

[SQUEAKING]

[RUSTLING]

[CLICKING]

ROBERT
TOWNSEND: As usual, let's find out where we are. We did the first two sections of general equilibrium on welfare and existence. And now, we're going to do Gorman aggregation today, and then finish this segment with identification. We have two lectures to finish out the class.

So that's where we are. In terms of the reading list for today, there's only one thing there, and it's the place from where I drew a lot of the material for the welfare theorem for the Gorman aggregation and my colleague Daron. It is starred, but it is also pretty extensively covered today as well. Then we have the study guide. So last time, we did concepts of existence in various distinct ways.

There aren't too many review questions here. I didn't ask the first set of them. But we get down to Negishi. So how does Negishi use the first and second welfare theorems to generate a mapping in lambda weights, the fixed point of which is a competitive equilibrium? And since I was calling on everyone last time, let me just take volunteers. Hopefully, there are volunteers.

AUDIENCE: I was--

ROBERT
TOWNSEND: Go on.

AUDIENCE: You go ahead.

AUDIENCE: I was about to say that I was a little unclear on the proof overall, but I remember that we use the fact that for any lambda, the corresponding period of optimal allocation would be in equilibrium with transfers by the second welfare theorem, and we use that to prove Negishi's theorem.

ROBERT
TOWNSEND: Yeah, that's right. We used to have the second one heavily. The first one is that any competitive equilibrium is Pareto optimal, and hence, it corresponds to some lambda weights in the Pareto problem. So Negishi's mapping is to find a way to find the particular special lambdas of the competitive equilibrium with private ownership that is the fixed point of his mapping.

I realized after class, in part talking with one of you who came to office hours, that I never-- I think I neglected to say the obvious. What I did say is to find the lambda such that the valuation of expenditures under the Pareto of optimum uses using the second welfare theorem is equal to the valuation of the endowment plus shares on profits that corresponds to the competitive equilibrium.

But what I think I neglected to say was-- maybe I did say it. Bears repeating. That if you initially guess a lambda which is high than the value for a particular household, then the valuation of that household's expenditures will be large relative to the valuation of wealth in the private ownership economy.

So sort of intuitively, then, you lower that lambda weight. That links up to something we learned when we did this the second welfare theorem, which is, those lambdas are inversely proportional to the marginal utility of wealth. So if the lambda is too high, the guy has too much-- the evaluation of expenditures is too high. So the marginal utility of wealth with concave utility is low.

So you want to move in the direction of increasing the marginal utility of wealth, which effectively removes wealth by lowering lambda. So anyway, thank you. Next question. Define a mixed-strategy Nash equilibrium.

AUDIENCE: I think I can do this one. Like, a mixed-strategy Nash equilibrium is-- it's a type of Nash equilibrium where at least some of the players are-- choose some of the options that are random? So there's some weighted probabilities where you can have some chance of doing option A, some chance of doing option B, and so on.

And in order for it to work, the player who is doing mixed stuff has to be indifferent from all of the things with non-zero between all of the different options with non-zero probabilities, and it also has to be in such a way that, again, no one has any incentive to deviate from the strategy.

ROBERT TOWNSEND: Right. The last thing, especially, is key. So the mixed-strategy Nash equilibrium is taking-- each player takes as given a configuration of mixed strategies of the others-- sort of star strategy, so to speak-- and then feels free to choose the best one for his or herself. But that will end up being exactly the star strategy to begin with. So it's, like, self-fulfilling in the space of strategies. Thank you. That was a good answer.

And then the last one. How do we give an incentive for traders to live up to their promises and achieve over raising an equilibrium? So I did this in the context of the market for trades in US treasuries in New York. Which is a trade agreed to by treasury for liquidity.

I mentioned that, in fact, there were trade fails where they don't live up to their agreement because the party selling the treasury doesn't own it, can't transfer it at the time, when a settlement is due or the person who promised to pay liquidity doesn't have sufficient liquidity.

And then I mapped that into the Walrasian equilibrium framework. So let me ask it again. How do we give an incentive for trades-- trade-- for traders, I guess it should be, to live up to their promise and achieve a Walrasian equilibrium? Anybody remember?

AUDIENCE: I remember there was a penalty for liquidity, right? If the creditor was negative, then it will have this utility. And so you're in-- so in the end, in order to achieve the [? variation ?] allocation, we need to let all of the downwards-- or there was a-- we need to let the [? operator ?] equal to 0 so that there was no exercise demand, no-- yeah. I mean--

ROBERT Yeah. Yeah, that's good.

