

[SQUEAKING]

[RUSTLING]

[CLICKING]

ROBERT TOWNSEND: So let me catch us all back up. We have one more week of classes consisting of two lectures. They're pretty exciting because in effect, everything comes together at the end. The one today is the failure of the welfare theorems, but not just that a theorem doesn't work, but why it doesn't work and also remedies that are going to make its way into proposals for new market structures that get rid of the externalities, existing and new.

And then the one on Thursday is on monetary theory where we will actually have money and cryptocurrency in the model. And we can address things like the price of Bitcoin, which yesterday hit an all-time record, but I haven't looked at it today. And that's kind of the point.

So on the reading list, today is Failure of the Well Welfare Theorems. There's two starred articles here. It's a little bit unconventional relative to other lectures in the sense that I'm not going to cover either one of these in enormous amount of detail, and yet they are a starred. So it's fair game in terms of things that I expect you to know.

I only have four or five slides at the end on the Jain-Townsend paper, but it relates to cryptocurrency and platforms for exchange. So it is very consistent with some of the themes we keep hitting on in the course. And this Arrow article, very famous one, I have only today extracted two or three-- really two slides from a much larger and much more interesting paper. So take a look at those.

Are there-- before I ask you all, are there any questions from last time, about the lecture from last time? So last time was identification and falsification with data. There is a question that's actually not written down here, but I feel I want to ask anyway. Namely, what was the point of the lecture? Why do we care about being able to falsify a model with data? And I'll take volunteers. Why is it important to be able to reject a model?

STUDENT: I mean, I guess it makes it more testable.

ROBERT TOWNSEND:

STUDENT: If we have empirical data, we can tell whether it's even feasible for it to be that model.

ROBERT TOWNSEND: Yeah, yeah. That's good. Or a more brutal way to put it is, if it's not falsifiable, then it can't be rejected. And if it can't be rejected, what is the point of taking it to the data in the first place? That's a little bit harsh in the sense that you could accept the model as an abstract version of reality, use it to estimate parameters, conduct policy experiments-- and we did that several times in the lectures on trade and so on.

But this lecture embraces the full empirical experimental science point of view that we want to make some progress by ruling things out. So you really want to if you're in the wrong class of models. And the way to do that is to at least have the possibility of data that would be inconsistent with the model so that you would know. OK. So that wasn't asked. Thank you for volunteering an answer.

So apropos along that line, how do we test whether an observed demand function came from a consumer-maximizing utility? Or I guess more specifically, does anyone remember one, two, or all three of the implications that come from the premise of consumer optimization for data?

STUDENT: It should have homogeneity of degree 0. It has to satisfy while reaching equilibrium. And the Slutsky matrix has to be symmetric and negative [INAUDIBLE].

ROBERT I'm not actually hearing you very well. Can you repeat that?

TOWNSEND:

STUDENT: Oh yeah. The demand function has to be homogeneous of degree 0. It has to satisfy while reaching equilibrium. And the Slutsky matrix has to be negative, [? semi-definite. ?]

ROBERT Yeah, perfect. So satisfy-- I guess we call it Walras' law, but essentially means that they're spending all their money and they're not in the interior of the budget set. And yeah, homogeneity of degree 0, no inflation illusions. And the Slutsky thing is the thing we derived pretty carefully the last time.

Does anyone remember what-- and can state the weak axiom of revealed preference?

STUDENT: Let's say we have two pairs of observations containing prices and consumption bundles. And for example, for like p_1, p_2 , the x_2 is preferred over x_1 . And then for p_1 , the x_2 should be outside of budget. Otherwise, x_1 would not be preferred under p_1 .

ROBERT Yep. Yeah, you said it right I think except for the very end. We start with two pairs, p_2x_2 and p_1x_1 , and-- so
TOWNSEND: maybe I have it here. I'm not finding the slide I wanted that I thought was going to be right in front of me, but--

STUDENT: It's in a previous page, the-- yeah, yeah. It was just a statement, yeah.

ROBERT So we have these two pairs and they're not the same, so there's a real experiment here. And we start with the
TOWNSEND: second pair. Under prices p_2 , they choose x_2 . They could have chosen x_1 , but did not. They could have chosen because x_1 is in the budget and they did not. So they revealed to prefer x_2 , quote-unquote.

But if that's true, then when at price is p_1 when they're choosing x_1 , x_2 should not be available. Because if it were available at prices p_1 , having already deduced that x_2 is preferred to x_1 , they would likewise choose x_2 in that situation and that would be a contradiction. So thank you, yeah, you said it fine.

OK. So again, those are just sample-- somewhat not-so-randomly-sampled questions, but there are others there I could have asked. I was tempted to go over the thing about the Slutsky matrix again and the concavity of the expenditure function and so on, but I-- and how we end up being able to deduce restrictions on data from the Hicksian demand which we don't observe. But anyway, I didn't ask. So OK.

