[SQUEAKING]

[RUSTLING]

[CLICKING]

JEAN TIROLE: OK. So today, we are going to cover a topic which is quite unusual for 14.271. Nonetheless, it's quite important. It's about the views of people about-- that people have about the morality of markets. [14.]271 is about antitrust mainly. And of course, antitrust is about market. And if people don't believe in markets, then that raises issues about what we are doing in the first place.

Furthermore, there are lots of development nowadays-- actually, it's broader than antitrust-- about what independent agencies should be doing. So that has started with central banks. Their mission of enlarge is no longer just inflation targeting and financial stability, but it goes way beyond. And now everybody talks about green central banks. And there is a issue about whether fighting the climate change should be the responsibility of the central bank or it should be the responsibility of ministries.

With antitrust, it's a little bit the same because there are more and more discussions-- in particular, in Europe, but all over the world-- about whether the consumer standard is the right way to proceed. So the consumer standard is basically that the antitrust mission is to defend the consumer.

Maybe not the short-term consumer. Maybe an intertemporal consumer because sometimes you have tradeoffs between today and tomorrow. But by and large, the mission of antitrust authorities everywhere in the world is to defend the consumer. It's a pretty clear mission. It's difficult in practice, but on paper, it's simple. You just stand as an advocate for the consumer.

But now people say, no, no, no, maybe we should relax antitrust first, for example, when we choose cases. So if there is a caviar cartel and a milk cartel, then if you have unlimited resources, you first go after the milk cartel. And the caviar cartel, you just say, OK, this is not that important.

So we should take into account redistribution. We should take into account climate change. So if we have a cartel which actually promotes the fight against climate change, maybe we should tolerate the cartel and so on and so forth.

There are lots, lots of discussions about that. Actually, next week, I'm going to a conference organizing by Margrethe Vestager. I'm giving a keynote there. And the panel just after me is exactly about this topic, with people like Thomas Philippon and others. It's really about this topic.

So I will know more in a week from now. It's not going to help very much for today. But that's really important because, in the end, that's the mission. In my own view, I've worked on mission and advocacy with Mathias Dewatripont. And my own view is that it's not always a good idea to have mission creep.

And if the issues are real-- so inequality is real, climate change is real, of course-- maybe it's not, say, to the antitrust authority to tackle them, but actually more political bodies should be doing that as opposed to-- but that's my own view on those issues.

But there is a broader thing about whether we like markets. And there is this bizarre thing that almost every economy in the world is a market economy, and nonetheless, nobody likes markets. So it's kind of paradox somehow. And I will come back to the big picture about how we think about the organization of society. But it's kind of strange.

So what I'm going to do is to have a brief introduction about this and then switch to a paper that I'm finishing-- the revision I'm finishing with Mathias Dewatripont on the morality of markets. But I just want to have a broad picture first.

I will go relatively fast because it's developed, in some detail, in*Economics for the Common Good,* if you have seen *Economics for the Common Good,* the book. But feel free to ask questions, actually. Try to ask more questions than you did on Monday. That's your mission.

There are two kinds of attacks against markets. The first is that a number of markets are what's called "repugnant." So we don't feel very much at ease with those markets. And almost every well-known philosopher has a book on that.

I put here Debra Satz and Michael Sandel, but there are many other well-known philosophers who also have books on-- with roughly the same title, not the same content. Because if you compare those two books, for example, they start from the same point of view, but they have very different analyses.

The most well-known one-- I mean, this was a huge bestseller of Michael Sandel-- talks about all those markets that make us uneasy. Do we have markets for babies or whatever? All those things are actually creating some conflicts in our brain about whether we should have those markets.

Now, this is more specific. It's not the market in general. It's specific markets that we don't quite like. But the frontier is unclear. So a number of people think, for example, that tradable emission permits market, the ETS systems for carbon, are immoral. They are immoral because you can pay for the right to pollute. Or carbon tax is the same. You pay the carbon tax, and you pay to pollute.

Of course, you can ask, is that more moral to not have to pay to pollute? But then there is a big silence in front. But lots of people, in my experience in the last 30 years, is that when you talk about a carbon price, it always sounds bad. It's immoral. Talk to a minister. Yeah, you're right, but it's immoral. And at that point of time, you say, OK, you prefer people polluting for free? Is that more moral? But I don't know.

And I think-- we are economists. And we have to realize that, in civil society, our ideas are kind of strange-- views are strange to some extent. And the ideas are really about market incentive. It's true that much of the work we have done in the last 30 or 40 years are about, basically, when those markets don't work and when the incentives don't work.

So there is a huge number of papers on why low-powered incentive schemes might be desirable because you have noisy performance measurement; team work, which is a special case-- I think about banks' paper-- collusion with monitors might work with-- might work on capture; repeated interactions; multitasking; asymmetric information; and so on and so forth.

So there is a huge number of pieces of theory which basically says, low-powered incentive scheme can be good sometimes, in specific circumstances. And same thing. There are a number of papers about the fact that incentives can crowd out intrinsic motivation.

So the two papers with Bénabou are along those lines. The first line is basically that, if I give you an incentive, I may reveal something about your-- the trust they have in you or about the task. And that may actually backfire because it conveys information to you that discourages you.

And the second one, which is more familiar from psychology, is the idea of-- that if I pay you to do something, like giving your blood, then people don't know whether you're giving your blood because you are generous or because you are greedy and you want the money. So that kind of stuff. So those are the two papers I had with Roland Bénabou. But I could go on and on.

So much of our work is actually about cases in which markets don't work, in which incentives don't work. But nonetheless, it is the case that we are still more pro-market and pro-incentive than the average person in the population. And the same thing for the empirical work. There's a lot of empirical work showing that actually incentives may be needed. And that's something that is not very well received.

