UNITED STATES
INTERNATIONAL TRADE COMMISSION

In the Matter of:	)	
	)	
SLICING THE PIE: QUANTIFYING THE AGGREGATE AND DISTRIBUTIONAL EFFECTS OF TRADE	)	Inv. No. 332-599

Pages: 1 through 59

Place: Washington, D.C.

Date: October 16, 2023

HERITAGE REPORTING CORPORATION
Official Reporters
1220 L Street, N.W., Suite 206
Washington, D.C. 20005
(202) 628-4888
contracts@hrccourtreporters.com

THE UNITED STATES INTERNATIONAL TRADE COMMISSION

In the Matter of:	)	
	)	
SLICING THE PIE: QUANTIFYING THE AGGREGATE AND DISTRIBUTIONAL EFFECTS OF TRADE	)	Inv. No. 332-599

Monday,
October 16, 2023

Teleconference
U.S. International
Trade Commission
500 E Street, S.W.
Washington, D.C.

The seminar commenced, pursuant to notice, at 9:31 a.m., before the United States International Trade Commission.

PARTICIPANTS:

USITC:

TAMAR KHACHATURIAN, Moderator
BILL POWERS, Chief Economist, Director, Office of Economics
SAAD AHMAD, Economist

SHARON BELLAMY, Supervisory Hearings and Information Officer
TYRELL BURCH, Management Analyst

Seminar #1

PAGE

P R O C E E D I N G S

(9:31 a.m.)

MS. KHACHATURIAN: We'll go ahead and get started as it's one minute past 9:30 a.m. here in Washington, D.C. Hello, and welcome to the U.S. ITC Distributional Effects Seminar series, Seminar 1. I'm going to turn it over to Bill Powers, Chief Economist and Director of the U.S. ITC Office of Economics. Bill will provide a short background on the past and ongoing USITC distributional effects of trade and trade policy on U.S. workers report, as well as the format for today's session, and he will introduce our speaker for today.

Bill?

MR. POWERS: All right. If you've just joined us, good morning, good afternoon, or good evening. Welcome to the first session of the Commission's seminar week for our report on the Distributional Effects of Trade and Trade Policy on U.S. Workers. Some of you may have joined us in our similar events last year. If so, great seeing you back. For those of you joining us for the first time, the Commission put out our first report last year on the effects of trade on U.S. workers and underserved communities.

The USTR has since asked for a series of five more reports on the topic over the next 15 years. The first of those reports will be due in January 2026. We're planning a

series of annual outreach events for those reports, the first one being this week. And so, Simon, you are our first of the first, so thank you for joining us. And the next annual event will be scheduled for July 2024. Please see our website. If you have the chat feature enabled by your organization, Saad has put the website in the chat for more details on the report and the seminars that we have going on for the rest of the week.

Some ground rules before I introduce Simon, our speaker today, I wanted to go over the ground rules. Simon will present for 60 minutes. Saad Ahmad, who you see on screen there, is our discusser from the U.S. ITC, and he will follow for 10 minutes. That will leave about 20 minutes at the end for questions and answers. Please hold your questions until the end, but please also put any questions you have into the chat while Simon and Saad are talking and we'll get to them in order during the Q&A session. Also during the Q&A, particularly for those of you who don't have the chat function enabled, please use the Hand-Raise function to flag your interest in asking a question.

All right. On to the introduction. Our speaker today is Simon Galle, an associate professor at the BI Norwegian Business School. He'll be discussing his paper, Slicing the Pie, Quantifying the Aggregate and Distributional Effects of Trade. The paper was published this year in the

Review of Economic Studies and is co-authored with Andres Rodriguez-Clare and Moises Yi. The authors incorporate heterogenous labor into a multisector gravity model and they examine the welfare effects of the China Shock. This paper is one of a series of papers that Simon is writing that explore the impacts of trade and technology on local labor markets and demographic groups, so right in the focus of our report. I am very pleased now to turn over the virtual podium to Simon.

Simon, please take it away.

MR. GALLE: Thank you, Bill, and thank you, Tamar, for the very kind and generous introduction, so hope to not disappoint. Let me start by sharing my slides if that works. There we are.

Okay. So, as Bill was saying, I'm going to present on Slicing the Pie, Quantifying the Aggregate and Distributional Effects of Trade, and the paper is co-authored with Andres and Moises. As you were saying, the paper has been published. We're very grateful for that. It's been a little bit of a long road at least for me as a relatively junior researcher. So we started to work on this in 2014 and back in the day, what we were doing was quite novel, but now there's more papers, luckily, who are working in a similar way.

When we were starting working on this, our main

motivation was the following. We have gravity models which we all know are very useful for quantifying the aggregate welfare effects of trade. But these models are fairly silent or are typically silent on the distributional effects of trade, and from a large body of empirical research, we know that these distributional effects can be very large. So our basic goal in this paper is to bridge these two literatures, so combine gravity models with a role for distributional effects of trade on workers, and in particular, we're going to focus on the aggregate and distributional effects of the China Shock on U.S., where we're going to refer mainly to the Autor, Dorn & Hanson work on the China Shock.

So two components, two basic components of the paper are gravity on one hand and the distributional effects of trade on the other hand. So, first, a few words on gravity models. So gravity models, for those working in trade, everyone knows that these are trackable and structural frameworks to predict trade flows so we can have many countries, many sectors. So this is basically a generalization of the standard Ricardian framework to many countries and many sectors.

Why are these models useful? Well, one reason is that with fairly few elements we can say something useful about, meaningful about, the gains from trade. So, if we just have trade data and the trade elasticity, then we can

perform counterfactual analysis, say something about the gains from trade, or, in our case, say something about the aggregate effects of the China Shock. So we're going to use this gravity model. We know that it's a stylized model, but it has a reasonable empirical performance, and there's a lot of ongoing work in making these models better perform to the data to the extent that they're not doing so already.

That was about the gravity part, so the trade part, the aggregate, again, effects of trade. Then on to the distributional effects on workers. So, there, we're going to focus in particular on the China syndrome and on the work by Autor, Dorn & Hanson in particular. For those of you following this seminar series, I'm sure many of you are fully on top of this literature and so, as you know, Autor, Dorn & Hanson, they focus on local labor markets and they find that in response to the increase in import competition that the U.S. was facing from China, there was a large relative decline in earnings as well as unemployment for the commuting zones most exposed to this increase in import competition coming from China.

There's a lot of other findings as well. So there's an increase in federal transfers. There's a decrease in merit trades. There's an increase in suicide and drug overdose. Electoral polarization might be attributed to this increase in import competition. So it does seem like

something major has happened related to the China Shock and its impact on U.S. society and U.S. economy.

There's a caveat, though, to these findings, which is that this empirical methodology is related to a difference-in-difference methodology, which means that you can only identify relative effects. You can identify the effects on one commuting zone compared to another compared to the typical commuting zone, but there's a limitation to what you can say about the absolute effects, and that's important since there's an increase in import competition, and the reason that there's an increase in import competition is precisely because there are these lower prices that are associated with the increased imports coming from China. That's the reason why we're buying more from China. So we know that workers' wages are going down, for instance, but prices are also going down. So then the question is, what is the net effect or, better, what is the absolute effect of both of these forces combined and are groups, namely, commuting zones, better or worse off? That's the type of question that we are going to try to answer in our paper.

