much about the outcome, per se. They care about the fairness or the intention of the other person. But you're right. There's other potential considerations at play.

Now, what do things look like in our class? In fact, this looks quite similar to what people do in general. I know you're very special, but when it comes to the ultimatum game, you're actually pretty ordinary. So what you see is, essentially, most people offer 40%, 50%. Some people offer somewhat more. You see in-- you can't see this that well. You see it in red are the rejections.

So rejections tend to be on the left side, in particular when people offer between 1% and 10%. There seem to be around some rejections, so 50%, or maybe it's just below 50%. But usually, if you look at, essentially, about 50%, or 40%, 50%, 60%, people tend to not reject any offers. Any questions on that? Yeah?

AUDIENCE: Could it also be that people are anticipating future games? Or maybe not [INAUDIBLE], but sending a signal, like there's any similar [INAUDIBLE] be better off [INAUDIBLE]?

PROFESSOR: Yes, that's a great question. So I kind of-- all of what I was doing, and I should emphasize this a bit more-- the way I implemented some of these games in class, some experimental economists would get a heart attack in terms of how sloppy that was in a sense of, like, a lot of these games, usually when I show you these results, these are usually very careful games, where, essentially, usually it's anonymous.

Usually it's a one-shot game. Usually it's very clear that it's a one-shot game. You will just never see that person. You never even interact with that person. It'll be clear whether it's public versus private, and so on. So usually in those games, things are extremely clear. There is no strategic interactions, and so on. All of these things are essentially shut down.

And then you can look at, randomly vary, some of these characteristics and look at deviations of that. So you're right. In our game, in fact, you might sort of think that, let's just be nice now, in part, because the other person, maybe I play with that person again. And I wasn't really very precise about that either. I think actually they're all like randomly allocated. So in fact, that was not a consideration, but it wasn't entirely clear. Or for example, some people were selected, and it happened to be that it was 50/50 selection. That may have affected others in the class, and so on.

So a lot of these things were fairly sloppy in terms of the actual implementation. In the real games, that would be all shut down and done much more carefully. There were some other questions, I think? No. OK. Yeah?

AUDIENCE: Was this just for the chocolate?

PROFESSOR: I think this is-- I looked at all of these. The ultimatum games all looked very similar. So this is, I think, for-- this is the money one. But even for others, it looked fairly similar.

I'll show you some variation in the dictator game in a second, where we have, like, do the stakes matter? Does it matter whether it's private versus public? Does it matter whether it's chocolate or money? And perhaps, why might it matter? We'll get back to that in a second. Turns out, it does matter. I'll say a little bit about that in a second. Yeah?

AUDIENCE: So the fraction of rejection in the green [INAUDIBLE] others, and then [INAUDIBLE] is much higher. The number of rejections in, like, \$1 to \$10 much higher than [INAUDIBLE]?

PROFESSOR: Yeah.

AUDIENCE: [INAUDIBLE] explain the notion of fairness or the equitable distribution of the pie you just mentioned, because [INAUDIBLE] accept when they are being given 90% of the pie, whereas they are just not ready to accept [INAUDIBLE].

PROFESSOR: Right. So that's a good point. So you were saying earlier, maybe some people have some issues about inequity or inequality across people. So if that were the only explanation, then you would have rejections not only at, like, if I give you \$1, and I keep 9, but also if they give you \$9, and I keep 1. So that means there's something else.

I think the fairness, people don't object to inequality. If I say, OK, I'll give you a lot, and you have more than I do, people don't object to that. People object to, if I keep a lot and give you very little, presumably because I'm selfish, and I'm not nice to you. So that's kind of the sense of fairness. So it's a fairness that goes beyond the outcomes.

So we're going to talk about this for a bit, which is to say you could just look at the outcomes and say, how much is the inequality of outcomes? And I care about that. But that doesn't seem to be the case. What seems to be rather the case is that people care about intentions one way or the other. It's like, what was I thinking, or was I a mean person when I was doing this?

And if you get the sense that I was not very nice, you might reject almost regardless of what the outcome will be eventually. And of course, you will reject less if it's more costly for you to do so. But perhaps it's also the case if I do, like, 70-30, you might say, well, that's maybe not nice, but it's also not that unfair either. OK.

Right. So now, one thing you can ask is how do these games look like across countries? And as I said, this game has been extensively played. It's a game that like lots and lots of people have played in all sorts of situations. And the patterns are remarkably stable across countries. And this is also the case-- like, every time I play this in class, these patterns look exactly the same. So whether people are playing in Pittsburgh, in Ljubljana, in Yogyakarta, Tokyo, play all essentially the same way.

And even increasing people's stakes to, like, several months pay by running essentially the experiment with poorer populations, where you can pay, essentially, days' or weeks' or even a month's pay, people seem to play very similar across places? Why is that? It seems to be like somehow rejecting becomes a lot more costly when I offer you, like, a day's wage or the like. But why do you think it looks fairly similar as well? Why do people still reject? Yes?

AUDIENCE: If your goal is to get back at the other person, if the pool is just bigger, then you're hurting them more by denying them.