TOWNSEND:

AUDIENCE: We should we want to promise that if the endowment is equal to the allocations, the extent of the allocation.

ROBERT TOWNSEND: Right. That's good. So the problem is, they're submitting bids, and we're looking for a Nash equilibrium. The bids consist of orders to buy and also orders to sell. We're shooting for the Walrasian equilibrium, in particular, the budget constraints.

So we want, for each trader, that the valuation of what they're buying is at least matched by the valuation of what they're selling. In other words, they have sufficient liquidity. And if we didn't put in a penalty, everybody would overspend, because why not? So yeah, so then the penalty that you mentioned-- so beta is negative.

When beta is negative, that means they're spending more than they have in revenue. That's a bad thing. So we effectively make the penalty beta-- sorry, I can't remember the notation. I think it was lambda times beta-- sufficiently high so then they do the right trade-off. They look at how much they would get in utility from buying this stuff directly and weigh that against the disutility of the penalty.

The magnitude of the penalty had to do with the marginal utility of income, that Lagrange multiplier we keep focusing on in the competitive equilibrium and the budget constraint. So the margin utility of income represents the incremental gain from having epsilon more liquidity. And they're tempted to spend more than they have to go in that epsilon direction, but the penalty is large enough so as to dominate that gain, and they don't deviate.

Well, in earmarking this question, I recognized that we rushed a bit at the end. In fact, I went over the class time by a couple of minutes. So anyway, hopefully, that was a useful review. Any other questions from last time? OK.

So now, I want to get to aggregation. This is kind of an exciting lecture in the sense that it goes back and reviews some of the things we've been doing before-- certainly, the basics of consumer theory with utility maximization, some things about Lagrangians and envelope theorems. It also goes back to the first application we have had of consumers maximizing utility subject to budgets, and what happens when income changes, and so on.

So we already talked about angles, curves, and Giffen goods, and so on. And so you have the micro background already. So you can judge the plausibility of the assumptions we're about to make in order to allow the aggregation. And the aggregation was used when we did, like, free trade. We would talk about a representative consumer, not just Robinson Crusoe, one man on the island, but someone representing the entire US population. So that is a big leap of faith.

Today, we'll spell out what it takes to do that. So this is going to be Gorman aggregation, and we're going to do it in the positive sense of predicting what will happen and in the normative sense of making statements about consumer welfare. And then we're going to have indirect utility functions, Gorman Polar forms, and the rest I already mentioned.

So some micro, but mostly macro economists, assume the representative consumer or representative agents. It kind of comes naturally, because they're worrying about aggregates, not distributional aspects. So that gives them and us a tractable way of understanding market demand-- and supply, for that matter-- and allows pretty easy welfare analysis because there's only one person. So those are all the advantages, but we'll see how realistic it is.

So here's the utility maximization problem. Maximizing-- we'll suppress the labels here for households i and h and so on. So this household has a utility function represented by u of x , where x is in the consumption set and the valuation of expenditures at prices p cannot exceed wealth w .

So we call the solution to this problem, after maxing it out-- the achieved maximized utility is called the value. And this function v is called the indirect utility function, v for value. So it takes as given the parameters that the household takes as given, namely, the price vector p and wealth w . So it is the maximal attainable utility in the budget set as a function of those parameters.

And the bottom line here-- you can just see, if x of p and w is the maximizer, we just stick it back into the utility function, and that would be the optimized value. So there are very useful properties of this indirect utility function, and they are reviewed on this slide.

Let's assume the consumption set is a subset of finite dimensional, L -dimensional Euclidean space. Utility functions are continuous, representing a preference relation that is rational and continuous. The indirect utility just defined. First of all, this v as a function of p and w is homogeneous of degree 0 in p and w , because as you already know, if you, for example, double prices and wealth, it doesn't change the picture of consumer maximization and you get the same solution.

It's like, there's no money illusion here. If price is double, but wealth loss also doubles, you're just as well off as you were before, and not worse off. This indirect utility function has intuitive derivatives with respect to p and w . Increasing p makes you worse off, or at least not better off. Increasing w makes you better off. w is wealth. Prices are what you have to pay.

In fact, that indirect utility will be strictly increasing in w when we have local non-satiation, which is what we're always going to be assuming here. The indirect utility function is quasi-convex, not concave-- convex in p and w . So again, when we did consumer theory, we had concave utility.

Upper contour sets where quasi-convex, but the utility function was concave. This maximized utility over the budget is convex, or at least quasi-convex. Continuous, which is easy to assume, and differentiable. And when the indirect utility function is differentiable, we can use it to recover demand.

