Recovering Credible Trade Elasticities from Incredible Trade Reforms*

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

November 2025

Abstract

We recover the elasticity of trade in response to an unanticipated, once-and-for-all tariff reform by estimating a dynamic heterogeneous-firm model on data from anticipated, uncertain reforms. Using 1974–2017 U.S. import and tariff data, we show that within-schedule tariff changes are frequent and transitory, while transitions across schedules are rare, persistent, and elicit far larger trade responses. We calibrate the model to U.S. trade with China and Vietnam—the countries that experienced the most substantial and persistent tariff reductions—and jointly estimate the parameters governing tariff expectations and firm-level adjustment dynamics. Simulating an unanticipated, once-and-for-all reform yields a short-run elasticity of about three and a long-run elasticity of about 15. Our results imply that reduced-form estimates based on typical transitory tariff variation substantially understate the long-run effects of permanent liberalizations.

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, 2024 Canadian International Trade Study Group, 2025 NBER Spring ITI, 2025 SED, and 2025 EIIT meetings for comments. We thank Yan Bai, Mark Bils, Brian McCaig, Fernando Parro, Anson Soderberry, Jim Tybout, 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 and Jiri Havel 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 and NBER

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 the key input into quantitative studies that 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 how trade responds to a *canonical reform*: an unanticipated, once-and-for-all tariff change. We argue that canonical reforms do not exist, and that reduced-form evidence on observed reforms is uninformative about this response. Instead, we recover the elasticity of trade to a canonical reform by using data from two unique episodes—the United States granting Normal Trade Relations (NTR) to Vietnam and China—as inputs to a structural model with time-varying policy risk. In the estimated model, a canonical trade liberalization generates a short-run trade elasticity of about three and a long-run trade elasticity of about 15. One-half of the gap between the short and long-run elasticity closes in five years.

The dynamic trade literature emphasizes that export participation is a forward-looking decision due to up-front investments (Baldwin and Krugman, 1989; Roberts and Tybout, 1997; Das et al., 2007) and back-loaded returns (Ruhl and Willis, 2017; Alessandria et al., 2021). This means that aggregate trade responds gradually to policy changes and depends on expectations about future policy as well as current policy (Ruhl, 2011; Handley and Limão, 2015), which makes it difficult to measure

[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.

the trade elasticity using data alone; one must disentangle the lagged effects of past reforms from the effects of expectations about future reforms. Our approach, which builds on Alessandria et al. (2025a), 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 US trade data for 1974–2017 to document tariff dynamics and estimate reduced-form responses of trade to tariff changes. Throughout this period, the United States maintained two main tariff schedules: Non-Normal Trade Relations (NNTR) and Normal Trade Relations (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 or 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 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, likely because these switches are more anticipated than other changes.[2]

[2]PTAs are the product of years of public negotiations and involve a gradual phase-out of tariffs. Countries typically leave the UTPP when they cross predetermined income per capita thresholds, making these departures predictable.

We focus on two important transitions between the NNTR and NTR schedules, those of China and Vietnam, which followed similar tariff paths at different points in time. China gained conditional NTR status in 1980 subject to annual renewal, and was later granted permanent NTR status (PNTR) after joining the World Trade Organization (WTO) in 2001. Vietnam went through the same steps, gaining conditional NTR status in 2002 and PNTR status when it joined the WTO in 2007. Ex post, these two reforms were the largest and most persistent in US trade history. Ex ante, there was a great deal of uncertainty about whether these reforms would occur and how long they would last. Importantly, both countries had also been subject to a complete trade embargo before gaining access to US markets at NNTR tariff rates, which allows us to study the transition from autarky and cleanly control for initial conditions.

The responses of trade to the Vietnam and China liberalizations share some similarities but also have some important differences. Starting a few years before gaining NTR and ending several years after gaining PNTR, we find that trade grows by about 10-11 times the change in tariffs in both countries, which is substantially greater than the estimate for the average tariff schedule change. The two countries' transition paths differ, though, 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 believe this is consistent with the historical evidence. China was one of the first non-market 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. (2025a) to estimate a path of elasticities from a surprise once-and-for-all 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 speed at which new exporters grow and continue in export markets, which both determines the rate at which firms discount future profits and contributes to the gradual adjustments in aggregate trade. These "technological" parameters, which are typically disciplined using firm-level data on entry, exit, and growth over the course of exporter life cycles (See, e.g., Alessandria et al., 2021), determine how trade responds to a canonical reform.

The key identification challenge is that these firm-level data moments are generated from an uncertain trade-policy environment and thus cannot separately be identified from the parameters that govern the stochastic process for tariffs. Our approach is to jointly estimate both sets of parameters simultaneously to match firm-level data and the path of aggregate trade. This is where the well-specified nature of U.S. trade policy towards China and Vietnam during our sample period works to our advantage: the policy process can be represented simply by a time-varying probability of moving between the NNTR and NTR tariff schedules.

The identification of the policy-regime transition probabilities comes from the variation in the NNTR-gap elasticities. The NNTR gap measures 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 schedule transitions, 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 just before the switch actually occurs, whereas in China the probability remains stable. 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 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 unanticipated, once-and-for-all liberalizations. Instead, they gradually gained credibility over time, and in the case of Vietnam there was a material degree of anticipation.

We also find similarities and differences in the technological parameters across the two countries. The costs of entering and continuing to export are larger in Vietnam than in China, reflecting a lower level of export participation, and the potential increase in export capacity that can occur over an exporter's life cycle is larger, reflecting a greater degree of dispersion in export sales between entrants and incumbents.

Using our calibrated model, we study the canonical unanticipated, once-and-for-all reform. We start the model in a steady state with NNTR tariffs and no chance of moving to the NTR regime, and switch regimes unexpectedly with no chance of reversal, and let the model converge to the new steady state. Despite the material differences in technological parameters, we recover very similar trade-elasticity dynamics for both countries: a long-run trade elasticity of about 15 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 shows that reduced-form estimates of the long-run trade elasticity are biased downward. Our empirical estimates for China and Vietnam are biased downward by about 25 percent and the biases for other reforms are much larger.

2 Relation to the literature

Our primary contribution is to estimate the canonical trade elasticity using a dynamic structural model with stochastic trade policy. While there is a vast literature measuring how trade responds to changes in trade policy, previous studies in this literature either use reduced-form empirical methods or static structural models.[3] We argue that neither of these approaches can recover the canonical trade elasticity.

A key challenge in measuring how trade responds to changes in policy is the most-favored-nation principle, which implies that there is a dearth of cross-country variation in tariffs that can be used to identify these responses. 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. 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.[4]

A potential concern with this approach is that the path of tariffs 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. Their trade-elasticity estimates are much lower than ours: less than one in the short run and less than two in the long run. However, when we focus on the same within-NTR tariff changes, we obtain similarly low estimates. This highlights an

[3] There are several recent estimations of the trade elasticity using static trade models such as Brancaccio et al. (2020) and Fieler and Eaton (2025).

[4] 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. Similarly, Besedes et al. (2020) find delayed effects of import growth in Mexico and Canada. Caliendo and Parro (2014), also studying NAFTA, estimate a trade elasticity of 4.5 for Mexico.

important external validity concern: estimates 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.

Numerous papers estimate the dynamic response of trade to changes in the exchange rate or other relative prices at the aggregate- or industry-level (e.g. Houthakker and Magee, 1969; Gallaway et al., 2003). Much of this work also finds a delayed response, although these estimates tend to be much lower than those for tariffs. For example, Alessandria and Choi (2021) use US net trade flows, relative prices, and expenditures to estimate a quarterly short-run elasticity 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.

More broadly, our paper also relates to earlier work that uses dynamic trade models to estimate the structure of trade costs using firm-level data on export participation (Roberts and Tybout, 1997; Das et al., 2007; Alessandria et al., 2021; Steinberg, 2023; Morales et al., 2019). Our contribution to this literature is to recognize that these data are generated from an environment in which trade policy is stochastic, and that trade costs must be jointly estimated alongside the policy process itself. We are the first to attempt such an estimation.[5]

3 Empirical evidence

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.

[5] In Alessandria et al. (2025a) and Alessandria et al. (2025b), we estimated the technological parameters that govern firm-level trade dynamics separately from the aggregate policy process, in a steady state rather than along the transition.

We emphasize that unanticipated, once-and-for-all trade reforms—the canonical reform studied in most models—do not exist in the data. Most tariff changes are highly transitory, and more persistent reforms are often characterized by anticipation due to phase-outs or long negotiation periods. We identify two particular tariff changes that were extremely large, highly persistent ex-post, and have several other features that make them well-suited to use in our structural approach to recovering canonical elasticities.

3.1 Overview and History

The United States has a unique tariff structure that creates tariff variation that facilitates the estimation of trade elasticities. Imports into the US generally enter under one of four tariff regimes: Normal Trade Relations (NTR, or Column 1); Non-Normal Trade Relations (NNTR, or Column 2); the Unilateral Trade Preference Program (UTPP); and Preferential Trade Agreements (PTA).[6]

Most US imports fall under the NNTR and NTR regimes. The NNTR tariff schedule is largely determined by the Smoot-Hawley Act of 1930 and is often treated as a source of exogenous variation (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 phase-out 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.[7] With the 1951 Trade Act, all non-market economies, except Yugoslavia, were transitioned from the NTR to the NNTR schedule. Of the 29 countries, including successor countries, subject to NNTR under the 1951 Trade Act or subsequent legislation, only

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

[7]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.

North Korea and Cuba still have NNTR status at the end of our sample, although trade with both countries is subject to a complete embargo.

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. Countries graduate from the program upon reaching development or market-share thresholds, acceding to other trade agreements, and other circumstances.[8,9] The program has lapsed several times, transitioning all participating countries to other regimes. 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 most notable being the 1989 Canada-US Free Trade Agreement and the 1994 North American Free Trade Agreement (NAFTA). By 2017, the last year in our sample, the United States had 20 PTAs. PTAs often involve domestic content requirements, leading some goods to be imported under both PTA and NTR rates during the same year.[10]

3.2 Data

We use annual customs data from the US Census Bureau from 1974 to 2017.[11] 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

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

[9]While GSP is the best-known temporary trade program, there are several other programs, which we describe in the appendix.

[10]See Head et al. (2024) for a discussion of this issue in the context of NAFTA's rules of origin for the automotive sector.

[11]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 dealing with the complications stemming from tariffs imposed during the 1st and 2nd Trump administrations, although we show some aggregated series for tariffs and trade through 2023 for completeness. The Trump tariffs fall outside of the schedules that had been set legislatively by Congress and do not fit into our "within-regime vs. across-regime" classification system. In Alessandria et al. (2025b) we study the evolution of trade and trade policy expectations in the window around the first US-China trade war that started in 2018.

a long horizon, we aggregate the data to the SITC 5-digit level.[12] 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. As our data features imports of the same good-country under different rates of provisions and country sub-codes, we assign goods to regimes if more than 50 percent of the value of the triplet is imported under it.[13]

Our final sample, summarized in Table 1, is an unbalanced panel with 44 years, 163 countries, 2,032 goods, and 2,105,521 observations. 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.

[12] 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).

[13] We describe our classification scheme in detail with examples in the appendix.

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.[14] 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 takes effect. The share of trade with NNTR countries remains near zero throughout, and there are no NNTR imports after 2005.[15]

Panel (c) plots imports coming from countries that were initially in the NNTR and UTPP regimes in the 1970s relative to NTR imports. Imports from initially-UTPP countries rise from two percent of NTR imports in 1976 to almost ten percent in 2017, while imports from initially-NNTR countries rise from zero to nearly 25 percent. Panel (d) shows the decline in tariffs for these groups. Tariffs for initially-NNTR countries started substantially above NTR rates and then fell in several discrete steps, resulting from a country or set of countries being transitioned to NTR or UTPP. The rise in trade with countries that transitioned away from NNTR is a key component of US import growth that we seek to capture.

3.3 Tariff dynamics across vs. within regimes

We begin our empirical analysis by comparing the properties of tariff changes across versus within regimes. Table 2 reports the annual persistence of regimes at the country (j) and country-product (jg) levels. All four regimes are highly persistent, but we observe multiple transitions away from each— no regime should be viewed as permanent.[16] Table 3 reports summary statistics for tariff changes. Most tariff

[14] The statutory NNTR rates are essentially constant over time. The movements shown in the figure are driven by changes in the composition of goods that are imported under that regime.

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

[16] If our data allowed us to go back further in time, we would observe more transitions. For example, starting in 1949 would add 20 more country-level transitions from NTR to NNTR and at least two from NNTR to NTR (Yugoslavia and Poland).

changes are small, with a one-year mean of -0.17 percentage points and a five-year mean of -0.74 percentage points. 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.

Tariff changes involving a regime change are much more persistent than within-regime tariff changes. To illustrate this, we measure the autocorrelation of tariff changes by estimating a 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 standard. In this setting, δ_{jt} captures common tariff changes within a country, absorbing the average effects of regime changes, and δ_{gt} captures common tariff changes across countries, absorbing tariff cuts resulting from GATT rounds and other multilateral negotiations.[17]

Figure 2(a) plots the paths of $\beta_h^{\tau,\text{within}}$ and $\beta_h^{\tau,\text{across}}$ for horizons $h = [-5, 10]$, along with estimates from a pooled specification without the interaction terms. Tariff changes across regimes are significantly more autocorrelated than within-regime changes. The within-regime autocorrelation is very similar to the pooled autocorrelation because the sample mostly consists of within-regime changes, which implies that the pooled estimates are poor predictors of the dynamics of across-regime tariff changes.

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

3.4 Case studies: China and Vietnam NTR access

We now describe the dynamics of US trade policy towards China and Vietnam, the case studies in our quantitative analysis. Figure 3(a) shows how tariffs on China and Vietnam changed over time by plotting the mean and interquartile range of the ratio of the NTR rate to the two countries' applied rates,

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

This measure is bounded between zero, during the embargo, and one, after NTR status is gained. In both countries, we see a long embargo followed by substantial tariff dispersion during the NNTR regime, which then subsides upon gaining NTR access.[18] When scaled this way, these two NNTR-to-NTR transitions look very similar to once-and-for-all trade liberalizations.

