

The Consumption Response to Trade Shocks: Evidence from the US-China Trade War

Michael E. Waugh
New York University and NBER

December 2019

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. I estimate the elasticity of consumption growth to Chinese retaliatory tariffs to be around minus one. This implies that counties in the upper quartile of the retaliatory-tariff distribution experienced a 3.8 percentage point decline in consumption growth. The fall in consumption corresponds with decline in both tradeable and retail employment. These results suggest that Chinese retaliation is leading to concentrated welfare losses in the US.

Email: mwaugh@stern.nyu.edu. 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, CEMFI, the University of Cambridge, the Federal Reserve Bank of Minneapolis, 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, and for financial support from NYU Stern's Center for Global Economy and Business. Code to reproduce aspects of the paper is posted at https://github.com/mwaugh0328/consumption_and_tradewar.

1. Introduction

What are the distributional impacts of trade shocks? In the context of the US-China trade war—the unprecedented tit-for-tat increase in tariffs by the US and China—this paper provides new evidence on the consumption effects of trade shocks. Both visually and through formal econometric specifications, I find that the trade war is inducing concentrated losses in consumption and employment for American communities most exposed to Chinese retaliatory tariffs.

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.

The focus on consumption stems from the desire to measure welfare. Consumption, rather than labor market outcomes, probably better reflects how economic welfare is allocated across those who are differentially impacted by trade. Moreover, the consumption response assists in revealing the opportunities households have to adjust to these shocks.¹ While studying consumption has many appealing features, 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 2016-2018 (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 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 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.

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

¹From a general perspective, see elaborations on this issue in, for example, Krueger and Perri (2006) or the survey in Attanasio and Pistaferri (2016). As a specific example, in Lyon and Waugh (2019), we find that while trade has harsh consequences for labor markets, but in welfare terms very few lose because of households’ ability to smooth out the shock.

are the most sensitive to trade. In particular, I estimate a 0.75 percentage point decline in employment growth for high tariff counties relative to low tariff counties and this estimate doubles for goods-producing employment. This evidence suggests that the decline in consumption is related to these negative labor market consequences.

I also find some evidence that retaliatory tariffs are affecting retail employment. Retail employment is another proxy for consumption expenditures (see, e.g., Guren, McKay, Nakamura, and Steinsson (2018)) and as a way to detect demand driven changes in local economic activity (Mian and Sufi (2014)). The evidence is mixed because my baseline specifications provide unclear results. But after controlling for a county's characteristics in a time-varying way, a sizeable negative effect of retaliatory tariffs on retail employment emerges. The decline in retail employment suggests that consumption, above and beyond autos, is falling and that effects of the trade war are spilling over into economic activity not directly exposed to retaliatory tariffs.

The final part of the paper looks for preexisting trends and announcement effects in the data. For both auto sales and employment, I find no evidence that highly exposed counties were growing in any different way relative to less exposed counties prior to the announcement of the trade war (see Figure 4 and Figure 5a). For auto sales, there is some evidence that consumption is reacting to the announcement of the US's intention to impose tariffs and China's retaliation in March and April of 2018. This and other evidence suggest that the decline in consumption may not be fully accounted for by the current changes the labor market, but driven by expectations and increases in uncertainty about future labor market opportunities.

Two different back-of-the envelope calculations suggest these effects are economically meaningful. Under an assumption that the estimated relative effects are the same as absolute effects, the auto estimates imply a decrease in aggregate consumption of up to \$54 billion. Moreover, this effect is concentrated with high tariff counties experiencing a \$1,600 per worker decrease in aggregate consumption. As a lower bound, the retail employment estimates imply a decrease in aggregate consumption of \$20 billion and \$630 per worker for those in high tariff counties.

Two ideas motivate this paper. The first 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 little empirical evidence that measures the labor-market-induced consumption effects of trade. Barrot, Loualiche, Plosser, and Sauvagnat (2018) is most closely related by studying debt and the China shock of the early 2000s.

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 wel-

fare 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 to understand the economic consequences of the US-China trade war. From a theoretical perspective, we should expect the trade war to affect welfare through two hardships. The first hardship is that US tariffs on Chinese goods will impose hardship on *all* US consumers through higher prices and a reduction in variety. This prediction is becoming apparent as Fajgelbaum, Goldberg, Kennedy, and Khandelwal (2019), Amiti, Redding, and Weinstein (2019), and Cavallo, Gopinath, Neiman, and Tang (2019) find that the US tariffs are leading to higher prices and a reduction in welfare for all consumers.

There is a second hardship of the trade war: Retaliatory Chinese tariffs should affect labor income and/or production opportunities for those directly impacted and lead to reductions in consumption and welfare. Unlike the price effects—which are spread widely across the US population—this hardship is concentrated. That is for those who had a position of comparative advantage for the Chinese market and lost it due to tariffs—they bear this burden of the trade war *alone*. The contribution of this paper is to use consumption data, at a narrow geographic detail and high frequency, to directly measure the size of this burden.