PROFESSOR: Exactly. So it's exactly as you said. There's two things going on. On the one hand, rejection becomes really expensive to me. So if there's a month's wage, and you give me 10% of that and keep 90% of that, rejecting 10% off a month's wage, that's three or four days or whatever. That's really costly for me to do.

However, it's also really mean of you to do that, and I can really screw you over now by sort of rejecting that offer. So those things essentially seem to be more or less [INAUDIBLE]. They go into opposite directions. Empirically it seems to be that cancels each other out. That's exactly right.

So now, there's a very interesting paper then that does this ultimatum game in hunter-gatherer societies and asks the question about, like, while it seems to be that people behave fairly similar in these games in industrialized societies-- whether they're relatively rich or not, people seem to behave very similarly. But perhaps if you look at hunter-gatherer society that are quite different in various ways, perhaps there we can look at interesting differences across places.

And so on the one hand, we see there's lots of uniformity of play. In lots of places, people seem to behave very similarly. But there are some important differences. And perhaps what this line of research is doing-- it's often called cultural economics-- is to try to see what is the role of culture. And there's Nathan Nunn and other people at Harvard are doing this in particular, inspired by Joe Henrich, who is the leader of that, is trying to look at, like, what is the role of culture in shaping people's preferences?

And so one thing they had done-- and at the time, they looked at 15 different hunter-gatherer societies. You can see them here on the map. Some of them are in South America. Some are in South America. Some are in Africa. And some are in Asia, Oceania. And so you could look at these different societies. And by looking at the way they are living and producing, you can try to think about potential predictions of how they might behave and social preferences behavior.

And there's two examples here that we have. One is the society that you can see here, the Machiguenga, which are, I guess, in-- which country is this?

AUDIENCE: Peru.

PROFESSOR: In Peru?

AUDIENCE: Peru.

PROFESSOR: Who are in Peru in Latin America. You can see them. They essentially are independent families. They do a lot of cash cropping. Their way of living appears to be slash-and-burn agriculture, gathering food, fishing, and hunting.

So it's overall-- and you can correct me if that's wrong. It seems to be fairly individualistic in the sense of the way people produce. Essentially, they live individualistically. They are not really relying on a lot of cooperation in the way they produce their food or their living.

In contrast, the Lamaleras in Indonesia, these are whale hunters. Essentially, when you try to hunt a whale-- I don't quite know. I imagine, at least-- it requires lots of cooperation in the sense of, like, people have to work together by nature of essentially finding food or hunting these whales. And so they on a daily basis are really relying on cooperation.

And so now this line of research by Henrich and others is trying to see, by looking at people's way of living and way of producing on any given life, is that predictive of people's behavior in ultimatum games? And what you see is the following. I'll just show you what results are.

So the society in Peru, the slash-and-burn horticulturalists, seem to have little concerns for others outside of the family or for social status. What you see essentially is relatively low offers. Most of these offers are being accepted. There's essentially very low rejection rates. And that seems to be essentially-- and the rejections of amounts that are below 20% is only, like, 1 out of 10.

So it seems to be the case, this society that looks a lot more individualistic is, in fact, much closer to perhaps the assumption, the neoclassical assumptions of no social preferences mattering in the sense of they seem to be-- in the way, at least, they play the ultimatum game, they seem to be happy to offer low amounts, and they seem to also be not upset or annoyed or getting back to others who offer them low amounts.

In contrast, if you look at the Lamaleras, the whale hunting culture based on strong and meticulously detailed cooperation, what you see is the mean offer is 57%, or the mode is 50%. There are quite a few rejections, in particular, rejections of small amounts. And that's to say, like, there's really an expectation of people cooperating. And if you don't cooperate, if you are not playing by those kinds of rules, people are being punished accordingly.

And there are some others. There's sort of another culture of the Au/Gnau, which is a culture where gift giving is an avenue to status and gives the right to ask for reciprocity at any time. So that's the thing that you often see in India as well, that people tend to reject offers that are very large in part because they think that's inappropriate. They think that's just a weird thing to do in a sense of, like, I offer you, like, 50% or 60%, they just think, like, if I accept this offer, now I owe you something and have to give you stuff back. You are nodding.

AUDIENCE: Oh, yes.

PROFESSOR: Yes. So that's in India as well to some degree as far as-- yeah. Any comments on this? Yeah.

AUDIENCE: Is that last experiment anonymous? Because if it's anonymous, then there's no sense of this specific person [INAUDIBLE].

PROFESSOR: That's a great question. I think it may not have been. I need to check. I think-- so there's some question on how closely in some sense to the-- these are 15 societies. So in some sense, these are not strictly speaking statistical tests of, like, does culture cause this behavior? Because you only have essentially 15 observations of certain ways people behave.

