NBER WORKING PAPER SERIES

THE CONSUMPTION RESPONSE TO TRADE SHOCKS: EVIDENCE FROM THE US-CHINA TRADE WAR

Michael E. Waugh

Working Paper 26353
http://www.nber.org/papers/w26353

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
October 2019

I thank Chase Coleman for many conversations about this project; Spencer Lyon with whom I first discussed the ideas behind this project; participants at the 2019 SED meetings; and for CEMFI for hosting me during the 2019-2020 academic year and the support from Spain's State Research Agency MDM-2016-0684 under the María de Maetzu Unit of Excellence Programme. The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

The Consumption Response to Trade Shocks: Evidence from the US-China Trade War
Michael E. Waugh
NBER Working Paper No. 26353
October 2019
JEL No. A0,E0,E21,E65,F0,F13,F4,F6

ABSTRACT

This paper provides evidence on the consumption effects of trade shocks by exploiting changes in US and Chinese trade policy between 2017 and 2018. The analysis uses a unique data set with the universe of new auto sales at the US county level, at a monthly frequency, and a simple difference-in-difference approach to measure the effect of changes in trade policy on county-level consumption. As a lower bound, I estimate the elasticity of consumption growth to Chinese retaliatory tariffs to be around –1. This implies that counties in the upper quartile of the retaliatory-tariff distribution experienced a 3.8 percentage point decline in consumption growth. I further show that the consumption response corresponds with a decline in employment growth. These results suggest that Chinese retaliation is leading to concentrated welfare losses in the US.

Michael E. Waugh
Stern School of Business
New York University
44 West Fourth Street, Suite 7-160
New York, NY 10012
and NBER
mwaugh@stern.nyu.edu

A Code Repository is available at https://github.com/mwaugh0328/consumption_and_tradewar

1. Introduction

What are the distributional impacts of trade shocks? Evidence suggests that trade has had important distributional effects on the labor market.[1] From a welfare perspective, however, labor market outcomes may not reflect how economic welfare is allocated across those who are differentially impacted by trade.[2] Moreover, abrupt changes in trade policy will have different welfare consequences, depending on the opportunities households have to adjust to these shocks.[3] One solution to these issues is to study how consumption responds to trade shocks and learn about welfare and households' ability to adjust to them.

This paper moves beyond the labor market and provides new evidence on the consumption effects of trade shocks. I exploit changes in Chinese trade policy towards the US between 2017 and 2018 and study how a measure of consumption—new auto sales—responds to these trade shocks. Both visually and through formal econometric specifications, I find that changes in trade policy had large effects on consumption, with high-tariff counties experiencing at least a 3.8 percentage point decline in new auto sales growth relative to low-tariff counties.

The research design is simple: I exploit variation in a county's exposure to Chinese retaliatory tariffs on US products between 2017 and 2018 and correlate it with changes in consumption at the county level. The focus on Chinese tariffs on US products stems from a desire to measure a trade-induced change in labor income/production opportunities; e.g., soybean farmers in Iowa lose the ability to sell their product due to Chinese retaliation. Measuring consumption at the microeconomic level is, in general, difficult. My approach to measuring consumption is to use a unique data set with the universe of new auto sales at the US county level, at a monthly frequency, over the period 2017-2018 (other years will be added as data become available).

Results using data up to January 2019 show that changes in US-China trade policy had large effects on consumption. Visual and simple comparisons of means show that auto sales growth is about 2.5 percentage points lower in high-tariff counties relative to low-tariff counties (see Figure 2 for example) after the start of the trade war in July 2018. In my most basic econometric specification, I find that the elasticity of consumption growth to tariffs to be around -1—i.e., one percentage point increase in a county's exposure to Chinese retaliatory tariffs leads to a 1 percentage point decrease in auto sales growth. In terms of magnitudes, this elasticity implies

[1] Autor, Dorn, and Hanson (2013) show that exposure to Chinese import competition has led to losses in labor income and reductions in labor force participation for import-competition-exposed workers in the United States. Krishna and Senses (2014) show that increases in import penetration are associated with increases in labor income risk. Pavcnik (2017) surveys the growing body of evidence regarding trade's affect on earnings and employment opportunities.

[2] From a general perspective, see elaborations on this issue in, for example, Krueger and Perri (2006) or the survey in Attanasio and Pistaferri (2016).

[3] For example, in Lyon and Waugh (2019), we find that while trade has harsh consequences for labor markets, in welfare terms very few lose because of households' ability to smooth out the shock.

a nearly 3.8 percentage point decline in auto sales growth for counties in the upper quartile of the tariff distribution relative to counties in the lower quartile. Depending on controls and inclusion of various fixed effects, these magnitudes can be much larger with elasticities as large as -1.4 which translates to a 5.5 percentage point decline in auto sales growth for counties in the upper quartile of the tariff distribution.

I connect the decline in consumption with exports and employment growth. Using the same empirical strategy, I find that Chinese retaliatory tariffs reduced a county's total exports and negatively affected the labor market, and especially so for segments of the labor market who are the most sensitive to trade. In particular, I estimate a 1.50 percentage point decline in goods-producing employment growth for high tariff counties relative to low tariff counties. This evidence suggests that the decline in consumption is related to the negative labor market consequences caused by Chinese retaliatory tariffs.

Motivating this paper is the desire to measure how trade-induced changes in labor income or production opportunities feed into consumption. While prior work has traditionally focused on labor market outcomes, there is no empirical evidence that measures the labor-market-induced consumption effects of trade. Evidence on the response of consumption is important for evaluating the consequences of and appropriate policy responses to trade exposure for the following reason: The consumption response reveals the extent to which households can adjust to trade shocks. For example, if consumption does not change much—even though trade negatively affects the labor market—this suggests that these shocks are insurable, and hence the distributional consequences and welfare losses associated with exposure to trade are small. In contrast, if consumption changes a lot, this suggests that the labor market consequences are passing through to consumption, and hence there are important distributional consequences for welfare associated with trade. This paper's main result is more consistent with the latter interpretation: Chinese retaliation is leading to welfare losses.

