

[SQUEAKING]

[RUSTLING]

[CLICKING]

**ROBERT
TOWNSEND:**

All right. So let's take a look at where we are. On the calendar, we're at lecture 17 today.

As I said last time, lectures 16, 17, 18 and 19 are all about Walrasian equilibria and their properties. So we did the welfare theorems, which we can review again momentarily today about existence of all raising equilibria. And while we're at it, how we could interpret the Walrasian equilibrium as the outcome of the game using a Nash equilibrium notion.

And then, we get to these two juxtaposed lectures, 18 and 19. Gorman aggregation has to do with when it looks as if there were far fewer households than there actually are, because we're able to add them up. That takes some pretty strong assumptions that we've already gone through in terms of linear income expansion paths. So we'll review that.

And then, identification is, if you go the other way and assume almost nothing other than rationality, are you capable of rejecting the model? And then, the last two lectures are about failures of the welfare theorems and relatedly, monetary economics. We're moving right along.

In terms of what you should be reading, ideally, I don't if you guys ever do read these start articles. But let me just point out that you should be. And in particular, today, lecture 17 today-- so on the course posting under Canvas, there'll be the Kreps section 6.4 on existence and number of equilibria.

Some of the diagrams from today come from that chapter. Although obviously, there's more in there in the book than there is in the lecture. And also, this course material-- Asu's stuff from a different MIT course. She's got a really great lecture on existence of Nash equilibrium.

And I borrowed that quite a bit. The entire lecture is posted on our website, but we will cover a big chunk of it today in class. So those are the two starred readings.

And then, we get to a bit of a review. So this is a bit challenging, but I'll ask it anyway. It really has to do with stating the first welfare theorem, and then outlining for us the steps necessary to prove it. So Daniel, you want to take a crack at that by any chance? You remember the steps for proving the first welfare theorem?

AUDIENCE:

If I recall correctly, the first welfare theorem was that under the assumption of everybody being rational and [INAUDIBLE]. I think that might have been it. There is a lambda-weighted-- then, the economy will create a lambda-weighted-- it will maximize the lambda-weighted sum of the utilities of the people. And the proof-- yeah, I don't think I remember the steps of the proof.

AUDIENCE:

Can I take a shot at it?

ROBERT Sure.

TOWNSEND:

AUDIENCE: If I remember correctly, the proof goes something like, for each person, because of the budget constraint, any outcome that would be strictly preferred by that person would require a larger budget. And so in order for something to be or not to-- for the allocation, whatever everyone ends up getting not to be free to optimal, it would require everyone getting-- it would require at least-- it would require all the budgets to either be the same or increased. But you can't-- and if you also factor in the price, the profit maximization part, you can do the same thing on the production side. And you basically end up getting an impossible situation, where you have to have more stuff being consumed than being produced because of market [INAUDIBLE].

ROBERT Yep, that's very good. That was a very, very good outline. So the first welfare theorem says that any Walrasian equilibrium is optimal. And the proof is by contradiction.

TOWNSEND:

Suppose it's not optimal. Then, there's something that's weakly preferred by some households and strictly preferred by others. And then, as was just said, the strictly preferred by others means it could not have been available, because it otherwise would have been purchased in Walrasian equilibria.

And then, we had a little lemma that was quite subtle even for those guys who weakly prefer the alternative allocation, they too would be spending weekly more. And so we get the total expenditure summing up over all the households that would be necessary at the competitive prices to achieve the alternative allocation to be strictly greater than what is being spent in the competitive allocation. Then, we add that to profit maximization.

And we get a relationship, a strict inequality relationship, which at the end of the day said that the allocation in the alternative-- allocation that's supposed to be Pareto dominating is not resource-feasible. But the definition of Pareto optimality is that there is a dominating-- if something's not optimal, it can be dominated by something that is feasible. So the feasibility part will be contradicted, and that was the proof.

OK, so you'll see this theorem again when we get to that last lecture on monetary economics. It's going to turn out that under certain conditions like having overlapping generations of people so that no person lives forever, but time goes on forever, that that's enough to make the first welfare theorem fail, surprisingly, when we get into monetary economics. And the proof then begs the issue of where this proof is breaking down. And it's going to turn out that the valuation of the alternative bundle at the competitive equilibrium prices costs an infinite amount.

So we implicitly assume everything is finite in this finite dimension of Euclidean space. Anyway, just to anticipate. Second thing I want to ask about is similar.

Outline the steps of the proof of the second welfare theorem that we did in class last time. So I don't expect a lot of detail on this, but let's see who would like to. Yuen, can you volunteer for that?