So in some sense, the mapping between the description of what I was telling you, which surely was inaccurate in terms of how people really live in those societies, and their behavior is more a descriptive mapping. So in some ways, it could be that even if it were anonymous, people behave in certain ways that reflects a culture of cooperation is to say, it's just inappropriate-- or reflect, essentially, when they say, if somebody gives you a large amount or a large gift, you're not supposed to accept it. You might sort of not accept it even in an anonymous game because you feel bad about it afterwards because you can't reciprocate, if that makes sense.

But I have to look it up, what the specific details were. The paper is also-- and it's a very short paper. So it's also on the course website if you want to look. Any other comments? Are you from Peru? Do you know about the society? Is that remotely accurate what I was saying?

AUDIENCE: I'm not really sure.

PROFESSOR: All right, well, you can look into this and report back. OK. But anyway, it's a very interesting application of the ultimatum game because it essentially tries to say, we're trying to look at how do people relate in the real world. And looking at these different societies gives you some sense of, like, it does seem to capture something in the real world of how people are, in fact, behaving. OK.

So now we can sort think about what's going on in ultimatum games. And we talked about this a little bit, but I want to be a little bit more precise. It's like, why do responders reject low offers? What's going on here? What are people doing? Yes?

AUDIENCE: [INAUDIBLE] spite.

PROFESSOR: Yes, and so the spite is-- so what exactly are you objecting to? I asked this before, but let me ask again.

AUDIENCE: So I think a lot of things have already been said. If you will reject, you think it's unfair, so much so that you would rather that they get nothing, you both get nothing, than [INAUDIBLE] take this pity offer.

PROFESSOR: Right. So there's two things. Let me put them out here. There's two things here. You're talking about the second thing, I think, which is the fairness part. So one is about a procedural thing, which is to say, I object to procedure. I object to you not treating me nicely or fairly. I'm just mad at you for doing that. Therefore, I reject.

Hypothesis number one-- this is what we were talking about earlier, which is to say, I just dislike being behind. You offer me-- you keep 6. I have 4. I just don't like having less than somebody else, and I might reject that as well. Can we disentangle those explanations? From dictator game data?

Yes?

AUDIENCE: You can have some [INAUDIBLE]. I guess, reject [INAUDIBLE].

PROFESSOR: Right. So from the data that I just showed, or from the pure, simple dictator game, it's very hard to disentangle these because it looks the same. But what you're saying now is we can have different versions of that where you could have either the computer implement it, or you could choose for other people, and so on. I think that's what you were saying.

And then from that, you could say, well, if you're also rejecting that, then perhaps it's not about the procedure per se. If the computer chose 6 and 4 and you reject it anyway, then it's not to do with fairness, because the computer was perhaps just randomizing. But it's rather to do with, like, I don't like to have less than somebody else in the same game. Right?

AUDIENCE: Yes. So if you used a computer, you would be able to [INAUDIBLE] just disliking [INAUDIBLE] less. If you have an individual respond to [INAUDIBLE] accept or reject an offer that somebody else made, it's quite [INAUDIBLE]--

PROFESSOR: Right.

AUDIENCE: Then you would see whether they were just trying to punish.

PROFESSOR: I see. Yes, I see. I see. So yes, you can do that. There's a bit of an issue, a little bit like-- another version of that would be just I don't like inequality, but I just might want to-- I just really like 5-5 or over 50-50 outcomes. There, I think we can look at what you were saying earlier, which is to say I could see if there's a third party that looks at somebody choosing 60 for themselves and 40 for somebody else, versus somebody choosing 40 for themselves and 60 for somebody else.

If you really dislike inequality, you would just reject both of those things. And if you sort of instead think this is like-- you dislike inequality. If the person only does that for themselves in some selfish way, you would only reject the 60-40, but not the 40-60. Great.

Right. So the answer is no, we cannot distinguish these motives based on the ultimatum game. So we need some additional evidence to do that. We're going to do that mostly next time. We're going to try to make some progress towards that.

Now, if you now think about the proposer's motive-- if a proposer gives a large or a small amount, what can we learn from that? Or why is it difficult to figure out what the proposal is doing? If you see somebody gives a large amount of money, what have you learned from that? Yeah?

AUDIENCE: It could be that they're just being optimistic, or it could be that they have beliefs about what the responder would reject versus what the responder would not reject. And they could reject [INAUDIBLE].

PROFESSOR: Exactly. So it could be that I'm just trying to be nice. I'm trying to be friendly. I care about you. I offer you a 50-50 because I want you to have half. Or it could be that I just actually don't care about you at all. I just think if I offer you only 30 or 20 or whatever, there's a good chance that you might reject it. Therefore, I just optimize over my-- so it depends on my beliefs about what the rejection rates are.

And if the beliefs are that below 30% or 40% people's rejection rates are positive and sort of higher, I might essentially just maximize depending on my risk preferences or whatever. Suppose I'm risk-neutral. I might just maximize expected payout. It might be that essentially 40% or 50% is just the optimal solution to that maximization problem, and I just don't care about you whatsoever.

So if I'm looking at this behavior essentially, if I'm looking at an ultimatum game and looking at a proposer offering large amounts or positive amounts, you can learn nothing about their-- or it's very hard to learn anything about their social preferences, because essentially there's strategic motives potentially involved as well. Is that clear? Or any questions?

