Recovering Credible Trade Elasticities from Incredible Trade Reforms*

George Alessandria†, Shafaat Yar Khan‡, Armen Khederlarian§,
Kim J. Ruhl¶, and Joseph B. Steinberg‖

February 2025

Abstract

We study how trade-policy dynamics affect the dynamics of trade volumes and the implications of these effects for estimates of the trade elasticity. We use data on US imports and trade policy from 1974–2017 for China and Vietnam—the countries with the largest import growth and the largest tariff reductions over the last fifty years—and a heterogeneous-firm dynamic trade model to recover the dynamic path of the trade elasticity following an unanticipated, permanent tariff change. We estimate a short-run trade elasticity of about three and a long-run trade elasticity of about 14, and find that it takes about five years to close half of the gap between the current and long-run levels of trade. We argue that the expected dynamics of future trade policy before and after these reforms biases down reduced-form estimates of the long-run trade elasticity and biases up estimates of the short-run elasticity. We argue that these measurement issues are even more problematic for other trade reforms, especially those within the Normal Trade Relations (NTR) regime that constitute the majority of the data.

JEL Classifications: F12, F13, F14

Keywords: Trade elasticity, trade dynamics, trade policy uncertainty

*We thank seminar audiences at the University of Rochester, St. Louis Federal Reserve, Syracuse University, 2024 ASSA Meetings, and the 2024 Canadian International Trade Study Group for comments. We thank Yan Bai, Fernando Parro, Anson Soderberry, Mike Waugh, and Kei Mu Yi for helpful discussions. We thank Trang Hoang for help with the Vietnamese trade data and Dan Lu and Lianming Zhu for help with the Chinese firm level data. Kazuma Inagaki provided excellent research assistance. Alessandria and Ruhl are grateful for support from the National Science Foundation (award #2214852). Ruhl completed his work on this project prior to joining the Council of Economic Advisers in the Executive Office of the President of the United States.

†george.alessandria@rochester.edu, University of Rochester and NBER
‡shafaatyar.k@gmail.com, Syracuse University
§armen.khederlarian@gmail.com, Hunter College, CUNY
¶ruhl2@wisc.edu, University of Wisconsin–Madison and NBER
‖joseph.steinberg@utoronto.ca, University of Toronto

1 Introduction

The elasticity of trade volumes to trade policy is one of the most important concepts in international economics. It determines how trade reforms affect substitution between domestic and foreign goods, reallocation of workers across sectors, terms of trade effects, and the income generated from tariffs. It is a valuable input into quantitative studies to evaluate the aggregate consequences of changes in policies or fundamentals. Despite its importance, there is no clear consensus on the magnitude of the trade elasticity for two key reasons.[1] First, trade responds gradually to policy changes, so the elasticity depends on the time horizon over which it is measured. Second, the response depends on expectations about future policy, which means the same observed policy change can generate many different outcomes.

Ideally, one would like to estimate a sequence of trade elasticities following an unanticipated, once and for all tariff change, but we argue that such "canonical" reforms do not exist in the data. We recover this path by using data from two unique trade reforms—the United States granting Normal Trade Relations (NTR) to Vietnam and China—as inputs to a structural model with a time-varying trade policy process. In the model, a canonical trade liberalization generates a short-run trade elasticity of about three and a long-run trade elasticity of about 14. One-half of the gap between the short and long-run elasticity closes in five years.

Building on the insights of Baldwin (1986) and Baldwin and Krugman (1989), a large literature finds the structure of trade costs makes firm-level export participation a forward-looking decision (Roberts and Tybout, 1997; Das et al., 2007). A common finding is entering and growing in export markets involves repeated up-front investments with back-loaded returns, which implies trade responds gradually to policy changes (Ruhl and Willis, 2017; Alessandria et al., 2021). It also implies trade depends on the underlying trade policy process—and changes to that process—as well as the observed sequence of realizations. For example, Ruhl (2008) argues that trade responds less to temporary shocks than permanent shocks, and Handley and Limão (2015) argue that uncertainty about future policy depresses

[1] Hillberry and Hummels (2013) provide a survey of trade elasticity estimates that range from 1 to 35. Head and Mayer (2014) review 435 elasticities from 32 papers: they obtain a median estimate of 5.03 with a standard deviation of 9.3.

trade in the present. Thus, it is difficult to measure the trade elasticity using data alone, as one must separately identify the lagged effects of past reforms and the effects of expectations over future reforms. Our approach, which builds on Alessandria et al. (2025), is to measure the trade elasticity by simulating a canonical reform in an estimated structural model.

In the first part of the paper, we use data on US imports and tariffs from 1974–2017 to document the dynamics of trade policy and estimate reduced-form trade responses to these tariff changes. Throughout this period, the United States has maintained two main tariff schedules: Non-Normal Trade Relations (NNTR) and NTR. The NNTR schedule was largely set by the 1930 Smoot-Hawley Tariff Act and is considered exogenous to any specific trading partner (Handley and Limão, 2017; Pierce and Schott, 2016). The NTR schedule has evolved gradually through multilateral negotiations that did not involve China nor Vietnam, making it largely exogenous to them as well. Other US tariff schedules are the result of Preferential Trade Agreements (PTAs) and the Unilateral Trade Preference Program (UTPP). This structure implies there are two kinds of tariff changes: (i) those that occur when countries shift from one schedule to another, leading to simultaneous tariff changes across many goods at once; (ii) and those that occur when the schedules themselves change, leading to good-specific changes in tariffs that are common across the countries that trade under that schedule.

We compare the two types of tariff changes in terms of the dynamics of the tariffs themselves and in terms of their effect on trade. Schedule transitions are rare, but cause large, highly-persistent changes in tariffs and have significantly higher long-run trade elasticities than typical estimates in the literature. Within-schedule tariff changes are common, but small and transitory outside of the GATT rounds, and have low long-run trade elasticities, similar to typical estimates. We also find that entering or exiting PTAs or the UTPP have higher short-run trade elasticities than other tariff changes, which is likely because these switches are more anticipated than other changes.[2]

We study two important transitions between the NNTR and NTR regimes: China and Vietnam. These transitions followed similar paths at different points in time. China gained

[2]PTAs are the product of years of public negotiations and involve a gradual phaseout of tariffs. Countries typically leave the UTPP when they cross predetermined income per capita thresholds, which is also predictable.

conditional NTR status in 1980 subject to annual renewal. In 2000, Congress passed a bill, that was signed by the President, to take China off of the annual renewal cycle upon joining the World Trade Organization (WTO). Thus, China gained Permanent NTR in 2002. Vietnam went through a qualitatively similar experience. It gained conditional NTR status in 2002 and permanent status when it joined the WTO in 2007. Importantly, both countries had been subject to a trade embargo allowing us to study the transition from autarky, and more importantly, control for initial conditions in each good.

The response of trade to the changes in policy for Vietnam and China share some similarities and have some important differences. Taking a window a few years before gaining NTR to several years after gaining PNTR, we find trade grows about 10-11 times the change in tariffs in both countries, substantially greater than the estimate when we consider all other tariff regime changes. The transition paths though are different, suggesting different degrees of anticipation and uncertainty. Vietnam's short-run elasticity is larger than China's and converged to its long-run level faster. This suggests that Vietnam's NTR access was more anticipated and was initially viewed as more credible. We argue this is consistent with the historical evidence. China was one of the first ex-communist countries to gain NTR status, and its relationship with the United States was fraught with political tension. Vietnam, however, was one of the last countries to gain this status, and, by the time it did so, only one other country had ever lost it. Moreover, with the fall of the Soviet Union, many of the cold war concerns that drove the uncertainty around trade policy had faded.

In the second part of the paper, we use the approach developed in Alessandria et al. (2025) to estimate a path of elasticities from a surprise once and for trade liberalization. We construct a model with two key ingredients: sunk exporting costs and export capacity that grows stochastically with a firm's tenure as an exporter. The first ingredient leads firms to make forward-looking export participation decisions that depend on the entire future path of trade policy. The second ingredient allows the model to accurately capture the rate at which firms discount future profits and contributes to the gradual adjustments in aggregate trade. For each country, we calibrate the model to match important empirical moments, the most of which are the policy-regime transition probabilities.

The identification of the policy-regime transition probabilities comes from the variation in the NNTR-gap elasticities. The NNTR gap is the difference between a good's NNTR tariff and its NTR tariff. The NNTR-gap elasticity measures the differential response of imports to the NNTR gap. When the country trades under the NNTR schedule, an increase in the probability of gaining NTR status has a larger positive effect on exports of goods with high NNTR gaps than goods with low NNTR gaps, which raises the gap elasticity. When the country trades under the NTR schedule, an increase in the probability of switching to the NNTR regime has a larger negative effect on high-gap goods, which lowers the gap elasticity.

We find similarities in the beginning- and end-of-sample probabilities, but some noticeable differences around the transition to NTR. In both countries, we find a low initial probability of gaining NTR access, between five and 15 percent. In Vietnam, the likelihood of switching to the NTR regime gradually rises to almost 30 percent several years before the switch occurs. In China, the probability drifts down slightly. In Vietnam, after gaining NTR access, the probability of reverting to NNTR begins to fall immediately. In China, the probability of reverting to NNTR temporarily increases before falling. At the end of the our sample, the probability of losing NTR access is low but nonzero, even after WTO accession. These probability estimates indicate the reforms were not perceived as surprise once and for all liberalizations. Thus, the reduced form elasticities estimated in the data are biased.

To evaluate the magnitude of these biases, we study the canonical surprise once and for all reform in our model. We solve the model under the counterfactual assumption that each country was initially in a steady state with NNTR tariffs and no chance of moving to the NTR regime. NTR status is then granted unexpectedly with no chance of reversal. We recover very similar trade-elasticity dynamics for both countries: a long-run trade elasticity of about 14, a short-run elasticity of nearly three, and a speed of adjustment such that half the gap between the current and long-run levels of trade is closed every three to five years. This indicates our long-run reduced form estimates for China and Vietnam are biased downward by about 25 percent, and that our short-run estimate for Vietnam is biased upwards by about 100 percent.

Our paper relates to several papers that estimate the dynamics of trade substitution around changes in trade policy. A key challenge in much of this work is the most-favored-nation principle, which implies that there is a dearth of cross-country variation in tariffs that can identify these dynamics. Thus, several papers have leveraged tariff variation from PTAs. Baier and Bergstrand (2007) use a panel of PTAs to identify the dynamics of trade growth that start prior to the signing of the PTA and stretch 10-15 years following the PTA. Kehoe and Ruhl (2013), Baier et al. (2014), and Alessandria and Avila (2023) show that trade agreements lead to large increases in the extensive margin of trade, as measured by the range of goods or firms exporting. Focusing more narrowly on NAFTA, Khan and Khederlarian (2021) estimate a short-run elasticity of 2.7 and a seven-year response that is close to nine for trade with Mexico. Similarly, Besedes et al. (2020) find delayed effects of import growth in Mexico and Canada for goods that have immediate tariff cuts or changes in trade policy risk.[3]

A potential concern with this approach is that the path of tariffs (either within-schedule changes or movements of countries across schedules) may be endogenous. While we do not think this is an issue with China and Vietnam, it could be an issue in general. To account for this endogeneity, Boehm et al. (2023) focus on the effects of within-NTR tariff changes on minor trading partners, which are unlikely to play a material role in determining countries' strategies in setting NTR tariff schedules. Using data on trade and applied tariffs, they estimate trade elasticities that are much lower than our estimates: less than one in the short run and less than two in the long run. However, when we focus on these within-NTR tariff changes, we recover similarly low estimates. This highlights an important external validity concern: estimates of trade elasticities from small, transitory tariff changes within the NTR schedule are likely to be poor predictors of the effects of unanticipated, large, and persistent reforms like PTA formation and NTR access grants.

More broadly, several papers estimate the dynamic response of trade volumes to changes in the exchange rate or relative prices at the aggregate- or industry-level (Junz and Rhomberg, 1965; Houthakker and Magee, 1969; Alaouze et al., 1977; Gallaway et al., 2003). Much of this work also finds a delayed response, although estimates of the trade response to relative prices tend to be much lower than tariffs. For example, Alessandria and Choi (2021) use US net trade flows, relative prices, and expenditures to estimate a quarterly short-run elasticity

[3]Romalis (2007) uses tariff variation from NAFTA to estimate a trade elasticity of substitution between 6.2 and 10.9, based on a ten year change. The estimates do not explicitly model dynamics. Caliendo and Parro (2014), also studying NAFTA, estimate a trade elasticity of 4.5 for Mexico.

of 0.2 and a long-run elasticity of 1.1. They also show that purely empirical approaches that do not control for the source of the shock can substantially bias these estimates downward. Our paper also relates to a growing literature on the effects of future trade policy on the current level of trade reviewed by Handley and Limão (2022). This literature attributes some trade growth to changes in the expected persistence of a trade policy, but largely abstracts from measuring how these changes bias estimates of the trade elasticity.

The remainder of the paper is organized as follows. In Section 2, we study the dynamic properties of trade policy changes in the data. In Section 3, we estimate reduced-form specifications of trade's response to tariff changes. Section 4 introduces the structural model and section 5 uses the model to illustrate the problems inherent in the reduced-form estimates. Section 6 contains the quantitative analysis and section 7 concludes.

2 Empirical evidence on trade-policy dynamics

We begin by documenting several facts about the dynamics of trade policy. We focus on the United States, owing to the availability of high-quality data over a long period.[4] We emphasize that surprise once and for all trade reforms—the canonical reform studied in most models—are extremely rare. Most tariff changes are highly transitory, and more persistent reforms are often characterized by anticipation due to phaseouts or long negotiation periods.

2.1 Overview and History

The United States has a unique tariff structure that creates tariff variation that facilitates estimating trade responses to tariffs. Imports into the US generally enter under one of four tariff regimes:

1. Normal Trade Relations (NTR, or Column 1)
2. Non-Normal Trade Relations (NNTR, or Column 2)
3. Unilateral Trade Preference Program (UTPP)
4. Preferential Trade Agreements (PTA)

For most of its trading partners, the United States uses two tariff schedules, NNTR and NTR. The NNTR schedule is largely determined by the Smoot-Hawley Act of 1930 and is often

[4]Recently, Teti (2023) demonstrates that the tariff data in World Integrated Trade Solutions is plagued by measurement errors that introduce some important biases in elasticity estimates using cross country data.

treated as a source of exogenous variation in product-level tariffs (Pierce and Schott, 2016; Handley and Limão, 2017). The NTR schedule, formerly the Most Favored Nation schedule, has been adjusted over time through multilateral GATT rounds involving successively larger groups of developed countries. Three major GATT rounds—Kennedy, Tokyo, and Uruguay—had large impacts on tariffs, involved long negotiations, and were implemented in phaseout periods of 5–8 years. Prior to the Trade Agreements Extension Act of 1951, trade policy followed the most-favored-nation principle, and all US trading partners, save those for which the United States did not trade, received the NTR rate.[5] With the 1951 Trade Act, all non-market economies, except Yugoslavia, were transitioned from the NTR to the NNTR schedule. Of the 29 countries, which includes successor countries, subject to NNTR under the 1951 Trade Act or subsequent legislation, only North Korea and Cuba have NNTR status at the end of our sample, but trade with both countries are subject to an embargo.[6]

With the Trade Act of 1974, the UTPP regime came into existence with the General System of Preferences (GSP). This program provided tariff-free access to the United States for some goods from less-developed countries. The program has lapsed several times, and countries may be graduated from the program upon reaching thresholds in level of development or market share in a product, their accession to other trade agreements, or other circumstances.[7] While GSP is the best-known temporary trade program, there are several other programs, which we describe in Table A1. Another key provision of the 1974 Trade Act was the Jackson-Vanik Amendment, which created a path for non-market economies to gain access to NTR rates subject to Congressional approval and annual renewal. PTAs came into being in 1985, with the largest US FTA being the Canada-US FTA in 1989 followed by the North American FTA in 1994. In 2017, the last year we consider, the United States had 20 FTAs, as listed in Table A2. PTAs often involve domestic content requirements, leading some goods to be imported under both PTA and NTR rates during the same year.

[5] Under the Trading with the Enemy Act of 1917, the United States imposed a trade embargo on China and North Korea in December of 1950.

[6] For a summary of US trade policy regimes see CRS Report for Congress (2005).

[7] For instance, former non-market economies in the European Union lost access to GSP rates when they joined the European Union.

2.2 Data

We use annual customs data from the US Census Bureau from 1974 to 2023. 1974 is the first year that applied tariffs are available digitally from the National Archives. We end most of our analysis in 2017 to avoid the effects from the US-China Trade War but present some data for the period of 2018–2023 for completeness.[8] The data for 1974–1988 are from the National Archives and Records Administration, while the data for the remaining years are from Schott (2021). To study trade over such a long horizon, we aggregate the data to the SITC 5-digit level.[9] For these SITC goods, we use applied tariffs rather than statutory tariffs. To account for changes in country definitions, particularly with the fall of the Soviet Union and its satellites, we merge countries that at some point in our sample were united, e.g. the Czech Republic and Slovakia after 1989, and East and West Germany before 1990.

Our unit of observation is the triplet jgt, where j denotes the exporting country, g denotes an SITC 5-digit good, and t denotes the year. Our main variables of interest are the log FOB import value, denoted by v_{jgt}, and the applied tariff rate, denoted by τ_{jgt}. We compute the latter as the log of one plus the ratio of total duty charges to the FOB import value. To classify triplets into different statutory tariff regimes, we use information on the rate of provision and country sub-code. The rate of provision and the country sub-code indicate the type of duty and special trade agreement under which a good from a country is imported, respectively.[10] As our data combines goods that come in under different rates of provisions, we assign products to rates of provision based on the most common provision by value in a country-good-year.[11] Our final sample is an unbalanced panel with 44 years, 163 countries,

[8]In Alessandria et al. (2024) we study the evolution of trade and trade policy expectations in the window around the trade war. A unique feature of the trade war was that the Trump administration in 2018 created a new schedule of tariffs.

[9]A challenge to analyzing trade over a long time period is the classification of goods into tariff lines changed in 1989, when the United States switched from the 7-digit TS-USA classification to the 8-digit Harmonized System (HS). Our approach is to use data aggregated to the 5-digit Standard International Trade Classification (SITC) level, for which concordances with the TS-USA and HS schedules are both readily available. The concordance between the TS-USA schedule and the SITC schedule is from Feenstra (1996) and the concordances with the HS schedule and its successive revisions are from United Nations Trade Statistics (2017). Other papers that use this product classification are, for instance, Broda and Weinstein (2006) and Kehoe and Ruhl (2013).

[10]For the current list of rates of provision and country sub-codes visit `https://www.census.gov/foreign-trade/reference/codes/rp.html` and `https://www.census.gov/foreign-trade/reference/codes/csc.html`, respectively.

[11]In appendix B.1 we give an example of our aggregation and consider other ways of aggregating goods.

2,032 goods, and 2,105,521 observations (Table 1). The median tariff is about 1.7 percent but this hides substantial differences across regimes. In the NTR regime, the median tariff is about 3 percent while, in the NNTR regime, it is 35 percent. In the other regimes, the median tariff is very close to zero. There is much dispersion in the NNTR and NTR states, but the dispersion in tariffs in the NTR state appears to be driven by some extreme outliers as when we winsorize at the 0.01 percent level it falls by more than half. We work with this winsorized sample.

