Credibility is Trade Policy*

Joseph B. Steinberg
University of Toronto and NBER

May 11, 2026

Abstract

Export participation is a forward-looking decision that involves large up-front investments with back-loaded returns. This implies that trade responds to expected future policy as well as current policy. The same observed policy change can trigger a wide range of responses depending on how transitory or persistent that change is expected to be, and trade responds to changes in expectations even when there are no observed changes in policy at all. This article combines empirical evidence with structural modeling to show that these lessons have important implications for how researchers interpret trade data and how policymakers should use trade policy to accomplish their objectives.

JEL Classifications: F12, F13, F14
Keywords: trade elasticity, trade dynamics, tariff risk, trade policy uncertainty, China shock, trade war

*This article is based on my State-of-the-Art Lecture at the 2025 Canadian Economics Association annual conference. It draws extensively on joint work with George Alessandria, Shafaat Y. Khan, Armen Khederlarian, and Kim J. Ruhl. I thank conference participants and the editors of the *Canadian Journal of Economics* for helpful comments. Any errors or interpretations are mine alone. I used OpenAI's ChatGPT 5.4 and Anthropic's Opus 4.6 to copy-editing and assistance with preparing the tables and figures, and

1 Introduction

On December 11, 2001, China formally acceded to the World Trade Organization (WTO). Seven years later, U.S. imports from China had nearly quadrupled, while the tariffs on these imports had barely changed. How could trade have grown so much without any changes in applied tariffs? The conventional story is that China's WTO accession eliminated a risk that had been depressing trade for decades: the possibility that the United States would revoke China's Normal Trade Relations (NTR) status and revert to the punitive Non-Normal Trade Relations (NNTR) tariffs set by the 1930 Smoot-Hawley Act. Pierce and Schott (2016) show that after 2001, U.S. imports of the goods most exposed to the risk of losing NTR status—those with the largest gaps between NNTR and NTR tariffs—grew faster than imports of less-exposed goods. Numerous studies have interpreted this pattern as evidence that permanent NTR access eliminated policy uncertainty and triggered the "China shock" that reshaped U.S. labor markets, investment, and asset prices.[1]

This interpretation of the empirical evidence is partly right, in that resolution of trade-policy uncertainty played an important role in the growth of U.S. trade with China, but as my coauthors and I show in Alessandria et al. (2025b), it misses several important details about the economic mechanisms at play. In two subsequent papers (Alessandria et al. 2025c,a), we show that these details matter systematically for how we interpret the responses of trade to tariff changes in general. This article reviews our research on the joint dynamics of trade policy and trade flows and discusses the implications of our findings for trade economists and policymakers.

In Alessandria et al. (2025b), we document how the elasticity of U.S. imports from China to the NNTR gap evolved over time, starting when the U.S. first lifted its embargo on China in 1971 and ending just before the Great Recession in 2008, and use the estimated series to discipline a dynamic trade model with time-varying policy expectations. We find that the probability of losing NTR status was about sixty percent shortly after the United States first granted NTR status to China in 1980, fell to less than ten percent by the late 1990s, and changed little thereafter. The calibrated model lets us decompose the observed growth in imports of high-gap goods into an "expectational" component driven by the changes in the regime-switching probabilities that we recover, and a "technological" component that captures how trade would adjust if there was no trade-policy uncertainty at all (i.e., if switching regimes was impossible). More than half of the

[1] The term "China shock," which originates with Autor et al. (2013), refers to the effects of rapid growth in Chinese exports on the economies of the United States and other developed countries at the beginning of the 21st century. See also Autor et al. (2014) and Acemoglu et al. (2016) on labor-market adjustment, Handley and Limão (2017) on trade-policy uncertainty as the operative channel, Pierce and Schott (2018) on investment, and Bianconi et al. (2021) on asset prices.

observed growth in imports of high-gap goods after 2001 is expectational, but virtually all of this component reflects delayed adjustment to rising credibility of China's NTR status during the late 1980s and early 1990s, which occurred long before China joined the WTO; the contribution of WTO accession itself is virtually nil in our model decomposition. The fact that trade adjusts gradually to changes in expectations as well as changes in tariffs is one of the most important insights of our research agenda.

In Alessandria et al. (2025c), we apply this framework to study the 2018 U.S.-China trade war. The trade war introduced a third tariff schedule alongside the pre-existing NTR and NNTR schedules, which we use to construct a second exposure measure—the trade-war gap, defined as the difference between the trade-war tariff and the NTR tariff. Because the trade-war gap and the original NNTR gap are approximately orthogonal, we can identify beliefs about the probability of a trade-war reversal separately from beliefs about the probability of a return to the Smoot-Hawley schedule. We find that during Trump's first term, the trade-war tariffs were widely perceived as likely to be short-lived, but by the time Biden left office, they were perceived as likely to persist. Our findings imply that the expected future U.S. tariff on China rose more under President Biden, who largely maintained the existing trade-war tariff schedule, than under President Trump, who was responsible for initiating the trade war.

In Alessandria et al. (2025a), we ask whether these case studies are unique, or whether trade-policy expectations play a systematically important role in determining the effects of trade reforms in general. We first document that nearly all of the tariff variation in postwar U.S. import data consists of transitory changes within the NTR regime; large, persistent reforms that switch countries from one statutory regime to another, like China's 1980 NTR grant, are rare. We show that "within-regime" tariff changes are associated with much smaller trade elasticities, especially in the long run, than "across-regime" changes, and argue that this difference can be explained by our dynamic trade model with forward-looking firms: transitory tariff changes induce smaller changes in export participation than persistent changes. We then use the transition probabilities recovered from our case study of China's NTR access, as well as a similar case study of Vietnam, to simulate a canonical reform—an unanticipated, once-and-for-all, perfectly credible tariff change—and the long-run trade elasticity is several times larger than conventional reduced-form estimates.

This article discusses three broader lessons from our research agenda. First, a positive lesson for empirical trade economists: there is no single trade elasticity estimate that applies in all settings. The elasticity that governs the response of trade to a tariff change depends both on the horizon at which this response is measured, and on the context surrounding that change. Elasticities measured from typical transitory tariff variation should not be used to predict the effects of major trade liberalizations, which tend to be much more persistent and generate much larger trade responses. Second, a normative lesson for welfare and optimal-policy analysis: conventional reduced-form estimates of the trade elasticity should not be used to measure the welfare gains from major trade liberalizations or to draw inferences about the optimal tariff implied by terms-of-trade reasoning. Our canonical trade elasticity estimates, which are several times larger than many estimates in the literature, imply smaller welfare gains and smaller optimal tariffs. Third, a policy-design lesson: credibility is itself a policy instrument with its own unique tradeoff. Tariffs expected to be transitory are likely to raise substantial tax revenue in the short run but produce little substitution from foreign goods toward domestic ones. Governments that seek to use trade policy to revitalize domestic industry must find ways to commit to durable policies, but must also accept lower tariff revenues.

These lessons are particularly salient at the time of this writing, as policymakers prepare for an imminent joint review of Canada's free-trade agreement with the United States and Mexico, which has a built-in sunset mechanism that terminates the agreement unless all three countries agree to extend it.[2] Viewed through the framework of our research agenda, this mechanism has a side effect that was presumably unintended: it reduces the agreement's perceived durability, which in turn reduces the extent to which the agreement fosters economic integration among the three member countries. A preferential trade agreement whose continued existence must be affirmatively renewed every few years is not the same economic object as one whose continued existence is taken for granted, even if the tariffs are identical. In an environment of rising trade-policy uncertainty, the economic value of North American market access depends more than ever on how durable that access is perceived to be.

2 Why expectations matter for trade

A firm deciding whether to export to a foreign market is making an investment decision. Entering the market requires paying costs that are at least partly sunk: identifying buyers, navigating customs procedures, establishing distribution relationships, and, for many products, adapting the good itself to foreign standards. These costs are not recovered if the firm exits, and the returns to paying them are earned gradually, as the firm builds a customer base and learns how to serve the

[2] Canada refers to this agreement as CUSMA, while the United States and Mexico call it USMCA and T-MEC, respectively.

market. A substantial literature going back to Das et al. (2007) and documented most sharply by Ruhl and Willis (2017) shows that new exporters start small, grow slowly over many years, and account for a modest share of total exports until well after entry. Because the investment is sunk and the returns are backloaded, a firm contemplating entry cares not only about the tariff it will face this year, but about the entire path of tariffs it will face in the future—and about the uncertainty surrounding that path.