So what do we need to do then? Well, now we need to sort of have some other evidence. So we talked about this already. When you think about social preferences, you can think about three broad categories of preferences. And we're going to talk about this a lot more.

So one is distributional preferences. This is to say that you care about outcomes. You care about how much each person gets. So that's to say we can sort of represent this to, like, I put some weight on my utility, on my outcomes. I put some weight on your outcomes. And then depending on what that weight is, I'm going to decide about how to distribute outcomes. These are sort of distributional preferences.

Then there are things like what we want to call face-saving concerns, which is people don't want to look bad in front of others. You can also think about this like social image, which is to say, I want to look like I'm a nice guy. So in particular, if I'm in public making certain choices, if I'm being observed, I'm going to behave in a very friendly way. If I could secretly be mean, I would do that. But as long as my social image is at stake, I might look like a nice person. We're going to talk about some experimental evidence looking at that.

And then there's some which we want call intentions-based preferences. These are things like reciprocity, procedural justice, and other facets of social preferences. Broadly speaking, you can think of this as fairness, which is to say people don't care about the outcomes, per se. They care about essentially the process in which the outcomes were generated. So they might be perfectly happy with one person getting a lot and another person getting nothing as long as the computer randomized and [INAUDIBLE], the probability of getting money was the same for everybody.

So that would be somebody who said-- these would be intention-based preferences, where people say, this is procedural justice. This is a fair thing to do. And people are motivated by that, or essentially, they might reciprocate. In cases when they think things are unfair, they might not do so in case they think things are fairly generated, even if things are unequal ex-post.

OK. So now we're going to talk about all three of those. So let's start with distributional preferences. This is, again, the simplest kinds of social preferences. This is very much a natural extension of how economists think about the world. This is going back to Gary Becker, I think, in the '60s, where essentially, instead of saying, does a utility function that just has my own outcomes, as my own consumption as an argument, instead, have a utility function that has my own utility, how much I get, a function of that, plus how much does the other person get, and a weight on that other person.

And so again, we sort of look at, like, the person cares not about only how much they get, but also the other. There's two versions of that, and this is what you were saying earlier. There are interested distributional preferences and disinterested distributional preferences. The interested one is essentially-- the version is, like, where I'm involved, and I get a certain outcome, and there's another person involved, and I put some weight on that other person. So it's about me versus another person, and how these resources are being divided. Could be also multiple people, of course.

And then there's a disinterested part, which is to say, how do I want resources to be divided across others? This is questions like, how do we want society to look? How do people feel about inequality and the like? Or, like what public goods-- how much should the government engage in redistribution and so on?

These would be often-- to some degree, of course, it's interested because people have to pay taxes or receive benefits from the government. But often, it's you can ask questions people have opinions about. Suppose you live in the US. You could have opinions about how should Europe's tax system look like. What is fair and what is unfair? Those would be, essentially, disinterested distributional preferences? Any questions on that?

So we're going to start with this and see how far can you get distributional preferences. What kinds of things can be explained? And then we're going to move towards saving face or social image type of concerns.

So the dictator game, I already explained to you what it looks like. So here's what the dictator game looked like in class with no stakes. This is very typical to how people behave in the real world. What do we see here? What are patterns of behaviors? This is no stakes, and it's just hypothetical questions here.

Yeah?

AUDIENCE: There are two large groups. One group that doesn't give anything, and a group that gives half.

PROFESSOR: Exactly. And this is very typical, again, in many situations. There's a group of people, often, like, 30%, 40% of people who just essentially keep the money. There's a group of people that thinks 50-50 is the right thing to do. That's often also something, like, 30% of people.

And then there's people often in between, often between 0 and 50. And there are some people who tend to give a lot. That's often a minority of people. Now, what can we infer from those choices? What can we infer from people who choose 0? Yes?

AUDIENCE: They give themselves [INAUDIBLE].

PROFESSOR: And do we have any objections to that? Or what other interpretation could we have?

Yes?

AUDIENCE: In the case of money, we can't necessarily assume they're self-interested, I guess, because they could be giving it to charity.

PROFESSOR: Yeah, exactly. So in the case of money, you could keep the money. You could say, I'm choosing the \$10. And instead, I'm giving it to somebody who is in higher need, who has higher marginal value of this money. That could be you could give it to charity or give it to somebody on the street or any other person you think is in need. So that person who looks like they're fairly selfish might, in fact, be quite generous in the sense of giving the money to others.

That's a common issue with a lot of these games, is essentially it's often an outside option in a sense of, like, we can look at these games as if this is the only game. That's the only thing that's happening in the world, and nothing else, and look at-- try to explain behavior through that. But often the issue is there's other things that happen in the world.

So when I think about, am I behind, or do I get more money versus somebody else, when you play with somebody who is a lot richer or a lot poorer than you are, in fact, giving them \$10 versus not or \$5 or not or will not change that fact that you're a lot richer than they are. So perhaps you should give a lot more, a lot less money, and so on.