In fact, these two reforms were the most persistent in US trade history. To illustrate this, we estimate their autocorrelations using the following specification,

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

Figure 3(b) shows that the changes in tariffs on China and Vietnam that occurred when these countries gained NTR status were much more persistent than the typical across-regime tariff change; 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.

In addition to being the closest to canonical reforms in recent US trade history, China and Vietnam's NNTR-to-NTR transitions were also two of the most important in terms of aggregate trade growth. Figure 3(c) shows that each country's share of US

[18] 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.

imports starts to rise sharply after gaining conditional NTR access and continues to rise all the way until the first US-China trade war begins in 2018, when China begins to lose US market share and Vietnam's growth accelerates. One other important feature of these reforms is that because both countries began in embargo, all cumulative trade growth is driven by the extensive margin. Figure 3(d) illustrates this by plotting the share of SITC goods with positive trade. This is another reason that these case studies are particularly well-suited to study using a model of export participation dynamics.

Ex post, China and Vietnam's NNTR-to-NTR transitions are the closest empirical counterparts to canonical unanticipated, once-and-for-all reforms in our sample, but this was not how these reforms were perceived ex ante. Uncertainty about the duration of China's NTR status has been the focus of numerous studies. In Alessandria et al. (2025a) we show that this status was initially expected to be quickly revoked, and we show below that the same is true for Vietnam, although there are some material differences in the trade-policy risks the two countries faced. Thus, even for these reforms, reduced-form empirical estimates of trade responses may be quite different than the response to a truly canonical reform.

4 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 3, 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.

4.1 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 the local-projections approach of Boehm et al. (2023),

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

where $\Delta_h \tau_{jgt}$ is instrumented by $\Delta_0 \tau_{jgt}$ for $h > 0$ in order to control for subsequent tariff changes that are correlated with the change in t. The vector of controls, $\mathbf{Z}_{jgt}$, includes the preceding tariff change, $\Delta_{-1}\tau_{jgt}$, as well as the preceding import growth, $\Delta_{-1}v_{jgt}$. We include the same set of fixed effects as in section 3, 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. 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 2(b). 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.

We then contrast these pooled estimates with 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). We allow for the pre-trend controls to vary with the interaction of interest. Figure 2(b) also plots 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.

4.2 Case studies: China and Vietnam NTR access

We now turn our attention to the effects of the NNTR-to-NTR transitions that China and Vietnam underwent in 1980 and 2002, which we have argued are the most canonical reforms in US trade history. Leveraging the heterogeneity in tariff changes across goods shown in Figure 3(a) , we estimate the dynamics of trade following these reforms, as well as in the years leading up to them.

Following Alessandria et al. (2025a), 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*:

$$X_g = \log(1 + \tau_g^{\text{NNTR}}) - \log(1 + \tau_g^{\text{NTR}}), \tag{7}$$

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 to measure tariff risk, as it represents the good-specific tariff hikes that would occur if a country was shifted back to NNTR. However, it also measures the tariff reductions that occurred when NTR status was first granted. 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. Figure 3(e) plots the NNTR gap distribution.

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

$$\begin{aligned} v_{jgt} = &\sum_{t'=1974}^{2008} \beta_t^{v,\text{CN}} \mathbb{1}\{t = t'\} \mathbb{1}\{j = \text{China}\} X_g \\ &+ \sum_{t'=1994}^{2017} \beta_t^{v,\text{VN}} \mathbb{1}\{t = t'\} \mathbb{1}\{j = \text{Vietnam}\} X_g + \delta_{jt} + \delta_{jg} + \delta_{gt} + u_{jgt}. \end{aligned} \tag{8}$$

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}, country-good fixed effects, δ_{jg}, and good-year fixed effects, δ_{gt}. This fixed-effects structure is common in the literature on China's NTR access (Pierce and Schott, 2016; Handley and Limão, 2017).

It is important to emphasize that in these event studies, the country-year fixed effects, δ_{jt}, absorb not only aggregate shocks in exporting countries, but also the responses to the end of the US embargoes on China and Vietnam. The lifting of the embargoes was effectively an infinite tariff reduction that triggered large, gradual

[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.

trade adjustments that were still well underway when these countries later gained NTR status. As a result, our empirical approach measures the effect of NTR access on exports to the United States above and beyond this ongoing post-embargo growth, comparing products that experienced tariff reductions from NTR access with those that did not. In the appendix, we present an alternative specification that explicitly controls for aggregate fluctuations using data on GDP and exports to non-U.S. markets. That approach allows the gap elasticities ($\beta_t^{v,CN}$ and $\beta_t^{v,VN}$) to capture the combined effects of NTR access and embargo removal. Our model is consistent with both approaches, but our baseline estimation targets the specification with country-year fixed effects.

Figure 3(f) shows the estimates of β_t^{CN} and β_t^{VN}.[20] 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 11.8 for China and 11.9 for Vietnam. 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, for China, high- and low-gap goods were growing 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).

[20]The appendix compares our event-study approach with a local-projections specification like (6) that focuses specifically on NTR access as in (3). While the latter works well for China, it significantly understates the long-run trade response for Vietnam due to the anticipatory trade growth in the years leading up to NTR access.

5 Model

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). 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 a constant-returns to-scale technology, $y_t = z_t \ell_t$. Productivity, z_t, evolves over time according to an AR(1) process in logs: $\log z_t = \rho_z \log z_{t-1} + \sigma_z \varepsilon_t$, 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: $d_t(p_t, \tau_t) = (p_t \tau_t)^{-\theta}$, 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. 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:

$$\pi_t(z_t, \xi_t, \tau_t) = \max_p \left\{ p d_t(p, \tau_t) - w_t \frac{\xi_t d_t(p, \tau_t)}{z_t} \right\}, \tag{9}$$

which is equal to zero if $\xi_t = \infty$. 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

$$V_t^1(z_t, \xi_t, \tau_t) = -f(\xi_t) + \frac{\delta(z)}{1+r} \mathbb{E}_{z,\xi,\tau}\left[V_{t+1}(z_{t+1}, \xi_{t+1}, \tau_{t+1})\right], \tag{10}$$

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

$$V_t^0(z_t, \xi_t, \tau_t) = \frac{\delta(z)}{1+r} \mathbb{E}_{z,\tau}\left[V_{t+1}(z_{t+1}, \infty, \tau_{t+1})\right], \tag{11}$$

where

$$V_t(z_t, \xi_t, \tau_t) = \pi_t(z_t, \xi_t, \tau_t) + \max\{V_t^0(z_t, \xi_t, \tau_t), V_t^1(z_t, \xi_t, \tau_t)\}. \tag{12}$$

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.

Aggregation. Aggregate exports are given by

$$X_t = \sum_{\xi \in \{\xi_L, \xi_H\}} \int_z p_t(z, \xi, \tau_t) d_t(p_t(z, \xi, \tau_t), \tau_t) \, \mathrm{d}\varphi_t(z, \xi), \tag{13}$$

where $\varphi_t(z, \xi)$ denotes the joint distribution of firms indexed by their productivity and variable trade cost. This distribution evolves according to the following law of motion:

$$\varphi_t(\mathcal{Z}, \infty) = \sum_{\xi} \left[\int_0^{\bar{z}_{t-1}(\xi)} h(\mathcal{Z}, z) \varphi_{t-1}(z, \xi) dz + \int_0^{\infty} \bar{h}(\mathcal{Z}) \varphi_{t-1}(z, \xi) dz \right], \tag{14}$$

$$\begin{aligned} \varphi_t(\mathcal{Z}, \xi_H) = & \int_{\bar{z}_{t-1}(\infty)}^{\infty} h(\mathcal{Z}, z) \varphi_{t-1}(z, \infty) dz + \rho_\xi \int_{\bar{z}_{t-1}(\xi_H)}^{\infty} h(\mathcal{Z}, z) \varphi_{t-1}(z, \xi_H) dz \\ & + (1 - \rho_\xi) \int_{\bar{z}_{t-1}(\xi_L)}^{\infty} h(\mathcal{Z}, z) \varphi_{t-1}(z, \xi_L) dz, \end{aligned} \tag{15}$$

$$\varphi_t(\mathcal{Z}, \xi_L) = (1 - \rho_\xi) \int_{\bar{z}_{t-1}(\xi_H)}^{\infty} h(\mathcal{Z}, z) \varphi_{t-1}(z, \xi_H) dz + \rho_\xi \int_{\bar{z}_{t-1}(\xi_L)}^{\infty} h(\mathcal{Z}, z) \varphi_{t-1}(z, \xi_L) dz, \tag{16}$$

where $\mathcal{Z}$ is a typical subset of the space of productivity draws, $h(\mathcal{Z}, z)$ is the probability of surviving and drawing a new productivity in $\mathcal{Z}$ conditional on current productivity z, and $\bar{h}(\mathcal{Z})$ is the probability of dying and being replaced by a new firm with productivity in $\mathcal{Z}$.

6 Illustrating expectational biases

We now use the model to illustrate how trade-policy expectations complicate trade-elasticity estimation. First, we show that less-persistent tariff changes have lower long-run trade elasticities, both because the level of trade is higher beforehand and because the response afterward is smaller, and that reduced-form estimation methods do not recover canonical trade elasticities from data generated by non-canonical reforms. Second, we show that the short-run trade elasticity is higher when tariff changes are preceded by reductions in persistence. Throughout this section, we use the model parameters calibrated in the case study of China in the next section, which are listed in Table 4. This may seem out of order to the reader, but we believe that developing these ideas before explaining our identification strategy is useful, and that illustrating their quantitative significance as well as their qualitative significance is important.

6.1 Tariff persistence and the long-run trade elasticity

Consider a symmetric Markov process for tariffs with two values, $\tau_L = 0$ and $\tau_H = 30$ percent, and a symmetric probability of switching states, $1-\omega$. 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.

Figure 4 shows the dynamics of trade before and after a once-and-for-all tariff change for different values of ω. In all versions of this experiment, the economy is in the high-tariff state for many periods, switches to the low-tariff state in period $t = 0$, and remains there forever. While this hypothetical is unlikely, especially when ω is low and tariffs are expected to fluctuate frequently, it is useful to illustrate the economic forces at play; we discuss trade-elasticity estimation using simulated data with persistent vs. transitory tariff processes below. Panel (a) shows the discounted expected tariff,

$$\bar{\tau}_t^E = (1-\beta)^{-1}\mathbb{E}_t\left[\sum_{s=t+1}^{\infty}\beta^{s-t-1}\tau_s\right]. \tag{17}$$

Note that this measure excludes the current tariff rate, as only expectations about future tariffs matter for the export participation decision. Panel (b) shows export participation, panel (c) shows aggregate exports, and panel (d) shows the cumulative trade elasticity, ε_h, measured as in (4).

The $\omega = 1$ case is the canonical trade reform found in virtually all quantitative and theoretical trade studies: an unanticipated, permanent reform. 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: $\varepsilon_0 = \theta = 3.17$. 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 3.5 log points higher than before the reform. Trade follows a similar trajectory, and the long-run trade elasticity, ε_∞, is about 15. This *long-run canonical trade elasticity* is a function of θ as well as the parameters that govern firm-level trade dynamics, f_0, f_1, ξ_L, ξ_H, and ρ_ξ, that we will recover in our quantitative exercise.

When $\omega < 1$, the discounted expected tariff is lower than the applied tariff before the reform. This raises export participation and trade in the high-tariff state relative to the canonical case. Conversely, the discounted expected tariff falls less than applied tariffs once the economy switches to the low-tariff state, which depresses the response of export participation and trade. To highlight these effects, figure 4 plots export participation and trade in each case relative to their values in the $\omega = 1$ case. Both effects push the measured long-run trade elasticity below the canonical long-run elasticity. The $\omega = 0.65$ case, which is very close to i.i.d., is particularly instructive. In this case, where the discounted expected tariff is essentially constant, export participation barely reacts to tariff shocks and the trade elasticity is virtually identical to the demand elasticity, θ, at all horizons.

These once-and-for-all reforms illustrate how tariff persistence affects measured long-run trade elasticities, but they are inconsistent with the statistical properties of the assumed tariff process, especially for low levels of persistence where tariffs change frequently. To illustrate that this issue matters "in sample" under rational expectations, we apply the Boehm et al. (2023) empirical approach we used in section 4 to simulated model data. For each value of ω shown in Figure 4, we simulate a panel of 1,000 goods for 1,000 periods, with tariff shocks drawn from the assumed distribution independently across goods. We then estimate tariff autocorrelations using (1) (with indicators for ω values instead of regime switches) and trade-elasticity dynamics using the local-projections specification (5).

Panel (a) of Figure 5 shows the simulated tariff autocorrelations. Despite its simplicity, the two-state Markov process exhibits tariff dynamics that are similar to the data; within-regime and across-regime changes are approximated well by Markov processes with $\omega = 0.6$–0.7 and $\omega = 0.8$–0.9, respectively. Panel (b) shows our simulated trade-elasticity estimates, benchmarked against the trade elasticity path for the canonical reform as in the previous figure. Overall, the "in-sample" elasticities estimated using local projections are similar to the calculations from the simple one-off reform episodes. Even when $\omega = 0.99$, which is virtually indistinguishable from the canonical case in the data unless one has an extremely long sample, the long-run elasticity estimate is lower than the canonical long-run elasticity.

Panel (c) of Figure 5 summarizes these results by plotting the ratio of the long-run trade elasticity to the short-run trade elasticity against tariff persistence ω. For low values of ω, this ratio is barely above one and essentially flat. Once ω exceeds 0.8, it starts to rise sharply, reaching a value of 4.5 in the canonical ($\omega = 1$) case. This confirms that reduced-form empirical approaches yield estimates of the long-run trade elasticity that are far below the canonical trade elasticity unless they are applied to extremely persistent reforms—much more persistent than the typical regime switch in our sample. This underscores the value of the China and Vietnam episodes as starting points for recovering canonical elasticities.

