

[SOUND LOGO]

ROBERT TOWNSEND: Thank you all for coming. So we're entering the last one-third of the class. I guess the best way to talk about that would be schedule. So we're down to these-- you just had your second exam. And the last and final problem set was released a bit late over the weekend. But anyway, you should have it now.

So what we're going to do is the welfare theorems today. And then things like existence of competitive equilibria and Nash equilibria, Gorman aggregation, when we're allowed to look at that representative household or not. And thinking about whether models identify data. In fact, whether you could reject a model based on data.

So these first four lectures are all in general equilibrium but nevertheless do link to other themes that we've been covering in the class. And then, the last few lectures go back to the welfare theorems and talk about how they can fail but, in particular, also how you can construct market remedies for failures. And then, we end with monetary theory. So about, I guess, seven or so lectures left.

And then, if we look on the reading list, optimality and competitive equilibria. This was the welfare theorem lecture. And there's only one starred reading, which is Kreps, section 6.3. So again, I do recommend that you look at-- this is more textbook style. But it is written at a fairly straightforward level. And I think you will find it useful.

So let me start in on lecture 16. So the fundamental welfare theorems, there are two of them. That's going to be the focus today. First one, competitive equilibria or while raising equilibria are Pareto optimal.

We've defined both concepts earlier. We'll raise the equilibrium Pareto optimality. So that's review. I've mentioned that these welfare theorems are true, but we've never gone through exactly under what conditions and how the proofs go. That's the first welfare theorem.

The second welfare theorem, any Pareto-optimal allocation can be supported as a competitive equilibrium with transfers. And again, we'll go through the proofs. Actually, two kinds of proofs of the second welfare theorem. One, I would say is constructive. It's matching up first-order conditions. And the second one is using separating hyperplanes in convex spaces.

The second is more general than the first, in the sense that it applies to many environments. But the first constructive proof is providing a lot of intuition. So in some sense, when these welfare theorems hold, you can go back and forth between optimal allocations and competitive allocations.

A particular competitive equilibrium is Pareto optimal, but that doesn't mean that you couldn't contemplate others. It's associated with a particular distribution of wealth from private ownership. And if you want to get to other redistribute wealth and benefit, some at the expense of others, find another Pareto-optimal allocation, there's also a way to support any additional one like that. But it is going to require these kinds of transfers. So that's the overall theme for today.

So let me remind you, paste in some slides from earlier. This is village India, with income over time and over households. And you may remember this gradient, that in some sense, average income is going up as you go from the landless to the big landowners. That's an income, but on the same scale, so a bit harder to see, it is nevertheless present in consumption as well, with this upsloping.

So in some sense, this cross-sectional gradient reflects the lambda weights, so to speak, when we solve for the optimal allocation of risk sharing in this village India environment. We maximize the lambda weighted sum of utilities and intuitively as well as, in particular, for particular functional forms, the higher are the lambdas, the higher would be average consumption. So that gradient reflects the lambda, so to speak.

Now, we only looked at Pareto optimality when we did the optimal allocation of risk sharing. We didn't really link it up to competitive equilibria, but we're going to do that today. And in particular, we could ask whether the implicit lambdas in that gradient are related to wealth.

Think about the village achieving its allocation of resources to competitive markets and, in particular, through private ownership. Then, we should see a correlation between the lambda weights and other variables that represent wealth.

So I don't think I showed you this slide earlier. It is in the paper risk and insurance village India. And here, you can see a regression of that gradient, the lambdas in particular, that are backed out from the empirical work, from the optimal allocation of risk sharing. Those lambdas are correlated with things like landholding, where the value of the bull that they own and plow with and the value of inheritance.

So that would seem to suggest there is some correlation between wealth and the lambda weights, although it's not uniform and not true at all in Shirapur. So it's mostly apparent in only one of the three villages, in Aurepalle. And what you don't know from this table is that the ownership of bulls and so on moves around over time. So it's not so representative of wealth anyway.

So this might suggest that the allocation in those villages, even if achieved through competitive market forces, also includes some kind of redistribution from the upper-caste landowners to the poor.

So let's turn to the welfare theorems. First one, suppose we had a competitive allocation, denoted x^* , y^* with price p and wealth w . And actually, this slide treats it more generally. A competitive equilibrium with transfers. A special case, of course, of a competitive equilibrium would be no transfers and just private ownership.

But let's start with a competitive equilibrium with transfers. And if preferences are rational and locally nonsatiated, that allocation x^* , y^* is Pareto optimal or, for sure, competitive equilibrium or Pareto optimal.