2.1 Formalizing the mechanism

The workhorse modeling framework used in all three of the studies described in this article is a version of the dynamic exporter model of Alessandria et al. (2021b) where tariffs follow a time-varying stochastic process. As in Chaney (2008), there is a fixed continuum of firms that produce differentiated varieties and choose in each period whether to sell those varieties abroad. Foreign sales depend on the tariff, τ, which is an aggregate state common to all firms, and on a firm's own productivity, z, and its export capacity, ξ, which is modeled as an "iceberg" transportation cost.[3] The profits that a firm earns from exporting can be written simply as $\pi(z, \xi, \tau)$. A firm that wants to begin exporting must pay a sunk entry cost, f_0, whereas a firm that is already an exporter must pay a smaller continuation cost, f_1, in order to continue serving the foreign market. New exporters begin with a high iceberg cost, ξ_H, and in each subsequent period face a fixed probability λ of transitioning to a low iceberg cost, $\xi_L < \xi_H$; incumbent exporters face a symmetric probability of switching back and forth between high and low iceberg costs. Tariffs follow a Markov chain with transition matrix $P_t(\tau'|\tau)$.[4] Crucially, the transition matrix P_t can change over time (hence the t subscript) as institutions, geopolitics, and expectations evolve. Firms know the current regime, τ, and the probability of switching regimes dictated by P_t, but not the regime they will face in the next period, τ'.

A firm's problem is to compare the discounted expected value of exporting to the value of waiting (if the firm is not yet an exporter) or exiting (if it is one already). Let a firm's idiosyncratic state be $x = (z, \xi, m)$, where z is productivity, $\xi \in \{\xi_H, \xi_L\}$ is the iceberg trade cost, and $m \in \{0, 1\}$ indicates whether the firm is already an exporter. Conditional on (x, τ), the firm chooses whether or not to export in the next period, which requires paying the relevant fixed cost today. This problem boils down to a binary decision $d \in \{0, 1\}$. If $d = 0$, the firm begins the next period as a

[3] In order to deliver one unit of goods to the foreign market, the firm must produce and ship an additional ξ units which "melt" before arriving.

[4] In Alessandria et al. (2025b) and Alessandria et al. (2025a), there are two possible tariff regimes: NNTR and NTR. In Alessandria et al. (2025c), there are three: NNTR, NTR, and the trade-war tariffs that were introduced in 2018.

non-exporter with $m' = 0$. If $m = 0$ and $d = 1$, the firm pays the sunk entry cost f_0 and begins next period with $m' = 1$ and $\xi' = \xi_H$. If $m = 1$ and $d = 1$, the firm pays the continuation cost f_1 and begins the next period with $m' = 1$ and has a chance λ of switching iceberg trade costs. The Bellman equation that represents this problem is

$$V_t(z, \xi, m, \tau) = \max_{d \in \{0,1\}} \Bigg\{ m\pi(z, \xi, \tau) - d[(1-m)f_0 + mf_1] \\ + \beta \sum_{\tau'} P_t(\tau' \mid \tau) \int \sum_{\xi' \in \{\xi_H, \xi_L\}} H(\xi' \mid \xi, m)\, V_{t+1}(z', \xi', d, \tau')\, G(dz' \mid z) \Bigg\}, \tag{1}$$

where $G(dz' \mid z)$ is the process for productivity and $H(\xi' \mid \xi, m) = m\left[(1-\lambda)\mathbb{1}_{\{\xi'=\xi\}} + \lambda \mathbb{1}_{\{\xi' \neq \xi\}}\right] + (1-m)\mathbb{1}_{\{\xi'=\xi_H\}}$ is the process for iceberg trade costs.

The optimal policy $d_t(x, \tau)$ is characterized by an entry threshold, $\underline{z}_t(\tau)$, which depends on τ and P_t (for which t is a sufficient statistic), and an exit threshold, $\bar{z}_t(\xi, \tau)$, that also depends on ξ. Non-exporters with productivity above the entry threshold become new exporters, and incumbents with productivity below the exit threshold become non-exporters. The sunk nature of the entry cost implies that $\underline{z}_t(\tau) > \bar{z}_t(\xi, \tau)$ for $\xi \in \{\xi_H, \xi_L\}$, which generates the classic hysteresis pattern described by Baldwin and Krugman (1989): some firms that are not sufficiently productive to initially enter are willing to remain as incumbents.[5] The most important property of these thresholds for our research, however, is that they depend on the current tariff τ only to the extent that it contains information about future tariffs. The thresholds are both increasing in the expected present value of future tariffs: when future tariffs are expected to be high, fewer non-exporters enter and more incumbents exit. Thus, holding fixed the current tariff τ, changes in the policy transition matrix P_t today will affect the volume of trade tomorrow by shifting the entry and exit thresholds.

2.2 Illustrating the mechanism

To illustrate how trade-policy expectations influence the dynamics of trade, Alessandria et al. (2025a) use the model to conduct a series of simple numerical experiments. Figure 1 reproduces the first set of experiments, which is intended to illustrate the role of tariff persistence. Tariffs follow a Markov chain with two values, $\tau_H = 30\%$ and $\tau_L = 0\%$, and switch back and forth between these values with probability $\omega \in [0, 0.5]$. When $\omega = 0$, tariff changes are zero-probability events ("MIT

[5] See Alessandria et al. (2021a) for a detailed characterization of sunk-cost exporting models.

shocks"). When $\omega = 0.5$, tariffs are i.i.d. and the current tariff contains no information about future tariffs. In between, tariffs are stochastic but exhibit some degree of persistence. The figure shows the dynamics of trade around an isolated tariff cut, where tariffs are high for many periods and then fall and remain low for many periods, for several values of ω.

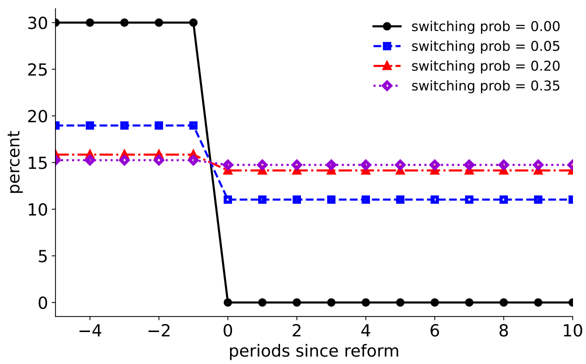

(a) Present value of tariffs

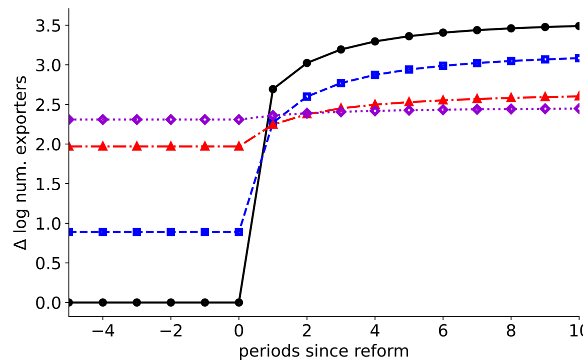

(b) Number of exporters

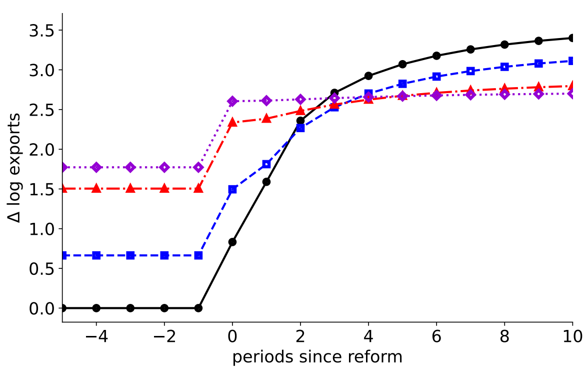

(c) Exports

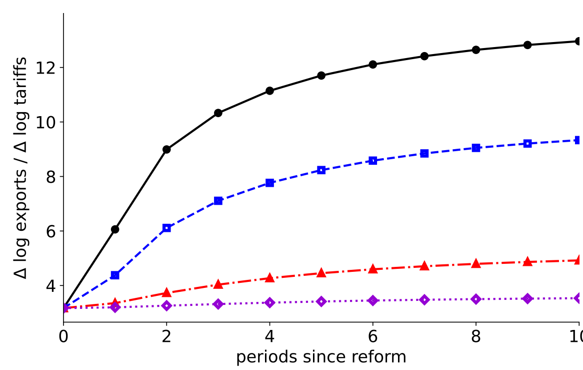

(d) Trade elasticity

Figure 1: Model dynamics following tariff cut for Markov switching probabilities ranging from zero to 35%. (a) Discounted present value of future tariffs. (b) Number of exporters. (c) Exports. (d) Trade elasticity. Data in (a) and (b) normalized relative to initial values from zero-probability tariff cut experiment. Reproduced from Alessandria et al. (2025a).