So that's-- oh. And then the one other thing, which is actually the reason I had this open-- wow. I'm having trouble manipulating the slide deck. This Echenique. So one of you asked me last time, I guess in my office hour on Thursday, what's going on here?

It's actually all on the slide. I took the time this morning to look up the paper on the-- I don't have the book. The book is in my MIT econ office and I have been in there for nine months. But portions of the paper are online. And it's algorithmic and just saying, if you have the data, like the kind of data we have in revealed preference-- so we see prices and quantities and we're trying to-- in integrability, we're trying to deduce the utility function or a utility function that would generate the data.

So the issue is, if you had a utility function the household would be maximizing it to generate the data, how hard is it to maximize the utility function? And the spirit of this was, there may be a whole set of utility functions. There may be some that would generate the data for which it would be actually hard to compute the solution. Like Lagrangians are iterating over possibilities and taking local derivatives and so on and so forth.

When you only have a couple of commodities and so on, it's obvious graphically what the maximum is. But in higher dimensional spaces, it's not so obvious. And the point of this article was only to say that if you choose a utility function that is consistent with the data that is hard to compute, there's another one that would generate the same data that for which the algorithm increases in complexity as a polynomial with the number of goods.

So it's definitely what you would have thought all along. It's a computer science paper about algorithms. And to get in into it in any more detail, one would actually have to look at the algorithm and deduce why it's only increasing in polynomial complexity. Anyway, I just wanted to clarify that since one of you asked me this last time.

OK. So for today, the failure of the welfare theorems. Sounds ominous. OK. So here's an outline of the lecture. We'll do failure of the second welfare theorem first, because ironically, it's kind of easier. Then we'll go to failure of the first welfare theorem, which is harder in some sense because it's so easy for it to be satisfied, but will feature local satiation which violates one of the sufficient assumptions.

And then when other reasons for failure of the welfare theorem have to do with incomplete markets where there is not a market in every good. Pollution is an example of that, and you may or may not be aware of markets for pollution rights as in cap and trade, and I'll go through that.

The generalization of that is when households are altruistic about each other or jealous of each other, they really care about what other households are consuming in a positive or negative way, that generates another externality, but in that Arrow article, there is a way to solve the problem, by creating markets for the missing goods.

Then I'll get into another failure of the welfare theorem. Go through the proof that we used and show you where it could fail if there's an infinite number of goods, which brings us to overlapping generations, nobody lives forever, but the economy hopefully goes on forever. And we have money in the economy. So that will take us to the lecture next Thursday.

And finally, I want to touch bases again on these platforms for trading cryptocurrency or anything else for that matter, which seem to be very far away from the welfare theorems and the frameworks we're using. It looks like there's an externality again. And the theorem ought to fail, but I'll show you how to design those markets in such a way as to recover optimality.

So at that point, we kind of crossed the line from positive to normative economics. We're using economic science to advise policymakers on what to do. So here's a statement of the second welfare theorem, which is literally pasted in from earlier lectures. So you know everything on this slide. An economy consists of consumption sets and preferences and endowments, production assets, ownership shares.

And for the second welfare theorem, we kind of needed a lot. We needed convexity of the consumption sets, we needed the production sets to be convex and admit this concave transformation. Preferences, also concave, locally non-satiated. And these kinds of interior points, so to speak.

And then we have a statement of the theorem that for any λ , hence, for any Pareto optimum, there exists a price vector and an assignment of wealth that would constitute a Walrasian equilibrium with transfers with those wealth assignments. So the lecture where we went through this featured two ways to attack the problem.

One, the Lagrangian problem both for the Pareto problem with the λ weights and for each of the firm and household maximization problems, then we just matched up all the Lagrange multipliers, but that's that only work-- typically only works for sure if you have all this convexity and concavity. And the second way we did it was just separating and supporting hyperplane theorem which, in turn, relied on concavity or convexity of the steps.

So on that note, here's an illustration when it works, suppose you had a single representative consumer, this could be our Gorman man, maximizing utility subject to production possibility set-- and this is obviously the overall maximum. So then we can create this budget line for the household and a isoprofit line for the producer which separates these two convex sets, the production set from the weak upper contour set of bundles that give you at least as much utility as that this red dot.

What if preferences were not convex? So we have a nice pink convex production possibilities set, but this is an illustrative indifference curve with all these waves. So it has concave, but convex portions. Or alternatively, the upper contour set of utility points which are greater than or equal to the utility achieved at the red dot, that's not a convex set as you can see by taking say these two points in the blue set in a weighted combination of those two points would not be in the blue set.

Another way to say that is at this budget line, which separated on the previous slide, would not separate here, because although we have profit maximization, we do not have utility maximization, and you can imagine doing better here, say, a tangency that would achieve a higher level of utility than a red dot. So the second welfare theorem would fail. This is a competitive equilibrium, but it's not Pareto optimum.

