

[SQUEAKING]

[RUSTLING]

[CLICKING]

ROBERT M. TOWNSEND: OK. Greetings, everyone. Thank you for coming today. I'll say a few things. So in terms of the calendar, and we're finishing off this section right before Thanksgiving break, when I get-- excuse me, on identification today. We're on that momentarily.

Then, we have two formal lectures after the break, on Tuesday, Thursday. So that's the calendar. Then, in terms of the readings today, it's important to point out, the section on identification and falsification we're about to do, there are three starred readings, and they all come from Varian, and they're all really, really good.

In fact, I had to abbreviate in spots, in order to cover all the topics. So we are drawing on this material in part, and you will find it useful. Again, it's not the whole book. It's just Varian, in these sections, 8.1 through 8.3, 8.5, and 8.7.

Then, from last time, let's have a little bit of a discussion. Let's see. Daniel, can you tell me, as best you can, what the definition of is the positive representative consumer and the normative representative consumer, explaining what they are Intuitively

AUDIENCE: So like the idea is you're trying to treat the entire economy as a single person. So a positive representative consumer is, at any given level of wealth and input, the way to maximize utility is to spend your wealth the way that customer would and then distribute that to the people like in some way. We don't really know which way from the consumer, but like it'll just be like the best distribution. And then the normative is very similar, except for instead of the utility maximizing, it's the Pareto optimal way to distribute the wealth.

ROBERT M. TOWNSEND: Yeah. I guess I would sharpen up a little bit the positive one. The idea is to get the excess demand curve for the economy, as if it were coming from this representative consumer maximizing utility subject to resource constraints. And then that representative consumer has all the resource constraints in the entire economy. In particular, the budget just says, don't spend more than the economy-wide wealth. And you said that, but maybe better for me to repeat it a little bit.

Yeah, and then the normative one is a way to rank order allocations, according to the Pareto principal. But ironically, rather than keeping track of all the individuals, and whether they're made better off or worse off, we can collapse it to as if we were maximizing utility for the whole economy through the lens of this representative consumer. And as you said, that determines the overall allocation. It doesn't pin down the potential reallocation of wealth that is necessary, if you go from something Pareto inferior to something Pareto superior or Pareto optimal, you may have to compensate the losers and by taxing the winners. OK, great.

So we made a lot of use of this Gorman form. Caleb, can you tell me in words, what is this Gorman polar form?

AUDIENCE: Yeah. So like it breaks down the equilibrium into like a linear form. I guess, I don't know exactly how to explain it.

ROBERT M. Do you remember any of the notation? What is it breaking down exactly?

TOWNSEND:

AUDIENCE: Like the equilibrium in the economy.

ROBERT M. Does the word indirect utility come to mind? OK. It's the indirect utility function, which is the maximized utility subject to the given vector of prices and wealth. And the statement about Gorman form is this form in which that that indirect utility function must take in order to do this positive and normative thing. Can someone volunteer to tell me what, in words if you can, what the green form looks like?

AUDIENCE: Sure. I can take this one.

ROBERT M. Great.

TOWNSEND:

AUDIENCE: So the Gorman form is basically an indirect utility function that is linear in W , the overall wealth of the economy. So it's some function A , which depends on prices, plus sum function B , which depends on prices, plus W .

ROBERT M. Yeah. That's exactly right. Well, it's at the level of the individuals.

TOWNSEND:

AUDIENCE: Yes.

ROBERT M. So if we had an individual I , where do we put the I in that equation?

TOWNSEND:

AUDIENCE: So it is on both the wealth that an individual person faces, and I think as well on the functions A and B . Each individual can have their own specific form for that.

ROBERT M. Two out of three. So it's on the wealth, because at the individual level, it's the indirect utility is a function of the individual's wealth, say W sub I . The intercept term, the A of P is the A_I of P . So that's where the heterogeneity is loaded.

TOWNSEND: But it's not on the B , and maybe you can remember why. What did we say about those linear expansion paths?

AUDIENCE: Oh, those need to be parallel linear expansion paths, and so they need to be identical for each consumer.

ROBERT M. Exactly. That's perfect. So there's no I for the B part. OK, and then finally-- oh, I almost answer this already, but we can discuss it some more. How realistic are the necessary and sufficient conditions for Gorman aggregation, based on what you've learned in previous lectures? In other words, do you believe it or not? Guanpeng?

AUDIENCE: Yeah. I believe one of the necessary conditions was that the aggregate demand doesn't depend on the distribution of the wealth. It has to depend only on the total wealth.

