The Curse of a Strong Dollar

PIER Research Exchange – preliminary, feedback welcome

Scott Sommers
USITC

Michael E. Waugh
Federal Reserve Bank of Minneapolis
NBER and @tradewartracker

Teerat Wongrattanapiboon
Bank of Thailand

Presented by Teerat Wongrattanapiboon

August 31, 2026

The views expressed herein are those of the authors and not necessarily those of the U.S. International Trade Commission, the Federal Reserve Bank of Minneapolis or the Federal Reserve System, or the Bank of Thailand. This project was developed with research support from the National Science Foundation (NSF Award number 1948800).

Big Picture — USD Appreciation 2014-2017

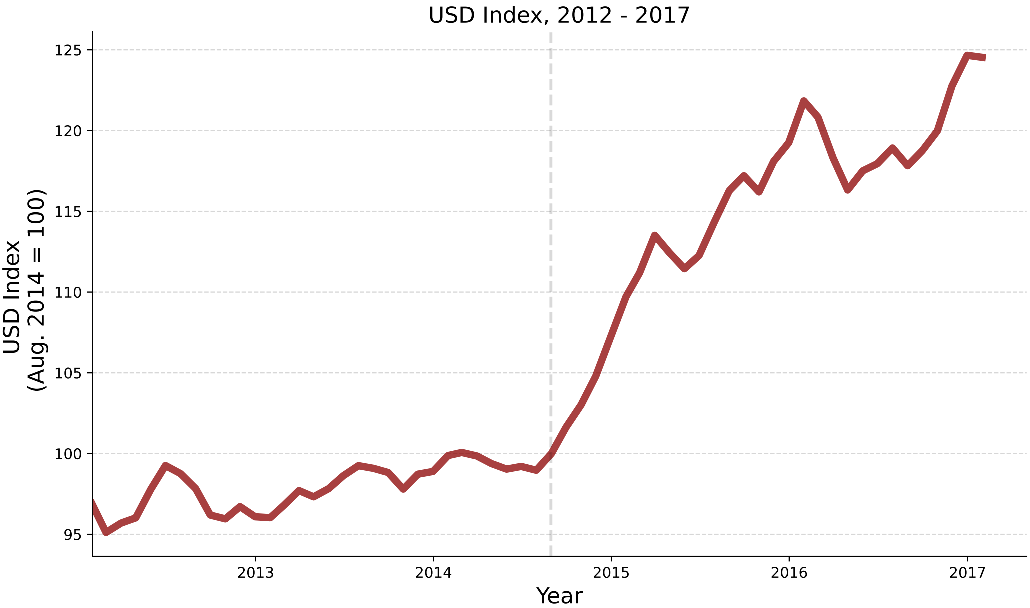

Preview of Main Result: USD Appreciation Reduced Employment

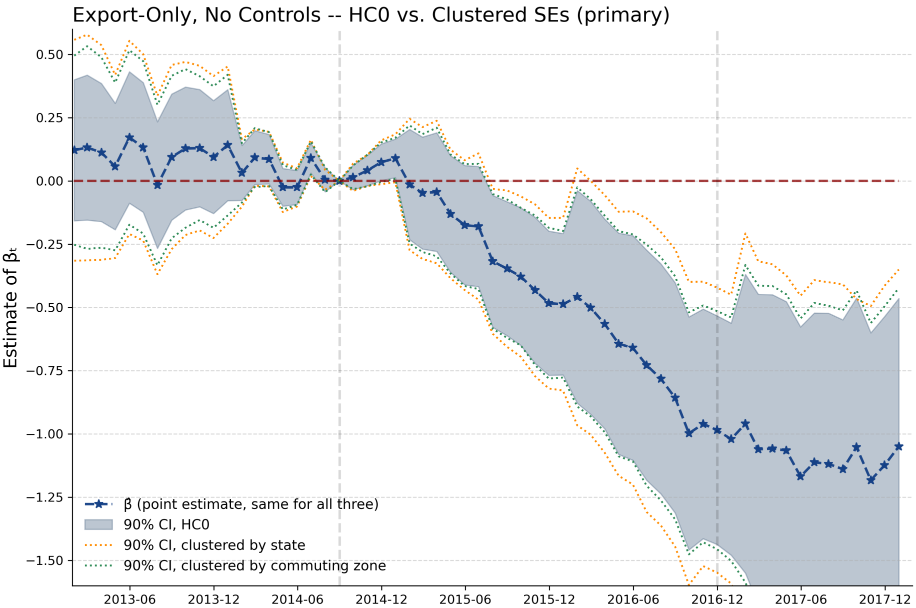

This Paper...

1. Construct dollar exposure measures at narrow geographic level.

- Similar to Autor, Dorn, and Hanson (2013), US-China trade war studies (Waugh (2019) or Autor et al. (2024)), effects of the financial crisis Mian, Rao, and Sufi (2013) or Yagan (2019) .
- This approach helps circumvent reverse causality issues that plague empirical FX-studies.
- Supplement this with narratives (Orchard, Ramey, and Wieland (2023)) from FRB Tealbooks showing that the USD appreciation was (i) unexpected and (ii) external to events in the US.

