

[SQUEAKING]

[RUSTLING]

[CLICKING]

ROBERT TOWNSEND: All right, so let's get started. Are there any questions from last time? So as you are no doubt well aware, this is the last lecture I'm going to give. The exam is a week from yesterday, on Monday. So what I want to do today is to review the material and see if there are questions.

So today, we're covering the material from lectures 16 on, 16 through 21. And normally, the exam would cover 16 through 21 and I could have written that in there. I am just trying to avoid confusion though because the course is cumulative and one thing builds on another. Essentially, it will focus on 16 onwards, but you should be, by now, very familiar with the early material. And some of it may come up.

So this is just to prevent me from misleading you. If I said it was only on the last one third of the class, then you may get lulled into thinking that it's quite narrowly about only those things, as opposed to the things on which it is building. Anyway, just one word of clarification. And again, it's only the third midterm. It doesn't count anything special in terms of your grade.

Let's see. So last time, we talked about money. Bitcoin, in particular. And I singled out two questions, one of which I kind of mentioned in class, but let's get you guys engaged in this. A central banker, namely, Dudley at New York Fed is heard to say that unlike Fiat money, Bitcoin's positive value in markets is just a bubble and with wild fluctuations. So Bitcoin should be regulated, if not banned altogether.

And then it says discuss this statement using the tools you have learned in the class. So could I get some volunteers to answer that question?

STUDENT: To be honest, I'm pretty unsure, but I do remember we talked about how you can have equilibrium with bubbles. So they're not inherently something that you have to regulate. It's possible to still have an equilibrium, despite Bitcoin being overvalued.

ROBERT TOWNSEND: Right. That's true. The question actually asks about Fiat money, I kind of slipped that in there. That unlike Fiat money, Bitcoin is just a bubble. Can you maybe respond to that part? Maria, or anyone else for that matter.

STUDENT: I mean, I guess that is somewhat true, as in it's not tied to a tangible asset underlying it. I'm not sure what exactly you're looking for here though.

ROBERT TOWNSEND: OK. Fiat money is a bubble. It's a bubble the way that Bitcoin is potentially a bubble. Fiat money used to be backed by gold. So there, if you didn't want to hold the notes, you could go to the Federal Reserve and demand a real commodity, so then it would not potentially be a bubble because it was just a certificate for a commodity that would have real value. But when Nixon broke the link to gold, then Fiat money is just pieces of paper and nevertheless, it has value.

So the point of the lecture was that Fiat money is a bubble. Now we use the word "bubble" is almost deliberately provocative because it sounds bad. But Dudley, who made the statement seem to have not let's say, forgotten-- seems to have forgotten that his domain-- Fiat money-- is potentially fragile, as well. Now, the difference is the Federal Reserve behaves itself, by and large-- or not to jump in with some parenthetical comments on that, but let's argue that that's the case.

So then with forward guidance and so on, no one is expecting high inflation, which would lead to speculations about how much Fiat money is really worth, and so on. So whereas with Bitcoin, we don't have a long history. We haven't set expectations, and so on. So there are differences between the two.

But in a model, it's pretty hard to make a distinction because they're both, essentially, worthless even though they have value. So that's what I was looking for. The second question, should currency pay interest? Or more specifically, what did we learn in the lecture about that?

STUDENT: Is it related to the non-zero price of a currency because I remember we say that cash has positive value because it can be used to pay taxes. And that's a cash advance. So if the price is equal to 0, then nobody will demand the currency, right?

ROBERT TOWNSEND: Yeah, you're speaking to the levels. So the point of having cash in advance or taxes and so on is that money would always have value because there's something you need to do with it, either buy stuff or pay taxes in it. This question interest implies something about how the price of money is changing over time, or that currency-- how can I put this-- the mechanical way to bear interest, even though you can't take it necessarily to a bank and exchange it for more money is to have the thing that money exchanges against be lowering in price. In other words, a deflation where the nominal price level is going down means that you can buy more goods with money tomorrow than you can buy with it today. And that's effectively like paying interest.

The way to remember that is inflation is bad, obviously, because the purchasing power of money is going down. It takes more and more money to buy goods. With deflation, the purchasing power of money is going up. So that's effectively a way to pay interest. And what did we learn in class about whether or not it's optimal to engineer that? Any volunteers? Yeah, sorry. I know it's the end of the term and you're all very, very busy.

So the answer was, in that model of money was spatially separated agents, there was a monetary equilibrium with a constant price level, but it was not Pareto Optimal. And to get to one of the Pareto Optimal allocations, it was necessary to engineer this kind of deflation. So that's an example of paying interest on money is actually something the Federal Reserve should be doing. I hope it's coming back to you now. Anyway, those were the two main ingredients of this very last lecture on bubbles.

