

[SQUEAKING]

[RUSTLING]

[CLICKING]

ROBERT M. TOWNSEND: Let's think about where we are in the class. You should have the calendar in front of you. So we're at lecture 7 today. This is kind of a big moment because we're fully embracing general equilibrium although you've seen versions of it before. The other thing would be the reading list.

For today, there is not a lot of new material on the reading is, not a lot of starred materials. We're going to do Pareto optimality. And the assigned starred reading is some of the sections in Kreps, which is available on the Canvas website, so not much additional stuff to read. We're going to do risk sharing at the end of it and then look-- get in more into the applications in lectures 8 and 9.

In terms of the study guide-- so, again, let me ask you a few questions. We did dynamic optimization last time. And the first question is really the motivation, but not everything-- the motivation and the economic problems we're trying to solve are as important to the class as the tools and techniques that we're learning. So let me ask you, what are the observations on storage on carryover of last year's crop to this year's harvest that motivated that lecture that we did last time? Wang Peng, can you answer that?

AUDIENCE: [INAUDIBLE] sorry.

ROBERT M. TOWNSEND: Guanpeng can you take a stab at this question? What are the observations on storage carryover of last year's crop to this year's crop that motivated the lecture from last time?

AUDIENCE: I believe the observation was that it wasn't necessarily that common for-- so I think it was like the crops that were carried over generally there's two ways to use both. One is carried over to be planted for next year and carry over for consumption.

ROBERT M. TOWNSEND: That's true. Those were the two ways to quote, "store the harvest." Did they carry it over for consumption very frequently?

AUDIENCE: I don't think so especially because the spoilage rate was relatively high.

ROBERT M. TOWNSEND: OK, yeah. So it was quite rare. Carryover was quite rare. And the purpose of the model and its parameters was to explain why that would be the case.

AUDIENCE: Right.

ROBERT M. TOWNSEND: The title of the lecture was storage seed and starvation. So do you remember how often were they starving to death?

AUDIENCE: It was once every decade or 12 years.

ROBERT M. TOWNSEND: Yes, that's right. Thank you. Good, all right. All right, so that was the motivation for that lecture. And then this question-- although I don't-- you won't be able to write it down for me, but maybe you can say it in words.

What was the notation we used for the return to seed? And in particular, in what way was it linear or at least up to a limit? And in which way did it depend on uncertainty? And let me see, Armando, can you take a stab at that? Are you there, Armando? Armando?

AUDIENCE: Yes. Sorry, my microphone was muted. Yeah, so there was-- I believe it was there was three multipliers, right. It was they each had a random chance of occurring. And it was like a poor harvest, medium harvest, and good harvest. And it was just linear. The harvest, the amount harvested is linear with how much they planted with either a 1, 2, or 4 multiplier, I believe, was it?

ROBERT M. TOWNSEND: That's right. Yeah, it was epsilon times alpha times k, where k was the amount of land planted, alpha was the linear constant, and epsilon capture those shocks you just mentioned low, medium, and high. So that's a perfect answer. Thank you.

So it's blending-- the point of why I'm asking is just to remind you that, oh, because it's linear, its constant returns to scale. And yet it also incorporated the uncertainty. So we're going to come back to that technology when we look at applications of risk sharing in a couple of lectures.

All right, and this next one, write down a flow equation for consumption in terms of the above notation, but I can't expect you to write it down. So could someone describe in words what the constraint is onto the problem in terms of this flow equation? I'll take a volunteer.

AUDIENCE: I remember the consumption is equal to the output of the last year input plus the output of the inventory and a minus how much you want to leave for next year as the seed and how much you leave for next year as the inventory. That was consumption.

ROBERT M. TOWNSEND: That's perfect.

AUDIENCE: And the consumption should be larger than 0, right?

ROBERT M. TOWNSEND: Yeah.

AUDIENCE: For each period?

ROBERT M. TOWNSEND: That's another constraint, yeah. This flow equation you just described so well. It's exactly what I was looking for. So to elaborate, we talked about resources available today, and that's the stuff coming from current harvest plus depreciated storage.

And then also, how much would it be available tomorrow? Well, that is co-determined with consumption. Because you have available resources, but you can use it for consumption or you can save it. And there were two different ways to save. So your answer was very good, and I'm just repeating in slightly different words. OK.