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

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

\mathbf{s} is the island-level state, \mathcal{S} is the aggregate state, and E is the expectation operator. The island-level state \mathbf{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 intertemporal 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

$$C(w, Ew'; \mathcal{S}, ES') \quad (2)$$

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 Es') 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.²

²This is imprecise. The precise statement would need to reflect a subtlety the idea that an aggregate increase

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 data set 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 leading up to the start of 2019.³ 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

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.

³This is obviously in flux and will be updated as data become available.

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.

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

US County Tariff Exposure to China (as of Dec 2018)

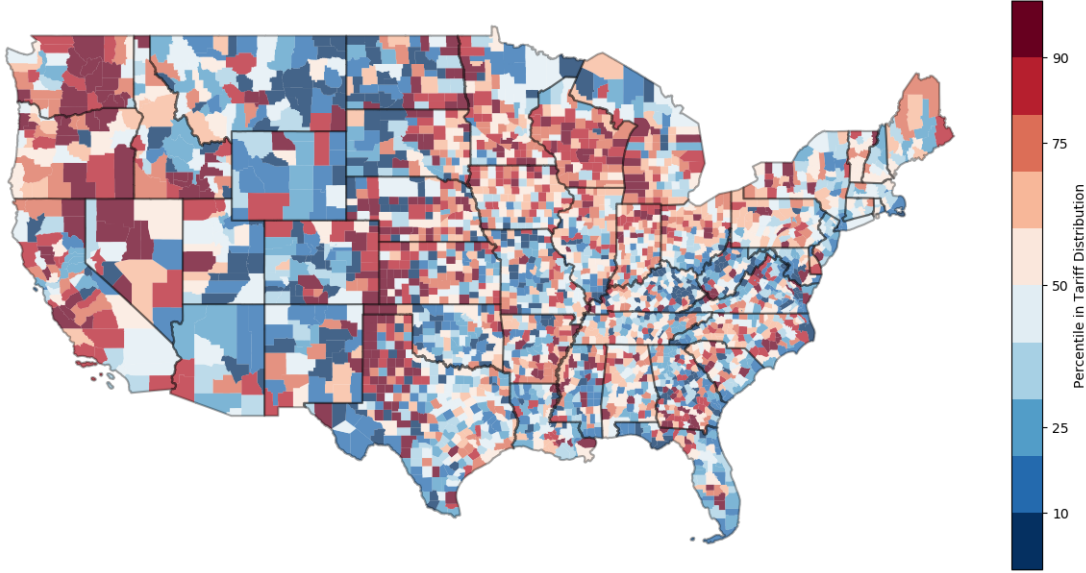


Figure 1: Tariff Exposure by County, Continental US (“lower 48”)

total employment within a NAICS code. Specifically, a county’s tariff is

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

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 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. I found that the share of employment in goods producing industries and the share of rural population account for a bit more than 17 percent of the variation in the change in tariffs. The importance of the goods share is consistent with the observations in Table 1 discussed below. In some ways, in that counties with the largest tariff increases are the most involved with the production things that the tariff could impact. Consistent with Figure 1, rural counties also happen to have experienced larger tariff increases. Surprisingly, the tariff change is uncorrelated with county-level income.⁴

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. For example, the Bureau of Economic Analysis reports 2017 annual sales of 17.55 million; in my data I have 17.58 million. In fact, NIPA documentation reveals that aspects of this from the same vendor are used in the construction of the new motor vehicular component of personal consumption expenditures.

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 2016 to January 2019. The monthly data for 2016 to 2019 form the core of the analysis. I focus only on lightweight vehicles; e.g., buses and semi-trucks are dropped. I do not exploit make and model variation. I simply aggregate counts of lightweight vehicles at the county level to proxy consumption.

There are several strengths of the data. An auto is a easy object to measure, and this data set provides variation at both a narrow geographic dimension and at high frequency. Geographic

⁴Casual observation also suggests that counties facing the most retaliation are also the counties which voted for President Trump in the 2016 US election. In a formal regression analysis, I do find that relationship, however, most the variation in tariff retaliation is simply explained by the export exposure of a county to China in the past.

variation is a prerequisite for the research design; high frequency is important in this context, due to the rapidly changing nature of trade policy during 2018. Moreover, this data set is essentially an “administrative” data set in the sense that it is the universe of auto sales and subject to little measurement error. In contrast, measurement error is a documented concern in standard survey data from the Consumer Expenditure Survey and the Panel Study of Income Dynamics. Moreover, these data sets have limited geographic variation and time windows which would restrict the research design.

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. As I discuss below, focusing on retail employment is a way to proxy for movements in non-durable consumption and supplement the analysis.

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 and 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.

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

Δ Tariff Quartile	Δ Tariff	Autos	Exports to China	Total Emp.	Goods Emp.	Retail Emp.	Population
Upper quartile	3.79	1,446	4,034	11,548	4,175	1,492	37,309
25th-75th quartiles	1.05	8,262	1,289	55,455	9,562	7,262	144,399
Bottom quartile	0.15	3,979	347	29,855	3,168	3,557	76,587
Average	1.51	5,525	1,879	38,144	6,624	4,896	101,000
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. Population is from the American Community Survey.

