

Recovering Credible Trade Elasticities from Incredible Trade Reforms—with AI!

Alessandria, Khan, Khederlarian, Ruhl, and Steinberg

T2M — June 2, 2026

Intro

- ▶ Unifying question: How much does trade change when policy changes?
- ▶ Trade elasticity: most important concept in int'l econ?
- ▶ Reduced form estimates: vary widely across time horizons & contexts
- ▶ Standard interpretation: response to unanticipated, once-and-for-all (“**canonical**”) reform

Intro

- ▶ Unifying question: How much does trade change when policy changes?
- ▶ Trade elasticity: most important concept in int'l econ?
- ▶ Reduced form estimates: vary widely across time horizons & contexts
- ▶ Standard interpretation: response to unanticipated, once-and-for-all (“**canonical**”) reform
- ▶ Main paper: canonical reforms don't exist in the data!
 - ▶ Empirical: compare “more-canonical” vs. “less-canonical” reforms
 - ▶ Quantitative: recover canonical elasticity by feeding data through structural model

Intro

- ▶ Unifying question: How much does trade change when policy changes?
- ▶ Trade elasticity: most important concept in int'l econ?
- ▶ Reduced form estimates: vary widely across time horizons & contexts
- ▶ Standard interpretation: response to unanticipated, once-and-for-all (“**canonical**”) reform
- ▶ Main paper: canonical reforms don't exist in the data!
 - ▶ Empirical: compare “more-canonical” vs. “less-canonical” reforms
 - ▶ Quantitative: recover canonical elasticity by feeding data through structural model
- ▶ This paper: can we recover canonical elasticity without a priori knowledge of underlying economic structure?
 - ▶ Or, can sufficiently complex neural networks learn that structure from observational data alone?

Main idea

Non-canonical reforms

Anticipation (PTA phaseouts)

Uncertainty (Brexit, Lib-Day)

+

Dynamic trade theory

Forward-looking export participation decisions due to front-loaded costs, back-loaded returns

Δ trade depends on observed policy + **expectations**

- ⇒ Same policy change leads to different trade responses under different expectations
- ⇒ Trade responds to changes in expectations even when policy doesn't change
- ⇒ **Trade responds more to large expected future tariff changes than small ones**

Main questions

1. How canonical is the typical trade reform? Which reforms are most/least canonical?
2. Do trade elasticity estimates differ across reform types?
 - ▶ How big are the differences?
 - ▶ Are the differences related to “canonicalness”?
3. **Can we recover canonical elasticities from non-canonical policy changes?**

Overview

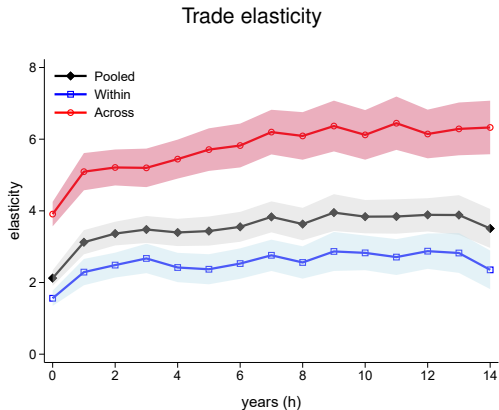
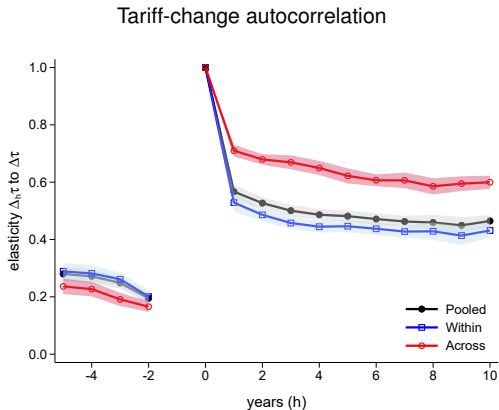
- ▶ Simulated data from our model provide a controlled testbed
 - ▶ Know the true structural (canonical) trade elasticity path
 - ▶ Can simulate data from non-canonical reforms with known tariff process
 - ▶ Ask: can an AI model trained *only* on non-canonical data recover the canonical elasticity?
- ▶ Experiment design:
 - ▶ Generate simulated model data using Markov tariff processes ranging from fairly transitory to highly persistent
 - ▶ Train LSTM (Long Short-Term Memory) network on simulated data, use trained network to predict outcome of canonical reform
 - ▶ Explore how results vary with persistence of reforms that network gets to see during training

Empirical motivation

- ▶ Sample: U.S. imports from 1974–2017
 - ▶ Captures transition from higher tariffs in 70s & 80s to low tariffs today
 - ▶ Covers major reforms: China's MFN grant, NAFTA, GATT rounds, etc.
- ▶ Compare responses to tariff changes that occur within statutory policy regimes (transitory) to changes triggered by regime shifts (persistent)
- ▶ Case studies of most persistent reforms in US history: MFN grants to China and Vietnam
- ▶ Our perspective:
 - ▶ Variation in trade responses can be explained by variation in expectations
 - ▶ Same underlying economic structure at work across empirical contexts