AUDIENCE: Sure, yeah. So the second welfare theorem states that any Pareto optimal allocation can be achieved by some price equilibrium, assuming that you can shift endowment point under exchange. And so in the Lagrangian maximization problem, basically, you're finding the actual endowment that will give you the correct price equilibrium corresponding to the desired Pareto optimal allocation.

And you can do that just by first solving the Pareto problem to get the first order conditions reallocations. And then, you can solve the Walrasian equilibrium problem-- so the consumer and the firm side problem separately. And then, you just match the first order of conditions. So the appropriate parameters, like the shadow prices, need to equal the prices in the price equilibrium-- stuff like that. And you can show that all these things exist in our system.

**ROBERT
TOWNSEND:**

Perfect, yep. The key is the Lagrangian method, which under convexity in sets and concavity in functions makes the first order conditions necessary and sufficient. So to find a solution, all you need to do is find variables and shadow prices that satisfy those first order conditions. We have the obvious candidate, because we start with a Pareto optimal allocation.

So we just compare what we need for the Walrasian allocation to what we already have from the Pareto problem-- allocation, and choose prices, and Lagrange multipliers, and so on accordingly. So your answer was perfect. Thank you very much.

All right. So that's so many review questions, because we really focused pretty narrowly on those two theorems in class. So all right. So then, the lecture today is existence.

Sounds ominous. [CHUCKLES] What does it mean to not exist? Well, as per the welfare theorems, I said with the first welfare theorem, if a competitive equilibrium exists and it's optimal, how do we know it exists in the first place? That's one way to put it.

Anyway, it's possible that many things we do are vacuous, because we assume a concept and assume an allocation that satisfies the concept. But we'd better go back and check. Along the way, there's more interesting stuff happening in this lecture, which has to do with reviewing the definition of a Walrasian equilibria and then seeing how sufficient assumptions for existence, but also in principle, how you might want to compute it, which brings us to modern day platforms and online markets.

So it's not just abstract. It's actually something computer scientists and fintechs worry about quite a bit. And relatedly, we'll compare the Walrasian equilibria to the Nash equilibria, which is a different concept of an equilibrium. But it takes the same sufficient conditions.

And again, that's a very, very useful tool to master. But we want to go beyond that. We want to use it. So I'm going to show you some pictures of what's going on in the New York financial markets and date, at least verbally, trade fails, and when those markets get messed up, as in the Coronavirus in March.

And then, we'll think about, are there ways to clean up the markets and avoid the failures that have to do with market design. So that's the agenda for today. Mostly, it's conceptual and notational. And I have deliberately taken out a lot of the notation necessary to look at the actual mapping, because it gets quite tedious, actually.

And it's not really the main point. OK, so here's the math tools. What do we mean by a fixed point of a function? We're going to want to solve nonlinear equations. We want to know whether or not there is a solution.

Those could be excess demand equations in the Walrasian case, or reaction functions in the Nash case. We'll get there momentarily. So a function, f , maps some domain, A , back into itself. And we're looking for a particular value of x^* in the domain such that when you stick it in, you get it back out.

And what I mean by that is, at x^* , f of x^* is x^* again. The more formal statement-- Brouwer's Fix Point Theorem-- is at the bottom of this slide. Let A be a subset of a finite dimensional L dimensional Euclidean space. And that quote "domain" being non empty, convex, and compact," and f maps that domain back into itself, A to A , f is continuous, that's enough conditions on f and on this set A such that there is a special point, x^* , in A , and when you stick it into the function f , you get x^* right back again.

So what's going on is pretty well captured by this classic picture. The set A is just the interval 0 to 1 . And that's the domain.

And the range is the same, A to A , 0 to 1 . And this function f is being drawn here as a continuous function. And the challenge would be, how could you possibly draw this function, which has to be continuous, in a way that it doesn't cross the green diagonal? This one clearly does.

So where it crosses, by definition of the 45 degree line we found, x^* , we get back what we put in. So, oh, no. I don't want to go down there. How about we hug the axis up here?

But even then, where f of x is 1 for all x , it would still be a fixed point at the end point. So anyway, that's the intuition. It's pretty clearly sufficient to have the fixed point.

And when you don't-- I'll show you a picture-- it's because it'll jump over. It will not be defined exactly where it would need to be defined to cross the 45 degree line. This is a generalization of Brouwer's theorem. It's Kakutani's Fixed Point Theorem.

A -- still the same, a non empty subset of a finite dimensional Euclidean space. f , however, doesn't have unique single values in the range. It could be a whole set of values of f . It's called a correspondence.

So we pick x in the domain. There's a plausible entire set of values-- f of x , which is a subset of A . So then, we just look at the conditions on A , which are already given to us to be non empty, compact, and convex, exactly as in Brouwer, f of x is non empty.

Well, that's been implicitly assumed that it has a value for every x . And here's some more-- convex-valued and closed. So although it wasn't about this, the question he asked at the very beginning-- so it's like this.