So in anything that we're assuming, we look at this through the lens of, like, there's only this game that happens, and nothing else. It turns out that's actually a pretty good approximation. People sort of, what we call, are narrowly framed. They look at essentially this game as if this is completely separate from anything else in the world. So people tend to not integrate these games with the rest of life or the rest of utility.

So expected utilities would say, you should integrate this game with everything else that happens in life. So you should not look at, like, 6 versus 4. I'm earning more than somebody else. You should look at, like, what is my income overall, a lifetime income? What is your lifetime income or the expected lifetime income of somebody else in the experiment? And depending on that, you should give money or not.

Or differently, you might say, well, should I give money to somebody else in class? Well, perhaps not because most people in class will eventually make a lot of money once they graduate. And instead, let's just use that money and give it to somebody else who has higher need of that. OK.

So here's no stakes. It turns out when you look at stakes, not much happens. So this is answers-- if you look at-- this is no stakes. Here's stakes. Answers look, in fact, very similar. There's not much of a difference.

Even when you have higher stakes-- I think this is, like, \$20, and the like-- people seem to behave very similarly. So it seems to be, in fact, asking hypothetical questions-- and this is how some of these games started, just asking hypothetically, what would you do if somebody gave you \$10 and so on? Turns out, you actually get pretty good answers for how people behave in reality.

That's true for this class, but that's also true in actual experiments. In actual experiments, if you just ask people, like, what would you do? And then when you actually implement it with the same people or with other people, you get actually pretty similar patterns across people.

Now, we do find that-- so we had a version of that that was essentially in private, which is to say you did not have to reveal if selected to the entire class, essentially, what you chose, but only to the other person. Notice that that's not really private in the sense of you still have to tell the other person, ha, I chose 10, and you get 0, so that's not very nice. There's other versions that would be in private, would just be the other person might never find out, or the other person might at least never find out that it was you who chose 10, and they got 0.

And so but what we do see, if you look at-- I think this is with monetary stakes. This is with \$10. This is with \$10 in private. And you do see, for example, that when you look at the fraction who gives essentially 100%, that fraction tends to go down. So it looks like people are somewhat less nice in private.

We're going to get back to that again in the sense of trying to see-- when we look at the question of, are people really nice when they give money, and so on. We're going to look at cleaner versions of that, where people have the ability to essentially hide in some ways that they're not nice. How might one do that? How could you-- what kind of game would you design to hide the fact that the person is not nice? Yes?

AUDIENCE: For example, if the recipient is only given the money and not told what percentage of the original money [INAUDIBLE].

PROFESSOR: Right. So one version of that would be the recipient just doesn't even know what the game is. They just get some money. Here's some money for whatever reason. You were just given it because you're in the experiment. And so you could just essentially hide the entire game from the recipient.

What's another thing you could do where you don't hide the entire game but sort of say, still, there's a game? But what other ways could you-- in what other ways could you hide the fact that people are not nice? What mechanisms could you implement? Yeah?

AUDIENCE: If you mix robots and [INAUDIBLE] people.

PROFESSOR: Exactly. And that's exactly-- I'm going to show you that, I think, on Wednesday. That's exactly what people have done, where are you say, with 50% chance there's a robot who is not very nice or very nice or whatever, and some chance the robot is really mean, with some chance the robot is really nice. And now we can look at-- so with 50%, the robot decides. With 50% chance, you decide. And the recipient will never know whether it was a robot versus you.

And so now you can essentially hide behind the robot. You can say, well, I really wanted to give you a lot. But sorry, the robot was really mean. Too bad. And the other person will never find out. And it turns out when you do that kind of game, people tend to be a lot less nice than when there's no robot.

OK. So then we have chocolate. Looks like people are giving somewhat more when it comes to chocolate. They tend to give a lot more when it comes to apples. So now, why is that? Why does that matter? Yes?

AUDIENCE: If you model algorithms, your utility plus the other person's utility [INAUDIBLE], it could be that people with marginal utility, in chocolate apples decline more than [INAUDIBLE].

PROFESSOR: Yes. It turns out getting the fifth or sixth or seventh apple of the day, maybe not as enjoyable as the first or second. So exactly. So well, it's really reasonable to think that the marginal utility of money is linear for small amounts. If I give you \$1, 2, or 3, or 4, or 5, in some sense, that shouldn't be very concave. We discussed that at length.

For apples, they're perishable. There's only so many apples you can eat on a given day. So really what we should look at is, like, what's the utility of apples? And if the marginal utility of apples is decreasing, that might matter quite a bit.

OK. So maybe let me summarize a little bit. What did we learn in those kinds of games? So first is people look fairly generous in dictator games, even when the game is largely private. We're going to challenge that assumption a little bit more, exactly as you suggested, with other ways in which either people can be more private, they can just hide the game entirely, or when they perhaps can hide behind a machine or a computer or in other ways, or when they perhaps can opt out of the game in certain ways.

People generally seem to give something like 30% of the total. That's quite common. 20% to 30% is quite common overall. It doesn't seem to matter that much whether there's hypothetical versus actual choices. The size of the stakes also seem not to matter that much. We seem to see somewhat different behaviors of chocolates and apples.