The first panel of the figure shows how the present value of future tariffs changes. In the MIT-shock version of the experiment, the present value changes one-for-one with the observed tariff level. As the switching probability increases, the present value changes less: it is lower than the observed tariff level before the tariff cut, as firms know there is a chance the cut could happen, and higher afterwards, as they know there is a chance it could be reversed. Even a small switching probability of 5% materially reduces the effect of the tariff cut on the present value of future tariffs, and a higher switching probability of 35% makes the present value barely change at all. The second panel shows how export participation evolves, the third shows the dynamics of exports, and the fourth shows the trade elasticity (the log change in exports relative to the log change in tariffs). In the MIT-shock version of the experiment, export participation and trade volumes are very low in the high-tariff state and grow dramatically after tariffs fall, ultimately generating a trade elasticity of more than 14 in the long run (much higher than typical estimates in the empirical literature). As the switching probability rises, there are more firms exporting and a greater volume of trade in the high-tariff state, and smaller adjustments in those variables after the tariff cut. When the switching probability is just 5%, the long-run trade elasticity falls to less than ten (a difference of almost a third). When the switching probability rises to 35%, the tariff cut induces essentially no entry into exporting at all, and the long-run trade elasticity is almost exactly equal to the short-run elasticity.

Figure 2 reproduces a second set of experiments from the appendix of Alessandria et al. (2025a), which is intended to illustrate the role of anticipation. As in the first set of experiments, tariffs fall from $\tau_H = 30\%$ to $\tau_L = 0\%$, but here, tariffs do not follow a stochastic process. Instead, firms may learn about tariff changes in advance, and these changes may be phased in gradually. In the unanticipated version, the tariff cut is announced in the period it occurs, whereas in the anticipated version, the cut is announced ten periods ahead of time. In the phased-in version, the cut occurs gradually over ten periods (this schedule is made known in the period the tariff change is announced). The figure shows all four combinations: $\{\text{unanticipated}, \text{anticipated}\} \times \{\text{immediate}, \text{phased in}\}$.

As before, the first panel shows how the present value of future tariffs moves in the different versions of the experiment. The unanticipated and immediate version corresponds to the zero-probability MIT shock version in the previous experiment, where the present value moves one-for-one with the observed tariffs. In the anticipated versions, the present value jumps down immediately when the reform is announced, and then falls linearly as the reform date nears. In the phased-in versions, the present value continues to fall until the tariff cut is fully implemented, ten years after the reform starts. The other panels show export participation, trade volumes, and the trade elasticity. When the reform is anticipated, firms begin to enter as soon as they learn about the tariff cut, especially when they know it will be implemented all at once rather than phased in. The elasticity of trade in these experiments is measured relative to the period before the reform announcement, so all four versions have the same long-run elasticity, but the short-run elasticity is higher when the reform is anticipated.[6] This is especially true in the anticipated, phased-in reform,

[6] If one measures relative to the period before the reform actually occurs, the anticipated versions have lower long-run elasticities.

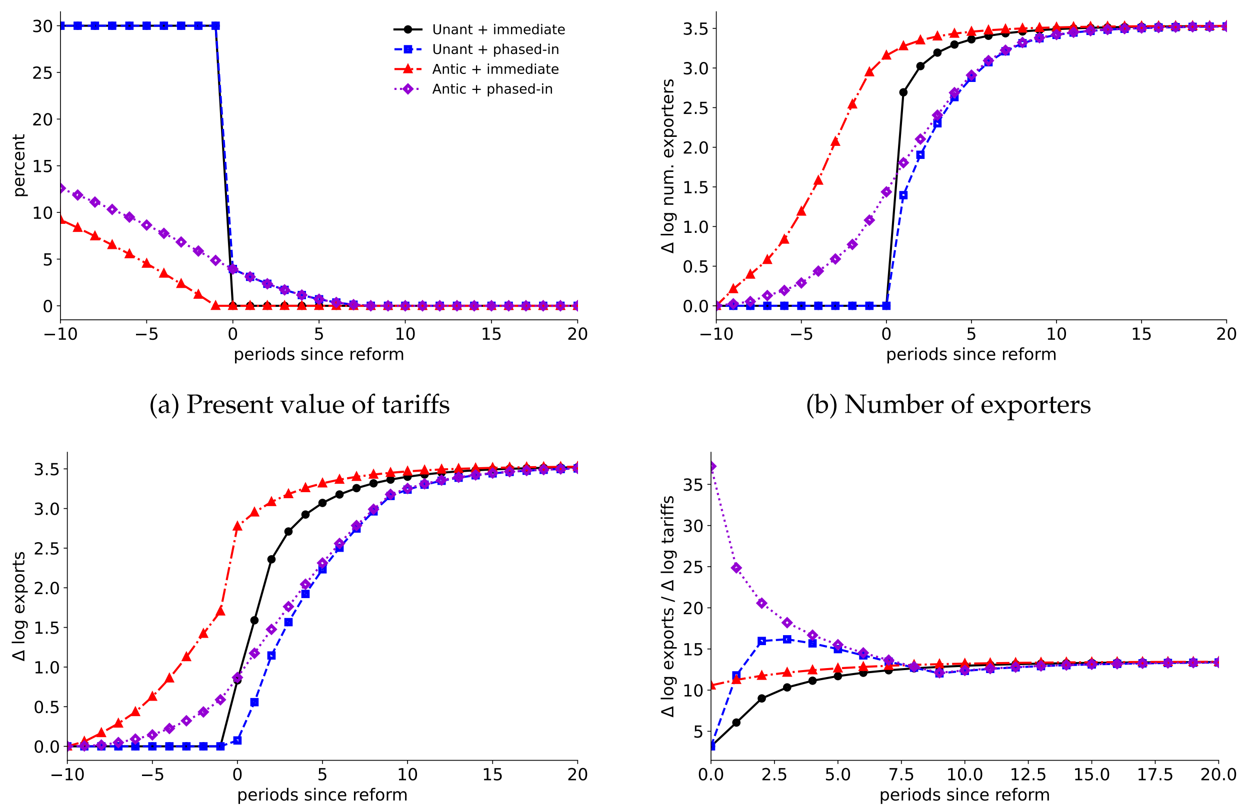

(a) Present value of tariffs

(b) Number of exporters

(c) Exports

(d) Trade elasticity

Figure 2: Model dynamics following tariff cut under alternative timing assumptions. (a) Discounted present value of future tariffs. (b) Number of exporters. (c) Exports. (d) Trade elasticity measured relative to the period before the reform. Reproduced from the appendix of Alessandria et al. (2025a).

where the growth in trade that begins after the reform is announced is very large relative to the small change in tariffs that has occurred one period into the reform.

Alessandria et al. (2025a) provide a clean way to see that these mechanisms have empirical salience by comparing two countries that followed the exact same path in terms of applied U.S. trade policy but followed quantitatively different paths in terms of trade dynamics: China and Vietnam. Both were subject to a complete U.S. trade embargo. Both were granted conditional NTR status after the embargo ended—China in 1980, Vietnam in 2002—and both subsequently received permanent NTR access upon acceding to the WTO, China in 2001 and Vietnam in 2007. The sequence of institutional events and the associated changes in applied tariffs were virtually identical, but the observed trade responses differed sharply. Vietnam's short-run trade response to NTR access was larger than China's, and Vietnamese exports converged to their long-run level much faster. Our interpretation is that since the observed tariff change was the same, the differences in trade dynamics must be due to differences in the expected path of policy, and our modeling reveals that Vietnam's NTR access was viewed as more likely ex ante and more credible ex post. We argue that these findings are consistent with the geopolitical narrative: the Cold War concerns that had dominated U.S. policy toward China in the 1970s had largely faded by the time Vietnam gained NTR access three decades later.

3 China's integration, reinterpreted

The modern history of U.S. trade policy toward China begins in 1971, when President Nixon lifted the two-decade-old embargo and Chinese goods began entering the U.S. market at the punitive Non-Normal Trade Relations tariffs set by the 1930 Smoot-Hawley Tariff Act. China faced these tariffs for the next nine years. In 1980, under the authority of the Jackson-Vanik amendment, President Carter granted China conditional NTR status, and U.S. tariffs on Chinese goods fell abruptly to the levels applied to most other trading partners. This status was not permanent: it had to be renewed annually by the U.S. President, and from 1990 onward by the U.S. Congress as well, in a procedure that put China's access to the U.S. market explicitly on the legislative calendar every summer.[7] In December 2001, China acceded to the World Trade Organization and gained Permanent Normal Trade Relations (PNTR), ending the annual renewal process.