And to that effect, we're going to combine the standard gravity model with a Roy-Frechet model of the labor market. What is a Roy-Frechet model? That's basically a model of labor reallocation across sectors, and it's going to -- when you have labor allocation or reallocation across

sectors, that's essentially a discrete choice for the workers, that the workers are facing. They need to choose in which sector to work. And the Roy-Frechet model in my eyes is or in our eyes is the standard discrete choice model to work with this, and I'll say a lot more about that as we go forward.

What do we gain from combining the gravity model with the Roy-Frechet discrete choice model of labor reallocation that we have? Well, if we start from the benchmark, which is a standard multisector gravity model, a standard multisector model of trade across many countries, many sectors, there, we have that workers are perfectly mobile. There's no friction whatsoever. They can -- there are no wage differences across sectors. So workers are perfectly mobile. That's going to be one extreme.

On the other extreme, we're going to have that we would have that we're in the specific factors world where workers are stuck in their sector. They cannot move across sectors. So multisector gravity model is one benchmark model. We also all know the specific factors model. What our model is going to do is it's going to nest these two extremes, which is then neat since we can sort of check what happens if we're closer to one extreme versus the other extreme.

How do we nest these two extremes? Well, first,

let me say what a Roy model is. Well, as I said, it's a discrete choice model where workers are going to self-select into sectors. How are they going to self-select into sectors? They're going to do this based on their comparative advantage. So, in a way, it's the efficient benchmark of worker allocation and reallocation that's going to happen in this model, workers selecting based on their comparative advantage based on which sectors it's best for them to work in. That's the Roy part, which is a standard framework in economics.

We then tag on to that the Frechet parameter, Kappa, which is going to determine the scope for reallocation, and this is where the two extremes come in. So, on one hand, when Kappa goes to infinity, we're going to have perfectly mobile workers and we're back in the standard multisector gravity world. When Kappa goes to one, we essentially have specific factors and so we're back in this other extreme. And then we're going to eventually estimate Kappa, and it turns out that we're going to be fairly close to the specific factors extreme, which is going to mean that there's going to be large distributional effects of trade.

So, methodologically, what this achieves is that we're going to generalize the specific factors intuition to a setting where there is some labor reallocation. But the specific factors intuition is going to be key in that what's

really going to be driving the distributional effects of trade is the initial sectoral specialization of a group of a commuting zone, and that's going to -- this setting is then going to allow us to examine the between-group distributional effects of trade. And so, when I say between-group, I mean between commuting zones or across commuting zones. So, on principle, I can also introduce different groups, but we are going to focus on commuting zones as the main group.

That's the baseline setup, gravity plus Roy-Frechet, and then, given that baseline setup, there's a bunch of extensions that one can introduce. The first one is intermediate goods input/output look. That's very important for the quantification, as we show in our paper, but it's not really changing the intuition of what's going on, so today I won't say much about that.

A similar comment on trade costs within countries, so, presumably, or we all know that there are trade costs for moving goods from San Francisco to D.C. or the other way around. In our benchmark model or our baseline model I should say, we're not going to focus on those trade costs. So we're going to abstract that. But, again, important for quantification, we address this in the paper, and it's essentially not changing much.

Very interesting from a theoretical point of view, an interesting extension is introducing imperfect

substitutability between different types of labor. So what we mean there is that in our baseline model, within each sector, there's only one type of labor. And so we're now distinguishing between, for instance, college workers versus non-college workers. But that's something that we can do in the next section, and, hopefully, I'll have time to talk about that towards the end of the talk.

What's really critical in my eye, in my eyes, especially from the point of view of the literature, the empirical literature on the China Shock, is allowing for changes in labor force participation and unemployment. So not a part of the baseline model, but we do spend a lot of time on it in the paper in the main extension to address these things, and, hopefully, I'll have time to go into that.

We also have an extension where there's labor mobility across commuting zones, so in the baseline, workers are stuck in their commuting zone, which is obviously a limitation. We have an extension where we allow for this mobility, but we don't address it since we lack the data. So it would be a huge data effort to bring this, to pursue this research avenue, and I hope someone else is doing it since I don't have the data.

Then some other limitations are strong parametric assumptions, no implications for within group inequality, so all the results are going to be for between groups, so across

commuting zones inequality, and also no transitional dynamics. So we're going to have an initial period the economy is going to be shocked by the China Shock, and then later on we're in a new equilibrium and we check the changes going from the initial to the new equilibrium, and we don't check the transitional adjustment path. That would also be important. So that's a limitation.

In terms of the literature, as everyone knows, there's a huge literature on the aggregate gains from trade. There's a huge literature on the distributional effects of trade, trade and sector reallocation. So I think we do a good job of discussing this in the paper, but if there are questions there, perhaps we can address them in the Q&A.

Then on to the model. So, in terms of the model, we have many countries and many sectors. In particular, we have N countries and we have S sectors. For the countries, there can be origin countries and destination countries, and those are going to be indexed by I and J. The sectors are then indexed by S or K in our case. In each country, there can be many groups based on predetermined variables, so one can have leeway in terms of how to define these groups. For our purposes, these groups are going to be commuting zones. That's going to be the main focus of today.

As I said, we have basically two blocks in the model, the gravity trade side and then the labor side, the

Roy-Frechet side. So, first, a few words about the trade side. We have many sectors and each of these sectors is modeled as in the benchmark gravity model, so as an Eaton-Kortum, which means that there are constant elasticity of substitution preferences over a continuum of goods, and then we have national constant terms technologies where these technologies for each good are drawn from an IID Frechet distribution where the level parameter of the Frechet is the TIS and theta IS is the shape parameter. And so this is going to be -- TIS is going to be important for determining the absolute advantage of a country in a certain sector, and then theta IS, as we'll see is, going to be important for the trade elasticity. It's going to determine the dispersion of comparative advantage across countries. There's either trade costs between any origin country and any destination country for every sector.

Finally, preferences across sectors are Cobb-Douglas, which is beta IS. So these are the expenditure shares of each country I for sector S, which is going to be important for the demand side. This was the gravity side, so bottom line, a standard gravity model for those who work with us.

On the labor side, things are in a way going to be very similar to what we have on the trade side. So it's also going to be a Ricardian model in a way of labor allocation

and reallocation. So we're going to have these groups. Group G, again, these are going to be commuting zones, and there's a mass LIG of workers of Group G in country I. These workers are going to have different skills in different sectors, and those skills are going to be determined by their Z-S, their efficiency units in each sector. And those skills or those efficiency levels are going to be again drawn from an IID Frechet distribution with a shape parameter Kappa and a level parameter AIGS. And in two slides, I'm going to show what -- or I guess on the next slide already, I'm going to show what that's going to be mean.

There's a wage in each sector in each country, and that's the wage per effective unit of labor in that sector and in that country. These wages are identical across space within a country. So identical across groups, so there's no trade cost and no national technologies. There's no trade cost and there are national technologies, which is going too imply that wages are the same across all groups.

But, importantly, these wages are going to be different across sectors, which is going to allow us to have an upward-sloping labor supply to each sector. So that's going to be the main role that our Roy-Frechet model here is playing. It's generating this upward-sloping labor supply to each sector. Because since these wages are different across sectors, workers are then going to choose the sector that is

maximizing their earnings, and their earnings are this product of their efficiency unit ZS and their wage WIS. Okay. That's the basic setup on the labor side. So workers have efficiency units in each sector, and then they sort across sectors based on their comparative advantage.

