A Temporary VAT Cut in Three Acts: Announcement, Implementation, and Reversal*

Tiago Bernardino†
IIES, Stockholm University

Ricardo Duque Gabriel‡
Federal Reserve Board

João Quelhas§
Banco de Portugal

Márcia Silva-Pereira¶
Nova SBE

March 2024

Abstract

We investigate the pass-through of a Value-Added Tax (VAT) decrease to consumer prices, using Portugal's temporary cut in VAT for a subset of food items in 2023 as a laboratory. Exploiting a novel high-frequency dataset of online retail prices, we use an event study approach to analyze price dynamics across the complete policy lifetime. We find that prices rose by around 1% upon announcement, that the pass-through was almost complete when the policy was implemented, and that the pass-through was approximately 70% at the reversal of the policy. The price reduction was highly persistent over the entire duration of the policy. We estimate that the policy decreased month-on-month inflation by 0.7 percentage points. We find evidence of deflation in producer prices around the implementation, which could be a potential mechanism driving the high pass-through.

JEL classification: E31, H20, H22

Keywords: Value-Added Tax; Pass-through; Public Policy; Inflation

*We thank Erwan Gautier and Elvira Prades, and the participants of seminars at Fed Board, Banco de Portugal, Nova SBE, and IIES for helpful comments. T. Bernardino acknowledges financial support from the Jan Wallanders & Tom Hedelius Foundation (H21-0020). R. D. Gabriel acknowledges financial support from the Peter G. Peterson Foundation and the NBER. M. Silva-Pereira acknowledges financial support from Fundação para a Ciência e a Tecnologia (2021.04898.BD). The views expressed herein are those of the authors and should not be attributed to any of the affiliated institutions. Any errors or omissions are the sole responsibility of the authors.

†Email: tiago.bernardino@iies.su.se
‡Email: ricardofilipeduquegabriel@gmail.com
§Email: jquelhas@bportugal.pt
¶Email: marcia.pereira@novasbe.pt

1 Introduction

In recent years, reductions in the Value-Added Tax (VAT) have been used by several European governments as a policy tool to mitigate the adverse effects of escalating inflation. These policies have targeted categories experiencing price increases above the average inflation rate and essential products disproportionately consumed by lower-income households. One of the categories included food products.

For most countries, consumption taxes represent the largest tax revenue item. In European countries, the VAT is the largest consumption tax. Thus, it is important to understand who benefits from a tax change, as a change in the policy rate has a non-trivial budget impact. In other words, it is crucial to understand whether the VAT change on prices is transferred towards buyers or accommodated by sellers.

This paper revisits this classical question in Economics. We provide empirical evidence on the consumer price dynamics across the complete lifetime of a temporary VAT cut. Using Portugal as a laboratory, we focus on a policy that reduced the VAT on a subset of food items from 6% to 0% during 9 months in 2023. Given the temporary and short-lived nature of the measure, we investigate its pass-through to prices over the entire lifespan: first, when the policy was announced; second, when it was implemented; and third, when it was reverted.

To do so, we use two identification strategies. First, using a novel dataset with daily online retail prices from the main Portuguese retailers, we compare the price dynamics of the food items included in the VAT cut basket with the other food items. Second, using monthly price data at the 5-digit level of the Classification of Individual Consumption by Purpose (COICOP), we compare Portugal's price dynamics with Spain's price evolution. We also discuss the potential mechanisms underlying our pass-through estimates. We further inspect one of them, the costs faced by retailers. For that, we use weekly wholesale price data for a subset of food products and analyze the producer price dynamics during the policy lifetime.

Our investigation into the effects of Portugal's temporary VAT reduction on 46 selected food items in 2023 yields a nuanced narrative that unfolds across the three policy acts. Upon the announcement, we find an anticipatory response, with relative prices of the treated items exhibiting an upward adjustment compared to the non-treated counterparts.

When the policy was implemented, reducing the VAT rate from 6% to 0% on the targeted set of goods, we observe a remarkably high pass-through rate to retail prices of 99%. This almost complete pass-through underscores the policy's efficacy in directly lowering consumer prices, closely aligning with its intended objective of providing immediate relief to consumers. Notably, this high pass-through persisted during the almost total life of the policy. Only by the end of it, the gap between treated and non-treated items starts to close.

Following the reversion to the 6% rate, our estimations indicate a 4.2% increase in prices compared to the week preceding the reversal. Upon the conclusion of the policy, the price index for treated items aligns with that of non-treated items. This observation suggests that the policy's effects were temporary, ceasing as soon as the policy ceased to be in effect.

We also explore the heterogeneity in the pass-through estimates across the various food categories. We find very small differences across food categories, suggesting the policy was effective independently of the product characteristics.

Given these price dynamics, we estimate the impact of the policy on the month-on-month inflation rate as measured by the Harmonized Index of Consumer Prices (HICP). Using the estimated impact that the policy had on treated food items in combination with the weight that this set of goods has in the consumption basket, we estimate that headline inflation fell around 0.7 percentage points (pp). This quantifiable impact on inflation highlights the potential of targeted fiscal policies as a tool for managing inflationary pressures.

The second empirical exercise at the macro-level further confirmed these results. Using the price evolution in Spain as counterfactual, we find that the consumer price indices of the treated categories were, on average, 5.48% lower than the control group. This finding suggests that a significant proportion of the VAT reduction was effectively transmitted to consumers. Moreover, it underscores the enduring nature of this effect, as this reduction persisted high over the entire duration of the policy.

Looking into the mechanisms underlying the observed pass-through dynamics, we find that wholesale prices of food items were in a descendant phase, at the policy's announcement and implementation. This deflationary trend in input costs may have facilitated the almost complete pass-through observed at the retail level, suggesting that the timing of fiscal

interventions relative to input price dynamics can significantly influence their effectiveness in transmitting to consumer prices. We also discuss other possible mechanisms that can explain such high pass-through such as the salience that the policy had for consumers and the dynamic interactions between supermarkets and government.

Related Literature. Despite the challenge of pinpointing the tax incidence of VAT changes, existing work estimates the pass-through to consumer prices across different sectors. Focusing on food products, studies offer a range of estimates from 60% (Benzarti et al. 2023) to 93% (Amores et al. 2023). Other studies examined this in sectors such as restaurants (Harju and Kosonen 2014), hairdressing (Benzarti et al. 2020), cinema (Arce and de Antonio 2020), gas (Gautier et al. 2023; Montag et al. 2023), and online retailing (Fedoseeva and Van Droogenbroeck 2024).

Our finding of an almost complete pass-through at the implementation of the VAT cut places us above the upper bound of previous studies, representing a novel result in this strand of the literature. The persistence of the effect found is also unusual. Recent papers using equivalent data show a reversion of the pass-through after several weeks (Forteza et al. 2023). This result is puzzling for classical public finance theory, where the pass-through depends on the relative elasticity of demand and supply. However, it resonates with studies by Poterba (1996) and Besley and Rosen (1999), which found that sales tax hikes are fully passed to customers.

By bringing a complete life cycle perspective, this analysis contributes to the existing literature, offering valuable insights into the dynamics of the pass-through from temporary tax changes. It emphasizes the relevance of considering the entire life cycle of a policy to accurately evaluate its effects on consumer prices. Moreover, we resort to daily online prices sourced from various retail supermarkets in our empirical analyses. This novel approach yet to be thoroughly investigated in the existing literature allows an examination of the dynamics of prices within a narrower time frame, wherein they are primarily influenced by the policy under examination.

Our findings also contribute to the literature examining how changes in consumer taxes shape different economic outcomes. Despite the prevalence of Value-Added Taxes in European countries, as they are the largest tax revenue category, there is limited research inspect-

ing their economic effects. While some exceptions exist, such as: Kosonen (2015) and Benzarti et al. (2020), which analyze the effects of a VAT cut in the hairdressing sector on prices, quantities, and profits; Freund and Gagnon (2017), which show that the real exchange rate tends to rise by the full amount of any consumption tax increase; and Benzarti and Tazhitdinova (2021) that find small elasticities of trade flows to VAT changes. A novel angle in our research is the assessment of whether producer prices and, therefore, a potential rise in retailers' margins played a role in the almost complete pass-through. We contribute to the recent public debate on sellers' inflation, in which firms further increase prices beyond the rise in costs during inflationary environments.

Finally, our analysis sheds light on potential asymmetries in the pass-through of consumption tax changes, as evidenced by the differential rates observed during the implementation and reversal phases of the policy. Understanding these asymmetries and their underlying drivers could have important implications for the design and timing of future tax policies aimed at managing inflation or achieving other economic objectives.

Outline. The paper is organized as follows. In Section 2 we describe the context in which the policy was taken and the moments of the policy lifetime. In Section 3 we discuss the data sources used, in particular the novel supermarket daily prices dataset. In Section 4 we present the results of the exercise coming from the first identification strategy that compares different groups of products. We estimate the pass-through of the VAT cut during the policy lifetime, as well as the contribution to reducing inflation. In Section 5 we present the results of the second identification strategy that compares Portugal with Spain. In Section 6 we investigate potential mechanisms that may be driving our novel result. Section 7 concludes.

2 Institutional Background

The VAT is a consumption tax levied on the value added to goods and services during each stage of the production and distribution chain. In the European Union, the VAT is included in consumer prices, making it less salient to consumers during purchases. Firms collect the VAT from consumers and remit it to the government, offsetting it with credits for the VAT paid on

input costs. This mechanism ensures that only the value-added portion is subject to taxation. Consumers who purchase goods and services for final consumption bear the full tax burden on the entire value of the final goods they buy. In Portugal in 2022, this tax represented 21.3% of the total government revenues (9.4% of GDP). There are three VAT regimes: the standard rate of 23% applied to the majority of goods and services, the reduced rate of 6% for certain essentials, and an intermediate rate of 13%.

In November 2021, Portugal's headline inflation rate, as measured by the year-on-year HICP, exceeded the European Central Bank's target of 2% and embarked on an upward trend, peaking at 10.6% in October 2022. Similar to other European countries, food price inflation in Portugal increased much more sharply than the total basket of goods and services. Figure E.4a in Appendix E illustrates the evolution of these two series over the recent inflation surge.

As food prices accelerated quickly, public pressure mounted for the government to take action. This pressure intensified when the Spanish government implemented a VAT reduction on a basket of essential goods. Nevertheless, several members of the Portuguese government, including the Minister of Finance, categorically rejected a VAT reduction, arguing that such a policy would have little effect on prices and that retailers would absorb a significant portion of the VAT reduction.

On March 24, 2023, the government announced a pact to stabilize and reduce food prices. Among the measures in the agreement was a cut in the VAT rate on a selected list of food items. These were chosen taking into account the Ministry of Health's healthy food basket and data from distribution companies on the most consumed products by the Portuguese population. The final list was announced on March 27, 2023, and voted on by the parliament.