2. Measure dollar-induced changes in labor market outcomes, consumption, and political outcomes.

- Evidence so far points to USD appreciation having reduced employment and consumption in high- relative to low-exposure counties.
- Evidence that USD appreciation shifted political support towards Republicans, in both the 2016 presidential and US House elections.

Today's goal: walk through the design and get your feedback on where it's weakest – not present this as a finished result.

Background — Why We Think This Move Was Exogenous

Read the publicly available Federal Reserve Tealbooks and FOMC transcripts, 2014-2016.

- Leading into summer 2014, the June Tealbook forecast a modest dollar **depreciation**; instead the dollar appreciated nearly 5% by October – a move in the 94th percentile of all equivalent-length historical windows.
- The move was broad-based (against nearly all currencies except the heavily-managed Renminbi) and driven by deteriorating growth/monetary conditions *outside* the US (UK, Euro Area, China, Japan) – not by US-specific news.
- Tealbooks were repeatedly surprised by the appreciation through 2015-2016; FOMC transcripts show real puzzlement, not a policy response to it.
- Oil prices collapsed over the same window ($100 mid-2014 → $25 late-2016) – a contemporaneous, potentially confounding shock we return to later.

Conceptually What We Want To Measure

Different channels through which the dollar might matter:

1. Direct exports — domestic exporters become **less** competitive abroad, employment falls.
2. Direct imports — imported varieties are cheaper, domestic firms become **less** competitive at home, employment falls.
3. Indirect imports — imported intermediates are cheaper, domestic firms become **more** competitive at home, employment rises – omitted entirely today.

- **Today we show exports only, for simplicity of exposition.**
- **New evidence:** exports respond sharply to FX exposure, imports barely at all – channel 2 looks quantity-muted, not just slow (see backup).

Step #1: Construct Sectoral Dollar Exposure

This is the Fed's Broad Dollar Index, but computed **separately for each sector** j (a NAICS 3-digit industry), using that sector's own export destinations n (one of the 26 countries in the Broad Dollar basket) as weights.

Step A – the weight: share of sector j's exports going to country n:

$$w_{j,n,t} = \underbrace{\frac{\texttt{exports}^{j}_{USA,n,t}}{\sum_{n' \in N} \texttt{exports}^{j}_{USA,n',t}}}_{\text{sector } j\text{'s export share to country } n} .$$

Step B – the sectoral index: chain last period's index forward using the export-weighted average bilateral exchange rate change:

$$\texttt{FX}^{x}_{j,t} = \underbrace{\texttt{FX}^{x}_{j,t-1}}_{\text{last period's index}} \times \underbrace{\prod_{n \in N} \left(\frac{\mathrm{ER}_{USA,n,t}}{\mathrm{ER}_{USA,n,t-1}} \right)^{w_{j,n,t}}}_{\text{export-weighted avg. bilateral ER change}} ,$$

- E.g., Computer & Electronic Products (NAICS 334) was the top U.S. export to Thailand in 2013 (~\$2.1B) – its sectoral index picks up some baht exposure too.

Step #2: Project Down to the County Level

We then project $\mathtt{FX}_{j,t}$ down to a geographic area (today is a county) and construct that county's exposure to fluctuations in the dollar.

The **Local Export-Weighted Dollar Index** is:

$$\mathtt{FX}^{x}_{c,t} = \sum_{j \in J} \left(\frac{L_{c,j,2013}}{L_{c,J,2013}} \right) \times \left(\frac{X_{j,2013}}{Y_{j,2013}} \right) \times \mathtt{FX}^{x}_{j,t} \ ,$$

- $\mathtt{FX}^{x}_{j,t}$ is our sectoral index,
- $X_{j,2013}$ is total exports of sector j and $Y_{j,2013}$ is gross output of NAICS code j,
- $L_{c,j,2013}$ is county c's 2013 employment in sector j and $L_{c,J,2013}$ is total employment.
- **Note the shares here are frozen at 2013** – deliberately, to dodge reverse causality from the dollar move itself.

Outcome Variables

Employment: We use the BLS's Quarterly Census of Employment and Wages (QCEW) as the source of labor market data.

- At a monthly frequency, county level,
- Results shown today use **total private employment**; goods-producing employment is in the merged data but hasn't been run through the event-study spec yet.

Consumption: Our proxy is counts of new auto sales (not values), from IHS Markit.

- At a monthly frequency, county level (by locale of registration, not sale), by make and model.
- **This data is licensed/confidential** – by design, Mike runs the auto-sales analysis locally rather than the raw data being checked into the shared repo. Results from that side aren't shown today; happy to discuss the construction conceptually.

Political Outcomes: Republican two-party vote share, county level – both the 2012 and 2016 presidential elections, and the 2012 and 2016 US House elections.