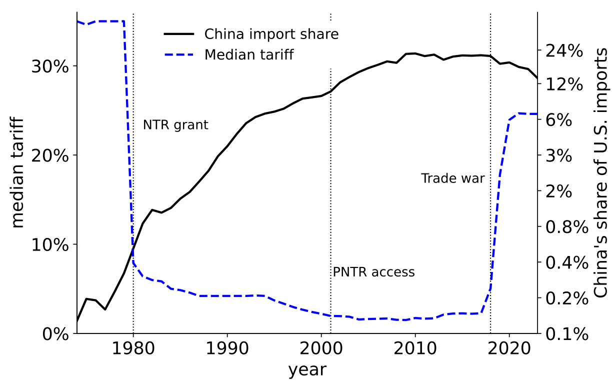

Figure 3: Median tariff on Chinese goods (left axis) and China's share of U.S. imports (right axis, log scale). Reproduced and extended from Alessandria et al. (2025b).

Over the next decade, China's share of U.S. imports more than doubled, going from about 8 percent to almost 20 percent as shown in Figure 3, despite the fact that PNTR access did not change U.S. tariffs on China at all. The conventional explanation for this pattern, due most influentially to Pierce and Schott (2016), is that PNTR access reduced uncertainty about U.S. tariffs on China going

[7] The House of Representatives actually voted to revoke China's NTR status three times in the early 1990s, following the events at Tiananmen Square, although the Senate never put it to a vote.

forward by eliminating the possibility of reverting back to the NNTR tariff schedule. To support this idea, Pierce and Schott (2016) point out that imports of goods with the largest gaps between NNTR and NTR tariffs, which had the most to lose from a hypothetical return to NNTR, grew faster between the early 1990s and the mid-2000s than imports of low-gap goods. Formally, they estimate the following equation:

$$\log \text{imports}_{cgt} = \beta \mathbb{1}_{\{t \geq 2000 \wedge c = \text{China}\}} \times \text{gap}_g \; + \; \delta_{cg} + \delta_{ct} + \delta_{gt} + u_{cgt}, \tag{2}$$

where c denotes the source country, g denotes the product, gap_g is the difference between the NNTR and NTR tariffs on good g, and δ_{cg}, δ_{ct}, and δ_{gt} are the standard country-good, country-year, and good-year fixed effects. The coefficient of interest, β, captures the relative growth in imports of high-gap goods from China following PNTR access relative to (i) low-gap goods from China and (ii) other trade partners. Pierce and Schott (2016) find that this "gap elasticity" is positive and statistically significant: imports of high-gap Chinese goods grew faster than imports of low-gap goods following China's WTO accession, and the same pattern is absent for other countries.

In Alessandria et al. (2025b), we argue that the conventional explanation for the post-PNTR growth of Chinese imports does not tell the whole story. The NNTR gap is indeed a measure of exposure to the risk of losing NTR status, but it is also a measure of something else: the size of the initial tariff reduction that occurred in 1980 when China first got NTR status. Goods with the largest NNTR gaps are precisely the goods that experienced the largest trade liberalizations in the first place. Because trade adjusts gradually to tariff changes, the gap elasticity in (2) captures the delayed effects of the 1980 liberalization as well as the potential effects of reducing tariff risk in 2001. More subtly, if trade also adjusts gradually to changes in trade-policy expectations, this elasticity will also capture the delayed effects of changes in expectations that occurred prior to China's WTO accession.

Building on the Pierce and Schott (2016) empirical approach, we study how the elasticity of Chinese imports to the NNTR gap evolved over time, from the early 1970s to just before the 2009 Great Recession. We use the following specification:

$$\log \text{imports}_{cgt} = \sum_{t'=1974}^{2008} \beta_t \mathbb{1}_{\{t = t' \wedge c = \text{China}\}} \times \text{gap}_g \; + \; \delta_{cg} + \delta_{ct} + \delta_{gt} + u_{cgt}, \tag{3}$$

The time-varying coefficient β_t captures the difference in U.S. imports of high-gap versus low-gap goods from China in a given year t, relative to that same difference in imports from other countries. The left panel of Figure 4 plots the resulting series (solid line with black dots). Three features stand out. First, the gap elasticity gets larger in magnitude (more positive in earlier years) the further back in time one goes. Second, it shrinks during the late 1980s and early 1990s, well before PNTR, and the cumulative change during this earlier period is much larger than the change around 2001. Third, the elasticity was largest in magnitude during the 1970s, before China was granted NTR access. During this initial period, the gap elasticity does not capture the effect of the risk of losing NTR status and experiencing a tariff hike; instead, it captures the effect of the possibility of gaining that status and experiencing a tariff cut.

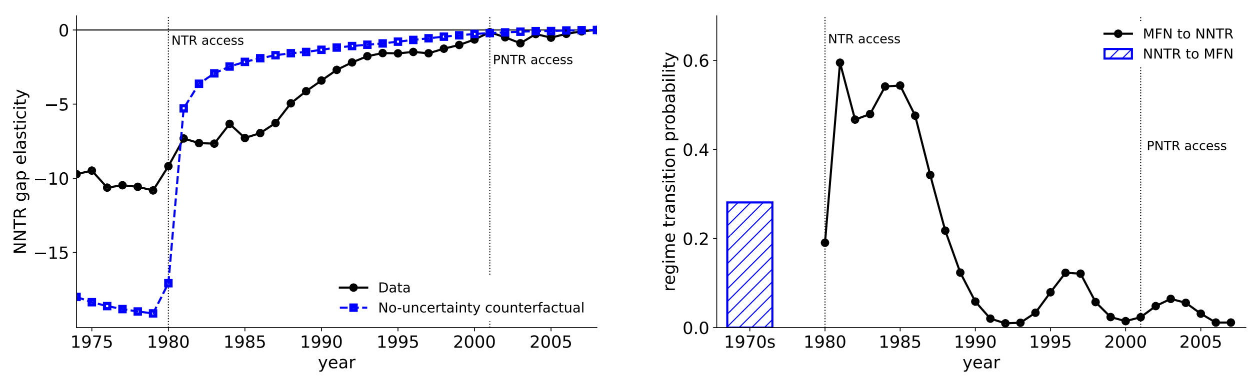

(a) NNTR gap elasticity

(b) Regime transition probability

Figure 4: (a) NNTR-gap elasticity dynamics in data (black dots) versus model counterfactual without policy uncertainty (blue squares), estimated as in (3). (b) Inferred probabilities of switching between NNTR and NTR tariff regimes. Reproduced from Alessandria et al. (2025b).

The change in the gap elasticity from one year to the next partly captures the delayed adjustment to the 1980 liberalization, and partly captures the effects of expectations about future U.S. trade policy toward China (and delayed adjustments to past changes in expectations). To separate these two channels, we use the Alessandria et al. (2021b) model described in Section 2.2 above to estimate the path of trade-policy expectations using indirect inference, and then analyze counterfactual scenarios with different expectations. We calibrate the model's technological parameters—sunk entry costs, continuation costs, and the parameters governing exporter life cycles—to match Chinese firm-level data from 2004–2007, and choose the year-by-year probabilities of switching between the NNTR and NTR tariff schedules so that the model reproduces the observed path of the gap elasticity in the left panel of Figure 4.

The results of this procedure are shown in the right panel of Figure 4. The probability of gaining NTR status during the 1970s was about 25 percent: it was not a complete surprise when this status was actually granted in 1980. The probability of losing NTR status and returning to NNTR was initially very high, about 60 percent, in the early 1980s, but began to fall in the late 1980s. It rose temporarily during the mid-1990s, but fell close to zero again several years before China acceded to the WTO and did not change materially afterwards. In other words, the credibility of China's access to the U.S. market was largely established during the Reagan and Bush administrations, not during the Clinton-era push for PNTR. China's 1986 application to the GATT, the gradual normalization of diplomatic relations, and the fading of Cold War concerns about trade with non-market economies all appear to have mattered more than the 2001 event itself. Thus, the conventional story—that U.S. imports from China surged in the 2000s because WTO accession and PNTR reduced uncertainty about U.S. trade policy—misdates the key credibility shift by more than a decade.

The calibrated model lets us ask: what would the growth of U.S. imports from China have looked like if the 1980 reform had been perfectly credible from the start, with no risk of reversal? The left panel of Figure 4 shows how the elasticity of trade to the NNTR gap would have evolved in this "no-uncertainty counterfactual." In the absence of any uncertainty about U.S. trade policy, imports of high-gap goods from China would have been substantially lower during the 1970s and would have grown substantially faster during the 1980s and early 1990s, but would have grown at about the same pace thereafter; the counterfactual gap elasticity tracks the observed path very closely after 2000. Our results indicate that of the total change in the gap elasticity between the early 1990s and the 2000s—the change originally documented by Pierce and Schott (2016)—about 45 percent is delayed adjustment to the 1980 tariff cut itself, and virtually all of the remaining 55 percent is delayed adjustment to the increase in policy credibility during the 1980s and 1990s; essentially none of this growth is attributable to the effect of WTO accession itself on policy expectations.