So f is convex-valued means take any two points in f of x , the convex combination is also in f of x . And closed simply means that if you have a sequence of points, x_n and y_n converging to x and y , with every point along the way in the sequence satisfying y_n , an element of f of x_n , then the limit y is an element of f of x . So again, it doesn't jump.

And then, Kakutani tells us that these are sufficient assumptions for us to have a fixed point. I wish this had been written with a star here. If x^* is an element of A , then x^* is an element of f of x^* , not equal. Because we have a correspondence, just an element of.

So here is the jumpy picture that I promised you-- the two exceptions. F has to be convex-valued and continuous, in some sense, or have a closed graph. Here, it's not convex-valued at this particular point in the domain. It is not true that any linear combination of points in the range are also in the range.

And in particular, there's nothing in the middle here. So it jumps over that 45 degree line. Likewise here, there's a sequence of points-- actually, unique points. Correspondence allows single value functions.

But when you get here, the limit point is no longer consistent with f of x being a point of x . And the function drops down. x is here. f of x is here.

That's fine. That jump is fine. But the problem is, this can be achieved by a limit. And the point in the limit is not in the pair. That's that little Cheerio with a hole.

That's not an argument of f of x . So anyway, this is what can go wrong to motivate the quote "sufficient assumptions." Now, we get to the first application-- existence of Walrasian prices.

We want sufficient conditions. Consider a pure exchange economy. Get rid of production. Have a strictly positive aggregate endowment.

Assume that preferences are rational, locally nonsatiated, and continuous. In fact, preferences can be represented by utility functions. The consumption set is, say, all of the L dimensional non-negative space. So we define three things here.

First of all, we're defining the maximizing demand. So we maximize utility function subject to consumption being in the consumption set, subject to not spending more than the valuation of the endowment at this price, p . That's the definition of demand for all prices p given parametrically.

Z will call excess demand-- namely, not total demand, but the difference between demand and the endowment. This is typically positive if they're buying stuff. It could be negative if they're selling stuff.

And yeah, we just call it excess demand. And then, we take the sum of the excess demands, summing over all the households. This is I , capital I . So this is the aggregate excess demand function.

Now, if we can find a price, p , such that this z of p is 0, then we're done. Why? Because we found the p such that when we march across the different households, we have excess demand being positive for some people. But since it has to sum to 0, that excess demand has to be negative. So whatever gets supplied by households, some households, gets demanded by other households.

By the way, all of this is written in vector notation, although it's not obvious. You have to be reminded of that. We want to be able to clear each and every one of the L markets with this price vector p .

OK, so here's a picture where it works, and a picture where it doesn't. The aggregate endowment is over here. And we'll look, despite what I just said, at good 1.

Take the price of good 1. Here's the demand for it, total demand. This is total supply. Excess demand hits 0 right here. We're done.

We found a price-- this price, where excess demand is 0, because supply is equal to demand. But if it were to happen that somehow those demand curves jump, then for some prices, excess demand is positive and for others, it's negative. And we would not find 0 of that function.

So when can this happen? I didn't draw the picture here, but if you take an Edgeworth box diagram, you can imagine that households might jump from wanting all of good x to wanting all of good y at a certain price. And maybe if you set it up even with more restrictions, they're not very happy with stuff in the middle, either-- for example, with not concave indifference curves, but convex indifference curves. The linear stuff won't work, because they'll be happy in the middle.

But when it's a concave function, or convex indifference curves, or concave upper contour sets, then they will only want to go to extremes. And that's going to create a hole. So I think it's not written out.

But you're going to see it indirectly in a minute. We need sufficient convexity of consumption sets, and concavity of utility functions, and if we had production, some convexity assumptions there in order to get the conditions to be able to use the fixed point. Remember, the fixed point assumes that functions are continuous and the domain of the underlying mapping is closed, and convex, and nonempty.

Now, I could have put many more slides in here. In earlier years, I used to have them. It's surprisingly tedious. This was the work done by Arrow, and Debreu, and McKenzie in the '50s and the '60s. And the mapping is not as straightforward as you might hope.

It's certainly nothing like these pictures. But I think the intuition is fine for today. Now, much more relevant since, we just did it, would be Negishi's mapping. Instead of looking for prices that clear markets, we're going to look for Pareto weights using the first and second welfare theorem.

Again, exactly what we did last time. The first welfare theorem tells us that the competitive equilibrium will be Pareto optimal. And we're going to be having enough assumptions so that that theorem is true-- local non satiation and rationality.

So we know, if it's optimal, there exists Pareto weights λ such that when we solve the λ weighted max problem, we will get that particular Pareto optimal allocation. Not any λ -- particular λ s that generate that particular Pareto optimal allocation corresponding to the Walrasian equilibrium.