I guess last one is principal of dynamic programming. This is a bit more technical. But, again, I don't want to hide anything from you, especially when we subsequently jump back and forth between finite horizon and infinite horizon problems. Does anyone remember how we converted into two period problems?

AUDIENCE: So we picked the last date and solved the maximization problem using that last date, and then iterated to the next last date until we got back to the first date.

ROBERT M. TOWNSEND: Excellent. That's perfect. And then the next part, the infinite horizon part, is we just keep going like that. We just keep iterating, so that the final terminal date is so far away that effectively they're acting the way they would in an infinite horizon problem.

All right, so that's probably enough time spent on those review questions. Let me then turn to the lecture today, lecture 7. All right, so a brief outline of what we're going to do-- we're going to define this concept of various concepts, Pareto optimality, then Pareto frontier. Then I'm going to show you a programming problem, a math problem so to speak, that allows us to determine, as its solution, all possible Pareto optimal allocations. And then we're going to go and apply it in the case of uncertainty and call the solution the optimal allocation of risk.

So although this slide says utility possibilities set, and that's going to be the bottom line of the slide, this is also a hugely important slide for the whole course because it's the first time, I think, that we've introduced, in notation, all the ingredients we've been covering so far and, in the notation, describing an economy. I went back and looked at lecture 1. It does say an economy is preferences, endowments, and technology. We called it a PET economy as a way to remember. But here it is in notation.

So an economy, script epsilon or economy, is the consumption set and preferences of each of finite number of agents, the production set of each of the finite number of firms, and some aggregated social endowment. So it's a very succinct list, preferences, endowments, and technology. Then an allocation for the whole economy now is an allocation x for each of the capital I consumers, an allocation in the production set for each of the capital J producers.

If we're in finite dimensional Euclidean space and there are capital L goods, then we have L times I plus J dimensional vectors here belied by this relatively brief notation. In other words, we're going to have to specify a consumption vector X_i in the consumption set for each consumer i , a production plan Y_j in the production set. It's capital Y_j for each firm j of the capital J . And we say an allocation is feasible if and only if the consumptions are feasible given resources available, which would be coming from the endowment and from production.

So a couple of comments-- sometimes people will write consumption less than or equal to as if you could throw something away. That's fine. It's just more useful today to talk about it as an equality.

The second thing is this has buried behind it-- and it's easy to write-- all of the business we've been working with respect to negative and positive elements. So for example, some of these Y_j 's are negative. I refer to it as production as if it were all positive, like grain or something. But if an input is negative over here, then there has to be some offsetting element.

Namely, consumers supply it, or it could be more evident coming from their endowment. So this is where the accounting quote of all the positive and negative stuff is kind of buried. But it is easier to remember this way, because it looks like you just have resources available from endowments and production meeting up with resources used in consumption.

OK, now given the same economy, we can define the set of all possible utilities, utility possibilities set, which is this script U and which is a specification of utility for each and every household. So this has the dimension capital I , which is the number of households. And it's a vector of that dimension. And we're going to say that a vector \bar{U} is in this set. And if two things are true, first of all, there has to be a feasible allocation generating it or generating something with at least as much utility.

So a vector over here, say, the i -th component of this vector, has to be able to be generated by part of a feasible allocation, which is specifying X_i for household i , which, when substituted in the utility function, generates a utility level at least as great as the target, namely \bar{U}_i . Again, it would be a little more intuitive just to put equality here. But what this allows us to say is, if you can generate a utility point through a feasible allocation, then all the utility points which are southwest of it are also feasible. So it's as if you could throw things away.

So in words, this utility possibility set is a set up all utility profiles for households in the economy that can be attained by feasible allocation, a feasible allocation with greater than or equal-- small typo-- so you can generate v . In other words, if you can generate U and v is less than U , you can have v in the production possibility set as well. Because whatever that generated U or something generating utility greater than or equal to U will generate utility greater than or equal to v , all right?

So given a utility possibility set, script U , we define the Pareto frontier, P of U , as the boundary. So it's just the set of subset of utility vectors such that there is nothing northeast. If U is in the utility possibility set, there's no U' in there, which would generate at least as much utility for all households i and strictly greater utility for some households j . So if the right hand side were true, then you could quote, "Pareto dominate" the vector U . And that would eliminate it from being in what we want to call the Pareto frontier.

