The Causal Effects of Korea's Quota Tariff Policy on Agricultural Product Retail Prices

Deokjae Jeong* Youngmi Kim†

Korea Customs and Trade Development Institute

22, Magokjungang 5-ro, Seoul 07788, South Korea

Abstract

The Korean government operates the Quota Tariff system with the objective of stabilizing prices. This study examines whether the policy causally lowered agricultural retail prices as intended. The sample is a daily panel of retail prices for 41 agricultural products from January 2021 through March 2025. Eleven of these products (six leafy and root vegetables and five fruits) experienced tariff reductions under the Quota Tariff; the remaining 30 did not. Quota Tariff designations begin at different times across products (staggered adoption), the treatment is suspended and later resumed (non-absorbing), and the depth of the tariff reduction also differs across products (heterogeneous treatment intensity). Since existing estimation methods that presuppose absorbing treatment cannot be applied to this structure, we adopt local projection difference-in-differences as the best-suited method. We then estimate two effects: the effect while the tariff reduction remains in force (the persistent-treatment specification) and the total effect of a Quota Tariff designation (the impulse-response specification). For the vegetable group, retail price declines were insignificant in both specifications, whereas the fruit group showed declines of about -0.78% per one percentage point of tariff reduction at 150 days after Quota Tariff onset (persistent treatment) and $-0.49 \sim -0.55\%$ at about 250∼273 days (impulse response). The Quota Tariff entails large policy costs. Concentrating designations on fruits, rather than on vegetables for which no effect is identified, can therefore reduce administrative costs and stabilize prices simultaneously.

Keywords: Quota Tariff, Price Stabilization, Local Projection Difference-in-Differences, Agricultural Trade Policy

JEL Classification: F13, F14, Q17, C23

*ubuzuz@gmail.com; www.jayjeo.com

†youngmiaa@kctdi.or.kr; Corresponding Author.

1 Introduction

This paper studies Korea's Quota Tariff system, grounded in *Article 71 of the Customs Act*, focusing on the tariff reductions aimed at stabilizing agricultural prices. The Quota Tariff is a flexible tariff instrument: by Presidential Decree, the government can adjust both the reduced tariff rate and the volume of imports eligible for that rate, according to its policy objectives. This allows the government to charge a rate below the Basic Tariff Rate, supporting the smooth supply of goods and promoting domestic price stability. Moreover, a quantity limit is not an essential element of the system, so a temporary tariff reduction without any volume restriction is also possible. The Quota Tariff is distinct from the Tariff-rate Quota (TRQ) operated under the WTO and free trade agreements (FTAs). The TRQ is an institution that guarantees the market access a country has committed to under the WTO Agreement on Agriculture or an FTA. It applies a low tariff rate up to a specified volume and a higher rate to imports beyond it, aiming to achieve market opening and the protection of domestic industry at the same time. Korea's Quota Tariff, by contrast, adjusts tariffs flexibly for domestic policy purposes independently of such international obligations, and a few countries besides Korea operate similar schemes. In this respect, Korea's Quota Tariff is a far more flexible instrument than the TRQ. To keep the two distinct, this paper uses the term 'Quota Tariff' to refer to Korea's system.

Prior studies have focused mainly on the trade and protection effects of the TRQ (Loginova et al., 2021; Schaefer and Wolf, 2025). Empirical analyses of the Quota Tariff, by contrast, are scarce. This study fills that gap, extending the literature by analyzing the causal effect of the Quota Tariff on consumer prices.

This gap extends beyond the choice of subject to the methodology. The tariff measures imposed by the TRQ family and by the Quota Tariff are temporary, so they generate a non-absorbing treatment: one that starts and ends repeatedly within the sample period. Yet studies that estimate the causal price effects of such institutions directly on non-absorbing data are, to the best of our knowledge, limited to a single one, Loginova et al. (2021), and even that study goes no further than a static contrast between protected and unprotected phases. This study estimates the effects dynamically from the non-absorbing treatment path itself, using a daily treatment record that reflects even the lapses caused by quota exhaustion; a detailed review of this methodological gap is presented in section 2.

Finally, the study also has important implications for Korea's Quota Tariff policy on agricultural products. In Korea, debate over the price-stabilizing effect of the Quota Tariff has persisted for years. As of 2024, Quota Tariff coverage of agricultural products had expanded considerably: the number of covered products rose from 20 in 2020 to 72 in 2024, and total tariff reductions reached about one billion dollars, a substantial fiscal cost (Jang, 2025). Since about 66% of these measures were deployed to stabilize retail prices, it is important to assess whether the policy actually achieves its original goal.

To answer this question, this paper constructs a daily panel of retail prices covering 41 agricultural products from January 2021 through March 2025. The sample comprises eleven products that had the Quota Tariff in force at least once during the sample period (six leafy and root vegetables and five fruits) and 30 products that never did. Quota Tariff designations begin on different dates across products (staggered adoption), the same product can be designated again after a designation lapses (non-absorbing treatment), and the depth of the tariff reduction differs widely across products (heterogeneous treatment intensity). Conventional difference-in-differences (DiD) estimators are either biased under this structure or presuppose absorbing treatment, so this paper adopts local projection difference-in-differences (LP-DiD), which accommodates non-absorbing treatment (Dube et al., 2025). For each of the two treatment definitions, treatment status (binary) and treatment intensity (continuous), the estimation targets the effect of treatment while the tariff reduction remains in force (persistent treatment) and the total effect of a Quota Tariff designation (impulse response). Throughout, the analysis presupposes the Stable Unit Treatment Value Assumption (SUTVA): one product's outcome is unaffected by another product's treatment.

To summarize the results, for the vegetable group, composed of leafy and root vegetables, the estimated retail price reductions are insignificant in every specification, whereas for fruits a significant and persistent reduction is found. The estimated response of fruit retail prices is about -0.78% per one percentage point (%p) of tariff reduction at 150 days after onset (the effect while the tariff reduction remains in force) and $-0.49 \sim -0.55\%$ at 250∼273 days after onset (the total effect of initiation). These magnitudes lie inside the level expected under full pass-through (0.77∼0.94% per 1%p), implying that a substantial share of the tariff reduction was passed through to consumers. These results suggest that refraining from Quota Tariff designations on vegetables, for which no price-reducing effect is identified, would substantially reduce policy costs, and that the Quota Tariff is effective when concentrated on products with a high import share, such as fruits.

The data and code needed to replicate this study are publicly available.[1] The remainder of the paper is organized as follows. Section 2 presents the institutional background and related literature. Section 3 describes the construction of the dataset. Section 4 estimates the causal effect of the Quota Tariff on retail prices using LP-DiD. Section 5 examines the robustness of the causal results. Finally, section 6 summarizes the main findings and discusses their policy implications.

2 Institutional Background and Related Literature

Following the Uruguay Round negotiations, the WTO Agreement on Agriculture converted non-tariff barriers into tariffs and, in exchange, introduced the TRQ to guarantee minimum market access. The TRQ has served as a policy instrument that applies a low tariff rate up to a specified volume and a higher rate to imports beyond it, thereby pursuing expanded market access and the protection of domestic agriculture at the same time. It has since been used as a principal instrument for opening agricultural markets not only under the WTO Agreement on Agriculture but also under FTAs.

Meanwhile, tariffs have been used not only as instruments of market opening under international agreements but also as tools for achieving domestic policy objectives. Although institutional designs differ across countries, temporary tariff adjustments to relieve supply shortages or stabilize prices can be found in a number of countries. Korea's Quota Tariff, the subject of this study, is one such instrument, operated under *Article 71 of the Customs Act* separately from the TRQs of the WTO and FTA frameworks. Since the system operates by Presidential Decree rather than through the legislative process of the National Assembly, the tariff rate and the volume of imports eligible for that rate can be adjusted relatively quickly and flexibly in response to changing market conditions.

The existing literature has mainly analyzed the economic effects of TRQs administered within the WTO and FTA frameworks. Loginova et al. (2021) analyze how

[1] https://github.com/jayjeo/github_TRQ

Switzerland's administration of agricultural TRQs affects domestic price formation. The Swiss government sets annual quotas for agricultural products (excluding meat and dairy) at levels exceeding import demand, thereby compensating for insufficient domestic production capacity while fulfilling its WTO market access quota obligations. During the seasons when domestic produce comes to market, however, the TRQ system is administered so as to protect domestic producers. The authors find that during these protection periods, Swiss vegetable prices were on average more than 20% higher than those in neighboring countries, with price gaps exceeding 50% for some products. These price effects, however, were not uniform across all agricultural products.

In addition, Schaefer and Wolf (2025) analyze the reform of Canada's dairy TRQ system. A dispute-settlement panel under the United States–Mexico–Canada Agreement ruled that Canada's procedures for allocating dairy TRQs were inconsistent with the market-access obligations of the agreement, and Canada revised its TRQ administration accordingly. The study finds that this reform generated additional trade expansion for some dairy products.

The literature has thus focused mainly on the market-access and trade effects of TRQs. By contrast, empirical evidence on the consumer-price effects of tariffs used as an instrument of domestic price stabilization remains insufficient. For Korea's Quota Tariff, Song (2023) analyzes the relationship between the import price excluding tariffs and consumer prices, but this approach falls short of directly identifying the policy effect of the Quota Tariff itself. This study therefore addresses this limitation of the prior literature by causally identifying the consumer-price stabilization effect of the Quota Tariffs deployed to stabilize agricultural prices.

The gap just described has a methodological dimension as well, one that concerns the temporal structure of the treatment that institutions of this kind, the Quota Tariff among them, generate. Institutions in the TRQ family and the Quota Tariff differ in purpose and legal basis, but they share the feature that matters most for causal evaluation: the tariff measures they impose are temporary, starting and ending repeatedly within the sample period. Protection phases of seasonal TRQs open and close every year with the domestic harvest calendar (Loginova et al., 2021; Hillen, 2019); safeguard tariffs under trade agreements are triggered when cumulative imports cross a threshold volume and are lifted afterwards (Livingston et al., 2025); and a Korean Quota Tariff designation likewise lapses when its notified end date arrives or its quota is exhausted, with the same product often designated again later.

The DiD literature distinguishes two temporal structures of treatment. A structure in which treatment, once switched on, remains on until the end of the sample is called absorbing, and one in which treatment can switch off and later switch back on is called non-absorbing (Bach et al., 2025; Dube et al., 2025). The formal definitions appear in section 3, but it is already clear that the treatments administered by the institutions above are all non-absorbing. Even so, empirical studies that use non-absorbing data directly when estimating the causal price effects of these institutions are scarce: to the best of our knowledge, the study of Loginova et al. (2021) reviewed above is the only one. Studies of causal price effects have, conversely, excluded this switching at the design stage.[2] The other studies that have exploited the switching structure of the treatment in estimation had outcome variables other than prices, such as market integration under the Swiss seasonal TRQs (Hillen, 2019) and bilateral trade flows under the beef safeguard of the United States–Japan trade agreement (Livingston et al., 2025). Direct estimation of causal price effects on non-absorbing data thus remains limited to that single study, and it is by adopting this approach that the present paper hopes to contribute.

Whether a study engages the non-absorbing path directly is not a technical detail: it determines which questions about this class of policies can be answered at all. In evaluating a temporary instrument, the key questions are two: whether prices respond while the tariff reduction remains in force, and whether the response outlives the designation itself. These correspond exactly to the two estimands developed in section 4, persistent treatment and treatment initiation. An absorbing design that codes treatment as 'ever designated', by contrast, counts quota-exhausted and expired days

[2]For example, Flaaen et al. (2020) estimate the retail-price effect of the 2018 U.S. washing-machine safeguard, a TRQ charging 20% within a 1.2-million-unit quota and 50% beyond it in its first year; because their sample ends while the policy is still in place, the design treats the policy's adoption as a single absorbing event. Fajgelbaum and Khandelwal (2026) apply the local projection event-study design of Dube et al. (2025), devised for non-absorbing treatments, to the 2025 U.S. tariffs, and yet state explicitly that they define treatment as "an absorbing state that remains on even if the applied rate falls in subsequent months." Even Loginova et al. (2021), the sole exception, uses the switching only statically: it pools the weekly observations of the protected and unprotected phases across years into a contrast between the two phases, estimates no dynamics by horizon, and does not consider the possibility that the effects of the preceding protection phase persist into the following unprotected phase, which serves as the control. That static design is workable because the treatment switching of the Swiss policy follows a regular seasonal calendar; Korea's Quota Tariff, by contrast, switches irregularly, driven by emergency designations and quota exhaustion, so such phase pooling is not available in the first place.

as treated, thereby conflating the effect while the reduction is in force with the effect after its lapse, and so answers neither question. Ignoring the switching structure of the policy treatment also carries an econometric cost: observations still carrying the residual effects of a lapsed episode blend into the control side and contaminate the estimates, and the *clean-control* and *washout* discipline that this study adopts is the device built precisely to prevent this. The contributions of this study therefore extend to (1) the institutional character of its subject, (2) the policy implications of its findings, and (3) the econometric methodology. As the review above shows, no study of which the authors are aware has yet combined these elements.

3 Data Construction

The dataset constructed for this study is a product-level daily panel covering 41 agricultural products from January 1, 2021 through March 31, 2025. Daily retail prices, obtained from *Nongnet*, serve as the dependent variable of the causal analysis. The data on Quota Tariff status, designation start and end dates, and treatment intensity by product (the most critical explanatory variables in the regressions) were compiled from the public records published under the *Regulations on the Application of Quota Tariffs under Article 71 of the Customs Act.*

This section focuses on how the dependent variable and the main explanatory variables are constructed, and on two important confounders for which covariates are built. The remaining covariates are described briefly where they enter the analysis in section 4.

Appendix C presents all summary statistics underlying the LP-DiD analysis, along with additional statistics provided for reference. Table 5 reports summary statistics for every variable entering the LP-DiD regressions. Tables 6 and 7 report the annual import volumes and values of the 41 agricultural products covered in this study, showing which products rely heavily on imports and which do so only marginally.

3.1 Selection of Treated and Controls

Since the analysis relies on LP-DiD, selecting the treated and control groups is essential. A structural feature of LP-DiD, however, is that a treated item can also serve as a control at some horizons (h), as section 4.2 discusses in detail. The selection is therefore stated more precisely in terms of an ever-treated and a never-treated group. An ever-treated item is one that had the Quota Tariff in force at least once during the sample period; depending on the horizon, it functions sometimes as treated and sometimes as a control. A never-treated item is one that never had the Quota Tariff in force during the sample period, and it functions only as a control.

The reasons for including or excluding particular items deserve a somewhat lengthy account, because it matters for the credibility of the design that group membership is determined by observable criteria, not by arbitrary selection on our part. The first point is that the analysis is confined to agricultural products: livestock, fishery, and forest products lie outside its scope, so foods such as chicken and eggs are excluded. Mushrooms, although grouped with agricultural products in official statistics, are legally classified as forest products and are excluded on that ground.[3] Items whose daily retail price series (the dependent variable) are unavailable or mostly missing are likewise excluded: broccoli, cherry, Korean melon, mustard greens, orange, peach, strawberry, and sweet persimmon. Potato is excluded for a different reason. Its Quota Tariff applies only to potatoes for chip manufacturing, whereas retail prices are observed only for table potatoes. The treatment and the dependent variable would therefore refer to different goods, so potato is dropped from the analysis. Grape is likewise excluded: it was designated once in May 2024, but its daily retail price series is unobserved from June 2023 onward, so no dependent variable exists anywhere over its event window. Rice, although an agricultural product, is also excluded: as the staple grain, it is managed under a separate TRQ operation plan as part of the government's priority food policy, and its import and supply decisions therefore differ from those of ordinary agricultural products.

Among the 41 products so selected, any item for which the Quota Tariff was in force at least once during the sample period (January 1, 2021 through March 31, 2025) is classified as ever-treated, and the rest as never-treated. The eleven ever-treated items are listed by group. The leafy and root vegetables are napa cabbage, cabbage, radish, carrot, onion, and green onion; the fruits are mango, kiwifruit, banana, avocado, and pineapple. The remaining 30 products are all never-treated: adzuki bean, bean, bell pepper, cherry tomato, cucumber, dried red pepper, garlic, ginger, green pepper, lettuce,

[3]Forestry and Mountain Villages Development Promotion Act.

mung bean, peanut, perilla leaf, red pepper, scallion, sesame, soybean sprout, spinach, squash, sweet pepper, sweet potato, tomato, water dropwort, young napa cabbage, young summer radish, apple, lemon, melon, pear, and watermelon.

3.2 Quota Tariff Treatment Status for Ever-Treated Items

Table 1 reports, on a daily basis, when each ever-treated item was under the Quota Tariff. Border lines indicate monthly boundaries; white cells denote days on which the Quota Tariff was not in force (*not-treated status*), whereas black cells denote days on which it was in force (*treated status*). Due to space constraints, data prior to 2022 are not presented; however, daily treatment status was constructed for all products from 2015 onward. Designations were exceedingly rare from 2015 through 2022 and increased markedly from 2024 onward. The treatment path used in the estimation is the status recorded in this table with a single adjustment: administrative year-end gaps of at most 30 days within a notice sequence (nine gaps) are bridged, as detailed in section 4.3.[4] Gaps longer than 30 days are never bridged and are used exactly as recorded.