ROBERT M. That's true, and then what does that imply about the underlying households in the economy?

TOWNSEND:

AUDIENCE: I believe it implies some uniformity.

ROBERT M. Yeah. It's uniform in the way that wealth impacts marginal changes.

TOWNSEND:

AUDIENCE: Right.

ROBERT M. So yeah. Those linear expansion paths are a shorthand way of saying different people can consume different amounts, because they have different wealth. But if you start changing wealth distribution, by definition, holding the aggregate fixed, if you're going to change the distribution, you're increasing wealth for some households and decreasing it for others.

And those marginal changes-- and not just local changes, global changes-- these could be big redistributions, offset each other in the aggregate. The amount that consumption of a given good is going down for the people who are taxed is exactly equal to the amount that consumption of that good is going up for the people who get the subsidies. And that's why the redistribution doesn't matter for the aggregate. That's why we can do this representative consumer thing to get the aggregates, because only aggregate wealth matters for aggregate demand, as you said.

However, it matters a whole lot for the individuals. I mean if you're paying taxes, lump sum to the government for it to be redistributed, you get hurt. So it's not a welfare statement.

And I'll just repeat, that's why, depending on what we're doing with free trade or something else, in principle, we may be able to make everyone better off, as we are told by situations where there is a representative in indifference curve and we're using implicitly Gorman's form for the indirect utilities. But that doesn't mean everyone is actually better off without the intervention, because the factor prices are moving, and some people own more labor than capital and vice versa. So we really do need to be actively redistributing wealth in order to reap the gains for everyone to benefit from the gains to trade. OK.

So this lecture from 18 was about a tool, namely Gorman aggregation, which you saw in the form of the earlier lectures without me elaborating on it. So this was an opportunity to do that, but it does take underlying assumptions, and in part, the microdata seem to contradict. Yeah. I didn't push you guys very hard, but do you believe it or not? Then, it's do we see these differentiable wealth effects in practice? We have normal goods, necessary goods, and so on and even given goods, as I said.

To know if it's a bad approximation, we'd have to take a stand on the underlying non-Gorman indirect utility functions and simulate the economy, and see if it looks almost linear, even though it's not. In which case, the redistribution might be small. Sorry. I said that wrong.

The literal form of Gorman aggregation as an analytic technique might get us close to what we want as a benchmark, and since it's so tractable, it's very tempting to use it. But then again, for other utility functions, we could be really, really far off and be misleading policymakers.

Hence the microdata. Where my pitch is, if you're going to do macro policy, you have to have micro data, period, but that's partly a belief that I've acquired over time. It's a matter of experience. Although, it's not always true. The approximations can be very helpful.

OK. So the lecture today is kind of the reverse of that. Instead of putting structure on indirect utility functions, we're going to see if we can get by with almost no structure at all on preferences and still be able to make predictions which are testable or rejectable. So this part, this lecture is all about economic science.

I featured the science of experiments in Lecture 1 as a motivation for the class. We went through [Koopmans] and Lucas and Josh and Matzkin and so on. So this is very much the science of doing experiments, and I'll tell you exactly what we're going to do with micro data, depending on the amount of data available, and also macro data, trying to put as little structure as possible and still get rejectable restrictions. Said that.

So we're going to go through consumer optimization again, which is, as always, a bit of a review. We're going to approach that as if we had an infinite amount of data, and that will take us back to the Slutsky matrix, which you may have forgotten by now, but I'll try to remind you.

That's with an infinite amount of data, we get restrictions that we can reject consumer rationality, in principle. And then we'll do the same thing with finite data and still derive algorithms running on finite data sets that would allow us to not reject or alternatively reject the rationality of the household generating the data. Then, we'll come back to Lucas a bit in computational considerations, which I think has been clear throughout the lectures that we visit from time to time.

And I have dug up some really interesting computer science material to share with you about how hard or easy it is to solve certain problems and whether it matters, be provocative. And then we'll go to general equilibrium theory and kind of do the same thing all over again. And that's going to again come back to this you need micro to do macro statement that I just mentioned.

OK. To get in the content of it, this slide appears to come out of nowhere. It was actually almost part of the lecture we had on income and substitution effects, and I'll remind you of that in a minute. In particular, it has to do with the Hicksian demand and the expenditure function. Well, I'll remind you now.