So there's a trade side and there's a labor side. On the trade side, we're going to generate labor demand, and I've already argued that the labor side is going to generate labor supply. The labor demand is determined by the trade side. That's because the demand for efficiency units in each sector is going to be given by this guide here.

So what is this? Well, in every origin country I and sector S, how much demand there is for their goods is determined by the expenditure level in the destination country, the expenditure share, like how much the destination country is spending on Sector S, and then on the trade shares along the IGS, which are telling us for Destination Country J how much are they sourcing from Origin Country I. What is the share of their spending in that sector that they are sourcing from Origin Country I?

So these guys are very important. The alumnus, the trade shares for Destination Country J, Origin Country I. What are those determined by? Those are determined by the level parameter of the Frechet, which is essentially determining the absolute advantage of a country in a certain

Sector I. So TIS is how the average -- let's say it's related to the average productivity of Origin Country I in Sector S. And as an aside, that's going to be important for the China Shock. That's going to be driven by these TIS, with China becoming more productive in the sectors where it's going to grow in. So China is going to grow its productivity in textiles, for instance, and then supply more textiles to all countries in the entire world.

So TIS is related to average productivity of the origin country. Where is the downward-sloping labor supply coming from? That's coming from the wage here, so the higher is the wage in your origin country's sector, the lower it's going to be demand for goods in that sector, and so the lower is also going to be the demand for workers in that origin country sector.

So what is the trade side doing? It's generating this downward-sloping labor demand function. As the wage is higher, there's going to be less demand for efficiency units of labor in that sector. So this is essentially a multisector Eaton-Kortum model. What's the only difference? Well, we said so already. The only difference is that now we have sector-specific wages instead of national-level wages, which is what Eaton and Kortum used to use.

From the labor demand side, we got this -- from the trade side, we got this downward-sloping labor demand

function. Now, from the labor supply side, what we're going to do is we're going to generate an upward-sloping labor supply. So now the number of efficiency units in a sector is going to be increasing in the wage, and why is that? Well, that's coming from our Roy-Frechet framework, where the employment shares of a certain group in a certain sector are going to be increasing in the wage here. So the higher is the wage, the more attractive it is to work in the sector and so the more workers, the more efficiency units sort into that sector.

Here, you can see the critical role for Kappa, which is going to determine the -- which is going to play the role of reallocation elasticity. So this was a dispersion parameter of the Frechet and it turns out that this turns into a reallocation elasticity. So the higher is Kappa, the less heterogeneity in workers there is and the more mobile workers are across sectors. So, when Kappa goes to infinity, we have perfect elastic -- perfectly elastic labor supply. When Kappa goes to one, in terms of efficiency units, we're going to have specific factors.

So these are the employment shares of a certain group in Sector S, a function of the wage, as well as a function of the AIGS, which were the level parameters of the Frechet which were determining the absolute advantage of a group in a certain sector. And so the higher is the AIGS,

the more workers are going to select into that sector and also the more important that sector is going to be for a group. And that's what we see here in this FI-IG. This FI-IG is going to be important. That's going to be a wage index which will determine the earnings potential of a group.

Why is it the wage index? Well, we see that the wages appear here, and roughly speaking, we can think of it as a way this FI-IG is a weighted average of the wages where the weights are determined by the IIG. So Detroit is specialized in car manufacturing, let's say. That's going to mean that relatively speaking, Detroit is going to have a high AIG in car manufacturing and so the wage in car manufacturing is going to be important for Detroit's commuting zone.

So employment shares and the wage index FI-IG, and this wage index, as I've indicated already, is going to be important for the earnings of a group. Why is that? Well, AIGS is going to be total efficiency units supplied by Group G in Sector S and, in addition, the average income that workers from a certain Group G are going to get in certain S is going to be determined by this wage index FI-IG. That's going to be the key and not just variable given our Roy-Frechet framework for the average income of the workers.

So this average income that workers get in a certain group in a certain sector is constant across sectors.

It only depends on the wage index, and that means that group-level income is also going to be -- the only endogenous, relevant endogenous variable there is going to be FI-IG again. So, again, this wage index where the employment shares in the sectors are going to be determining the importance of the wage in that sector for a certain group.

So, here, we have the equilibrium of the labor market. From the trade side, we were getting our downward-sloping labor demand, and from the Roy-Frechet part, we were getting the upward-sloping labor supply to each sector. The key thing here is that this is mainly methodologically a technical point, is that all these endogenous variables are a function of the whole matrix of wages W, and so we're going to have a system of excess labor demand functions, so the difference between labor demand and labor supply. So we're going to have a system of excess labor demand functions for all origin countries in all sectors, and that system is going to determine the matrix of wages. So exactly identify a system of wages, a system of equations for the wages.

This is going to be the expression for the initial equilibrium, and that we can take it to the data and then feed it with a shock and look at the counterfactual equilibrium. So, for our purposes, we're going to be interested in the China Shock, so there's going to be a

foreign shock. Let's say China grows its productivity in textiles, and that's going to affect the entire trade equilibrium for the entire world. What are these hat variables? So those hat variables are the percentage changes. They're the value of -- they're the new value for the variable defined by the initial value of the variable. So we're going to work with this system in hats.

In particular, we're going to write a counterfactual equilibrium, so the equilibrium after the shock, in the following way. And so, here, it's not important to be fully on top of all the equations. I know that there's a few too many on this slide. The key thing is that we can take this to the data and that with the publicly available data, we can solve for the counterfactual equilibrium, say something meaningful about the aggregate effects of the shock and about the distributional effects of the shock.

How do we do that? Well, all the blue guys here, those are data, so those we can observe in the data. What we then need to say something about the counterfactual effect is the values for the shocks, which are the red guys. So we need to say something about how much productivity in China is changing or how much trade costs or tariffs are going down between China and the U.S.

So, given the data and given the shocks, we can say

we're going to be able to solve for the endogenous variables, which are the black variables. The only thing that we still need are the green guys, which are the Kappa and the theta. So this was the reallocation elasticity, the Kappa and theta is the trade elasticity, so telling us how responsive trade flows are to changes in trade cost or domestic production costs.

So, bottom line, given the blue guys, which are the data, and given the green guys, which are the elasticities, the reallocation elasticity and the trade elasticity, and given values for the shock, we can solve for the wages. We can solve for the wage changes and then thereby find what are the changes in the wage indices for the different groups, what are the changes in the trade flows, and what are the change in the employment shares in each of the groups in all of the sectors.

And so this system of equations here is sufficient to say something meaningful about both aggregate and the distributional effects in all countries in the world, so to speak. So we think this is useful for applied research.

So let's say we found the wage changes. Then we can say something about what are going to be the distributional effects of the shock. So distributional effects of the shock, we're going to talk about real income. We're going to talk about real income changes across groups,

and those are defined as the change in nominal income divided by the price index.

The price index reappears here, and the change in nominal income, we know that it's a function of the FI-IG, the wage index, so the change in nominal income is just going to be given by the change in the wage index. Given Cobb-Douglas preferences, we can substitute the price index for Cobb-Douglas preferences for the overall price index in the economy, so that's the first thing that we do.