Figure 1 plots some features of US imports and tariffs by tariff regime. Panel (a) presents the share of US non-oil imports, by current regime of the importer, relative to the share in the NTR regime. Panel (b) presents the path of tariffs for those same countries.[12] In 1974, nearly all US imports were under NTR. In 1976, with the introduction of the GSP program, there is a discrete jump in the share of countries that gain UTPP and its associated zero tariffs. In 1989 and 1995, the PTA share jumps as US PTAs with Canada and Mexico went into effect. The share of trade with NNTR countries is close to zero throughout and no US imports fall under NNTR after 2005.[13]

Panel (c) plots the share of US NTR imports, identified by the country's modal regime in the 1970s. The UTPP countries' share rises from two percent in 1976 to almost ten percent in 2017, while the former NNTR countries' share rises from zero to nearly 25 percent of US NTR imports. Panel (d) shows the decline in tariffs for these groups. Tariffs for former NNTR countries were substantially above NTR rates and then fell almost to the NTR rates. These drops are discrete and involve a country or set of countries being transitioned to NTR or UTPP. The long-run differences in the tariffs of countries that started with NNTR and NTR rates reflect differences in the composition of goods being imported. In contrast to the large drops in tariffs among countries initially assigned NNTR rates, the NTR rates fell only about five percent. The rise in trade with countries that transitioned from NNTR to NTR is a key component of US import growth that we seek to capture.

[12] The movements in the NNTR rate are caused by composition effects. The statutory NNTR rates have changed little.

[13] In 2022, Belarus and Russia were returned to the NNTR schedule, but our sample ends in 2017.

2.3 Tariff dynamics across vs. within regimes

We now document key features of trade-policy dynamics, with a focus on comparing tariff changes across versus within regimes. We begin by discussing the dynamics of these regimes themselves. Table 2 reports the annual persistence of regimes at the country (j) and country-product (jg) levels, and Table 3 lists the five largest (in terms of the number of affected goods) countries for each type of transition. While all four regimes are highly persistent, there are several transitions away from each of them, indicating that no regime should be viewed as permanent. If our data allowed us to go back further in time, we would observe more transitions. For example, going back to 1949 would add about 20 more country-level transitions from NTR to NNTR and at least two from NNTR to NTR (Yugoslavia and Poland).

We turn now to tariff dynamics. Table 4 reports summary statistics for tariff changes. For the median good, the annual and five-year change is zero although the standard deviation is close to 9 percentage points. The mean tariff change is relatively small, with a mean of -0.17 percentage points over one year and -0.74 percentage points over five years. However, tariff changes are much larger when we condition on regime changes. For example, the mean absolute tariff change during a transition into or out of NNTR is almost 30 percentage points. The large tariff changes that occur during regime changes will likely generate longer transitions, which we allow for in our modeling.

Tariff changes that occur during regime changes are also much more persistent. To illustrate this, we measure the the autocorrelation of tariff changes by estimating the local projection of the horizon-h tariff change on the initial tariff change following Boehm et al. (2023), but allow for different responses for changes across versus within regimes:

$$\begin{aligned}\Delta_h \tau_{jgt} = {} & \beta_h^{\tau,\text{within}} \Delta_0 \tau_{jgt} \mathbb{1}\{\text{regime}_{jgt} = \text{regime}_{jgt-1}\} \\ & + \beta_h^{\tau,\text{across}} \Delta_0 \tau_{jgt} \mathbb{1}\{\text{regime}_{jgt} \neq \text{regime}_{jgt-1}\} + \mathbf{Z}_{jgt} + \delta_{jt} + \delta_{gt} + u_{jgt},\end{aligned} \tag{1}$$

where $\mathbf{Z}_{jgt}$ is a vector that includes controls for pre-trends in tariff changes which we allow to vary for across and within regime changes, that is, we include $\mathbb{1}\{\text{regime}_{jgt} = \text{regime}_{jgt-1}\}\Delta_{-1}\tau_{jgt}$ and $\mathbb{1}\{\text{regime}_{jgt} \neq \text{regime}_{jgt-1}\}\Delta_{-1}\tau_{jgt}$. The set of fixed effects is relatively standard and applies to tariff changes across and within regime. The δ_{jt} fixed effect captures common changes in tariffs to a country and ends up picking up changes across regimes.

The δ_{gt} fixed effect captures tariff-change variation that is common across countries, e.g. tariff cuts during GATT rounds.[14] We estimate (1) for horizons $h = [-5, 10]$. Figure 3(a) plots the paths of $\beta_h^{\tau,\text{within}}$ and $\beta_h^{\tau,\text{across}}$, as well as estimates from a pooled specification without the interaction terms. Tariff changes across regimes are more autocorrelated than within-regime changes. As we will show, trade responds very differently to transitory tariff changes than to unanticipated once and for all reforms. Note the within-regime autocorrelation is very similar to the pooled autocorrelation because the sample is mostly comprised of within-regime changes. The pooled estimates are poor predictors of the dynamics of across-regime tariff changes.

There is substantial variation in tariff dynamics within the NTR regime. Specifically, tariff changes following GATT rounds are very different than other within-NTR changes. The Tokyo round, which was negotiated from 1973–1979, involved gradual tariff cuts from 1980–1987, and the Uruguay round, which was negotiated from 1986–1994, involved gradually tariff cuts from 1995–2000. When we restrict attention to tariff changes during these windows, the five-year tariff change is roughly five times the one-year change, while in other periods the one- and five-year changes are similar (Table 4). The gradual tariff changes during GATT rounds are similar to the tariff "phaseouts" that occur during NTR-to-PTA transitions.[15] In these transitions, the one-year change is -2.7 percentage points, compared to -4.0 percentage points for the five-year change. As we will argue, trade responds very differently to tariff changes with long negotiations and pre-announced tariff phaseouts compared with tariff changes from unanticipated once and for all reform.

2.4 Case studies: China and Vietnam NTR access

We now describe the dynamics of tariffs and imports from China and Vietnam, which are two important transitions from NNTR to NTR. Figure 2(a) shows how tariffs on China and

[14] In the appendix, we show that excluding these fixed effects magnifies the differences in tariff dynamics across versus within regimes.

[15] We plot the path of tariffs separately for Tokyo and Uruguay rounds in Figure A2

Vietnam have changed relative to NTR tariffs. We plot the mean and interquartile range of the ratio of the NTR rate to the tariff on China and Vietnam,

$$\tilde{\tau}_{jgt} = \frac{1+\tau_{gt}^{\text{NTR}}}{1+\tau_{jgt}}. \tag{2}$$

This measure is bounded between zero, during the embargo, and one, when a country gains NTR status. In both countries, we see a long embargo, followed by substantial dispersion during the NNTR regime, which then subsides when each country shifts to NTR.[16] When scaled this way, it is clear the transitions from NNTR are most similar to once and for all tariff changes. However, as we have seen, regime changes are still possible. The reasons for the transitions are unique and closely tied to the state of geopolitical relations between each country and the United States.[17]

The share of US imports accounted for by the two countries rises after conditional NTR and continues until the US-China trade war (Figure 2(b)). China loses US market share during the trade war, but Vietnam gains market share. Since both countries began in embargo, as a whole, all trade growth is driven by the extensive margin. In Figure 2(c), we plot the share of SITC goods with positive trade.

The tariff changes for China and Vietnam during their transitions from NNTR to NTR were much more persistent than most other tariffs changes, including other across-regime changes. To illustrate this, we estimate the autocorrelation of these specific reforms using the specification,

$$\begin{aligned}\Delta_h \tau_{jgt} = {}& \beta_h^{\tau,\text{CN}} \Delta_0 \tau_{jgt} \mathbb{1}\{j = \text{China}, t = 1980\} \\ &+ \beta_h^{\tau,\text{VN}} \Delta_0 \tau_{jgt} \mathbb{1}\{j = \text{Vietnam}, t = 2002\} \\ &+ \beta_h^{\tau,\text{other}} \Delta_0 \tau_{jgt} \mathbb{1}\{\text{other } jt\} + \mathbf{Z}_{jgt} + \delta_{jt} + \delta_{gt} + u_{jgt},\end{aligned} \tag{3}$$

which considers the same vector of pre-trends in tariff changes interacted with the three indicators and fixed effects of (1). Figure 3(b) plots the paths of these estimates. Tariffs

[16]The trade embargo on North Vietnam started in 1964 and was extended to all of Vietnam in 1975. The trade embargo on China was in place from 1950 to 1971.

[17]The transition probabilities are likely time-varying and cannot be inferred from other countries.

on China and Vietnam are extremely persistent compared to tariffs on other countries; statistically, we cannot reject the hypothesis that there are no further changes in the next ten years for Vietnam while there are some minor further increases for China owing to the GATT rounds. Thus, the transitions of China and Vietnam from NNTR to NTR appear to be the closest empirical counterparts to surprise once and for all reforms.

3 Reduced-form evidence on trade dynamics

We now turn to measuring the dynamics of the trade elasticity. Our conceptual object of interest is the path of the cumulative trade elasticity,

$$\varepsilon_h = -\frac{\Delta_h v_{jgt}}{\Delta_h \tau_{jgt}}. \tag{4}$$

We first estimate the dynamics of the trade elasticity to the average tariff change, and show that the responses are small and grow little over time. Next, we condition on regime changes and find larger, more gradual responses. Finally, we estimate trade dynamics for China and Vietnam following their NTR grants and find even larger responses. It is important to emphasize that these estimates are reduced-form responses that reflect agents' beliefs about the stochastic process for tariffs as well as the tariff changes observed in the data. In light of the evidence in section 2, we interpret the differences in trade dynamics for tariff changes within versus across policy regimes as evidence of differences across regimes in their underlying tariff processes.

3.1 Trade dynamics across vs. within regimes

As we have shown, most tariff changes, especially those that occur within policy regimes, are transitory. To account for this, we estimate (4) using differences in differences,

$$\Delta_h v_{jgt} = -\beta_h^v \Delta_h \tau_{jgt} + \mathbf{Z}_{jgt} + \delta_{jt} + \delta_{gt} + u_{jgt}, \tag{5}$$

but using $\Delta_0 \tau_{jgt}$ as an instrument for $\Delta_h \tau_{jgt}$ for $h > 0$, following the approach of Boehm et al. (2023). The vector of controls, $\mathbf{Z}_{jgt}$, includes the preceding tariff change, $\Delta_{-1}\tau_{jgt}$, as well as the preceding import growth, $\Delta_{-1} v_{jgt}$. Note in what follows we allow for the pre-trend controls to vary with the interaction of interest. We include the same set of fixed effects as in section 2, which is standard in the trade-elasticity estimation literature. The country-year fixed effects δ_{jt} control for aggregate shocks such as exchange-rate movements or aggregate cyclical fluctuations. The good-year fixed effects control for changes in multilateral resistance, such as good-specific fluctuations in US demand as well as multilateral policy changes.[18] By taking differences of the dependent variable, constant good-level effects such as trends are eliminated. We provide estimates for horizons $h = [0, 14]$.

We first estimate (5) for the full sample, shown by the black line in Figure 4. The short-run trade elasticity, β_0^v, is 2.14, and the long-run elasticity, β_{14}^v, is about 3.48. These pooled estimates not only restrict trade responses to be the same for within- and across-regime tariff changes, but also impose the same autocorrelation structure for tariffs in the first stage. Next, we estimate a version of (5) that allows for different responses to tariff changes across and within regimes:

$$\begin{aligned} \Delta_h v_{jgt} = &-\beta_h^{v,\text{within}} \Delta_h \tau_{jgt} \{\text{regime}_{jgt} = \text{regime}_{jgt-1}\} \\ &- \beta_h^{v,\text{across}} \Delta_h \tau_{jgt} \{\text{regime}_{jgt} \neq \text{regime}_{jgt-1}\} + \mathbf{Z}_{jgt} + \delta_{jt} + \delta_{gt} + u_{jgt}. \end{aligned} \tag{6}$$

The first stage in this extended specification is (1). In Figure 4, we plot the estimated $\beta_h^{v,\text{within}}$ and $\beta_h^{v,\text{across}}$ coefficients. The estimates for within-regime trade elasticities are lower than the pooled estimates and there is little adjustment over time, with a short-run elasticity of 1.61 and a long-run elasticity of 2.11. The across-regime elasticities are noticeably higher and grow more over time, with a short-run value of 3.87 and a long-run value of 5.86. As with our estimates of tariff dynamics, these results indicate that the pooled estimates are not predictive of how trade responds to large reforms that involve regime changes because within-regime tariff changes constitute the vast majority of the sample.

Finally, we estimate a version of (5) that compares trade dynamics for China and Vietnam

[18] Our results are robust to the inclusion of more restrictive 1- and 2-digit SITC-country-year fixed effects.

to other countries' trade dynamics:

$$\Delta_h v_{jgt} = -\,\beta^{v,\mathrm{CN}}\Delta_h \tau_{jgt}\{j = \text{China}\} - \beta_h^{v,\mathrm{VN}}\Delta_h \tau_{jgt}\{j = \text{Vietnam}\} \\ - \beta_h^{v,\text{others}}\Delta_h \tau_{jgt}\{j = \text{other}\} + \mathbf{Z}_{jgt} + \delta_{jt} + \delta_{gt} + u_{jgt}. \tag{7}$$

Here, the first-stage is similar to (3), except we group China and Vietnam together and do not condition on years. Figure 4 plots the path of $\beta_h^{v,\mathrm{CN+VN}}$. The short-run elasticity is 3.16, higher than in the pooled and within-regime samples, but lower than those in the across-regime sample. The long-run elasticity is 6.91, which is significantly higher than in the across-regime sample. Our interpretation of this result is that other regime changes, which are mostly PTA and UTPP transitions, feature a greater degree of anticipation due to long negotiations and tariff phaseouts. As we show in the next section, anticipation increases measured trade elasticities in the short run but also reduces them in the long run. Again, this suggests that the cases of China and Vietnam are better-suited for estimating the how trade responds to surprise once and for all reforms.

3.2 Case studies: China and Vietnam NTR access

We continue our analysis of trade dynamics by homing in on the effects of the transitions from NNTR to NTR that China and Vietnam underwent in 1980 and 2002. Leveraging the heterogeneity in tariff changes across goods, we estimate the dynamics of trade following these reforms, as well as in the years leading up to them. We follow Alessandria et al. (2025), and estimate the relationship between trade flows and the difference between the NNTR tariff and the NTR tariff, which we refer to as the *NNTR gap*:

$$X_g = \log(1 + \tau_g^{\mathrm{NNTR}} - \tau_g^{\mathrm{NTR}}), \tag{8}$$

where the tariffs are measured as the average over the 1999 scheduled NNTR and NTR rates at the HS-8 tariff line from Feenstra et al. (2002). The NNTR gap has been widely used in the literature as a measure of tariff risk, as it represents the tariff hike on good g that would occur if a country lost NTR status and moved to NNTR. However, it is also the tariff reduction that the good received when the country originally gained NTR status. This dual meaning complicates the interpretation of the estimates discussed below, but does not preclude them from being used as inputs to a quantitative analysis, which is the strategy we ultimately follow. We plot the distribution of the NNTR gap in Figure 5(a).

We use an event-study design to estimate the responses of imports from each country to NTR access:

$$\begin{aligned} v_{jgt} = & \sum_{t'=1974}^{2008} \beta_t^{v,\mathrm{CN}} \mathbb{1}\{t = t'\} \mathbb{1}\{j = \text{China}\} X_g \\ & + \sum_{t'=1994}^{2017} \beta_t^{v,\mathrm{VN}} \mathbb{1}\{t = t'\} \mathbb{1}\{j = \text{Vietnam}\} X_g + \delta_{jt} + \delta_{jg} + \delta_{gt} + u_{jgt}. \end{aligned} \tag{9}$$

The dependent variable is the log import value of good g from country j in year t. Our coefficients of interest, β_t^j, measure the elasticity of trade with respect to the NNTR gap, relative to the other countries with NTR status, and relative to a country-specific base year. We focus on a window from 1974 to 2008 for China to avoid effects from the Great Recession, while for Vietnam we focus on the full sample ending in 2017.[19] As before, we include country-year fixed effects (δ_{jt}) to capture aggregate shocks to exporting countries, country-good fixed effects (δ_{jg}) to capture the average level of exports and time-invariant bilateral trade barriers, and good-year fixed effects (δ_{gt}) to capture good-level US demand shocks and good-specific trade barriers common to all exporters.[20] The elasticities are normalized to zero in 2008 for China and 2017 for Vietnam.

Figure 5(b) shows the estimates of β_t^{CN} and β_t^{VN}. There are four key observations. First, both countries have similar long-run NNTR-gap elasticities that are substantially higher than the average long-run trade elasticities we previously estimated. From trough-to-peak, the elasticity is about 11 for Vietnam and China. Second, the Chinese response, on impact, was about half as big as the Vietnamese response. Third, the two countries have different trade dynamics in the lead up to the tariff change. Trade with Vietnam in high-gap goods was growing faster than low-gap goods while, China, high- and low-gap goods were growing

[19] As a major US trade partner at that time, the Great Recession had a significant impact on US imports from China. Trade with Vietnam was still in the very early stages at this time, and was not significantly impacted.

[20] This fixed-effects structure has been used frequently in the literature on China's NTR access, including by Pierce and Schott (2016) and Handley and Limão (2017).

at about the same pace. Fourth, Chinese exports grew much slower than exports from Vietnam; it took more than twelve years to reach China's long-run level, whereas Vietnam's adjustment was more than 90 percent complete after four years (note that this is precisely the length of the period in which the Vietnamese reform remained conditional).

While the results of the event study specification of Figure 5 appear to yield much larger reduced-form long-run elasticities than those under the local projection approach of (7), three considerations reconcile them. First, our event study specification adds the anticipation to the long-run response by defining the long-run response with respect to the trough, which in both cases occurs fours years before the NTR access. This is especially relevant in the case of Vietnam, for which almost one third of the long-run response occurs prior to the NTR access. Second, the local projections of (7) average over all years including tariff changes that are much less persistent and similar to the "pooled" in Figure 3(a). Thus, it is not surprising that the reduced-form elasticities when averaging over all years are smaller. Third, when we restrict the local projections to the year of the NTR access, the 15-year elasticity are 20 and 25 percent larger for China and Vietnam, respectively. In effect, given the absence of significant anticipation for China, the local projections following the NTR access look almost identical to the gap-elasticities under the event study specification. Figure A4 illustrates these points by plotting the results of the event study when the reference period is the year before the NTR access (thus eliminating the role of anticipation), together with the local projections of (7) and those when focusing on the NTR access only; as well as the autocorrelation of tariff changes as in (3) in comparison to those that exclude the NTR access years.

Our results on tariff and trade dynamics within and across regimes and for China and Vietnam are robust to a range of sample construction and empirical methodologies discussed in Appendix B.2. Specifically, we find alternative aggregations of good-level regimes and tariffs yields similar trade dynamics. Likewise, the differences in trade elasticities across regime, and for China and Vietnam, are similar using other measures of trade elasticities such as the cumulative change in trade over the cumulative change in tariffs or an error-correction specification. Alternative measures of tariff risk, controls for pre-trends or country samples yields similar outcomes.

The large response in trade to large, persistent country- and good-specific changes in tariffs is consistent with the idea that trade responds more to more persistent tariff change. However, the fact that the tariff change was ex-post more persistent does not tell us about the expected path of tariffs at each moment in time, which depended on complex geopolitical risks. We will use a model to sort out the source of these differences.