You can ask why. Why is that the case? I mean, part of the thing is that we are the bearer of bad news. Because when we say we need incentives, we implicitly say people are not as nice as we would like them to be. That's exactly the point, which is, we would like people to be nice, but they are not that nice. They keep on polluting, et cetera, et cetera. And that triggers a reaction.

In France, every time you talk about the carbon price, people say, the environment should not be-- protecting the environment should not be punitive. It's always the same thing. You don't trust people. I mean, sure. I mean, I don't trust people. I mean, for over the last 30 years, nothing has happened in the fight against climate change. Come on.

But still, they still say, oh, no, no. People will do what it takes. And come on. I mean, the answer is, no, it's not going to happen. The planet is going to burn. So I don't know. I mean, it's just like--

But so, still, despite all the caveats we have been working on for decades, it's still the case that we are pro-market and pro-incentive on the whole, compared to the population.

And then the psychologists-- I already talked about those ones. But the legal scholars actually have another line of attack, which is interesting. It's about the expressive content of the law.

So basically, the idea is the law is not just an incentive scheme. I mean, we tend to see the law as an incentive scheme. You have to pay a fine or you go to jail. I mean, that's going to incentivize you to behave well.

Legal scholars, they go beyond that. They say, no, but the law also sends a signal to society about the norm in society, about the views, what is acceptable or not. So it goes beyond that. And then, philosophers I mentioned [INAUDIBLE]. So they are unhappy with markets.

Now, I'm not going to go into that. But in *Economics for the Common Good,* I spent a fair amount of time trying to explain that a number of repugnant markets are actually markets with externalities, with internalities, with whatever, asymmetric information. They can be explained in the framework of economics. There is no contradiction between economics and the repugnant market.

I mean, just take babies for adoption. I mean, we will not want the babies for adoption to be sold to the highest bidder. Well, I mean, you might say, well, after all, if someone is willing to pay a lot for the baby, then maybe that person is going to benefit more for the baby.

But there is a third party who is not part of the contract there. It's the baby. You have a huge externality on the baby. You have no guarantee that actually the highest bidder is going to be the best father or mother. But it goes beyond that.

Also, something which is interesting, which is image externalities. So that I discuss. Sometimes, you could have two parties who are actually part of a contract, but their contract actually may shed bad light on the rest of the population. And image externalities is the thing.

Then there are all kinds of internalities, but internalities like present bias, for example, is just an externality on yourself, on your future self. So it's the same thing. So we don't like people who sell drugs or opioids or things like that. And of course, we are against involuntary slavery because that's a huge externality.

But we are even against voluntary slavery. And you might say, why not voluntary slavery? If everybody agrees, why not, to a contract? But of course, there are present biases which could be growing those markets.

Asymmetric information. And [? Sander, ?] for example, says, how? Those economists will say, OK, just pay to get friends. I mean, but if you pay to get a friend, then how do you know this person is your friend? I mean, it's called asymmetric information. It makes no sense. So there is at least a number of those things that can be. It's just a consequence of standard economics.

But the thing I want to insist on is moral postures. I don't like moral postures. I mean, that's my own bias, which is they can be useful as a warning signal. So I see something, I get some emotion, and that might suggest that something might be wrong.

So it's a warning signal. It actually can be useful. But I don't like people who just stay with the emotion and say, I don't like it and this is immoral, without explaining why. And that's something-- because if you say, in the old time, for example, life insurance was immoral. Everybody thought life insurance was immoral.

Now, why? I mean, so they would say, oh, yeah, it has to do with life. Sure. Come on. Same thing in many countries today and in most countries yesterday, having sex among people of the same sex or different races, for example, was thought of as "immoral" by a majority of people. And that was just a way of actually infringing on the freedom of people.

So those moral postures, I'm always very nervous. And when I hear philosophers telling me, oh, this is immoral, period, I just-- I just don't like it. Again, the warning signal is pretty fine. And our emotions actually can reveal real issues. But you have to go the next step and say, let's think a little bit more about this. Is there an externality, is someone hurt, and so on and so forth. So that's just my own view on this. OK. Questions? Yes.

AUDIENCE: How could we differ or think about the difference from someone having moral cost of taking some action or consuming some thing and then someone having blockage, where there's a cutoff after which they just stop even considering that-- considerations that it's a possible thing that we could do, for instance?

JEAN TIROLE: So there's this saying that the very act of thinking about something actually reveals already stuff. There is this movie. I don't remember the title. It was not such a good movie. But there was this movie in which a millionaire was offering-- there was a young couple, very good, united couple. And then the millionaire comes. But they are poor.

And the millionaire comes in and say, I pay $1 million to sleep with your wife. And then they start entertaining the thought. And I actually don't remember what happens in the end. [CHUCKLES] But they start entertaining the thought, and that destroys her. Even entertaining the thought destroys that couple. Not doing it, just thinking about it.

And same thing. If you tell a Catholic you have to think about whether Mary could be a virgin or whatever, just starting to question things may actually-- already the signal. Is that what you had in mind or--

AUDIENCE: Yeah.

JEAN TIROLE: Yeah. And sometimes, we economists actually are guilty of that in a sense. You may remember Larry Summers talking about taking-- polluting stuff to less-developed countries because they are willing to get less money in order to take care of the stuff. And that was a big scandal.

I don't remember if he was recommending it, which is a different thing. Just thinking about it was just already thought of as being very strange. And we often-- actually, when we do research, we often think about various other things which might be thought as not right. And I will say, it's part of the freedom of the researcher.