We did utility maximization subject to the budget, and we got the Walrasian or Marshallian demand. We also decomposed that into income and substitution effects, and the way we got the substitution effects was to do compensation. Prices, say if they decrease, are associated with increases in income. We take the income effects away, and just look at how demands change, as we change prices.

In other words, we minimize the expenditure necessary to achieve a certain level of utility, and we got the Hicksian demand. So with the Hicksian demand and the expenditure function, we didn't actually go through these properties in Lecture 3, but they are intuitive. If the utility function is continuous represents locally non-satiated preferences.

We might as well define the consumption set to be the non-negative L dimensional orthant. Then, we have the following four properties the expenditure function, this minimized expenditure at prices p to achieve a certain utility level u , is homogeneous of degree what? 1 in those prices p .

I think I slipped up the other day, when I said 0, and then I switched back to 1. It depends on what object. When we do utility maximization to budgets, if we increase income and prices, things are homogeneous of degree 0, because nothing changes.

But this is a property of the expenditure function. So logically, if prices go up, and you're hitting the same level of utility, it's going to cost you more. So the expenditure is just double, if you're doubling prices. It's homogeneous of degree 1.

The expenditure function for any given price vector p is increasing in utility. That makes sense, when you think about having nice concave indifference curves, or even linear indifference curves. If you do increase u by moving further and further north out, hitting higher and higher indifference curves, then to minimize the expenditure of achieving those points, the expenditure has to go up. Likewise, as prices go up, expenditures go up, which I kind of already said.

And this is a key property. In addition, the expenditure function is concave in the prices. So we're given that the utility function represents local non-satiated preferences. We're not given much else.

So this might be a bit surprising. This allows for it to be linear, however, and in a minute, we'll make it strictly concave. Anyway, you remember how these proofs go. Pick two prices, pick a linear combination, prove that an intermediate combination is also minimizing expenditure, et cetera. And then the utility expenditure function is continuous in both objects, but what I want to focus on is number three, that the expenditure function is concave in prices for fixed utility.

OK. Relatedly, the Hicksian demand is, if you're going to achieve a target utility u , then you will achieve it exactly, essentially. That is to say, the x 's, which are part of the solution to the minimizing expenditure problem when substituted into the utility function, generate u exactly. If preferences are, number two, convex, then these minimized Hicksian demands are convex valued, entertaining multiple solutions that minimize expenditure, for given p and u . But if preferences are strictly convex, then this h function is single-valued and continuous.

So here, it's where the strict convexity of the preferences comes in, which we're going to get to again in 4. 3 says this Hicksian demand is single value if it's single valued, which it would be with strictly convex preferences. Then, it's differentiable, and moreover, when you take the derivative of it, with respect to say a price p of L , for the L th good, you get the minimized-- you get the quantities back. So the intuition for this is expenditure is just $p_1 x_1 + p_2 x_2 + \dots + p_L x_L$.

You take the derivative with respect to p , you're going to get the x . That's a bit sloppy, and we did the envelope theorem earlier and reviewed it last time. But basically, the derivative at the optimizing minimized expenditures, at the parametric price, p sub L , is you get the solution back, the Hicksian demand.

So these two pieces, 3 and 3, on each of the two slides generate this statement. That if you take the Hicksian demands and differentiate with respect to prices, then-- because level h is on the right-hand side in 8, but the derivative is already on the left hand side. So if you're differentiating the right all over again with respect to p , we're double differentiating, so to speak, the left-hand side. So we get these second-order partial derivatives, [? own ?] end cross partial derivatives, of the expenditure function, when we differentiate the Hicksian demands.

This is really compact notation. Obviously, there's a whole vector of prices, p_1 through p_L , et cetera. We're writing it down almost as if it were single dimension, but it's meant to capture the entire L by L matrix.

That is to say, derivative of the first Hicksian demand with respect to the first price, the first Hicksian demand good with respect to the second, price and so on, going across the row. Fill out all the rows. So it has all the own H_{II} , H_{IPI} derivatives, as well as H_{IJ} , I th good with respect to the J th price. It's all loaded into this compact notation.

Now, I'm saving the best for last, which is this guy here is the matrix of second-order derivatives of the expenditure function, and three already told us the expenditure function is concave. So that means, if you think about first-order conditions and second-order conditions, that second-order conditions are going to be satisfied for a concave function. And in fact, if you look it up-- stuff does slip memory a bit-- it means that this matrix of own and cross-partial derivatives is symmetric and negative semidefinite.