My measure of exports at the county level is

$$EX_{c,t} = \frac{1}{L_{c,S,2017}} \sum_{s \in S} \frac{L_{c,s,2017}}{L_{n,s,2017}} EX_{s,t}, \quad (4)$$

where $\frac{L_{c,s,2017}}{L_{n,s,2017}}$ is a county's share of national employment in industry s and $EX_{i,t}$ 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 several measures of employment. The first is total private employment; this excludes government employment. The second measure is private, goods-producing employment. This measure is a super-sector component reported in the QCEW. The broad sectoral components within the goods-producing super sector are natural resources and mining, construction, and manufacturing.⁵ Because this sector primarily concerns production of goods which are tradable, this category is probably the most susceptible to changes in trade exposure and tariffs.

The third measure is retail trade employment (NAICS 44-45) which consists of activities such as grocery stores, food and beverage stores, clothing stores, etc. This measure of activity is of interest for two reasons. First, Guren, McKay, Nakamura, and Steinsson (2018) make a compelling case that retail employment is a good proxy for consumer expenditures and, thus, its study complements the narrow measure of consumption (auto sales) that I'm using. Second, following the arguments of Mian and Sufi (2014), changes in non-tradable employment like retail reveal changes in local demand conditions. Thus, if retail employment falls with tariff exposure, this suggest that there are important local general equilibrium forces at play with the implication that the tariff is causing a fall in non-tradable consumption.

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

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.

⁵Doing something more disaggregate within this super-sector at the monthly frequency, county-level geography is problematic as the QCEW masks employment in instances (e.g., mass layoffs in small counties) where employers may be revealed. This gives rise to situations with non-random missing observations. Hence, I stick with the more aggregated super sector.




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

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., $\text{month} \times \text{county}$ effect) which clearly stand out 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 2019.⁶ 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 second half of 2018, high-

⁶Visually, there is no discernable trend if the window were pushed back to 2017, see Figure 4 and the discussion regarding pre-trends.

Table 2: Auto Sales Growth, Pre and Post Trade War

Tariff Quartile	Pre-Trade War	Post-Trade War
Upper quartile	0.013 [0.005]	-0.0269 [0.005]
Bottom quartile	0.010 [0.006]	-0.0052 [0.005]

Note: Values are 12-month log differences averaged across counties and time periods. Pre-Trade War is the period from January 2017 to June 2018; Post-Trade War is July 2018-January 2019. Standard errors are reported in brackets.

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 of this observation using data starting in January 2018. 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 (which is consistent with rising interest rates in the US). But the important observation is that, for those counties most affected by the trade war, auto sales growth fell by 2 percentage points more.

Overall, Figure 2 and Table 2 provide strong, prima facie evidence that (i) prior to the trade war, 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. Below, section 6 formally validates the observation that there is no differential trend across counties prior to the trade war and Section 4.2 formally quantifies the reduction brought about by tariffs.

Figure 3 illustrates the heterogeneity hiding behind simple means. Figure 3 is a scatter plot of a county's difference in auto sales growth from January 2018 to July 2018 and after July 2018 versus its tariff as of December 2018. The size of each county's bubble represents its population 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 -0.88 . 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.