6.2 Persistence shocks and the short-run trade elasticity

We now use our model to illustrate what happens when changes in the expected discounted tariff precede changes in the applied tariff. This is a different form of anticipation than in the previous analysis. We now suppose that the economy rests in a perfectly persistent high-tariff steady τ_H with $\omega = 1$ for many years and then experiences a tariff cut to τ_L that may be coupled with an unanticipated decline in persistence to $\omega = 0.8$. We present four versions of this experiment with different timings in Figure 6.

The first version (labeled $\tau \downarrow$) is a drop in tariffs from τ_H to τ_L at $t = 0$ without a change in persistence. This is the canonical surprise, once-and-for-all reform we have already discussed; we repeat it here to use as a benchmark against which to compare the other versions. The second version (labeled $\omega \downarrow$) is an unanticipated drop in persistence from $\omega = 1$ to $\omega = 0.8$ without a tariff cut. Export participation begins to rise starting in the following period, generating an eventual increase in trade of about 1.5 log points, but since there is no change in tariffs, we cannot measure the trade elasticity.

The third version (labeled $\tau, \omega \downarrow$) is a transition from τ_H to τ_L coupled with a simultaneous decline in persistence from $\omega = 1$ to $\omega = 0.8$ in $t = 0$. This lowers the current tariff, but since this change is expected to be less persistent than the canonical reform, the expected discounted tariff falls much less. The long-run response of trade is about halfway between the first two reforms' responses. Note that there is no anticipatory effect since the change in persistence occurs alongside the change in tariffs, so the short-run trade elasticity is still equal to the demand elasticity $\theta = 3.17$.

The last version (labeled $\omega \downarrow \rightarrow \tau \downarrow$) is an unanticipated drop in persistence from $\omega = 1$ to $\omega = 0.8$ in $t = -1$ followed by a cut in tariffs from τ_H to τ_L in $t = 0$. The long-run response to this two-step reform is the same as in the third reform, but the short-run trade elasticity is now higher than the demand elasticity. This is because some firms make investments in market access in $t = -1$ when they learn that a future tariff cut is more likely, pushing up export participation in $t = 0$ when the tariff cut indeed materializes.

This kind of anticipatory effect is the mechanism that we believe to be at play in generating larger short-run responses to regime switches in section 4.1 and Vietnam's NTR access in section 4.2. Indeed, tariff-schedule shifts 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. We show how phased-in reforms can dramatically increase short-run trade elasticities in the appendix.[21,22] 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. Our quantitative analysis interprets the data for Vietnam and China as indicating that the likelihood of NTR access increased during the former's NNTR period but not for the latter's.

7 Quantitative analysis

We now take our model to the data, which we have shown are contaminated by anticipation of and uncertainty about tariff changes, and use it to measure how trade would respond in the absence of these expectational biases. Our strategy is to jointly discipline the model's technological parameters and its stochastic process for trade policy with empirical evidence from two trade reforms that we have argued are uniquely well-suited to this task: the United States granting NTR status to China in 1980 and Vietnam in 2002. We then use the model to simulate surprise once-and-for-all reforms to measure the canonical trade-elasticity path.

7.1 Environment and calibration

To exploit the heterogeneity in tariff and trade dynamics across goods documented above, we consider a multi-good version of the model in section 5. There are G goods, which correspond to the 5-digit SITC goods in our data. In each good $g = 1, \ldots, G$,

[21] Khan and Khederlarian (2021) document evidence for these effects in NAFTA.

[22] Likewise, our analysis in this section assumes the economy starts from a steady state, but initial conditions can also complicate trade-elasticity estimation, particularly if the economy is still adjusting to an earlier reform like the lifting of an embargo. We discuss this issue in the appendix as well.

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.[23] The basic structure of the economy is the same for both countries but we allow parameter values and timing to differ. For clarity, we omit country labels.

Tariffs, $\tau_g(s)$, differ across goods and trade-policy regimes, $s \in \{\text{NNTR}, \text{NTR}\}$. We take the tariff rates directly from the data. A country's tariff regime, which is an aggregate state, follows a time-varying Markov process with transition probabilities

$$\Omega_t = \begin{bmatrix} \omega_t^{\text{NNTR}} & 1 - \omega_t^{\text{NNTR}} \\ 1 - \omega_t^{\text{NTR}} & \omega_t^{\text{NTR}} \end{bmatrix}, \tag{18}$$

where ω_t^s is the persistence of state s in year t.[24]

Our calibration strategy builds on Alessandria et al. (2025a). 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 6. However, we set the demand elasticity, θ, to the average estimate in Soderbery (2018) for US imports from China, which is 3.17. Soderbery's estimate comes from price and quantity variation and is therefore closely related to our demand elasticity parameter.[25] We use this estimate rather than the reduced-form short-run trade elasticity, which depends on changes in trade policy expectations (section 6.2).

Second, we calibrate the other technological parameters, σ_z, f_0, f_1, and ξ_H, to

[23] We explore heterogeneity across goods in these parameters for China in Alessandria et al. (2025a). There are too few firms in many sectors in Vietnam to reliably estimate exporter-dynamics.

[24] We follow Alessandria et al. (2025b) and assume that firms believe the current transition probabilities will remain in force forever, and that firms are surprised each period when these probabilities change. In Alessandria et al. (2025a), we studied a model where firms have perfect foresight over the entire path $\{\Omega_t\}_{t=0}^{\infty}$ and showed that the results are similar under both approaches.

[25] We have tried estimating θ directly, but it cannot be separately identified from the probability of gaining NTR status $1 - \omega_t^{NNTR}$, because both parameters affect the initial jump in trade when that occurs in the same direction. Note that this externally-assigned θ value works well for both countries despite the fact that the observed short-run elasticity is much higher for Vietnam, which our model interprets as a growing likelihood of NTR access, generating anticipatory growth in export participation. Moreover, as we discuss in the appendix, θ must be close to Soderbery (2015)'s estimate for the model to generate the observed NTR-gap elasticity dynamics.

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. We measure these moments in the model in the same years that we observe them in the data (2004–2007 for China and 2010–2014 for Vietnam).[26] 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 4 lists the target moments and calibrated values of the technological parameters for each country.

Third, we calibrate the sequence of trade-policy transition matrices $\{\Omega_t\}_{t=0}^{\infty}$ to match the NNTR-gap elasticity dynamics described in section 4.2. ω_t^{NNTR} is identified by the dynamics in 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 indicates an increase in the likelihood of retaining NTR status. The main difference between our strategy and that of Alessandria et al. (2025a) 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 7(a) shows the fit of the model against these target moments.[27]

Fourth, we start the model from the embargo, i.e., we set export participation to zero, and then lift the embargo unexpectedly in the appropriate year for each country.[28]

One material difference between our calibration strategy in this paper and in our

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

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

[28] Even if firms believe there is a chance the embargo may be lifted, they do not enter before it is lifted unless this probability is extremely high.

previous work Alessandria et al. (2025a,b) is that here we jointly calibrate the technological parameters and the policy process, whereas in those papers, we separately calibrated the technological parameters in the long-run steady state with NTR tariffs and no risk of returning to the NNTR regime. In the case of China, which was the focus of those papers, this approach was relatively innocuous because the data moments are measured more than two decades after NTR access was granted, at a time when policy risk was widely believed to be minimal. In the case of Vietnam, it would be more problematic because the data moments are measured at a time when the economy was still adjusting to NTR access and, as we estimate, the risk of returning to NNTR was higher. We view this new approach as a methodological improvement.

Figure 7(b) plots the calibrated regime transition probabilities, $1 - \omega_t^s$. For China, the 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. (2025a) 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.

7.2 Canonical trade elasticities

We now use the calibrated model to measure the trade elasticity following a canonical tariff reform that features neither anticipation nor uncertainty and is not affected by initial conditions. For each country, we first solve for the steady-state with permanent NNTR tariffs. We then introduce a surprise once-and-for-all switch to the NTR tariff schedule and solve for the transition to the new steady state. The canonical trade elasticity at horizon h is the difference in the canonical reform between the NTR-gap elasticity h periods after NTR access and its initial value. The long-run canonical trade elasticity is the difference between the gap elasticities in the new steady state with NTR tariffs and the initial steady state with NNTR tariffs. We find that this elasticity is 14.67 for China and 15.29 for Vietnam. We report the sensitivity of these estimates to the model parameters in the appendix.

Figure 7(c) compares the NTR-gap elasticity dynamics in the canonical NTR reforms to the baseline dynamics. There are three key differences. First, the pre-NTR gap elasticities in the baseline model are 4–6 log points smaller than the gap elasticities in the NNTR steady state. 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, the economy is not in a steady state but transitioning from the embargo, biasing 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 more 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.[29] 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 7(b); note there is no

[29] Alessandria et al. (2025a) discuss the geopolitical background of the "stalled" period.

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 and raises the short-run elasticity we recover from Vietnam.

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

$$\lambda_h = \frac{\log V_\infty - \log V_{t+h}}{\log V_\infty - \log V_{t-1}}, \tag{19}$$

where V_∞ is the level of trade in the new steady state with no uncertainty. We plot this gap in Figure 7(d) and fit a geometric function on the latter half of the sample. 22–23 percent of the long-run change occurs in the first year, which captures the intensive-margin effect driven by the demand elasticity θ. Another 40–45 percent occurs in the second year, reflecting the one-period delay on the extensive margin. After that initial two-year burst, we find that half the distance to the long-run closes every five years.

Table 5 compares the long-run canonical trade elasticities to our reduced-form estimates, both for US data from section 4.1 and the model simulation in section 6.1. The reduced-form gap-elasticities for China and Vietnam NTR access are 20 and 22 percent smaller than their canonical counterparts, respectively. The estimates for other regime changes are about sixty percent smaller, and the within-regime estimates are about 80 percent smaller. In our simulations, even extremely persistent tariff changes ($\rho = 0.99$) have long-run trade elasticity estimates ten percent smaller than our canonical values. We conclude that the typical transitory trade reforms observed in the data are uninformative about the canonical long-run trade elasticity, and that even highly-persistent reforms yield underestimates unless filtered through a structural model.

8 Conclusions

We estimate the long-run trade elasticity to a canonical unanticipated, once-and-for-all tariff change is 15. This is five times what one finds when applying standard reduced-form empirical methods to US data and about three times the trade elasticity common in quantitative work (Simonovska and Waugh, 2014; Caliendo and Parro, 2014). This large response is hard to find in the data due to the stochastic properties of trade policy and forward-looking decisions of firms related to anticipation, uncertainty, and initial conditions that are not captured by existing methods.

Our estimated elasticities are recovered from a model-based analysis that leverages key elements of the structure of US trade policy. We consider two case studies, China and Vietnam, which underwent extremely large liberalizations that were highly persistent ex-post but were characterized by considerable uncertainty ex ante owing to these countries' unique geopolitical relations with the United States. We model these reforms as two-state Markov processes with time-varying transition probabilities, and jointly calibrate these probabilities together with the technological parameters that govern trade adjustment dynamics.

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 tariff variation in the data is highly transitory, so the elasticities identified using this variation are much lower than the elasticities that should be used to model canonical 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 to account for the stochastic processes that tariffs follow in the data. Our approach is one example of such a strategy.

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.[30] Of course, perhaps it's not possible to credibly commit to a permanent trade reform, in which case researchers need to explicitly incorporate trade policy expectations into their analyses.

References

Alessandria, G. and O. Avila (2023): "Trade Integration, Industry Reallocation, and Welfare in Colombia," *IMF Economic Review*, 71, 649–687.

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

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

Alessandria, G., H. Choi, and K. J. Ruhl (2021): "Trade Adjustment Dynamics and the Welfare Gains from Trade," *Journal of International Economics*, 131, 1034–1058.

Alessandria, G., S. Y. Khan, A. Khederlarian, K. J. Ruhl, and J. B. Steinberg (2025a): "Trade-Policy Dynamics: Evidence from 60 Years of U.S.-China Trade," *Journal of Political Economy*.

——— (2025b): "Trade War and Peace: U.S.-China Trade and Tariff Risk from 2015–2050," *Journal of International Economics*, 155, 104066.

Baier, S. and J. Bergstrand (2007): "Do Free Trade Agreements Increase Members International Trade?" *Journal of International Economics*, 71, 72–95.

Baier, S. L., J. H. Bergstrand, and M. Feng (2014): "Economic Integration Agreements and the Margins of International Trade," *Journal of International Economics*, 93, 339–350.

Baldwin, R. and P. Krugman (1989): "Persistent Trade Effects of Large Exchange Rate Shocks," *The Quarterly Journal of Economics*, 104, 635–654.

Besedes, T., T. Kohl, and J. Lake (2020): "Phase Out Tariffs, Phase In Trade?" *Journal of International Economics*, 127, article 103385.

[30]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."

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

Brancaccio, G., M. Kalouptsidi, and T. Papageorgiou (2020): "Geography, Transportation, and Endogenous Trade Costs," *Econometrica*, 88, 657–691.

Broda, C. and D. E. Weinstein (2006): "Globalization and the Gains From Variety," *The Quarterly Journal of Economics*, 121, 541–585.

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

CRS Report for Congress (2005): "Normal-Trade-Relations (Most-Favored-Nation) Policy of the United States," Tech. Rep. RL31558, Congressional Research Service.

Das, S., M. J. Roberts, and J. R. Tybout (2007): "Market Entry Costs, Producer Heterogeneity, and Export Dynamics," *Econometrica*, 75, 837–873.

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

Feenstra, R. C., J. Romalis, and P. K. Schott (2002): "U.S. Imports, Exports, and Tariff Data, 1989–2001," *NBER Working Paper No. 9387*.

Fieler, A. C. and J. Eaton (2025): "The Margins of Trade," *Econometrica*, 93, 129–160.

Fitzgerald, D., S. Haller, and Y. Yedid-Levi (2023): "How Exporters Grow," *The Review of Economic Studies*, 91, 2276–2306.

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

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

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

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

——— (2022): "Trade Policy Uncertainty," *Annual Review of Economics*, 14, 363–395.