So I probably won't, if history is any guide, I won't have time to open up lecture, this last lecture 21 and review it, but those were the two principal components. Namely, money as a bubble and this policy aspect. So again, the study guide should help you prepare for the exam. And now we're ready to review the lectures.

There's only six of them. So this material is probably pretty fresh in your minds. I think what I'll do rather than obviously I wouldn't have time anyway to go over all the individual slides of each lecture, I will try to remind you what we did and maybe draw some connections to the other lectures and how things are kind of fitting together. So this was the lecture on the fundamental welfare theorems, namely, that competitive equilibria under certain assumptions are Pareto Optimal. And any optimum can be supported with transfers.

So this lecture was about sufficient assumptions and the corresponding proofs. And it ended with extensions to more general commodity spaces. So you may remember that I did this. I showed this-- back to the earlier lectures, the first part of the course on risk sharing, I showed the income movements of households over time and compared to each other and including the consumption, we noted this gradient here that the average level of consumption is increasing as one moves from landless households to larger landowners.

In the risk sharing framework, that was related to λ , the Pareto weight. And very clearly, for certain functional forms, it was exactly that. Pareto weight of one household relative to the others. But the motivation here was to introduce the welfare theorems. So then the issue was whether the λ s that are implicit in that diagram are related to wealth, like landholdings and the number of plow animals, and so on.

So I was anticipating a connection between λ s as in the Pareto problem of Pareto Optimal allocations versus the wealth that may be needed in a competitive equilibrium. You'll see that again, or be reminded of it momentarily. So the first welfare theorem is that any price equilibrium-- competitive equilibrium, including a price equilibrium with transfers is Pareto Optimal under minimal sufficient assumptions, namely rational preferences-- which we haven't reviewed for a while but it was in lecture two-- complete transitive and continuous preferences and local non-satiation. With those two assumptions, we get this sufficient for this theorem.

So this was a picture in our familiar Edgeworth box of a competitive equilibrium being Pareto Optimal. And the proof, you may remember, relied on starting with a given allocation, the star allocation and then conjuring up, by contradiction, that there might be something Pareto improving. We prove this little Lemma, that anything which is weakly better has to cost us at least as much. And then we just went through the math, summing up over all the households. Obviously, if it's strictly better for some households i prime, then it must have costed more, or otherwise, they would have chosen it.

So we get, in equation one, the statement about total expenditures, the value of total expenditures summing up. All the households-- I'm going quickly now-- we used profit maximization. We coupled equations one and two to get a statement about resource constraints. And at the bottom line, concluded that this alternative Pareto dominating allocation would not be feasible in contradiction to what we assumed. So that was QED. That was the end of the proof of the first welfare theorem.

Now later, even today, we'll review it. This thing fails with if these sums are infinite sums and don't have finite value. Then the second welfare theorem was illustrated here that we have Pareto Optimum characterized by tangency of indifference curves. But that nevertheless, it's not reachable if we start from the endowment and start drawing potential budget lines through the endowment point, as you would for candidate budget lines in a competitive equilibrium.

So this was meant to illustrate that to get to this particular Pareto Optimum that's in the second welfare theorem, we're going to have to take some value away from agent A over here and give it to Agent B. And that's what this said. The tax on A is equal to the subsidy on B. T being transfer, so negative is a tax. So then we could get, as in the second welfare theorem, to any allocation along that contract curve where the indifference curves are tangent by reallocating wealth.

Now to prove it, we went through this constructive argument. We defined an economy. And this slide, or a version of it, shows up at least three times in these six lectures. So an economy consisting of preferences endowments production sets and profit shares here, assuming a lot more. Concavity and convexity in preferences, production sets, et cetera. And then we get the second welfare theorem, that any Pareto Optimum can be supported. For some wealth, any of optimum associated with some lambda vector of maximized weighted utilities can be supported by a suitable redistribution of wealth.

The way we prove this was to go back to what you know about Lagrangians. So this was, I think, also in the second lecture of the class. And you've used it over and over again. Maximizing concave functions on convex sets, generating all the shadow prices. So we used it here three times. First we use it for the Pareto problem. So we're starting with a Pareto Optimal allocation.

So we know, again, from things you did earlier in the class, when we introduced the notion of Pareto optimality-- I don't remember the lecture number-- we maximize weighted sums of utilities, lambda's being these weights, subject to resource constraints and production constraints. That's all written down here. The solution to any one of these Pareto problems is Pareto Optimal, and all the Optima can be described as solutions for some Lambdas.

So we start thinking about these so-called Pareto problems as equivalent with Pareto Optimal allocations. Use the Lagrangian, set it up this way and get the first order conditions, including the primal, dual feasibility on the shadow prices, complementary, slackness and all of that. The point is that with all that concavity and convexity as appropriate, these first order conditions from differentiating the Lagrangian are not only necessary, they're sufficient.