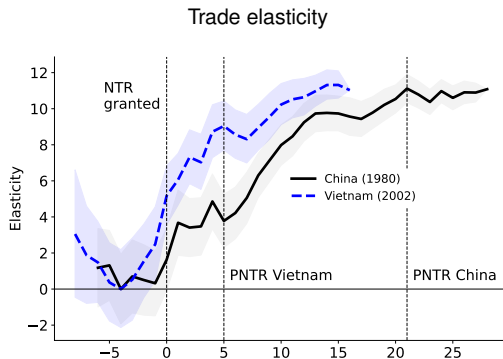

Across-regime tariff changes: more persistent, higher trade elasticities

- ▶ Determine each observation's statutory tariff regime (e.g. MFN, NNTR, PTA)
- ▶ Split sample into “across-regime” and “within-regime” observations
- ▶ Measure tariff persistence and trade responses to within- versus across-regime reforms



CHN& VNM NTR grants: even more persistent, even higher elasticities

- ▶ Same observed policy trajectory: embargo → NNTR → MFN
- ▶ Ex post, “most canonical” reforms in US trade history. Ex ante, lots of uncertainty.
- ▶ Clearly-defined policy risk, no phase-in, embargo start allow for clean quantitative analysis





Model in one slide

- ▶ Partial equilibrium version of Alessandria, Choi & Ruhl 2021 (ACR 2021)
 - ▶ Slow adjustment due to exporter life-cycle, large gap between SR and LR response
 - ▶ Expectations about future trade policy, not current policy, drive export participation
- ▶ Firms
 - ▶ Heterogeneous in productivity (z), iceberg trade cost (ξ)
 - ▶ Pay sunk cost to export next period, smaller fixed cost to continue
 - ▶ New exporters start with high iceberg trade cost
 - ▶ Longer tenure as exporter \Rightarrow greater chance of low iceberg cost
- ▶ Trade policy shocks
 - ▶ Current tariffs (τ)
 - ▶ Expectations about future tariffs ($\mathbb{E}\tau'$)
- ▶ Decision problem

$$V(z, \xi, \tau) = \pi_{gt}(z, \xi, \tau) + \max \left\{ \underbrace{-f(\xi) + \frac{\delta(z)}{1+r} \mathbb{E}_{z', \xi', \tau'} V(z', \xi', \tau')}_{\text{export}}, \underbrace{\frac{\delta(z)}{1+r} \mathbb{E}_{z', \xi', \tau'} V(z', \infty, \tau')}_{\text{don't export}} \right\}$$

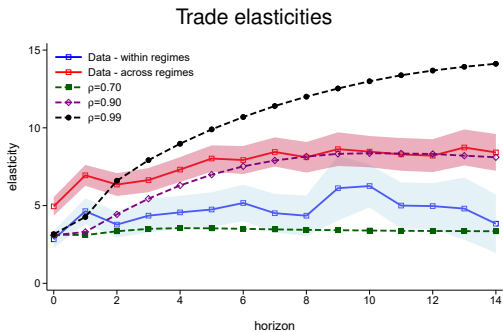

Model simulations

- ▶ Simulate model under two-state Markov process ($\tau_L = 0\%$, $\tau_H = 30\%$) with switching probability $1 - \omega$ for $\omega \in (0.5, 1.0)$
- ▶ 1,000 goods in each simulation, each receive idiosyncratic tariff changes
- ▶ Estimate trade elasticities w/local projections method (Boehm et al., 2023)



Eyeballing model versus data


- ▶ Within-regime transition $\approx \rho = 0.70$, across $\approx \rho = 0.90$.




LSTM architecture & training

- ▶ Observation (a.k.a. “anchor”) tariff-change event for good g at time t
 - ▶ “Past” window: log tariffs, Δ log tariffs, log exports, Δ log exports, regime indicator, consecutive periods in current state
 - ▶ “Future”: path of tariff changes $\{\Delta \log \tau_{t+k}\}_{k=1}^h$
 - ▶ Target: cumulative log change in exports $\Delta_h \log X_t = \log X_{t+h} - \log X_{t-1}$
- ▶ Key: network never observes ω directly; must infer from joint dynamics of tariffs and trade
- ▶ Five training runs, successively expanding range of ω :
 1. $\omega \leq 0.85$ (similar to within-regime MFN changes)
 2. Add $\omega = 0.90$ (similar to average across-regime switch)
 3. Add $\omega = 0.95$ (similar to most persistent observed reforms)
 4. Add $\omega = 0.98$
 5. Add $\omega = 0.99$ (essentially do not exist in historical data)
- ▶ “Empirical budget”: fixed fraction of anchors per ω file, so training set reflects dominance of transitory changes even when persistent changes are included


Results




Results




Results



Results



Results



Summary

- ▶ LSTM *does* extrapolate meaningfully beyond its training support
 - ▶ E.g., $\omega \leq 0.85$ model never sees LR elasticities > 4.2 but predicts ~ 6.1
 - ▶ $\omega \leq 0.90$ model never sees LR elasticities > 5.8 but predicts > 10
- ▶ But accuracy depends critically on proximity of training data to the target policy
 - ▶ Within-regime tariff changes ($\omega \leq 0.85$) provide very poor basis for extrapolation to canonical reform
 - ▶ Need extremely persistent reforms ($\omega \geq 0.95$) in training data to get close—but these are rare/nonexistent empirically
- ▶ Next step: train network on actual data!
 - ▶ Use transfer-learning proposed by Chen et al. (NBER WP 34713)
 - ▶ Pre-train network on model simulations, then use actual data to fine-tune