On April 18, 2023, the policy came into effect for 46 essential food items, including some fruits, vegetables, legumes, meat, fish, dairy products, cereals, and oils. All of these products, except vegetable oils, were previously taxed at a 6% rate.[1] The complete list of products can be found in Appendix A.

Importantly, the policy garnered significant attention from the media and the general public. Consumer associations, as well as journalists closely monitored the policy's implementa-

[1] In our baseline analysis, we exclude vegetable oils and analyze them separately in Appendix E.6.

tion, assessing whether supermarkets were raising prices. Figure E.4b in Appendix E shows the search intensity on Google in Portugal for ”Imposto sobre o Valor Acrescentado” (Value-Added Tax in Portuguese) and ”IVA” (VAT in Portuguese) over 2023.

The policy was initially announced as a temporary price relief measure that would last until the end of October. However, in September, the government announced an extension until the end of 2023. On January 5, 2024, the policy was reversed, and the VAT rate applied to the targeted products returned to 6%.

Given the policy timeline, which is summarized below, we structure our empirical analysis around three crucial moments – which we name acts – in the lifespan of this temporary VAT cut. The first one is the unexpected announcement. The second act is the implementation of the policy. The third one is the reversal of the temporary VAT cut. Analyzing these three acts allows for a more comprehensive understanding of the price dynamics throughout the policy's life cycle.

The timeline of the policy with the acts identified is the following:

- Mar 24, 2023: A VAT cut is announced for “essential products” **(Act I)**
- Mar 27, 2023: Official announcement of the list of products covered
- Apr 18, 2023: Implementation of the policy **(Act II)**
- Sep 6, 2023: Announcement of an extension until the end of 2023
- Oct 27, 2023: Announcement of the official ending date of the policy
- Jan 5, 2024: Reversal of the VAT cut on all the products **(Act III)**

3 Data and Descriptive Statistics

3.1 Data

This study draws upon three distinct datasets to explore the effects of the VAT cut on price dynamics. The first dataset comprises daily online supermarket prices from Portugal, offering a detailed look at consumer price movements. Complementary wholesale price data and aggregate consumer prices further enrich the analysis, providing a comprehensive perspective on the VAT cut's impact across different market levels. Next, we describe each dataset.

Supermarket Daily Prices. The empirical analysis uses a novel and comprehensive dataset on online supermarket prices. This dataset was compiled by the Banco de Portugal Microdata Research Laboratory (BPLIM) and contains granular data on products from the main retailers in Portugal. These collectively represented around 60% of the retail market in the country in 2022. The dataset is particularly unique as it encompasses a comprehensive array of goods typically found in supermarkets and covers all the products affected by the policy.

The BPLIM dataset is built using automated web scraping algorithms that retrieve, on a daily frequency, relevant product information, including the product name, brand, capacity, COICOP 5-digit category, and price (regular price, price before promotion, and price per unit). For our analysis, we use data from January 2023 until January 2024. BPLIM has also identified, for each supermarket, the list of products subject to the VAT cut. This comprehensive dataset encompasses around 60,000 daily supermarket-product combinations of observations, among which 12% had a reduced VAT rate between April 18, 2023, and January 5, 2024. In Appendix C we describe the data cleaning procedures performed on this dataset.

Using this dataset, we can construct a series that closely tracks the official food inflation series. Figure C.1 in Appendix C plots the food inflation rate reported by Statistics Portugal against the food inflation rate implied by this dataset. For our sample period, they coincide except in the month when the policy was implemented due to the timing of data collection. As prices are collected in the first two weeks of each month and the policy was implemented in the third week of April, the official inflation rate of April does not include the VAT reduction.

Wholesale Producer Prices. We complement our analysis with weekly wholesale price data from the Agricultural Markets Information System of the Planning, Policy and General Administration Office.[2] Data is available at a weekly frequency for several product categories, species, regions, and local markets. Some products do not have available data for selected weeks due to product seasonality or the timing of the harvesting season. The product categories for which we analyze wholesale prices are selected based on the existence of a match to a 5-digit COICOP category, which we also have in the supermarket daily prices dataset,

[2]Originally, in Portuguese, the *Sistema de Informação de Mercados Agrícolas* of the *Gabinete de Planeamento, Políticas e Administração Geral*. The data is publicly available in this link.

yielding a total of 9 categories of treated goods.

Aggregate Consumer Prices. The third data source used is the HICP at the 5-digit COICOP level, provided by Eurostat. In the second empirical exercise, we use data for Portugal and Spain, with the latter being used as a control group for the price evolution in the former.

3.2 Descriptive Statistics

The newly assembled supermarket daily prices dataset provides interesting insights about the behavior of retail prices in Portugal. Panel (A) of Figure 1 shows the price evolution of the treated basket (food products that were included in the VAT cut basket, except vegetable oils) and the control basket (the other food products that were not included). The two lines show the unweighted average of the daily price index between January 1, 2023, and January 31, 2024. The index is normalized to 100 on the day before the policy was implemented, March 17, 2023. The VAT cut was announced on March 24, 2023, implemented on April 18, 2023, and ended almost 9 months later on January 5, 2024 – the three vertical dashed lines indicate these dates. Treated products were taxed at the rate of 6% before the policy implementation and returned to this rate at the reversal day.

Four patterns appear in this plot. First, the food products not subject to a VAT cut seem to be a good control group for the treated basket as prices follow a similar trend before the policy was announced, i.e. there are parallel trends between these two groups of products. Second, the largest response of the treated group when the policy started occurred on April 18, 2023, the day that the policy was implemented. Third, the difference between the two food products remained relatively constant for 6 months, and in the last 2 months that the policy was in action, the difference started shrinking. Lastly, on the reversal day, the average price index of the treated products returns to the same level as the control group.

Panel (B) of Figure 1 shows the frequency of price adjustments on each week between January 2023 and February 2024. We consider that a product had a price change in a certain week if its last price is different from the last price of the previous week. The left stacked bars graph shows these frequencies for the control basket, whereas the right stack bar graph

Figure 1: Price indices during the lifetime of the Temporary VAT Cut in Portugal

(A) Price indices for treated and control products

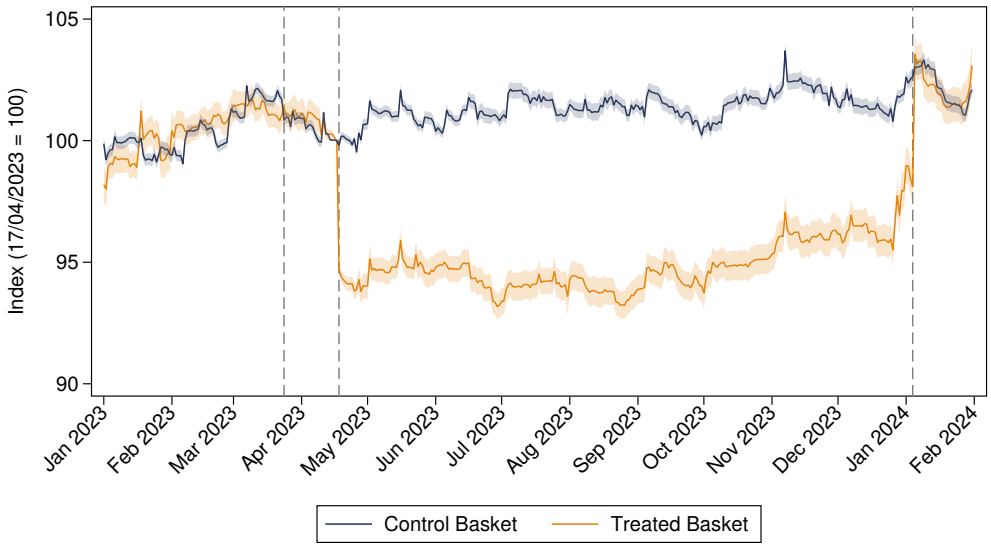

Notes: This graph plots the price evolution of the basket of food products that had a VAT reduction (treated basket in orange) against the basket that kept the same VAT rate (control basket in blue). We use a daily unweighted average of the product-level price index, normalized to 100 on the day before the implementation of the policy. The graph also includes a band defined by two times the standard error around the mean. The first dashed line indicates the announcement of the measure (March 24, 2023), the second dashed line indicates the implementation day (April 18, 2023), and the third dashed line indicates to the reversion date (January 5, 2024).

(B) Frequency of positive and negative price changes per week

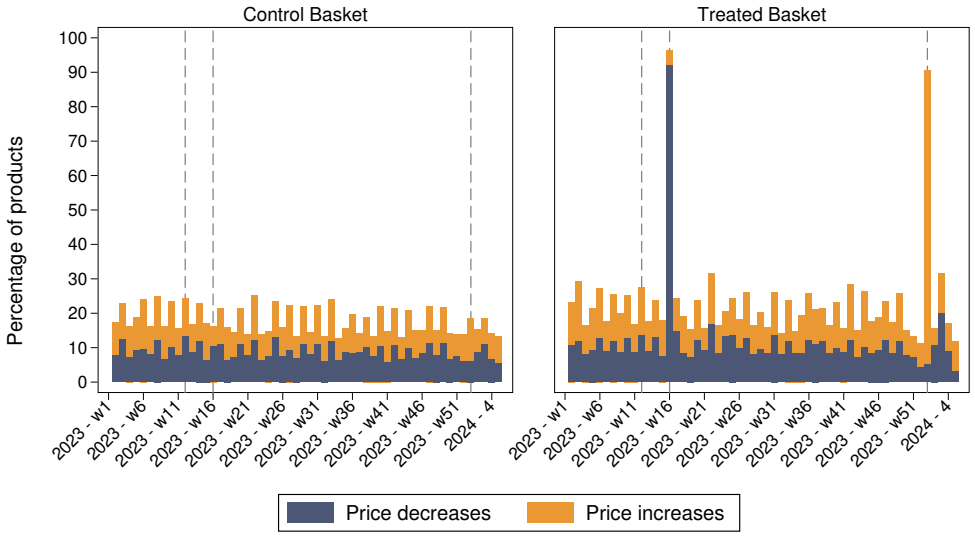

Notes: These plots show the percentage of products with an increase (orange bars) and a decrease (blue bars) in the last price level of a given week, with respect to the last price from the previous week. The left plot shows these statistics for the food products not affected by the VAT reduction – control basket. The right plot shows the same statistics for the food products affected by the measure – treated basket. The dashed lines indicate the announcement (week 12, 2023), implementation (week 15, 2023), and reversal (week 1, 2024) dates, respectively.

shows the treatment basket of goods. Each of the graphs splits the frequencies by direction of adjustment, i.e., price increases (orange bars) and price decreases (blue bars).