Over here, we have non-convexity of the technology. So the upper contour set is convex, but you see these non-concave portions. So there's not a concave transformation representing this production possibility set of-- production possibility set is not convex. So again, the candidate here would be the budget line for the household for which we have utility maximization. At that budget line, we would not have profit maximization because the producer could be doing better.

I think I misstated the result over here. We have a Pareto optimum. In the second welfare theorem we start with Pareto optimal points. And then we're trying to decentralize it. In this case, we would have profit maximization, but we would not have utility maximization. In this case, we would have utility maximization but not profit maximization. So these are perfectly fine Pareto-optimal points, they're just not decentralizable and the second welfare theorem fails.

Now it doesn't just because we don't have the convexity doesn't mean it always fails. This is illustrative. I've shown you this picture before. We have this upper contour set with various wiggles in it and the production possibility set with portions of increasing returns, but nevertheless, at this optimum and the associated supporting hyperplane, we would have utility maximization and profit maximization.

So we might get lucky even with the non-convexities, but it's just not guaranteed. Those assumptions in the second welfare theorem are sufficient assumptions, not necessary.

First welfare theorem. And again, I've kind of, quote-unquote, just pasted in this slide, let's start with a price equilibrium with transfers, including, say, Walrasian equilibrium with x^* , y^* , p , and w . And then if preferences are rational and locally non-satiated, x^* is a Pareto optimum. That's the theorem. Start with a competitive equilibrium, it will be Pareto optimum if it's sufficient-- that is to say, to have rational locally non-satiated preferences.

I remember in presenting this, I emphasized how weak the assumptions were, because you might think they're always satisfied, hence, every competitive equilibrium or equilibrium with transfers would be Pareto-optimal. However, I'll show you a picture of where it's violated and we're going to pick on this local-- I'm going to show you an example of locally-satiated preferences, and hence, failure the first welfare theorem, and here we are.

So this is the Edgeworth box. One household's over here in the southwest another here in the northeast. Southwest one's easy to see what's going on. There's an endowment point, the budget line runs through the endowment point. This household southwest is maximizing utility according to the red indifference curves and has achieved the highest possible utility level within his or her budget set.

Household over here, northeast, has indifference curves that are, in general, monotone-increasing as you move southwest. However this shaded-in blue region is the region of local satiation. So the idea is that that household has an indifference band, indifference curve band, not just curves, but this whole interval of points are such that the utility level is the same for this northeastern household.

And maybe just in that region, maybe the rest doesn't matter. OK, so what's the first welfare theorem? It says that the competitive equilibrium is optimal. Here's the competitive equilibrium. Household southwest is certainly maximizing utility. Household northeast is maximizing utility because this point gives at least as much utility and maybe more than any other point in household's northeast budget set in this whole region. So it is a competitive equilibrium.

But is it Pareto-optimal? No. Why? Because we could move in this northeast direction and increase the utility of the southwest household without hurting the utility of the northeast household. So there you go. That's an illustration. So we definitely use implicitly this assumption about local non-satiation. And I could have drawn pictures that had to do with discrete points and consumption sets which also violates local satiation because there's nothing nearby.

OK, so let me go back. Another thing that's easy to forget in this definition of an economy and the associated welfare theorems is we have L goods, say, finite dimensional, and households have endowment and preferences over consumption for each and every item in that L dimensional space.

So when we're coming up with these prices p , we're pricing each and every commodity in the entire commodity space. Another language we could use is incomplete markets rather than complete markets. There could be items that households care about and that firms can produce for which there is no price system. You can't trade it. And that will also cause the first welfare theorem to fail. So I'm going to give you a well-known example, which is pollution.

So in the example there are two goods, one representative consumer, and one firm. The firm can transform one of the goods, good 1, into the second good, good 2, according to this production function. And remember, for firms, inputs are negative. So minus y_1 is the positive quantity of the input of the first good mapped through the production function F to the output of the second good.

Now the idea here is that the firm generates pollution when it produces a second good. The act of transforming good 1 to good 2 produces this level of pollution, but for the minute we're not going to have a market in pollution. It's a genuine good-- actually, a bad, but for the moment no prices on it.

Preferences and production both specify this bad pollution good. Households care about their consumption of good 1, good 2, but also this pollution P . And I should have said, down here in red, consumer 1 begins with 1 unit of good 1 and no units of good 2. So if the consumer is going to get a hold of good 2, it has to come from firms producing it with the input surrendered by good-- of the household of good 1 as an input. OK.

So the utility maximization problem is just maximized consumption of good 1 and good 2 with the fallout that there's pollution, P , as a consequence of y_2 . y_2 , the second good is exactly what the consumer is eating, but y_2 comes from production using the input of the first good. And the input of the first good, being negative, subtracts off from the unit endowment that the household has of that good, the residual being household's consumption of good 1.