Summary Statistics

Δ FX Quantile	$\%\Delta\mathtt{FX}$	Total Emp.	Goods Emp.	Autos	Rep. Share (2016, in %)
Upper quartile	1.55	22,537	6,579	3,003	59
25th-75th quartiles	0.49	52,299	8,668	7,471	46
Bottom quartile	0.07	19,550	2,150	1,990	47
Average	0.61	35,210	6,177	4,934	49
Number of Counties	3,192				

Note: All values are for the year 2013, except for Republican two-party vote share; Δ FX is the change in the trade-weighted exchange rate index between January 2013 and December 2016. Employment is data from BLS's Quarterly Census of Employment and Wages.

Change in Export FX Exposure by County

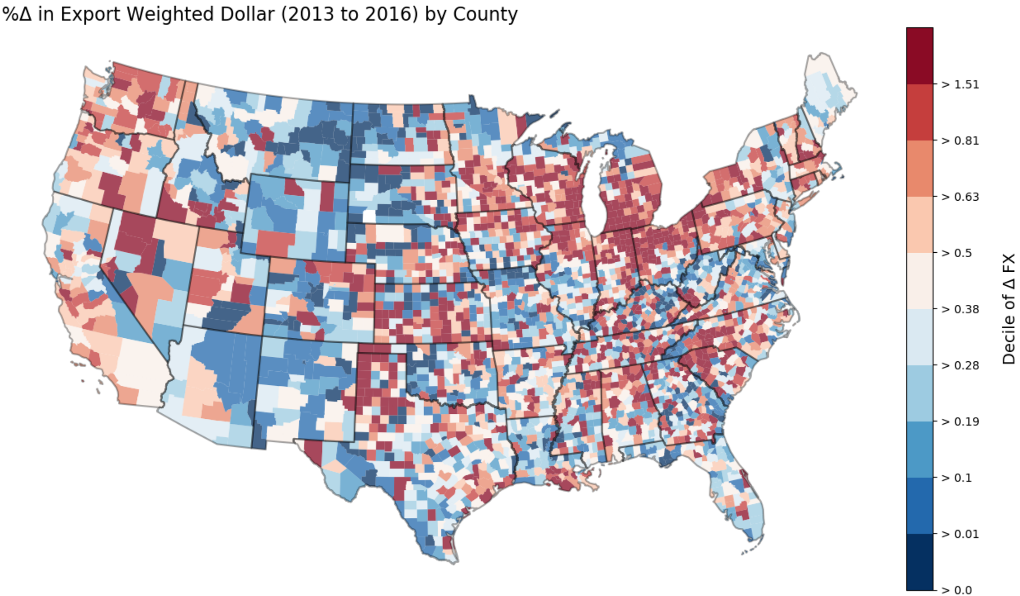

Research Design

Explore different permutations of the following specification:

$$\log Y_{c,t} - \log Y_{c,2014-8} \;=\; \alpha_t + \beta_t \Delta \mathtt{FX}_{c,2016-12} + X_c^{'} \lambda_t + \epsilon_{c,t},$$

- $\Delta \mathtt{FX}_{c,2016-12}$ is the percent change in the index between August 2014 and December 2016 – a *frozen*, single "dose" per county.
- β_ts are the coefficient of interest...
  answers how the change in $\mathtt{FX}$ exposure affected employment and consumption.
- λ are coefficients on county-level characteristics; rural share, log median hh income.
- And we do this for every date t between 2013 and 2017.

A standard specification, similar to Autor et al. (2024) or Yagan (2019) – but *not* staggered-adoption DiD, so the classic TWFE critique doesn't directly apply. Its continuous-treatment analogue might.

No Controls — USD Appreciation Reduced Employment

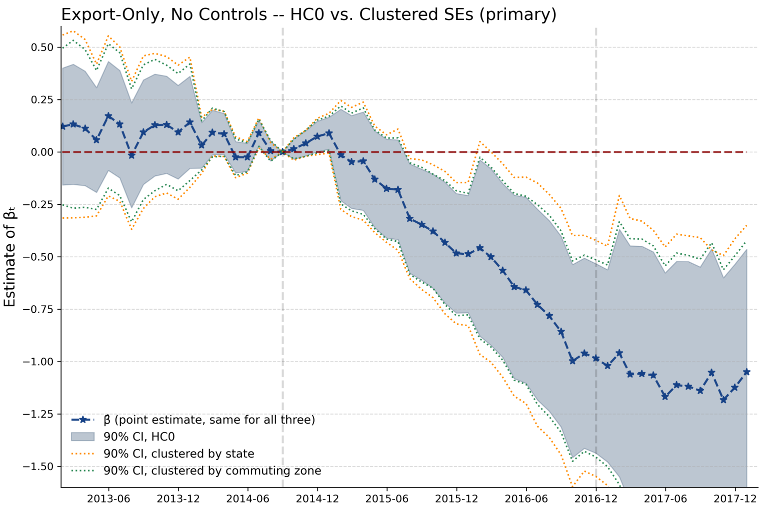

- **Magnitude:** a 10% dollar appreciation is associated with roughly **0.6 pp** less cumulative employment growth in the most export-exposed counties relative to the least-exposed.