And the proof is pretty easy, but here, it looks kind of amazing, I suppose, which is if we want the maximized quantity of good L for this consumer facing price vector p and wealth w , the answer will be take the derivative of the indirect utility function with respect to the L -th price, and divide by the derivative of the indirect utility function with respect to wealth, and stick a negative sign in front of all of that.

So here's a little proof of it. We have the Lagrangian, which represents a utility maximization problem, with 0 being the shadow price of the budget and the L -- sorry, are λ_0 , and λ_L being potential shadow prices on non-negativity constraints. So we're going to use the envelope theorem, which to remind you, states that you can take a derivative with respect to the parameters of the optimization problem, in this case, p and w .

And if you take the derivatives at the optimizing solution to get sort of the increment in utility associated with changing those parameters. So the increment in utility from solving the max problem at those parameters is also the derivative of the Lagrangian with respect to those parameters.

So there's, like, two steps there-- how to represent the maximum value, namely, as a solution to the Lagrangian, and also the envelope theorem that tells you-- reminds you that the derivative of the optimized solution with respect to the parameter is, like, differentiating at the maximizing quantities. So too many words.

So the first derivative here in equation 3, the derivative of the value function with respect to a particular price p_l is the derivative of the Lagrangian with respect to-- at price p_l evaluated at the maximum, namely, x_{p_l} and w is the demand. And going back upstairs here, derivative of the Lagrangian with respect to p_l , where does it appear? It only appears here in the budget.

So you get a λ_0 , and p is linear in x . So you get the x_{p_l} back, and that's this object. Likewise, the derivative with the indirect utility with respect to w , derivative of the Lagrangian with respect to w wealth, evaluated at the maximizing solution. There's only one place that w appears. It's up here. So you just get λ_0 . So λ_0 is now in 4 and in 3.

Substitute the expression in 4 into 3, and you get $dv dp_l$ equal to-- well, it's all written out here. Just basically taking the λ_0 , looking at the left-hand side, sticking it in here, and then repeating the x_{p_l} . So now, solve this equation for x_{p_l} and you get what we wanted, namely, the ratio of the derivative with respect to p_l divided by-- with respect to w with the minus sign in front of it. So we get what we wanted. Questions?

So that's an interesting problem of the indirect utilities. Now, we're ready to define a particular version of the indirect utility function called Gorman Polar form, or Gorman form. This doesn't always happen, but it can happen. For particular utility functions, the indirect utility function in p and w has an intercept and a slope and. The intercept is a function of vector p .

The slope b with respect to w is a function of vector p , and it's linear in w . And I'll show you an example right now, because this seems a bit mysterious. So let's go back to our favorite utility function, this sort of Cobb Douglas, with two goods, x_1 and x_2 , raised to the power α and $1 - \alpha$. You may remember that these α s and $1 - \alpha$ represent the expenditure share in the optimized solution.

So p_1 times x_1 is α times w . So this household spends α fraction of its wealth, w , on good 1, and spends $1 - \alpha$ times w on good 2. I'm just shifting these prices over here to the left-hand side so it's not just quantity. It's quantity times price or value. So with each Cobb Douglas, it's constant expenditure shares. More on that momentarily.

And now, the definition of indirect utility says, just substitute in the maximized value. So we take $U(x_1)$, but use this expression for x_1 to the α x_2 to the $1 - \alpha$. That's this business. And when you solve it out, you get of a of p , looking for it to be in this form. a of p is 0, and b of p is this coefficient, [INAUDIBLE] $\alpha^{1-\alpha} p_1^\alpha p_2^{1-\alpha}$ to these powers α and $1 - \alpha$.

That's b of p . It depends on p and on nothing else. And it is indeed linear in w . The linearity in w just comes from the property that $\alpha + 1 - \alpha = 1$. w the α times w to the $1 - \alpha$. Just add up the power coefficients. So this is all, so far, for a single household, hence neglecting to put in notation h or i , or whatever label the household is.

However, if this Cobb Douglas function were the same over all of the households, then we have a Gorman Polar form, where all of the households have the same intercept, a of p , and the same slope term, b of p .

Likewise, if it were to be the case that the alpha coefficient varies across over the households, then this BSP is not going to be common over all of the households. It's going to depend on particular preferences. So maybe that's illustrative just how tenuous, already, the assumption is going to be. But it is going to be very powerful if it's true.

So let's go at it and talk about aggregate demand. So we have an economy with I households, I agents. Their wealths are w_1 through w_I . A wealth profile. Let's define \bar{w} to be the aggregate total wealth, just by summing up all of the individual wealths. And retain this notation that x_I of p and w_I is the demand function for agent I .