So any solution to these first order conditions, allocations and shadow prices that satisfy the first order conditions will be Pareto Optimal, will be a solution to the Pareto problem. So we did that. Then we went to the second part, reminded ourselves of the definition of a competitive equilibrium with transfers. Namely, consumer optimization, firm maximization of profits, and so on, along with a feasible distribution of wealth, as I was just saying. So here's the second and third time. We're going to use the Lagrangian for the consumer optimization problem and use it for the firm profit maximization.

Here is the consumer maximization problem. Maximizing utility subject to a budget, we get this first order condition. And again, with a concavity in the constraint setting in the objective function and convexity in the constraint set, these are not only necessary, but they're sufficient condition for an optimum. All we need to do is find the μ_i at given prices, p , and in X_i^* , and it will be maximizing for the consumer. Likewise, for the firm maximization, we have necessary and sufficient conditions to find candidates for the decentralized competitive equilibrium.

And then all we did was match them up. Now I think the tricky thing here is to remember what the given is and what the objective function. The given was a Pareto Optimal allocation. So for the consumer, we started out with these conditions, which are the necessary and sufficient conditions from the Pareto problem being satisfied. So we have the γ_L for commodity L as a given. The lambdas were given because we had a given certain candidate, Pareto Optimum.

These were all satisfied. We're looking for this thing. We're trying to see if consumers are going to be maximizing. So the candidate for the price would be the Lagrange multiplier. And the candidate for the marginal utility of income, the shadow price and the budget constraint would be one over lambda. So here, I go back to the slides of the Rocky Mountain and Kansas. This establishes a direct relationship between the Pareto weights and the Pareto problem and the marginal utility of wealth. And hence, the level of wealth. It's quite general.

It's super robust. And so then the rest of this was just the tedious part of constructing the wealth, has the evaluation of expenditure and making sure that they all add up. We refer to this way of finding the prices from the second welfare theorem as marginal utility pricing. I mean, more accurately, it's the lambda weighted margin utilities which are equated all over the households, not the levels.

But when you take ratios across different goods, L and K, the marginal rates of-- the lambdas cancel out. So the marginal rates of substitution are equated over all the households, as they are in this picture, along this orange, green, orange, green? A yellow, orange line. And this made the wealth's add up. And we did a second proof of the welfare theorem, which is namely one nuanced thing, which is instead of just maximizing utility subject to budget, we minimize expenditures subject to achieving a certain level of utility.

That conjured up this. So this is the target and different curve, as it were, as if it were coming from one consumer-- interesting comment. But it's just the aggregated better than set associated with the Pareto Optimal allocation. The other part of this diagram is two production functions, different from each other, but aggregated up to an aggregated production function. And then you can see, as if it were a Robinson Crusoe representative consumer, the optimum is characterized by a maximal utility level subject to the single production set.

And then my point is that we get this separating hyper-plane. So all these definitions supporting hyperplane, et cetera, that's really where given all the convexity, that's all you need to do. These convex sets are separated from each other. So this is a quite general proof and it works in many spaces, not just finite dimensional Euclidean spaces. And we're about to review the representative consumer, as I said, this looks all the world like a representative consumer because I've seemingly added up the indifference curves.

But one should be aware that it is not the Gorman consumer that's going on here. And this particular Pareto Optimum is sensitive to the distribution of wealth. You just can't see that. I did not go through that when I gave this lecture because we hadn't gotten to Gorman and it wouldn't have made any sense. But I might as well say it now, if you don't have a Gorman representative consumer, then as you increase this guy's budget line and decreases this one, you may change the shape of this seemingly aggregated up in indifference curve and, therefore, change the whole optimal allocation. So hopefully that helps you think about the hidden way in which heterogeneity is entering into this picture.

And this lecture, I finished it off in an abstract way by saying the proofs of the welfare theorems are quite general, that does not require a finite dimensional commodity spaces. I gave you the sufficient assumptions for the two welfare theorems and kind of quit there. If we'd had more time, I could have gone into greater detail. I only asserted without proving it that if we took what we did earlier in the course about Pareto Optima with obstacles to trade with incentive constraints and potentially non-convex consumption sets because of discrete choices, we could nevertheless introduce lotteries. And that would have turned the space into a convex set and the theorems, as in Debreu, would apply to that.

So although this lecture just proves the welfare theorems carefully for finite dimensional commodity spaces and only asserted that the theorems extend, it would cover competitive equilibria in contract space where households can buy whatever contracts they want, granted at certain prices. And firms as intermediaries will create contracts that maximize profits. Tying some of the loose ends together, but without actually going through much of the details there. So that was lecture 16. Questions about 16?

So we'll come back to it again because we're going to go through the failures, or review them quickly. Then this was the existence lecture. The point being that, for example, the first welfare theorem assumes existence of an equilibrium and then assures everybody that it would be optimal under sufficient conditions. But if the model could be vacuous if there is no competitive equilibrium. So we're dealing with things like that directly in this lecture.