The remainder of this subsection details the criteria by which the black and white cells in table 1 were determined. This determination matters because the treatment status is the foundation of the LP-DiD analysis. In particular, it is what identifies the clean controls and clean treated units at the core of the design, so the status must be constructed accurately from the factual record. As the table shows, the treatment status in this study is 'non-absorbing.' A configuration in which white cells appear first chronologically and, once transitioning to black, never revert to white is termed 'absorbing' (Bach et al., 2025). In contrast, configurations permitting reversion from black to white are termed 'non-absorbing' (Dube et al., 2025), which allows for indefinite white-black-white-black repetition.

A first distinction is between designations announced with a quota limit and those

[4]The bridge is dictated by the design. Without it, each of the nine gaps would end with a new onset on the day the renewed decree takes effect. Since the gaps run 13 to 23 days, far shorter than the 120-day washout, every such re-onset would be dirty by construction and discarded from the treated side. In the persistent-treatment specifications the stay-on condition would fail on the first day of the gap, so the spell of the ongoing event would end because of an administrative formality rather than an economic withdrawal. A single continuous policy episode would thus be split into a short event plus an unusable orphan segment, distorting the very object of measurement, namely the effect of treatment while it remains in force. A detailed account of the washout and stay-on conditions is given in section 4.2.3.

Table 1: Quota Tariff Treatment Status

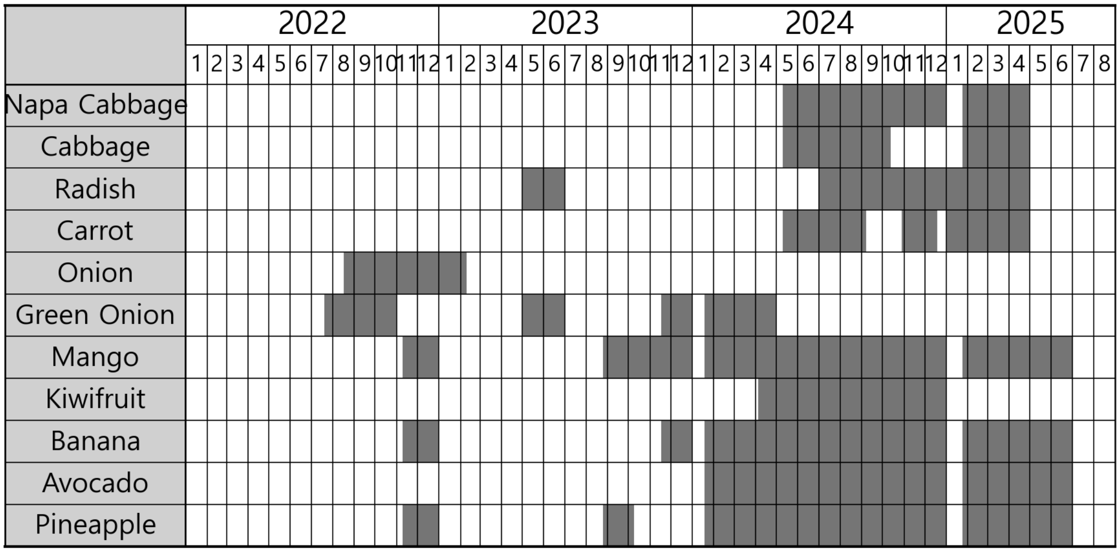

announced without one. Under a limited designation, the tariff relief is available only until cumulative imports, counted from the start of the designation, reach the quota; once the quota is exhausted, the relief lapses even though the announced period has not ended, and the item must be regarded as reverting to not-treated status. Two notions of duration must therefore be kept apart: the *designated period*, announced in advance in the government notice, and the days on which the Quota Tariff is actually *in force*. It is the latter that the treated status in table 1 records. Under an unlimited designation the two coincide, and the relief runs uninterrupted to the end of the designated period.

Determining the exhaustion date requires care, because the cumulative imports counted against the quota are not the item's total imports but only the imports cleared through customs under the Quota Tariff. The Korea Customs Service discloses this narrower series by item, but at monthly rather than daily frequency, so the data reveal the month in which the quota was exceeded but not the exact day. Apart from this limitation, it is the most authoritative source available (Method A). Alternatively, using the monthly total import volume of the item, the exact crossing day can be computed under the assumption that each month's imports are distributed uniformly across its days (Method B). The drawback is that Method B rests on total imports rather than only those cleared under the Quota Tariff, and it therefore overstates the relevant volume. This study combines the two methods to arrive at the most defensible exhaustion date for each episode.

Table 2 illustrates the calculation for pineapple, for which six Quota Tariff designations were announced during the sample period. In the table, Start date and End date delimit the designated period announced in advance by the government, irrespective of any subsequent quota exhaustion, and Quota (ton) is the quota announced alongside it. For No. 1, Method A shows cumulative imports falling short of the quota through the End date, whereas Method B places the crossing on December 21, 2022. Since Method B overstates volumes, the authoritative Method A record prevails, and pineapple is taken never to have exhausted its quota during the No. 1 period. For No. 2, Method A establishes that the quota was exceeded at some point in October 2023, because no imports cleared under the Quota Tariff are recorded for November or December; Method B dates the crossing to October 9, 2023, and this date is adopted. Reflecting this, the cells from October 9, 2023 onward within the No. 2 period are white in table 1 (not-treated status).

Table 2: Pineapple Treatment Detail

No	Start date	End date	Quota (ton)
1	Nov. 10, 2022	Dec. 31, 2022	8600
2	Aug. 25, 2023	Dec. 31, 2023	5000
3	Jan. 19, 2024	June 30, 2024	40000 $\rightarrow$ unlimited
4	July 1, 2024	Sept. 30, 2024	unlimited
5	Oct. 1, 2024	Dec. 31, 2024	unlimited
6	Jan. 24, 2025	June 30, 2025	46000

For No. 3, Method A records imports cleared under the Quota Tariff all the way to the End date. Method B, computed against the initially announced quota of 40,000 tons, places the crossing of cumulative total imports on June 29, 2024, only one day before the End date. Moreover, that quota had been revised to unlimited on April 5, 2024. Under either method, therefore, pineapple never exhausted its quota during No. 3, and the entire designated period appears as black cells in table 1 (treated status). For No. 4 and No. 5, the government announced the quota as unlimited from the outset, so no exhaustion calculation is needed and both periods are black cells throughout. Finally, No. 6 shows no exceedance under either method and is likewise black throughout.

3.3 Anticipation of the Quota Tariff Policy

Identification of the policy effect must contend with anticipation by economic agents. If importers and distributors had fully anticipated a Quota Tariff designation, they would have cut import volumes (deferring imports until the lower tariff took effect) and run down inventories ahead of the expected price decline, and such behavior would affect the estimates.

For the agricultural Quota Tariff studied here, however, the scope for such anticipation is limited. Unlike industrial raw materials and similar items, for which the government announces an annual Quota Tariff operation plan in advance, agricultural designations are mostly emergency responses to current market conditions. Indeed, some measures took effect with no advance legislative notice at all, and where notice was given, a lead of about one month occurred in only a few cases. For most measures the interval between notice and effect was shorter still, and once the decree was promulgated, enforcement typically began immediately or within one to five days.

The agricultural Quota Tariff thus leaves economic agents no practical window in which to anticipate the policy and adjust before it takes effect. Accordingly, the analysis does not treat the legislative notice as a separate policy shock and identifies the policy effect from the date each designation takes effect.

3.4 Treatment Intensity

The LP-DiD analysis is conducted for two treatment definitions. The first is treatment status: the effect on consumer prices of a Quota Tariff being in force, captured in the regressions by a dummy variable for treated status. The second is treatment intensity: the effect on consumer prices of each 1%p of tariff reduction generated by a designation, captured by a continuous variable measuring how far the designation lowers the tariff rate relative to its pre-designation level. Computing this variable requires two pieces of information: the tariff rate charged while the Quota Tariff is in force, which is read directly from the government notice, and the tariff rate that would prevail in its absence, a rate we denote the 'applied tariff rate.'

Deriving the applied tariff rate requires complex calculations across tariff agreements and Basic Tariff Rates for each country, product, and month, a task demanding specialized expertise in customs administration. Multiple rates coexist for each Harmonized Tariff Schedule of Korea (HSK)[5] code, including Basic Tariff Rates and Preferential Tariffs under an FTA, and the priority order of their application is governed by *Article 50 of the Customs Act*. Even for identical products, the rate actually applied may vary with the country of origin, product specifications, and importer circumstances. Since no single statutory rate can represent the tariff burden on a product, an indicator such as the applied tariff rate is needed.

In this study, actual import records were used to determine the 'most favorable tariff rate' applicable to each country's monthly imports, following the tariff rate priority hierarchy established under the *Customs Act*. The monthly applied tariff rate was then derived as a weighted average across countries, with each country's monthly import volume as the weight. Although this approach does not permit direct verification of whether Preferential Tariffs were actually applied in individual transactions, it incorporates both the structural characteristics of the tariff system and observed import patterns, and thus provides a meaningful estimate that closely approximates the tariff burden actually incurred.[6]

3.4.1 Derivation of the Applied Tariff Rate

To calculate the applied tariff rate, this study used country-specific monthly import volumes by product[7] from the *Korea Agricultural Trade Information* (KATI) as weights. Since the data are organized by country, different tariff rates may apply to the same product depending on the exporting country. Accordingly, the applicable tariff rate for each country was first determined following the 'Tariff Rate Priority Structure'

[5]The Harmonized Tariff Schedule of Korea (HSK) is Korea's national tariff nomenclature, built on the Harmonized System (HS) maintained by the World Customs Organization. The HS harmonizes product classification across countries only down to the six-digit subheading level, beyond which each country may create its own finer subdivisions for tariff and administrative purposes. The HSK extends the six-digit HS subheadings to ten-digit tariff lines, at which the Korean government defines individual products and assigns tariff rates. The first six digits of an HSK code are therefore internationally comparable, whereas the remaining four digits are specific to Korea and do not correspond to the national tariff lines of other countries.

[6]A further limitation concerns items for which the law prescribes charging whichever of the ad valorem duty and the specific duty is larger in amount: because the price and weight of individual import transactions are unobservable, the duty actually charged cannot be determined in such cases, and the applied tariff rate is therefore computed from ad valorem rates only.

[7]Ten-digit HSK code.

presented in Appendix A. The monthly applied tariff rate was then calculated as a weighted average of these country-specific rates, using each country's import volume in the corresponding month as the weight. Repeating this calculation for every month from January 2021 through March 2025 yields a product-level panel of monthly applied tariff rates.

For example, as of July 2024, the import shares for banana under HSK code 0803.90-0000 were highest for the Philippines, followed by Vietnam and Ecuador. Ecuadorian banana is not subject to first-priority rates (such as Anti-Dumping Duty, Retaliatory Duty, or Countervailing Duty), and Ecuador is not an FTA partner country; therefore, the review begins with the third-priority rates. Since the fourth- through sixth-priority rates do not apply and the third-priority WTO General Tariff Concession rate (90%) exceeds the seventh-priority Basic Tariff Rate (30%), the Basic Tariff Rate is ultimately applied. Vietnamese banana is subject to a 0% tariff under the Korea-Vietnam FTA, while Philippine banana, which accounts for approximately 61% of imports, is subject to the 30% Basic Tariff Rate. Weighting these country-specific rates by import volume shares yields an applied tariff rate of approximately 23.3% for banana in July 2024. With the Korea-Philippines FTA entering into force on December 31, 2024, the applied tariff rate for banana declined to approximately 14∼15% as of the first quarter of 2025, as illustrated in figure 1.

Figure 1: Applied Tariff Rate (Banana, Onion)

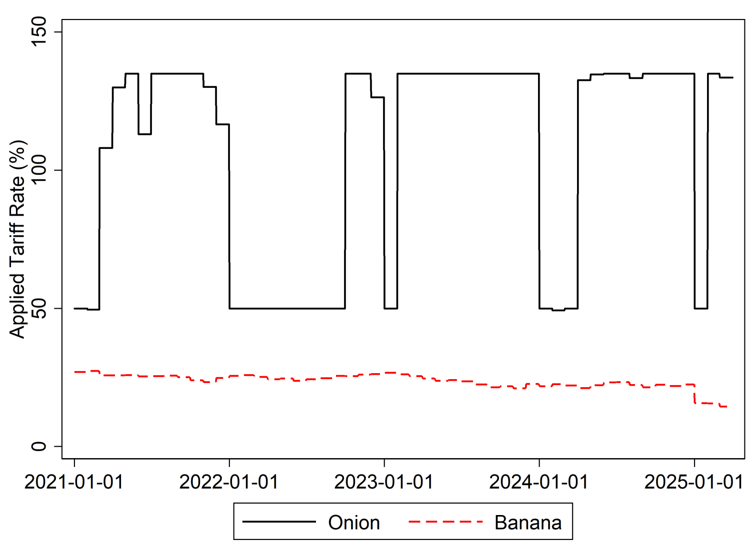

Fresh unpeeled onion is not subject to first-priority rates; thus, the review begins with the second-priority rates. Notably, Chinese onion accounts for the majority of total import volume, but it derives no practical benefit from the FTA Tariff Concessions, and therefore the third-priority rates apply. The third-priority rate is the WTO Tariff Concession applicable to agricultural, forestry, and livestock products with an established Market Access Quota. Imports within the Market Access Quota are subject to the Recommended Tariff Rate of 50%, while imports exceeding the quota are subject to the non-recommended rate of 135%. Consequently, the monthly applied tariff rate for onion varies depending on whether the Market Access Quota is exceeded. This explains the substantial variation, between 50% and 135%, of the applied tariff rate for onion in figure 1. Values that deviate from exactly 50% or 135% reflect onion imports originating from countries other than China.

3.5 Removing Seasonality

This paper uses deseasonalized prices as the dependent variable, a choice justified as follows. Seasonality is removed through a regression-based adjustment using calendar dummy variables and Fourier terms (Pierce et al., 1984), designed to preserve the treatment effects. For daily series, standard official seasonal-adjustment tools (e.g., X-12/X-13-ARIMA-SEATS) are not directly applicable because they are designed for monthly or quarterly data and assume fixed 12-month or 4-quarter cycles (Mollins and Lumb, 2024). An alternative is Seasonal-Trend decomposition based on Loess (STL), a non-parametric method that has been applied extensively to weekly and daily data (Cleveland et al., 1990). Since STL relies on non-parametric smoothing, however, part of a trend or policy shock may be absorbed during the decomposition. This limitation can be critical for policy evaluation with DiD models. If treatment effects are removed at the seasonal-adjustment stage, the subsequent DiD analysis is distorted. The regression-based adjustment with calendar dummies and Fourier terms is chosen precisely because it removes the seasonal repetition while leaving the policy effect intact.

One might nevertheless object that, because the Quota Tariff is deployed precisely to suppress prices when they surge seasonally, the analysis should rely on the original prices rather than the deseasonalized ones. This objection can be countered on two grounds. (a) The pure effect of the Quota Tariff survives deseasonalization. Although prices fluctuate seasonally, the timing of Quota Tariff designations in Korea is far from a regular seasonal pattern, as shown in table 1. Moreover, the adjustment with calendar dummies and Fourier terms removes only recurring patterns, not genuine shocks; the pure effect of the Quota Tariff therefore becomes even more transparent once the seasonal repetition is stripped away. (b) Using the original prices within the LP-DiD framework would obscure the pure effect of the Quota Tariff, because each item has its own seasonal frequency and amplitude, making treated and control items not directly comparable.

The detailed implementation of deseasonalization is documented in Appendix B. As an example, figure 2 presents the original and seasonally adjusted series of spinach retail prices.

Figure 2: Seasonally Adjusted Retail Price (Spinach)

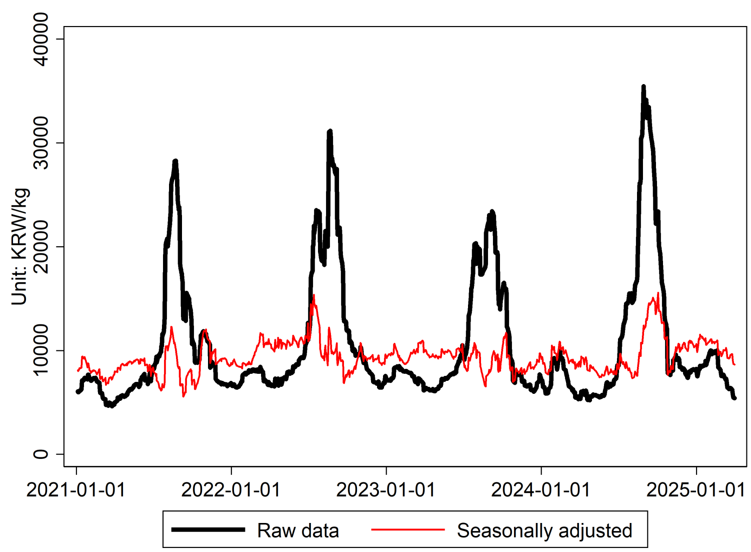

3.6 Confounding Effects

The Quota Tariff is not the only instrument the government deploys against agricultural price surges. Two further instruments share both its trigger (supply shortfalls and surging prices) and its purpose of stabilizing consumer prices: temporary expansions of WTO Market Access Quotas and releases of government-stockpiled produce. Either can move retail prices at the same time as a Quota Tariff designation, so leaving them uncontrolled could conflate their effects with the effect under study. The estimation therefore includes covariates for both, collected in the blocks Q_{it} and R_{it} of the covariate vector in section 4.3.