The aggregate demand, \bar{x} , by definition, will be the sum of the individual demands. And we write that notationally as \bar{x} of the vector-- the parameters of the problem, namely, the common price vector and individual wealths. So that's the definition of aggregate demand. Nothing so far other than the definition.

But the first question about aggregation is, can we find an aggregate demand function that is special in that it only depends on total wealth in the population, not how the wealth is distributed in the population? So the idea here is this \bar{x} of L -- which, again, was just defined. But now, for the L -th good, \bar{x} of L at p and \bar{w} is the aggregate demand.

We're looking for this function, capital X , now, of L -- capital, I guess, for big aggregate, which is a function of p , of course, but now \bar{w} . So here's the content of it. Or looking at it this way, vector p , yes, but the wealth component in the aggregate demand depends on the sum of individual wealths only. It doesn't depend on the distribution.

On this slide, wealth enter each agent's demand function. And there's no reason why it should be representable to depend only on the aggregate. But we're going to look for specifications of utility where that's true. And we haven't even asked whether that aggregate demand function, that capital X , comes from some underlying utility function as a solution.

We're just looking for whether it depends only on aggregate wealth, or rather, might depend on the distribution. But this is a necessary condition for there to be a representative household that's maximizing some pseudo-utility function that's going to represent aggregate demand. So we're kind of doing one thing at a time.

So what is a necessary condition for us to be able to aggregate up like that? And it is that the marginal impact of wealth changes on demands, say, for the L -th good, is the same over all of the households, and that is also true for all of the goods. So whatever good L you pick, whatever pairwise households I and J you pick, they need each to have the same income effect of how demand changes with wealth.

And here's the picture, which you've kind of seen before. These are linear wealth expansion paths. We're holding prices fixed of the two goods, in this case. And for consumer-- one of these consumers, we're tracing out the optimized solution as we change wealth-- optimizing in the sense that the household has achieved maximal utility for any given wealth characterized by a tangency between the indifference curve and the budget line.

The other households up here also has a linear expansion path. Now, it's true that this household-- I , it is-- has higher levels of demand for good 2 relative to good 1 than is the case for household J . So they're not identical. But when you think about what happens when you redistribute wealth with these parallel linear expansion paths, the derivatives are the same.

We could, for example, take some wealth away from household I. We would move southwest. We'd have the change in x_1 , x_2 , and the change in x_1 little triangle over there, and give that wealth to household J, who would be moving northeast. But again, because these lines are parallel, we're going to have exactly the same changes [INAUDIBLE].

So the levels can be different, but the derivatives are the same. And that picture represents this condition 10, that the income effects are the same over all of the households. And that's a necessary condition for only aggregate wealth to matter. Individual demands are moving as you redistribute the wealth. But what one guy gives up is what the other guy gains. So the total demand isn't moving. Who gets what is moving, but not the total demand, as we redistribute the wealth.

The next step is to find a utility function that's going to give that aggregate demand. We're going to call this guy the positive representative household. So let's have the consumption set be a subset of this L-dimension Euclidean space. The individual demands, x_i of vector p and w_i for households i over all of them to be the individual demands.

We're looking for a preference ordering with a star on it on the aggregate consumption set, and a particular utility function, u^* , that represents this preference ordering. And that object, consumption set and the preference ordering in particular, is said to represent or correspond with a positive representative household if and only if the aggregate demand, that capital X of p and w bar, comes from the utility maximization problem where the utility function is of the form u^* .

So again, we're used to utility maximization. u^* is some utility function. We maximize it. Oh, subject to the budget, what is the budget? It's like giving this, quote, "person," this representative household, the entire wealth of the economy, w bar. And this X thing, in case you've forgotten, this capital X of p and w bar, was the object way back here that we were looking for, the aggregate demand as a function only of the sum of the wealths, or aggregate wealth.

So under what conditions can we find this amazing u^* ? Yep, Gorman Polar form. So a necessary and sufficient condition to get that equation tend to be true is that everybody's preferences can be represented by 16, namely, now with the i 's to be explicit, the indirect utility function for household or agent i has this intercept and slope.

The intercept term can actually depend on i . We do allow that degree of heterogeneity. But the slope with respect to w and w_i for household i does not depend on i . Common slope terms. So again, to write it out, the way the indirect utility functions for i and j respond to wealth is the same. Namely, it's b_i , which is not quite the picture we were looking at, but it's getting close, as you can imagine.

So if 16 is true over all households i , then we can find our amazing u^* , or in particular, the indirect utility function corresponding to u^* . And it's going to look a lot like the individual indirect utility functions in the sense that it's written down as a_i times b of w .