Head, K. and T. Mayer (2014): "Chapter 3 - Gravity Equations: Workhorse,Toolkit, and Cookbook," in *Handbook of International Economics*, ed. by G. Gopinath, E. Helpman, and K. Rogoff, Elsevier, vol. 4 of *Handbook of International Economics*, 131–195.

Head, K., T. Mayer, and M. Melitz (2024): "The Laffer curve for rules of origin," *Journal of International Economics*, 150, 103911.

Hillberry, R. and D. Hummels (2013): "Trade Elasticity Parameters for a Computable General Equilibrium Model," in *Handbook of Computable General Equilibrium Modeling*, ed. by P. B. Dixon and D. W. Jorgenson, Elsevier, vol. 1, 1213–1269.

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

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

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

Morales, E., G. Sheu, and A. Zahler (2019): "Extended Gravity," *The Review of Economic Studies*, 86, 2668–2712.

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

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

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

Ruhl, K. J. (2011): "Trade Dynamics under Policy Uncertainty," *American Journal of Agricultural Economics*, 93, 450–456.

Ruhl, K. J. and J. L. Willis (2017): "New Exporter Dynamics," *International Economic Review*, 58, 703–726.

Schott, P. K. (2021): "Data," https://sompks4.github.io/sub_data.html, accessed: 06-09-2021.

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

Soderbery, A. (2015): "Estimating import supply and demand elasticities: Analysis and implications," *Journal of International Economics*, 96, 1–17.

——— (2018): "Trade Elasticities, Heterogeneity, and Optimal Tariffs," *Journal of International Economics*, 114, 44–62.

Steinberg, J. B. (2023): "Export Market Penetration Dynamics," *Journal of International Economics*, 145, article 103807.

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

Figure 1: Import shares and tariffs by policy regimes

(a) Shares (regime in each year)

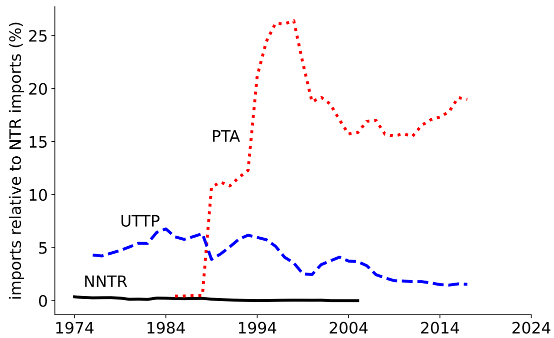

(b) Tariffs rate (regime in each year)

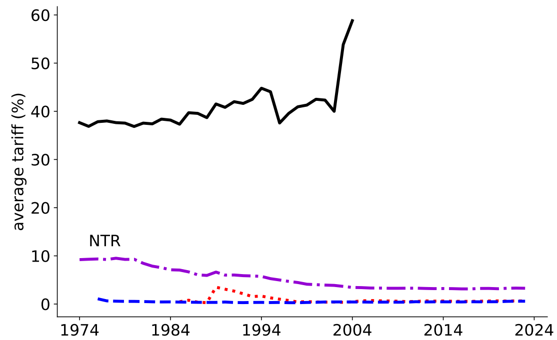

(c) Shares (initial regime)

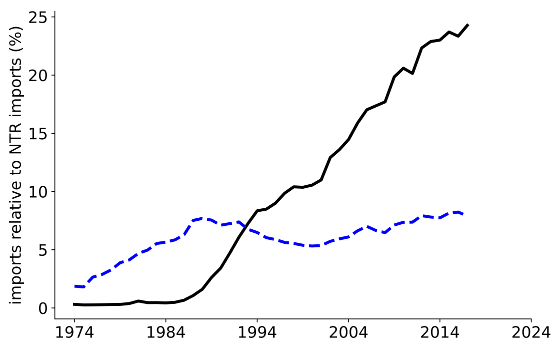

(d) Tariff rate (initial regime)

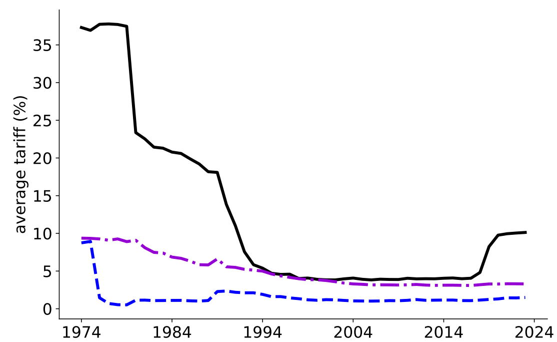

Notes: (a) Import share of each trade policy regime described in section 3 after dropping unclassified variety-years. (b) Average tariff rate by regime. (c) and (d) Analogues of (a) and (b) with each variety's regime fixed at its modal regime during 1970s.

Figure 2: Tariff and trade dynamics: across vs. within regimes

(a) Tariff autocorrelation

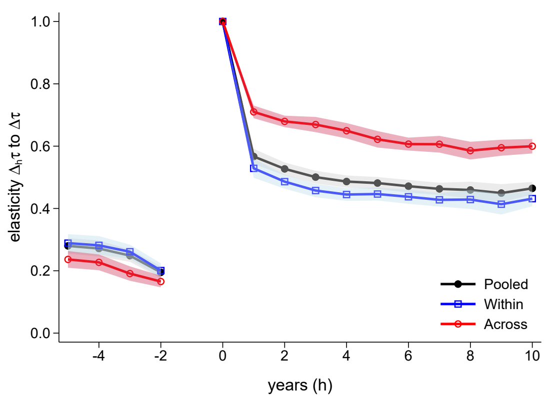

(b) Trade elasticities

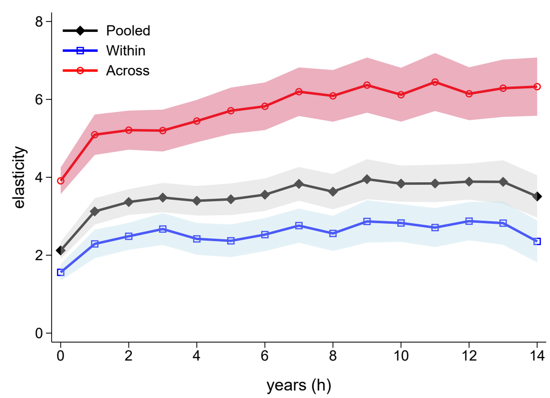

Notes: (a) Autocorrelation estimates from (1). (b) Elasticity estimates from (6). 95-pct. confidence intervals constructed using standard errors clustered at the jg level.

Figure 3: Tariff and trade dynamics: China and Vietnam

(a) Tariff dynamics **(b)** Tariff autocorrelation

(c) Share of US imports (percent) **(d)** Extensive margin dynamics

(e) Gap between NNTR and NTR tariffs **(f)** NNTR gap elasticities

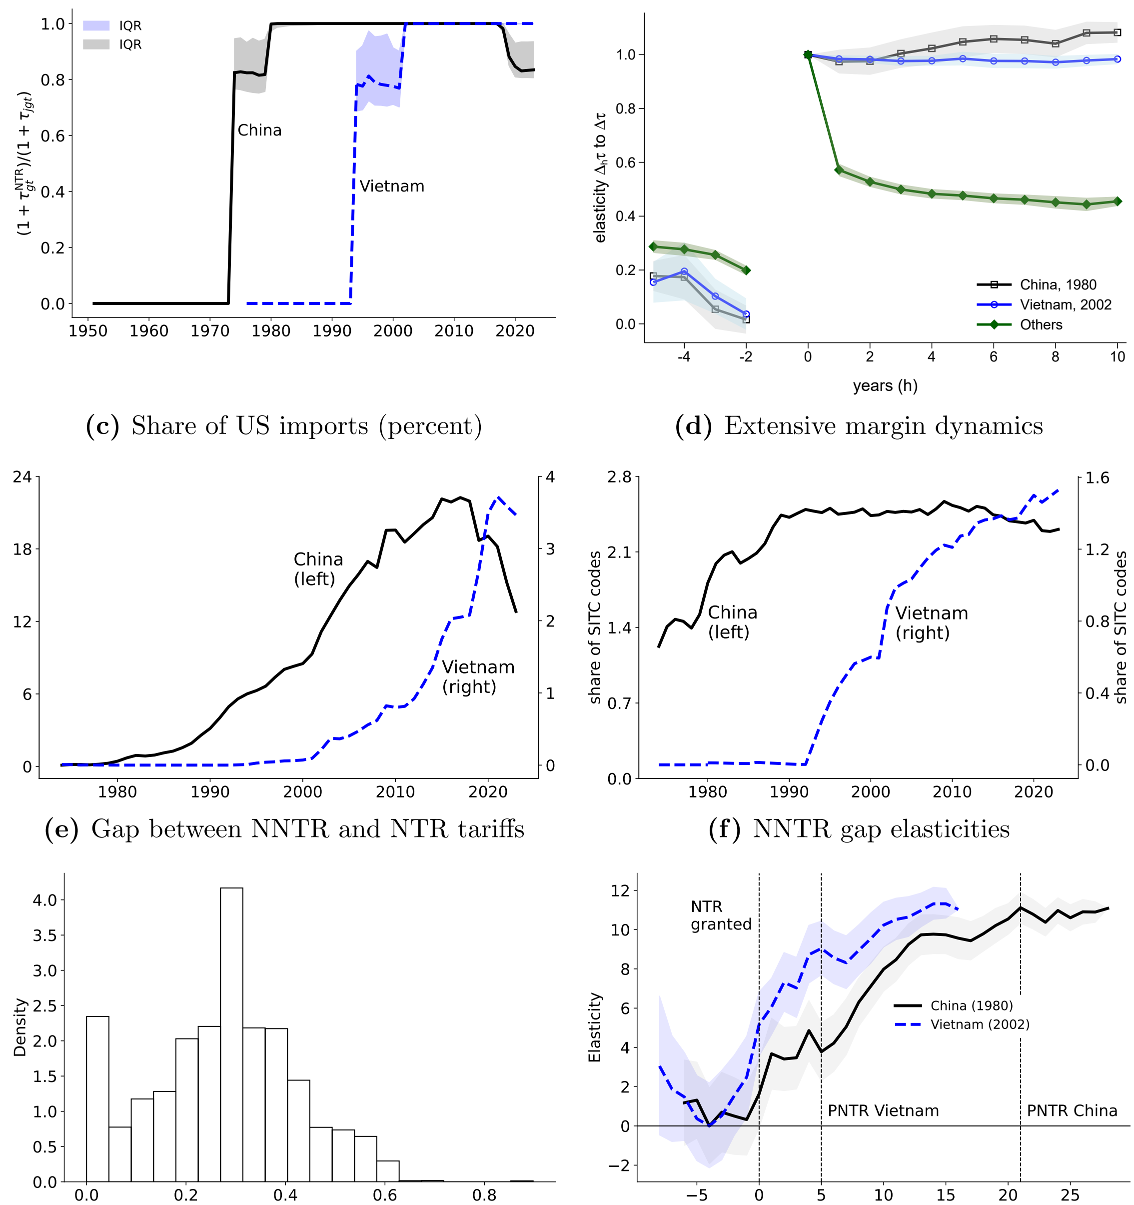

Notes: (a) Tariffs on NTR countries relative to tariffs on China and Vietnam. (b) Autocorrelation estimates from (3). (c) Share of US imports from China and Vietnam. (d) Share of SITC goods with positive imports from China and Vietnam. (e) Distribution of gap between NNTR and NTR tariffs in 1999. (f) Elasticity estimates from (8) with 95-pct. confidence intervals constructed using standard errors clustered at the jg level.

Figure 4: Model responses to persistent vs. transitory once-and-for-all reforms

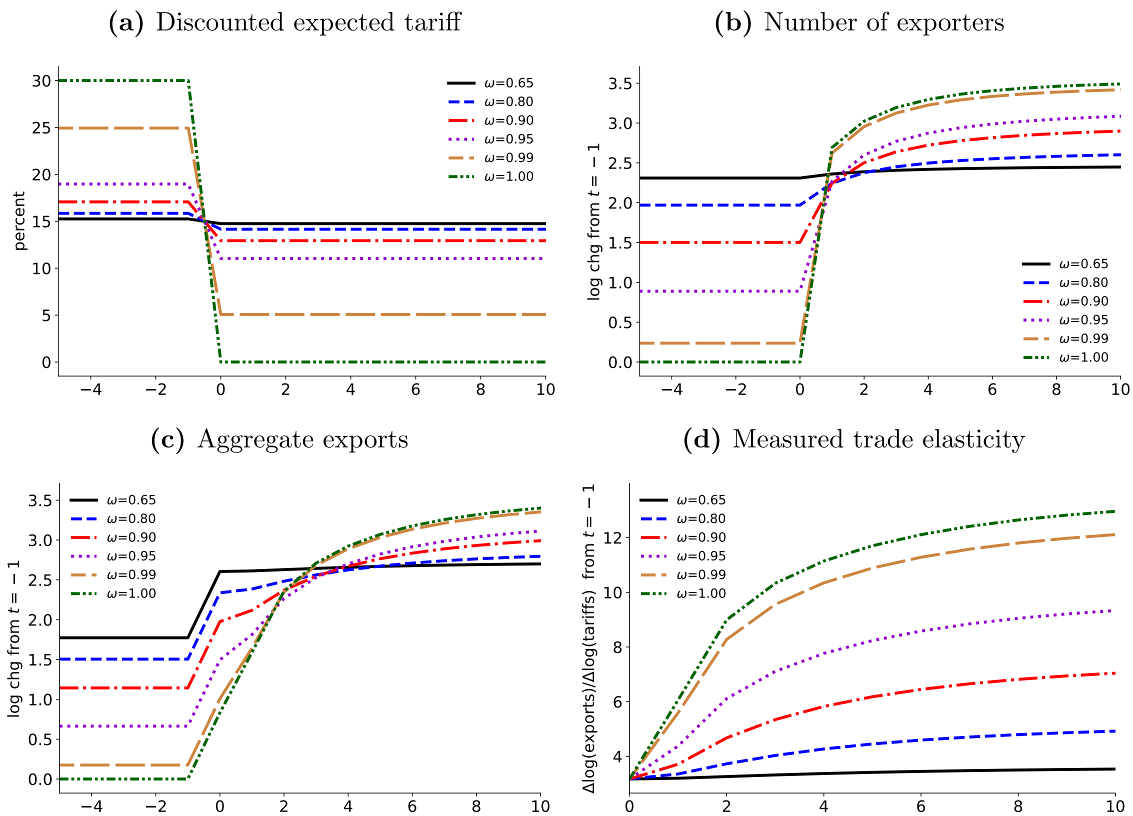

Notes: The figure plots the transition dynamics from the trade reform experiments described in section 6.1, in which the persistence of tariffs (ω) varies. In panel (d), the measured trade elasticity is defined in (4). In panels (b) and (c), outcomes are measured relative to period $t = -1$ in the calibration with perfectly persistent tariffs ($\omega = 1$).