Both groups considered have a similar behavior of price adjustments, except in the weeks when there was a policy change. Within the same week, around 20% of the goods changed their prices, on average. In the week when the VAT cut was implemented, week 15 of 2023, around 92% of the goods had a price decrease, among the goods included in the treated basket. On the other hand, in the week when the reversal took place, week 1 of 2024, about 85% of the goods saw their prices increase. In comparison with the control basket, in these weeks the prices followed a similar behavior as in the other weeks of the year, suggesting again that the basket of food products not included in the policy is a good control group.

4 Micro-Level Analysis

4.1 Empirical Approach

We employ an event study approach to compare the evolution of the prices of food items affected by the VAT cut to the prices of other food items not affected by the policy in the three moments in time. Making use of the unique policy setting in Portugal, we study the consumer price behavior when the policy is announced (March 24, 2023), when the policy is implemented (April 18, 2023), and when the policy is reverted (January 5, 2024). More specifically, we consider the following linear panel model with dynamic policy effects:

$$P_{i,t} = \alpha_i + \gamma_t + \sum_{m=-7}^{7} \beta_m Z_{i,t-m} + \varepsilon_{i,t}, \tag{1}$$

where P_{it} corresponds to the price index (normalized to 100 on the day before the respective act) of each item i, which corresponds to a product-supermarket combination, in day t, α_i denotes an item fixed effect, γ_t a time fixed effect, and $Z_{i,t}$ is an indicator function that is equal to one for products included in the zero VAT basket after the announcement, implementation or reversal, depending on the specification which we separately estimate for each policy act.

The $\{\beta_m\}_{m=-7}^{7}$ is our set of coefficients of interest that summarize the magnitude of the

dynamic effects. They can be interpreted as the cumulative average treatment effects (ATT) at each horizon m. As such, a positive value means that prices increased and, conversely, a negative value means that prices decreased. For $m < 0$, the β_m coefficients estimate the lagged response to the policy change, whereas, for $m > 0$, the coefficients estimate the dynamic response after the policy. We use a 7-day window around each act in the estimation.

The main identification assumption of our coefficients is that absent the tax change, there would have been no change in prices of the treated relative to the control products, i.e. there would be parallel trends. We formally test this hypothesis resorting to a Wald test on whether the coefficients $\{\beta_m\}_{m=-7}^{-2}$ are different than zero (Freyaldenhoven et al. 2021). Moreover, both sets of fixed effects contribute to our identification strategy: on the one hand, the item fixed effects control for unobserved, time-invariant characteristics that might influence the outcome variable such as the brand value or the typical validity. On the other hand, time fixed effects account for shocks or trends that simultaneously affect all products such as overall economic conditions, seasonal effects, or common inflationary trends.

We use the computational implementation of the event study for the panel data estimator by Freyaldenhoven et al. (2021). Among the key modeling assumptions, we normalize β_{-1} to be 0 such that the plotted coefficients can be interpreted as estimated effects relative to the effect of the policy adoption one period ahead and choose a specific control group that only includes other food items sold in the retailers, as described in Section 3.

4.2 Price pass-through of VAT cut to consumer prices

Figure 2 reports our main estimates of Equation 1. It displays the price dynamics, $\{\beta_m\}_{m=-7}^{7}$, for each of the three moments that we analyze. In all three panels, we thus plot the percentage change of prices around the announcement (Act I), the implementation (Act II), and the reversal (Act III) of the temporary VAT cut. The shaded region in each panel corresponds to the 95 percent confidence interval around the point estimates.

Panel (A) depicts the estimated coefficients around the announcement day of the VAT cut policy, the first Act. The coefficients fluctuate around zero in the days leading up to the announcement, indicating no significant price changes in anticipation. However, there is a

noticeable increase in the coefficient immediately after the announcement, suggesting that prices rise in response to the policy announcement, potentially due to anticipatory behavior from retailers or producers. The gap between the control group and the treatment group is approximately 1.27% and remains roughly at the same level for the 7 days post-event.

Three weeks after the policy was announced and after the parliamentary approval the policy started. In panel (B), we show the price dynamics in the week before and after the policy implementation. Before the event, the coefficients are relatively stable around zero, implying no significant pre-implementation price difference changes. Immediately after, there is a sharp and substantial drop in the coefficients indicating that prices declined substantially in response to the VAT reduction taking effect. On average, prices were 5.56% lower the week following the VAT cut implementation which corresponds to a striking pass-through of 99.2%.[3]

Taking the evidence in the first two panels into account and the anticipation effect at face value, we can estimate a lower bound for the VAT cut pass-through. In this case, the pass-through of the policy would be about 76%. Even adjusting for the announcement effect, such a pass-through is still considerably higher than most estimates in the literature. We explore potential mechanisms that can account for this behavior in Section 6.

The policy was in effect for almost nine months. In panel (C), we plot the event-study coefficients of the policy around the date when the policy was reverted and the VAT rate returned to 6%. The plot shows an increasing trend in the price difference between the two groups in the week before the reversal. In fact, this dynamic likely started even before as Figure 1 hints. The event-study plot on this third act also shows that in the week after the reversal, prices increased, on average, 4.18%. This corresponds to a pass-through of 69.7% if prices were to rise by the full amount of the tax rate.

Overall, the event study analysis provides compelling evidence that the temporary VAT cut policy had a tangible and immediate effect on consumer prices, both in the downward direction upon implementation and in the upward direction upon reversal. It provides a unique perspective on the pass-through of the temporary VAT reduction to consumer prices across the full life cycle of the policy – from announcement to implementation and reversal. While

[3]Note that the VAT rate is decreasing from 6% to 0% corresponding to a $-0.06/1.06 \times 100 = -5.66\%$ decrease in prices in the case of full pass-through.

Figure 2: The Three Acts of the Temporary VAT Cut in Portugal

(a) (A) Act I: Policy Announcement

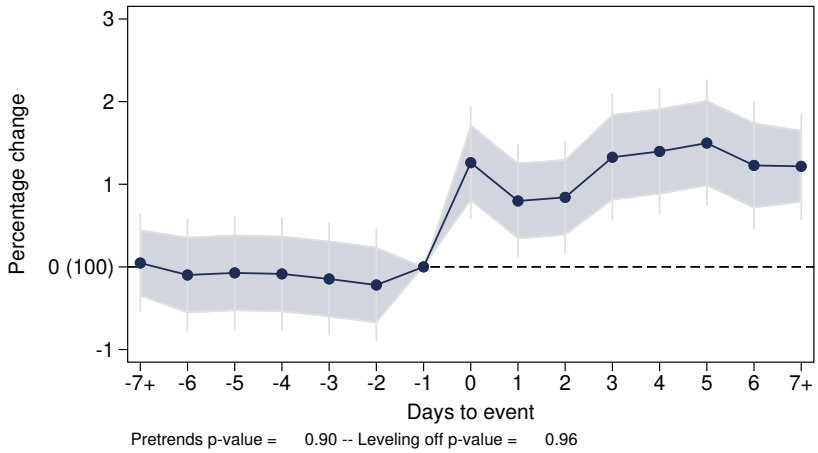

(b) (B) Act II: Policy Implementation

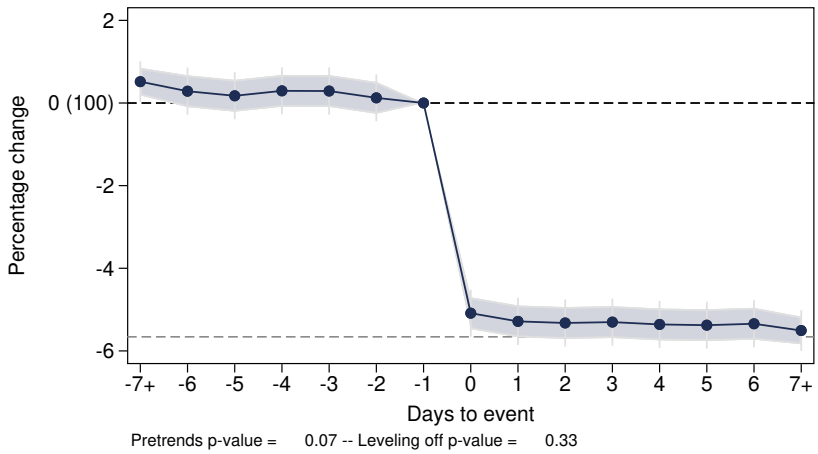

(c) (C) Act III: Policy Reversal

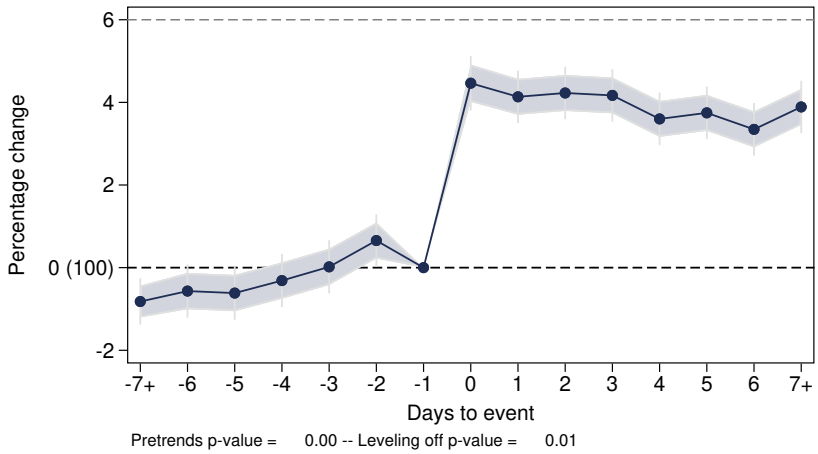

Notes: The panels report the event study estimates from Equation (1) in three moments of the policy lifetime. Panel (A) plots the estimated coefficients for the announcement day (March 24, 2023), panel (B) plots for the implementation day (April 18, 2023), and panel (C) plots for the reversal day (January 4, 2024). The shaded areas correspond to 95% confidence intervals.

previous studies have examined the pass-through of tax changes, few have addressed the complete impact of a temporary tax policy, from its initial announcement to its final reversal.

Contrary to the common wisdom, these results highlight a flipped asymmetric incidence of a VAT reduction. The higher pass-through at implementation, which contrasts with some earlier evidence of price stickiness to tax reductions as observed by Benzarti et al. (2020), was not fully reverted at the time of the policy reversal. This may be attributable to the temporary and widely announced nature of this policy which incentivized retailers to pass on the VAT reduction to consumers on impact, even when considering their reaction at the announcement. Moreover, as the media coverage and time passed by, the persistence of the effect waned out thus leading to a 69.7% pass-through when the policy was reverted.

Heterogeneous effects. The almost complete pass-through of the VAT cut to consumer prices conceals that not all product categories responded uniformly. It is plausible that the degree of pass-through varied across products, potentially influenced by factors such as price elasticity, market structure, and consumer behavior. While the unexpected nature of the VAT announcement may have limited the scope for such factors to influence initial price responses, retailers ultimately have the discretion to set final posted prices. Therefore, an analysis of heterogeneous effects across product categories is warranted in our setting.