Now, when we go public, it's a different ball game. So I would try to distinguish between my life as a researcher, where I want to enjoy maximum freedom of thought, and then maybe when I go public, I may offend some people and so on. And that's a different thing. But, of course, there's always a thin line between the two.

Now, if you think about how our societies are organized-- and you know all of that, but I just want to give you the framework. I will say, it's Adam Smith plus Pigou. It's Adam Smith plus Pigou.

So the very bizarre thing is that, by and large, we have for-profit corporations. So for-profit corporations are entities in which the decision-making obeys the shareholders. The shareholders decide. Of course, it's not quite true. You have agency issues with the managers. I mean, we have worked on all those things. But basically, it's the shareholders who decide.

Now, we know why it's the case, because that's where the shareholders make sure they get their money back. Otherwise, they might not get their money back, or investors, more generally. But it still raises the issue that shareholders, of course, are interested in the decision. They want their money back. But what about workers and suppliers and customers and so on? There are many stakeholders who are actually left-- left aside in the decision prime, right?

So we have this situation where one of the stakeholders takes a decision. Those are the investors. And the others are basically left-- they are decision externalities. So that speaks to the corporate social responsibility.

Friedman had this famous column in which he said, the responsibility of business is actually to maximize profit, not to take into account the stakeholders' welfare. And his colleague at Chicago actually said, well, if you want to protect shareholders-- this is Ronald Coase-- the first one was Milton Friedman. Sorry. I realize your generation may never have seen a picture of Milton Friedman or Ronald Coase.

So Coase basically said, well, let's have contracts to protect shareholders. And to some extent, that's being done. So the contracts which insulate the most the workers or the debt holders and so on against the decisions of shareholders are-- say, for example, the worker is a fixed wage plus some severance pay if they are laid off.

Same thing if you look at debt, although they get often some fixed debt repayment plus some collateral pledging in case the firm goes bankrupt. And you could also have exit options and the like.

So basically the idea, OK, shareholders decide-- that's Ronald Coase in a sense-- shareholders decide, but you protect the stakeholders through contracting. The next step is, oh, yeah, but what about the quality of contracting, of course?

Imperfect contracts, collective action prime, and the environment and competition policy are a good example. You have millions of people who are affected by the decision of the firm, and, of course, it's very hard for them to contract with the firms and so on. And there are many other issues with Coase.

So in a sense, what we have converged to is some kind of mixture of-- you know which one was Smith and was Pigou, right-- to mix basically a market, the invisible hand of the market together with regulation.

So in a sense-- the other way you can protect stakeholders is to have a regulation. You have an environmental agency. You have a competition policy. You have a banking regulator. You have someone in government who is going to protect the stakeholders. So that sort of thing.

Now, in this world, where you have the market which is pretty efficient but has lots of market failures, then the government corrects a market failure. It's a simple world. We are "market fixers" in a sense.

Why should there be any corporate social responsibility or social responsibility as an investor? And the answer is there is none. If the government does their job, you don't have to do their job. The problem is that, in practice, there is a twin failure between-- so there is a market failure plus a government failure.

And the carbon case is a typical example of a market failure. The firms and the individuals can emit carbon. They have no incentive not to. And the government doesn't want to have a carbon tax, period. So you have a twin failure.

Now we can ask, why is that the government fails? And there are many reasons for that. So then, given that the market fails, the government fails, what is left? Civil society. And that's why we have a social responsibility.

I told you a little bit about the repugnant market, which is only a subset of markets which are attacked. But there is a broader condemnation of markets-- in a sense, all markets-- on the basis that the social value of activities may differ from private gain. That's something that economists worked on all the time.

In a sense, we know about externalities and the like. And something I'm going to discuss in the-- I discuss in the work with Mathias Dewatripont is the idea, which is more subtle, that markets can actually frame ethical choices. So the idea is that the managers confronted with this twin market plus government failure. The managers actually could substitute and say, OK, I'm going to behave morally.

Now, there's the issue about whether they will be allowed by shareholders to do that. But same thing-- the shareholders could behave morally. If they are socially responsible, they could behave morally.

But still, the market might be creating issues with moral behavior, especially if you are a consequentialist, which I hope you are. I don't know. I'm not going to impose values on you. That doesn't make sense. But consequentialist means you care about consequences of things. And actually, this is, I think, the most potent criticism of consequentialism.

So the issue about the replacement logic is basically, I don't like to pollute. But if I don't pollute, somebody else will. I don't like to sell weapons to a dictator. That's not nice. But what the hell? If I don't do it, somebody will sell weapons to that dictator, which is correct.

I don't like to bribe officials. It's not very moral. But if I don't bribe, I don't get the contract, and someone who bribes is going to get it. So in the end, the outcome is the same. So I might as well do it myself because I'd make money out of it.

The doctors who overprescribe opioids were saying that, saying, I don't like to overprescribe opioids. But if I don't do it, some other doctor will prescribe opioids.

And the same thing happens with sick leave. So in France, you go to any doctor. You want a sick leave for two weeks, the doctor signs. Even so, you're in perfect health. But the story is that, of course, if the doctor doesn't do that, then the next doctor will do it anyway.

And there are many other examples. I mean, athletes, for example. Many cyclists don't like to take those drugs. I mean, a football player and so on, they don't like to take those drugs. But they say, it's the only way I can actually compete and so on.

So you see the idea-- in all those examples, there is a common idea that competitive pressure is going to make it hard to actually-- to actually compete. So the "conclusion," in quotes, is that you should relax competition. And therefore, that has deep implication for [14.]271. So that's something we're going to discuss.

Now, I already mentioned that that might have implications for competition policies, which are being discussed right now, the promotion of open platform ecosystems, for competitive procurement, for trade, and so on. And that might also have-- there's also a discussion about whether you should actually either favor or just completely allocate a market to firms which are not for profit. So that kind of thing.