So here's an example and then a picture. Let's simplify our economy and get rid of production, although we will still have endowments. So this is a pure exchange economy. There are only two goods. And there are two consumers. L and I are equal to 2.

And we're going to say, for simplicity of notation, first good is denoted X . The second good is denoted Z . So the first household has the utility function over X and Z , which is just linearly additive in the 2. And the second household is very similar, except it is multiplied by 2, but it's still linear, additive, and so on.

Now, each has equal endowment of both goods. And the endowment of household one is equal to the endowment of household two. So $1/2, 1/2$ everywhere means there's 1 unit of good 1 in total, although it's split initially equally across the two households, and 1 unit of good 2 available, also split equally. So a feasible allocation over these two goods X and Z has to sum up properly.

So if we take the first good and add up what household 1 is getting, what household 2 is getting, it's got to equal the endowment, which is 1, likewise for the second good, Z . What the first guy gets and the second guy gets in allocation-- it's not utility-- has to add up to 1, the sum of the endowments, OK.

So it can be shown-- and I will show you the picture-- that the utility possibility set consists of all points U_2 , which are less than 4 minus $2U_1$. And the frontier is just the outer portion of that. Without doing all the algebra right now, you can just imagine, for example, what would happen if we gave the entire social endowment to agent 2.

Agent 2 would then get 1 here, 1 here. And you get 2 plus 2 equal 4. And you see this 4 down here, which is the maximum utility that agent 2 could get. And the fact that it's minus means moving away from there. We're giving agent 1 some of the social endowment and taking away from household 2. And since 2's utility level's twice that of household 1, then every unit utility that 2 gives up counts twice for household 1.

OK. So again, that's just to be motivational without actually going through all the algebra. Here is what the utility of possibilities frontier looks like. It's linear, again, with 4 over here as the maximal point. This is a slope of minus 2. It's negative because you're giving up the social endowment away from 2 toward 1 as you move along this line. What's not drawn in here very carefully are the boundaries of consumption and so on, which if X_1 and X_2 , Z_1 , Z_2 cannot go negative, then we shouldn't be drawing in this thing that extends beyond the first orthant.

Pareto optimality, this is the big one, but it's almost obvious already. An allocation x and y for consumption and production is Pareto optimal if it's feasible and there's no other feasible allocation x' , y' such that the prime is at least as good from every agent i 's point of view and there's at least one i' household agent that really likes prime better. And we call these allocations the Pareto set, the set of Pareto optimal allocations.

Well, at one level, it's kind of intuitive that those should be called optimal because you can't make anyone better off without making someone worse off. If you could make everyone better off, then it wouldn't be Pareto optimal to begin with. A little more telling, this definition doesn't give a darn about the distribution of income. So there's no normative statement here about equality and the income distribution or anything like that.

So that's a little bit jarring. We could give all the allocation to household 2 and none to household 1, and that would still be Pareto optimal. We can look at the distribution of income, and we surely will often in the subsequent lectures and in some of the pictures about to follow. That means, if you want to talk about the distribution of income and put some normative criteria for that, you're going to need more than this definition of Pareto optimality.

OK, so let's do a two-person economy. This is literally putting the pieces together. We did for households, utility functions and preference orderings. And this puts consumer 1 with two goods, changing the names to a and b now, nice, smooth, regular decreasing marginal rates of substitution. Oh, and household 1's endowment is here, mark e , which means a certain amount of good b and a certain amount of good a .

So there are two households, but we're going to make this guy stand on his head. So the 0 point for household b is here. And more of good b moves downward. More of good a moves leftward.

So these are regular, nice, smooth, concave, quasi-concave indifference curves preference sets with more preferred to less. And here's the endowment of agent 2, boom, boom-- certain amount of good 2, a certain amount of b , a certain amount of good a . You with me so far? I think this came from Kreps, but, you know, the Edgeworth box is a big part of economics.

Now, I'm going to take this point e and start moving this whole figure for consumer 2 to the point that these two e 's correspond and are identically lying on top of one another. So that's like moving this edge of what's about to become the box over in this direction. And here, although e should have been ideally put on this picture, this would be the Edgeworth box.

So now, we've got household 1 down here, consumer 2 here, and the indifference curves of both. The endowment, say, here, is denoting the good a assigned initially as an endowment to consumer 1, but the other direction here is good b being assigned to household 2. So things add up. The total adds up except you're moving-- the amounts are signifying different directions for the two consumers because consumer 2 is upside down.