You may not remember that. You think about one good, then you just have first the ii derivative. You think about two, then you're going to have a 2 by 2 matrix. Then, you have the determinant of that matrix.

If you have three goods, then it's a bit tricky. To factor the matrix, got to figure out the determinant of each of the subcomponents and so on. But an equivalent way to check and see whether all the conditions are satisfied for concavity is just to pre-multiply and post-multiply the matrix by a prime and a, as arbitrary vectors, and make sure the result is negative.

So that's a bit of a review. Don't worry about it if you never learned it. It's the same thing as getting concavity of the expenditure function. OK. So what? OK.

Well, it's a bit problematic. I said the motive today was to get restrictions on observables, but we don't see this guy. We don't see the Hicksian demand. We just see the Marshallian demand, so we're in trouble.

Except this-- this is a statement about the Slutsky matrix. What we want to find, which is the derivative of that thing we just looked at, turns out, it is associated with observables. This is the L th good changing with respect to the K th price is equal to-- the way the whole demand is changing of the L th good, as we change the price p_k . But we adjust for the income effect, which is how the L th good is changing, as we change income or wealth, post multiplied by the quantity of the good price of which is changing, namely $p_k x_k$ for p_k . OK?

So then, you're probably still struggling. I went back to check. You may remember this. Remember this lecture where we did the income and substitution effects, these guys?

And then we decomposed the total demand change, as a consequence of the price change into the substitution effect and the income effect. And we had to clarify that the income effect was negative, because in that diagram we had p on the wrong axis, given the strange way economists invented demand curves. So then, what we're trying to do today is to get this substitution effect matrix, and we're getting it from observables, how demand is changing as price changes plus how demand is changing as income is changing.

So this simple example equation is elaborated in this Slutsky matrix. The sign changes, only because we changed what's on the left-hand side, from the total demand to the substitution effect. Substitute effect equals total demand plus income changes. Substitution effect equals change in total demand plus income changes. So in fact, the Slutsky thing is really familiar to you from that entire lecture that we did on income and substitution effects.

And now back to the point, to reiterate. This thing under the assumptions on the previous slide has to be, if we enumerated it out for all l and k , a symmetric negative semidefinite matrix. So we can fill in the l , the row column elements from all these observables, and just check and see whether it's true.

OK. So I should have said, this is with infinite data, because we're talking about having derivatives, which is infinitesimally small data. But it's still good to know that, in principle, there's a way to do it. If you can't identify something with infinite data, you're not going to identify something with finite data.

So let's review where we are. The goal here is to see whether we have testable restrictions. In particular, ask the question whether what we see, in terms of the demand function, came from a utility maximizing rational consumer or not. So I'm just going to stick you with a big database, in this case an infinite dimensional database, and have you check and see whether it could have come from a rational maximizing-- maximizing some utility function, consumer. OK?

So what are the three properties homogeneity of degree 0. Again, multiplying income and prices by a scalar, you just get-- nothing changes. That's a statement about not being naive about inflation, for example. Although, I guess arguably, it could be violated, but it seems pretty reasonable for a thoughtful consumer. In two different situations with prices and income just being multiples of one another, they shouldn't change what they do.

The second property is called Walras Law, which is a fancy name for spending all your money. Namely, if you have wealth w , and you add up all the expenses observed to take place at prices p , that's also equal to your wealth. So you don't underspend.

You can't overspend, because that would violate the budget. You don't have the money. But you could, in principle, underspend. However, a rational consumer with local non-satiated preferences would never do that. They go right out to the budget.

Technically, Walras Law is used to mean you only need to solve for the equilibrium-- if there are capital L goods, you only need to solve for the equilibrium in L minus 1 of them. Because the L th is pinned down by this budget equation. So it's like no extra restrictions.

And finally, the Slutsky matrix, if the demand function is differentiable and single-valued, then that matrix, s sub p and w , is symmetric and negative semidefinite. OK. Who cares? Here's the content.

If the other two things are true-- which one would argue is likely to be true in any generated data set-- if you can find anything that's not symmetric or not negative semidefinite, then preferences could not have come from a locally-- from locally non-satiated preference, as a solution to the max problem. So this is the necessity part. If we have utility max, then all three properties have to hold. Otherwise, can't be true that it came from utility max problem.

Now, one thing-- we'll get to convexity in a minute, but I'll just earmark that, that there is no statement here about convexity, just locally non-satiated preferences, and it's not quite fair. The Slutsky thing we derived did rely on-- at least the way I presented it, it relied on convex preferences. The question here is whether one can test for non-convex preferences with the data, which is a bit of a different question, and I'll come back to that.