4 Model

To illustrate how policy dynamics complicate the measurement of trade dynamics, we use a model in which firms make forward-looking export participation decisions to show how trade responds to non-canonical reforms. Our model is a partial-equilibrium version of Alessandria et al. (2021) in which a continuum of heterogeneous firms choose to start and stop exporting in the face of idiosyncratic technology shocks and aggregate shocks to trade policy. The model has two key ingredients that generate gradual adjustment to policy: sunk exporting costs and export capacity that grows stochastically with a firm's tenure as an exporter. The first ingredient makes firms consider the entire future path of trade policy, not just its current state (Das et al., 2007; Alessandria and Choi, 2007). The second ingredient generates exporter life-cycle dynamics (Alessandria et al., 2021; Fitzgerald et al., 2023; Steinberg, 2023), which captures the behavior of marginal exporters and non-exporters. The model allows for potentially large differences between short- and long-run trade elasticities.

Production and demand. Firms use labor, ℓ_t, to produce output according to constant returns to scale technologies:

$$y_t = z_t \ell_t. \tag{10}$$

A firm's productivity, z_t, evolves over time according to an AR(1) process in logs:

$$\log z_t = \rho_z \log z_{t-1} + \sigma_z \varepsilon_t, \tag{11}$$

where ε_t is i.i.d. across firms and time. Firms produce differentiated goods and operate in monopolistically competitive markets. Foreign demand for a firm's good is a downward-sloping function of the price the firm charges, p_t, and the import tariff, τ_t:

$$d_t(p_t, \tau_t) = (p_t \tau_t)^{-\theta}, \tag{12}$$

where θ is the price elasticity of demand.

Trade costs. There are three types of trade costs: import tariffs; iceberg-type variable trade costs; and fixed trade costs. Import tariffs apply to all firms equally and are, at this stage, allowed to follow an arbitrary stochastic process. We index model objects by the time period, t, which is a parsimonious way to encode firms' expectations about future tariffs. Variable trade costs can take three values, $\infty > \xi_H > \xi_L$, that encode a firm's status as an exporter. A firm with $\xi = \infty$ is a non-exporter, a firm with ξ_H is a high-cost exporter, and a firm with $\xi = \xi_L$ is a low-cost exporter. To begin exporting, a non-exporter must pay a fixed cost f_0. To continue exporting, a low- or high-cost exporter must pay a fixed cost f_1. Thus, the fixed cost of exporting can be expressed as a function, $f(\xi)$, of the variable cost, where $f(\infty) = f_0$ and $f(\xi_H) = f(\xi_L) = f_1$.

Exporter life cycles. Firms are born as non-exporters with $\xi = \infty$. In each period, firms die with probability $1 - \delta(z) = \max\left[0, \min\left(e^{-\delta_0 z} + \delta_1, 1\right)\right]$, in which case they are replaced by newborn firms with productivities drawn from the ergodic distribution associated with the process (11). When a non-exporter chooses to enter the export market, its variable trade cost falls to ξ_H in the next period. When a high-cost exporter chooses to continue exporting, its variable exporting cost falls to ξ_L in the next period with probability $1 - \rho_\xi$. Symmetrically, when a low-cost exporter chooses to continue exporting, the probability of switching to ξ_H is also ρ_ξ.

Firm's problem. The firm's state variables are its productivity, z, its variable trade cost, ξ, and its tariff, τ. The firm's problem has a static component and a dynamic component. The static problem entails choosing a price to maximize the profits from exporting in the current period:

$$\pi_t(z_t, \xi_t, \tau_t) = \begin{cases} \max_p \left\{ p d_t(p, \tau_t) - w_t \frac{\xi_t d_t(p, \tau_t)}{z_t} \right\} & \text{if } \xi \in \{\xi_L, \xi_H\} \\ 0 & \text{if } \xi = \infty \end{cases} \tag{13}$$

The dynamic problem entails deciding whether to export in the next period. The value of a firm that chooses to export in $t+1$ is

$$V_t^1(z_t, \xi_t, \tau_t) = -f(\xi_t) + \frac{\delta(z)}{1+r} \mathbb{E}_{z,\xi,\tau} \left[V_{t+1}(z_{t+1}, \xi_{t+1}, \tau_{t+1}) \right], \tag{14}$$

and the value of a firm that chooses not to export is

$$V_t^0(z_t, \xi_t, \tau_t) = \frac{\delta(z)}{1+r} \mathbb{E}_{z,\tau} \left[V_{t+1}(z_{t+1}, \infty, \tau_{t+1}) \right], \tag{15}$$

where

$$V_t(z_t, \xi_t, \tau_t) = \pi_t(z_t, \xi_t, \tau_t) + \max\{V_t^0(z_t, \xi_t, \tau_t), V_t^1(z_t, \xi_t, \tau_t)\}. \tag{16}$$

The solution to the export participation problem is characterized by a break-even level of productivity, $\bar{z}_t(\xi)$, such that $V_t^0(\bar{z}_t(\xi), \xi, \tau_t) = V_t^1(\bar{z}_t(\xi), \xi, \tau_t)$. Importantly, $\bar{z}_t(\xi)$ depends on firms' expectations about future tariffs, not just the current tariff rate. We will consider several processes for tariffs.

Aggregation. Aggregate exports are given by

$$X_t = \sum_{\xi \in \{\xi_L, \xi_H\}} \int_z p_t(z, \xi, \tau_t) d_t(p_t(z, \xi, \tau_t), \tau_t) \, \mathrm{d}\varphi_t(z, \xi), \tag{17}$$

where $\varphi_t(z, \xi)$ denotes the joint distribution of firms indexed by their productivity and variable trade cost.

Parameters. For the numerical examples in the next section, we follow Alessandria et al. (2021). We assign several values directly from that paper: $\theta = 4$, $\rho_z = 0.65$, $\sigma_z = 1.32$, $\rho_\xi = 0.92$, $\delta_0 = 21$, $\delta_1 = 0.02$. We set the interest rate, r, to 4 percent. We normalize the low variable trade cost, ξ_L, to one and choose values for the fixed costs, f_0 and f_1, and the high variable trade cost, ξ_H, so that the model's free-trade steady state reproduces three moments from Alessandria et al. (2021): (i) 22.3 percent of firms export; (ii) 17 percent of exporters exit annually; and (iii) new exporters sell half as much as incumbent exporters on average. We discuss our parameterization of the quantitative model for the case studies of China and Vietnam in section 6.1.

5 Illustrative examples

We now use the model to illustrate some issues in estimating trade elasticity dynamics. We conduct three kinds of hypothetical reforms in the model: unanticipated once and for all reforms, stochastic reforms, and changes to beliefs about the policy's stochastic process. Each reform highlights a different issue. We also consider how the initial conditions and the window of observation influence the trade elasticity.

When tariffs follow a Markov process, the anticipation of future tariff changes and persistence of past changes simultaneously affect trade volumes. Consider a tariff process with two states, $\tau_L = 0$ and $\tau_H = 30$ percent, and a symmetric probability of switching states, $1-\omega$. We conduct the same reform—switching from the high-tariff state to the low-tariff state—for values of $\omega \in [0.5, 1]$. When $\omega = 0.5$, tariffs are i.i.d. across time, and the current tariff contains no information about future tariffs. When $\omega = 1$, tariffs are constant and the tariff cut is permanent. In between these extremes, tariffs are persistent but firms know there is a chance of switching states. In all versions of this experiment, the economy begins in the risky steady state with high tariffs, switches to the low-tariff state in period $t = 0$, and remains there forever.

It is useful to summarize the future path of tariffs by the *mean discounted expected tariff*,

$$\bar{\tau}^E_{t+1} = (1-\beta)^{-1} E_t \sum_{s=t+1}^{\infty} \beta^{s-t-1} \tau_s. \tag{18}$$

In Figure 6(b), we show how the path of $\bar{\tau}^E_{t+1}$ changes as we vary the persistence of the tariff.

5.1 Canonical reform: surprise once and for all

As a benchmark, we consider the canonical trade reform found in virtually all quantitative and theoretical trade studies: an unanticipated, immediate, and permanent reform. This is the tariff process with $\omega \to 1$. Figure 6 shows the results of this reform in our model (line $\omega = 1.0$). In period zero, the distribution of firms over trade costs and trade participation is predetermined, so trade only responds on the intensive margin, and the short-run trade elasticity is equal to the CES demand elasticity: $\varepsilon_0 = \theta$. We refer to ε_0 as the *short-run* *canonical trade elasticity*. In period one, export participation begins to rise, and converges to its new long-run value, which is about four log points higher than before the reform (Figure 6(c)). Trade follows a similar trajectory, and the long-run trade elasticity, ε_∞, is about 17.5 (Figure 6(e)). We refer to ε_∞ as the *long-run canonical trade elasticity*. It captures the long-run change in trade relative to the long-run change in tariffs.

5.2 Stochastic reforms

When ω is not one, the mean discounted expected tariff differs from the current tariff, and the measured trade elasticities will deviate from their canonical values. When tariffs are less persistent, exporting is more valuable in the high-tariff state because a tariff cut is more likely to occur in the future. This increases export participation and aggregate exports in the initial period. To highlight this effect, we plot changes in export participation and trade in Figures 6(c) and 6(d), relative to the pre-reform period in the model in which tariffs are believed to be permanent ($\omega \to 1$). Conversely, when tariffs are less persistent, exporting is less valuable in the low-tariff state because a tariff increase is more likely to occur in the future. This depresses export participation and trade after the tariff cut.

When tariffs are less persistent, the pre-reform level of trade is higher and post-reform trade growth is lower. Both of these effects push the measured long-run trade elasticity below the canonical long-run elasticity. The i.i.d. case is particularly instructive. In this case, the current tariff state does not influence the mean discounted expected tariff and the value of exporting in the future is constant. Consequently, export participation does not react to tariff shocks at all, and the trade elasticity at all horizons is equal to the demand elasticity, θ.

These experiments illustrate why within-regime tariff changes, which are transitory, have lower long-run trade elasticities than across-regime tariff changes, which are more persistent. Trade is higher before within-regime changes occur because trade participants understand that these changes are likely to occur, and does not grow as much following these changes because they are not expected to last very long.

5.3 Shocks to expectations

In the previous section, when tariffs changed, the mean discounted expected tariff could also change. Here, we consider changes in the mean discounted expected tariff that may not coincide with a change in tariffs. To do so, we shift from a two- to a four-state Markov process that allows for different transition probabilities from the same tariff levels. States HP (high permanent) and HT (high temporary) have a tariff of 30 percent and states LT (low temporary) and LP (low permanent) have a zero tariff. The persistence of states HP and LP are close to one while states HT and LT are less persistent. Transitioning between HT and LT is much more likely than transitioning from HT or LT to either HP or LP. With this structure, we can separate changes in the mean discounted expected tariff from changes in the current tariff. Our four-state transition matrix is

$$\begin{array}{c} \begin{array}{cccc} \quad HP & \quad\quad HT & \quad\quad LT & \quad\quad LP \end{array} \\ \Omega = \begin{bmatrix} 1-3\epsilon & \epsilon & \epsilon & \epsilon \\ \epsilon & \omega-\epsilon & 1-\epsilon-\omega & \epsilon \\ \epsilon & 1-\epsilon-\omega & \omega-\epsilon & \epsilon \\ \epsilon & \epsilon & \epsilon & 1-3\epsilon \end{bmatrix}. \end{array} \tag{19}$$

In Figure 7, we plot the response to four reforms. In all reforms, we set $\omega = 0.8$ and $\epsilon \approx 0$. We start in HP with high tariffs that are expected to be permanent. As before, we assume we have been in this state for many years. Our first reform is a transition from HP to LP at $t = 1$. This is the canonical unanticipated once and for all reform we have already discussed.

Our second reform is a transition from HP to HT. The tariff remains constant, but the probability of transitioning to a low-tariff state increases. The mean discounted expected tariff falls to 15.8 percent. We assume the shift occurs in period $t = -1$. Trade jumps in period zero and grows persistently to about 1.58 log points above the initial level. Without a change in the observed tariff, we can not construct an h-on-1 trade elasticity.

Our third reform is a transition from HP to LT. This removes the tariff, but, as the change is less persistent than our canonical reform, the mean discounted expected tariff only falls to 14.2 percentage points. We find a large and gradual increase in trade. The short-run trade elasticity is four and the long-run is 13.5, reflecting trade growth of 3.54 log points. The much larger long-run response, compared with our second reform, which reduced the expected tariff by 1.6 percentage points less, arises from the strong effect of a lower current tariff.

Finally, we consider a transition from HP to HT followed by a transition to LT. This uncertain two-step reform has the same long-run change in trade as our third reform, but our measured h-on-1 trade elasticity is larger in the first period, about 7.5. The larger response to the change in tariff in $t = 0$ occurs because some firms have made investments to increase market access in $t = -1$. Thus, trade is responding to the past and current tariff shocks. Our model can sort out the difference in prior expectations of tariff changes, which results in different trade response on impact, like the ones we found in Section 3.2.

These experiments illustrate why across-regime tariff changes tend to have higher short-run trade elasticities than within-regime changes. Across-regime changes often involve changes in expectations that precede changes in tariffs. For example, free-trade agreements typically involve long negotiation periods, are formally announced before being actually implemented, and phased in gradually over several years.[21] Of course, the case of China's 1980 NTR access, which has a similar short-run trade elasticity to within-regime tariff changes, indicates that isn't always true.

5.4 Initial conditions

In our prior analysis, the economy was in a steady state at the time of reform. In China and Vietnam the reform takes place while the economies are transitioning from the end of the embargo. To capture the role of initial conditions, or being out of steady state, we start from an embargo and then consider a surprise once and for all change in tariffs. To align with our empirical approach, we do this for two goods that permanently go to different tariff levels of zero and 30 percent.

Figure 8(a) depicts trade dynamics starting from the period following the embargo's end.

[21] We show how phased-in reforms can dramatically increase short-run trade elasticities in the appendix. Khan and Khederlarian (2021) document evidence for these effects in the case of NAFTA.

We also estimate the trade elasticity from these two goods using a panel approach,

$$v_{gt} = \sum_{t'=1}^{25} \beta_t \mathbb{1}\{t = t'\} \log(\tau_{gt}) + \delta_t + u_{gt}. \tag{20}$$

Figure 8(b) shows that the good that goes to zero tariffs starts with more trade and exporters than the good that goes to 30 percent tariffs and the gap between the two goods grows over time. On impact, we find a very large trade elasticity from comparing the differences in trade by tariffs of -13.8. Over time, we see that this elasticity gradually falls to about -17.5, which is about the long-run trade elasticity we found in the model for an surprise once and for all liberalization. Importantly, five years after lifting the embargo we have only closed about two-thirds of the gap between the first year differences and the long-run differences across goods.

We have seen that the faster export growth of low tariff goods following the lifting of the embargo is a feature of the trade dynamics in China and Vietnam prior to gaining NTR tariffs. This experiment makes clear that both initial conditions and the length of the observation window may bias empirical estimates of the trade elasticity downward.

6 Quantitative analysis

We take our model to the data—which are contaminated by anticipation of and uncertainty about tariff changes and the gradual effects of past reforms—and use it to measure how trade would respond in the absence of these expectational biases and initial conditions. Our strategy is to discipline the model's technological parameters and the stochastic process for trade policy with empirical evidence from two trade reforms: the United States granting NTR status to China in 1980 and Vietnam in 2002. We use the model to simulate surprise once and for all reforms that generate estimates the dynamics of the trade elasticity.

The Chinese and Vietnamese reforms have three key advantages. First, the distribution of tariffs and uncertainty are well-understood owing to the structure of trade policy. Second, given the move from the embargo we know the initial conditions in each country and in each sector. Third, the countries receiving the reforms become very large trading partners.

6.1 Environment and calibration

In what follows, we describe the model of firms in China and Vietnam. The structure of the economy is the same in the two countries, but parameter values and timing differ. For clarity, we omit country labels. To exploit the heterogeneity in trade dynamics across goods documented above, we consider a multi-good version of the model in section 4. There are G goods, which correspond to the 5-digit SITC goods in our data. In each good $g = 1, \ldots, G$, a continuum of firms sell differentiated varieties. The technological parameters—the productivity process and the variable and fixed costs of exporting—are the same across goods.[22]

Tariffs, $\tau_{gt}(s)$, differ across goods and trade-policy regimes, $s \in \{\text{NNTR}, \text{NTR}\}$. We take the tariff rates directly from the data. The NTR rates change as new rates are negotiated. We allow for this within-regime variation, and firms have perfect foresight over the path of NTR tariffs. A country's tariff regime, which is an aggregate state, follows a time-varying Markov process with transition probabilities

$$\Omega_t = \begin{bmatrix} \omega_t^{\text{NNTR}} & 1-\omega_t^{\text{NNTR}} \\ 1-\omega_t^{\text{NTR}} & \omega_t^{\text{NTR}} \end{bmatrix}, \tag{21}$$

where ω_t^s is the persistence of state s in year t. Our aim is to estimate a path $\{\omega_t^{\text{NNTR}}, \omega_t^{\text{NTR}}\}_{t=0}^{\infty}$ that is consistent with the dynamics of trade flows observed for each country.

Our calibration strategy builds on Alessandria et al. (2025). First, we set several parameters externally. We use the same values for the interest rate, r, the firm survival function parameters, δ_0 and δ_1, and the iceberg-cost transition probability, ρ_ξ, from section 5. However, we set the demand elasticity, θ, to the average estimate in Soderbery (2018) for US imports from China, which is 3.13. Soderbery's estimate comes from price and quantity variation and, thus, is closely related to our demand elasticity parameter.[23] We use this estimate rather than the reduced-form short-run trade elasticity, which depends on changes

[22] We have explored heterogeneity across goods in these parameters in Alessandria et al. (2025). While this works well for China, there are too few firms in many sectors in Vietnam to reliably estimate exporter-dynamics.

[23] We also estimated θ directly in our system and recovered a similar estimate.

in trade policy expectations (section 5.3).

Second, we calibrate the other technological parameters, σ_z, f_0, f_1, and ξ_H, to match four moments computed from firm-level panel data: the export participation rate, the export exit rate, the incumbent size premium (the average sales of incumbent exporters relative to the average for new exporters), and the coefficient of variation of log exports. In both the model and the data, we measure these moments many years after NTR status has been granted (2004–2007 for China and 2010–2014 for Vietnam).[24] The entry cost, f_0, is mainly identified by the export participation rate. The continuation cost, f_1, is mainly identified by the exit rate. The initial iceberg trade cost, ξ_H, is mainly identified by the incumbent premium. Productivity dispersion, σ_z, is mainly identified by the coefficient of variation of log exports. Table 5 lists the target moments and calibrated values of the technological parameters for each country.

Third, we calibrate the trade-policy persistence parameters, ω_t^{NNTR} and ω_t^{NTR}, to match the NNTR-gap elasticity dynamics described in section 3.2. ω_t^{NNTR} is identified by the dynamics during the period before the country was granted NTR status. When the NNTR-gap elasticity grows more negative during this period, i.e., exports of high-gap goods grow faster than exports of low-gap goods, it indicates an increase in the likelihood of gaining NTR status. ω^{NTR} is identified by the NNTR-gap elasticity dynamics after the NTR grant. A decrease in this elasticity during this period indicates an increase in the likelihood of retaining NTR status. The main difference between our strategy and that of Alessandria et al. (2025) is that we allow ω_t^{NNTR} to vary over time, allowing us to capture changes in the probability of the NTR grant occurring in advance of the policy. Figure 9(a) shows the fit of the model against these target moments.[25]

Fourth, we start the model from the embargo so that there are no exports to the United States from either country. We assume that firms do not expect the embargo to be lifted. Changes in the probabilities of gaining and losing NTR status are believed to be permanent. Firms are surprised when these probabilities change.