And this was a listing of the things covered. So the main mathematical object tool really was fixed point theorem, this was an example that some function has across the 45 degree line, in which case whatever you put in, you get back out. The argument in the range is the argument in the domain under the function, F . And it was sufficient for this to have the set of bases of values in the domain and the range being non-empty convex and compact and for the function F to be continuous.

And that's Brouwer's Fixed Point Theorem. I should say though, well, we're about to get there, that this is all about existence. It's not about how to find it. And at least one of you has asked me questions about that. Anyway, this is Brouwer's Fixed Point Theorem. It's paired with Kakutani's Fixed Point Theorem. And the way in which they differed, primarily, was not the set A , which is still compact and convex and non-empty, but the function F , which is no longer single, necessarily single value.

It could be multiple values. Or there are multiple values in the range for any particular input x in the domain. And of course, not anything goes and had to be convex value and continuous in a certain sense. And then with Kakutani, again, we have a fixed point. So these are very powerful tools and they're used for competitive equilibria in one way or another and used in games. And we actually combined those at the end of this lecture to have a game that would be achieved under a Nash equilibrium as a Walrasian equilibrium.

So here are what happens if you don't assume continuity of the function, you get these jumps. So if the aggregate demand for an economy somehow had jumps like this, then you could not be assured that there would exist a competitive equilibrium. And we're looking for the existence of prices, such that you put prices in to each of the agents maximization problems, look at their demands, see if demands add up to supply. And if we find the price that makes that happen and all of this can be formulated as a fixed point.

Now I did not go through the math of this because it's surprisingly tedious. However, let me say that the existence proof is the existence of a competitive equilibrium as in at least one fixed point. Herbert Scarfe did a lot of work on this subject. And in his proof of the Walrasian equilibrium, he actually did more, which was to go back to Brouwer and actually show how to find the fixed point. That is to say, in an algorithm, take a guess, look at directional derivatives, formulate a new guess and iterate until you get to the end or approximately the end, however close the computer can make it.

And he would have constructively found the Brouwer's fix point, which works much more generally than just finding the existence of a competitive equilibrium. This comes up later when we talked about restrictions on data having to do with algorithms and what the algorithms actually look like when computing solutions to individual problems and defining a competitive equilibrium. Anyway, I just wanted to draw your attention to the distinction as an overview-- that's my job today-- between existence of a competitive equilibrium and how to find the competitive equilibrium, algorithmically.

What we did do in this lecture was to prove the existence of a competitive equilibrium using the second welfare theorem, which we just reviewed in lecture 16. So the idea is instead of iterating on candidate prices to find where supply is equal to demand, we iterate on these lambda weights. And the idea is very much, here's that economy again, I promised you it would keep coming back. The statement of the economy with the convexity and concavity.

The idea is to go back to the Pareto problems, this is now the second time today that the Pareto problem has come up. We just used it in 16 for the second welfare theorem. Here it is again, stated rather succinctly to remind us of these lambda weights, such that when we max, we get Pareto Optimal. So we start with those lambdas and then we essentially look at the solution to the Pareto problem, these x^* 's, for those particular lambdas. Look at the prices, which you now know and remember are the shadow prices on the resource constraints and we generate the wealth associating with this purchasing power.

So what we're going to look for to get a standard Walrasian equilibrium his particular lambdas, such that the valuation of privately owned resources for a given household, I , is equal to the wealth coming from the lambdas. Here this is a better statement of it. A particular lambda star, the fixed point-- find me, please-- the fixed point would be such that when we substitute lambda into the Pareto problem and generate the optimal allocations, we'll have some valuation at those shadow prices of the expenditures, which when we sum up over-- and summed up over all the commodities, L , would be exactly equal to the right hand side of the budget, the valuation of purchasing power that the consumer has as a result of owning the endowment and getting shares in the profits of the firms when the firms are maximizing at that same lambda star.

So this is the fixed point. I think I mentioned in a question and answer at the beginning of lecture 18 when we reviewed 17, that I felt I hadn't done a good enough job explaining the intuition for this mapping. But it does come from the second welfare theorem. Start with a given guess about the lambda that pins down which particular Pareto Optimum were targeting and suppose that were the case that the valuation of expenditure at that optimum for consumer I were strictly greater than the valuation of his or her privately owned resources.

Well in this case, we gave this consumer too much. They can't afford it. So the obvious way to iterate-- if indeed you were even searching-- is to reduce the lambda for this household and increase the lambdas for the other. So this is very much like that Edgeworth Box diagram for the second welfare theorem, where I said household A has to give up stuff. This would be household A if the inequality is greater than or equal to. And household B would be the recipient and be getting the transfers.