So should we, for example, make-- in France, there was a scandal a year ago about a retirement home. They charged a lot of money and provided a very low quality, at least some of them. Same thing has happened with hospitals and so on.

I mean, should they all be not for profit? Should universities be not for profit? Actually, most of them are. Well, they are not not-for-profit for everyone. Should they be subsidized and so on? So those are important issues. I mean, in the last decades, nobody was talking about that that much, but now it's resurfacing.

Just a last caveat before I go into more technical stuff. I'm not going to invent anything on morality. I mean, people have been discussing this for millennia. So I'm not going to say anything interesting on morality, but I just want to say, we economists have a view of morality which tends to be, in a sense, no externality or cutoff.

Basically, if I stab you, then you'll find it's not very nice. It's not quite moral. The avoidance of harm. And I would say that this kind of conception of morality is more stable across culture and across time, but it's restrictive. Many people have a much broader view of morality.

So if you look at Jonathan Haidt's books, for example, you get lots of examples of immoral behavior without victims in a sense. So one of his-- anything that violates a convention or a taboo or generates discuss-- one of his favorite examples, if you haven't read the book, is about your favorite dog runs out of the house and is run over by a car. The dog is dead. It's your favorite dog. So you are very sad. You take the dog, you cook the dog, and you eat your dog.

Now, you might say, there is no harm. The dog is dead. Why don't you eat it? Now, many people will view that as being totally disgusting. And that's the kind of thing. And many things about-- the more serious examples are things like duty, authority, loyalty to the in-group and so on, which are deep values for many people.

Now, an economist will not think of those as being deep values. They might think of that as being an instrument. So, sure, loyalty may be important, but that's an instrument to sustain cooperation, for example.

Same thing-- authority might be an efficiency enhancing stuff, but we would not value authority per se as an economist. But many people do. And that's the kind of thing that-- and that raises the more general issue about, what do we mean by harm? So and that's something which is very important now in universities.

If I tell you something you don't like, is that aggressive? So of course, if it's an insult, something racist or something like that, of course it is. But if it's an opinion, which is not your opinion, is that aggressive? So that's really a big thing and huge debate now on exactly those things.

So let me-- I mean, we could talk for hours about that. Do you have any views on that or any questions?

AUDIENCE: How much of morality do you think is formed [INAUDIBLE].

JEAN TIROLE: I'm sorry. How much?

AUDIENCE: How much is determined by society and how much is by some innate nature?

JEAN TIROLE: I guess it's pretty unknown. I mean, society plays a big role, of course, because of social norms. And as we see from some of these examples, they fluctuate a lot across time, for example, or across culture. So the social norms are very important.

Now, it's hard to say. Of course, we learn that actually punching the other child is not a nice thing to do. Now, is that innate or is that society? Well, you probably would have learned it by yourself anyway, but I don't know.

I mean, we still are a social species and a cooperative species, by and large. So how much do we learn through epigenetics? I don't know. I mean, anyone knows here? No.

AUDIENCE: So I guess that kind of changes how you think of the intervention. Because if it's innate then--

JEAN TIROLE: Yeah, it becomes more and more difficult. But I would guess much of it is not innate because you see the difference across culture and across time, that actually what is being considered moral in one society may not be in another society.

And that's why I was saying, do no harm. I'm more kind of universal. Everybody thinks that committing a crime, for example, is not moral. I mean, even pollution. Nobody will say it's great to pollute. Even those who don't care, they will not claim the moral high ground and say it's wonderful to pollute. No, no, no.

Those are relatively consensual issues in a sense. We may disagree about the importance, but the sign is clear. Yes.

AUDIENCE: So going back to the point on competition. So is there a sense in which having more competition can actually be a bad thing if the result is--

JEAN TIROLE: That's what we're going to study now. Because that's the claim, is that competition basically takes the morality out of people. So the managers or the shareholders, they will behave better if they were in a non-compete world.

AUDIENCE: But that doesn't include the possibility that, if there was no competition, that firms would still do immoral things? it's just to purely seek profit.

JEAN TIROLE: That would be a possibility. So we are going to look at an environment where people have moral preferences. So they care about their money, but they also care about other things. Which, I think we all are like that to some extent, to different extents, of course. But we care about the money, but we also care about doing the right thing, with different intensities of course.

So what we are going to see is that the nature of competition affects market morality in many ways. So it's going to depend on the intensity of competition, so market structure, the number of firms and substitutability, which are standard IO things. But it's going to depend also on whether prices are flexible or fixed.

You are going to tell me, prices are always flexible. Well, not quite, because sometimes your price is fixed. So if you are a doctor in many countries or a notary or whatever, you get some fixed amount of money-- or a taxi or-- you don't have price flexibility.

In franchising environment, your price may be determined by the franchiser anyway. So sometimes the prices are fixed, but the normal thing is flexible prices. Then we'll see it's important to know whether cutting ethical corners lowers cost or boosts demand.

By the way, we are going to look at a situation where all stakeholders can have moral preferences. There can be zero, but there can be also positive. So the stakeholders are, of course, you have the managers, but you also have the consumers, the investors, the workers.

So here is a picture. You will have an oligopoly with n firms. Those firms will set a price, except if the prices are regulated. And they will choose a morality of their-- a morality technology, if you want.

So they can choose something that pollutes a lot or pollutes little. They may be nice with their workers. They may misrepresent their product. They may-- all kinds of issues which creates anxiety on people and raises a moral issue.

They will sell to the consumers. And there are input suppliers, like workers and investors, that bring labor and capital to the firm. Now, those are stakeholders. So they may care, actually, about the moral action or the ethical actions which are chosen by the firms.