Figure 9(b) plots the calibrated regime transition probabilities, $1 - \omega_t^s$. For China, the

[24]The appendix contains more details about the firm-level datasets.

[25]We HP-filter the NNTR-gap coefficients to smooth out temporary spikes. The calibrated model exactly matches this smoothed series.

probability of gaining NTR status in the 1970s was low and did not change materially. When China gained NTR status in 1980, the pattern of trade growth is consistent with the view that this status was unlikely to be maintained, and only started to become credible in the mid-1980s. Alessandria et al. (2025) discuss these considerations in detail. For Vietnam, the initial probability of gaining NTR status was similar to China's, but during the late 1990s and early 2000s this probability more than tripled. The model interprets the fast growth of high-gap exports relative to low-gap exports in advance of the Vietnamese NTR grant as evidence of an increase in the expectation that this reform would occur. After Vietnam gained NTR status in 2002, the fast growth in trade is consistent with the probability of losing this status falling faster than in China. As the long-run changes in the NNTR-gap elasticity are similar in the two cases, we estimate similar initial and final transition probabilities.

While we estimate our model to match the trade response in a limited window, our model yields a forecast of the trade beyond that window. For both countries, we see that the trade elasticity continues to increase by one percentage point in the five years following our estimation period. The continued trade growth so late in the transition from the tariff cut reflects the time it takes for the reform to become credible.

6.2 Trade response to canonical reforms

We now use the calibrated model to measure the trade elasticity following a canonical tariff reform that features neither anticipation nor uncertainty, is not affected by initial conditions, and for which we can observe the entire transition. For each country, we solve for the steady state good level imports with NNTR tariffs that are permanent. We then introduce a surprise once and for all reduction to the 1999 good-level NTR tariffs for each country, and solve for the transition to the new steady state. In Table 6 we report the short-run, 10-year, and long-run trade elasticity for China and Vietnam for the canonical reform along with our estimates from the data using the h-on-1 and an error correction specification described in the appendix. Figure 9(c) compares the NNTR-gap elasticity dynamics in the canonical NTR reforms to the baseline dynamics. The transition dynamics to the canonical reform are quite similar for Vietnam and China, with very small differences owing to the different parameterization of firm dynamics.

There are three key differences between the trade elasticity from the canonical reform and the estimated model with stochastic tariffs. First, the pre-NTR gap elasticities in the baseline model (which range from about -8 to -11) are smaller than the gap elasticities in the NNTR steady state (about -14). This is largely because firms in the baseline model have a greater incentive to begin exporting in advance of gaining NTR status, particularly firms that produce high-gap goods which have the most to gain from the reform. Moreover, in the data, we are not starting from a steady state but on the transition from the embargo, which we have shown biases our long-run estimates downward.

Second, the gap elasticities converge more slowly after the baseline NTR reforms than after the canonical reforms. This is because the probability of keeping NTR status is initially low in both countries, and rises gradually over time. It is particularly noticeable in the case of China, where the gap elasticity stalls for several years after the baseline NTR reform and only starts to rise again several years later.[26] This implies that trade should continue to grow beyond our estimation window for both countries. In our canonical reform there is very little continued growth outside our estimation period.

Third, in the case of Vietnam, the gap elasticity begins to rise several years before the baseline NTR reform, whereas there is no change in advance of the canonical reform. This is an anticipatory effect driven by the increase in the probability of Vietnam gaining NTR status during 1999–2001 shown in Figure 9(b); note there is no anticipatory growth in the case of China, where the likelihood of gaining NTR status was more or less constant throughout the 1970s. This anticipation effect lowers the long-run trade elasticity we recover from Vietnam. It also leads us to estimate a higher short-run trade elasticity with the reform of 4.44 compared to 3.17 in the model, which is the demand elasticity. With China, the combination of no anticipation to the NTR reform and subsequent shocks to tariffs lead us to estimate a short-run trade elasticity that is close to half of the true short-run elasticity.

Finally, we calculate the speed of adjustment of the canonical reform. Specifically, along the transition, we measure the gap between the final level of trade and the current level of

[26] See Alessandria et al. (2025) for a detailed discussion of the geopolitical background of the "stalled" period.

trade in each period relative to the long-run change,

$$\lambda_h = \frac{\log V_\infty - \log V_{t+h}}{\log V_\infty - \log V_{t-1}}, \tag{22}$$

where V_∞ is the long-run level of trade in the new steady state with no uncertainty. We plot this gap in Figure 9(d) and fit a geometric function of the gap on the latter half of the sample. As we have seen, the model predicts very strong growth in the first two years owing to the intensive margin and delayed extensive margin response from the one period export lag. 22–23 percent of the long-run change occurs in the first year (the intensive-margin effect), and another 40–45 percent in the second year (the one-period delay on the extensive margin). After that initial burst, we find that half the distance to the long-run closes about every five years.

7 Conclusions

We estimate the long-run trade elasticity to a surprise once and for all tariff change is about 14 times the tariff change. This is about five times what one finds when applying standard empirical methods to US tariff and trade changes and about three times the trade elasticity commonly used in quantitative work (Simonovska and Waugh, 2014; Caliendo and Parro, 2014; Sposi et al., 2021). This large response is hard to find in the data, owing to the stochastic properties of trade policy and forward-looking decisions of firms related to anticipation, uncertainty, and initial conditions that are not captured in existing empirical methods. We find a short-run elasticity of about three, and that half of the gap between the current level and long-run level of trade is closed in five years.

Our estimated elasticities are recovered from a model-based analysis of trade and trade policy dynamics that leverages key elements of the structure of US trade policy and its relationships with China and Vietnam. These reforms were incredibly large, moving countries from embargo to NTR rates, and characterized by a high probability of reversal, owing to the unique geopolitical relations between each of these countries and the United States. Importantly, they are easily modeled as a two-state Markov transition matrix with a common time-varying transition probability across goods but different known good-specific tariffs and allow us to easily control for initial conditions owing to the embargoes in place. Our approach can be extended to consider alternative regime transitions.

Our results imply that there is a mismatch between empirical estimates of trade elasticities and the structural parameters used in quantitative analyses. The vast majority of the tariff variation in the data is highly transitory, so the long-run elasticities identified using this variation are much lower than the elasticities that are appropriate to use in models of surprise once and for all trade reforms. Progress needs to be made on two fronts. First, the empirical literature should develop methods to estimate trade elasticities that control for anticipation and uncertainty. Our approach uses a structural dynamic model and we suspect integrating empirics with theory will help. Second, the quantitative literature should develop calibration strategies that explicitly account for the stochastic processes that tariffs follow in the data.

If trade policy aims to foster deeper integration, then it is important to find ways to make trade policy more credible. One possibility would be to adopt Milton Friedman's suggestion of a constitutional amendment that sets import tariffs to zero (Friedman and Friedman, 1980), much like the proscription on states taxing interstate commerce or imposing export taxes.[27] Of course, perhaps it's not possible to credibly commit to a permanent trade reform, in which case one needs to explicitly incorporate trade policy expectations into their analysis.

[27]The commerce clause (Article 1, Section 8, Clause 3) gives Congress the power to regulate commerce between states. The export clause (Article 1, Section 9, Clause 5) states "No tax or duty shall be laid on articles exported from any state."

References

Alaouze, Chris M., John S. Marsden, and John Zeitsch, "Estimates of the Elasticity of Substitution Between Imported and Domestically Produced Commodities at the Four Digit ASIC Level," Centre of Policy Studies/IMPACT Centre Working Papers o-11, Victoria University, Centre of Policy Studies/IMPACT Centre July 1977.

Alessandria, George and Horag Choi, "Do Sunk Costs of Exporting Matter for Net Export Dynamics?," *The Quarterly Journal of Economics*, 2007, *122(1)*, 289–336.

_ and _ , "The Dynamics of the U.S. Trade Balance and Real Exchange Rate: The J Curve and Trade Costs?," *Journal of International Economics*, 2021, *132.* Article 103511.

_ and Oscar Avila, "Trade Integration, Industry Reallocation, and Welfare in Colombia," *IMF Economic Review*, September 2023, *71* (3), 649–687.

_ , Horag Choi, and Kim J. Ruhl, "Trade Adjustment Dynamics and the Welfare Gains from Trade," *Journal of International Economics*, 2021, *131*, 1034–1058.

_ , Shafaat Y. Khan, Armen Khederlarian, Kim J. Ruhl, and Joseph B. Steinberg, "Trade War and Peace: U.S.-China Trade and Tariff Risk from 2015–2050," Working Paper 32150, National Bureau of Economic Research February 2024.

_ , Shafaat Yar Khan, Armen Khederlarian, Kim J. Ruhl, and Joseph B. Steinberg, "Trade-Policy Dynamics: Evidence from 60 Years of U.S.-China Trade," *Journal of Political Economy*, 2025.

Baier, Scott L. and Jeffrey H. Bergstrand, "Do Free Trade Agreements Actually Increase Members International Trade?," *Journal of International Economics*, 2007, *71* (1), 72–95.

_ , _ , and Michael Feng, "Economic Integration Agreements and the Margins of International Trade," *Journal of International Economics*, 2014, *93* (2), 339–350.

Baldwin, Richard, "Hysteresis in Trade," *MIT mimeo prepared for 1986 NBER Summer Institute, April 1986*, 1986.

_ and Paul Krugman, "Persistent Trade Effects of Large Exchange Rate Shocks," *The Quarterly Journal of Economics*, 1989, *104* (4), 635–654.

Besedes, Tibor, Tristian Kohl, and James Lake, "Phase Out Tariffs, Phase In Trade?," *Journal of International Economics*, 2020, *127.* Article 103385.

Boehm, Christoph E., Andrei A. Levchenko, and Nitya Pandalai-Nayar, "The Long and Short (Run) of Trade Elasticities," *American Economic Review*, 2023, *113(4)*, 861–905.

Broda, Christian and David E. Weinstein, "Globalization and the Gains From Variety*," *The Quarterly Journal of Economics*, 05 2006, *121* (2), 541–585.

Caliendo, Lorenzo and Fernando Parro, "Estimates of the Trade and Welfare Effects of NAFTA," *The Review of Economic Studies*, 11 2014, *82* (1), 1–44.

CRS Report for Congress, "Normal-Trade-Relations (Most-Favored-Nation) Policy of the United States," Technical Report RL31558, Congressional Research Service 2005. Retrieved from https://www.everycrsreport.com/files/20051215_RL31558_8600c4dcebded883cd8887acd107db3bcaa49f84.pdf.

Das, Sanghamitra, Mark J. Roberts, and James R. Tybout, "Market Entry Costs, Producer Heterogeneity, and Export Dynamics," *Econometrica*, 2007, *75* (3), 837–873.

Feenstra, Robert C., "U.S. Imports, 1972-1994: Data and Concordances," *NBER Research Working Paper No. 5515*, 1996.

_ , John Romalis, and Peter K. Schott, "U.S. Imports, Exports, and Tariff Data, 1989–2001," *NBER Working Paper No. 9387*, 2002.

Fitzgerald, Doireann, Stefanie Haller, and Yaniv Yedid-Levi, "How Exporters Grow," *The Review of Economic Studies*, 07 2023, *91* (4), 2276–2306.

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

Gallaway, Michael P., Christine A. McDaniel, and Sandra A. Rivera, "Short-run and Long-run Industry level Estimates of U.S. Armington Elasticities," *The North American Journal of Economics and Finance*, 2003, *14(1)*, 49–68.

Handley, Kyle and Nuno Limão, "Trade Investment under Policy Uncertainty: Theory and Firm Evidence," *American Economic Journal: Policy*, 2015, *7* (4), 189–222.

_ and _ , "Policy Uncertainty, Trade, and Welfare: Theory and Evidence for China and the United States," *American Economic Review*, 2017, *107(9)*, 2731–2783.

_ and _ , "Trade Policy Uncertainty," *Annual Review of Economics*, 2022, *14* (1), 363–395.

Head, Keith and Thierry Mayer, "Chapter 3 - Gravity Equations: Workhorse,Toolkit, and Cookbook," in Gita Gopinath, Elhanan Helpman, and Kenneth Rogoff, eds., *Handbook of International Economics*, Vol. 4 of *Handbook of International Economics*, Elsevier, 2014, pp. 131–195.

Hillberry, Russell and David Hummels, "Trade Elasticity Parameters for a Computable General Equilibrium Model," in Peter B. Dixon and Dale W. Jorgenson, eds., *Handbook of Computable General Equilibrium Modeling*, Vol. 1, Elsevier, 2013, pp. 1213–1269.

Hooper, Peter, Karen Johnson, and Jaime Marquez, "Trade Elasticities for the G-7 Countries," *Princeton Studies in International Economics*, 2000, *87.*

Houthakker, H. S. and Stephen P. Magee, "Income and Price Elasticities in World Trade," *The Review of Economics and Statistics*, 1969, *51* (2), 111–125.

Junz, Helen B. and Rudolf R. Rhomberg, "Prices and Export Performance of Industrial Countries, 1953-63)," *IMF Staff Papers*, July 1965, *12* (2), 224–271.

Kehoe, Timothy J. and Kim J. Ruhl, "How Important is the New Goods Margin in International Trade?," *Journal of Political Economy*, 2013, *121* (2), 358–392.

Khan, Shafaat Y. and Armen Khederlarian, "How Does Trade Respond to Anticipated Tariff Changes? Evidence from NAFTA," *Journal of International Economics*, 2021, *133*, Article 103538.

Pierce, Justin R. and Peter K. Schott, "The Surprisingly Swift Decline of US Manufacturing Employment," *American Economic Review*, 2016, *106(7)*, 1632–1662.

Roberts, Mark J. and James R. Tybout, "The Decision to Export in Colombia: An Empirical Model of Entry with Sunk Costs," *American Economic Review*, 1997, *87(4)*, 545–564.

Romalis, John, "NAFTA's and CUSFTA's Impact on International Trade," *The Review of Economics and Statistics*, 2007, *89* (3), 416–435.

Ruhl, Kim J., "The International Elasticity Puzzle," 2008. Mimeo.

_ **and Jonathan L. Willis**, "New Exporter Dynamics," *International Economic Review*, 2017, *58* (3), 703–726.

Schott, Peter K., "Data," https://sompks4.github.io/sub_data.html 2021. Accessed: 06-09-2021.

Simonovska, Ina and Michael E. Waugh, "The Elasticity of Trade: Estimates and Evidence," *Journal of International Economics*, 2014, *92* (1), 34–50.

Soderbery, Anson, "Trade Elasticities, Heterogeneity, and Optimal Tariffs," *Journal of International Economics*, 2018, *114*, 44–62.

Sposi, Michael, Kei-Mu Yi, and Jing Zhang, "Deindustrialization and Industry Polarization," Working Paper 29483, National Bureau of Economic Research November 2021.

Steinberg, Joseph B., "Export Market Penetration Dynamics," *Journal of International Economics*, 2023, *145.* Article 103807.

Teti, Feodora, "Missing Tariffs," 2023. Mimeo.

United Nations Trade Statistics, "Correspondence Tables," https://unstats.un.org/unsd/trade/classifications/correspondence-tables.asp 2017. Accessed: 06-09-2021.

Figure 1: Import shares and tariffs by policy regimes

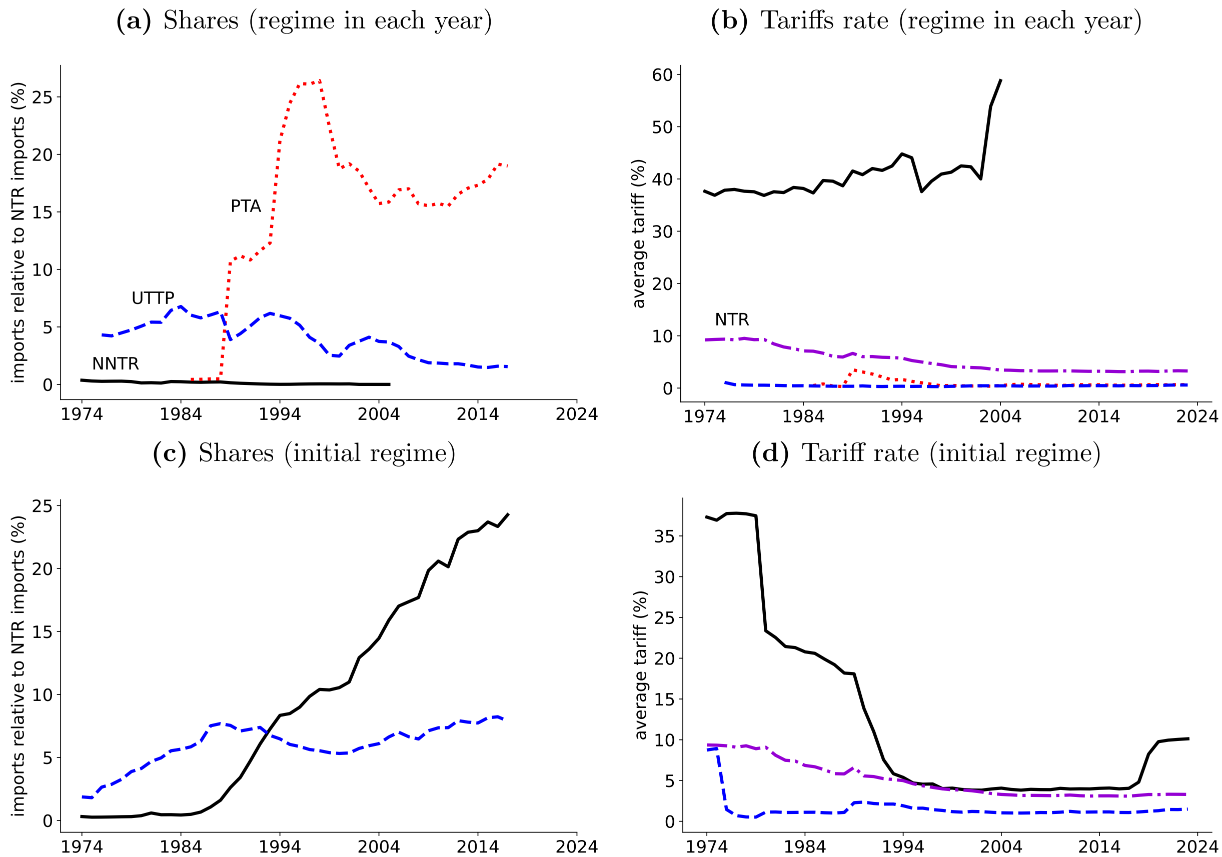

Notes: Panel (a) plots the import share of each of our four trade policy regimes described in section 2, after dropping unclassified variety-years. Panel (b) plots the average tariff rate of across goods by the regime in that year. Panels (c) and (d) are the respective analogues of panels (a) and (b) where the regime for each variety is fixed across the years as the modal regime in 1970s.