With Controls — USD Appreciation Reduced Employment

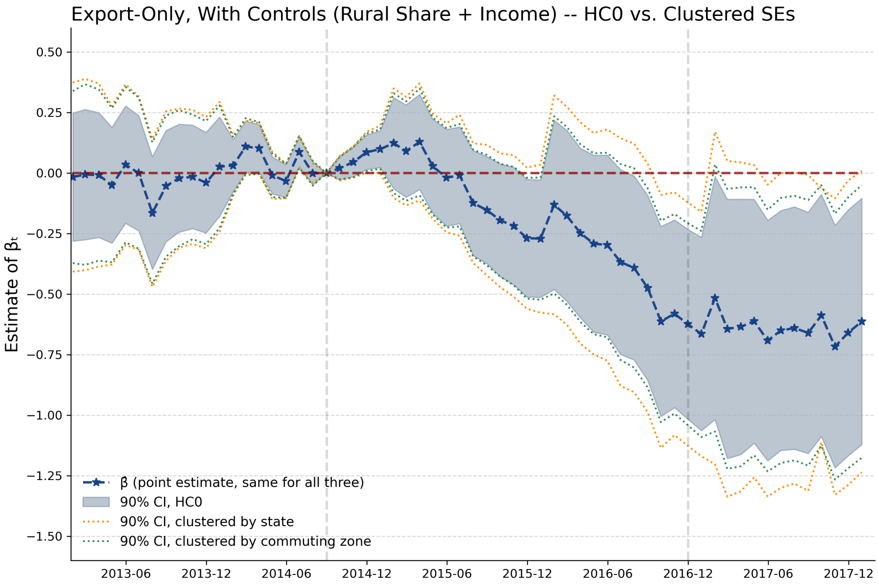

Research Design — Political Outcomes

Explore different permutations of the following specification:

$$R_{c,2016} - R_{c,2012} \;=\; \alpha + \beta \Delta \mathtt{FX}_{c,2016-12} + X_c^{'} \lambda + \epsilon_{c,t},$$

- $R_{c,t}$ is the Republican two-party vote share at date t – estimated separately for two outcomes: the presidential vote and the US House vote.
- β is the coefficient of interest...
  answers how the change in `FX` exposure affected the Republican vote share, in both races.
- λ are coefficients on county-level characteristics; oil exposure (built the same way as `FX` exposure – detail in backup), rural share, log median hh income.

Δ Republican Vote Share and FX Exposure

	Presidential		House	
	(1)	(2)	(3)	(4)
ΔFX	**1.77***** [0.41]	**0.74**** [0.29]	**2.61***** [0.72]	**1.97***** [0.61]
log Income		**-5.33***** [0.78]		**-6.55***** [1.89]
Rural Share		**11.86***** [0.69]		**7.22***** [1.38]
Oil Exposure		**0.26**** [0.12]		**-0.14** [0.17]
R^2	0.02	0.47	0.01	0.08
# Observations	3,108			

Note: Dependent variable is the change in the Republican two-party vote share between 2012 and 2016 (columns 1–2: presidential; columns 3–4: House). County-level observations are weighted by a county's 2013 population.

Identification and Robustness — Where Things Stand

#	Concern	Status
1	Shift-share exogeneity	**Checked – passes**
2	Heterogeneous dose-response	*Derisked* (next slide)
3	Inference / SE clustering	**Checked – survives**
4	Export vs. import-competition exposure	**Checked – too collinear to separate (r=0.84–0.87), see item 7**
5	Omitted shocks (oil, China, demand, Fed)	**Checked – survives (Fed funds tier-1 only)** (next two slides)
6	Sample/weighting	**Checked – real finding** (following slide)
7	Trade-flow response (exports vs. imports)	**Checked – exports respond, imports don't (cf. item 4)**

Where we'd value your input: item 2 and how to frame item 6's large-county result.

Dose-Response Linearity — Prerequisite for Item 2

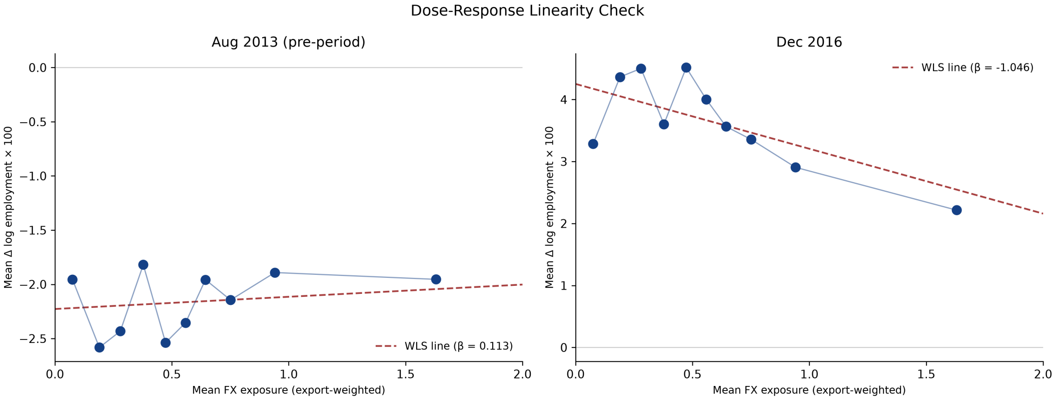

Note: Counties sorted by Dec 2016 export-weighted FX exposure and binned into 10 equal-population bins. Dots show population-weighted mean employment growth vs. mean exposure within each bin. Dashed line is population-weighted OLS (WLS) fit on individual county data, no controls. Employment growth measured as log change relative to Aug 2014 base (×100).