So a consumer, for example, might want to buy fair trade coffee or whatever, might actually want to have more production for their consumptions. Same thing-- the workers may care. They may want to work for an ethical firm.

By the way, if you look at-- there is a premium, actually, which is paid to workers working in, say, the tobacco industry. So if you work in an industry which is considered as immoral, you have a higher wage-- or casino or whatever. Same thing for investors. The investors, they might want to engage in socially-responsible investment.

Now, the key question, as we are going to see, is whether unethical behavior raises or lowers demand. Because if social responsibility, as we're going to see, goes through the cost structure, then it's not going to have any competition. It's not going to have an impact because competition is going to operate through the demand curve, not on the cost side. So we are going to see that's going to be. This is pretty unclear at this stage.

So the model is the following. Oligopoly of n firms serving unit-demand consumers, say. The firm chooses a price. Firm i chooses a price, p i, and an ethical action, ai. So think of ai equals 0 as the most immoral action, so you pollute maximally, for example. You sell to any dictator. You sell weapons to any dictator. And that creates some-- that has a welfare impact-- think of that as being an externality-- a welfare impact, wi of ai per unit of--

So you might think, for example, that might be the pollution damage from emitting carbon. It might be the damage done by the dictator with a weapon. So that kind of thing. And it's increasing ai. So [? i or ?] ai would be a more moral action just by-- it's just a normalization. ai will belong to some interval, with 0 being the lower bound. And the welfare impact will be-- so it will be increasing, and it will be concave.

The moral action may actually also affect the unit cost. So the unit cost of production, which will be a constant, will be ci of ai. Now, you can consider two cases, one in which this unit cost is increasing or decreasing with immoral action. So if you think about the moral action is not using child labor, not using fossil fuel, then it's expensive. I mean, people use children, they use carbon, they use coal because it's cheap. Otherwise, they will not be doing it. They use it because it's cheap.

So in that case, if you behave more morally, your unit cost goes up. But there are also cases in which, when you behave morally, your unit cost goes down. So that's a case-- coming back to the previous slide-- if you have responsible input suppliers-- so basically, workers or investors with social preferences-- then they are willing to actually have a lower wage or a lower return on their capital in order to be associated with a firm which behaves morally.

So in that case, if you behave more morally, your input supplier actually offers you a discount because they have social preferences. So in practice, you really have both. So we are going to be agnostic at this stage about whether the moral action actually increases or decreases the cost. That depends on the situation. Is that clear?

Now, a key concept for what's going to follow is the net price. Consumers care about the price, and they consider about the ethical behavior. So they also will have social preferences. I'm going to come to that in a minute. So everything for them will be as if they care about the real price p i, the nominal price. And then they will have some preferences over the action-- the morality of the action.

Now, this phi i can be also increasing or decreasing with the morality of the action. So usually consumers prefer a more moral action. They are socially responsible consumers. They want fair trade coffee. They want workers' rights to be respected and so on.

But they're also-- so in that case, phi i is decreasing in ai because that's as if the price were going down. But in the other example I gave you of bribing a foreign official, selling weapons to a dictator, giving more opioids, and so on, there, basically, phi i is increasing in ai. Because if you choose more moral actions, that means you are reluctant to sell weapons to the dictator, you are more reluctant to give opioids, and that's going to reduce demand.

So the same thing. The moral action can actually increase demand or reduce demand, but not in the same circumstances. So for fair trade coffee, more action is going to increase demand. For the example I gave you with replacement logic, actually, immoral action is going to reduce demand. And that's going to be completely crucial for what's going to happen.

There will be an outside option. So the consumer chooses between one of the firms and the outside option. The outside option will have an a0 and a p0. And p0 is often 0. So if you don't consume it, it might be 0. But still, if your option is actually to reopen the coal mines, then p0 is small, but a0 is also very small.

You have a vector of net prices. And those net prices are going to determine the demand. So the demand for firm i will depend on that price of firm i and the net prices of the rivals. There are substitute goods. So basically, my demand increases with your price and decreases with my price, the standard assumption.

So you can think about the discrete choice model, in which, basically, there are valuations which are drawn-- including for the outside option, by the way-- and you choose which option you prefer. And that's one way of generating such a demand.

So we are going to call eta i the elasticity of demand for firm i. So it's dDi over dpi, or dpi hat-- it's the same because p i hat is equal to p i plus phi i. So that's the standard density of demand. We make the standard assumption that when prices go up, the elasticity of demand decreases.

It's almost the same assumption as saying the goods are strategic complements. So it's strategic complementarity assumption, basically.

So you assume that for higher prices, the elasticities are lower, which is standard. But it's also standard just to assume that, basically, the goods are strategic complements. So the assumptions I make are--

So what's going to play a big role is this. I'm going to say that when demand increases with the morality of the action, that means the consumers are socially responsible. When demand decreases with the moral action, then the consumers are socially irresponsible.

So all of the examples I gave you on the first slide were examples with socially irresponsible consumers-- the dictator, the consumer of the opioids, the official who gets bribed, the athlete who actually wants a drug-- a performance-enhancing drug, and so on.

So socially irresponsible consumer, phi i prime is negative. Socially responsible consumer means that phi i prime is positive. So, for example, one of the things that might happen is that you internalize-- and this is going to be our first example of social preferences-- you internalize the welfare impact of your action.

So think about that as being the cause of pollution, for example. Then, you put a weight, alpha c, on the cost of your pollution, which, in general, would be less than 1 because you never fully internalize the impact on others.

Now, you have this strange case in which-- I mean, they look like Homo economicus, of course-- in which the consumers don't care about the moral action or, I should say, rather they may not observe the moral action. Because there are also cases in which you don't observe the moral action of the firm. In that case, their demand doesn't react to immoral action because they don't observe it. And they are kind of socially neutral. Do you have questions about the framework? Yes.