So that's Negishi's method. Instead of trying to find prices that clear markets, it's going to try to find Pareto weights-- in particular, Pareto weights such that we're searching as we vary the Pareto weights over the space of Pareto optimal allocations in such a way that the fixed point of a mapping will be the particular Pareto weights that generate the Walrasian allocation. And in particular, the valuation of the expenditures associated with the competitive equilibrium is exactly the valuation of wealth under the private ownership economy.

So those are the words and the notations coming. It's three or four slides away. OK, so what does Negishi do in-- he actually adds up over firms. So we act like there's just one single firm.

That's perfectly fine to do. It just makes the notation easier. And these assumptions are pretty much what we did last time when we did the Lagrangian and the second welfare theorem.

The consumption sets are closed and convex. Let's say they contain 0. That's a little bit more special.

Utility functions are continuous concave functions with local non satiation. We define an aggregate consumption set, the aggregate social endowment, the aggregated production of the representative firm, which is also convex and has a concave transformation, possibility of an action in the production set, and feasibility-- in particular, the existence of an interior point, and that the feasible set of allocations is compact. This compactness may seem strange. And I can come back to that if you want.

Certainly, if you had an Edgeworth box economy, the set of aggregate resources is fixed and finite. Production creates the possibility of creating more, but you have a finite number of inputs, and so on. So in the end, you end up with a compact set.

All right. So Negishi's theorem, 1960-- suppose we have a private ownership economy with this string of elements that you're now very familiar with that satisfy A.1 through A.4, all the convexity, and concavity, and continuity, and so on. Then, voila. There exists a Walrasian equilibrium.

x^* , y^* , p^* . Let's assume utility functions are differentiable. F is differentiable. We did that anyway with the second welfare theorem.

So the key here is that Pareto optimal allocations are indexed by λ . This simple-looking problem, which is, pick a λ vector, maximize λ -weighted sums of utilities by choices of consumption and production that are in consumption sets, production sets, and satisfy the resource constraint for each of the L goods. You saw a slide like this last time with the second welfare theorem. Here, we're going to explicitly index the solution of a given λ as x and y .

So for a given λ , a particular target Pareto optimal allocation is the solution to this Pareto problem. Now again, said this earlier. We know there is a λ , of course, that generates a Pareto optimal allocation that corresponds to the Walrasian equilibrium.

And more generally, from the second welfare theorem, any Pareto optimal allocation hence has a particular λ such that it can be achieved as a competitive equilibrium with transfers. So let's just search in the space of Pareto optimal allocations till we find the right one, find the Walrasian one.

And in particular, the prices are going to happen along the way. They are the Lagrange multipliers associated with the resource constraints. This is exactly, in fact, the same notation from last time, except that this γ , which is the price vector, is indexed by λ when we matched up, remember, the consumer problem with the first order conditions in the Walrasian equilibrium with the first order conditions of the Pareto problem.

And we also will get, from the Pareto problem, the consumer's allocation, the firm's allocation, these prices, again, as I just mentioned, and the wealth allocation. And the wealth for a given λ for a given Pareto optimal allocation is just the valuation at these Lagrange multiplier prices of the target Pareto optimal x allocation. This is a given.

Here is a household I . We're summing over all the L goods, finite number of them. And so again, add λ \hat{x} . Add λ of L -- is the demand for the L th good.

This is the price of the L th good. We sum over all of the L goods. We get total expenditures. And we assign that wealth level.

You may remember that slide. I said, why do we need this? And I kept reminding you of the global algorithm, which is, we have to make sure the assignments of wealth that correspond with supporting prices of the Pareto problem sum up to the total social endowment.

So you saw this before. What are we looking for? We're looking for an assignment of wealth such that the wealth is equal to the valuation of wealth in the private ownership economy. So this is just the endowment of household I plus household I 's claim on the shares of the firm.

You used to have θ_j summation over j , but because for simplicity we assumed 1 representative aggregated firm, we only need θ_I , not θ_j . It's I 's claim on the profits of the firm, the single firm. Anyway, that's private ownership.

So we're looking for λ s somehow. And again, I didn't go much further with the mapping. We're going to iterate on λ s in such a way that eventually, we define the mapping in such a way that the λ that we want will be a fixed point of the mapping and will satisfy the valuation of expenditures at that particular Pareto optimum is exactly the valuation of wealth in the private ownership economy.

Again, the mapping is a bit more tedious than you might hope for. On the other hand, you can see clearly from the second welfare theorem how we were using convexity, and continuity, and so on. So perhaps it's not too surprising, at the end of the day, that we end up with sufficient assumptions to-- oh, here it is.

Here I am talking away. Sufficient assumptions to get this particular λ^* as the fixed point we want. I said this in words, and I should have had this slide on the board.