Strikingly, you don't need much. All you need is essentially rational preferences and local nonsatiation. And intuitively, the reason you need so little is that we're given a lot to begin with.

We're given the existence of a competitive equilibrium with transfers. So we don't have to have all the assumptions that go with that that ensure that equilibrium. That's kind of taken as a given in the statement of the problem.

In particular, we don't have to talk about the consumption sets of the agents, other than they have to be constructed in a way that is consistent with local nonsatiation. We don't have to say anything about production sets, other than evidently they're nonempty because there's a y^* up here, although in principle, that could be 0. So very little is assumed.

So here's an illustration in the Edgeworth box of the first welfare theorem. We start out with the endowment, private ownership. No production. And we find prices such that each of the two agents here are maximizing. And along their budget line, it's the same line for each, in some sense.

Household 1 giving up Good 1, getting Good 2. Household 2 going in the opposite direction, giving up Good 2 and getting Good 1. And here's the equilibrium. And you can see how beautiful it is because the marginal rates of substitution match up with the price ratio.

And we've achieved the competitive equilibrium. So competitive equilibrium is Pareto optimal. Of course, it's not the only one. There are many, many others, as you can see. But the first welfare theorem just starts with a particular competitive equilibrium, in this case, private ownership.

Now, how do we prove the theorem? The picture makes it seem obvious. But we want the proof to be more general.

So if preferences are locally nonsatiated, then, in particular, if the x_i^* of the competitive equilibria is maximizing under Agent i 's preferences, maximizing in the sense that it's the utility maximizing object in h_i and i 's budget set, this is expenditures less than or equal to wealth, then for any other point x_i , which is at least as good in the preference as x_i^* , it must be the case that the valuation of that other allocation, evaluated price's p , is not less than, is greater than or equal to the valuation of expenditures at the competitive equilibrium.

That's kind a lot to take in. The key part of it is this implication, and I used to have the proof all written out. But it takes like six or seven lines. Let me just give you the intuition. Suppose it weren't true? Suppose that if $p x_i$ were strictly less than $p x_i^*$. Well, they're spending everything they have. So $p x_i^*$ is their wealth, and they could spend less if they chose to. If $p x_i$ were less than $p x_i^*$, that would be a point where they're spending less.

However, finally, we come back to the nonsatiation. Draw a little ball or circle around x_i . If you remember the statement, there has to be a bundle that's strictly preferred to x_i that's close by. Since it's close by, it is also spending less than x_i^* . So now, there's a preferred point because x_i is at least as good as x_i^* .

This other point is going to strictly dominate x_i and hence dominate in the preference at x_i^* . So there's a point in the carrier of the budget set which is strictly preferred to the chosen point x_i^* , which contradicts utility maximization. Even took me quite a while to say it, but it's quite intuitive when you think about it.

So this innocuous little Lemma is about to get used with the proof of the first welfare theorem, namely that a competitive equilibrium is Pareto optimal. Proof by contradiction. No, it's not. You say it is. But I'm going to assume it's not Pareto optimal.

Well, any point that's not Pareto optimal means, of course, you guys all remember, that there's a feasible allocation, which has the property that it's at least as good as the star competitive allocation for all i and strictly dominates for at least one agent, in this case, denoted i' .

So that's what we mean by X_i^* not being Pareto optimal. Now, clearly, for the guys, X_i^* , these bundles are strictly preferred to the one chosen. So they cannot be in the budget set. So we have a strict inequality in expenditures for those i^* guys.

But what about these guys that were only weakly indifferent? That could be a problem. But no, we've got it covered. By the previous Lemma, we already know that expenditures are not less.

So we have strictly greater, weakly greater expenditures of the alternative supposed dominating allocation with the given competitive allocation. Let's sum up the values over all the households I , and we're going to get a strict inequality. So if it were true that there exists an allocation, which Pareto dominates the competitive equilibrium, it would be associated at the equilibrium prices with the higher aggregate expenditure.

Then, the next step is to review the definition of profit maximization in a competitive equilibrium, meaning at y_j^* , the valuation of profits is not less than any other feasible-- technologically feasible vector y_j . And this is true for each firm j , hence true for all of them.

So we can sum up over them. And now we take this statement for firms with the statement on the previous slide for households and start using them both simultaneously. So from the previous slide, we have the summation over I $p \cdot x_i$ strictly dominates $p \cdot X_i^*$. But we also had a competitive equilibrium given to us as an initial condition, therefore feasible. Therefore, X_i^* is equal to $\bar{\omega} + y^*$.