Omitted Shocks — FX Effect Survives Oil Controls

Built the same way as the `FX` index – WTI oil price in place of the dollar – across three channels: **direct** (oil & gas employment), **upstream** (suppliers to oil & gas), **downstream** (users of oil & gas as an input). Full formulas in backup.

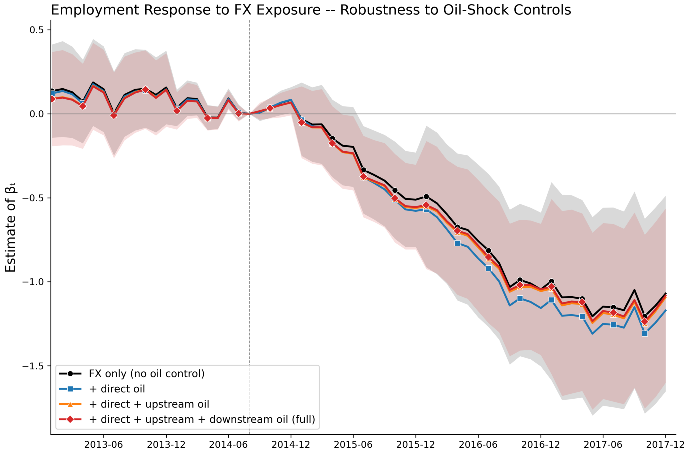

Shaded bands: 90% CI (HC0) on the no-control and full-control bookend specs.

Omitted Shocks — FX Effect Survives China-Shock Controls

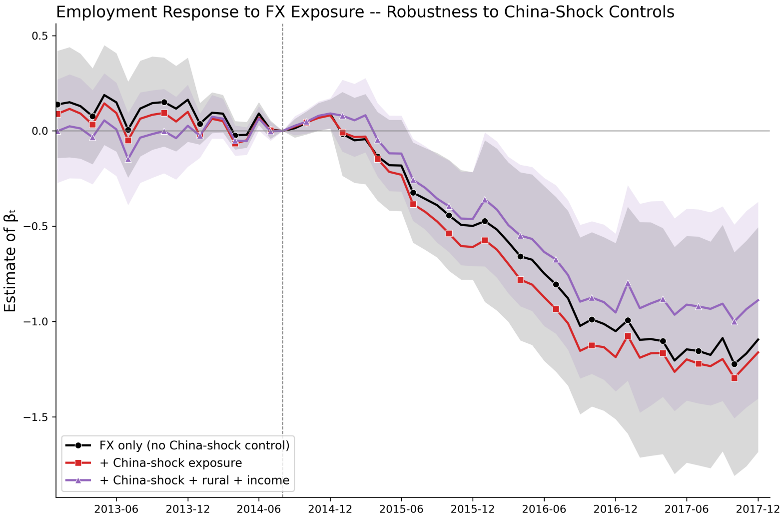

Shaded bands: 90% CI (HC0) on the no-control and full-control bookend specs.