Figure 2: Trade and tariffs in China and Vietnam

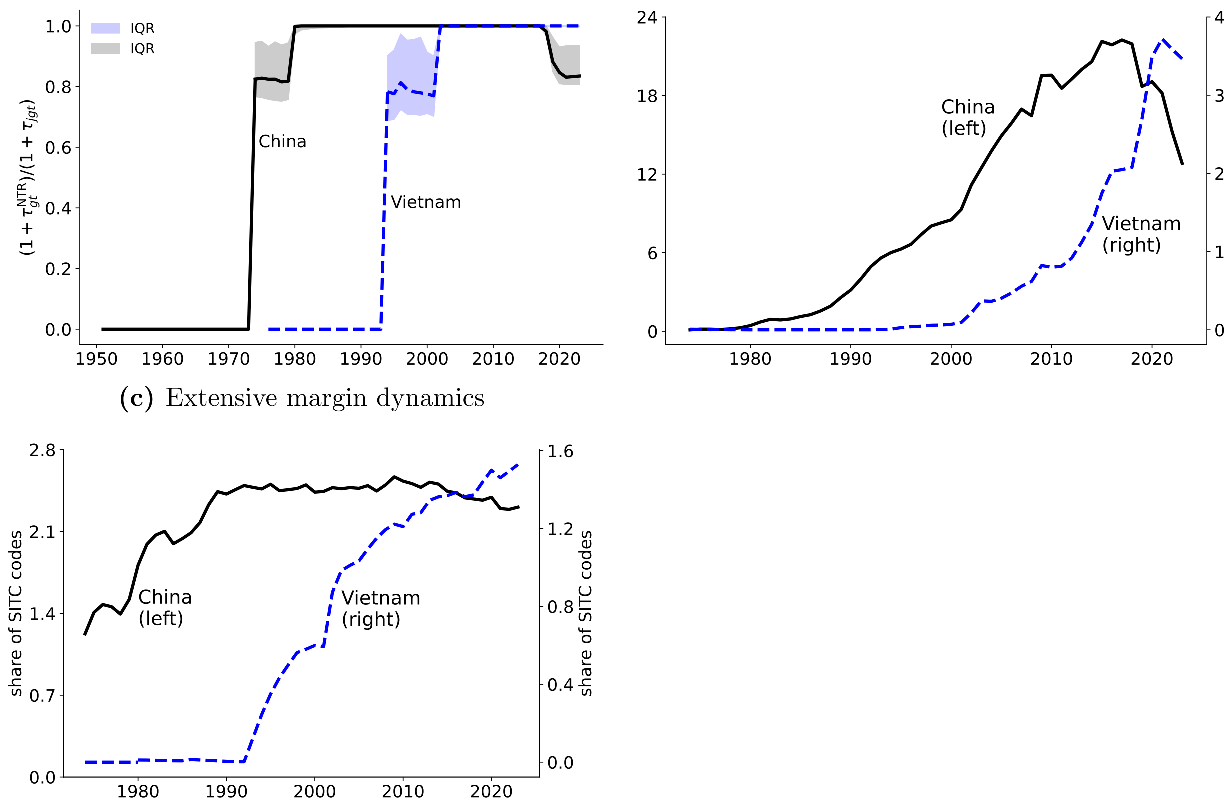

(a) Tariff dynamics

(b) Share of US imports (percent)

(c) Extensive margin dynamics

Figure 3: Autocorrelation of tariff changes

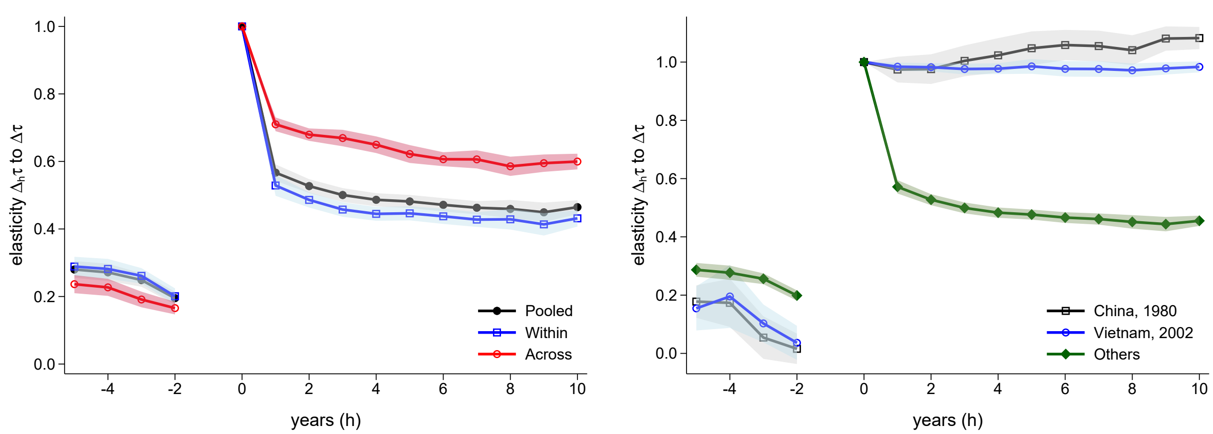

(a) Across and Within Regimes

(b) China and Vietnam Regime Change

Notes: Panel (a) plots the results of (1) and panel (b) the results of (3). The standard errors that construct the 95 percent confidence intervals are clustered at the jg level.

Figure 4: Reduced-form trade elasticities

(a) Across and Within Regimes **(b)** China and Vietnam (all years)

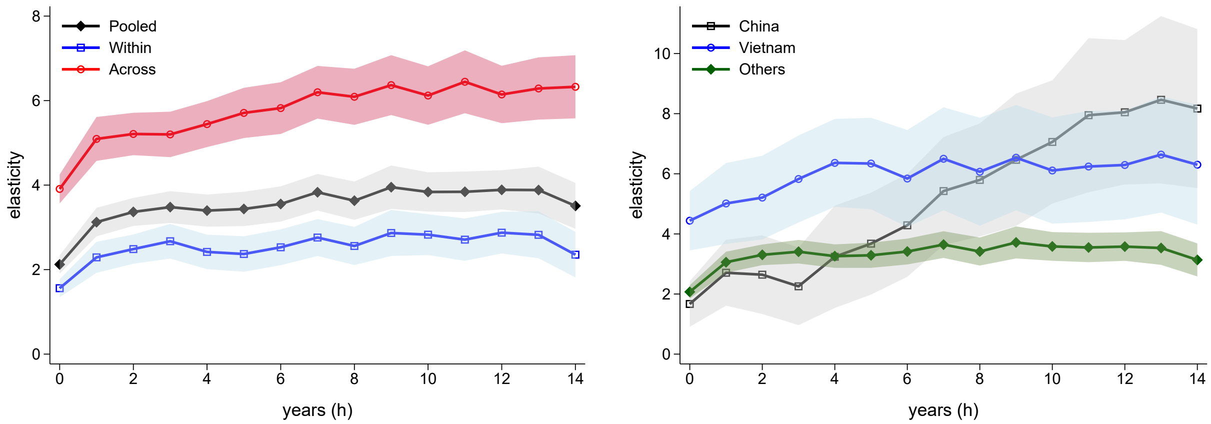

Notes: Panel (a) plots the results of (6) and panel b the results of (7). The standard errors that construct the 95 percent confidence intervals are clustered at the jg level.

Figure 5: The NNTR liberalizations of China (1980) and Vietnam (2002)

(a) Gap between NNTR and NTR **(b)** NNTR gap elasticity

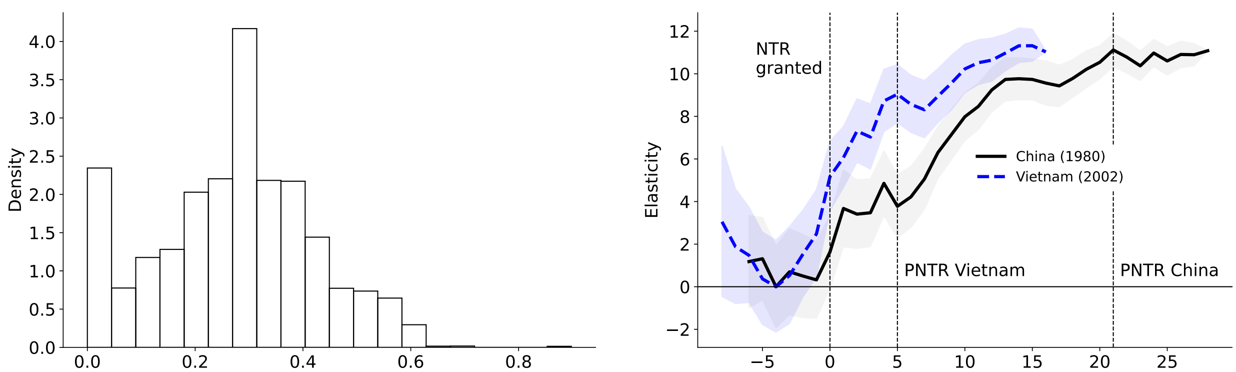

Notes: Panel (a) plots the distribution of the gap between the NNTR and the NTR rate for 1999. Panel (b) plots the elasticity coefficients from (9). The standard errors that construct the 95 percent confidence intervals are clustered at the jg level.

Figure 6: Model responses to Markov reforms

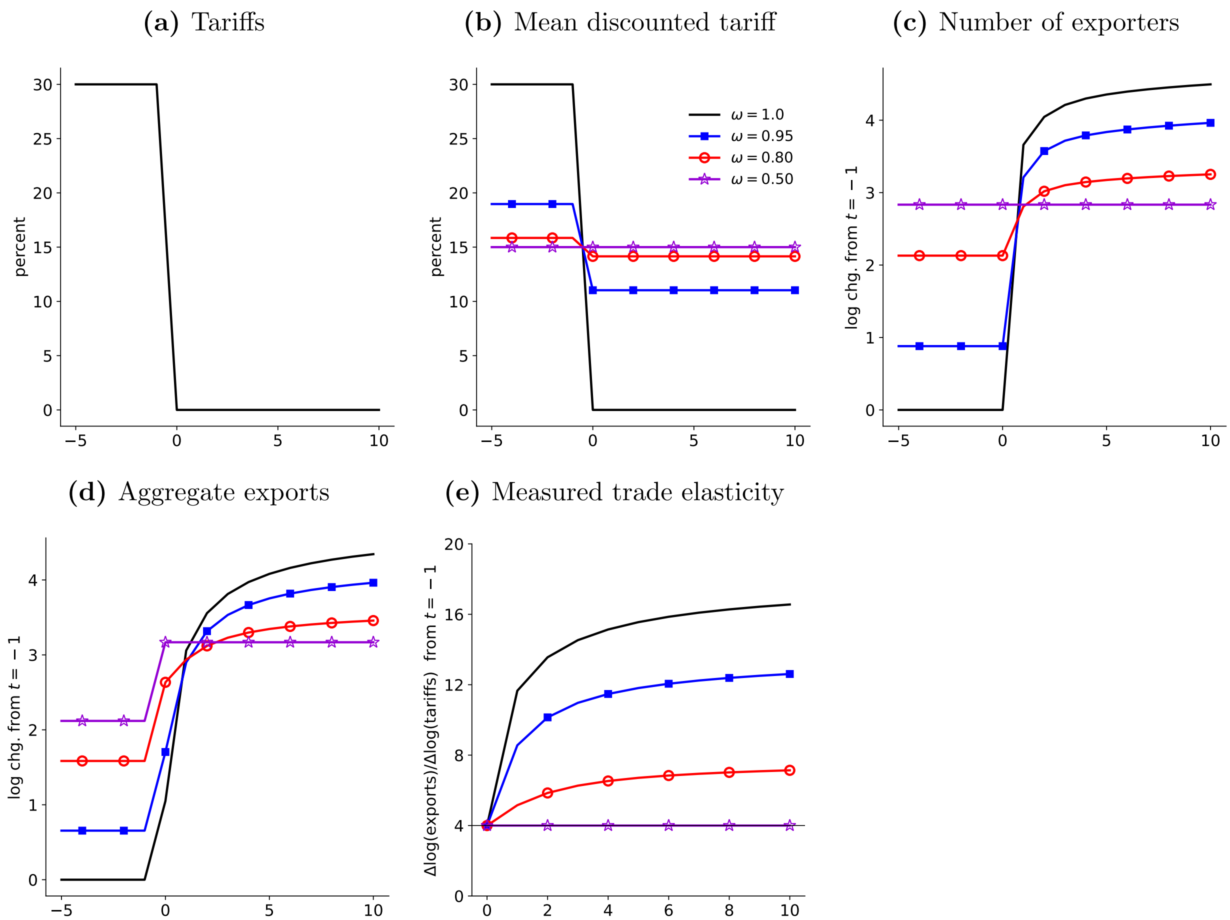

Notes: The figure plots the transition dynamics from the trade reform experiments described in section 5.2, in which the persistence of tariffs (ω) varies. In panel (e), the measured trade elasticity is defined in (4). In panels (c) and (d), outcomes are measured relative to period $t = -1$ in the calibration with perfectly persistent tariffs ($\omega = 1$).

Figure 7: Model responses to change in persistence

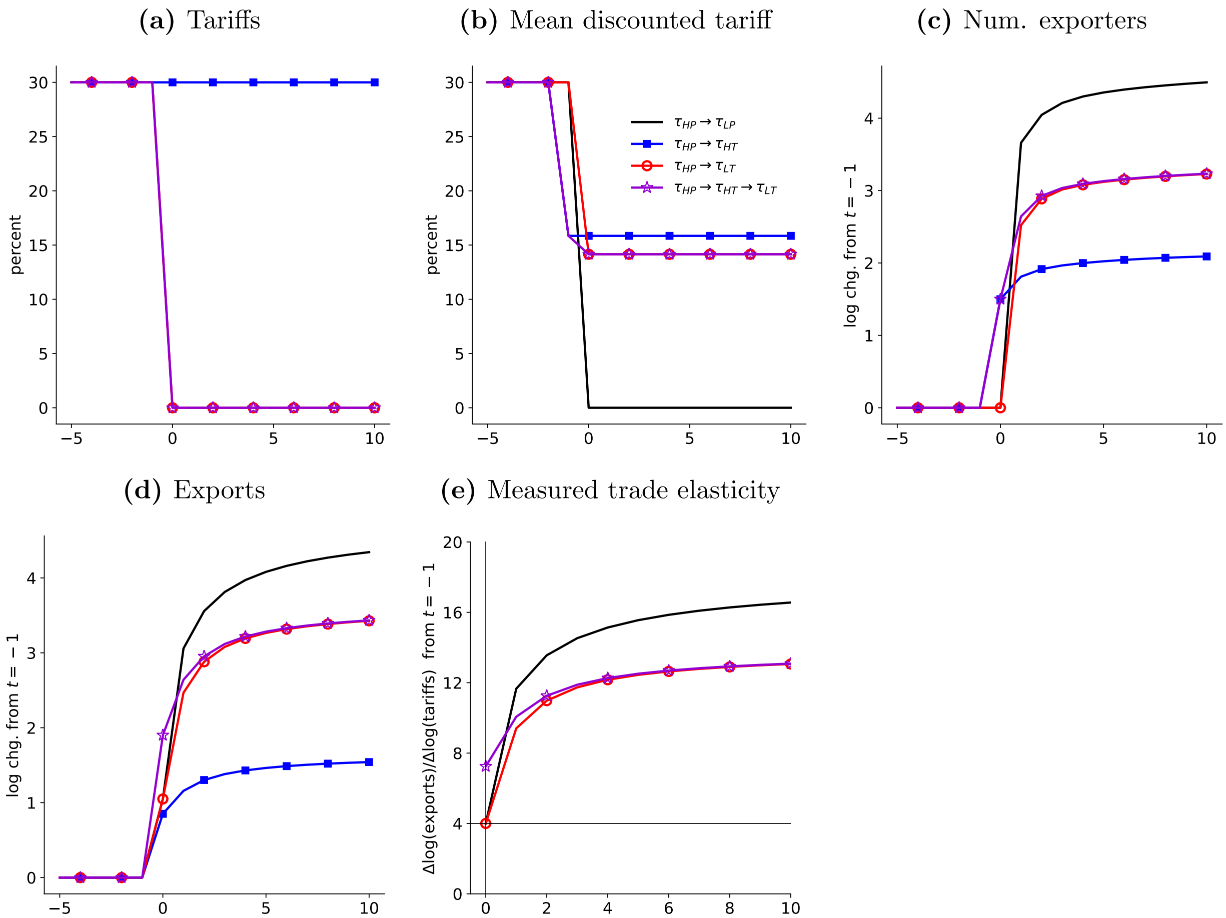

Notes: The figure plots the transition dynamics from the trade reform experiments described in section 5.3, in which the economy faces a high-persistence regime (P) and a low-persistence regime (T). In each regime, the tariff rate can be high (H) or low (L). In panel (e), the measured trade elasticity is defined in (4). In panels (c) and (d), outcomes are measured relative to period $t = -1$ in the canonical case $\tau_{HP} \to \tau_{LP}$.

Figure 8: Role of initial conditions: Dynamics from an embargo's end

(a) Exports **(b)** Elasticity

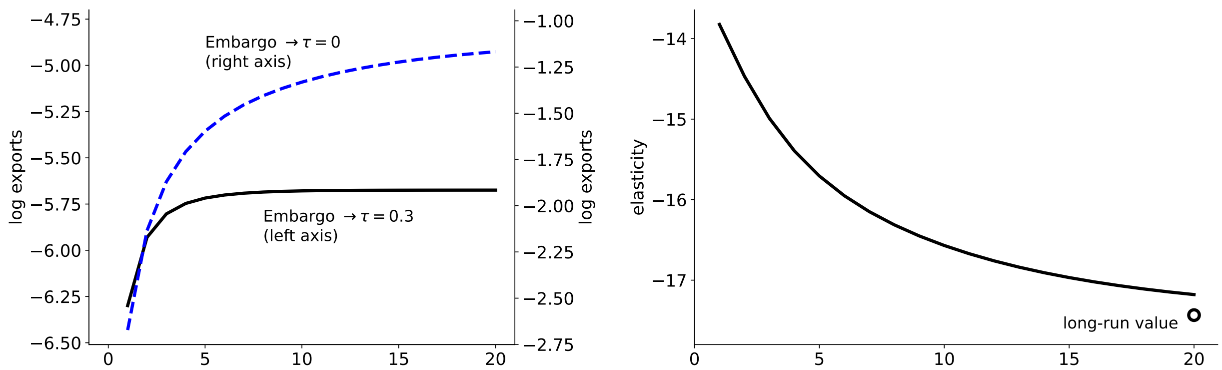

Notes: The figure plots the transition dynamics from the experiments described in section 5.4. Panel (a): Aggregate exports after unexpectedly ending an embargo and applying either a 30-percent tariff or no tariff. Panel (b): The "gap" elasticity estimates from (20).