And I should have pointed out. I think I just skipped over it too fast. It's easiest to think about prices as being strictly positive. So the only algebra going on here is to take dot products, take prices, put it outside the parentheses, then collect terms, put it back inside.

So $p \cdot x_i$, summation over I , but X_i is equal to $\bar{\omega} + y_j^*$. So if we put p outside and then multiplied it inside, we get price times the odd quantity object in each of the two terms.

Now, this last inequality is the statement of profit maximization. Holding the valuation of the aggregate endowment of fixed $p \cdot y_j^*$ is no less than $p \cdot y_j$. And so now, we again do this trick. We take prices, put in on the outside, and look at all terms inside the parentheses, and we've got the X_i over here and the $\bar{\omega}$ and the y_j here. So we collect those terms together, and we get this statement. But if this is greater than this, then when you subtract them, it's greater than 0.

So we finally arrive at the contradiction. Why? Because this supposedly Pareto-dominating point has to be feasible in the sense of the resource constraint, which means this has to be 0 for each component. But by the argument achieved, it can't be 0 for every component or this could not be greater than 0. So it is not true that it's feasible, and that's the contradiction. So it's a big proof by contradiction.

So something's happening along the way. We're going to go to the second welfare theorem now. Any Pareto-optimal allocation can be implemented as a competitive equilibrium. We can come back to this slide. But here's the picture.

So we pick another Pareto-optimal point in that Edgeworth box. Not this one that ran through the endowment, but something southwest of it. So here's the endowment. It's still there. But now we're targeting this other Pareto-optimal allocation. And the theorem says there's a way to redistribute wealth in such a way to achieve it.

Note the marginal rates of substitution all match up for Agent A and Agent B. They're equal to sum price. The price, p_1 over p_2 , is captured by the slope of that, quote, "budget line." But I really want to say wealth line because it is not a line that runs through the privately owned endowments.

So the previous slide, which I skipped, just says that those star allocations for A and B have the same slopes, which are their marginal rates of substitution. And the marginal rates of substitution match up since they're equal to each other. Just something called gamma. So that's the picture here. Margin rates of substitution are equal, equal to the slope of this wealth line, which has a slope of gamma.

Now, the point is that if they had started from the endowment, this thing, agent B would not have enough wealth to get to this Pareto-optimal allocation. Conceptually, imagining somehow the slopes were the same if we had drawn at those prices gamma line, a budget line going through the endowment, Agent B has to be further out. He has to have more wealth.

And where does that wealth come from? Well, Mr. A over here gets to make the sacrifice because he has gone southwest of where he would be in terms of the budget line through the endowment if we insisted that didn't go through the endowment.

So all of that is summarized on this slide. If we start out with that price ratio-- by the way, gamma 1, why? Because the slope of this thing is p_1 over p_2 . So we might as well set p_2 equal to 1 as the numerator because we're not going to get any more determination than that.

So it looks weird when you write it like this, but p equal p_1 equal gamma p_2 equal to 1 are the prices we need. Agents are going to maximize at those prices, depending on what their wealth is. But the allocation at those prices for Agent B costs more than the valuation of Agent B's endowment at the price gamma 1.

So what we propose to do to Agent B is to give them the difference, just give them enough extra wealth so that Agent B can afford to buy the star bundle, which requires more than the valuation of his wealth at his previous endowment. So we call this a transfer to B. And it's positive, as you could see from that diagram. And A makes the sacrifice.

There seems to be a typo here. It's probably-- anyway, you can see it. The wealth transfer is the same. It goes from B to A, and this is written oddly as if it were B's endowment and B's expenditure. But it's really referring to the point on the box here, which is simultaneously the point for B and for A.

And so we just move southwest, taxing A, and giving the tax to B as a physical transfer. So we found a price equilibrium with transfers that will support that target Pareto-optimal allocation. And it's possible to do any one of them.

So if you go back here, this was the private ownership Pareto optimum. There are all these other Pareto-optimal allocations. The theorem says, give me any Pareto-optimal allocation, and I can redistribute wealth in such a way as to support it as a competitive equilibrium with transfers. Who gets taxed and who gets the transfers depends on which optimum we have under consideration.

So this one takes more. We're going to have to start assuming stuff for the theorem to be-- for sufficient conditions for the theorem to be true. We start with the notation of our economy. Consumption sets for the agents, utility functions, endowments, production sets, shares in ownership of profits, which is a quick review of the notation you've seen before.