A second motivation of this paper is simply to understand the economic consequences of the US-China trade war. A unique feature of the data used is the combination of geographic detail and high frequency. Given the many abrupt changes in policy (and announced policy), high-frequency data provide a unique perspective on how communities and households are quickly reacting to these changes. The geographic and high-frequency nature of my paper complements related analyses of the US-China trade war that study the aggregate affects of tariffs on US imports and prices of goods from China, e.g., see, e.g., Amiti, Redding, and Weinstein (2019); Fajgelbaum, Goldberg, Kennedy, and Khandelwal (2019); or Flaaen, Hortaçsu, and Tintelnot (2019). This paper focuses on a different mechanism: how Chinese tariffs on US exports affect consumption, depending on changes in a county's exposure to the retaliation. While the focus is on these distributional issues, a simple back-of-the-envelope calculation suggests that this mechanism—reduced consumption in response to Chinese retaliation—is as large as the aggregate effects found by Fajgelbaum et al. (2019).

2. The Economic Model

Motivating the research design is the quantitative framework developed (in collaboration with Spencer Lyon) in (Lyon and Waugh, 2018, 2019). Like existing work in the trade and labor markets literature, the framework builds on the idea that labor market adjustment is costly, and hence labor is exposed to changes in a market's trade orientation. That is, labor is not free to move and escape the negative effects of trade. Specifically, real wages within a labor market can be expressed as

$$w(\mathbf{s}, E\mathbf{s}'; \mathcal{S}, E\mathcal{S}'). \tag{1}$$

s is the island-level state, $\mathcal{S}$ is the aggregate state, and E is the expectation operator. The island-level state **s** depends on the tariff a labor market faces, world prices, and local productivity shocks. This formulation embeds the idea that a labor market may be depressed for several reasons: unfavorable trade exposure or unfavorable (local) productivity shocks. Also these reasons may be interrelated through the nature of comparative advantage. The aggregate state $\mathcal{S}$ would embed aggregate demand and productivity conditions. It would also embed the distribution of asset holdings across markets, which would affect wages through wealth effects in labor supply. Expectations about future states are made explicit here as wages today may depend upon the expectations about the future through intertermporal labor supply motives.

Wages in a labor market are connected with consumption through the households' consumption savings decision. Aggregating within a labor market, consumption per capita is

$$C(w, Ew'; \mathcal{S}, E\mathcal{S}') \tag{2}$$

and depends on labor earnings w in (1), the aggregate state $\mathcal{S}$, and expectations about future states. Earnings today and expectations about the future will determine consumption, depending on the extent of insurance, smoothing, and precautionary motives. Aggregate states and, in particular, the distribution of asset holdings within a market would also influence consumption.

Thinking through (1) and (2) motivates the empirical approach. The idea is that Chinese retaliatory tariffs on US products are shifting **s** (and $E\mathbf{s}'$) differently across labor markets. For example, counties in Iowa that produce soybeans and pork products will have their state variable **s** shifted by Chinese retaliation; in contrast, service-oriented markets such as New York City are not treated. The primary aim of the project is to directly measure C and see how these tariff-induced shifts pass through to consumption C.

Evidence on the response of consumption is important for evaluating the consequences of and the appropriate policy responses to trade exposure for the following reason: The consumption response reveals the extent to which households can adjust or are insured against the trade shock. Take, for example, a complete markets setting. There we would expect that a trade shock would result in no differential change in (2) across labor markets—even if the trade shock is inducing differential changes in the labor market in (1)! In this case, the change in trade policy and trade shocks would have *no distributional impact* on welfare.

In contrast, consider the polar opposite case in which households have no insurance opportunities and/or limited abilities to adjust to the shock (e.g., by moving to a new location). In this case, we would expect to find differential changes in consumption depending on a labor markets' exposure to the shock. In this case, the change in trade policy would have distributional impacts on welfare.

The consumption response discussed above differs from those typically considered in the trade literature. The standard mechanism through which trade affects consumption is more along the lines whereby US tariff increases during the trade war raised prices and lowered consumption (see, e.g., Amiti et al. (2019), Fajgelbaum et al. (2019), or Flaaen et al. (2019)). In the context of the discussion above, US tariff increases are aggregate effects and would affect all households.[4] In contrast, the goal below is to measure how Chinese tariffs on US exports differentially feed into consumption depending on changes in a county's exposure to the retaliation.

3. Data Overview

I combine multiple data sources to investigate how Chinese retaliatory tariffs affected consumption and then explore how they operate through trade and employment effects. The code and (when possible) the data is publicly posted at `www.github.com/mwaugh0328/consumption_and_tradewar`.

3.1. Tariff Data

Per the discussion above, my primary focus is on the Chinese government's retaliation for the tariffs the US imposed on Chinese goods beginning in the spring of 2018. US actions and Chinese reactions played out in several stages. Below, I provide a brief summary of the relevant events up to the start of 2019, which is when the auto sales dataset ends (until the 2019 data are available).

Timeline of the trade war. Below I outline the main sequence of events I consider—i.e., those

[4]This is imprecise. The precise statement would need to reflect a subtlety the idea that an aggregate increase in the price index would not lead to a uniform change in consumption simply because households differ in their asset holdings and, hence, their marginal propensity to consume. See, e.g., Carroll and Hur (2019), who make a related observation.

leading up to the start of 2019.[5] Bown and Kolb (2019) provides an excellent resource for understanding and tracking various aspects of the trade war.

In April 2017, the United States opened an investigation under Section 232 of the Trade Act of 1974 to ascertain whether steel and aluminum imports constitute a national security threat. Then, in August 2017, another investigation was opened under Section 301 of the Trade Act of 1974 to investigate whether Chinese trade practices are discriminatory and harmful to US intellectual property rights. These investigations were resolved in early 2018, with findings that steel and aluminum imports do pose a national security threat and that the Chinese government is conducting unfair trade practices related to technology transfer, intellectual property, and innovation. This finding set off the sequence of events outlined below:

- **March, 2018.** The US government increases tariffs on steel and aluminum products as a result of the Section 232 investigation.
- **April 2, 2018.** The Chinese government retaliates with tariffs on select products in response to the Section 232 tariffs.
- **April 3, 2018.** The US government releases a $50 billion list of Chinese products under consideration for 25 percent tariffs as a result of the Section 301 investigation. Before implementation, the list is revised in June.
- **April 4, 2018.** The Chinese government responds with its own $50 billion list of US products under consideration for 25 percent tariffs. Like the US list, the list is subsequently revised prior to implementation.
- **July 6, 2018: Phase 1.** Both the US and China impose tariffs on approximately $34 billion of their respective $50 billion lists.
- **August 23, 2018: Phase 2.** Both the US and China impose tariffs on the remaining $16 billion of their respective $50 billion lists
- **September 18-19 2018: New lists.** The US government finalizes its $200 billion list with tariffs ranging from 5 to 10 percent (and threats to raise the rate to 25 percent by January 2019). China finalizes its retaliation in the form of a $60 billion list, with tariffs also ranging from 5-10 percent.
- **September 24, 2018: Phase 3.** Both the US and China impose tariffs on their new lists.
- **December 1, 2018: US-China Tariff Truce.** Presidents Trump and Xi agree to halt any further escalation of the trade war and work toward a negotiated settlement with a deadline of March 2019.