Figure 9: The NTR liberalizations of China (1980) and Vietnam (2002): Model

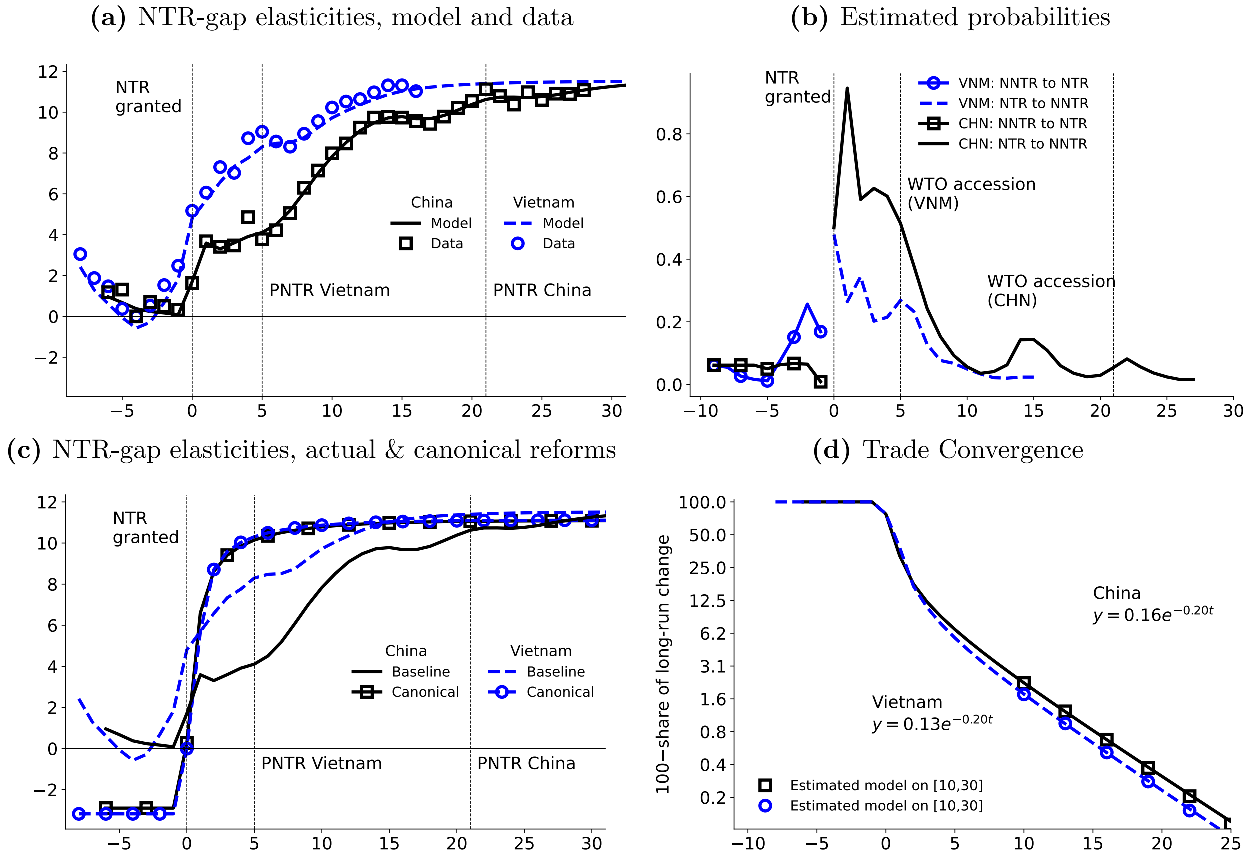

Notes: Panel (a): NTR-gap elasticities in data and baseline model. Panel (b): Estimated trade-policy transition probabilities. Panel (c) NTR-gap elasticities baseline model and the counterfactual canonical reform. Panel (d) The share of total adjustment remaining in the NTR-gap elasticities from the counterfactual canonical reform, as defined in (22).

Table 1: Tariff levels by regimes

	Baseline Sample				Winsorized (99.99)			
Regime	N # jgt	Mean (p.p.)	Median (p.p.)	SD (p.p.)	N # jgt	Mean (p.p.)	Median (p.p.)	SD (p.p.)
NTR	1,758,012	4.87	2.90	14.09	1,757,870	4.82	2.90	6.66
NNTR	19,507	38.61	35.00	22.56	19,423	38.16	35.00	21.07
PTA	97,548	0.78	0.05	2.21	97,545	0.78	0.05	2.08
UTPP	230,683	0.40	0.00	1.11	230,683	0.40	0.00	1.11
Unclassified	34,374	2.06	0.17	7.20	34,366	2.03	0.17	6.92
Total	2,105,750	4.50	1.70	13.58	2,105,521	4.46	1.70	7.38

Notes: This table reports some moments of the tariff levels by regime. Tariffs are calculated as the duties collected over FOB import values at the 5-digit SITC aggregation level, our definition of a good. Our classification of regimes is described in detail in Appendix B.1.

Table 2: Trade policy regime transition frequencies (percent)

$t-1/t$		NNTR	NTR	PTA	UTPP
NNTR	jg	80.02	18.66	0.00	1.31
NNTR	j	89.00	9.29	0.00	1.70
NTR	jg	0.05	96.93	0.65	2.37
NTR	j	0.33	88.62	0.15	10.90
PTA	jg	0.00	8.58	91.42	0.01
PTA	j	0.00	0.00	100.00	0.00
UTPP	jg	0.01	16.60	0.93	82.46
UTPP	j	0.03	9.97	0.49	89.51

Notes: The transition frequencies are calculated as the average of the annual transitions between and within regimes over the full sample period at the j level. Each cell in the annual transition matrix is calculated as the ratio of sum of countries that moved from regime R' to regime R between years $t-1$ and t to the total number of countries in regime R in year $t-1$, where $R = \{\text{NTR, NNTR, PTA, UTPP}\}$.

Table 3: Top five country-year transitions across regimes

From	To	jg (# g)
NTR	NNTR	Poland-1983 (232), Poland-1984 (78), Poland-1985 (43), Afghanistan-1986 (46), Romania-1989 (119)
NTR	PTA	Canada-1989 (889), Mexico-1994 (387), South Korea - 2012 (325), Australia-2005 (241), Israel-1986 (203)
NTR	UTPP	Taiwan-1976 (280), Hong Kong-1976 (214), Israel-1976 (204), South Korea-1976 (189), Brazil-1976 (177)
NNTR	NTR	China-1980 (273), Vietnam-2002 (347), Poland-1989 (253), Soviet Union-1992 (226), Soviet Union-1993 (215)
NNTR	UTPP	Romania-1994 (32), Czechoslovakia-1992 (31), Czechoslovakia-1991 (28), Bulgaria-1992 (26), Poland-1990 (23)
PTA	NTR	Canada-1999 (205), Mexico-1999 (179), Israel-1999 (165), Australia-2009 (135)
UTPP	NTR	South Korea-1989 (403), Taiwan-1989 (400), Hong Kong-1989 (265), Malaysia-1997 (262), Portugal-1986 (214)
UTPP	NNTR	Romania-1989 (7), Romania-1990 (6), Romania-1992 (5), Romania-1993 (5), Yugoslavia-1996 (5)
UTPP	PTA	Israel-1985 (354), Mexico-1994 (342), Peru-2007 (241), Colombia-2001 (229), Dominican Republic-2007 (176)

Notes: The transition are based on the most common regime transitions on traded products. These regime transitions are based on consecutive good-country observations but need not be in consecutive years. For instance, when Poland was transitioned back to NNTR in October of 1982 and several goods traded under NTR where not traded in 1983 or 1984.

Table 4: Tariff changes across and within regimes

		1-Year Changes				5-Year Changes			
From	To	N # jgt	Mean (p.p.)	Median (p.p.)	SD (p.p.)	N # jgt	Mean (p.p.)	Median (p.p.)	SD (p.p.)
NTR	NNTR	466	27.36	26.96	19.02	913	28.03	29.26	14.90
NTR	PTA	10,281	−3.01	−1.80	4.56	16,838	−4.21	−2.70	5.64
NTR	UTPP	29,978	−4.02	−2.90	14.50	26,923	−5.22	−3.83	6.32
NTR	NTR	1,352,360	−0.15	0.00	9.47	901,653	−0.66	0.00	15.80
Tokyo	Tokyo	180,246	−0.45	−0.17	9.66	125,494	−1.81	−0.80	34.46
Uruguay	Uruguay	197,315	−0.25	−0.04	2.46	158,845	−0.84	−0.30	2.87
Non-GATT	Non-GATT	974,799	−0.07	0.00	10.29	442,584	−0.24	0.00	12.18
NNTR	NNTR	10,542	−0.25	0.00	9.25	3,012	−0.44	0.00	10.02
NNTR	NTR	1,523	−27.63	−26.17	24.04	3,343	−31.48	−30.03	21.46
NNTR	UTPP	72	−29.81	−32.93	16.20	735	−36.25	−34.99	18.12
PTA	NTR	8,432	2.47	1.30	3.95	8,911	1.39	0.00	3.57
PTA	PTA	75,876	−0.12	0.00	1.34	48,431	−0.50	0.00	2.07
PTA	UTPP	1	0.00	0.00					
UTPP	NNTR	12	23.04	24.08	17.77	75	34.67	34.96	16.24
UTPP	NTR	30,373	2.89	2.08	3.80	38,307	2.33	1.59	3.41
UTPP	PTA	1,848	0.05	0.00	1.36	7,057	−0.18	0.00	1.56
UTPP	UTPP	150,464	−0.03	0.00	1.04	82,515	−0.06	0.00	1.14
	Total	1,672,228	−0.17	0.00	8.92	1,138,713	−0.74	0.00	14.40

Notes: Tariff changes are calculated as the average over $\tau_{jgt} - \tau_{jg,t-h}$ with $h = \{1, 5\}$. For the 1-year change, GATT-Tokyo are averages for the years 1980-1987 and GATT-Uruguay 1995-2000. For the 5-year change, GATT-Tokyo are averages for the years 1984-1992 and GATT-Uruguay 1999-2004.

Table 5: Exporter-dynamics statistics and sector-level model parameters

Panel a: Common assigned parameters

	Parameter	Value	Target/Source
θ	Demand elasticity	3.17	Soderbery (2018)
r	Interest rate	0.04	Common Value
δ_0	Constant exit rate	21	Alessandria et al. (2021)
δ_1	Elasticity of exit to productivity	0.02	Alessandria et al. (2021)
ρ_ξ	Trade cost transition persistence	0.92	Alessandria et al. (2021)

Panel b: Country-specific jointly calibrated parameters

	Parameter	China	Vietnam	Target/Source	China	Vietnam
f_0	Sunk cost	0.73	1.57	Export part. (%)	28	11
f_1	Export cost	0.342	0.657	Exit rate (%)	11	15
ξ_H	High iceberg cost	3.92	5.89	Incumbent prem.	2.9	4.41
σ_z	Productivity shock dispersion	1.50	1.69	Log CV exports	2.27	2.91

Table 6: Estimates of Trade Elasticity Dynamics

	Canonical Reform			h-on-1			ECM		
Country	ε_0	ε_{10}	ε_∞	ε_0	ε_{10}	ε_{14}	ε_0	ε_{10}	ε_∞
China	3.17	13.74	13.99	1.67	7.06	8.17	2.69	9.69	9.91
Vietnam	3.17	14.09	14.29	4.44	6.11	6.30	4.85	8.39	8.40

Notes: Canonical reforms are based on transition from permanent NNTR to permanent NTR tariffs. The h-on-1 estimates and ECM (error correction mode) are based on the entire sample.

Appendix (For online publication)

We study several features of the model and the data in this appendix. In section A we use the model to study alternative reforms without uncertainty but include other forms of anticipation. In section B, we discuss the robustness of our empirical approach.

A Model

Aside from the changes between NNTR and NTR/UTPP, many tariff changes involve a phaseout. These are a key feature of GATT rounds and PTAs. We now discuss how trade evolves with these types of tariff reforms through the lens of our model and show that short-run and long-run trade elasticities can differ quite substantially from those from a once and for all trade reform.

A.1 Anticipated reforms

To study how anticipation affects trade adjustment dynamics, we consider several variations on the canonical scenario. We consider two kinds of variations: anticipated versus unanticipated and immediate versus phaseouts. Unanticipated reforms are announced in the period in which they occur, whereas anticipated reforms are announced ten periods in advance. In immediate reforms, tariffs change by 30 p.p. in one period, and in tariff phaseout reforms, tariffs fall by 3 p.p. per year for 10 years. The linear phaseout is a key future of trade agreements, including the GATT and FTAs. In each variation, firms learn the entire path of tariffs when the reform is announced. Thus, phased-in reforms always feature some anticipation of future tariff cuts, even when the start of the reform is a surprise. None of the reforms considered in this section feature uncertainty about future tariffs. Figure A1 shows the results of these experiments. There are three main takeaways.

First, anticipation causes exports to rise before tariffs fall, increasing the short-run trade elasticity, ε_0. When the reform is anticipated, the future decrease in tariffs increases the value of exporting when the reform is announced, which causes export participation to rise. When tariffs fall, aggregate exports increase from the intensive-margin effect and the pre-reform export entry. Compared with an unanticipated reform, anticipated reforms feature more trade in $t = 0$. Thus, the measured short-run trade elasticity is higher than the short-run canonical elasticity (Figure A1(e)).

Second, phased-in reforms cause exports to grow slower than immediate reforms, but lead to higher trade elasticities in the short and medium run. When tariffs fall gradually, the value of exporting and the export participation rate rise gradually. However, exports rise faster, relative to the rate at which tariffs change, than in immediate reforms. Export entry depends on the entire path of future tariffs, not just the current tariff rate, and firms anticipate future tariff cuts in phased-in reforms. In unanticipated phased-in reforms, the trade elasticity in period $t = 0$ is equal to the demand elasticity, θ (as in the canonical reform), but in period $t = 1$ it jumps above its long-run value. In anticipated and phased-in reforms, the trade elasticity is very high even in period $t = 0$ and exhibits similar, albeit more muted, non-monotonic behavior.

Third, anticipation does not affect the long-run level of trade but does affect the measured long-run trade elasticity. In all four reforms, export volumes converge to the same level (Figure A1(d)), but each reform has a different long-run trade elasticity. This is because the level of trade in the period before the reform begins, X_{t-1} in (4), is higher when the reform is anticipated. This is particularly true in the anticipated, immediate reform, where more than three-quarters of the long-run growth in trade occurs before tariffs change at all, which results in a much lower long-run trade elasticity than in the other versions of the experiment. Of course, if we measure the change in trade relative to the announcement date, as in Figure A1(g), the long-run elasticity is the same across experiments. Empirical approaches that attribute this anticipatory growth to pre-trends will generate even smaller trade elasticities.

B Data

We describe our data construction in greater detail. We demonstrate our reduced for elasticities are robust to several

B.1 Regime classification

Here we describe our regime classification used in sections 2 and 3. To classify each triplet jgt into one of the four regimes, we incorporate information from the rate of provision and country subcode aggregation level of the census trade data. In particular, we classify jgt's into each of the regimes when at least 50 percent of the import value of the triplet is attributed to that regime. We define regimes at the jgt-rate of provision-country subcode level as follows:[28]

- *NTR*: Rates of provision 10, 61 or country subcodes C, L, and K.
- *NNTR*: Rates of provision 62.
- *PTA*: Country subcodes specified in Table A2 or rates of provision 18, 19, or 64 for country-years in Table A2.
- *UTPP*: Country subcodes specified in Table A1 or rates of provision 18, 19, or 64 for country years with at least one observation under country subcode A, A*, A+ or 4 (GSP).

This procedure allows us to classify 98.4 percent of the jgt triplets (2,587,514/2,629,684). Panel a of Figure A3 plots the import share and median applied tariffs for the unclassified jgt's. It indicates that these represent a small fraction of imports and are mostly characterized by low tariffs until 2017, when the United States raised some tariffs on a set of goods and, especially, on imports on China, was mostly carried out under provision rates 69 and 79, which are not used by any of our 4 regimes. This justifies our focus on the data until 2017.

The imports of most of the jgt triplets that are classified according to one of our four regimes are completely characterized by their regime. This can be seen in Panel b of Figure A3, which is a plot of the import share of the attributed jgt regime. The import share

[28] Note the provision rate codes have changed over time. Using the Annual Import Data Bank Files between 1974 and 1989 from the National Archives we concord previous provision rates to the current ones. For instance, before 1989, provision rates distinguished ad valorem, specific, compound, and minimum rates, e.g. dutiable rates prescribed by column 2 tariffs were coded as 22 for specific rates, 32 for ad valorem rates, 42 for compound rates, and 52 for minimum rates. Currently all of these would fall under the provision rate 62 and are classified as the NNTR regime under our regime classification.

of the regime is always very close to 1, especially in the case of the NTR and NNTR regime. The fact that the regime of some jgt's does not account for all imports is mostly due to product aggregation.

To clarify our classification, consider for example electric motors incl. universal A.C./D.C. motors (SITC code 71621) imported from Canada in 1990. We classify this triplet as "PTA". However, the "PTA" import share of this jgt is only 65 percent, with the "NTR" share being 34 percent and 1 percent unclassified. Table A3 shines some light on this: There are 20 HTS-8 goods associated in this SITC code, of which some provide lower than NTR rates under the US-Canada FTA while others already have zero tariffs under the NTR regime.

B.2 Robustness: Reduced-form trade elasticities

Our estimates of the reduced-form trade elasticities reported in section 3.1 are robust to a range of alternative estimation strategies and data samples that we describe below.

By regime transition While in the baseline we focused on the distinction of across and within-regime transition, here we report the results of (1) and (6) when we consider specific regime transitions. In particular, we show the estimates for transitions from NNTR to NTR/UTPP, from NTR to PTA, and from NTR to UTPP, as well as the estimates for jg's that were always entering under the NTR regime. Panel (a) of Figure A5 plots the autocorrelation of tariffs and panel (b) the results of the reduced-form trade elasticities. It confirms that within NTR tariffs are less persistent and lead to less substitution.

Definition of regimes Our baseline classification into trade policy regimes considers that jgt's are classified into one of the four regimes if 50 percent of its imports entered under the corresponding rate of provision and/or country subcode, as described in B.1. Here we consider two robustness checks. First, we restrict the import share to be 90 percent instead of 50 percent. Second, we consider only the first transition of any jg over our sample period, addressing potential concerns of pre-trends. The results, reported in Figure A6, indicate that across-regime elasticities increase slightly when we restrict the imports to 90 percent, while focusing on the first transition only has a minimal impact on the estimates, although it yields more precise estimates (standard errors are not shown).

Estimation specification Our baseline estimation of the reduced-form trade elasticities largely follows the local projection approach of Boehm et al. (2023). While our results are similar to their when we neglect regime transitions and pool over the full sample (black line in Figure 1, the differences between in elasticities from across versus within-regime transitions remain similar if we consider alternative specification of the dynamic trade elasticities. In particular, we now show two additional common methods to estimate the dynamic pattern of trade adjustment yield quite similar results. Our first approach, which we call the the h-on-h method, measures the ratio of the cumulative change in trade and tariffs. Our second approach is a partial adjustment model known as the error-correction model. Partial adjustment models have been used widely to account for trade dynamics since at least Houthakker and Magee (1969).

First, we consider a direct application of (4) to the data, regressing changes in trade after h periods on changes in tariffs after h periods while using fixed effects to control for confounding covariates. The estimating equation, which we call the *h-on-h* specification, is:

$$\Delta_h v_{jgt} = -\beta_h^{hh} \Delta_h \tau_{jgt} + \mathbf{Z}_{jgt} + \delta_{jt} + \delta_{gt} + u_{jgt}. \tag{23}$$

where $\Delta_h x_t = x_{t+h} - x_{t-1}$ and β_h^{hh} is the estimator of the h-horizon trade elasticity under the h-on-h approach.