Valuation of expenditures equals the valuation of wealth in the private ownership economy. So there's really two ways to search for a competitive equilibrium. We can search in price space, or we can search in λ space.

It's curious that in both cases, we need a simplex. To remind you, a simplex is the space where each of the elements is non-negative greater than or equal to 0, and the sum of the elements adds up to 1. So prices add up to 1, because we have a degree of freedom for the normalization of prices, because of homogeneity of degrees 0, and homogeneity of degree 1, I guess.

And also in the λ weights, the λ s are normalized. They're non-negative. And they sum up to 1.

So that's another simplex. But on the other hand, the dimensions may be easier in the λ space. Maybe there are more goods than there are types of households.

We have inputs and outputs of all kinds. There may be, at least for economic modeling-- we can think about a finite number of types of households. Maybe when you think about wealth, we have rich and poor.

We have medians. We have quartiles. When we thought about international trade, we had households differing by their endowments.

Some are laborers. Some own the capital, and so on. So typically, the number of λ s corresponding to the number of types of households is less than the number of prices we need. And it's just easier, therefore, to search and compute in the λ space.

Anyway, I thought you might appreciate this, because we just got done with the second welfare theorem. And this existence proof uses that. Now, lest we think it's just elegant, I turn to what's going on in computer science-- namely, how to find the Walrasian equilibrium.

So Echenique and Wierman at Caltech, actually, have some really cool stuff. It can actually be pretty hard to find the competitive equilibrium. We draw simple pictures, right? But you have lots of households.

And in particular, households could be posting information on websites. And we want to run some auction or something. And hopefully, there are enough households that we're trying to get to the Walrasian allocation.

It's not easy to solve these things. In fact, in general, it looks like a very hard problem. And look, the examples of how to find it look relatively easy. So there's this big gap in the literature.

Anyway, these guys discovered Negishi. Everybody goes back to 1960. And they have an algorithm that works to find the competitive equilibria using Negishi's algorithm that I just showed you.

And this one is related. Computing Walrasian equilibria fast. [CHUCKLES] So again, there's stuff here. Although it's cool in the sense of writing a code that iterates. And you want to know the time that it takes to find the solution.

And hopefully, that's not going up exponentially with the number of households, or the number of goods, or something. So they want it to work fast. So they don't want to evaluate derivatives all the time.

So they figured out this algorithm that makes a call to the oracle. And the oracle, when consulted, will tell them the derivative of the excess demand curve. And so they've got this down.

They only have to consult the oracle relatively few times. And it's polynomial. And so they can-- this algorithm works pretty quickly. This language of consulting oracles is used in computer science any time you want to go to an external data source.

And could be rainfall, or stock prices, or whatever. And then, bring that into the code-- like *The Matrix* oracle. OK, let's go to another concept of equilibrium, Nash equilibria, and define a strategic form game.

Let me back up a second. The Walrasian equilibrium is fine-- say, prices such that when everyone takes those prices as given and maximizes, then we look at the resource constraint, we find the prices such that excess demands are 0. But that does beg the issue of where the heck the prices are coming from.

Even if we have algorithms for finding them, that's a step along the way. The strategic form games are ways to be very explicit about the strategies and outcomes that traders are using. And again, I'll link this up to the Walrasian price momentarily.

OK, so here are some definitions. A strategic form game has a finite number of players, capital I , script I . There there's a set of possible actions for each player, S sub I . A particular action, little s_i , is in big S_I . So that's a particular action in a set of all possible actions that define the game.

And the payoff functions-- look like utility functions-- map the actions into real numbers. But this vector, capital S here, is actually the actions of all the players, unlike, say, in a Walrasian equilibrium, where prices are given and you choose demands.

Here, we need to map explicitly the actions of each and every player in order to determine the outcome, and hence determine the utility of each player. Here's some convenient notation. S_{-I} means everybody but I . So it's the particular actions of all players j excluding player I .

Likewise, capital S_I is the set of all possible actions over all players other than I . And finally, this notation $S_I \times S_{-I}$ is the strategy profile. It's particularly relevant for player I , because it has his or her own action in there, S_I , but also as a summary, the actions of all the other players, S_{-I} .

OK, so these are the definitions of a strategic form game-- actions, utilities, and so on. So a strategy is how to play the game, how to choose particular actions. It's as if you were writing code to tell yourself what to do as a function of what everyone else is doing. So this can be easy or hard.

Nowadays, you could program computers to play chess. But that's thought to be a pretty hard problem. But this is the way we think about agents choosing what to do in a strategic form game.