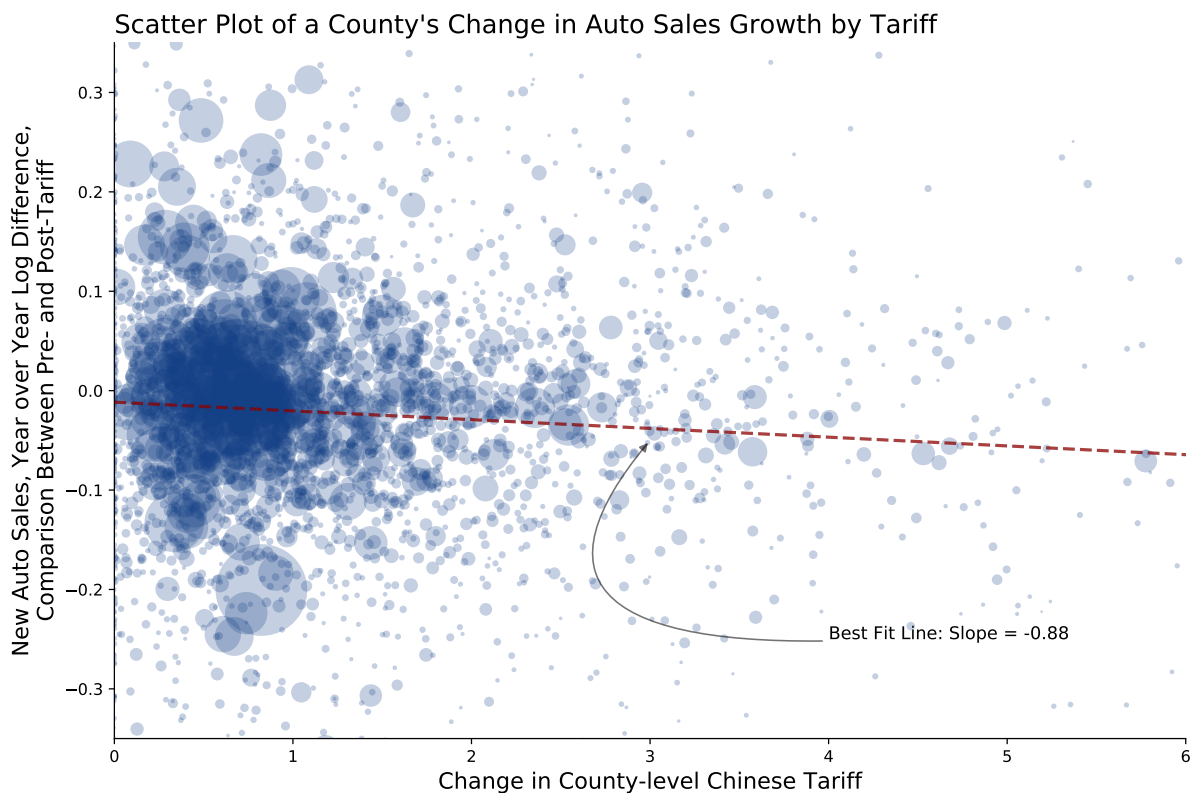


Figure 3: Auto Sales and Chinese Tariffs

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 main empirical specification I explore is

$$\Delta \log C_{c,t} = \alpha_0 + \alpha_t + \beta \Delta \log(1 + \tau_{c,t}) + X'_{c,t} \delta + \epsilon_{c,t}, \quad (5)$$

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.

The time fixed effect controls for any common, aggregate changes in the economy. Given the economic events of 2018, aggregate effects are important. The Federal Reserve increased US interest rates substantially over the course of 2018 which probably reduced aggregate auto demand and is consistent with the overall decline in growth seen in Table 2 in the second half of 2018. A second force are the aggregate effects of US tariffs against China or the aggregate effects of trade uncertainty which might also reduce overall demand for goods and services. Again,

the time effect here will control for these (and other) aggregate forces.

My main county-by-time control variable, $X'_{c,t}$ is employment growth from the QCEW discussed in Section 3.

The specification in (5) 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.

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, consumption in two counties will grow, in expectation, at the same rate for specifications. Empirically, Figure 2 and Table 2 suggest that high- and low-tariff counties are growing at similar rates; Section 6 explores the parallel trends assumption in more depth.

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 \mathbf{s} contains both 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). Thus, the 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 and county-level observations are weighted by a county’s population in 2010. Results with un-weighted observations (or weighted with employment) give similar results and are briefly discussed below; full results are posted in the code repository.

4.3. Auto Sales Results

Table 3 shows that Chinese retaliatory tariffs had an economically and statistically meaningful impact on consumption.

The first column reports the raw projection of county-level tariffs on consumption. Columns

Table 3: Auto Sales Growth and Chinese Tariff Exposure

	(1)	(2)	(3)	(4)	(5)
$\Delta \log(1 + \tau_{c,t})$	-0.97*** [0.22]	-1.06*** [0.24]	-1.03*** [0.24]	-1.00*** [0.33]	-1.01*** [0.33]
$\Delta \log L_{c,t}$			0.20*** [0.06]		0.07 [0.11]
Time Effects	N	Y	Y	Y	Y
County Fixed Effects	N	N	N	Y	Y
# Observations	40383				
Time Period	Jan 2017 - Jan 2019				

Note: Dependent variable is 12 month, log differenced auto sales. County-level observations are weighted by a county's 2010 population. Standard errors are clustered at the county level and are reported in brackets. $\Delta \log L_{c,t}$ is employment.

two and three report the results for the main specification in (5) with and without employment growth as a control. The point estimate for β , in both cases, is around -1.0 , and it is statistically different from zero at the one percent level. When monthly employment growth is included as a control, the point estimate for β remains essentially the same and the coefficient on employment is positive meaning that faster employment growth is associated with faster consumption growth.

The fourth and fifth columns in Table 3 include county fixed effects. This specification is the most stringent in that it controls for any time-invariant, county-specific differences in growth rates. Here the point estimate is essentially the same hovering around -1.0 and all are statistically significant at the 1 percent level with or without employment as a control.

These are economically large effects. Starting from the point estimate in column (3), a move from the lower quartile to the upper quartile of the tariff distribution implies a $-1.03 \times (3.79 - 0.15) \approx 3.8$ percentage point decline in auto sales growth. 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' consumption is absorbing the shock. The implication is that Chinese retaliatory tariffs are harming segments of the US population.

There are several robustness issues worth mentioning. One regards weighting. Alternative weighting schemes, such as total employment, generally lead to larger estimates than reported in Table 3. For example, the coefficient in specification (2) or (3) rises to -1.4 . Non-weighting,

i.e. treating each county identical, leads to smaller estimates in only the specification without county-fixed effects. Expanding the time window back to 2016 leads to similar or larger results. In particular, the point estimates doubles to -2 for the specifications with county-fixed effects which is consistent with the event-study/pre-trend results in Figure 4.