In Appendix D.1, we present preliminary results focusing on heterogeneity across the three stages of the policy's life cycle. Figure D.2 provides evidence of heterogeneous pass-through patterns across 18 product groups. While the majority of categories did not exhibit statistically significant announcement effects, we uncover that categories with the stronger anticipatory price increases typically experienced the most substantial price declines upon implementation (e.g., bread, fresh fruit, frozen vegetables, and yogurts). Nonetheless, there are instances where anticipatory price hikes occurred, but the implementation had a muted impact (rice, frozen fish, and fresh fish). The latter groups' price dynamics partially account for the almost complete pass-through.

What was the impact on food inflation? By combining the coefficients estimated from Panel (B) of Figure 2 with the HICP weights associated to the group of treated products, we

quantify the contribution of the VAT cut to the decline in food inflation. The 5-digit COICOP categories affected by the VAT cut ranged from 12.95% to 13.3%, corresponding to the 2022 and 2023 weights computed by Statistics Portugal based on the Household Budget Survey, respectively. This implies that the effect on aggregate inflation at implementation is estimated to be between 0.72 and 0.74 pp.[4] Looking at the policy reversal, and using the same reasoning, we estimate that the end of the zero VAT policy led to an increase between 0.54 and 0.56 pp. This corresponds to the beginning of the policy-induced reduction if we consider the anticipatory effect of the policy announcement. Note that these calculations can bias the estimated contribution of the policy in reducing inflation. We assume that consumption patterns are the same before and after the policy, which may not be the case.

In our setting, the abnormally high pass-through and the high inflation environment are important factors. Consistent with the recent review by Checherita-Westphal et al. (2023), we provide robust evidence that a reduction in indirect taxes can effectively curb inflation. Nonetheless, it is crucial to recognize that lowering VAT rates to ease inflationary pressures may lead to a reduction in government revenues at a critical juncture when fiscal resources might be required to address other economic challenges or fund essential public services. Therefore, policymakers should exercise caution when considering indirect tax cuts to reduce inflation, as they might undermine the sustainability of public finances.

4.3 Robustness Exercises

In an event study, causal impacts are inferred by comparing treated units to control units, in our case, the VAT cut basket of goods with other food products sold by retailers. Time-invariant level differences between units are allowed, as we have a 7-day window around each policy event, as well as general common trends, which one could expect due to the high inflation environment in Portugal during 2023. Nevertheless, drawing causal inferences requires three pivotal prerequisites: the existence of suitable control groups, which we argue we have by focusing only on food items; a parallel trend assumption, which we discuss and explicitly test, finding no evidence of violation at the announcement and implementation stages;

[4]With the anticipation effects from the announcement day, the effect would be 0.56-0.57 pp.

and that no potentially unobserved confounding factor is correlated with the policy, which is highly unlikely given our focus on a small window around the policy events. In this subsection, we present a battery of robustness checks ranging from an alternative methodology to different data-cleaning approaches, to strengthen our analysis and confidently provide an accurate range of estimates for each pass-through estimate.

Synthetic Difference-in-Difference. Although we provided evidence in favor of these three assumptions, the reader might still not be convinced. One particular solution to this challenge in the literature has been the application of synthetic control methods (Abadie 2021). Specifically, when integrated with difference-in-differences methodologies (Arkhangelsky et al. 2021), this approach accommodates potential time-varying confounds by treating them as latent variables. This strategic recourse aligns seamlessly with our study, given the expectation that any latent confounding element would emanate from aggregate factors, such as macroeconomic shocks, exerting disparate impacts across all units. In the Appendix D.2, we provide a set of results based on the Synthetic DiD methodology. We find evidence in favor of a decrease in the prices of the treated basket of around -5.55%, which overlaps with the baseline estimates presented in Figure 2 and confirms the result of an almost complete pass-through. For the other acts, the results are also in line with the ones presented above: an impact of 0.7% in the announcement day and 4.2% for the reversal of the policy.

Other Specifications. In addition, we explore other specifications leveraging the richness of our available data. Initially, we conduct analyses using different control groups. This entails examining two distinct scenarios: firstly, including all products sold across the online marketplace of the retailers under scrutiny, and secondly, restricting the analysis solely to food (i.e., excluding drinks) and then only non-food products. This last maneuver aims to discern the nuanced demand effects stemming from the VAT cut. Our findings reveal a lack of distinct substitution effects surrounding the events, as evidenced by the similarity in outcomes across these varied control group configurations.

Subsequently, we broaden our analytical scope by diversifying the outcome variables under investigation. One such variable involves the assessment of price per unit, a metric that ac-

counts for fluctuations in product quantities sold. We also scrutinize the regular price of items, disregarding any promotional discounts. This dual approach allows us to discern whether supermarkets strategically adjusted their pricing in response to the VAT changes.

Furthermore, we acknowledge the presence of missing data points across several observation days. These gaps may arise from factors such as product unavailability or inherent measurement errors inherent in our web scraping algorithms. To mitigate the impact of these gaps, we employ diverse data imputation techniques. Moreover, to ensure consistency and comparability in our analysis, we maintain a constant basket of goods for which complete information is available throughout the entire sample period. A summary of these results is available in Appendix D.4.

5 Macro-Level Analysis

In this section, we employ an alternative identification strategy to address the core question of how much of the VAT reduction was passed through to consumer prices. Drawing upon less granular data, this exercise does not allow an analysis of the three distinct policy stages as conducted in Section 4. Instead, we focus on the second act, the policy implementation.

By using an alternative empirical approach and a different data source, this section aims to provide a complementary perspective on the degree of pass-through of the VAT cut to retail prices. While the granularity of the data may be coarser, this analysis offers an additional layer of robustness to the findings presented earlier, shedding light on the policy's impact during the critical implementation period.

5.1 Empirical approach

To estimate the impact of the temporary VAT cut on consumer prices, we employ a difference-in-differences strategy. For this approach, we require a suitable control group, which in this exercise is Spain. Spain shares similar production and consumption structures with Portugal, and both countries were affected comparably by recent economic shocks. Furthermore, the price evolution of the food baskets in Portugal and Spain follow similar trends, making the

parallel pre-trends assumption reasonable when using Spain as a control for Portugal. Figure E.5 in the Appendix illustrates the price evolution of the food HICP in both countries.

The underlying idea is that Portugal and Spain were following the same food price trend when the Portuguese government implemented the VAT cut on a set of food products. More formally, we estimate the following equation:

$$P_{i,t} = \mu + \alpha C_i + \gamma T_t + \tau C_i \times T_t + \varepsilon_{i,t}, \tag{2}$$

where $P_{i,t}$ is the price index (normalized to 100 in April 2023) of category i in period t, C_i is an indicator variable that takes the value of one for Portugal, and T_t is an indicator variable that takes the value of one after the VAT cut reduction.

We use monthly HICP series by COICOP 5-digit level, as described in Section 3, for the two countries. As Spain also implemented a VAT change for some food products in December 2022, we restrict our analysis period to start from January 2023. The treatment period is set to begin after May 2023. Note that the price data for April 2023 was collected in the first two weeks of the month, hence the series for April is not affected by the VAT cut.

5.2 Results and Discussion

Table 1 describes the estimation results of Equation (2) using ordinary least squares.

Table 1: Impact of the VAT Cut on food inflation rate over the entire lifetime of the policy

	(1)	(2)	(3)
C_i	-1.881***		1.769***
	(0.000)		(0.000)
T_t		0.445	3.21***
		(2.738)	(0.000)
$C_i \times T_t$			**-5.475*** **
			(0.000)
$N \times T$	1 188	1 188	1 188

Notes: Robust standard errors clustered at country-level. * p $<$ 0.10, ** p $<$ 0.05, *** p $<$ 0.01

Our coefficient of interest is presented in column (3). During the period when the policy was in effect, the consumer price indices of the affected product categories were, on average, 5.48% lower than the control group. This finding suggests that a significant proportion of the VAT reduction was effectively transmitted to consumers, implying that the benefit of this reduction was passed on to them. Moreover, it underscores the enduring nature of this effect, which persisted consistently over the entire duration under examination.

This result is consistent with the high degree of pass-through described in Section 4. At the time of policy implementation, consumer prices declined, but not by the full amount of the VAT reduction. In this alternative identification strategy setting, we estimate a consumer price pass-through rate of 96.8%. The small discrepancy between the two exercises may be attributable to the coarser granularity of these data, which does not allow us to distinguish the first two moments of the temporary VAT reduction.

Nevertheless, we can still approximate the impact of the VAT reduction on food inflation. Using the same back-of-the-envelope computations as in Section 4, the temporary VAT cut contributed to a reduction in food inflation between 0.71 and 0.73 pp, aligning with the numbers reported in the previous exercise. As a benchmark, one may refer that the theoretical impact, the one expected with a pass-through of 100%, would be 0.8 pp.

6 Inspecting the Mechanism

In this section, we investigate the potential mechanisms underlying the almost complete pass-through of the VAT cut upon implementation.

First, as mentioned in Section 2, the discussion around the policy gathered significant media coverage and public attention due to several factors, including the rapid escalation of food prices, anecdotal claims of price gouging by supermarkets which raised prices by more than necessary to cover increases in their costs, pressure from retailers for a VAT reduction, and public officials denying that such a policy would be enacted. Before the announcement, the government had instructed food safety inspectors from the Economic and Food Safety Authority to monitor prices, which was frequently mentioned in the media. This degree of

media attention and public scrutiny is one potential mechanism that provided an incentive for food retailers to pass through the VAT cut to consumers. Figure E.4b provides evidence of the high level of public attention during this period.

Second, before the official policy announcement on March 24, 2023, the government met with stakeholders across the food supply chain and established several formal agreements with collective groups such as farmers, distributors, and industrial associations. These meetings aimed to ensure that the agents represented by these groups were committed to reducing and stabilizing prices and would therefore fully pass through any reduction in indirect taxes targeting goods produced, distributed, or sold by these agents.

Third, this policy was implemented during an inflationary environment. Existing literature shows that, during high inflation periods, attention to prices increases (Binder and Kamdar 2022; Pfäuti 2023). This can justify heightened attention to prices, thus motivating the increased pass-through.

Fourth, shortly before the implementation, producer prices embarked on a deflationary path. As movements in producer prices for the set of treated goods in wholesale markets could serve as a proxy for the cost changes faced by retailers, a decrease in costs around the time of the policy can also be a potential mechanism.

Of the set of previously discussed possible mechanisms that could justify the very high pass-through of the VAT cut, we empirically analyze the last one, namely the deflation in producer prices, using the wholesale price data described in Section 3. We do this for a selection of product categories for which data is available, and which have a direct mapping to a COICOP 5-digit category, also included in the supermarket data previously used. The list of the 9 product categories of treated items considered in the following exercise can be found in Appendix B.