[5] This is obviously in flux and will be updated as data become available.

County-level Tariffs. The source of the tariff data is Bown, Jung, and Zhang (2019). This list starts from China's most-favored-nation (MFN) tariff leading into the start of 2018 and then incorporates changes in tariffs arising from the finalized Chinese tariff lists (**Phases 1—3**), all at the HS10 product level. In addition to **Phases 1—3**, the data also incorporate Chinese retaliation against the US's Section 232 actions (steel and aluminum tariffs) and Chinese reductions in their MFN rates for various products.

Given the time-by-product variation in tariff rates, I construct a county-level measure of Chinese tariff exposure by month over the period 2017 to early 2019 using the following procedure. First, I merge the tariff lists with 2017 US exports to China at the HS6 level. Tariff lists and 2017 trade values are then assigned a three-digit NAICS code using the concordance published by the US Census. I then aggregate the tariff data to the three-digit NAICS level by taking a trade-weighted average of the tariff using 2017 trade values as weights. This procedure yields a tariff measure at the three-digit NAICS level, $\tau_{s,t}$, for NAICS code s at time t.

I apportion the tariff measure at the three-digit NAICS level to a county based on that county's total employment within a NAICS code. Specifically, a county's tariff is

$$\tau_{c,t} = \sum_{s \in S} \frac{L_{c,s,2017}}{L_{c,S,2017}} \, \tau_{s,t}, \tag{3}$$

where $L_{c,s,2017}$ is a county's 2017 employment in NAICS code s and $L_{c,S,2017}$ is total employment in the set of NAICS codes S. I use 2017 employment weights to avoid any impact the change in tariff may have on a county's employment structure, yet accurately reflect its the industrial composition when they are imposed. The set of NAICS codes S are those associated with private employment—i.e., government activities are excluded. The basic idea behind (3) is that if a county has a large share of employment in a high-tariff sector, then my county-tariff measure will reflect the high tariff. As an extreme example, if a county's employment is all in soybeans, then the county's tariff is the soybean tariff.

The first column of Table 1 reports summary statistics for the change in this tariff measure between December 2017 and December 2018. Across all counties, the average tariff increased by about 1.5 percent. The top panel of Table 1 breaks down the variation in tariff exposure by quartile of the tariff exposure distribution. Within the top quartile of the distribution, the tariff increased by about 4 percent, while there was essentially no change at the bottom of the distribution. While the imposed tariffs are large, their incidence in a county is much smaller. A large reason these values are small is because most employment within a county is not engaged in tradable producing activities.

Figure 1 provides a sense of the spatial variation in Chinese retaliation. It plots the change (not the level) in a county's tariff between December 2017 and December 2018. In this map, a

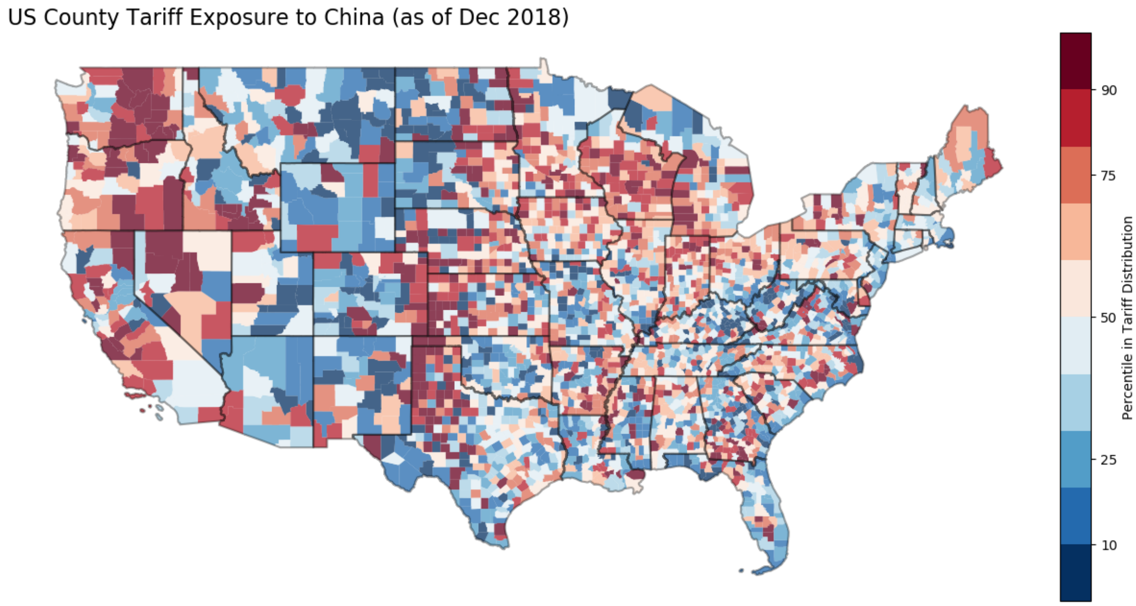

Figure 1: Tariff Exposure by County, Continental US ("lower 48")

county is colored according to its position within the distribution across counties; red indicates a county's tariff increased a lot and blue indicates that a county's tariff did not increase that much. Consistent with the notion that much of the Chinese tariff retaliation targeted agriculture commodities, much of the US midwest is heavily exposed to Chinese retaliation.

3.2. Auto Data

My measure of consumption at the county level is new auto sales. As provided by IHS Markit, the data set contains counts of new auto registrations (not values) by make (e.g., Ford) and model (e.g., F-150) and is geographically identified at the county level as determined by the locale of the entity registering the vehicle, not of the purchase. These data are derived from registration data purchased from State DMVs. Complete data are critical for the data vendor, as the data are sold/used in manufacturers' recall campaigns.

The use of data of this nature to proxy consumption expenditures is not unprecedented. Mian, Rao, and Sufi (2013) use similar data (from the same provider, but with different specifications) to study the consumption response to changes in home values between 2006 and 2009. Almunia, Antràs, Lopez-Rodriguez, and Morales (2018) use Spanish auto sales to study changes in local demand conditions and local firms' exporting behavior.