So those three properties not only have to be true, if we maximize utility, but if they're true in the data, then we can always find a utility function that generated the observations, and that's called integrability. Let me give you the formal statement of it. So this is like the sufficient part. If preferences of a household are strictly increasing, strictly quasiconcave, then the Walrasian demand function is homogeneous of degree 0, satisfies Walras Law, and its Slutsky matrix is symmetric and negative semidefinite. That's the first part.

The second part is, if the demand function is homogeneous of degree 0, satisfies Walras's Law, and the Slutsky matrix is symmetric and negative semidefinite, then there exists a utility function that is increasing quasiconcave that would generate the observed data. So let me backtrack. This first part A is what we did initially. The second part B is the integrability theorem. It says we can work backwards from the data set to an underlying utility function, which again has to be increasing quasiconcave and would generate the data as if a solution to the max problem.

So this is the bit about not testability of convex preferences, and the idea is we've deliberately drawn underlying utility function, which is not concave or these upper contour sets are not strictly convex. You see the wiggle here. But if we start rotating a budget line or along the indifference curve and picking out the tangencies, we would generate the Hicksian demand, et cetera.

Then, we get to this particular budget line, where we're no longer pivoting. We jump. For a moment, we would be indifferent between this point and this point, and as we keep rotating the line and making it have less and less slope, we will start finding pivoting along this curve indifference curve. That shouldn't be bending back up again, but anyway.

So the point is we never end up in this non-concave portion. So we don't generate observables in there. What we would infer is they have strictly concave, or weakly linear, preferences over certain ranges of prices, which is consistent with an underlying weakly quasiconcave utility function. Now this doesn't mean that we can't test for convexity of preferences in other ways, but this whole lecture is about restrictions on data that come from market behavior, maximizing utility subject to budgets. And if that's all we have in the data, we cannot test for convexity.

Again, the statement, there exists a utility function that is increasing in quasiconcave. That's probably confusing, because it looks like we just said, the utility function is quasiconcave. No, this one is not. That's the true underlying function, but we can't test for that.

All we see in the observables are the one that as if generated the dashed line, and that utility function is quasiconcave. So that's the content of this statement. It's as if there were a utility function which is quasiconcave. Questions?

All right. So operationally, we might as well assume convex preferences, because we're not going to be able to reject it anyway. That's an odd-sounding corollary to this whole first part of the lecture, which is the goal of placing restrictions on data, whether data place restrictions on theory or not. And finally, let me say again, why are we trying so hard to see what we could potentially reject?

The answer is, if say it's not Slutsky, then it could not have come from a rational consumer maximizing utility, and that's good news. If for any data set, we can always find a consumer maximizing utility that could have generated the data, then there's no content to the theory. It might be fun to do the theory, but it's a vacuous empirical exercise. It has no content, if it's always true.

Fortunately, it's not always true, and these are the conditions we can check. It's an odd thing, because then you could end up with, oh my god, where did this data come from? It couldn't be coming from a rational consumer, but anyway. We actually like it if we can reject things, strange science.

So now, let's go to finite data. We had observed, a minute ago, we had the entire Walrasian demand function for arbitrarily small changes in prices and wealth and so on, but typically, we don't see that. We don't have an infinite data set. We may have a household, and if we're lucky we can observe that household over long periods of time and see what decisions are being made, if those things are recorded on Alipay, for example.

And if it were true that everything that was bought was acquired on Alipay, then we would see-- and if we also saw the wealth, because they're spending everything on Alipay, then we would actually, potentially, have a very long data set, but not infinite. So we want to see whether we can, nevertheless, test the hypothesis of rationality, and the answer is going to be positive. There is a way to take the data and run an algorithm and see whether the results are consistent but potentially inconsistent. And in the latter case, we can reject.

So this is the definition of the Weak Axiom of Revealed Preference. Let's start with just having two observations with different prices and different choices. (p_1, x_1) (p_2, x_2) is a pair of observations on prices and consumption bundles, implicitly for a given household.

t doesn't mean, necessarily, t over time, although you could imagine we observed at t equal 1 the first choice and a t equal 2 the second choice. But that's confusing, because it raises the dynamic issues and everything else. So t is really only meant here to index the data point, and there's two data points and hence two values for t . And let's suppose it's not trivial, so these choices are different.