In short, China's 2001 WTO accession was not the moment the United States committed credibly to market access for Chinese goods. That commitment was already most of the way in place. The post-2001 boom in Chinese exports to the United States is best understood as the continuation of a trajectory set in motion two decades earlier.

4 Disintegration: the 2018 trade war

After China acceded to the WTO and gained PNTR access in 2001, its exports to the United States were treated exactly the same as any other country's for most of the first two decades of the 21st century. Beginning in 2018, however, President Trump imposed a new tariff schedule on China following an investigation into its treatment of U.S. intellectual property, in a series of actions that brought the median tariff on Chinese goods from 3 percent to about 25 percent by mid-2019, as shown in Figure 3. In Alessandria et al. (2025c), we build on the approach developed in Alessandria et al. (2025b) to recover market participants' beliefs about how long this unprecedented trade war would last.

The new tariff schedule introduced in the 2018 trade war featured a similar average rate as the old NNTR tariff schedule from the Smoot-Hawley era, but as Figure 5a shows, it featured less dispersion and the two schedules were largely orthogonal: on average, goods exposed to trade-war risk were not the same as goods exposed to NNTR risk. This orthogonality allows us to identify beliefs about the two risks separately. Our first step is to estimate a version of (3) with two gap-elasticity coefficients, one for the same NNTR gap as in Alessandria et al. (2025b), and another for the trade-war gap, which we define as the gap between the trade-war tariff and the NTR tariff. Figure 5b plots the results of this estimation.

Both gaps were stable up until the start of the trade war, indicating that the trade war was largely unanticipated, despite President Trump's campaign pledge to raise tariffs on China. Afterwards, the trade-war gap fell below zero: imports of goods with higher trade-war tariffs fell relative to imports of goods with lower tariffs, as one would expect. This drop took place gradually, with most of the adjustment taking place after 2019. The pace of adjustment to the trade-war tariffs was almost identical to the pace of the adjustment to the 1980 tariff cut that occurred when China was initially granted NTR status. Given the results in our prior work, this suggests that the trade war was initially viewed as a transitory aberration that would end quickly. Interestingly, the NNTR gap elasticity rose above zero, indicating that imports of goods that would experience larger tariff hikes if China were returned to NNTR status grew relative to imports of other goods. This growth—about one log point—was about the same as the increase in the NNTR-gap elasticity around China's 2001 WTO accession documented by Pierce and Schott (2016).

Our second step is to use our structural model to recover the probabilities at each point in time of switching between tariff schedules: of the trade war starting (switching from NTR to the trade-war schedule); of the trade war ending (switching back from the trade-war schedule to NTR); and of reverting back to the old NNTR schedule. We use the same basic approach as in Alessandria et al. (2025b): calibrating the model to simultaneously match firm-level data on the behavior of Chinese exporters from the pre-trade-war era, and the dynamics of the two gap elasticities shown in Figure 5b. The regime-switching probabilities that we recover are shown in Figure 5c. Until the

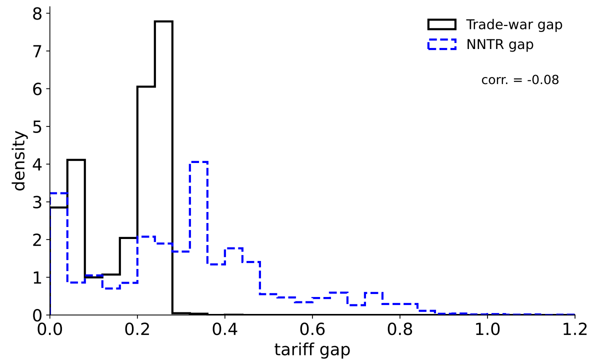

(a) Tariff-gap distributions

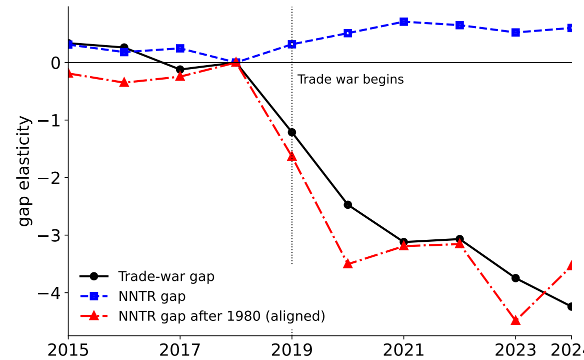

(b) Trade-war and NNTR gap elasticities

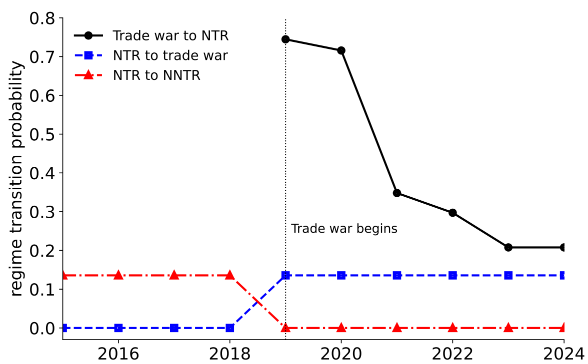

(c) Inferred regime transition probabilities

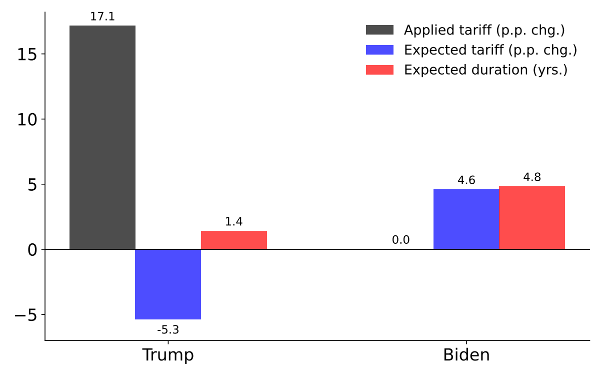

(d) Trade-policy innovations by administration

Figure 5: (a) Trade-war gap and NNTR gap distributions. (b) NNTR-gap and trade-war gap elasticity estimates. Estimates for the NNTR-gap elasticity following the 1980 NTR grant from Alessandria et al. (2025b) are also shown, with sign flipped and year shifted to align the reform year of 1980 with the start of the trade war in 2018. (c) Inferred probabilities of switching between NNTR, NTR, and trade-war tariff regimes. (d) Changes in applied and expected trade policy under Trump (2017–2020) and Biden (2021–2024) administrations. Years start in July, rather than January, to synchronize with the timing of the trade war. Reproduced from Alessandria et al. (2025c).

trade war actually started, market participants believed it to be a zero-probability event, and that switching back to NNTR status was a material possibility. Once it started, the probability that it would end was initially quite high, almost 80 percent. Starting in 2021, however, the likelihood of the trade war ending fell sharply, and continued to fall throughout the Biden administration. Moreover, the probability of going back to NNTR status fell to zero, indicating that the trade war did not just increase applied tariffs and uncertainty about future tariffs, but also fundamentally changed the nature of the policy risk that Chinese exporters faced.

These results imply a striking finding: while President Trump sharply increased applied tariffs, expected future tariffs rose much more under his successor, President Biden. As Figure 5d shows, when President Trump left office in January 2021, the trade war was expected to last only 1.4

years, but by the time President Biden's term ended, it was expected to last for about another five years. Surprisingly, the expected future tariff actually fell under President Trump: the decline in the probability of switching to NNTR, which has a higher average tariff rate than the trade-war schedule, outweighs the trade war itself, which was initially believed to be transitory. Under President Biden, who largely maintained the existing trade-war tariff schedule, the expected future tariff rose by almost five percentage points due to the perception that the trade war had become more persistent.

5 A systematic role for expectations

Our case studies of the history of U.S.-China trade show that trade-policy expectations can be an important determinant of trade growth. In Alessandria et al. (2025a), we show that these cases are not isolated—expectations play a key role in determining how trade responds to changes in tariffs in general—and we argue that this should change our understanding of how conventional trade-elasticity estimates relate to the major trade liberalizations emphasized in our theories and quantitative analyses. The headline finding is that the response of trade to a credible reform is roughly six times larger than the response to a typical tariff change.