I currently have access to these data at the monthly frequency from January 2017 to January 2019 and for the year 2010. The monthly data for 2017 to 2019 form the core of the analysis. I focus only on lightweight vehicles; e.g., buses and semi-trucks are dropped. I (currently) do not exploit make and model variation. I simply aggregate counts of lightweight vehicles at the

Table 1: Summary Statistics: Tariffs, Autos, Trade, Employment

Δ Tariff Quartile	Δ Tariff	Autos	Exports to China	Total Emp.	Goods Emp.
Upper quartile	3.92	1,463	3,720	11,631	4,125
25th-75th quartiles	1.04	8,316	1,321	55,941	9,657
Bottom quartile	0.15	3,963	337	29,668	3,187
Average	1.54	5,525	1,805	38,144	6,624
Number of Counties	3,122				

Note: All values are for the year 2017; Δ Tariff is the change in the tariff between December 2017 and December 2018. Exports to China are on a per worker basis.

county level to proxy consumption.

There are several strengths of the data. First, an auto is a easy object to measure, and this dataset provides variation at both a narrow geographic dimension and at high frequency. High frequency is important in this context, due to the rapidly changing nature of trade policy during 2018. Moreover, there is more scope than I am currently using to exploit make-model variation and differential substitution patterns in response to shocks.

There are also weaknesses with the data. One is that an auto is a durable consumption good. Thus, there is a disconnect between the flow-consumption measure in economic models and the measure I observe in the data. Moreover, because of its durability, expectations of future outcomes may play a strong role and have less to say about the shock today. A second issue is that I only have access to counts and not purchase prices (or net purchase prices after trade-ins). Future work will use national average prices and then aggregate make/model variation at the county level based on sales. A third issue is that a broad array of entities (beyond households) register autos, e.g. business and governments. Restricting attention to lightweight vehicles helps on this dimension; presumably, households do not buy full size buses, semi-trucks, etc. Further work on other aspects of make and model variation can alleviate these concerns.

The second column of Table 1 reports annual auto sales for 2017. On average, a county has about 5,500 new auto registrations. Foreshadowing the employment numbers and differences in size across counties, large tariff increase counties are also smaller in employment, and hence have fewer new auto registrations. However, as a percentage of employment, high-tariff counties and low-tariff counties are quite similar (between 13 and 14 percent).

3.3. Trade Data

Trade flow data are important to examine the channels through which tariffs would affect production opportunities, income, and then consumption. I use US Census Monthly International Trade Data to measure trade flows. This data set provides monthly totals of imports and exports, at varying HS-code levels, and by source and destination (and more). Consistent with the auto data, I focus on the period from January 2017 onward. As with the tariff data, I start at the HS6 level and then aggregate to three-digit NAICS codes. This is done for US exports to China and total US exports.

My measure of exports at the county level is

$$\mathrm{EX}_{c,t} = \frac{1}{L_{c,S,2017}} \sum_{s \in S} \frac{L_{c,s,2017}}{L_{n,s,2017}} \mathrm{EX}_{s,t}, \tag{4}$$

where $\frac{L_{c,s,2017}}{L_{n,S,2017}}$ is a county's share of national employment in industry s and $\mathrm{EX}_{i,t}$ are exports associated with industry s. This measure is then put on a per worker basis by dividing through by a county's total employment, $L_{c,S,2017}$. The basic idea behind (4) is that if a nation's soybean employment is all in county c, then all soybean exports are apportioned to that county.

The third column of Table 1 reports summary statistics for exports to China for the year 2017. Of interest in Table 1 is the observation that high-tariff counties were also the most oriented toward Chinese trade. On a per worker basis, a high-tariff county has more than twice the level of exports to China relative to the average county.

3.4. Employment Data

The other channel to explore is how changes in tariffs affect labor market outcomes and, in turn, consumption. I use the BLS's Quarterly Census of Employment and Wages (QCEW) as the source of labor market data. The QCEW provides county-level employment and breaks county-level employment down by sector and by month for the US. The data primarily comes from the reporting of employment and wages to the Unemployment Insurance (UI) programs of the US. The QCEW covers about 97 percent of all wage and salary civilian employment in the country.

Exploiting the disaggregate nature of employment in the QCEW, I focus on two measures of employment. The first is total private employment; this excludes government employment. The second measure is private, goods-producing employment. This measure, because it concerns goods production, presumably is more tradable and, thus, susceptible to changes in trade exposure and tariffs.

While the strength of this dataset is its geographic and high-frequency coverage, it does have