From a measurement perspective, these estimates are probably lower bounds because autos sales are 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.

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.

5.1. Exports Results

The top panel of Table 4 reports the effects of running similar specifications as with autos 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 -21 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.

Table 4: Exports and Chinese Retaliatory Tariffs

	US Exports to China			US Exports		
$\Delta \log(1 + \tau_{c,t})$	-21.7*** [0.80]	-13.7*** [0.82]	-9.74*** [0.89]	-4.36*** [0.20]	-1.76*** [0.15]	-1.61*** [0.13]
Time Effects	N	Y	Y	N	Y	Y
County Fixed Effects	N	N	Y	N	N	Y
# Observations	51,552					
Time Period	Jan 2017 - Jun 2019					

Note: County-level observations are weighted by a county's 2017 population. Standard errors are clustered at the county level and are reported in brackets.

Because these estimates combine both geographical variation and sectoral variation, comparing them to a traditional gravity-consistent trade elasticity is not straight forward. However, in Waugh (2019), I use the same exact tariff and trade data at the sectoral level (before it's projection down to the county-level) and estimate a trade elasticity. My estimates are between -4 and -3 which are consistent with previous estimates of the trade elasticity (see, e.g., Simonovska and Waugh (2014)) and in the same ball park as Fajgelbaum et al. (2019). This suggests is that while the trade war is unprecedented, the response of trade flows to changes in tariffs is not.

The next three columns are for total US exports. Here I 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, 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.

5.2. Employment Results

As discussed in Section 3, three measures of county-level employment are considered: total, goods-producing, and retail employment. And recall that goods-producing employment comprises mining, agriculture, manufacturing and construction activities. Outside of construction, all the other activities are directly exposed to changes in tariffs. Retail employment comprises employment in grocery stores, food and beverage establishments, etc. The presumption is that employment in these activities is only indirectly exposed due to local changes in demand for their services.

Table 5 reports the results with employment.

For total employment, these point estimates are between -0.21 and -0.15 and are statistically different from zero at either the one or five percent level. Using the value of -0.21 , a move from the lower quartile to the upper quartile of the tariff distribution implies a $0.21 \times (3.79 - 0.15) = 0.75$ percentage point decline in employment growth.

For goods-producing employment, the estimates are twice as large with elasticity of -0.47 in the specification with time effects. All are statistically different from zero at the 1 percent level.

Table 5: Employment Growth and Chinese Retaliatory Tariffs

	Total Employment			Goods Employment			Retail Employment		
$\Delta \log(1 + \tau_{c,t})$	-0.24*** [0.04]	-0.24*** [0.06]	-0.18** [0.06]	-0.44*** [0.08]	-0.50*** [0.10]	-0.36*** [0.13]	-0.23** [0.06]	0.07 [0.07]	0.14* [0.08]
Time Effects	N	Y	Y	N	Y	Y	N	Y	Y
County Fixed Effects	N	N	Y	N	N	Y	N	N	Y
# Observations	55,053								
Time Period	Jan 2017 - June 2019								

Note: County-level observations are weighted by a county's 2010 population. Standard errors are clustered at the county level and are reported in brackets.

This estimate of -0.47 implies 1.70 percentage point decline in goods producing employment growth for high tariff counties relative to low tariff counties.

For retail employment, the estimates are all over the place. Absent any controls, the estimate is negative and significant; with time-effects, it is essentially zero; with time and county-level effects it becomes positive and statistically significant at the ten percent level.

What is going on with retail employment? The issue is that there are time-varying county effects not controlled for which are correlating with a county's exposure to retaliatory tariffs. To explore this issue, I estimated the following specification

$$\Delta \log L_{c,t} = \beta \Delta \log(1 + \tau_{c,t}) + \sum_{y=B1,2018}^{B3,2019} \left(1 \{t = y\} X'_c \delta_y \right) + \alpha_t + \alpha_0 + \epsilon_{c,t} \quad (6)$$

where the second term interacts a time-dummy with fixed county characteristics. The county characteristics that I consider are the ones that I found correlate with a county's tariff exposure: a county's goods employment share (in 2016) and an interaction between time and a county's rural population share (as measured by the Census in 2010). The time dummies are at the bimonthly frequency and was chosen to preserve the high-frequency fidelity of the data, yet reduce some of the noise that arises at the monthly frequency. The idea here is to control for any time-varying shocks that are specific to rural or goods-producing areas. Also note that because of the interaction terms, I am unable to include county fixed effects.

Table 6 reports the results from (6) for all measures of employment. Focusing on retail employment, the point estimate is negative, large, and statistically significant after including the

Table 6: Employment and Tariffs, Specification (6)

	Total	Goods	Retail
$\Delta \log(1 + \tau_{c,t})$	-0.26*** [0.06]	-0.47*** [0.12]	-0.20** [0.08]
Time Effects	Y	Y	Y
# Observations	54,927		
Time Period	Jan 2017 - June 2019		

Note: Results from specification (6). County-level observations are weighted by a county's 2010 population. Standard errors are clustered at the county level and are reported in brackets.

flexible interaction terms. Goods employment remains essentially the same as in the time-effect specification in Table 5. Consistent with the negative retail employment response, the elasticity of the tariff to total employment increases by a third. For goods-producing employment and auto-sales, estimates from (6) are little changed relative to those in Table 5 and Table 3.