The first instrument operates within the conventional WTO-style TRQ, from which section 2 distinguishes the Quota Tariff: imports within a scheduled Market Access Quota enter at a low in-quota rate, while imports beyond it face a substantially higher over-quota rate. Coverage is limited to the items for which Korea's WTO concessions schedule a quota (onion, garlic, and ginger among them, but not banana, avocado, or napa cabbage), and only covered items can therefore experience an expansion. For covered items the government may temporarily enlarge the quota when supply–demand imbalances intensify, as it has done for sesame and onion. Since an expansion admits a larger import volume at the low in-quota rate, it is a tariff-relief measure deployed on the same grounds as the Quota Tariff. The expansion history is compiled from the enactment records and coded by a simple rule: the quota is administered by calendar year, with cumulative imports counted from January 1st, so each expansion takes effect on its enactment date and remains in force through December 31st of that year, and a further expansion enacted later in the same year supersedes the earlier one from its own effective date.[8]

The block Q_{it} accordingly contains three covariates: a coverage dummy for items under a WTO TRQ; the expansion ratio in force on day t (the additional volume granted by the active expansion relative to the scheduled quota, equal to zero when no expansion is in force); and the expansion ratio interacted with the tariff gap between the over- and in-quota rates, which scales each expansion by the size of the tariff gap it relieves.[9] Meanwhile, the two kinds of zero are not equivalent: for an item outside the WTO TRQ system a zero expansion is forced (there is no quota to enlarge), whereas for a covered item it records the government's choice not to deploy an available instrument. The coverage dummy is what keeps these two zeros distinct, which is why it enters the covariates alongside the expansion terms.

The second instrument is the release of government stockpiles: the government

[8] The one exception is ginger, whose 2023 expansion lapsed on 30 September 2023 by its own terms.

[9] All three are defined on every product-day, zero rather than missing, so no observation drops from the regressions.

purchases and stores staple produce in normal times and releases it onto the market when supply tightens or prices surge, with napa cabbage, radish, and dried red pepper among the prominent cases. A release presupposes a stockpile, so a recorded zero means different things for different products. For most items no stockpile program exists, so the release value is necessarily zero. However, for stockpiled items a zero records a month in which the program chose not to release. The release records are monthly totals in tons; mapped onto the daily panel, each monthly total applies to every day of its month, and products absent from the records are coded zero. The block R_{it} contains two covariates built from these monthly volumes: the inverse hyperbolic sine of the average release over the trailing window $[t-130,\ t-31]$, which ends before the month containing day t, and an incidence dummy equal to one when any release falls in that window. The inverse hyperbolic sine transformation serves to compress the heavy right tail of the volume distribution while keeping the many zero values in the sample (a simple log transformation would turn the zeros into missing values) (Bellemare and Wichman, 2020). The incidence dummy is introduced to separate whether releases were under way (a binary indicator) from how large they were (a continuous variable). Closing the window before the current month ($[t-130,\ t-31]$) is meant to ensure that only release information from before the Quota Tariff treatment date is taken into account. Since the release information is monthly, without this measure, predetermined release information could spill over past the Quota Tariff treatment date and distort the estimated effect of the treatment.

The daily panel of the eleven ever-treated products (January 2021 through March 2025) shows how the covariates derived from the WTO TRQ system (Q_{it}) and from government stockpile releases (R_{it}) correlate with the timing of Quota Tariff designations. The quota-expansion channel barely touches the treated set: onion is the only treated product covered by a WTO TRQ, and its single expansion episode (July∼December 2023) does not intersect its Quota Tariff designation (August 2022∼February 2023), so no product-day in the sample carries both instruments at once. The expansions of the period fall instead on never-treated items (garlic, bean, mung bean, sesame, peanut, and ginger), whose prices enter the comparison group, and the Q_{it} covariates absorb those movements where they occur.

The stockpile channel, by contrast, overlaps the treatment substantially: the three stockpiled products among the eleven (onion, napa cabbage, and radish) all began their longest designation episodes in months with active releases (3,527, 955, and 1,959 tons, respectively), and releases were under way in the same month on 70% (573 of 815) of their days under the Quota Tariff. Within-product correlations between daily Quota Tariff status and the two release covariates are positive though moderate, 0.11∼0.16 over the eleven products and 0.23∼0.33 within the stockpiled three. The two instruments evidently respond to the same market conditions, which is exactly why the release covariates belong in the regressions: omitting them would risk crediting the Quota Tariff with the stabilizing effect of the releases. Section 4.4 examines how prices themselves move around designations, and how much of that movement the covariates absorb, once the estimating equations are in place.

4 Causal Analysis

This section estimates the causal effect of the Quota Tariff on retail prices. The argument proceeds in six steps. Section 4.1 explains why the treatment structure documented in section 3 calls for LP-DiD rather than conventional event-study estimators. Section 4.2 distinguishes the two estimands the design delivers (the effect of persistent treatment and the total effect of treatment initiation) and develops the two sample restrictions on which both rest, clean controls and clean treated units. Section 4.3 translates this design into estimating equations. Section 4.4 then pauses, before any estimate is shown, to probe the conditional parallel-trends assumption on which the design rests, a step the targeted nature of the policy makes necessary. Section 4.5 presents the estimates as dynamic response graphs, and section 4.6 condenses them into a regression table at selected horizons.

4.1 Why Local Projection Difference-in-Differences

Three features of the Quota Tariff, established in section 3, jointly determine the choice of estimator. First, adoption is staggered: products enter treatment at different dates. Second, treatment is non-absorbing. A treatment is called *absorbing* when a unit, once treated, remains treated for the rest of the sample, and *non-absorbing* when treatment can switch off and on again. Section 3.2 sets out this distinction for our data. The Quota Tariff is of the latter kind: a designation lapses at its notified end date or, for quota-limited notices, as soon as the quota is exhausted, and the same product may be designated again later. Pineapple, for example, cycles through three separate treatment episodes within the sample period. Third, treatment intensity is heterogeneous: the implied tariff reduction ranges from 1.5%p to 69.9%p across products (section 3.4). The analysis accordingly asks two questions on the same panel: whether Quota Tariff adoption as such lowers retail prices (a binary treatment-status question), and by how much retail prices fall per 1%p of tariff reduction (a continuous treatment-intensity question).

Conventional two-way fixed effects (TWFE) event-study regressions are ill-suited to this structure. When treatment turns on and off and effects are heterogeneous, such regressions implicitly use previously treated observations as comparisons, although lagged effects of an earlier episode may still be unfolding. The resulting contamination can severely bias the estimates, including through negative weighting (Goodman-Bacon, 2021; De Chaisemartin and D'Haultfœuille, 2020). In this application the problem is concrete. Retail prices of pineapple in early 2024 are a poor counterfactual for anything, because the effects of its 2023 episode may persist well after quota exhaustion ended that episode.

The recent heterogeneity-robust estimators confront this contamination through a shared route: they secure valid comparisons by restricting the controls to observations that are not yet treated or never treated, rather than relying on the implicit comparisons that a conventional TWFE regression forms on its own (Roth et al., 2023; de Chaisemartin and D'Haultfœuille, 2023). The interaction-weighted estimator of Sun and Abraham (2021) is constructed for the absorbing case: it characterizes each unit by a single initial-treatment date and aggregates cohort-by-relative-time effects on that basis, paralleling the irreversibility-of-treatment assumption of Callaway and Sant'Anna (2021). However, the Quota Tariff is non-absorbing, and this structure is common to *both* the treatment-status and the treatment-intensity specifications. Forcing the design into an absorbing framework by recoding treatment as an ever-treated indicator would change the estimand to the effect of *having ever* been designated and would misclassify terminated and quota-exhausted periods as treated. Relabeling each onset as a fresh absorbing cohort fails as well, because several episodes adjoin so closely that any inter-episode "pre-period" would be contaminated by the preceding episode.[10]

[10] Sun and Abraham (2021), Callaway and Sant'Anna (2021), Borusyak et al. (2024), and Gardner (2022)

This study therefore employs LP-DiD (Dube et al., 2025), which is constructed for non-absorbing treatments: each horizon-specific regression compares newly treated observations only with contemporaneously[11] clean comparisons. Three further properties make it the natural choice here. It conditions directly on pre-determined time-varying covariates, which the specification below requires. It nests the binary and the intensity treatments in one framework, so that differences between the two sets of estimates reflect the treatment definition rather than a change of estimator family. And it extends naturally to a treatment that is, as here, simultaneously non-absorbing and continuous: Dube et al. (2025) develop the clean-control machinery for binary non-absorbing treatments and flag continuous treatment as a natural extension, which this study implements through a pre-determined intensity scalar (see Dube and Lindner, 2024 for an LP-DiD application with non-absorbing treatment). LP-DiD does not rest on weaker identifying assumptions than the cohort estimators (both families require parallel trends and no anticipation), and in canonical absorbing designs it performs comparably; the issue is scope, not general superiority. In other words, LP-DiD is necessary here because it can accommodate non-absorbing treatments.

Throughout, the analysis maintains the potential-outcomes framework under SUTVA acknowledged in the Introduction: one product's potential outcome is unaffected by another product's treatment status. Meanwhile, the credibility of LP-DiD rests not on the regression equation itself but on the discipline it imposes: strict rules about which observations may serve as comparisons. The next subsection first distinguishes the two estimands this design delivers, and then develops that discipline.

4.2 Core Design Elements of Local Projection Difference-in-Differences

The LP-DiD design has two components. The first is a choice of estimand: the same LP-DiD framework can measure either the effect of a designation while it remains in force (*persistent-treatment*) or the total effect of having initiated one (*treatment initiation*). The second is a pair of sample-selection rules (clean controls and clean treated units).

propose improved estimators for staggered designs, but their scope conditions presume that treatment, once switched on, stays on. Once a product's status switches on and off, the initial-treatment date no longer characterizes its treatment path. The binary coding of the treatment-status specification is thus necessary but not sufficient for those estimators to apply.

[11] At the same horizon.

This subsection introduces the two estimands, then develops the two sample-selection rules, and explains how these restrictions interact with each estimand at horizon h.

4.2.1 Two Estimands: Persistent Treatment versus Treatment Initiation

The *persistent-treatment* scheme asks what happens to prices while a designation is actually in force. It keeps a treatment event in the horizon-h regression only while the treatment persists, and it stops tracking the event once the designation lapses. The *treatment initiation* scheme asks a different question: what is the total effect of having initiated a designation, including whatever happens after the episode ends or the product re-enters? It follows every initiation across all horizons, letting the post-initiation treatment path unfold as it actually did. The resulting coefficient path is an impulse response function (IRF): the average effect of treatment *initiation* h days later. This is the impulse-response-type estimand discussed by Dube et al. (2025), and it corresponds to the actual-versus-status-quo event-study effects formalized by De Chaisemartin and D'Haultfœuille (2026).

Neither estimand is more correct than the other. Dube et al. (2025) regard the choice as a matter of the research question; the two are thus complements here. The persistent-treatment estimates speak to whether the policy lowers prices while it operates. The impulse-response estimates speak to whether initiation effects outlive the designations themselves, a question of first-order policy interest for a temporary instrument. As shown below, the two schemes differ in exactly one sample condition on the treated side. Both rest on the same pair of restrictions, to which the discussion now turns.

4.2.2 Clean Controls

For a regression at horizon h, the comparison group at date t consists of product–day observations that have never been treated up to t and remain untreated through $t + h$. This "clean control" design originates in the stacked event-study designs of the minimum-wage literature, which coined the term (Cengiz et al., 2019), and is formalized as a flexible clean control condition in Dube et al. (2025). Its purpose follows directly from the contamination problem above. Units whose own treatment dynamics are still unfolding cannot reveal the untreated counterfactual, so they are excluded from the comparison set at every horizon at which they could contaminate it. The same comparison group serves both estimation schemes: *persistent-treatment* and *treatment initiation*.

The implementation here is deliberately stricter than the theory requires. The window-based condition of Dube et al. (2025) would readmit a previously treated product once its dynamics have stabilized. This study instead removes a product from the control pool permanently after its first treatment, so that an ever-treated product contributes control observations only before its first onset. Since 30 of the 41 products are never treated (section 3.1), a deep pool of clean controls remains available at every horizon, and identification never relies on recycled post-treatment observations.

4.2.3 Clean Treated

A parallel discipline applies to the treated side, and it is on this side that the two schemes (*persistent-treatment* and *treatment initiation*) diverge. Two separate conditions below govern which treatment events enter the estimation, and they play different roles.

The first is the *washout* condition: the product must have been untreated throughout the 120 days preceding the onset, that is, over $[t-120,\ t-1]$ (the washout length $L = 120$). This condition is what makes an onset *clean*. It ensures that the outcome dynamics attributed to the event reflect that event alone rather than the residual dynamics of a recent previous episode, and it implements the assumption of Dube et al. (2025) that the dynamic effects of a past episode stabilize within L periods. Since cleanliness concerns the past, the washout condition applies to all specifications, the impulse-response specifications included.

The second is the *stay-on* condition: a treated event remains in the horizon-h regression only if the product is treated on every day of $[t,\ t+h]$. This condition is not about cleanliness. It is what defines the persistent-treatment estimand (the effect of treatment that is still in force at $t+h$), and it therefore applies to the persistent-treatment specifications only. The impulse-response specifications drop exactly this condition and nothing else: a clean onset stays in their sample at every horizon, whatever happens to the designation afterwards.

Taken together, the two conditions reproduce the treated-side sample restriction that Dube et al. (2025) impose for non-absorbing treatments. Their "clean treated units" are, stated precisely, units that are untreated throughout $[t-L,\ t-1]$, enter treatment at t, and remain treated through $t+h$. The same design logic runs through the stacked event-study literature: it originates in Cengiz et al. (2019), is formalized for stacked estimation by Wing et al. (2024), and is applied to non-absorbing minimum-wage events by Dube and Lindner (2024).

The treated-side and control-side conditions differ in a way worth stating explicitly. A treated event needs a clean *recent history* (a 120-day washout before its own onset), whereas a control observation needs a clean *entire history* (no treatment ever, up to the comparison window). Pineapple illustrates both the washout rule and its practical force. Of the six notices in table 2, bridging of administrative gaps yields three distinct onsets: November 10, 2022; August 25, 2023; and January 19, 2024. The first onset is clean by construction (no prior designation). The second is clean because roughly eight months separate it from the end of the first episode, well beyond the 120-day washout. The third is *dirty*: the preceding episode's last treated day was October 8, 2023, after its quota was exhausted (section 3.2), only 102 days before the new onset, so the washout condition fails and the January 2024 onset is excluded from the treated side under both schemes. Nor can any post-2022 pineapple observation serve as a control, since the product's treatment history is no longer clean. Two further onsets (carrot in October 2024 and cabbage in January 2025) fail the washout in the same mechanical way (Appendix D). Each classification follows mechanically from table 1; no discretion enters.

4.2.4 Treated-Unit Attrition over Horizons

The stay-on condition has a mechanical consequence that shapes how far the persistent-treatment estimates can be read. An event whose treatment spell lasts s days satisfies the stay-on condition only for horizons $h < s$; as h grows, events with shorter spells drop out of the treated set one by one. Among the eleven ever-treated products, nine vegetable events and eight fruit events survive the washout screen at $h = 0$. The exits occur as the corresponding spells end. The vegetable treated set shrinks step by step at $h = 61$, 104, 120, 156, 166, and 177, and the fruit set at $h = 45$, 52, and 271 (Appendix D provides the details). Each such change means that estimates at distant horizons compare a different collection of events, a "composition effect" that Dube et al. (2025) flag as inherent to horizon-specific estimation.

The data locate the point at which this composition effect becomes severe enough to disqualify the estimates. At exactly $h = 156$ (the day the cabbage event's spell ends and the event exits the treated set), the vegetable coefficient in the treatment-status specification jumps from $+0.079$ to $+0.237$ and turns statistically significant on impact; the surviving far-horizon treated set comes to be dominated by the radish and napa cabbage events of 2024, whose post-treatment windows coincide with the severe 2024 supply disruptions in exactly those markets. A discontinuity that arrives on the day of a compositional break, rather than growing out of the preceding dynamics, is a composition artifact, not a treatment dynamic. The persistent-treatment graphs are therefore truncated at $h = 150$, a conservative margin below the break at 156. The underlying estimates for all horizons up to 273 are retained in the replication files. Attrition, however, is not a defect of the design but a direct consequence of the stay-on condition, which is precisely what defines the persistent-treatment estimand. The impulse-response specifications, which drop that condition, are free of attrition altogether: their treated counts stay at nine and eight across all horizons, and the impulse graphs extend to $h = 273$, the maximum horizon the sample supports.