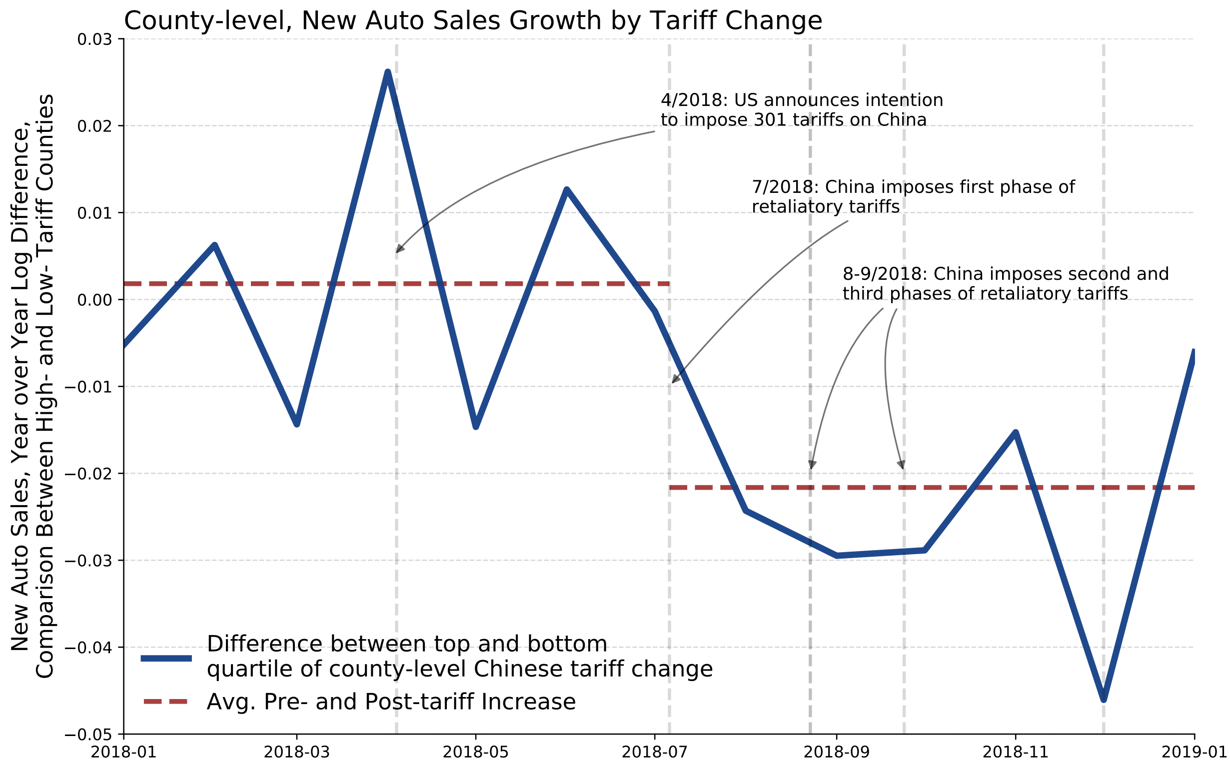

Figure 2: US County-level Auto Sales and Chinese Retaliatory Tariffs

several limitations. First, wages/earnings are only reported at quarterly frequency rather than monthly like the employment data. Thus, I do not have a direct measure of labor income at monthly frequency to match with monthly changes in tariffs and auto-sales. A second issue is that some employment figures are not reported within county-industry cells due to confidentiality concerns.

The final two columns in Table 1 report total and goods-producing employment on average and by position in the tariff distribution. First, notice that high-tariff counties are distinctly different from other counties in size. For example, the average county is almost four times larger than a county in the upper quartile of the tariff distribution. Second, high-tariff counties have a larger share of employment in goods-producing activities. Here high-tariff counties have about one-third of employment versus the average county, with a bit less than 20 percent.

4. Auto Sales and Chinese Retaliatory Tariffs

This section explores the impact of Chinese tariff retaliation on consumption as proxied by county-level auto sales. The analysis progresses through several steps, from simple visualizations and tabular representations to more formal regression analysis.

Table 2: Auto Sales Growth

Tariff Quartile	Pre-Trade War	Post-Trade War
Upper quartile	0.0129	-0.0269
	[0.005]	[0.005]
Bottom quartile	0.0111	-0.0052
	[0.005]	[0.005]

Note: Values are 12-month log differences averaged across counties and time periods. Pre-Trade War is January 2018 to June 2018; Post-Trade War is July 2018-January 2019. Standard errors are reported in brackets.

4.1. Difference-in-Difference by Visualization

I first visually illustrate the impact of Chinese tariff retaliation on county-level auto sales. I take 12-month log differences of auto sales which is my "first difference." This controls for any time-invariant county-level effects. This approach also and addresses county-specific seasonality issues (e.g., month$\times$county effect) which clearly standout when plotting the data in levels.

The "second difference" compares high Chinese-tariff counties versus low Chinese-tariff counties. Here high versus low is a comparison of counties in the upper quartile of the Δ tariff distribution and those in the lower quartile, as of December 2018 (see Table 1 or Figure 1). In other words, I compare auto sales growth across counties that had large increases in tariffs due to Chinese retaliation versus those that had small increases in tariffs.

Figure 2 plots this comparison between January 2018 and January 2019. Dashed vertical lines (with annotation) indicate important events during the trade war. Units on the y-axis are in log points, so an interpretation of the value of 0.01 is a 1 percentage point difference. Prior to the implementation of tariffs in July 2018, Figure 2 shows that there is no difference in auto sales growth between high- and low-tariff counties. A difference immediately emerges after the implementation of the first round of tariffs in July 2018. For the later half of 2018, high-tariff counties grew slower relatively to low-tariff counties. The magnitude is large, with a 2 percentage point difference between high and low tariff exposure counties.

Table 2 provides a tabular illustration. Prior to the trade war, auto sales in both county types were essentially growing at the same rate: about 1 percentage point. After the trade war, growth in both county types fell, but those in the upper quartile of the tariff distribution fell by 2 percentage points more.

Overall, Figure 2 and Table 2 provide strong, prima facie evidence that (i) prior to the trade war,

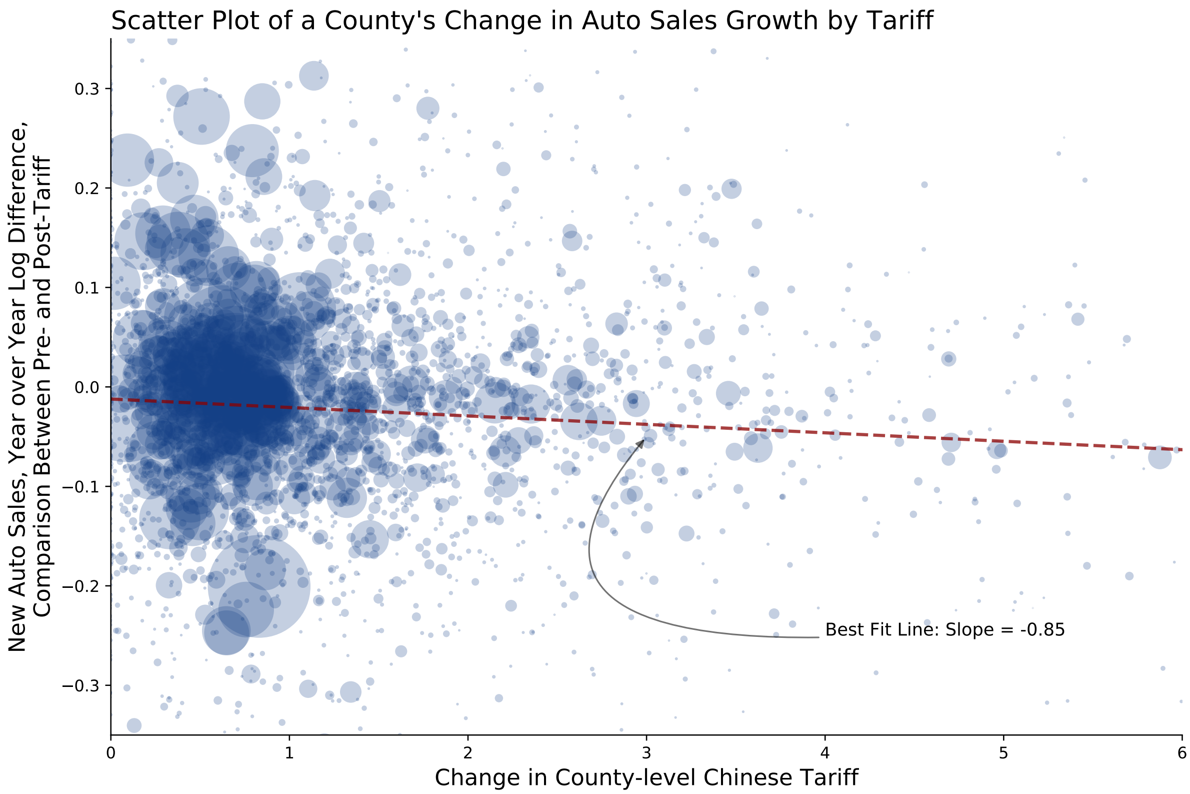

Figure 3: Auto Sales and Chinese Tariffs

both high- and low-tariff counties are essentially the same in terms of auto sales growth, but (ii) after the trade war there is a divergence, with high-tariff counties growing systematically slower.