Our second alternative specification, which we call *ECM*, is a version of the error correction model (Houthakker and Magee, 1969; Hooper et al., 2000; Gallaway et al., 2003; Alessandria and Choi, 2021). This approach allows us to recover a parametric estimate of the long-run trade elasticity by specifying trade flows as an autoregressive process of order one and allowing for lagged effects of tariffs. We estimate:

$$\Delta_0 v_{jgt} = \beta_0^{ecm} \Delta_0 \tau_{jgt} + \gamma(v_{jg,t-1} - \beta_\infty^{ecm} \tau_{jg,t-1}) + \delta_{gt} + \delta_{jt} + \delta_{jg} + u_{jgt}, \tag{24}$$

where the dependent variable is the one-year log difference in import values between t and $t-1$, and the right-hand side includes the one-year change in tariffs, lagged tariffs, lagged log imports and a set of fixed effects. The one-year, or short-run, elasticity, β_0^{ecm}, is the coefficient on the contemporaneous tariff change. The long-run elasticity, β_∞^{ecm}, is pinned down by the ratio of the coefficient on lagged tariffs to the coefficient on lagged imports or speed of adjustment. The fixed effects are the same as in (5) except for the inclusion of variety fixed effects, δ_{jg}, which capture each variety's average import level. This is warranted because, in contrast with the previous two specifications, jg fixed effects do not cancel out with the use of changes on the left hand side since once can rewrite the ECM specification in levels.[29] One important difference between the two local projection specifications, h-on-1 and h-on-h, is that the former can fix the reform period at t, while the "ECM" approach averages over all changes.

Figure A7 compares the estimates under the h-on-h and the ECM approach to our baseline estimates. Panel (a) reports the results for the across- and within-regime elasticities. In the case of the across-regime elasticities the estimates are very similar. In contrast, in the case of the within-elasticities, especially the h-on-h elasticities are much larger, indicating that the autocorrelation properties of within-regime tariff changes are critical for the identification with local projections. Nevertheless, under all three specifications the differences between across- and within-regime elasticities remain sizeable. Panel (b) plots the results for China and Vietnam. For China, the 1-year and 14-year elasticities are very similar under the three specification, although the ECM-specification misses out on the very slow adjustment path captured by the local projections. In the case of Vietnam, the long-run h-on-h and ECM elasticities are slightly larger.

Alternative sample designs Our baseline sample design excludes jg's that were targeted by temporary trade barriers, such as antidumping duties, at some point throughout our sample period, as well as observations that are not classified by our regime classification. Here we show that our results are robust to this choice by reestimating (6) and (7) with the full sample. We also show that our results are unchanged if we drop main trading partners from the sample, as is done in Boehm et al. (2023). Note, as in Boehm et al. (2023), we define main trading partners the jg's that account for the largest import share of good g in period t and $t-1$ as well as the largest trading partner j overall in periods t and $t-1$.[30] Figure A8

[29]This is, (24) can be written as $v_{jgt} = \beta_0^{ecm}\tau_{jgt} + (1+\gamma)v_{jg,t-1} - (1+\gamma\beta_\infty^{ecm})\tau_{jg,t-1} + \delta_{gt} + \delta_{jt} + \delta_{jg} + u_{jgt}$.

[30]The rationale for this exclusion discussed in Boehm et al. (2023) is that major trading partners were

illustrates that both considerations have minimal effects on the reduced-form elasticities.

Tariff measurement Our baseline measurement of tariffs considers the weighted average of applied tariffs of the HS-8/TS-USA tariff lines at the 5-digit SITC level. Here we show that our results are very similar if we instead consider the simple average or the median over the HS-8/TS-USA tariff lines. Figure A9 plots the results of (6) and (7). The differences between the across- and within-regime elasticities become slightly larger, as do Vietnam's, while China's remain mostly unchanged.

Controlling for pre-trends. In the baseline specification we control for pre-trends of import growth and tariffs. Figure A10 plots the results of (6) and (7) when we only control for lagged import growth and when we don't include any pre-trends at all. While the point estimate tend to decrease slightly in both cases and for all groups, their differences with respect to the pooled estimates and the differences wit respect to each other remain largely unchanged.

B.3 Robustness: China and Vietnam NTR access

Our approach to estimate the gap-elasticities laid out in section 3.2 and which we use in section 6 are robust to a range of alternative estimation strategies which we describe below.

Measurement of the gap In the baseline we consider a common gap for China and Vietnam, $X_g = \log(1 + \tau_g^{NNTR} - \tau_g^{NTR})$, where we consider the simple average over the HS-8 1999 scheduled NNTR and NTR rates. We consider the following robustness checks to this choice. First, we consider the median over the HS-8 tariff lines. Second, we consider the 2001 rates. Third, we consider the applied NNTR and NTR rates at the time of the NTR access of each country, that is 1980 for China and 2002 for Vietnam. To calculate the average applied NNTR rates we consider the years 1974–1979 for China and 1994–2001 for Vietnam. We calculate the NTR rates using the two years after the access.[31] Figure A11 indicates that the gap-elasticities are virtually unchanged under these alternative definitions of the gap.

generally the ones negotiating the MFN rate reductions in the WTO's multilateral negotiation rounds.

[31]Note, this definition of the gaps requires goods to be traded both before and after the access, thus we separate the gap-elasticities for China and Vietnam separately under this specification.

Alternative sample designs In the baseline we exclude goods subject to the Multi Fibre Agreement quotas from the sample since these goods faced very different trade barriers. As robustness checks we further exclude goods that at some point throughout our sample period were affected by temporary trade barriers, such as antidumping duties. We also consider an approximate version of a balanced sample of goods, that is goods that were traded before the the two countries' respective NTR access. Finally, we consider a sample with only NTR countries (defined at the jgt-level) to allow a closer control of changes in MFN rates. Figure A12 illustrates that none of these changes in the sample design significantly affect our estimates.

C Additional figures

Figure A1: Model responses to anticipated vs. unanticipated reforms

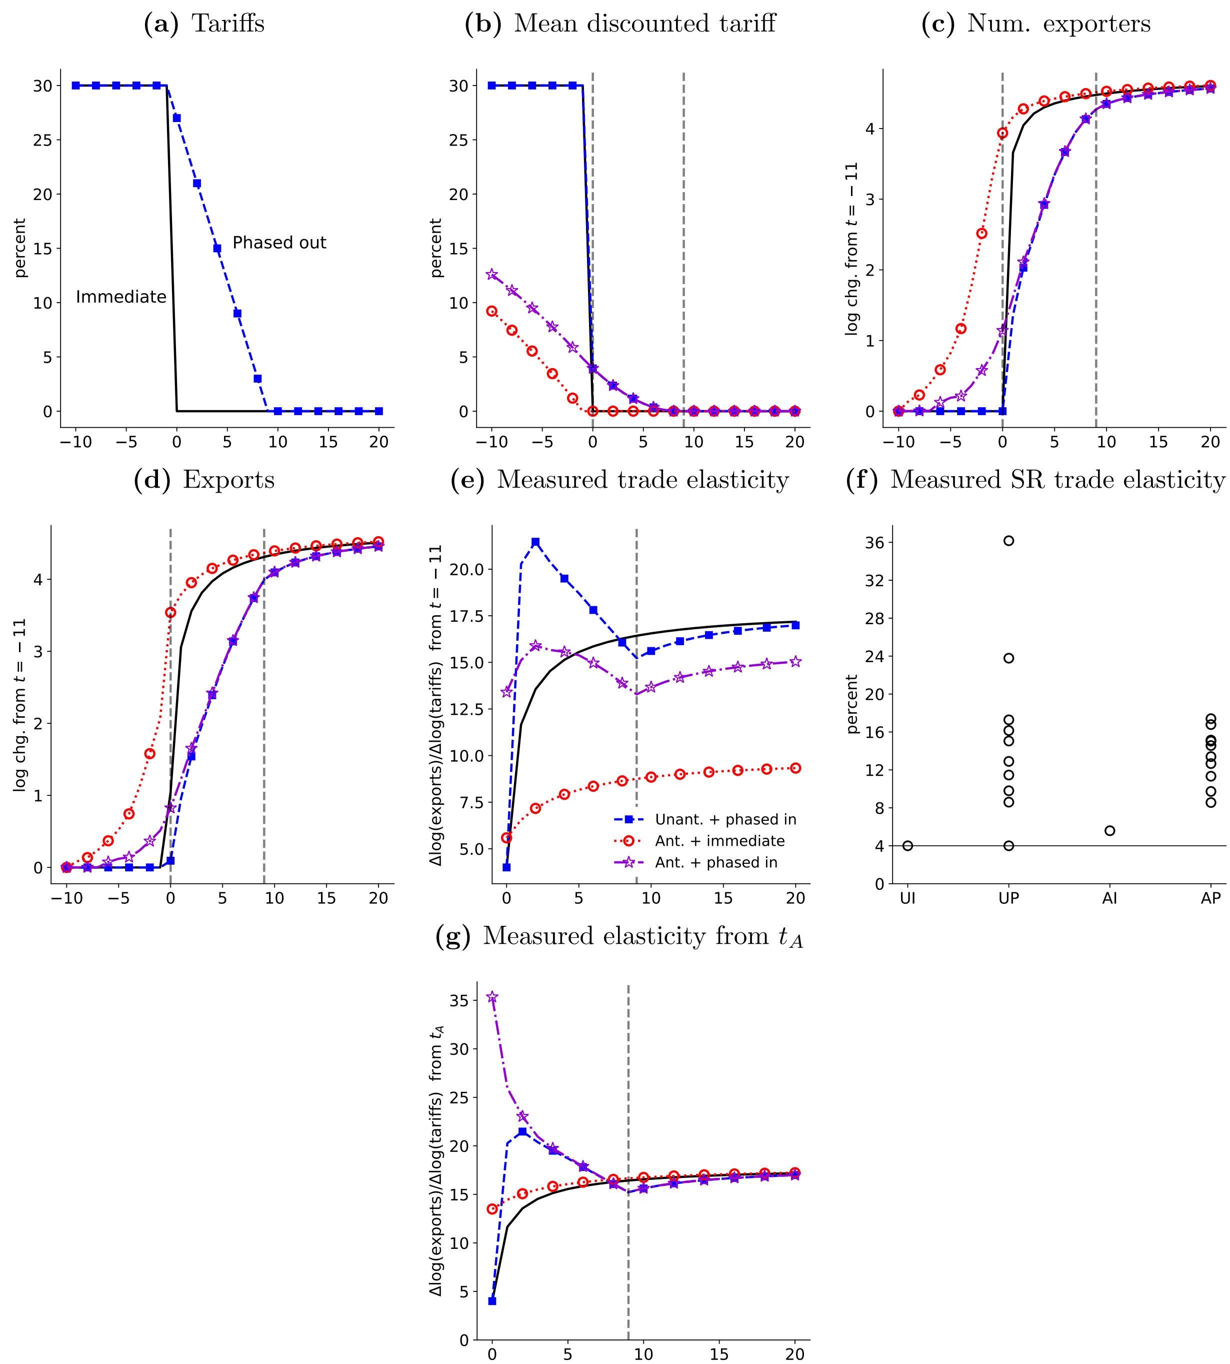

Notes: Figure shows transition dynamics in experiments described in section A.1. Panel (a): observed path of tariffs in immediate and phased-in reforms. Panel (b): Mean discounted expected tariff in unanticipated/immediate, anticipated/immediate, unanticipated/phased-in, and anticipated/phased-in reforms. Panel (c): number of exporters. Panel (d): aggregate exports. Panel (e): measured trade elasticity defined in (4). In panels (c) and (d), outcomes measured relative to period $t = -11$ to illustrate how they begin to respond after reform is announced, but before tariffs actually change.

Figure A2: Tariff Phaseouts from GATT rounds

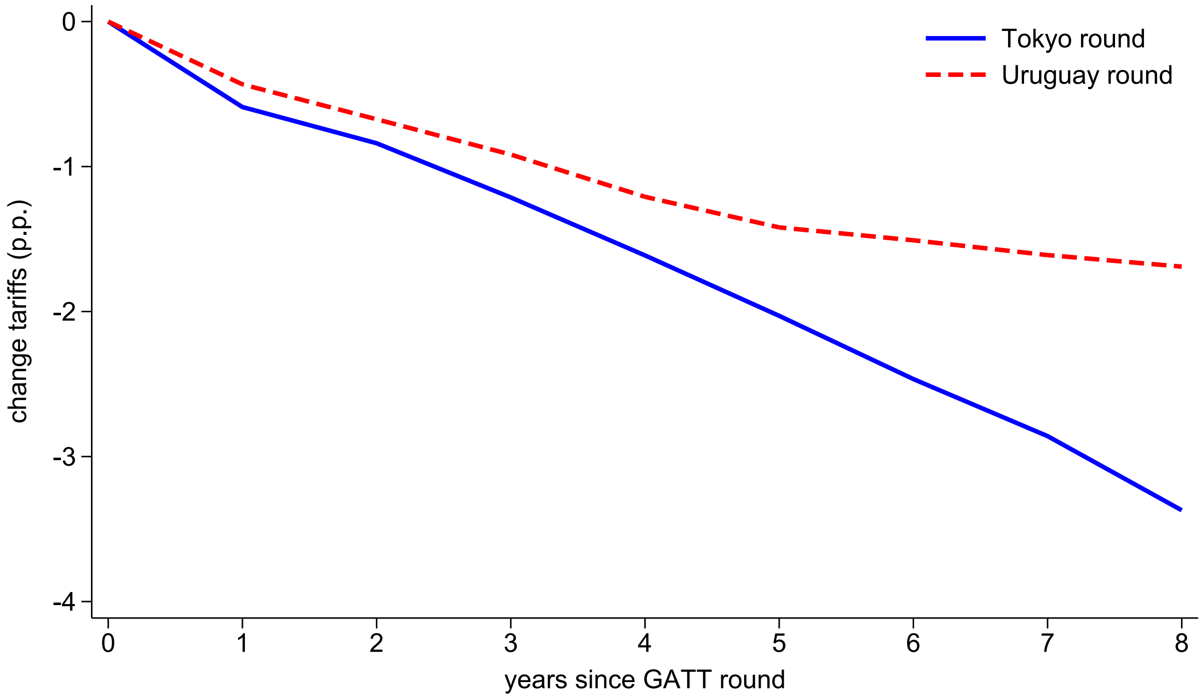

Notes: The figure shows the average tariff cut across varieties for each year, relative to the period before the Tokyo (1979) and Uruguay (1994) GATT rounds.

Figure A3: Regime Classification

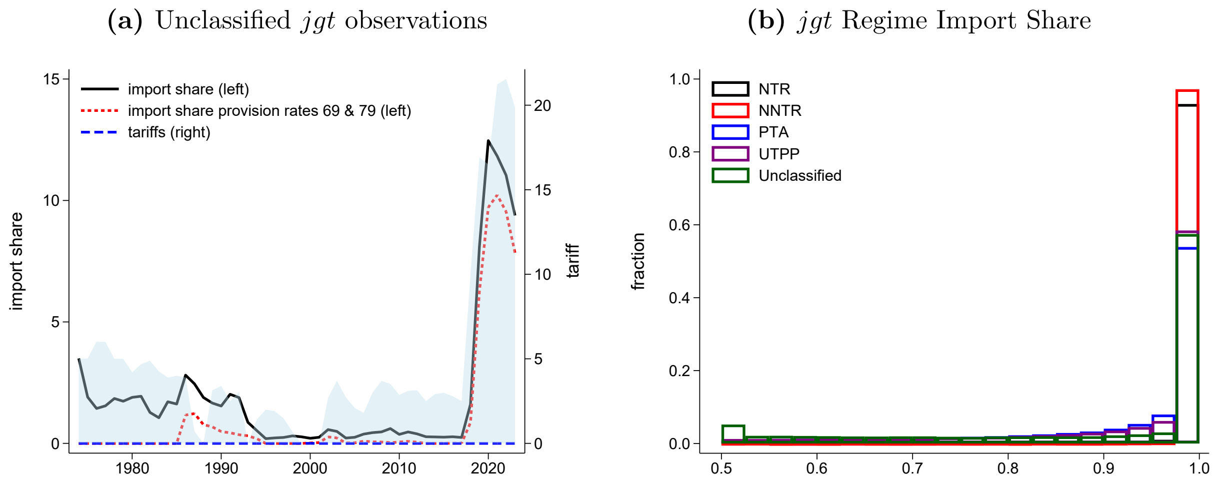

Notes: Panel (a) plots the import share and tariffs of observations (jgt's) that our regime classification described in Appendix B.1 is not able to classify into one of the four regimes. Panel (b) plots the distribution of the import share of the classified regime over the total imports (solid black line) of the corresponding jgt. Note by definition this has to be at least 50 percent for our four regimes. For the illustrative purposes, we truncate unclassified jgt's at 0.5. It also plots the share specific to provision rates 69 and 79 excluded by our classification (dotted red line). The blue dashed line plots the median tariff and the shaded area is the 25^{th} and 75^{th} percentiles of the tariff distribution.

Figure A4: Empirical trade elasticity - Reconciling specifications

(a) Autocorrelation Tariff Changes

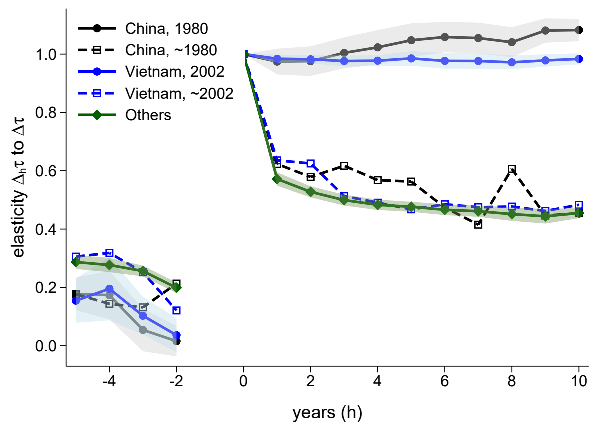

(b) China

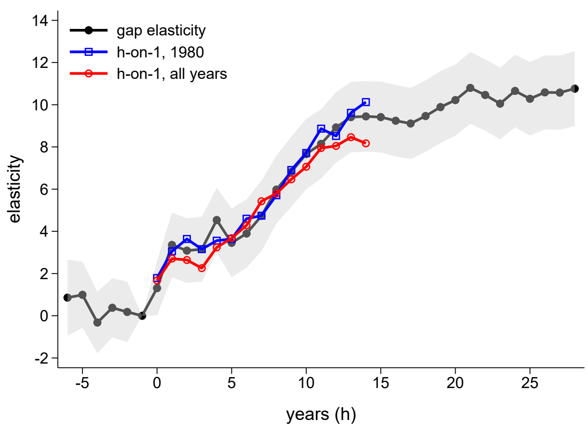

(c) Vietnam

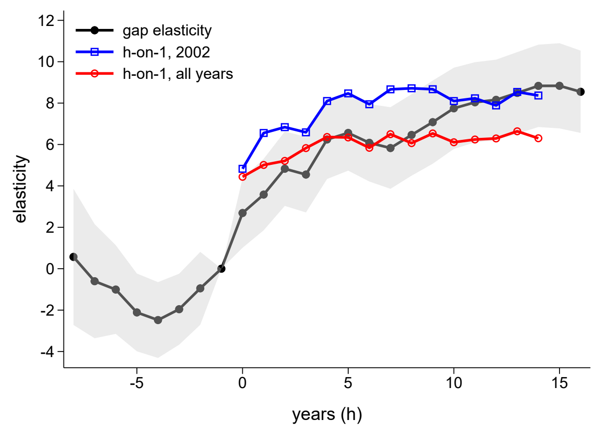

Notes: Panel (a) plots the estimates of (3) plotted in Figure 3(b) together with the estimates when excluding the NTR access years, namely 1980 for China and 2002 for Vietnam. Panels (b) figure plots the estimates of the trade elasticity under the two approaches in section 3, namely (7) and (9). The black solid line (and the gray shaded area) plot the gap-elasticity when considering the period before the NTR access ($h = -1$) as the reference period. The blue line plots the h-on-1 results when focusing on the changes from the NTR access; and the red line plots when averaging over all years. In panel (b) period 0 is 1980 (China's NTR access) and in panel (c) it is 2002 (Vietnam's access).