And we say the individual choices, these choices, satisfy the weak axiom of revealed preference, if whenever at the second price is p_2 with x_2 chosen and x_1 in the interior of the budget, we must then have, at prices p_1 , spending (p_1, x_1) , that x_2 is not attainable.

So let me say it again. It's an if-then thing. If this is true in the data that, if they had chosen x_1 , they would not have been spending all their money-- if that configuration is true in the data, then we better find that, at prices p_1 , the valuation of expenditures on x_2 is strictly greater than x_1 , and again, the idea is revealed preference. So at prices p_2 , they are revealed to prefer x_2 over x_1 , because x_1 is available, and they didn't choose it.

So it better not be the case that, in some other way of looking at the data, this inequality is violated. Because if at prices p_1 , x_2 were available, because it's interior to the budget at prices p_1 , then to be consistent, they should have chosen it, because it's already shown to be weakly preferred to x_1 . So that's the idea. If this is true, then this is true. So it's like a statement of an algorithm running over the data, and we're going to generalize that.

So what happens if this isn't true? This is true, but this is not. So let's look at a picture. At prices p_2 , they choose x_2 , so on the budget. And x_1 is interior, so that's this condition. x_1 is interior to the budget at prices p_2 .

Then, we look at prices p_2 , p_1 , are they? Here. So here's x_1 . It's on the budget at price p_1 . The question is, where is x_2 ? In this case, x_2 is interior to the p_1 budget line. That violates. We didn't want to find that. That violates what ought to be true, and so it violates the weak axiom.

Not only that, try to draw indifference curves from a quasiconcave or linear or strictly convex, you can't do it. You'd get a tangency here, you'd get a tangency here, or looking at the bottom, the indifference curves would cross. There's no way you can draw this picture without crossing the indifference curves and have a tangent at the two choices.

Well, we learned right in Lecture 2 that indifference curves can cross for rational consumers. A is preferred to B. B is preferred to C. A is preferred to C. You may remember that diagram. So that's meant to be an intuitive proof of the picture of the proposition.

So in conclusion, for the weak axiom if the agent has preferences which are locally non-satiated and rational, then given observations on prices, expenditures, and wealth, you must have that the pairs of data points satisfy this Weak Axiom of Revealed Preference, WARP. OK?

So now, we want to generalize. We have here still two data points, just slightly different notation. The data points are indexed by t and s , and we write that the t data point of x is revealed directly preferred to the data point indexed by s , if x is available at prices p_t .

So again, at prices p_t , they chose x_t . That's the way the data are organized. They could have chosen x_s , because it's in the budget, but they didn't do it. So we say directly reveal preferred.

This might imply strict preference, but it doesn't really mean that. It just means, as with the axioms of reveal preference, what we see is what they want. That's all we're saying. They reveal to prefer it, then they continue to prefer it. In this case, they prefer it over other things they could have chosen and didn't. It's pretty simple-minded.

Then, we come to I almost want to say indirectly revealed preferred. Actually, it's in parentheses here to contrast with the directly revealed preferred. So now we're going to have a bigger data set. We're going to have a whole sequence, and we say that x_t R without the d is just revealed preferred not directly preferred. Or if you like, I for indirectly preferred, x_t is revealed preferred to x_s , if we have a sequence of data points, such that pairwise we have direct revelation.

So we want x_t to, at the beginning, to be compared to x_s at the end. But we do that, we never see, say in the data, that they're directly comparable. Instead, we see x_t in a situation where, within the budget, x_t is revealed directly preferred to x_{r1} . In another piece of the data, we see x_{r1} is revealed directly to x_{r2} , and so on down the chain, with x_{rk} being revealed directly preferred to x_s .

So this is like the transitivity axiom, which is part of the axioms of rational consumer. Right? We have a chain in which we have directly revealed preference. Hence, the one at the beginning ought to be indirectly revealed preferred to the one at the end. Next x_t to x_s .

So then we come to the Generalized Axiom of Revealed Preference, called GARP so I don't know there was a famous movie "The World According to Garp." I don't know if you-- was a Robin Williams movie. They didn't mean this.

So generalized revealed preference is building on this notion of indirectly revealed preference. If we have the data points of the dimensionality-- the data set is capital T , and we have prices and demands for each T , then the general axiom of revealed preference means that, if x_t is indirectly revealed to x_s , then if we're able to make the comparison, we have the strict inequality. Namely, x_t is preferred to x_s , it better be the case that, at prices p_s , x_t is outside the budget, relative to the expenditures on x_s .