For household B, you'd be under valuing evaluation of the resources that the assigned optimum could be less than the valuation of their privately owned resources. Now again, Negishi didn't actually show that one could find the fixed point by iterating on the lambdas, but in practice, it typically works. So as a numerical algorithm, you could just be trying out candidates, lambda checking out these kinds of budget equations and so on. So here again, this links up to lambda equal marginal utility of income. That's the Kansas Rocky Mountain pictures at the beginning and I reinforce that when we went to the second welfare theorem.

So this proof of the existence is working off the same equivalence between the lambda and the inverse marginal utility of wealth. This is the computer science part that it looks like it might be hard to find, actually constructively. So these guys algorithmically want to find the Walrasian equilibrium, not just be happy with the notion that there's a fixed point so that it exists. And this literature works on the same, they kind of discovered, again, Negishi. So that's why I told you about it.

But you get a better sense here and I gave you a hint of that, with respect to Herb Scarf, that the way to find it is to periodically look at the aggregate demand oracle, as it were, and take a local derivative of it to see which way you ought to move. So it's a little bit harder than what I said about just iterating on Negishi's algorithm. But their point is that they designed an algorithm that finds the competitive equilibrium using Negishi and it's not that hard in the sense that it's polynomial of degree three and the number of goods, which is N . So it's polynomial, rather than exponentially hard.

So then we went to the second part of the lecture, which was Strategic Form Games. This looks quite different. You have utility functions again, but they can be a function of other people's strategies, not just your own. Finite set of players, that's kind of similar. And we talked about strategy profiles. What household i is doing, or trader i , versus what all the others are doing. And we defined mixed strategies, where the players are randomizing, potentially, over a finite set of discrete actions.

So this is one place in the course where, again, these lotteries come up to convexity things in this case. And indeed, the convexification is what allows us to prove the existence of this mixed strategy equilibrium. A mixed strategy equilibrium is a set of randomization routines for each player stacked up in a vector, such that given what all the other players " I " are doing, minus " I ", all the others. But player " I " wants to do is the I -star, signa-star vector component of that vector, better than any other component.

So this is the definition of a mixed strategy equilibrium. And then we had this interesting side note here. And this is a third related, but not identical statement, which is you have a candidate for an equilibrium, how do it really is an equilibrium? It's related to the idea of the fixed point, momentarily I'll show you that. But here, the question is, how many things do you have to check at the candidate mixed equilibrium Nash, mixed strategy Nash equilibrium. And it turns out you only need to check pure strategy deviations here, rather than all possible mixed strategies. Thank goodness because that would be hard.

And the reason for that is, if you can find a degenerate action which dominates the supposed mixed strategy maximum, that's enough. That's a particular degenerate mixed strategy to show that you could do better. And if there were multiple ones that were doing better, you pick the best one that dominates to break up the solution. And if there were two or more that generated the same utility, you could randomize over them, but there's no need because they all generate the same utility so you might as well pick one.

So that's the intuition for why it's enough to look to check a given candidate profile against these degenerate actions, not mix strategies. The rest of it is remembering exactly these kinds of responses. The better-than sets, or best response sets for each player. So given the others are doing, you want a best response for player I, obviously is going to maximize. Hopefully, generating the same mixed strategy as a candidate, but maybe not, if it's not a mixed strategy equilibrium.

But all the other guys are doing the same. So we're going to look for a fixed point in the space of best response strategies. And the proof of the existence of a strategy Nash equilibrium had to do with the convexity of these best responses, as in Kakutani, for example. In particular, all this stuff. I'm a little worried I'm going to get bogged down in the weeds, but I did want you to, on the other hand, see the big picture.

So now we have two parts. The Nash equilibrium for strategic games and the Walrasian equilibria, they're all using fixed points, but the economic problem looks kind of different. So I finished this lecture off by drawing the connection, like how do you actually set up a market to find a Walrasian equilibria. And we did that with respect to these limit order strategies on the supply and demand side and then wrote it down as the game. So once we wrote this down as the game, where players decide what to bid, given what others are bidding it is a strategic form game and we characterize the equilibrium.

Turned out, you needed some penalties, otherwise, they would spend more than they have and that can't be consistent with maximization-- well, it would be maximization, but there'd be no limit. So it would not be consistent with finding the Walrasian equilibrium. So we did have to find the optimal penalty structure in order to prove the theorem that any Walrasian equilibrium could be supported as a mixed strategy equilibrium in this market game. And the other direction is also true.

So that was lecture 17. Are there any questions about what we were doing in lecture 17? 18 was aggregation. And again, this is getting more recent. So hopefully, a little more fresh in your minds. We want to do this positive and normative. And I would say the content of this lecture is, what do we have to assume, like it or not, in order to be in the land of macro, where you have a representative consumer? And is usable in the sense that what the representative consumer would do, which from a positive economic standpoint, generate aggregate demand as if you had all the underlying households, but you don't need to do it. You just use the representative consumer.