The a^* for the representative consumer is the sum of the a_i of p s over i . So we just add them up. And the b^* is the common b of p over all of the households i . So then this guy, represented by the indirect utility function, is the positive representative household.

So it's an if and only thing, but let's just go one way. If and only if means if they have these polar forms, then it's Gorman-agreeable and we have a representative consumer. And also, if we have a representative consumer, then the underlying households have to have utility represented by 16.

So let's just go one way. Suppose that preferences have indirect utility functions that satisfy 16. We'll go from the micro to the macro, so to speak. Now, we use Roy's identity. You were probably wondering why we ever had that. So demand of household i for good l at price vector p and wealth w_i is just this ratio of derivatives with the minus sign out front.

And now, we use the particular forms, namely, the derivative of the indirect utility function with respect to w_i is namely just b of p , because that's that coefficient there. And the derivative of the indirect utility function with respect to p is what? Well, go back here. It's indirect utility function for household i . Take derivatives with respect to the intercept and the slope term with respect to p .

So that's what's written out here, the derivative of the intercept and slope terms with respect to p . And of course, that w is hanging on there as a coefficient. So this must be true. Now, having derived the demands of household i , we can take the derivative with respect to wealth. And it's just going to appear here. So we get this times this as the way demand is changing.

Well, one thing you can see, on the right-hand side here, there's no i . So the change in demands with respect to wealth are the same across all of the households, and equal to this common term on the right-hand side. So this gives us that with the indirect utility function, we have these constant income effects over all of the households.

Now, the second thing we want to show is not just that, but that we can find this amazing representative consumer. So we want-- when we maximize whatever this v star is with respect to-- when we get this maximized solution for this star representative consumer and aggregate wealth w , it should represent the sum of the individual demands.

So you have this expression already for-- let me remind you, the v star, the candidate, is 17. So we're going to prove that it works. Take v star indirect utility to be this a star, b star, w bar term. Then by Roy's identity, that's just another indirect utility function. So the demand that comes from it at the aggregate wealth w bar and price vector p would just be the ratio of these marginal effects with a negative sign out in front of it.

And of course, we know the form of v star. So at an aggregated level, it's very much like before, 1 over b_p , because that's just the wealth coefficient. And p enters into the intercept and slope term in this aggregated equation-- aggregate indirect utility 17. So we get this.

Remember that the a star is the sum of the a_i 's by construction. So now, that is all written out very explicitly. Nothing going on with wealth. But now, we get to just take this summation and pull it outside. So we have the demand that we want is the summation over i of what's left over after pulling out the summation.

We have the individual terms. But of course, that is nothing other than, again, from Roy's identity, the underlying demands, which, when we sum them up, is aggregate demand. So we've used this particular utility function, 17, to derive the aggregate demand, which is now proved to be the sum of the individual demands, and by this construction, depends only on aggregate wealth. The w bar is missing there. Questions? All right.

So let me give you an example of the underlying utility functions and consumption sets. Suppose, being explicit about the heterogeneity, each household i 's consumption set is a subset of this L -dimensional space. Namely, they're, like, subsistence or minimal bundles. So the amount of good l cannot be lower than some \bar{x}_l per household i .

This is the heterogeneity. It can depend on household i . And let's have utility of household i over vector x just be the difference to a power between x_l and those minimal subsistence bundles. So the heterogeneity is in the subsistence bundles, which can vary over households. But the sigmas are the same. It's not σ_i . It's σ , common, across all of the households.

So if we had an economy represented by these consumption sets and these preferences, we could look at their indirect utility. And if you do the math, you get this ugly thing. But it does have the right form, $a_i p + b_i$ times w_i for household i . And then you add it up-- namely, add up this guy over households i . And you get the expression at the bottom of the slide.

And in particular, the heterogeneity in the subsistence bundles is right here. So they're all different. They have different minimums. But miraculously, we're able to add up those minimums and call it kind of a pseudo-subsistence bundle for an aggregated representative household who would have the indirect utility that's written out here at the bottom of the slide.

I don't have too many more examples in front of me today, but I mean, I do want to give-- I did want to give you one example, because this indirect utility space is a bit different and a bit abstract. So here, we actually start with underlying utility functions. I said all that in words.

So now, let's go to the normative representative household. I should say, it's the same guy. It's the same person. It just has two characteristics. It can be used for positive economics in the sense that it's predictive of what the entire economy is doing, and it can also be used to make welfare statements by a rank ordering over the utility function.