4.2.5 Scope and Limitations of the Impulse Estimand

The persistent-treatment and impulse-response specifications share two further design details. First, as noted above, the washout is retained both in the persistent-treatment and impulse-response specifications: a past episode is a separate event, not a reason to backdate the current one. The pineapple onset of January 2024 is therefore excluded from both specifications. The 2022 pineapple event, by contrast, enters both estimation samples, though over different spans: the persistent-treatment regressions include it from $h = 0$ through $h = 51$, whereas the impulse-response regressions include it from $h = 0$ through $h = 273$. Both specifications admit the event from $h = 0$ because it satisfies the washout condition. The persistent-treatment span ends at $h = 51$ because the stay-on condition fails beyond that horizon (the tariff cut under the pineapple Quota Tariff ends on the 51st day after the designation takes effect), and the impulse-response span runs through $h = 273$ because that specification drops the stay-on condition by construction. Second, the control-side condition is identical in the two specifications: comparisons are still required to remain untreated through $t + h$. The impulse-response estimand is thus the total effect of initiation *under future-untreated* *controls*, not a fully unconditional impulse response.

The impulse estimand also has a limitation that should be stated plainly. Treatment durations differ widely across events (spell lengths range from 45 to 585 days, the longest right-censored by the end of the sample), and the impulse coefficient (denoted β^h_{imp} below) averages over initiations without adjusting for the number of treated days each event accumulates within the window. It is therefore a per-initiation effect, not a per-treated-day effect. The same convention governs the intertemporal framework of De Chaisemartin and D'Haultfœuille (2026), whose event-study effects are likewise defined by the change in treatment alone, with no duration term entering the estimation. Read this way, β^h_{imp} answers the question a policymaker faces at the moment of designation: what initiating a Quota Tariff sets in motion for retail prices h days ahead.

4.3 Regression Specification

Following the pre-mean-differenced (PMD) LP-DiD of Dube et al. (2025), the dependent variable at horizon h is the log retail price h days after the event, measured against the average log price over the 120 days preceding the event:

$$\bar{y}_{i,t-1} \equiv \frac{1}{120}\sum_{s=1}^{120} \ln P_{i,t-s}, \qquad \widetilde{\Delta}^h \ln P_{it} \equiv \ln P_{i,t+h} - \bar{y}_{i,t-1}, \qquad -500 \le h \le 273, \tag{1}$$

where i and t denote product and calendar day and P_{it} is the seasonally adjusted retail price (KRW/kg) of section 3.5. The pre-mean baseline is fixed before the event and is not an ad hoc choice: Dube et al. (2025) develop this specification, their PMD LP-DiD, on the efficiency ground articulated by Borusyak et al. (2024), namely that the pre-event outcome level is estimated more precisely from many pre-treatment days than from the single day $t-1$ (the `lpdid` package in *STATA* and *R* implements PMD as a built-in option (Busch and Girardi, 2024)). The argument has particular force in daily retail price data, where anchoring the baseline to a single day would transmit that day's idiosyncratic movement into the dependent variable at every horizon.[12] Since equation

[12] The 120-day window matches the washout length and was fixed ex ante, in line with the recommendation of Dube et al. (2025) that the reference period be chosen on the characteristics of the application, independently of the estimation results. The size of the efficiency gain depends on the serial correlation of untreated outcomes (Borusyak et al., 2024; Dube et al., 2025). PMD does not by itself guarantee parallel trends; the pre-event coefficients reported in section 4.5 remain the relevant diagnostic.

(1) differences each observation against its own pre-event level, product fixed effects are differenced out by construction and do not appear in the regressions.

For each horizon h, the treatment-status (binary) specification is the cross-sectional[13] regression

$$\widetilde{\Delta}^h \ln P_{it} = \beta^h \Big(D_i \cdot \Delta D_{it} \Big) + T_t + X_{it} + \varepsilon_{it}^h, \tag{2}$$

estimated, for $h \geq 0$, on the sample

$$\begin{aligned} \mathcal{S}^h_{\text{base}} = &\underbrace{\Big\{ (i,t) :\ \Delta D_{it} = 1,\ \ D_{i,t-s} = 0 \; \forall s \in [1,120],\ \ D_{i,t+j} = 1 \; \forall j \in [0,h] \Big\}}_{\text{clean treated: fresh entry, remaining treated through } t+h} \\ &\cup \underbrace{\Big\{ (i,t) :\ D_{is} = 0 \; \forall s \leq t,\ \ D_{i,t+j} = 0 \; \forall j \in [1,h] \Big\}}_{\text{clean control: no history, untreated through } t+h}. \end{aligned} \tag{3}$$

Here $D_{it} \in \{0,1\}$ is the daily Quota Tariff treatment status of table 1. D_i indicates ever-treated products, and T_t are calendar-day fixed effects, so that every event-time comparison is simultaneously a same-calendar-date comparison, absorbing market-wide movements. The two sets in equation (3) are constructed separately rather than by matching treated to control observations one-to-one, so the sample definition does not require the dates t of the two sets to coincide; the horizon-h regression pools all observations from both sets into a single sample, and the same-calendar-date comparisons are carried out not by the sample definition but by the calendar-day fixed effects just described. For $h < 0$ the dependent variable is the corresponding lag, and the forward conditions in equation (3), namely the stay-on condition $D_{i,t+j} = 1 \; \forall j \in [0,h]$ on the treated side[14] and the future no-treatment condition $D_{i,t+j} = 0 \; \forall j \in [1,h]$ on the control side, are dropped:

$$\begin{aligned} \mathcal{S}^h_{\text{base}} = &\underbrace{\Big\{ (i,t) :\ \Delta D_{it} = 1,\ \ D_{i,t-s} = 0 \; \forall s \in [1,120] \Big\}}_{\text{clean treated: fresh entry, washout only}} \\ &\cup \underbrace{\Big\{ (i,t) :\ D_{is} = 0 \; \forall s \leq t \Big\}}_{\text{clean control: no history through } t}, \qquad h < 0. \end{aligned} \tag{4}$$

[13] Although the data used in this study form a panel, at each horizon h LP-DiD first selects every observation satisfying the clean-control and clean-treated conditions and then estimates the equation on that sample as a cross-sectional rather than a panel regression.

[14] The $j = 0$ component is already implied by $\Delta D_{it} = 1$.

The washout condition on the treated side requires only that a product be untreated during the 120 days immediately preceding treatment onset; it imposes no restriction on the product's treatment history before that window. At horizons $h < -120$, therefore, an event is not excluded from the sample merely because $t + h$ falls within an earlier episode of the same product. On the control side, the no-history condition $D_{is} = 0 \, \forall s \leq t$ already covers every date up to t, so no additional condition is needed at horizons $h < 0$.

The key regressor ΔD_{it} equals one only on the day product i *enters* treatment ($D_{i,t-1} = 0$ & $D_{it} = 1$). It is an onset indicator, not a treatment level: the regression compares the price path following a fresh designation with the contemporaneous paths of clean controls, while everything concerning the duration of treatment is handled by the sample restriction, not by the regressor. For $h \geq 0$, β^h therefore measures the average log price gap h days after a clean onset, among events still under treatment; for $h < 0$, β^h is the pre-onset price gap that underlies the pre-trend diagnostics of section 4.4 below.

The treatment-intensity (continuous) specification scales the onset indicator by the product-specific, pre-determined intensity G_i, in %p:

$$\widetilde{\Delta}^h \ln P_{it} = \beta^h \Big(G_i \cdot D_i \cdot \Delta D_{it} \Big) + T_t + X_{it} + \varepsilon^h_{it}, \qquad \text{sample } \mathcal{S}^h_{\text{base}}, \tag{5}$$

so that β^h is interpreted per 1%p of tariff reduction. Relative to the binary specification, reading β^h as a per-1%p causal effect adds a linear dose–response restriction (the horizon-h effect of an onset is taken to be proportional to G_i); absent strict proportionality, β^h retains its interpretation as an intensity-weighted average of event-level effects expressed per 1%p. For each onset, G_i is the difference between the applied tariff rate and the Quota Tariff rate, the tariff reduction the designation delivers.[15] Table 3 reports the resulting intensities.

The impulse specifications estimate the same equations (2) and (5); only the sample changes. The stay-on condition is removed from the treated side, while the control side

[15] Products with multiple onsets receive the average across their onsets, yielding a single time-invariant scalar per product. Averaging over the pre-onset year smooths short-term fluctuations in the applied tariff rate, and using only pre-onset information keeps G_i pre-determined: tariff rates realized after the onset, which could respond to the treatment itself, never enter. Section 3.4 details the construction of the applied tariff rate.

Table 3: Quota Tariff Treatment Intensity

Product	Applied Tariff Rate (%)	Quota Tariff Rate (%)	Treatment Intensity (%p)
Napa cabbage	27.0	0	27.0
Cabbage	26.9	0	26.9
Radish	29.4	0	29.4
Carrot	28.7	0	28.7
Onion	79.9	10	69.9
Green onion	27.0	0	27.0
Mango	15.7	0	15.7
Kiwifruit	6.5	5	1.5
Avocado	8.0	0	8.0
Pineapple	29.6	0	29.6
Banana	24.5	0	24.5

is unchanged:

$$\begin{aligned} \mathcal{S}^h_{\text{imp}} = &\underbrace{\Big\{(i,t):\ \Delta D_{it} = 1,\quad D_{i,t-s} = 0\ \forall s \in [1,120]\Big\}}_{\text{clean treated: fresh entry, subsequent path unrestricted}} \\ &\cup \underbrace{\Big\{(i,t):\ D_{is} = 0\ \forall s \leq t,\quad D_{i,t+j} = 0\ \forall j \in [1,h]\Big\}}_{\text{clean control: as in (3)}}, \qquad h \geq 0, \end{aligned} \tag{6}$$

with coefficients denoted β^h_{imp}. For $h < 0$, the future no-treatment condition on the control side, $D_{i,t+j} = 0\ \forall j \in [1,h]$, is dropped, as in the persistent-treatment specification; the treated-side condition remains unchanged, since it contains no forward-looking element to begin with:

$$\begin{aligned} \mathcal{S}^h_{\text{imp}} = &\underbrace{\Big\{(i,t):\ \Delta D_{it} = 1,\quad D_{i,t-s} = 0\ \forall s \in [1,120]\Big\}}_{\text{clean treated: unchanged}} \\ &\cup \underbrace{\Big\{(i,t):\ D_{is} = 0\ \forall s \leq t\Big\}}_{\text{clean control: no history through } t}, \qquad h < 0. \end{aligned} \tag{7}$$

This coincides exactly with the persistent-treatment sample in equation (4). At horizons $h < 0$, the two specifications therefore share the same sample and the same regression, so their pre-event coefficient paths are identical; this coincidence extends to $h = 0$ as well.

The treated-side conditions (the washout condition, common to all four specifications, and the stay-on condition, specific to the persistent-treatment ones) and the dependent variable are used exactly as prescribed by Dube et al. (2025) and as implemented in the `lpdid` package for *STATA* and *R* (Busch and Girardi, 2024). The control-side condition is stricter than the window-based clean control condition that Dube et al. (2025) and Busch and Girardi (2024) employ: their implementation readmits a once-treated product to the clean control pool after a sufficient interval, whereas the implementation in this study never readmits a treated product, regardless of how much time has passed. This setting places no strain on the design, because the large majority of products (30 of 41) are never treated and clean controls therefore remain abundant at every horizon.

The one element that goes beyond their binary framework is the intensity scaling in equation (5): they develop the non-absorbing conditions for a binary treatment and flag their application to a continuous treatment as a direction for future extension. The intensity specification in this study implements precisely that extension, and its interpretation follows the dose framework of De Chaisemartin and D'Haultfœuille (2026).

Standard errors are clustered by product throughout, and 95% confidence bands use the normal critical value, matching the construction of the figures. Two caveats are in order. First, the number of *treated* clusters is small (six vegetable and five fruit products), so cluster-robust intervals at long horizons deserve caution (MacKinnon and Webb, 2018); section 5.1 examines this issue and the available alternatives in detail. Second, at far horizons of the persistent-treatment specification the treated set itself thins (section 4.2). Both cautions inform the reading of the results.

The composition of the covariate vector X_{it} differs between the two treated groups, and each group is estimated in a separate regression. The comparison pools of the two regressions share a common base, the 30 never-treated products, supplemented in each case with the pre-onset observations of the group's own ever-treated products.[16] For the vegetable group,

$$X_{it}^{\text{veg}} = A_{it} + I_i \cdot O_t + I_i \cdot W1_{it} + I_i \cdot W2_{it} + R_{it} + Q_{it},$$

[16] Treated Group 1 comprises the six leafy and root vegetables (napa cabbage, cabbage, radish, carrot, onion, and green onion); Treated Group 2 comprises the five fruits (mango, kiwifruit, banana, avocado, and pineapple). When one group is estimated, the other group's products are excluded from the sample. Estimating the groups separately, rather than pooling all treated products, is not cosmetic: the two groups turn out to respond in opposite directions (section 4.5), so a pooled coefficient would average a negative fruit response against a positive vegetable response and conceal both.

where A_{it} is the applied tariff rate (the tariff that would prevail absent the Quota Tariff; section 3.4.1), O_t is the oil price in KRW interacted with product dummies I_i to allow heterogeneous transport-cost channels, and $W1_{it}$ and $W2_{it}$ are climate covariates in the spirit of Roberts and Schlenker (2013): station averages of daily temperature, humidity, precipitation, and sunshine over the window $[t-100,\ t]$ ($W1_{it}$), and the same averages one year earlier (over $[t-465,\ t-365]$, $W2_{it}$), each interacted with product dummies. Roberts and Schlenker (2013) state that current climate shifts supply directly; the previous year's climate shapes planting decisions and hence current-year supply. R_{it} and Q_{it} collect the government stockpile-release covariates and the WTO market-access-quota covariates, constructed and motivated in section 3.6. For the fruit group the climate blocks are omitted:

$$X_{it}^{\text{fruit}} = A_{it} + I_i \cdot O_t + R_{it} + Q_{it}.$$

The treated fruits are predominantly imported, so Korean weather is not the relevant supply shifter for them. Empirically, fruit pre-event coefficients remain flat without the climate covariates, whereas removing them from the vegetable regressions re-opens the pre-event run-up documented in section 4.4 below (figure 4). Domestic weather drives both vegetable supply and the timing of vegetable designations, making the climate blocks essential covariates there.

Meanwhile, identification rests on the two assumptions standard for LP-DiD: no anticipation, and parallel trends of treated products relative to clean controls, conditional on X_{it}, the calendar-day effects, and product fixed effects.[17] First, anticipation has little room to operate: most designations take effect either with no advance legislative notice at all or with a lead of only one to five days, and a lead approaching one month occurred in just a few cases (section 3.3), leaving importers and distributors no practical window to adjust before enforcement. The pre-event coefficients of the four estimating specifications, reported in section 4.5, speak directly to both assumptions.

Second, the LP-DiD results can be interpreted causally only if the conditional parallel-trends assumption holds. A convincing verification of this assumption must therefore be presented. The Quota Tariff is, by statutory design, aimed at products whose prices

[17] Product fixed effects are absent from equations (2) and (5) by construction rather than omission: the dependent variable of equation (1) already subtracts each product's own pre-event mean, differencing out any time-invariant product component.

have surged, so treated products may be on price paths unlike those of the comparison pool precisely when treatment begins. The next subsection therefore takes the conditional parallel-trends assumption to the data.

4.4 Pre-Designation Price Dynamics

The concern with the conditional parallel-trends assumption is an Ashenfelter dip in reverse (Ashenfelter, 1978). If Quota Tariff designations are systematically triggered by rising prices, a subsequent price decline occurring for reasons unrelated to the tariff cut could be mistaken for an effect of the policy. This subsection first documents how retail prices actually move around designations, and then shows that the pre-designation run-up in vegetable prices is adequately absorbed by the climate covariates.

Figure 3 aligns every designation event on its onset date and traces average retail prices in event time, in the metric the estimation itself uses: the log retail price minus its own average over the 120 days preceding the onset.[18] The comparison series applies the same transformation to the 30 never-treated products, aligned on the same calendar dates, so market-wide movements enter both curves symmetrically. For vegetables the pattern is stark. Vegetable prices climb from about 7 log points below the never-treated comparison four months before designation to 9$\sim$14 log points above it in the final three weeks, and the climb appears in eight of the nine events. For fruits there is no counterpart. The two curves stay within about 4 log points of each other throughout.