There's also some people who told me that they're lactose intolerant and so on. So essentially saying, like, their marginal utility of chocolate is quite low. So it's much easier-- it's much less costly to give to somebody else. Or put differently, it's actually efficient potentially because the marginal utility of chocolate is much higher for that other person than it is for you. Of course, there's a problem now, again, that you could just choose the chocolates for yourself and then just give it to your friends. So you look quite selfish, but in fact, you might not be.

So anyway, but sort of what really matters here-- and this is perhaps-- in some sense, when we talk about apples and chocolates, it's a little bit silly. But when you think about a dictator game where you play with somebody in Kenya whose income or lifetime income is, like, orders of magnitude lower than yours, then it matters, actually, quite a bit what's the marginal utility of the \$10 that you could keep for yourself versus for that person. For that person, that's a lot of money. For you, a lot less. And so you might be more generous when playing with a person whose marginal utility of money or consumption is a lot higher.

This is what I was mentioning as well. You want to think about, what is the outside option? So some people might give low amount and decide to give the money or the apples or whatever to somebody else. Those people will look selfish in dictator games, while, in fact, they're quite generous.

That's a problem with these games that are hard to deal with. We're not going to be able to deal with that overall. You can think about it a little bit and sort of see. Can we give people options to opt out and so on? But it will always be an issue.

It's less of an issue than you perhaps think about it. Conceptually, that's a huge problem. But when we look at people's actual behavior, it looks like people really behave as if this was the only thing in the world. And then, in some ways, it's quite predictive of actual behavior in the world, even perhaps, in some ways, it shouldn't be. Any other comments or things we learned perhaps from the game or that you observed? Yeah?

AUDIENCE: Is there a possibility that the people that give 0, it's just that they're interested distribution [INAUDIBLE] is the same as your disinterested one, and they actually think, like, oh, the best thing for society is for me to have those \$10 and not the other person?

PROFESSOR: So because you put a lot of weight on yourself or?

AUDIENCE: Sort of like, regardless if it's correct or not, but you think that the way you'll use these \$10 is better for society than the other person?

PROFESSOR: Yeah, I think you have to justify that somehow in the sense of, like, if you are very similar to the other students, then you might sort of think that dividing it-- so depends what do you think is the marginal utility of money. So if you think the marginal utility of money is essentially the same for everybody, in some sense then, it doesn't matter very much whether you keep it to yourself versus somebody else, in some sense because, in the grand scheme of things, it doesn't look very different.

It could be that, for example, you think, for whatever reason, you just really had a bad day, and you can use the \$10 to buy something nice. It could be that you think you're relatively poor to others in class. For whatever reason, you just don't have a lot of money right now. The marginal utility of money for you might be really high. So again, and in some sense, that's socially efficient. You would actually-- even if it weren't you, you would give the money to yourself.

Now, I think it is the case that when you look at these distributions, it tends to be that 30% of people keep all of that money. So it tends to be that it's hard to believe that that's really-- I mean, you'd have to look at, like, how does that correlate with other characteristics? It probably does, and we'll get back to that in the sense of when we see now, when people have the chance to hide money and so on, when you see now the fraction of people who choose \$10 for themselves and 0 for others goes up, that, in some ways, compared to when you can't hide, that perhaps is an identifying selfishness and so on.

Here, when you see people keeping 10, you can come up with lots of explanations, including the one that you mentioned, but also including the ones you keep it for yourself and then give it to your friends and share it with them and so on. And we can't really infer very much. But in these cleaner experiments, where you can really say, OK, when you don't have the outside option, when you don't have the chance to hide, you look quite nice. That goes away when you have the chance to hide. Presumably, that is because you just are selfish and not for other reasons. OK. Yes?

AUDIENCE: Have people looked at if things change if you actually get the money or get the apples beforehand? Like, [INAUDIBLE]?

PROFESSOR: Surely they-- surely they have. What I don't know-- so surely [INAUDIBLE]. So [INAUDIBLE] is very robust in various situations. So my guess is they were less likely to give. I don't have-- I can look it up, but I don't have like a great experiment that I can tell you about.

But my sense is that-- so the question is, if you give people the \$10 into their hand, are they going to choose differently? My guess is they will look a lot less selfish. In the experiment that you played, in some sense you are sort of endowed with a game. So in some ways-- but you could do deviation of that. You could look at, like, here's \$10. It's neither yours nor the other person's. How would you like to divide it?

My guess is that person will be more generous when they do that compared to saying, like, here's \$10. They give it to you in \$1 bills, and now you have to give me back some. People will be more likely to keep the money when they actually have the money in hand, when they feel like it's yours or theirs, as opposed to have to divide money that doesn't really belong to any person.

So my guess is that's true, but I don't-- there's literally thousands of experiments of people doing these kinds of variations. So I'm sure there is somebody who has done this. I can look up and see whether I find that. OK. Any other thoughts of what we learned?