But there's not, like, one representative normative household and one representative positive household. It's a given representative household with both positive and normative characteristics, and we are about to go through the normative part.

So let's have the individual demand functions-- this is-- the first little paragraph here is the positive part, that these demands can be derived as if through an aggregated demand function given the total wealth of the economy. Now, we want to ask the normative question.

Given a set of individual demand functions, can we find this fictitious household i^* such that Pareto optimal allocations in the economy can be thought of as coming from this representative household maximizing utility? I keep emphasizing fictitious, amazing, because I don't want you to think that this is a real person. It's a construct.

So let's take our underlying economy, which we've been reading this line for the last two or three lectures. Underlying economy is script E characterized by consumption sets, utility functions over all of the households, production possibility sets over all of the firms, some aggregate endowment \bar{w} .

Now, finally, we define a normative representative household with two characteristics. If we started with an allocation x and y and we found an alternative allocation x' , y' , which Pareto dominates x , y , meaning at least as good for some and strictly better for others. Then under u^* , evaluated that to some of the prime bundle is greater than u^* evaluated at the original bundle.

In other words, it's kind of intuitive, but it's not always true. Here, for the definition of Pareto dominance, we're looking at individual allocations over all the households and trying to find something that's at least as good for some and strictly better for the others.

If we can find this amazing u^* , then all we need to do is evaluate the aggregates, namely, the aggregate under x' and aggregate under x . And what we need-- and apart from condition 2, the only thing we need-- is that u^* assigns a higher value. In other words, it's easier to look for Pareto improvements just by looking at the value under the pseudo-utility function evaluated at the aggregates.

So the second thing is similar, but not quite. Instead of starting with the underlying micro distributions, suppose that we had a feasible allocation x and y , as before, and an alternative feasible allocation, x' , y' . And we're given that under u^* , the utility is higher at the aggregates.

Then we want to claim that there is an alternative feasible allocation which Pareto dominates. In other words, is this x allocation Pareto optimal or not? Not, if we can find another allocation, x' , which dominates, but dominates in the sense of u^* . And u^* is just evaluating these things at the summed up aggregates. And actually, there's a subtlety here which is important. You can't just stop. You may have to redistribute. Well, here's what I mean by that.

If this is true, there exists an alternative feasible allocation, \tilde{x} , \tilde{y} such that the summation over \tilde{x} of the demands and so on allocations is exactly the same aggregate as under the prime, and the y 's aren't changing. So this kind of subtle step here, it doesn't mean, necessarily, that you can stop with x' and y' and claim victory, because if you stop there, it's not necessarily true that everyone's better off-- or no worse off.

Some people may be damaged. In fact, maybe this is the time to talk about free trade. When we were looking at that lecture, we represented a country by some indifference curve as if from some u^* representative household for the whole country to look at the direction of trade. So we were getting the country to agree we get demand in a positive sense by looking at that representative agent indifference curve.

And it might have appeared that the whole country could be made better off by free trade. On the other hand, you saw in detail, in two different lectures-- three, really-- that not everyone is made off under free trade. Remember, the households that are relatively well endowed with a factor that's abundant are helped by exports. The other guys are hurt. Et cetera.

So how can we claim that free trade makes everyone better off? We have to redistribute, which is not always done. Hence, people complain about free trade and they want tariffs. Anyway, I'm belaboring the fact that this lecture is consistent with the other ones, except now, we're finally in a position to be able to talk explicitly about redistribution as a necessary part of making everyone better off.

In other words, to sum up, even when we have preferences aggregating, it is not necessarily the case that everyone is made better off under some alternative allocation x_i^* . You have to redistribute the underlying bundles.

So let's do the same definition, but do it in the space of indirect utility instead of direct utility. Suppose we're given a vector of wealth, as in the distribution of income, and some aggregate wealth that corresponds with that. p is a price vector. Define the indirect utility functions, again, as we have been doing.

Now, we want to find this miraculous star person, this normative representative household, i^* , that has an indirect utility function v^* , evaluated at the price vector p and the aggregate wealth w . And there are two conditions, 1 and 2, that correspond with the one in the allocation space x_i on the previous slide. Namely, if you have two configurations, you start with a p and w wealth.

You're looking at an alternative, p^* , w^* , price vector in wealth, such that looking at the actual real indirect utilities of all the households, some people are no worse off and other people are strictly better off under the prime allocation. So that's the statement of Pareto dominance in the space of indirect utility functions.

And what we want is a way to do it more easily through the star, the indirect utility of the normative representative household, which is just evaluating the alternative scenario, p^* and w^* . But it's w^* added up over all of the households. So we're just using aggregate wealth here on the left-hand side, and also aggregate wealth on the right-hand side, when we're thinking about trying to, quote, unquote, "Pareto dominate."