Whether this movement threatens identification depends on whether it survives the conditioning that the estimation applies. Figure 4 reports event-study coefficients at pre-onset horizons $h \in [-120, -1]$ from the binary specification in equation (2) (the same baseline-differenced outcome, calendar-day fixed effects, clean-control comparisons, and covariates), estimated with and without the climate blocks $I_i \cdot W1_{it}$ and $I_i \cdot W2_{it}$ on a common sample. Without the climate covariates the vegetable path reproduces the run-up of figure 3 in coefficient form: a significantly negative cluster deep in the pre-onset window, followed by significantly positive coefficients on the 39 consecutive days preceding onset, between $+5$ and $+12$ log points. With the climate

[18] The events are those that enter the estimation, namely onsets preceded by at least 120 treatment-free days (section 4.2): nine vegetable onsets and eight fruit onsets. By construction each series averages to zero over its own 120-day pre-onset window, here and in figure 4; the diagnostic content lies in the path within the window, not its level.

Figure 3: Retail Prices around Quota Tariff Onsets in Event Time

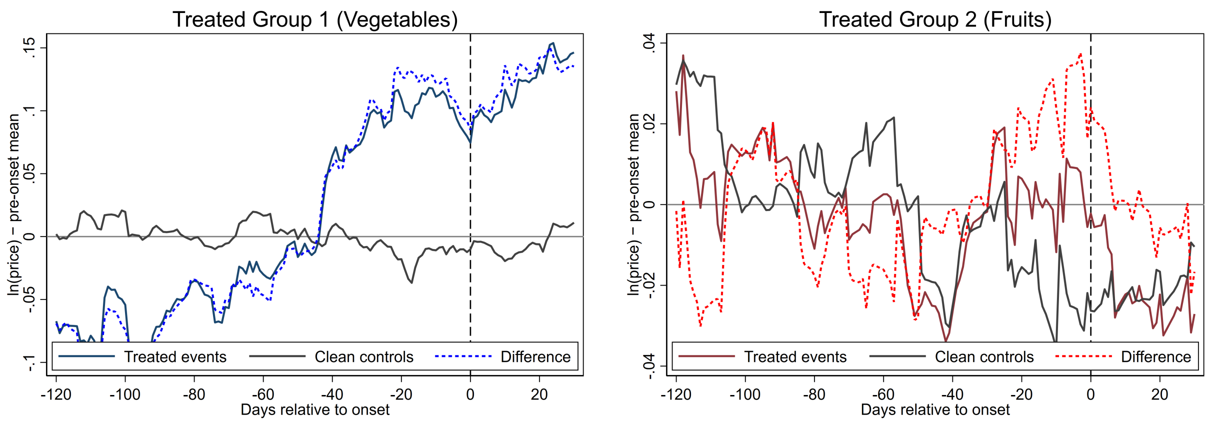

blocks (adopted on supply-shifter grounds (Roberts and Schlenker, 2013) rather than in response to this diagnostic), the same path collapses toward zero: 11 of 120 coefficients significant, none in the 66 days preceding onset, and near-onset magnitudes of 2∼5 log points. The coefficients themselves move toward zero rather than merely losing precision, and the gap between the two paths (the vegetable pre-onset coefficient paths estimated with and without the climate blocks) is statistically significant rather than sampling noise. Since both versions are estimated on the same sample, their coefficients can be differenced horizon by horizon; on a grid of horizons fixed in advance ($h = -120, -105, \ldots, -30$), the standard error of this difference is computed by a cluster jackknife that re-estimates both paths leaving out one product at a time. The resulting test rejects the hypothesis that the two paths coincide ($t = -2.0$ for the average difference across the grid, $t = -2.9$ at $h = -75$). For the fruit group, the two versions (with and without the climate blocks) are indistinguishable ($t = 0.00$), and both are flat, with 1∼2 of 120 coefficients significant. This confirms the validity of the decision to exclude the climate blocks from the covariates of the LP-DiD fruit specification: the treated fruits are predominantly imported, so Korean weather is not their supply shifter. The parallel-trends assumption stated in section 4.3 is *conditional on the covariates*. The graphical results in this subsection show that the systematic pre-designation differences between the vegetable group and the comparison group are absorbed by the covariates. With this diagnostic in hand, and subject to the two caveats of section 4.3, the analysis turns to the estimates.

Figure 4: Pre-Onset Coefficients with and without Climate Covariates

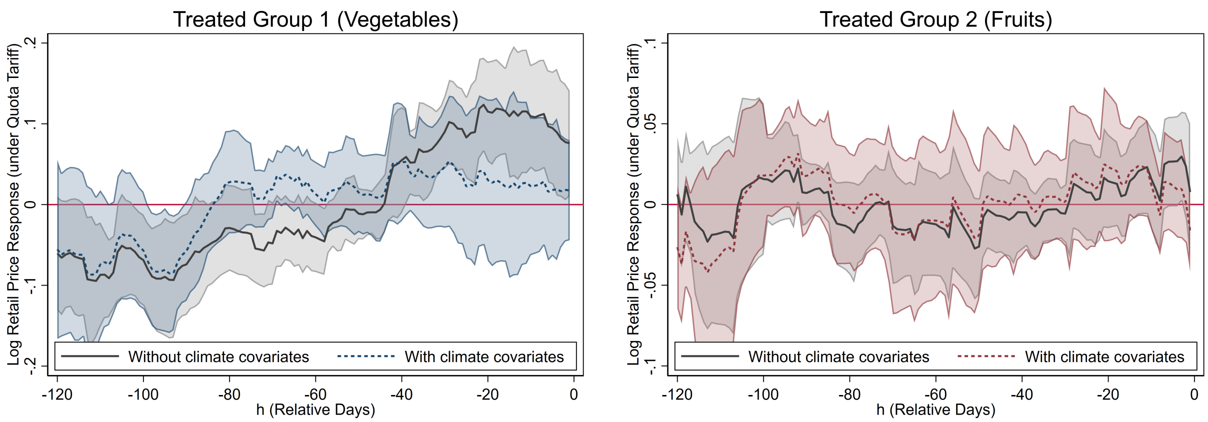

4.5 Graphical Results

Figures 5∼8 present the four sets of estimates: the treatment-status specification in its persistent-treatment and impulse-response forms, followed by the treatment-intensity specification in the same order. In all four figures, solid blue lines denote Treated Group 1 (vegetables) and dashed red lines Treated Group 2 (fruits), and the shading shows 95% confidence intervals (clustered by product). The treatment-intensity estimates are the main results of this study. This is because the onset of a Quota Tariff lowers the tariff rate by 1.5%p to 69.9%p depending on the product, and the intensity specification accounts for that depth of reduction instead of treating every designation as an identical shock (the approach of the treatment-status specification). The treatment-status estimates capture the effect of Quota Tariff adoption regardless of the depth of the cut.

Figure 5 shows the persistent-treatment estimates for treatment status. The pre-event window supports the identifying assumptions for both groups: for vegetables (Treated Group 1), 11 of the 120 pre-event coefficients are significant at the 5% level, modestly above the six expected by chance, but concentrated in a single nine-day negative cluster at $h \in [-99, -91]$ plus two isolated positive days at $h = -69$ and $h = -67$, with no significant coefficient in the 66 days preceding adoption; for fruits (Treated Group 2), only 5 of 120 are significant. That count is below the six expected by chance, and all five coefficients form a single short positive run at $h \in [-7, -3]$, of roughly $+3$ to $+4$ log points. Since the baseline equation (1) averages over 120 pre-event days, a bump of this size and duration leaves the reference level essentially

Figure 5: Treatment Status (Binary), Persistent-Treatment Specification

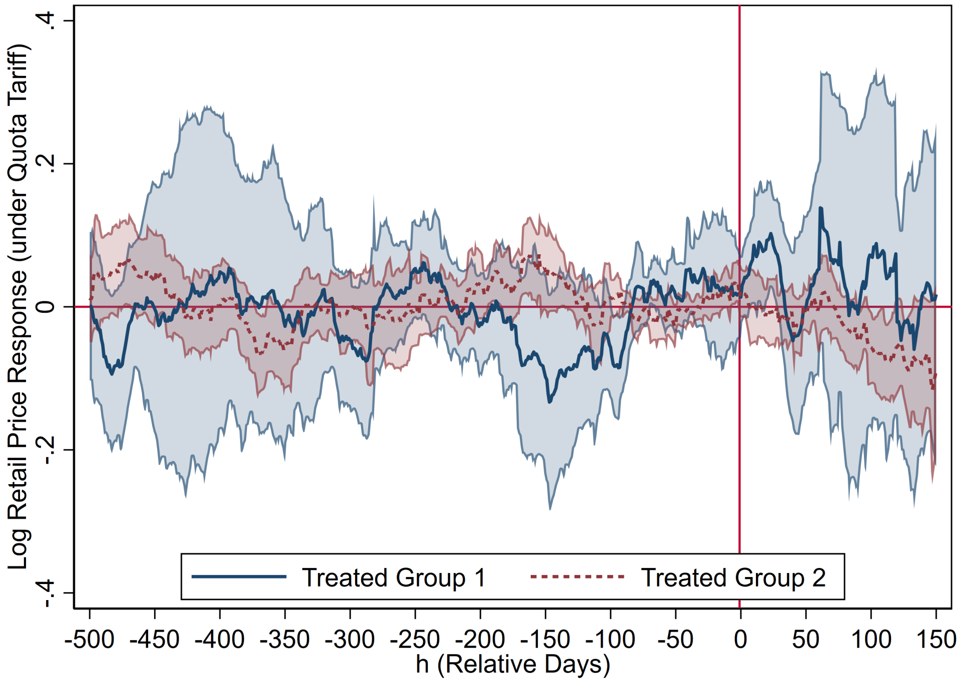

untouched.

Post-onset, the two groups of vegetables and fruits diverge. Vegetable prices do not fall. The point estimates are never significantly negative, and are significantly positive in the first weeks after adoption (at $h \in [8, 10]$, $h = 14$, and $h \in [20, 27]$, reaching about $+0.10$ log points at $h = 23$), then insignificant through $h = 150$. The early positive coefficients should not be read as the tariff cut raising prices: if the tariff cut were the cause, the coefficients would be larger for products with deeper cuts, yet they disappear entirely in the intensity specification below.

Fruit prices, in contrast, begin to decline about three months after the Quota Tariff adoption: significant negative pockets appear over $h \in [90, 131]$ (about -5.7 log points, or -5.5 percent, at $h = 100$), and in the underlying estimates (retained in the replication files beyond the plotted window) they recur through $h \approx 200$.[19] The treatment-status specification, which treats shallow-cut events such as kiwifruit (1.5%p) on the same footing as deep-cut events such as banana (about 25%p), makes it difficult to

[19]However, the estimate at $h = 150$ (-8.9 log points) is imprecise.

estimate the fruit response precisely, whereas the treatment-intensity specification reveals this response far more clearly.

Figure 6: Treatment Status (Binary), Impulse-Response Specification

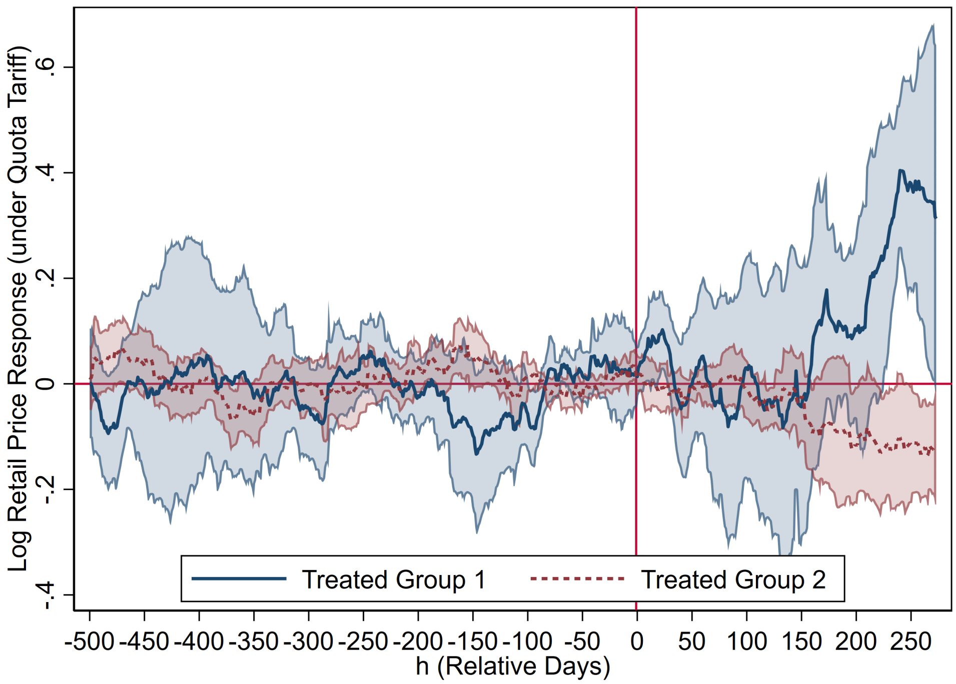

Figure 6 extends the results of the treatment-status specification to the impulse-response form, the total effect of initiation, and tracks them to $h = 273$ with no attrition of treated products. For fruits (Treated Group 2), the initiation effect turns significantly negative in pockets from $h = 118$ onward (about -4.5 log points over $h \in [118, 124]$) and settles into a sustained decline at far horizons: -11.6 log points at $h = 200$ and -11.5 at $h = 250$, remaining there through $h = 273$, a decline of roughly 11 percent. Since many fruit episodes have ended well before these horizons, the persistence of the decline indicates that the effect of Quota Tariff initiation outlives the period during which the treatment remains in force. For vegetables, apart from the early positive coefficients observed in the persistent-treatment specification,[20] the impulse-response results are insignificant through $h \approx 225$ and then turn sharply positive ($+0.38$ log points, or $+46\%$, at $h = 250$). This far-horizon vegetable estimate should not be inter-

[20]The two samples are mechanically identical at short horizons, before any spell has ended.

preted causally. It is driven by the 2024 onsets of the Quota Tariff policy for vegetables, whose 200∼270-day windows overlap with the severe late-2024 supply disruptions in the kimchi-vegetable markets (radish, napa cabbage). Given treatment intensities of roughly 27∼30%p, the tariff-reduction mechanism could not generate a coefficient of this size even at full pass-through. The estimate thus picks up a contemporaneous supply disturbance specific to these products, not an effect of the designations; the far-horizon vegetable coefficients are therefore reported for transparency but excluded from the causal conclusions.

Figure 7: Treatment Intensity (per 1%p), Persistent-Treatment Specification

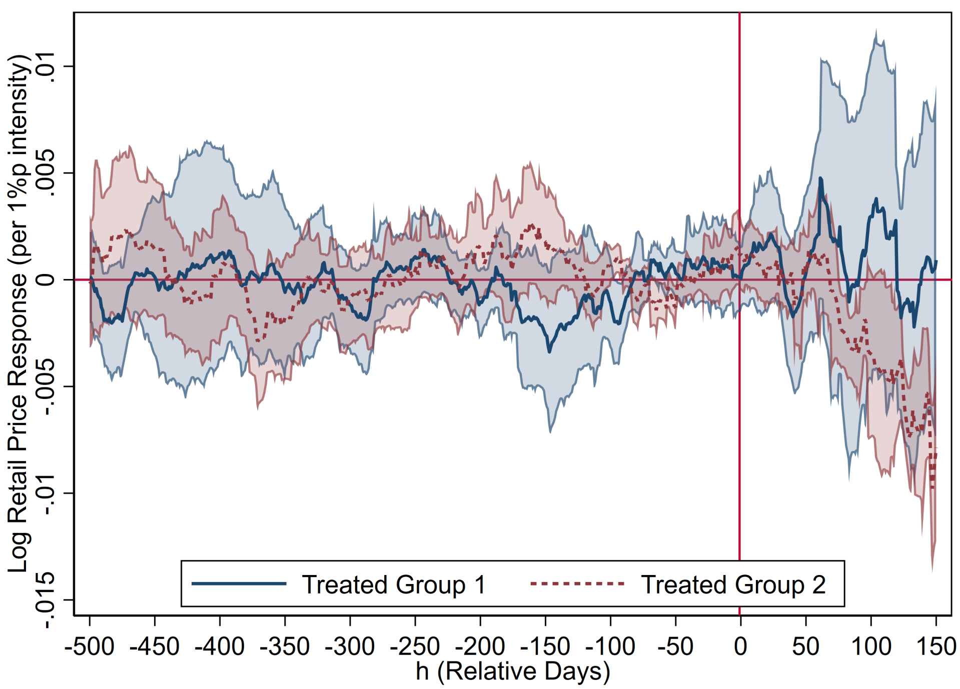

Figure 7 turns to the main specification of this study: the persistent-treatment response per 1%p of the tariff reduction induced by a Quota Tariff onset. For vegetables (Treated Group 1) the pre-event record is clean (5 of 120 significant, none within 67 days of onset), and no significant response appears anywhere in the post window through $h = 150$. Under ongoing treatment, vegetable prices simply do not fall per unit of tariff reduction. For fruits (Treated Group 2), 17 of 120 pre-event coefficients are significant, above the chance rate; this is the one pre-event diagnostic in the four figures that warrants scrutiny. The significant days fall in two clusters (positive around $h \in [-99, -92]$, negative around $h \in [-65, -50]$), both seven weeks or more before adoption, with no significant coefficient in the 49 days preceding the onset. Since adjacent horizons are serially correlated with one another, clusters of significant coefficients like these can arise. However, at the horizons adjacent to the Quota Tariff onset, there is no systematic pattern of significant coefficients. Post-onset, the response of the fruit treated group hovers near zero for the first ten weeks and is negative essentially throughout thereafter. The response is significant at $h \in [76, 96]$ (roughly -0.2 to -0.35 per 1%p, log points $\times 100$) and reaches -0.79 at the plotted window's edge, $h = 150$. Beyond the plotted window the fruit treated set is unchanged through $h = 270$ (no further event exits until $h = 271$), and in the replication files (which provide the estimates beyond $h = 150$) the significant band continues over $h \in [112, 209]$, reaching -1.5 per 1%p at $h = 182$. The graphs stop at $h = 150$ only because, as described earlier, the horizon is capped in response to the attrition of treated products caused by the composition effect in the vegetable treated group. The significant fall in the retail prices of the fruit treated group between 90 and 210 days after the Quota Tariff onset is the core causal finding of the paper. Fruit retail prices fall by economically meaningful magnitudes per 1%p of tariff reduction, and the sign of this response is robust across all four specifications.