OK, so this is like our Pareto problem with a couple of simplifications and one complication. There's only one household, no need for the λ weights. You can think of this as a Gorman person if you want. And there's only one firm, although we've done that before. And now we've had we have this non-traded good, this pollution aspect jointly produced along with y_2 . OK.

So this is a Pareto problem. So I'm about to determine-- characterize the set of all Pareto-optimal allocations. We do that by maximizing utility subject to effectively the production set and the resource constraints. I've lined up the Lagrange multipliers on equations 5, 6, and 7 and ordered them 1, 2, 3. So those are the γ s, those are the shadow prices.

OK so everything's convex in this world. So we can do the first-order conditions for the Lagrangian differentiating with respect to x_1 , x_2 ; consumption for the households, y_1 and y_2 , inputs and outputs for the firm. You can check me. At this point, you should be familiar with the Lagrangians. You may want to go back and review it. The only subtlety here is to get the signs right. When we write down these constraints, we should write them as greater than or equal to inequalities so that we get the interpretable sign on the shadow prices.

Well, they should look kind of familiar. This says the marginal utility of good 1 is equal to λ_1 . The marginal utility of good 2 equal to λ_2 . And the household's utility function. This is a statement about the marginal product of good 1 in production. And this is a statement about how the Lagrangian is changing as we change y_2 .

Here and only here you can see that the Pareto optimum is taking into account the extent to which pollution is created by production. And if you go back and look over here, the y_2 is written-- P is pollution written as a function of y_2 . So when we differentiate with respect to y_2 , we pick up that partial P , partial y_2 in the household's utility function. And of course, we pick up y_2 and the two other spots where it enters in equation 6 and 7. So that's where this thing is coming from.

OK, so the rest is algebra. You can see that we could solve for the inverse. If we took df to dy_1 and set it equal to λ_1 over λ_3 , we could then go over here and find λ_3 and λ_1 and make some substitutions. Actually, we flipped it upside-down. Instead of df dy_1 , we take 1 over that, but we do go and look for those substitutions, which would be γ_3 , which picks up this mess over here, and γ_1 , which picks up this guy. So that's the algebra behind this expression.

OK, so let's interpret it. The marginal rate of substitution for the household had a Pareto-optimal allocation. And here, we kind of do something unconventional. This is the marginal rate of substitution of good 2 for good 1 and not the other way around. Normally if you go back and look at the consumer lecture 2, I guess it was, we had the marginal rate of substitution of good 1 for good 2 being u_1 over u_2 . And also the price ratio we'll get to being P_1 over P_2 .

So here, it's kind of flipped the other way around. It's not wrong, you just gotta remember, we're redefining the marginal rate of substitution to be the opposite of what we did before. OK. So the marginal rate of substitution for the household is the marginal utility of good 2, marginal utility divided by the marginal utility of good 1, but this is equal to, from this expression, the first part, numerator-denominator plus numerator-denominator. So this plus part numerator-denominator is here, it's just rewritten.

So this is the full characterization of the Pareto-optimal allocation. And it kind of looks like the marginal rate of substitution for the consumer is equal to something having to do with the marginal product or marginal rate of transformation in production, but we have this extra term and this extra term reflects the fact that to produce y_2 , in order for the household to get it, x_1 gets surrendered, put in the production function, but it also produces pollution.

So we're taking into account-- subtracting, as it were, the bad thing that happens, the extra pollution that happens as a consequence of producing good 2 which generates pollution P . I'm not taking you through all the-- I'm probably taking you through too much of the algebra. You've got to be careful about these signs, negative versus positive. I was just pondering that, but you were just subtracting the P 's from the left-hand side. OK. You can check the algebra later.

Now we want to decentralize this optimum like we've been doing before in the first welfare theorem. But we're going to do it in two steps. In the first step, we have a market in the two goods, but not in pollution. So the firm would maximize profits. Let's assume that the price of the first good, good 1, is the numeraire. So I'm not going to be writing p_1 anymore, but I will be writing p_2 . So this is the price of good 2 relative to good 1 because a good one is the numeraire.

This is the output, this is the revenue of the firm. This is the input cost. p_1 to y_1 . y_1 is negative, that's why it's being subtracted. Subject to this transformation set which we've talked about before. Now I'm going to make a substitution. If y_2 is F of minus y_1 , then take the inverse function, F inverse on both sides, and we get F inverse of y_2 equals minus y_1 . Multiply it by a minus sign and we get y_1 equal minus F_1 inverse of y_2 .

So I'm going to substitute this expression into this term y_1 in the objective-- in the profit objective function. Then when we differentiate with respect to how much y_2 to produce, we've already taken into account the amount of x_1 that has to be higher to get there. The derivative of this whole expression is going to be p for the first term, that's the revenue, and then the derivative of the second term, which is how this inverse production function is changing as we change y_2 .