Methodology. We employ an event study approach using a linear panel model with dynamic policy effects such as the one described in Section 4.1. While the setting differs, the key distinction between the empirical approach used here and the one previously described is that we do not have a control group due to insufficient available data on product categories. Thus, we pool all product categories for which data is available.

Since the products considered are all agricultural products, and the wholesale prices have information on the region and local market where the products are sold, the panel variable used is a combination of category-region-market. We employ high-dimensional fixed effects to estimate it. The policy variable used to test whether wholesale prices were statistically significantly different before or after a given date takes three different values: the announcement date, the implementation date, and the reversal date, corresponding to the three acts initially mentioned.

Additionally, we allow the policy variable for each of the acts to vary by one week before and after each act, as we do not require nor expect the wholesale price deflation to necessarily coincide with the exact date of each of these three moments but rather to occur around the event. We select a time window of 6 weeks before and after each event.

Results. The results with wholesale prices are presented in Figure 3. Comparing the differences between pre- and post-event averages from the linear regression, we estimate that the post-announcement average is -15.81% lower than the pre-announcement, statistically significant at the 10% confidence level. Upon implementation, we find that the difference increases to -26.71%, also statistically significant. Both these results align with the ongoing deflation in producer prices coinciding with the timing of the policy. Upon reversal, we do not see any statistically significant differences between the pre-event and post-event averages for the considered window of 6 weeks before and after the reversal.

The obtained result around the announcement and implementation of the policy provides evidence that the costs faced by supermarkets were on a downward path around the time of the first two acts. This could contribute to explaining why the pass-through of the VAT cut was almost complete at implementation, signaling a less-than-complete pass-through. Since these are wholesale prices in agricultural markets, thus not necessarily representing the price at which supermarkets purchase these goods, we cannot further quantify how much of the pass-through of the VAT cut could have been attributed to this channel. Nonetheless, this can signal that measures such as a VAT cut are more likely to experience a higher pass-through if implemented during a deflationary phase of input prices.

Figure 3: The Three Acts with Producer Prices in Portugal

(A) Act I: Policy Announcement

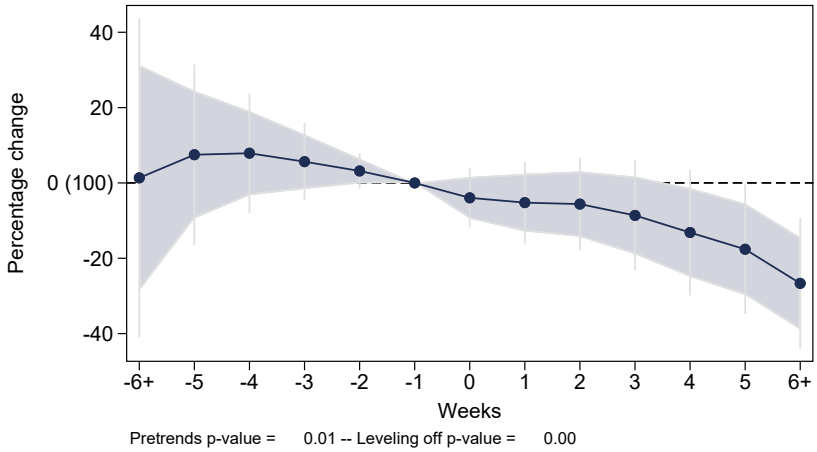

(B) Act II: Policy Implementation

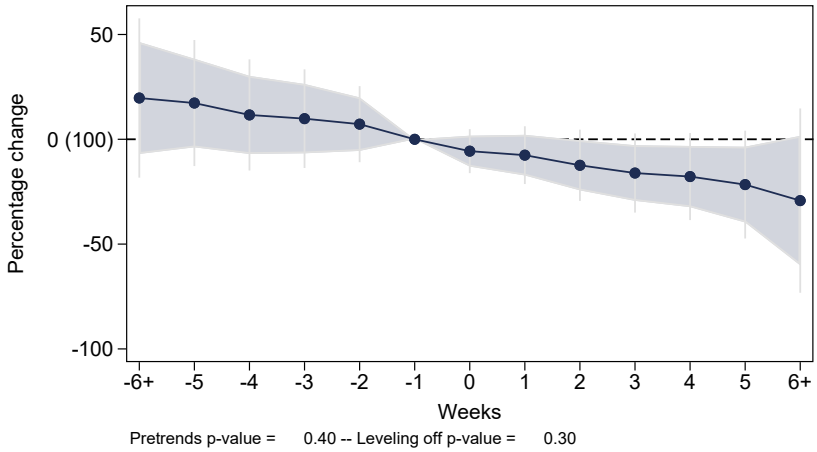

(C) Act III: Policy Reversal

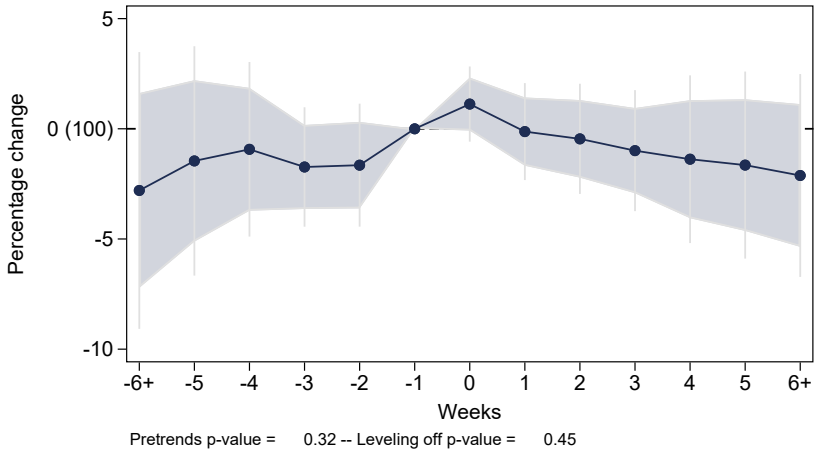

Notes: The panels report the event study estimates from Equation (1) in the three moments of the policy lifetime with wholesale producer prices of treated goods. Panel (A) plots the estimated coefficients for the announcement week (week 12, 2023), panel (B) plots for the implementation week (week 15, 2023), and panel (C) plots for the reversal week (week 1, 2024). The shaded areas correspond to 95% confidence intervals.

7 Conclusion

In this paper, we investigate the pass-through of a temporary cut in VAT to consumer prices, using a novel high-frequency dataset of online retail prices in Portugal. The decrease in VAT affected a subset of food items in 2023 for around 9 months and we analyze the price dynamics across the complete policy lifetime.

Initially, the policy announcement led to an anticipatory price adjustment, where relative prices of treated items increased by 1.27% relative to non-treated ones. In the second act, the policy implementation, we observe an almost complete pass-through to consumer prices, amounting to 99%, underscoring the effectiveness of such fiscal measures in directly influencing retail prices and, by extension, alleviating consumer cost burdens. Notably, we estimate that the VAT cut reduced inflation by 0.7 pp. This high pass-through exhibits a high persistence and only by the end of the policy do we see the pass-through decreasing. When the policy is reverted and VAT returned to the original rate, the price increase implies a smaller pass-through than expected by about 70%.

Collectively, these findings not only shed light on the direct impacts of VAT adjustments on price levels but also highlight the critical importance of timing, communication, and market expectations in maximizing the efficacy of public policies aimed at managing inflation. Building upon the foundations laid in this manuscript, future avenues for exploration emerge, particularly in understanding the underlying mechanisms that drive differential pass-through rates across various product categories and market conditions. Investigating the role of market competition, consumer awareness, and retailer strategies during the different phases of tax adjustments could provide deeper insights into the optimal design and timing of such tax policies. Additionally, comparative studies across different economies and tax regimes could shed light on contextual factors that influence the efficacy and efficiency of VAT adjustments as a tool for managing inflation and supporting economic stability.

References

Abadie, A. (2021). Using synthetic controls: Feasibility, data requirements, and methodological aspects. *Journal of Economic Literature 59*(2), 391–425.

Amores, A. F., S. Barrios, R. Speitmann, and D. Stoehlker (2023). Price effects of temporary vat rate cuts: Evidence from spanish supermarkets. Technical report, European Commission.

Arce, I. M. and M. G. de Antonio (2020). Una evaluación sobre los efectos de una reducción del tipo de gravamen del iva para los bienes y servicios culturales. *Papeles de trabajo del Instituto de Estudios Fiscales. Serie economía* (4), 1–44.

Arkhangelsky, D., S. Athey, D. A. Hirshberg, G. W. Imbens, and S. Wager (2021). Synthetic difference-in-differences. *American Economic Review 111*(12), 4088–4118.

Benzarti, Y., D. Carloni, J. Harju, and T. Kosonen (2020). What goes up may not come down: Asymmetric incidence of value-added taxes. *Journal of Political Economy 128*(12), 4438–4474.

Benzarti, Y., S. Garriga, and D. Tortarolo (2023). Can vat cuts dampen the effects of food price inflation?

Benzarti, Y. and A. Tazhitdinova (2021). Do value-added taxes affect international trade flows? evidence from 30 years of tax reforms. *American Economic Journal: Economic Policy 13*(4), 469–489.

Besley, T. J. and H. S. Rosen (1999). Sales taxes and prices: an empirical analysis. *National Tax Journal 52*(2), 157–178.

Binder, C. and R. Kamdar (2022). Expected and realized inflation in historical perspective. *Journal of Economic Perspectives 36*(3), 131–155.

Checherita-Westphal, C. D., N. Leiner-Killinger, and T. Schildmann (2023). Euro area inflation differentials: the role of fiscal policies revisited.

Clarke, D., D. Pailañir, S. Athey, and G. Imbens (2023). Synthetic difference in differences estimation. Technical report, arXiv.

Fedoseeva, S. and E. Van Droogenbroeck (2024). Temporary vat rate cuts and food prices in e-commerce. *Journal of Retailing and Consumer Services 77*, 103693.

Forteza, N., E. Prades, and M. Roca (2023). Analysing the vat cut pass-through in spain using supermarkets' web-scraped data and machine learning techniques. Technical report.

Freund, C. and J. Gagnon (2017). Effects of consumption taxes on real exchange rates and trade balances. *Peterson Institute for International Economics Working Paper* (17-5).

Freyaldenhoven, S., C. Hansen, J. P. Pérez, and J. M. Shapiro (2021). Visualization, identification, and estimation in the linear panel event-study design. Technical report, National Bureau of Economic Research.

Gautier, E., M. Marx, and P. Vertier (2023). How do gasoline prices respond to a cost shock? *Journal of Political Economy: Macroeconomics*.

Harju, J. and T. Kosonen (2014). The inefficiency of reduced vat rates: Evidence from restaurant industry. *Government Institute for Economic Research*.