And then we get fairly easily to our main welfare expression, which is going to be fairly intuitive, on the next slide. So, based on the link between domestic expenditure shares, the openness of the economy, an inverse measure of the openness of the economy, and the price changes, we can obtain the -- we can substitute the price changes here for our country level ACR 2, and I'll explain in a moment what that means.

And based on the link between the FI-IG, the wage index, and the employment shares, we can substitute those by an observable term, which is the reallocation term, the Roy term, that we're going to have here. So this is going to be our main welfare expression, and now let me tell you what it means.

So this is telling us what the real income changes are for a certain group, and those are going to be coming

from a country's consumer gains and the group's labor income gains. The consumer gains are also the ACR term, and, here, we have the quote. So the consumer gains, those measure gains from specialization at the country level, whereas the labor income gains, those are going to measure specialization at the group level, so at the commuting zone level.

And so why is that? So the first term here, the ACR term, is measuring gains from specialization at the country level. Why is that? Well, that's because the lower is your lambda hat. The more specialized is your economic data. So the lower is your lambda hat, which means that you're sourcing less domestically, you're buying less domestically, your economy is more open, so you're going to be more specialized in the goods in which you have a core comparative advantage. So the lower is lambda hat, the more specialized you're going to be, and then the ACR term is telling you what the consumer gains are associated with that specialization at the country level.

It's a similar intuition for the pie hat term here, for the Roy term. So these pie GS, these employment shares, so the specialization pattern of a group across sectors, those pie GS, they sum to one, which means that the lower is your pie GS, the lower is your pie hat IGS means that your initially larger sectors are growing more. So, if your average pie hat is lower, it means that that's because your

initially larger sectors are growing more, which, in comparative advantage lingo, means that the sectors in which you have a comparative advantage, you are specializing more according to your comparative advantage.

So this is written in this bullet point. So workers in Group G gain more if sectors of their comparative advantage can expand. So the average lambda hat is a matter of specialization at the country level. The average pie hat is a measure of specialization at the group level.

I haven't told you anything yet about the elasticities here. So theta is the standard trade elasticity, so the lower is theta, the more dispersion in comparative advantage you're going to have and so the higher the gains from specialization are going to be.

Similarly, for Kappa, the lower is Kappa, the more dispersion in comparative advantage you're going to have and so the higher the gains from specialization on the labor market are going to be. Okay. So this is our main result on the theory side.

We can then take this to the data as well, or we can provide additional intuition for what's going on, namely, that the distributional effects in our model are going to be driven by changes in input competition at the group level, and that's what we're going to see here, that the changes in relative income are going to be driven by changes in import

competition at the group level.

So why is that? So what does this IG term mean? Well, the beta is your expenditure share on a sector, so it's telling you how the share of expenditure that's going to Sector S in a certain country. RS is then the output share of a sector in a certain country.

So why is this related to import competition? Well, the higher is your expenditure share and the lower is your output level, the more you're competing with imports in that sector. And the other way around: the higher is your output and the lower is your expenditure share, the more you're exporting, so the less import competition there's going to be.

So this guy here is a measure of import competition at the sector level. Then we bring it to the group level with these employment shares weight, so then we have a shift-share measure, so to speak, of import competition at the group level, where the shares are the employment shares and the shift is the measure of import competition at the sector level.

Bottom line is that we can use this group-level import competition measure to say something about the distributional effects of trade, the changes in relative income across groups. So the stronger your increase in import competition, the lower your relative income is going

to be. And the strength of those distributional effects is governed by Kappa, so the lower is Kappa, the stronger are going to be the distributional effects because the lower is Kappa, the closer we are to a specific factors rule.

Finally, if the shock is back through autarky, then a group is going to gain proportionally to its current level of import competition. So, if you currently face a lot of import competition, then going to autarky is going to be relatively beneficial for you.

All right. Final slide on the theory. So we want to say something about aggregate and distributional effects of trade, and, ideally, we would want to do so in just one measure to sort of account for the inequality adjusted welfare effects of trade.

In order to do that, we're going to work with CS preference, or constant relative risk-aversion preferences, where rho here is going to be the measure of inequality aversion. So, when rho is equal to zero, we don't have any inequality aversion. When rho goes to infinity, we're in the Rawlsian case where we only care about the people at the bottom of the income distribution.

So, given those preferences for or preferences against inequality, we can say something about the inequality adjusted welfare effects. So now we have an inequality adjusted welfare effect, so the U hat, which is where we're

aggregating the group-level welfare effects, the group-level changes in real income, and the rho is going to determine the weight that we're going to put on the people at the bottom of the or the groups at the bottom of the income distribution.

So, again, when rho goes to zero, we're in the standard case where there's no inequality aversion. When rho goes to infinity, we're in the Rawlsian case, where we only care about people at the bottom of the income distribution.

So we set up a little bit of machinery, and now we want to take it to the data. We're going to mainly focus on the U.S. for the estimation, so U.S. labor market with 722 commuting zones. We use IPEM's data for the labor income and the employment shares. Estimation is going to closely follow Autor, Dorn & Hanson trade data from UN Comtrade, time period 1990 until 2007, so a standard ADH.

For the simulations, things are a little bit different, so there, we need to have the full trade network across all countries, so we use the World Input Output Database, and due to data limitations, we focus on the period 2000 until 2007. We're going to have 14 sectors, 13 manufacturing sectors and one non-manufacturing sector.

We need to estimate Kappa, or better, we need two elasticities, theta and Kappa. Estimation of theta has been done many times before. We are going to focus on the estimation of Kappa. How are we going to estimate Kappa?

Well, we're going to combine empirical and theoretical elements to estimate it.

On the empirical side, we know from Autor, Dorn & Hanson that higher exposure to the China Shock is going to imply that your manufacturing employment is going to go down, so there's going to be fewer people employed in manufacturing. On the theoretical side, we can fairly easily show in the model that when there's a decline in manufacturing employment, there's going to be a decline in relative income also, and that's going to depend on Kappa.

So, when your Kappa is high, it's going to be easy to reallocate or it's going to be fairly costless to reallocate workers from manufacturing to non-manufacturing. When Kappa is low, it's going to be more costly, and so there's going to be a higher relative decline in income associated with this decline in manufacturing employment.

And that's essentially what our main estimation specification here is saying. So we have 14 sectors, 13 manufacturing, one non-manufacturing, so, here, we focus on the expansion of non-manufacturing. We can do it different ways. In my paper, I do it in a different way, but this is a good way of doing it.

So we're checking what is the income change that's associated with a certain expansion of the non-manufacturing sector, and then, as this non-manufacturing sector is

expanding more, Kappa is going to determine how costly this expansion of the non-manufacturing sector is.

We need an instrument to have an exogenous change in this non-manufacturing sector, and that's going to be the standard Autor, Dorn & Hanson instrument where we use import penetration from China to other countries and then weight it by initial employment shares in the sector.

There's a long debate, extensive debate, about the exclusion restriction. We do all the robustness checks that were asked of us in the paper. For the purpose of today, we can just focus on the identification strategy in ADH, so use this instrument combined with the same set of control variables as in Autor, Dorn & Hanson.

Okay. This sets up our estimation. So IV estimation of Kappa, our main parameter that's determining the size of the distributional effects. So, on the left-hand side, we have change in commuting zones income per worker, and on the right-hand side, we have the expansion of the non-manufacturing sector.