We consider the entire sample of U.S. imports from all countries, not just China, over 1974–2017. During this period, which includes most major trade reforms since World War II but excludes the more recent trade war episodes, U.S. trade policy featured four distinct statutory regimes. Two of these regimes, NNTR and NTR, are familiar from our China case studies. The other two regimes are Preferential Trade Agreement (PTA), such as the North American Free Trade Agreement (NAFTA) and its successor, the Canada-United States-Mexico Agreement (CUSMA), and the Unilateral Trade Preference Program (UTPP), a form of preferential tariff treatment that the United States has, at various times, afforded to developing countries. We assign each observation—a country, product, and time triplet—in our dataset into one of these four regimes. As Panel A of Table 1 shows, the vast majority of the sample comprises observations in the NTR regime, which features low tariffs on average but a fair amount of dispersion. The fewest observations are in the NNTR regime, which features the highest and most widely dispersed tariffs. The PTA and UTPP regimes, which both feature approximately zero tariffs, are more common than NNTR but still occur far less frequently than NTR.

Our focus in this paper is on the properties and effects of tariff changes. Having assigned each observation to a statutory regime, we can classify observations into those where the regime was

Table 1: Tariff regimes and transitions in U.S. import data

Panel A: Tariff levels by regime				
Regime	Num. obs.	Mean (p.p.)	Median (p.p.)	SD (p.p.)
NTR	1,758,012	4.87	2.90	14.09
NNTR	19,507	38.61	35.00	22.56
PTA	97,548	0.78	0.05	2.21
UTPP	230,683	0.40	0.00	1.11
Panel B: Average regime transition probabilities (percent)				
$t-1 \backslash t$	NNTR	NTR	PTA	UTPP
NNTR	80.02	18.66	0.00	1.31
NTR	0.05	96.93	0.65	2.37
PTA	0.00	8.58	91.42	0.01
UTPP	0.01	16.60	0.93	82.46

Notes: Reproduced from Alessandria et al. (2025a).

the same from one period to the next ("within-regime" observations) and those where the regime changed ("across-regime" observations). Panel B shows that all four regimes are quite persistent, with transitions out of the NTR regime especially rare. Together, these summary statistics in the two panels imply that the vast majority of tariff changes occur within the NTR regime, and that tariff changes triggered by statutory regime shifts are extremely rare. This raises potential external validity concerns about conventional trade-elasticity estimates that use pooled samples. Such estimates are identified primarily from tariff variation within the NTR regime, which means that they may not be suitable for predicting the effects of major trade liberalizations that switch countries from one regime to another—especially if these two types of tariff changes have different expectational properties.

Figure 6a shows that this concern has bite: within-regime tariff changes do appear to follow a different stochastic process than across-regime changes. The figure shows the results of estimating the following conditional autocorrelation specification:

$$\begin{aligned}\Delta_h \text{tariff}_{cgt} = \beta_h^{\text{within}} \mathbb{1}_{\{\text{regime}_{cgt}=\text{regime}_{cgt-1}\}} \times \Delta_0 \text{tariff}_{cgt} \\ + \beta_h^{\text{across}} \mathbb{1}_{\{\text{regime}_{cgt}\neq\text{regime}_{cgt-1}\}} \times \Delta_0 \text{tariff}_{cgt} + \delta_{ct} + \delta_{gt} + u_{cgt},\end{aligned} \tag{4}$$

where the Δ_h operator is the h-period log change, i.e., $\Delta_0 x_t = \log x_t - \log x_{t-1}$ and $\Delta_h x_t = \log x_{t+h} - \log x_{t-1}$. The blue line with square markers shows the autocorrelation for within-regime tariff changes, and the red line with triangle markers shows the across-regime autocorrelation. The black line with circles shows a version estimated on the pooled sample that treats both types of tariff changes equally. Across-regime tariff changes are significantly more persistent than within-regime changes; in terms of the Markov process analyzed in Figure 1, the former corresponds to a switching probability of about 30 percent and the latter to a value of only about 10 percent. Moreover, the pooled autocorrelation is virtually identical to the within-regime autocorrelation. While this is not surprising, it shows that the average tariff change in the U.S. trade data is much more transitory than the major liberalizations that most trade economists care about.

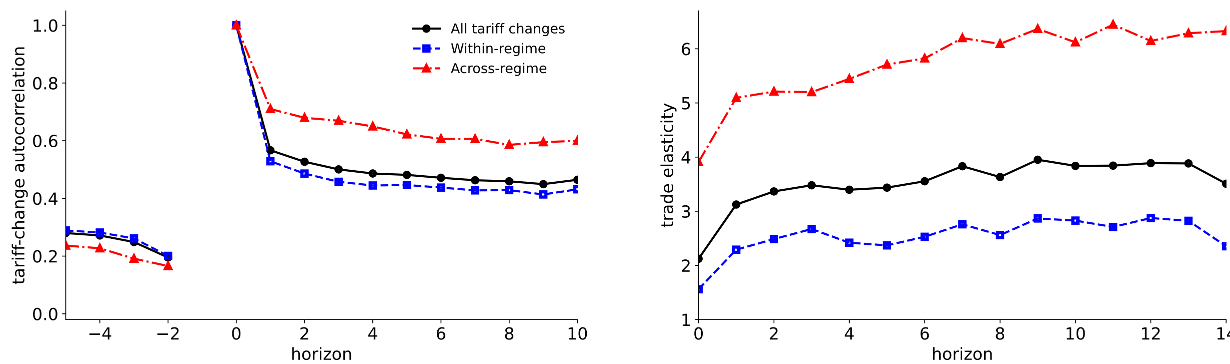

(a) Tariff-change autocorrelation

(b) Trade elasticity

Figure 6: (a) Autocorrelation of tariff changes estimated on pooled sample of U.S. imports versus estimates for within-regime and across-regime tariff changes separately using (4). (b) Corresponding trade elasticities estimated using (5). Reproduced from Alessandria et al. (2025a).

Having documented this difference in expectational properties between within- and across-regime tariff changes, we ask whether it matters for trade dynamics, and find that it does. We estimate the trade elasticity (again, the log change in trade relative to the log change in tariffs), again conditioning on the type of tariff change as in our autocorrelation specification:

$$\Delta_h \text{imports}_{cgt} = \beta_h^{\text{within}} \mathbb{1}_{\{\text{regime}_{cgt} = \text{regime}_{cgt-1}\}} \times \Delta_h \text{tariff}_{cgt} + \beta_h^{\text{across}} \mathbb{1}_{\{\text{regime}_{cgt} \neq \text{regime}_{cgt-1}\}} \times \Delta_h \text{tariff}_{cgt} + \delta_{ct} + \delta_{gt} + u_{cgt}. \tag{5}$$

Following the local-projections approach of Boehm et al. (2023), we instrument the h-period tariff change, $\Delta_h \text{tariff}_{cgt}$, with the one-period change, $\Delta_0 \text{tariff}_{cgt}$ in order to control for subsequent tariff changes that are correlated with the initial change. Our version of this approach uses (4) as the instrument, which allows us to account for the differences in autocorrelations across types of tariff changes documented in Figure 6a.

Figure 6 shows the results of this estimation. The responses of trade to across-regime tariff changes are substantially larger than the responses to within-regime changes, especially at longer horizons. This is precisely what one would expect from our modeling in Section 2.2, where we showed that trade responds more in the long run to persistent tariff shocks than transitory ones. We argue that the larger short-run response to across-regime tariff changes is likely due to anticipation effects. PTAs are publicly negotiated over extended periods and announced ahead of time, and movement in and out of the UTPP regime is predictably related to a country's level of development. We show in the appendix to Alessandria et al. (2025a) that the short-run response to those types of regime changes is larger than the response to movements out of NNTR, which are fairly similar to the short-run response to within-regime tariff changes. This is also consistent with the modeling in Section 2.2. As in the autocorrelation analysis, the within-regime trade response is similar to the unconditional average response across all tariff changes, and the latter is materially smaller than the across-regime response. This confirms our external validity concern: major liberalizations tend to generate much larger trade responses than the typical tariff change we usually observe.

Theoretical models and quantitative analyses typically focus on unanticipated, perfectly persistent (what we call "canonical") trade reforms, which are the appropriate benchmark for measuring things like the welfare gains from trade. Clearly, the observed response to the average transitory tariff change is a poor predictor of the response to a hypothetical canonical reform, but we argue that even our estimated across-regime responses substantially understate that response; as Figure 6a shows, the average across-regime tariff change is persistent, but nowhere close to perfectly so.