We're going to have to assume that these consumption sets essentially can be formed by combinations of quasi-concave functions. This first line is the most awkward. And if you wish, you could just think about these being the non-negative orthant, that any positive consumption bundle is fine.

I only put the most general version because I want to emphasize you need restrictions on things. Here's a statement about the production sets. Production sets Y_j admit a concave transformation function called F_j .

What the heck does that mean? Well, it means that for any vector Y in the production set, it can be represented through some function F_j , which is non-negative when evaluated at that Y .

An example, because this is really obscure, suppose we have an input z mapped through production function f to get output q . And we throw away some of the output. So we're not on the frontier, we're in the interior. Then if you rewrite this as $f(z) - q$, it's non-negative. So any input-output combination, z, q , when evaluated under F_j , which takes this form, is non-negative.

So it really is quite innocuous. But it's very convenient just to talk about this concave transformation. I should be emphasizing again, these are sufficient assumptions. Something about consumption sets being represented by quasi-concave, this thing being a concave transformation, preferences to get down to 3 are going to be concave and locally nonsatiated.

Locally nonsatiated we had before. Concave is new, which you know from one of the early lectures is equivalent to preferences being rational, convex, continuous, and nonsatiated.

And finally, we have an interior point so that none of these constraints or representations are binding at 0, namely non-negative, non-negative, and even in terms of the resource constraint through existing allocation, which doesn't use up all of the endowment plus production. That's a technical interior condition.

If we have all of these things for this economy, then we get the second welfare theorem stated in terms of the λ s. For any λ , which is a Pareto weight vector, there exists a price vector p in this L dimensional Euclidean space.

And our vector of wealth, one for each of the i agents such that x^*, y^*, p and w^* is a Walrasian equilibrium with transfers. That's a formal statement of the theorem.

Now, we're going to prove it. And again, in the proof, it's kind of easier to assume that those λ s are strictly positive. And it will generalize to some of them being 0. But it just gets in the way. So let's make it a little bit easier on ourselves. The proof is constructive and a bit tedious, but it is straightforward enough. I'll try to guide you.

Step 1, characterize the Pareto-optimal allocation. That's the given that we want to achieve with a competitive equilibrium of transfers. 2, characterize the conditions for a competitive equilibrium with transfer.

And then, 3, kind of match up those conditions and show that if we have the conditions being satisfied for the optimal allocation, they're going to be satisfied for the competitive allocation, if we choose properly.

This was the Pareto problem, written out now for the fully specified economy. To find a Pareto-optimal allocation, we're going to find the solution to maximizing a lambda-weighted sum of utilities, subject to resource feasibility and production constraints. Max, subject to 15 and 16. You've seen this many times before, in like six or seven applications already.

This is a full statement. So the so-called Pareto problem. Sometimes I call it the math problem. And you may remember that solutions to this problem generate Pareto-optimal allocations and that any Pareto-optimal allocation can be achieved as a solution to this problem. So we're not losing anything by restricting our attention to these Pareto problems.

Second step is to use the Lagrangian. And the reason I was emphasizing on the previous slide all the concavity and convexity is those, it's not just getting the necessary conditions for an optimum from the Lagrangian method, those conditions are also sufficient when we have this kind of concavity. I'll remind you momentarily.

So we've got production constraints here. We've got resource constraints. So we're looking for an allocation consumption in the consumption set, positivity from those concave transformations, and resource clearing.

Here's the review slide. We did this, the first time it came up was consumer optimization subject to a budget. And I said, oh, now we're going to learn a cool tool, the tool of constrained optimization.

So we maximize the function subject to a series of constraints. We form the Lagrangian, took derivatives, and looked at the additional first-order conditions, which I'm about to show you again in the context of the problem. So that was a review slide.

This is what we want to solve. And the Lagrangian for that is this problem. So we want to solve to maximize a lambda weighted sum of utilities by choice of an allocation X_i vector, with now we start writing out the L resource constraints. But we put a Lagrange multiplier γ_l in front of them, as we were instructed to do from the algorithm.

We also put on to the Lagrangian these positivity-- these F's that represent the production function frontier. There are j of those, and we come up with some notation for the Lagrange multiplier. And then, when we differentiate this whole system, say with respect to the L-th good for the l-th household, we want to set the Lagrangian derivative to 0, which is the same thing as $\lambda_i U_i' + X_i^*$.