Figure 5: Local-projections estimates model simulations with Markov tariffs

(a) Tariff autocorrelation

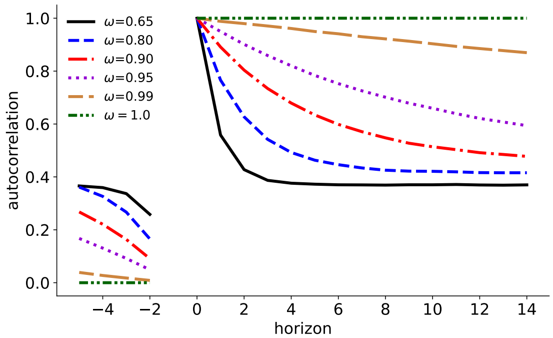

(b) Trade elasticity estimates

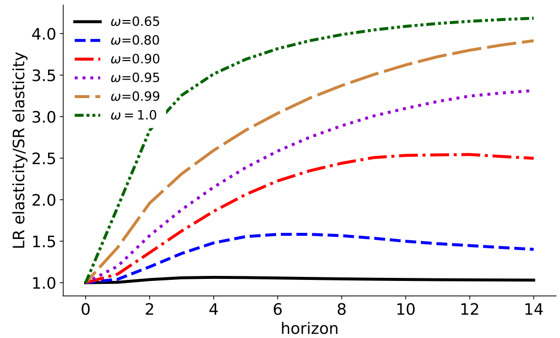

(c) Long-run response vs. tariff persistence

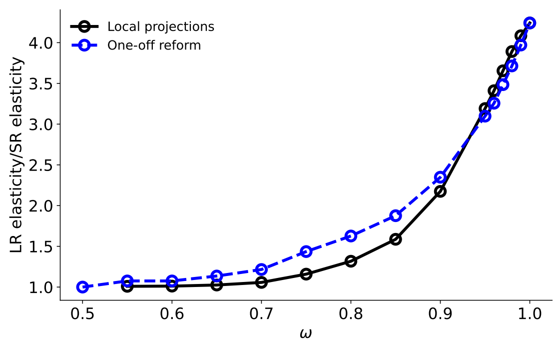

Notes: The figure plots the estimates of tariff and trade dynamics in simulated data from the model for different values of tariff persistence ω. Panel (a): tariff autocorrelation estimated using (1). Panel (b): trade elasticities estimated using (5). Panel (c): ratio of long-run trade elasticity to short-run trade elasticity as function of ω, in simulated data vs. once-and-for-all reforms.

Figure 6: Model responses to changes in persistence vs. changes in tariffs

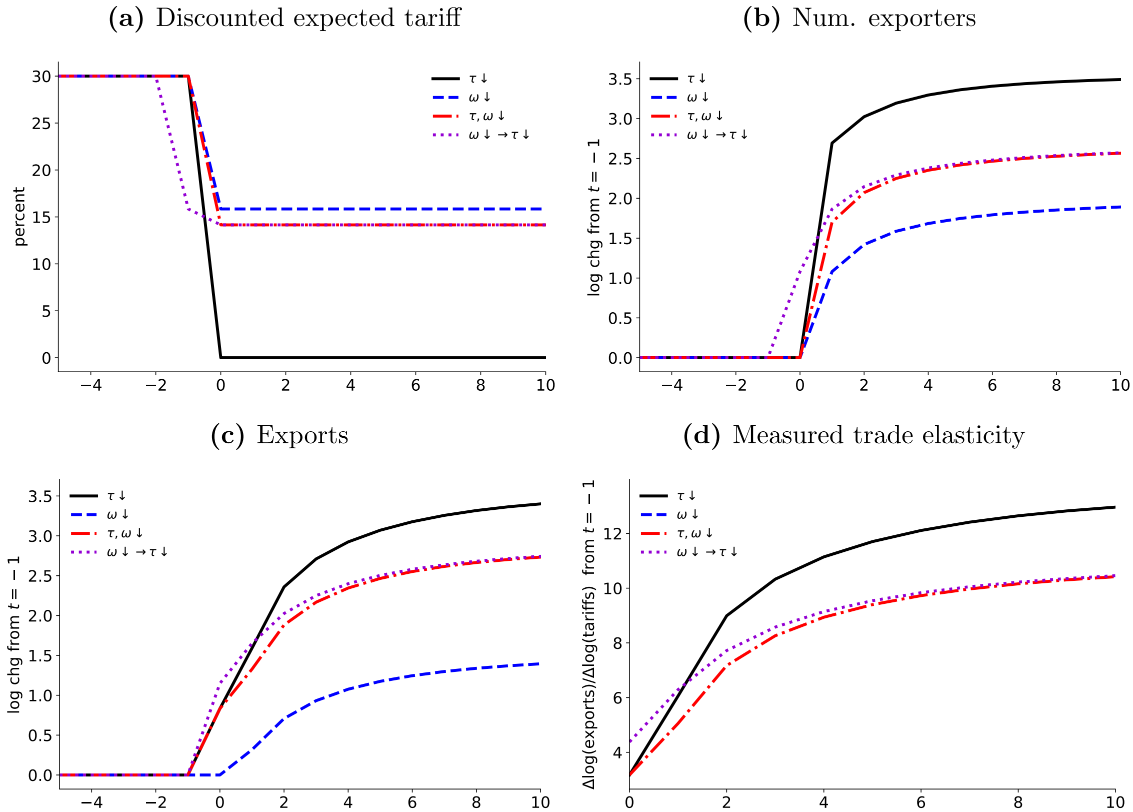

Notes: The figure plots the transition dynamics from the trade reform experiments described in section 6.2. $\tau \downarrow$: a drop in tariffs in $t = 0$. $\omega \downarrow$: a drop in persistence in $t = 0$. $\tau, \omega \downarrow$: a drop in tariffs and persistence in $t = 0$. $\omega \downarrow \rightarrow \tau \downarrow$: a drop in persistence in $t = -1$ followed by a drop in tariffs in $t = 0$. In panel (d), the measured trade elasticity is defined in (4). In panels (b) and (c), outcomes are measured relative to period $t = -1$ in the canonical case $\tau_{HP} \rightarrow \tau_{LP}$.

Figure 7: The NTR liberalizations of China (1980) and Vietnam (2002): Model

(a) NTR-gap elasticities, model and data

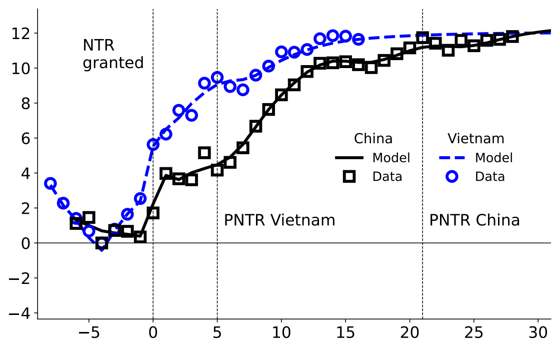

(b) Estimated probabilities

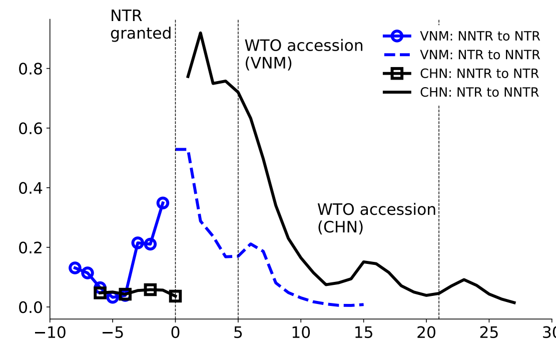

(c) NTR-gap elasticities, actual & canonical reforms

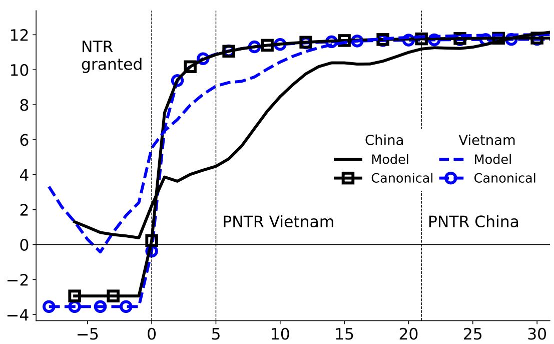

(d) Trade Convergence

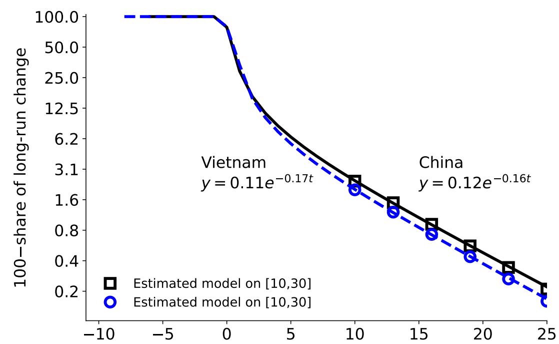

Notes: Panel (a): NTR-gap elasticities in data and baseline model. Panel (b): Estimated trade-policy transition probabilities. Panel (c) NTR-gap elasticities baseline model and the counterfactual canonical reform. Panel (d) The share of total adjustment remaining in the NTR-gap elasticities from the counterfactual canonical reform, as defined in (19).

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 A.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: Tariff changes across and within regimes

From	To	1-Year Changes N # jgt	Mean (p.p.)	Median (p.p.)	SD (p.p.)	5-Year Changes 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 4: Exporter-dynamics statistics and sector-level model parameters

(a) Common assigned parameters

	Parameter	Value	Target/Source
θ	Demand elasticity	3.17	Soderbery (2018)
r	Interest rate	0.04	Standard
ρ_z	Productivity persistence	0.65	Alessandria et al. (2021)
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)

(b) Country-specific jointly calibrated parameters

	Parameter	China	Vietnam	Target/Source	China	Vietnam
f_0	Sunk cost	0.49	1.29	Export part. (%)	28	11
f_1	Export cost	0.28	0.53	Exit rate (%)	11	15
ξ_H	High iceberg cost	4.34	6.63	Incumbent prem.	2.9	4.41
σ_z	Productivity shock dispersion	1.47	1.67	Log CV exports	2.27	2.91

Table 5: Long-run trade elasticities: canonical vs. reduced-form

Setting	Estimate
(a) Canonical reforms (section 7.2)	
China	14.76
Vietnam	15.26
(b) Event studies (section 4.2)	
China	11.80
Vietnam	11.86
(c) Across vs. within regimes (section 4.1)	
Across regimes	6.33
Within regimes	2.35
Pooled	3.51
(d) Numerical examples (section 6.1)	
$\omega = 0.65$	3.24
$\omega = 0.80$	4.13
$\omega = 0.90$	6.84
$\omega = 0.95$	10.07
$\omega = 0.99$	12.93

Notes: Table lists long-run trade elasticities in model vs. data. (a) Estimates for canonical reform based on transition from permanent NNTR to permanent NTR tariffs in model. (b) Gap-elasticities from (8). (c) Local-projections estimates from (6) using full US data sample. (d) Local-projections estimates from model simulations with two-state Markov processes.

Appendix (For online publication)

In section A, we discuss the robustness of our empirical approach. In section B we first use the model to explore various forms of anticipation, the role of initial conditions, and uncertainty over the likelihood of re-entering embargo, then discuss the sensitivity of the canonical long-run elasticity to model parameters. In section C, we study in more detail the role of time fixed effects in absorbing the transition from embargo, both empirically and in the model.

A Data

A.1 Regime classification

Here we describe our regime classification used in sections 3 and 4. 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:[31,32]

- *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

[31]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.

[32]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.

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 or 2,587,514 of the jgt triplets. Panel a of Figure A1 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 A1, which is a plot of the import share of the attributed jgt regime. The import share 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.

A.2 Robustness: Within vs. across tariff and trade dynamics

Our estimates of tariff autocorrelations in section 3.3 and reduced-form trade elasticities in section 4.1 are robust to a range of alternative estimation strategies and data samples that we describe below. The results are shown in Figures A2–A3.

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 A2 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.

Tariff autocorrelations with less restrictive fixed effects In our baseline estimation of tariff autocorrelations, we include both country-time (jt) fixed effects that absorb average tariff changes within countries over time, and good-time (gt) fixed effects that absorb average tariff changes within goods across countries. Figure A3(a) shows tariff autocorrelations estimated with only good-time (gt) fixed effects and no fixed effects at all. In both alternatives, the within-regime autocorrelations barely change, while the across-regime changes become more persistent. Thus, our more restrictive fixed-effects structure yields a conservative estimate of the difference in persistence between across vs. within tariff changes.

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 A.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 A3(b), 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), but the differences between across-regime vs. within-regime elasticities remain under alternative specifications.