It is the inverse function, but there's an inverse function theorem which you may not remember, which is the derivative of the inverse function is 1 over the derivative of the direct function. So this is the derivative of the direct function mapping input y_1 into y_2 . So in a competitive equilibrium, we get exactly what I was alluding to before, that the marginal rate of-- well, I haven't said it yet. The marginal revenue product of a good 1 into production is equal to its price.

And it's also going to be true that this price is going to be equal, from the consumer side you'll see momentarily right here, equal to the marginal rate of substitution in utility terms for the household. So it looks like you get the right thing when you decentralize it, that the marginal rate of substitution is equal to the rate of product transformation, but it won't be the right thing because it's neglecting the indirect impact on pollution from using good 1 in production.

Here's the other half of the Walrasian definition. This is the household's maximization problem. To maximize utility, including acknowledging the disutility of pollution subject to having 1 unit of good 1 valued at p_1 equal to 1 plus the profits-- this profit-- that the firm produces that's handed back lump sum to the household. We only have one household, one firm. This household has complete ownership share of profits of the firm. And we get this first-order conditions, which is what I've said. The marginal rate of substitution is equal to the price ratio.

So then if we take 11 and go back and remember 10, 10 was this thing from production, and substitute in, we'll say the marginal rate of substitution in an equilibrium is equal to the marginal rate of product transformation, but that rate of product transformation is less than what you would get if you take the marginal rate of transformation and add on a positive term.

It may not be obvious in the algebra right away that this term is positive, but it is, because the marginal utility is decreasing with P . So this is a negative object. And you've got a minus sign here. So negative-negative equals positive. This whole thing is a positive term, so that's why we have this strict inequality.

This thing on the right-hand side, though, is, as I said, from 10, the marginal rate of substitution under the optimum, this guy here. Actually, it was 8. We used 10 to get this. So the conclusion is the competitive equilibrium is not Pareto-optimal, because we want this expression to hold as an equality, it's holding as an inequality.

OK. So by the way, here's the economic intuition for what's going on. This marginal rate of substitution is too small. How do we get it to be bigger? We'd like to increase the numerator. How do you get the marginal utility with respect to good 2 be higher? Less of good 2. And/or get the denominator to be smaller, that would also make the marginal rate of substitution higher. How do you get marginal utility of good 1 to be smaller? Increase x_1 .

So the economics here is telling us we want to take the competitive equilibrium and move to the direction of increasing x_1 , eat it rather than use it as an input to produce x_2 of which we have too much. So it's intuitive in the economics that if you don't have a market in the pollution rights, you're going to have too much pollution.

But we can fix it by having the pollution rights. So what do I mean by that? Suppose the household can sell rights to pollute to the firm and that we can enforce those rights in the sense that if the firm doesn't have the rights, then it can't pollute. The rights will specify the number of units of P . Firms are going to have to buy those rights. It comes with y_2 . So the level of P that comes from y_2 has to be consistent with the rights that the firm purchases, and likewise, consistent at a price of rights-- with the rights that the household is willing to surrender.

The household is going to get revenue from selling the rights to pollute to the firms. So that's a positive thing. But on the other hand, the household knows that the firm will produce with those rights and it will suffer that pollution. So there's a trade-off for the household, but this is the right trade-off now, because we have the market for the pollution rights.

OK. In notation, we maximize the utility of the household over the two goods, x_1 , x_2 , and pollution level \bar{P} subject to the budget, which is get revenue from selling your endowment of good 1, you get profits from the firm. You also get revenue from selling pollution rights to the firm, the price being q , the per-unit price. Per unit of pollution is q . First time we've used that.

And if you look at the solution to this problem through Lagrangian's, we now have three goods, good 1, good 2, and p , \bar{p} . If you thought about μ as the Lagrange multiplier on this budget constraint 12, you get the obvious that the marginal utilities of the goods are equal to the marginal utility of income times the price.

Again, p_1 is equal to 1, so there's no price in the first term. Third term is kind of interesting, because it's got a negative sign. But again, pollution is a bad from the point of view of the household. So this marginal utility is negative, the negative sign makes it positive. Positive equal positive. The marginal utility of reducing pollution is equal to the marginal utility of income times the price of pollution rights.

So now in an equilibrium, we've got the marginal rate of substitution over the two goods from the household point of view equal to p . We could also look at the marginal rate of substitution of good 1 for pollution \bar{P} -- \bar{P} , actually, in the household problem, and that's equal to q . So we have three goods and we have two ratios.

For the firm's problem, they're going to maximize profits. This is revenue. This is costs. They have to buy those rights, same q . And rights are an input. So we put a negative sign here, because they've got to buy them as an input, the way we've been signing inputs all along. And anyway, it's a cost, so it makes a lot of sense.