So instead of jumping to a notion of a Nash equilibrium right away, let's think about probability measures over pure strategies. But let me just say from the beginning, why are we doing this? If you thought there's a finite discrete number of actions for two players, then player one could do one thing, and conditioned on that, player two would choose his thing. And again, conditioned on what player two is doing, player one wants to change his mind and do something else.

So in principle, it could jump all around and never find a set of consistent actions. But if instead we think about each player choosing his or her action at random, then we're back to the simplex. Because the set of probability numbers are going to be non-negative and add up to 1.

So you can see where we're going, right? We're going to map this into Kakutani. So that's why, right away, we start talking about probability measures over pure strategies. So for example, if you had two actions, you could think about this probability measure as assigning a probability of doing one thing or the other thing, or the probability over any finite number of things.

And we call these probability measures mixed strategies, σ_I being the mixed strategy of player I in the set of all probability measures that put mass over actions. So σ_I for player I , σ_{-I} is a set of mixed strategies over everybody, over all the players. And it is being assumed here that there's no coordination, that they're randomizing independently. σ_{-I} is the mixed strategies of every player other than I .

And you can think about the utility function of player I as a function of the vector of σ that everyone is playing. Why? Because σ includes not just I , but all the other players. It's a random way of choosing actions. And a function of a given set of configuration of actions will generate a payoff, and hence generate the utility of each agent I .

So this is just expected utility. This is just a Von Neumann Morgenstern expected utility explicitly taking into account the random way that allocations are achieved. And I've been a bit vague about the strategy space, S . Sometimes, I said it was finite, like two actions.

This bottom line specification here is like, there could be a continuum of possible actions. So we integrated rather than summing things up. OK, so let's finally define a Nash equilibrium.

Mixed strategy Nash equilibrium is a strategy profile σ^* for each and every player such that player i here takes as given the mixed strategy, the strategy σ_{-i}^* of all the other players, and hopefully will find the σ_i^* for him or herself is, in fact, not worse in utility terms than doing any other mixed strategy. So Nash is always, take as given what everyone else is doing. And then, you're fully free to optimize.

And you can see where this is going. It's a bit more complicated. We're going to look for some kind of fixed point in these mixed strategies so that when we put σ^* in, we get σ^* back out. One curious thing here, not that it matters, but it's interesting to think about the out of equilibrium problem of a given agent.

The agent takes as given the σ_{-i}^* of all the other players less himself, and then would in principle, be looking for a mixed strategy that might do better than σ_i^* . But the mixed strategy would be putting mass on potential non-trivial number of actions. But if one action in the mixed strategy provides higher utility than the other, then player i would always choose that action.

Why randomize over other things that give you less utility? So in fact, to try to guard against alternatives that do better, you only need to evaluate over pure strategies, not mixed strategies. But anyway, the definition is all on Nash that σ_i^* here is going to weakly dominate any other pure strategy σ_i given that the other players are doing σ_{-i}^* .

So that's the definition and a proposition. Now finally, we get the dividend. Every finite game has a mixed strategy Nash equilibrium.

It's huge, hugely important. And again, you could write down a game. But how do we know that it even has Nash equilibrium? Now we know that if it's a game with a finite number of actions for every player, it does have a Nash equilibrium.

It may not be a pure strategy equilibrium. It may involve randomization. But it exists.

And to come to the fixed point part of it, we're looking for this profile σ^* such that for every player i , σ_i^* weakly dominates any other strategy σ_i , including pure strategies, taking as given what the other players are doing. In other words, if you start with a σ^* , player i would be reacting to what the other players are doing and choose the best response given what the other players are doing.

So indeed, this B_i is the best response set. There may be more than one utility maximizing response. If so, pick one. It should be utility maximizing per player i given the configuration of what everyone else is doing.

So now, we can talk about the best response for player i , taking as given what other players are doing. It's a bit tricky notation here, because σ_{-i}^* is what we want to find. And this is the best response given the strategies of the others σ_{-i}^* .

And to be consistent with the notation, this B_i has a σ_{-i}^* on it. But that does not mean that these are the responses of the other players j . It is just a way of denoting the best response for player i conditioned on what all the other guys are doing.

So $B_i(\sigma_{-i}^*)$ and B_i both being consistent with each other. It's a bit weird in terms of notation. But so this B_i of σ_{-i}^* means the best response for player i . We do that for all players i and define it here.

This is B . Should have been a minus I . I'll have to email Ashley. Anyway, so we have B of σ to be the configuration of best responses for all of the players.

And so we're looking for a Nash equilibrium to be a fixed point of this mapping. In particular, we can find a σ such that given that σ for each player, now knows what the other guys are doing in the mixed strategy sense, comes up with his or her own best mixed strategy as part of a best response. Each player is doing that and ends up doing what they all took as given to begin with-- namely, σ^* . So we just need this B thing to have good properties.