Figure A5: Across- versus Within Regimes: Per regime transitions

(a) Tariff autocorrelations **(b)** Reduced-form elasticities

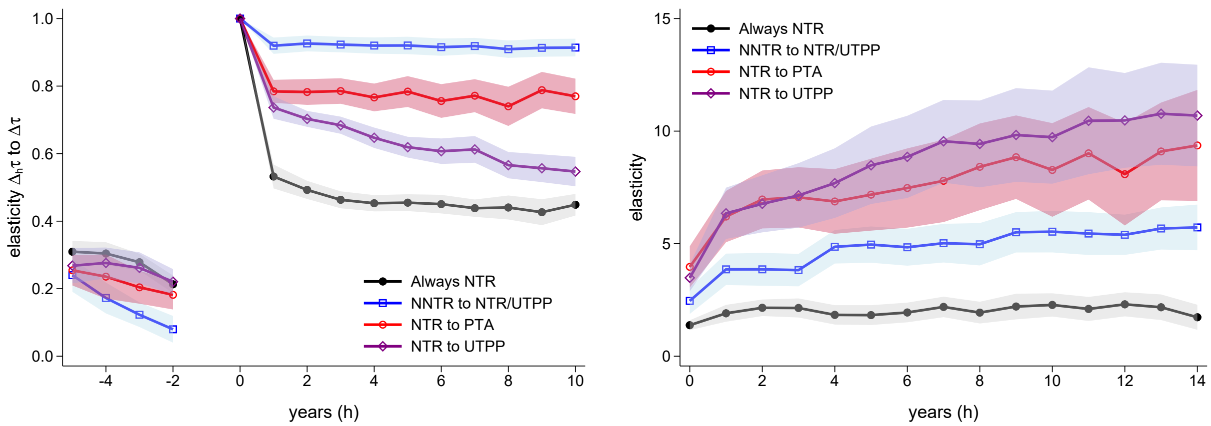

Notes: Panel (a) plots the results of (1) and panel (b) the results of (6) when considering the regime transitions specified in the legend. In both figures the shaded area the 95 percent confidence interval, with standard errors clustered at the jg level.

Figure A6: Reduced-form trade elasticities: Alternative Regime Classifications

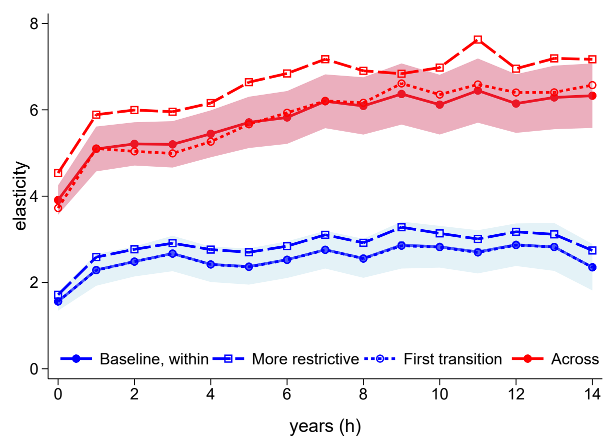

Notes: This figure plots the results of (6)when considering alternative regime classifications. The solid line and the shaded area are the baseline estimates and their 95 percent confidence interval, with standard errors clustered at the jg level. The long dashed line considers 990 percent import share of the corresponding rate of provision and country subcodes instead of 50 percent; and the short dashed line considers only the first transition undergone by a jg.

Figure A7: Reduced-form trade elasticities: Estimation specification

(a) Across- vs. Within **(b)** China & Vietnam

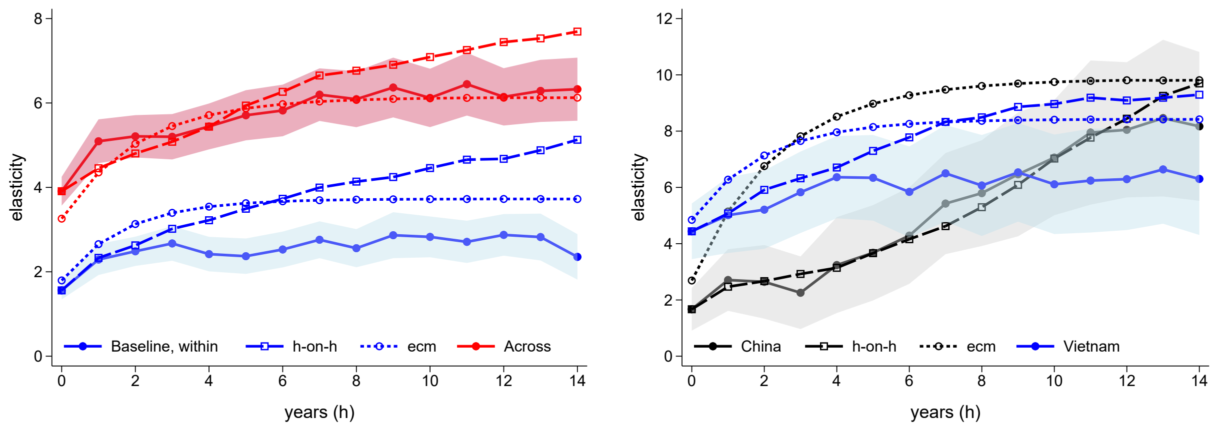

Notes: Panel (a) plots the results of (6) and panel (b) the results of (7). In both figures the solid line and the shaded area are the baseline estimates of those equations and their 95 percent confidence interval, with standard errors clustered at the jg level. The other lines are estimates from alternative dynamic specifications of the trade elasticity, namely the h-on-h approach in (23) (long dash) and the ECM approach in (24) (short dash).

Figure A8: Reduced-form trade elasticities: Sample design

(a) Across- vs. Within **(b)** China & Vietnam

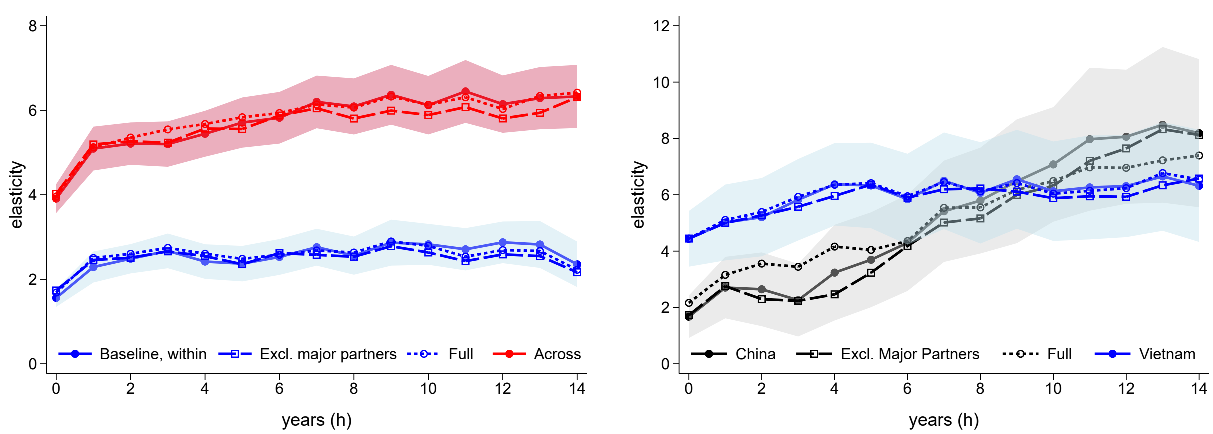

Notes: Panel (a) plots the results of (6) and panel (b) the results of (7). In both figures the solid line and the shaded area are the baseline estimates of those equations and their 95 percent confidence interval, with standard errors clustered at the jg level. The other lines are estimates from modified versions of the baseline in which consider a sample that excludes major trading partners (long dash) and a sample that includes jg's affected by temporary trade barriers and which are not classified into any regime (short dash).

Figure A9: Reduced-form trade elasticities: Tariff measurement

(a) Across- vs. Within **(b)** China & Vietnam

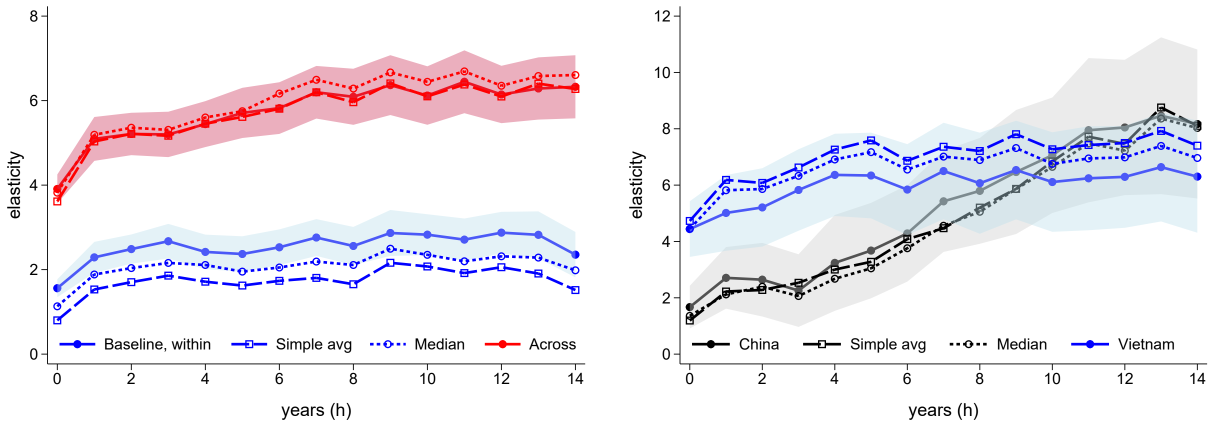

Notes: Panel (a) plots the results of (6) and panel (b) the results of (7). In both figures the solid line and the shaded area are the baseline estimates of those equations and their 95 percent confidence interval, with standard errors clustered at the jg level. The other lines are estimates from modified versions of the baseline in which consider the simple average over the HS-8/TS-USA tariff lines to measure tariffs (long dash) and their median (short dash).

Figure A10: Reduced-form trade elasticities: Pre-trends

(a) Across- vs. Within **(b)** China & Vietnam

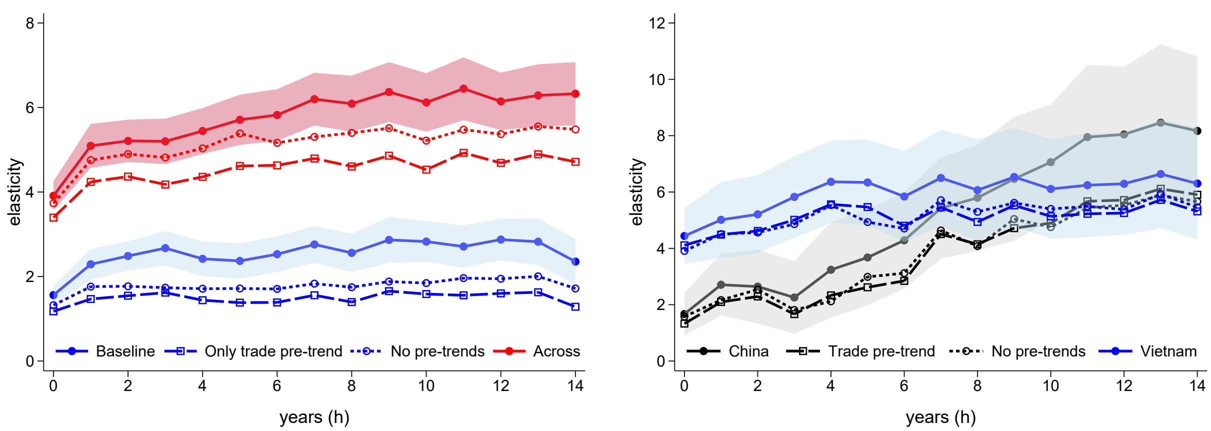

Notes: Panel (a) plots the results of (6) and panel (b) the results of (7). In both figures the solid line and the shaded area are the baseline estimates of those equations and their 95 percent confidence interval, with standard errors clustered at the jg level. The other lines are estimates from modified versions of the baseline in which consider we only include lagged import growth (long dash) and neither lagged import growth nor lagged tariff changes (short dash).

Figure A11: Gap-elasticities: Robustness Gap Measure

(a) China

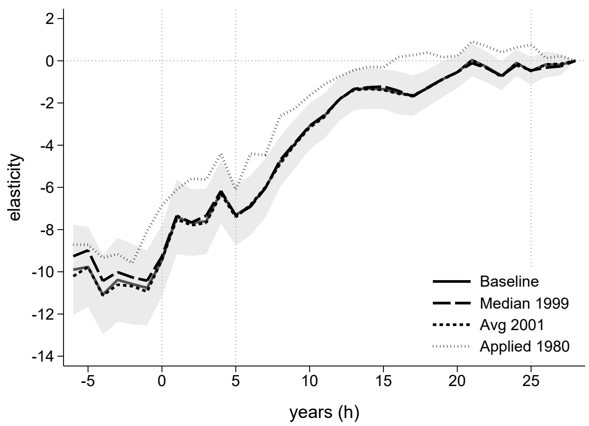

(b) Vietnam

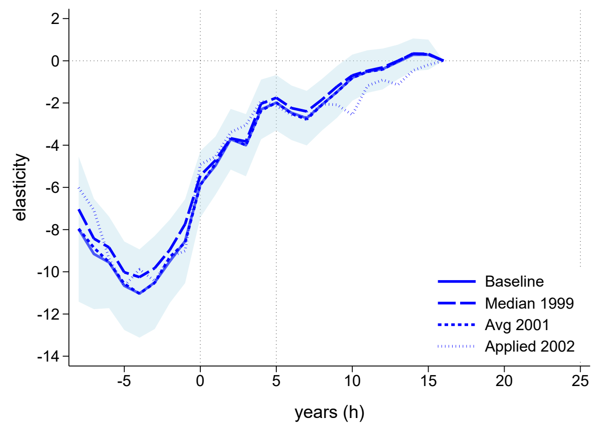

Notes: This figure plots the results of (9) under the baseline approach (solid line and shaded area for the confidence interval) and alternative measurements of the gap, as indicated in the legend. For illustrative purposes we separate the estimates for China (panel a) and Vietnam (panel b). The period $h = 0$ is 1980 for China and 2002 for Vietnam.

Figure A12: Gap-elasticities: Robustness Sample

(a) China

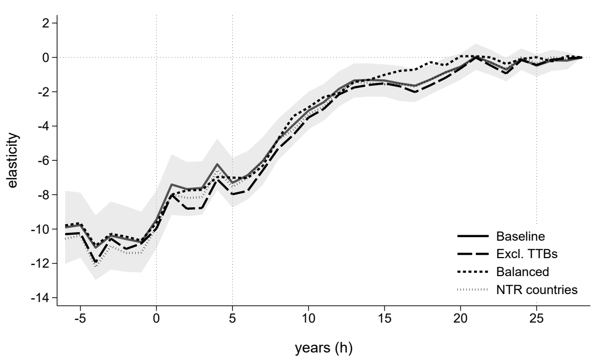

(b) Vietnam

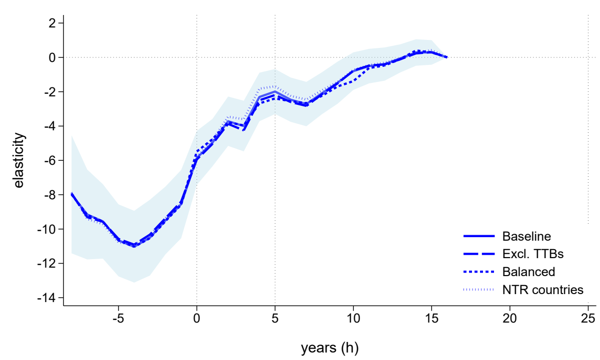

Notes: This figure plots the results of (9) under the baseline approach (solid line and shaded area for the confidence interval) and alternative sample design, as indicated in the legend. For illustrative purposes we separate the estimates for China (panel a) and Vietnam (panel b). The period $h = 0$ is 1980 for China and 2002 for Vietnam.

D Additional tables

Table A1: List of Unilateral Trade Preference Programs

Year	Program	Code
1976	Generalized System of Preferences (GSP)	A, A*, A+, 4
1982	Caribbean Basin Economic Recovery Act (CBERA)	E, E*, 9
1992–2014	Andean Trade Preferences Act (ATPA)	J, J*, J+
2001	African Growth and Opportunity Act (AGOA)	D
2004	Caribbean Basin Trade Partnership Act Initiative (CBTPA)	R
1999	Duty-free Code for West Bank and Gaza Strip	N
2017	Nepal Preference Program	NP

Notes: The *Year* column is the first (and last) year in which we observe the corresponding country sub-code in the US Census trade data, reported in the *Code* column.

Table A2: List of countries in a PTA with the United States

Year	Country	Code	Year	Country	Code
1985	Israel	IL	2006	Singapore	SG
1988	Canada	X, CA	2007	Peru	PE
1994	Mexico	MX	2010	Costa Rica	P, P+
2001	Colombia	CO	2010	El Salvador	P, P+
2001	Jordan	JO	2010	Dominican Republic	P, P+
2004	Australia	AU	2010	Guatemala	P, P+
2004	Chile	CL	2010	Honduras	P, P+
2006	Bahrain	BH	2010	Nicaragua	P, P+
2006	Morocco	MA	2010	South Korea	KR
2006	Oman	OM	2012	Panama	PA

Notes: The *Year* columns report the first full year the agreement went into effect. The *Code* columns are the country sub-codes reported in the US Census trade data.

Table A3: Example of the Regime Classification

Panel a: Census Data for SITC good

Provision Rate	Country Subcode	Tariff (%)	Import Share (%)	Regime
10	0	0.00	5	NTR
13	0	0.00	1	.
16	0	0.00	0	.
18	B	0.00	10	PTA
18	C	0.00	0	PTA
61	0	4.12	28	NTR
64	0	3.87	23	PTA
64	X	3.00	31	PTA
79	0	0.00	0	.
79	X	0.00	1	PTA

Panel b: 1990 HTS Tariff Schedule for corresponding HTS8 Products

HTS-8 Product	NTR rate (%)	PTA rate (%)
85011020	10	8
85011040	6.6	5.2
85011060	4.2	3.3
85012020	4.2	2.5
85012040	5	3
85012050	5	3
85012060	3.7	2.2
85014020	4.2	3.3
85014040	5	4
85014050	5	4
85014060	3.7	2.9
85015120	4.2	3.3
85015140	5	4
85015150	5	4
85015160	3.7	2.9
85015240	3.7	2.9
85015280	Free	
85015340	Free	
85015360	4.2	2.5
85015380	4.2	2.5

Notes: Based on imports of electric motors incl. universal A.C./D.C. motors (sitc code 71621) imported from Canada in 1990.