So we rewrite this problem, substituting in this inverse function for y_1 , and take derivatives, and we get, with respect to y_1 , where does it enter? It enters here and here. y_1 -- I should have said y_2 . By making the substitution, we don't have y_1 in the expression anymore. We have y_2 in terms of revenue in y_2 in terms of production and the associated pollution.

So then just manipulating this equation-- well, sorry, looking at the first-order condition we get this thing, which is the marginal revenue on the left-hand side and the derivative of these two terms on the right-hand side. And now we can go back and do our checks of the first-order condition.

The marginal rate of substitution for households of the two goods is p . The p is equal to this thing from 14. So substituted in, leave the first term the same, and then find the expression for q up here from the consumer's problem, and keep writing dP/dy_2 . So that's the algebra.

And now, guess what? We've got a full optimum. This is exactly the expression we wanted over here in the full optimum problem, not the decentralize problem, but in the optimum problem. The marginal rate of substitution is equal to the rates of production taking into account pollution and the adverse utility consequences of pollution.

This equation 8 is exactly what we just derived here at the bottom of this slide. So we've achieved it, but we achieved it by creating the missing market, which was, in this case, was pollution. So OK, so there's a general lesson here. Now you may have heard of pollution rights before under the terminology cap and trade. So that's kind of related.

Particularly you let firms pollute all they want, but they have to buy the rights to do it. So you-- maybe endow firms with arbitrary assignments of rights, and amazingly, it's not going to matter who-- it's not going to matter for production who is initially assigned a high right to pollute and who's assigned a low right to produce and pollute. And then you'll see that now.

So we have the two firms, call them a and b. They're both endowed with pollution rights, \bar{P}_a and \bar{P}_b . And then we look at the profit maximization problem of each firm j . To produce revenue from output less the input required of good 1, that's this inverse function, which is what we've been doing before, except now we have a j everywhere-- this is firm type j .

And then we have this revenue or cost. Q is the price of pollution rights, they're endowed with pollution rights, \bar{P} . So we could stop there and just say they get that revenue. But not if they're producing, because they need to have the rights corresponding to the amount of y_2 that they're producing.

So we subtract off the corresponding rights that they need for a chosen level of output y_2 . This term could be-- the whole term could be positive or negative. If they're overly endowed with rights, they'll be selling some and be left with a positive residual term. If they're underendowed with rights, they will be buying more rights than they were endowed with.

But the economics is the same. Regardless of what the \bar{P} is, you want to evaluate the price of the permits as an opportunity cost. That is to say, take the case where they're overly endowed with rights, they could hoard them, and then they could effectively sell them in the market. That would be fine. Or they could use them up.

So the opportunity cost of using it up is what they could have done with it, which is just sell it outright. So this opportunity cost argument works for both types of firms. And it is the logic of why the endowments don't matter. What matters is what they do on the margin. That's what they take under consideration when they're thinking about production.

A way to see that in the algebra is just to take the derivative of these profit statements for each firm j with respect to the choice variable y_2 , taking into account that pollution is also a function of y_2 , and we would get, again, the derivative with respect to the revenue with p minus-- or put it on the right-hand side, the derivative of production through the 1 over the direct inverse-- the derivative of the direct production function, making that substitution plus q times the way pollution rights are changing-- pollution permit requirements are changing as we're changing production.

OK. So this equation 16 holds for both firms j of type a or b. Can we see who's polluting more? Well, let's suppose that this coal mine in West Virginia has really bad consequences for pollution as opposed to something using clean energy. The marginal effect of pollution is higher for type a firms relative to type b firms.

But if this term is higher for a than for b, and for both a and b, whatever else is going on the right-hand side, the two terms have to sum up to p, then for a, if this is higher, q is the same, this term has to be lower. How do we get for a firm this term to be lower? We make the marginal product of the input higher. How do we do that? We do that by reducing the input.

So again, the economics make sense. The firm that-- type that pollutes the most will use less of the input and therefore pollute less. So this is another example of pollution rights and maybe a more familiar one, which is cap and trade. Which, I think you know, is implemented in the US. California has various kinds of permits for polluting and so on. They were increasingly popular and then things happened over the last four years, but we won't go into that. OK.

Walrasian equilibrium with pollution rights. So there's a more general idea here, which is that be externalities all over the place. Households can care positively or negatively about what other households are eating. Households could care about what firms are producing-- you just saw an example of that.

So the consumption set would be the standard one over the L dimensional commodity space, but in addition, it would specify what other households are eating of each of the goods over all the other households, i minus 1 households. And again, what the firm is doing. Or if they're j firms, what each and every one of them is doing potentially.