Omitted Shocks — FX Effect Survives Global Demand Controls

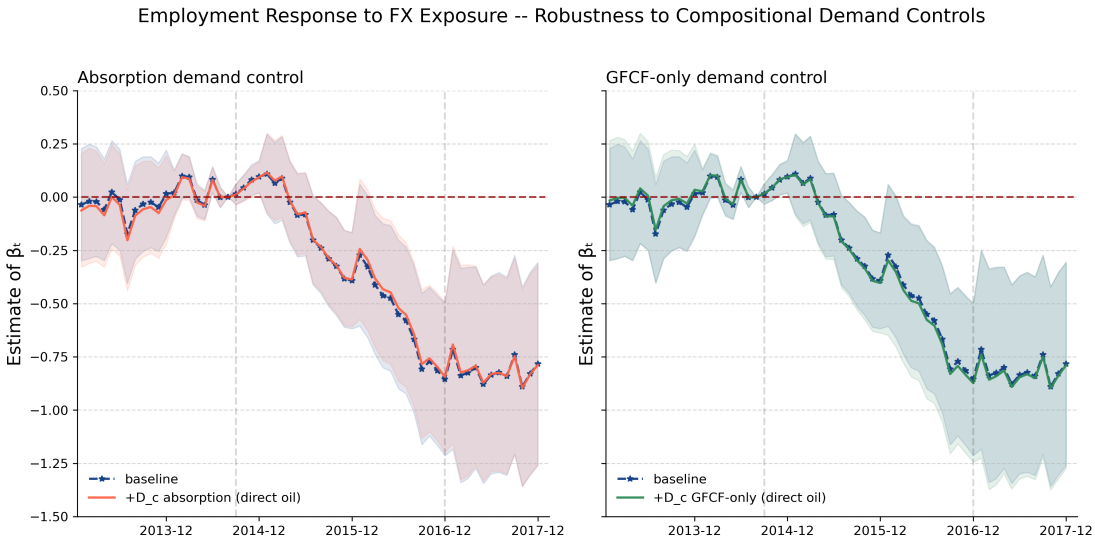

Sample/Weighting — Effect Is Driven by Large Counties

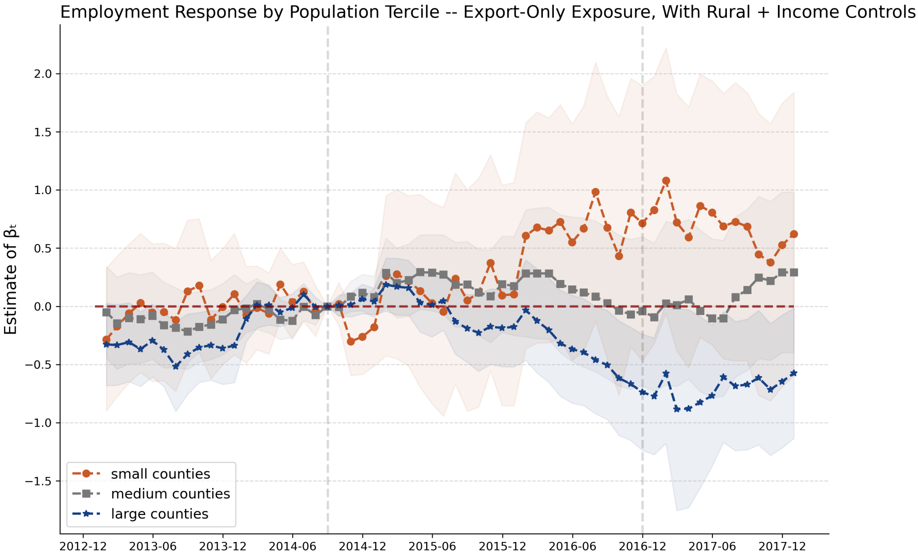

Where We Are...

- A 25% USD appreciation (2014-2016) is a missing episode in a literature dominated by the China Shock and the US-China trade war.
- Employment and political outcomes move as that story predicts, in more-exposed counties.
- Trade-flow evidence confirms the story: exports respond, imports don't – the export channel is doing the work.
- Most identification checks hold up; two remain genuinely open (previous slides).
- **Where we'd value your input:** framing the large-county result and dose-response heterogeneity.

References I

Autor, D., A. Beck, D. Dorn, and G. H. Hanson (2024): "Help for the Heartland? The Employment and Electoral Effects of the Trump Tariffs in the United States," Tech. rep., National Bureau of Economic Research.

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.

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

Orchard, J., V. A. Ramey, and J. F. Wieland (2023): "Micro MPCs and macro counterfactuals: the case of the 2008 rebates," Tech. rep., National Bureau of Economic Research.

Waugh, M. E. (2019): "The Consumption Response To Trade Shocks: Evidence From The US-China Trade War," Tech. rep., National Bureau of Economic Research.

Yagan, D. (2019): "Employment hysteresis from the great recession," *Journal of Political Economy*, 127, 2505–2558.

Backup — Items 1-3, Identification and Inference

1. **Shift-share exogeneity.** Is $\texttt{FX}(c, \text{Dec } 2016)$ – a share $\times$ shift construction – as-good-as-randomly assigned, on the *shares* (Goldsmith-Pinkham, Sorkin & Swift 2020) or the *shocks* (Borusyak, Hull & Jaravel 2022) side? Both sufficient conditions pass; no disqualifying problem, though the shocks-side test is inherently low-powered ($\sim$10 effective sectors out of 27).
2. **Heterogeneous dose-response.** Continuous-dose event studies (Callaway, Goodman-Bacon & Sant'Anna 2024) can carry *negative* comparison weights when dose-response is heterogeneous. Prerequisite now checked: the unconditional dose-response is approximately linear at peak exposure with a clean flat pre-period placebo – doesn't resolve the weighting question, but removes the most likely reason it would bite.
3. **Inference.** `HC0` robust SEs vs. clustering by state/commuting-zone: clustering widens SEs 4–34%, but the headline result stays significant at 90% throughout.