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 *h-on-h* specification):

$$\Delta_h v_{jgt} = -\beta_h^{hh} \Delta_h \tau_{jgt} + \mathbf{Z}_{jgt} + \delta_{jt} + \delta_{gt} + u_{jgt}. \tag{20}$$

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 (the *ECM* specification), is a version of the error correction model (Gallaway et al., 2003; Alessandria and Choi, 2021) which recovers a parametric estimate of the long-run trade elasticity by specifying trade flows as an AR(1) process allowing for lagged effects of tariffs:

$$\Delta_0 v_{jgt} = \beta_0^{ecm} \Delta_0 \tau_{jgt} + \gamma(v_{jg,t-1} - \beta_\infty^{ecm} \tau_{jg,t-1}) + \delta_{gt} + \delta_{jt} + \delta_{jg} + u_{jgt}, \tag{21}$$

The dependent variable is the one-year log difference in import values. The right-hand side includes the one-year change in tariffs, lagged tariffs, lagged log imports, and a set of fixed effects. The 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. We add variety fixed effects, δ_{jg}, to capture each variety's average import level.[33]

[33] 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 as $v_{jgt} = \beta_0^{ecm} \tau_{jgt} + (1+\gamma) v_{jg,t-1} - (1+\gamma\beta_\infty^{ecm}) \tau_{jg,t-1} + \delta_{gt} + \delta_{jt} + \delta_{jg} + u_{jgt}$.

Figure A3(c) 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, the alternative specifications give higher estimates, especially the h-on-h specification. This indicates that Boehm et al. (2023)'s approach of controlling for the autocorrelation in the tariff process is critical for identification when that process has a low degree of persistence. Nevertheless, the differences between across- and within-regime elasticities, which is our main result in this part of the paper, are large under all specifications. 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 re-estimating (6) 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 $t-1$.[34] Figure A3(d) 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

[34]The rationale for this exclusion discussed in Boehm et al. (2023) is that major trading partners were generally the ones negotiating the MFN rate reductions in the WTO's multilateral negotiation rounds.

average or the median over the HS-8/TS-USA tariff lines. Figure A3(e) plots the results of (6) with these different tariff measures. 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 A3(f) plots the results of (6) 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 with respect to each other remain largely unchanged.

A.3 Robustness: China and Vietnam NTR access

Our approach to estimate the gap-elasticities laid out in section 4.2 and which we use in section 7 are robust to a range of alternative estimation strategies which we describe below. The results are shown in Figures A4–A5.

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.[35] The results are shown in Figure A4(a) and A4(b). The gap-elasticities are virtually unchanged under these alternative definitions of the gap.

[35] 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 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. The results are shown in Figure A4(c) and A4(d). None of these changes in the sample design significantly affect our estimates.

Gap elasticities vs. local projections In the main text we focus our case-study analysis on elasticities of trade to the NNTR gap following Alessandria et al. (2025a). Here we show estimates for China and Vietnam using two local projections specifications similar to (6). The first specification estimates local projections specifically to NTR access:

$$\begin{aligned}\Delta_h v_{jgt} = &-\beta^{v,\mathrm{CN}}\Delta_h \tau_{jgt}\{j=\text{China}, t=1980\} - \beta_h^{v,\mathrm{VN}}\Delta_h \tau_{jgt}\{j=\text{Vietnam}, t=2002\}\\ &-\beta_h^{v,\mathrm{others}}\Delta_h \tau_{jgt}\{j=\text{other}\} + \mathbf{Z}_{jgt} + \delta_{jt} + \delta_{gt} + u_{jgt}.\end{aligned} \tag{22}$$

These estimates are normalized to zero in the year before NTR access, and we renormalize our gap-elasticities the same way to be consistent. Figure A5(a) shows that in the case of China, these LP estimates are closely aligned with the gap elasticities.

In the case of Vietnam, however, Figure A5(b) shows the LP estimates significantly understate the extent of long-run trade growth due to the anticipatory growth that occurs before NTR access. The value of the long-run LP estimate from (22) is similar to the long-run value of the gap elasticity, but the former only captures the growth in trade from the period before NTR access, which is two full log points higher than several periods earlier; recall that the long-run response of trade in the event-study specification is the difference between the highest gap elasticity (at the end of the sample) and the lowest (which occurs in period -4). This shows that our event-study approach does a better job of dealing with this anticipatory behavior than the local-projections specification.

The second specification estimates local projections for the average tariff change for China and Vietnam:

$$\begin{aligned}\Delta_h v_{jgt} = & -\beta^{v,\mathrm{CN}} \Delta_h \tau_{jgt} \{j = \mathrm{China}\} - \beta_h^{v,\mathrm{VN}} \Delta_h \tau_{jgt} \{j = \mathrm{Vietnam}\} \\ & - \beta_h^{v,\mathrm{others}} \Delta_h \tau_{jgt} \{j = \mathrm{other}\} + \mathbf{Z}_{jgt} + \delta_{jt} + \delta_{gt} + u_{jgt}. \end{aligned} \tag{23}$$

This specification yields lower long-run trade elasticities than the first, especially for Vietnam. This is because the average tariff change for China and Vietnam is much less persistent than the changes associated with NTR access. Figure A5(c) illustrates this by plotting the autocorrelation of the NTR-access tariff changes for China and Vietnam against the autocorrelation of these countries' other tariff changes. The latter look very similar to within-regime tariff changes discussed in the main text.

B Model

We now use the model to illustrate some additional principles related to anticipation and uncertainty, and measure the sensitivity of our canonical elasticity to parameters.

B.1 Initial conditions and Temporary Lifting of Embargo

In the numerical illustrations in section 6, the economy was in a neighborhood of the steady state at the time of reform, where tariffs had been high (but finite) for many periods. In China and Vietnam the transition from NNTR to NTR takes place while the economies are transitioning from the end of the embargo. To illustrate how the economy adjusts to lifting an embargo, and the implication of this adjustment for trade-elasticity measurement, we consider a scenario where the economy begins in embargo and then opens up permanently. We also consider a case in which there is a constant probability of returning to the embargo.

To align with our empirical approach, we apply these reforms for two goods that permanently go to different tariff levels of zero and 30 percent. We also estimate the trade elasticity from these two goods using a panel approach,

$$v_{gt} = \sum_{t'=1}^{25} \beta_t \mathbb{1}\{t = t'\} \log(\tau_{gt}) + \delta_t + u_{gt}. \tag{24}$$

Initial Conditions Figure A6(a) depicts trade dynamics starting from the period following the embargo's end. Figure A6(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 large trade elasticity of about 10 from comparing the differences in trade across the two goods to the difference in tariffs, but this elasticity is smaller than the canonical elasticity of about 15 that we estimated in the paper. As trade grows further in the subsequent periods, the trade elasticity eventually converges to that canonical value. 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.

Return to Embargo Possibility. We also do a version of this experiment where there is a 10% chance of returning to embargo. For both goods, the level of trade when the embargo is lifted is lower and there is less adjustment over time. The trade elasticity still starts about at about 10, but does not grow in magnitude as much as in the case without uncertainty.

To connect with the empirical discussion about the how the standard set of country-year fixed effects absorbs part of the adjustment from the end of the embargo, the figure also shows the path of the trade elasticity without the time effect δ_t. Now, it starts out much larger in magnitude when the embargo is lifted, and shrinks over time instead of growing. As discussed above, this is similar to what we observe in the data for China and Vietnam and in our modeling of those episodes.

B.2 Anticipation to phased-in reforms

Aside from the changes between NNTR and NTR/UTPP, many tariff changes involve a phase-out. 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.

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 phase-outs. 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 phase-out reforms, tariffs fall by 3 p.p. per year for 10 years. The linear phase-out 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 A7 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.

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, but each reform has a different long-run trade elasticity. This is because the level of trade in the period before the reform begins 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, shown in panel (e), 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.3 Sensitivity of canonical long-run elasticity

Our estimate of the canonical long-run elasticity is one of the main outputs of our analysis. We now investigate the sensitivity of this estimate to the model parameters, both those that were externally assigned and internally calibrated. For each parameter, we simulate the canonical reform when that parameter is set 10% higher and 10% lower than the baseline value. Table A5 shows the results of this sensitivity analysis.

Demand elasticity θ. The canonical elasticity is increasing in θ. An increase in θ

causes the long-run elasticity to increase by about twice as much. The response is greater than one-for-one for the following reason. Mechanically, if firms all had the same export technology and did not respond at all, the response would be one-for-one, since each firm's sales would rise by the same amount as demand. But firms with low iceberg costs benefit disproportionately, since the term $\xi^{1-\theta}$ enters in the firm's profit function. Moreover, the increase in θ makes the benefit of being a high-capacity exporter greater, increasing entry and survival.

While we have externally assigned the value of θ based on Soderbery (2015)'s estimates, the path of the NTR-gap elasticity in the data actually contains some identifying information. Specifically, the value of this elasticity at the end of the sample imposes some bounds on what θ could be. If θ is too low, the NTR-gap elasticity in the model would be too small even if NTR access was a complete surprise; the canonical reform would generate a long-run elasticity that is smaller than the observed change in the NTR-gap elasticity which is generated from a world with uncertainty. When we set θ only 10% smaller (to a value of about 2.9, which generates a canonical long-run elasticity of 11.87) we are already near that region of the parameter space.

Conversely, if θ is too high, the canonical reform generates such a large long-run response that the model cannot account for the NTR-gap elasticity dynamics even with a probability of losing NTR access equal to one (at least during the years right after the NTR reform occurs). Thus, θ must be fairly close to Soderbery (2015)'s estimate for the model to be able to match the NTR-gap elasticity trajectory. We view this as a validation of sorts of our approach to calibrating this parameter.

Entry cost f_0. The canonical elasticity is increasing in the entry cost, but the effect is quite small. An increase in the entry cost reduces export participation, but the effect is greater when the profits from exporting are lower, when the country is in the NNTR regime. Consequently, an increase in the entry cost amplifies the extensive-margin effect of a tariff cut.

Continuation cost f_1. The canonical elasticity is decreasing in the continuation cost, but again the effect is small. The reason is as follows. When the continuation cost is higher, the exit rate is higher, and firms are less likely to survive to reach the low-iceberg state ξ_L. A smaller mass of low-iceberg exporters means a smaller trade response.

High iceberg cost ξ_H. The canonical elasticity is increasing in ξ_H. The sensitivity here is a bit greater than to the fixed costs. The logic is that this parameter increases the firm-level long-run response to a tariff reform, conditional on making it to the low-iceberg state.

Iceberg switching probability ρ_ξ. This parameter can have a non-monotonic effect on the canonical elasticity: increasing or decreasing this parameter can lower the long-run response of trade. The source of the non-monotonicity is as follows. There are two opposing forces:

- "Arrival" factor: how soon you start enjoying ξ_L after entering with ξ_H. Decreasing in ρ_ξ.
- "Durability" factor: how long you enjoy ξ_L once you attain it. Increasing in ρ_ξ.

At very low value of ρ_ξ, arrival is fast, which makes exporting more attractive, but ξ_L spells are short and frequently interrupted, which makes exporting less attractive. Conversely, at a very high value of ρ_ξ, ξ_L spells are very long, making entry more attractive, but also take a long time to begin, making entry less attractive. Intermediate values of ρ_ξ balance these two effects.

For both countries, ρ_ξ is one of the parameters for which the canonical long-run elasticity is quite sensitive. For China, the effect is indeed non-monotonic: increasing ρ_ξ lowers the canonical long-run elasticity quite a lot, and reducing it lowers the elasticity a small amount. For Vietnam, the effect is monotonic: increasing ρ_ξ lowers the elasticity, and increasing ρ_ξ raises it.

Productivity shock dispersion σ_z. The canonical elasticity is decreasing in σ_z. A more dispersed productivity shock distribution puts less mass near the entry threshold, making export participation less sensitive to tariff changes. The effect here is also fairly large.

Productivity persistence ρ_z. The canonical elasticity is decreasing in productivity persistence as well. The effect here is very large. The reason is the same as with σ_z. A higher ρ_z makes the ergodic productivity distribution more dispersed.

Fixed death rate δ_0. This parameter has a negligible effect on the canonical elasticity. For China, the effect is positive but very small. For Vietnam, it is non-monotone and slightly larger, but still quite small.

Elasticity of death to productivity δ_1 This parameter has a negligible effect on the canonical elasticity. For China, the effect is negative but very small. For Vietnam, it is again non-monotone and slightly larger, but again quite small.

C Model & data: Adjustment from embargo

As described in the main text of the paper, the country-year fixed effects, δ_{jt}, are typically included in trade-elasticity estimation frameworks. In the context of the China and Vietnam case studies, they do not just absorb aggregate shocks in exporting countries, they contain some useful information about the response of trade to the lifting of the U.S. embargoes on the two countries. To demonstrate this, we estimate versions of (8) without these effects, instead including other independent variables that more directly control for aggregate shocks without absorbing the effects of these reforms.

Specifically, we estimate:

$$v_{jgt}/AS_{jt} = \sum_{t'=1974}^{2008} \beta_t^{v,\mathrm{CN}} \mathbb{1}\{t=t'\}\mathbb{1}\{j=\text{China}\}X_g \tag{25}$$
$$+ \sum_{t'=1994}^{2017} \beta_t^{v,\mathrm{VN}} \mathbb{1}\{t=t'\}\mathbb{1}\{j=\text{Vietnam}\}X_g + \delta_{jg} + \delta_{gt} + u_{jgt}.$$

where AS_{jt} are aggregate supply factors which are measured as total exports excluding exports to the U.S.[36] Note that with respect to (8) the jt fixed effects are dropped. In our baseline estimation (8) these fixed effects control for aggregate supply factors; however, as emphasized here they also capture the gradual adjustment to the lifting of the embargo. By dividing good-level exports by AS_{jt} we intend to control for the former but not for the latter.

Figure A8 shows the results. When country-year fixed effects are excluded and no other controls for aggregate shocks are added (the blue lines), the gap elasticities grow much larger in magnitude at the beginning of the sample. This is because these elasticities now include the effects of productivity growth, population growth, etc. on exports to the United States, in addition to the effects of trade reforms. When we directly control for aggregate shocks that are unrelated to U.S. trade policy but still exclude country-year fixed effects (the red lines), the gap elasticities shrink in magnitude substantially. In both cases, the red line starts well below the baseline elasticity (in black) and rises over time throughout the NNTR period. This is because the red line includes the growth of trade generated from the end of the embargo.