Kosonen, T. (2015). More and cheaper haircuts after vat cut? on the efficiency and incidence of service sector consumption taxes. *Journal of Public Economics 131*, 87–100.

Montag, F., R. Mamrak, A. Sagimuldina, and M. Schnitzer (2023). Imperfect price information, market power, and tax pass-through. *George J. Stigler Center for the Study of the Economy & the State Working Paper* (337).

Pfäuti, O. (2023). The inflation attention threshold and inflation surges. *Working Paper*.

Poterba, J. M. (1996). Retail price reactions to changes in state and local sales taxes. *National Tax Journal 49*(2), 165–176.

A Temporary VAT Cut in Three Acts: Announcement, Implementation, and Reversal

Online Appendix

Tiago Bernardino, IIES – Stockholm University

Ricardo Duque Gabriel, Federal Reserve Board

João Quelhas, Banco de Portugal

Márcia Silva-Pereira, Nova SBE

Appendix A Consumer Prices Data

Table A.1: List of items covered by the VAT cut

Category	Items	VAT cut size
Cereals and Tubers	Bread, Potato, Pasta and Rice	6%
Dairy Products	Cow's Milk, Yogurt or Fermented Milk, Cheese	6%
Fruits	Apple, Banana, Orange, Pear, Melon	6%
Legumes	Red Beans, Black-Eyed Peas, Chickpeas	6%
Vegetables	Onion, Tomato, Cauliflower, Lettuce, Broccoli, Carrot, Zucchini, Leek, Pumpkin, Turnip Tops, Portuguese Cabbage, Spinach, Turnip, Peas	6%
Meat and Fish	Pork, Chicken, Turkey, Beef, Codfish, Sardine, Hake, Horse mackerel, Sea Bream, Mackerel	6%
Fats and Oils	Olive Oil, Butter	6%
	Vegetable Oils	23%
Other Products	Canned Tuna, Chicken Eggs, Plant-Based Drinks and Yogurts, Gluten-Free Products	6%

Appendix B Producer Prices Data

Table B.2: Product categories available in the wholesale price dataset

COICOP 5	Category
CP01112	Flour and other cereals
CP01121	Beef
CP01122	Pig meat
CP01124	Poultry
CP01147	Eggs
CP01153	Olive oil
CP01161	Fresh or refrigerated fruit
CP01171	Fresh or refrigerated vegetables except potatoes and other tubers
CP01174	Potatoes

Appendix C Data Cleaning

Our data consists of web-scrapped prices from 5 different supermarkets. The exact start date of the price collection varies by supermarket but it started in the second semester of 2021. The raw data consists of 60,445 items (which correspond to product $\times$ supermarket – a product sold in different supermarkets has a different identifier) and 44,779,544 observations over time.

We restrict to observations after December 2022 leaving us with a total of 24,272,188 observations. During the first weeks of the policy, some goods raised doubts regarding their inclusion in the zero VAT basket. We exclude those items from our sample decreasing the total number of observations to 19,399,736. As Table B.2 indicates the VAT cut size was 6% for all goods except vegetable oils that were taxed at the rate of 23% before the policy. In our baseline analysis, we remove them from our sample to our estimates to be comparable. This reduces the total number of observations to 19,382,934.

We remove outlier observations, in particular, the last day of each month from a specific supermarket, when prices consistently experience sudden hikes exclusively. For our baseline analysis, we only consider food products for the control group. This leaves us with a total of 27,780 products and 10,589,024 observations, corresponding to almost a quarter of the total observations in the raw data.

Table C.3 summarizes the above information.

Table C.3: Data cleaning steps

Step	Description	Products	Observations	% Obs.
0	Web scraped products	60,445	44,799,544	100%
1	Keep observations after Dec-22	60,445	24,272,188	54.2%
2	Drop products that raised doubts	48,474	19,399,736	43.3%
3	Drop vegetable oils	48,433	19,382,934	43.3%
4	Drop outlier observations	48,433	18,499,204	41.3%
5	Drop non-food product	27,780	10,589,024	23.6%

The cleaned dataset of supermarket data tracks well the official food inflation series. Figure C.1 plots the food CPI index reported by Statistics Portugal (INE) and the CPI index implied by our dataset. The latter series is the monthly average of the price indexes of a given month.

Figure C.1: Comparison of the Food Price Index

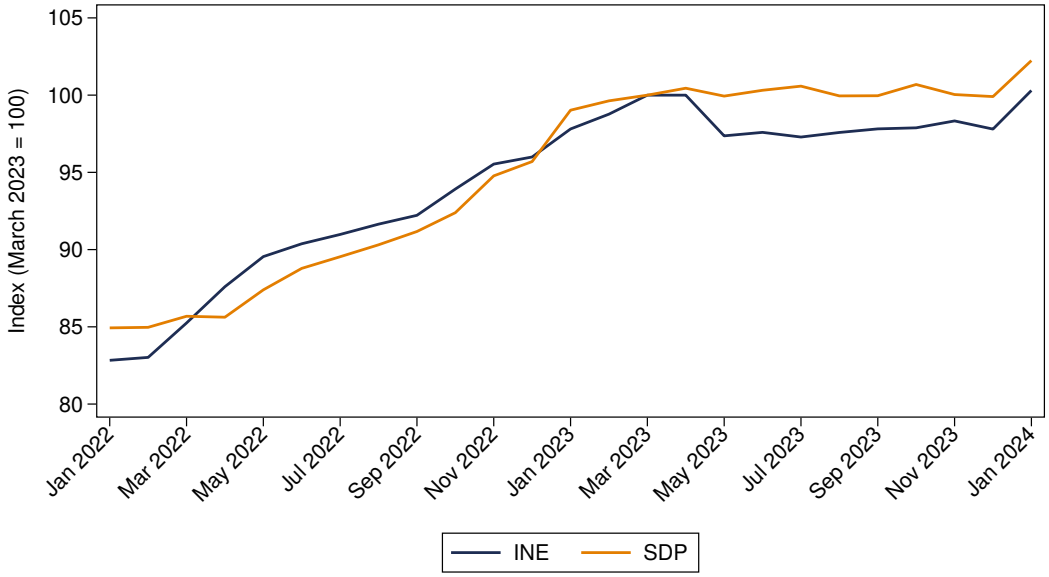

Notes: The two series report the food consumer price index. The blue line corresponds to the series computed by Statistics Portugal and it is the official series that is available in Eurostat. The orange line corresponds to the unweighted monthly average food price index implied by our data. Both series were normalized such that both price indices in March 2023 correspond to 100.

Appendix D Further Analysis

D.1 Heterogeneous pass-through

We analyze the pass-through heterogeneity across food categories. Figure D.2 reports the results across the different categories. More specifically, we look at the price change of products in the same 5-digit COICOP category, using the same control group as in the main event study exercise. In practice, we estimate the coefficients in Equation (1) but restrict the treatment group to products of one category. We look at the price dynamics in the same three acts of the policy lifetime, i.e., 1 week around the announcement, the implementation, and the reversal.

In the announcement day, for the majority of the food categories, the estimated coefficient is not different from the average effect of 1.27%. Yogurt and rice seem to have a price increase slightly higher than the average price increase. On the other hand, fresh skimmed milk, fresh vegetables, butter, and potatoes have a smaller change at the announcement such that one cannot reject a zero price change for these categories at the announcement moment. Looking at the reversal, we also see that the majority of categories have a price adjustment that is not statistically different from the average effect of 5.56%. The exceptions are bread, fresh fruit, frozen vegetables, and yogurt which have a larger price fall and fresh fish, frozen fish, rice, and cheeses which have a smaller price change. Importantly, all categories but one have a coefficient statistically positive. For a 95% confidence interval, the fresh fish coefficient includes zero. Similar to the previous two acts, at the policy reversal most of, the categories had a price increase close to the average increase of 4.18%. Olive oil, poultry, and cheese had a statistically higher increase, whereas beef had a smaller increase. In fact, for beef and for eggs we cannot reject the hypothesis that prices did not increase at all. In sum, different products have different levels of pass-through at different moments of the policy life cycle. In other words, there is no specific pattern that emerges across categories.

Figure D.2: Heterogeneity by food category in the VAT Cut pass-through during the Policy Lifetime