Figure 8, the last of the four figures, shows the impulse responses of the treatment-intensity specification. For fruits (Treated Group 2), the estimates show the same pattern as the impulse-response results of the treatment-status specification: significant negative pockets appear from $h = 119$ onward (-0.23% per 1%p over $h \in [119, 124]$), and a sustained significant decline follows at far horizons, -0.49% per 1%p at $h = 250$ and -0.55% at $h = 273$. For vegetables (Treated Group 1), the estimates are insignificant until $h = 226$ and then jump abruptly positive at far horizons ($+1.35\%$ at $h = 250$); the caution on this result was already noted above for the impulse-response results of the treatment-status specification.

Two additional diagnostics complete the review of the four graphical results above. The first returns to the reverse Ashenfelter-dip concern examined in section 4.4. Three features of the evidence and design speak against it. Across all four specifications, the pre-event coefficients are flat in the window adjacent to adoption.[21] Equation (1)

[21]The sole exception is the five-day blip for fruits discussed above.

Figure 8: Treatment Intensity (per 1%p), Impulse-Response Specification

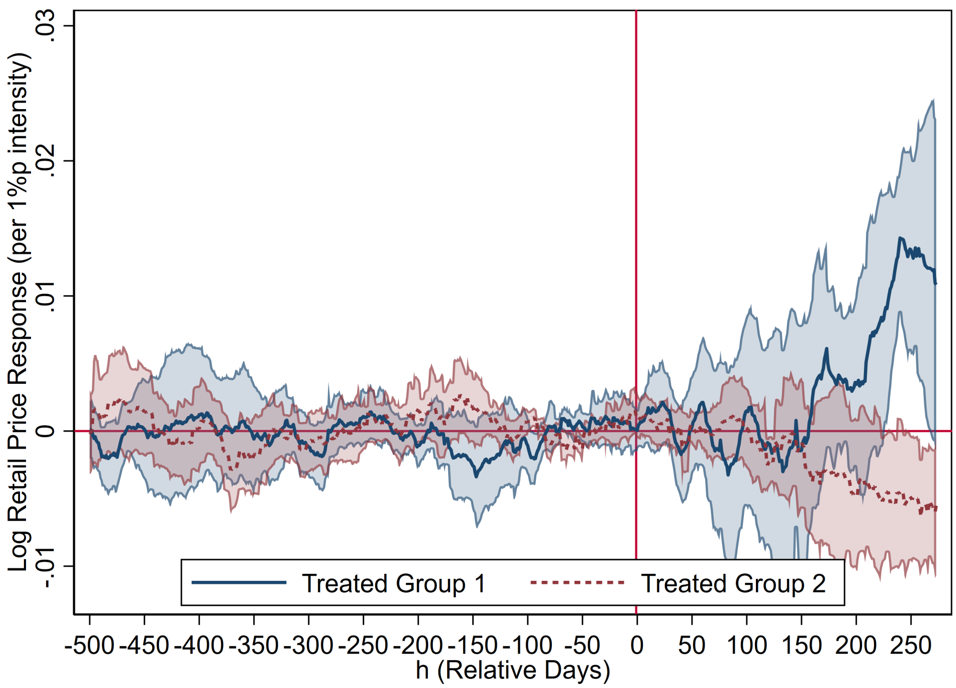

anchors the reference level to a 120-day average, so a surge in the last few days barely moves it, and the calendar-day fixed effects absorb market-wide surges. In addition, the regressions for the vegetable treated group use current and lagged climate as covariates, absorbing the supply shocks that drive both price surges and Quota Tariff designations (figure 4). A classic negative dip is likewise absent. The only case that cannot be fully excluded is that anticipation happens to offset a dip exactly, but such a coincidental offset is highly unlikely because, as described earlier, anticipation has little room to operate.

The second diagnostic is economic magnitude. Under full and immediate pass-through, a 1%p tariff reduction lowers the consumer price by approximately $100/(100+\bar{\tau})$ percent, where $\bar{\tau}$ is the prevailing tariff level; for the treated fruit products, the counterfactual applied tariff rate $\bar{\tau}$ ranges from about 6.5 to 29.6 percent (table 3), so the full-pass-through benchmark lies between roughly 0.77 and 0.94 percent per 1%p. The estimated fruit responses (-0.78% per 1%p at $h = 150$ under ongoing treatment, and -0.49% to -0.55% per 1%p at $h \geq 250$ for initiations) therefore imply substantial, at some horizons near-complete, pass-through of the tariff reduction to consumers.[22] The sign of the fruit response is robust across all four specifications; its magnitude varies with horizon and specification, and the estimates above bracket the plausible range rather than pin down a single pass-through rate. The figures trace the full dynamic response; section 4.6 then condenses it into a regression table that reports magnitudes, standard errors, and fit at selected horizons.

4.6 Local Projection Difference-in-Differences Estimates at Selected Horizons

Table 4 reports, in regression form, the four graphical results presented in section 4.5. The reported horizons h are chosen by the design rather than by the results. For the persistent-treatment specification, $h = 150$ is the truncation point of its graphs, set to leave a conservative margin below the compositional break in the treated products at $h = 156$, and is thus the longest horizon at which its results are reported. The impulse-response specification is reported at the same horizon for comparability with the persistent-treatment specification, and also at $h = 250$ to cover the long horizon. This horizon falls roughly eight months after Quota Tariff initiation, in the far-post window where the impulse-response treated sets are complete (nine vegetable and eight fruit events), and lies 23 days inside the sample's maximum horizon of $h = 273$. Coefficients and standard errors are multiplied by 100 for readability.

The table restates the findings of the four figures in compact form. At $h = 150$ in the persistent-treatment specification (panel A), the only significant estimate is the fruit response per 1%p of intensity, -0.79 (log points $\times 100$; panel A, column 4). The fruit estimate in the treatment-status specification is negative but insignificant (panel A, column 2), and both vegetable estimates (panel A, columns 1 and 3) are near zero. At the same horizon, the impulse-response estimates (panel B) are insignificant across the board. Contrasting this result with panel C reveals how long the effect of Quota Tariff initiation takes to materialize. By $h = 250$, the initiation effects emerge as significant fruit declines (-11.5 binary, panel C, column 2; -0.50 per 1%p, panel C, column 4). As

[22] Computed as $(e^{\beta^h} - 1) \times 100$: -0.783% at $h = 150$ (persistent treatment), -0.494% at $h = 250$ and -0.548% at $h = 273$ (impulse). The benchmark treats the quota as non-binding within the window and abstracts from inventory and contracting lags, so it is an upper bound on the speed, not the existence, of pass-through.

Table 4: Local Projection Difference-in-Differences Estimates at Selected Horizons

	(1)	(2)	(3)	(4)
	Quota Tariff Treatment Status		Quota Tariff Treatment Intensity (1%p)	
	Group 1 (Vegetables)	Group 2 (Fruits)	Group 1 (Vegetables)	Group 2 (Fruits)
Panel A: Persistent treatment, $h = 150$				
Quota Tariff Shock	1.740	-8.950	0.092	-0.786***
	(12.260)	(6.166)	(0.420)	(0.224)
Observations	33678	45026	33678	45026
R^2	0.492	0.159	0.492	0.159
Adjusted R^2	0.471	0.131	0.471	0.131
Panel B: Initiation (impulse), $h = 150$				
Quota Tariff Shock	-4.360	-2.338	-0.143	-0.088
	(13.521)	(4.400)	(0.476)	(0.247)
Observations	33682	45030	33682	45030
R^2	0.492	0.159	0.492	0.159
Adjusted R^2	0.470	0.131	0.470	0.131
Panel C: Initiation (impulse), $h = 250$				
Quota Tariff Shock	37.966***	-11.519**	1.340***	-0.495**
	(10.634)	(5.653)	(0.397)	(0.237)
Observations	30154	41530	30154	41530
R^2	0.593	0.191	0.593	0.191
Adjusted R^2	0.576	0.164	0.576	0.164

Coefficients and standard errors multiplied by 100 for readability; standard errors, in parentheses, clustered by product (36 clusters in columns 1 and 3, 35 clusters in columns 2 and 4; 6 treated vegetable products, 5 treated fruit products).
Significance is based on the normal critical values used for the confidence bands in the figures.
Panel A conditions on treatment persisting through h; Panels B and C track initiations regardless of the subsequent path.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

detailed above, the far-horizon vegetable positives in the same panel (panel C, columns 1 and 3) reflect the attrition of treated products caused by the composition effect, as well as contemporaneous supply disruptions, rather than causal effects. The explanatory power differs mechanically across groups.[23] The modest R^2 values themselves have no bearing on causal identification.[24]

Taken together, the causal analysis yields one robust policy-relevant effect and one clearly labeled non-effect. Quota Tariff designations on imported fruits lower retail prices per 1%p of tariff reduction, significantly and persistently, with the decline outliving the designations themselves; designations on domestic leafy and root vegetables, by contrast, show no detectable price-reducing effect at any interpretable horizon, under either estimand. The next section probes the robustness of these conclusions.

[23] The vegetable regressions, with their product-specific climate blocks, absorb roughly half of the variance; the leaner fruit specification absorbs less.

[24] The modest R^2 values presumably arise because each dependent variable is a long-horizon daily deviation from a 120-day pre-event mean.

5 Robustness

5.1 Inference with Few Treated Clusters

The number of treated clusters is set by the policy itself: over the sample period the Quota Tariff reached six vegetable and five fruit products, so any estimator applied to this setting must conduct inference with few treated clusters, whatever its other merits. The econometric difficulty is well documented. With few treated clusters, cluster-robust t-statistics tend to over-reject, because the treated clusters contribute too few residual blocks for their variance to be estimated reliably (Conley and Taber, 2011; MacKinnon and Webb, 2020). Corrections push in the opposite direction and overshoot: the restricted wild cluster bootstrap is known to under-reject in exactly this configuration (even its most recent variants still tend to under-reject severely when the number of treated clusters is small (MacKinnon et al., 2023)). Therefore its wide intervals mainly reflect the documented failure of the procedure, not sampling uncertainty in the estimates. MacKinnon and Webb (2018) conclude that when the restricted and unrestricted bootstrap p-values diverge, as they systematically do with few treated clusters, neither can be trusted.

A literature has developed alternatives, but each restores validity by imposing a homogeneity assumption that the present setting violates by construction. Roth et al. (2023) summarize that "all of the available approaches require the researcher to impose some potentially strong additional assumptions." Randomization and placebo procedures (Conley and Taber, 2011; Hagemann, 2019; MacKinnon and Webb, 2020) infer the distribution of the treated clusters' shocks from the control clusters, which is valid only when treated and control clusters are exchangeable. MacKinnon and Webb (2018) show that such procedures fail when error variances differ across clusters in unknown form. Thus exchangeability is untenable here. Quota Tariff designation itself selects products prone to price surges, and daily retail price volatility differs sharply between the domestic vegetables and the imported fruits. Ferman and Pinto (2019) relax exchangeability, but their correction requires that the source of the heteroskedasticity be known: if each cluster's error variance is a known parametric function of an observable,[25] the control residuals can be re-scaled by that function to stand in for the treated clusters'

[25] In their leading case, the number of underlying observations averaged into each group×time cell.

shocks. The present panel offers no such observable: it is balanced, each product-day observation is a single price rather than a cell average, and the volatility differences across products therefore follow no known parametric structure that such a re-scaling could exploit.

Canay et al. (2021) derive conditions under which the wild bootstrap remains valid with a small number of clusters, but these require homogeneity of the covariate design across clusters and are, by the authors' own account, difficult to satisfy when treatment is assigned at the cluster level. These frameworks are, moreover, all developed for a binary treatment that turns on once and stays on; even the recent extensions to staggered timing and heterogeneous effects (Alvarez and Ferman, 2023) maintain a binary absorbing treatment and a modeled variance structure. A non-absorbing, continuous-intensity design of the kind the Quota Tariff requires falls outside the scope of every one of them.

The analysis therefore reports conventional product-clustered intervals and disciplines their reading in three ways. First, the documented distortion of cluster-robust t-statistics with few treated clusters is over-rejection, a bias toward finding effects. The conclusion for vegetables is reached against this bias and is therefore all the more robust: the null of no effect was not rejected even by tests strongly predisposed to reject it. The far-post vegetable positives (the one place where a liberal test could mislead) are already excluded from the causal conclusions on composition and confounding grounds (section 4.5). Second, the fruit effect does not rest on an isolated significant coefficient of the kind small-sample noise produces. It appears as contiguous bands of horizons, with the same sign and timing across all four specifications, at magnitudes that lie below what a tariff pass-through rate of 100% would imply and are therefore reasonable. Third, following the reporting guidance of MacKinnon et al. (2023), the cluster structure is disclosed in full for each horizon h in Appendix D (the number of clusters, the treated-cluster counts, and the event-level composition of the treated sets).

6 Conclusion and Policy Implications

Korea's Quota Tariff is an institution distinct from the TRQ of the WTO and FTA frameworks. The TRQ guarantees the market access committed under the WTO Agreement on Agriculture and FTAs, applying a low tariff rate up to a specified volume and a higher rate to imports beyond it. The Quota Tariff, by contrast, is a flexible tariff instrument operated under *Article 71 of the Customs Act*: the tariff rate and the covered volume can be adjusted quickly by Presidential Decree, and temporary tariff reductions without any quantity limit are possible. This paper has examined whether the Quota Tariff designations aimed at stabilizing agricultural prices causally lowered retail prices as the government intended.

Prior studies have mainly analyzed the trade and price effects of TRQs (Loginova et al., 2021; Schaefer and Wolf, 2025), and the study of Korea's Quota Tariff by Song (2023) likewise analyzed the relationship between the import price excluding tariffs and consumer prices; the causal effect on consumer prices of a Quota Tariff designation itself had therefore never been identified. Meanwhile, Quota Tariff coverage of agricultural products has expanded sharply in recent years, entailing substantial fiscal losses, and much of it has aimed at price stabilization (Jang, 2025), so an assessment of whether the policy attains its goal is urgent. To anticipate the results: no retail price reduction was found for the vegetables, composed of leafy and root vegetables, whereas a significant and persistent reduction was found for the fruits.

The analysis used a daily panel of retail prices for 41 agricultural products from January 1, 2021 through March 31, 2025. The treated and control groups were selected by objective criteria alone, with no researcher discretion involved. After excluding items that are not agricultural products or whose daily retail prices are missing, potato, for which the Quota Tariff and the retail price refer to different goods (chip-manufacturing versus table potatoes), and grape, whose retail prices are unobserved over its event window, the eleven products that had the Quota Tariff in force at least once during the sample period (six leafy and root vegetables and five fruits) were classified as ever-treated and the remaining 30 products as never-treated. An ever-treated product functions as treated at some horizons and as a control before its first onset, whereas a never-treated product functions only as a control. Treatment was defined not by the designated period announced in advance but by the days on which the tariff reduction was actually in force. The treatment status so constructed is non-absorbing. A structure in which treatment, once switched on, remains on until the end of the sample is called absorbing (Bach et al., 2025), and one in which treatment can switch off and on again is called non-absorbing (Dube et al., 2025); the Quota Tariff is of the latter kind because a designation lapses when its announced end date arrives or its quota is exhausted, and the same product can be designated again later.

Quota Tariff designations begin on different dates across products (staggered adoption), the treatment is suspended and later resumed (non-absorbing), and the depth of the tariff reduction differs widely across products (heterogeneous treatment intensity); this structure effectively dictates the choice of estimator. Conventional TWFE event-study regressions implicitly use previously treated observations as comparisons, so their estimates can be severely biased under this structure (Goodman-Bacon, 2021; De Chaisemartin and D'Haultfœuille, 2020); the recent heterogeneity-robust estimators presuppose an absorbing structure in which treatment, once switched on, stays on, and are therefore inapplicable (Sun and Abraham, 2021; Callaway and Sant'Anna, 2021; Borusyak et al., 2024; Gardner, 2022). LP-DiD is the method that accommodates non-absorbing treatment, admits covariates directly in the regression, and estimates the treatment-status (binary) and treatment-intensity (continuous) specifications within the same framework (Dube et al., 2025). LP-DiD compares only clean treatment events, untreated throughout the 120 days before onset, with clean controls that have never been treated up to the comparison window (Cengiz et al., 2019; Dube et al., 2025). The intensity specification of this paper implements the continuous, non-absorbing LP-DiD model proposed by Dube et al. (2025).