Figure 3 illustrates the heterogeneity hiding behind simple means. Figure 3 is a scatter plot of a county's difference in auto sales growth before July 2018 and after July 2018 versus its tariff as of December 2018. The size of each county's bubble represents its total employment in 2017. Not surprisingly, the figure shows that there is a lot of variation in auto sales growth across counties. However, there is a systematic, downward-sloping relationship between the change in growth and a county's tariff exposure. Plotted on this scatter chart is the simple best fit line, which has a slope of -0.85. This is consistent with the magnitude of the change found in Figure 2 and Table 2. Moreover, it foreshadows the regression results, which I turn to now.

4.2. Formal Regression Analysis

This section moves beyond simple visualizations and tabular representations and explores the effect of tariff exposure on consumption in a more formal regression setting. The most basic empirical specification I explore is

$$\Delta \log C_{c,t} = \alpha_0 + \beta \Delta \log(1 + \tau_{c,t}) + \epsilon_{c,t}, \tag{5}$$

where $\Delta \log C_{c,t}$ is the 12-month log differences of auto sales in county c, and $\Delta \log(1 + \tau_{c,t})$ is the 12-month log differenced tariff rate. Again, differencing of this nature differences out any time-invariant county-level effects and any month-by-county effects. The parameter of interest is β, which measures how (in relative terms) a county's exposure to Chinese tariffs affects their consumption.

I also explore specifications that take the form of

$$\Delta \log C_{c,t} = \alpha_t + \beta \Delta \log(1 + \tau_{c,t}) + \epsilon_{c,t}, \tag{6}$$

which incorporates a time fixed effect. This specification has a long history as it is essentially the same as those in Townsend (1994), Cochrane (1991), and Mace (1991) as tests of risk sharing (see, e.g., Ljungqvist and Sargent (2012) for a textbook treatment). For example, in a complete markets allocation (with power utility) consumption growth should be orthogonal to any county-specific factors after conditioning on the aggregate state (which the time effect absorbs). Thus, an implication of a complete markets allocation is that the estimate of β should be zero. Using auto data as in this paper, Mian et al. (2013) employ a similar specification but with a measure net worth on the right-hand side rather than tariffs.

The final specification that I consider is more general and is

$$\Delta \log C_{c,t} = \alpha_c + \alpha_t + \beta \Delta \log(1 + \tau_{c,t}) + X^{'}_{c,t}\delta + \epsilon_{c,t}. \tag{7}$$

Here this second specification includes county fixed effects, time fixed effects, and county-by-time control variables. My main county-by-time control variable, $X^{'}_{c,t}$ will be employment growth from the QCEW discussed in Section 3. Note that because the data are already differenced, the inclusion of county fixed effects is controlling for county-specific growth. The inclusion of county fixed effects also changes the test of risk sharing interpretation since a complete markets allocation would imply that the county fixed effects should be zero.

A couple of comments are in order regarding the interpretation and specification of the error term $\epsilon_{c,t}$. From a purely econometric standpoint, the identifying assumption is that Chinese retaliatory tariffs are orthogonal to unobserved factors, i.e., $\epsilon_{c,t}$. Given that the outcome variable is differenced, this is the standard "parallel trends" assumption in difference-in-difference research designs. That is, absent the treatment from China, two counties will grow, in expectation, at the same rate for specifications (5 and 6) or at different rates but are county-specific in (7). Empirically, Figure 2 and Table 2 suggest that high- and low-tariff counties are growing at similar rates.

From an economics standpoint, the unobserved shock in these specifications are county-level productivity shocks. In the theory discussed in Section 2, the labor market state **s** contains both

Table 3: Auto Sales Growth and Chinese Tariff Exposure

	(5)	(6)	(7)	(7′)
$\log(1+\tau_{c,t})$	-0.96^{***} $[\ -1.63,\ -0.29\ ]$	-1.30^{***} $[\ -1.97,\ -0.63\ ]$	-1.44^{***} $[\ -2.36,\ -0.53\ ]$	-1.44^{***} $[\ -2.36,\ -0.52\ ]$
$\Delta \log L_{c,t}$				0.05 $[\ -0.40,\ 0.50\ ]$
Time Effects	N	Y	Y	Y
County Fixed Effects	N	N	Y	Y
# Observations	40,438			
Time Period	Jan 2017 - Jan 2019			

Note: Dependent variable is 12 month, log diffferenced auto sales. County-level observations are weighted by a county's employment in 2017. Standard errors are clustered at the county level; 95-5 confidence intervals are reported. $\Delta \log L_{c,t}$ is total employment in a county and 12 month log difference.

the tariff and the local productivity shock. From the theory's standpoint, the local productivity shock is the key unobserved (and possibly confounding) factor (see, e.g., the discussion in Lyon and Waugh (2019) in the context of the Autor et al.'s (2013) research design). The theory-consistent interpretation of the parallel trends assumption is that Chinese retaliatory tariffs are uncorrelated with the change in the local productivity shock.

In all specifications, the standard errors are clustered at the county level. In all specifications, county-level observations are weighted by a county's employment in 2017. Results with unweighted observations (or weighted with tradable employment) give similar results and they are posted in the code repository.

Table 3 reports the results.

Specification (5) in Table 3 shows that Chinese retaliatory tariffs had an economically and statistically meaningful impact on consumption. Table 3 report the results associated with (5) and 95-5 confidence intervals are reported below the point estimate. The point estimate is -0.96, and it is statistically different from zero. This point estimate can be interpreted as follows: A move from the lower quartile to the upper quartile of the tariff distribution implies a $-0.96 \times (3.92 - 0.15) = 3.6$ percentage point decrease in auto sales growth. This is larger than the magnitudes predicted by Figure 2 or Table 2.