AUDIENCE: Sure. How restricted, if at all, is it to have the morality term entering through the price like that?

JEAN TIROLE: Yeah, so the additivity property. So the idea is that, in a sense, there is no income effect. So you pay a price, and then you get some benefit or some cost of the immoral action. So it's true if this particular purchase were really big compared to the income of the person, that might not hold, for example.

So think about a bribe, for example, for an official. So if the official gets lots of small bribes, that's not an issue. But if that was a huge bribe, which then there might be some income effect, which-- because then, that means the marginal utility of income, and therefore the price, in a sense-- the impact of the price will not be the same.

But it's pretty reasonable, I would say, in general. That's a good question.

OK, the cost side impact of ethical behavior. So we're going to assume-- it's more general than that, actually, but we're just going to assume that people have social preferences. So they internalize Wi to some extent.

So just to normalize, imagine that you need one unit of capital and one unit of labor to produce one unit of good. And then imagine that investors are willing to forego alpha i-- i for investor-- Wi of ai in their return. And same thing for the workers, with a different weight. So they are willing to work for lower wage if you are more moral.

Then, basically, your marginal costs, the way it changes with the action depends on the cost of doing business. So here, you might have the impact of child labor or pollution, if you use a polluting-- a cheap but polluting input. But then on the other side, you might also have the fact that workers and investors are willing to actually get a lower rate and a lower return on their capital.

So I already said that in a sense. Demand impact of supplier's moral choice. You have irresponsible consumers. Now, in the paper, or in the appendix, we have some foundations for the demand-side benefits of cutting ethical corners. It could be an externality. That's the case of those examples. It could be an internality, like an opioid. It could also be all kinds of shrouded attributes also. That could also be the case. And then responsible consumers.

I'm coming to the firm's objective function because that's-- so they too have social preferences. So they care about the profit. The profit is price minus unit cost times demand. That's nothing new here. But the difference with a standard firm is that they will have social preferences, so they care about the welfare.

Now, we can discuss that. What is welfare? Now, in the paper, we have general consequentialist preferences. And I'm just going to give you one example of that, which can be criticized, but it's much broader than that.

So I'm going to assume, for the sake of the argument, that you internalize basically the industry pollution or the industry's amount of opioids or something like that. So you take into account the sum-- so basically, firm j will get demand Dj. A fraction of consumer Dj will go to firm j.

The welfare consequence of firm j is Wj of aj. And you sum over all firms and also the outside option, possibly, of those welfare consequences. So you care, say, about the total pollution in the industry or the total amount of the opioids that are sold in the industry.

The class of consequentialists' preferences is much broader. So you might internalize, for example, just your own Wi Di so your own pollution, in a sense, which is harder to defend morally. Because you have to take into account also that if you sell one more-- if you get one more consumer, that consumer will have purchased elsewhere.

So morally, you will say, no, no, just looking at your own pollution makes no sense morally. But that's OK. For the theory, you can allow it. And there are other social consequences that you might want to consider. The important thing is that they are proportional to quantity. And that's something which also can be debated.

So if you are a consequentialist, you say, if I kill 1 bird-- if I kill 1,000 birds, it's 1,000 times the cost of killing 1 bird, for example, if it's not an endangered species. You might say, if I kill 1,000 birds, the damage is 1,000 times killing 1 bird.

Some people don't think in those ways. Actually, people would try to measure, say, the damage incurred when there is an oil spill, for example. You see that sometimes people are willing to pay almost as much to save 1 bird than to save 20 birds or 50 birds, which morally makes no sense, but it's a reality. I mean, we could discuss why it's the case, actually. Why do people react in those ways?

So you can criticize, but I'm just going to assume that people are consequentialists. So this is an important slide here. Do you have any questions? OK, so social preferences.

So the key assumption-- [? Ian, ?] and we might add your comment on separability-- but the key assumption in consequentialism says, all players'-- the suppliers, the consumers, the workers, the investors-- perception of social impact of their trade is proportional to the size of that trade. So if you pollute twice as much, it's twice as bad.

The second assumption, which is going to be crucial and-- read that carefully. What is a flexible price? So I will say that prices are locally flexible at equilibrium, configuration p hat. By saying that, if I change my ethical behavior, delta ai, a little bit, then I can compensate through a price change. Whatever sign it has, I can compensate through a price change delta pi, which has to be, of course, minus phi prime of ai. So it has to keep the net price constant.

So if I change the morality of my action, I can have an upsetting effect in the price. So that overall, I don't change the attractiveness of my offering. And therefore, demand remains constant in that case.

And there will be two important counter examples. The first is the case where p i is regulated, of course, because then you cannot do that. The second, which is more interesting in a sense, is the case where you have limited liability.

So you see, if you have this constraint which is binding, so if your price is equal to your marginal cost, then you are in the limited liability region. And if you say, for example, you increase the morality of your action, that costs you more money, for example. Then you have to raise your price accordingly. And that may not actually be neutral with respect to the demand curve.

So there is a link here that introduces a direct link between the price and the moral action. And same thing-- if you are not-for-profit, basically, you have to obey p i equals ci of ai. And again, you don't have flexible prices. Otherwise, you do.

And finally, we are going to assume constant returns to scale of a certain kind. So I'm going to assume that the cost function-- I already did that-- is basically-- I mean, you could have decreasing or increasing returns to scale. But at least with respect to the moral action, that has to be-- basically, you have constant return to scale.

Actually, there are a bunch of cases where you can generalize. So this assumption is not as strong as it looks. Those are the assumptions.