And where else does that X_{il} enter? It enters over here in the resource constraint. So it picks up a gamma. And now, I start putting stars on the solution because the solution will satisfy these first-order conditions. So there are special allocations and special Lagrange multipliers because they are solutions.

Likewise, for the L-th good, for the J-th firm, set the derivative of the Lagrangian with 0. Where does that stuff happen? It happens here. So there's a picking up a gamma l. And over here, should we pick up a derivative.

So this would be positive, positive equal to 0. So one of those terms has to be negative. We subtract it off, and we get this expression. The gamma l is the derivative of the function F with respect to y_l . This is for the J-th firm, and we have this star on the Lagrange multiplier.

So this equation 10 is a set of conditions. We need something, a technical condition that there can be something that these things are all-- sorry. Not technical, that the constraints are all satisfied, and the Lagrange multipliers are not non-negative.

And then finally, this complementary slackness condition that the Lagrange multipliers times the respective constraints equals 0. Probably a lot to review, but it is in all this notation of the environment exactly what you've learned before.

The point is that equations 10 and these others represent a solution to the maximum Pareto problem. And they're not just necessary, they're sufficient because of all the concavity. And in particular, these utility functions are strictly increasing.

And I said earlier, everybody has a positive lambda. It's enough that at least one person has a positive lambda, then we're going to have to have the gammas being positive y .

That's this thing because here's the marginal utility, which is positive, a lambda, which is positive. So if the left-hand side is positive, the right-hand side is positive. So that's where the positive gammas are coming from.

So we know a little bit more about the solution. It's positive prices. And from this complementary slackness condition, if the prices are all positive, and this whole thing has to be 0, then the interior thing in brackets has to be 0. So the resource constraints are satisfied at equality, as is natural.

So let's summarize. We've characterized the Pareto-optimal allocations, x^* , y^* . Also denoted x^* , y^* at λ because we were given a particular Pareto-optimal allocation, hence a particular lambda weighted sum was used in the objective function. And it's the solution to 14, which is that Pareto problem.

So it's Pareto optimal. It's feasible, and there are these gammas, which are strictly positive for every good l . And these Lagrange multipliers on the production sets such that x^* , y^* , γ^* , and ϕ^* satisfy 19 and 21. 19 and 21 were, yeah, something happened in the process. This is 10. And 21 is the feasibility constraint. No. It's right here. 21.

So much for characterizing the Pareto-optimal allocation. Let's look at the second part, characterizing the equilibrium with transfers, just separately. Forget the first part for a minute. We have a given while raising equilibrium with transfers, meaning what?

Is x^* , y^* , p and w^* which must be consistent with consumer optimization, profit maximization, resource constraints being satisfied. And since it's an equilibrium with transfers, not necessarily private ownership, we have to make sure that the distribution of wealths are feasible, that the wealths add up to the valuation of the endowment and profits.

So let's just take this one at a time. Consumer op, firms op, and start look at the first-order conditions. Now, the third part of this is putting these pieces together. But at this point, we're lost, right? We don't remember where we started.

We started from the second welfare theorem. So we were given a Pareto-optimal allocation. And we want to show that it can be achieved as a price equilibrium with transfers. But we want to find the components x^* , y^* , p and w^* of the competitive equilibrium with transfers that's going to give us back exactly the Pareto-optimal allocations we started from.

So let's start with the households. We have a set of conditions for the Pareto-optimal allocation, and which are given because we started with the Pareto optimal-- now, you're not used to seeing it like this. We're usually writing λ_i times the marginal utility is equal to γ . Here it is.

All I did was divide by λ_i and put it on the right-hand side. That was part of the necessary and sufficient conditions for the Pareto problem. That's this.

So we're given this. And we want to find Lagrange multipliers μ and prices p so as to satisfy this optimizing first-order condition per household i in the competitive equilibrium with transfers.

Now, what would you guess? Well, we want-- this is true, and we want this. So let's guess that the price that we want is the γ , Lagrange multiplier on the resource constraint for good l . p_l equal γ_l .

And let's guess that the μ that we want to find is simply 1 over λ_i . That's down here summarized at the bottom. And the final guess would be that the wealth allocation that we need in the Pareto-optimal allocation is just the valuation of expenditures at that Pareto-optimal allocation.

I use my words carefully here. The valuation, though, is at the shadow prices γ_l . If you go back to the Pareto problem, these γ s were arbitrary Lagrange multipliers on the resource constraint. And we went to the math of all the first-order conditions.

But really, what it is, is how much the solution that shadow price reflects the incremental gain in the objective function to weakening the resource constraint by a teeny amount.