In the case of China (left panel), the red line remains below the baseline elasticity (black line) throughout the sample. The red dashed line in the figure shows the gap elasticity estimate for China without country-time fixed effects in the model. It qualitatively mimics the pattern seen in the data, but it is not shifted downward as much relative to the black line as in the data. Experimenting with our model, we found that one way to explain this is that with a non-zero probability of returning to

[36] Results are similar if we use GDP or total exports including to the U.S.

embargo from the NNTR state. This reduces the growth of trade stemming from the end of the embargo, and while this affects all goods, it disproportionately affects goods with low (and especially zero) NNTR tariff rates. These goods are not meaningfully exposed to the risk of losing NNTR status, and therefore do not contribute to the baseline gap elasticity.

Our experiments show that the larger the probability of returning to embargo, the larger the gap between the black and red lines. We can estimate this probability to match the red line from the data. The green line in the figure shows the result of this effort. We can match the empirical estimate more or less exactly. The last panel in the figure shows the resulting probabilities. The probability of embargo is low during the NNTR period, but rises to almost 20 percent shortly after NTR access is granted. It falls back near to zero throughout most of the late 1980s and early 1990s, before rising temporarily around WTO accession. The effects on the other probabilities (of going between NNTR and NTR) are minor. Thus, our main results are not sensitive to allowing for a chance of returning to embargo.

In the case of Vietnam (right panel), the red line is further below the black line during the NNTR period relative to China. This is because the end of the embargo is closer in time; for Vietnam, our first observation is the year after the embargo ended, whereas for China it is three years afterward. However, the red line rises more swiftly than in the Vietnamese case and is statistically indistinguishable from the black line starting a a few years before NTR access. Even with no chance of returning to embargo, our model's version of the red line is slightly below the empirical estimate. This tells us that returning to embargo was less likely to have been a material risk for Vietnam than for China, which we believe makes sense given the geopolitical context surrounding the two countries and their relationships with the United States.

D Additional figures

Figure A1: Regime Classification

(a) Unclassified jgt observations

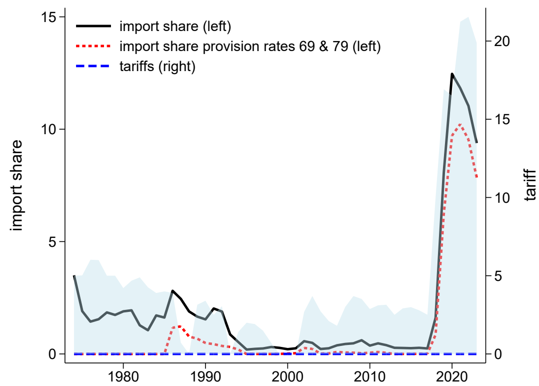

(b) jgt Regime Import Share

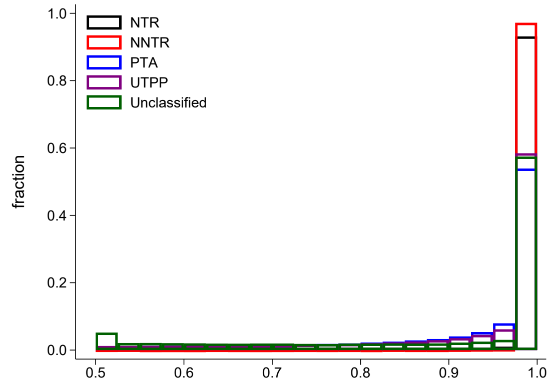

Notes: (a) import share and tariffs of unclassified observations (jgt's). Also plots share specific to provision rates 69 and 79 excluded by our classification (dotted red line). The blue dashed line plots the median tariff and the shaded area is the 25$^{\text{th}}$ and 75$^{\text{th}}$ percentiles of the tariff distribution. (b) Distribution of import shares of the classified jgt's (by regime) vs. unclassified observations. For illustrative purposes, we truncate unclassified jgt's at 0.5.

Figure A2: Tariff and trade dynamics across specific regime changes

(a) Tariff autocorrelations

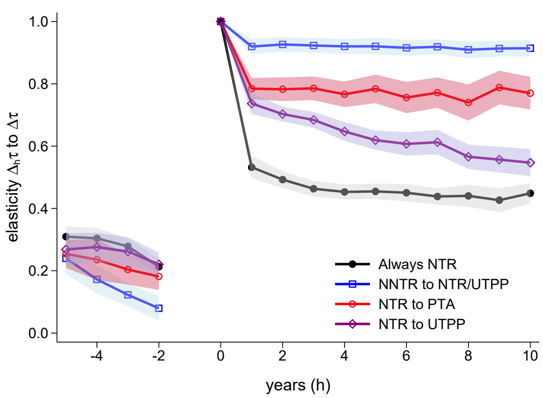

(b) Reduced-form elasticities

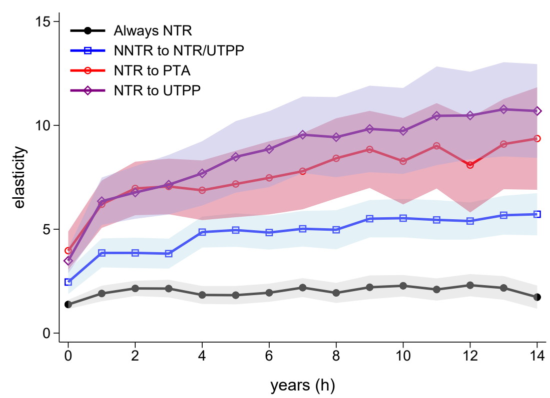

Notes: (a) Results of (1) when considering the regime transitions specified in the legend. (b) Same, but for (6). Shaded areas show 95-pct. confidence intervals constructed using standard errors clustered at jg level.

Figure A3: Robustness: within-vs. across tariff and trade dynamics

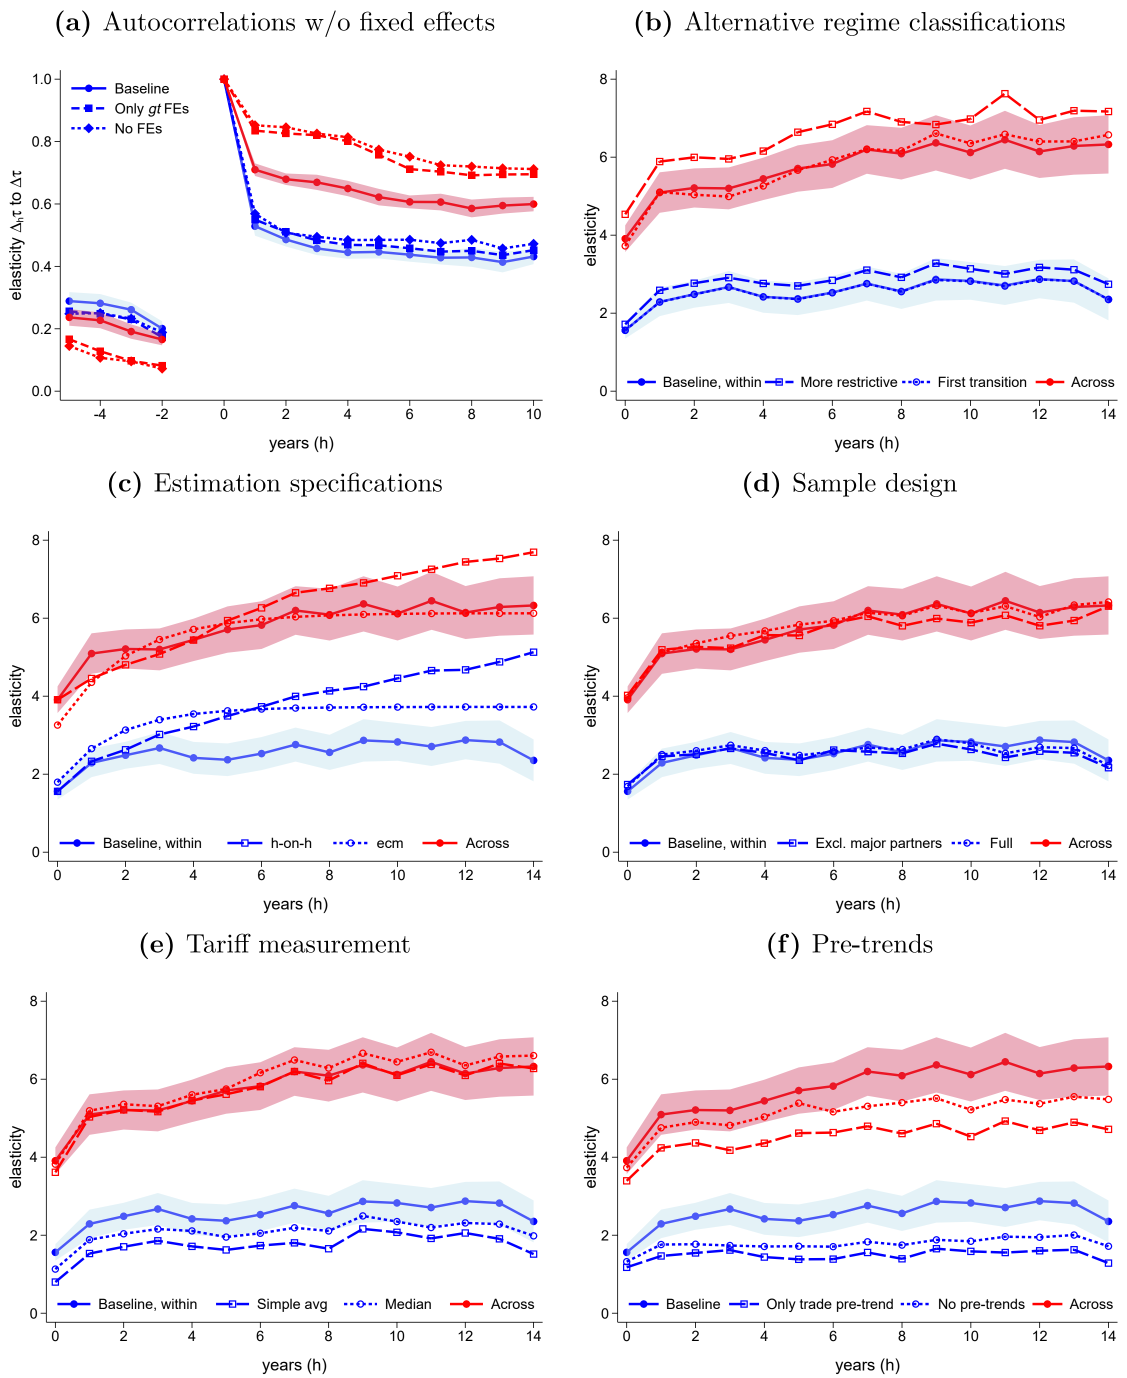

(a) Estimates of (1) with different fixed-effect specifications. (b) Estimates from (6) using alternative regime classifications. (c) Estimates using alternative regression specifications. (d) Estimates using alternative sample designs. (e) Estimates using alternative tariff measures. (f) Estimates with fewer pre-trend controls.

Figure A4: Robustness: Gap elasticities

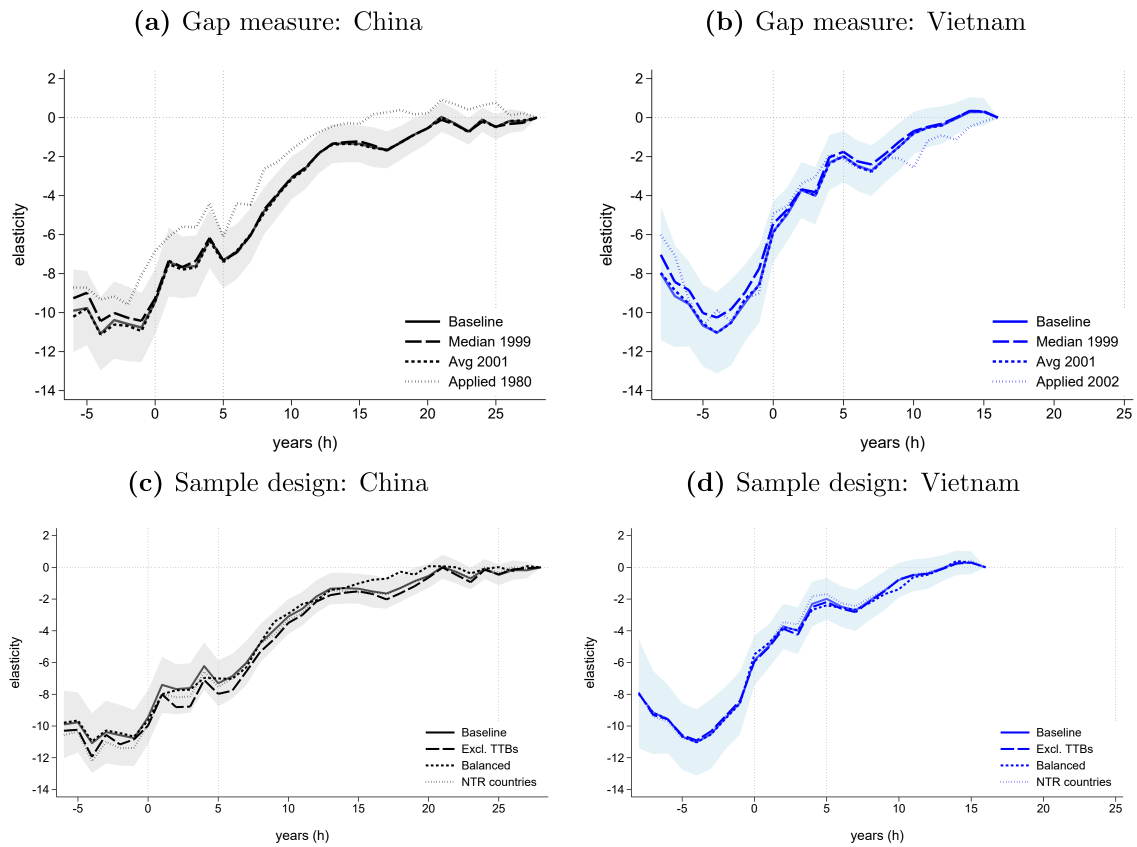

Notes: (a) and (b) Estimates from (8) under baseline approach (solid line and shaded area for the confidence interval) vs. alternative gap measurements. (c) and (d) Estimates of (8) under baseline approach vs. alternative sample designs.The period $h = 0$ is 1980 for China and 2002 for Vietnam.