So just like before, it's very intuitive, but amazing when it's true, that the way to look for Pareto dominating allocations is just to evaluate the aggregated allocations and see whether you get a higher value under v^* . Likewise, going the other way, start with the aggregate wealth and price vector p .

Look at an alternative scenario, p^* and w^* , such that the utility under v^* is higher in the alternative allocation. If that's true, then there exists an alternative income distribution w^* which adds up to the aggregate wealth in the alternative-- the dominating allocation with the prime on it.

Now, again, let's pause here. This alternative income means an income distribution that works under the prime scenario relative to the initial income distribution in the baseline without the prime. So we have to have the freedom to distribute wealth, not necessarily according to private ownership, when we change the prices and so on, but potentially, with transfers of wealth across the agents-- giving to some, taking away from others, as in the second welfare theorem, where we've seen this kind of redistribution before.

So we find this alternative such that we're getting Pareto dominance in the underlying indirect utility functions. Not worse for anyone and strictly better for some. So how do we find this representative household, or the representative household with this normative property? Same guy, Gorman Polar form, although this is getting stated as a proposition.

If we have this Gorman Polar form with this particular way that heterogeneity is entering into the indirect utility functions, then when we look at the indirect utility function under the aggregate, being in particular the sum of the a_i plus b times w , that is the normative household.

We already proved the positive part earlier. Now, it's normative. I put positive and normative both on the slide to make the point that, as I've said, we're not looking for a normative household that's different from the positive household. We're looking for this representative household with both attributes. We proved the positive thing earlier. Now, we need to prove the normative part.

So suppose we have those Gorman Polar form, in particular, to figure out which direction we're going from the condition 1, this guy here. Condition 1. So we have a rank ordering under the actual indirect utility functions being explicit about household i . So under the v_i 's, not worse for some and strictly better for others. So we're given that.

Now, we want to prove that when we evaluate under v^* using the aggregates, we get strict dominance. So let's look at it. v^* under the aggregate alternative prime allocation is-- and this is just now writing things out. This is the definition of the indirect utility function v^* , the one that's going to work.

And then as we keep going across these equations, we recognize, again, the taking the summation over l out in front of all of it, that we're just adding up the indirect utility functions. But again, we're given not worse for some and strictly better for others. So when we add up the indirect utility functions under the prime allocation and compare it to the original allocation, it must be strict dominance.

And then if you reverse the ordering of these v_i 's, collecting terms and going back toward the left-hand side-- not under prime, but under the baseline-- you would get this v^* indirect utility for the representative household. So the bottom line is the v^* under prime strictly greater than v^* under the baseline. And we proved one.

The second thing we want to prove is this-- that we have an initial allocation and we're looking at a prime allocation that gives higher indirect utility for the aggregated household. And then we want to find a way to assign wealth that makes everyone better off, essentially.

So again, from 2, we have this as a given, that v^* returns a higher value under the prime relative to the baseline. We want to construct an alternative income distribution-- seems like that might be hard to do-- that makes everyone, in fact, I'm going to show you, strictly better off. Actually, strict dominance for everyone.

And it's actually very constructive and intuitive. We're already given here that we have this equation 30 at the aggregates, the left-hand side being greater than the right-hand side. So take the difference and call it c . It's a positive constant now, the rest of the thing is just writing out what v^* is under prime and v^* is under the baseline, and then collecting terms.

The only reason to do this part using the definition of v^* is to go searching for the wealth levels that we want, namely, let's give each household their proportion of c . There are i households in the economy. Give each one c over i . So we start with the baseline in terms of the indirect utility function. We add on to the baseline this positive constant.

And we want that to be the indirect utility per household i under the alternative. So now, we can just solve this equation for w_i . It's almost mechanical. These are, though, the required new distributions of wealth that will work. Why does it work? Well, working backwards, if you give them this, then we have this equation.

And you can also see, because we've just added a constant to the indirect utility function, that each household i under the alternative is having more utility than under the baseline. That's this v_i of p' w' , v_i of p and w_i , plus this proportional constant. So it's almost trivial, because now, we've gone the other way. We proved 2, this thing. So we've succeeded, if they have these Gorman Polar forms, in constructing a pseudo-utility function over the indirect utility function over the aggregates. Questions?

OK. So now let's go back to the Pareto problem. In particular, we did a version of this slide when we did the second welfare theorem. We have this underlying economy with consumption, preferences, production functions, aggregate, endowments, satisfying non-negative orthant, and differentiable convex production sets or concave transformations, non-negative aggregate resources, consumption sets having interior points, utility functions being strictly concave, and also an interior point with respect to aggregate resources.