Not everything needs to be in here, but this is-- externalities written down in its most general form. OK. So we did a version of this just now with pollution, and we can do a version with consumption. So this is that part of Arrow I was mentioning when we went over the reading list. We have a pure exchange economy to make it a bit easier. We have the usual. Household cares about consumption good k. There are n households. So i goes from 1 to n. There are m consumer goods, so k goes from 1 to m.

So this part so far is pretty standard. However, suppose that the utility function of the household depends on the consumption of the others. So if we started with household j, for example, if they could choose it, this would be the rights to eat that j would assign to household i for commodity k. In other words, household j cares about what household i is doing. We're going to try to decentralize this. So we want to create a market in those rights.

Mathematically we're just adding another subscript and you have to memorize that it goes what j wants to assign to i for commodity k. Well, that's household j, here's household i. It's the same thing. Household has utility function U_i over what i wants to assign of the first good for the first household all the way up to what i wants to assign for the n-th household of the m-th good. So this is a great, big, long vector running over all goods and all households, including, actually, the own terms, i i terms, as if assigning to your goods to themselves. Anyway.

Then we have-- if we're going to solve the Pareto problem, we have the typical constraint over here that if household i is consuming good k, and we sum up overall households, we'll get the total amount of good k being allocated like a resource constraint. But in addition, we would have that what household j experiences in terms of what i is eating of household-- of good k, whatever household i is doing of good k, it's impacting household j and all the other households.

So x_{ik} , this guy here, what i is really doing of k is what j experiences that i is doing of k over all the other households j. So this is an externality, you could think, positively or negatively inflicted on all the other households. Whatever household i is doing, it has consequences for all the other households.

Well anyway, it looks like a standard Pareto problem. We could maximize lambda-weighted sums of utilities and we would have these two resource constraints, 17 being conventional, but 18 being the new thing that we need. So let's think about shadow prices. This guy here, 17, is going to produce a Lagrange multiplier for every good k . So we can call that q_k .

This guy here, 18, is going to produce a Lagrange multiplier for every j because there-- j goes from 1 to n for every household i and for every good k . So there's lots of Lagrange multipliers and we can call them P_{ijk} for all these constraints. These are shadow prices. These are going to be the shadow prices of the rights. And with these shadow prices, we're going to be able to decentralize to get to Pareto-optimal competitive equilibrium.

If you looked at the first-order conditions, where does j_{ik} enter? It enters into the utility function of j . There's a λ_j for the Pareto weight of household j , because we'd be maximizing lambda-weighted sums of utilities in the Pareto problem. And where else does x_{jik} enter? It enters here.

So then we pick up that P_{ijk} Lagrange multiplier because everything else is linear. So that's where this is coming from. And then this second first-order comes from going back to look at the x_{ik} . Now the x_{ik} appears here in 17, so we pick up the conventional Lagrange multiplier there. But that x_{ik} also enters into a multitude of constraints here on 18, namely one for each of the j 's. So we have a P_{ijk} down here, summing over all those because it appears in each and every one of them, and that's this expression.

So there's our price system. We have a conventional price system q_k for each of the underlying goods, but in addition to that, we have these price of rights to assign consumption for other households. And I deleted the rest of the slides because this is the intuition. It's enough. You've already seen us take prices as Lagrange multipliers and decentralize. You can write down the problem of the household and the decentralized problem with these prices through the Lagrangian, take the derivatives, and you will get these expressions back.

So that's the way we proved the second welfare theorem, that any Pareto-optimal allocation can be achieved. Now there's a lot of externalities here. Where did they all go? The answer is we created a market in all the goods. Another way to view this on the production side is that when a household is deciding what to eat, it is aware of the utility consequences and also the expense in the budget.

But it's also getting revenue, because it's producing the j_{ik} goods that the other households j care about. So it's like as a firm, it's selling those for revenue, which is kind of like combining the examples that we've had so far. OK. So let me get to the failure of the first welfare theorem in a different way.

Instead of having missing markets, I'm going to have too many markets. An infinite number of goods. So you may remember how the proof went of the first welfare theorem, but just in case, I'll review it for you. So the first welfare theorem was that any competitive equilibrium would be Pareto-optimal. We did it by contradiction. We said no. Suppose it were not Pareto-optimal, then this baseline x allocation would be weakly dominated for some households, the baseline would be the star. The alternative is without the star.

So the dominating one weakly dominates for some households and strictly dominates for others relative to the baseline. And then we kind of did this algebra and decided that clearly for the households for which this dominates, it would not be affordable. We used a lemma that even for the households for which they would be indifferent, it costs at least as much. We then add it up over all the households and got that the value of expenses in the alternative, suppose Pareto-dominating allocation is strictly greater than the valuation of expenses in the baseline of the competitive equilibrium.