And the intuition, as I wrote in red at the bottom of the slide, is simply imagine this inequality is reversed. Then, we would have, at prices p_s , x_t is within the budget, and x_s ought to be at least weakly if not strictly preferred, because x_s was chosen, and x_t was available. So x_s would be directly revealed to be preferred to x_t , according to this definition.

But it goes the other way here, where x_t is revealed, granted indirectly, to be preferred to x_s . So that if this were a strict inequality going the other way, we would intuitively violate. These two pieces of information would not be consistent with each other.

Now, again, you should think about this as an algorithm. Some classes do this. I found one on the web today. They give students a very large data set of prices and quantities and just ask them is this consistent with a rational household?

And how do you approach that? Well, you start looking for these chains. You have to construct them algorithmically and get, as a result when it's true, a statement about x and t indirectly reveal preferred.

And then in the same data set, go looking for an instance where there was a direct comparison, and it better be that the inequality is weakly going this direction. So it's an algorithm, I'm trying to say. Because again, if it doesn't go in this direction, then it contradicts the whole notion of revealed preference and transitivity.

So again, good news, ironically. We are potentially able to reject. There are comparisons in the data which, if this inequality goes the other way, would tell us there's no way that this data came from a rational maximizing consumer, and that's good for us.

Now, it doesn't mean that, in any given data set, you will necessarily find a violation, but you end up with a weak statement that the data set look as if there was nothing inconsistent in this data, with respect to the statement that they came from a rational maximizing consumer. So we should check and make sure that that's true, and that's the spirit of it. Questions? OK.

So let me get to the computational part. These guys at Caltech, consumer theory assumes consumers possess infinite computational abilities. Like they can solve max problems and potentially hard problems.

Proponents of bounded rationality want to require that any reasonable model of consumer behavior incorporate computational constraints. In other words, they may not be able to solve a hard problem. They may choose something, but it wouldn't be the full solution, because they're not able to compute it.

And Echenique and his coauthors are saying, that's false. Very much in the spirit of revealed preference, any consumption data set that is compatible with a rational consumer-- excuse me-- is also compatible with a rational and computationally bounded consumer. That is to say, the data set in hand can always be rationalized by a utility function that's pretty easy to solve, meaning solved in polynomial time.

So this is very much like not being able to test the convexity of preferences. You start with something non-convex, but it generates a data set that could also be generated by something convex. We start with something that could have come from the household solving a hard problem, but it could also have come from a household solving a simpler problem.

I'm not giving you any intuition on this right now. I'm telling you what these computer scientists have figured out. And again their metric is, how hard is it to solve the problem? Looking at derivatives and so on, how long does it take, as we, say, increase the number of goods? If it's exponential, it's bad. If it's polynomial, it's good.

All right. Another version of this, instead of having a single household maximizing, or otherwise, we have multiple agents interacting in a market economy. We want to say, if possible, that the observed market outcomes are consistent not only with individual maximization but with Walrasian equilibrium. So what is the empirical content?

Is it really hard to solve the general equilibrium computationally? In which case, it's not likely we're going to be able to accept in the data that it came from the underlying economy, or it's a version of the individual maximization problem. That there's another economy, which is actually easier to solve, that could have generated the same data. This is totally remarkable stuff.

Because it goes against the grain of the way one would like to think about it, as rational as consumers can possibly be, it's unlikely they can solve in their minds problems that prove intractable for computer scientists equipped with the latest technology. It's a cool statement. These are really famous theorists, by the way. Gilboa, Schmeidler, and Postlewaite solve in their minds-- get it-- what you would need for a supercomputer.

If an equilibrium is not efficiently computed-- this is the market version-- much of the credibility of using the predicate that we have rational agents is lost, and a shorthand version of that-- if your laptop can't find it, neither can the market. But these guys are saying, that's just wrong. It's just completely wrong.

Because they misunderstand the role of models in economics, which is as if the model were a reality, then there are restrictions on data. It is not a statement that we have all of the reality captured in the model. So the reality behaves as if the theory were true.

A consumption data set is either not rationalizable at all-- so that's the rejection part. They're not going to overturn that. There's still content. It's either not rationalizable at all, or it's rationalizable by utility function that's fairly easy to maximize. When I show you these slides, I want to show you more, but I guess we have limited time. So I probably should have put this paper on the reading list as optional.