There are two estimands. The persistent-treatment specification imposes the stay-on condition, which retains only events treated on every day of $[t, t+h]$, and thus measures the price effect only while the Quota Tariff keeps the tariff reduction in force. The impulse-response specification removes exactly that stay-on condition and leaves the post-onset treatment path as it actually unfolded, measuring the total effect of a Quota Tariff designation, that is, the IRF (Dube et al., 2025; De Chaisemartin and D'Haultfœuille, 2026). The two estimands are complementary: the former speaks to whether the policy lowers prices while it operates, and the latter to whether the effect of Quota Tariff initiation outlives the very period during which the tariff reduction remains in force.

The treatment-intensity specification, which reflects the differences in the depth of the tariff reduction across products, is the main specification of this paper, and its results are summarized by a pronounced heterogeneity between vegetables and fruits. In the persistent-treatment specification (figure 7), the estimated price declines for vegetables (Treated Group 1) were insignificant at every horizon through $h = 150$, the end of the plotted window. Fruits (Treated Group 2) showed a significant decline from $h \in [76, 96]$ onward, reaching about -0.78% per 1%p of tariff reduction at $h = 150$. In the impulse-response specification (figure 8), the initiation effect for fruits continued to fall at far horizons, reaching -0.49% per 1%p at $h = 250$ and -0.55% at $h = 273$; because many fruit episodes had lapsed well before these horizons, the price-reducing effect persists longer than the designations themselves. The estimated fruit effects lie inside the benchmark of 0.77∼0.94% per 1%p expected under full and immediate pass-through, implying that the tariff reduction was passed through to consumers to a substantial degree, and at some horizons almost completely.

The policy implications are clear. First, Quota Tariff designations on leafy and root vegetables did not achieve their original goal of lowering retail prices. Since the Quota Tariff entails a direct fiscal cost (the forgone tariff revenue), simply refraining from designations on vegetables, for which no price-reducing effect is identified, would substantially reduce policy costs. Second, Quota Tariff designations on fruits produced significant and persistent declines in retail prices; as a price-stabilization instrument, the Quota Tariff is therefore effective when deployed with a focus on products with a high import share, such as fruits. Consistent with this contrast, the treated fruits are predominantly imported, so the channel from a tariff reduction to retail prices is direct, whereas for the mostly domestically produced vegetables that scope is structurally limited. Moreover, the price-reducing effect for fruits emerges with a lag of several months after onset and persists after the designation ends, so the instrument should be operated and evaluated as a medium-term price-stabilization tool. In sum, the agricultural Quota Tariff is a selective instrument whose success depends on a product's supply structure: concentrating it on fruits and reconsidering designations on vegetables would deliver a larger price-stabilization effect at the same fiscal cost.

References

Alvarez, L., Ferman, B., 2023. Extensions for Inference in Difference-in-Differences with Few Treated Clusters. doi:10.48550/arXiv.2302.03131, arXiv:2302.03131.

Ashenfelter, O., 1978. Estimating the Effect of Training Programs on Earnings. *The Review of Economics and Statistics* 60, 47. doi:10.2307/1924332, arXiv:1924332.

Bach, P., Klaassen, S., Kueck, J., Mattes, M., Spindler, M., 2025. Sensitivity analysis for treatment effects in Difference-in-Differences models using Riesz representation. *arXiv preprint* .

Bandara, K., Hyndman, R.J., Bergmeir, C., 2025. MSTL: A seasonal-trend decomposition algorithm for time series with multiple seasonal patterns. *International Journal of Operational Research* 52, 79–98. doi:10.1504/IJOR.2025.143957.

Bellemare, M.F., Wichman, C.J., 2020. Elasticities and the Inverse Hyperbolic Sine Transformation. *Oxford Bulletin of Economics and Statistics* 82, 50–61. doi:10.1111/obes.12325.

Borusyak, K., Jaravel, X., Spiess, J., 2024. Revisiting Event-Study Designs: Robust and Efficient Estimation. *Review of Economic Studies* 91, 3253–3285. doi:10.1093/restud/rdae007.

Busch, A., Girardi, D., 2024. Lpdid: Stata module implementing local projections difference-in-differences (lp-did) estimator. *Statistical Software Components* .

Callaway, B., Sant'Anna, P.H., 2021. Difference-in-Differences with multiple time periods. *Journal of Econometrics* 225, 200–230. doi:10.1016/j.jeconom.2020.12.001.

Canay, I.A., Santos, A., Shaikh, A.M., 2021. The Wild Bootstrap with a "Small" Number of "Large" Clusters. *The Review of Economics and Statistics* 103, 346–363. doi:10.1162/resta00887.

Cengiz, D., Dube, A., Lindner, A., Zipperer, B., 2019. The Effect of Minimum Wages on Low-Wage Jobs*. *The Quarterly Journal of Economics* 134, 1405–1454. doi:10.1093/qje/qjz014.

Cleveland, R.B., Cleveland, W.S., McRae, J.E., Terpenning, I., 1990. STL: A seasonal-trend decomposition. *Journal of Official Statistics* 6, 3–73.

Conley, T.G., Taber, C.R., 2011. Inference with "Difference in Differences" with a Small Number of Policy Changes. *Review of Economics and Statistics* 93, 113–125. doi:10.1162/RESTa00049.

De Chaisemartin, C., D'Haultfœuille, X., 2020. Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects. *American Economic Review* 110, 2964–2996. doi:10.1257/aer.20181169.

de Chaisemartin, C., D'Haultfœuille, X., 2023. Two-way fixed effects and differences-in-differences with heterogeneous treatment effects: A survey. *The Econometrics Journal* 26, C1–C30. doi:10.1093/ectj/utac017.

De Chaisemartin, C., D'Haultfœuille, X., 2026. Difference-in-Differences Estimators of Intertemporal Treatment Effects. *Review of Economics and Statistics* 108, 863–880. doi:10.1162/resta01414.

Dube, A., Girardi, D., Jordà, Ò., Taylor, A.M., 2025. A Local Projections Approach to Difference-in-Differences. *Journal of Applied Econometrics* 40, 741–758. doi:10.1002/jae.70000.

Dube, A., Lindner, A., 2024. Minimum wages in the 21st century. *Handbook of Labor Economics* 5, 261–383. doi:10.1016/bs.heslab.2024.11.004.

Fajgelbaum, P., Khandelwal, A., 2026. Tariffs in 2025: Short-Run Impacts on the U.S. Economy. *SSRN Electronic Journal* doi:10.2139/ssrn.6566859.

Ferman, B., Pinto, C., 2019. Inference in Differences-in-Differences with Few Treated Groups and Heteroskedasticity. *The Review of Economics and Statistics* 101, 452–467. doi:10.1162/resta 00759.

Flaaen, A., Hortaçsu, A., Tintelnot, F., 2020. The Production Relocation and Price Effects of US Trade Policy: The Case of Washing Machines. *American Economic Review* 110, 2103–2127. doi:10.1257/aer.20190611.

Gardner, J., 2022. Two-stage Differences in Differences. *arXiv preprint* .

Goodman-Bacon, A., 2021. Difference-in-differences with variation in treatment timing. *Journal of Econometrics* 225, 254–277. doi:10.1016/j.jeconom.2021.03.014.

Hagemann, A., 2019. Placebo inference on treatment effects when the number of clusters is small. *Journal of Econometrics* 213, 190–209. doi:10.1016/j.jeconom.2019.04.011.

Hillen, J., 2019. Market Integration and Market Efficiency under Seasonal Tariff Rate Quotas. *Journal of Agricultural Economics* 70, 859–873. doi:10.1111/1477-9552.12355.

Jang, S.h., 2025. Tariff-Rate Quota Implementation: Current Status and Policy Improvement Agenda (할당관세 운용 현황과 개선과제). *National Assembly Budget Office (국회예산정책처)*.

Livingston, A., Schaefer, K.A., Malone, T., 2025. Market Access Under a Dynamic Safeguard Mechanism: Evidence From the United States–Japan Trade Agreement. *Agricultural Economics* 56, 895–904. doi:10.1111/agec.70039.

Loginova, D., Portmann, M., Huber, M., 2021. Assessing the Effects of Seasonal Tariff-rate Quotas on Vegetable Prices in Switzerland. *Journal of Agricultural Economics* 72, 607–627. doi:10.1111/1477-9552.12424.

MacKinnon, J.G., Nielsen, M.Ø., Webb, M.D., 2023. Cluster-robust inference: A guide to empirical practice. *Journal of Econometrics* 232, 272–299. doi:10.1016/j.jeconom.2022.04.001.

MacKinnon, J.G., Webb, M.D., 2018. The wild bootstrap for few (treated) clusters. *The Econometrics Journal* 21, 114–135. doi:10.1111/ectj.12107.

MacKinnon, J.G., Webb, M.D., 2020. Randomization inference for difference-in-differences with few treated clusters. *Journal of Econometrics* 218, 435–450. doi:10.1016/j.jeconom.2020.04.024.

Mollins, J., Lumb, R., 2024. Seasonal adjustment of weekly data. *Bank of Canada Staff Discussion Paper, No. 2024-17* .

Pierce, D.A., Grupe, M.R., Cleveland, W.P., 1984. Seasonal Adjustment of the Weekly Monetary Aggregates: A Model-Based Approach. *Journal of Business & Economic Statistics* 2, 260–270. doi:10.1080/07350015.1984.10509393.

Roberts, M.J., Schlenker, W., 2013. Identifying Supply and Demand Elasticities of Agricultural Commodities: Implications for the US Ethanol Mandate. *American Economic Review* 103, 2265–2295. doi:10.1257/aer.103.6.2265.

Roth, J., Sant'Anna, P.H., Bilinski, A., Poe, J., 2023. What's trending in difference-in-differences? A synthesis of the recent econometrics literature. *Journal of Econometrics* 235, 2218–2244. doi:10.1016/j.jeconom.2023.03.008.

Schaefer, K.A., Wolf, C.A., 2025. Trade protection via tariff rate quota administration. *Food Policy* 131, 102801. doi:10.1016/j.foodpol.2025.102801.

Song, Y.g., 2023. 2022 Tariff-Rate Quota implementation results and impact analysis (2022년도 할당관세 지원실적 및 효과분석). *Ministry of Economy and Finance (기획재정부)*.

Sun, L., Abraham, S., 2021. Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics* 225, 175–199. doi:10.1016/j.jeconom.2020.09.006.

Wing, C., Freedman, S.M., Hollingsworth, A., 2024. Stacked Difference-in-Differences. Working Paper 32054. National Bureau of Economic Research.

A Appendix: Tariff Rate Priority Structure

To determine the final applicable rate among multiple tariff rates that may apply to a single HSK code, the priority order of rate application was examined in accordance with *Article 50 of the Customs Act*. The priority order ranges from first to seventh, with the application method as follows.

The rates applied with the highest priority are Special Tariff Rates designed to protect domestic industries or correct trade imbalances, including Anti-Dumping Duty, Retaliatory Duty, and Countervailing Duty. The second-priority rate is the FTA rate under the *Act on Special Cases of the Customs Act for the Implementation of Free Trade Agreements*. If this rate is equal to or lower than rates under the *Customs Act* (excluding first-priority rates), the FTA rate is applied preferentially upon the importer's request. Third-priority rates include International Cooperation Tariffs, such as WTO General Tariff Concessions, and Beneficial Tariffs. Among third-priority rates, Tariff Concessions for agricultural, forestry, and livestock products (conceded at rates equivalent to the domestic-foreign price differential or at rates exceeding Basic Tariff Rates in conjunction with domestic market opening under the *Regulations on Tariff Concessions in the Framework of the Agreement Establishing the World Trade Organization*) take precedence over Basic Tariff Rates and Provisional Tariff Rates. Fourth-priority rates include Adjusted Duty[26], Quota Tariff, and Seasonal Tariffs. Quota Tariff takes precedence when it is lower than Preferential Tariffs for least developed countries, and it also supersedes Basic Tariff Rates and Provisional Tariff Rates. The fifth-priority rate is Preferential Tariffs for least developed countries. Provisional Tariff Rates and Basic Tariff Rates correspond to the sixth and seventh priority rates, respectively. However, due to data availability limitations, Specific Duties were not considered in this study.

B Appendix: Removing Seasonality

This paper removes seasonality via a regression-based adjustment employing calendar dummy variables and Fourier terms (Pierce et al., 1984). For daily data, we include day-of-week dummies, model annual seasonality using Fourier (sine/cosine) terms, and control for major and moving holidays with dedicated dummies; the seasonally

[26] *Article 69, Subparagraphs 1, 3, and 4 of the Customs Act*

adjusted series is obtained by subtracting the estimated seasonal components from the raw data.

Two considerations motivated this choice. The first is the need to keep policy effects out of the seasonal component: Quota Tariff designations often coincide with periods of seasonal price fluctuation, so the adjustment must distinguish price movements attributable to the policy from recurring seasonal patterns. A natural candidate for daily deseasonalization is STL, a non-parametric method that extracts seasonality at any specified period and has been applied extensively to weekly and daily data (Cleveland et al., 1990).[27] Since STL relies on non-parametric smoothing, however, trends or policy shocks in the data may be partially absorbed during the decomposition, and treatment effects removed at the seasonal-adjustment stage cannot be recovered in the subsequent DiD analysis.[28]

The regression-based adjustment eliminates this risk. The calendar dummies and Fourier terms enter as deterministic regressors, independent of the treatment variable, so the seasonal component is estimated separately rather than confounded with trends or policy effects. No seasonal pattern remains in the adjusted series, while trend changes and the effects of Quota Tariff designations are fully preserved.

The second consideration is the structure of the raw series itself. The daily data exhibit strong, multi-layered seasonality: week-of-year and day-of-week patterns, holiday effects, and an annual cycle that is not even fixed in length, some years comprising 52 weeks and others 53. Standard official seasonal-adjustment tools such as X-12/X-13-ARIMA-SEATS are not directly applicable to such data, because they are designed for monthly or quarterly series and assume fixed 12-month or 4-quarter cycles (Mollins and Lumb, 2024).

The regression-based approach is established practice for exactly this situation: the U.S. Bureau of Labor Statistics (BLS), when seasonally adjusting weekly unemployment insurance claims, developed and has employed `MoveReg`, a moving weighted

[27] Multiple STL (MSTL) extends the method to decompose several seasonal factors simultaneously, such as 7-day and annual cycles (Bandara et al., 2025).

[28] STL-based methods also assume fixed seasonal cycles, so they cannot fully accommodate the week-53 problem or cycle variation attributable to leap years: imposing an annual 52-week cycle presumes that exactly the same cycle repeats every year, which can leave subtle seasonal residuals in year-end weeks. Nor do they offer adjustments for moving holidays or trading-day effects, so variation from shifting holiday dates is not adequately removed (Mollins and Lumb, 2024).

regression based on weekly dummies and Fourier terms, instead of traditional X-13 techniques (Mollins and Lumb, 2024).

C Appendix: Summary Statistics

This appendix summarizes the estimation sample and variables in three tables. Table 5 reports descriptive statistics for the variables entering the LP-DiD regressions, while Table 6 and Table 7 report the annual import scale of the 41 products. Table 5 is organized in three panels. Panel A covers the daily panel variables: the seasonally adjusted retail price from which the dependent variable is built, the daily Quota Tariff treatment status, and the covariates for the applied tariff rate, climate, oil prices, and WTO Market Access Quotas. For the two variables in panel B, summary statistics are computed not at the daily level but at the unit at which each variable enters the LP-DiD regressions. Treatment intensity G, the tariff reduction induced by a designation, varies across ever-treated products but not over time,[29] and its summary statistics are therefore computed at the product level. Government stock releases are recorded as monthly totals in the raw data, so their summary statistics are computed at the product–month level. Panel C reports the seasonally adjusted retail price of each of the 41 products in four blocks: treated leafy and root vegetables, treated fruits, never-treated vegetables and other crops, and never-treated fruits. The reported statistics (N, mean, standard deviation, minimum, and maximum) are the standard set in the empirical pass-through literature.

All entries in Table 5 are level values in their raw units, prior to any transformation for use in the LP-DiD regressions. In the estimation itself, the dependent variable is the log of the seasonally adjusted retail price differenced against its average over the 120 days preceding the event, as defined in equation (1), and the covariates enter in the forms defined in sections 3.6 and 4.3. Reporting raw levels keeps the scale and range of each variable visible rather than hidden behind those transformations.

[29]Treatment intensity G—the difference between the tariff rate that would apply in the absence of the designation (the applied tariff rate) and the rate under the designation—is indeed time-varying, since the applied tariff rate changes monthly. However, by construction the LP-DiD model requires a time-invariant treatment variable, and because G in this study is defined as the tariff reduction at the time the designation takes effect, we hold G fixed at that value.