So are there questions about the construction of the box? All right, so without putting the endowment in there, this would be red indifference curves for household agent 1, bluish indifference curves for consumer 2. I call them consumers, households, and agents. I'm not being consistent. They're all called consumers in this diagram. And you can see these indifference curves cross.

And I want to claim, going back to the definition, that this cannot be Pareto optimal. Why? Because there are plenty of other feasible allocations that make at least one household better off and the other no worse off. And in fact, there are many allocations that make them both better off. This whole region here lies above the red indifference curve for household consumer 1 and above the blue indifference curve for consumer 2.

Here, it's a different picture where the two baseline indifference curves are tangent. And in contrast, this is Pareto optimal because I challenge you to find an allocation that makes one better off without making the other one worse off. If you're going to make consumer one better off, you have to keep track of these red indifference curves. So we would be on or above the red indifference curve depending on whether we're making them no worse off or strictly better off. And any of those points other than the starting point actually makes consumer 2 worse off.

So now, we can talk-- this is still called Edgeworth Box, although it's starting to look more like a rectangle. That's because there's more of good 1 than there is of good 2. It's wider than it is high. And what's being traced out here, like this is from the previous slide, that there are many other allocations which are also Pareto optimal, where those indifference curves are tangent. So this dotted orange yellow brick road, so to speak, dotted line, is a set of all Pareto optimal allocations in the whole economy.

On the one hand, it's a lot less, a lot fewer, allocations than any other feasible allocation. Any allocation in this box is feasible because resource allocations add up to the total endowment. So [INAUDIBLE] worse off. And again, you can see these extreme distributions, where, if we could go that far to 0,0, consumer 1 over here is in a terrible situation. And consumer 2 just got everything. And that's feasible, as is at this point-- feasible, also optimal.

AUDIENCE: I have a question.

ROBERT M. TOWNSEND: Yes.

AUDIENCE: Does this [INAUDIBLE] have some special properties?

ROBERT M. TOWNSEND: Yeah, it's where, in this sense with these all smooth indifferent curves, the marginal rates of substitution that we learned about are equal across the two households.

AUDIENCE: OK.

ROBERT M. TOWNSEND: You know, the reason this isn't optimal is because the marginal rates of substitution are not equal. So the amount that consumer 2 requires to maintain indifference, giving up good a and getting good b, is greater than what consumer 1 requires to maintain indifference in terms of giving up good a and getting good b. So there's plenty of quote, "gains to trade." Both can be made better off, OK?

All right, now, I mean, one thing we could have done is put the endowment point here and then looked at the set of-- well, maybe imagine this is the endowment point. Then we could look at the set of allocations, which are Pareto optimal. And that yellow brick road would run like that. And so these allocations in the lens shaped region are a subset of the Pareto optimal allocations, which make one or both agents strictly better off than at the endowment. So that's a more refined notion of optimality called the core, which I didn't create a slide to show you. But that's sort of beginning to take into account that you can't make someone worse off than their endowment, whereas Pareto optimality, as I keep saying, does not do that. It's much more general.

OK, Pareto optimality and the utility possibility set, let's define, if it's not obvious already. An allocation of consumption and production is Pareto optimal-- actually, it's a theorem. We defined Pareto optimality already. And we defined the Pareto set already.

So this puts them together in a rather obvious way, namely an allocation x, y is Pareto optimal if and only if the utilities that are generated under x are on the frontier of the Pareto set, which we denoted P script U . So you can go both ways here. It's if and only if the direction from Pareto optimal to being on the frontier is left as an exercise. It's pretty much as trivial as the one we're going to do, but we'll do one of the many way.

So let's go in the opposite direction. That is we start with an allocation on the Pareto frontier, then we want to show that it's a Pareto optimum. So this is really almost just a test of remembering the definitions. Suppose that it were not Pareto optimal, then we'll do a proof by contradiction. So if it's not Pareto optimal, then there has to exist, by the definition, some other feasible allocation denoted with a tilde such that utility is at least as high under the tilde allocation as it was under the original and strictly greater in utility for some household j .