OK. So when you now think about how should you model these distributional preferences-- and there's a classic paper by Charness and Rabin. Rabin is the guy that you know from the calibration theorem about risk preferences. And the way their model is very simple.

They essentially have the preferences over outcomes x_1 and x_2 . Think of this as money. It could be also apples or something else. Money is perhaps a better approximation because the marginal utility of money-- think of this as rather constant.

And now player 2 is the dictator. So x_2 is how much player 2 gets. So player 2 has a utility function u_2 . The utility of player 2 has a function of x_1 . This is how much the other person gets, and x_2 , how much the player 2 gets.

And the utility function looks like ρ times x_1 plus $1 - \rho$ times x_2 . If x_2 is larger or equal than x_1 -- that is to say that the person is ahead, if x_2 gets more than x_1 -- and it's σ times x_1 plus $1 - \sigma$ times x_2 if x_2 is smaller or equal than x_1 . So what do these parameters measure? What a ρ and σ measure here?

Why might they be different? Yeah?

AUDIENCE: How nice they are?

PROFESSOR: Yes, exactly. What does it do? How much weight do you put on yourself or the other person?

AUDIENCE: So if it's a lower value, then you don't really care about the other person, so you put more weight on yourself.

PROFESSOR: Right. So in the extreme case-- I guess, in one case, I guess, if you choose-- if this is 0, then essentially, you're back in the neoclassical case of not caring about others at all. So if you choose, like, $\rho = \sigma = 0$, then it's just, like, how much do you get? You should get 0, and you should accept any offer, right?

And now as ρ and σ goes up, 0.5 or whatever, you give more weight on the other person. So 0.5, for example, is the case where you care equally about the other person and yourself, right? Because the weights are then the same. Now, why might ρ and σ be different? Yes?

AUDIENCE: So probably ρ is larger than σ .

PROFESSOR: And why is that?

AUDIENCE: So if you're person 2 and you already know you're getting more than the other person, then you don't feel such a need to do so much-- you don't feel a need to have so much more than them.

PROFESSOR: Yeah, exactly. It's just easier. That seems to be a very intuitive feature of the world. It seems to be much easier to be generous when you're ahead anyway. If you're very rich, you might as well give some money to the poorer person, and so on.

If you are poorer than the other person, it's much harder to be generous in some ways, perhaps because you feel like you deserve to get as much as the other person, so why would you give the other person more? You put essentially less weight on that other person if you're behind. What's the case of sigma smaller than 0?

Yeah?

AUDIENCE: Would it be then the person who want to hurt the other person, so then they put more weight on decreasing their value?

PROFESSOR: Yeah, exactly. So if you're behind-- so it's the case where you are behind, you're the x_2 , the player 2. If player 2 is behind, has less than player 1. Now, you might get positive utility from reducing player 1's outcome. So if you're behind, you'd rather have the other person get less, and you'll feel better about that, presumably because you're then less behind, even holding constant your own outcomes.

So if you keep your own outcome, your own payment, the same-- suppose you get 5. The other person gets 20-- you feel better if the other person gets 15 or 10 because you're less behind now than the other person. OK, so you're willing to hurt the other person. You might be even willing to pay to hurt the other person, which is kind of what happens in the dictator game when the person rejects the dictator game. Any questions on that?

So I have an example of sigma smaller than 1. Suppose you have sigma being minus $1/3$. You have to choose between 0, 0 and 9, 1. What are you going to choose? So this is what this looks like, sigma times x_1 plus 1 minus sigma x_2 .

Yes?

AUDIENCE: You choose 0, 0.

PROFESSOR: And why is that?

AUDIENCE: Because it multiplies the other numbers.

PROFESSOR: Yeah.

AUDIENCE: [INAUDIBLE].

PROFESSOR: Right, so essentially-- the reason is that essentially, if you look at this situation, and you put minus $1/3$ weight, essentially, on the other person. So you really don't like the other person getting 9. So you're really unhappy about that. Despite the fact that you value getting some money yourself-- you have positive value on that as well. It's, I guess, $4/3$ times 1. But sort of putting together, that's negative.

So essentially, you're willing to pay-- and this is, I guess, one explanation potentially why people might reject. And this is, I guess, what I was saying earlier. You might reject the ultimatum game simply for distributional reasons. You might just say, I don't like being behind in this game. If I'm behind, I put negative weight on this other person, and then I'm now rejecting it because I feel happier if the person who is in front of me gets less, even if that comes at the cost of me having to pay some money.

Now, notice, you are not going to reject all unequal offers. So I think 6-4, you would probably accept. But you will reject really uneven offers. So in some sense, if I ask you questions using the strategy method, as we did-- if I asked you questions about at what threshold would you reject an offer, I can essentially-- if I assume that your preferences look like this, I can back out, essentially, your σ . Questions on that? Yeah?

AUDIENCE: What's the utility [INAUDIBLE] for the dictator, not the person who gets to reject the offer?