Table 5: Summary Statistics of Regression Variables

	N	Mean	SD	Min	Max
Panel A. Daily panel variables					
Retail price, seasonally adjusted (KRW/kg)	63,013	9,435.9	7,494.7	753.5	36,797.0
Quota Tariff treatment status (0/1)	63,591	0.062	0.240	0.000	1.000
Applied tariff rate (%)	63,591	35.5	42.5	0.0	587.9
Mean temperature (°C)	63,550	13.3	10.8	-14.9	31.8
Mean relative humidity (%)	63,591	65.2	13.8	22.9	99.3
Precipitation index	63,345	2.2	0.9	0.1	3.8
Sunshine duration (hours)	63,345	6.2	4.0	0.0	13.5
Real crude oil price ($\times$1,000)	63,591	20.7	3.9	14.2	34.9
WTO market access quota, applicable (0/1)	63,591	0.341	0.474	0.000	1.000
WTO quota expansion (ratio)	63,591	0.276	1.420	0.000	9.548
WTO expansion intensity	63,591	1.441	8.258	0.000	56.334
Panel B. Treatment intensity and policy variables					
Treatment intensity G (%p; ever-treated items)	11	26.2	17.3	1.5	69.9
Government stock releases (tons/month; item$\times$month)	2,091	70.0	388.9	0.0	5,214.0
Panel C. Seasonally adjusted retail price (KRW/kg), by product					
Treated: leafy and root vegetables					
cabbage	1,548	2,180.8	331.4	1,504.5	3,061.4
carrot	1,548	4,060.7	820.0	2,703.9	6,072.3
green onion	1,548	3,130.0	825.9	1,711.7	6,361.5
napa cabbage	1,548	2,253.6	404.8	1,496.8	3,758.7
onion	1,548	2,146.6	363.4	1,505.0	3,088.1
radish	1,548	1,222.6	250.4	753.5	1,974.4
Treated: fruits					
avocado	1,093	8,359.5	988.9	6,385.9	11,981.9
banana	1,548	2,851.3	300.3	1,956.4	3,346.3
kiwifruit	1,548	4,714.2	323.4	3,877.7	5,558.7
mango	1,548	13,505.1	2,057.3	8,666.4	19,104.1
pineapple	1,548	4,625.3	342.0	3,663.3	5,500.4
Never-treated: vegetables and other crops					
adzuki bean	1,548	16,149.5	2,725.4	12,131.5	24,999.6
bean	1,548	9,937.2	1,107.9	7,743.0	12,224.3
bell pepper	1,548	8,176.9	781.8	6,014.7	11,121.2
cherry tomato	1,548	8,345.8	836.9	6,072.9	10,715.3
cucumber	1,548	4,071.2	555.6	2,817.7	5,941.3
dried red pepper	1,548	26,393.4	3,793.7	20,620.9	36,797.0
garlic	1,548	10,297.2	1,983.8	6,741.8	13,866.9
ginger	1,548	11,304.9	2,363.3	7,155.3	16,390.4
green pepper	1,548	11,641.5	1,826.2	7,999.8	18,859.5
lettuce	1,548	11,134.9	1,917.4	7,419.2	18,435.2
mung bean	1,548	20,505.6	1,211.9	18,071.3	23,180.6
peanut	1,548	27,732.3	2,102.0	24,094.2	33,784.4
perilla leaf	1,548	23,060.0	2,337.6	17,495.5	29,706.7

continued on next page

Table 5 (continued)

	N	Mean	SD	Min	Max
red pepper	1,548	18,386.9	2,860.9	11,551.6	29,229.2
scallion	1,548	7,672.0	1,171.0	5,059.3	10,891.7
sesame	1,548	31,001.2	2,340.7	26,306.4	36,493.6
soybean sprout	1,548	4,572.9	487.7	3,136.4	5,585.0
spinach	1,548	9,404.9	1,544.8	5,566.8	15,590.7
squash	1,548	5,410.5	892.2	3,612.6	9,152.9
sweet pepper	1,548	12,274.1	1,731.0	8,625.9	18,976.1
sweet potato	1,548	5,159.0	716.9	3,927.9	6,941.1
tomato	1,548	5,516.2	682.9	3,919.8	7,626.1
water dropwort	1,548	11,250.7	3,194.3	6,435.8	17,156.8
young napa cabbage	1,548	2,941.1	405.6	2,063.8	4,214.5
young summer radish	1,548	3,220.3	726.4	1,956.0	6,398.6
Never-treated: fruits					
apple	1,548	8,779.5	1,114.7	6,599.0	12,383.1
lemon	1,548	6,747.0	457.0	5,584.8	7,976.3
melon	1,548	6,352.3	584.9	5,112.6	7,785.4
pear	1,548	5,875.6	1,180.7	3,200.4	9,583.2
watermelon	1,548	4,192.1	351.5	3,211.7	4,899.9

The estimation panel covers 41 products from January 1, 2021 to March 31, 2025. The retail-price series is seasonally adjusted (Appendix B); the panel is unbalanced because avocado's series begins on April 4, 2022, lowering its N. Treatment status is the raw daily Quota Tariff indicator, whereas the treatment path used in estimation bridges administrative gaps of at most 30 days (Appendix D). Treatment intensity G is a single time-invariant value per ever-treated product ($N = 11$), whose per-product values appear in Table 3; government stockpile releases are summarized at the product–month level. The real crude oil price is shown multiplied by 1,000 for readability; the underlying value equals real WTI divided by the KRW/USD exchange rate. In the regressions the climate variables enter as $[t - 100, t]$ averages and their one-year lags, and stockpile releases enter through an inverse-hyperbolic-sine transform with an occurrence dummy (Section 3.6); the estimation sample varies by horizon under the clean-treated and clean-control conditions, and horizon-specific treated and control counts accompany the estimates.

Table 6 reports the annual import volume of the 41 products in metric tons, obtained by aggregating the monthly import records in the customs data to product–year totals. The rows follow the same four blocks and ordering as panel C of Table 5. The salient contrast in the two import tables lies between the treated groups. The five treated fruits show substantial imports in every year of the sample period (banana, for example, at roughly 400,000 tons in 2024), whereas the treated vegetables, apart from carrot and onion, import on a far smaller scale and much more erratically. The never-treated products range from items with no recorded imports at all to large importers such as bean. This contrast underpins the interpretation in the main text that the treated fruits are predominantly imported, giving a tariff reduction a direct channel into retail prices, whereas the mostly domestically produced vegetables offer little structural scope for that channel.

Table 6: Annual Import Volume by Product

	2021	2022	2023	2024	2025 (Jan.–Mar.)
Treated: leafy and root vegetables					
cabbage	4,450.1	4,333.7	6,741.6	23,527.3	19,989.6
carrot	91,719.9	96,243.6	110,396.1	117,695.2	29,917.4
green onion	8,272.3	1,135.5	7,748.9	23,388.6	3,924.9
napa cabbage	66.7	2,030.1	164.2	4,135.2	5,383.4
onion	38,513.4	70,650.9	117,075.1	78,625.2	42,678.2
radish	7,693.8	2,461.1	1,307.3	14,973.3	8,700.2
Treated: fruits					
avocado	17,313.1	13,479.2	12,671.2	11,251.6	2,450.1
banana	351,902.9	319,852.0	329,025.1	400,769.4	95,363.4
kiwifruit	41,000.8	44,350.8	37,191.8	49,649.9	2,280.4
mango	21,618.4	22,471.4	24,829.7	32,094.9	11,312.3
pineapple	61,688.0	67,104.1	67,047.7	81,573.9	16,720.9
Never-treated: vegetables and other crops					
adzuki bean	18,523.7	20,587.5	21,653.7	23,153.4	5,288.6
bean	236,141.2	277,082.0	242,931.3	251,862.6	78,334.5
bell pepper	0.0	4.4	0.0	0	5.8
cherry tomato	0.0	0	0	0	0
cucumber	0	0	0	0.0	0
dried red pepper	3,809.0	3,937.0	6,458.9	5,184.7	1,012.0
garlic	1,022.1	1,428.9	330.2	425.0	122.0
ginger	5,640.7	2,799.4	4,081.2	4,462.9	1,358.0
green pepper	2.0	0.0	11.1	5.5	6.4
lettuce	6,676.2	3,365.9	5,043.2	11,080.7	2,849.1
mung bean	16,213.3	12,193.6	14,001.4	14,067.6	3,424.4
peanut	30,949.7	31,352.6	30,695.1	33,303.6	10,276.6
perilla leaf	10,458.1	0	0.0	0	0
red pepper	2.0	0.0	11.1	5.5	6.4
scallion	2,320.8	22.7	38.2	116.3	34.7
sesame	39,948.8	37,841.6	49,419.1	53,651.7	15,019.1
soybean sprout	10,458.1	9,125.0	9,515.5	11,170.9	3,376.8
spinach	0	0.0	0.1	0.1	0
squash	25,203.6	19,808.5	12,067.6	22,110.4	7,906.1
sweet pepper	0.0	4.4	0.0	0	5.8
sweet potato	0	0	0	0.0	0
tomato	0.0	0	0	0	0
water dropwort	10,458.1	9,125.0	9,515.5	11,170.9	3,376.8
young napa cabbage	66.7	2,030.1	164.2	4,135.2	5,383.4
young summer radish	22,918.0	22,383.3	23,229.9	21,597.4	4,603.6
Never-treated: fruits					
apple	0.0	0	0	0.0	0
lemon	17,348.3	20,466.7	19,710.4	28,037.2	6,624.2
melon	1,944.6	1,727.1	3,150.3	3,434.6	0.0
pear	18.5	12.8	0.0	10.4	13.6

continued on next page

Table 6 (continued)

	2021	2022	2023	2024	2025 (Jan.–Mar.)
watermelon	0	0	0	0	0

Volume in metric tons. "0" denotes no recorded imports, whereas "0.0" denotes positive imports below 50 kg. The 2025 column covers January–March only and is not comparable to full calendar years.

Under the same blocks and notation, Table 7 reports the corresponding annual import values in thousands of USD, complementing the volume figures of Table 6. Because unit values differ widely across products, the two tables do not rank products alike: kiwifruit, for example, imports less by weight than pineapple yet records substantially larger import values in each year of 2021∼2024.

Table 7: Annual Import Value by Product

	2021	2022	2023	2024	2025 (Jan.–Mar.)
Treated: leafy and root vegetables					
cabbage	1,607	1,485	1,953	8,542	6,083
carrot	42,528	44,096	38,877	51,065	12,831
green onion	8,952	1,120	4,496	15,163	2,216
napa cabbage	45	1,140	77	2,290	1,824
onion	17,111	32,760	46,869	25,788	12,827
radish	1,521	673	435	3,227	3,079
Treated: fruits					
avocado	50,386	42,240	32,267	36,885	9,466
banana	290,149	283,876	304,728	382,031	77,219
kiwifruit	145,453	158,945	131,100	175,129	6,615
mango	90,655	97,045	98,373	138,524	42,950
pineapple	52,639	62,218	65,872	85,603	14,941
Never-treated: vegetables and other crops					
adzuki bean	37,175	43,531	34,537	32,895	7,759
bean	168,983	234,705	219,768	207,219	57,934
bell pepper	0	11	0	0	27
cherry tomato	0	0	0	0	0
cucumber	0	0	0	0	0
dried red pepper	10,654	11,842	23,566	16,418	2,876
garlic	1,836	2,342	636	934	257
ginger	7,915	1,378	6,900	5,628	1,592
green pepper	11	0	11	5	6
lettuce	5,337	2,967	4,969	9,738	2,113
mung bean	32,555	25,106	30,955	31,324	7,567
peanut	60,473	61,984	64,073	65,134	18,543
perilla leaf	18,476	0	0	0	0
red pepper	11	0	11	5	6
scallion	2,953	8	14	47	14
sesame	82,433	86,444	116,143	122,388	32,713

continued on next page

Table 7 (continued)

	2021	2022	2023	2024	2025 (Jan.–Mar.)
soybean sprout	18,476	18,562	18,456	16,898	3,715
spinach	0	0	1	1	0
squash	17,120	14,606	13,696	12,318	5,279
sweet pepper	0	11	0	0	27
sweet potato	0	0	0	0	0
tomato	0	0	0	0	0
water dropwort	18,476	18,562	18,456	16,898	3,715
young napa cabbage	45	1,140	77	2,290	1,824
young summer radish	12,474	12,729	12,144	16,834	4,146
Never-treated: fruits					
apple	1	0	0	0	0
lemon	38,678	43,260	42,271	58,915	13,167
melon	2,920	3,149	5,152	5,752	0
pear	90	71	0	49	56
watermelon	0	0	0	0	0

Value in thousands of USD, rounded to the nearest thousand; “0” denotes no recorded imports or a positive value below USD 500. The 2025 column covers January–March only and is not comparable to full calendar years.

D Appendix: Additional Details of the Estimation Design

Clean treatment events

Table 8 lists every treatment onset among the ever-treated products after bridging administrative gaps of at most 30 days. For each onset, the table reports spell lengths, the gap since the previous episode, and the clean/dirty classification implied by the 120-day washout condition. Three onsets are dirty: carrot (October 29, 2024; 52-day gap), pineapple (January 19, 2024; 102-day gap), and cabbage (January 24, 2025; 103-day gap). All three are therefore excluded from the treated side in every specification, leaving nine vegetable and eight fruit events.

Treated-set attrition in the persistent-treatment specification

Under the stay-on condition, an event remains in the sample that LP-DiD selects for the horizon-h regression only while h is smaller than its spell length, so the treated set surviving at each horizon is exactly the one recorded by the spell lengths in table 8.

Table 8: Quota Tariff Onsets and Washout Classification

Group	Product	Onset date	Spell (days)	Gap (days)	Classification
1	Cabbage	2024-05-10	156	—	clean
1	Cabbage	2025-01-24	$67^{\dagger}$	103	dirty (washout)
1	Carrot	2024-05-10	120	—	clean
1	Carrot	2024-10-29	$154^{\dagger}$	52	dirty (washout)
1	Green onion	2022-07-20	104	—	clean
1	Green onion	2023-05-01	61	181	clean
1	Green onion	2023-11-17	166	139	clean
1	Napa cabbage	2024-05-10	$326^{\dagger}$	—	clean
1	Onion	2022-08-17	177	—	clean
1	Radish	2023-05-01	61	—	clean
1	Radish	2024-07-01	$274^{\dagger}$	366	clean
2	Avocado	2024-01-19	$438^{\dagger}$	—	clean
2	Banana	2022-11-10	52	—	clean
2	Banana	2023-11-17	$501^{\dagger}$	320	clean
2	Kiwifruit	2024-04-05	271	—	clean
2	Mango	2022-11-10	52	—	clean
2	Mango	2023-08-25	$585^{\dagger}$	236	clean
2	Pineapple	2022-11-10	52	—	clean
2	Pineapple	2023-08-25	45	236	clean
2	Pineapple	2024-01-19	$438^{\dagger}$	102	dirty (washout)

† Spell right-censored at the end of the sample period (March 31, 2025).

Spells are measured after bridging administrative gaps of at most 30 days; gaps are days since the previous episode ended.

The composition of this surviving treated set shrinks step by step, at $h = 61$, 104, 120, 156, 166, and 177 for vegetables and at $h = 45$, 52, and 271 for fruits. The persistent-treatment LP-DiD graphs in the main text are truncated at $h = 150$ because, on exactly the day of the compositional change at $h = 156$ (the end of the cabbage spell), the coefficient of the vegetable treatment-status specification rises abruptly from $+0.079$ to $+0.237$. This discontinuity is not a natural effect of the Quota Tariff treatment but the effect of a break in the sample composition caused by events failing the stay-on condition. The sharp rise of the vegetable coefficients beyond $h = 156$ must therefore not be interpreted as a causal effect of the Quota Tariff. In the impulse-response specification, there is no stay-on condition in the first place, so no event ever exits the sample and the composition of the treated set is identical at every horizon.

The washout window and long-horizon tracking

The washout condition treats a Quota Tariff onset as a clean treated event when the product has never been treated during the preceding $L = 120$ days, under the assumption that the dynamic effects of a previously treated episode stabilize within 120 days (Dube et al., 2025). The impulse-response specification, however, tracks each onset for up to 273 days from the day it begins, so an episode that ended 120∼273 days before a new onset gives rise to the following overlap of episodes. The new episode begins 120∼273 days after the previous one and thus satisfies the washout condition, so it is classified as treated. In addition, the past episode, which had begun 120∼273 days earlier, also remains classified as treated, because the impulse-response specification requires no stay-on condition in the first place. As a result, the new episode and the past episode remain treated at the same time, and both enter the sample of the LP-DiD regressions. In the data of this study, four Quota Tariff onsets are in this state. Raising L to 273 days would remove this overlap at the cost of discarding clean events; the results of this paper retain $L = 120$. The clean treated-event sets implied by (3) coincide with those produced by the `nonabsorbing()` option of the `lpdid` Stata module (Busch and Girardi, 2024).