But then these vectors of allocations Pareto dominate the original one and, hence, the original one could not have been in the Pareto frontier. That's one way to remember that is that, if it's on the Pareto frontier, there's nothing northeast of it. We just constructed something northeast of it. That's the contradiction. So this direction is now proved. And separately after class, you can take a stab at the other direction.

OK, so the moral of the story here is we can go back and forth between the Pareto frontier and Pareto optimal allocations. And even I have to be careful not to intermingle them because I know in principle they're the same. And here's the proof that they are.

Next step, we defined a welfare function for the entire economy, and we make it linear. So this big welfare comes from the utility vector across the individual agents. And this welfare function is linear when it's just a weighted average of the underlying utility components with weights λ_i . Weights good in principle be 0, somebody doesn't count at all. Otherwise, the weights are positive.

And when I say we take a weighted average, each of the weights might as well sum up to 1. Because if they didn't, we could add up whatever they do total to and divide through the same objective function. And if you're going to maximize one welfare function, you maximize the scalar multiple of it.

So the idea of this welfare function is a way I don't like, but then I'll tell you what I do like. It's a way of aggregating utilities across individuals in society as if there were some social planner who determines who counts more than others. And particularly under the linear welfare function, the lambda weights are the weights on the individuals.

Now, what I do like about this is this math problem, as it were, which, if there were a social planner, we would be maximizing the weighted sum of utilities subject to allocations being in the utility possibilities set. So it looks like we've ignored feasibility and everything, but no. The feasibility of an allocation is loaded into the definition of the utility possibilities set script U , and we're constrained to pick utility vectors which are in that set. And we're going to pick them in such a way as to maximize this lambda weighted sum of utilities. And we can call that the planner's problem.

So the big result is going to be-- let me say it in words-- that, with certain relatively weak assumptions, any Pareto optimal allocation that we pick will be a solution to the planning problem for some weights lambda. And any solution to the planning problem for given weights lambda will be Pareto optimal. So the conceptualization thing here, before we get lost in the notation, is that we can now go back and forth between the notions of Pareto optimality and the idea of maximizing the social welfare function, although the weights can vary.

So let's look at the math. Suppose we add a solution to the planner's problem seven. And just to make it easier, suppose the lambda weights are all strictly positive. Then the first claim is that solution to seven must be in the Pareto set. And if it's on the Pareto frontier, sorry, then it must be Pareto optimal.

Going the other way, if the utility possibility set is convex and closed, then picking any utility on the Pareto frontier and, hence, Pareto optimal, we can find particular lambda weights that makes that utility vector a solution to the planner's problem. All right, so let's do the two parts a bit.

For the first part, suppose we're given some U^* that solves the planning problem at positive weights, but claim by contradiction that it's not on the Pareto frontier. Well, if it's on the Pareto frontier, we can dominate. There's a U' vector, which is no worse for all i and strictly better for j . But since the lambda weights are strictly positive, we've now found a new vector, the prime vector, such that the weighted average of the prime utilities is strictly greater than the weighted average of the starting point.

Hence, in more cryptic notation, the dot product λU is strictly greater than λU^* , but we started with the claim that we've got a solution to seven. And we just found a feasible solution to seven, which gives a higher welfare number. So that's a contradiction. You with me so far? OK.

So suppose that the utility possibility set is convex and closed, and take some solution U^* , which is on the Pareto frontier and, hence, Pareto optimal. Take a vector U^* , which is Pareto optimal. It's got to be on the boundary.

And because it's on the boundary of the utility possibility set, in turn is convex, we can apply something called-- and I'll show you this momentarily-- the supporting hyperplane theorem, which is going to give us, guarantee the existence of, a lambda vector such that at that lambda, $\lambda \cdot U^*$, is strictly greater than $\lambda \cdot U$ for any other U in the utility possibility set. In other words, λ^* at λ , the particular λ , $\lambda^* \cdot U^*$ is the highest possible, cannot be dominated by in some senses weakly, the highest possible lambda weighted utility that you can get running over all possible other feasible utilities in the utility possibility set. That means that U^* is a solution to the planning problem by definition.

So that completes the proof of going both ways, but it kind of begs the issue of what the heck is this supporting hyperplane theorem. So let me show you some pictures. And this is, along the way, an extremely useful tool. It's just the first time that we've needed it.