So it is like a shadow price. We often call the Lagrange multiplier shadow prices. It's really marginal utility pricing because we're incrementing the weighted sum of utilities a bit, hence like a derivative, as we weaken the constraint.

So it's natural really that these Lagrange multipliers λ_l for good l , should be shadow prices. And the actual competitive equilibrium prices in the corresponding welfare raising in equilibrium with transfers.

This thing down here, very simply written, is that μ_i is 1 over λ_i . Oh, who cares? No, no, no, no, no, no. Don't think of it that way. Think of those beautiful graphs of Kansas and so on at the beginning. OK?

So what this is saying is μ_i , which is the marginal utility of income, the marginal utility of wealth in the target competitive-- we're raising an equilibrium with transfers that we're trying to find.

The marginal utility of wealth at that location is inversely related to the λ weights in the Pareto problem. So the higher the λ weights, the lower would be the marginal utility of wealth.

Remember the first two slides, I told you that the gradient in the consumption profile had something to do with the lambda weights. And we could look in the data to see if it was related to wealth. So this may look upside down, but that's only because μ_i is the marginal utility of wealth. As wealth goes up, marginal utility goes down.

So as lambda goes up, and they're more favored in the Pareto problem, the marginal utility of wealth is going down. And how does that happen? It happens because they get more wealth. So there's kind of a 1 to 1 mapping, a monotone mapping, between the lambdas in the Pareto problem and the wealth that corresponds to the competitive equilibrium with transfers. And here it is falling out in a very elegant way.

And this other thing at the bottom, that W_i^* is the valuation of expenditures. Again, that's pretty much what we were doing in the box. We took a target allocation here, that it wasn't a competitive equilibrium with private ownership.

But we could tax A and transfer to B so these stars here represented the target Pareto-optimal allocation. And we're valuating those expenditures at the prices corresponding to the budget line.

So you can start them out with wealth being at carefully chosen prices, the valuation of the star Pareto-optimal allocation. And that's this part here at the bottom, that the wealth we give to Agent I is just the valuation of the Pareto-optimal allocation. Now also is utility maximizing allocation at the prices γ , γ^* .

We did all that for consumption. Something else happens automatically along the way, namely when you look at those first-order conditions, we can define the marginal rate of substitution of good I for good k for agent I. The lambdas cancel out, right? It's λ_i for a household i for any good I.

So if we take the derivative of utility with respect to the two goods, in this case I and k, the lambda is in the numerator and denominator it cancels out. So this is the marginal rate of substitution corresponding to these optimizing first-order conditions for household I's, obviously equal to-- the I's, as I said, are canceling out, equal to the ratio of prices, back to some other household h.

So this is the math of what you already saw in the picture that I've shown you three times, the marginal rates of substitution are equated across households and equal to the common price ratio. All of that coming from the first-order conditions that we've now formally derived.

One more step, firms. So we have profit maximization. Oh, are you lost again? I say that a bit tongue in cheek because everything is equivalent to everything else. And it's easy to get turned around and forget our starting point.

So we started with a Pareto-optimal allocation. We looked at the first-order conditions, which are necessary and sufficient. And we got this thing. So this is true.

Now, we want to find the conditions for profit maximization in a competitive equilibrium with transfers. So we want this equation to be satisfied. As long as we can choose the gammas and the piece appropriately, we'll be done, and it's easy to choose them. P , we already pinned down already. $p_{sub I}$.

The price of good l is the shadow price derived previously γ_l . And it just remains to figure out what this δ_l is. Can we find that Lagrange multiplier somehow? Searching? Sure. Let's let it equal to the associated Lagrange multiplier and the solution to the Pareto problem.

And again, because all these-- if you can satisfy this because it's necessary and sufficient, then we have profit max. So we just pick the δ and the p to be what-- obviously, it would be true so that we have profit maximization.

So we're almost there. We have a solution. We started with a target Pareto-optimal allocation. We have a solution to the household problem in the competitive equilibrium with transfers. We have a solution to the first problem. So what do we know about feasibility, and what do we know about the allocation of wealth?

So we know what householdized wealth has to be, as I've been saying. It's just the valuation of the assignment, the X_i^* assignment in the Pareto problem at the prices γ equal to p . So that's the wealth of household i . Let's sum them up and make sure it's feasible.

So we sum up the valuation of the target at prices p . But now we do this trick. Bring p outside. It's common. I suppose it's strictly positive in all the components too, just to make it easy. Then, this summation of-- because of we started with a Pareto optimum, so it had to be resource feasible.