The first thing is that, if you fix prices, which will be the case when prices are regulated-- if you fix prices, then the ethical choices are strategic complements under weak conditions. So that's going to be true if consumers are socially irresponsible. So that means that demand decreases with the morality of the action. That's always going to be true if the economy is symmetrical. So if you have a symmetric economy, and you-- or if the social preferences are not too strong.

I mean, you have a first factor for strategic complementarity in strategy [INAUDIBLE], which is that we have assumed strategic complementarity in the first place on prices. But since prices are like ethical choices, you see, that translates into strategic complementarity for--

So that tells you that you may easily have multiple equilibria, or multiple social norms in ethical behavior, if prices are fixed exogenously, if they are regulated, which is not-- that's going to make things interesting-- which is not going to be the case at all if prices are flexible.

Actually, we are going to see that there are always uniquely determined if prices are flexible. So if prices are fixed, you can have multiple social norms about moral behavior. But when they are flexible, we will see there is only one possible norm.

Now, here is the first order condition with respect to ai. Actually, I don't know if-- you have seen that. You may not remember, but you have seen that many times.

If you have a standard oligopolist with alpha i equals 0-- so there is no more preferences, which is what you study in [14.]271 usually-- then you have Lerner index equals to 1 over the elasticity of demand, where the Lerner index is p i minus phi i over p i, the price cost margin.

That's true for a monopolist, but that's true also for an oligopolist. Of course, the elasticity is not the same-- elasticity of demand is not the same for an oligopolist and for a monopolist. But this formula is always true. That's just the first order condition with respect to price.

Now, for the moment, I'm not maximizing with respect to price. I'm maximizing with respect to the action. Just do the computation from the objective function. And what is that?

So the left-hand side-- so the right-hand side is something which is pretty well known-- product of the Lerner index times eta i, except we are going to see the Lerner index is a bit different. And this is basically, when you change ai, you have this internalization of your own actions' impact, Wi of a prime, minus your marginal cost.

And this phi i prime is just a normalization, in a sense, that transforms moral action into money. That's a marginal rate of substitution between the moral action and money for the consumer. So that's the equation.

But an interesting thing is the expression of the Lerner index. So as I mentioned, the Lerner index is always price minus cost of a price. If you have social preferences, the cost is a modified cost, which includes your social responsibility. Multiply it by alpha i, which is your internalization of social welfare.

And the social responsibility index, in a sense, has a very simple expression, which reflects what happens when you change your moral action or you change your price.

Here, I forgot to say that wi is equal to Wi of ai and wj is equal to Wj of aj. And sigma ij is basically the displacement. So for example, if I acquire one consumer, the question is, where does this consumer come from? It can come from-- either from j or from the outside option, because this consumer was not consuming before. That's why we have j greater than 0 to include the outside option.

And what matters is a difference in terms of welfare consequences between i and j. And the sum of sigma ij is equal to 1. So it's a weighted-- you look at-- you ask-- it's really about the replacement effect. You look at, if I acquire one more consumer, am I going to increase pollution or reduce pollution? Well, that depends where I get this consumer from.

So if I get it from someone who was producing using coal, that's a good thing, even if I use gas. But if I get this consumer from someone who was actually using renewables, that's a bad thing if I'm using gas. So basically, you look at the impact on welfare. But for that, you have to know, whom are you stealing consumers from?

So that's why, in your course, you also take into account your social responsibility, which means that if you are highly polluting-- actually, your social responsibility is actually to leave the activity to other firms. If you are highly ethical, your social responsibility is to take market, to increase your market share, which kind of makes sense.

Now, if you have a symmetric equilibrium-- if you have a symmetric equilibrium, then wi is equal to wj. Let's assume that the market is covered, so the outside option is not-- nobody chooses the outside option. Everybody needs a doctor. And basically, then wi is equal to wj, and then your social responsibility is 0 because you're all the same. In that case, the Lerner index is a standard Lerner index if you have a symmetric equilibrium.

But the basic point is that, under a pretty weak assumption, then you get multiple-- for given prices, you get, often, multiple ethical norms in equilibrium. That's easy to get because of the strategic complementarities. Whereas, when you look at flexible prices, you get a very different picture, which is that, if you have flexible prices, your moral behavior is never going to depend on the demand curve.

So it's going to be the same for a monopolist or for an oligopolist. It's the same for an oligopoly and a cartel. The demand curve, whether they are highly substitutable or not, doesn't matter.

So if prices are flexible, actually, then competitive pressure doesn't matter. Which is important because it tells you, just go ahead with competition policy. Those are two separate things. Moral behavior will depend on how moral the stakeholders are but will not depend on the demand curve, so will not depend on the intensity of competition.

So actually, in that case, I mean-- so put differently, if you have a price optimization, you know that you're going to get this. Price will be equal to the Lerner index, perhaps modified by the social responsibility. But it's always the case that if you have price competition, the Lerner index is equal to the inverse elasticity of demand.

So that will be equal to 1. And basically, this entity will be equal to 1 as well. And basically, your moral behavior, indexed by [INAUDIBLE], will be given by this equation, which means that it's completely independent of the demand curve.

So whenever prices are flexible, the intensity of competition is completely irrelevant. And then therefore, there is no point not doing antitrust because you want to increase morality, because you're not going to increase morality anyway. It's going to remain constant.

OK. What is intuition? Let me give you two intuitions for whatever they are worth. The first is to understand there are two offsetting effects. So takes a particular case. So we are going to take a particular case, which is a worry of many people-- the fact-- the replacement effect. So the idea is that I might want to bribe an official because, otherwise, I will lose the market. So irresponsible consumers.