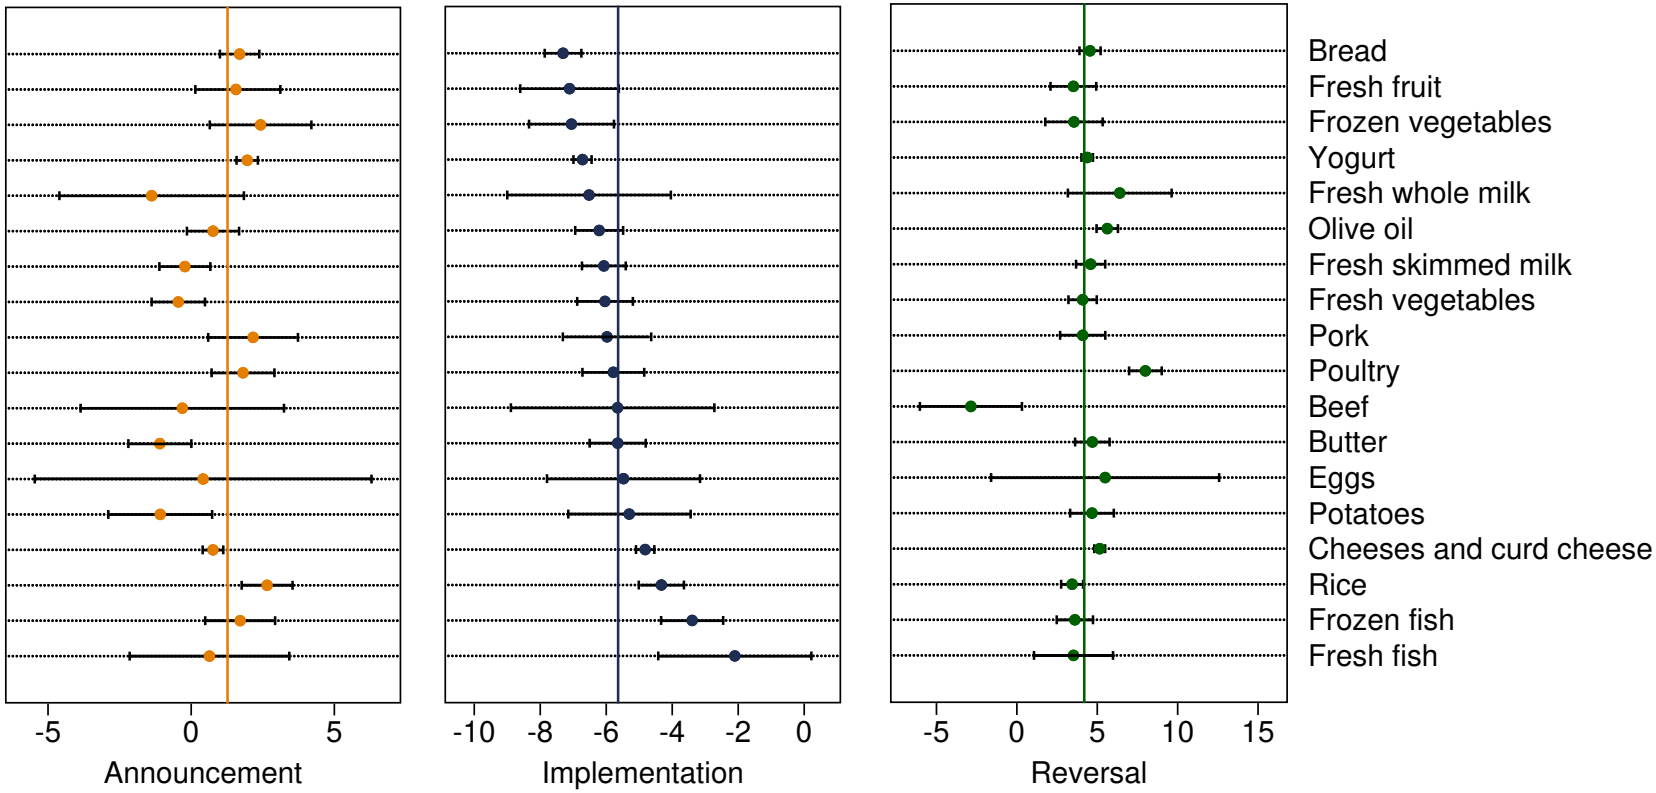

Notes: The panels report the event study estimates from Equation (1) in three moments of the policy lifetime for each food category (COICOP 5-digit level). The left panel reports the coefficients on the announcement day (March 24, 2023), the central panel on the implementation day (April 18, 2023), and the right panel on the policy reversal day (January 4, 2024). Categories are ordered according to the size of the implementation coefficients. The vertical bars on each panel correspond to the aggregate coefficient for each of the three acts.

D.2 Synthetic Difference-in-Differences

We use the Synthetic Difference-in-Differences (SDiD) method to estimate the causal effects of the VAT cut in Portugal as a robustness check. This method re-weights and matches pre-exposure trends to weaken the reliance on parallel trend-type assumptions (Synthetic Control – SC) and is invariant to additive unit-level shifts, allowing for valid large-panel inference (Difference-in-Differences – DiD).

In a conventional DiD analysis, changes in an outcome between a treatment group receiving the intervention and a control group are compared over time. The assumption is that any differences between the two groups can be attributed to the treatment. The SDiD introduces a more flexible approach. Like DiD models, SDiD allows for treated and control units to be trending on entirely different levels prior to a reform of interest. And like SC methods, SDiD seeks to optimally generate a matched control unit which considerably loosens the need for parallel trend assumptions. Correspondingly, SDiD avoids common pitfalls in standard DiD and SC methods — namely the inability to estimate causal relationships if parallel trends are not met in the case of DiD, and a requirement that the treated unit be housed within a "convex hull" of control units in the case of SC.

A major advantage of using SDiD in this particular event is that we "only" need to match the price dynamics of each treated food item with other food items. By using one full year as a pre-treatment period we are able to capture seasonality patterns that are very important to accurately estimating a proper counterfactual for each food item. We follow 4 steps:

1. **Data collection**: assemble price data on the treatment and control groups before and after the intervention takes place.

2. **Construct the synthetic control**: use statistical methods to create a synthetic control group that resembles the treatment group's pre-treatment characteristics. This involves selecting and weighing appropriate control group units with non-negative weights.

3. **Estimation**: compare the changes in the outcome variable over time between the treatment group and the synthetic control group after the intervention. The estimated difference between these is the causal effect of the treatment.

4. **Sensitivity analysis**: study potential anticipation effects of the policy at the announcement and use different donor pools to construct the control group.

We use the computational implementation of a SDiD estimator of Arkhangelsky et al. (2021) for Stata by Clarke et al. (2023), the $sdid$ command, and focus on the case where there are multiple treatment units and a single treatment period.

We thus start with a balanced panel of N units observed over T time periods. In this we use a constant basket with products available for the whole time window in analysis. The outcome variable, price index, denoted P_{it}, is observed for each unit i in each period t. Some, but not all, of these observations are treated with a specific binary variable of interest, denoted W_{it}. This treatment variable $W_{it} = 1$ if observation i is treated by time t, otherwise, $W_{it} = 0$ indicates that unit i is untreated at time t. Here, we assume that there is a single adoption period for treated units, which Arkhangelsky et al. (2021) refer to as a 'block treatment assignment'. Once treated, units are assumed to remain exposed to treatment forever thereafter. Not always treated units can be included in estimation as the method requires at least two pre-treatment periods off of which to determine control units (Clarke et al. 2023).

The goal of SDiD is thus to consistently estimate the causal effect of the VAT cut for the treated food items W_{it}, (an average treatment effect on the treated, henceforth ATT) even if we do not believe in the parallel trends assumption between all treatment and control units on average.

The estimation of the ATT proceeds as follows:

$$\left(\hat{\tau}^{\text{sdid}}, \hat{\mu}, \hat{\alpha}, \hat{\gamma}\right) = \arg\min_{\tau,\mu,\alpha,\gamma} \left\{ \sum_{i=1}^{N} \sum_{t=1}^{T} \left(P_{it} - \mu - \alpha_i - \gamma_t - W_{it}\tau\right)^2 \hat{\omega}_i^{\text{sdid}} \hat{\lambda}_t^{\text{sdid}} \right\}, \tag{D.3}$$

where the estimand is the ATT, generated from a two-way fixed effect regression, with optimally chosen weights ω_i^{sdid} and λ_t^{sdid} discussed below. This procedure flexibly allows for shared temporal aggregate factors given the estimation of time-fixed effects γ_t and time-invariant unit-specific factors given the estimation of unit-fixed effects α_i. As is standard in fully saturated fixed-effect models, one α_i and one γ_t fixed effect are normalized to zero to avoid mul-

ticollinearity. The presence of unit-fixed effects implies that SDiD will seek to match treated and control units on pre-treatment trends, and not necessarily on both pre-treatment trends and levels, allowing for a constant difference between treatment and control units.

As laid out in Arkhangelsky et al. (2021), the selection of unit weights, ω seeks to ensure that the comparison is made between treated units and controls which were approximately following parallel trends prior to the adoption of treatments. The selection of time weights, λ seeks to draw more weight from pre-treatment periods that are more similar to post-treatment periods, in the sense of finding a constant difference between each control unit's post-treatment average and pre-treatment weighted averages across all selected controls. Given that we have a large number of treated and control units, we can use the estimator proposed by Arkhangelsky et al. (2021) without reservation and estimate its variance using a block bootstrap approach.

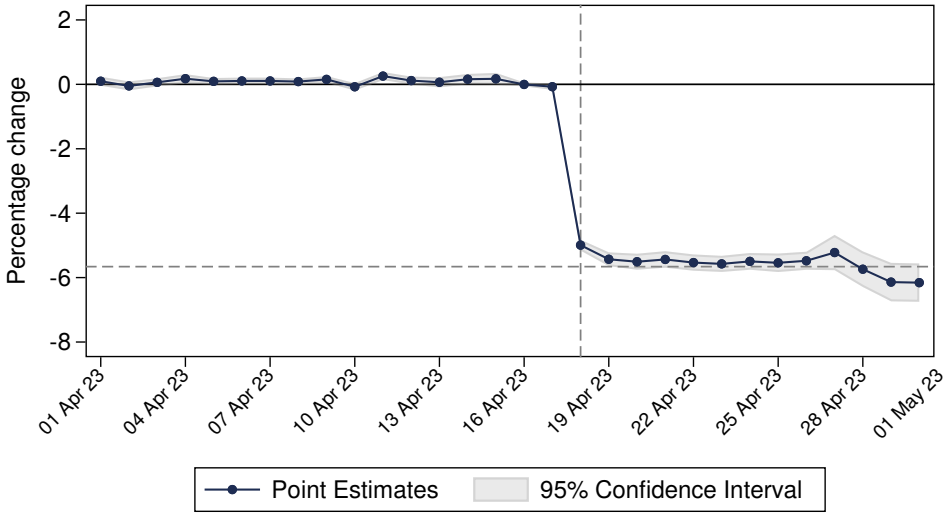

Figure D.3: Policy Implementation using the SDiD

Notes: The figure reports the SDiD estimates from Equation (D.3) for the implementation day (April 18, 2023). The shaded areas correspond to 95% confidence intervals and are computed using bootstrap inference. The vertical gray dashed line corresponds to the implementation of the VAT cut in April 18, 2023. The horizontal gray dashed line corresponds to the percentage change associated to a full pass-through of the VAT cut.

D.3 Other Robustness Exercises

Table D.4: Robustness results for the Three Acts of the Temporary VAT Cut in Portugal

Test	Description	Announcement	Implementation	Reversal
1	Including all products	1.20	-5.86	4.29
2	Incl. only food (no drinks)	0.94	-5.38	3.68
3	Incl. only non-food products	1.43	-6.29	4.69
4	Price per unit	0.97	-5.37	4.22
5	Regular Price	0.27	-4.68	4.04
6	Data Imputation (1)	1.19	-5.41	3.53
7	Data Imputation (2)	0.97	-5.76	3.39
8	Constant Basket	1.27	-5.87	5.15
	Average pass-through	**18.7%**	**98.2%**	**68.5%**

Notes: The coefficients presented correspond to the difference between post- and pre-treatment averages of the event study estimates from Equation (1). Each test is described in Section 4.3. The first imputation method carry-forwards the information from the previous day for each product for a maximum of 7 days if the price in the previous day with available information is equal to the price in the next day with available information until the product exits. The second imputation method carry-forwards the information from the previous day for each product until the product exits without any time limitation for the missing gap.