Specification (6) includes the just time effects to capture aggregate changes over the time period. The point estimate rises now to -1.30.

Specification (7) and (7)′ in Table 3 further show that Chinese retaliatory tariffs had a economically and statistically meaningful impact on auto sales. Specification (7) includes both county and time fixed effects, and the point estimate for β increases to -1.44 and is statistically significant at the 5 percent level. Specification (7)′ includes monthly employment as a control variable and the point remains the same at -1.44 and is statistically significant at the 1 percent level. Starting from this point estimate, a move from the lower quartile to the upper quartile of the tariff distribution implies a $-1.44 \times (3.92 - 0.15) = 5.42$ percentage point decline in auto sales growth.

Across all specifications, the point estimates are significant and of similar orders of magnitude. As a rough ballpark, I find that a 1 percentage point increase in a county's exposure to Chinese retaliatory tariffs causes a 1.40 percentage point decrease in auto sales growth.

A couple of comments regarding these estimates. First, these estimates are probably lower bounds because of how autos sales are measured, i.e., in counts not values. One would suspect that there would be important intensive margin moments in the type of car purchased in response to the shock. That is, some consumers still purchase a new car, but purchase a less expensive car than they would have had the trade war not taken place. Future work could aggregate auto make and model based on national average prices to better reflect this effect.

The second comment regards interpretation. Per the discussion in Section 2, the evidence in Table 3 strongly supports the notion that the trade war is having distributional impacts on welfare. Obviously, this response is inconsistent with a complete markets, benchmark which would predict a β of zero, no change in consumption, and no distributional consequences. In contrast, the evidence in Table 3 is consistent with the notion that households have limited insurance or ability to adjust, and consumption is absorbing the shock. The implication is that Chinese retaliatory tariffs are harming segments of the US population.

In aggregate, are these effects large or small? My instinct is that this is the role of an economic model to infer and interpret these estimates—especially when thinking about the distributional impact. With that said, from an aggregate perspective, some simple tabulations suggest that the impacts—just from auto sales—are in the same ballpark relative to estimates of other effects from the trade war. On an annual basis, for the most exposed counties, this amounts to about 82 fewer cars sold based on 2017 values, $1,463 \times (.0144 * 3.92) \approx 82$ and about 64,000 fewer cars sold across all high-exposed groups. With an average purchase price of about $36,000 (see Kelly Blue Book), this is about $2.3 billion in lost sales. If we add in the moderately impacted groups, this is an additional $7 billion in lost sales. Lost auto sales alone, in response to Chinese retaliation (9.3 billion), is as large as the aggregate effects ($7.8 billion) found by Fajgelbaum et al. (2019).

5. Trade and Employment Effects

The previous section provides evidence that US county-level consumption responded to Chinese retaliatory tariffs. This section examines some channels that may explain as to why. I walk through this in several steps. First, I examine changes in exports to China and in total. Second, I explore the employment effects.

The top panel of Table 4 reports the effects of running similar specifications as with autos (equations (5), (6), and (7)) but with county-level US exports to China and county-level US exports in total on the left-hand side.

First, focusing on exports to China, the change in tariff exposure had a huge effect, with elasticities in the range of -20 to -10 depending on the structure of fixed effects. This simply verifies that Chinese retaliatory tariffs are doing what they were assumed to be doing. It might be tempting to interpret these coefficients as trade elasticities, and the structure in Lyon and Waugh (2019) is supportive of this interpretation, but more careful analysis is needed for this interpretation.

The next three columns are for total US exports. Here we find that Chinese retaliatory tariffs did have an effect on a county's ability to export, in total. Depending on the structure of the fixed effects, here the elasticity ranges from about -4 to a bit more than -1 and all are statistically different from zero at the 1 percent level.

These estimates are important because they suggest that exporters in high-tariff counties did not have the ability to simply redirect exports to other destinations. For example, one might suspect that Chinese retaliatory tariffs induced exporters to sell their products to other destinations; e.g., soybean farmers in Iowa sold their soybeans to Japan rather than to China. In this case, the bilateral Chinese retaliatory tariffs would have no effect on production, employment, and consumption. In contrast, these estimates suggest that for counties relatively more exposed to Chinese tariffs, it was hard for them to replace these lost export opportunities. And these lost export opportunities are one force that would lead to the reductions in consumption found in Section 4.

The bottom panel of Table 4 reports the results with employment. Across all specifications, the coefficients are negative, implying that relatively more exposed counties experienced reductions in employment growth.[6] For total employment, these point estimates are all around -0.25. Starting from this value, a move from the lower quartile to the upper quartile of the tariff distribution implies a $0.25 \times (3.92 - 0.15) = 0.94$ percentage point decline in employment growth. For goods-producing employment, the estimates are twice as large with an elasticity of

[6] These results are the one aspect which is sensitive to the weighting of county-level observations. Unweighed results are still negative, but the magnitudes are smaller, and statistical significance depends a lot on the structure of fixed effects.

Table 4: Trade, Employment, and Chinese Retaliatory Tariffs

	US Exports to China			US Exports		
$\log(1+\tau_{c,t})$	-21.12^{***}	-12.65^{***}	-10.12^{***}	-4.33^{***}	-1.64^{***}	-1.37^{***}
	[-22.27, -19.97]	[-13.92, -11.39]	[-11.50, -8.75]	[-4.56, -4.11]	[-1.77, -1.50]	[-1.56, -1.18]
Time Effects	N	Y	Y	N	Y	Y
County Fixed Effects	N	N	Y	N	N	Y
# Observations	53,226					
Time Period	Jan 2017 - Jun 2019					

	Total Employment			Goods Employment		
$\log(1+\tau_{c,t})$	-0.22^{***}	-0.29^{***}	-0.24^{***}	-0.35^{***}	-0.50^{***}	-0.40^{***}
	[-0.27, -0.19]	[-0.34, -0.23]	[-0.30, -0.18]	[-0.45, -0.26]	[-0.63, -0.38]	[-0.54, -0.26]
Time Effects	N	Y	Y	N	Y	Y
County Fixed Effects	N	N	Y	N	N	Y
# Observations	47,061					
Time Period	Jan 2017 - March 2019					

Note: County-level observations are weighted by a county's employment in 2017. Standard errors are clustered at the county level; 95-5 confidence intervals are reported.

around -0.40 across all specifications and statistically different from zero at the 1 percent level. This estimate implies 1.50 percentage point decline in goods producing employment growth for high tariff counties relative to low tariff counties.