To determine just how big the response to a canonical reform would be, we draw once again on our case study of China from Alessandria et al. (2025b), and extend it with a similar study of Vietnam, which went through a qualitatively similar policy trajectory. Ex post, the grants of NTR status to these two countries were the most persistent trade reforms in U.S. history, but as we know from Alessandria et al. (2025b), they were certainly not perceived that way initially. We use broadly the same procedure for jointly estimating technological parameters and trade-policy expectations described in Section 3, with some improvements made to better handle the quantitative differences between the two countries' experiences integrating into the U.S. market.[8] Then, we use the calibrated model for each country to run the canonical reform, where there is no anticipation of gaining NTR status and no risk of losing it once granted, and compare the response of trade to what we see in the data.

Figure 7 compares the "canonical trade elasticity" we recover from this exercise to our empirical

[8]Section 2.2 describes some of these differences and how they translate into different expectations about U.S. trade policy.

estimates. In the canonical reforms simulated in our calibrated models, trade grows by about 15 times the tariff change. This is about 25 percent larger than the responses to China and Vietnam's NTR grants that we measure in the data, more than twice as large as the response to the typical regime change, and more than six times larger than the response to a typical within-regime tariff change. The figure also shows the simulated elasticities in response to the uncertain tariff changes described in Section 2.2, where tariffs follow a two-state Markov process. Even compared to a reform with a seemingly tiny 1% chance of being reversed, the response to a canonical reform is 16% larger. A switching probability of 10% yields a similar response to across-regime tariff changes, and within-regime changes yield a response even smaller than the most transitory reform we consider in our model experiments. To sum up, the typical tariff change in the data is nothing like the canonical reform that many economists care about, and even the most persistent reforms in U.S. trade history are materially different.

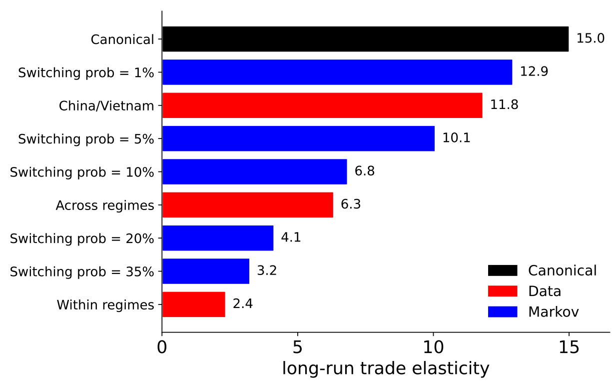

Figure 7: Long-run trade elasticities in model versus data. Black: canonical reform in calibrated model. Blue: reforms in model with two-state Markov tariff process from Section 2.2. Red: empirical estimates. Values for China and Vietnam are averaged. Reproduced from Alessandria et al. (2025a).

6 Broader implications

The research agenda in Alessandria et al. (2025a,b,c) was initially motivated by specific historical questions about U.S.-China trade, but we ultimately discovered that the mechanism it exposes—that trade depends on the entire expected path of policy, not just the current tariff—has implications that reach beyond any particular episode. I organize these implications into three themes, aimed at three audiences: empirical trade economists, theoretical and quantitative researchers studying welfare and optimal policy, and the governments that design and apply trade policy in practice.

Measurement. There is a great deal of disagreement in the empirical literature about the magnitude of the trade elasticity. On the low end, Boehm et al. (2023) estimate an elasticity of 0.7 in the short run and about two in the long run using panel data on within-NTR tariff changes. In the middle of the range, Simonovska and Waugh (2014) and Caliendo and Parro (2015) estimate elasticities of about 4–6 using cross-sectional data on bilateral trade between many country pairs. At the top end, Romalis (2007) and Khan and Khederlarian (2021) estimate long-run elasticities of 6–11 using panel data around the creation of NAFTA. Our perspective is that much of this disagreement can be resolved by recognizing that these estimates come from different empirical contexts with different expectational properties. The response of trade to a tariff change is not an invariant object—it depends on whether it was anticipated and how long firms expect it to last. An elasticity like the one in Boehm et al. (2023), estimated from the transitory tariff variation that dominates the postwar U.S. data is a legitimate statistic when used in the context of that kind of variation, but it is not the same object as the response of trade to a persistent reform, especially at longer horizons.[9] We argue that structural models are needed to disentangle the effects of observed tariff changes themselves from the effects of expectations—and changes to those expectations, which we showed are quantitatively important in Alessandria et al. (2025b) and Alessandria et al. (2025c).

Welfare and optimal policy. If the long-run response of trade to a credible tariff change is several times larger than the reduced-form estimates routinely used in calibrations, then welfare calculations that substitute the latter for the former will misstate the welfare consequences of major reforms. Lower elasticity estimates imply that given changes in trade flows (e.g. autarky to free trade, which is a common benchmark in theoretical work) generate larger welfare gains (Arkolakis et al., 2012). At the same time, however, lower elasticities also imply that given changes in tariffs (e.g. the "Liberation Day" tariffs announced by the United States on April 2, 2025) will generate smaller changes in trade flows. To measure welfare gains from trade accurately, we argue that one cannot avoid explicitly modeling expectations using a dynamic model. One needs to match both the changes in tariffs and trade flows, and reconciling the two will generally require some degree of policy uncertainty. If the canonical trade elasticity is high, as we have argued, but tariff changes often generate small responses because they are believed to be transitory, the welfare consequences of those transitory tariff changes are likely to be very small.

[9] A related, but conceptually distinct issue pertains specifically to cross-sectional specifications. Their estimated trade elasticities mix short-run responses to tariff changes that occurred more recently (which are also more likely to be transitory) with the long-run responses to reforms that occurred further in the past. We argue that cross-sectional approaches are unlikely to be useful in recovering trade elasticities at longer horizons.

On a related point, our findings have important implications for optimal tariffs. The standard optimal-tariff formula says that the lower the trade elasticity, the higher the optimal tariff, because the terms-of-trade effects are stronger (Johnson, 1950). Our estimates of the canonical trade elasticity suggest that the optimal perfectly-credible tariff is likely to be very small—much smaller than recent estimates such as Pujolas and Rossbach (2024), Ignatenko et al. (2025), and Itskhoki and Mukhin (2025), which range from roughly 9 to 34 percent. At the same time, however, tariffs that are not perfectly credible reduce trade less, and thus raise larger amounts of tariff revenue. Depending on the government's objective—consumer welfare, tariff revenues, or some combination of the two—it might be optimal to levy higher, but more uncertain, tariffs. Of course, policy uncertainty can have other adverse side effects (Steinberg, 2019; Caldara et al., 2020)—yet another reason that dynamic trade models are needed to properly evaluate these tradeoffs.

Credibility as a policy instrument. The fact that trade depends on the expected future path of tariffs implies that credibility is as much a policy instrument as the headline tariff rate itself. One somewhat perverse application of this idea is that trade-policy uncertainty could actually be a tool for generating revenue: high tariffs that are expected to be reversed quickly will raise more revenue, at least in the short run, than tariffs that are believed to be permanent, because the former have a smaller effect on trade volumes than the latter. However, the reverse is true when it comes to other goals of trade policy related to real economic activity. Whether the aim is to lower tariffs to foster cross-border integration or to raise tariffs to reshore manufacturing production, the less credible tariffs are, the less of the desired effect the government will get. A corollary is that a government need not change the current tariff at all to move trade. A credible announcement about the future path of trade policy can shift export participation today. This is the trade-policy counterpart of the "forward guidance" that central banks use to move investment and inflation without changing current interest rates. This mechanism may be more powerful in the trade context than in the monetary one, given the highly sunk nature of the investment decisions at stake and the degree to which the returns to these investments are backloaded. A government that cannot change applied tariffs in the present—because of international commitments or domestic political constraints—can still move trade flows meaningfully by signaling credibly about its future intentions.

As a concrete example of this principle, consider the problem of a government that wants to use higher tariffs as a form of industrial policy to build up domestic manufacturing capacity. Tariffs high enough to induce firms to relocate production may invite painful retaliation from trading partners, political pressure at home as input prices rise, and—in response to both—a growing expectation that the tariffs may be reversed. The resulting decline in credibility undermines the goals of the policy, and this effect is likely to be stronger the more ambitious the policy is; inducing more reshoring requires higher tariffs, which, in the political equilibrium outlined above, will have a greater likelihood of reversal. This is particularly likely if the reshoring benefits take a long time to materialize but the other downsides come more quickly, which I argue is likely to be the case in other research (Steinberg, 2025a). This is not an argument against tariffs-as-industrial-policy in principle, but it outlines a tight constraint on how such a strategy must be designed in practice. A strategy that does not come with a credible mechanism for locking in the tariffs is likely to be self-defeating by construction.