And in particular, this B thing turns out to satisfy the conditions of Kakutani's theorem. It's operating over compact, convex, nonempty sets. Itself being nonempty, has a convex value correspondence and has a closed graph. So then, we invoke Kakutani and we're done.

So the second or third illustration, now, of Kakutani's theorem, in this case, to establish the existence of a mixed strategy Nash equilibrium in a finite game-- are there questions? All right. So let's apply this artillery to US markets and talk about trade fails.

So a trade fail is where a buyer of a treasury agrees to hand over liquidity-- money, excess reserves-- to the seller, who has a treasury. And likewise, the other guy, the seller, is happy to supply the treasury in return for the liquidity. So buyer and seller meet. They meet in over the counter market.

It's not a centralized Walrasian market. And then, they're supposed to show up and settle the trade. Well, technically, the seller of the treasury may not yet own the treasury, but is planning to get a hold of it in one way or the other so as to be able to pass the treasury from him or herself to the buyer. And likewise, the buyer may or may not have liquidity at the time they strike the deal, but is planning, hopefully, to get a hold of the liquidity in order to be able to buy to pay for what they've already agreed to do.

But it doesn't always work out that way. And there are settlement fails in the system. This happened, for example, with a computer outage. How can you settle trades when it's all electronic and there's nothing on the screen?

9/11 was another problem where there are a lot of disruptions. Now, they all have backup systems in New Jersey and at Sanford, but not at the time. So as I said, the seller may not have the securities. One possibility is that he was planning to get them from someone else, but that someone else who was selling the securities to him for him to sell it to the third party himself is failing.

So we get what we call this daisy chain, a cascade of failures. So you really getting into the weeds in terms of-- it's not Walrasian. It's not simple Nash. It's a much more complicated institutional arrangement.

But we're going to try to use the Nash equilibrium. Now, here's a picture of the trade fails. Although this was between 1990 and 2014, this paper from the New York Fed is intended to be illustrative of this surge of fails that happened in the great financial crisis of 2007, '07, '08.

So OK. It's a spike. So what? We'll go over here, say, higher than 9,000.

9,000 what? 9,000 billion. We're talking trillions here. This is unbelievable. OK?

It doesn't get bigger than this. In fact, the treasury market is the world's highest value market in the whole world. And now, we have a substantial fraction of trades that have been agreed to as failing.

I wasn't able to update this slide, but I did see a paper just two days ago, which is basically, this year-- starting, say, January, February, March-- in March, the virus hit. And trade fails just skyrocketed. So the Fed had convinced itself that having enough reserves on its books would provide ample liquidity to the system sufficient so that we would never see trade sales again.

They got it wrong. I'm not sure how much background you have. But the Fed is buying securities. They're buying mortgage backed securities that other people wouldn't hold. And for buying them, they credit buyers, broker dealers, and banks with accounts at the Federal Reserve.

So those things are called reserves. They're well over excess reserves. That's where banks keep excess value. And they pay for stuff with those reserves.

Banks don't borrow from one another very often anymore because of all the excess reserves, and so on. But in March, it all fell apart. And the Fed had to inject even more liquidity.

So we're not quite ready for monetary economics yet, but we are ready to talk about how to model this market for securities. So let's think about it as a game, a strategic game played among traders. So there's going to be various goods, like liquidity, treasuries. We're back to our usual notation here.

There are j goods, j equal 1 through not L , but k . And there's going to be a trading post-- a market, so to speak-- for each of the goods. And they're going to buy by paying unit of account to get the good, or sell the good, and return get unit of account into their accounts.

So agent I , trader I , plays a strategy, which has four objects. Basically, can buy or sell, but even more than that, you could buy and specify a price p_i for good j . You're willing to buy if the price is less than or equal to that price that you specified.

These are consumers, right? So lower prices are a good thing. If you're willing to buy at p_{ij} for good j , then you're willing to buy if the price were less. But you limit your exposure. Regardless of whether it's p_i or less, you're only going to agree to buy at most q_{ij} units of the goods.

So these are called limit prices. Then, that's the way the New York Stock Exchange works, et cetera. Not all of the exchanges work like this. The tildes here denote limit prices for sellers.

So a seller of good j will say, well, I need a minimum price, \tilde{p}_{ij} . I'm willing to sell at that price. I'm willing to sell if the price is higher, but I'm not going to sell an infinite amount. I'm going to limit my sales to q_{ij} .

So in this game, all the traders are submitting these limit orders. And they effectively could be on one side or the other, or be willing to do both depending on the prices that they need. So a strategy, S , consists of the set of possible actions, all these configurations of vectors for each of the capital N players.