So it's a very powerful tool. But again, for you to decide how strong or potentially unrealistic are the underlying assumptions that allows us to do this? So we reviewed the notion of the indirect utility function, which is the maximized utility after solving max utility subject to budget, given the parameters, the price and the wealth of a household. And we quickly reviewed some properties of the indirect utility functions. In particular, Roy's identity, which miraculously says that if you start differentiating the indirect utility function with respect to prices and respect to wealth and start taking ratios, you'll get the underlying demands.

And this was a tool. And it's heavily used. Again, now for the third time, it just came from the Lagrangian problem for the individual household. And Roy's identity came from that. And then we got to Gorman, which is a particular representation of these indirect utility functions as a function of prices and wealth, namely, this thing in six. So it's separable, it has an intercept of the function of p , slope term on wealth as a function of p .

I noted that the Cobb Douglas would satisfy this. We actually did the math. This math is never very hard to do, as long as you remember that when you have these Cobb Douglas utility function, these coefficients represent expenditure shares. So $p_1 x_1$ is α share of wealth, w -- likewise, $p_2 x_2$ is $1 - \alpha$ share of wealth, W , for the second good. So you can plop them down right away and just substitute into the given utility function and you end up with this functional form.

I also noted that although this is special in that the intercepts are 0, it also seemingly assumed that the coefficients, B , are the same. But no, this was a particular indirect utility function for a household-- a given household. We have to start putting " I " in the utility function to decide. And if these alphas were different across different households, we would need an αI on them and then this thing, the Gorman form, would be violated.

So the first part of the positive representative consumer was just saying a definition of aggregate demand, namely, it's the sum of the demands over all the individual households at price vector, p , and Wealth's W . And the Gorman part is to find that demand-- this guy here-- depends on the vector of wealth's in a particular way. Namely, it's just the sum of them.

Or in other words, it's enough. We should be able to find a utility function and give this mythical representative household all the wealth in the economy and have that household maximize utility, generating aggregate demand, which would be the sum of the individual demands if we'd bothered to keep track of the micro economy and the individual wealth levels. So this statement was necessary, but it's about to be sufficient if the individual households have this particular indirect utility function of the Gorman form with the I only on the A intercept piece, although w depends on I , but that's because it's individual households wealth. Then the utility function we're looking for is just going to be an aggregated version, aggregating up the intercept terms and imposing the same coefficient on the wealth across all the households.

Now do you like it or not? It is going to require that we have these linear expansion paths, if we go back to the micro economy and see how they're behaving. But giving wealth to one guy would increase utility and consumption in exactly the same way that consumption would be decreased for the other one. So moving wealth around because the slopes of these lines are the same, they're parallel lines, moving wealth's around has no effect on aggregate demand. It's a restatement of the positive part of the representative consumer.

If you're listening to macro economists, worried a lot about marginal propensity to consume out of wealth and out of transfers. They're trying to find transfers to particular constrained households in a way that they would eat a lot as if there was a special virtue of pumping up aggregate demand. Well they don't have Gorman in mind because aggregate demand is not determined by the distribution of wealth under Gorman. It's determined by the aggregate wealth.

So it wouldn't matter who you gave the money to, if it were Gorman world. Positive representative consumer. The proof of it I'm not going to go through. It's basically using Roy's identity over and over again and exploiting the linearity. Adding up and then taking it back down. There's examples of Gorman aggregate preferences, allowing variation in the intercepts. So households do not have to be all alike.

And the normative household, normative in addition to being positive, there's another definition, namely we're looking-- we have the underlying micro economy with the heterogeneity across all the households. We're looking for this fictitious utility function, U^* , which we label the utility function of the normative representative household to satisfy two properties. Namely, if we want to see if something x and y is Pareto Optimal, then we can search for alternative allocations that might Pareto dominate. In the normal micro land of all these different households, you have to check out each and every one to see if they're better off under the proposed alternative.

With Gorman normative, you don't have to do that. You just sum up all the demands in the baseline and in the alternative and evaluate that sum under this mythical utility function, U^* . So if something predominates, then it has to have a higher return under U^* . So that makes the search a whole lot easier.

And the flip-side is the same with an important caveat. We start with a baseline and you find a prime allocation, which is dominant under U^* , after you sum up. Then we should be able to do better. So there is a way in the micro space to assign allocations that add up to the prime bundle that are not necessarily the actual prime allocation. A way to do that in such a way that everyone is at least as well off and some people are better off.

So the point of this is, it's very strong when it works, it's a way of evaluating allocations. But it is not to say that you don't have to worry about the distribution. You could make people better off, as I've said before, I think when I gave this lecture, free trade could make everyone better off. But you shouldn't stop there because even under Gorman, you have to redistribute wealth in such a way as to let it Pareto dominate the tariff equilibrium. Maybe that's a less that trade theorists kind of forgot to emphasize

So we did this in-- we actually did the proof in the space of indirect utilities. And again, I don't want to go through that too much. So at the end, we have the application of those positive and normative representative consumers to this slide. This is the third time that I've shown you this economy. And at least maybe the third time, we talked about the Pareto problem in this economy. This Pareto problem, of course, has the lambda weights, as well it should. But if we have a positive and normative representative consumer, then amazingly, we can drop the Pareto weights and just maximize U^* .