In terms of the first stage, this is always very high, so the lowest F stat is -- the first stage is always very powerful. The lowest F stat is 24, so we're good there. And standard errors are clustered at the state level for those entries. So a strong first stage, which is one reason why we're focusing on the non-manufacturing sector.

What is the impact of an exogenous increase in the share of workers employed in non-manufacturing? Well, it's a negative impact. As the theory predicts, there's going to be a decline in relative income when the non-manufacturing sector expands exogenously.

The Kappa is then determining -- so, here, we have an inverse measure of Kappa. So one over Kappa is the coefficient that we're estimating here. So what is the impact on Kappa? Well, it ranges from 1.4 up until 2.8. So some uncertainty, but, I mean, we can rule out very large values of Kappa, is the main message, and we're going to show results both for our baseline estimate. say, 1.5, as well as for higher values of Kappa.

So we have our Kappa estimate, 1.5. We have our trade model and our labor market model, so now we can take it to the data to perform counterfactual analysis. The only thing that's left -- so we have our data, we have the blue guys. We have our elasticities, the green guys. The only thing that's left is the red guys, which are the shocks.

So we need values for the China Shock in our case, so the T hat here. How are we going to calibrate the China shock? Well, we're going to follow Caliendo, Dvorkin, Parro essentially, and we're going to use a predicted first -- in the first step, we're going to obtain predicted values of increases in China's exports to the U.S., and then we're

going to calibrate the increase in China's productivity to match the predicted values of U.S. imports from China.

How do we get the predicted values of U.S. imports from China? Well, this is based on a variation of ADH first-stage regression, where we regress the share of U.S. imports from China on how these trade shares between China and other countries are changing during that period. So, here, we get predicted values, and then we calibrate the changes in Chinese productivity such that our model matches the predicted values in U.S. imports from China.

So now we have everything. We have our model, we have the trade and labor data, we have estimates for Kappa, and we have our calibrated China Shock. What are the numbers that come out of the model? Well, the aggregate gains from the China Shock are fairly low. They're just .22 percentage points.

What's important is that the coefficient of variation, so the standard deviation divided by the mean, this is fairly sizable, so that's a hundred percent. So, even though the aggregate effects are modest, the distributional effects are sizable. And so, in particular, some groups are going to lose up to -1.18 percentage points of their income, whereas other groups are going to gain 1.65. So, relative to the aggregate gains, the distributional effects are sizable.

We can also see the role of Kappa here in that, when Kappa is growing, the CV, the coefficient of variation, the dispersion in the group-level effects is going down. So, when Kappa is one, we have the largest distributional effect. When Kappa goes to infinity, distributional effects disappear.

Where are the losses from the China Shock concentrated? Or better, what is the geographical distribution of the real income gains from the rise of China? So, perhaps not surprisingly, the Rust Belt is not doing great, as well as the broader eastern heartland, and so southern and central Appalachia is not doing great in response to the China Shock, and in particular, in these regions -- and this is not a surprise if you know the work by Autor, Dorn & Hanson -- the sort of strong geographic concentration of low-to-negative income effects associated with the China Shock. So that's important for policy, I would assume.

A little bit more about import competition. So I said that the basic channel governing the distributional effects is this local exposure to sector-level changes in import competition. And so that's what we see here. Here, we have our measure for group-level import competition, so local exposure to import competition, and that's determining the relative changes in income of a commuting zone. And we

can see again that when Kappa goes to infinity, there are no distributional effects, whereas, when Kappa goes to one, the distributional effects are very strong, so we have a very steep slope for this function. This was import competition.

Now on to our inequality adjusted welfare effects. So we have our aggregate gains. We have our distributional effects. Let's say we are interested in just one number: what are the overall welfare effects of the China Shock, adjusting for inequality, for the impact of inequality? Well, they're still going to be positive and, moreover, they're often going to be larger than the benchmark effect.

So, here, we have our parameter for inequality aversion, and we see that as rho is growing as we're putting more weight on the people at the groups at the bottom of the income distribution, the inequality adjusted welfare effects are rising slightly. The question then is why is this, and that's what I'm going to answer here.

So the distributional effects are determined by the change in import competition at the group level, and, here, we have the initial income level on the horizontal axis. The first thing to note is that for any income level, there's huge dispersion in the change in import competition. So just the initial income level is not a good predictor for how exposed you're going to be negatively or positively to the China Shock. So, basically, there's a lot of conditional

income levels. There's a lot of variation in import competition, so income is not a good predictor for how you're going to be affected by the China Shock.

But what we do see in our version of the model is that here, for the groups at the bottom of the income distribution, there's a slight decline in import competition. And so that slight decline in import competition at the bottom of the income distribution, that's going to imply that as rho becomes fairly sizable as the inequality aversion becomes fairly sizable, the inequality adjusted gains from trade are going to be larger.

So that's sort of concluding the chapter of the China Shock for now, except for the extensions that are still coming. We can also say something about return to autarky, which tells you something about the gains from trade. First, some numbers. So the gains from trade are 1.56 percent according to our model for the U.S., with a coefficient of variation of 63 percent. So still sizable distributional effects of trade. And we again see in terms of geographical distribution of the gains from trade that the broader eastern heartland is negatively affected by trade. So, overall, there are positive gains from trade, but there is a strong spatial concentration of low-to-negative effects of trade in here around Michigan, as well as in southern and central Appalachia. So, presumably, this can be interesting for

helping to understand why there might be some anti-trade sentiment in these communities.

I'm not saying that anti-trade sentiment is justified for the U.S. as a whole since there are positive gains from trade even accounting for the inequality adjusted effects, but there are important regions in the U.S., large regions in the U.S., where these gains are on the low side.

I still have a few minutes to talk about extensions, I think, so, first, I'll say a few words about introducing a Heckscher-Ohlin-type mechanism. So, as I said, in the baseline model, we just have one type of worker in each sector. What we do in this extension is we're going to differentiate between two types of workers, so college workers and non-college workers, and those are going to be imperfect substitutes in production with the analysis data.

So this is sort of combining specific factors, Roy-Frechet, with our Heckscher-Ohlin-type mechanism. So sector-specific wages and employment shares now differ by worker type, and there's an important new variable that comes in, kai, which is the cost share of labor of Type M in Sector S.

And given that kai, we can write our updated welfare expression. So this is the welfare expression that we had before, so ACR term, Roy term, and now let's call this maybe the Heckscher-Ohlin term or the skill premium term,

college premium term. That's because this final term here is capturing the average change in the college premium across sectors.

So why is that? Well, we see that the college premium, this term is inversely related to the cost share, so for an eta larger than one, which we typically assume to be the case. So, when your college premium is going up, let's say, that's going to reduce the cost share of -- that's going to lead to a lower cost share of college workers in a given sector, and so that's going to be reflected in this term here, which is capturing the reduction in the cost share for this worker type.

What else do we learn from this extension? So we can play around with the model. We can turn off the Roy part or we can turn off the Heckscher-Ohlin part, so to speak. When we turn off the Roy part, so the reallocation elasticity across sectors goes to infinity, we find that there's a reduction in the college premium, and that's because there's one particular sector -- I think it's the electrical and optical equipment sector -- that is employing a lot of college workers, high degree college workers, that actually has to contract due to the China Shock. So there's a reduction in the college premium, but really not major.