OK. So let's go to general equilibrium in a paper of Brown and Matzkin, which all things considered, is relatively recent. Suppose we're given data set on a number of households, capital I in the economy, and we see their individual endowments-- private ownership economy, ω_i , across households capital I .

Then, there's something that I'll elaborate in words now called Anything Goes. General equilibrium theory does not put restrictions on data. Anything goes. And that was a killer. That killed the whole subject for a long time, because of what I'm saying.

Now, the object in question was our famous excess demand function that we have used before, the aggregate demand function that we've been focusing on with aggregation and so on. What shape does that aggregate demand function have, as we vary prices? And the answer was whatever wiggle and turn it has, including upsloping portions and so on. You can always find an underlying economy that would have generated that excess demand.

So this part of it is the seeming answer to the question. What happens if we don't put any restrictions on utility functions, other than things like rationality? And the answer is anything could happen. That aggregate demand function could look any way you want it to. For any aggregate demand function, we can populate an economy with a finite number of households, give them certain preferences and ownership of resources in such a way that we can get that aggregate demand to shift and turn and do whatever we want.

So again, the shorthand is anything goes, and because we could never reject the model, people thought that general equilibrium had no content, and they stopped doing it. It was discredited. This was not just one person. It was Hugo Sonnenschein, Debreu, and Mantel. OK.

But let's just change this. First, what is this seemingly reasonable but actually crazy thing that we actually see the excess demand? We never see that. If we see market economy, we only see market clearing prices. We don't see what happens out of equilibrium. So we only see where it crosses. We see the fixed points.

Well, let's give us a little more data then. Given a set of observations on prices and individual endowments, does there exist-- OK. So T now indicates the economy. It's as if we're going to go in the cross section, across several or many economies that vary in the individual ownership of endowments.

So it's ω^I for the I th economy observation. For any T that varies over I , we have a finite number of households, but T is varying in the sample. That's the number of economies in the sample. Those economies are all alike for preferences, that's assumed. But they differ in the ownership structure, the ω^I s.

So then does there exist a sequence of these economies, each with its associated aggregate excess demand, but such that those prices that we see stuck into the aggregated excess demand equals 0. This last statement is just a statement that the price has come from a market. In other words, the data set we're going to see would be the prices that because it's a market or associated with zero excess demand, equilibrium prices. So we're going to see equilibrium prices, and we're also going to see individual endowments, assuming preferences are all the same.

So I'll give you the punch line. That's Brown and Matzkin. The punch line is going to be what we've been saying throughout the lecture today. Namely, if we're given data like this, we can actually reject that the data could have come as a market equilibrium. That is to say, that the market the data couldn't potentially never have come from individual household maximization and market clearing. So it has content.

Here's the differentiable version of it. As abstract as general equilibrium theory is, this Sonnenschein, Debreu, and Mantel is overturned. They're applying it to an exchange economy, where we see prices and endowments. They derive necessary and sufficient conditions for the equilibrium prices as a function of initial endowments, and in their language, they show that the economy can be generically identified.

What they mean by that is there is an underlying economy that would have generated the data, but also that, in principle, one could have not been able to find such an economy. And then to summarize where we are, if you only had aggregate data, you're sunk. You're never going to get conditions that allow you to reject. But when you have individual data, you can get testable restrictions and bring back the content of general equilibrium theory. So again, this summarizes my deliberately provocative statement that you need micro data to do macro.

And then they go on to talk about something you're familiar with, because we did it in the class. We're sharing in Indian villages. It's not uncommon to observe endowments, like crops and prices. We did a little bit of that, less so with the risk sharing, in which case their content works directly.

If you really believed you were going from one village to another, each were self-contained. They had the same preferences and the premise would be we're seeing market outcomes, then their theory applies directly. You can even go to large economies, not villages, and think about types of individuals, finite number of types, but a large number of individuals of a given type and again you can use their artillery.

Now, interestingly, lest you think economists have this all figured out, an interesting question is how to extend this to production. They refer to changes in fact or endowments, which have observable impact on factor prices. So that's an allusion to the trade lectures we did.

And they haven't figured it out yet. They haven't figured out whether there are rejectable restrictions. Remember, when we did the model, the 2 by 2 case, we had constant returns to scale and a few other assumptions. So they're saying like drop all assumptions. Could we put restrictions on data and say we'd never see violations of some kind, and they haven't figured it out. So we're on the research frontier at this point.

So that's what I have for today, again, very much part and parcel of the whole spirit of the class. Which is to think about economists as running experiments, and you saw three of them today.