So it's as if there's a single representative consumer with an aggregated production technology, like Robinson Crusoe. And this is a much, much easier problem to solve. Hence the power of it for macro because we're solving it in the space of macros. And again, typically, people forget to think about the underlying distribution and the underlying wealth distribution of welfare gains and maybe even losses so stick with the aggregates. So are there any questions about lecture 18?

19 went the other direction. And I paired them deliberately. Instead of saying, look at all the special things you can do if you assume special structure, like linear income expansion paths, let's try to avoid any restrictions whatsoever on utility and production sets and so on and still see whether we can say anything at all. So the spirit of this is very different from the previous lecture. And this is about not assuming Gorman aggregation, or not even assuming anything much at all about expansion. Could have given goods.

What, if anything, can we say or given data, can we just always accept that there's an economy that would have generated the data. And we've talked about this at the beginning of the following class. So we want to put as little structure as possible. And we did this a couple of ways with an infinite amount of data, with a finite amount of data. We did it in partial equilibrium under consumer choice and then we made some connections to general equilibrium theory.

The Slutsky thing came from a property, namely, that we minimized the expenditure necessary to achieve a certain level of utility. That expenditure function is concave. And then the derivative of the expenditure function-- a bit like Roy's identity, not quite-- is the Hicksian demands. So the second derivative of the expenditure function, a matrix, really of cross partial derivatives has to be symmetric and negative semi-definite.

This statement is what led to the Slutsky equation. This is a version of that equation. I guess for the purposes of the lecture in terms of what's the point of it all, one is we get restrictions from the data. Assuming preferences are regular, we get restrictions on something we never see. Namely, these compensated Hicksian demand functions, as I think we went through in lecture 3. But fortunately, we're able to write out the left hand side, these Slutsky derivative, cross partial derivatives, in terms of observables on the right hand side, which are directly from the consumers optimization problem.

How demand changes with prices, how demands change with wealth. And of course, we would be observing the actual allocations. This is with infinite data. So again, we're allowed conceptually to think about having all these derivatives for infinitesimally small changes available to us. And again, as I went through in class from lecture 3, when we did the income and substitution effects we had written with the S on the right hand side and $D_x D_p$ on the left hand side. So there was a negative sign. But it's all the same. It's all equivalent.

So then we're saying Slutsky would be satisfied and a couple of other things are satisfied if we have the premise of an underlying rational consumer maximizing utility. And we reviewed this in the discussion of the subsequent class. We have homogeneity of degree 0, and Walras law. Spending your budget, no price solution and spending and being on the budget line that plus Slutsky are the restrictions implied by having a utility maximizing rational household. And it turned out that it goes the other way, as well.

That is to say, if you had data that satisfied these three properties, then there would be an underlying utility function that's quasi-concave increasing and so on that would have generated the data. So this seemingly bears some similarity to Gorman in the sense of could we generate the aggregate demand by maximizing a utility function? Could we find the utility function, U^* , such that given the household all the wealth, the solution to that problem would be the underlying demands.

Here though, we're giving the data, including demands, for a given household and trying to find whether or not there exists a utility function that will generate demands with all these properties. And the answer is, yes, that it's sufficient and necessary to have these properties. So the theory has content, you just need to scroll through this list in the data and see if you can violate one of them. And it could not have come from a maximizing household.

Convexity was not testable, not with this demand theory because you generate the same data and we actually made a connection to computer science. So this part, this is testing with finite data, not an infinite amount of data. And the key here was the beginning, at least, this Weak Axiom of revealed preference, which you should all know-- and again, we reviewed it in the subsequent lecture at the beginning-- that pairwise, if we have two observations prices and demands and situation one and situation two, then a price P_2 , if they're observed choosing X_2 , it better be the case.

And X_1 was available and not chosen, it better be the case that its price is P_1 , where they're choosing X_1 , they don't have X_2 available. Otherwise, they were supposed to have chosen it. If it had been interior in the budget, according to the first statement, it would be utility maximizing. So this is putting restrictions on data. The Weak Axiom Revealed Preference is looking at pairwise observations as we move prices and demands around to see if it's quote, "rational". Weak Axiom of Revealed Preference.

And then we generalize that. We call those Directly Revealed Preferred pairwise comparisons. And then we went on to talk about Indirectly Revealed Preferred, or just Revealed Preferred. And the content of this was, we never see all the pairwise comparisons. We never see X_T directly compared to X_S , but we do see it indirectly because we have X_T , in this case, Reveal Directly Preferred to X_{R1} , X_{R1} reveal preferred to X_{R2} . And so on down the chain.