So now, we use that fact, that the summation over i X_i^* must be equal to the summation over i of the endowment plus output. So demand is equal to supply. Demand is equal to endowments plus production.

And then, you can rewrite the $p \cdot y_j$'s to be consistent with the claim on those profits under these ownership shares, θ_{ij} . It probably would have been clearer if p were also still multiplying each part. But it's there. It's just outside the parentheses.

And then, we group-- finally, we do put the p 's on the inside. And we have what we want because this is total wealth in the economy. It's the total value. It's the valuation of the aggregate endowment plus profits. And it adds up to the sum of the wealths that we need, according to the solutions to the consumer maximization problem. So now we're done.

This last slide is always a bit mysterious. It's like, well, why do we have to do anything else? Well, again, just remember the definition of a competitive equilibrium with transfers. Because it's transfers, it's specified not in terms of private ownership but in terms of wealth. So whatever those wealths are, they have to add up to the total. That's what this slide is all about.

So that's what we call the constructive proof, is just using finite dimensional Euclidean spaces, putting in a lot of concavity, convexity, et cetera, then sort of matching up first-order conditions. But the theorem holds more generally.

The second welfare theorem is true, even without making special assumptions about the commodity space and so on. A step along the way is something called the quasi-equilibrium with transfers, quasi meaning it's not quite fully consumer maximization. But it's going to be close.

So here's the definition. An allocation, x^* , y^* and a price vector p , constitute a price quasi-equilibrium with transfers if we can find wealth levels that add up to the total wealth if each firm J is profit maximizing. So far so good. Nothing different. Here's the difference. Household i , this condition has to be true.

If x_i^* is the allocation we want in the definition of the equilibrium with transfer price quasi-equilibrium with transfers, if you start with that x_i^* and there were an allocation that's strictly better for household i , then it costs at least as much as the wealth assignment in the price quasi-equilibrium with transfers. It's kind of a mouthful.

One way to put this, though, is that w_i^* is the valuation of expenditures at the equilibrium, at the equilibrium x_i^* . So then this says $p x_i$, the valuation of this arbitrary x_i , which is strictly preferred, the valuation is not less than the valuation of $p x_i^*$.

That means that $p x_i^*$ costs no more or is cost minimizing over the upper contour set of allocations that are associated with strictly more utility. So you've seen this before, way back when.

When we did the consumer utility maximization problem, I pointed out that this is kind of a dual thing. Instead of maximizing utility subject to budgets, we minimized the budget necessary to achieve a given level of utility.

In fact, we did a lot with that. We did Hicksian substitutions versus Walrasian substitutions and so on. We had a whole lecture on it. So here it is again, except so instead of assuming utility maximization, we're going to assume cost minimization. That's the only difference between a regular equilibrium with transfers and this quasi-equilibrium with transfers. The feasibility of the allocation is the same.

So I'll show you the theorems now. But just to say, when we do the second welfare theorem, we're going to say that it's a price quasi-equilibrium with transfers. And then, given the structure of the economy, you may need a little bit more to establish that cost minimization corresponds with utility maximization.

So suppose that the production sets are strict or convex, the preference relations are rational convex and satisfy local nonsatiation, then any Pareto-optimal allocation has a price p such that x^* , y^* , p is a price quasi-equilibrium with transfers.

So it's looking pretty similar, but we're going to prove this in a different way. And it's going to be a proof that works for quite arbitrary spaces. And the proof that's going to work has to do with separating hyperplanes.

So let me remind you what that was about. Actually, this is not quite the right slide. This was a review slide about a supporting hyperplane. The better slide was, we had two convex sets that were tangent at a point. And then we were able to draw this hyperplane that would be tangent to the common slope of the two convex sets.

So all you needed for the existence of supporting and separating hyperplanes is a presumption that the sets are convex. Here's a picture of it. We've seen versions of this before. Here's one production set. Here's another production set. Here's the endowment somewhere here, I think it was.

And so we add up the production sets and the endowments, and we get this sort of outer frontier. The point being, it's defining a convex set. Things on the frontier but also interior. Any linear combinations of things in the interior is also in the interior.

What's this other monster here? This thing looks like an indifference curve. But, in fact, we have many, many households. So what the heck is it? Well, it's this thing. It's the upper contour set. It's the set of allocations, which are weakly preferred to the allocation x^* and summed up over all the households.