If credibility is a lever, the question is how to pull it. Central banks faced an analogous problem in the late twentieth century—how to persuade markets that the inflation target announced today would still be honored when political incentives changed tomorrow—and the solution was institutional: delegated authority to independent bodies, statutory constraints on discretionary action, and transparency mechanisms that raised the political cost of deviation. The counterparts in trade policy include binding international commitments under the WTO and preferential trade agreements, statutory limits on executive tariff-setting authority, and domestic institutional arrangements that make reversal politically expensive. As we have seen over the past decade, none of these devices fully ties the government's hands, and none is a substitute for political willingness to honor commitments. Friedman and Friedman (1980) advocated solving this problem with a Constitutional amendment that sets import tariffs to zero, mirroring the Constitution's treatment of export taxes. In any case, the experience of monetary policy suggests that credibility is not something governments possess by virtue of their preferences alone. It is something they build by making their preferences harder to abandon.

7 Conclusion

Unanticipated, perfectly credible tariff changes—canonical reforms—are the primary focus of the trade literature's theoretical analyses and quantitative experiments, but such hypotheticals do not exist anywhere in the data. The typical reform that we observe in the data is highly transitory and generates a much smaller trade response than what we believe a canonical reform would generate. Even the few highly persistent reforms that we do observe, such as China's NTR grant or the formation of PTAs, are not perfectly credible and generate substantially smaller responses. In Alessandria et al. (2025a,b,c), my coauthors and I account for these patterns using a dynamic trade model in which firms make forward-looking investments in market access, putting expectations about future trade policy at the forefront.

Our insights have important implications for both researchers and policymakers. For researchers aiming to predict the effects of trade reforms and analyze their economic implications, conventional reduced-form elasticities may be reliable for studying small, transitory tariff changes precisely because this sort of change occurs frequently, but these estimates should not be used to study large, credible reforms, especially in the long run. For policymakers, credibility is an instrument that matters just as much as applied tariff rates. A policymaker who wants to use trade policy to achieve large, durable economic effects—whether the goal is integrating supply chains across borders or reshoring production—needs to deliver a policy that is perceived to be credible. Tariffs imposed with a visible expectation of reversal, or cut with a visible threat of reimposition, will produce smaller effects than their headline magnitudes suggest.

Uncertainty about trade policy has been one of the defining characteristics of the current geopolitical era since the United Kingdom voted to leave the European Union in June 2016, but it reached an unprecedented level of intensity with the start of U.S. President Trump's second term in 2025. Viewed through the lens of the research program described in this article, these developments are bittersweet. On the one hand, they vindicate the perspective that my coauthors and I have long held: trade policy is not a a permanent fixture of the economic landscape, and even free-trade agreements such as CUSMA should not be viewed as perfectly credible. On the other hand, these developments have made it significantly harder to continue to advance this agenda. Most of our work thus far has leveraged the well-defined structure of U.S. trade policy during the pre-Trump period, when fixed tariff schedules were set legislatively and the relevant risk was movement between regimes. Today, the very idea of a fixed tariff schedule seems almost quaint; the relevant risk is now how tariff schedules themselves will continue to evolve.[10] Trade-policy uncertainty is more important than ever, but measuring it and quantifying its effects is also harder than ever.

References

Acemoglu, Daron, David Autor, David Dorn, Gordon H. Hanson, and Brendan Price (2016) 'Import competition and the great US employment sag of the 2000s.' *Journal of Labor Economics* 34(S1), S141–S198

Alessandria, George A., Shafaat Yar Khan, Armen Khederlarian, Kim J. Ruhl, and Joseph B.

[10] I explore the effects of ambiguity (a.k.a. Knightian uncertainty) about trade policy in Steinberg (2025b).

Steinberg (2025a) 'Recovering credible trade elasticities from incredible trade reforms.' NBER Working Paper 33568, National Bureau of Economic Research, March

Alessandria, George, Costas Arkolakis, and Kim J. Ruhl (2021a) 'Firm dynamics and trade.' *Annual Review of Economics* 13(1), 253–280

Alessandria, George, Horag Choi, and Kim J. Ruhl (2021b) 'Trade adjustment dynamics and the welfare gains from trade.' *Journal of International Economics* 131, 103458

Alessandria, George, Shafaat Yar Khan, Armen Khederlarian, Kim J. Ruhl, and Joseph B. Steinberg (2025b) 'Trade policy dynamics: Evidence from 60 years of US-China trade.' *Journal of Political Economy* 133(3), 713–749

_ (2025c) 'Trade war and peace: U.S.-China trade and tariff risk from 2015–2050.' *Journal of International Economics* 155, 104066

Arkolakis, Costas, Arnaud Costinot, and Andrés Rodríguez-Clare (2012) 'New trade models, same old gains?' *American Economic Review* 102(1), 94–130

Autor, David H., David Dorn, and Gordon H. Hanson (2013) 'The China syndrome: Local labor market effects of import competition in the United States.' *American Economic Review* 103(6), 2121–2168

Autor, David H., David Dorn, Gordon H. Hanson, and Jae Song (2014) 'Trade adjustment: Worker-level evidence.' *The Quarterly Journal of Economics* 129(4), 1799–1860

Baldwin, Richard, and Paul Krugman (1989) 'Persistent trade effects of large exchange rate shocks.' *The Quarterly Journal of Economics* 104(4), 635–654

Bianconi, Marcelo, Federico Esposito, and Marco Sammon (2021) 'Trade policy uncertainty and stock returns.' *Journal of International Money and Finance* 119, 102492

Boehm, Christoph E., Andrei A. Levchenko, and Nitya Pandalai-Nayar (2023) 'The long and short (run) of trade elasticities.' *American Economic Review* 113(4), 861–905

Caldara, Dario, Matteo Iacoviello, Patrick Molligo, Andrea Prestipino, and Andrea Raffo (2020) 'The economic effects of trade policy uncertainty.' *Journal of Monetary Economics* 109, 38–59

Caliendo, Lorenzo, and Fernando Parro (2015) 'Estimates of the trade and welfare effects of NAFTA.' *The Review of Economic Studies* 82(1), 1–44

Chaney, Thomas (2008) 'Distorted gravity: The intensive and extensive margins of international trade.' *American Economic Review* 98(4), 1707–1721

Das, Sanghamitra, Mark J. Roberts, and James R. Tybout (2007) 'Market entry costs, producer heterogeneity, and export dynamics.' *Econometrica* 75(3), 837–873

Friedman, Milton, and Rose D. Friedman (1980) *Free to Choose: A Personal Statement* (New York: Harcourt Brace Jovanovich)

Handley, Kyle, and Nuno Limão (2017) 'Policy uncertainty, trade, and welfare: Theory and evidence for China and the United States.' *American Economic Review* 107(9), 2731–2783

Ignatenko, Anna, Ahmad Lashkaripour, Luca Macedoni, and Ina Simonovska (2025) 'Making america great again? the economic impacts of liberation day tariffs.' *Journal of International Economics* 157, 104138

Itskhoki, Oleg, and Dmitry Mukhin (2025) 'The optimal macro tariff.' NBER Working Paper 33839, National Bureau of Economic Research, May

Johnson, Harry G. (1950) 'Optimum welfare and maximum revenue tariffs.' *The Review of Economic Studies* 19(1), 28–35

Khan, Shafaat Yar, and Armen Khederlarian (2021) 'How does trade respond to anticipated tariff changes? evidence from NAFTA.' *Journal of International Economics* 133, 103538

Pierce, Justin R., and Peter K. Schott (2016) 'The surprisingly swift decline of U.S. manufacturing employment.' *American Economic Review* 106(7), 1632–1662

— (2018) 'Investment responses to trade liberalization: Evidence from U.S. industries and establishments.' *Journal of International Economics* 115, 203–222

Pujolas, Pau, and Jack Rossbach (2024) 'Trade wars with trade deficits.' Available at SSRN

Romalis, John (2007) 'NAFTA's and CUSFTA's impact on international trade.' *The Review of Economics and Statistics* 89(3), 416–435

Ruhl, Kim J., and Jonathan L. Willis (2017) 'New exporter dynamics.' *International Economic Review* 58(3), 703–726

Simonovska, Ina, and Michael E. Waugh (2014) 'The elasticity of trade: Estimates and evidence.' *Journal of International Economics* 92(1), 34–50

Steinberg, Joseph B. (2019) 'Brexit and the macroeconomic impact of trade policy uncertainty.' *Journal of International Economics* 117, 175–195

Steinberg, Joseph B (2025a) 'Tariffs, manufacturing employment, and supply chains.' Working Paper 34236, National Bureau of Economic Research

_ (2025b) 'Trade policy ambiguity.' Working Paper 34408, National Bureau of Economic Research