Inspecting the interaction terms confirms the discussion above. Both the time-interacted rural-share dummies and goods-producing-share dummies are often significant. However, the most important set of interactions are the ones with goods-producing share. What appears to be occurring is that retail (and total) employment in counties oriented towards goods-producing actives (and to an extent rural) performed relatively better over this time period and *increasingly so* over the second-half of 2018. And because counties oriented towards goods-producing actives are also more likely to be exposed to Chinese tariffs, the negative tariff effect is being masked unless these events are controlled for.

Returning to the insignificant or positive point estimates on retail employment in Table 5, the issues are easier to understand now. The interactions reveal non-time-invariant factors and, hence, county effects do not capture these changes. And because these factors are specific to counties oriented towards goods-producing activities, a common time effect does not capture this either.

The point estimate of -0.19 implies a $0.19 \times (3.79 - 0.15) = 0.70$ percentage point decline in retail employment growth. This point estimate also has an interpretation for overall consumption growth as well. Guren et al. (2018) find (both in the time series and the cross-section) an elasticity of retail employment to personal consumption expenditures to be about one. Thus, using this elasticity, the point estimate in Table 6 suggests that overall consumption fell by 0.70

percentage points in the most exposed counties relative to the least exposed counties.

Unlike the auto-sales results, these point estimates are far more sensitive to weighting issues. In the specifications for Table 5, not weighting county-level observations pushes the point estimates for total employment and retail employment essentially towards zero; goods-producing employment estimates are reduced by about half and still significant. The specifications in Table 6 with the interaction terms are less sensitive to weighting issues. The point estimates for total and goods-producing employment reduced by about half and (surprisingly) the retail employment point estimate is virtually the same as reported in Table 6.

6. Preexisting Trends and Announcement Effects

The key threat to the validity of the interpretation that Chinese retaliation is leading to reductions in consumption and employment is some violation of the parallel trends assumption. That is, absent the treatment from China, consumption or employment in high versus low tariff counties would have grown, in expectation, at the same rate.

There are several plausible scenarios that could violate this assumption. One scenario relates to structural change and regional divergence issues. That is high tariff counties were already experiencing slower consumption and employment growth relative to low tariff counties because of, say, rural-urban migration or broader structural change (e.g., shift away from goods producing activities). As discussed in Section 3, counties facing more severe Chinese retaliation were those with larger rural populations and more employment in goods producing activities.

A second scenario relates to unobserved shocks that spuriously correlate with the change in tariff exposure. An example would be a negative agriculture shock (e.g., widespread flooding or poor weather conditions) or a negative demand or productivity shock concentrated on US manufacturing.⁷ In this example, because tariff retaliation correlates with agriculture and manufacturing intensive areas, the results may reflect the unobserved shock and not the effect of Chinese tariffs. Per the discussion around the retail employment effects above, this is an important concern.

To explore these issues, my strategy is to project consumption and employment growth on (i) interactions between time and the change in the tariff as of the final available observation (January 2019 for autos, March 2019 for employment) and (ii) interactions between time and a county's goods employment share and an interaction between time and a county's rural population share. The interactions between time and the final tariff is designed to check for any preexisting trends and announcement effects. The later set of interactions is designed to pick up trends and concentrated shocks that are associated with relatively more rural or agricul-

⁷These are not hypothetical scenarios. In the 2015-2016, the QCEW shows very clearly that areas intensive in goods producing activities were severely depressed. And in 2019 parts of the US midwest faced severe flooding.




Figure 4: Auto Sales and Chinese Tariffs, Jan 2017 - Jan 2019

ture and manufacturing intensive counties. Moreover, these are the same set of interactions explored in Table 6 and specification (6).

Specifically, the empirical specification for both consumption and employment growth is

$$\Delta \log Y_{c,t} = \sum_{y=B1,2017}^{B1,2019} (1 \{t = y\} \beta_y \Delta \log(1 + \tau_{c,2019})) + \sum_{y=B1,2018}^{B1,2019} (1 \{t = y\} X'_{c,t} \delta_y) \quad (7)$$

$$+ X'_{c,t} \lambda + \alpha_t + \alpha_0 + \epsilon_{c,t},$$

where the first sum on the right-hand side interacts a time dummy with the final tariff $\tau_{c,2019}$ as of January 2019 (in the case of autos). The second sum interacts the time dummy variables with observable county characteristics: a county's goods employment share (in 2016) and a county's rural population share (as measured by the Census in 2010). The time dummies are at the bimonthly frequency and was chosen to preserve the high-frequency fidelity of the data, yet reduce some of the noise that arises at the monthly frequency. As before, standard errors are clustered at the county level, observations are weighted by a county's population.