So under one of those properties of a rational consumer, namely transitivity, if we had a sequence of Reveal Directly Preferred and under transitivity, X_T should be revealed preferred, directly preferred to X_S too. So when you can do the pairwise comparison, a TNS, it better be the case they don't violate that X_T is preferred to-- weakly preferred X_S . So this second part is, assuming underlying consumer rationality in order to set up a situation where we can actually impose the test.

Then we did that in computational economics, which I mentioned I elaborated on in the subsequent lecture. And we took it to general equilibrium. So are there any questions about risk sharing? Any questions about lecture 19? These reviews get faster and faster as we approach the end of the hour. But on the other hand, it's more and more recent, so not so worried. 21, we already did. So that just leaves us with 20. And 20 is pretty familiar, which was the failure of the welfare theorem. So we kind of covered sufficient assumptions already.

And then for the fourth-- third or fourth time-- I'm showing you this economy again in all of its glory, with the convexity and concavity in the appropriate spots. Statement of the second welfare theorem, you actually saw this slide before in the review today. And the second welfare theorem says that any Pareto Optimum can be supported as of Walrasian equilibrium. So here's a case where everything's working out fine. There is a Pareto Optimum, obviously.

This is the supporting hyperplane, again. So it's going to be a competitive equilibrium. But the second welfare theorem has sufficient assumptions to convexity. So as I said, I think last time-- or in answer to this second to last time-- the starting point here is something that's Pareto Optimal. And it's easy to lose track of the direction of the argument. Is this Pareto Optimal? Yes, it certainly is, or at least you can imagine so. This is the indifference curve that's tangent, locally. And you could try to move away from there, globally. But whatever you do, it has to be on this production set to be feasible and that's going to be associated with a lower level of an indifference curve.

So this is the best level of utility achievable within and on the production set, and likewise over here. Oh, but then the second, I'm forgetting the second part of the theorem is, can it be supported as a competitive equilibrium? And the answer is no because given this imagine supporting hyperplane, the household would be better off moving away to some extreme. Have a higher utility. So it's not going to be consistent with utility maximization. And this one is not consistent with profit maximization, even though it is still a Pareto Optimal point. As you could construct, graphically.

So I showed you this picture earlier. This is kind of an aggregated version of it. It says those assumptions are sufficient, they're not necessarily necessary. The first welfare theorem as a competitive equilibrium is optimal. This one assumes the existence of a competitive equilibrium. All we need is rationality and non-satiation. But this was a picture showing that if we don't have non-satiation locally, we're not necessarily going to show that the competitive equilibrium is Pareto Optimal because we're given the competitive equilibrium we can find something Pareto dominating.

Then we went through this pollution example, which is not allowing trade in all the goods. And again, the intuition for this miraculous price system when it works means that households only have to evaluate things on the margin at those prices. They internalize through the prices and make the right decision. The prices reflect the social trade off of increasing or decreasing demands. And of course, they're much aware of their utility consequences, their marginal rate of substitution.

So if we get it right in a competitive equilibrium in complete markets, and the price system conveys all the right things on the margin and all the people have to do is take from God this price system from Adam Smith, you might say, as if by an invisible hand, namely, the price system be guided to Pareto Optimal allocations. But likewise, if we screw up the markets, then the whole thing falls apart. And the solution example was where we didn't have a price for all the goods, so we don't have a rate of substitution of one good for the other in the market and households and firms make efficient choices.

So we did that with the pollution as a leading example. We went to arrow to generalize it, where households care about other's consumption-- either they're altruistic or envious. And the same thing kind of works. This is the fourth or fifth time today that we talk about a Pareto problem and associated first order condition generating shadow prices. So we're following exactly the same steps as in the second welfare theorem and finding a way, despite the externalities and utility functions, to support the Pareto Optimal allocation as a competitive equilibrium.

We already talked about failure when we have an infinite number of goods. This was an example of how to Pareto dominate an allocation by, essentially, giving the first household some goods the second had and then the third giving to the second, and so on. This is exactly what's going on in the overlapping generations, monetary economy of lecture 21. Money is worthless. It's just pieces of paper. There's no intrinsic value, but if household two here came up with money and gave it to household-- took money, sorry, from household one and exchanged it for goods, then one gets the goods and two gets the money and then two can pass the money along to household three.

So there is a connection between this artificial example and the fact that we get a non-optimal autarky equilibrium dominated by a monetary equilibrium in the overlapping generations model, which we talked about at the beginning of class. OK. So essentially, that was it for lecture 23. And we already reviewed lecture 21. So questions?

I enjoyed teaching the class. I really enjoyed the interaction with you guys. You're really doing well. It's a great joy to me. I hope you like this way of learning economics and you're really quite well trained now. Thank you very much.