PROFESSOR: Sorry, this is just-- sorry, I should have been clearer. So this could be for-- these are preferences that are generic for any types of games that are being played. So you could apply this to any game. So Charness and Rabin, essentially what they do is-- in their paper, and it's a long paper that covers a lot of ground. But essentially, what they're saying is, like, here's what distributional preferences might look like. And you can then choose different parameters.

And then the question is, like, what kinds of behavior now can we explain? Being the dictator game, potentially you can try to explain sort of people's behavior. But not I'm trying to apply it to the ultimatum game. Or you could also apply it to the trust game and say, for what kinds of behavior, what kinds of preferences might people do certain behavior? So what kinds of behaviors can be explained with that?

So now what's the experimental evidence on ρ and σ ? Just to remind you, ρ is the parameter, the weight that you put on the other person if you're ahead. σ is the parameter on the other person when you're behind in a situation.

So we think that when people are ahead, most people seem to have a positive ρ . So people usually tend to-- when ahead, they're willing to sacrifice some money to increase the other person's payout. So this is essentially-- when you think about dictator games, when they have 10, essentially you think of this as being, like, 10-0, you're ahead. So you're willing to sacrifice at least some money, like 7-3 or 6-4, even 5-5. You like that better. You're happy to give the other person some amount.

A minority are even willing to sacrifice money to give the other person an equal amount of money or even less-- sorry, even more. When deciding to split \$10, subjects tend to give about 20% to 25% on average. In class, that was actually pretty close. The 28%, I think, are about the \$10 amount. You guys gave it something like 28%. That's actually pretty close to what people would do in normal dictator games. The estimated ρ tends to be about 0.4. Yeah?

AUDIENCE: How can ρ be estimated from the dictator game? The linearity of the preferences here, shouldn't it predict that player 2 will always either give 0 or give 50%?

PROFESSOR: No, it's just to say-- I think this is just to say you have the same marginal utility than I have. So depending on when I ask you, like, how much do you give, that pins down your ρ .

AUDIENCE: But if x_1 is 10 minus x_2 , and we want to maximize $\rho x_1 + 1 - \rho x_2$, and you're maximizing that over x_2 . [INAUDIBLE] equation, shouldn't it be [INAUDIBLE]?

PROFESSOR: No, no, but the unknown here is the ρ . It's essentially to say, for example, you choose 6-4, that means you prefer 6-4 over 7-3, and you prefer 6-4 over 5-5. And those two inequalities essentially give you a bound on ρ . We can talk afterwards, but I'm pretty sure that's true. Yeah. Yeah.

But yeah, that's essentially just saying, like, essentially what you do is you have-- since you chose a certain option, now essentially you'd prefer not to choose a higher option. You prefer not to choose a lower option. When you write that down, that should give you inequalities as a function of ρ in this case. But we'll talk about it in two minutes.

What's the experimental evidence of σ ? Only about 10% to 20% of players have σ strong enough to pay a non-trivial amount to hurt the other player. These are people essentially who reject offers. About 30% of people will sacrifice to help a player ahead of them. So even if they're behind and have the option of, like, can you sort of help the person in front of you, 30% of people will do that.

If a person is behind you, often neither wants to help or hurt the other person by much. Essentially, sort of like, whatever, I don't care about this other person. I don't feel inclined to do very much. But also, I'm not necessarily hurting them.

And now I want to just preview what we're going to talk about next time, which is to say it looks like people are pretty nice in various situations. It looks like in the dictator game, they give money. It looks like in the ultimatum game, they give money, and so on.

Looks like people give money to charity. This is what we discussed earlier about 2% of GDP. People also do a bunch of volunteering or other contributions. So they give some time, about 15 hours per month on average when you do surveys. So it looks like people seem fairly generous overall.

Now, when you think about social recognition or social image, perhaps that picture gets less rosy overall. It seems to be that essentially, people care a lot about what others think of them, and not just about the outcomes, per se. So one example would be gifts to organizations.

If you look at thresholds of-- like, when you look at the Boston Symphony Orchestra, and you look at the distribution of people giving, it tends to be there's these thresholds where you can give a certain amount, and then you recognize them. Like, when you go to the Boston Symphony, in this booklet, you can who's like a certain donor of which category. People tend to bunch just above the thresholds for these categories. And presumably, that's because they want to look good.

That's just one example. So I think the next question we're going to ask on Wednesday is to say, is social recognition a major motivation for giving? So people seem to not only care about what others get, but what others think of them when they give or not. So in some ways, that's a philosophical discussion in saying, like, are people really nice versus not? And you might say, it's sort of disappointing if they're not nice in certain situations compared to others.

The flip side of that is, well, if you know what kinds of situations generate altruistic or prosocial behavior, well, if you know that, then you can essentially also create situations to foster prosociality. So once we kind of know under which structure or in which kind of circumstances people are nice, well, that gives us some lever or some ways for policy and saying, like, well, you have to design your institutions or your organizations or your company in certain ways to get people to behave nice to each other.

In contrast, if you thought people are either generally nice or not, well, then in some sense, that's good to know. But then in some sense, we can't do very much about that. And that's what we're going to talk about next time. Thank you.