Figure 4 plots the results for auto sales. For each bi-monthly time period, the coefficients β_y and 90-10 confidence intervals associated with them are plotted. For all of 2017 and the first two months of 2018, the estimates of β_y are near zero and not statistically different from zero. This

means that the increase in a county's tariff from Chinese retaliation as of 2019 is uncorrelated with auto sales growth for all of 2017 and the first two months of 2018. This evidence supports the parallel trends assumption and the conclusion that Chinese retaliatory tariffs are causing the decline in auto sales growth.

Starting with in March/April of 2018, the β_y coefficient becomes statistically different from zero and essentially remains that way until the end of the window in January 2019. And the overall magnitudes are consistent with those found in Table 3, i.e. an elasticity of around minus one. The fact that the coefficient becomes negative and statistically different from zero around the announcement but prior to the implementation of the tariffs in July 2018, suggests that consumers are reacting to the news of the forthcoming trade war. That is forward-looking consumers in March/April stop purchasing cars in anticipation of the negative consequences associated with the tariff war.


Very similar patterns emerge when looking at employment growth. Figure 5a - 5c show the same analysis for total employment, goods-producing employment, and retail employment.

For both total employment (Figure 5a) the estimates of β_y are near zero and not statistically different from zero for all of 2017 and the first several months of 2018. Then starting in March/April of 2018, the β_y coefficient becomes negative and increasingly so for the rest of 2018 and into 2019.

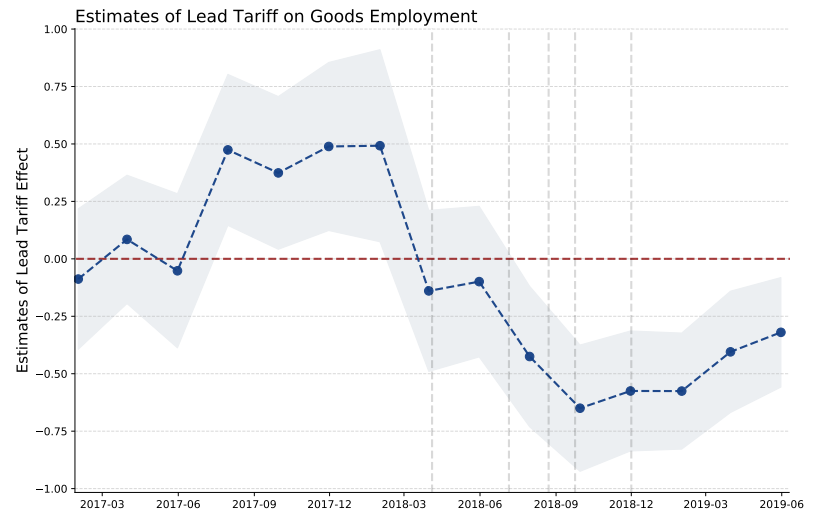
Goods-producing employment Figure 5b displays a very similar pattern with one exception. There is a brief period that counties more exposed to retaliatory tariffs in 2019 experienced faster employment growth at the end of the end of 2017 and 2018. The point estimates are meaningful and just statistically significant at the ten percent level. Per the discussion around retail employment, this is another piece of evidence that Chinese retaliatory tariffs seemed to hit well performing areas in the US economy.

Figure 5c displays the results for retail employment. These are a bit less clear relative to the previous findings, but supportive of the interpretation that retaliatory tariffs are hurting retail employment. For all of 2017 and the beginning of 2018, the estimates of β_y are negative and, with the exception of the very first period, all are statistically indistinguishable from zero. After the announcement of the start of the Trade war, the coefficients systematically become more negative and become statistically different from zero.


Overall, the behavior of these employment measures prior to the start of the trade war is supportive of the parallel trends assumption. And the behavior of employment around the announcement suggests some anticipatory effects in response to just the news of the trade war, though this is less clear relative to auto-sales growth.



(a) Employment (Total) Growth, Jan 2017 - Jun 2019



(b) Employment (Goods) Growth, Jan 2017 - Jun 2019



(c) Retail Employment Growth, Jan 2017 - Jun 2019

7. Discussion

Given the results, this section discusses interpretation, open questions, and provides a simple, back of the envelope calculation behind the distributional and aggregate effects of Chinese retaliatory tariffs.

7.1. Interpretation

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 exports (Table 4), it reduced employment (Table 5) and this mechanism reduced consumption (Table 3). Moreover, because retail employment is affected (Table 6) the consumption effects are likely broader than changes in the decision to purchase a car.

With that said, these results do leave open some questions.

One question regards magnitudes and if the results in Table 3 and Table 5 jointly make quantitative sense. My instinct is that this is the role of an economic model to infer and interpret these estimates together. Given the discussion below about the role of expectations and uncertainty in shaping these results this is an important, but complicated question which I leave for 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. One piece of evidence that suggests earnings must be falling is the result of Cavallo et al. (2019) who find that export prices (excluding the tariff) to China are falling which suggests that producers are bearing the incidence from the retaliatory tariffs.

A final issue is the role of expectations and uncertainty. In the context of the economic model discussed in equations (1) and (2), a key force concerns E_s and E_w , not about the current states. The story would proceed along lines like this: Consumers foresee future negative consequences of the trade war and react to the loss in future income by reducing consumption. They could also be reacting through precautionary motives due increased uncertainty about future income.