When we turn on the Roy part of the model or the specific factors part of the model, we see that the college

premium becomes lower and the coefficient of variation, so the distributional effects, is really growing a lot. And what that is saying is that this pattern of sectoral specialization is really important. It's much more important than the college premium for explaining the distributional effects of trade in response to the China Shock at least according to our model.

So sectoral specialization is much more important than standard skill premium type stories for the distributional effects of trade. I think I'm running out of time, so you can tell things are a little bit different for the gains from trade, but we can discuss that later maybe because I want to make sure to have time to discuss involuntary unemployment and voluntary non-employment.

So things become a little bit more complex here compared to the baseline model. So some technical remarks, I guess. So, here, we introduce a three-stage model. So, in the first stage, workers are going to choose between home production, so voluntary non-employment, versus employment in the formal sector. And there, the key elasticity is going to be mu. Mu is going to determine how easy it is for workers to reallocate from formal employment to voluntary non-employment.

In the second stage, conditional on seeking formal employment, you're going to choose your sector where you want

to be employed, and, again, that's determined by the old Kappa elasticity, so the reallocation elasticity, across formal sectors.

Then, finally, in the third stage, there's going to be random draws about who is going to be employed or unemployed in the chosen sector. So all this machinery is essentially developed by Kim, Vogel in their 2021 AR Insights paper, and the key or one thing that's going on there is that there's going to be frictional matching which leads to involuntary unemployment.

How is this working? Well, the match probability is increasing with the real wage, so when the real wage is increasing in the sector when the wage index is increasing in a group, it's going to be more profitable for employers to hire workers, and so that's going to lead to a higher match elasticity, so higher employment rate, lower unemployment rate in that group.

And what this unemployment adjustment is doing, it's going to amplify the effect of a shock because now we have an updated welfare expression, so this is the real income change in a group. Here, we have our old change in the wage index. So ACR plus Roy is what we had before, and, here, we see that this is amplified by the changes in the employment rate.

The changes in the employment rate are going to

amplify the changes in the wage index, whereas the share of employment in home production, so the share of voluntary non-employment, that's going to dampen the welfare effects of the trade shock since this is sort of an outside option that workers can go to when they are negatively affected in the formal sector.

What do we find? Well, we estimate these parameters. I won't go into how we estimate them; we can leave that to the paper. We find alpha is the employment elasticity or matching elasticity. So this ranges from 1.19, which is on the low side, to a fairly high 1.5. And so what would be great for follow-up research is pinning down just how large this employment elasticity is, so how sensitive frictional unemployment is to local changes in import competition.

Interestingly, there's an interesting difference between the Kappa and the mu. So the Kappa is slightly lower compared to what we found before -- I mean, still 1.5 or one at the lowest -- so fairly low Kappa, difficult to reallocate within formal sectors or across formal sectors, whereas the mu can be a lot higher. So we find values up to 3, which means that it's a lot easier or workers have a much higher propensity to move into home production, to move into non-employment when they are negatively affected by a shock instead of finding employment in a different formal sector.

So that's quite striking there.

In terms of welfare implications, so when there's no Sam and no H, no home production, so no frictional unemployment and no home production, these are our baseline values. When we add frictional unemployment, there's amplification of the shock. When we add home production, there's a neat outside option, so that's going to lower the welfare impact of the shock.

Then I think I need to conclude. So we wrote down a framework to study the aggregate and distributional effects of trade. The welfare effects are summarized in a parsimonious equation that nests the multisector ACR result. There's a key additional parameter, Kappa, that governs the strength of the distributional effects. We find that it's equal to or around 1.5.

The counterfactual analysis is revealing that the China Shock increases average welfare, but there are strong distributional effects. So some groups experience losses more than six times the average gains. So adjusted for plausible measures of inequality aversion, the gains in social welfare do remain positive still.

And if I can say one more word. So, in terms of improving the model fit, there's more work to be done there. So our model does -- I didn't say anything about this until now, but this is an avenue for future research. So our model

appears to understate the distributional effects compared to the data, and there's four complementary approaches to addressing this.

One can add intensive and extensive margin. One can look at the joint role of automation technical change driving down the labor share. Or these, so this is what I'm doing in one paper, and then other people are looking at downward nominal wage rigidity or agglomeration forces and spatial spillover. So all of these are important to improving the model fit of this model. So thank you very much for sticking with me.

MS. KHACHATURIAN: Thank you, Simon, so much for your presentation and for kicking off our distributional effects seminar series. I am going to hand it over to Saad Ahmad, an economist in the Research Division of the U.S. ITC. He will be the discussant for your presentation. He'll be sharing some slides. And after his presentation, we will move to a broader Q&A.

I see that, Bill, you've already posted a question in the chat. I'll try to keep track of the questions that are posted there, and as you said earlier, we'll also take questions from people that just want to raise their hand. So hand it over to you, Saad.

MR. AHMAD: Thanks, Tamar. Let me share my slideshow. And, again, I want to thank Simon again for

presenting his work. Let me see if this works. Can you guys see it?

MR. POWERS: It looks fine here, Saad.

MR. AHMAD: Okay. So, yeah, I'm going to be discussing this interesting paper by Simon and his co-authors looking at the distributional effects. Okay. I think I've got it. Can everyone see Slide 2?

MR. POWERS: Yes.

MS. KHACHATURIAN: We can see it.

MR. AHAMAD: Okay.

MR. POWERS: You can try Control L for full screen. But it's fine.

MR. AHMAD: Okay. Excellent. Yes. So, as Simon mentioned, this is basically the interesting or the novel contribution of this paper is like combining kind of two different strands of modeling to kind of get to what distributional effects.

The first part is kind of the standard kind of workhorse trade model that a lot of us are familiar with and is becoming kind of very useful in applied applications as well where we have a multisector Eaton-Kortum kind of correlated trade model, and to kind of apply this model, you need trade shares, trade costs, and trade elasticity.

So having these kind of observable values, as well as the trade elasticity, allows you to run a lot of

counterfactual simulations on what happens if you increase trade costs or if you increase productivity as well. So this is kind of the standard model for trade, you know, foreign trade these days, and what they had is a Roy-Frechet model of labor allocation across sectors, and the interesting thing about that model is that it allows workers to kind of select the sectors that they are going to earn the most wages on.

And so some features of this model is that each country has some groups of workers, and then, within each group, a certain share of workers are employed in a sector, and then there's a degree of heterogeneity in terms of labor productivity that is kind of driven by the Kappa parameter, and that's going to impact how heterogenous workers are. And so, again, kind of adding this kind of heterogenous workers element to this kind of very practical framework of the Frechet distribution allows them to kind of have a very practical overall model that can be used to look at distributional effects.

And so, from the trade side, they get labor demand in each sector. On the Roy-Frechet side, they get the labor supply in each sector, and then you can get equilibrium. And once you derive this, you get a multisector gravity model with heterogenous workers, right, and that is something that is a new contribution here because, previously, the Eaton-Kortum kind of framework does not have heterogenous

workers, does not really say much about distributional effects. It's looking more at the aggregate welfare effects. Now, with this framework, you can also make some, you know, statements about how groups are faring when there's a change in trade cost, so that is kind of the key contribution of this paper.

I'd just point out that I think Simon mentioned as well there are some limitations to keep in mind with this framework. It is a static model, so there's no dynamics. There's no, you know, transition taking place from one, you know, equilibrium to another. And so we can think of this as more of a long-run model.