OK, so here's the utility possibility set. Here's the U^* that we picked. It's on the Pareto frontier, so it's Pareto optimal. And that theorem is true, so I should be able to find some lambda weight such that, when I maximize the lambda weighted sum of utilities, I would be at a maximum at U^* , which visually you can see is true. There are other points that are Pareto optimal, but they would not be maximal under the lambdas, which are implicitly the slope of this red line. These are level curves. We could call them iso-welfare curves.

These are level curves of the social planner at particular lambda weights lambda. And it's kind of tradition to denote the slope by a line, which is right angled to the line in question. So this lambda going up this way is just a way to say that the slope of both the Pareto frontier and the iso-welfare curve are the same and equal to λ_1 over λ_2 .

All right, so that's a picture of why it's true that, when you pick a point on the Pareto frontier, you can find lambdas which makes it a solution to the programming problem. What is this hyperplane? So we'll do everything in R^2 . Here's a line in R^2 .

And in higher dimensions, this is called a hyperplane, fancy language. You could probably draw it as a plane in R^3 . It gets harder to draw, unless you're better than I am, in R^4 and so on. But we'll stick to R^2 .

So here's a line. It's called a hyperplane. And we can obviously talk about the half space above the line and the half space below the line. So above the line means, if we specify the intercept c and the slope p then the sets above the line, the half space above, is a set of all values z which gives no lower value than c at slopes p and vice versa for the lower half space.

Here's another line that separates these two sets. We could call one of these set A and the other one set B . Note that they are both convex sets. Pick a set. Pick the pink one, pick any two points, then weighted average of those points are, in turn, in the set.

This is a separating line in the sense that the half space above this line contains the pink convex set is contained in that. It's within the half space above. And the blue convex set is part of the half space below. And that's what all this stuff up here says.

And the culmination of this is called the supporting hyperplane. So again, we run from hyperplanes to separating hyperplane theorem having to do with convex sets to supporting hyperplane theorem. So this pick this set B . And let's require that it be a convex set.

Pick this point x in the set and on the boundary. Then the claim is we can find a hyperplane, in particular a supporting hyperplane, which just touches the set B at x and is tangent to the boundary at x . Are you with me?

So the slopes of this hyperplane are determined by p . And we don't need a constant c anymore. This is said to be a supporting hyperplane in the sense that, if you picked any other x in the interior other than the one on the boundary, any other y in this case in terms of the notation, it will have a return which is less than-- $p \cdot y$ will be less than $p \cdot x$.

So this is what we were using a moment ago, this separating hyperplane theorem. Because I kind of asserted we could draw this picture, and it seems really natural that we ought to be able to. It is using the convexity of the utility possibility set and picking a point, like x was three slides away, and then drawing this hyperplane which is tangent. And the separating hyperplane theorem says we can always do that under these weak assumptions like convexity.

That also means we can pick another one, right? If we pick another point on the frontier, we can find a separating hyperplane. But it's clearly, in this diagram at least, going to have a different slope. So we can support it. We're just going to have to change the λ s. And that's what's going on here.

This was the initial point. We could have picked another point. If we did that, we'd have to find new slopes for the supporting hyperplane, so that it would kind of pivot and allow it to be tangent at this new predetermined point.

So again, not to get lost, this is how we're going back and forth between Pareto optimal allocations, which are on the Pareto frontier, and finding a λ such that at those λ s we would be maximizing a linear welfare function. And any Pareto optimal point which is on the boundary has to have that property. And likewise, the first part, which was anything that is maximal under the welfare function, must also be Pareto optimal. And the proof there that we did was by contradiction.

All right, so in notation, we can parameterize Pareto optimal allocations by choosing Pareto weights. And this λ and δ sort of means there in the simplex, which, again, without loss of generality, the λ s are non-negative and they add up to 1. We used simplex earlier for lotteries. Here's the second time that we're using it.

OK, so if we want to maximize a particular λ weighted sum of utilities, λ_i weighting U_i , we can do that and find the maximum to it and call the solution X^* and Y^* . But here, we're going to earmark them by λ because it was a solution at particular λ s. This is just a repeat of consumption and production sets and feasibility.

So any solution that maximizes we earmark by λ , this is a particular Pareto optimum. And also, given any Pareto optimal allocation, we can find a λ which makes it a solution. So this is just restating in words what we already determined.