Backup — Items 4-6, Specification and Scope

4. **Exposure construction.** The baseline sums export + import-competition exposure into one scalar dose. Turns out this is empirically necessary, not just a simplifying choice: the raw channels are too collinear ($r = 0.84$–0.87) to separately identify. The imported-input channel – which works in the *opposite* direction – remains omitted from the regression entirely; its true sign is still unresolved. This is a statistical limitation of the *constructed exposure indices*; it is a different question from item 7's finding that import *volumes* themselves barely move with FX exposure – the two are complementary, not in tension: one says we can't statistically separate the channels, the other says there may be little economic signal in the import channel to separate out in the first place.
5. **Omitted contemporaneous shocks.** Oil- and FX-exposed counties are largely different geographies, but the FX effect still shrinks $\sim$18% in magnitude once all three oil channels (direct/upstream/downstream) are controlled for. China-shock exposure (Autor-Dorn-Hanson) is only modestly correlated with our index and *strengthens* the FX effect slightly when added, not a confound. Global demand now has a full county-level compositional-demand control (slide 20, built from destination-mix tilt in absorption and GFCF growth) that the FX effect survives. The Fed funds rate/2yr yield correlates with the dollar index at the trend level but far more weakly month-to-month – mildly reassuring, but still only an aggregate-level pass; a real county-level

Backup — Item 5, With-Oil-Controls Specification

Slide 18 extends the baseline no-controls spec (Research Design slide) by adding the three oil exposure channels directly, as a robustness check rather than a change to the headline spec:

$$
\begin{aligned}
\log Y_{c,t} - \log Y_{c,\text{2014-8}} &= \alpha_t + \beta_t \Delta\mathtt{FX}_{c,\text{2016-12}} \\
&\quad + \ \gamma_t^{dir}\mathtt{Oil}_c^{dir} + \gamma_t^{up}\mathtt{Oil}_c^{up} + \gamma_t^{down}\mathtt{Oil}_c^{down} + \epsilon_{c,t},
\end{aligned}
$$

- Same β_t loop as the baseline: one population-weighted WLS fit per month t, `HC0` SEs.
- $\mathtt{Oil}_c^{dir}$, $\mathtt{Oil}_c^{up}$, $\mathtt{Oil}_c^{down}$ are the three channels on the next slide, each frozen at their Dec-2016 value – same single-dose logic as $\Delta\mathtt{FX}$.
- Nested one channel at a time (direct $\to$ +upstream $\to$ +downstream) – that's the three added lines on slide 18, versus the black no-oil baseline.
- No rural/income here, matching the baseline's primary (no-controls) spec – adding rural+income shrinks β_t further but carries the same reproducibility caveat flagged for any rural/income spec.

Backup — Item 5, Constructing the Three Oil Exposure Channels

All three freeze 2013 employment/input-output shares (same reverse-causality logic as FX) and are scaled by P_t^{Oil}, the WTI spot price.

Direct (employment within oil & gas itself):

$$\texttt{Oil}_{c,t}^{dir} = \sum_{j \in \Omega^{Oil}} \left(\frac{L_{c,j,2013}}{L_{c,J,2013}} \right) \times P_t^{Oil}$$

Upstream (industries supplying inputs *to* oil & gas):

$$\texttt{Oil}_{c,t}^{up} = \sum_{j \in J \setminus \Omega^{Oil}} \left(\frac{L_{c,j,2013}}{L_{c,J,2013}} \right) \times \sum_{i \in \Omega^{Oil}} \left(\frac{use_{j,i,2013}}{Q_{j,2013}} \right) \times P_t^{Oil}$$

Downstream (industries using oil & gas *as* an input):

$$\texttt{Oil}_{c,t}^{down} = \sum_{j \in J \setminus \Omega^{Oil}} \left(\frac{L_{c,j,2013}}{L_{c,J,2013}} \right) \times \sum_{i \in \Omega^{Oil}} \left(\frac{use_{i,j}}{Int_j + Comp_j} \right) \times P_t^{Oil}$$

Ω^{Oil} = NAICS 211/213/324 (oil & gas extraction, mining support, petroleum & coal products); $use_{j,i}/Q_j$ from the BEA input-output "use" table – same tables and same functional form as the imported-input FX channel (Step #2, omitted from today's talk). Oil & gas itself excluded from up/downstream due to collinearity (its own intra-industry input shares are very high). Full derivation and citations: paper, §"Robustness – Exposure to Oil Price Fluctuations."

Backup — Item 5, With-China-Shock-Control Specification

Slide 19 nests the same β_t loop against historical China-shock exposure, on its own and then with rural + income added:

$$\begin{aligned}\log Y_{c,t} - \log Y_{c,2014\text{-}8} &= \alpha_t + \beta_t \Delta\mathtt{FX}_{c,2016\text{-}12} + \theta_t\, \mathtt{China}_c \\ &\quad + \lambda_t^{rural}\mathtt{Rural}_c + \lambda_t^{inc} \log \mathtt{Income}_c + \epsilon_{c,t},\end{aligned}$$

- Same β_t loop as before: one population-weighted WLS fit per month t, `HC0` SEs.
- $\mathtt{China}_c$ is each county's cumulative 1990-2007 China-shock import-competition exposure (Autor-Dorn-Hanson 2013), broadcast from the commuting zone to every county in it – a *fixed*, pre-determined regressor, not a contemporaneous shock like oil.
- θ_t (on $\mathtt{China}_c$) is not the object of interest; β_t is, and it gets slightly *larger* in magnitude once $\mathtt{China}_c$ is added, not smaller – the opposite direction from oil or rural.
- Because $\mathtt{China}_c$ is fixed 15+ years before the study window starts, the relevant validity check is a pre-trend: does it predict 2013 (pre-period) growth holding $\Delta\mathtt{FX}$ fixed? It doesn't – clean null throughout.