Appendix E Additional Figures

Figure E.4: Policy implementation context

(A) HICP of food and all items in Portugal

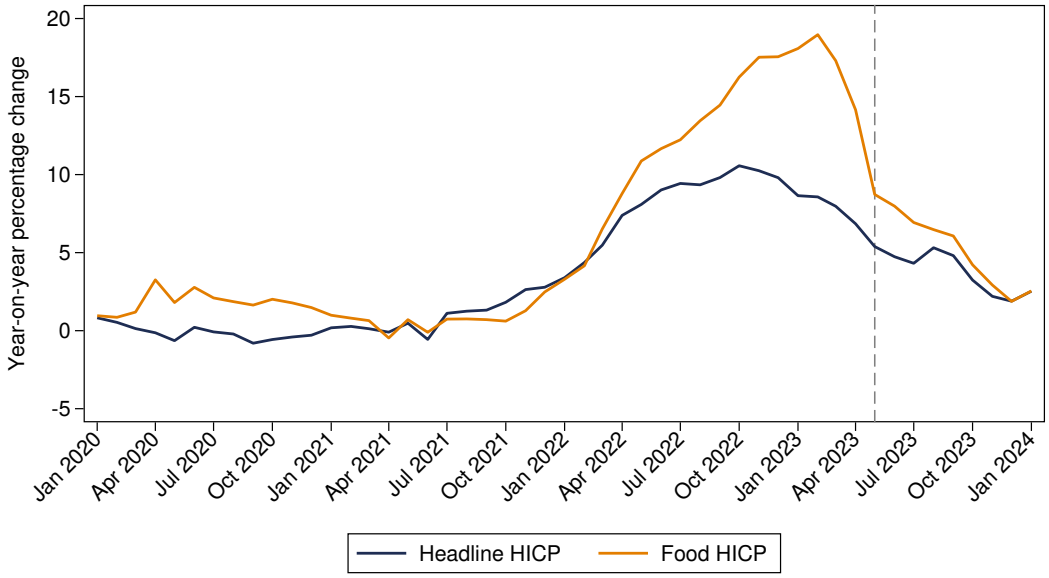

Notes: The gray dashed line corresponds to the implementation of the VAT cut in April 2023.

(B) Search intensity on Google for Value-Added Tax and VAT in Portugal

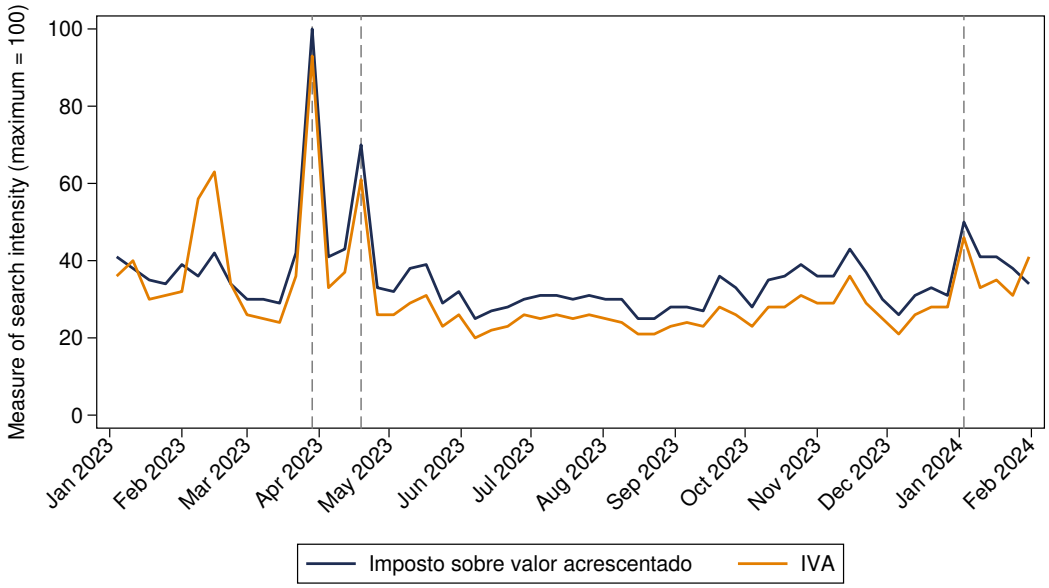

Notes: The first gray dashed line corresponds to the announcement of the VAT cut in week 12 of 2023 and the second to the actual policy implementation, week 15 of 2023 and the third line corresponds to the reversal in week 1 of 2024. *Imposto sobre o valor acrescentado* translates to Value-Added Tax and *IVA* to VAT in Portuguese. Data retrieved from Google trends website.

Figure E.5: Price evolution of the basket of treated items in Portugal and Spain

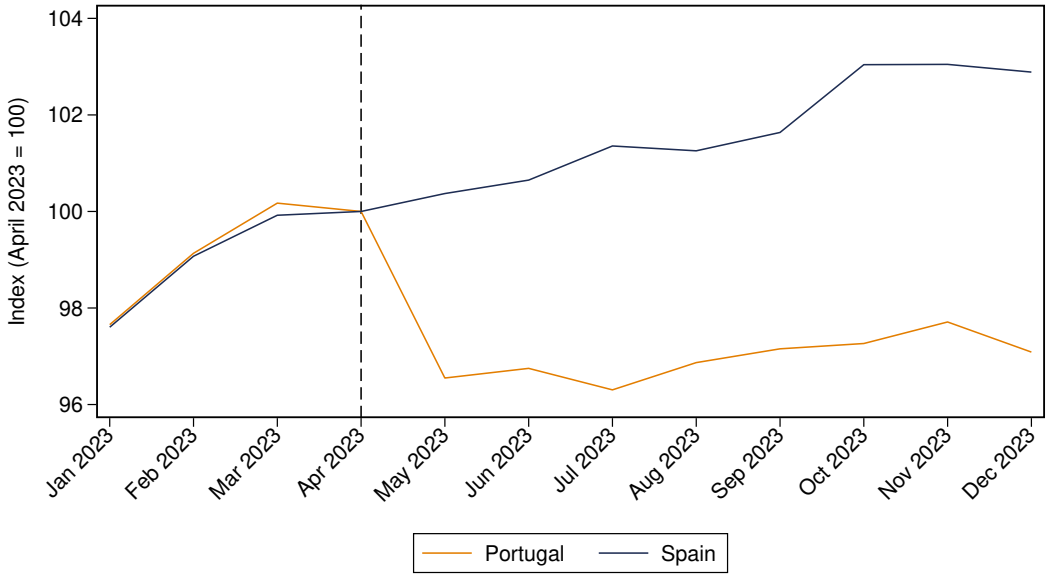

Notes: The two series were obtained by aggregating the price indices of the treated COICOP 5-digit categories for both countries using their corresponding weights for 2023 as defined by the Eurostat. The gray dashed line corresponds to the implementation of the VAT cut in April 2023.

Figure E.6: The Three Acts of the Temporary VAT Cut in Portugal for Vegetable Oils

(A) Act I: Policy Announcement

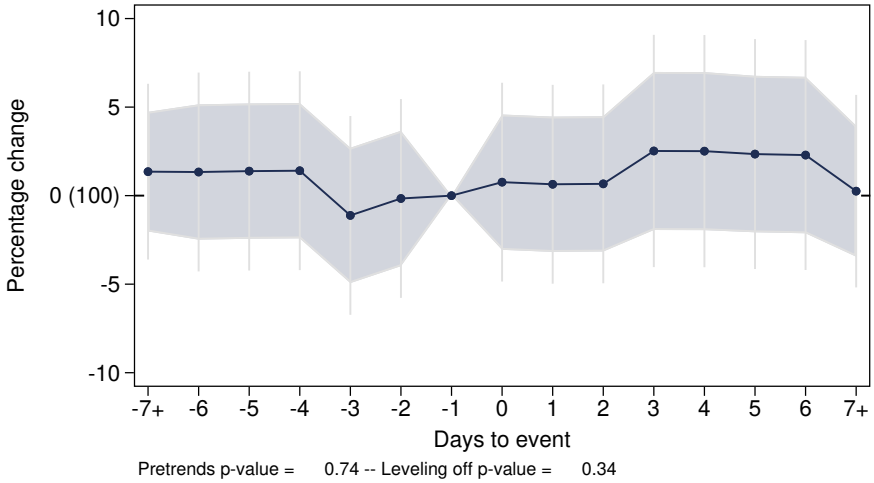

(B) Act II: Policy Implementation

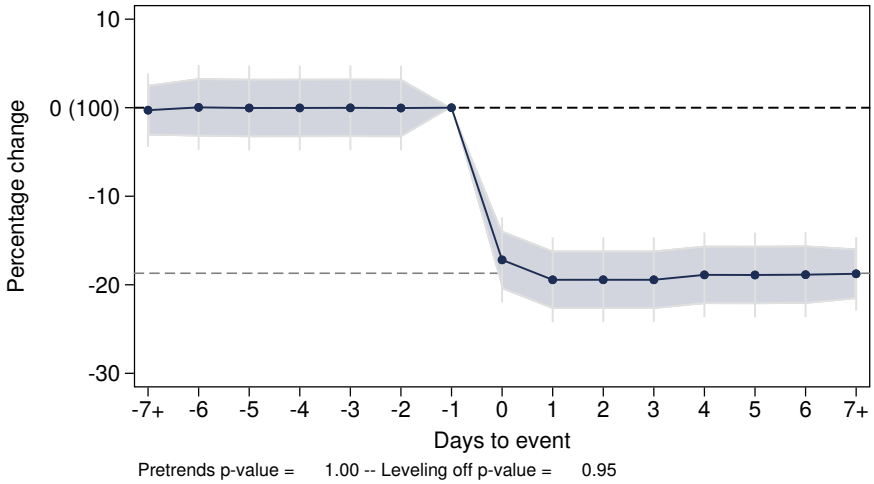

(C) Act III: Policy Reversal

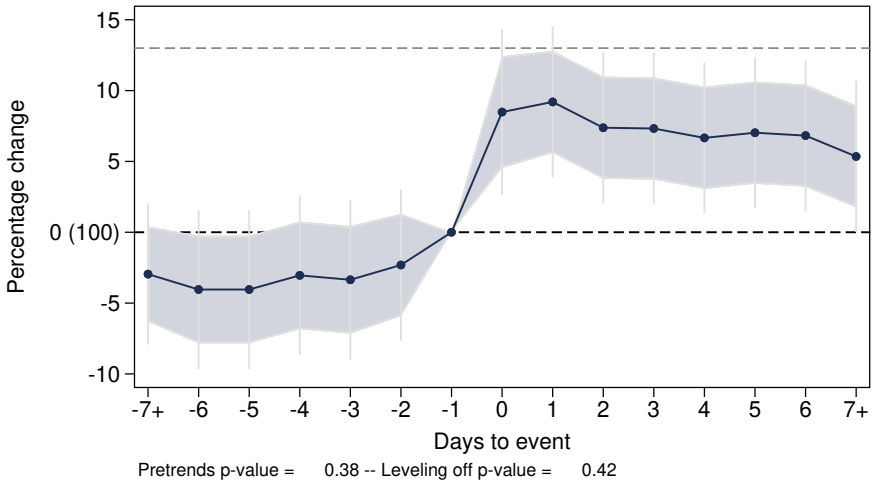

Notes: The panels report the event study estimates from Equation (1) in three moments of the policy lifetime. Panel (A) plots the estimated coefficients for the announcement day (March 24, 2023), panel (B) plots for the implementation day (April 18, 2023), and panel (C) plots for the reversal day (January 4, 2024). The shaded areas correspond to 95% confidence intervals.