So with irresponsible consumers, if I cut ethical corner, so if I reduce my morality, I gain market share because that's exactly what the corrupt official, the dictator, or the opioid consumer wants. So if I behave less morally, I gain market share.

Now, if you change the elasticity of demand, so you increase the elasticity of demand by having more competition, then that means that I'm going to get more market share from-- I will get a bigger increase in market share if I behave less morally.

So that seems to confirm the conventional wisdom that, actually, with the responsible consumer, markets will lead to immoral behavior. Because then, by cutting ethical corners, you actually gain a lot of market share.

But at the same time, if you have more competition, you also have lower markups. And if you have lower markups, that means you gain less in terms of money whenever you acquire market share.

And that makes-- you have a tradeoff between behaving morally and making money. But when you behave immorally, you make less money per unit on more units in a sense. And those two effects actually perfectly offset. That, of course, I'm not showing you. I'm just showing you there are two effects which go in the opposite directions.

I mean, if everybody made 0 at the margin, nobody would have incentive to behave immorally. If you don't make any money from misbehaving, then you'll behave morally if you have any social preferences, even weak ones. And competition is going to reduce the stake, in a sense. So it's going to push toward more moral behavior. So you have two effects, and they offset.

The other way to see it is some kind of cost minimization. And that explains why they perfectly offset, in a sense. So when you look at your costs as a function of ai, ci of ai, but then you also have to add phi i of ai. Because phi i of ai is basically the perceived price increase due to ai for the consumer. So if you want to keep your demand constant, you have to reduce your price by phi i of ai.

And then there is your social responsibility, which we wrote here. So remember, sigma ij is the fraction of consumers that you steal from firm j whenever you acquire new consumers. And then I rewrote, slightly differently, the objective function. So I have to take into account the difference between the welfare that I offer and the welfare which is offered by my competitors.

And when you look at this, given that sum of sigma ij is equal to 1, you can-- you can rewrite it this way, in a part which-- as a part which depends on my own action and another part which depends on the prices, the net prices, but also on the action of the rivals, including the outside option. And it's completely separable.

So that means that, OK, that I cannot do anything about. For given net prices, I can just minimize costs. And that gives me the formula, which is on the previous page.

So that's a proof. Now, how robust is this conclusion of irrelevance? It's not fully robust, but it's actually pretty robust-- more than we would have thought.

So for example, the relevance of-- again, the relevance is that the intensity of competition has no impact on the choice of morality if prices are flexible. You can use Cournot. In that case, you don't choose a price. You choose a quantity. But that's still the case that works. You can use a search model. It also works.

If you have imperfect consumer information, that works. If you have different internalizations, I give you a specific internalization, you can generalize to nonlinear pricing. So I've assumed you need demands, but you could have multiunit demands. That will also work and so on.

Non-constant returns to scale is important for the result, but you can relax it somewhat. So, for example, if ethics impacts only the demand side but not the cost side, it doesn't matter. You could have decreasing or increasing return to scale and so on. So it's pretty robust.

Where things are different-- and probably I will stop here because-- is if you have-- if you don't have flexible prices, then it's a different ball game. So actually, it suggests lots of empirical tests.

There are actually a couple of experiments which have been run which fit relatively well with the theory, but it's only two experiments. We need more experiments, and we need also some empirical work, in general, to check all those things.

So what we do is-- OK, let's assume, first, you have a regulated price. If you have a regulated price, you have no reduced-markup effects of intense competition. Competition reduces markup, and therefore reduces the stakes in behaving immorally. Whereas if you have fixed prices, that's not going to happen. Competition is not going to reduce the price.

In that case, the critiques of markets are right, which is, if you have irresponsible consumers, then more competition is going to reduce moral behavior. That's correct. So if you make the assumption there is a unique equilibrium, you can show that actually more behavior decreases with more competition.

So there's a grain of truth in the criticism. It's not like it's completely wrong. There's a grain of truth. But you have to qualify it substantially, not only because of the irrelevance result, but even when the prices are fixed, then you also get the reverse conclusion when you have responsible consumers.

It's pretty clear what happens is that, if you have a responsible consumer-- so we all want fair trade coffee. Then, actually, more competition is going to force the firms actually to buy fair trade coffee because they want to please the customer. So it's going to be exactly the reverse.

So if you have responsible consumers, then competition is a good formality. So if you add those things, which is the relevance result plus this slide, you see that, first, it's not the case that everything goes. You have predictions, but they are context-specific. Is a price flexible or not? Are the consumers irresponsible or responsible? Those are not the same applications.

And you get a picture, which you say, OK, we cannot conclude. And just modifying antitrust and all those things just on the basis of morality doesn't seem like the right thing to do.

Now, if the policymakers use the [INAUDIBLE], maybe that's the right thing to do. But you cannot just say, we don't want markets. That's not the right stance.

Now, the last thing we do in the paper, I think, is to look at what happens when you introduce either heterogeneity in social preferences. Imagine one of the competitors is Mother Teresa. What happens? Well, she might have a limited liability constraint. Because at some point, she will run out of cash. Or if you have a for-profit, then it becomes more interesting.

And there might be an argument in some industries, actually, to protect the not-for-profit from the competition of for-profit. Because what we show is that, when you have very intense competition, you tend to have mimicking-- the behavior-- mimic between for-profit and not-for-profit, the distinction disappears.

Basically, the for-profit to compete with the not-for-profit has to choose prices close to marginal cost. But the not-for-profit have to behave unethically to keep consumers if consumers are irresponsible. And you see that you lose a little bit the distinction between not-for-profit and for-profit in that case.

And that raises the issue about whether you want to reserve the hospital market or the education market, where moral issues can be very important to not-for-profit. And we don't solve this issue, but it's actually a very interesting issue to study.