Overall, the trade and employment effects connect well with the reductions in consumption. That is, for counties who were more exposed to Chinese retaliatory tariffs, these tariffs reduced a county's ability to export (top panel of Table 4), it feed into the labor market (bottom panel of Table 4) and this mechanism reduced consumption (Table 3).

With that said, these results do leave open some questions. A key issue regards magnitudes and if the results in Table 3 and Table 4 jointly make quantitative sense. Again, my instinct is that this is the role of an economic model to infer and interpret these estimates together, for which I leave to future work.

Another issue is that these results do not provide a complete picture about the labor market effects. I measure employment along the extensive margin which may not provide a complete picture of the labor market outcomes. Intensive margin changes (e.g. reductions in hours) could be taking place in the background that would induce changes in consumption as well. Similarly, reductions in earnings through wages cuts, loss of bonuses, etc, are other changes that may be taken place for which I do not observe.

A final issue is the role of expectations on the consumer side. In the context of the economic model discussed in equations (1) and (2), a key force concerns Es and Ew, not about the current states. The story would proceed along lines like this: Consumers foresee future negative consequences of the trade war for their county and react to the loss in future income (or increased uncertainty) by reducing consumption, as in Table 3. There are two reasons to be mindful of this story. First, returning to specification $7'$ in Table 3 (which included tariffs and labor market outcomes on the right-hand side), an interpretation of it is that something above and beyond tariff-induced changes in the labor market are affecting consumption. If the labor market mechanism were the only force through which changes in tariffs are operating, then I would have suspected that the coefficient on the tariff change would have gone to zero (which it did not).

The second reason concerns the nature of consumption that I am measuring. Autos are a durable good and typically financed though long-term arrangements, thus autos may be particularly sensitive to changes in expectations about future economic conditions. Contrasting these results with nondurable consumption may be one way to learn more about this mechanism.

6. Conclusion

These results may raise more questions than answers. Let me pose some that I think are interesting. The most obvious question is: What is going on now? The US-China trade war has, if anything, been escalating. The results in this paper extend to January 2019; the new data should be available by the beginning of 2020. The additions of these data should enrich the already intriguing results.

The second question concerns a more precise interpretation. As noted in the text, a formal economic model is needed—in particular, one that takes into account (i) the durable nature of consumption in the data I am using and (ii) can explore the joint relationship between changes in trade, employment, and consumption and (iii) can examine the idea that expectations play an important role. I leave this for future work.

These results also have several important policy implications. In the context of the current economic environment in the US, these results have policy implications for short-run demand management policy in the US and the appropriate response to the trade war. The conventional wisdom is that the trade war is a negative, aggregate supply shock with declines in output and inflationary pressure. In contrast, the trade-war-induced declines in consumption that I am finding suggest that there are important demand-side effects from the trade war for consideration in the formulation of monetary policy in the US.

Another policy implication concerns the design of policies to address the distributional impacts of trade. In particular, this paper's main result provides new evidence that changes in Chinese trade policy are leading to concentrated welfare losses. While the current situation in the US is self-induced, the results of this paper validate a broader point: policy should be cognizant of the distributional effects associated with changes in trade exposure and trade policy.[7]

[7] Lyon and Waugh (2018), Hosseini and Shourideh (2018), and Costinot and Werning (2018) are recent papers exploring the design of trade and tax policy in the presence of these distributional concerns.

References

ALMUNIA, M., P. ANTRÀS, D. LOPEZ-RODRIGUEZ, AND E. MORALES (2018): "Venting Out: Exports During a Domestic Slump," Tech. rep., National Bureau of Economic Research.

AMITI, M., S. J. REDDING, AND D. WEINSTEIN (2019): "The impact of the 2018 trade war on US prices and welfare," Tech. rep., National Bureau of Economic Research.

ATTANASIO, O. P. AND L. PISTAFERRI (2016): "Consumption inequality," *Journal of Economic Perspectives*, 30, 3–28.

AUTOR, D., D. DORN, AND G. H. HANSON (2013): "The China Syndrome: Local Labor Market Effects of Import Competition in the United States," *The American Economic Review*, 103, 2121–2168.

BOWN, C., E. JUNG, AND E. ZHANG (2019): "Trump Has Gotten China to Lower Its Tariffs. Just Toward Everyone Else," *PIIE Report*.

BOWN, C. AND M. KOLB (2019): "Trump's Trade War Timeline: An Up-to-Date Guide," *PIIE Report*.

CARROLL, D. AND S. HUR (2019): "On the Distributional Effects of International Tariffs," .

COCHRANE, J. H. (1991): "A simple test of consumption insurance," *Journal of political economy*, 957–976.

COSTINOT, A. AND I. WERNING (2018): "Robots, Trade, and Luddism," Working paper.

FAJGELBAUM, P. D., P. K. GOLDBERG, P. J. KENNEDY, AND A. K. KHANDELWAL (2019): "The return to protectionism," Tech. rep., National Bureau of Economic Research.

FLAAEN, A. B., A. HORTAÇSU, AND F. TINTELNOT (2019): "The Production Relocation and Price Effects of U.S. Trade Policy: The Case of Washing Machines," Working Paper 25767, National Bureau of Economic Research.

HOSSEINI, R. AND A. SHOURIDEH (2018): "Inequality, Redistribution and Optimal Trade Policy: A Public Finance Approach," Working paper.

KRISHNA, P. AND M. Z. SENSES (2014): "International Trade and Labour Income Risk in the U.S." *The Review of Economic Studies*, 81, 186–218.

KRUEGER, D. AND F. PERRI (2006): "Does income inequality lead to consumption inequality? Evidence and theory," *The Review of Economic Studies*, 73, 163–193.

LJUNGQVIST, L. AND T. J. SARGENT (2012): *Recursive macroeconomic theory*, MIT press.

Lyon, S. and M. Waugh (2019): "Quantifying the Losses from International Trade," .

Lyon, S. G. and M. E. Waugh (2018): "Redistributing the gains from trade through progressive taxation," *Journal of International Economics*, 115, 185–202.

Mace, B. J. (1991): "Full insurance in the presence of aggregate uncertainty," *Journal of Political Economy*, 99, 928–956.

Mian, A., K. Rao, and A. Sufi (2013): "Household Balance Sheets, Consumption, and the Economic Slump*." *Quarterly Journal of Economics*, 128.

Pavcnik, N. (2017): "The Impact of Trade on Inequality in Developing Countries," Working Paper 23878, National Bureau of Economic Research.

Townsend, R. M. (1994): "Risk and insurance in village India," *Econometrica*, 539–591.