So it's not capturing the short-term transition costs that might be associated as workers move from, you know, one equilibrium to another, and that has been shown in the empirical literature to be sometimes substantial because the transition is sometimes, you know, sluggish, and so it takes time for workers to go from one job to another or from one sector to another. So that is something that this model will not be able to capture.

Another kind of thing to keep in mind with this framework is that there's more restrictions on the movement across groups. So, you know, workers, the main groups are commuting zones, so workers are not able to move from one group to another, so that might be, you know, limiting their

movement from, you know, going from one zone to another, although there is a theoretical part that they have. Their kind of labor mobility is included.

Another margin of adjustment could be workers could switch from, you know -- they could acquire more skills, you know, increase their skill level, but here again, you know, the workers are either divided to, you know, non-college or college. So, again, that is something that, you know, we should keep in mind with this model.

And the last point I'll make about limitations -- this is kind of something newer these days -- a lot of, you know, these Eaton-Kortum models have exogenous trade balances, but there's recent work by Raphael Descardero and co-authors which have shown that, you know, the trade balance can also be a mechanism to which, you know, sector allocations can happen.

And so, you know, if we allow the trade balance to be endogenous, it allows you to kind of have more amplification of, you know, sectoral effects because, you know, the dynamic, it allows countries who are running a trade deficit to be able to continue to support, you know, certain sectors more than that might be possible if trade balances are exogenous.

Given this, as I said, this is very interesting work, and one of the things that they use is to, like,

calibrate the China Shock that has been studied a lot in the literature, both on the empirical side as well as, you know, kind of more on the theoretical side, and so they use this model to look at the effect of the China Shock on different groups of workers.

And one of the key findings they find is that 30 peer groups representing 13 percent of the U.S. population actually experience welfare losses from the China trade shock. And, you know, some groups actually see losses in real income up to four times as high as the average case.

And I think there's a very strong geographical component of these losses. They are mostly concentrated in the Rust Belt, Midwest area, as well as the, you know, inland eastern region of the U.S., you know, and they have a very nice map showing which areas are, which, you know, commuting zones, you know, are affected negatively.

One thing that I found that was really interesting is that, you know, the main groups that they had in their baseline is commuting zones, right? So they had, like, groups are divided by commuting zones, but they have extensions. But they also include, you know, groups can be distinguished by, you know, education and gender in a commuting zone and also education and age in a commuting zone.

But it doesn't really change their main findings,

and so the takeaway I got was, you know, to understand distributional effects in their model, that commuting zones are actually driving it, right? So, if you are a worker in, say, you know, the Midwest/Rust Belt area, you're going to be affected negatively even though you might be a high-educated worker than say if you're a worker in, you know, another county that's not being affected negatively by the trade shock. And so, you know, it seems like you can't separate out the location effects from the gender or education effects in this framework, which I thought was really interesting.

And as Simon mentioned, they have done a lot of extensions on the baseline model. One particular extension is that they also allow the model to incorporate changes in labor force participation and unemployment as to other margins of adjustment from the trade shock, but they don't see big effects, big differences from their main finding. But it's something that this model cannot incorporate, which I think is, you know, kind of pretty neat for us because those are some aspects that we in the policy circle focus a lot on, is unemployment and labor force participation.

And then I'll wrap up kind of thinking more, you know, just the general idea of, like, you know, taking some papers of a literature and trying to apply them, you know, for policy work, which is what we do at the Trade Commission. And one of the, you know, nice kind of derivations they do in

their paper is they have, like, a simple kind of Bartik approximation for how, you know, relative income changes across groups.

And this can be -- you know, it's an approximation, but it's shown, you know, with the China counterfactual that it's a pretty good fit for these novel values of Kappa. And the benefit is that now you can connect, you know, these changes in relative income to observable data, right?

You only need for this formula, you only need the share of workers, you know, by IGK to the share of workers in Group G that are employed in Sector K in Country I, and that can be obtained from the data, as well as the change in the share, you know, the output share of sector --

(Technical interference.)

MR. AHMAD: -- parameter Kappa, which, again, is, you know --

(Technical interference.)

MS. KHACHATURIAN: I believe Saad's camera froze, but I also believe that that was sort of wrapping up his final slide.

MR. POWERS: I don't see or hear him at the moment either.

MS. KHACHATURIAN: Yeah. So, with that, I don't know, Simon, if you have any initial reactions to sort of the questions that Saad posed on this slide that he has shared

right now or any other -- oh, you want to go, Saad?

MR. AHMAD: Can everyone hear me?

MS. KHACHATURIAN: Yes.

MR. AHMAD: Okay. I don't know how much you got, but I just wanted to say that this Bartik approximation was, you know, one of the useful features in this paper. And so, if you want to move towards a, you know, ex-ante analysis, like looking at the, you know, forward-looking effects of trade policy changes, some of the questions I had was, you know, more in line with, okay, how do we estimate, you know, more consistent estimates of, you know, output share changes, right? Like, ex-ante can we use, like, a standard kind of, you know, Eaton-Kortum multisector model, just run it with the change in trade cost and, you know, use that change in output as input in this formula.

Also, you know, the parameter Kappa, you know, matters a lot in this, you know, analysis, and so they have done a really good of estimating it using the China trade shock. But, you know, can we be certain that that value is, like, stable over time, so can we just use that for, you know, ex-ante analysis or should we, you know, think about re-estimating it using more recent trade shocks and U.S. trade shocks. You know, there's the COVID shock. There's also been a lot of trade policy actions, you know, China-related. So could we use that as a way to estimate

Kappa?

And also, you know, this was a very U.S.-centric, you know, analysis. How can we extend it to other countries, regions? And so I'll stop here. Sorry for the technical issues, but, hopefully, you know -- as I said, this was a very interesting paper, and, hopefully, we can continue our understanding of this paper in the discussions.

MR. GALLE: Should I respond or -- okay. Yeah, thank you, Saad, for the very kind and insightful discussion. Like, it was really great to see it, so thank you. So, to respond to your points, so starting with the final question that you had, so how to extend the analysis to other countries beyond the U.S., that's not too difficult. So, as we try to argue, you need -- you just need trade data and, for instance, some employment data across regions or employment data across groups, and then you can say something about the distributional effects of those groups.

So we have in the replication package, we have the data available for 2000 and 2007, but it's easy to incorporate it with more recent World Input-Output data and, for instance, IPEMS or another data source with group-level data, and then you can sort of get started there.

How to determine model consistent estimates of the changes in the sectoral shares, so the Rs? Well, ex-ante is always difficult, and I don't have experience with policy

work, so you might have a better sense there. But your suggestion seems correct. Like, if you have a measure of, for instance, tariff costs, how tariffs or tariff rates are going down, you can feed them, for instance, in a multisector Eaton-Kortum model, obtain them and then just plug them in there. Or you can also just plug them into our model right away.

One of the neat things of the Bartik approximation is that in the background, there's this entire machinery of calculating counterfactual trade equilibrium and labor market equilibrium, but then, for the distributional effects, it all boils down to this local exposure to changes in import competition, so I think that's a need for communicating the counterfactual or ex-ante analysis to policymakers potentially.