I'm rushing a bit, but only because you've seen this slide before. If all of that is true, then we can solve this Pareto problem to get Pareto optimal allocations. This is just maximizing λ weighted sums of utilities subject to consumption being feasible, production being feasible, and the resource constraints being satisfied. Demand equals supply.

I should say, allocations are feasible given endowments in the production set. Problem? What are the λ s? We move the λ s around, we get different Pareto optimum. That's what we did with the second welfare theorem. And we've used this with risk sharing in many applications previously. But we had to keep track of the λ s. Here's what happens with the representative consumer. No more λ s.

Now, we just maximize utility. Well, this x is now aggregate consumption. Production sets are the same as before. Resource constraint's the same in the aggregate. I mean, x is aggregate. Aggregated x is entering into the pseudo utility function u^* , and we want to find the highest u . Remember why-- because we already proved this, right? It's enough to look at u^* in the aggregates in the sense of satisfying both these conditions.

And that's what we're doing here. We're just looking in the space of aggregates under this pseudo-utility function u^* . If we're not maximizing it, then when we move in a direction that improves it, we're actually improving the welfare of all of the individual households. We'll, in a sense, Pareto dominate. So simple to write, but extremely powerful. And this-- so this is contemporary macro.

This is what is usually assumed in macroeconomics or aggregative economics. No reference to redistributions or wealth, underlying wealths, or whatever. So mesmerized by the individual household that we forget to spell out the underpinnings, and in particular, the assumptions that are needed. Are they satisfied? No, typically not.

That previous lecture on micro data, we saw that income expansion paths are a special case. In particular, there are non-linear expansion paths. As we change income, expenditure shares go up for some goods. We call them luxury goods. Expenditure shares go down for other goods. We call them necessities. And so on. So that's what's in the data. Engel found this very early on.

And we presented other applications, Giffen goods being an extreme version of the case when the income effects, through necessities, are so strong that they dominate the price effects. Well, maybe one can hope that as an approximation, even though we would not literally satisfy the assumptions, that those income effects come close to canceling out.

That's kind of the best case. Macro with micro underpinnings. It's a reference to the trade lecture. I've already given you my spiel on that this afternoon about the gains to trade and the fact that we have losers, although the theorem is telling us that if we have representative consumer, then we're safe in looking at the max, both the positive and the normative. But it will require redistribution of a particular kind.

And finally, the economic impact of COVID-19. I put that in there to remind you of things that we found in that lecture when I was reviewing what the impact has been. And we had relatively poor households who were unemployed cutting back on expenditure, or potentially increasing expenditure only if they got the government money, and also looked at higher-income households through JP Morgan's data set.

And those guys had different expenditure patterns. So it does not seem as though these effects are canceling out when we actually look at the underlying micro data. I mean, to cancel out we wouldn't see anything happening-- well, I should be careful here. There is an aggregate adverse shock, obviously. We're all suffering from it, at least indirectly. So the aggregates are down. But then the question is, do other changes balance out at the aggregate? And the data I presented to you suggest that's not true. So questions?

AUDIENCE: Can I ask a question about the slide 22?

ROBERT 22?

TOWNSEND:

AUDIENCE: Yeah.

ROBERT Yes?

TOWNSEND:

AUDIENCE: Yeah, here. I'm a little bit confusing that. Do we do assume that the sum of x prime will equal to the sum of x here? Because in the slide 21, we say that, at the [? center ?] of the allocation should be the same as original. And the sum of production-- the production is the equal to the original. So when you go from 21 to 22, do we still assume this?

ROBERT It's only implicit in the sense that we're looking at individual wealths and aggregating it up. So I agree, it's probably not obvious, but there are just two representations, one in this space of indirect utility functions with the parameters p and w , and the other and the actual underlying allocations.

Another way to put that is that if we-- where is it? If we went with-- can't find it. Yeah. Take, as an example, this utility function and the associated indirect utilities, these guys. You could go back and forth. You could do it one way or do it the other way. You're going to get the same answer.

AUDIENCE: Oh, I get it.

ROBERT OK. OK. So that's Gorman aggregation, with a lot of the pieces coming together. Next time, we're going to go back in the other direction. We're going to ask, if we don't do anything special like this parametrically, does economic theory put any restrictions on data at all?

And the answer is going to be yes, so we'll talk about theory that can be rejected in the data, which is a good thing. Otherwise, there's no content. So that's all for today. Thank you for coming.