So the task of finding Pareto optimal allocations in the Edgeworth box is simply solving this math problem. So let's apply it. Suppose we had this date space with uncertainty and this sort of, Debreu tree about the set of possible outcomes at date 1, date 2, and date 3.

And we're going to index consumption by the history of all shocks leading up to that particular date. And this was expected utility. And this is a review slide. We went over this last time or at least earlier-- not last time, but the time before that.

So let's determine an optimal allocation of risk in this economy that has states determined at random. So we're going to want to-- let me say it in words. We're going to maximize a lambda weighted sum of utility subject to resource constraints. Note we're not going to do this in the space of utilities. We're going to do it in the space of underlying allocations, but it's the same thing. Instead of going from allocations to utilities and from utilities to maximizing a lambda weighted sum, we'll just kind of directly search over the underlying allocations.

It's small typeset. You may be squinting to try to see this. Here was a repeat from the previous slide, the utility function, with a little more notation. So this is expected utility as we studied in the second lecture on preferences and utility.

This counted by beta, as we kind of did in the storage problem, and weighting outcomes by probabilities. The summation here is over all dates and all states, as in this tree. And then we take those expected utility outcomes and pre-multiply by the lambda weights, as we've now been trained to do. And because we're not in utility space and we're in space of underlying allocations, we have to make sure that consumptions add up to the total.

Now, here's, again, a bit of a switch of notation. These Y's are now-- these right-hand sides are using the Y for endowment, not for production. But it's not that misleading. So we have an exogenously given set of endowments depending on the data and the state. And then consumption has to add up to that or certainly could not be greater than that because it would not be feasible. So we call this the Pareto problem or the planner's problem for, here, particular weights lambda.

What does the solution look like? First of all, let me tell you intuitively what's going on, and then we can generalize it. We wanted to maximize the lambda weighted sum of expected utilities. All I've done here is pull out in front of those expressions the thing that is common across all the agents i , namely the beta and the probabilities. Otherwise, it's the same expression. But maximizing one is equivalent to maximizing the other because this is just some scalar constant.

We also know that consumptions are going to add up to the total social endowment, which is the sum of individual incomes, which was I think I forgot to mention at the bottom of this slide. This is the aggregate. Social endowment is just the sum of the individual endowments. So there's lots of dates and lots of states, but for every date and state there's a particular social endowment.

So we can simplify this so that it looks like a static problem. We can maximize each component one at a time and just respect the resource constraint. In the objective function, a component one at a time is just the lambda weighted sum of utilities. And this would be the resource constraint that consumptions have to add up appropriately. So amazingly, we've sort of been able to abstract away from dates and states and about to determine an optimal risk sharing rule as long as we keep track of the aggregate social endowment. Questions?

All right, so another tool we can use that you already know is the Lagrangian. So we'll maximize this lambda weighted sum of expected utilities, and we'll throw in the constraint sets with Lagrange multipliers where y is greater than and is going to end up being equal to the sum of consumptions. So these thetas are the quote, "Lagrange multipliers." And you can think of them as shadow prices, like the marginal utility of income. But there are lots of them because we have a theta for every possible date and all possible histories of states leading up to that date. So there's a large number.

Anyway, max away, we get, for a particular date and a particular set of states, st . We have the derivative of the utility function with respect to its argument preceded by beta and lambda. So we get this thing, theta to the t because it's at beta t on the left-hand side. And on the right-hand side, we get-- where else does it enter? Here. So we pick up the Lagrange multiplier.

And then this is just a repeat of the constraints. You could write it out more formally with theta times the sum being 0 and so on, but the problem is implicitly assuming that more is preferred to less, resources are scarce. So you're not going to throw anything away.

These first order conditions look like lambda times the marginal utility of consumption for household 1 is equal to lambda i times the marginal utility of consumption for all other households i . And this is true overall, for every i . And this characterizes solution for all possible dates and states where the aggregate social endowment is a parameter.

So there are how many households? Capital I . So we have I minus 1 equations here. We pick up one more equation here. We can solve this thing and denote the consumption of household I to be some function g of the aggregate social endowment. This is a Pareto optimal allocation. It's certainly a subset of the set of all feasible allocations, and it has to do with these weighted marginal utilities being equated at particular lambda weights.

Because utility functions are concave, strictly concave, if they were strictly concave, then this g is going to be positive with a positive derivative. So there's a monotonicity argument here that, if the aggregate social endowment is going up, then the consumptions of all the agents are going up. How fast they go up as the social endowment goes up is a function of the g .