Figure A5: Gap elasticities vs. local-projections estimates

(a) China **(b)** Vietnam **(c)** Tariff autocorrelations

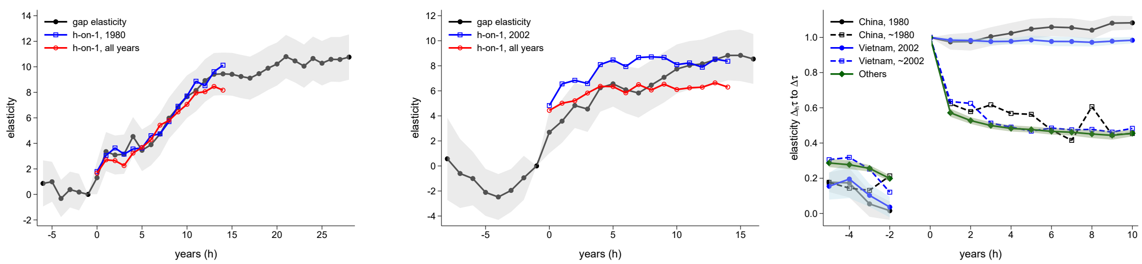

Notes: (a) and (b) plot the gap-elasticities estimated by (8) when normalizing to the period prior to the NTR access ($h = -1$). Blue line plots local-projections estimates from (22) focusing on changes from the NTR access. Red line plots local-projections estimates from (23) that average over all years. Period 0 is 1980 in (a) and 2002 in (b). (c) plots estimates of (3) from Figure 3(b) together with estimates when considering only the NTR access years (1980 for China and 2002 for Vietnam) and when excluding those.

Figure A6: Role of initial conditions: Dynamics from an embargo's end

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

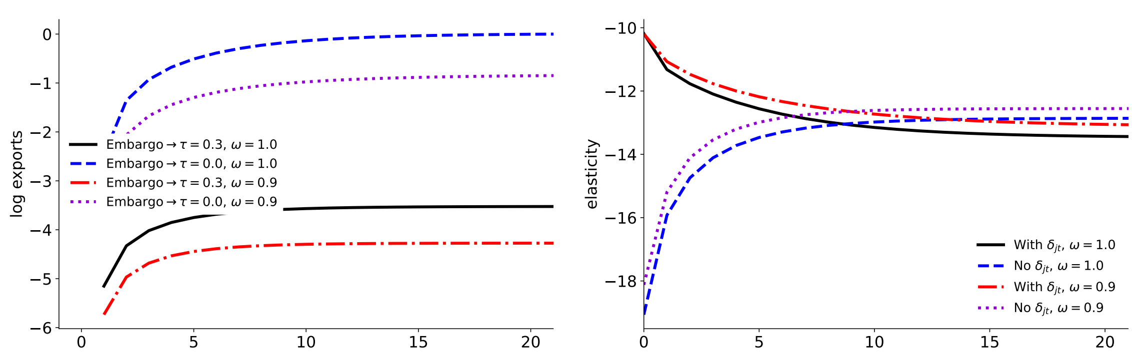

Notes: The figure plots the transition dynamics following the end of an embargo for two goods, one that goes to low tariffs and another to high tariffs. There are two cases, one where the embargo ends for good and another where there is a 10% chance of it restarting. Panel (a): Log exports. Panel (b) Time-varying elasticity of trade to tariffs, with and without time fixed effects.

Figure A7: Responses to anticipated and/or phased-in reforms

(a) Tariffs

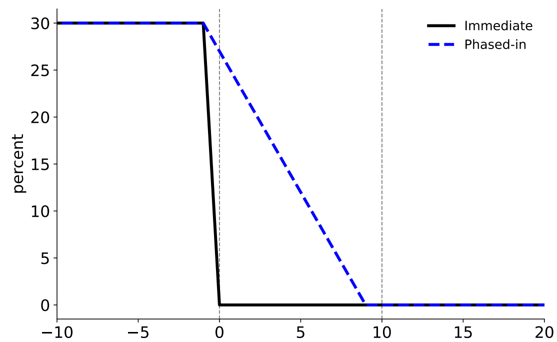

(b) Discounted expected tariffs

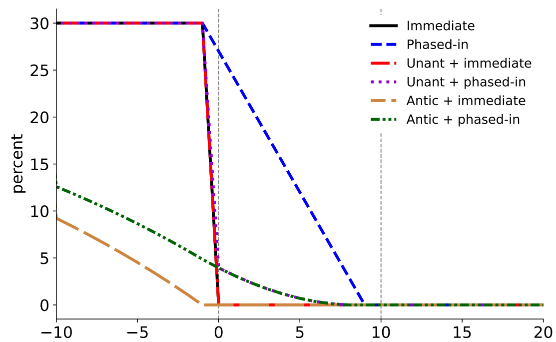

(c) Num. exporters

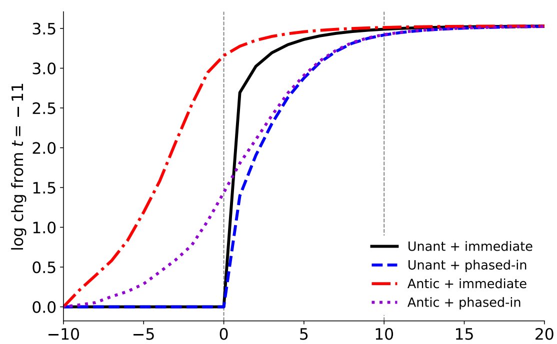

(d) Exports

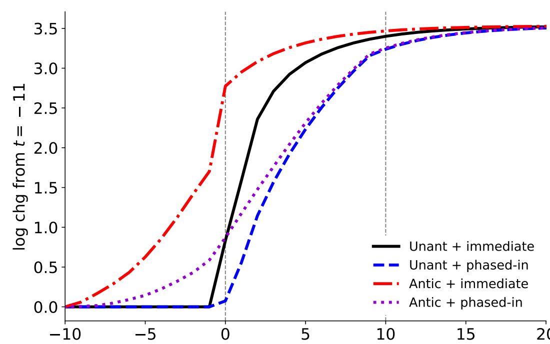

(e) Trade elasticity from $t = -1$

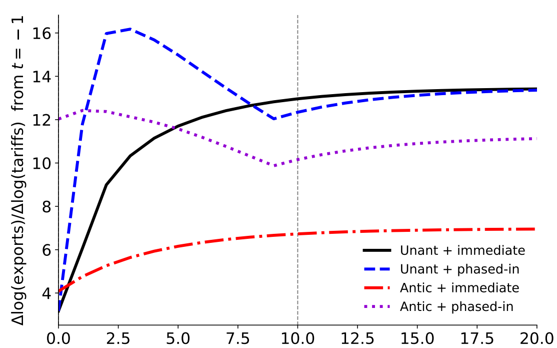

(f) Trade elasticity from $t = t_A$

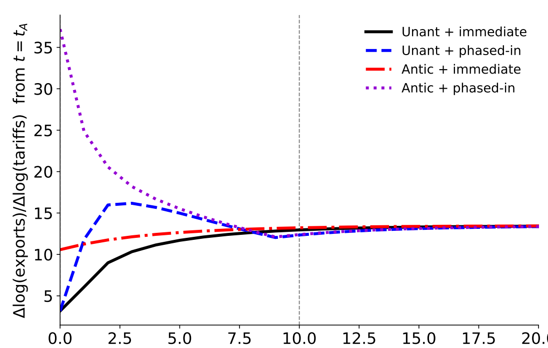

Notes: (Figure shows transition dynamics in experiments described in section B.2. (a): observed path of tariffs in immediate (solid black) and phased-in (dashed blue) reforms. (b): discounted expected tariff in unanticipated/immediate (solid black), unanticipated/phased-in (dashed blue), anticipated/immediate (dash-dotted red), and anticipated/phased-in (dotted purple) reforms. (c): number of exporters with same legend as (b). (d): aggregate exports with same legend as in (b). (e): trade elasticity measured relative to period before reform with same legend as (b). (f): trade elasticity measured relative to period firms learn about reform with same legend as (b).

Figure A8: Gap-elasticities: Robustness country-year fixed effects

(a) China

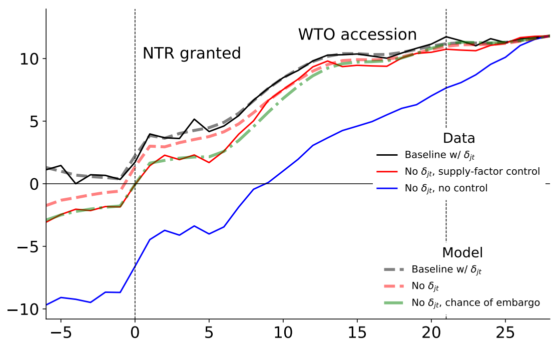

(b) Vietnam

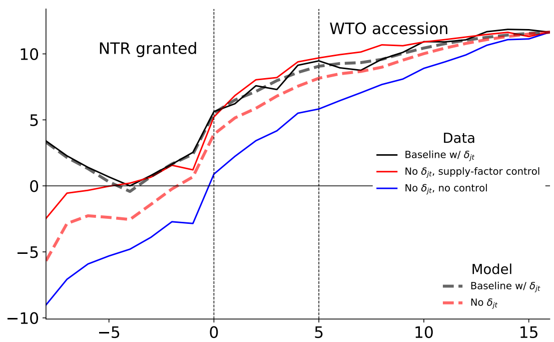

(c) Transition probabilities for China

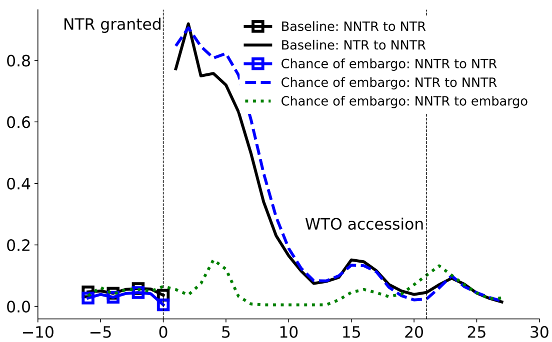

Notes: (a) and (b) plot the results of (8) under the baseline approach (black line), without country-year fixed effects but with aggregate export control (red line), and without country-year fixed effects or control (blue line). Solid lines are data and dashed lines are model. For China, estimates without country-year fixed effects are also shown for the model with a probability of embargo (green line). The period $h = 0$ is 1980 for China and 2002 for Vietnam. (c) plots the transition probabilities in the model with a chance of embargo versus the baseline probabilities. Black is the baseline, blue (NNTR to/from MFN) and green (NNTR to embargo) are the embargo model.

E 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.

Table A4: Top five country-year transitions across regimes

From	To	*jg* (# *g*)
NTR	NNTR	PLD-1983 (232), PLD-1984 (78), PLD-1985 (43), AFG-1986 (46), ROU-1989 (119)
NTR	PTA	CAN-1989 (889), MEX-1994 (387), KOR-2012 (325), AUS-2005 (241), ISR-1986 (203)
NTR	UTPP	TWN-1976 (280), HKG-1976 (214), ISR-1976 (204), KOR-1976 (189), BRA-1976 (177)
NNTR	NTR	CHN-1980 (273), VNM-2002 (347), PLD-1989 (253), USSR-1992 (226), USSR-1993 (215)
NNTR	UTPP	ROU-1994 (32), CZE-1992 (31), CZE-1991 (28), BGR-1992 (26), PLD-1990 (23)
PTA	NTR	CAN-1999 (205), MEX-1999 (179), ISR-1999 (165), AUS-2009 (135)
UTPP	NTR	KOR-1989 (403), TWM-1989 (400), HKG-1989 (265), MYS-1997 (262), PRT-1986 (214)
UTPP	NNTR	ROU-1989 (7), ROU-1990 (6), ROU-1992 (5), ROU-1993 (5), YUG-1996 (5)
UTPP	PTA	ISR-1985 (354), MEX-1994 (342), PER-2007 (241), COL-2001 (229), DOM-2007 (176)

Notes: Regime transitions are based on consecutive good-country observations but need not be in consecutive years. For instance, Poland transitioned back to NNTR in October of 1982 and several goods traded under NTR were not traded in 1983 or 1984.

Table A5: Sensitivity of canonical long-run elasticity to model parameters

	China		Vietnam	
Baseline	14.67		15.29	
Parameter	Increase	Decrease	Increase	Decrease
θ	18.02 (2.16)	11.87 (2.22)	18.68 (2.10)	12.68 (1.97)
f_0	14.99 (0.23)	14.39 (0.20)	15.66 (0.25)	15.29 (0.01)
f_1	14.24 (−0.31)	15.10 (−0.30)	15.10 (−0.14)	15.94 (−0.44)
ξ_H	15.52 (0.59)	13.86 (0.59)	16.05 (0.51)	14.85 (0.31)
ρ_ξ	14.13 (−1.50)	14.61 (0.16)	14.82 (−1.28)	15.45 (−0.40)
σ_z	12.48 (−1.70)	16.50 (−1.24)	13.81 (−1.07)	17.15 (−1.20)
ρ_z	9.99 (−4.03)	18.28 (−2.31)	10.56 (−3.89)	18.88 (−2.21)
δ_0	14.70 (0.03)	14.63 (0.03)	15.53 (0.16)	15.47 (−0.12)
δ_1	14.58 (−0.06)	14.76 (−0.06)	15.39 (0.06)	15.55 (−0.18)

Notes: Each row shows what happens when one parameter is increased or decreased by 10%, except for ρ_ξ where the change is 2.5% due to how close the parameter is to 1.0. Numbers in parentheses are elasticities of the canonical long-run elasticity with respect to the parameter.