So that's this thing that looks like an indifference curve. But because we have convexity, this upper contour set is also convex. So now, we have two convex sets. We have a supporting hyperplane for both of them, separating supporting hyperplane. And we're going to define the price system.

So this is kind of a graphical way to think about that constructive proof without taking first-order conditions and matching up Lagrange multipliers. We're defining using the convexity and the separating supporting hyperplane theorems to get the prices.

So in 1954, Debreu worked all this out. He called his equilibrium a valuation equilibrium, as distinct from a competitive equilibrium. Pareto optimality was defined the same way.

The environment is similar. We have a finite number of households getting allocations in their consumption sets. But when it says the consumption set is a subset of this L , it does no longer mean a finite dimensional Euclidean space.

It's an arbitrary linear space, conveniently denoted L for linear. Preferences have a complete ordering. That's rationality. Producers J have a production set, choosing objects in the set. Also a subset of the linear space.

Market clearing looks pretty similar. $x = y + \sum e_i$ or $x - y = \sum e_i$, where x is the summation over households, y is the summation over producers, and e_i is the endowment. So we actually call a state of the economy something that's attainable being in the production sets, consumption sets, and satisfying the resource constraint.

So then, we get to the theorems. First and second welfare theorems are similar, although the assumptions are a little bit different. Valuation equilibrium, i.e. competitive equilibrium, is Pareto optimal. What does the price system look like? In these general linear spaces, it's just called a linear functional, meaning that it's retaining additivity.

If you take the valuation of a weighted average of two bundles, it's the weighted average of the valuation of the two bundles. Seems quite natural. Obviously, for finite dimensional Euclidean spaces and the dot product, this holds automatically. This is a generalization of that idea.

So a valuation equilibrium is attainable, maximizes utility, subject to the budget, and maximizes profits, where the price system is this valuation function v . This is literally utility maximization because it says that if it's in the budget, it is not preferred to x_i . If that weren't true, if you had a point in the budget that were strictly preferred to the allocation x_i^0 , then it would be chosen. So that would contradict utility maximization.

The second welfare theorem of valuation equilibrium is a Pareto optimum welfare theorem. What we need, we need convexity of the consumption sets and a little bit more about convexity of preferences. We did not need this previously. When we had finite dimensional Euclidean spaces, I made a big deal of saying all we needed was local nonsatiation.

We need no satiation here. We also need a little bit more about convexity of consumption sets and preferences. So you're trying for something bigger, something that works for a very, very large number of economic environments, you need to assume a little bit more.

The second welfare theorem, any Pareto optimum can be achieved as a valuation equilibrium. So again, more things about consumption sets. When your combinations of points in the set are also in the set, sets closed. Production sets, when aggregated up, the aggregated production set is convex.

And here, either L is finite dimensional or if infinite or other dimensional, Y has an interior point. I know this is really, really abstract. I'm going to show you a couple of applications momentarily.

And the theorem is, assuming the 3, 4, and 5, every Pareto-optimal allocation, which is not a satiation point, is associated with a nontrivial continuous linear function, where we have a condition for households and a condition for firms. So here, the condition for households is that if X_i , we have the Pareto optimal allocation Pareto X_{i0} .

We have an allocation that is strictly weakly preferred to X_{i0} and is in the consumption set, then this implies the valuation of that other allocation X_i is not less than the valuation of X_{i0} . This corresponds to cost minimization. And again, just reverse it. If the valuation were less than X_{i0} , then that other point is, say, better, maybe indifferent.

But because we have local nonsatiation, just like the Lemma, we could construct something close by which is strictly better and in the budget set for X_{i0} . So that would be a contradiction. So this is the cost minimization part. This is why it's going to be a quasi-equilibrium, not a full equilibrium.

So I feel like it took a lot of work to get through all of this. The main takeaways are, there's a way to do things more generally with separating and supporting hyperplanes, as long as we have sufficient convexity. I took the time to take you through Debreu's article.

It's really very similar, apart from changes in notation, to what we did with the finite dimensional Euclidean spaces. But it applies to other economies, in particular, competitive equilibria in the space of contracts, including the possibility of randomized allocations that we've used before.

So I think I mentioned contracts with options, like CDs where you could withdraw early or pensions where you could decide to take a lump sum or take an annuity. I gave you those as examples of contracts with options. And then we had a couple of lectures talking about contracts.

I'm sparing you working out through all of that notation how the contracts with lotteries and indeed without grids, so measures, can be mapped into Debreu's valuation equilibrium, and hence we get the welfare theorems. OK. So that's all for today.