And then, we need to denote the outcome. The outcome function is this function g -- we'll put it over here-- that if everyone is S , being S_1, S_2, S_n for each of the players, this vector S of actions limit orders submitted by all of the end players maps into a particular outcome for player I , trader I , which is x , the excess demand or supply, and β . And β is a new object here.

It's going to be the net credit that agent I receives when strategy profile S played by all the agents. Beta is going to be the total unit of account an agent spends. And if positive, if positive, beta means you spend less than what you get as a seller. You're a seller, you're getting revenues in unit of account from selling.

How much depends on what everyone else put in. If you're buying, then you owe unit of account. How much depends on what everyone else is doing.

Beta could reflect the fact that they sold more in value in unit of account than they demanded. That would be a good thing sort of. They're just underspending. The opposite is a disaster, where they've committed to spend unit of account they don't have.

That's when beta goes negative. That's a fail. That's going to be a trade fail. They can't come up with the liquidity.

Or well, we want to prevent that from happening. By the way, here's the simple-looking supply and demand schedules. These are consistent with the limit orders.

For example, there's a maximum price that a household is willing to pay. We go over all the traders. The max over the max is the highest demand.

And as price starts dropping, you'll pick up a few other traders who are willing to buy, but they limited their order in quantities. And the price drops some more. So you get these step functions.

And likewise, for the suppliers, you get a step function going up. As the price gets higher and higher, you attract more sellers. But they have all submitted limit orders. So we're looking for this equilibrium in this single dimensional market.

There could be ties. Supplies equal demand. We have to have a way of rationing or allocating the whatever is short or long here. In this case, there's more demand than supply. So we have to allocate that across the demanders.

Anyway, these are meant to be intuitive pictures. So that's just describing that picture. Now, we want to make sure there are no fails. So we introduced this lambda.

So lambda is going to be a penalty function. And how does it work? If beta is positive, then under their strategies and the strategies of everyone else, there's more revenue than costs, in which case, beta is positive. Positive's higher than 0.

This whole term is 0. There's no penalty. But conversely, if beta is negative because their committed net expenditure is higher than their net revenue, then the minimum of this function is a negative number. And it's amplified by lambda.

So lambda is the penalty cost. It incurs if you overspend. And of course, ex-ante, they're going to choose their strategies given the strategies that other people are adopting in such ways to maximize this total payoff, not just the utility of getting goods if they succeed in doing that, but also the disutility if they violate the budget.

So there's a curious lemma here along the way that the betas end up being 0. In other words, in an active Nash equilibrium, everyone's on their budget. Oh, yeah, like the Walrasian thing, right? That's what we're looking for.

We're going to try to support the Walrasian allocation momentarily. But one thing-- not the only thing which is true of the Walrasian allocation-- is that excess demands are 0 for each household. The valuation of what is sold is equal to the valuation of what is demanded.

They're in their budget. Data I is 0. It turns out, the way this game is being defined, βI must be 0. Why?

Well, it's deceptively simple. Let's sum over the β s of all the households. That has to be 0. Why?

Because it's all in unit of account if whatever is spent by me ends up in the lap of whoever sold me the goods. So the unit of account liquidity is just passing around the system. In other words, liquidity always flows so that it's balanced.

So the sum of the liquidity deficits must be 0, because someone always ends up holding the liquidity. And then, you could try out what would happen if some βI were not 0. Well, if this sum is 0 and sum βI were positive, then if there's a βI which is not 0-- say, less than 0-- then there's some other βL which is greater than 0.

Now, that trader didn't spend all of his unit of account valuation. So that couldn't be maximal, given the strategies all the other players are adopting. Player L would purchase more. So the β s have to be 0.

And the final slide here is, how do we choose λ ? Or alternatively, given a competitive equilibrium, how do we find the corresponding Nash equilibrium in this game that achieves it? But we're given the prices and allocations of the competitive equilibrium.

And in a competitive equilibrium, each trader's maximizing utility subject to budget defined by those given prices, vector p . That will define the marginal utility of income, μ , which coincidentally is exactly what we called it on Tuesday-- μI , the marginal utility of income of trader I , OK?

So what we want is to choose these penalties high enough so that the penalty is higher than the marginal utility of income. So μI is the marginal utility of income. We're going to scale it by something called α in such a way that we get this inequality.

So you contemplate not honoring the trade. That's going to be associated with this penalty cost λ . And that's greater than the margin utility of income.

You should have avoided that situation by not going over your budget. And so given the α , we can redefine the limit prices and the limit quantities in a way as to achieve a Walrasian equilibrium. So we've actually defined Walrasian equilibrium, Nash equilibrium, defined a game here in a market that potentially has a problem.

Evidently, the penalties are not high enough. If we follow this model, there's a way to implement penalties to avoid trade fails so that the Nash equilibria of that institutional setup would correspond with the Walrasian equilibrium. OK, so thank you for coming today.