And the g has an i on it. That i reflects potential heterogeneity in the underlying utility functions. We did not require that everyone has the same utility levels or even the same utility functions. And I'll show you some solutions momentarily.

So we'll eventually get to an application, risk and insurance in village India and then one with production. We will rederive this equation and use it. But the restrictions on the data come from equations like 17. Remember, economic science is about making predictions.

So we have a whole economy. Feasibility is one way to restrict the prediction. Optimality is a further restriction on the prediction. It's a criterion function, a metric we're going to use. We're now in the position to make predictions about what allocations ought to look like at least in this risk sharing setting. And they have to satisfy equation seven. So there are lots of things that are inconsistent with equation seven. And that creates a test. We can go to an actual application and start testing.

So let me just give you two examples. If we started with constant absolute risk aversion-- and that concept was introduced when we talked about choice under uncertainty under the consumer preferences. It's sort of negative exponential, which makes the marginal utility positive. This γ_j , which is a weighting term, is close to absolute risk aversion.

Actually, $1/\gamma_j$ is risk aversion. The higher this γ_j here, lower this thing is. So γ_j itself is the inverse of risk aversion. It's like risk tolerance.

And I'll spare you the algebra. But if you maximized a λ weighted sum of utilities subject to the aggregate social endowment being e , you would get this solution. So this is saying consumption of household 1 is a linear function of the endowment. This is the intercept, and this is the slope.

The intercept has to do with those λ weights. So very naturally, in this case, the higher is λ_1 , then the higher is the log of this ratio. And the higher will be the level of consumption of household 1.

And the only thing left is how those consumptions will change as we change the aggregate social endowment determined over time or randomly or whatever. And it's saying that consumption of household 1 will go up with e at a faster rate. And for household 2, the higher is γ_1 relative to γ_2 .

γ_1 , as I said, is like risk tolerance. It's the inverse of risk aversion. So now, we're introducing some randomness in the economy's GDP. And the issue is, who's going to bear that risk? Who's going to suffer in recessions? Who's going to bear the bonus of being in a boom?

We have risk aversion. So the guy that's most risk averse with the $1/\gamma$ would be the one where this is close to 0, meaning that his concerns over consumption is pretty stable against the aggregate social endowment. And the other guy is bearing all the risk, but that's a Pareto optimal allocation. The riskiest person should bear the brunt of the aggregate fluctuations.

Here's another example. This is constant relative risk aversion, a version of which you saw before. We're going to raise consumption to the power d . This is actually consumption over and above subsistence, as here. So the consumption set has a positive lower bound. And this is the distance from that subsistence consumption raised to the power d .

And when you solve these equations, they're linear. And again, you get an intercept and a slope. However, it's a bit different from before because the intercept contains not only the λ weights, but it contains the a thing, which is related to the b , related to the power in the utility function.

It looks like higher λ is going to lower consumption, but then you've got to remember λ is a number between 0 and 1. We're taking it to a power so you can make it smaller instead of bigger, et cetera. So when you do it right and think about it, it does mean that higher λ_1 is associated with higher levels of consumption, but it's weighted, effectively, by risk aversion and by the subsistence level of consumption.

And now, if we vary the aggregate social endowment, the allocations will change, but they're not going to go up 1 to 1 the way they did before. That's going to be determined by both the λ weights as well as risk aversion. Both those pictures, constant absolute risk aversion and constant relative risk aversion, look like this with different slopes and different intercepts.

And then this is-- I'm not and I probably should have written out the underlying utility functions. But you can think of this as one person having constant relative risk aversion, the other one having constant absolute risk aversion. And then you get non-linear schedules, but they are both monotone increasing as predicted. But now, the slopes vary and the levels are clearly different depending on the aggregate social endowment e . But we can still put restrictions on the data.

So those are three pictures. And we're going to come back to these pictures when we do one of the applications, which is dividing up land in medieval England. We've referred to the several times. We're actually now empowered with the tools to be able to analyze that application.

All right, are there questions? OK. So again, next Thursday we're going to not have another new lecture. We're going to have a review session over the first seven of them, including this one. So review the material, take advantage of the opportunity. Feel free to ask me questions, and I will try to provide a summary of what we've done so far in the class. OK, thank you very much.