Backup — Item 5, With-Demand-Control Specification

Slide 20 extends the same nested logic to global demand: baseline plus rural, income, and direct oil, then one compositional demand-exposure term added at a time:

$$\begin{aligned}\log Y_{c,t} - \log Y_{c,\text{2014-8}} &= \alpha_t + \beta_t \Delta\texttt{FX}_{c,\text{2016-12}} + \gamma_t\,\texttt{Oil}_c^{dir} \\ &+\ \delta_t\, D_c + \lambda_t^{rural}\texttt{Rural}_c + \lambda_t^{inc}\log\texttt{Income}_c + \epsilon_{c,t},\end{aligned}$$

- Same β_t loop as before: one population-weighted WLS fit per month t, `HC0` SEs.
- D_c: export-share-weighted destination-demand growth, demeaned to isolate *which* partners a county leans toward from how trade-exposed it is overall. Two variants: absorption-based, GFCF-only.
- δ_t (on D_c) isn't the object of interest – small, sign-unstable to outliers. β_t is, and it barely moves when D_c is added.
- Direct oil included throughout (matches slide 18), so this nests that spec rather than replacing it.

Change in Import-Competition FX Exposure by County

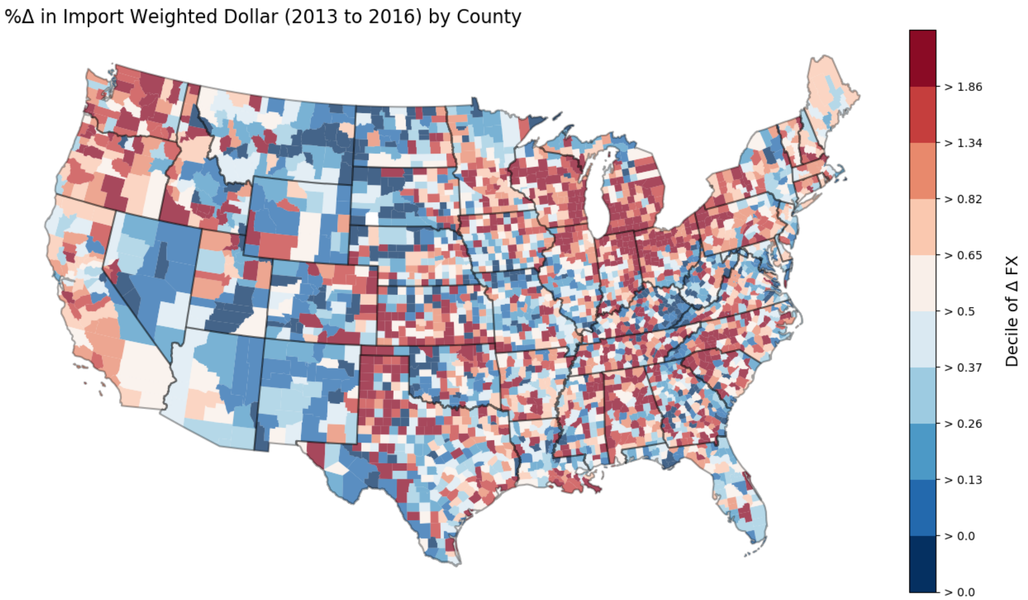

Trade-Flow Response — Exports Fall, Imports Don't Move

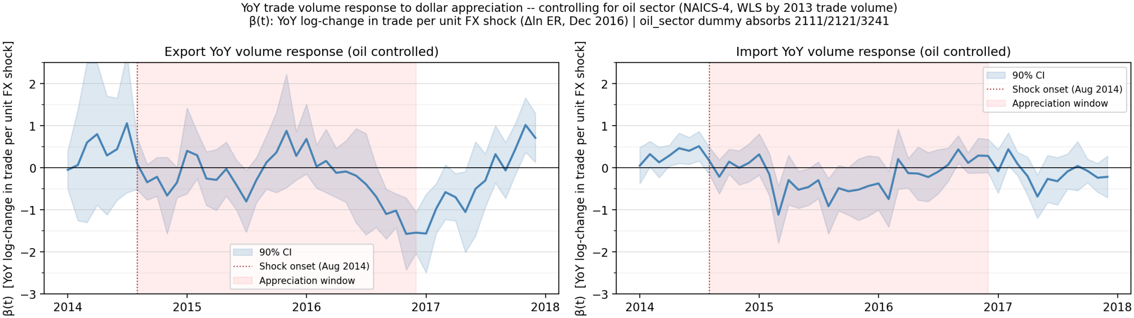

7. Uses actual NAICS-4 trade-*quantity* data, independent of the constructed exposure indices behind item 4's collinearity finding. Export volumes fall with FX exposure; import volumes barely move – substantive, not just statistical, support for treating the export channel as the one doing the work.