There are two reasons to be mindful of this story. First, returning to specification (3) 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 fallen (which it did not).

The second piece of evidence that expectations are playing a role are the responses seen in Figure 4 and Figure 5a - 5c. For auto sales, the mere announcement of a trade war appears to get consumers to respond. And there is also some evidence that employers are responding through employment decisions.

7.2. The Aggregate and Distributional Effects: Two Back-of-the-Envelope Calculations

The estimates in Table 3 and Table 5 are relative effects across counties. The often asked question is how these results map into aggregate outcomes. There are well known difficulties in answering the aggregate question with the research design that I employ in this paper. One aggregation approach is to operate under the assumption that the relative effects are the same as the absolute effects (see, e.g., Autor et al. (2013) as one example of this approach). This is what I do below.

Operating from a point estimate of -1.04 for the effect of tariffs on auto sales growth in specification (3) of Table 3, I calculate the lost car sales for each month of the trade war, for each county and then sum across all counties. Using this approach, I find that this results in $\approx 67,000$ lost autos. On a base of about 17.5 million autos sold in 2017, this is a decline of 0.39 percent. With an average purchase price of \$36,000 (see Kelly Blue Book), this works out to about \$2.5 billion in lost sales.

One can infer the aggregate consumption response with an estimate of the relationship between total expenditure and auto sales. Aguiar and Bils (2015), using the Consumer Expenditure Survey, estimate an elasticity of vehicle expenditure to overall expenditure to be between 0.72 and 1 (see their Table 2). Starting from a base of \$14,000 billion in personal consumption expenditures for 2017, these expenditure elasticities imply that Chinese retaliatory tariffs reduced aggregate consumption by between 0.28 and 0.39 percent or between \$40 billion and \$54 billion.

A second way to estimate the aggregate effect is to use the retail employment estimate and then connect it with consumption. Operating from a point estimate of -0.19 (Table 6), I calculate the forgone retail jobs for each month of the trade war (for employment this goes to June 2019). Summing across all counties, I find that this results in about 22,500 lost retail jobs. On a base of about 15.8 million retail jobs in 2017, this maps into a 0.14 percent decline in retail employment. I can then use Guren et al.'s (2018) finding that a one percent change in retail employment maps into a one percent change in personal consumption expenditures. Starting from a base of \$14,000 billion, this implies that Chinese retaliatory tariffs reduced aggregate consumption by 0.14 percent or \$20 billion.

For both these numbers, it's important to remember that these aggregate effects are highly

concentrated on a small few. Using the estimates from retail employment, counties in the upper quartile of the tariff distribution are experiencing a \$630 decrease in consumption per worker. In contrast, the estimate for the least exposed is fourteen times smaller—counties in the lower quartile of the distribution experience a \$46 decrease in consumption per worker. Estimates from auto sales result in numbers which are twice as large—counties in the upper quartile of the tariff distribution are experiencing between a \$1,204 and \$1,673 decrease in consumption per worker vs. \$91 and \$126 decrease in the least exposed communities. These calculations suggest that Chinese retaliation is leading to large and concentrated welfare losses in the US.

8. 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 auto 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.⁸

⁸Antràs, De Gortari, and Itskhoki (2017), 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

- AGUIAR, M. AND M. BILS (2015): "Has consumption inequality mirrored income inequality?" *American Economic Review*, 105, 2725–56.
- 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 U.S. Prices and Welfare," Working Paper 25672, National Bureau of Economic Research.
- ANTRÀS, P., A. DE GORTARI, AND O. ITSKHOKI (2017): "Globalization, inequality and welfare," *Journal of International Economics*, 108, 387–412.
- 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.
- BARROT, J.-N., E. LOUALICHE, M. C. PLOSSER, AND J. SAUVAGNAT (2018): "Import Competition and Household Debt," .
- 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," .
- CAVALLO, A., G. GOPINATH, B. NEIMAN, AND J. TANG (2019): "Tariff Passthrough at the Border and at the Store: Evidence from US Trade Policy," Working Paper 26396, National Bureau of Economic Research.
- 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," Working Paper 25638, 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.
- GUREN, A. M., A. MCKAY, E. NAKAMURA, AND J. STEINSSON (2018): "Housing wealth effects: The long view," Tech. rep., National Bureau of Economic Research.
- HOSSEINI, R. AND A. SHOURIDEH (2018): "Inequality, Redistribution and Optimal Trade Policy: A Public Finance Approach," Working paper.
- 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.
- MIAN, A. AND A. SUFI (2014): "What explains the 2007–2009 drop in employment?" *Econometrica*, 82, 2197–2223.
- SIMONOVSKA, I. AND M. E. WAUGH (2014): "The Elasticity of Trade: Estimates and Evidence," *Journal of International Economics*, 92, 34–50.
- TOWNSEND, R. M. (1994): "Risk and insurance in village India," *Econometrica*, 539–591.
- WAUGH, M. E. (2019): "Measuring Trade Elasticities," Tech. rep.