And then using that fact along with the other conditions, we ended up-- the bottom line of the proof was, this alternative allocation was not feasible, because we had this inequality strictly greater than 0 going on here. But the alternative Pareto-dominating allocation is supposed to be feasible, so this would be the contradiction.

Now it all seems pretty straightforward, but what happens if there's an infinite number of households? We start summing up over infinity here. It has no meaning that-- not necessarily having a meaning. If it's infinity, greater than infinity, well that's where the logic of the proof breaks down. You can't have one infinity being bigger than another one.

So this can happen in overlapping generations models where we have a first generation, a second generation, a third one, and it goes on without end. Everyone has a finite life, but the economy goes on forever. So how many goods are there? An infinite number of goods? How many households are there? An infinite number of households.

As an example, but this isn't exactly overlapping generations, it's very much in that spirit and it's easier to understand what goes wrong with the welfare theorem. Let's just outright have an infinite number of agents and only two goods, 1 and 2. Each household has an endowment of each of the two goods, 1 unit of the first and 1 unit of the second.

Each household i also has a common utility function, which is this log utility. Additive Cobb-Douglas. So let's call the price vector of the two goods p_1, p_2 , let the first would be the numeraire, second one have a relative price of p . You may remember that these Cobb-Douglas-like utility functions generate constant expenditures.

So p , the price of the second good, times the quantity of the second good is equal to $1/2$ of the endowment, because the weights here are equal on the two goods. Endowment, they have the first good at a price of 1, 1 unit of the second good at a price of p . So one half of the valuation of their endowment is spent on good 2, $1/2$ is spent on good 1. These are the demand equations at prices p .

What's a possible solution here? Well, let's try p equal to 1. If the price of the second good was equal to 1, then the solution to these equations are x_1 and x_2 is equal to 1. $1/2$ of 1 plus 1, et cetera. So that is a candidate-- autarky, actually-- equilibrium and price is p , they choose not to trade. If the welfare theorem were true, then that competitive equilibrium should be Pareto-optimal, but it's not.

Why not? OK, now I'm going to pull this rabbit out of my hat. It's kind of amazing. Let's have the second household transfer all of its endowment to the first one. The first one started out with 1, 1, now he gets to 2, 2. And you say, oh, this guy's 0, 0. No, no, no, no.

So household 2 gets compensated from household 3. Transfers his or her endowment to household 2. Household 2 is no worse than before. Household 1 is better off. If we had a finite number of households, the last guy would get screwed. But this goes on forever, so you can dominate. You're like pulling this good out of infinity.

Another way to say this is what's happening to the welfare theorem? As I said, what we need is, even with an infinite number of goods, that the valuation over all of them, over all the households is a real number. For the proof to go through, it has to dominate the valuation of the endowment.

We need to be able to evaluate the endowments in an infinite good economy, and in that equilibrium and autarky, the valuation of wealth is infinite. Why? Because we have the endowment of household 1 equal 1, 1. The price equal 1, 1. So this dot-product is 1, 2-- sorry. And we add up to overall the households and we get infinity.

So the valuation of wealth in that very special economy was infinite, and that's why the proof broke down. You could see that it broke down because we dominated the supposed Pareto-optimal autarky allocation by something better.

Now when we go to an infinite horizon model, something like this could happen, because we've got an infinite number of goods. Even without overlapping generations, if we imagine households live forever, like in a medieval storage economy or whatever-- well, that's not-- they died every 13 years. But with an infinite dimensional commodity space, it could be that the valuation of wealth is infinite and competitive equilibrium would not be Pareto-optimal.

However, typically those prices are declining at rate β , and β 's the discount rate and utility. So you have an infinite sum, but that infinite sum converges to something finite as we add more and more terms. So that's why this doesn't always happen. The necessary and sufficient condition is this one, 23.

So decision time. I'm going to skip the last five slides. I will tell you what the content of it is, which is a bit frustrating, because I don't have time to go over it. I anticipated this thinking about the lecture this morning, there's a lot of material here. But this is the starred paper on the reading list. So between these five slides and that paper, you should be able to understand.

The economics of this problem is, what do I get from owning Bitcoin? Hopefully I have the right to buy things from someone else, but only if they want the coin. So my problem is defined as in a network. If I'm a buyer, I want sellers. If I'm a seller, I want buyers. My utility is a consequence of who else is on the platform, whether it's a cryptocurrency platform or something else.

So there's a natural externality here that these households care about who else is populating the platform. You would think with an externality, competitive equilibrium is going to go bad and we're not going to get to a Pareto optimum, but there is a remedy, and that is to price the right to participate in the platform according to your agent utility type.

So different types pay different amounts. They buy the rights to participate on the platform, and that's going to internalize the externality, much in the same way that we were talking about with pollution rights and so on.

But anyway, I'm already a minute over. So I'm going to skip these slides. You can take a look at them. And if you have questions, maybe we can talk a little bit more about it next time. And that's all for today.