Then, finally, can the estimate of Kappa be identified from the China Shock, can that be used for more recent U.S. trade shocks and trade policies. So there, I have two answers. One is that in my other work, for instance, I looked at a different shock, so the automation shock, and so, there, the value for Kappa was fairly stable, so, again, around between 1 and 1.5, which is on the low side.

But that's still for the same period, like, in the 2000s, let's say, and so I think it would indeed be great,

and it's not that costly actually, to re-estimate it for more recent periods. So the approach is fairly straightforward. And it's also not limited to trade shocks, so any type of national-level sector shock can be used to identify this cost of reallocation.

Whether the old values of Kappa are still relevant today, that's a great question, and I think the broader question is what are sort of the -- you can think of Kappa as a reduced-form parameter capturing the cost of reallocation or the scope for reallocation, and there's sort of a lot of micro-foundations that are not in the paper, that are not in the model, that might be driving this.

And so more research on that would really be great. Some people are arguing for downward nominal wage rigidity, agglomeration effects, lack of social insurance, unions potentially. So I don't have a -- I don't take a stand there, but what I do know is that, for instance, if you compare results of the China Shock in the U.S. versus the China Shock in Germany, so it seems much more benign in Germany.

So people are arguing -- I'm not a full expert there -- that the labor market is functioning in a better way in Germany than the U.S., and so why that is the case would be interesting, and more recent, Peter Morrow and Brian Kovak are making a similar argument for the impact of trade

liberalization in Canada than in the U.S., that trade liberalization was more benign in Canada than in the U.S. And so the main question is why. I don't have an exact answer there, but it would be great to think more about it.

MR. AHMAD: Thank you. No, that is really helpful for us as we work forward.

MR. GALLE: Thank you.

MS. KHACHATURIAN: Thank you. Thank you, Saad, and thank you, Simon. I think we can take these next several minutes to sort of open the floor to more general questions from the audience. If the chat feature is not enabled, please feel free to raise your hand and ask a question.

Bill, since your camera is on and you had a question in the chat, do you want to just read the question out for Simon?

MR. POWERS: Yes, absolutely. I don't know, Simon, if you had a chance to look at it earlier. So you're taking a look at technology, changes in the technology level in China T, which makes a ton of sense, you know, 20 years ago when they were flooded with investment, which really is no longer the case. And I think Saad just mentioned we've had a number of other recent shocks which might have more to do with trade costs shipping between China and the United States.

If you introduce trade shocks into the model, do

you get an answer that's qualitatively different than introducing technologies? Obviously, you can introduce them, but do you get something that's qualitatively different, or does it all kind of boil down to the same change? Can you talk us through how -- obviously, there's great interest in expanding these types of models and their effects to more recent years. Any thoughts on how we can do that?

MR. GALLE: So, yes, we can use both technology shocks and trade shocks. My intuition -- so we haven't done a formal analysis for this model for comparing them. My intuition is that technology shocks and trade shocks are going to have similar effects on the receiving country, on the destination country, but different effects in the origin country.

So, if China is growing its technology, that's making the entire workforce in China more productive, so its PPF is expanding, so the aggregate gains of that would be larger, I believe, than if simply the trade costs between the U.S. and China that are changing. So the impacts on Chinese workers would be different, whereas the impacts on U.S. workers would be different. We can learn from that.

What we are focusing is essentially targeting the lambda hat, so the changes in the trade flows from China to the U.S., and so the ACR term is telling us that that's the key statistic to focus on for the aggregate effects. And

then I believe the intuition goes through for the distributional effects. Does that answer the question?

MR. POWERS: Yeah, that's great. I think that the difference across countries would also make a big difference if this were like an input-output structured model where you'd get things flowing across borders very differently. But, yes, absolutely, that makes a lot of sense.

MS. KHACHATURIAN: Thank you. Thank you, Simon. We have another question in the chat about capital mobility. So the question is, it seems that capital mobility is not explicitly addressed in the model's setup. Based on the assumption that labor is immobile across groups, does that mean that capital is assumed as fixed in this model?

MR. GALLE: So a few thoughts there. On the theory side, so the model does not have capital. That's a strong assumption, of course. That's the standard in these Eaton-Kortum-type models, gravity models, where there's only one input in production, which is labor.

And for standard for a Cobb-Douglas production function, adding capital would not add much, I believe, or it might amplify things, but the distributional effects might stay similar.

Another thing, so, also, production in a way is taking things at the U.S. level, so all workers in U.S. are making -- are having access to the same capital, so there's

no difference in -- I guess this is not the right way of thinking about it, but one can write a different model, which I do in my paper, where I look at both trade and technology, and there, workers in Detroit have access to the same level of capital as workers in New York or workers in San Francisco.

So, once you introduce capital, the location of capital in that particular model is not determined. One can make this more realistic, of course, and split U.S. in different countries and then have different capital levels in different sub-regions in the U.S.

As a final note there, in my other work, I do show that changes in the labor share, so when there is automation and you're going to use more capital, more machines in production in favor of workers, that that can also generate distributional effects, and those are important. That's outside a Cobb-Douglas setting. Sorry.

MS. KHACHATURIAN: Thank you. I don't see any other questions posted in the chat, and so I'm not sure I'm seeing -- I don't see the full screen here, so if you have your hand raised and I don't see it, please feel free to either turn on your camera and ask a question or just go ahead and ask the question. I'm scrolling through to see if there are any raised hands that I am missing.

MR. AHMAD: You're not missing any questions. I

don't see any raised hands.

MR. POWERS: We might want to call it because we are actually at time.

MS. KHACHATURIAN: Okay. I just wanted to make sure I didn't miss anything. But thank you so much, Simon, again, for kicking off our distributional effects seminar series. We thank you for your time and your participation. Thank you for engaging with our questions and the application to the policy work that we do here at the U.S. International Trade Commission. So, on behalf of the U.S. ITC, thank you, and we look forward to your future work in this area.

MR. GALLE: Thank you very much for having me and for the very insightful discussion. Thank you.

(Whereupon, at 11:00 a.m., the session in the above-entitled matter adjourned.)

//

//

//

//

//

//

//

//

//

CERTIFICATION OF TRANSCRIPTION

TITLE: Seminar 1: Slicing the Pie: Quantifying the Aggregate and Distrubtional Effects of Trade
INVESTIGATION NO.: Inv. No. 332-599

HEARING DATE: October 16, 2023

LOCATION: Washington, D.C.

NATURE OF HEARING: Seminar

I hereby certify that the foregoing/attached transcript is a true, correct and complete record of the above-referenced proceeding(s) of the U.S. International Trade Commission.

DATE: October 16, 2023

SIGNED: _____
Signature of the Contractor or the Authorized Contractor's Representative
1220 L Street, N.W. - Suite 206
Washington, D.C. 20005

I hereby certify that I am not the Court Reporter and that I have proofread the above-referenced transcript of the proceeding(s) of the U.S. International Trade Commission, against the aforementioned Court Reporter's notes and recordings, for accuracy in transcription in the spelling, hyphenation, punctuation and speaker-identification, and did not make any changes of a substantive nature. The foregoing/attached transcript is a true, correct and complete transcription of the proceeding(s).

SIGNED: _____
Signature of Proofreader

I hereby certify that I reported the above-referenced proceeding(s) of the U.S. International Trade Commission and caused to be prepared from my tapes and notes of the proceedings a true, correct and complete verbatim recording of the proceeding(s).

SIGNED: _____
Signature of Court Reporter